

Optimal Deterministic Auctions with Correlated Priors

Christos Papadimitriou*
UC Berkeley
christos@cs.berkeley.edu

George Pierrakos*
UC Berkeley
georgios@cs.berkeley.edu

Abstract

We revisit the problem of designing the profit-maximizing single-item auction, solved by Myerson in his seminal paper for the case in which bidder valuations are independently distributed. We focus on general joint distributions, seeking the optimal deterministic incentive compatible auction. We give a geometric characterization of the optimal auction, resulting in a duality theorem and an efficient algorithm for finding the optimal deterministic auction in the two-bidder case and an NP-completeness result for three or more bidders.

1 Introduction

Myerson’s paper [21] on optimal auction design is remarkable in several ways. It is not the first important paper on auctions of course [30], but it pioneers the point of view of its title: *auction design*, that is, the exploration and evaluation of a large design space in a mindset that is very much one of computer science. It is elegant and methodologically original and powerful, completely resolving this important and difficult problem and providing insights, lemmata, and methodology that would be useful in related contexts for decades. Myerson considers the auction of an item to bidders with private valuations whose prior distributions are independent and known, and seeks incentive compatible mechanisms that maximize auctioneer revenue. His simple and elegant solution involves the *virtual valuation*, a quantity each bidder can compute from their private valuation and the valuation distribution through a mathematical maneuver called *ironing*.

Myerson leaves open the case in which the valuations are correlated; in subsequent work, Crémer and McLean [8, 7] consider correlated valuations and solve the problem for the case where auctions are only required to be *interim* individually rational. In fact, in this framework the uncorrelated case is a singularity in the sense that in most cases (when the correlation has “full rank” in a certain precise sense) full surplus can be extracted in expectation through appropriate offers of lotteries to the bidders. Despite the elegance of their result, the fact that bidders may be charged merely for participating in the auction–lottery (including losers) has been criticized as rendering the auction impractical [20], especially for settings where agents may easily cancel their participation after the auction is conducted. It is therefore of tantamount importance to consider the question of designing the optimal *ex post* individually rational auction for correlated valuations. This is the question we consider, and in a sense completely answer, in the current paper.

This past decade saw the advent and rise of Algorithmic Game Theory [24, 22], a research tradition which can be seen as a complexity-theoretic critique of Mathematical Economics, with Internet in the backdrop. This point of view has yielded a host of important results and new insights, for example related to the complexity of equilibria [6, 9], the trade-offs between complexity, approximation, and incentive-compatibility in social welfare-maximizing mechanism design [22, 25], and (in an extended sense that includes on-line algorithms as a part of complexity theory) the price of anarchy [19, 29, 28]. However, there has been little progress in looking at

*Research (partially) supported by NSF grant CCF-0964033 and by a Google University Research Award.

Bayesian auctions à-la Myerson from this point of view. Ronen [26] came up with a mechanism for the correlated case that achieves half of the optimum revenue, while Ronen and Saberi [27] showed that no “ascending auction” can do better than $7/8$, and they conjectured that all relevant auctions are ascending (incidentally, we disprove this conjecture by showing that the optimal two-bidder auction may not be ascending). Missing from these two papers, however, is a concrete sense of the ways in which this is a difficult problem. We provide this here.

There were several follow-up papers [4, 17, 10], examining the extent to which simple mechanisms can achieve good approximations of Myerson’s optimal mechanism, and motivated by Myerson’s astonishing result that the optimal mechanism for the regular i.i.d. setting is simply a second-price auction with reserve prices. For the most part, these are positive results yielding constant approximations by simple mechanisms in a variety of settings. Another line of work is prior-free mechanism design, where the goal is to design mechanisms that achieve profits comparable to that of some well-behaved benchmark [13]. This direction became especially interesting after Hartline and Roughgarden developed a framework in [16] that is grounded in Bayesian optimal mechanism design allowing one to design mechanisms that simultaneously approximate all Bayesian optimal mechanisms. The intermediate approach of having bidders’ valuations coming from a distribution that is nonetheless *unknown* to the auctioneer has also been considered [10]. Finally the important open problem of designing optimal multi-parameter mechanisms has been addressed for some distributions of preferences, such as additive valuations with budget constraints [2] and unit-demand settings [5] and connections of this problem to the algorithmic problem of optimal pricing have been made [4, 14]; economists have also made some attempts to extend Myerson’s results to multi-item auctions [1, 18].

In this paper we take a complexity-theoretic look at the general, correlated valuations case of Myerson’s single-item auctions. We point out that the optimal auction design problem can be reduced essentially to a maximum weighted independent set problem in a particular graph whose vertices are all possible tuples of valuations (an uncountable set, of course, in the continuous case). If the distribution is discrete, this is an ordinary graph-theoretic problem; no such combinatorial characterization had been known, and this had been the main difficulty in developing an algorithmic and complexity-theoretic understanding of the problem. For discrete distributions, this leads directly to an efficient algorithm in the case of two bidders, where the graph is bipartite, while in the case of three or more bidders NP-completeness (in fact, inapproximability) prevails. For continuous distributions, we prove a duality characterization through a Monge-Kantorovich-like problem [12], and from this a fully polynomial approximation scheme for two bidders when the distribution is continuous enough and accessible through an oracle — plus, as an aside, a $2/3$ approximation for three bidders, improving the previously best known approximation of [26].

Our results rest on a geometric characterization of optimal deterministic auctions. An important element of our proof is the so-called *marginal profit contribution* function; it bears some similarities to Myerson’s virtual valuation function [21], the most important of them being that they both admit a marginal revenue interpretation in the spirit of [3]. However, despite their similarities and their somewhat common derivation, marginal profit contribution functions are different from Myerson’s virtual valuations in a number of ways: they only take positive values, they are not necessarily monotone and they do not admit a natural interpretation as valuations in some modified domain. One important ingredient of Myerson’s approach to the design of optimal auctions is an analytical maneuver he calls *ironing*; Myerson uses ironing to transform a potentially non-monotone allocation rule into a monotone one, without hurting revenue. Our approach circumvents ironing by restricting the space of mechanisms explored; we achieve that by imposing an additional technical condition which limits the design space into a subset of all mechanisms, but one which still contains all the optimal mechanisms.

The work most related to ours is that of Dobzinski, Fu and Kleinberg [11], who also study the problem of designing the optimal auction for the correlated setting. They obtain a collection of very interesting results, which however are quite complementary to ours: Based, among others, on insights from [26], they arrive at efficient algorithms for computing the optimal randomized mechanism that is truthful in expectation, and a constant factor approximation of the optimal deterministic auction for any number of bidders.

2 Preliminaries

Imagine n bidders seeking an indivisible good offered in auction. We assume that each bidder has a private valuation v_i for the item and that bidders' valuations are drawn from some joint distribution over $[0, 1]^n$ whose density function we denote by $\phi(\mathbf{v})$. We consider both discrete and continuous ϕ . Discrete distributions are the source of the combinatorial insights underlying our approach, while continuous distributions provide continuity with the spirit and methodology of Myerson's paper, another important source of inspiration. In the continuous case, we follow Myerson in making the analytically convenient assumption that $\phi(x) > 0$ for all $x \in [0, 1]^n$. This is hardly a loss of generality, since a small minimum value on every point can be achieved by changing ϕ very little. In stating an algorithm for the two-dimensional continuous case (Section 5), we shall also assume that ϕ is Lipschitz-continuous and accessible through an oracle (in such a way that, for example, it can be approximately integrated over nice regions).

In the discrete case, let $Sup(\phi)$ denote the finite support of the joint discrete distribution ϕ . Then ϕ is presented as a finite set of $|Sup(\phi)|$ $(n + 1)$ -tuples of the form $(v_1^i, v_2^i, \dots, v_n^i, \phi^i)$, one for each point $v^i \in Sup(\phi)$, where ϕ^i is the probability mass concentrated at the point (v^i) of the support.

We are interested in designing auctions that maximize the auctioneer's profit. Formally, an auction consists of an allocation rule denoted by $x_i(\mathbf{v})$, the probability of bidder i getting allocated the item, and a payment rule denoted by $p_i(\mathbf{v})$ which is the price paid by bidder i . In this paper we focus our attention on deterministic mechanisms so that $x_i(\mathbf{v}) \in \{0, 1\}$. Our goal is to maximize the auctioneer's expected profit $\sum_{i=1}^n \mathbb{E}[\phi(\mathbf{v})p_i(\mathbf{v})]$.

We want our auction to be ex-post incentive compatible (IC) and individually rational (IR), with no positive transfers (NPT). These notions are defined as follows:

$$\text{IC: } \forall i, v_i, v'_i, v_{-i} : v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i}) \geq v_i x_i(v'_i, v_{-i}) - p_i(v'_i, v_{-i})$$

$$\text{IR: } \forall i, v_i, v_{-i} : v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i}) \geq 0$$

$$\text{NPT: } \forall i, v_i, v_{-i} : p_i(v_i, v_{-i}) \geq 0$$

We say that an allocation function $x_i(v_i, v_{-i})$ is *monotone* if $v_i \geq v'_i$ implies $x_i(v_i, v_{-i}) \geq x_i(v'_i, v_{-i})$ for all i, v_{-i} . For the case of deterministic mechanisms that we are interested in monotonicity implies that $x_i(v_i, v_{-i})$ is a step-function. The *threshold value* of such a step-function allocation rule is set to be the minimum winning valuation for every player given the valuations of the other bidders: $t_i(v_{-i}) = \inf\{v_i \in [0, 1] \mid x_i(v_i, v_{-i}) = 1\}$.

The following theorem provides a characterization of mechanisms for the setting we are interested in; a proof can be found in [23].

Theorem 2.1. *A deterministic mechanism satisfies IC, IR and NPT if and only if the following conditions hold.*

1. $x_i(v_i, v_{-i})$ is monotone.
2. For all i, v_i, v_{-i} we have

$$p_i(v_i, v_{-i}) = \begin{cases} t_i(v_{-i}) & \text{if } x_i(v_i, v_{-i}) = 1; \\ 0 & \text{if } x_i(v_i, v_{-i}) = 0. \end{cases}$$

3 The Geometry of Optimal Auctions

Here we focus on the two-bidder case, and provide an alternative geometric interpretation of the mechanism design problem as a space partitioning problem. Our characterization holds for any number of bidders, with the appropriate generalizations and modifications; however we only address the multi-dimensional case in Section 6, where we will use our geometric characterization to establish the inapproximability of the problem for 3 or more bidders.

