

An Expressive Mechanism for Auctions on the Web

Paul Dütting* Monika Henzinger† Ingmar Weber‡

October 29, 2010

Abstract

Auctions are widely used on the Web. Applications range from internet advertising to platforms such as eBay. In most of these applications the auctions in use are single/multi-item auctions with unit demand. The main drawback of standard mechanisms for this type of auctions, such as VCG and GSP, is the *limited expressiveness* that they offer to the bidders. The General Auction Mechanism (GAM) of [1] is taking a first step towards addressing the problem of limited expressiveness by computing a bidder optimal, envy free outcome for linear utility functions with identical slopes and a single discontinuity per bidder-item pair. We show that in many practical situations this does not suffice to adequately model the preferences of the bidders, and we overcome this problem by presenting the first mechanism for *piece-wise linear* utility functions with *non-identical slopes* and *multiple discontinuities*. Our mechanism runs in polynomial time. Like GAM it is incentive compatible for inputs that fulfill a certain non-degeneracy requirement, but our requirement is more general than the requirement of GAM. For discontinuous utility functions that are non-degenerate as well as for continuous utility functions the outcome of our mechanism is a *competitive equilibrium*. We also show how our mechanism can be used to compute approximately bidder optimal, envy free outcomes for a general class of continuous utility functions via piece-wise linear approximation. Finally, we prove hardness results for even more expressive settings.

1 Introduction

Auctions are widely used on the Web. They are, e.g., used by Google, Microsoft, and Yahoo! for search advertising [37]. They are also used on platforms such as eBay for selling a broad variety of goods and services [43, 8]. In most of these applications the auctions in use are single/multi-item auctions with unit demand. The problem solved by

*Ecole Polytechnique Fédérale de Lausanne (EPFL), Station 14, CH-1015 Lausanne, Switzerland. Email: paul.duetting@epfl.ch.

†University of Vienna, Universitätsstrasse 10/11, A-1090 Vienna, Austria. Email: monika.henzinger@univie.ac.at.

‡Yahoo! Research, Avinguda Diagonal 177, E-08018 Barcelona, Spain. Email: ingmar@yahoo-inc.com.

these auctions is essentially a *matching* and *pricing* problem. In this problem n bidders have to be matched to k items. Each bidder has a utility function $u_{i,j}(p_j)$ that expresses his utility for being matched to item j at price p_j . An outcome (μ, p) consisting of a matching μ and prices p is said to be *envy free* if at the current prices every bidder prefers the item that he is currently matched to over every other item.¹ An envy free outcome (μ, p) is called *bidder optimal* if the utility of every bidder is at least as high as in every other envy free outcome (μ', p') . Mechanisms that compute a bidder optimal, envy free outcome are desirable because they not only guarantee that everyone is “happy with what he gets”, but also that everyone is “as happy as possible”. From an economic point of view a bidder optimal, envy free outcome in which all unsold items have price zero is desirable because it represents a *competitive equilibrium* (or *Walrasian equilibrium*) [42].

1.1 Limitations of Current Mechanisms

Standard mechanisms for auctions on the web, such as First- and Second Price for single-item auctions and Vickrey Clarkes Groves (VCG) [49, 14, 32] and Generalized Second Price (GSP) [26, 47] for multi-item auctions, nicely fit into the above model. For *linear* utility functions with *identical* slopes, i.e., utilities of the form $u_{i,j}(p_j) = v_{i,j} - p_j$, where $v_{i,j}$ denotes bidder i 's valuation for item j , the outcome of VCG, for example, is envy free and bidder optimal [38].

The main drawback of these standard mechanisms is the limited *expressiveness* that they offer to the bidders. In particular, (1) they typically limit what *functions* the bidders can use to express their utility for receiving a given item at a given price. In ad auctions, for example, some of the bidders may have per-click valuations, while others have per-impression valuations. Mechanisms that can handle either type exist (see, e.g., [26, 47] or [40, 28]), but mechanisms that can handle both types simultaneously are still rare [31]. Moreover, (2) they do not allow to express *budgets*. Budgets can either be *soft* (a limited amount of cash after which the bidder has to take out a loan) or *hard* (an upper bound on the price the bidder is willing to pay). Budgets are considered an important and difficult problem that recently has received a lot of attention (see, e.g., [11, 23]).

The General Auction Mechanism (GAM) of [1] is taking a first step towards addressing the problem of limited expressiveness by allowing the bidders to specify a *maximum price* for each item. That is, it finds a bidder optimal, envy free outcome for *linear* utilities with *identical slopes* and a *single discontinuity* per bidder-item pair. More specifically, the utility functions are of the form $u_{i,j}(p_j) = v_{i,j} - p_j$ if $p_j \leq m_{i,j}$ and $u_{i,j}(p_j) = -\infty$ otherwise, where $m_{i,j}$ denotes the bidder-item specific maximum price. GAM requires the input to be in *general position* (see [1] and Section 4 for a generalization of this concept). In general position it also preserves a very desirable property of the original model. Namely, no bidder can misreport his valuations and/or maximum prices to achieve a strictly higher

¹In the related literature (see, e.g., [1]) such outcomes are also referred to as *stable*.

utility. This makes it a (weakly) dominant strategy for the bidders to report truthfully. Mechanisms with this property are said to be *incentive compatible* (or *strategy proof*) [42].

Despite its generality GAM has three major limitations: (1) It can only handle linear utility functions with identical slopes. (2) It can only handle a single discontinuity with a jump to $-\infty$ per bidder-item pair. (3) It cannot handle non-linear utility functions. We illustrate why and when these shortcomings are problematic by means of three examples.

Example 1: Per-click vs. per-impression valuations. (This example motivates linear utilities with non-identical slopes.) Consider an ad auction with bidders with per-click valuations $v_{i,j}^{click}$ and bidders with per-impression valuations $v_{i,j}^{imp}$. Suppose that the mechanism collects per-click valuations and charges per-click prices. That is,

$$u_{i,j}(p_j^{click}) = v_{i,j}^{click} - p_j^{click}. \quad (1)$$

A bidder with per-impression valuations can translate his valuations into per-click valuations using the click through rate $ctr_{i,j}$ as follows: $v_{i,j}^{click} = v_{i,j}^{imp}/ctr_{i,j}$. That is, he reports $u_{i,j}(p_j^{click}) = v_{i,j}^{imp}/ctr_{i,j} - p_j^{click}$. Now suppose that given the per-click valuations, the mechanism computes an envy free outcome (μ, p^{click}) consisting of a matching μ and per-click prices p^{click} . That is, for every matched bidder-item pair $(i, j) \in \mu$ and all items $j' \neq j$ we have:

$$v_{i,j}^{click} - p_j^{click} \geq v_{i,j'}^{click} - p_{j'}^{click}. \quad (2)$$

What we actually want for bidders i with per-impression valuations $v_{i,j}^{imp}$ is that for $(i, j) \in \mu$ and all $j' \neq j$:

$$v_{i,j}^{imp} - p_j^{imp} \geq v_{i,j'}^{imp} - p_{j'}^{imp} \quad (3)$$

But if we take (2), replace $v_{i,j}^{click}$ with $v_{i,j}^{imp}/ctr_{i,j}$, p_j^{click} with $p_j^{imp}/ctr_{i,j}$, and multiply by $ctr_{i,j}$, then we get

$$v_{i,j}^{imp} - p_j^{imp} \geq C \cdot (v_{i,j'}^{imp} - p_{j'}^{imp}), \quad (4)$$

where $C = ctr_{i,j}/ctr_{i,j'}$. That is, if $C < 1$, then (4) is *not* strong enough to guarantee envy freeness for per-impression bidders. With non-identical slopes this can be sidestepped by having bidders with per-impression valuations report

$$u_{i,j}(p_j^{click}) = v_{i,j}^{imp} - ctr_{i,j} \cdot p_j^{click}. \quad (5)$$

In this case the outcome (μ, p^{click}) computed by the mechanism will be envy free for both types of bidders, i.e., the above problem does *not* arise.

Example 2: Soft and hard budgets. (This example motivates piece-wise linear utilities with non-identical slopes and multiple discontinuities.) Suppose that bidder i wants to buy

a car on eBay. In the current system it would be dangerous for i to bid on more than one car at the same time due to the risk of winning and having to pay for several cars when one is already enough. With a GAM-like auction i could bid on many cars at the same time while still being guaranteed that he gets at most one.

With expensive items, such as cars, i 's valuation $v_{i,j}$ for item j may exceed the amount c of cash that he possesses. In this case i might be willing to take out a loan if the price p_j of item j exceeds c . Assume that i is offered a loan with a maximum amount of a , an interest rate of r , and a fixed fee of f . Then i 's utility for item j has the following form: (1) Because no interest is due for the first c dollars the utility function drops linearly with a slope of -1 from 0 to c . (2) At c it drops by the fixed fee f . (3) Afterwards, due to the interest rate r , every dollar spent causes $1 + r$ dollars in actual cost. Hence the utility function drops linearly with a slope of $-(1 + r)$ from c to $c + a$.

In addition to the soft budget constraint c , bidder i may have a hard budget constraint $b < c + a$, which is typically modeled by a jump to $-\infty$ at $p_j = b$. Hence bidder i 's utility function for item j ultimately looks like this:




Figure 1: Bidder i 's utility function for item j .

Note that import duties that apply when bidder i wants to purchase car j abroad and its price p_j exceeds a certain threshold give rise to similar utility functions.

Example 3: Risk aversion. (This example motivates arbitrary, non-linear utilities.) In the previous example bidder i may also be faced with costs for uncertain events after the purchase, such as potential car repairs. A risk-averse buyer would have a utility function that drops super-linearly in the price as a higher price is associated with a higher risk of not having sufficient money to pay for the unforeseen events.

1.2 Our Contributions

We overcome the limitations of GAM as follows: (1) We present the first mechanism for *piece-wise linear* utility functions with *non-identical slopes* and *multiple discontinuities*. The problem solved by our mechanism is *more difficult* than that solved by GAM as

non-identical slopes require the prices to be increased by different amounts, and multiple discontinuities can cause previously matched bidders to become unmatched several times. Our mechanism is *more expressive* than GAM as it can simultaneously auction off items to bidders with per-click and per-impression valuations (Example 1), and it can handle a variety of soft and hard budget constraints (Example 2). The running time of our mechanism is *polynomial* in the number of bidders n and items k .² It is linear in the number of discontinuities D and constant-slope intervals T of the utility functions. (2) We precisely characterize under which conditions *any* mechanism that computes a bidder optimal, envy free outcome in this setting is *incentive compatible*. This characterization involves a generalization of the general position concept of [1]. For discontinuous utility functions that are non-degenerate as well as for continuous utility functions the outcome of our mechanism is a *competitive equilibrium*. (3) We show how to use our mechanism to compute a γ -bidder optimal, envy free outcome for a general class of continuous utility functions via piece-wise linear approximation in time linear in $1/\sqrt{\gamma}$. (4) Finally, we prove hardness results for two natural generalizations of our problem. In the first, the utilities may depend on the matching. In the second, they may depend on the whole vector of prices. In both cases computing a bidder optimal, envy free outcome is \mathcal{NP} -hard.

1.3 Related Work

Continuous utility functions. For *linear* utility functions with *identical* slopes the underlying matching and pricing problem was first studied by Shapley and Shubik [45]. They formulated the matching problem as a linear program and observed that the dual program yields envy free prices. With the help of this formulation they also proved the existence of an outcome with smallest prices/largest utilities, which is the bidder optimal outcome. Later Leonard [38] examined the incentives for misreporting and found that the bidder optimal outcome is identical to the outcome of VCG [49, 14, 32] and therefore incentive compatible. The classic mechanism for linear utility functions with identical slopes is the so-called Multi-Item Auction of Demange et al. [20], which is a variant of the so-called Hungarian Method by Kuhn [36]. The basic idea of this mechanism is to start with prices all zero and to repeatedly raise the prices of overdemanded items by the same amount. This idea was generalized to continuous, *piece-wise linear* utility functions with *non-identical slopes* by Alkan [3, 4], who showed that the prices of overdemanded items need to be raised by different amounts and that these amounts can be computed by solving a primal/dual problem. The existence of a bidder optimal outcome for more general, *non-linear* utility functions was shown by Demange and Gale [18] using a lattice-theoretic argument. They also proved that for continuous utility functions any mechanism that finds a bidder optimal outcome is incentive compatible. Recently, Alaei et al. [2] presented a novel, inductive characterization of the bidder optimal utilities/prices in this setting, which yields a constructive proof of existence. Although hardness results have been

²For $n > k$ it is linear in n .

established for related problems (see, e.g., [22, 48]), it is not clear whether or under which conditions a bidder optimal outcome can be found efficiently for such general continuous utility functions.

Discontinuous utility functions. The first to add a *single discontinuity* to otherwise *linear* utility functions with *identical slopes* were Aggarwal et al. [1]. They also gave a mechanism, which - for inputs in general position - is incentive compatible and finds a bidder optimal outcome in polynomial time. Similar results to that of Aggarwal et al. [1] were obtained by [5, 6] and [25]. In [35] it was shown how to find the smallest envy free prices for a *given* matching. Recently, Chen et al. [12] gave a polynomial-time mechanism for *consistent* utility functions. Note that all these results either assume *identical slopes* [1, 6, 25, 35], just a *single discontinuity* [1, 6, 25, 35, 12], or both. Also note that the piecewise linear utility functions with non-identical slopes and multiple discontinuities that we study here are *not* consistent. The existence of a bidder optimal, envy free outcome for more general, *non-linear* utility functions with *multiple discontinuities* was established in [24], but just as in the continuous case no polynomial-time mechanism is known for such general discontinuous utility functions.

Externalities. Our hardness results rely on the fact that bidder i 's utility for being matched to item j may depend on (a) who is assigned which item or (b) the prices of the other items. Such dependencies are referred to as *externalities*. The “classic” result here is due to Jehiel et al. [34], who presented a revenue maximizing auction for the sale of a single item when bidders who do not acquire the item for sale incur a bidder-dependent externality. Other related results are [7] and [29]. The former analyzes Nash equilibria of so-called social context games in which utilities are computed based on an undirected neighborhood graph among players and aggregation functions. The latter proves that the following winner determination problem, which, e.g., arises in the pay-per lead model, is \mathcal{NP} -hard and hard to approximate:³ Choose a set S of at most k advertisers from a set of n advertisers, each with a private value v_i and a random quality variable q_i to maximize the expected value $v(S) = \sum_{i \in S} v_i \cdot \Pr[\forall j \in S \cup \{0\} : q_i \geq q_j]$, where q_0 is the quality of the best outside option. The externality is thus encoded in the choice of S and the fact that the distributions from which the q_i are drawn need not be independent. Our problem is different as there is not a single winner and the externality is either encoded in the matching or the prices.

1.4 Overview of Techniques

Our main result, the generalization of GAM, draws from the rich literature on matching and pricing problems in various ways: First, it uses the same terminology (e.g., strict overdemand, alternating path, alternating tree) and exploits the same link to Hall's theorem

³www.iab.net/lead_generation

[33] that was already used in [20] and in [36]. It also borrows from [3, 4], in which the computation of price increases for continuous, piece-wise linear utility functions with non-identical slopes is formulated as a primal/dual problem, although we refine this approach to significantly improve upon the running time⁴ and extend it to discontinuous utility functions. Finally, our analysis follows a similar approach as [1] to cope with the fact that in the presence of discontinuities bidder optimality and incentive compatibility no longer coincide by characterizing inputs for which this is still the case.

On a conceptual level our main achievement is a decoupling of the argument for bidder optimality from the argument for incentive compatibility. On a technical level it is that we can show bidder optimality even if discontinuities require previously matched bidders to be unmatched. We achieve this through a refinement of Alkan’s technique for increasing the prices when slopes are non-identical (Lemma 7) with a novel argument that shows that all price increases by the mechanism are necessary, even if discontinuities are reached and bidders become unmatched (Lemma 8). Together these lemmata allow us to prove bidder optimality even for inputs which make it necessary to unmatch previously matched bidders. Note that neither [3, 4] nor [1] have to unmatch bidders, because they either assume continuity or restrict themselves to inputs in general position.

Our new insight for incentive compatibility is that if for each price increase at most one utility function reaches a discontinuity, then bidders never get unmatched (Lemma 9). We also observe that in this case (a) all items with price $p_j > r_j$, where r_j denotes an item-dependent reserve price, are matched, and (b) at least one item j that is matched has $p_j = r_j$ (Lemma 9). We then use a variant of Hwang’s Lemma (Lemma 10) to prove that whenever (a) and (b) are satisfied any mechanism that computes a bidder optimal outcome is incentive compatible (Theorem 2). Finally, we formulate a condition on the input that guarantees that for each price increase at most one discontinuity is reached. Following Aggarwal et al. [1] we say that inputs that satisfy this condition are in *general position*, but our condition on the input is more general than that of [1] as it applies to piece-wise linear utility functions with non-identical slopes and multiple discontinuities.

Concerning general continuous utility functions it is intuitively obvious that the solution to a close enough piece-wise linear approximation cannot be far worse than the solution to the original problem. The difficulty here is to get a running time that is polynomial in $1/\gamma$, where γ is the maximal distance between the resulting bidder optimal utilities for the approximate utility functions and for the original utility functions. In fact, it is not difficult to show that $O(1/\sqrt{\epsilon})$ linear segments suffice to ensure that the maximum distance between the approximate utility functions and the original utility functions is at most ϵ (Lemma 11). The hard part here is to show that γ scales linearly in ϵ (Lemma 12).

We establish the hardness result for utility functions that depend on the matching through a reduction from 3-SAT [16] and the hardness result for utility functions that

⁴In [3, 4] the running time is stated as $O(n^2 \cdot k^4 \cdot \prod_{i,j} t_{i,j})$, where $t_{i,j}$ denotes the number of linear segments of $u_{i,j}(\cdot)$.

depend on the whole vector of prices through a reduction from MAX-2-NASH [30, 15].⁵ The reduction from MAX-2-NASH represents a novel and interesting link between the matching and pricing problem studied here and the problem of computing Nash equilibria, whose computational complexity has been settled only recently [17, 13].

All proofs are given in the Appendix.

