

Efficiency-Revenue Trade-offs in Auctions

Ilias Diakonikolas^{1*}, Christos Papadimitriou^{1**}, George Pierrakos^{1**}, and Yaron Singer^{2***}

¹ UC Berkeley, EECS, {ilias, christos, georgios}@cs.berkeley.edu

² Google, Inc., yaron@cs.berkeley.edu

Abstract. When agents with independent priors bid for a single item, Myerson’s optimal auction maximizes expected revenue, whereas Vickrey’s second-price auction optimizes social welfare. We address the natural question of *trade-offs* between the two criteria, that is, auctions that optimize, say, revenue under the constraint that the welfare is above a given level. If one allows for randomized mechanisms, it is easy to see that there are polynomial-time mechanisms that achieve any point in the trade-off (the *Pareto curve*) between revenue and welfare. We investigate whether one can achieve the same guarantees using *deterministic* mechanisms. We provide a negative answer to this question by showing that this is a (weakly) NP-hard problem. On the positive side, we provide polynomial-time deterministic mechanisms that approximate with arbitrary precision any point of the trade-off between these two fundamental objectives for the case of two bidders, even when the valuations are correlated arbitrarily. The major problem left open by our work is whether there is such an algorithm for three or more bidders with independent valuation distributions.

1 Introduction

Two are the fundamental results in the theory of auctions. First, Vickrey observed that there is a simple way to run an auction so that social welfare (efficiency) is maximized: The second-price (Vickrey) auction is optimally efficient, independently of how bidder valuations are distributed. However, the whole point of the Vickrey auction is to deliberately sacrifice auctioneer revenue in order to achieve efficiency. If auctioneer revenue is to be maximized, Myerson showed in 1980 that, when the bidders’ valuations are distributed independently, a straightforward auction (essentially, a clever reduction to Vickrey’s auction via an ingenious transformation of valuations) achieves this.

These two criteria, social welfare and revenue, are arguably of singular and paramount importance. It is therefore a pity that they seem to be at loggerheads: It is not hard to establish that optimizing any one of these two criteria can be very suboptimal with respect to the other. In other words, there is a substantial *trade-off* between these two important and natural objectives. *What are the various intermediate (Pareto) points of this trade-off? And can each such point be computed — or all such points summarized somehow — in polynomial time?* This is the fundamental problem that we consider in this paper. See Figure 1 (a) for a graphical illustration.

* Supported by a Simons Postdoctoral Fellowship.

** Research supported by NSF grant CC-0964033 and by a Google University Research Award.

*** Research supported by a Microsoft Graduate Fellowship and a Facebook Graduate Fellowship.

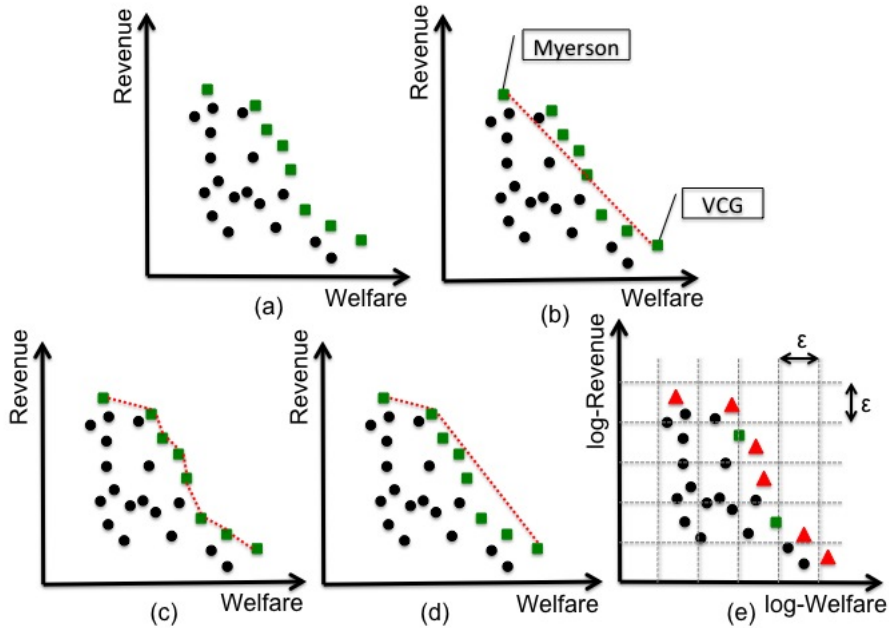


Fig. 1: The Pareto points of the bi-criterion auction problem are shown as squares (a); the Pareto points may be far off the line connecting the two extremes (b), and may be non-convex (c). The Pareto points of randomized auctions comprise the upper boundary of the convex closure of the Pareto points (d). Even though the Pareto set may be exponential in size, for any $\epsilon > 0$, there is always a polynomially small set of ϵ -Pareto points, the triangular points in (e), that is, points that are not dominated by other solutions by more than ϵ in any dimension. We study the problem of computing such a set in polynomial time.