We start by noting that the allocation function can be described in terms of a partition of the unit square (the space of all possible valuation pairs) into three regions: In region A bidder 1 gets the item, in region B bidder 2 gets the item and in region C neither gets the item. The shape of those regions is restricted by monotonicity as follows (see Figure 1): Region

A is *rightward closed*, meaning that $(x, y) \in A$ and $x' > x$ implies $(x', y) \in A$, while region B is, analogously, *upward closed*. These regions are captured by their boundaries: Region A 's boundary is a function $\alpha : [0, 1] \mapsto [0, 1]$ where for all $y \in [0, 1]$ $\alpha(y) = \inf\{x : (x, y) \in A\}$, and similarly for region B and its boundary $\beta(x)$.




Figure 1: A pair of valid allocation rules.

Definition 3.1. A **valid allocation pair** (α, β) is a pair of functions from $[0, 1]$ to itself satisfying the non-crossing property: for all points $(x, y) \in [0, 1]^2$ we have $y \geq \beta(x) \Rightarrow x \leq \alpha(y)$.

Notice that it is not necessary for the functions to be monotone; the monotonicity property of the allocation is ensured through the fact that α and β are proper functions and therefore regions A and B are rightward and upward closed respectively. The non-crossing property ensures that for any bid pair (x, y) at most one bidder is allocated the item.

In our proof we will make extensive use of the following notion of marginal profit.

Definition 3.2. Let f (respectively, g) be the marginal profit contribution of a bid pair (x, y) for player 1 (resp. 2) defined as:

$$f(x, y) = -\frac{\partial}{\partial x} \left[\max_{x' \geq x} x' \cdot \int_{x'}^1 \phi(t, y) dt \right], g(x, y) = -\frac{\partial}{\partial y} \left[\max_{y' \geq y} y' \cdot \int_{y'}^1 \phi(x, t) dt \right]$$

wherever the derivative is defined, and is extended to the full range by right continuity.

Intuitively, $f(x, y) dx dy$ is the added expected profit obtained from including the infinitesimal area $dx dy$ to A , that is, deciding to give the item to the first bidder if the valuations are x, y ; in Remark 3.5 we further discuss the intuition behind these functions and their relation to Myerson's virtual valuation functions.

Definition 3.3. Call a valid allocation pair **proper** if it satisfies the following condition:

$$\left. \begin{aligned} \alpha(y) \cdot \int_{\alpha(y)}^1 \phi(t, y) dt &= \max_{x' \geq \alpha(y)} x' \cdot \int_{x'}^1 \phi(t, y) dt, \quad \text{for all } y \\ \beta(x) \cdot \int_{\beta(x)}^1 \phi(x, t) dt &= \max_{y' \geq \beta(x)} y' \cdot \int_{y'}^1 \phi(x, t) dt, \quad \text{for all } x \end{aligned} \right\} \quad (1)$$

Marginal profit contribution functions provide us with an alternative way to express the objective of expected profit.

Lemma 3.4. Let (α, β) be a proper valid allocation pair. Then the expected profit of any auction with payments defined as in Theorem 2.1 is:

$$\int_0^1 \int_{\alpha(y)}^1 f(x, y) dx dy + \int_0^1 \int_{\beta(x)}^1 g(x, y) dy dx$$

Proof. Let $p_1(x, y), p_2(x, y)$ be the payment functions induced by the allocation rule (α, β) according to Theorem 2.1. Then the expected profit of our auction is:

$$\begin{aligned}
& \int_0^1 \int_0^1 p_1(x, y) \phi(x, y) dx dy + \int_0^1 \int_0^1 p_2(x, y) \phi(x, y) dx dy \\
&= \int_0^1 \left[\alpha(y) \cdot \int_{\alpha(y)}^1 \phi(t, y) dt \right] dy + \int_0^1 \left[\beta(x) \cdot \int_{\beta(x)}^1 \phi(x, t) dt \right] dx \\
&= \int_0^1 \left[\max_{x' \geq \alpha(y)} x' \cdot \int_{x'}^1 \phi(t, y) dt \right] dy + \int_0^1 \left[\max_{y' \geq \beta(x)} y' \cdot \int_{y'}^1 \phi(x, t) dt \right] dx \\
&= \int_0^1 \int_{\alpha(y)}^1 f(x, y) dx dy + \int_0^1 \int_{\beta(x)}^1 g(x, y) dy dx
\end{aligned}$$

where in the first equality we used the characterization of truthful payments as every player's critical value, in the second equality we made use of condition (1) and in the last equality we made use of the definition of marginal profit contribution functions. \square

Remark 3.5. *There is an intuitive connection between the marginal profit contribution and Myerson's (ironed) virtual valuations [21]. As pointed out in the introduction, even though these functions are not identical, the goal in both cases is the same: to capture some notion of marginal revenue. Indeed the quantity $x \cdot \int_x^1 \phi(t, y) dt$ corresponds to the expected profit of the auctioneer from an agent when she is offered a price of x , keeping the value of the other player fixed at y ; equivalently one can write the expected revenue as a function $R(q)$ of the probability of sale $q = 1 - \int_x^1 \phi(t, y) dt$ (for a thorough discussion of this maneuver the reader is referred to [15]). Myerson's virtual valuations then correspond to the derivative of the function $R(q)$, and ironing corresponds to taking the convex hull $\hat{R}(q)$ of $R(q)$ and then differentiating. The "ironing" in our case corresponds to the action of taking the derivative of $R'(q) = \max_{q' \leq q} R(q')$ instead of $R(q)$, which is intuitively the "skyline" one would see from $q = 0$. In Figure 2 we show an example of a revenue curve $R(q)$ and its corresponding $\hat{R}(q)$ and $R'(q)$ curves.*

It has been already noted by Myerson [21] that the allocation rule will remain invariant across ironed regions: in his case this is a consequence of his proof technique involving ironing. In our case we explicitly demand that the allocation is invariant by imposing Condition 1. Since the invariance of the allocation rule across ironed regions follows from Myerson's analysis as well, Condition 1 is indeed not a loss of generality in the sense that all optimal auctions satisfy it. We prove this formally next.

The next lemma establishes that without loss of generality we can restrict ourselves to proper allocation pairs. Let $Profit(\alpha, \beta, \phi)$ denote the profit of an auction with allocation curves (α, β) , when the joint distribution of valuations is ϕ .

Lemma 3.6. *For any ϕ and for any valid allocation pair (α, β) there is a proper valid pair (α', β') such that $Profit(\alpha', \beta', \phi) \geq Profit(\alpha, \beta, \phi)$*

Proof. For the sake of contradiction suppose this is not true, i.e. $Profit(\alpha', \beta', \phi) < Profit(\alpha, \beta, \phi)$ for any proper valid allocation pair (α', β') .

We start by defining the following sets of points:

$$\begin{aligned}
\mathcal{Y} &= \left\{ y \in [0, 1] \mid \alpha(y) \cdot \int_{\alpha(y)}^1 \phi(t, y) dt \neq \max_{x' \geq \alpha(y)} x' \cdot \int_{x'}^1 \phi(t, y) dt \right\} \\
\mathcal{X} &= \left\{ x \in [0, 1] \mid \beta(x) \cdot \int_{\beta(x)}^1 \phi(x, t) dt \neq \max_{y' \geq \beta(x)} y' \cdot \int_{y'}^1 \phi(x, t) dt \right\}
\end{aligned}$$

as the set of all coordinates y (resp. x) where condition (1) is violated by function α (resp. β). Consider now the auction defined by the following allocation curves:

$$\alpha'(y) = \begin{cases} \alpha(y) & , \text{ if } y \notin \mathcal{Y}; \\ \arg \max_{x'} \{x' \int_{x'}^1 \phi(t, y) dt \mid x' \geq \alpha(y)\} & , \text{ if } y \in \mathcal{Y}. \end{cases}$$

$$\beta'(x) = \begin{cases} \beta(x) & , \text{ if } x \notin \mathcal{X}; \\ \arg \max_{y'} \{y' \int_{y'}^1 \phi(x, t) dt \mid y' \geq \beta(x)\} & , \text{ if } x \in \mathcal{X}. \end{cases}$$

where –in the case of ties– $\arg\{\cdot\}$ returns the largest y or x respectively. By construction, the new pair (α', β') satisfies condition (1). In what follows we claim that the resulting allocation pair (α', β') is also valid and moreover it has greater revenue than (α, β) , thus reaching a contradiction.

The monotonicity property of the allocation is satisfied since α' and β' are proper functions of y and x respectively. The non-crossing property follows from the non-crossing property of α and β and the fact that $\alpha'(y) \geq \alpha(y)$ for all $y \in [0, 1]$ and $\beta'(x) \geq \beta(x)$ for all $x \in [0, 1]$. Finally, for the profit of the two auctions defined by the allocation curves (α, β) and (α', β') we have:

$$\begin{aligned} & Profit(\alpha', \beta', \phi) \\ &= \int_0^1 \left[\alpha'(y) \cdot \int_{\alpha'(y)}^1 \phi(t, y) dt \right] dy + \int_0^1 \left[\beta'(x) \cdot \int_{\beta'(x)}^1 \phi(x, t) dt \right] dx \\ &\geq \int_0^1 \left[\alpha(y) \cdot \int_{\alpha(y)}^1 \phi(t, y) dt \right] dy + \int_0^1 \left[\beta(x) \cdot \int_{\beta(x)}^1 \phi(x, t) dt \right] dx \\ &= Profit(\alpha, \beta, \phi), \end{aligned}$$

where the inequality follows from the definition of α' and β' . The lemma now follows. \square

Denote now by \mathcal{AB} the set of all proper valid allocations (α, β) . The problem of finding the optimal auction can now be restated as the following variational calculus-type problem:

Definition 3.7. [Problem A]

$$\sup_{(\alpha, \beta) \in \mathcal{AB}} \left\{ \int_0^1 \int_{\alpha(y)}^1 f(x, y) dx dy + \int_0^1 \int_{\beta(x)}^1 g(x, y) dy dx \right\}$$




Figure 2: The revenue curves R , \hat{R} and R' .

