

An Expressive Mechanism for Auctions on the Web

PAUL DÜTTING, London School of Economics and Political Science
 MONIKA HENZINGER, University of Vienna
 INGMAR WEBER, Qatar Computing Research Institute

Auctions are widely used on the Web. Applications range from sponsored search to platforms such as eBay. In these and in many other 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 Aggarwal et al. [2009] takes a first step toward 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 *piecewise linear* utility functions with *nonidentical slopes* and *multiple discontinuities*. Our mechanism runs in polynomial time. Like GAM it is incentive compatible for inputs that fulfill a certain nondegeneracy assumption, but our requirement is more general than the requirement of GAM. For discontinuous utility functions that are nondegenerate 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 piecewise 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

Additional Key Words and Phrases: VCG, GSP, general auction mechanism, expressiveness, envy freeness, bidder optimality, competitive equilibrium

ACM Reference Format:

Paul Dütting, Monika Henzinger, and Ingmar Weber. 2015. An expressive mechanism for auctions on the Web. *ACM Trans. Econ. Comput.* 4, 1, Article 1 (November 2015), 34 pages.
 DOI: <http://dx.doi.org/10.1145/2716312>

1. INTRODUCTION

Auctions are widely used on the Web. They are, for example, used by Google and Microsoft for selling sponsored search results and on platforms such as eBay for selling

A preliminary version of this work appeared in the *Proceedings of the 20th International Conference on World Wide Web (WWW 2011)*, Hyderabad, India, March 28–April 1, 2011.

The research leading to these results has received funding from the Vienna Science and Technology Fund WWTF grant ICT10-002 and from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement n. [340506].

Authors' addresses: P. Dütting, Department of Mathematics, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, United Kingdom; email: p.duetting@lse.ac.uk; M. Henzinger, Faculty of Computer Science, University of Vienna, Währinger Straße 29, A-1090 Vienna, Austria; email: monika.henzinger@univie.ac.at; I. Weber, Qatar Computing Research Institute, P.O. Box 5825, Doha, Qatar; email: iweber@qf.org.qa.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2015 ACM 2167-8375/2015/11-ART1 \$15.00

DOI: <http://dx.doi.org/10.1145/2716312>

a broad variety of goods and services. In these and in many other 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 i 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 (weakly) prefers the item that he is currently matched to over every other item.¹ An envy-free outcome (μ, p) is said to be *bidder optimal* if the utility of every bidder is at least as high as in every other envy-free outcome (μ', p') . From an economic point of view an envy-free outcome in which the price of every unsold item is identical to the item's reserve price is desirable because it represents a *competitive equilibrium* (or *Walrasian equilibrium*) [Nisan et al. 2007].² Under the additional requirement on the prices of unsold items the bidder optimal, envy-free outcome is the competitive equilibrium with maximum payoffs to the bidders. Maximizing the payoffs to the bidders is also a reasonable objective for the auctioneer as it guarantees participation and thus revenue in the long run.

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) [Vickrey 1961; Clarke 1971; Groves 1973] and Generalized Second Price (GSP) [Edelman et al. 2007; Varian 2007] for multi-item auctions, nicely fit into the preceding model. For linear utilities of the form $u_{i,j}(p_j) = v_{i,j} - p_j$, where $v_{i,j}$ denotes a bidder-item-dependent valuation, the VCG mechanism finds a bidder optimal, envy-free outcome [Leonard 1983]. (Locally) envy-free equilibria of the GSP mechanism for linear utilities of the form $u_{i,j}(p_j) = \alpha_j \cdot v_i - p_j$, where α_j is an item-dependent constant and v_i is a bidder-dependent valuation, are discussed in Edelman et al. [2007] and Varian [2007].

A main drawback of these mechanisms, however, is the *limited expressiveness* that they offer to the bidders: (1) They typically restrict the utilities to be *linear* in the price, with *identical slopes* across the items. (2) They often require the utilities to be *continuous* in the price. The former is a limitation in settings where the preferences naturally lead to utilities that are nonlinear in the price, such as sponsored search auctions in which part of the bidders are interested in clicks and others are interested in impressions (see, e.g., Goel and Munagala [2009]). It is also a limitation in settings where bidders have *soft budget constraints* (i.e., budget constraints that can be circumvented by taking out a loan). The continuity requirement in turn limits the bidders ability to express *hard budget constraints*. Budget constraints are considered an important aspect in practice and have recently received a lot of attention (see, e.g., Borgs et al. [2005], Dobzinski et al. [2012], Dütting et al. [2012], and Colini-Baldeschi et al. [2012]).

The General Auction Mechanism (GAM) of Aggarwal et al. [2009] takes a first step toward addressing the problem of limited expressiveness. It applies to *linear* utilities with *identical slopes* and a *single discontinuity* per bidder-item pair. More specifically, the utilities 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 $m_{i,j}$ is a so-called bidder-item-dependent *maximum price*. GAM requires the input to be in *general position* (see Aggarwal et al. [2009] and Section 4 for a generalization of this concept). For inputs in general position it also preserves a desirable property of the original model. Namely, no bidder can misreport his valuations and/or maximum prices to achieve a higher utility. This makes it a (weakly) dominant

¹In the related literature (see, e.g., Aggarwal et al. [2009]) such outcomes are also referred to as *stable*.

²The additional requirement regarding the prices of unsold items is needed to ensure that the auctioneer prefers not to sell these items.

strategy for the bidders to report truthfully. Mechanisms with this property are said to be *incentive compatible* (or *strategy proof*) [Nisan et al. 2007].

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 nonlinear 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 nonidentical slopes.) Consider an ad auction with bidders with per-click valuations $v_{i,j}^{click}$ and bidders with per-impression valuations $v_{i,j}^{imp}$. The former are envy free if $u_i^{click} \geq v_{i,j}^{click} - p_j^{click}$ for all j and the latter are envy free if $u_i^{imp} \geq v_{i,j}^{imp} - p_j^{imp}$ for all j . 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 nonidentical 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, that is, the preceding problem does *not* arise.

Example 2: Soft and hard budgets. (This example motivates piecewise linear utilities with nonidentical slopes and multiple discontinuities.) Suppose that bidder i wants to buy a car on eBay. In the current system it would be dangerous for bidder 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 bidder 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, bidder i 's valuation $v_{i,j}$ for item j may exceed the amount c of cash that he possesses. In this case, bidder i might be willing to take out a loan if the price p_j of item j exceeds c . Assume that bidder i is offered a loan with a




Fig. 1. Bidder i 's utility function for item j .

maximum amount of a , an interest rate of r , and a fixed fee of f . Then bidder 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) Afterward, 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 appears as shown in Figure 1. 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, nonlinear 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 superlinearly 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 *piecewise linear* utility functions with *nonidentical slopes* and *multiple discontinuities*. The problem solved by our mechanism is *more difficult* than that solved by GAM as nonidentical 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 .³ 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 outcome in this setting is *incentive compatible*. This characterization involves a generalization of the general position concept of Aggarwal et al. [2009]. For discontinuous utility functions that are nondegenerate as well as for continuous utility

³For $n > k$ it is linear in n .

functions the outcome of our mechanism is a *competitive equilibrium*. (3) We show how to use our mechanism to compute a γ -bidder optimal, γ -envy-free outcome for a general class of continuous utility functions via piecewise 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

Expressiveness. On a conceptual level our work is related to work that seeks to measure expressiveness and design expressive mechanisms for concrete applications.

Benisch et al. [2008] proposed a measure for quantifying the expressiveness of a mechanism and showed that (an upper bound on) the expected efficiency of the most efficient Nash equilibrium of any mechanism is strictly increasing in the mechanism's expressiveness. In a subsequent paper, Benisch et al. [2009] presented a method for comparing mechanisms with different degrees and forms of expressiveness. They applied this method to sponsored search auctions and found that the GSP mechanism can be arbitrarily inefficient, and proposed a more expressive version of the GSP mechanism that overcomes the shortcomings of the original mechanism.