The problem of exploring the revenue/welfare trade-off in auctions turns out to be a rather sophisticated problem, defying several naive approaches. One common-sense approach is to simply randomize between the optima of the two extremes, Vickrey’s and Myerson’s auctions. This can produce very poor results, since it only explores the straight line joining the two extreme points, which can be very far from the true trade-off (Figure 1 (b)). A second common-sense approach is the so-called *slope search*: To explore the trade-off space, just optimize the objective “revenue + $\lambda \cdot$ welfare” for various values of $\lambda > 0$. By modifying Myerson’s auction this objective can indeed be optimized efficiently, as it was pointed out seven years ago by Likhodedov and Sandholm [19]. The problem is that the trade-off curve may not be convex (Figure 1 (c)), and hence the algorithm of [19] can miss vast areas of trade-offs:

Proposition 1. *There exist instances with two bidders with monotone hazard rate distributions for which the Pareto curve is not convex; in contrast the Pareto curve is always convex for one bidder with a monotone hazard rate distribution.*

The proof is deferred to the full version. It follows that the slope search approach of [19] is incorrect. However, the correctness of the slope search approach is restored if

one is willing to settle for randomized mechanisms: The trade-off space of randomized mechanisms is always convex (in particular, it is the convex hull of the deterministic mechanisms, (Figure 1 (d)). It is easy to see (and it had been actually worked out for different purposes already in [23]) that the optimum randomized mechanism with respect to the metric “revenue + λ · welfare” is easy to calculate.

Proposition 2. *The optimum randomized mechanism for the objective “revenue + λ · welfare” can be computed in polynomial time. Hence, any point of the revenue/welfare trade-off for randomized mechanisms can be computed in polynomial time.*

1.1 Our results

In this paper we consider the problem of exploring the revenue/welfare trade-off for deterministic mechanisms, and show that it is an intractable problem in general, even for two bidders (Theorem 2). Comparing with Proposition 2, this result adds to the recent surge in literature pointing out complexity gaps between randomized and deterministic mechanisms [26, 12, 13, 11]. Randomized mechanisms are of course a powerful and useful analytical concept, but it is deterministic mechanisms and auctions that we are chiefly interested in. Hence such complexity gaps are meaningful and onerous. We also show that there are instances for which the set of Pareto optimal mechanisms has exponential size.

On the positive side, we show that the problem can be solved for two bidders, even for correlated valuations (Theorem 4). By “solved” we mean that any trade-off point can be approximated with arbitrarily high precision in polynomial time in both the input and the precision — that is to say, by an FPTAS. It also means (by results in [28]) that an approximate summary of the trade-off (the ϵ -Pareto curve), of polynomial size (Figure 1(e)), can be computed in polynomial time. The derivation of the two-bidders algorithm (see Section 4.1) is quite involved. We first find a pseudo-polynomial dynamic programming algorithm for the problem of finding a mechanism with welfare (resp. revenue) *exactly* a given number. This algorithm is very different from the one in [26] for optimal auctions in the two bidder case, but it exploits the same feature of the problem, namely its planar nature. We then recall Theorem 4 of [28] (Section 2) which establishes a connection between such pseudo-polynomial algorithms for the exact problems and FPTAS for the trade-off problem. However, the present problem violates several key assumptions of that theorem, and a custom reduction to the exact problem is needed.

Unfortunately for three or more bidders the above approach no longer works; this is not surprising since, as it was recently shown in [26], just maximizing revenue is an APX-hard problem in the correlated case. The main problem left open in this work is whether there is an FPTAS for three or more bidders with *independent* valuation distributions.

We also look at another interesting case of the n -bidder problem, in which the valuation distributions have support two. This case is of some methodological interest because, in general, n -dimensional problems of this sort in mechanism design have not been characterized computationally, because of the difficulty related to the exponential size of the solution sought; binary-valued bidders have served as a first step towards

the understanding of auction problems in the past, for example in the study of optimal *multi-object* auctions [3]. We show that the trade-off problem is in PSPACE and (weakly) NP-hard (Theorem 5).

