

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

The Complexity of Optimal Auction Design

Permalink

<https://escholarship.org/uc/item/0132p6nt>

Author

Pierrakos, Georgios

Publication Date

2013

Peer reviewed|Thesis/dissertation

The Complexity of Optimal Auction Design

by

Georgios Pierrakos

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Christos H. Papadimitriou, Chair
Professor Satish Rao
Associate Professor David Ahn

Spring 2013

The Complexity of Optimal Auction Design

Copyright 2013
by
Georgios Pierrakos

Abstract

The Complexity of Optimal Auction Design

by

Georgios Pierrakos

Doctor of Philosophy in Computer Science

University of California, Berkeley

Professor Christos H. Papadimitriou, Chair

This dissertation provides a complexity-theoretic critique of Myerson’s theorem [57], one of Mechanism Design’s crown jewels, for which Myerson was awarded the 2007 Nobel Memorial Prize in Economic Sciences. This theorem gives a remarkably crisp solution to the problem faced by a monopolist wishing to sell a single item to a number of interested, rational bidders, whose valuations for the item are distributed *independently* according to some given distributions; the monopolist’s goal is to design an auction that will maximize her expected revenue, while at the same time incentivizing the bidders to bid their true value for the item. Myerson solves this problem of designing the *revenue-maximizing* auction, through an elegant transformation of the valuation space, and a reduction to the problem of designing the *social welfare-maximizing* auction (i.e. allocating the item to the bidder who values it the most). This latter problem is well understood, and it admits a deterministic (i.e. the auctioneer does not have to flip any coins) and simple solution: the Vickrey (or second-price) auction. In the present dissertation we explore the trade-offs between the plausibility of this result and its tractability:

First, we consider what happens as we shift away from the simple setting of Myerson to more complex settings, and, in particular, to the case of bidders with arbitrarily *correlated* valuations. Is a characterization as crisp and elegant as Myerson’s still possible? In Chapter 2 we provide a negative answer: we show that, for three or more bidders, the problem of computing a deterministic, ex-post incentive compatible and individually rational auction that maximizes revenue is NP-complete –in fact, inapproximable. Even for the case of two bidders, where, as we show, the revenue-maximizing auction is easy to compute, it admits nonetheless no obvious natural interpretation à-la Myerson.

Then, motivated by the subtle interplay between social welfare- and revenue-maximizing auctions implied by Myerson’s theorem, we study the trade-off between those two objectives for various types of auctions. We show that, as one moves from the least plausible auction format to the most plausible one, the problem of reconciling revenue and welfare becomes less and less tractable. Indeed, if one is willing to settle for randomized solutions, then auctions that fare well with respect to both objectives simultaneously are possible, as

shown by Myerson and Satterthwaite [56]. For deterministic auctions on the other hand, we show in Chapter 3 that it is NP-hard to exactly compute the optimal trade-off (Pareto) curve between those two objectives. On the positive side, we show how this curve can be approximated within arbitrary precision for some settings of interest. Finally, when one is only allowed to use variants of the simple Vickrey auction, we show in Chapter 4 that there exist auctions that achieve constant factor approximations of the optimal revenue and social welfare *simultaneously*.

to my mother

Contents

Contents	ii
List of Figures	iv
List of Tables	vi
1 Introduction	1
1.1 Internet and the battle of incentives	1
1.2 Algorithmic Mechanism design	2
1.3 Single item auctions and Myerson's theorem	3
1.4 Beyond Myerson's theorem: our contribution	7
2 Optimal Deterministic Auctions with Correlated Priors	10
2.1 Introduction	10
2.2 The geometry of optimal auctions	11
2.3 Two bidders: the discrete case	19
2.4 Two bidders: the general case	22
2.5 Three bidders: NP-completeness	31
2.6 Discussion and open problems	45
3 Efficiency-Revenue Trade-offs in Auctions	47
3.1 Introduction	47
3.2 Preliminaries	52
3.3 The complexity of Pareto optimal auctions	53
3.4 An FPTAS for 2 bidders	61
3.5 The case of n bidders	63
3.6 Open questions	65
4 Simple, Optimal and Efficient Auctions	66
4.1 Introduction	66
4.2 The regular, independent case	69
4.3 The non-regular, i.i.d. case	76