Also working in the context of internet advertising Parkes and Sandholm proposed an optimize-and-dispatch architecture consisting of an (offline) optimizer module and an (online) dispatch module that allows for expressive ad auctions [Parkes and Sandholm 2005]. Boutilier et al. [2008], following the optimize-and-dispatch architecture, described different techniques for the offline optimization and the online dispatching phase for the special case of banner ad auctions. Walsh et al. [2010] presented a method to circumvent the exponential nature of the offline optimization problem, while pertaining approximate optimality of the solution. Finally, Ghosh et al. [2009] proposed a solution for the problem faced by a publisher of display ads that can sell his ads via guaranteed contracts and a spot market.

Expressive mechanisms have also been proposed for settings other than advertising. Sandholm [2007], for example, proposed the “expressive commerce” paradigm that combines the advantages of human negotiation with the advantages of electronic reverse auctions and demonstrated its effectiveness by applying it to the sourcing problem. Another concrete example is work by Conitzer and Sandholm [2011], who considered the problem of expressing matching offers, that is, offers to donate when other parties donate as well, in the context of charities.

Matching and pricing. On a technical level, our work is most closely related to work on matching and pricing problems with different utility functions.

Linear utility functions with *identical* slopes were studied by Shapley and Shubik [1971]. 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 [1983] examined the incentives for misreporting the utility functions and found that the bidder optimal outcome is identical to the outcome of VCG [Vickrey 1961; Clarke 1971; Groves 1973] and therefore incentive compatible. The classic mechanism for linear utility functions with identical slopes is the so-called Multiitem Auction of Demange et al. [1986], which is a variant of the so-called Hungarian Method by Kuhn [1955]. 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, *piecewise linear* utility functions with *nonidentical slopes* in Alkan [1989, 1992], which 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. Demange and Gale [1985] established the existence of a bidder optimal outcome for general, *nonlinear* utility functions 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. [2011] 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., Deng and Du [2008] and Vazirani and Yannakakis [2010]), it is not clear whether or under which conditions a bidder optimal outcome can be found efficiently for such general continuous utility functions.

The first to add a *single discontinuity* to otherwise *linear* utility functions with *identical slopes* were Aggarwal et al. [2009]. 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. [2009] were obtained by Ashlagi et al. [2009], Ashlagi et al. [2010], and Dütting et al. [2013b]. In Kempe et al. [2009] it was shown how to find the smallest envy-free prices for a *given* matching. Recently, Chen et al. [2010] gave a polynomial-time mechanism for *consistent* utility functions. Note that all these results either assume *identical slopes* [Aggarwal et al. 2009; Ashlagi et al. 2010; Dütting et al. 2013b; Kempe et al. 2009], just a *single discontinuity* [Aggarwal et al. 2009; Ashlagi et al. 2010; Dütting et al. 2013b; Kempe et al. 2009; Chen et al. 2010], or both. Also note that the piecewise linear utility functions with nonidentical slopes and multiple discontinuities that we study here are *not* consistent. The existence of a bidder optimal, envy-free outcome for more general, *nonlinear* utility functions with *multiple discontinuities* was established in Dütting et al. [2013a] using a lattice-theoretic argument, but just as in the continuous case no polynomial-time mechanism is known for such general discontinuous utility functions.

Externalities. Our hardness results for even more expressive settings 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 for settings with externalities is due to Jehiel et al. [1996], 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 Ashlagi et al. [2008] and Ghosh and Sayedi [2010]. 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, for example, arises in the pay-per lead model, is \mathcal{NP} -hard and hard to approximate: 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 [Hall 1935] that was already used in Demange et al. [1986] and in Kuhn

[1955]. It also borrows from Alkan [1989, 1992], in which the computation of price increases for continuous, piecewise linear utility functions with nonidentical slopes is formulated as a primal/dual problem, although we refine this approach to significantly improve upon the running time and extend it to discontinuous utilities.⁴ Finally, our analysis follows a similar approach as Aggarwal et al. [2009] 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 items to be unmatched. We achieve this through a refinement of Alkan's technique for increasing the prices when slopes are nonidentical (Lemma 3.8) 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 3.7). Together these lemmata allow us to prove bidder optimality even for inputs that make it necessary to unmatched previously matched items. Note that neither in Alkan [1989, 1992] nor in Aggarwal et al. [2009] do items have to be unmatched, 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 items never get unmatched (Lemma 4.2). 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 4.2). We then use a variant of Hwang's Lemma (Lemma 4.3) to prove that whenever (a) and (b) are satisfied any mechanism that computes a bidder optimal outcome is incentive compatible. 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. [2009], we say that inputs that satisfy this condition are in *general position*, but our condition on the input is more general than that of Aggarwal et al. [2009] as it applies to piecewise linear utility functions with nonidentical slopes and multiple discontinuities.

Concerning general continuous utility functions it is intuitively obvious that the solution to a close enough piecewise 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 5.1). The hard part is to show that the distance γ between the resulting bidder optimal utilities scales linearly in ϵ (Lemma 5.2).

We establish the hardness result for utility functions that depend on the matching through a reduction from 3-SAT [Cook 1971] and the hardness result for utility functions that depend on the whole vector of prices through a reduction from MAX-2-NASH [Gilboa and Zemel 1989; Conitzer and Sandholm 2008].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 [Daskalakis and Papadimitriou 2005; Chen and Deng 2006].

⁴In Alkan [1989, 1992] the running time is stated as $O(n^2 \cdot k^4 \cdot \prod_{i,j} t_{i,j})$, where $t_{i,j}$ is the number of linear segments of $u_{i,j}(\cdot)$.

⁵MAX-2-NASH: Find a Nash equilibrium of a two-player game that maximizes the sum of the utilities.

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 nondummy 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 *piecewise 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 nondummy 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 .⁷ 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 article 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 FOR PIECEWISE LINEAR UTILITIES

In this section, we describe and analyze our mechanism for piecewise 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 [Hall 1935]. Our mechanism, which can be seen as an ascending price auction, starts with prices all zero and iteratively raises

⁶This definition together with the requirement that in every feasible outcome the price of the dummy item j_0 is $p_{j_0} = 0$ ensures that in every envy-free outcome every bidder i has utility at least o_i .

⁷Representing the fact that a bidder is unmatched through a match to a dummy item unifies the treatment of matched and unmatched bidders. Using a single dummy item rather than one dummy item per bidder has advantages in terms of computational complexity.

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 without loss of generality 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 Alkan [1989]. 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., Demange and Gale [1985]), 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 3.1. (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.*

PROOF. 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, that is, 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): By contradiction assume that there is an envy-free outcome (μ', p') with $u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu^*(i)}(p^*_{\mu^*(i)})$ for some bidder i . Since (μ^*, p^*) is envy free, we have $u_{i,\mu^*(i)}(p^*_{\mu^*(i)}) \geq u_{i,\mu'(i)}(p^*_{\mu^*(i)})$. 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! \square

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. For each bidder $i \in I$ we use $F_p(i) = \{j : \exists (i, j) \in F_p\}$ to denote the items demanded by this bidder at prices p and for each item $j \in J$ we use $F_p(j) = \{i : \exists (i, j) \in F_p\}$ to denote the bidders demanding this item at prices p . Analogously, for every subset of bidders $T \subseteq I$ we use $F_p(T) = \cup_{i \in T} F_p(i)$ to denote the set of items demanded by bidders in T at prices p and for every subset of items $S \subseteq J$ we use $F_p(S) = \cup_{j \in S} F_p(j)$ to denote the set of bidders demanding items in S at prices p . We say that a set of nondummy 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) each of the bidders in T demands only items in S , that is, $F_p(T) \subseteq S$, and (b) for every nonempty subset R of the items in S there

is more demand from bidders in T than there are items, that is, $|F_p(R) \cap T| > |R|$. We say that S is strictly overdemanded at prices p if it is strictly overdemanded at prices p with respect to some set of bidders T . Our definition of strict overdemand is *stronger* than the definition of *overdemand* [Demange et al. 1986], 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 definition of *minimal overdemand* [Demange et al. 1986], which requires that the set itself but no proper subset is overdemanded. It also differs from the definition of *directional overdemand* in Alkan [1992], which is identical to our definition except that it applies to the first choice graph *after* the envisioned price update. The advantage of our definition will become clear in the next subsection. The following lemma is established using Hall's Theorem [Hall 1935] and the fact that a strictly overdemanded set of items exists *if and only if* an overdemanded set of items exists.