1.2 Related work

Although [19] appears to be the only previous paper explicitly treating optimal auction design as a multi-objective optimization problem, there has been substantial work in studying the relation of the two objectives. The most prominent paper in the area is that of Bulow and Klemperer [4] who show that the revenue benefits of adding one extra bidder and running the efficiency-maximizing auction surpasses those of running the revenue-maximizing auction. In [2] the authors show that for valuations drawn independently from the same monotone hazard rate distribution, an analogous theorem holds for efficiency: by adding $\Theta(\log n)$ extra bidders and running Myerson’s auction, one gets at least the efficiency of Vickrey’s auction. This paper also shows that for these distributions both the welfare and the revenue ratios between Vickrey and Myerson’s auctions are bounded by $1/e$: in our terms this implies that the extreme points of the Pareto curve lie within a constant factor of each other and so constant factor approximations are trivial; we note that no such constant ratios are known for more general distributions (not even for the case of regular distributions), assuming of course that the ratio between all bidders’ maximum and minimum valuation is arbitrary. This kind of revenue and welfare ratios are also studied in [29] for keyword auctions (multi-item auctions), and in [24] for single-item english auctions and valuations drawn from a distribution with bounded support. In [1] the authors present some tight bounds for the efficiency loss of revenue-optimal mechanisms, which depend on the number of bidders and the size of the support. Finally, and very recently, [7] gives simple auctions (in particular, second-price auctions with appropriately chosen reserve prices) that simultaneously guarantee a 20% fraction of both the optimal revenue and the optimal social welfare, when bidders’ valuations are drawn independently from (possibly different) regular distributions: in multiobjective optimization parlance, their auctions belong to the *knee* of the Pareto curve. In this work (Section 4) we provide an algorithm for approximating *any* point of the Pareto curve within arbitrary precision, albeit sacrificing the simplicity of the auction format.

2 Preliminaries

2.1 Bayesian Mechanism Design

We are interested in auctioning a single, indivisible item to n bidders. We assume every bidder i has a private valuation v_i for the item and that her valuation is drawn from some discrete probability distribution over support of size h_i with probability density function $f_i(\cdot)$. We use v_i^k and f_i^k , $k = 1, \dots, h_i$, to denote the k -th smallest element in the support of bidder i and its probability mass respectively.

Formally an auction consists of an allocation rule $x_i(v_1, \dots, v_n)$, the probability of bidder i getting allocated the item, and a payment rule $p_i(v_1, \dots, v_n)$ which is the price paid by bidder i . In this paper we focus our attention on deterministic mechanisms

so that $x_i(\cdot) \in \{0, 1\}$. We demand from our auctions to satisfy the two standard constraints of ex-post incentive compatibility (IC) and individual rationality (IR); it is well known [25] that any such auction has the following special form: if we fix the valuation of all bidders except for bidder i , then there is a threshold value $t_i(v_{-i})$, such that bidder i only gets the item for values $v_i \geq t_i(v_i)$ and pays $t_i(v_{-i})$. In particular one can show that, for the discrete setting and for the objectives of welfare and revenue we are interested in, we can wlog assume that the threshold values t_i of any Pareto optimal auction will always be on the support of bidder i .

Relying on the above characterization, we will describe our mechanisms using the concept of an *allocation matrix* A : a $h_1 \times \dots \times h_n$ matrix where entry (i_1, \dots, i_n) corresponds to the tuple $(v_1^{i_1}, \dots, v_n^{i_n})$ of bidder's valuations. Each entry takes values from $\{0, 1, \dots, n\}$ indicating which bidder gets allocated the item for the given tuple of valuations, with 0 indicating that the auctioneer keeps the item. In order for an allocation matrix to correspond to a valid (ex-post IC and IR) auction a necessary and sufficient condition is the following *monotonicity constraint*: if $A[i_1, \dots, i_j, \dots, i_n] = j$ then $A[i_1, \dots, k, \dots, i_n] = j$ for all $k \geq i_j$. Notice that the payment of the bidder who gets allocated the item can be determined as the least value in his support for which he still gets the item, keeping the values of the other bidders fixed; moreover, when there is only a constant number of bidders, the allocation matrix provides a polynomial representation of an auction.

2.2 Multi-Objective Optimization

Trade-offs are present everywhere in life and science — in fact, one can argue that optimization theory studies the very special and degenerate case in which we happen to be interested in only one objective. There is a long research tradition of *multi-objective* or *multi-criterion optimization*, developing methodologies for computing the trade-off points (called the *Pareto set*) of optimization problems with many objectives, see for example [18, 14, 20]. However, there is a computational awkwardness about this problem: Even for simple cases, such as bicriterion shortest paths, the Pareto set (the set of all undominated feasible solutions) can be exponential, and thus it can never be polynomially computed. In 2000, Papadimitriou and Yannakakis [28] identified a sense in which this is a meaningful problem: They showed that there is *always* a set of solutions of polynomial size that are *approximately* undominated, within arbitrary precision; a multi-objective problem is considered tractable if such a set can be computed in polynomial time. Since then, much progress has been made in the algorithmic theory of multi-objective optimization [30, 10, 9, 17, 6, 5, 8], and much methodology has been developed, some of which has been applied to mechanism design before [16]. In this paper we use this methodology for studying Bayesian auctions under the two criteria of expected revenue and social welfare.