2 Problem Statement

We are given a set I of n bidders and a set J of k items. The set of items J contains a dedicated dummy item that we denote j_0 . For each bidder i we are given a constant o_i , called the *outside option*, which is the utility that bidder i derives from not getting any non-dummy item. For each item j we are given a constant $r_j \geq 0$, called the *reserve price*, which is a lower bound on p_j . Finally, for each bidder-item pair (i, j) we are given a *utility function* $u_{i,j}(p_j)$, where p_j denotes the price of item j . The utility functions are *piece-wise linear*. That is, each $u_{i,j}(\cdot)$ is composed of $t_{i,j}$ linear segments

$$u_{i,j}^{(t)}(p_j) = v_{i,j}^{(t)} - c_{i,j}^{(t)} \cdot p_j \text{ for } p_j \in [s_{i,j}^{(t)}, e_{i,j}^{(t)}], \quad (6)$$

where $t \in \{1, \dots, t_{i,j}\}$, $s_{i,j}^{(1)} = r_j$, $e_{i,j}^{(t_{i,j})} = \infty$, $s_{i,j}^{(t)} < e_{i,j}^{(t)}$ ($\forall t$), and $e_{i,j}^{(t)} = s_{i,j}^{(t+1)}$ ($\forall t \neq t_{i,j}$). Where possible we omit (t) to improve readability. We make the following assumptions concerning the utility functions: (1) They are strictly monotonically decreasing. (2) They need *not* be globally continuous. (3) For every bidder-item pair (i, j) there exists a threshold value $\bar{p}_{i,j}$ such that $u_{i,j}(\bar{p}_{i,j}) \leq o_i$. (4) The utility functions $u_{i,j_0}(\cdot)$ for the dummy item j_0 are of the form $u_{i,j_0}(p_{j_0}) = o_i - p_{j_0}$ for $p_{j_0} \in [0, \infty)$ and $r_{j_0} = 0$.⁶

We want to compute (1) a subset $\mu \subseteq I \times J$ of the bidder-item pairs, in which (a) every bidder i appears in *exactly one* pair $(i, j) \in \mu$ and (b) every non-dummy item $j \neq j_0$ appears in *at most one* pair. We refer to this set as *matching*, even though multiple bidders i can be matched to the dummy item j_0 . We also want to compute (2) per-item *prices* $p = (p_1, \dots, p_k)$. We refer to the pair (μ, p) as *outcome*. An outcome (μ, p) is *feasible* if

$$p_{j_0} = 0 \text{ and } p_j \geq r_j \text{ for all } j \neq j_0. \quad (7)$$

We say that prices with this property are *feasible*. In the remainder of this paper whenever we refer to prices we mean feasible prices. An outcome is *envy free* if it is feasible and for all i and $(i, j) \in I \times J$,

$$u_{i,\mu(i)}(p_{\mu(i)}) \geq u_{i,j}(p_j), \quad (8)$$

⁵MAX-2-NASH is the problem of finding a Nash equilibrium of a 2-player game, which maximizes the sum of the utilities.

⁶Note that this definition guarantees that in every envy free outcome every bidder i has utility at least o_i .

where $\mu(i)$ denotes the item bidder i is matched to. An outcome (μ, p) is *bidder optimal* if it is envy free and for every bidder i and every envy free outcome (μ', p') we have

$$u_{i, \mu(i)}(p_{\mu(i)}) \geq u_{i, \mu'(i)}(p'_{\mu'(i)}). \quad (9)$$

Our goal is to compute a bidder optimal outcome.

3 Mechanism

In this section we describe and analyze our mechanism for piece-wise linear utilities. We proceed as follows: First, we show how to reduce the problem of finding a bidder optimal outcome for an input with reserve prices to the problem of finding such an outcome for a different input in which the reserve prices are all zero. Second, we prove that the bidder optimal outcome has minimal prices among all envy free outcomes. We then formulate the problem as a graph problem. This allows us to define strict overdemand and to prove that an envy free outcome exists if and only if no set of items is strictly overdemanded using Hall's Theorem [33]. Our mechanism, which can be seen as an algorithmic version of Hall's Theorem [33], starts with prices all zero and iteratively raises the prices of strictly overdemanded items. To ensure minimality of the resulting prices it does this in an envy free and overdemand preserving manner.

Standard form. We say that the input is in *standard form* if $r_j = 0$ for all j . The following lemma shows that we can w.l.o.g. assume that the input is in standard form as for any problem instance that is not in standard form there is a linear-time reduction to an instance in standard form. This reduction is similar to the reduction described in [3]. The lemma also shows that a sufficient condition for an outcome (μ^*, p^*) to be bidder optimal is that the prices p^* are the minimum prices at which an envy free outcome exists. This was already known for *continuous* utility functions (see, e.g., [18]), but it is a novel observation for *discontinuous* utility functions. Moreover, unlike in the continuous case, the reverse need *not* be true for discontinuous utility functions.

Lemma 1. *We have:*

1. *If the outcome (μ, p) is bidder optimal for $u'_{i,j}(p_j) = u_{i,j}(p_j + r_j)$ and $r'_j = 0$, then the outcome (μ, p') with $p'_j = p_j + r_j$ is bidder optimal for $u_{i,j}(p_j)$ and r_j .*
2. *If the outcome (μ^*, p^*) is envy free and $p_j^* \leq p_j$ for all items j and every envy free outcome (μ, p) , then (μ^*, p^*) is bidder optimal.*

Graph-theoretic formulation. Next we formulate the problem of computing an envy free outcome as a graph problem. Central to this formulation is the *first choice graph* $G_p = (I \cup J, F_p)$ at prices p , which consists of 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 at the current prices.

That is, $u_{i,j}(p_j) \geq u_{i,k}(p_k)$ for all k . 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 $F_p(S) = \cup_{j \in S} F_p(j)$. Based on the first choice graph we define *strict overdemand*: A set of non-dummy items $S \subseteq J \setminus \{j_0\}$ is *strictly overdemanded at prices p* with respect to the set of bidders $T \subseteq I$ if (a) $F_p(T) \subseteq S$ and (b) for all $R \subseteq S$ with $R \neq \emptyset$: $|F_p(R) \cap T| > |R|$. A set of non-dummy items $S \subseteq J \setminus \{j_0\}$ is *strictly overdemanded*, if it is strictly overdemanded with respect to some set of bidders $T \subseteq I$. Our definition of strict overdemand is *stronger* than the definition of *overdemand* [20], which only requires that the number of bidders T demanding only items in the set S is greater than the number of items in the set. It is different from the notion of *minimal overdemand* [20], which requires that no subset is overdemanded. It is also different from the notion of *directional overdemand* in [4]. The advantage of our definition will become clear in the next subsection. The following lemma is established using Hall's Theorem [33] and the fact that a strictly overdemanded set of items exists *if and only if* an overdemanded set of items exists.

Lemma 2. *The following statements are equivalent:*

1. *The outcome (μ, p) is envy free.*
2. *There exists a matching μ in G_p .*
3. *No set of items $S \subseteq J \setminus \{j_0\}$ is strictly overdemanded at prices p .*

Alternating paths and trees. To identify strictly overdemanded items our mechanism makes use of alternating paths and trees: Let μ be a *partial matching*. That is, a matching in which not all of the bidders have to be matched. An *alternating path* \mathcal{P} with respect to μ in the first choice graph G_p at prices p from an unmatched bidder i_0 to some item or bidder j is a sequence of edges that alternates between unmatched and matched edges and in which all items except j are non-dummy items. An *alternating tree* \mathcal{T} with respect to μ with root i_0 is a tree in the first choice graph G_p at prices p which is rooted at an unmatched bidder i_0 and in which all paths from the root i_0 to a leaf j are alternating. An alternating tree is *maximal* if the first choice items of all bidders in the tree are contained in the tree and all matched items in the tree are matched to bidders in the tree. Formally: If $T \subseteq I$ and $S \subseteq J$ are the bidders and items in the tree \mathcal{T} , then $F_p(T) \subseteq S$ and $\mu(S) \subseteq T$. The fact that a partial matching can be augmented along an alternating path from an unmatched bidder to an unmatched item has been used before (see, e.g., [20]). The new insight of the following lemma is thus that there is a close correspondence between maximal alternating trees and our definition of strict overdemand.

Lemma 3. *For any maximal alternating tree \mathcal{T} with respect to μ with root i_0 in G_p , we have:*

1. *If the dummy item j_0 or some unmatched item $j \neq j_0$ is contained in \mathcal{T} , then the matching μ can be augmented along an alternating path \mathcal{P} from i_0 to j_0 resp. j .*

2. If all items S in \mathcal{T} are non-dummy items and matched, then S is strictly overdemanded with respect to the bidders T in the tree and $|T| = |S| + 1$.

Envy free price increase. Once we have identified a strictly overdemanded set of items we need to determine how to increase the prices in the set: A *price increase* d is a k -dimensional vector with entries d_j for $j \in \{1, \dots, k\}$. A price increase is *envy free* with respect to a set of first choice edges $E \subseteq F_p \setminus (I \times \{j_0\})$ at prices p if (a) $d_j > 0$ for all j such that there is a bidder i with $(i, j) \in E$ and $d_j = 0$ otherwise, and (b) $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ for all $(i, j) \in E$, all $(i, k) \in F_p$, and all sufficiently small $\lambda > 0$. Note that it is sufficient to require (b) for all $(i, k) \in F_p$ and not all $(i, k) \in I \times J$, because for every $(i, k) \in (I \times J) \setminus F_p$ we have $u_{i,j}(p_j) > u_{i,k}(p_k)$ and, thus, $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ holds already. Also note that because for no i we have $(i, j_0) \in E$ any envy free price increase has $d_{j_0} = 0$. Our definition of envy free price increase is similar to the definition of a *competitive direction* in [3]. The next two lemmata are proved in [3] for competitive directions and continuous utility functions, we generalize them to envy free price increases and discontinuous utility functions. The first lemma is an immediate consequence of the definition of envy free price increases. The second lemma gives a sufficient and necessary condition for a price increase $d \neq 0$ to be envy free for a set of first choice edges $E \subseteq F_p$. It shows that a price increase d is envy free for a first choice edge $(i, j) \in F_p$ if and only if the “utility drop” $c_{i,j} \cdot d_j$ on this edge is minimal across the first choice edges $(i, k) \in F_p$ incident to i . We exploit this characterization in the computation of price increases described in the next subsection.

Lemma 4. *If d is an envy free price increase with respect to the set of first choice edges $E \subseteq F_p \setminus (I \times \{j_0\})$ at prices p , then E belongs to the set of first choice edges at prices $p + \lambda \cdot d$ for all sufficiently small $\lambda > 0$.*

Lemma 5. *A price increase $d \neq 0$ is envy free for the set of first choice edges $E \subseteq F_p$ at prices p if and only if $c_{i,j} \cdot d_j \leq c_{i,k} \cdot d_k$ for all $(i, j) \in E \subseteq F_p$ and all $(i, k) \in F_p$.*

Overdemand preserving price increase. It is not difficult to see that envy free prices are not enough to guarantee minimum prices. To achieve this goal we define a stronger notion of price increases, which exploits the correspondence between maximal alternating trees and strict overdemand: An *overdemand preserving price increase* d for a maximal alternating tree \mathcal{T} with respect to μ with root i_0 in G_p with item set $S \subseteq J \setminus \{j_0\}$ and bidder set T in which all items are matched, is a price increase d such that (a) there is some partial matching μ' that matches the same bidders and items as μ and that is identical to μ on $I \setminus T \times J \setminus S$, (b) there is a maximal alternating tree \mathcal{T}' with respect to μ' with root i_0 that has the same item and bidder set as \mathcal{T} , and (c) d is envy free for the edges of the maximal alternating tree \mathcal{T}' . We say that μ' is the matching that *corresponds* to d . Note that μ' can be different from μ on $T \times S$. The crucial and new fact is that by (b) all items in the tree, i.e., all items whose price is increased, remain overdemanded for any small enough price increase.

Lemma 6. *If d is an overdemand preserving price increase for a maximal alternating tree \mathcal{T} with respect to μ with root i_0 in G_p with item set $S \subseteq J \setminus \{j_0\}$ and bidder set T in which all items are matched, then S is strictly overdemanded with respect to T in $G_{p+\lambda \cdot d}$ for all sufficiently small $\lambda > 0$.*

Next we present a subroutine that computes an overdemand preserving price increase d and a corresponding matching μ' for a maximal alternating tree \mathcal{T} with respect to μ with root i_0 in G_p with item set $S \subseteq J \setminus \{j_0\}$ and bidder set T in which all items are matched. The computation consists of three steps: (1) The subroutine computes a matching σ between $T \setminus \{i_0\}$ and S consisting of first choice edges, which minimizes $\prod_{(i,j) \in \mu} c_{i,j}$, or equivalently, $\sum_{(i,j) \in \mu} \log(c_{i,j})$. It also computes an envy free price increase d for σ . This can be accomplished by solving a linear program (LP) and its dual (DP), e.g., by using a primal-dual algorithm [36]. The duality between slopes and utility drops exploited here is reminiscent of the duality between value-maximizing matchings and envy free prices in [45]. (2) The subroutine extends d to an envy free direction for a maximal alternating tree \mathcal{T}' with respect to σ with root i_0 in G_p with bidder set T and item set S . (3) The subroutine extends σ to μ' by adding to it the bidder-item pairs from $I \setminus T \times J \setminus S$ that were matched in μ . While (1) is essentially an application of Lemma 5 (and has been used in a similar form in [3, 4]), (2) and (3) exploit the newly established correspondence between maximal alternating trees and strict overdemand.

Subroutine for price increases

Input: maximal alternating tree \mathcal{T} with respect to μ with root i_0 in G_p with item set S and bidder set T in which all items are matched

Output: overdemand preserving price increase d for \mathcal{T} with corresponding matching μ'

1 compute x as optimal solution to the following LP and

$$\text{let } \sigma = \{(i, j) \in T \setminus \{i_0\} \times S \mid x_{i,j} = 1\}$$

$$\begin{aligned} \min \quad & \sum_{i,j} x_{i,j} \cdot \log(c_{i,j}) \\ \text{sb} \quad & \sum_{j \in F_p(i)} x_{i,j} = 1 \quad (\forall i \in T \setminus \{i_0\}) \\ & \sum_{i \in F_p(j)} x_{i,j} = 1 \quad (\forall j \in S) \\ & x_{i,j} \geq 0 \quad (\forall (i, j) \in F_p \cap (T \setminus \{i_0\} \times S)) \end{aligned}$$

2 compute ω, ρ as optimal solution to the following DP

$$\begin{aligned} \max \quad & \sum_i \omega_i + \sum_j \rho_j \\ \text{sb} \quad & \omega_i + \rho_j \leq \log(c_{i,j}) \quad (\forall (i, j) \in F_p \cap (T \setminus \{i_0\} \times S)) \end{aligned}$$

3 extend ω from $T \setminus \{i_0\}$ to T by setting

$$\omega_{i_0} = \min_{j \in S} \log(c_{i_0,j}) - \rho_j$$

4 let $H_\rho = (T \cup S, E_\rho)$, where

$$E_\rho = \{(i, j) \in F_p \cap (T \times S) \mid \omega_i + \rho_j = \log(c_{i,j})\}$$

5 let \mathcal{T}' be a maximal alternating tree in H_ρ with
respect to σ with root i_0
6 let $S' \subseteq S$ and $T' \subseteq T$ denote the items and
bidders in \mathcal{T}'
7 **while** $T' \neq T$ or $S' \neq S$ **do**
8 let $\delta = \min_{(i,j) \in F_p: i \in T', j \in S \setminus S'} \log(c_{i,j}) - \omega_i - \rho_j$
9 set $\rho_j = \rho_j + \delta$ for all $j \in S \setminus S'$,
set $\omega_i = \omega_i - \delta$ for all $i \in T \setminus T'$
10 recompute \mathcal{T}' , T' , and S'
11 **end while**
12 set $d_j = e^{-\rho_j}$ for all $j \in S$ and $d_j = 0$ otherwise
13 set $\mu' = \sigma \cup (\mu \cap (I \setminus T \times J \setminus S))$
14 output d and μ'

Lemma 7. *This subroutine finds an overdemand preserving price increase and a corresponding matching. It can be implemented to run in time $O(\min(n, k)^3)$.*

The following lemma – our key lemma and main technical improvement over [3, 4] – shows that if overdemand preserving price increases are used, then the resulting prices will be minimum over all envy free outcomes.

Lemma 8. *Let d be an overdemand preserving price increase for a maximal alternating tree \mathcal{T} in G_p with item set S and bidder set T in which all items are matched. Let $\lambda > 0$ be the smallest scalar such that at $p + \lambda \cdot d$ (a) a bidder-item pair $(i, j) \in T \times J \setminus S$ enters $G_{p+\lambda \cdot d}$ or (b) the end point $e_{i,j}^{(t)}$ of some interval $t > 0$ is reached. Then for any envy free outcome (μ'', p'') with $p'' \geq p$ we have $p'' \geq p + \lambda \cdot d$.*

Bidder optimal outcome. Our mechanism starts with an empty matching $\mu = \emptyset$ and prices $p = 0$. It then matches one bidder after the other until eventually all bidders are matched. For this it computes a maximal alternating tree \mathcal{T} with respect to μ with root i_0 , where i_0 is the bidder to be matched, in the first choice graph G_p . If the alternating tree contains the dummy item j_0 or an unmatched item j , then by Lemma 3 the current matching μ can be augmented along an alternating path from i_0 to j_0 resp. j . If this is not the case, then – again by Lemma 3 – the items S in the tree are strictly overdemanded with respect to the bidders T in the tree. In this case the mechanism computes an overdemand preserving price increase d together with a corresponding matching μ' (using the subroutine from the previous subsection) and raises the prices in compliance with d until (a) a bidder-item pair $(i, j) \in T \times J \setminus S$ enters the first choice graph $G_{p+\lambda \cdot d}$ or (b) the end point $e_{i,j}^{(t)}$ of some interval $t > 0$ is reached. In either case the current matching μ is replaced with μ' and the minimality of the new prices is guaranteed by Lemma 8. If at least one of the new prices $p_j + \lambda \cdot d_j$ corresponds to a discontinuity, then one or multiple edges might drop out of the first choice graph. The mechanism corrects for this by removing such edges from the

matching if necessary. If no discontinuity is reached, then the maximal alternating tree \mathcal{T} rooted at i_0 grows by at least one item.