4 Two Bidders: The Discrete Case

In this section we present an algorithm for computing the optimal mechanism when there are two bidders with a *discrete* joint distribution ϕ with support $Sup(\phi) = [N] \times [N]$.

The marginal profit contribution functions defined in the previous section can be appropriately modified for the discrete setting in hand:

Definition 4.1. *The discrete analogues of the marginal profit contribution functions (Definition 3.2) for each player are defined as follows:*

$$f(i, j) = \max \left\{ i \cdot \sum_{i' \geq i} \phi(i', j) - \sum_{i' > i} f(i', j), 0 \right\} \quad \text{for player 1}$$

$$g(i, j) = \max \left\{ j \cdot \sum_{j' \geq j} \phi(i, j') - \sum_{j' > j} g(i, j'), 0 \right\} \quad \text{for player 2}$$

As in the continuous case, we will represent the auction through a pair of functions (α, β) :

Definition 4.2. *A valid allocation pair (α, β) for the discrete setting is a pair of functions from $[N]$ to itself satisfying the non-crossing property: for all points $(i, j) \in [N] \times [N]$ we have $j \geq \beta(i) \Rightarrow i < \alpha(j)$. Such a pair partitions the set $Sup(\phi)$ into three sets A, B , and C , where $A = \{(i, j) \in Sup(\phi) : i \geq \alpha(j)\}$, $B = \{(i, j) \in Sup(\phi) : j \geq \beta(i)\}$, and $C = Sup(\phi) - A - B$. We say that (α, β) **induces** the partition (A, B) , where C is implicit.*

*Suppose that for a valid pair (α, β) representing an optimal auction and $j \in [n]$, we have that $f(\alpha(j), j) = 0$. Then the auction represented by the same valid pair, except that $\alpha(j)$ is increased by one, is also an optimal auction (since it entails the same set of positive marginal profit contributions). We can therefore consider, without loss of generality, only valid pairs with $f(\alpha(j), j), g(i, \beta(i)) > 0$. We call such valid pairs **proper**. Finally, we define $\Pi(\phi)$ to be the set of all partitions A, B of $Sup(\phi)$ induced by proper valid pairs.*

Now, looking back at the formulation of the optimum-revenue auction problem in Problem A at the conclusion of the last section, it is immediate that obtaining the optimum-revenue auction in the discrete case is tantamount to solving a discrete optimization problem:

$$\max_{(A, B) \in \Pi(\phi)} \sum_{(i, j) \in A} f(i, j) + \sum_{(i, j) \in B} g(i, j).$$

All that remains now is to provide a useful characterization of the set $\Pi(\phi)$. To this end, notice first that any such pair of sets (A, B) has the following two additional properties, the first one inherited from incentive compatibility, and the second one following from the fact that the allocation curves form a proper valid pair:

Definition 4.3. *We call a pair of disjoint subsets (A, B) of $Sup(\phi)$ **monotone** if the following holds:*

- *If $(i, j) \in A$ and $(i', j) \in Sup(\phi)$ with $i' > i$, then $(i', j) \in A$.*
- *If $(i, j) \in B$ and $(i, j') \in Sup(\phi)$ with $j' > j$, then $(i, j') \in B$.*

*We call such a pair **proper** if*

- *If $(i, j) \in A$ and $(i', j) \notin A$ for all $i' < i$, then $f(i, j) > 0$.*
- *If $(i, j) \in B$ and $(i, j') \notin B$ for all $j' < j$, then $g(i, j) > 0$. That is, in proper partitions, all lower boundary points of the regions A and B have positive marginal profit contributions (the intuition being that otherwise, either this is not an optimal auction in the case of a negative marginal contribution, or there is another optimal auction with this property in the case of a zero marginal contribution).*

Let us define now a bipartite graph $G^\phi = (U, W, E)$:

- $U = \{u_{i,j} \mid i, j \in [N]\}$ and $W = \{w_{i,j} \mid i, j \in [N]\}$;
- $(u_{i,j}, w_{i',j'}) \in E$ if and only if $i \leq i'$ and $j \geq j'$. In other words, there is an edge between u and w if and only if, informally, (i', j') “lies to the southeast” of (i, j) (see Figure 3); (Notice that, in our informal sense, a point “lies to the southeast” of all points to its north and to its west, including the point itself.)
- The weight of any node $u_{i,j}$ is $f(i, j)$ and the weight of any node $w_{i,j}$ is $g(i, j)$.

Intuitively, the bipartite graph captures impossibilities in constructing the optimal auction: an edge (u, w) signifies that it is not possible that both $u \in A$ and $w \in B$ (slightly abusing notation).




Figure 3: The bipartite graph for 2 players with $N = 2$.

We can now prove the sought combinatorial characterization of $\Pi(\phi)$:

Lemma 4.4. *Let (A, B) be a pair of disjoint subsets of $\text{Sup}(\phi)$. Then $(A, B) \in \Pi(\phi)$ if and only if (A, B) is monotone and proper, and $\{u_{i,j} : (i, j) \in A\} \cup \{w_{i,j} : (i, j) \in B\}$ is an independent set of G^ϕ .*

Proof. (If.) Since $\{u_{i,j} : (i, j) \in A\} \cup \{w_{i,j} : (i, j) \in B\}$ is an independent set of G^ϕ it follows that the sets A and B are disjoint. Now, since (A, B) is a pair of monotone, proper and disjoint subsets of $\text{Sup}(\phi)$ the following pair of functions $(\alpha(j), \beta(i))$ is a proper valid pair, immediately implying that $(A, B) \in \Pi(\phi)$: $\alpha(j) = \min\{i \mid (i, j) \in A\}$ and $\beta(i) = \min\{j \mid (i, j) \in B\}$.

(Only if.) If $(A, B) \in \Pi(\phi)$ then by the definition of $\Pi(\phi)$ the sets (A, B) have to be monotone and proper and they also need to form a partition, i.e. be disjoint. To show that $\{u_{i,j} : (i, j) \in A\} \cup \{w_{i,j} : (i, j) \in B\}$ is an independent set of G^ϕ assume towards contradiction that there are nodes $u_{i,j}$ and $w_{i',j'}$ such that $(i, j) \in A$ and $(i', j') \in B$, with an edge between $u_{i,j}$ and $w_{i',j'}$; from the construction of G^ϕ , it follows that $i \leq i'$ and $j \geq j'$. Since A and B are both monotone, it follows that $(i', j) \in A \cap B$, contradicting the disjointness of A and B . \square

The next Lemma shows that the additional assumptions that (A, B) is a proper and monotone pair of subsets are not necessary, if one restricts attention to the optimal solution (i.e. the solution of maximum weight).

Lemma 4.5. *Let (A, B) be a pair of subsets of $\text{Sup}(\phi)$ such that the set $\{u_{i,j} : (i, j) \in A\} \cup \{w_{i,j} : (i, j) \in B\}$ is a maximum weight independent set of G^ϕ , of minimum cardinality among all independent sets of the same weight. Then (A, B) is monotone and proper, i.e. $(A, B) \in \Pi(\phi)$.*

Proof. It suffices to show that the set (A, B) is monotone and proper, and the lemma follows from Lemma 4.4. Indeed, the monotonicity of (A, B) follows from the fact that the set $\{u_{i,j} : (i, j) \in A\} \cup \{w_{i,j} : (i, j) \in B\}$ is a maximum weight independent set, and all weights are non-negative. Moreover, since the independent set has minimum cardinality among all independent sets of the same weight, it follows that it does not contain any node $u_{i,j}$ of zero weight, i.e. corresponding to some valuation (i, j) such that $f(i, j) = 0$, unless it also includes a node $u_{i',j}$ corresponding to some valuation (i', j) with $f(i', j) > 0$ for some $i' < i$ (and analogously for $w_{i,j}$). Hence, by definition, (A, B) is proper as well. \square

Theorem 4.6. *Given a discrete joint valuation distribution ϕ for two bidders, the optimal ex-post IC and IR deterministic auction can be computed in time $O(|\text{Sup}(\phi)|^3)$.*

Proof. It follows from Lemma 4.5 that computing the optimal auction for two bidders with a joint valuation distribution ϕ reduces to computing a maximum weight independent set on the induced bipartite graph G^ϕ . In particular the optimal allocation rule corresponds to a partition (A, B) such that $\{u_{i,j} : (i, j) \in A\} \cup \{w_{i,j} : (i, j) \in B\}$ is a maximum weight independent set of G^ϕ , with minimum cardinality among all independent sets of the same weight. Finding the maximum weight independent set by running a min-cost-flow algorithm yields the desired running time. \square

5 Two Bidders: The General Case

In this section we return to the continuous two-bidder problem of Section 3. Our main result is an efficient algorithm that approximates the optimum solution of Problem A within an additive ϵ . Our main tool is a duality theorem, generalizing the duality between the maximum-weight independent set problem in a bipartite graph and a minimum-cost flow in an associate network. In particular, we show that the maximization Problem A defined in Section 3 is equivalent to a certain mass-moving minimization problem (reminiscent in some aspects of the classic Monge-Kantorovich [12] mass-transfer problem), namely the following:

Definition 5.1. [Problem B]

$$\begin{aligned} \inf_{\gamma} & \left\{ \int_0^1 \int_0^1 \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 \right\} \\ \text{s.t.} & \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 \geq f(x_1, y_1), \forall (x_1, y_1) \in [0, 1]^2 \\ & \int_{y_2}^1 \int_0^{x_2} \gamma(x_1, y_1, x_2, y_2) dx_1 dy_1 \geq g(x_2, y_2), \forall (x_2, y_2) \in [0, 1]^2 \\ & \gamma(x_1, y_1, x_2, y_2) \geq 0, \forall (x_1, y_1), (x_2, y_2) \in [0, 1]^2 \end{aligned}$$

Let us denote by Γ the set of all functions $\gamma : [0, 1]^4 \mapsto \mathfrak{R}$ satisfying the above constraints.

