

An Expressive Mechanism for Auctions on the Web ^{*}

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

Monika Henzinger
University of Vienna
Universitätsstrasse 10/11
A-1090 Vienna, Austria
monika.henzinger@univie.ac.at

Ingmar Weber
Yahoo! Research
Av. Diagonal 177 80
E-08018 Barcelona, Spain
ingmar@yahoo-inc.com

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.

Categories and Subject Descriptors

F.2.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity—*Nonnumerical Algorithms and Problems*

General Terms

Algorithms, Economics, Theory

Keywords

VCG, GSP, General Auction Mechanism, Expressiveness, Envy Freeness, Bidder Optimality, Competitive Equilibrium

^{*}A full version of this paper with all proofs is available from: <http://infoscience.epfl.ch/record/153929>

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.
WWW 2011, March 28–April 1, 2011, Hyderabad, India.
ACM 978-1-4503-0632-4/11/03.

1. INTRODUCTION

Auctions are widely used on the Web. They are, e.g., used by Google, Microsoft, and Yahoo! for search advertising [36]. They are also used on platforms such as eBay for selling a broad variety of goods and services [41, 8]. In most of these applications the auctions in use are single/multi-item auctions with unit demand. The problem solved by 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*) [40].

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 Clarke Groves (VCG) [45, 14, 31] and Generalized Second Price (GSP) [26, 43] 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 [37].

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, 43] or [38, 27]), but mechanisms that can handle both types simultaneously are still rare [30]. 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 consid-

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

ered 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 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*) [40].

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 [42]. 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 [37] examined the incentives for misreporting and found that the bidder optimal outcome is identical to the outcome of VCG [45, 14, 31] 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 [35]. 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 established for related problems (see, e.g., [22, 44]), 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 poly-

nomial time. Similar results to that of Aggarwal et al. [1] were obtained by [5, 6] and [25]. In [34] 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, 34], just a *single discontinuity* [1, 6, 25, 34, 12], or both. Also note that the piece-wise 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. [33], 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 [28]. 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 [32] that was already used in [20] and in [35]. 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

³www.iab.net/lead_generation

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 $u_{i,j}(\cdot)$.

²For $n > k$ it is linear in n .

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 unmatched previously matched bidders. Note that neither [3, 4] nor [1] have to unmatched 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 depend on the whole vector of prices through a reduction from MAX-2-NASH [29, 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].

A full version of this paper with all proofs is available from: <http://infoscience.epfl.ch/record/153929>

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 $(i, j) \in I \times J$,

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

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 overdemand using Hall’s Theorem [32]. Our mechanism, which can be seen as an algorithmic version of Hall’s Theorem [32], starts with prices all zero and iteratively raises the prices of strictly overdemand 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

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 .

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 overdemand* 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 overdemand*, if it is strictly overdemand 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 overdemand. 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 [32] and the fact that a strictly overdemand set of items exists *if and only if* an overdemand 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 overdemand at prices p .*

Alternating paths and trees. To identify strictly overdemand 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 overdemand with respect to the bidders T in the tree and $|T| = |S| + 1$.*

Envy free price increase. Once we have identified a strictly overdemand 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 [35]. The duality between slopes and utility drops exploited here is reminiscent of the duality between value-maximizing matchings and envy free prices in [42]. (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 let $\sigma = \{(i, j) \in T \setminus \{i_0\} \times S \mid x_{i,j} = 1\}$

$$\min \sum_{i,j} x_{i,j} \cdot \log(c_{i,j})$$
sb
$$\sum_{j \in F_p(i)} x_{i,j} = 1 \quad (\forall i \in T \setminus \{i_0\})$$

$$\sum_{i \in F_p(j)} x_{i,j} = 1 \quad (\forall j \in S)$$

$$x_{i,j} \geq 0 \quad (\forall (i, j) \in F_p \cap (T \setminus \{i_0\}) \times S)$$
- 2 compute ω, ρ as optimal solution to the following DP
$$\max \sum_i \omega_i + \sum_j \rho_j$$
sb
$$\omega_i + \rho_j \leq \log(c_{i,j}) \quad (\forall (i, j) \in F_p \cap (T \setminus \{i_0\}) \times S)$$

- 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 subroutine 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 $j \in \{1, 2\}$.

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.

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

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

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

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

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 [39], 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/>).

8. 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. Feldman, M. Henzinger, N. Korula, V. Mirrokni, and C. Stein. Online stochastic packing applied to display ad allocation. *ESA*, pages 1–13, 2010.
- [28] A. Ghosh and A. Sayedi. Expressive auctions for externalities in online advertising. *WWW*, pages 371–380, 2010.
- [29] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- [30] A. Goel and K. Munagala. Hybrid keyword search auctions. *WWW*, pages 221–230, 2009.
- [31] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [32] P. Hall. On representatives of subsets. *London Mathematical Society*, 10:26–30, 1935.
- [33] P. Jehiel, B. Moldovanu, and E. Stacchetti. How (not) to sell nuclear weapons. *American Economic Review*, 86(4):814–29, 1996.
- [34] D. Kempe, A. Mu’alem, and M. Salek. Envy-free allocations for budgeted bidders. *WINE*, pages 537–544, 2009.
- [35] H. W. Kuhn. The Hungarian Method for the assignment problem. *Naval Research Logistics*, 2:83–97, 1955.
- [36] S. Lahaie, D. M. Pennock, A. Saberi, and R. V. Vohra. Sponsored search auctions. *Algorithmic Game Theory*, pages 699–716, 2007.
- [37] H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Political Economy*, 91:461–479, 1983.
- [38] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized online matching. *Journal of the ACM*, 54(5):1–19, 2007.
- [39] P. Milgrom. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior*, 70:62–70, 2010.
- [40] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [41] 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.
- [42] L. S. Shapley and M. Shubik. The assignment game: The core I. *Game Theory*, 29:111–130, 1971.
- [43] H. Varian. Position auctions. *Industrial Organization*, 25(6):1163–1178, 2007.
- [44] V. Vazirani and M. Yannakakis. Ppad-completeness of fisher markets under piecewise-linear, concave utility functions. *ICTS*, pages 156–166, 2010.
- [45] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Finance*, 16:8–27, 1961.