Mechanism for piece-wise linear utility functions

Input: bidders I , items J , piece-wise linear utility functions $u_{i,j}(\cdot)$ with non-identical slopes and multiple discontinuities, reserve prices $r_j = 0$, outside options o_i

Output: bidder optimal outcome (μ, p)

- 1 set $p_j = 0$ for all j and set $\mu = \emptyset$
- 2 **while** there exists an unmatched bidder i_0 **do**
- 3 compute maximal alternating tree \mathcal{T} wrt μ in the first choice graph G_p with root i_0
- 4 let T and S be the bidders and items in \mathcal{T}
- 5 **while** all items in S are matched and S does not contain the dummy item j_0 **do**
- 6 compute overdemand-preserving price increase d for \mathcal{T} and corresponding matching μ' (using the sub-routine from the previous subsection)
- 7 let $\lambda > 0$ be the smallest scalar such that at prices $p + \lambda \cdot d$
 - (a) a bidder-item pair $(i, j) \in T \times J \setminus S$ enters the first choice graph $G_{p+\lambda \cdot d}$ or
 - (b) the end point $e_{i,j}^{(t)}$ of some interval $t > 0$ is reached
- 8 set $p_j = p_j + \lambda \cdot d_j$ for all $j \in J$ and set $\mu = \mu'$
- 9 remove bidder-item pairs from μ that do not belong to the first choice graph G_p
- 10 compute maximal alternating tree \mathcal{T} wrt μ in the first choice graph G_p with root i
- 11 let T and S be the bidders and items in \mathcal{T}
- 11 **end while**
- 12 augment μ along alternating path P from i_0 to unmatched item j or dummy item j_0
- 13 **end while**
- 14 output (μ, p)

Theorem 1. *This mechanism finds a bidder optimal outcome. It can be implemented to run in time $O((n \cdot \min(n, k) + D \cdot \min(n, k) + T) \cdot \min(n, k) \cdot (\min(n, k)^2 + k))$, where $D = \sum_{i,j} d_{i,j}$ and $T = \sum_{i,j} t_{i,j}$ denote the total number of discontinuities and constant-slope intervals.*

4 Incentive Compatibility

In this section we precisely characterize under which conditions *any* mechanism that computes a bidder optimal outcome is incentive compatible. Intuitively, a mechanism is incentive compatible if for every bidder i , independently of all other bidders, reporting his true utility functions yields an outcome, which gives him the highest possible utility. This can be formalized as follows: A mechanism is incentive compatible if for every bidder i with utility functions $u_{i,j}(\cdot)$ and every two sets of utility functions $u'_{i,j}(\cdot)$ and $u''_{i,j}(\cdot)$, where $u'_{i,j}(\cdot) = u_{i,j}(\cdot)$ for i and all j and $u'_{k,j}(\cdot) = u''_{k,j}(\cdot)$ for all $k \neq i$ and all j , and corresponding outcomes (μ', p') and (μ'', p'') of the mechanism we have

$$u_{i,\mu'(i)}(p'_{\mu'(i)}) \geq u_{i,\mu''(i)}(p''_{\mu''(i)}). \quad (10)$$

Note that this definition does not involve the reserve prices r_j or outside options o_i . This makes sense because the reserve prices r_j are typically set by the seller and misreporting o_i is never beneficial to i .⁷

Example: Lying pays off. (This example shows that bidder optimality does not imply incentive compatibility.) There are two bidders $i \in \{1, 2\}$ and two items $j \in \{1, 2\}$. The utility functions for $i \in \{1, 2\}$ are:

$$u_{i,1}(p_1) = \begin{cases} 20 - p_1 & \text{for } p_1 \in [0, 5), \text{ and} \\ -\infty & \text{otherwise,} \end{cases}$$

$$u_{i,2}(p_2) = 1 - p_2 \quad \text{for } p_2 \in [0, \infty).$$

The reserve prices are $r_j = 0$ for $j \in \{1, 2\}$ and the outside options are $o_i = 0$ for $i \in \{1, 2\}$. A bidder optimal outcome for this input is $\mu = \{(1, 2)\}$ with $p_1 = 5$ and $p_2 = 1$. For this outcome both bidders have a utility of zero. Bidder 1 can improve his utility by pretending to have $u_{1,1} = 0 - p_1$ for $p_1 \in [0, \infty)$. In this case the bidder optimal outcome is $\mu = \{(1, 2), (2, 1)\}$ with $p_1 = p_2 = 0$. The utility of bidder 1 improves from 0 to 1. The crucial point – as we will show below – is that in the computation of the bidder optimal outcome on this example two first choice edges, namely $(1, 1)$ and $(2, 1)$, simultaneously break away from the first choice graph at price $p_1 = 5$.

Price-independent formulation. We will define next a condition on the input that implies that *never* during the execution of the mechanism two edges will break away from the first choice graph during the same price increase. Which edges break away depends on the current prices and the price increases. However, using the following idea we can write down a condition that does *not* depend on the current prices: Suppose that the edges (i, j) , (i', j) , and (i', j') belong to the first choice graph G_p at prices p . It follows that

$$v_{i',j} - c_{i',j} \cdot p_j = v_{i',j'} - c_{i',j'} \cdot p_{j'}. \quad (11)$$

⁷Over-reporting can only lead to a missed chance of being assigned an item and under-reporting can only lead to a utility below the true outside option.

Suppose further that d is an envy free price increase for the set of first choice edges $E = \{(i, j), (i', j), (i', j')\}$, then

$$v_{i',j} - c_{i',j} \cdot (p_j + \lambda d_j) = v_{i',j'} - c_{i',j'} \cdot (p_{j'} + \lambda d_{j'}). \quad (12)$$

By subtracting (11) from (12), dividing by $\lambda > 0$, and after rearranging we get

$$d_j = c_{i',j'}/c_{i,j} \cdot d_{j'}. \quad (13)$$

Now suppose that the discontinuities $D_{i,j}$ and $D_{i',j'}$ are reached simultaneously. Then by (12):

$$v_{i',j} - c_{i',j} \cdot D_{i,j} = v_{i',j'} - c_{i',j'} \cdot D_{i',j'}. \quad (14)$$

Using (13), subtracting $1/d_j \cdot v_{i,j}/c_{i,j}$ from both sides, and after rearranging we get

$$\frac{1}{d_j} \left(D_{i,j} - \frac{v_{i,j}}{c_{i,j}} \right) = -\frac{1}{d_j} \frac{v_{i,j}}{c_{i,j}} + \frac{1}{d_j} \frac{v_{i',j}}{c_{i',j}} + \frac{1}{d_{j'}} \left(D_{i',j'} - \frac{v_{i',j'}}{c_{i',j'}} \right).$$

Below we will define a multigraph such that the left and right hand side of this equation correspond to the weights of two alternating walks in the graph, namely $P = (i, j)$ and $Q = (i, j, i', j')$. Note that neither the weight of P nor the weight of Q depend on the prices.

General position. For a given input we define a multigraph, called *input graph*, as follows: There is one node per bidder $i \in I$ and one node per item $j \in J$. There are three types of edges: (a) there is one forward edge from i to j for each linear segment of $u_{i,j}(\cdot)$, (b) a backward edge from j to i for each linear segment of $u_{i,j}(\cdot)$, and (c) a discontinuity edge from i to j for each discontinuity $D_{i,j}$ of $u_{i,j}(\cdot)$.

Let $P = (i_0, j_1, \dots, i_s, j_s)$ be a walk in the input graph that alternates between forward and backward edges, and ends with a discontinuity edge. Let d be a price increase such that $d_j = (c_{i,j'}/c_{i,j}) \cdot d_{j'}$ for any two edges (i, j) and (i, j') on P . Define the weight of each forward edge (i, j) on P with respect to d as $(-1/d_j) \cdot (v_{i,j}/c_{i,j})$, of each backward edge (j, i) as $(1/d_j) \cdot (v_{i,j}/c_{i,j})$, and of the discontinuity edge (i, j) as $(1/d_j) \cdot (D_{i,j} - v_{i,j}/c_{i,j})$. Here $v_{i,j}$ and $c_{i,j}$ are the constants of the corresponding linear segments. Define the weight $w_d(P)$ of P with respect to d as the sum of these weights. We say that the input is in *general position* if for no two walks P and Q that start with the same bidder and end with a distinct discontinuity edge and for no price increase d such that $d_j = (c_{i,j'}/c_{i,j}) \cdot d_{j'}$ for any two edges (i, j) and (i, j') on P resp. Q we have $w_d(P) = w_d(Q)$. Note that this definition of general position is more general than that in [1]. In particular, it takes into account that the utility functions have *non-identical slopes* and *multiple discontinuities*.

Lemma 9. *We have:*

1. *An input is in general position if and only if the associated input in standard form is in general position.*
2. *Let (μ, p) denote the outcome of the mechanism in Section 3. If the input is in general position, then*
 - (i) *no two discontinuities are reached simultaneously,*
 - (ii) *if an item gets unmatched, it gets matched again in the subsequent iteration,*
 - (iii) *if $p_j > r_j$, then item j is matched under μ , and*
 - (iv) *the last item, say j , that gets matched has $p_j = r_j$.*

Characterization. We already know that bidder optimality does not imply incentive compatibility, if the input is *not* in general position. With the help of the following lemma – which is a generalization of Hwang’s Lemma (see, e.g., [19, 21]) – we can show that any mechanism that computes a bidder optimal outcome is incentive compatible if conditions (iii) and (iv) from Lemma 9 are satisfied. An easy corollary is that if the input is in general position, then any mechanism that computes a bidder optimal outcome is incentive compatible. Note that while we do not have a polynomial-time algorithm to check whether an input is in general position, we can easily check whether conditions (iii) and (iv) are satisfied using our mechanism from Section 3.

Lemma 10. *If conditions (iii) and (iv) from Lemma 9 are satisfied, then:*

1. *If the outcome (μ^*, p^*) is bidder optimal, then for no feasible outcome (μ', p') we can have $u_{i, \mu'(i)}(p'_{\mu'(i)}) > u_{i, \mu^*(i)}(p^*_{\mu^*(i)})$ for all i .*
2. *If the outcome (μ^*, p^*) is bidder optimal, the outcome (μ', p') is feasible, and $I^+ = \{i \in I \mid u_{i, \mu'(i)}(p'_{\mu'(i)}) > u_{i, \mu^*(i)}(p^*_{\mu^*(i)})\} \neq \emptyset$, then there exists a bidder-item pair $(i, j) \in I \setminus I^+ \times J$ s.t. $u_{i, \mu'(i)}(p_{\mu'(i)}) < u_{i, j}(p'_j)$.*

Theorem 2. *If conditions (iii) and (iv) from Lemma 9 are satisfied, then any mechanism that computes a bidder optimal is incentive compatible.*

We conclude our discussion of the incentives involved in computing bidder optimal outcomes with several interesting open questions concerning the general position concept: First, what does it take (time-wise) to evaluate whether an input is in general position? Second, what does it take to change an input that is not in general position to be in general position? Finally, under which conditions is a bidder optimal solution to an input that has been brought into general position also bidder optimal for the original input?

5 Approximation

In this section we show how our mechanism for piece-wise linear utility functions can be applied to compute approximately bidder optimal outcomes for a general class of continuous utility functions $u_{i,j}(\cdot)$. The idea is as follows: Approximate each utility function $u_{i,j}(\cdot)$ by a piece-wise linear utility function $\tilde{u}_{i,j}(\cdot)$. Then solve the problem for this approximated input exactly and use the outcome $(\tilde{\mu}, \tilde{p})$ obtained as an approximate solution to the original problem.

More specifically, we say that an outcome $(\tilde{\mu}, \tilde{p})$ is γ -*envy free* for the input $u_{i,j}(\cdot)$ if it is feasible and $u_{i,\tilde{\mu}(i)}(\tilde{p}_{\tilde{\mu}(i)}) + \gamma \geq u_{i,j}(\tilde{p}_j)$ for all $(i, j) \in I \times J$. An outcome $(\tilde{\mu}, \tilde{p})$ is γ -*bidder optimal* if it is γ -envy free and for any envy free outcome $(\tilde{\mu}', \tilde{p}')$ we have $u_{i,\tilde{\mu}(i)}(\tilde{p}_{\tilde{\mu}(i)}) + \gamma \geq u_{i,\tilde{\mu}'(i)}(\tilde{p}'_{\tilde{\mu}'(i)})$.

Apart from the assumptions that the utility functions $u_{i,j}(\cdot)$ are strictly monotonically decreasing and continuous we make the following mild assumptions concerning the first and second derivatives $\dot{u}_{i,j}(\cdot)$ and $\ddot{u}_{i,j}(\cdot)$ of $u_{i,j}(\cdot)$, which allow us to bound the number of linear segments needed for piece-wise linear approximation:

- (A.1) The utility functions $u_{i,j}(\cdot)$ are twice differentiable on $[r_j, \bar{p}_{i,j}]$, i.e., $\ddot{u}_{i,j}(p_j)$ exists on this interval.
- (A.2) There exists a constant B such that for $\forall (i, j) \in I \times J$:

$$\max_{p_j \in [r_j, \bar{p}_{i,j}]} |\ddot{u}_{i,j}(p_j)| \leq B.$$
- (A.3) There exist constants m and M such that $\forall (i, j) \in I \times J$ and $\forall p_j \in [r_j, \bar{p}_{i,j}]$: $0 < m \leq |\dot{u}_{i,j}(p_j)| \leq M$.

Note that we use $\dot{u}_{i,j}(\cdot)$ and $\ddot{u}_{i,j}(\cdot)$ to denote the first and second derivative of $u_{i,j}(\cdot)$ to avoid confusion with $u'_{i,j}(\cdot)$ and $u''_{i,j}(\cdot)$, which were previously used in a different context.

Piece-wise linear approximation. Given $\epsilon > 0$ we construct a piece-wise linear, continuous approximation $\tilde{u}_{i,j}(\cdot)$ for $u_{i,j}(\cdot)$ with error at most ϵ , i.e., $|u_{i,j}(p_j) - \tilde{u}_{i,j}(p_j)| \leq \epsilon$ for all $p_j \in [r_j, \bar{p}_{i,j}]$, as follows: First, as $u_{i,j}(\cdot)$ can extend from r_j to ∞ and could potentially require an infinite number of segments to approximate, we limit the approximation to the range $[r_j, \bar{p}_{i,j}]$ as follows: Since $p_j > \bar{p}_{i,j}$ cannot correspond to a match as then $u_{i,j}(p_j) < o_i = u_{i,j_0}(p_{j_0})$, we can extend $u_{i,j}(p_j)$ for prices $p_j > \bar{p}_{i,j}$ in a continuous and differentiable way by the line $\tilde{u}_{i,j}(p_j) = \dot{u}_{i,j}(\bar{p}_{i,j}) \cdot p_j - \dot{u}_{i,j}(\bar{p}_{i,j}) \cdot \bar{p}_{i,j}$. This limits the “interesting” domain to $[r_j, \bar{p}_{i,j}]$. Next we split the range $[r_j, \bar{p}_{i,j}]$ into $S = \lceil (\bar{p}_{i,j} - r_j) / \sqrt{8\epsilon/B} \rceil$ intervals of equal width $w = (\bar{p}_{i,j} - r_j) / S$. On any interval with endpoints $[\tilde{e}^{(t)}, \tilde{e}^{(t+1)}]$, where $\tilde{e}^{(t+1)} = \tilde{e}^{(t)} + w$, the line $\tilde{u}_{i,j}(\cdot)$ used to approximate $u_{i,j}(\cdot)$ is defined by

$$\tilde{u}_{i,j}(p_j) = \frac{u_{i,j}(\tilde{e}^{(t+1)}) - u_{i,j}(\tilde{e}^{(t)})}{\tilde{e}^{(t+1)} - \tilde{e}^{(t)}} p_j + \frac{u_{i,j}(\tilde{e}^{(t)}) \tilde{e}^{(t+1)} - u_{i,j}(\tilde{e}^{(t+1)}) \tilde{e}^{(t)}}{\tilde{e}^{(t+1)} - \tilde{e}^{(t)}}.$$

We call this kind of approximation *point-to-point approximation* as the piece-wise linear approximation agrees with the original function at the end points of each interval. The following lemma shows that the above algorithm does indeed give a close approximation.

Lemma 11. *For every bidder-item pair $(i, j) \in I \times J$ the algorithm described above gives a point-to-point approximation using $O(|\bar{p}_{i,j} - r_j| \cdot \sqrt{B} \cdot 1/\sqrt{\epsilon}) = O(\sqrt{1/\epsilon})$ segments that (i) is piece-wise linear, (ii) continuous, and (iii) has error at most ϵ .*

Approximately bidder optimal outcome. Given a set of continuous, piece-wise linear, point-to-point approximations $\tilde{u}_{i,j}(\cdot)$ with error at most ϵ , we can use the mechanism from Section 3 to obtain a bidder optimal outcome $(\tilde{\mu}, \tilde{p})$ for the approximated input. The following lemma bounds how “far” away from the bidder optimal outcome (μ, p) for the original utility functions $u_{i,j}(\cdot)$ this outcome can be depending on how “close” the piece-wise linear approximations are, i.e., depending on how small ϵ is. The idea is as follows: Suppose we knew the bidder optimal prices p for $u_{i,j}(\cdot)$. We could use them to “jumpstart” the mechanism from Section 3. That is, starting from p we could find the smallest prices \tilde{p}' such that $(\tilde{\mu}', \tilde{p}')$ is envy free for $\tilde{u}_{i,j}(\cdot)$. From Lemma 1 we know that for the bidder optimal prices \tilde{p} we have that $\tilde{p} \leq \tilde{p}'$ so that any upper bound on \tilde{p}' also applies to \tilde{p} . To bound the difference between p and \tilde{p}' we first bound the ratio between the biggest and the smallest non-zero entry of an overdemand preserving price increase by $O((M/m)^{\min(n,k)})$. We then argue that between any two consecutive executions of Step 7(a) in the mechanism from Section 3 this difference is increased by a multiplicative factor of $O((M/m)^{\min(n,k)})$. The crucial point here is that the number of executions of Step 7(a) is $O(\min(n, k)^2)$ and thus independent of ϵ , while the number of executions of Step 7(b) depends on the number of linear segments and thus on ϵ . We use the resulting bound on the difference between the prices to bound the difference between the utilities. Theorem 3 follows from this bound for ϵ small enough.

Lemma 12. *For every envy free outcome (μ, p) for $u_{i,j}(\cdot)$ and continuous, piece-wise linear, point-to-point approximation $\tilde{u}_{i,j}(\cdot)$ with error at most ϵ the mechanism from Section 3 finds an envy free outcome $(\tilde{\mu}, \tilde{p})$ for $\tilde{u}_{i,j}(\cdot)$ with $\tilde{p}_j \leq p_j + (2M/m)^{(\min(n,k)+1)^3} \cdot \epsilon/M$ for all j and $\tilde{u}_{i,\tilde{\mu}(i)}(\tilde{p}_{\tilde{\mu}(i)}) + (2M/m)^{(\min(n,k)+1)^3} \cdot \epsilon \geq u_{i,\mu(i)}(p_{\mu(i)})$ for all i .*

Theorem 3. *Given strictly monotonically decreasing, continuous utility functions $u_{i,j}(\cdot)$ satisfying assumptions (A.1) to (A.3), we can compute a γ -bidder optimal outcome with the running time specified in Theorem 1, where $D = 0$, $T = O(\sqrt{1/\epsilon})$, and $\epsilon = \gamma \cdot (2M/m)^{-(\min(n,k)+1)^3}$.*

6 Externalities

In this section we show that two natural generalizations of our problem are \mathcal{NP} -hard: In the first, the utility functions are allowed to depend on the matching. In the second, they are allowed to depend on the vector of prices.

Utilities that depend on the matching. Our first hardness result is for utility functions that depend on the matching. These utility functions allow, for example, to express that a given advertiser (e.g. Coca-Cola) achieves a higher utility if he gets an ad slot that is above the slot of her competitor (e.g. Pepsi). For utility functions of this kind it is \mathcal{NP} -complete to decide whether there exists an outcome in which the sum of the utilities is above a certain threshold.

Theorem 4. *Given utility functions $u_{i,j}(\mu)$ that depend on the matching μ , and given a constant K , deciding whether there exists an envy free outcome with sum of the utilities at least K is \mathcal{NP} -complete.*

Utilities that depend on the vector of prices. Our second hardness result concerns situations in which the utility that a bidder derives from being matched to an item depends on the *whole* vector of prices. These utility functions allow, for example, to express that a bidder i achieves a higher utility when the other bidders have to pay more, i.e., when items that i does not get become more expensive. For these utility functions deciding whether there exists an envy free outcome in which the sum of the utilities is above a certain threshold is \mathcal{NP} -complete.

Theorem 5. *Given utility functions $u_{i,j}(p)$ that depend on the whole vector of prices $p = (p_1, \dots, p_k)$, that are monotonically decreasing in p_j and monotonically increasing in p_k for all $k \neq j$, and given a constant K , it is \mathcal{NP} -complete to decide whether there exists an envy free outcome (μ, p) in which the sum of the utilities is at least L .*

7 Future Work

The demand for more expressive mechanisms is reflected in the richness of preferences offered by web applications as diverse as matchmaking sites, sites like Amazon and NetFlix, and services like Google’s AdSense. Standard mechanisms often do not meet this demand. Providing mechanisms that do meet this demand and that at the same time (1) guarantee the existence of a stable solution, (2) are computationally tractable, and (3) have good incentive properties is one of the major challenges that the field of computational mechanism design is currently faced with.