LEMMA 3.2. *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 .*

PROOF. (1) \Leftrightarrow (2): An outcome (μ, p) is envy free if and only if (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 $(i, j) \in I \times J$ if and only if 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 [Hall 1935] there exists a matching μ in the first choice graph at prices p with $p_{j_0} = 0$ and $p_j \geq 0$ for $j \neq j_0$ if and only if $\forall T \subseteq I: |F_p(T)| \geq |T|$ or $j_0 \in F_p(T)$.

(2) \Rightarrow (3): 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' , that is, (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!

(2) \Leftarrow (3): 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 [Hall 1935]. 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 $|F_p(T'')| > |\mu'(T'')| = |\mu'(T'' \setminus \{i\})| = |T'' \setminus \{i\}| = |T''| - 1$, that is, $|F_p(T'')| \geq |T''|$. Hence, Lemma 3.3 shows that S' is strictly overdemanded with respect to T' . Contradiction! \square

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 nondummy 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 that 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., Demange et al. [1986]). The new insight of the following lemma is that there is a close correspondence between maximal alternating trees and our definition of strict overdemand.

LEMMA 3.3. *For any maximal alternating tree \mathcal{T} with respect to μ with root i_0 in G_p :*
 (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 ; j , respectively. (2) *If all items S in \mathcal{T} are nondummy items and matched, then (a) S is strictly overdemanded with respect to the bidders T in the tree and (b) $|T| = |S| + 1$.**

PROOF. To (1): The path P is the path in the maximal alternating tree \mathcal{T} that leads from i_0 to j_0 ; j , respectively. 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)(a): 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 $F_p(R) \cap T$.

To (2)(b): By maximality of the alternating tree \mathcal{T} the items that the bidders in T are matched to must be contained in S . The fact that all items in S are nondummy items and matched implies that the only unmatched bidder in T is i_0 . Hence, $|T| = |T \setminus i_0| + 1 = |S| + 1$. \square

Envy-free price increase. Once we have identified a strictly overdemanded set of items we need to determine how to increase the prices of the items 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 it satisfies the following conditions:

- (a) The entry d_j is strictly larger than zero for items j that are the end point of at least one edge in E and it is zero otherwise.
- (b) At prices $p + \lambda \cdot d$, for small enough $\lambda > 0$, each bidder i that is the end point of at least one edge in E weakly prefers each item j to which he has an edge in E to each item k that is currently among his first choice items.

Note that the dummy item j_0 is *not* the end point of an edge in E and so (a) implies that $d_{j_0} = 0$. Also note that it is sufficient to require (b) for all items that are currently among the first choice items of bidder i as bidder i will prefer each item j to which he has an edge in E to every item that was *not* among his first choice items at prices $p + \lambda \cdot d$ as long as $\lambda > 0$ is small enough. Our definition of envy-free price increase is similar to the definition of a *competitive direction* for continuous utility functions in Alkan [1989]. The next two lemmata are proved in Alkan [1989] 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 $E \subseteq F_p \setminus (I \times \{j_0\})$. It shows that a price increase d is envy free for a first choice edge $(i, j) \in F_p$ with $j \neq j_0$ 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 3.4. *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 $\lambda > 0$.*

PROOF. 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 $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ 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 $(i, k) \in I \times J$ and all $\lambda > 0$ sufficiently small. Since $d_j \geq 0$ for all j such that 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 $j \neq j_0$. \square

LEMMA 3.5. *A price increase $d \neq 0$ is envy free for the set of first choice edges $E \subseteq F_p \setminus (I \times \{j_0\})$ at prices p if and only if $c_{i,j} \cdot d_j \leq c_{i,k} \cdot d_k$ for all $(i, j) \in E$ and all $(i, k) \in F_p$, where $-c_{i,j}$ and $-c_{i,k}$ are the slopes of the utility functions $u_{i,j}(\cdot)$ and $u_{i,k}(\cdot)$ at prices p .*

PROOF. 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 $(i, j) \in E$, and all $(i, k) \in F_p$. Consider arbitrary edges $(i, j) \in E$ and $(i, k) \in F_p$. By piecewise 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 $(i, j) \in E$ 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 $(i, j) \in E$ and all $(i, k) \in F_p$. Consider arbitrary edges $(i, j) \in E$ and $(i, k) \in F_p$. By piecewise 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 $(i, j) \in E$ 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 $\lambda > 0$. \square

Overdemand-preserving price increase. It is not difficult to see that envy-free price increases 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* for a maximal alternating tree 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 T' with respect to μ' with root i_0 that has the same item and bidder set as T , and
- (c) d is envy free for the edges of the maximal alternating tree 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, that is,

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 the intervals to be of the form $[\cdot, \cdot)$.

all items whose price is increased, remain overdemanded for any small enough price increase.

LEMMA 3.6. *If d is an overdemand-preserving price increase for a maximal alternating tree 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 $\lambda > 0$.*

PROOF. Denote the partial matching and the maximal alternating tree corresponding to d by μ' and T' . Since d is envy free for T' , Lemma 3.4 shows that all edges in 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 T' is a maximal alternating tree with item set S and bidder set T in which all items are matched, Lemma 3.3 shows that (1) the set of items S is strictly overdemanded with respect to the set of bidders T and (2) $|T'| = |S| + 1$. \square

The following lemma—our key lemma and main technical improvement over Alkan [1989, 1992]—shows that if overdemand-preserving price increases are used, then the resulting prices will be minimum over all envy-free outcomes.

LEMMA 3.7. *Let d be an overdemand-preserving price increase for a maximal alternating tree 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 $p'' \geq p + \lambda \cdot d$.*

PROOF. 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 T' and $p_j \leq p_j + (\lambda - \epsilon) \cdot d_j < p_j + \lambda \cdot d_j \leq e_j$ for all $j \in S$, that is, 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 $T' \subseteq F_{p+(\lambda-\epsilon) \cdot d}$. Since 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

$$p''_k - p_k > (\lambda - \epsilon) \cdot d_k, \quad \text{and} \quad (10)$$

$$u_{i,k}(p''_k) \geq u_{i,j}(p''_j) \quad \text{for all } j. \quad (11)$$

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

$$p''_j - p_j \leq (\lambda - \epsilon) \cdot d_j, \quad \text{and} \quad (12)$$

$$u_{i,j}(p_j + (\lambda - \epsilon) \cdot d_j) \geq u_{i,k}(p_k + (\lambda - \epsilon) \cdot d_k). \quad (13)$$

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

$$\begin{aligned} u_{i,j}(p''_j) &\geq u_{i,j}(p_j + (\lambda - \epsilon) \cdot d_j) && \text{(from (12))} \\ &\geq u_{i,k}(p_k + (\lambda - \epsilon) \cdot d_k) && \text{(from (13))} \\ &> u_{i,k}(p''_k). && \text{(from (10))} \end{aligned}$$

Since this would give a contradiction to (11), we must have $F_{p''}(B) \subseteq A$, that is, $|A| \geq |F_{p''}(B)|$. 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! \square

Computing an overdemand-preserving price increase. 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), for example, by using a primal-dual algorithm [Kuhn 1955]. The duality between slopes and utility drops exploited here is reminiscent of the duality between value-maximizing matchings and envy-free prices in Shapley and Shubik [1971]. (2) The subroutine extends d to an envy-free price increase 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 3.5 (and has been used in a similar form in Alkan [1989, 1992]), (2) and (3) exploit the newly established correspondence between maximal alternating trees and strict overdemand.