One intuitive interpretation of Problem B is this: We are given two landscapes in the unit square, captured by two functions $f, g : [0, 1]^2 \mapsto \mathfrak{R}$. We are seeking a plan for transforming landscape f to landscape g , where the following operations are allowed:

- take material away from any point (x, y) ;
- add material to any point (x, y) ;
- transfer material from any point to any other point *in the southeast direction* (if some material is not moved, we think of it as having moved in the southeast direction zero distance).

We want the plan in which the total amount of material moved (irrespective of distance moved, here is where this problem differs significantly from Monge-Kantorovich) is minimized. We next show that this problem coincides, at optimality, with the optimal auction:

Theorem 5.2. [Duality Theorem]

For any joint density function ϕ on $[0, 1]^2$:

$$\begin{aligned} & \sup_{(\alpha, \beta) \in \mathcal{AB}} \left\{ \int_0^1 \int_{\alpha(y)}^1 f(x, y) dx dy + \int_0^1 \int_{\beta(x)}^1 g(x, y) dy dx \right\} \\ & = \inf_{\gamma \in \Gamma} \left\{ \int_0^1 \int_0^1 \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 \right\} \end{aligned}$$

5.1 Proof of the Duality Theorem

General Plan. The proof of the theorem is by discretizing the unit square into domains of small size, proving a duality result for the discrete version, establishing upper bounds for the discretization error, and taking the limit for finer and finer discretization. In the course of the proof we will introduce a number of auxiliary problems.

Discretization. We start by discretizing the continuous functions f and g defined on $[0, 1] \times [0, 1]$ by two discrete functions f_d and g_d defined on the $[n] \times [n]$ grid, where n is an integer greater than one and $[n] = \{0, 1, \dots, n-1\}$; we denote $\frac{1}{n}$ by ϵ . We subdivide the $[0, 1] \times [0, 1]$ square into $\epsilon \times \epsilon$ little squares; we are mapping a little square with southwest coordinate (x, y) to the grid point $n \cdot (x, y)$. The discrete functions are now obtained by assigning to each point in the grid the aggregate mass of its corresponding square on the plane.

$$f_d(i, j) = \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} f(x, y) dy dx, \text{ for } i, j = 0, \dots, n-1$$

and

$$g_d(i, j) = \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} g(x, y) dy dx \text{ for } i, j = 0, \dots, n-1.$$




Figure 4: The graph consisting of two grids, one for each player, for $n = 3$.

The Graph. We next create a weighted bipartite graph $G = (U, V, E)$ as follows: $U = V = [n] \times [n]$. We use $u(i, j)$ and $v(i, j)$ to denote the vertices of U and V respectively, and sometimes use the shorthand u and v to refer to nodes of each grid respectively. Vertex $u(i, j)$ of U has a weight equal to $w_u = f_d(i, j)$, and similarly vertex $v(i', j')$ of V has a weight $w_v = g_d(i', j')$. A pair $(u(i, j), v(i', j'))$ is in E if and only if $i < i'$ and $j > j'$, that is, if grid point v is in the (strictly) Southeast direction from grid point u (see Figure 4 for an example with $n = 3$).

Consider now the following problem, familiar from the previous section:

Definition 5.3. [Problem C: MAXIMUM WEIGHT INDEPENDENT SET]

Given the weighted bipartite graph $G = (U, V, E)$ above,

$$\begin{aligned} \max_{\{x_u \in \{0,1\}, x_v \in \{0,1\}\}} & \sum_{u \in U, v \in V} x_u w_u + x_v w_v \\ \text{s.t.} & x_u + x_v \leq 1, \forall (u, v) \in E \end{aligned}$$

The dual of the above problem is the following:

Definition 5.4. [Problem D: MINIMUM COST TRANSSHIPMENT]
Given the weighted bipartite graph $G = (U, V, E)$ above,

$$\begin{aligned} \min_{\{y_{uv} \in \mathfrak{R}\}} \quad & \sum_{(u,v) \in E} y_{uv} \\ \text{s.t.} \quad & \sum_{v:(u,v) \in E} y_{uv} \geq w_u, \forall u \in U \\ & \sum_{u:(u,v) \in E} y_{uv} \geq w_v, \forall v \in V \\ & y_{uv} \geq 0, \forall (u,v) \in E \end{aligned}$$

The Inequalities. The crux of the proof is a sequence of results relating the various solutions and optimum solutions of these four problems. In what follows we use $SOL(\cdot)$ to denote the cost of any feasible solution of a problem among (A), (B), (C), and (D) defined above, and $OPT(\cdot)$ to denote the cost of the optimum solution of a problem (sometimes SOL and OPT also denotes the actual solutions). The first such inequality establishes a form of weak duality between Problems A and B, while the next two show that the discretization error is small.

Lemma 5.5. For any feasible solutions of A and B we have: $SOL(A) \leq SOL(B)$.

Proof.

$$\begin{aligned} & SOL(A) \\ = & \int_0^1 \int_{\alpha(y_1)}^1 f(x_1, y_1) dx_1 dy_1 + \int_0^1 \int_{\beta(x_2)}^1 g(x_2, y_2) dy_2 dx_2 \\ \leq & \int_0^1 \int_{\alpha(y_1)}^1 \int_0^{y_1} \int_{x_1}^{x_2} \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 + \int_0^1 \int_{\beta(x_2)}^1 \int_{y_2}^1 \int_0^{x_2} \gamma(x_1, y_1, x_2, y_2) dx_1 dy_1 dy_2 dx_2 \end{aligned}$$

where we used the inequality constraints of Problem B to upper bound the values of $f(x, y)$ and $g(x, y)$. We next notice that:

$$\int_0^1 \int_{\beta(x_2)}^1 \int_{y_2}^1 \int_0^{x_2} \gamma(x_1, y_1, x_2, y_2) dx_1 dy_1 dy_2 dx_2 \leq \int_0^1 \int_0^{\alpha(y_1)} \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1$$

This inequality follows from the non-negativity of γ and the fact that the (x_1, y_1, x_2, y_2) included in the integral of the LHS are the following set:

$$\begin{aligned} & \{(x_1, y_1, x_2, y_2) \in [0, 1]^4 \mid y_2 \geq \beta(x_2), x_1 \leq x_2, y_1 \geq y_2\} \\ = & \{(x_1, y_1, x_2, y_2) \in [0, 1]^4 \mid y_2 \geq \beta(x_2), x_1 \leq x_2, y_1 \geq y_2, x_1 \leq \alpha(y_1)\} \\ \subseteq & \{(x_1, y_1, x_2, y_2) \in [0, 1]^4 \mid x_2 \leq \alpha(y_2), x_1 \leq x_2, y_1 \geq y_2, x_1 \leq \alpha(y_1)\} \\ \subseteq & \{(x_1, y_1, x_2, y_2) \in [0, 1]^4 \mid x_1 \leq x_2, y_1 \geq y_2, x_1 \leq \alpha(y_1)\} \end{aligned}$$

which is exactly the set of (x_1, y_1, x_2, y_2) included in the integral of the RHS. The first equality above follows from the fact that the inequality $x_1 \leq \alpha(y_1)$ follows from the inequalities $\{y_2 \geq \beta(x_2), x_1 \leq x_2, y_1 \geq y_2\}$, because we have $y_1 \geq \beta(x_2)$ so the non-crossing property implies that $x_2 \leq \alpha(y_1)$ and therefore $x_1 \leq \alpha(y_1)$. The first set inclusion follows from the fact that $y_2 \geq \beta(x_2) \Rightarrow x_2 \leq \alpha(y_2)$ from the non-crossing property, while the last inclusion is trivial.

We have therefore concluded that the cost of any feasible solution of A is upper bounded by:

$$\begin{aligned} & \int_0^1 \int_{\alpha(y_1)}^1 \int_0^{y_1} \int_{x_1}^{x_2} \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 + \int_0^1 \int_0^{\alpha(y_1)} \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 \\ = & \int_0^1 \int_0^1 \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 \\ = & SOL(B) \end{aligned}$$

□

Lemma 5.6. *For the optimal solutions of A and C we have: $OPT(A) \geq OPT(C) - \epsilon$.*

Proof. Consider the optimal solution of Problem C ; we will use it to come up with a feasible solution for Problem A such that $SOL(A) \geq OPT(C) - \epsilon$. We start with the following solution: we allocate the item to player 1 for all valuations (x, y) such that $x_{u(\lfloor \frac{x}{\epsilon} \rfloor, \lfloor \frac{y}{\epsilon} \rfloor)} = 1$; we allocate to player 2 for all valuations (x, y) such that $x_{v(\lfloor \frac{x}{\epsilon} \rfloor, \lfloor \frac{y}{\epsilon} \rfloor)} = 1$, and $x_{u(\lfloor \frac{x}{\epsilon} \rfloor, \lfloor \frac{y}{\epsilon} \rfloor)} = 0$; and finally, we allocate to nobody for all valuations (x, y) such that $x_{u(\lfloor \frac{x}{\epsilon} \rfloor, \lfloor \frac{y}{\epsilon} \rfloor)} = 0$ and $x_{v(\lfloor \frac{x}{\epsilon} \rfloor, \lfloor \frac{y}{\epsilon} \rfloor)} = 0$.

We next show that the resulting allocation regions have the shape of Figure 5, meaning that the borders of those regions consist a valid allocation pair. First notice that for any pair of valuations (x, y) –including those for which $x_{u(\lfloor \frac{x}{\epsilon} \rfloor, \lfloor \frac{y}{\epsilon} \rfloor)} = 1$ and $x_{v(\lfloor \frac{x}{\epsilon} \rfloor, \lfloor \frac{y}{\epsilon} \rfloor)} = 1$ – only one player gets allocated the item, so the non-crossing property is satisfied. To see why the regions are rightward and upward closed consider two nodes $u(i, j)$ and $u(i', j)$ on player 1's grid, where $i' > i$. Notice that the set of nodes on player 2's grid that node $u(i, j)$ of player 1's grid is connected to, is a strict superset of the nodes that node $u(i', j)$ is connected to. Hence, if the *maximum* weight independent set includes node $u(i, j)$ on the grid of player 1, it should also include $u(i', j)$ for all values $i' > i$.

This gives us two stairwise curves which –although being a valid allocation pair– may fail to satisfy Condition 1, and hence may not be a feasible solution for Problem A . To turn them into a proper valid pair, we can follow the same procedure as the one in the proof of Lemma 3.6 and come up with a feasible solution $SOL(A)$ for Problem A .