Bibliography

List of Figures

2.1	A pair of valid allocation rules.	12
2.2	The revenue curves R , \hat{R} and R'	15
2.3	The allocation regions of an optimal auction for a product distribution.	16
2.4	The allocation curves discussed in Lemmas 3 and 4.	18
2.5	The allocation regions of an optimal auction for i.i.d. bidders.	18
2.6	The bipartite graph for 2 bidders with $N = 2$	20
2.7	The graph consisting of two grids, one for each bidder, for $n = 3$	23
2.8	The solution of Problem A (red and purple curves) resulting from the solution of Problem C (blue and light-blue shaded regions) as described in Lemma 8.	27
2.9	The reduction of Lemma 10, shown here for 2 bidders for ease of depiction.	32
2.10	The gadget of 3SEGMENTS resulting from a single clause of TYPED3SAT. This picture is not accurate –especially with respect to the location of the segments’ apices– and is solely meant to convey the basic idea of the reduction. We give a precise example in the next section.	35
2.11	The usage of scaffolding segments.	35
2.12	The plane $z = k + 1$, $k = 1, \dots, n_z$, 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 $\hat{n} + 2$ and column $\hat{n} + 2$; also, since every variable appears at most 7 times, there can be at most 7 entries in line and column $\hat{n} + 2$ for each such level.	37
2.13	The plane $z = 1$. 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 at $(l, l, 1)$, so that it intersects with the non-dummy literal segment. The only reason why a c_2 point might be missing is because of a clause of the form $(x_i \vee y_j)$	38
2.14	The gadgets resulting from $(x_1 \vee y_1 \vee \bar{z}_1) \wedge (\bar{x}_1 \vee \bar{y}_1) \wedge (x_2 \vee y_1 \vee z_1)$	39
2.15	The full set of segments for $(x_1 \vee y_1 \vee \bar{z}_1) \wedge (\bar{x}_1 \vee \bar{y}_1) \wedge (x_2 \vee y_1 \vee z_1)$	40
2.16	The plane $z = \hat{n} + 2m + 4$	43

- 3.1 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. 48
- 3.2 The objective space of all two-bidders auctions when $(v_1^1, v_1^2) = (11, 20)$, $(v_2^1, v_2^2) = (2, 5)$ and $(f_i^1, f_i^2) = (1/3, 2/3)$ for $i = 1, 2$ 49

List of Tables

3.1	An instance with an exponential size Pareto set	61
4.1	(α, β) stands for α -approximation for welfare and β -approximation for revenue. Notice that our result for regular distributions gets a handle on the whole Pareto boundary achieved by the Vickrey auction with non-anonymous reserve prices. .	67

Acknowledgments

There is very little I could say about Christos, his caring approach as an advisor and a teacher, his taste in research problems, and his ability to focus on what really matters, that could top off the accolades of his colleagues and the kind words of his many students before me. So, instead, I want to thank him from my heart for teaching me the importance of giving other people the freedom to develop their own course of action, in research, work and life in general; and of respecting others' choices, even when you are unsure or disagree with them. I can only aspire to be more like Christos in that respect, as a team member and a mentor to others, as a citizen, and ultimately as a parent.

I want to thank the theory faculty at Berkeley for the relaxed atmosphere they have chosen for the group; in particular, Satish Rao and David Ahn for serving on my dissertation and quals committees, Richard Karp for serving on my quals committee, Umesh Vazirani for his wit and Alistair Sinclair for his unforgettable lectures; thanks to all the theory students and post-docs too, for their collegiality and upbeat attitude. Back in Greece, I'd like to thank Elias Koutsoupias for being a friend and a role-model as a researcher and a teacher, and for sending me on this journey; Aris Pagourtzis for his help through my first research attempts, and Stathis Zachos for inspiring me to focus on theoretical computer science.

I am grateful to my co-authors, Gagan Aggarwal, Evangelos Bampas, Yang Cai, Constantin Daskalakis, Ilias Diakonikolas, Rafael Frongillo, Elias Koutsoupias, Aranyak Mehta, Aris Pagourtzis, Christos Papadimitriou, Katerina Potika, Yaron Singer, Vasilis Syrgkanis and Greg Valiant, for all the things they taught me; special thanks to Yaron, for helping me get out of my academic comfort zone and take a first peek into the business world. I'd also like to thank Christian Borgs, Jennifer Chayes and Adam Kalai at Microsoft Research New England, as well as Gagan Aggarwal and Aranyak Mehta at Google Mountain View, for two very enjoyable summer internship experiences. The research in this dissertation was made possible through a UC Regents' Fellowship, NSF grant CCF-0964033, a Google University Research Award and a scholarship by the Onassis Public Benefit Foundation; I gratefully acknowledge their support.

Life in Berkeley would not have been the same if it was not for all the tenants (both short- and long-term) of the legendary 1780 Spruce St apartment: thanks to Antonis, Christos, Dimitris, Eleftherios and Kostas for making this apartment a home. Also, thanks to the fierce warriors of the 300 and the Tzatziki Turbo, the two greek soccer teams that made weeknights something worth looking forward to.

Above all, I'd like to thank my mother for everything she has done for me, and my extended big-fat-greek family for all their love and support. And Elli, for making those nine hour, red-eye, coast to coast flights totally worth it.