In this paper we contributed to this general agenda by considering the domain of *multi-item auctions* with *unit demand* and by providing the most expressive mechanism for this setting so far. This mechanism, which can be seen as a generalization of the General Auction Mechanism of [1], can handle *piece-wise linear* utility functions with *non-identical slopes* and *multiple discontinuities*. These utility functions allow the bidders to explicitly specify *conversion rates* (enabling, e.g., per-click auctions that are simultaneously envy free for bidders with per-click *and* per-impression valuations) and a variety of *soft* and *hard budget constraints* (which, e.g., arise when bidders have a limited amount of cash and have to take out loans). An interesting direction for future work would be to push the

“expressiveness frontier” even further. This is particularly true for more general domains, e.g., one-to-many and many-to-many domains.

On a more abstract level it would be desirable to have a “theory of expressiveness”, which helps to find the right degree of expressiveness. More expressiveness might be good (e.g., because more efficient outcomes are obtainable), but it might also be too much (e.g., stable outcomes may no longer exist, may be hard to compute, or may be easy to manipulate). A first step towards such a general theory, although under very different premises, was recently undertaken by [9, 10], but, especially in the light of [41], any refinement of this theory would be highly interesting.

Acknowledgement

This project has been funded by the Vienna Science and Technology Fund WWTF grant ICT10-002 and a European Young Investigators Award (<http://www.esf.org/euryi/>).

References

- [1] G. Aggarwal, S. Muthukrishnan, D. Pál, and M. Pál. General auction mechanism for search advertising. *WWW*, pages 241–250, 2009.
- [2] S. Alaei, K. Jain, and A. Malekian. Walrasian equilibrium for unit demand buyers with non-quasi-linear utilities. *Working paper*, pages 1–19, 2010.
- [3] A. Alkan. Existence and computation of matching equilibria. *Political Economy*, 5:285–296, 1989.
- [4] M. Alkan. Equilibrium in a matching market with general preferences. *Equilibrium and Dynamics. Essays in Honour of David Gale*, pages 1–16, 1992.
- [5] I. Ashlagi, M. Braverman, and A. Hassidim. Ascending unit demand auctions with budget limits. *Working paper*, pages 1–11, 2009.
- [6] I. Ashlagi, M. Braverman, A. Hassidim, R. Lavi, and M. Tennenholtz. Position auctions with budgets: Existence and uniqueness. *The B.E. Journal of Theoretical Economics*, 10 (1):(Advances), 2010.
- [7] I. Ashlagi, P. Krysta, and M. Tennenholtz. Social context games. *WINE*, pages 675–683, 2008.
- [8] P. Bajari and A. Hortacsu. The winner’s curse, reserve prices, and endogenous entry: Empirical insights from ebay auctions. *The RAND Journal of Economics*, Vol. 34(2):329–355, 2003.

- [9] M. Benisch, N. Sadeh, and T. Sandholm. A theory of expressiveness in mechanisms. *AAAI*, pages 17–23, 2008.
- [10] M. Benisch, N. Sadeh, and T. Sandholm. Methodology for designing reasonably expressive mechanisms with application to ad auctions. *IJCAI*, pages 46–52, 2009.
- [11] S. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget constrained bidders. *EC*, pages 44–51, 2005.
- [12] N. Chen, X. Deng, and A. Ghosh. Competitive equilibria in matching markets with budgets. *Working paper*, pages 1–34, 2010.
- [13] X. Chen and X. Deng. Settling the complexity of two-player nash equilibrium. *FOCS*, pages 261–272, 2006.
- [14] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
- [15] V. Conitzer and T. Sandholm. Complexity results about nash equilibria. *IJCAI*, pages 765–771, 2003.
- [16] S. A. Cook. The complexity of theorem-proving procedures. *STOC*, pages 151–158, 1971.
- [17] C. Daskalakis and C. H. Papadimitriou. Three-player games are hard. *ECCC*, pages 1–10, 2005.
- [18] G. Demange and D. Gale. The strategy structure of two-sided matching markets. *Econometrica*, 53(4):873–888, 1985.
- [19] G. Demange, D. Gale, and M. Sotomayor. Some remarks on the stable marriage problem. *Discrete Applied Mathematics*, 11:223–232, 1985.
- [20] G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *Political Economy*, 94(4):863–72, 1986.
- [21] G. Demange, D. Gale, and M. Sotomayor. A further note on the stable matching problem. *Discrete Applied Mathematics*, 16(3):217–222, 1987.
- [22] X. Deng and Y. Du. The computation of approximate competitive equilibrium is ppad-hard. *Information Processing Letters*, 108(6):369–373, 2008.
- [23] S. Dobzinski, R. Lavi, and N. Nisan. Multi-unit auctions with budget limits. *FOCS*, pages 260–269, 2008.
- [24] P. Dütting, M. Henzinger, and I. Weber. Bidder optimal assignments for general utilities. *WINE*, pages 575–582, 2009.

- [25] P. Dütting, M. Henzinger, and I. Weber. Sponsored search, market equilibria, and the hungarian method. *STACS*, pages 287–298, 2010.
- [26] 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.
- [27] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM*, 19(2):248–264, 1972.
- [28] J. Feldman, M. Henzinger, N. Korula, V. Mirrokni, and C. Stein. Online stochastic packing applied to display ad allocation. *ESA*, pages 1–13, 2010.
- [29] A. Ghosh and A. Sayedi. Expressive auctions for externalities in online advertising. *WWW*, pages 371–380, 2010.
- [30] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- [31] A. Goel and K. Munagala. Hybrid keyword search auctions. *WWW*, pages 221–230, 2009.
- [32] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [33] P. Hall. On representatives of subsets. *London Mathematical Society*, 10:26–30, 1935.
- [34] P. Jehiel, B. Moldovanu, and E. Stacchetti. How (not) to sell nuclear weapons. *American Economic Review*, 86(4):814–29, 1996.
- [35] D. Kempe, A. Mu’alem, and M. Salek. Envy-free allocations for budgeted bidders. *WINE*, pages 537–544, 2009.
- [36] H. W. Kuhn. The Hungarian Method for the assignment problem. *Naval Research Logistics*, 2:83–97, 1955.
- [37] S. Lahaie, D. M. Pennock, A. Saberi, and R. V. Vohra. Sponsored search auctions. *Algorithmic Game Theory*, pages 699–716, 2007.
- [38] H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Political Economy*, 91:461–479, 1983.
- [39] C. Limongelli and R. Pirastu. Exact solution of linear systems over rational numbers by parallel p-adic arithmetic. *CONPAR*, pages 313–323, 1994.
- [40] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized online matching. *Journal of the ACM*, 54(5):1–19, 2007.

- [41] P. Milgrom. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior*, 70:62–70, 2010.
- [42] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [43] A. E. Roth and A. Ockenfels. Last-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon auctions on the internet. *The American Economic Review*, 92(4):1093–1103, 2002.
- [44] A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, 1998.
- [45] L. S. Shapley and M. Shubik. The assignment game: The core I. *Game Theory*, 29:111–130, 1971.
- [46] N. Tomizawa. On some techniques useful for solution of transportation network problems. *Game Theory*, 1:173–194, 1971.
- [47] H. Varian. Position auctions. *Industrial Organization*, 25(6):1163–1178, 2007.
- [48] V. Vazirani and M. Yannakakis. Ppad-completeness of fisher markets under piecewise-linear, concave utility functions. *ICTS*, pages 156–166, 2010.
- [49] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Finance*, 16:8–27, 1961.

A Proof of Lemma 1

We prove the claim in two steps:

To 1.: The outcome (μ, p') is feasible for $u_{i,j}(p_j)$ and r_j because $p'_{j_0} = p_{j_0} + r_{j_0} = p_{j_0} = 0$ and $p'_j = p_j + r_j \geq r_j$ for all $j \neq j_0$. It is envy free because $u_{i,\mu(i)}(p'_{\mu(i)}) = u_{i,\mu(i)}(p_{\mu(i)} + r_{\mu(i)}) = u'_{i,\mu(i)}(p_{\mu(i)}) \geq u'_{i,j}(p_j) = u_{i,j}(p_j + r_j) = u_{i,j}(p'_j)$ for all $i \in I$ and all $j \in J$. To see that it is also bidder optimal, suppose for a contradiction that it is not, i.e., there exists an outcome (μ'', p'') that is envy free for $u_{i,j}(p_j)$ and r_j and has $u_{i,\mu''(i)}(p''_{\mu''(i)}) \geq u_{i,\mu(i)}(p'_{\mu(i)})$ for all $i \in I$ and all $j \in J$; with at least one inequality strict. But then the outcome (μ'', p''') with $p'''_j = p''_j - r_j$ for all j is (a) feasible for $u'_{i,j}(p_j)$ and $r'_j = 0$ because $p'''_{j_0} = p''_{j_0} - r_{j_0} = 0$ and $p'''_j = p''_j - r_j \geq 0$ and (b) envy free for $u'_{i,j}(p_j)$ and $r'_j = 0$ because $u'_{i,\mu''(i)}(p'''_{\mu''(i)}) = u'_{i,\mu''(i)}(p''_{\mu''(i)} - r_{\mu''(i)}) = u_{i,\mu''(i)}(p''_{\mu''(i)}) \geq u_{i,j}(p''_j) = u'_{i,j}(p''_j - r_j) = u'_{i,j}(p'''_j)$ for all $i \in I$ and all $j \in J$. Hence $u'_{i,\mu''(i)}(p'''_{\mu''(i)}) = u_{i,\mu''(i)}(p''_{\mu''(i)}) \geq u_{i,\mu(i)}(p'_{\mu(i)}) = u'_{i,\mu(i)}(p_{\mu(i)})$ for all $i \in I$ and all $j \in J$; with at least one of the inequalities strict. This contradicts the bidder optimality of (μ, p) for $u'_{i,j}(\cdot)$ and $r'_j = 0$.

To 2.: For a contradiction suppose that there is an envy free outcome (μ', p') such that $u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu^*(i)}(p^*_{\mu^*(i)})$ for some bidder i . Since (μ^*, p^*) is envy free, we have that

$u_{i,\mu^*(i)}(p_{\mu^*(i)}^*) \geq u_{i,\mu'(i)}(p_{\mu'(i)}^*)$. It follows that $u_{i,\mu'(i)}(p_{\mu'(i)}') > u_{i,\mu'(i)}(p_{\mu'(i)}^*)$, which implies $p_{\mu'(i)}' < p_{\mu'(i)}^*$. Contradiction!

B Proof of Lemma 2

We prove the claim in two steps:

1. \Leftrightarrow 2.: An outcome (μ, p) is envy free iff (a) $p_{j_0} = 0$ and $p_j \geq 0$ for $j \neq j_0$ and (b) $u_{i,\mu(i)}(p_{\mu(i)}) \geq u_{i,j}(p_j)$ for all $(i, j) \in I \times J$ iff all edges $(i, j) \in \mu$ belong to the first choice graph G_p at prices p with $p_{j_0} = 0$ and $p_j \geq 0$ for $j \neq j_0$.

2. \Leftrightarrow 3.: By Hall's Theorem [33] there exists a matching μ in the first choice graph at prices p with $p_{j_0} = 0$ and $p_j \geq 0$ for $j \neq j_0$ iff $\forall T \subseteq I: |F_p(T)| \geq |T|$ or $j_0 \in F_p(T)$.

\Rightarrow : Assume that $\forall T \subseteq I: |F_p(T)| \geq |T|$ or $j_0 \in F_p(T)$. For a contradiction suppose that there exists a set of items $S' \subseteq J \setminus \{j_0\}$ that is strictly overdemanded with respect to the set of bidders T' , i.e., (a) $F_p(T') \subseteq S'$ and (b) $\forall R \subseteq S': |F_p(R) \cap T'| > |R|$. Let $T'' = F_p(S') \cap T'$. From (a) we know that $|F_p(T')| \leq |S'|$. From (b) we know that $|T''| = |F_p(S') \cap T'| > |S'|$. Since $T'' = F_p(S') \cap T' \subseteq T'$, we have $F_p(T'') \subseteq F_p(T')$ and, thus, $|F_p(T'')| \leq |F_p(T')|$. It follows that $|T''| > |F_p(T'')|$. Since $F_p(T'') \subseteq F_p(T') \subseteq S' \subseteq J \setminus \{j_0\}$, we have $j_0 \notin F_p(T'')$. Hence for T'' neither $|F_p(T'')| \geq |T''|$ nor $j_0 \in F_p(T'')$. Contradiction!

\Leftarrow : Assume that no set of items $S' \subseteq J \setminus \{j_0\}$ is strictly overdemanded. For a contradiction suppose that there exists $T'' \subseteq I: |F_p(T'')| < |T''|$ and $j_0 \notin F_p(T'')$. Consider the smallest such T'' and some $i \in T''$. For all proper subsets $T''' \subset T'': |F_p(T''')| \geq |T'''|$. Hence all bidders in $T'' \setminus \{i\}$ can be matched to items in $F_p(T'' \setminus \{i\})$ by Hall's Theorem [33]. Let μ' be such a matching. Let $\mu'(T'' \setminus \{i\})$ denote the items matched to bidders in $T'' \setminus \{i\}$ under μ' . Compute a maximal alternating tree \mathcal{T} with respect to μ' with root i . Denote the bidders and items in this tree by $T' \subseteq T''$ and $S' = F_p(T') \subseteq F_p(T'')$. It follows that (a) $j_0 \notin S'$ because $S' = F_p(T') \subseteq F_p(T'') \subseteq J \setminus \{j_0\}$ and (b) all items in $F_p(T'') \supseteq F_p(T') = S'$ are matched because otherwise $|F_p(T'')| > |\mu'(T'' \setminus \{i\})| = |\mu'(T'' \setminus \{i\})| = |T'' \setminus \{i\}| = |T''| - 1$, i.e., $|F_p(T'')| \geq |T''|$. Hence Lemma 3 shows that S' is strictly overdemanded with respect to T' . Contradiction!

C Proof of Lemma 3

We prove the claim in two steps:

To 1.: The path P is the path in the maximal alternating tree \mathcal{T} that leads from i_0 to j_0 resp. j . All bidders on this path except i_0 are incident to two edges, one matched and one unmatched, and they are indifferent between the two. Hence we can swap the matched and unmatched edges along P to augment the size of the matching by one.

To 2.: We know that $j_0 \notin S$. From the maximality of the tree \mathcal{T} we get $F_p(T) \subseteq S$. We still have to show that for all $R \subseteq S: |F_p(R) \cap T| > |R|$. For every item set $R \subseteq S$ we know that there exists a node $x \in R$ such that no other node of R lies on the path P from x to

the root i_0 . Note that x is not the root because the root does not belong to R . Let y be the neighbor of x on P . Then y belongs to $F_p(R) \cap T$, but it is not matched to any node in R . Thus, counting the nodes matched to nodes in R and y , there are at least $|R| + 1$ nodes in $F_p(R) \cap T$.

D Proof of Lemma 4

Consider an arbitrary bidder-item pair $(i, j) \in E$. If $(i, k) \in F_p$: Since d is envy free for E and $(i, j) \in E$, we have $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot p_k)$ for all $\lambda > 0$ sufficiently small. If $(i, k) \in (I \times J) \setminus F_p$: Since $u_{i,j}(p_j) > u_{i,k}(p_k)$, we have $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ for all $\lambda > 0$ sufficiently small.⁸ We conclude that $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ for all $(i, k) \in I \times J$ for all $\lambda > 0$ sufficiently small. Since $d_j \geq 0$ for all j s.t. there is a bidder i with $(i, j) \in E$ and $d_j = 0$ otherwise, we have $p_{j_0} + \lambda \cdot d_{j_0} = p_{j_0} = 0$ and $p_j + \lambda \cdot p_j \geq 0$ for all $j \neq j_0$.

E Proof of Lemma 5

For the *if*-part assume that $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ for all sufficiently small $\lambda > 0$, all $(i, j) \in S \subseteq F_p$, and all $(i, k) \in F_p$. Consider arbitrary edges $(i, j) \in S \subseteq F_p$ and $(i, k) \in F_p$. By piece-wise linearity, $u_{i,j}(p_j + \lambda \cdot d_j) = u_{i,j}(p_j) - c_{i,j} \cdot \lambda \cdot d_j$ and $u_{i,k}(p_k + \lambda \cdot d_k) = u_{i,k}(p_k) - c_{i,k} \cdot \lambda \cdot d_k$ for all sufficiently small $\lambda > 0$. From this and the fact that $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ it follows that $u_{i,j}(p_j) - c_{i,j} \cdot \lambda \cdot d_j \geq u_{i,k}(p_k) - c_{i,k} \cdot \lambda \cdot d_k$. Since $(i, j) \in S \subseteq F_p$ and $(i, k) \in F_p$, we have $u_{i,j}(p_j) = u_{i,k}(p_k)$ and, thus, $c_{i,j} \cdot d_j < c_{i,k} \cdot d_k$.

For the *only if*-part assume that $c_{i,j} \cdot d_j \leq c_{i,k} \cdot d_k$ for all $(i, j) \in E \subseteq F_p$ and all $(i, k) \in F_p$. Consider arbitrary edges $(i, j) \in E \subseteq F_p$ and $(i, k) \in F_p$. By piece-wise linearity, $u_{i,j}(p_j + \lambda \cdot d_j) = u_{i,j}(p_j) - c_{i,j} \cdot \lambda \cdot d_j$ and $u_{i,k}(p_k + \lambda \cdot d_k) = u_{i,k}(p_k) - c_{i,k} \cdot \lambda \cdot d_k$ for all sufficiently small $\lambda > 0$.⁹ Since $(i, j) \in E \subseteq F_p$ and $(i, k) \in F_p$, we have $u_{i,j}(p_j) = u_{i,k}(p_k)$. It follows that $u_{i,j}(p_j + \lambda \cdot d_j) = u_{i,j}(p_j) - c_{i,j} \cdot \lambda \cdot d_j \geq u_{i,k}(p_k) - c_{i,k} \cdot \lambda \cdot d_k = u_{i,k}(p_k + \lambda \cdot d_k)$ for all sufficiently small $\lambda > 0$, i.e., $d \neq 0$ is envy free for $E \subseteq F_p$.

F Proof of Lemma 6

Denote the partial matching and the maximal alternating tree corresponding to d by μ' and \mathcal{T}' . Since d is envy free for \mathcal{T}' , Lemma 4 shows that all edges in \mathcal{T}' belong to the first choice graph $G_{p+\lambda \cdot d}$ at prices $p + \lambda \cdot d$ for all $\lambda > 0$ sufficiently small. Since \mathcal{T}' is a maximal alternating tree with item set S and bidder set T in which all items are matched,

⁸There is no discontinuity in the utility function $u_{i,j}(\cdot)$ within the range $[p_j, p_j + \lambda \cdot d_j]$ for all $\lambda > 0$ sufficiently small because the utility function $u_{i,j}(\cdot)$ is locally right-continuous.

⁹For this we need that the intervals are of the form $[..)$.