PSEUDOCODE 1: 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

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

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

2 compute ω, ρ as optimal solution to

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

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**

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

PROOF. Let x , ω , and ρ be defined as in the mechanism. The constraint matrix of LP is totally unimodular, that is, $x_{i,j} \in \{0, 1\}$ for all $(i, j) \in F_p \cap (T \setminus \{i_0\} \times S)$ [Schrijver 1998]. 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 $\sigma = \{(i, j) \in T \setminus \{i_0\} \times S \mid x_{i,j} = 1\}$ matches every bidder $i \in T \setminus \{i_0\}$ and every item $j \in S$ exactly once. From duality:

- (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 line 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 , T' , T' and S' be defined as in lines 4–6. It is not difficult to see that the while loop in lines 7–11 has the following properties:

- (a) For all $i \in T'$ and all $j \in S'$: If we had $\omega_i + \rho_j < (=, \text{ respectively}) \log(c_{i,j})$ before the update, then we have $\omega_i + \rho_j < (=, \text{ respectively}) \log(c_{i,j})$ after the update.
- (b) For all $i \in T \setminus T'$ and all $j \in S \setminus S'$: If we had $\omega_i + \rho_j < (=, \text{ respectively}) \log(c_{i,j})$ before the update, then we have $\omega_i + \rho_j < (=, \text{ respectively}) \log(c_{i,j})$ 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 (e) we get that no edge from σ and 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 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, T' will cover the same bidders and items as T .

Let d be defined as in line 12. Then $d \neq 0$ is envy free for T' by Lemma 3.5 because

- (1) For all $(i, j) \in T'$ 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 $(i, j) \in \mu' \setminus T'$ 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 line 13, then μ' matches the same bidders and items as μ because (1) μ' is identical to σ on $T \times S$ and (2) μ' is identical to μ on $I \setminus T \times J \setminus S$.

The LP and the DP can be solved in time $O(\min(n, k)^3)$ [Edmonds and Karp 1972; Tomizawa 1971]. The maximal alternating tree T' can be computed in time $O(\min(n, k)^2)$ using a breadth-first search approach. The while loop in lines 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 line 8 can be computed in time $O(\min(n, k))$. When the ω_i 's and ρ_j 's are updated in line 9, the δ_j 's are adapted in time $O(\min(n, k))$ by subtracting δ from each δ_j . Thus, lines 9 and 10 take time $O(\min(n, k))$ per iteration for a total of

$O(\min(n, k)^2)$. Instead of recomputing the maximal alternating tree T' in line 10 from scratch we can keep the old one and add the required edges. Thus, maintaining 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 $O(\min(n, k)^2)$. \square

Computing a 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 3.3 the current matching μ can be augmented along an alternating path from i_0 to j_0 (j , respectively). If this is not the case, then—again by Lemma 3.3—the items S in the tree are strictly overdemanding 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 3.7. 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.