The BI-CRITERION AUCTION problem. We want to design deterministic auctions that perform favorably with respect to (expected) social welfare, defined as $SW = \mathbb{E}[\sum_i x_i v_i]$ and (expected) revenue, defined as $Rev = \mathbb{E}[\sum_i p_i]$. Based on the aforementioned characterization with allocation matrices, we can view an auction as a feasible solution to a combinatorial problem. An instance specifies the number n of bidders

and for each bidder its distribution on valuations. The size of the instance is the number of bits needed to represent these distributions. We map solutions (mechanisms) to points (x, y) in the plane, where we use the x -axis for the welfare and the y -axis for the revenue. The objective space is the set of such points.

Let $p, q \in \mathbb{R}_+^2$. We say that p dominates q if $p \geq q$ (coordinate-wise). We say that p ϵ -covers q ($\epsilon \geq 0$) if $p \geq q/(1+\epsilon)$. Let $A \subseteq \mathbb{R}_+^2$. The Pareto set of A , denoted by $P(A)$, is the subset of undominated points in A (i.e. $p \in P(A)$ iff $p \in A$ and no other point in A dominates p). We say that $P(A)$ is *convex* if it contains no points that are dominated by convex combinations of other points. Given a set $A \subseteq \mathbb{R}_+^2$ and $\epsilon > 0$, an ϵ -Pareto set of A , denoted by $P_\epsilon(A)$, is a subset of points in A that ϵ -cover all vectors in A . Given two mechanisms M, M' we define domination between them according to the 2-vectors of their objective values. This naturally defines the Pareto set and approximate Pareto sets for our auction setting.

As shown in [28], for every instance and $\epsilon > 0$, there exists an ϵ -Pareto set of polynomial size. The issue is one of efficient computability. There is a simple necessary and sufficient condition, which relates the efficient computability of an ϵ -Pareto set to the following *GAP Problem*: given an instance I , a (positive rational) 2-vector $b = (W_0, R_0)$, and a rational $\delta > 0$, either return a mechanism M whose 2-vector dominates b , i.e. $\text{SW}(M) \geq W_0$ and $\text{Rev}(M) \geq R_0$, or report that there does *not* exist any mechanism that is better than b by at least a $(1 + \delta)$ factor in both coordinates, i.e. such that $\text{SW}(M) \geq (1 + \delta) \cdot W_0$ and $\text{Rev}(M) \geq (1 + \delta) \cdot R_0$. There is an FPTAS for constructing an ϵ -Pareto set iff there is an FPTAS for the GAP Problem [28].

Remark 1. Even though our exposition focuses on discrete distributions, our results easily extend to continuous distributions as well. As in [26], given a sufficiently smooth continuous density (say Lipschitz-continuous), whose support lies in a finite interval $[v, \bar{v}]$,³ we can appropriately discretize (while preserving the optimal values within $O(\epsilon)$) and run our algorithms on the discrete approximations.

From exact to bi-criterion. We will make essential use of a result from [28] reducing the multi-objective version of a linear optimization problem A to its exact version: Let A be a discrete linear optimization problem whose objective function(s) have *non-negative* coefficients. The *exact version* of a A is the following problem: Given an instance x of A , and a positive rational C , is there a feasible solution with objective function value *exactly* C ? For such problems, a pseudo-polynomial algorithm for the exact version of implies an FPTAS for the multi-objective version:

Theorem 1 ([28]). *Let A be a linear multi-objective problem whose objective functions have non-negative coefficients: If there exists a pseudo-polynomial algorithm for the exact version of A , then there exists an FPTAS for constructing an approximate Pareto curve for A .*

To obtain our main algorithmic result (Theorem 4), we design a pseudo-polynomial algorithm for the exact version of the BI-CRITERION AUCTION problem and apply Theorem 1 to deduce the existence of an FPTAS. However, it is not obvious why BI-CRITERION AUCTION satisfies the condition of the theorem, since in the standard representation of the problem as a linear problem, the objective functions typically have

³ This is the standard approach in economics, see for example [22].

negative coefficients. We show however (Lemma 2) that there exists an alternate representation with monotonic linear functions.

3 The complexity of Pareto optimal auctions

Our main result in this section is that – in contrast with randomized auctions – designing deterministic Pareto optimal auctions under welfare and revenue objectives is an intractable problem; in particular, we show that, even for 2 bidders⁴ whose distributions are independent and regular, the problem of maximizing one criterion subject to a lower bound on the other is (weakly) NP-hard.