Because of the aforementioned transformation the cost of this solution is *greater or equal* to the cost of the optimal solution $OPT(C)$ *minus* the contribution to the weight of the independent set by those nodes $v(i, j)$ for which the corresponding node $u(i, j)$ on the grid of player 1 is also included in the independent set. The reason for that is that for valuations (x, y) such that $x_{v(\lfloor \frac{x}{\epsilon} \rfloor, \lfloor \frac{y}{\epsilon} \rfloor)} = 1$ and $x_{u(\lfloor \frac{x}{\epsilon} \rfloor, \lfloor \frac{y}{\epsilon} \rfloor)} = 1$, our solution explicitly allocates the item only to player 1, therefore losing the weight contribution of node $v(\lfloor \frac{x}{\epsilon} \rfloor, \lfloor \frac{y}{\epsilon} \rfloor)$. In what follows we argue that this results in the loss of an ϵ -additive factor, so that the cost of the resulting solution is at least:

$$\sum_{u \in U} x_u f_d(u) + \sum_{v \in V} x_v g_d(v) - \epsilon = OPT(C) - \epsilon$$

To show this we first argue that the number of nodes $v(i, j)$ for which this happens is small, in particular there can only be at most $1/\epsilon$ such nodes. To see this notice that in the constructed feasible solution to Problem A , these nodes lie on the boundary between regions where player 1 gets the item and player 2 gets the item; any such boundary has to be monotone, since it corresponds to the overlap of the two allocation curves α, β , and it can therefore contain at most $1/\epsilon$ nodes. Next notice that the value of g at any point (x, y) is at most 1: indeed, $g(x, y)$ is defined as $y\phi(x, y) - \int_y^1 \phi(x, t) dt$, wherever $\frac{\partial}{\partial y} \left[\max_{y' \geq y} y' \cdot \int_{y'}^1 \phi(x, t) dt \right]$ is defined, and extended to full range by right continuity. It follows immediately that $g(x, y) \leq 1$ and $w_{v(i, j)} = g_d(i, j) \leq \epsilon^2$; therefore the total weight loss is at most $\frac{1}{\epsilon} \cdot \epsilon^2$ and the lemma follows. □

Lemma 5.7. *For the optimal solutions of B and D we have: $OPT(D) \geq OPT(B)$.*

Proof. Given a feasible solution for Problem D , we will come up with a feasible solution of the same cost for Problem B . The optimum solution for Problem B will have at most that cost and the lemma follows.

We start by defining $\gamma(x_1, y_1, x_2, y_2)$, for any pair of points $(x_1, y_1), (x_2, y_2)$ where the ϵ^2 -area square containing (x_2, y_2) lies in the (strict) southeast orthant of the ϵ^2 -area square containing (x_1, y_1) , as follows:

$$\gamma(x_1, y_1, x_2, y_2) = f(x_1, y_1) \cdot g(x_2, y_2) \cdot \frac{y_{uv}}{w_u w_v}, \quad \text{if } \left\lfloor \frac{x_1}{\epsilon} \right\rfloor < \left\lfloor \frac{x_2}{\epsilon} \right\rfloor \quad \text{and} \quad \left\lfloor \frac{y_1}{\epsilon} \right\rfloor > \left\lfloor \frac{y_2}{\epsilon} \right\rfloor$$

exactly $OPT(D)$:

$$\begin{aligned}
& \int_0^1 \int_0^1 \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 \\
&= \sum_{i,j} \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 \\
&= \sum_{i,j} \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} f(x_1, y_1) \cdot \sum_{i' > \lfloor \frac{x_1}{\epsilon} \rfloor, j' < \lfloor \frac{y_1}{\epsilon} \rfloor} \frac{y_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor), v(i', j')}{w_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor)} dx_1 dy_1 \\
&= \sum_{i,j} \sum_{i' > i, j' < j} \frac{y_u(i,j), v(i', j')}{w_u(i,j)} \cdot \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} f(x_1, y_1) dx_1 dy_1 \\
&= \sum_{i,j} \sum_{i' > i, j' < j} \frac{y_u(i,j), v(i', j')}{w_u(i,j)} \cdot w_u(i,j) \\
&= \sum_{i,j} \sum_{i' > i, j' < j} y_u(i,j), v(i', j') \\
&= OPT(D)
\end{aligned}$$

where in the first equality we split the integration of (x_1, y_1) over discretized square regions of area ϵ^2 and in the second equality we plugged in the expression for $\int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2$ that we had derived from our previous proof establishing that the first constraint of Problem B was satisfied. In the third equality we once again rearranged the order of summation and integration, noticing that the weights w and flows y remain constant across the discretized squares (independently of the actual value of (x_1, y_1)), in the fourth equality we used the definition of the weight w and in the last equality we replaced with the objective function of Problem D . \square

Proof of the Duality Theorem. First notice that, by strong duality and since the constraint matrix of Problem C is totally unimodular, we have that $OPT(C) = OPT(D)$; combining this with Lemmas 5.6 and 5.7 we get $OPT(A) \geq OPT(B) - \epsilon$. From Lemma 5.5 we get $OPT(A) \leq OPT(B)$. By having $\epsilon \rightarrow 0$ we get the result. \square

5.2 The Algorithm

The proof of the Main Theorem suggests a fully polynomial-time approximation scheme (FP-TAS) for the continuous case, that is, a mechanism that approximates the optimal profit within additive error ϵ , and runs in time polynomial in $\frac{1}{\epsilon}$. In the algorithm and the correctness proof, we assume that the continuous joint distribution ϕ is Lipschitz continuous, and that it is presented through oracle access. It is easy to see that these assumptions are essentially necessary, in that no approximation (or meaningful solution of any other nature) is possible when the function can be arbitrarily discontinuous, or is inaccessible for large parts of the domain.

Algorithm 1 OPTIMALAUCTION for two bidders with continuous distributions

- 1: **Input:** probability distribution $\phi \in [0, 1]^2$
 - 2: Compute $f(\cdot, \cdot), g(\cdot, \cdot)$
 - 3: Discretize the unit plane and construct the bipartite graph G as described in Section 5.1
 - 4: Compute a MAXIMUM WEIGHT INDEPENDENT SET for G
 - 5: **Output:** (α, β) , the valid allocation pair corresponding to the stair-like curves of Figure 5
-

The following theorem establishes that our algorithm has the desired properties.

Theorem 5.8. *Algorithm 1 returns a truthful mechanism that approximates the optimal profit within ϵ additive error; moreover the algorithm runs in time polynomial in $1/\epsilon$.*

Proof. Algorithm 1 returns a valid allocation pair so it is truthful by construction. However the allocation pair (α, β) returned may well not satisfy condition (1) and may consequently not constitute a feasible solution to Problem A. This is problematic since it does not allow us to use Lemma 3.4 to compute the profit of the auction returned. To that end we need to establish that the violation of condition (1) is—in some sense—negligible; we do that next.

Suppose that curve $\alpha(y)$ violates condition (1) for some $y^* \in [0, 1]$ and let

$$x^* = \arg \max_{x \geq \alpha(y^*)} x \cdot \int_x^1 \phi(t, y^*) dt,$$

be the minimum x for which we could create a new solution α' by setting $\alpha'(y^*) = x^*$ and have condition (1) restored, while not altering the profit of our auction.¹ We will argue that

$$\alpha(y^*) \cdot \int_{\alpha(y^*)}^1 \phi(t, y^*) dt \geq x^* \cdot \int_{x^*}^1 \phi(t, y^*) dt - \Theta(\epsilon) \quad (2)$$

To do that we consider the node corresponding to the little square on the unit plane containing $(\alpha(y^*), y^*)$. Since this node belongs to the boundary of the allocation region of player 1, we can assume wlog that it has non-zero weight; hence there must exist some point (x_1, y_1) in the corresponding square on the unit plane with $f(x_1, y_1) > 0$. By the definition of f this immediately implies that

$$x_1 \cdot \int_{x_1}^1 \phi(t, y_1) dt \geq x \cdot \int_x^1 \phi(t, y_1) dt, \quad \text{for all } x \geq x_1$$

and therefore in particular that

$$x_1 \cdot \int_{x_1}^1 \phi(t, y_1) dt \geq x^* \cdot \int_{x^*}^1 \phi(t, y_1) dt. \quad (3)$$

Since the l_1 -distance of points $(\alpha(y^*), y^*)$ and (x_1, y_1) and of points (x^*, y^*) and (x^*, y_1) is at most 2ϵ , we get that:

$$\alpha(y^*) \cdot \int_{\alpha(y^*)}^1 \phi(t, y^*) dt \geq x_1 \cdot \int_{\alpha(x_1)}^1 \phi(t, y_1) dt - 2\lambda\epsilon \geq x^* \cdot \int_{x^*}^1 \phi(t, y_1) dt - 2\lambda\epsilon \geq x^* \cdot \int_{x^*}^1 \phi(t, y^*) dt - 4\lambda\epsilon$$

where in the first and third inequalities we used the fact that $x \cdot \int_x^1 \phi(t, y) dt$ is Lipschitz-continuous for some constant λ and in the second inequality we used inequality (3). The exact same argument applies for $\beta(x)$ as well.

We are now ready to prove a lower bound on the profit of the auction returned by our algorithm; in what follows we use $\alpha'(y)$ and $\beta'(x)$ to denote the allocation curves that would result by the aforementioned transformation. Note that by construction it holds that

$$\alpha'(y) \cdot \int_{\alpha'(y)}^1 \phi(x, y) dx = \int_{\alpha'(y)}^1 f(x, y) dx \quad \text{and} \quad \beta'(x) \cdot \int_{\beta'(x)}^1 \phi(x, y) dy = \int_{\beta'(x)}^1 g(x, y) dy \quad (4)$$

so the profit of the algorithm is:

$$\begin{aligned} & \int_0^1 \left[\alpha(y) \cdot \int_{\alpha(y)}^1 \phi(x, y) dx \right] dy + \int_0^1 \left[\beta(x) \cdot \int_{\beta(x)}^1 \phi(x, y) dy \right] dx \\ & \geq \int_0^1 \left[\alpha'(y) \cdot \int_{\alpha'(y)}^1 \phi(x, y) dx - \Theta(\epsilon) \right] dy + \int_0^1 \left[\beta'(x) \cdot \int_{\beta'(x)}^1 \phi(x, y) dy - \Theta(\epsilon) \right] dx \\ & = \int_0^1 \int_{\alpha'(y)}^1 f(x, y) dx dy + \int_0^1 \int_{\beta'(x)}^1 g(x, y) dy dx - \Theta(\epsilon) \\ & = OPT(C) - \Theta(\epsilon) \\ & \geq OPT(A) - \Theta(\epsilon) \end{aligned}$$

¹The reader is referred to Lemma 3.6 for further discussion on this point.

where in the first inequality we used inequality (2), in the first equality we used (4) and in the last inequality we used our Main Theorem from the previous section.