Lemma 3 shows that (a) the set of items S is strictly overdemanded with respect to the set of bidders T and (b) $|T| = |S| + 1$.

G Proof of Lemma 7

Let x , ω , and ρ be defined as in the mechanism. The constraint matrix of LP is totally unimodular, i.e., $x_{i,j} \in \{0, 1\}$ for all $(i, j) \in F_p \cap (T \setminus \{i_0\} \times S)$ [44]. Hence $\sum_{j \in F_p(i)} x_{i,j} = 1$ ($\forall i \in T \setminus \{i_0\}$) and $\sum_{i \in F_p(j)} x_{i,j} = 1$ ($\forall j \in S$) ensure that $\sigma = \{(i, j) \in T \setminus \{i_0\} \times S \mid x_{i,j} = 1\}$ is a matching that matches every bidder $i \in T \setminus \{i_0\}$ and every item $j \in S$ exactly once. From duality we get that:

1. For all $i \in T \setminus \{i_0\}$ and all $j \in S$ with $(i, j) \in \sigma$: $\omega_i + \rho_j = \log(c_{i,j})$.
2. For all $i \in T \setminus \{i_0\}$ and all $j \in S$: $\omega_i + \rho_j \leq \log(c_{i,j})$.

If we extend ω from $T \setminus \{i_0\}$ to T as described in l. 3, then we also have:

3. There exists a $j \in S$: $\omega_{i_0} + \rho_j = \log_{c_{i_0,j}}$.
4. For all $j \in S$: $\omega_{i_0} + \rho_j \leq \log_{c_{i_0,j}}$.

Let $H_\rho = (S \cup T, E_\rho)$, E_ρ , \mathcal{T}' , T' and S' be defined as in ll. 4-6. It is not difficult to see that the while-loop in ll. 7-11 has the following properties:

- a. For all $i \in T'$ and all $j \in S'$: If we had $\omega_i + \rho_j < (\text{resp. } =) \log(c_{i,j})$ before the update, then we have $\omega_i + \rho_j < (\text{resp. } =) \log(c_{i,j})$ after the update.
- b. For all $i \in T \setminus T'$ and all $j \in S \setminus S'$: If we had $\omega_i + \rho_j < (\text{resp. } =) \log(c_{i,j})$ before the update, then we have $\omega_i + \rho_j < (\text{resp. } =) \log(c_{i,j})$ after the update.
- c. For all $i \in T \setminus T'$ and all $j \in S'$: If we had $\omega_i + \rho_j \leq \log(c_{i,j})$ before the update, then we have $\omega_i + \rho_j < \log(c_{i,j})$ after the update.
- d. For all $i \in T'$ and $j \in S \setminus S'$: If we had $\omega_i + \rho_j < \log(c_{i,j})$ before the update, then we have $\omega_i + \rho_j \leq \log(c_{i,j})$ after the update.
- e. There exist at least one $i \in T'$ and $j \in S \setminus S'$: $\omega_i + \rho_j = \log(c_{i,j})$ after the update.

From a. to d. we get that no edge from σ and \mathcal{T}' in E_ρ is lost. From e. we get that at least one edge from some $i \in T'$ to some $j \in S \setminus S'$ is added to E_ρ . Since this item j was matched under σ along an edge in E_ρ to an item $i' \in T \setminus T'$, we know that after each iteration of the while-loop the maximal alternating tree \mathcal{T}' with respect to σ with root i_0 in H_ρ will at least cover the bidders and items in $T' \cup \{i'\}$ and $S' \cup \{j\}$. Hence, after at most $|S| \leq k$ iterations, \mathcal{T}' will cover the same bidders and items as \mathcal{T} .

Let d be defined as in l. 12. Then $d \neq 0$ is envy free for \mathcal{T}' by Lemma 5 because:

1. For all $(i, j) \in \mathcal{T}'$ and all $(i, k) \in F_p$ we have that $c_{i,j} \cdot d_j \leq c_{i,k} \cdot d_k$ because $e^{\omega_i} = e^{\log(c_{i,j}) - \rho_j} = c_{i,j} \cdot e^{-\rho_j} = c_{i,j} \cdot d_j$ and $e^{\omega_i} \leq e^{\log(c_{i,k}) - \rho_k} = c_{i,k} \cdot e^{-\rho_k} = c_{i,k} \cdot d_k$.
2. For all $(i, j) \in \mu' \setminus \mathcal{T}'$ and $(i, k) \in F_p$ we trivially have that $c_{i,j} \cdot d_j \leq c_{i,k} \cdot d_k$ because $d_j = 0$, $d_k \geq 0$, and $c_{i,k} \geq 0$.

Let μ' be defined as in l. 13, then μ' matches the same bidders and items as μ because (a) μ' is identical to σ on $T \times S$ and (b) μ' is identical to μ on $I \setminus T \times J \setminus S$.

The LP and the DP can be solved in time $O(\min(n, k)^3)$ [27, 46]. The maximal alternating tree \mathcal{T}' can be computed in time $O(\min(n, k)^2)$ using a breadth-first search approach. The while-loop in ll. 8-12 can be implemented using “slack variables” $\delta_j = \min_{(i,j) \in F_p, i \in T'} (\log(c_{i,j}) - \omega_i - \rho_j)$ for each item $j \in S \setminus S'$ so that *all* iterations of the while loop take total time $O(\min(n, k)^2)$: The initialization of the δ_j 's takes time $O(\min(n, k)^2)$ as for each of the up to $\min(n, k)$ items in $S \setminus S'$ the minimum is computed over the up to $\min(n, k)$ bidders in T' . In each iteration of the while-loop at least one bidder-item pair $(i, j) \in T' \times S \setminus S'$ is added to $T' \times S'$. Since $|S \setminus S'| \leq \min(n, k)$ it follows that there are at most $\min(n, k)$ iterations. Using the δ_j 's for $j \in S \setminus S'$ the δ in l.8 can be computed in time $O(\min(n, k))$. When the ω_i 's and ρ_j 's are updated in l. 9, the δ_j 's are adapted in time $O(\min(n, k))$ by subtracting δ from each δ_j . Thus, ll. 9 and 10 take time $O(\min(n, k))$ per iteration for a total of $O(\min(n, k)^2)$. Instead of re-computing the maximal alternating tree \mathcal{T}' in l. 10 from scratch we can keep the old one and add the required edges. Thus maintaining \mathcal{T}' takes only time $O(\min(n, k)^2)$ for all iterations of the while loop. Additionally the δ_j 's must be updated. For each bidder that is added to T' all δ_j 's must be updated. This takes time $O(\min(n, k))$ per bidder. But this happens at most once for each of the up to $\min(n, k)$ bidders that are added to T' , since no bidder is ever removed from T' . Thus, all the updates to the δ_j 's that are required when bidders are added to T' take total time $O(\min(n, k)^2)$.

H Proof of Lemma 8

For a contradiction suppose that $p_s'' < p_s + \lambda \cdot d_s$ for some $s \in S$. Choose $\epsilon > 0$ such that $p_s'' = p_s + (\lambda - \epsilon) \cdot d_s$. Note that $\epsilon \leq \lambda$ because $p_s'' \geq p_s$. Let $A = \{j \in S \mid p_j'' - p_j \leq (\lambda - \epsilon) \cdot d_j\}$ and let $B = F_{p+(\lambda-\epsilon) \cdot d}(A) \cap T$. Note that $A \neq \emptyset$ because $s \in A$.

Since d is envy free for \mathcal{T}' and $p_j \leq p_j + (\lambda - \epsilon) \cdot d_j < p_j + \lambda \cdot d_j \leq e_j$ for all $j \in S$, i.e., for no $(i, j) \in T \times S$ there is a discontinuity in $u_{i,j}(\cdot)$ within the range $[p_j, p_j + (\lambda - \epsilon) \cdot d_j]$, we have $\mathcal{T}' \subseteq F_{p+(\lambda-\epsilon) \cdot d}$. Since \mathcal{T}' covers all bidders in T and items in S and $A \subseteq S$ and S is strictly overdemanded with respect to T , we have that $|B| = |F_{p+(\lambda-\epsilon) \cdot d}(A) \cap T| > |A|$.

Next we show that $|A| \geq |F_{p''}(B)|$. For this it suffices to show that $F_{p''}(i) \subseteq A$ for all $i \in B$. For a contradiction assume that there exists an $i \in B$ and a $k \notin A$ with $k \in F_{p''}(i)$.

It follows that

$$p_k'' - p_k > (\lambda - \epsilon) \cdot d_k, \quad \text{and} \quad (15)$$

$$u_{i,k}(p_k'') \geq u_{i,j}(p_j'') \quad \text{for all } j. \quad (16)$$

But by the definition of B , and since $i \in B$, there must be a $j \in A$ such that $j \in F_{p+(\lambda-\epsilon)\cdot d}(i)$. It follows that

$$p_j'' - p_j \leq (\lambda - \epsilon) \cdot d_j, \quad \text{and} \quad (17)$$

$$u_{i,j}(p_j + (\lambda - \epsilon) \cdot d_j) \geq u_{i,k}(p_k + (\lambda - \epsilon) \cdot d_k). \quad (18)$$

Using the fact that the utility functions are strictly monotonically decreasing we get

$$u_{i,j}(p_j'') \geq u_{i,j}(p_j + (\lambda - \epsilon) \cdot d_j) \quad (\text{from (17)})$$

$$\geq u_{i,k}(p_k + (\lambda - \epsilon) \cdot d_k) \quad (\text{from (18)})$$

$$> u_{i,k}(p_k''). \quad (\text{from (15)})$$

Since this would give a contradiction to (16), we must have $F_{p''}(B) \subseteq A$, i.e., $|A| \geq |F_{p''}(B)|$. It follows that $|B| > |A| \geq |F_{p''}(B)|$. But this shows that in (μ'', p'') not all bidders can be matched in an envy free manner. Contradiction!

I Proof of Theorem 1

We prove the claim in two steps:

Correctness. The matching μ is a subset of the first choice edges at prices p . Hence Lemma 2 shows that (μ, p) is envy free. By Lemma 1 (μ, p) is bidder optimal if $p_j \leq p_j''$ for every item j and every envy free outcome (μ'', p'') . Let $p^{(t)}$ denote the prices after the t -th update. We prove that $p_j^{(t)} \leq p_j''$ for every item j , every envy free outcome (μ'', p'') , and all time steps t by induction over t .

For $t = 0$: The claim follows from the fact that any envy free outcome (μ'', p'') has $p_j'' \geq 0$ for all j .

For $t-1 \rightarrow t$: Assume that the claim is true for $t-1$. Let \mathcal{T} be the maximal alternating tree with respect to the current matching $\mu^{(t-1)}$ with root i_0 right before the t -th update. Let S and T denote the items and bidders in \mathcal{T} . Let d be an envy free price increase for \mathcal{T} with corresponding alternating tree \mathcal{T}' and matching μ' . Let λ be defined as in the mechanism. Note that λ can be computed in time $O(\min(n, k) \cdot k)$ by iterating over all bidders in T , of which there are at most $\min(n, k)$, and all items in J . The mechanism sets $p_j^{(t+1)} = p_j^{(t)}$ for $j \notin S$ and $p_j^{(t+1)} = p_j^{(t)} + \lambda \cdot d_j$ for $j \in S$. Lemma 8 shows that any envy free outcome (μ'', p'') with $p'' \geq p^{(t)}$ must have $p_j'' \geq p_j^{(t)} + \lambda \cdot d_j$. It follows that $p_j'' \geq p_j^{(t+1)}$ for all j .

Running time. We bound the total time required by (1) the outer while-loop *without* the inner while-loop (ll. 2-4 & 12-13) and the inner while-loop corresponding to Case (a) (ll. 5-11) separately from the total time required by (2) the inner while-loop corresponding to Case (b) (ll. 5-11).

To bound the total time required by (1) observe that: (i) In each iteration of the outer while-loop exactly one bidder gets matched. (ii) Bidders can get unmatched only if the boundary of a box is reached that corresponds to a discontinuity in at least one of the utility functions $u_{i,j}(\cdot)$. (iii) A discontinuity in $u_{i,j}(\cdot)$ can only unmatch bidder i . Since the prices are monotonically increasing at most $O(\sum_{i,j} d_{i,j})$ discontinuities are reached. From (i) to (iii) we deduce that there are at most $O(n + \sum_{i,j} d_{i,j})$ iterations of the outer-while loop without the inner while-loop and that the inner while-loop is executed at most $O(n + \sum_{i,j} d_{i,j})$ times. Each iteration of the outer-while without the inner while-loop takes time $O(\min(n, k)^2)$. Each execution of the inner-while loop consists of at most $\min(n, k)$ iterations corresponding to Case (a). Each of these iterations takes time $O(\min(n, k) \cdot (\min(n, k)^2 + k))$, namely $O(\min(n, k)^3)$ for computing the overdemand preserving price increase (see Lemma 7) and $O(\min(n, k) \cdot k)$ for computing the λ value. Hence the total time required by (1) is $O((n \cdot \min(n, k) + \sum_{i,j} d_{i,j} \cdot \min(n, k)) \cdot \min(n, k) \cdot (\min(n, k)^2 + k))$.

To bound the total time required by (2) observe that because the prices are monotonically increasing there are at most $O(\sum_{i,j} t_{i,j})$ iterations of the inner-while loop that correspond to Case (b). As argued above each iteration of the inner while-loop takes time $O(\min(n, k) \cdot (\min(n, k)^2 + k))$. Hence the total time required by (2) is $O((\sum_{i,j} t_{i,j}) \cdot \min(n, k) \cdot (\min(n, k)^2 + k))$.

J Proof of Lemma 9

We prove the claim in two steps.

To 1.: Denote the original input by $(u_{i,j}(\cdot), r_j, o_i)$ and the associated input in standard form by $(u'_{i,j}(\cdot), 0, o_i)$. Recall that $u'_{i,j}(p_j) = u_{i,j}(p_j + r_j)$ for all i , all j , and all $p_j \geq 0$.

Every pair of walks P and Q in the input graph for input $(u_{i,j}(\cdot), r_j, o_i)$ corresponds to a pair of walks P' and Q' in the input graph for input $(u'_{i,j}(\cdot), 0, o_i)$, and vice versa. Since $c_{i,j} = c'_{i,j}$ for every edge (i, j) on P and P' resp. Q and Q' , every price increase d for which $d_j = (c_{i,j}/c_{i,j'}) \cdot d_{j'}$ for any two edges (i, j) and (i, j') on P resp. Q has $d_j = (c'_{i,j}/c'_{i,j'}) \cdot d_{j'}$ for any two edges (i, j) and (i, j') on P' resp. Q' , and vice versa. Furthermore, since $v_{i,j} = v'_{i,j} + c_{i,j} \cdot r_j$ for every edge (i, j) on P and P' resp. Q and Q' and $D_{i,j} = D'_{i,j} + r_j$ for every discontinuity edge (i, j) on P and P' resp. Q and Q' (a) the weight of each forward edge (i, j) is $-(1/d_j) \cdot (v_{i,j}/c_{i,j}) = -(1/d_j) \cdot (v'_{i,j}/c_{i,j}) - (1/d_j) \cdot r_j$, (b) the weight of each backward edge is $(1/d_j) \cdot (v_{i,j}/c_{i,j}) = (1/d_j) \cdot (v'_{i,j}/c_{i,j}) + (1/d_j) \cdot r_j$, and (c) the weight of each discontinuity edge is $(1/d_j) \cdot D_{i,j} - (1/d_j) \cdot (v_{i,j}/c_{i,j}) = (1/d_j) \cdot D'_{i,j} - (1/d_j) \cdot (v'_{i,j}/c_{i,j})$. Since on each of the walks P and P' resp. Q and Q' there are as many forward as backward edges and every item j is incident to exactly one forward and one backward edge, the $-(1/d_j) \cdot r_j$

and $(1/d_j) \cdot r_j$ terms cancel each other. Hence $w_d(P)$ is identical to $w_d(P')$ and $w_d(Q)$ is identical to $w_d(Q')$. The claim follows.

To 2.: We prove each of the claims (i) to (iv) individually.

To (i): For a contradiction suppose that two discontinuities are reached at the same time. W.l.o.g. assume that the discontinuities are reached on (i_1, j_1) and (i_t, j_{t-1}) , and denote them by D_{i_1, j_1} and $D_{i_t, j_{t-1}}$. Consider the walks $P = (i_1, j_1)$ and $Q = (i_1, j_1, i_2, j_2, \dots, i_{t-1}, j_{t-1}, i_t, j_{t-1})$, both alternate between forward and backward edges and end with a discontinuity edge. Note that these walks always exist because (a) either j_1 lies on the path from i_1 to i_t or $i_2 = i_1$ and (b) either j_{t-1} lies on the path from i_1 to i_t or $i_t = i_{t-1}$. Since the discontinuities are reached at the same time there must be a price increase $d \neq 0$ that is envy free for some maximal alternating tree containing P and Q and a positive scalar $\lambda > 0$ such that $\lambda \cdot d_{j_1} = D_{i_1, j_1} - p_{j_1}$ and $\lambda \cdot d_{j_{t-1}} = D_{i_t, j_{t-1}} - p_{j_{t-1}}$. Rearranging this gives

$$\frac{1}{d_{j_1}} \cdot D_{i_1, j_1} - \frac{1}{d_{j_{t-1}}} \cdot D_{i_t, j_{t-1}} = \frac{1}{d_{j_1}} \cdot p_{j_1} - \frac{1}{d_{j_{t-1}}} \cdot p_{j_{t-1}} \quad (19)$$

From the fact that P and Q belong to the first choice graph G_p at prices p we get that for $s = 1..t-2$: $v_{i_{s+1}, j_s} - c_{i_{s+1}, j_s} \cdot p_{j_s} = v_{i_{s+1}, j_{s+1}} - c_{i_{s+1}, j_{s+1}} \cdot p_{j_{s+1}}$ (\star). Solving for p_{j_s} we get that for $s = 1..t-2$:

$$p_{j_s} = \frac{v_{i_{s+1}, j_s}}{c_{i_{s+1}, j_s}} - \frac{v_{i_{s+1}, j_{s+1}}}{c_{i_{s+1}, j_{s+1}}} + \frac{c_{i_{s+1}, j_{s+1}}}{c_{i_{s+1}, j_s}} \cdot p_{j_{s+1}}. \quad (20)$$

From the fact that P and Q also belong to the first choice graph $G_{p+\lambda \cdot d}$ at prices $p + \lambda \cdot d$ we get that for $s = 1..t-2$: $v_{i_{s+1}, j_s} - c_{i_{s+1}, j_s} \cdot (p_{j_s} + \lambda \cdot d_{j_s}) = v_{i_{s+1}, j_{s+1}} - c_{i_{s+1}, j_{s+1}} \cdot (p_{j_{s+1}} + \lambda \cdot d_{j_{s+1}})$. Subtracting (\star) and solving for $d_{j_{s+1}}$ we get that for $s = 1..t-2$:

$$d_{j_{s+1}} = \frac{c_{i_{s+1}, j_s}}{c_{i_{s+1}, j_{s+1}}} \cdot d_{j_s}. \quad (21)$$