Theorem 2. *For two bidders with independent regular distributions, it is NP-hard to decide whether there exists an auction with welfare at least W and revenue at least R .*

Proof (Sketch). Due to space constraints, in this version of the paper we only provide the reduction for the exact problem for the welfare objective; quite simple and intuitive, it also captures the main idea in the (significantly more elaborate) proof for the bi-criterion problem, which can be found in the full version.

The reduction is from the Partition problem: we are given a set $B = \{b_1, \dots, b_k\}$ of k positive integers, and we wish to determine whether it is possible to partition B into two subsets with equal sum. We assume that $b_i \geq b_{i+1}$ for all i . Consider the rescaled values $b'_i := b_i / (10k \cdot T)$, where $T = \sum_{i=1}^k b_i$, and the set $B' = \{b'_1, \dots, b'_k\}$. It is clear that there exists a partition of B iff there exists a partition of B' .

We construct an instance of the auction problem with two bidders whose independent valuations v_r (row bidder) and v_c (column bidder) are uniformly distributed over supports of size k . (To avoid unnecessary clutter in the expressions, we assume w.l.o.g – by linearity – that the “probability mass” of all elements in the support is equal to 1, as opposed to $1/k$.) The valuation distribution for the row bidder is supported on the set $\{1, 2, \dots, k\}$, while the column bidder’s valuation comes from the set $\{1 + b'_1, 2 + b'_2, \dots, k + b'_k\}$. Since $b'_i \geq b'_{i+1}$ and $\sum_{i=1}^k b'_i = 1/(10k)$, it is straightforward to verify that both distributions are indeed regular (the proof is deferred to the full version).

The main idea of the proof is this: appropriately *isolate* a subset of 2^k feasible mechanisms whose welfare values encode the sum of values $\sum_{i \in S} b'_i$ for all possible subsets $S \subseteq [k]$. The existence of a mechanism with a specified welfare value would then reveal the existence of a partition. Formally, we prove that there exists a Partition of B' iff there exists a feasible mechanism M^* with (expected) welfare

$$\text{SW}(M^*) = (2/3) \cdot (k-1)k(k+1) + (1/2) \cdot k(k+1) + \sum_{i=2}^k (i-1)b'_i + 1/(20k) \quad (1)$$

Consider the allocation matrix of a feasible mechanism. Recall that a mechanism is feasible iff its allocation matrix satisfies the monotonicity constraint. The main claim is that *all mechanisms that could potentially satisfy (1) must allocate the item to the highest bidder, except potentially for the outcomes $(v_r = i, v_c = i + b'_i)$ (i.e. the ones*

⁴ Note that for a single bidder, one can enumerate all feasible mechanisms in linear time.

corresponding to entries on the secondary diagonal of the matrix) when the item can be allocated to either bidder. Denote by \mathcal{R} the aforementioned subclass of mechanisms. The above claim follows from the next lemma, which shows that mechanisms in \mathcal{R} maximize welfare (see full version for the proof):

Lemma 1. *We have $\max_{M \notin \mathcal{R}} \text{SW}(M) < \min_{M \in \mathcal{R}} \text{SW}(M) < \text{SW}(M^*)$.*

To complete the proof, observe that all 2^k mechanisms in \mathcal{R} satisfy monotonicity, hence are feasible. Also note that there is a natural bijection between subsets $S \subseteq [k]$ and these mechanisms: we include i in S iff on input $(v_r = i, v_c = i + b'_i)$ the item is allocated to the column bidder. Denote by $M(S)$ the mechanism in \mathcal{R} corresponding to subset S under this mapping; we will compute the welfare of $M(S)$. Note that the contribution of each entry of the allocation matrix (input) to the welfare equals the valuation of the bidder who gets the item for that input. By the definition of \mathcal{R} , for the entries above the secondary diagonal, the row bidder gets the item (since her valuation is strictly larger than that of the column bidder – this is evident since $\max_i b'_i < 1/(10k)$). Therefore, the contribution of these entries to the welfare equals $\sum_{i=2}^k i(i-1) = (1/3)(k-1)k(k+1)$. Similarly, for the entries below the diagonal, the column bidder gets the item and their contribution to the welfare is $\sum_{i=2}^k (i+b'_i)(i-1) = (1/3)(k-1)k(k+1) + \sum_{i=2}^k (i-1)b'_i$. Finally, for the diagonal entries, if $S \subseteq [k]$ is the subset of indices for which the column bidder gets the item, the welfare contribution is $\sum_{i \in S} (i + b'_i) + \sum_{i \in [k] \setminus S} i = k(k+1)/2 + \sum_{i \in S} b'_i$. Hence, we have:

$$\text{SW}(M(S)) = (2/3) \cdot (k-1)k(k+1) + (1/2) \cdot k(k+1) + \sum_{i=2}^k (i-1)b'_i + \sum_{i \in S} b'_i \quad (2)$$

Recalling that $\sum_{i=1}^k b'_i = 1/(10k)$, (1) and (2) imply that there exists a partition of B' iff there exists a feasible mechanism satisfying (1). This completes the proof sketch. (See the full version of the paper for the much more elaborate proof of the general case.) \square

We can also prove that the size of the Pareto curve can be exponentially large (in other words, the problem of computing the entire curve is exponential even if $P = NP$). The construction is given in the full version.