In terms of running time, the discretized approximations of f and g are trivial (because of Lipschitz continuity, we can take $f(i, j) = f(i\epsilon, j\epsilon)$, and similarly for g). Solving the MAXIMUM WEIGHT INDEPENDENT SET problem is done exactly as in the previous section. \square

6 NP-completeness

We show that for 3 bidders the problem of designing an approximately optimal (deterministic) auction becomes NP-hard. The current proof establishes hardness for a small threshold around 0.05%; we believe that this will not be too hard to improve upon.

6.1 A geometric characterization for three bidders

We start by formally defining the discrete version of the problem we will prove to be NP-hard; to simplify the exposition of the problem, we assume that the support of the discrete distribution is $\mathcal{G} = [S] \times [S] \times [S]$. We are interested in this problem:

Definition 6.1 (3OPTIMALAUCTIONDESIGN). *Given a joint discrete probability distribution ϕ supported on \mathcal{G} , find the optimal, ex-post IC and IR deterministic auction, denoted by a 3-dimensional allocation matrix A , where $A[x, y, z] = i$, with i being the index of the bidder who gets the item when the bid vector is (x, y, z) , or 0 if the auctioneer keeps the item.*

As was the case with 2 bidders, the following notion of marginal profit contribution, appropriately modified for the discrete case, will be useful to our proof.

Definition 6.2. *The discrete analogues of the marginal profit contribution functions (Definition 3.2) for each player are the following:*

$$f(x, y, z) = \max \left\{ x \cdot \sum_{x' \geq x} \phi(x', y, z) - \sum_{x' > x} f(x', y, z), 0 \right\} \quad \text{for player 1}$$

$$g(x, y, z) = \max \left\{ y \cdot \sum_{y' \geq y} \phi(x, y', z) - \sum_{y' > y} f(x, y', z), 0 \right\} \quad \text{for player 2}$$

$$h(x, y, z) = \max \left\{ z \cdot \sum_{z' \geq z} \phi(x, y, z') - \sum_{z' > z} f(x, y, z'), 0 \right\} \quad \text{for player 3}$$

The Segments. Given a distribution $\phi(x, y, z)$ over the points of \mathcal{G} , any node (x, y, z) of \mathcal{G} with $f(x, y, z) > 0$ is the starting point of what we shall henceforth be calling an x -segment: an interval (sequence of points) starting at node (x, y, z) and including all nodes (x', y, z) with $x' \geq x$. The *weight* of this segment is $\sum_{x' \geq x} f(x', y, z)$. We define segments across the other dimensions analogously. The following problem² is essentially equivalent to the auction design problem:

Definition 6.3 (3SEGMENTS). *Given a joint discrete probability distribution ϕ supported on \mathcal{G} , which induces a set of segments on \mathcal{G} as described above, find a subset of non-intersecting segments with maximum sum of weights.*

Lemma 6.4. *The problem 3SEGMENTS(ϕ) is equivalent to 3OPTIMALAUCTIONDESIGN(ϕ) via approximation-preserving reductions.*

²The reader may notice that the 3SEGMENTS problem is a maximum weight independent set problem in disguise.

Proof. (Sketch:) The correspondence between solutions of the two problems is rather immediate, with the x -segment (respectively, y -segment, z -segment) at point (x, y, z) included in the output set of segments if and only if $A[x, y, z] = 1$ (respectively, 2, 3) and $A[x', y, z] = 0$ for all $x' < x$ (respectively, $A[x, y', z] = 0$ for all $y' < y$ and $A[x, y, z'] = 0$ for all $z' < z$). \square

6.2 The construction

It suffices to show that 3SEGMENTS is NP-hard to approximate for some constant; we do that by reducing from 3CATSAT, a special case of 3SAT:

Definition 6.5. *Let 3CATSAT be the 3SAT problem with the input formula restricted to be of the following form. We have three types (categories) of variables $\{x_i\}_{i=1\dots n_x}$, $\{y_i\}_{i=1\dots n_y}$ and $\{z_i\}_{i=1\dots n_z}$, i.e. a total of $n = n_x + n_y + n_z$ variables, and m clauses of the following form: every clause has at most one literal from every type (e.g. $(\bar{x}_2 \vee \bar{y}_3 \vee z_1)$ or $(\bar{x}_5 \vee z_7)$).*

Lemma 6.6. *3CATSAT is NP-hard to approximate better than $79/80$.*

Proof. We reduce from MAX3SAT. In order to turn an instance of MAX3SAT to an instance of 3CATSAT it suffices to create three copies for every occurrence of a variable: for variable x we create x_1, x_2, x_3 and include the clauses $(\bar{x}_1 \vee x_2), (\bar{x}_2 \vee x_3), (\bar{x}_3 \vee x_1)$. This increases the number of clauses by at most $3\bar{n}$, where \bar{n} is the number of literals in the formula. Given that MAX3SAT is hard to approximate better than $7/8$, it follows immediately that 3CATSAT is hard to approximate better than $\frac{7/8m+3\bar{n}}{m+3\bar{n}}$. Noticing that $\bar{n} \leq 3m$ we get the desired approximation factor as an upper bound on this expression by picking $\bar{n} = 3m$. \square

We start with some intuition about the reduction. The instance of 3SEGMENTS we create has three types of segments: literal segments, clause segments and scaffolding segments.

- *Literal segments* are used to model truth assignments on variables; they ensure that every variable is assigned exactly one of the two possible truth values and that this assignment is consistent across all appearances of literals of this particular variable.
- *Clause segments* model the truth assignment to literals of a particular clause; we create one such clause segment for every literal that appears in a clause and we make them intersect with the literal segments of those literals; the idea is that if the clause is satisfied we will be able to pick at least one clause segment per clause because the corresponding literal segment will not be picked. Moreover, we cannot pick two or more clause segments per clause, since they will all intersect with each other.
- *Scaffolding segments*³ ensure that, for some points, there are literal or clause segments that extend only to one of the three possible directions: In particular, given some point (x, y, z) with positive marginal profit contribution for more than one players (for example when $f(x, y, z) > 0$ and $g(x, y, z) > 0$) we may want to have only x or only y -segments starting from (x, y, z) ; scaffolding points make sure this is the case.

We show how to construct a probability distribution ϕ that serves as the input of 3SEGMENTS, given an instance of 3CATSAT. In what follows we use \hat{n} to denote $\max\{n_x, n_y, n_z\}$. The support⁴ of ϕ is: $\{h(i)|i = 1, \dots, \hat{n} + 2m + 4\}^3$, for an appropriate choice of the values $h(i)$ which we will fix later; for now all we assume is that h is an increasing function of i . The size of the support is at most $(\hat{n} + 2m + 4)^3$, so this is clearly a polynomial time construction.

We shall abuse notation and write $\phi(x, y, z)$ instead of $\phi(h(x), h(y), h(z))$ when there is no ambiguity. We shall also refer to the sub-matrices $\phi(i, \cdot, \cdot), \phi(\cdot, j, \cdot), \phi(\cdot, \cdot, k)$ as the “planes” $x = i, y = j$ and $z = k$ respectively.

The construction goes as follows: we start with an all-zero matrix ϕ of the above size. Consider an arbitrary ordering of the clauses 1 through m . Suppose the l -th clause is of the form $(x_i \vee y_j \vee z_k)$ where x_i, y_j and z_k can be either positive or negative literals, and i (resp. j, k) =

³Rather, scaffolding points, as these segments have zero length.

⁴The support is actually a subset of this; these are all the points (values of the players) with *potentially* non-zero probabilities. This will become clear in the actual construction.

$1, \dots, n_x$ (resp. n_y, n_z) (the same construction also works for clauses with less than 3 variables). For this clause we introduce the following literal segments (1,2,3) and clause segments (4). In what follows we will first set the probability mass of the point that is the apex of each segment, and will later show how to use scaffolding-points to ensure that there is only one segment starting at each such point, towards the appropriate direction.

1. If $pos(x_i)$ then $\phi(i+1, \hat{n}+2, \hat{n}+2+l) = c_1$ and $\phi(i+1, \hat{n}+2+m+l, \hat{n}+2) = c_1$
 (These points are intended to be the apices of a y -segment and a z -segment respectively.)
 else $\phi(i+1, \hat{n}+2+l, \hat{n}+2) = c_1$ and $\phi(i+1, \hat{n}+2, \hat{n}+2+m+l) = c_1$
 (These points are intended to be the apices of a z -segment and a y -segment respectively.)
2. If $pos(y_j)$ then $\phi(\hat{n}+2+l, j+1, \hat{n}+2) = c_1$ and $\phi(\hat{n}+2, j+1, \hat{n}+2+m+l) = c_1$,
 (These points are intended to be the apices of a z -segment and an x -segment respectively.)
 else $\phi(\hat{n}+2, j+1, \hat{n}+2+l) = c_1$ and $\phi(\hat{n}+2+m+l, j+1, \hat{n}+2) = c_1$
 (These points are intended to be the apices of an x -segment and a z -segment respectively.)
3. If $pos(z_k)$ then $\phi(\hat{n}+2, \hat{n}+2+l, k+1) = c_1$ and $\phi(\hat{n}+2+m+l, \hat{n}+2, k+1) = c_1$,
 (These points are intended to be the apices of an x -segment and a y -segment respectively.)
 else $\phi(\hat{n}+2+l, \hat{n}+2, k+1) = c_1$ and $\phi(\hat{n}+2, \hat{n}+2+m+l, k+1) = c_1$
 (These points are intended to be the apices of a y -segment and an x -segment respectively.)
4. $\phi(1, \hat{n}+2+l, \hat{n}+2+l) = \phi(\hat{n}+2+l, 1, \hat{n}+2+l) = \phi(\hat{n}+2+l, \hat{n}+2+l, 1) = c_2$
 (These points are intended to be the apices of an x -segment, a y -segment and a z -segment respectively.)