PSEUDOCODE 2: Mechanism for piecewise linear utility functions

Input: bidders I , items J , piecewise linear utility functions $u_{i,j}(\cdot)$ with nonidentical 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  $\mathcal{T}$  with respect to  $\mu$  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  $T$  and a corresponding
     matching  $\mu'$  (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  $\mu$  that do not belong to the first choice graph  $G_p$ 
10    compute maximal alternating tree  $\mathcal{T}$  w.r.t.  $\mu$  in the first choice graph  $G_p$  with root  $i_0$ 
11    let  $T$  and  $S$  be the bidders and items in  $\mathcal{T}$ 
12  end
13  augment  $\mu$  along alternating path  $P$  from  $i_0$  to unmatched item  $j$  or dummy item  $j_0$ 
14 end
15 output  $(\mu, p)$ 

```

In principle, it is possible that there is more than one alternating path along which the matching can be augmented. While selecting any such path guarantees that the outcome is bidder optimal, we bias our mechanism toward items that were already matched but got unmatched.

THEOREM 3.9. *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.*

PROOF. The matching μ is a subset of the first choice edges at prices p . Hence, Lemma 3.2 shows that (μ, p) is envy free. By Lemma 3.1 (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 timesteps t by induction over t .

Induction basis: For $t = 0$ the claim follows from the fact that every envy-free outcome (μ'', p'') has $p'_j \geq 0$ for all j .

Inductive step: To prove the claim for t assume that it holds 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 a strict overdemand-preserving price increase for 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 3.7 shows that any envy-free outcome (μ'', p'') with $p'' \geq p^{(t)}$ must have $p'_j \geq p^{(t)} + \lambda \cdot d_j$. It follows that $p'_j \geq p_j^{(t+1)}$ for all j .

We bound the total time required by (1) the outer while loop *without* the inner while-loop (lines 2–4 and 13–14) and the inner while loop corresponding to Case (a) (lines 5–12) separately from the total time required by (2) the inner while loop corresponding to Case (b) (lines 5–12).

To obtain a bound for (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 . (iv) Since the prices are monotonically increasing at most $O(D)$ discontinuities are reached. From (i) to (iv) we deduce that there are at most $O(n + D)$ iterations of the outer while loop without the inner while loop. Each iteration of the outer while loop without the inner while loop takes time $O(\min(n, k)^2)$. From (i) to (iv) it also follows that there are at most $O(n + D)$ iterations of the inner while corresponding to Case (a) such that right before their execution either the outer while loop was executed or the inner while loop was executed and a discontinuity was reached. Between any two such iterations there can be at most $O(\min(n, k))$ iterations of the inner while loop corresponding to Case (a) because each of these iterations adds at least one item to the maximal alternating tree under consideration. Each iteration of the inner while loop 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 3.8) and $O(\min(n, k) \cdot k)$ for computing the λ value. Hence, a bound for (1) is $O((n \cdot \min(n, k) + D \cdot \min(n, k)) \cdot \min(n, k) \cdot (\min(n, k)^2 + k))$.

To obtain a bound for (2) observe that because the prices are monotonically increasing there are at most $O(T)$ iterations of the inner while loop that correspond to Case (b). As argued previously, each iteration of the inner while loop takes time $O(\min(n, k) \cdot (\min(n, k)^2 + k))$. Hence, a bound for (2) is $O(T \cdot \min(n, k) \cdot (\min(n, k)^2 + k))$. \square

4. INCENTIVE COMPATIBILITY FOR PIECEWISE LINEAR UTILITIES

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

$$u_{i,\mu'(i)}(p'_{\mu'(i)}) \geq u_{i,\mu''(i)}(p''_{\mu''(i)}). \quad (14)$$

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 .¹⁰

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 $i \in \{1, 2\}$.

Given this input, our mechanism computes a bidder optimal outcome as follows: Starting from prices all zero, at which both bidders strictly prefer item 1 over item 2, it raises the price of the first item to 5. At this point both bidders lose their interest in item 1 and demand item 2 instead. The mechanism continues by raising the price of item 2 to 1. Now both bidders are indifferent between item 2 and the dummy item. Hence, the mechanism can assign bidder 1 the dummy item and bidder 2 item 2 (or vice versa). For this outcome both bidders have a utility of zero.

Next, suppose bidder 1 lies about his utility function for item 1 by reporting $u_{1,1} = 0 - p_1$ for $p_1 \in [0, \infty)$. In this case, bidder 1 (presumably) prefers item 2 over item 1 and bidder 2 prefers item 1 over item 2 when prices are all zero. Hence, our mechanism will leave the price of all items at 0 and assign item 1 to bidder 2 and item 2 to bidder 1. Bidder 1's utility for this outcome is 1 and, thus, strictly higher than the utility that he would get if he reported truthfully.

The crucial point—as we will show in the following—is that in the computation of the bidder optimal outcome for the original utility functions the first choice edges from bidder 1 to item 1 and from bidder 2 to item 1 break away during the same price increase.

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

¹⁰Overreporting can only lead to a missed chance of being assigned an item and underreporting can only lead to a utility below the true outside option.

follows that

$$v_{i',j} - c_{i',j} \cdot p_j = v_{i',j'} - c_{i',j'} \cdot p_{j'}. \quad (15)$$

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

$$v_{i',j} - c_{i',j} \cdot (p_j + \lambda d_j) = v_{i',j'} - c_{i',j'} \cdot (p_{j'} + \lambda d_{j'}). \quad (16)$$

By subtracting (15) from (16), dividing by $\lambda > 0$, and after rearranging we get

$$d_j = c_{i',j'}/c_{i',j} \cdot d_{j'}. \quad (17)$$

Now suppose that the discontinuities $D_{i,j}$ and $D_{i',j'}$ are reached simultaneously. Then by (16)

$$v_{i',j} - c_{i',j} \cdot D_{i,j} = v_{i',j'} - c_{i',j'} \cdot D_{i',j'}. \quad (18)$$

Using (17), subtracting $1/d_j \cdot v_{i,j}/c_{i,j}$ from both sides, and after rearranging we get

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

Next, we will define a multigraph such that the left- and right-hand sides 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 depends 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: (1) There is a forward edge from $i \in I$ to $j \in J$ for each linear segment of $u_{i,j}(\cdot)$. (2) There is a backward edge from $j \in J$ to $i \in I$ for each linear segment of $u_{i,j}(\cdot)$. (3) There is a discontinuity edge from $i \in I$ to $j \in J$ 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. Then 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 (Q , respectively), we have $w_d(P) = w_d(Q)$.

Note that this definition of general position is more general than the definition in the following sense: (a) It reduces to the definition in Aggarwal et al. [2009] for utility functions with identical slopes and at most one discontinuity. (b) It is also applicable to utility functions with nonidentical slopes and multiple discontinuities.

Properties of the bidder optimal outcome. Next, we show that for inputs in general position the bidder optimal outcome computed by our mechanism has a specific structure. To this end we first show that if an input is in general position, then the corresponding input in standard form is in general position. We then prove that for inputs in general position no two discontinuities are reached during the same price increase of our mechanism. A consequence of this is that if a matched item ever gets unmatched, then it gets matched again in the subsequent iteration. We use this to show that in the bidder optimal outcome computed by our mechanism the prices of all *unmatched* items and at least one *matched* item are equal to the reserve prices. Note

that the former is precisely what is needed for the bidder optimal, envy-free outcome to correspond to a competitive equilibrium.

LEMMA 4.1. *An input is in general position if and only if the associated input in standard form is in general position.*

PROOF. 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' (Q and Q' , respectively), 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 (Q , respectively), has $d_j = (c'_{i,j}/c'_{i,j'}) \cdot d_{j'}$ for any two edges (i, j) and (i, j') on P' (Q' , respectively), 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' (Q and Q' respectively), and $D_{i,j} = D'_{i,j} + r_j$ for every discontinuity edge (i, j) on P and P' (Q and Q' , respectively), (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' (Q and Q' , respectively), 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 $w_d(Q')$. \square

LEMMA 4.2. *Let (μ, p) denote the outcome of the mechanism from the previous section. 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 $p_j = r_j$.*

PROOF. We prove each of the claims (i)–(iv) individually.

To (i): For a contradiction suppose that two discontinuities are reached at the same time. Without loss of generality, 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 alternating walks $P = (i_1, j_1)$ and $Q = (i_1, j_1, i_2, j_2, \dots, i_{t-1}, j_{t-1}, i_t, j_{t-1})$ in the maximal alternating tree corresponding to the price increase. Note that these walks always exist because (1) either j_1 lies on the path from i_1 to i_t or $i_2 = i_1$ and (2) 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

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