Theorem 3. *There exists a family of two-bidder instances for which the size of the Pareto curve for BI-CRITERION AUCTION grows exponentially.*

4 An FPTAS for 2 bidders

In this section we give our main algorithmic result:

Theorem 4. *For two bidders, there is an FPTAS to approximate the Pareto curve of the BI-CRITERION AUCTION problem, even for arbitrarily correlated distributions.*

In the proof, we design a pseudo-polynomial algorithm for the exact version of the problem (for both the welfare and revenue objectives) and then appeal to Theorem 1. There is a difficulty, however, in showing that the problem satisfies the assumptions of Theorem 1, because in the most natural linear representation of the problem, the

coefficients for revenue, coinciding with the virtual valuations, may be negative, thus violating the hypothesis of Theorem 1.

We use the following alternate representation: Instead of considering the contribution of each entry (bid tuple) of the allocation matrix separately, we consider the revenue and welfare resulting from all the *single-bidder mechanisms* (pricings) obtained by fixing the valuation of the other bidder.

Definition 1. Let $r_1^{i_1, i_2}$ and $w_1^{i_1, i_2}$ be the (contribution to the) revenue and welfare from bidder 1 of the pricing which offers bidder 1 a price of $v_1^{i_1}$ when bidder 2's value is $v_2^{i_2}$: $r_1^{i_1, i_2} = \sum_{j \geq i_1} v_1^{i_1} \cdot f(j, i_2)$ and $w_1^{i_1, i_2} = \sum_{j \geq i_1} v_1^j \cdot f(j, i_2)$, where $f(\cdot, \cdot)$ is the joint (possibly non-product) valuation distribution. (The quantities $r_2^{i_1, i_2}$ and $w_2^{i_1, i_2}$ are defined analogously.)

Lemma 2. The BI-CRITERION AUCTION problem can be expressed in a way that satisfies the conditions of Theorem 1.

Proof. We consider variables $x_{ij}, y_{ij}, i \in [h_1], j \in [h_2]$. The x_{ij} 's are defined as follows: $x_{ij} = 1$ iff $A[i, j] = 1$ and $A[i', j] \neq 1$ for all $i' < i$. I.e. $x_{ij} = 1$ iff the (i, j) -th entry of A is allocated to bidder 1 and, for this fixed value of j , i is the smallest index for which bidder 1 gets allocated; symmetrically, $y_{ij} = 1$ iff $A[i, j] = 2$ and $A[i, j'] \neq 2$ for all $j' < j$. It is easy to see that the feasibility constraints are linear in these variables. We can also express the objectives as linear functions with non-negative coefficients as follows:

$$\begin{aligned} \text{Rev}(x, y) &= \sum_{i=1}^{h_1} \sum_{j=1}^{h_2} x_{ij} r_1^{i,j} + \sum_{i=1}^{h_1} \sum_{j=1}^{h_2} y_{ij} r_2^{i,j} \\ \text{SW}(x, y) &= \sum_{i=1}^{h_1} \sum_{j=1}^{h_2} x_{ij} w_1^{i,j} + \sum_{i=1}^{h_1} \sum_{j=1}^{h_2} y_{ij} w_2^{i,j} \end{aligned}$$

□

4.1 An algorithm for the exact version of BI-CRITERION AUCTION

The main idea behind our algorithm, inspired by the characterization of Lemma 2, is to consider the contribution from each bidder (fixing the value of the other) independently, by going over all (linearly many) single-bidder mechanisms for both bidders. The challenging part is to combine the individual single-bidder mechanisms into a single two-bidder mechanism and to this end we employ dynamic programming:

Assume that both bidders have valuations of support size h ; the subproblems we consider in our dynamic program correspond to settings where we condition that the valuation of each bidder is drawn from an upwards closed subset of his original support. Formally, let $M[i, j, W]$ be True iff there exists an auction that uses the valuations (v_1^i, \dots, v_1^h) and (v_2^j, \dots, v_2^h) and has welfare exactly W . In what follows $N_{i,j}$ is the normalization factor for valuations (jointly) drawn from (v_1^i, \dots, v_1^h) and (v_2^j, \dots, v_2^h) , namely $N_{i,j} = \sum_{k \geq i, l \geq j} f(v_1^k, v_2^l)$.