Every positive occurrence of a variable of type, say, x results in the following two segments: a *positive* literal segment starting at $(i+1, \hat{n}+2, \hat{n}+2+l)$ that intersects with the corresponding clause segment starting at $(1, \hat{n}+2+l, \hat{n}+2+l)$, and a literal segment starting at $(i+1, \hat{n}+2+m+l, \hat{n}+2)$ that does not intersect with any clause segment; this is also called a *dummy-negative* literal segment. Negative occurrences of variables analogously result in *negative* and *dummy-positive* literal segments. Dummy literal segments are introduced because –for reasons that will become apparent later– we want to ensure that we have an equal number of positive and negative literal segments (when dummies are included).

The reduction relies on the fact that the **only** intersections involving literal and clause segments will be between literal segments of literals that are negations of each other, between clause segments of the same clause and between clause segments and their corresponding literal segments. To ensure that we need the aforementioned scaffolding segments; these will ensure that the only segments starting from the points defined above as having probability masses c_1 and c_2 (henceforth called c_1 and c_2 points), are the desired literal and clause segments (in other words we want those c_1 and c_2 points to be the apices of *only* the segments mentioned above):

1. We first ensure that there is exactly one segment starting from every point c_2 , which is perpendicular to the plane $x, (or\ y, z) = 1$; in other words we ensure that there are no other segments starting at c_2 that lie on the plane, by introducing: $\phi(1, \hat{n}+l+2, \hat{n}+2m+3) = \phi(1, \hat{n}+2m+3, \hat{n}+l+2) = \phi(\hat{n}+l+2, 1, \hat{n}+2m+3) = \phi(\hat{n}+2m+3, 1, \hat{n}+l+2) = \phi(\hat{n}+l+2, \hat{n}+2m+3, 1) = \phi(\hat{n}+2m+3, \hat{n}+l+2, 1) = c_3$, for all $l = 1, \dots, m$, with the requirement that:

$$h(\hat{n}+l+2) \cdot (c_2 + c_3) < h(\hat{n}+2m+3) \cdot c_3, \quad \forall l = 1, \dots, m$$

The above requirement forces the corresponding marginal profit contribution function to be negative at any c_2 point, for some appropriately chosen player-direction: For example, we want point $(1, \hat{n}+2+l, \hat{n}+2+l)$ to be the apex of an x -segment only, and no y or z -segment should start at this point. We achieve this by including the points $\phi(1, \hat{n}+l+2, \hat{n}+2m+3) = \phi(1, \hat{n}+2m+3, \hat{n}+l+2) = c_3$; the above inequality then ensures that

$$g(1, \hat{n}+2+l, \hat{n}+2+l) = h(1, \hat{n}+2+l, \hat{n}+2+l) = 0$$

and therefore there are no y or z -segments starting at this point.

2. We next ensure that there are no segments starting from c_1 that go along row or column $\hat{n} + 2$, by introducing: $\phi(i + 1, \hat{n} + 2, \hat{n} + 2m + 3) = \phi(i + 1, \hat{n} + 2m + 3, \hat{n} + 2) = \phi(\hat{n} + 2, j + 1, \hat{n} + 2m + 3) = \phi(\hat{n} + 2m + 3, j + 1, \hat{n} + 2) = \phi(\hat{n} + 2, \hat{n} + 2m + 3, k + 1) = \phi(\hat{n} + 2m + 3, \hat{n} + 2, k + 1) = c_4$, for $i = 1, \dots, n_x, j = 1, \dots, n_y, k = 1, \dots, n_z$, with the requirement that:

$$h(\hat{n} + l + 2) \cdot (2mc_1 + c_4) < h(\hat{n} + 2m + 3) \cdot c_4, \quad \forall l = 1, \dots, 2m$$

because we can have at most m occurrences of any variable, and hence at most $2m$ c_1 -entries on any given level. The rationale behind the inequality above is the same as in case (1) above.

3. Finally, we ensure that there are no segments starting from a point c_1 on plane $x = i$ (resp. $y = i, z = i$), for $i = \hat{n} + 3, \dots, \hat{n} + 2m + 2$, perpendicular to the plane x (resp. y, z), by introducing: $\phi(\hat{n} + 2m + 4, i, \hat{n} + 2) = \phi(\hat{n} + 2m + 4, \hat{n} + 2, i) = \phi(\hat{n} + 2, \hat{n} + 2m + 4, i) = \phi(i, \hat{n} + 2m + 4, \hat{n} + 2) = \phi(i, \hat{n} + 2, \hat{n} + 2m + 4) = \phi(\hat{n} + 2, i, \hat{n} + 2m + 4) = c_5$, for all $i = \hat{n} + 3, \dots, \hat{n} + 2m + 2$, with the requirement (along the same lines as above) that:

$$\hat{n} \cdot h(\hat{n} + 1) \cdot c_1 < h(\hat{n} + 2m + 4) \cdot c_5$$

The following values for $h(\cdot)$ and the constants⁵ satisfy all of the constraints above:

$$h(i) = \begin{cases} 1 + \frac{i-1}{\hat{n}+2m+1} & \text{for } i = 1, \dots, \hat{n} + 2m + 2 \\ 4 & \text{for } i = \hat{n} + 2m + 3 \\ 5 & \text{for } i = \hat{n} + 2m + 4 \end{cases}$$

and

$$c_1 = \frac{1}{\hat{n}^2 m}, \quad c_2 = 1, \quad c_3 = 1, \quad c_4 = \frac{3}{\hat{n}^2}, \quad c_5 = \frac{2}{5\hat{n}m}$$

Our main result for this section is then the following:

Lemma 6.7. *It is NP-hard to approximate 3SEGMENTS better than 0.05%.*

Proof. We first note that regardless of the instance of 3CATSAT we are reducing from, we can always obtain a fixed profit for 3SEGMENTS from the scaffolding points; by picking the most profitable segments starting at each of the scaffolding points we get a total profit of:

$$F = 6m \cdot h(\hat{n} + 2m + 3)c_3 + 2n \cdot h(\hat{n} + 2m + 3)c_4 + 12m \cdot h(\hat{n} + 2m + 4)c_5$$

Let \bar{n} be the total number of literal occurrences in the formula. We then have the following:

- If the 3CATSAT formula is satisfiable then the profit of 3SEGMENTS is exactly:

$$m \cdot h(1)c_2 + \bar{n} \cdot h(\hat{n} + 2)c_1 + F \tag{5}$$

To see this first consider the following way to pick the literal segments according to the truth values assigned to the corresponding variables: if a variable is set to **true**⁶ we include its **negative** literal segments (and the corresponding dummy-positive literal segments). Notice that –thanks to the dummy literal segments– there is an equal number of positive and negative literal segments (dummies included), with totally $2\bar{n}$ of them; we include exactly half of them for every variable (either the positive or the negative ones), so the total profit from these segments is exactly $\bar{n} \cdot h(\hat{n} + 2)c_1$.

Moreover, since the formula is satisfiable, at least one literal per clause is satisfied; if this is a positive⁶ literal, then the variable has been set to true so the literal segments for this variable included in our 3SEGMENTS solution will be the negative ones. However since the variable appears as a positive literal at this clause, our construction ensures that the corresponding clause segment intersects only with the positive literal segment of this

⁵In order to have a proper probability distribution these constants need to be normalized by their sum.

⁶The other case is completely symmetrical.

variable. Therefore, since the clause segment does not intersect with the negative literal segment (that we have already included in our 3SEGMENTS solution), we can include the clause segment in our solution as well. Each one of these clause segments contributes $h(1)c_2$; noticing that we cannot include more than one clause segment from each clause (because they intersect) we get that their total contribution is exactly $m \cdot h(1)c_2$.

- If the optimal assignment for 3CATSAT satisfies at most ρ of the clauses, then the optimal profit for 3SEGMENTS is at most:

$$\rho m \cdot h(1)c_2 + \bar{n} \cdot h(\hat{n} + 2)c_1 + F \quad (6)$$

To prove this we show how to transform a solution of profit $\geq \rho m \cdot h(1)c_2 + \bar{n} \cdot h(\hat{n} + 2)c_1 + F$ to a truth assignment that satisfies more than ρ of the clauses. First notice that in order to achieve such a profit we must include exactly \bar{n} literal segments and exactly one clause segment per clause for ρ of the clauses: we cannot include more without having intersecting segments and we cannot include any less and achieve the same profit. These \bar{n} literal segments correspond to a truth value assignment to the variables of the formula as described above: we set every variable whose positive (resp. negative) literal segments are included to false (resp. true). The claim follows by noticing that this truth value assignment satisfies every clause for which a clause segment was included in the 3SEGMENTS solution, i.e. for a fraction ρ of the clauses.