Solving the recurrence (20) for p_{j_1} , substituting (21), and rearranging gives:

$$\begin{aligned} \frac{1}{d_{j_1}} \cdot p_{j_1} - \frac{1}{d_{j_{t-1}}} \cdot p_{j_{t-1}} &= \frac{1}{d_{j_2}} \cdot \frac{v_{i_2, j_1}}{c_{i_2, j_2}} - \frac{1}{d_{j_2}} \cdot \frac{v_{i_2, j_2}}{c_{i_2, j_2}} + \\ &\dots + \frac{1}{d_{j_{t-1}}} \cdot \frac{v_{i_{t-1}, j_{t-2}}}{c_{i_{t-1}, j_{t-1}}} - \frac{1}{d_{j_{t-1}}} \cdot \frac{v_{i_{t-1}, j_{t-1}}}{c_{i_{t-1}, j_{t-1}}}. \end{aligned}$$

We combine this with (19) to get

$$\begin{aligned} \frac{1}{d_{j_1}} \cdot D_{i_1, j_1} - \frac{1}{d_{j_{t-1}}} \cdot D_{i_t, j_{t-1}} &= \frac{1}{d_{j_2}} \cdot \frac{v_{i_2, j_1}}{c_{i_2, j_2}} - \frac{1}{d_{j_2}} \cdot \frac{v_{i_2, j_2}}{c_{i_2, j_2}} + \\ &\dots + \frac{1}{d_{j_{t-1}}} \cdot \frac{v_{i_{t-1}, j_{t-2}}}{c_{i_{t-1}, j_{t-1}}} - \frac{1}{d_{j_{t-1}}} \cdot \frac{v_{i_{t-1}, j_{t-1}}}{c_{i_{t-1}, j_{t-1}}}. \end{aligned}$$

We add $(-1/d_{j_1}) \cdot (v_{i_1, j_1}/c_{i_1, j_1})$ to both sides and $(1/d_{j_t}) \cdot (v_{i_t, j_{t-1}}/c_{i_t, j_t}) - (1/d_{j_{t-1}}) \cdot (v_{i_t, j_{t-1}}/c_{i_t, j_{t-1}}) = 0$ to the right hand side. After rearranging we get:

$$\begin{aligned} & \frac{1}{d_{j_1}} \cdot D_{i_1, j_1} - \frac{1}{d_{j_1}} \cdot \frac{v_{i_1, j_1}}{c_{i_1, j_1}} = -\frac{1}{d_{j_1}} \cdot \frac{v_{i_1, j_1}}{c_{i_1, j_1}} + \frac{1}{d_{j_2}} \cdot \frac{v_{i_2, j_1}}{c_{i_2, j_2}} \\ & - \frac{1}{d_{j_2}} \cdot \frac{v_{i_2, j_2}}{c_{i_2, j_2}} + \dots + \frac{1}{d_{j_{t-1}}} \cdot \frac{v_{i_{t-1}, j_{t-2}}}{c_{i_{t-1}, j_{t-1}}} - \frac{1}{d_{j_{t-1}}} \cdot \frac{v_{i_{t-1}, j_{t-1}}}{c_{i_{t-1}, j_{t-1}}} \\ & + \frac{1}{d_{j_t}} \cdot \frac{v_{i_t, j_{t-1}}}{c_{i_t, j_t}} + \frac{1}{d_{j_{t-1}}} \cdot (D_{i_t, j_{t-1}} - \frac{v_{i_t, j_{t-1}}}{c_{i_t, j_{t-1}}}). \end{aligned}$$

Since the left hand side corresponds to $w_d(P)$ and the right hand side to $w_d(Q)$, we get a contradiction to the fact that the input is in general position.

To (ii): Consider an arbitrary price update. Let i_0 be the bidder, and let μ be the matching under consideration. Denote the maximal alternating tree with respect to μ with root i_0 by \mathcal{T} . Let d be the overdemand preserving price increase for \mathcal{T} computed by our mechanism. Denote the corresponding maximal alternating tree and matching by \mathcal{T}' resp. μ' . We know that μ' matches the same set of bidders and items as μ . If the price update corresponds to a $\lambda_{\text{out}} < \lambda_{\text{box}}$, then none of the edges in \mathcal{T}' drops out of the first choice graph $G_{p+\lambda \cdot d}$, i.e., none of the edges has to be removed from μ' . If it corresponds to a $\lambda_{\text{box}} \geq \lambda_{\text{out}}$, then – due to the fact that the input is in general position – at most one edge in \mathcal{T}' drops out of the first choice graph $G_{p+\lambda \cdot d}$. If the edge is an unmatched edge, then no edge has to be removed from μ' . If it is a matched edge, then exactly one edge has to be removed from μ' . But since all items in \mathcal{T}' are closer to the root i_0 than the bidder they are matched to, there must be an alternating path P from the unmatched item, say j , to the root i_0 . It follows that in the subsequent iteration there is exactly one alternating path along which the matching can be augmented, namely the path from i_0 to j . Hence j gets matched again in the subsequent iteration.

To (iii): Our mechanism raises the price of an item strictly above its reserve price only if this item belonged to the set of items S in a maximal alternating tree \mathcal{T} in which all items are matched. From (a) we know that every item that ever got matched, will be matched in the end.

To (iv): Consider the iteration in which the last item j gets matched. Since j belongs to $F_p(T)$, we must have $p_j \geq r_j$. From (a) we know that item j was never matched before and from (b) we know that $p_j > r_j$ would imply that j was matched before. Thus we must have $p_j = r_j$.

K Proof of Lemma 10

We prove the claim in two steps:

To 1.: For a contradiction assume that there exists a feasible outcome (μ', p') with $u_{i, \mu'(i)}(p'_{\mu'(i)}) > u_{i, \mu^*(i)}(p^*_{\mu^*(i)})$ for all i . Let (μ, p) denote the outcome of our mechanism.

Since $u_{i,\mu(i)}(p_{\mu(i)}) = u_{i,\mu^*(i)}(p_{\mu^*(i)}^*)$ for all i , it follows that $u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu(i)}(p_{\mu(i)})$ for all i .

Consider any pair $(i, j) \in \mu'$. It follows that $u_{i,j}(p'_j) = u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu(i)}(p_{\mu(i)}) \geq u_{i,j}(p_j)$ and, thus, $p_j > p'_j \geq r_j$. Condition (iii) from Lemma 9 implies that item j is matched under μ . We conclude that (a) all the items that are matched under μ' are also matched under μ and (b) $p'_j < p_j$ for all of these items j .

Case 1: At least one bidder i is matched to the dummy item j_0 under μ' . By (b) $p_{j_0} > p'_{j_0} \geq 0$, which contradicts the feasibility of the outcome (μ, p) .

Case 2: All bidders i are matched to non-dummy items j under μ' . By (a) all bidders are matched to non-dummy items under μ . Condition (iv) from Lemma 9 shows that at least one item j is matched under μ at price $p_j = r_j$. By (b) $p'_j < p_j = r_j$, which contradicts the feasibility of the outcome (μ', p') .

To 2.: Let (μ, p) denote the outcome of our mechanism. Since $u_{i,\mu(i)}(p_{\mu(i)}) = u_{i,\mu^*(i)}(p_{\mu^*(i)}^*)$ for all i , we have $I^+ = \{i \in I \mid u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu(i)}(p_{\mu(i)})\} \neq \emptyset$. Let $\mu(I^+)$ resp. $\mu'(I^+)$ denote the set of items matched to bidders in I^+ under μ resp. μ' . From the first part of this lemma we know that $I^+ \neq I$.

Case 1: $\mu(I^+) \neq \mu'(I^+)$. There must be an item $j \in \mu'(I^+)$ such that $j \notin \mu(I^+)$. Let $i' \in I^+$ be the bidder that is matched to item j in μ' . Since $i' \in I^+$ and the outcome (μ, p) is envy free we have that $u_{i',j}(p'_j) = u_{i',\mu'(i')}(p'_{\mu'(i')}) > u_{i',\mu(i')}(p_{\mu(i')}) \geq u_{i',\mu'(i')}(p_{\mu'(i')}) = u_{i',j}(p_j)$ which shows that $p_j > p'_j$. Hence j is matched under μ by condition (iii) from Lemma 9. Let $i \in I \setminus I^+$ be the bidder that is matched to item j under μ . Since $i \notin I^+$ it follows that $u_{i,\mu'(i)}(p'_{\mu'(i)}) \leq u_{i,\mu(i)}(p_{\mu(i)}) = u_{i,j}(p_j) < u_{i,j}(p'_j)$.

Case 2: $\mu(I^+) = \mu'(I^+)$. Let $J^+ = \mu(I^+) = \mu'(I^+)$. Consider the following *restricted problem*: The set of bidders is I^+ , the set of items is J^+ , the utility functions are $u_{i,j}^+(\cdot) = u_{i,j}(\cdot)$ for all $(i, j) \in I^+ \times J^+$, the reserve prices are $r_j^+ = \max(r_j, \max_{i \notin I^+} (u_{i,j}^{-1}(u_{i,\mu(i)}(p_{\mu(i)}))), 0)$ for all $j \in J^+$, and the outside options are $o_i^+ = o_i$ for all $i \in I^+$.¹⁰ Since the outcome (μ, p) is envy free for the original input it is also envy free for the restricted input. It is even bidder optimal because the existence of an envy free outcome (μ'', p'') for the restricted problem in which at least one bidder $i \in I^+$ has a strictly higher utility would imply the existence of an envy free outcome (μ''', p''') for the original problem with this property and therefore contradict the bidder optimality of (μ, p) .

Case 2.1: The outcome (μ', p') is feasible for the restricted problem. From the first part of this lemma we know that there exists a bidder $i \in I^+$ such that $u_{i,\mu'(i)}(p'_{\mu'(i)}) \leq u_{i,\mu(i)}(p_{\mu(i)})$. This contradicts the definition of I^+ .

Case 2.2: The outcome (μ', p') is *not* feasible for the restricted problem. This can

¹⁰If $u_{i,j}(\cdot)$ is continuous then $u_{i,j}^{-1}(\cdot)$ is indeed the inverse function. More generally, it is defined for $u \in [o_i, \infty)$ by $u_{i,j}^{-1}(u) := \min_{p_j \in [r_j, \infty)} \{u_{i,j}(p_j) \leq u\}$, and is merely a one-sided inverse function satisfying $u_{i,j}^{-1}(u_{i,j}(p_j)) = p_j$. Note that the minimum is indeed contained in the set itself as we only consider right-continuous utility functions.

only happen if $p'_j < r_j^+$ for some item $j \in J^+$. Since the outcome (μ', p') is feasible for the original problem this can only happen if $r_j^+ > r_j$ and so $r_j^+ = \max_{i \notin I^+} (u_{i,j}^{-1}(u_{i,\mu(i)}(p_{\mu(i)})), 0)$. We cannot have $r_j^+ = 0$ as this would imply $p'_j < r_j^+ = 0$. Thus we must have $r_j^+ = u_{i,j}^{-1}(u_{i,\mu(i)}(p_{\mu(i)}))$ for some $i \in I \setminus I^+$. It follows that $p'_j < r_j^+ = u_{i,j}^{-1}(u_{i,\mu(i)}(p_{\mu(i)})) \leq u_{i,j}^{-1}(u_{i,\mu'(i)}(p'_{\mu'(i)}))$ and so $u_{i,\mu'(i)}(p'_{\mu'(i)}) < u_{i,j}(p'_j)$.

L Proof of Theorem 2

For a contradiction suppose some subset of bidders $I^+ \subseteq I$ strictly benefits from misreporting their utility functions. Denote the true input by $(u_{i,j}(\cdot), r_j, o_i)$, and the falsified one by $(u'_{i,j}(\cdot), r_j, o_i)$. Note that $u'_{i,j}(\cdot) = u_{i,j}(\cdot)$ for all $(i, j) \in I \setminus I^+ \times J$. Let (μ^*, p^*) resp. (μ', p') denote the bidder optimal outcome for the true resp. falsified input. Then $I^+ = \{i \in I \mid u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu^*(i)}(p^*_{\mu^*(i)})\}$. Note that (μ', p') is feasible for the true input $(u_{i,j}(\cdot), r_j, o_i)$ because $p'_{j_0} = 0$ and $p'_j \geq 0$ for all $j \neq j_0$.

Case 1: $I^+ = I$. The first part of Lemma 10 shows that if condition (iii) and (iv) from Lemma 9 are satisfied, then no feasible outcome (μ', p') can give all bidders a strictly higher utility than the bidder optimal outcome (μ^*, p^*) . This gives a contradiction.

Case 2: $I^+ \neq I$. The second part of Lemma 10 shows that if condition (iii) and (iv) from Lemma 9 are satisfied, then some feasible outcome (μ', p') gives only some of the bidders a strictly higher utility than the bidder optimal outcome (μ^*, p^*) , then there must be at least one bidder $i \in I \setminus I^+$ and an item $j \in S$ for which $u_{i,\mu'(i)}(p'_{\mu'(i)}) < u_{i,j}(p'_j)$. But since $i \notin I^+$, i.e., i is not among the liars, this implies that $u'_{i,\mu'(i)}(p'_{\mu'(i)}) = u_{i,\mu'(i)}(p'_{\mu'(i)}) < u_{i,j}(p'_j) = u'_{i,j}(p'_j)$ and contradicts the fact that (μ', p') is bidder optimal and therefore envy free for the falsified input $(u'_{i,j}(\cdot), r_j, o_i)$.

M Proof Outline of Theorem 3

The proof has the following structure. Lemma 11 shows that continuous functions with a bounded second derivative can be approximated using piece-wise linear functions with an error of ϵ using at most $O(1/\sqrt{\epsilon})$ segments. This is needed to guarantee that we can solve the approximated input exactly in an efficient manner. Lemma 13 shows that the individual slopes of such piece-wise linear approximations cannot exceed/be inferior to the maximum/minimum slope of the original function, independent of ϵ . This is needed as Lemma 14 then uses these maximum (resp. minimum) slopes to bound changes in the price (resp. utility) function in terms of changes in the utility function/price, i.e., small changes in one correspond to small changes in the other. Lemma 15 proves that the ratio between biggest and smallest non-zero entries of the overdemand preserving price increase d computed during the execution of our mechanism can be bounded by $(M/m)^{\min(n,k)}$, guaranteeing that if *any* price for an item in S is increased by some amount, then *all*

prices for items in S are increased by similar amounts. Lemma 17 lets us attribute the price increase of all consecutive case of 7(b) in the mechanism, where an end point $\tilde{e}_{i,j}^{(t)}$ is reached but no new bidder-item pair enters the first choice graph, to the following case of 7(a). This observation is needed as the number of cases 7(b) can be very large and will depend on ϵ , whereas the number of cases of 7(a) is comparatively small and independent of ϵ . Finally, Lemma 12 shows that if we use the prices p , or rather $p + \epsilon/m$, of an (unknown) bidder optimal outcome (μ, p) for $u_{i,j}(\cdot)$ to “jumpstart” our mechanism in Section 3, then the price increases attributed to each case 7 (a) are small and there at most $O(\min(n, k)^2)$ such small changes to the initial prices p and to the corresponding utilities. This guarantees that the outcome $(\tilde{\mu}, \tilde{p})$ found for a “fresh” run of the mechanism with input $\tilde{u}_{i,j}(\cdot)$ cannot have prices much higher than the unknown bidder optimal prices p for input $u_{i,j}(\cdot)$. The actual proof of Theorem 3, which then uses $\epsilon = \gamma(2M/m)^{-(\min(n, k)+1)^3}$, is given in Appendix R.

N Proof of Lemma 11

Point (i) is trivial as $\tilde{u}_{i,j}(p_j)$ is linear on each segment. Point (ii) holds as $\tilde{u}_{i,j}(\tilde{e}_{i,j}^{(t)}) = u_{i,j}(\tilde{e}_{i,j}^{(t)})$ and $\tilde{u}_{i,j}(\tilde{e}_{i,j}^{(t+1)}) = u_{i,j}(\tilde{e}_{i,j}^{(t+1)})$ so that for two adjacent segments the approximation agrees on the boundaries. To see (iii) we will argue as follows.

Given that the line segment $\tilde{u}_{i,j}(p_j) = v_{i,j} - c_{i,j} \cdot p_j$ intersects the function $u_{i,j}(p_j)$ at points $\tilde{e}_{i,j}^{(t)}$ and $\tilde{e}_{i,j}^{(t+1)} > \tilde{e}_{i,j}^{(t)}$, what is the maximum $\max_{p_j \in [\tilde{e}_{i,j}^{(t)}, \tilde{e}_{i,j}^{(t+1)}]} |u_{i,j}(p_j) - \tilde{u}_{i,j}(p_j)|$ for *any* continuous function $u_{i,j}(p_j)$ with a bounded second derivative? Let us define $g_{i,j}(p_j) = u_{i,j}(p_j) - \tilde{u}_{i,j}(p_j)$. Note that $g_{i,j}(p_j)$ is continuous, satisfies $g_{i,j}(\tilde{e}_{i,j}^{(t)}) = g_{i,j}(\tilde{e}_{i,j}^{(t+1)}) = 0$ and is twice differentiable. Denote the first and second derivative of $g_{i,j}(\cdot)$ by $\dot{g}_{i,j}(\cdot)$ and $\ddot{g}_{i,j}(\cdot)$. Suppose $g_{i,j}(p_j)$ reaches an extreme point at $p_j = \tilde{e}_{i,j}^{(t+1)} - \delta \in [\tilde{e}_{i,j}^{(t)}, \tilde{e}_{i,j}^{(t+1)}]$ with $\dot{g}_{i,j}(\tilde{e}_{i,j}^{(t+1)} - \delta) = 0$. Then, for $\tilde{e}_{i,j}^{(t+1)} - \delta \leq p_j \leq \tilde{e}_{i,j}^{(t+1)}$, $|\dot{g}_{i,j}(p_j)| \leq |\dot{g}_{i,j}(\tilde{e}_{i,j}^{(t+1)} - \delta)| + |\int_{\tilde{e}_{i,j}^{(t+1)} - \delta}^{p_j} \ddot{g}_{i,j}(z) dz| \leq \int_{\tilde{e}_{i,j}^{(t+1)} - \delta}^{p_j} |\ddot{g}_{i,j}(z)| dz \leq (p_j + \delta - \tilde{e}_{i,j}^{(t+1)}) \cdot B$. And so the maximum of

$|g_{i,j}(p_j)|$ between $\tilde{e}_{i,j}^{(t+1)} - \delta$ and $\tilde{e}_{i,j}^{(t+1)}$ is

$$\begin{aligned}
& g_{i,j}(\tilde{e}_{i,j}^{(t+1)}) + \int_{\tilde{e}_{i,j}^{(t+1)}}^{p_j = \tilde{e}_{i,j}^{(t+1)} - \delta} \dot{g}_{i,j}(p_j) dp_j \\
& \leq \left| \int_{p_j = \tilde{e}_{i,j}^{(t+1)} - \delta}^{\tilde{e}_{i,j}^{(t+1)}} |\dot{g}_{i,j}(p_j)| dp_j \right| \\
& \leq \left| \int_{p_j = \tilde{e}_{i,j}^{(t+1)} - \delta}^{\tilde{e}_{i,j}^{(t+1)}} (p_j + \delta - b) \cdot B dp_j \right| \\
& = \left((\tilde{e}_{i,j}^{(t+1)})^2 - (\delta - \tilde{e}_{i,j}^{(t+1)})^2 \right) / 2 + \delta(\delta - \tilde{e}_{i,j}^{(t+1)}) \cdot B \\
& = \delta^2 / 2 \cdot B.
\end{aligned}$$

Similarly, the maximum of $|\dot{g}_{i,j}(p_j)|$ in the interval $[\tilde{e}_{i,j}^{(t)}, \tilde{e}_{i,j}^{(t+1)} - \delta]$ is $(\tilde{e}_{i,j}^{(t+1)} - \tilde{e}_{i,j}^{(t)} - \delta)^2 B / 2$. As $g_{i,j}(p_j)$ is continuous, the maximum achieved at the point $\tilde{e}_{i,j}^{(t+1)} - \delta$ coming from the left, increasing from $g_{i,j}(\tilde{e}_{i,j}^{(t)}) = 0$, has to equal the maximum achieved at this point coming from the right, increasing from $g_{i,j}(\tilde{e}_{i,j}^{(t+1)}) = 0$. So $|g_{i,j}(p_j)|$ is bounded by $\min \left((\tilde{e}_{i,j}^{(t+1)} - \tilde{e}_{i,j}^{(t)} - \delta)^2 B / 2, \delta^2 B / 2 \right)$, which is maximized for $\delta = (\tilde{e}_{i,j}^{(t+1)} - \tilde{e}_{i,j}^{(t)}) / 2$, corresponding to a peak of at most $(\tilde{e}_{i,j}^{(t+1)} - \tilde{e}_{i,j}^{(t)})^2 / 8 \cdot B \leq w^2 / 8 \cdot B \leq \epsilon$.