Lemma 3. *We can update the quantity $M[i, j, W]$ as follows:*

$$\begin{aligned}
M[i, j, W] = & \bigvee_{k \geq j} M[i + 1, j, (W \cdot N_{i,j} - w_2^{i,k}) \cdot N_{i+1,j}^{-1}] \\
& \vee \bigvee_{k \geq i} M[i, j + 1, (W \cdot N_{i,j} - w_1^{k,j}) \cdot N_{i,j+1}^{-1}] \\
& \vee \bigvee_{\substack{k > i \\ l > j}} M[i + 1, j + 1, (W \cdot N_{i,j} - w_1^{k,j} - w_2^{i,l}) \cdot N_{i+1,j+1}^{-1}]
\end{aligned}$$

Proof. Let $A[i \dots h, j \dots h]$ be the allocation matrix of the auction that results from the above update rule, fixing i and j . We start by noting that any allocation matrix A can have one of the following four forms:

- F1:** There exist i' and j' such that $A[i, j'] = 1$ and $A[i', j] = 2$.
- F2:** There exists i' such that $A[i', j] = 2$ but there is no j' such that $A[i, j'] = 1$.
- F3:** There exists j' such that $A[i, j'] = 1$ but there is no i' such that $A[i', j] = 2$.
- F4:** There exist no i' and j' such that $A[i, j'] = 1$ or $A[i', j] = 2$.

Because of monotonicity it follows immediately that no allocation matrix of form F1 can be valid, and the other three forms correspond to the three terms of the recurrence; finally note that for any such form, say F2, the first term of the update rule for $M[i, j, W]$ runs over all possible pricings for bidder 1 (keeping the value of bidder 2 at v_2^j) and checks whether they induce the required welfare. \square

We omit the straightforward details of how the above recurrence can be efficiently implemented as a pseudo-polynomial dynamic programming algorithm. The algorithm for deciding whether there exists an auction with revenue exactly R is identical to the above by simply replacing R (the revenue target value) for W and $r_j^{i_1, i_2}$ for $w_j^{i_1, i_2}$.

5 The case of n bidders

When the number n of bidders is part of the input, the allocation matrix is no longer a polynomially succinct representation of a mechanism. In fact, it is by no means clear whether BI-CRITERION AUCTION is even in NP in this case: we next show that for the case of n binary bidders, the problem is NP -hard and in $PSPACE$:

Theorem 5. *For n binary-valued bidders BI-CRITERION AUCTION is (weakly) NP -hard and in $PSPACE$.*

Proof (Sketch). For simplicity, we prove both results for the exact version of the problem for welfare; the bi-objective case follows by a straightforward but tedious generalization.

The NP -hardness reduction is from Partition. Let $B = \{b_1, \dots, b_k\}$ be a set of positive rationals; we can assume by rescaling that $\sum_{i=1}^k b_i = 1/100$. We construct an instance of the auction problem as follows: there are k bidders, with uniform distributions (again we will assume unit masses for simplicity) over the following supports

$\{l_i, h_i\}, i = 1 \dots n$, where $l_i < h_i$. We set $l_i = b_i$ and demand that $\{h_i\}_{i=1, \dots, n}$ forms a super-increasing sequence (i.e. $h_{i+1} > \sum_{j=1}^i h_j$), with $h_1 > \max_i b_i$. The claim is that there exists a partition of B iff there exists an auction with welfare equal to $\sum_{i=1}^k h_i + (1/2) \sum_{i=1}^k b_i$. To see this notice that – since the sequence $\{h_i\}_{i=1, \dots, n}$ is super-increasing – any mechanism with the above welfare value must allocate to bidder i for *exactly* one valuation tuple (v_i, v_{-i}) where $v_i = h_i$; the corresponding contribution to the welfare from this case is h_i . Monotonicity then implies that this auction can allocate to bidder i for *at most* one valuation tuple (v_i, v_{-i}) where $v_i = l_i$; the corresponding contribution to the welfare from this case is b_i . We therefore get a bijection between subsets of B and mechanisms, by including an element b_i in the set S iff bidder i gets allocated the item for some valuation tuple (v_i, v_{-i}) where $v_i = l_i$, and the claim follows.