From the way we set the constants and $h(\cdot)$, it follows that expression (5) is $\geq 25m$ and expression (6) is $\leq (24 + \rho)m + \delta$ for some constant $\delta \ll 0.1$, for sufficiently large m and n . Using the fact from Lemma 6.6 that $\rho = 79/80$ we get an approximation ratio of $1999/2000$. \square

		negative literal-segments for variables of type z						dummy-positive literal-segments for variables of type z					
dummy-negative literal-segments for variables of type z	$\hat{n}+2m+4$	0	...	0	0	0	...	0	0	...	0	0	0
	$\hat{n}+2m+3$	0	...	0	c_4	0	...	0	0	...	0	0	0
	$\hat{n}+2m+2$	0	...	0	$[c_1]$	0	...	0	0	...	0	0	0
	...	\vdots	...	\vdots	\vdots	\vdots	...	\vdots	\vdots	...	\vdots	\vdots	\vdots
	$\hat{n}+m+3$	0	...	0	$[c_1]$	0	...	0	0	...	0	0	0
	$\hat{n}+m+2$	0	...	0	$[c_1]$	0	...	0	0	...	0	0	0
	...	\vdots	...	\vdots	\vdots	\vdots	...	\vdots	\vdots	...	\vdots	\vdots	\vdots
	$\hat{n}+3$	0	...	0	$[c_1]$	0	...	0	0	...	0	0	0
	$\hat{n}+2$	0	...	0	0	$[c_1]$...	$[c_1]$	$[c_1]$...	$[c_1]$	c_4	0
	$\hat{n}+1$	0	...	0	0	0	...	0	0	...	0	0	0
	...	\vdots	...	\vdots	\vdots	\vdots	...	\vdots	\vdots	...	\vdots	\vdots	\vdots
	1	0	...	0	0	0	...	0	0	...	0	0	0
		1	...	$\hat{n}+1$	$\hat{n}+2$	$\hat{n}+3$...	$\hat{n}+m+2$	$\hat{n}+m+3$...	$\hat{n}+2m+2$	$\hat{n}+2m+3$	$\hat{n}+2m+4$

Figure 6: The xy -plane for $z = k+1$, $k = 1, \dots, n_z$, contains the literal segments of variable z_k . By $[\cdot]$ we mean that this point does not appear at all levels z . Every such level must have an equal number of c_1 -entries in line $\hat{n} + 1$ and column $\hat{n} + 1$; also notice that it is not possible to simultaneously have non-zero entries both for row and column i at the same level z .

$\hat{n}+2m+4$	0	...	0	0	0	...	0	0	...	0	0	0
$\hat{n}+2m+3$	0	...	0	0	c_3	...	c_3	0	...	0	0	0
$\hat{n}+2m+2$	0	...	0	0	0	...	0	0	...	0	0	0
...	\vdots	...	\vdots	\vdots	\vdots	...	\vdots	\vdots	...	\vdots	\vdots	\vdots
$\hat{n}+m+3$	0	...	0	0	0	...	0	0	...	0	0	0
$\hat{n}+m+2$	0	...	0	0	0	...	$[c_2]$	0	...	0	c_3	0
...	\vdots	...	\vdots	\vdots	\vdots	...	\vdots	\vdots	...	\vdots	\vdots	\vdots
$\hat{n}+3$	0	...	0	0	$[c_2]$...	0	0	...	0	c_3	0
$\hat{n}+2$	0	...	0	0	0	...	0	0	...	0	0	0
$\hat{n}+1$	0	...	0	0	0	...	0	0	...	0	0	0
...	\vdots	...	\vdots	\vdots	\vdots	...	\vdots	\vdots	...	\vdots	\vdots	\vdots
1	0	...	0	0	0	...	0	0	...	0	0	0
	1	...	$\hat{n}+1$	$\hat{n}+2$	$\hat{n}+3$...	$\hat{n}+m+2$	$\hat{n}+m+3$...	$\hat{n}+2m+2$	$\hat{n}+2m+3$	$\hat{n}+2m+4$

Figure 7: The xy -plane for $z = 1$. The exact position of every c_2 depends on the clause in which the variable appears: every time a variable z_k appears in clause l , we introduce two literal segments lying on level $z = k + 1$, through c_1 -points, and a clause segment perpendicular to this plane, through a c_2 -point, so that it intersects with the non-dummy literal segment.

$\hat{n}+2m+4$	0	...	0	0	0	...	0	0	...	0	0	0
$\hat{n}+2m+3$	0	...	0	0	0	...	0	0	...	0	0	0
$\hat{n}+2m+2$	0	...	0	c_5	0	...	0	0	...	0	0	0
...	\vdots	...	\vdots	\vdots	\vdots	...	\vdots	\vdots	...	\vdots	\vdots	\vdots
$\hat{n}+m+3$	0	...	0	c_5	0	...	0	0	...	0	0	0
$\hat{n}+m+2$	0	...	0	c_5	0	...	0	0	...	0	0	0
...	\vdots	...	\vdots	\vdots	\vdots	...	\vdots	\vdots	...	\vdots	\vdots	\vdots
$\hat{n}+3$	0	...	0	c_5	0	...	0	0	...	0	0	0
$\hat{n}+2$	0	...	0	0	c_5	...	c_5	c_5	...	c_5	0	0
$\hat{n}+1$	0	...	0	0	0	...	0	0	...	0	0	0
...	\vdots	...	\vdots	\vdots	\vdots	...	\vdots	\vdots	...	\vdots	\vdots	\vdots
1	0	...	0	0	0	...	0	0	...	0	0	0
	1	...	$\hat{n}+1$	$\hat{n}+2$	$\hat{n}+3$...	$\hat{n}+m+2$	$\hat{n}+m+3$...	$\hat{n}+2m+2$	$\hat{n}+2m+3$	$\hat{n}+2m+4$

Figure 8: The xy -plane for $z = \hat{n} + 2m + 4$.




Figure 9: The construction in three dimensions. The red lines depict literal segments and the green lines clause segments as they are positioned in the space.

Combining Lemmas 6.4 and 6.7 immediately yields our main theorem for this section:

Theorem 6.8. *It is NP-hard to approximate 3OPTIMALAUCTIONDESIGN better than 0.05%.*

7 Discussion and Open Problems

Even though in this paper we focused on deterministic mechanisms, our geometric characterization has interesting consequences for randomized mechanisms. Remember that for the discrete case the optimal (deterministic) mechanism immediately follows from solving the integer program of Problem C in Section 5.1, i.e. computing a maximum weight independent set in the corresponding n -partite graph. Our first observation is that the linear programming relaxation of this integer program corresponds to computing the optimal randomized mechanism. For two players, where the graph is bipartite and the integer program is totally unimodular, the optimum integer solution is also the optimum of the relaxed linear program. Therefore, for two bidders, the program of Problem C computes a deterministic mechanism that is optimal among all randomized mechanisms: this is reminiscent of Myerson's original result, where the deterministic mechanism obtained is optimal for the (larger) class of *Bayesian truthful* randomized mechanisms [21]. For a constant number of three or more bidders, the generalization of our geometric characterization yields polynomial-time algorithms for computing the optimal randomized auction, in sharp contrast with the intractability of computing the optimal deterministic auction, even for three bidders; of course for a large number of bidders, the size of this linear program may become exponentially large and therefore this approach is infeasible. For an alternative linear program that computes the optimal randomized auction for any number of bidders, when the distribution is given explicitly, the reader is referred to [11].

An important open problem of this work is to close the gap between the best approximation algorithm known for the optimal auction problem (currently .60) [11] and the inapproximability bound (currently about .9995). We believe that progress there is attainable. Interestingly, our work implies an approximation of $2/n$ for n players: before having the bidders announce their bids, the auctioneer looks at their joint distribution and privately runs the optimal auction for all possible pairs of players. Since solving for the optimal auction is nothing but a maximum

weight independent set problem on the corresponding graph, it is easy to prove that the profit of the best of those $\binom{n}{2}$ auctions is at least $2/n$ of the overall profit. The auctioneer then rejects a priori all but the bidders who were part of the most profitable two-bidder auction and then runs it. The overall auction is obviously truthful as long as bidders are rejected before even submitting their bids. For 3 bidders this gives an approximation ratio of $2/3$, improving over Ronen’s auction [26], but for $n \geq 4$ the approximation ratio drops below $1/2$.

References

- [1] Mark Armstrong. Optimal multi-object auctions. *Review of Economic Studies*, 67(3):455–81, July 2000.
- [2] Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. Budget constrained auctions with heterogeneous items. In *STOC*, pages 379–388, 2010.
- [3] Jeremy Bulow and John Roberts. The simple economics of optimal auctions. *Journal of Political Economy*, 97(5):1060–90, October 1989.
- [4] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic pricing via virtual valuations. In *ACM Conference on Electronic Commerce*, pages 243–251, 2007.
- [5] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *STOC*, pages 311–320, 2010.
- [6] Xi Chen and Xiaotie Deng. Settling the complexity of two-player nash equilibrium. In *Focs*, pages 261–272, 2006.
- [7] Jacques Cremer and Richard P McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2):345–61, March 1985.
- [8] Jacques Cremer and Richard P McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–57, November 1988.
- [9] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. In *STOC*, pages 71–78, 2006.
- [10] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*, pages 129–138, 2010.
- [11] Shahar Dobzinski, Hu Fu, and Robert Kleinberg. Optimal auctions with correlated bidders are easy. In *STOC*, 2011.
- [12] Lawrence C. Evans. Partial differential equations and monge-kantorovich mass transfer (surveypaper. In *Current Developments in Mathematics, 1997, International Press*, 1999.
- [13] Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, Andrew Wright, and Michael Saks. Competitive auctions. In *Games and Economic Behavior*, pages 72–81, 2002.
- [14] Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *SODA*, pages 1164–1173, 2005.
- [15] Jason Hartline. Lectures on approximation in mechanism design. Lecture notes, Northwestern University, 2010.
- [16] Jason D. Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *STOC*, pages 75–84, 2008.
- [17] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*, pages 225–234, 2009.
- [18] Philippe Jehiel, Moritz Meyer-ter Vehn, and Benny Moldovanu. Mixed bundling auctions. *Journal of Economic Theory*, 134(1):494–512, May 2007.
- [19] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413, 1999.

- [20] R Preston McAfee and Philip J Reny. Correlated information and mechanism design. *Econometrica*, 60(2):395–421, March 1992.
- [21] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [22] Noam Nisan and Amir Ronen. Algorithmic mechanism design (extended abstract). In *STOC*, pages 129–140, 1999.
- [23] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [24] Christos H. Papadimitriou. Algorithms, games, and the internet. In *ICALP*, pages 1–3, 2001.
- [25] Christos H. Papadimitriou, Michael Schapira, and Yaron Singer. On the hardness of being truthful. In *FOCS*, pages 250–259, 2008.
- [26] Amir Ronen. On approximating optimal auctions. In *ACM Conference on Electronic Commerce*, pages 11–17, 2001.
- [27] Amir Ronen and Amin Saberi. On the hardness of optimal auctions. In *FOCS*, pages 396–405, 2002.
- [28] Tim Roughgarden. The price of anarchy is independent of the network topology. *J. Comput. Syst. Sci.*, 67(2):341–364, 2003.
- [29] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.
- [30] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.