O Bounding Changes in Continuous Approximations

The following two auxiliary lemmata are needed to obtain bounds, independent of ϵ , on the changes in utility in terms of changes in prices and, similarly, on the changes in prices that are sufficient to cause a certain change in utility.

Lemma 13. *Given any continuous, differentiable function $u_{i,j}(p_j)$ and a piece-wise linear approximation such that for both the left and the right end point of each segment $[\tilde{e}_{i,j}^{(t)}, \tilde{e}_{i,j}^{(t+1)}]$ the value of the linear function $\tilde{u}_{i,j}(p_j) = v_{i,j} - c_{i,j} \cdot p_j$ agrees with $u_{i,j}(p_j)$, we have that the slope $-c_{i,j}$ of $\tilde{u}_{i,j}(p_j)$ satisfies $\min_{p_j \in [\tilde{e}_{i,j}^{(t)}, \tilde{e}_{i,j}^{(t+1)}]} \dot{u}_{i,j}(p_j) \leq -c_{i,j} \leq \max_{p_j \in [\tilde{e}_{i,j}^{(t)}, \tilde{e}_{i,j}^{(t+1)}]} \dot{u}_{i,j}(p_j)$.*

Recall that we refer to such an approximation that agrees with the original function at both endpoints of each interval as a *point-to-point approximation*.

Proof. This is a generalization of the Intermediate Value Theorem. The proof is by contradiction. If $\min_{p_j \in [\tilde{e}_{i,j}^{(t)}, \tilde{e}_{i,j}^{(t+1)}]} \dot{u}_{i,j}(p_j) > -c_{i,j}$ then $u_{i,j}(\tilde{e}_{i,j}^{(t+1)}) = u_{i,j}(\tilde{e}_{i,j}^{(t)}) + \int_{p_j = \tilde{e}_{i,j}^{(t)}}^{\tilde{e}_{i,j}^{(t+1)}} \dot{u}_{i,j}(p_j) dp_j > \tilde{u}_{i,j}(\tilde{e}_{i,j}^{(t)}) + \int_{p_j = \tilde{e}_{i,j}^{(t)}}^{\tilde{e}_{i,j}^{(t+1)}} (-c_{i,j}) dp = \tilde{u}_{i,j}(\tilde{e}_{i,j}^{(t+1)})$, contradicting that $u_{i,j}(\tilde{e}_{i,j}^{(t+1)}) = \tilde{u}_{i,j}(\tilde{e}_{i,j}^{(t+1)})$. Similarly for the other direction. \square

Lemma 14. *For a monotonically decreasing, continuous, piece-wise linear function $\tilde{u}_{i,j}(\cdot)$ with flattest slope m and steepest slope M we have:*

$$|p_j - p'_j| \leq \rho \Rightarrow |\tilde{u}_{i,j}(p_j) - \tilde{u}_{i,j}(p'_j)| \leq M \cdot \rho \quad (22)$$

$$|\tilde{u}_{i,j}(p_j) - \tilde{u}_{i,j}(p'_j)| \leq \sigma \Rightarrow |p_j - p'_j| \leq 1/m \cdot \sigma \quad (23)$$

Proof. For a single line segment with slope $-c_{i,j}$ we have $\tilde{u}_{i,j}(p_j) - \tilde{u}_{i,j}(p'_j) = -c_{i,j}(p_j - p'_j)$. And so $|\tilde{u}_{i,j}(p_j) - \tilde{u}_{i,j}(p'_j)| \leq M|p_j - p'_j|$ and $|p_j - p'_j| \leq |\tilde{u}_{i,j}(p_j) - \tilde{u}_{i,j}(p'_j)|/m$. If the price difference $p_j - p'_j$ spans several segments, the result follows by applying the result to each segment individually and then adding the individual bounds. As the segments connect in a continuous manner and as $\tilde{u}_{i,j}(\cdot)$ is monotonically decreasing, the total utility difference equals the sum of the differences achieved over each segment. \square

P Analysis of Mechanism for Continuous Utility Functions

The following three auxiliary lemmata are needed to bound the price increase between two consecutive cases 7(a) in the mechanism in Section 3.

Lemma 15. *The overdemand-preserving price increase d computed in line 6 of our mechanism satisfies $\max_{j' \in S, j \in S} d_{j'}/d_j \leq (M/m)^{\min(n,k)}$.*

This first lemma is required so that we know that if any price for an element in S increases by some amount β , then *all* prices in S increase by at least $(m/M)^{\min(n,k)} \cdot \beta$. Conversely, if we have to wait for some price in S to increase by β , then this means that *all* other prices will have increased by at most $(M/m)^{\min(n,k)} \cdot \beta$ in the meantime.

Proof. Let j and j' be two distinct elements in S . As \mathcal{T} is a tree we know that j and j' are connected by a path, alternating between items in S and bidders in T , of length at most $2 \min(|S|, |T|) \leq 2 \min(n, k)$. Let this path be $j = j_0, i_0, j_1, i_1, j_2, \dots, j_{s-1}, i_{s-1}, j_s = j'$. From Lemma 5 we know that as d is envy free for \mathcal{T} we must have both $c_{t,t} \cdot d_t \leq c_{t,t+1} \cdot d_{t+1}$ and $c_{t,t} \cdot d_t \geq c_{t,t+1} \cdot d_{t+1}$, and hence $c_{t,t} \cdot d_t = c_{t,t+1} \cdot d_{t+1}$ for all $0 \leq t < s$. Applying this to the whole path we get $d_{j'} = d_j \prod_{t=0}^{s-1} c_{t,t}/c_{t,t+1} \leq (M/m)^{\min(n,k)} d_j$. \square

Lemma 16. *Let $d^{(1)}, \dots, d^{(s)}$ be a sequence of consecutive overdemand preserving price increases (line 6) such that the set S (line 4) is the same for all of them. Let $\lambda^{(1)}, \dots, \lambda^{(s)}$ be the corresponding scalars (line 7). Then for any pair of items $j \in S$ and $j' \in S$ we have $\left(\sum_{l=1}^s \lambda^{(l)} d_j^{(l)} \right) / \left(\sum_{l=1}^s \lambda^{(l)} d_{j'}^{(l)} \right) \leq (M/m)^{\min(n,k)}$. All line numbers are for the mechanism in Section 3.*

Proof. Lemma 15 ensures that for all $l \in \{1, \dots, s\}$: $d_j^{(l)} \leq (M/m)^{\min(n,k)} d_{j'}^{(l)}$. Hence, $\lambda^{(l)} d_j^{(l)} \leq \lambda^{(l)} (M/m)^{\min(n,k)} d_{j'}^{(l)}$ and the result follows by summing over all l . \square

Lemma 17. *From one case of 7(a), where a new bidder-item pair (i, j) enters the first choice graph, to the next, which can be separated by many cases of 7(b) and changes in the price increase d , the sets T and S do not change.*

Note that when both the conditions for 7(a) and 7(b) are satisfied we count this as a case of 7(a).

Proof. We will prove the claim by induction over the number of cases 7(b) separating two cases 7(a). Let b be this number. If $b = 0$ then the claim holds trivially. Denote the maximum alternating tree before the b -th update by $\mathcal{T}^{(b)}$. This tree is wrt to the matching $\mu^{(b)}$ and spans the items in S and the bidders in T and is computed in line 3 of the mechanism in Section 3. Computing the overdemand preserving price increase d in line 6 of the mechanism results in an implicitly computed new tree $\mathcal{T}'^{(b)}$, the tree corresponding to d , spanning the same items and bidders that corresponds to a new matching $\mu'^{(b)}$. Since d is envy free for $\mathcal{T}'^{(b)}$ and the utility functions are continuous all edges in $\mathcal{T}'^{(b)}$ (and in particular those in $\mu'^{(b)}$) belong to the first choice graph after the price update (hence no edges are removed from $\mu'^{(b)}$ in line 9 of the mechanism). It follows that the max alternating tree $\mathcal{T}^{(b+1)}$ that is explicitly computed in line 10 of the mechanism is identical to $\mathcal{T}^{(b)}$, which we know spans the same items and bidders as $\mathcal{T}^{(b)}$, namely S and T . \square

Q Proof of Lemma 12

Let X be the set of unmatched bidders under μ , i.e. $X = \{i \in I : \mu(i) = j_0\}$. We will show that there exists an envy free outcome $(\tilde{\mu}, \tilde{p})$ for $\tilde{u}_{i,j}(\cdot)$ that has prices not much higher than p and such that $X \subseteq \{i \in I : \tilde{\mu}(i) = j_0\}$. To ensure that bidders unmatched under μ are also unmatched under $\tilde{\mu}$ we initialize the prices p of the mechanism as $p_j = \max(p_j, \min\{\hat{p}_j : \forall i \in X : \tilde{u}_{i,j}(\hat{p}_j) \leq o_i\})$. If $u_{i,j}(p_j) \leq o_i$ then $\tilde{u}_{i,j}(p_j) \leq o_i + \epsilon$ and so $\tilde{u}_{i,j}(p_j + \epsilon/m) \leq o_i$. In other words, the initial prices satisfy $p_j \leq p_j + \epsilon/m$ for all j . With these initial prices, we know that bidders in X can all be matched to the dummy item j_0 in any envy free outcome. Hence, we can reduce the problem size and we will in the following only consider the remaining $\min(n, k)$ bidders.

Note that whenever we compute the sets T and S , we have $|T| > |S|$ by the definition of strict overdemand. And so there exists at least one bidder $\hat{i} \in T$ such that $\mu(\hat{i}) \notin S$, where μ is the matching in the (unknown) bidder optimal outcome (μ, p) . Regardless of whether \hat{i} turns out to be the bidder i of the (i, j) pair in the next case 7(a) or not, we will show that the required price updates for items in T to add the edge $(\hat{i}, \mu(\hat{i}))$ to the first choice graph are small. This will give us an upper bound on how much we have to increase the prices of any item in S until we reach case 7(a).

Let us write q for the number of cases of 7(a) the mechanism has encountered since the beginning of its execution. We write $p^{(q)}$ for the prices after q such cases and so $p_j^{(0)} = \max(p_j, \min\{\hat{p}_j : \forall i \in X : \tilde{u}_{i,j}(\hat{p}_j) \leq o_i\})$. Furthermore, we write $\delta^{(q)} = \max_j(p_j^{(q)} - p_j^{(0)})$.

Let $\hat{j} \in S$ be any of \hat{i} 's current first choice items before the price update, i.e. $\hat{j} \in F_{p^{(q)}}(\hat{i})$. Let us define the smallest utility gap for the current prices $p^{(q)}$ after q occurrences of case 7(a) as $\alpha^{(q)} = \min_{\hat{i} \in T: \mu(\hat{i}) \notin S} \alpha_{\hat{i}}^{(q)}$, where $\alpha_{\hat{i}}^{(q)} = \max(0, \tilde{u}_{\hat{i}, \hat{j}}(p_{\hat{j}}^{(q)}) - \tilde{u}_{\hat{i}, \mu(\hat{i})}(p_{\mu(\hat{i})}^{(q)}))$ with $\hat{i} \in T$ such that $\mu(\hat{i}) \notin S$. Lemma 17 shows that the sets T and S do not change between consecutive cases of 7(b) and so all of \hat{i} 's first choice items, which can change with each new case 7(b), remain in the set S until the next case 7(a). This means that if *all* items in S have their price increased by at least $\alpha^{(q)}/m$ then item $\mu(\hat{i}) \notin S$ will become a new first choice item for \hat{i} and so case 7(a) is reached. To ensure a price increase of at least $\alpha^{(q)}/m$ for all items in S , it is sufficient if *any* item in S has a price increase of $\alpha^{(q)}(M/m)^{\min(n,k)} \cdot 1/m$, using Lemma 16. So we know $\delta^{(q+1)} \leq \delta^{(q)} + \alpha^{(q)}(M/m)^k \cdot 1/m$. As, according to Lemma 17, the sets T and S do not change between consecutive cases of 7(b), we know that this increase can cover several such cases, even an exponential number, without affecting the argument above. We bound $\alpha^{(q)}$ as follows:

$$\begin{aligned}
\alpha^{(q)} &\leq \alpha_{\hat{i}}^{(q)} \leq \max_{\hat{j}} \left(\max \left(0, \tilde{u}_{\hat{i}, \hat{j}}(p_{\hat{j}}^{(q)}) - \tilde{u}_{\hat{i}, \mu(\hat{i})}(p_{\mu(\hat{i})}^{(q)}) \right) \right) \\
&\leq \max_{\hat{j}} \left(\max \left(0, \tilde{u}_{\hat{i}, \hat{j}}(p_{\hat{j}}) - \tilde{u}_{\hat{i}, \mu(\hat{i})}(p_{\mu(\hat{i})} + \delta^{(q)}) \right) \right) \\
&\leq \max_{\hat{j}} \left(\max \left(0, \tilde{u}_{\hat{i}, \hat{j}}(p_{\hat{j}}) - \tilde{u}_{\hat{i}, \mu(\hat{i})}(p_{\mu(\hat{i})}) \right) \right) + \delta^{(q)} M \\
&\leq \max_{\hat{j}} \left(\max \left(0, (u_{\hat{i}, \mu(\hat{i})}(p_{\mu(\hat{i})}) - \tilde{u}_{\hat{i}, \mu(\hat{i})}(p_{\mu(\hat{i})})) - (u_{\hat{i}, \hat{j}}(p_{\hat{j}}) - \tilde{u}_{\hat{i}, \hat{j}}(p_{\hat{j}})) \right. \right. \\
&\quad \left. \left. + (u_{\hat{i}, \hat{j}}(p_{\hat{j}}) - u_{\hat{i}, \mu(\hat{i})}(p_{\mu(\hat{i})})) \right) \right) + \delta^{(q)} M \\
&\leq 2\epsilon + \delta^{(q)} M.
\end{aligned}$$

Combining $\alpha^{(q)} \leq 2\epsilon + \delta^{(q)} M$ with $\delta^{(q+1)} \leq \delta^{(q)} + \alpha^{(q)}(M/m)^{\min(n,k)} \cdot 1/m$ gives $\delta^{(q+1)} \leq X \cdot \delta^{(q)} + Y$ where $X = 1 + (M/m)^{k+1}$ and $Y = 2\epsilon/m \cdot (M/m)^k$. Unrolling this then gives $\delta^{(q+1)} \leq X^2 \delta^{(q-1)} + XY + Y \leq \dots \leq X^{q+1} \delta^{(0)} + Y \cdot (X^{q+1} - 1)/(X - 1) = Y \cdot (X^{q+1} - 1)/(X - 1) \leq 2\epsilon/m \cdot (M/m)^{\min(n,k)} \cdot (X^{q+1})/(X - 1) \leq 2\epsilon/M \cdot (1 + (M/m)^{\min(n,k)+1})^{q+1}$.

So the maximum price increase with respect to the initial prices after q occurrences of case 7(a) is bounded above by $\delta^{(q)} \leq (1 + (M/m)^{\min(n,k)+1})^q \cdot 2\epsilon/M$. Now q is bounded above by the total number of cases of 7(a) that can occur. We claim that this number is bounded by $\min(n, k)^2$ and we argue as follows. The outer while loop matches a new bidder each time and gets executed $\min(n, k)$ times (as we have reduced the number of bidders from n to $\min(n, k)$). Note that in the continuous case bidders can never get unmatched. In each iteration of the inner while loop corresponding to 7(a) we add an edge to an item $j \notin S$ and the maximum alternating tree grows by one item and at most until all bidders are part of the tree. As there are only k items and n bidders this case can occur at most $\min(n, k)$ times for a particular bidder. As the item $j \notin S$ that gets added is either

unmatched, in which case the matching can be extended, or it is matched to a bidder, we can also bound $q \leq \min(n, k)^2$.

Given that the prices for the approximated input $\tilde{u}_{i,j}(\cdot)$ cannot be much higher than for the input $u_{i,j}(\cdot)$, we can bound the utility difference as follows:

$$\begin{aligned}
\tilde{u}_{i,\tilde{\mu}(i)}(\tilde{p}_{\tilde{\mu}(i)}) &\geq \tilde{u}_{i,\mu(i)}(\tilde{p}_{\mu(i)}) \\
&\geq \tilde{u}_{i,\mu(i)}(p_{\mu(i)} + \epsilon/m + 2\epsilon/M \cdot (1 + (M/m)^{\min(n,k)+1})^q) \\
&\geq \tilde{u}_{i,\mu(i)}(p_{\mu(i)}) - \epsilon M/m - \epsilon(2M/m)^{(\min(n,k)+1)(\min(n,k)^2)} \\
&\geq u_{i,\mu(i)}(p_{\mu(i)}) - \epsilon - \epsilon M/m - \epsilon(2M/m)^{(\min(n,k)+1)(\min(n,k)^2)} \\
&\geq u_{i,\mu(i)}(p_{\mu(i)}) - \epsilon(2M/m) - \epsilon(2M/m)^{(\min(n,k)+1)(\min(n,k)^2)} \\
&\geq u_{i,\mu(i)}(p_{\mu(i)}) - \epsilon(2M/m)^{(\min(n,k)+1)(\min(n,k)^2)} - \epsilon(2M/m)^{(\min(n,k)+1)(\min(n,k)^2)} \\
&\geq u_{i,\mu(i)}(p_{\mu(i)}) - \epsilon 2 \cdot (2M/m)^{(\min(n,k)+1)(\min(n,k)^2)} \\
&\geq u_{i,\mu(i)}(p_{\mu(i)}) - \epsilon(2M/m) \cdot (2M/m)^{(\min(n,k)+1)(\min(n,k)^2)} \\
&= u_{i,\mu(i)}(p_{\mu(i)}) - \epsilon \cdot (2M/m)^{1+(\min(n,k)+1)(\min(n,k)^2)} \\
&\geq u_{i,\mu(i)}(p_{\mu(i)}) - (2M/m)^{(\min(n,k)+1)^3} \cdot \epsilon.
\end{aligned}$$

Here we used the bound $2(1 + (M/m)^{\min(n,k)+1}) \leq 4(M/m)^{\min(n,k)+1} \leq (2M/m)^{\min(n,k)+1}$ for $\min(n, k) \geq 1$. We also used $M/m \geq 1$ and, at the last step, $(x+1)^3 = x^3 + 3x^2 + 3x + 1 \geq x^3 + x^2 + 1 = 1 + (x+1)x^2$ for $x > 0$. Note that exactly the same type of manipulations can be applied to show that $\tilde{p}_j \leq p_j + \epsilon/m + 2\epsilon/M \cdot (1 + (M/m)^{\min(n,k)+1})^q \leq p_j + (2M/m)^{(\min(n,k)+1)^3} \cdot \epsilon/M$.