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 increase in which a matched item gets unmatched. 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}' , $(\mu'$, respectively). We know that μ' matches the same set of bidders and items as μ . A matched item j can get unmatched only if the edge along which this item is matched under μ' drops out of the first choice graph $G_{p+\lambda \cdot d}$ for some $\lambda > 0$. From (i) we know that apart from this edge no other edge can drop out of the first choice graph at this point. But then, because 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 j to the root i_0 . Hence, since the mechanism is biased toward matching items that were already matched and got unmatched, item j will get 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 (ii) 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 $F_p(\mathcal{T})$, we must have $p_j \geq r_j$. From (ii) we know that item j was never matched before and from (iii) we know that $p_j > r_j$ would imply that j was matched before. Thus, we must have $p_j = r_j$. \square

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., Gale and

Sotomayor [1985] and Demange et al. [1987])—we can show that any mechanism that computes a bidder optimal outcome is incentive compatible if conditions (iii) and (iv) from Lemma 4.2 are satisfied. A straightforward corollary is that if the input is in general position, then any mechanism that computes a bidder optimal outcome is incentive compatible.¹¹

LEMMA 4.3. *If conditions (iii) and (iv) from Lemma 4.2 are satisfied, then (1) if the outcome (μ^*, p^*) is bidder optimal, then there does not exist a feasible outcome (μ', p') with $u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu^*(i)}(p^*_{\mu^*(i)})$ for all i . (2) If the outcome (μ^*, p^*) is bidder optimal, the outcome (μ', p') is feasible, and $I^+ \neq \emptyset$, where $I^+ = \{i \in I \mid u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu^*(i)}(p^*_{\mu^*(i)})\}$, then there exists a bidder-item pair $(i, j) \in I \setminus I^+ \times J$ such that $u_{i,\mu'(i)}(p'_{\mu'(i)}) < u_{i,j}(p'_j)$.*

PROOF. To (1): For a contradiction assume that there is 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 4.2 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 nondummy items j under μ' . By (a) all bidders are matched to nondummy items under μ . Condition (iv) from Lemma 4.2 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 $\mu(I^+)$ ($\mu'(I^+)$, respectively), denote the set of items matched to bidders in I^+ under μ (μ' , respectively). 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 4.2. 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 $i \in I^+$.¹² 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

¹¹Our argument, in fact, shows an even stronger claim. Namely, that no *group* of bidders can benefit from misreporting their preferences.

¹²If $u_{i,j}(\cdot)$ 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.

$i \in I^+$ 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 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 $u_{i,\mu'(i)}(p'_{\mu'(i)}) < u_{i,j}(p'_j)$. \square

THEOREM 4.4. *If conditions (iii) and (iv) from Lemma 4.2 are satisfied, then any mechanism that computes a bidder optimal is incentive compatible.*

PROOF. 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^*) ((μ', p') , respectively) denote the bidder optimal outcome for the true (falsified, respectively) 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 4.3 shows that if conditions (iii) and (iv) from Lemma 4.2 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 4.3 shows that if conditions (iii) and (iv) from Lemma 4.2 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 $i \notin I^+$, that is, 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 $(u'_{i,j}(\cdot), r_j, o_i)$. \square

Discussion. The general position concept serves as an informal argument to qualify inputs that are *not* in general position as degenerate. The intuition is that for generic inputs, for example, randomly generated ones, it is rather unlikely that two walks in the graph that alternate between forward and backward edges and end with a discontinuity edge have *exactly* the same weight.

It would still be interesting, though, to be able to test whether a given input is in general position and to know how to deal with inputs that are *not* in general position. Regarding the former we would like to point out that our mechanism can be used to test whether an input satisfies conditions (iii) and (iv), which are sufficient for incentive compatibility. Regarding the latter, a promising approach seems to be to first bring the input that is not in general position into general position (e.g., by random perturbation) and to then use the bidder optimal solution to the input that has been brought into general position as a solution to the original input. The problem with this approach, however, is that it is not clear under which conditions the computed solution will also be bidder optimal for the original input.¹³

¹³Note that we cannot simply assign nothing and charge nothing if we detect that the input is *not* in general position as this may not be bidder optimal.

5. GENERAL UTILITIES VIA PIECEWISE LINEAR APPROXIMATION

In this section, we show how our mechanism for piecewise 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 piecewise 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 piecewise linear approximation:

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

Piecewise linear approximation. Given $\epsilon > 0$ we construct a piecewise linear, continuous approximation $\tilde{u}_{i,j}(\cdot)$ for $u_{i,j}(\cdot)$ with error at most ϵ , that is, $|u_{i,j}(p_j) - \tilde{u}_{i,j}(p_j)| \leq \epsilon$ 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 end points $[\tilde{e}^{(t)}, \tilde{e}^{(t+1)}]$, 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 piecewise linear approximation agrees with the original function at the end points of each interval. The following lemma shows that the preceding algorithm does indeed give a close approximation.

LEMMA 5.1. *For every bidder-item pair $(i, j) \in I \times J$ the algorithm described previously 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 piecewise linear, (ii) continuous, and (iii) has error at most ϵ .*

PROOF. 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. For point (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

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

Similarly, the maximum of $|g_{i,j}(p_j)|$ 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 $\min((\tilde{e}_{i,j}^{(t+1)} - \tilde{e}_{i,j}^{(t)} - \delta)^2 B / 2, \delta^2 B / 2)$, which is maximized for $\delta = (\tilde{e}_{i,j}^{(t+1)} - \tilde{e}_{i,j}^{(t)}) / 2$, corresponding to a peak of at most $(\tilde{e}_{i,j}^{(t+1)} - \tilde{e}_{i,j}^{(t)})^2 / 8 \cdot B \leq w^2 / 8 \cdot B \leq \epsilon$ by the choice of w . \square

Approximately bidder optimal outcome. Given a set of continuous, piecewise 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 piecewise linear approximations are, that is, 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 “jump-start” 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 3.1 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 nonzero 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 , which leads to a difference between p and \tilde{p}' linear in ϵ . We use the resulting bound on the difference between the prices to bound the difference between the utilities. Theorem 5.3 follows from this bound for ϵ small enough.

LEMMA 5.2. *For every envy-free outcome (μ, p) for $u_{i,j}(\cdot)$ and continuous, piecewise 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 $\tilde{p}_j \leq p_j + (2M/m)^{\min(n,k)+1} \cdot \epsilon/M$ for all j and $\tilde{u}_{i,\tilde{\mu}(i)}(\tilde{p}_{\tilde{\mu}(i)}) + (2M/m)^{\min(n,k)+1} \cdot \epsilon \geq u_{i,\mu(i)}(p_{\mu(i)})$ for all i .*

PROOF. Consider an arbitrary envy-free outcome (μ, p) for $u_{i,j}(\cdot)$. Denote the set of bidders that are unmatched under μ by U . Let $p'_j = \max(p_j, \min\{\hat{p}_j : \forall i \in U : \tilde{u}_{i,j}(\hat{p}_j) \leq o_i\})$. 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, $p'_j \leq p_j + \epsilon/m$ for all j . Moreover, for every envy-free outcome for $\tilde{u}_{i,j}(\cdot)$ with prices higher or equal to p'_j for all j we can match the bidders in U to the dummy item.

Suppose that we use the mechanism from Section 3 to find a bidder optimal outcome for the bidders in $I \setminus U$, utility functions $\tilde{u}_{i,j}(\cdot)$, and reserve prices $r_j = p'_j$ for all j . Let us write q for the number of times Step 7(a) has been encountered by the mechanism since the beginning of its execution. We write $p^{(q)}$ for the prices after q such cases and so $p_j^{(0)} = p'_j$. Furthermore, we write $\delta^{(q)} = \max_j(p_j^{(q)} - p_j^{(0)})$. Whenever we compute the sets S and T we have $|T| > |S|$ by the definition of strict overdemand. Hence, there exists a bidder $\hat{i} \in T$ such that $\mu(\hat{i}) \notin S$. Regardless of whether \hat{i} turns out to be the bidder i of the (i, j) pair in the next Step 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 the price of any item in S has to be increased until we reach Step 7(a). Let $\hat{j} \in S$ be any of \hat{i} 's current first choice items before the price update, that is, $\hat{j} \in F_{p^{(q)}}(\hat{i})$. Let us define the smallest utility gap 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 A.3 in the Appendix shows that the sets S and T do not change between consecutive Steps 7(b) and so all of \hat{i} 's first choice items, which can change with each new Step 7(b), remain in the set S until the next Step 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 Step 7(a) is reached. Lemma A.2 in the Appendix shows that 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 $\alpha^{(q)}(M/m)^{\min(n,k)} \cdot 1/m$. So we know $\delta^{(q+1)} \leq \delta^{(q)} + \alpha^{(q)}(M/m)^k \cdot 1/m$. We bound $\alpha^{(q)}$ as follows:

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

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

price increase after q occurrences of Step 7(a) is bounded by $\delta^{(q)} \leq (1 + (M/m)^{\min(n,k)+1})^q \cdot 2\epsilon/M$.