For the PSPACE upper bound, we start by noting that the problem of computing an auction with welfare (or revenue) *exactly* W , can be formulated as the problem of computing a matching of weight exactly W in a particular type of bipartite graphs (first pointed out in [12], see also the full version of the paper) with a number of nodes that is exponential in the number of bidders. The EXACT MATCHING problem is known to be solvable in RNC [21]; since our input provides an exponentially succinct representation of the constructed graph, we are interested in the so-called *succinct version* of the problem [15, 27]. By standard techniques, the succinct version of EXACT MATCHING in our setting is solvable in PSPACE, and the theorem follows. \square

We conjecture the above upper bound to be tight (i.e. the problem is actually PSPACE-complete) even for n bidders with arbitrary supports.

6 Open Questions

Is there an FPTAS for 3 bidders? We conjecture that there is, and in fact for any constant number of bidders. Of course, the approach of our FPTAS for 2 bidders cannot be generalized, since it works for the correlated case, which is APX-complete for 3 or more bidders. We have derived two different dynamic programming-based PTAS's for the uncorrelated problem, but so far, despite a hopeful outlook, we have failed to generalize them to 3 bidders. Finally, we conjecture that for n bidders the problem is significantly harder, namely PSPACE-complete and inapproximable.

On a different note, it would be interesting to see if we can get better approximations for some special types of distributions; we give one such type of result in the full version of the paper. Are there improved approximation guarantees for more general kinds of distributions and n bidders?

References

- [1] V. Abhishek and B. E. Hajek. Efficiency loss in revenue optimal auctions. In *CDC*, 2010.
- [2] G. Aggarwal, G. Goel, and A. Mehta. Efficiency of (revenue-)optimal mechanisms. In *EC*, 2009.
- [3] M. Armstrong. Optimal multi-object auctions. *Review of Economic Studies*, 67(3):455–81.

- [4] J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–94, March 1996.
- [5] C. Chekuri, J. Vondrak, and R. Zenklusen. Multi-budgeted matchings and matroid intersection via dependent rounding. In *SODA*, 2011.
- [6] C. Daskalakis, I. Diakonikolas, and M. Yannakakis. How good is the chord algorithm? In *SODA*, 2010.
- [7] C. Daskalakis and G. Pierrakos. Simple, optimal and efficient auctions. In *WINE*, 2011.
- [8] I. Diakonikolas. *Approximation of Multiobjective Optimization Problems*. PhD thesis, Columbia University, 2010.
- [9] I. Diakonikolas and M. Yannakakis. Succinct Approximate Convex Pareto Curves. In *SODA*, 2008.
- [10] I. Diakonikolas and M. Yannakakis. Small approximate pareto sets for biobjective shortest paths and other problems. *SIAM J. Comput.*, 39:1340–1371, 2009.
- [11] S. Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In *STOC*, 2011.
- [12] S. Dobzinski, H. Fu, and R. D. Kleinberg. Optimal auctions with correlated bidders are easy. In *STOC*, 2011.
- [13] S. Dughmi, T. Roughgarden, and Q. Yan. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. In *STOC*, 2011.
- [14] M. Ehrgott. *Multicriteria optimization*. Springer-Verlag, 2005.
- [15] H. Galperin and A. Wigderson. Succinct representations of graphs. *Information and Control*, 56(3):183–198, 1983.
- [16] F. Grandoni, P. Krysta, S. Leonardi, and C. Ventre. Utilitarian mechanism design for multi-objective optimization. In *SODA*, 2010.
- [17] F. Grandoni, R. Ravi, and M. Singh. Iterative rounding for multi-objective optimization problems. In *ESA*, 2009.
- [18] E. J. Climacao. *Multicriteria Analysis*. Springer-Verlag, 1997.
- [19] A. Likhodedov and T. Sandholm. Mechanism for optimally trading off revenue and efficiency in multi-unit auctions. In *EC*, 2003.
- [20] K. M. Miettinen. *Nonlinear Multiobjective Optimization*. Kluwer, 1999.
- [21] K. Mulmuley, U. V. Vazirani, and V. V. Vazirani. Matching is as easy as matrix inversion. In *STOC*, 1987.
- [22] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73.
- [23] R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, April 1983.
- [24] Z. Neeman. The effectiveness of english auctions. *Games and Economic Behavior*, 43(2):214–238, May 2003.
- [25] N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. 2007.
- [26] C. H. Papadimitriou and G. Pierrakos. On optimal single-item auctions. In *STOC*, 2011.
- [27] C. H. Papadimitriou and M. Yannakakis. A note on succinct representations of graphs. *Information and Control*, 71(3):181–185, 1986.
- [28] C. H. Papadimitriou and M. Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *FOCS*, 2000.
- [29] T. Roughgarden and M. Sundararajan. Is efficiency expensive? In *3rd Workshop on Sponsored Search*, 2007.
- [30] S. Vassilvitskii and M. Yannakakis. Efficiently computing succinct trade-off curves. *Theoretical Computer Science*, 348:334–356, 2005.