R Proof of Theorem 3

We want to prove the existence of an envy free outcome $(\tilde{\mu}, \tilde{p})$ for input utility functions $\tilde{u}_{i,j}(\cdot)$, where the prices \tilde{p} are not much higher, at most $O(\epsilon)$, than the prices of a bidder optimal outcome (μ, p) for $u_{i,j}(\cdot)$. To prove this existence, Lemma 12 shows that if our mechanism is started with initial prices $p_j = \max(p_j, \min\{\hat{p}_j : \forall i \in \{i' : u_{i',\mu(i')}(\mu(i')) \leq o_{i'}\} : u_{i,j}(\hat{p}_j) \leq o_i\})$, instead of 0 for all items, it finds such an outcome which is envy free for $\tilde{u}_{i,j}(\cdot)$. This solution will generally not be bidder optimal for $\tilde{u}_{i,j}(\cdot)$ though, as the initial prices were non zero, but it proves the existence of some envy free solution for $\tilde{u}_{i,j}(\cdot)$ which is not much worse than the (unknown) bidder optimal outcome (μ, p) for $u_{i,j}(\cdot)$. If we now start the mechanism with initial prices set to 0 we will obtain a bidder optimal solution for $\tilde{u}_{i,j}(\cdot)$ which cannot be worse than the (unknown) solution $(\tilde{\mu}, \tilde{p})$. The details are as follows.

Let $\epsilon = \gamma(2M/m)^{-(\min(n,k)+1)^3}$. For each utility function $u_{i,j}(\cdot)$, find a piece-wise linear approximation $\tilde{u}_{i,j}(\cdot)$ with approximation guarantee ϵ . This generates $O(nk \cdot |\bar{p} -$

$r|\sqrt{B}(2M/m)^{(\min(n,k)+1)^3/2}/\sqrt{\gamma}$ linear segments by Lemma 11, where $|\bar{p}-r| = \max_{i,j} |\bar{p}_{i,j} - r_j|$. For this approximated input we then obtain a bidder optimal outcome $(\tilde{\mu}, \tilde{p})$ in time linear in the total number of segments times $\min(n, k)^3$ by Theorem 1, i.e., in time $O(nk \cdot |\bar{p} - r|\sqrt{B}(2M/m)^{(\min(n,k)+1)^3/2}\min(n, k)^3/\sqrt{\gamma})$. We claim that this outcome is then γ -bidder optimal for $u_{i,j}(\cdot)$.

We have $\tilde{p}_j \geq r_j$ and $\tilde{p}_{j_0} = 0$ as (i) we did not change the reserve prices and (ii) we did not approximate $u_{i,j_0}(p_{j_0})$. And so $(\tilde{\mu}, \tilde{p})$ is feasible for $u_{i,j}(\cdot)$. It is also γ -envy free as $\epsilon \leq \gamma/2$, and so we know that $\max_i(\max_j(u_{i,j}(\tilde{p}_j) - u_{i,\tilde{\mu}(i)}(\tilde{p}_{\tilde{\mu}(i)}))) \leq \max_i(\max_j(\tilde{u}_{i,j}(\tilde{p}_j) - \tilde{u}_{i,\tilde{\mu}(i)}(\tilde{p}_{\tilde{\mu}(i)})+2\epsilon)) \leq \gamma$. Finally, and most importantly, Lemma 12 shows that the utility losses with respect to the unknown bidder optimal outcome (μ, p) are bounded by γ .

S Proof of Theorem 4

Checking whether a given outcome (μ, p) is envy free and has sum of the utilities at least K takes polynomial time. This shows that the problem belongs to \mathcal{NP} . We prove \mathcal{NP} -completeness by providing a polynomial-time reduction from 3-SAT [16].

3-SAT: Given n Boolean variables x_k , m clauses c_r (= disjunction of up to three literals, i.e., x_k or $\neg x_k$), and a formula f (= conjunction of the m clauses) decide whether there exists a satisfying assignment x .

For a given instance of 3-SAT: Insert one bidder for every variable and one bidder for every clause. Insert two items, one labeled with TRUE the other with FALSE, for every variable and one item for every clause. All bidders derive a utility of zero from being matched to an item that corresponds to a variable. The bidders that correspond to a clause have a utility of one for the item corresponding to that clause iff the bidders that correspond to variables are matched to the items that correspond to variables in a way such that the labels of these items correspond to a satisfying assignment for the clause. Otherwise, they have a utility of zero. Choose $K = m$.

It is not difficult to see that the formula f is satisfiable iff there exists an envy free outcome with sum of the utilities equal to the number of clauses. This completes the proof.

T Proof of Theorem 5

We prove the claim for utility functions $u_{i,j}(p_1, \dots, p_k)$ that depend on the vector of prices and are decreasing in every component p_l of the price vector. The proof for utility functions $u_{i,j}(p_1, \dots, p_k)$ that are decreasing in p_j and increasing in p_l with $l \neq j$ is similar.

We can check whether a given outcome (μ, p) is envy free and has $\sum_i u_i \geq L$ in polynomial time. This establishes membership in \mathcal{NP} . To prove hardness we provide a polynomial-time reduction from MAX-2-NASH. MAX-2-NASH is \mathcal{NP} -complete even if

both players have the same number of strategies $|S_1| = |S_2| = N$ and all entries in the payoff matrices are integers between 0 and N [30, 15].

MAX-2-NASH: There are two players $i \in 1, 2$. For $i = 1$ we use $-i$ to denote player 2, and vice versa. Each player i has a finite set of *strategies* S_i and a *payoff function* $a_i(s, t)$ describing her payoff if he plays strategy $s \in S_i$ and the other player plays strategy $t \in S_{-i}$. A *mixed strategy* for player i is a probability distribution σ_i over S_i . A *strategy profile* $\sigma = (\sigma_1, \sigma_2)$ defines a mixed strategy for each of the two players. A strategy $s \in S_i$ is in the *support* of player i if and only if $\sigma_i(s) > 0$, i.e., player i plays strategy s with non-zero probability. A strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is a *mixed Nash equilibrium* (mixed NE) if for every player i and all mixed strategies $\sigma_i : \sum_{s \in S_i} (\sum_{t \in S_{-i}} a_i(s, t) \cdot \sigma_{-i}^*(t)) \cdot \sigma_i^*(s) \geq \sum_{s \in S_i} (\sum_{t \in S_{-i}} a_i(s, t) \cdot \sigma_{-i}^*(t)) \cdot \sigma_i(s)$. A pure strategy s is a *best response* if for all pure strategies $r : \sum_{t \in S_{-i}} a_i(s, t) \cdot \sigma_{-i}^*(t) \geq \sum_{t \in S_{-i}} a_i(r, t) \cdot \sigma_{-i}^*(t)$. It is well known that $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is a mixed NE iff every pure strategy in the support of σ^* is a best response. The goal is to compute a mixed NE $\sigma^* = (\sigma_1^*, \sigma_2^*)$ with maximum total payoff. The decision version of this problem is to decide whether there exists a mixed NE $\sigma^* = (\sigma_1^*, \sigma_2^*)$ with total payoff $\geq K$.

Claim: If all payoffs $a_i(s_i, s_{-i})$ in a 2-player game are non-negative integers, with the largest being N , then the probabilities in the mixed NE are rational numbers with the smallest being $1/(N! \cdot N^N)$.

Given the support of a mixed NE, the corresponding mixed NE $\sigma^* = (\sigma_1^*, \sigma_2^*)$ can be determined by solving a linear equation system (LES). Since the coefficients in this LES are non-negative integers, with the largest being N , any solution $\sigma^* = (\sigma_1^*, \sigma_2^*)$ to this LES consists of rational numbers with denominator at most $N! \cdot N^N$ [39]. Hence the smallest probability in the mixed NE is $\geq 1/(N! \cdot N^N)$.

Let $C = (N! \cdot N^N)^{2N}$. Note that C can be represented with $\leq 4N^2 \log(N)$ bits. Given an instance of MAX-2-NASH we construct an instance of our problem by means of the following ‘‘gadgets’’:

Best response gadget B. For each player i and strategy $s \in S_i$ one *best response gadget* B consisting of two bidders $is1$ and $is2$ and two items $is3$ and $is4$. Bidder $is1$ has a utility of $u_{is1, is3}(p) = \sum_{t \in S_{-i}} a_i(s, t) \cdot (1 - p_{-it3})$ for item $is3$ and a utility of $u_{is1, is4}(p) = \max_{t \in S_i} u_{it1, it3}(p)$ for item $is4$. Bidder $is2$ has a utility of $u_{is2, is4}(p) = 1/(8NC) - p_{is4}$ for item $is4$. Both bidders have an outside option of 0, i.e., $o_{is1} = o_{is2} = 0$. We will show below that for at least one strategy s that is a best response to the mixed strategy played by the other player $p_{is3} < 1$ and $p_{is4} = 0$, $is3$ is matched to $is1$, and $is2$ is matched to $is4$.

Probability gadget P1. For each player i and strategy $s \in S_i$ one *probability gadget* $P1$ consisting of two bidders isa and isb and two items isc and isd . Bidder isa has a utility of 0 for item isc , and a utility of $1 - p_{is3}$ for isd . Bidder isb has a utility of 1 for item isc , and a utility of $-\infty$ for isd . The outside options are $-\infty$. The purpose of this gadget is

to force $p_{is3} \leq 1$ or equivalently $1 - p_{is3} \geq 0$.

Probability gadget P2. For each player i and strategy $s \in S_i$ one *probability gadget P2* consisting of two bidders isa' and isb' and two items isc' and isd' . Bidder isa' has a utility of 0 for item isc' , and a utility of $1/(8NC) - p_{is4}$ for isd' . Bidder isb' has a utility of 1 for item isc' , and a utility of $-\infty$ for isd' . The outside options are $-\infty$. This gadget forces $p_{is4} \leq 1/(8NC)$.

Distribution gadget D1. For each player i one *distribution gadget D1* consisting of two bidders and one item. Both bidders have a utility of $[\sum_{s \in S_i} (1 - p_{is3}) * (1/(8NC) - p_{is4})] - 1/(8NC)$ for the item and an outside option of 0.

Distribution gadget D2. For each player i one *distribution gadget D2* consisting of two bidders and one item. Both bidders have a utility of $[\sum_{s \in S_i} (1 - p_{is3})] - 1$ for the item and an outside option of 0.

Claim: There exists a mixed NE $\sigma^ = (\sigma_1^*, \sigma_2^*) = ((1 - p_{1s3})_{s \in S_1}, (1 - p_{2s3})_{s \in S_2})$ with total payoff K iff there exists a feasible and envy free outcome (μ, p) with sum of the utilities strictly larger than $N \cdot (K + 4) + 2$, and less than $N \cdot (K + 4) + 2 + 1/C$.*

\Rightarrow : Assume such a mixed NE $\sigma^* = (\sigma_1^*, \sigma_2^*)$ exists. Denote the support of player i by $SUPP_i$. For each player i and strategy s : If s is in the support of player i , then match $is1$ to $is3$ and $is2$ to $is4$ and set $p_{is3} = 1 - \sigma_i^*(s) < 1$ and $p_{is4} = 0$. If it is not, then match $is1$ to $is4$ and leave $is2$ unmatched and set $p_{s3} = 1 - \sigma_i^*(s) = 1$ and $p_{is4} = 1/8NC$. In the former case this is envy free because s is one of i 's best responses and so $u_{is1, is3}(p) \geq u_{is1, is4}(p)$ and at $p_{is4} = 0$ we have $u_{is2, is4} \geq 0$. In the latter case this is envy free because $u_{is1, is4}(p) \geq u_{is1, is3}(p)$ and at $p_{is4} = 1/(8NC)$ we have $u_{is2, is4}(p) = 0$. Since $0 \leq p_{is3} \leq 1$ for all $is3$ items and $0 \leq p_{is4} \leq 1/(8NC)$ for all $is4$ items, we can match the bidders in the $P1$ and $P2$ gadgets in an envy free manner. Since (a) $\sum_{s \in S_i} (1 - p_{is3}) = 1$ and (b) $(1 - p_{is3}) = 0$ iff $(1/(8NC) - p_{is4}) = 0$, this is also true for the bidders in the $D1$ and $D2$ gadgets. The $is1$ vertices contribute a utility of exactly $N \cdot K$, namely K_1 each for the ones of player 1 and K_2 each for the ones of player 2 with $K_1 + K_2 = K$. The $is2$ vertices contribute exactly $(|SUPP_1| + |SUPP_2|)/(8NC)$. The $P1$ gadgets contribute exactly $2 \cdot N + 2$. The $P2$ gadgets contribute exactly $2 \cdot N + (|SUPP_1| + |SUPP_2|)/(8NC)$. The distribution gadgets contribute nothing. Hence the total utility is $N \cdot (K + 4) + 2 + (|SUPP_1| + |SUPP_2|)/(4NC)$. The claim follows from the fact that $1 \leq |SUPP_i| \leq N$ for both i and, thus, $2 \leq |SUPP_1| + |SUPP_2| \leq 2N$.

\Leftarrow : Assume the total utility equals L . Let $T_1 = \{s \in S_i : is1 \text{ is matched to } is3\}$ and let $T_2 = \{s \in S_2 : is1 \text{ is matched to } is3\}$. Let K_1 denote the utility of each of the $1s1$ players and let K_2 denote the utility of each of the $2s1$ players. Let $K = K_1 + K_2$. It follows that $L = N \cdot (K + 4) + 2 + (|T_1| + |T_2|)/(8NC)$. Since $|T_1| \leq N$ and $|T_2| \leq N$, this is at most $N \cdot (K + 4) + 2 + 1/C$. We show below in Step 1 that for each player i if $p_{is3} < 1$ then $is1$ is matched to $is3$. Due to the distribution gadget $D2$, $\sum_{s \in S_i} (1 - p_{is3}) = 1$. Thus

there is at least one strategy s s.t. $1 - p_{is3} > 0$, i.e., $1 - p_{is3} < 1$. It follows that for each player there is at least one $is1$ node that is matched to an $is3$ node. For this strategy the utility is maximized if the $is2$ node is matched to the $is4$ node at a price of 0. Thus, $L > N \cdot (K + 4) + 2$.

We still have to show that $\sigma^* = (\sigma_1^*, \sigma_2^*)$ with $\sigma^* = ((1 - p_{1s3})_{s \in S_1}, (1 - p_{2s3})_{s \in S_2})$ is a mixed NE, i.e., for each player i every strategy s with $1 - p_{is3} > 0$ is a best response to $\sigma_{-i}^* = (1 - p_{-is3})_{i \in S_{-i}}$, and the payoff is K .

Claim A. For $i \in \{1, 2\}$ the following two implications are equivalent: (1) If for a strategy $s : p_{is3} < 1$, then $is1$ is matched to $is3$. (2) If for a strategy s $is1$ is matched to $is4$, then $p_{is3} = 1$.

If $is1$ is matched to $is4$ then $p_{is4} \geq 1/(8NC)$ because otherwise $is4$ is not envy free. With $P2$ we know that $p_{is4} = 1/(8NC)$. Additionally, since $is4$ tries to maximize its utility, if $is1$ is matched to $is3$, then $p_{is4} = 0$. Now $D1$ guarantees: $\sum_{s \in S_i, is1} is1$ is matched to $is3(1 - p_{is3}) = 1$. And $D2$ guarantees: $\sum_{s \in S_i} (1 - p_{is3}) = 1$. We get $\sum_{s \in S_i, is1} is1$ matched to $is4(1 - p_{is3}) = 0$ (*). Since by $P1$ we know that $1 - p_{is3} \geq 0$ for all i and s , (*) is only possible if $(1 - p_{is3}) = 0$ for all $s \in S_i$ with $is1$ matched to $is4$, said differently if $is1$ is matched to $is4$, then $p_{is3} = 1$.

We now know: For $i = 1, 2$: If for a strategy $s : p_{is3} < 1$, then $is1$ is matched to $is3$. We also know that $p_{is3} < 1$ iff $1 - p_{is3} = \sigma_i^*(s) > 0$ iff s is in the support of the mixed NE. So we have: For $i = 1, 2$: If a strategy s is in the support of the mixed NE, then $is1$ is matched to $is3$.

Claim B. $s1$ is matched to $is3$ iff s is a best response to $\sigma_{-i}^* = (1 - p_{-is3})_{s \in S_{-i}}$.

\Rightarrow : Let $u_i(s, \sigma_{-i}^*) = \sum_{t \in S_{-i}} a_i(s, t) \cdot \sigma_{-i}^*(t)$. If $is1$ is matched to $is3$ then $u_i(s, \sigma_{-i}^*) + 1/(8NC) \geq \max_{s'} u_i(s', \sigma_{-i}^*)$, i.e., $\max_{s'} u_i(s', \sigma_{-i}^*) - u_i(s, \sigma_{-i}^*) \leq 1/(8NC)$. We also have $\max_{s'} u_i(s', \sigma_{-i}^*) - u_i(s, \sigma_{-i}^*) \geq 0$. Since all u_i values are sums of at most N rational numbers with denominator $\leq N! \cdot N^N$, they are themselves rational numbers with denominator $\leq (N! \cdot N^N)^N$. Hence their difference is either 0 or $\geq 1/((N! \cdot N^N)^N)^2 = 1/C$. Since $1/C > 1/(8NC)$ this shows that $\max_{s'} u_i(s', \sigma_{-i}^*) - u_i(s, \sigma_{-i}^*) = 0$, i.e., s is a best response to $\sigma_{-i}^* = (1 - p_{-is3})_{s \in S_{-i}}$.

\Leftarrow : If s is a best response to $\sigma_{-i}^* = (1 - p_{-is3})_{s \in S_{-i}}$ then both edges of $is1$ have the same weight. To maximize the utility of the whole gadget $is1$ will be matched to $is3$ and $is2$ to $is4$, giving $is2$ a utility of $1/(8nC)$.

Claim C. The total payoff of the mixed NE $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is K .

The payoff to player i in the mixed NE $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is $u_i(s, \sigma_{-i}^*) = \sum_{t \in S_{-i}} a_{s,t}(s, t) \sigma_{-i}^*(t)$, where s is some strategy in i 's support. For every such strategy, we have $\sigma_{-i}^*(t) = (1 - p_{-it3})$ and, thus, $u_i(s, \sigma_{-i}^*) = \sum_{t \in S_{-i}} a_{s,t}(s, t)(1 - p_{-it3}) = K_i$. Since this is true for $i \in \{1, 2\}$, the total payoff is $K_1 + K_2 = K$.