We claim that q is at most $\min(n, k)^2$ and we argue as follows. In the continuous case bidders never get unmatched and the outer while loop matches a new bidder each time it gets executed. Since there are at most $\min(n, k)$ bidders this shows that the outer while loop gets executed at most $\min(n, k)$ times. Between two consecutive executions of the outer while loop there can be at most $\min(n, k)$ executions of the inner while loop corresponding to Step 7(a) because each such step adds an edge from a bidder $i \in T$ to an item $j \notin S$. As a consequence, either bidder i_0 gets matched or the maximal alternating tree grows by at least one bidder-item pair. The former can happen exactly once. The latter can happen at most $\min(n, k)$ times because there are at most $\min(n, k)$ bidder-item pairs.

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:

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

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 $\min(n, k) \geq 1$ and that $M/m \geq 1$. Exactly the same type of manipulations can be applied to show that $\tilde{p}_j \leq p'_j + \delta^{(q)} \leq p_j + \epsilon/m + 2\epsilon/M \cdot (1 + (M/m)^{\min(n,k)+1})^q \leq p_j + (2M/m)^{(\min(n,k)+1)^3} \cdot \epsilon/M$. \square

THEOREM 5.3. *Given strictly monotonically decreasing, continuous utility functions $u_{i,j}(\cdot)$ satisfying assumptions (A)–(C), we can compute a γ -bidder optimal outcome with the running time specified in Theorem 3.9, where $D = 0$, $T = O(\sqrt{1/\epsilon})$, and $\epsilon = \gamma \cdot (2M/m)^{-(\min(n,k)+1)^3}$.*

PROOF. Let $\epsilon = \gamma(2M/m)^{-(\min(n,k)+1)^3}$. For each utility function $u_{i,j}(\cdot)$, find a piecewise linear approximation $\tilde{u}_{i,j}(\cdot)$ with error at most 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 5.1, 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 3.9, that is, 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 γ -bidder optimal for $u_{i,j}(\cdot)$.

We have $\tilde{p}_j \geq r_j$ and $\tilde{p}_{j_0} = 0$ as (1) we did not change the reserve prices and (2) 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 $\max_i(\max_j(u_{i,j}(\tilde{p}_j) - u_{i,\tilde{\mu}(i)}(\tilde{p}_{\tilde{\mu}(i)}))) \leq \max_i(\max_j(\tilde{u}_{i,j}(\tilde{p}_j) - \tilde{u}_{i,\tilde{\mu}(i)}(\tilde{p}_{\tilde{\mu}(i)}) + 2\epsilon)) \leq \gamma$ by the choice of ϵ . Finally, and most importantly, Lemma 5.2 shows that the utility losses with respect to the unknown bidder optimal outcome (μ, p) are bounded by γ . \square

6. HARDNESS RESULTS FOR 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 one, 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 6.1. *Given utility functions $u_{i,j}(\mu)$ that depend on the matching μ , and given a constant K , deciding whether there exists an outcome with sum of the utilities at least K is \mathcal{NP} -complete. This is true even if there is an ordering on the items, each bidder i has a set S_i of three other bidders, and the utility of bidder i depends on whether in the ordering the items to which the bidders in S_i are matched to appear before or after the item that bidder i is matched to.*

PROOF. Checking whether a given outcome (μ, p) 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 [Cook 1971]: Given n Boolean variables x_k , m clauses c_r (disjunction of up to three literals), and a formula f (conjunction of the m clauses) decide whether there exists a satisfying assignment x .

For a given instance of 3-SAT: For every variable create one bidder, called variable bidder, and two items, one called true item and one called false item. The variable bidder has a utility of zero for each of these items and minus infinity for all other items. For every clause create one bidder, called clause bidder, and one item, called clause item. In the ordering all true items are before all clause items, and all clause items are before all false items. The utility of a clause bidder for the clause item depends on whether the variable bidders corresponding to variables in that clause are matched to items before or after the clause item, that is, to a true item or a false item. It is 1 if and only if the matching is such that the clause is satisfied and 0 otherwise. The utility function of a clause bidder for all other items is minus infinity.

It is not difficult to see that there is an outcome with sum of the utilities at least $K = m$ if and only if the formula is satisfiable. This completes the proof. \square

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 one, for example, to express that a bidder i achieves a higher utility when the other bidders have to pay more, that is, 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 6.2. *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_l for all $l \neq j$, and given a constant L , 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 .

PROOF. We prove the claim for utility functions $u_{i,j}(p_1, \dots, p_k)$ that depend on the vector of prices $p = (p_1, \dots, p_k)$ and are decreasing in p_l for all l . 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.

Checking whether a given outcome (μ, p) is envy free and has sum of the utilities at least L takes polynomial time. This establishes membership in \mathcal{NP} . To prove hardness we provide a polynomial-time reduction from MAX-2-NASH [Gilboa and Zemel 1989; Conitzer and Sandholm 2008]: 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 $\sigma_i(s) > 0$, that is, player i plays strategy s with nonzero 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 can be shown that $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is a mixed NE if and only if 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 asks for a mixed NE with total payoff at least K .

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 [Gilboa and Zemel 1989; Conitzer and Sandholm 2008]. Let $C = (N! \cdot N^N)^{2N}$. Note that C can be represented with at most $4N^2 \log(N)$ bits. Given an MAX-2-NASH instance with $|S_1| = |S_2| = N$ and payoff matrices with integer entries between 0 and N we construct an instance of our problem by means of the following ‘‘gadgets’’:

Best response gadget B: For each player i and strategy $s \in S_i$ we add 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, that is, $o_{is1} = o_{is2} = 0$.

Probability gadget P1: For each player i and strategy $s \in S_i$ we add 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 $-\infty$.

Probability gadget P2: For each player i and strategy $s \in S_i$ we add 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 $-\infty$.

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

Distribution gadget D2: For each player i we add 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 1. *If all payoffs $a_i(s_i, s_{-i})$ in a two-player game are nonnegative integers, with the largest being N , then the probabilities in any mixed NE are rational numbers with the smallest being $1/(N! \cdot N^N)$.*

PROOF. 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 nonnegative integers, with the largest being N , any solution $\sigma^* = (\sigma_1^*, \sigma_2^*)$ to this LES consists of rational numbers with denominator at most $N! \cdot N^N$ [Limongelli and Pirastu 1994]. Hence, the smallest probability in the mixed NE is at least $1/(N! \cdot N^N)$. \square

CLAIM 2. *There exists a mixed NE $\sigma^* = (\sigma_1^*, \sigma_2^*)$ with total payoff K if and only if there exists an envy free outcome (μ, p) with sum of the utilities strictly larger than $N \cdot (K + 4) + 2$ and less than $N \cdot (K + 4) + 2 + 1/C$.*

PROOF. For the *only if* part assume that 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 $p_{is3} = 1 - \sigma_i^*(s) = 1$ 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 $u_{is2, is4}(p) \geq 0$. 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) $1 - p_{is3} = 0$ if and only if $1/(8NC) - p_{is4} = 0$, this is also true for the bidders in the $D1$ and $D2$ gadgets. The $is1$ vertices contribute a utility of $N \cdot K$, 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 $(|SUPP_1| + |SUPP_2|)/(8NC)$. The $P1$ gadgets contribute $2 \cdot N + 2$. The $P2$ gadgets contribute $2 \cdot N + (|SUPP_1| + |SUPP_2|)/(8NC)$. 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, $2 \leq |SUPP_1| + |SUPP_2| \leq 2N$.

For the *if* part consider an envy-free outcome that maximizes the sum of the utilities. Denote the sum of the utilities by L . Let $T_1 = \{s \in S_1 : is1 \text{ is matched to } is3\}$ 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 in the following that for each player i with $p_{is3} < 1$ bidder $is1$ must be 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 such that $1 - p_{is3} > 0$, that is, $p_{is3} < 1$. It follows that for each player there is at least one $is1$ node that is matched to an $is3$ node. In other words, $|T_1| + |T_2| \geq 2$ and, thus, $L > N \cdot (K + 4) + 2$.

Next, we show that $\sigma^* = ((1 - p_{1s3})_{s \in S_1}, (1 - p_{2s3})_{s \in S_2})$ is a mixed NE, that is, for each player i every strategy $s \in S_i$ with $1 - p_{is3} > 0$ is a best response to $\sigma_{-i}^* = (1 - p_{-is3})_{s \in S_{-i}}$. We do this by showing that (1) if $1 - p_{is3} > 0$, then $is1$ is matched to $is3$, and (2) if $is1$ is matched to $is3$, then s is a best response to $\sigma_{-i}^* = (1 - p_{-is3})_{s \in S_{-i}}$.

To (1): To show that if $(1 - p_{is3}) > 0$, then $is1$ is matched to $is3$ we show that if $is1$ is *not* matched to $is3$, then $(1 - p_{is3}) \leq 0$. Note that by envy freeness if $is1$ is not matched to $is3$, then it must be matched to $is4$. Also note that by P1 we have $(1 - p_{is3}) \geq 0$ so that $(1 - p_{is3}) \leq 0$ is equivalent to $p_{is3} = 1$. In other words, we have to show that if $is1$ is matched to $is4$, then $p_{is3} = 1$. If $is1$ is matched to $is4$, then $p_{is4} \geq 1/(8NC)$ because otherwise $is2$ is not envy free. With P2 we know that $p_{is4} = 1/(8NC)$.

Additionally, since the outcome maximizes the sum of the utilities, 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 $is4(1 - p_{is3}) = 0$ (\star). Since by $P1$ we know that $1 - p_{is3} \geq 0$ for all i and s , (\star) is only possible if $(1 - p_{is3}) = 0$ for all $s \in S_i$ with $is1$ matched to $is4$. This shows that if $is1$ is matched to $is4$, then $p_{is3} = 1$.

To (2): If $is1$ is matched to $is3$, then $u_{is1, is3}(p) \geq u_{is1, is4}(p)$ because otherwise the outcome would not be envy free. This shows that $\sum_{t \in S_{-i}} a_i(s, t) \sigma_{-i}^*(t) \geq \max_{r \in S_i} \sum_{t \in S_{-i}} a_i(r, t) \sigma_{-i}^*(t)$. Thus, for all $r \in S_i$: $\sum_{t \in S_{-i}} a_i(s, t) \sigma_{-i}^*(t) \geq \sum_{t \in S_{-i}} a_i(r, t) \sigma_{-i}^*(t)$. In other words, s is a best response to $\sigma_{-i}^* = (1 - p_{it3})_{t \in S_{-i}}$.

We conclude by showing that the payoff of the mixed NE $\sigma^* = ((1 - p_{1s3})_{s \in S_1}, (1 - p_{2s3})_{s \in S_2})$ is K . The payoff to player i 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 s , 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 both i the total payoff is $K_1 + K_2 = K$. \square

7. CONCLUSION AND FUTURE WORK

The demand for more expressive mechanisms is reflected in the ‘‘richness of preference expression offered by businesses as diverse as matchmaking sites, sites like Amazon and NetFlix, and services like Google’s AdSense’’ [Benisch et al. 2008]. 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 facing.

In this article, 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 Aggarwal et al. [2009], can handle *piecewise linear* utility functions with *nonidentical 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, for example, 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 toward such a general theory was recently undertaken by Benisch et al. [2008, 2009], Milgrom [2010], and Dütting et al. [2011]. Any refinement of this theory would be highly interesting.

APPENDIX

A. ANALYSIS OF MECHANISM FOR CONTINUOUS UTILITY FUNCTIONS

The following lemmata bound the price increase between two consecutive Steps 7(a) in the mechanism for piecewise linear utility functions when there are no discontinuities.

LEMMA A.1. *Let \mathcal{T} , S , and T be the maximal alternating tree, item set, and bidder set considered by the mechanism. Then the overdemand-preserving price increase d computed in line 6 satisfies $\max_{j' \in S, j \in S} (d_{j'} / d_j) \leq (M/m)^{\min(n, k)}$.*

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 3.5 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 A.2. *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 $j, j' \in S$ we have $(\sum_{l=1}^s \lambda^{(l)} d_j^{(l)}) / (\sum_{l=1}^s \lambda^{(l)} d_{j'}^{(l)}) \leq (M/m)^{\min(n,k)}$.*

PROOF. Lemma A.1 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 A.3. *From one Step 7(a), where a new bidder-item pair (i, j) enters the first choice graph, to the next, which can be separated by many Steps 7(b) and changes in the price increase d , the sets S and T considered by the mechanism do not change.*

PROOF. We will prove the claim by induction over the number of Steps 7(b) separating two Steps 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 with respect 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 correspond 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

ACKNOWLEDGMENTS

We would like to thank Veronika Loitzenbauer and the anonymous referees for their valuable feedback.

REFERENCES

- G. Aggarwal, S. Muthukrishnan, D. Pál, and M. Pál. 2009. General auction mechanism for search advertising. In *Proceedings of the 18th WWW*. 241–250.
- S. Alaei, K. Jain, and A. Malekian. 2011. Competitive equilibrium in two sided matching markets with general utility functions. *SIGecom Exchanges* 10, 2, Article 8 (2011).
- A. Alkan. 1989. Existence and computation of matching equilibria. *European Journal of Political Economy* 5, 2–3 (1989), 285–296.
- M. Alkan. 1992. Equilibrium in a matching market with General Preferences. In *Equilibrium and Dynamics. Essays in Honour of David Gale*, M. Majumdar (Ed.). The MacMillan Press Ltd, London, 1–16.
- I. Ashlagi, M. Braverman, and A. Hassidim. 2009. Ascending unit demand auctions with budget limits. *Working paper* (2009).
- I. Ashlagi, M. Braverman, A. Hassidim, R. Lavi, and M. Tennenholtz. 2010. Position auctions with budgets: Existence and uniqueness. *The B.E. Journal of Theoretical Economics (Advances)* 10, 1, Article 20 (2010).
- I. Ashlagi, P. Krysta, and M. Tennenholtz. 2008. Social context games. In *Proceedings of the 4th WINE*. 675–683.

- M. Benisch, N. Sadeh, and T. Sandholm. 2008. A theory of expressiveness in mechanisms. In *Proceedings of the 23rd AAAI*. 17–23.
- M. Benisch, N. Sadeh, and T. Sandholm. 2009. Methodology for designing reasonably expressive mechanisms with application to ad auctions. In *Proceedings of the 21st IJCAI*. 46–52.
- S. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. 2005. Multi-unit auctions with budget constrained bidders. In *Proceedings of the 6th EC*. 44–51.
- C. Boutilier, D. C. Parkes, T. Sandholm, and W. E. Walsh. 2008. Expressive banner ad auctions and model-based online optimization for clearing. In *Proceedings of the 23rd AAAI*. 30–37.
- N. Chen, X. Deng, and A. Ghosh. 2010. Competitive equilibria in matching markets with budgets. *SIGecom Exchanges* 9, 1, Article 5 (2010).
- X. Chen and X. Deng. 2006. Settling the complexity of two-player Nash equilibrium. In *Proceedings of the 47th FOCS*. 261–272.
- E. H. Clarke. 1971. Multipart pricing of public goods. *Public Choice Journal* 11, 1 (1971), 17–33.
- R. Colini-Baldeschi, M. Henzinger, S. Leonardi, and M. Starnberger. 2012. On multiple keyword sponsored search auctions with budgets. In *Proceedings of the 39th ICALP*. 1–12.
- V. Conitzer and T. Sandholm. 2008. New complexity results about Nash equilibria. *Games and Economic Behavior* 63, 2 (2008), 621–641.
- V. Conitzer and T. Sandholm. 2011. Expressive markets for donating to charities. *Artificial Intelligence* 175, 7–8 (2011), 1251–1271.
- S. A. Cook. 1971. The complexity of theorem-proving procedures. In *Proceedings of the 3rd STOC*. 151–158.
- C. Daskalakis and C. H. Papadimitriou. 2005. Three-player games are hard. *Electronic Colloquium on Computational Complexity* (2005), 1–10.
- G. Demange and D. Gale. 1985. The strategy structure of two-sided matching markets. *Econometrica* 53, 4 (1985), 873–888.
- G. Demange, D. Gale, and M. Sotomayor. 1986. Multi-item auctions. *Journal of Political Economy* 94, 4 (1986), 863–872.
- G. Demange, D. Gale, and M. Sotomayor. 1987. A further note on the stable matching problem. *Discrete Applied Mathematics* 16, 3 (1987), 217–222.
- X. Deng and Y. Du. 2008. The computation of approximate competitive equilibrium is PPAD-hard. *Information Processing Letters* 108, 6 (2008), 369–373.
- S. Dobzinski, R. Lavi, and N. Nisan. 2012. Multi-unit auctions with budget limits. *Games and Economic Behavior* 74, 2 (2012), 486–503.
- P. Dütting, F. Fischer, and D. C. Parkes. 2011. Simplicity-expressiveness tradeoffs in mechanism design. In *Proceedings of the 12th EC*. 341–350.
- P. Dütting, M. Henzinger, and S. Starnberger. 2012. Auctions with heterogeneous items and budget limits. In *Proceedings of the 8th WINE*. 44–57.
- P. Dütting, M. Henzinger, and I. Weber. 2013a. Bidder optimal assignments for general utilities. *Theoretical Computer Science* 478 (2013), 22–32.
- P. Dütting, M. Henzinger, and I. Weber. 2013b. Sponsored search, market equilibria, and the Hungarian method. *Information Processing Letters* 113, 3 (2013), 67–73.
- B. Edelman, M. Ostrovsky, and M. Schwarz. 2007. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97, 1 (2007), 242–259.
- J. Edmonds and R. M. Karp. 1972. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM* 19, 2 (1972), 248–264.
- D. Gale and M. Sotomayor. 1985. Some remarks on the stable matching problem. *Discrete Applied Mathematics* 11, 3 (1985), 223–232.
- A. Ghosh, R. P. McAfee, K. Papineni, and S. Vassilvitskii. 2009. Bidding for representative allocations for display advertising. In *Proceedings of the 5th WINE*. 208–219.
- A. Ghosh and A. Sayedi. 2010. Expressive auctions for externalities in online advertising. In *Proceedings of the 19th WWW*. 371–380.
- I. Gilboa and E. Zemel. 1989. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior* 1, 1 (1989), 80–93.
- A. Goel and K. Munagala. 2009. Hybrid keyword search auctions. In *Proceedings of the 18th WWW*. 221–230.
- T. Groves. 1973. Incentives in teams. *Econometrica* 41, 4 (1973), 617–631.
- P. Hall. 1935. On representatives of subsets. *London Mathematical Society* 10 (1935), 26–30.

- P. Jehiel, B. Moldovanu, and E. Stacchetti. 1996. How (not) to sell nuclear weapons. *American Economic Review* 86, 4 (1996), 814–829.
- D. Kempe, A. Mu’alem, and M. Salek. 2009. Envy-free allocations for budgeted bidders. In *Proceedings of the 5th WINE*. 537–544.
- H. W. Kuhn. 1955. The Hungarian Method for the assignment problem. *Naval Research Logistics* 2 (1955), 83–97.
- H. B. Leonard. 1983. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy* 91, 3 (1983), 461–479.
- C. Limongelli and R. Pirastu. 1994. Exact solution of linear systems over rational numbers by parallel p-adic arithmetic. In *Proceedings of the 3rd CONPAR* (1994), 313–323.
- P. Milgrom. 2010. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior* 70, 1 (2010), 62–70.
- N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press, New York.
- D. C. Parkes and T. Sandholm. 2005. Optimize-and-dispatch architecture for expressive ad auctions. In *Proceedings of the 1st Workshop on Ad Auctions*.
- T. Sandholm. 2007. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine* 28, 3 (2007), 45–58.
- A. Schrijver. 1998. *Theory of Linear and Integer Programming*. John Wiley & Sons, New York.
- L. S. Shapley and M. Shubik. 1971. The assignment game: The core I. *International Journal of Game Theory* 1, 1 (1971), 111–130.
- N. Tomizawa. 1971. On some techniques useful for solution of transportation network problems. *Networks* 1, 2 (1971), 173–194.
- H. Varian. 2007. Position auctions. *International Journal of Industrial Organization* 25, 6 (2007), 1163–1178.
- V. Vazirani and M. Yannakakis. 2010. PPAD-completeness of fisher markets under piecewise-linear, concave utility functions. In *Proceedings of the 6th ICTS*. 156–166.
- W. Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Finance* 16, 1 (1961), 8–27.
- W. E. Walsh, C. Boutilier, T. Sandholm, R. Shields, G. L. Nemhauser, and D. C. Parkes. 2010. Automated channel abstraction for advertising auctions. In *Proceedings of the 24th AAAI*. 887–894.

Received November 2011; revised June 2012; accepted January 2015