Chapter 1

Introduction

1.1 Internet and the battle of incentives

The rise of the Internet over the past couple of decades has profoundly transformed the way we build and analyze computer science systems. One of the biggest challenges so far has been to understand and address the issue of conflicting incentives, which unavoidably emerge when a large collection of rational individuals -each one with their own objective- interact with each other. Economics, and in particular Game Theory, which have long been studying interactions between rational agents, were instrumental in providing computer science with the right tools to analyze those large scale computer systems. Meanwhile, economists have also started using those systems as a promising arena on which to test their theories and predictions. It should therefore come as no surprise that the Internet sparked an ongoing research interaction between computer science and economics, resulting in a fruitful exchange of ideas, where each discipline is having a marked effect upon the other:

On one hand, the computer science community is reinventing its approach to *analyzing* its systems. In an attempt to acknowledge the fact that the input is no longer bound to be “worst case”, but rather “utility-maximizing”, and that it may follow some given distribution (the standard assumption in economics), computer science theory occasionally deviates from its standard worst case input assumptions. Economics have been all the more influential in the *design* of systems, where concepts such as auctions, that have originated within game theory, have become the workhorse and theoretical foundation of the multi-billion dollar industry of internet advertising.

On the other hand, computer science’s main intellectual export, namely complexity theory, has driven the development of the field of *Algorithmic Game Theory* [62, 59], 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 [16, 23], the trade-offs between complexity, approximation, and incentive-compatibility in social welfare-maximizing mechanism design [59, 64], and (in an extended sense that includes on-line algorithms as a

part of complexity theory) the price of anarchy [50, 71, 70].

The present dissertation advances the latter research agenda by focusing this “computational lens” on Myerson’s theorem [57], one of Mechanism Design’s most celebrated results. Myerson considers the problem faced by a monopolist who wishes to sell a single item to a number of interested buyers by means of an auction, in a way that will maximize her expected revenue. The difficulty is of course that the buyers are strategic and must therefore be appropriately incentivized in order to comply to the auction’s rules. Under certain assumptions, Myerson is able to provide a very clean and elegant solution to this problem; our goal in this work, is to explore what happens when we relax some of those assumptions.

Before diving into the specifics of this problem, in the next section we do a quick overview of Mechanism Design, and briefly mention some of the results from this past decade that gave rise to the (still nascent) field of *Algorithmic Mechanism Design*.

1.2 Algorithmic Mechanism design

Summarized in a single sentence, *Mechanism Design* is nothing more than inverse game theory: given a desired outcome, it is the problem of designing a game (the mechanism) which implements this outcome in equilibrium, i.e. in a state from which no agent has an incentive to deviate.

In its full generality, mechanism design is (literally) an impossible problem: In one of the field’s early results, Gibbard [39] and Satterthwaite [73] show that designing a mechanism for agents with arbitrary utilities cannot be accomplished in any meaningful way. Fortunately, the problem is much easier once we introduce money and assume that agents have quasi-linear utilities, i.e. their utility is linear in the payments, a very common assumption in economics. Therefore, most of the subsequent work in mechanism design studies settings with quasi-linear utilities. Even though mechanism design encompasses many problems, in this dissertation we focus on the problem of *auction design*, arguably the most extensively studied one. In an auction design problem we have a number of items for sale and a set of interested bidders, for which all we know (at best) is their distribution of valuations for the items; our goal then is to design an auction, namely an allocation rule (who gets which item), and a payment rule (how much we will charge each person). Thanks to the celebrated *revelation principle* [57], as long as the allocation and payment rules of the auction are such that it is always in every bidder’s best interest to bid her true value (a property called *incentive compatibility*), we can wlog restrict our attention to direct revelation mechanisms, i.e. mechanisms where the bidders will reveal their true value for the item.

There are of course many possible “desired outcomes” that the auction designer could have, but most of the research focuses on the natural objectives of social welfare and revenue. In social welfare-maximizing (or *efficient*) auction design the goal is to award the items to the bidders who value them the most, while in revenue-maximizing (or *optimal*) auction design the goal is to generate as much revenue as possible.

For an auctioneer who has a single item to sell and is interested in efficiency, the celebrated *second-price* or *Vickrey* auction [75] (arguably mechanism design’s most famous intellectual export) is the way to go: simply allocate the item to the highest bidder, and charge her the second highest price. Social welfare optimality is immediate, while it is not hard to see that, under this allocation and payment rules, no bidder has an incentive to lie about her value for the item. The Vickrey auction has two additional desirable properties: first of all, it does not rely on any distributional assumptions about the bidders’ valuations (which is not the case with the revenue optimal auction as we shall soon see), and moreover it is easy to implement. Vickrey, Clarke and Groves [75, 18, 43] used the idea behind the Vickrey auction –charging every bidder her negative externality– to come up with the VCG mechanism: an efficient, incentive compatible auction for multiple items, albeit one that cannot be implemented (or even reasonably approximated) in polynomial time. This tension between incentive compatibility and computational efficiency has been the focus of a recent surge in the computer science literature studying *combinatorial auctions*, commencing with [59] (which also started the field of algorithmic mechanism design) and culminating with [64] (where the first dichotomy between approximability and incentive compatible approximability was proved), to be fully resolved only quite recently in a series of papers [31, 35, 33].

Moving to the other important objective of revenue, as we mentioned earlier, the landmark paper here is undoubtedly that of Roger Myerson [57]. There were several follow-up papers by economists (for example [19, 20, 4, 48]), but it was not until very recently (around 2002) that computer scientists started following up on this literature. In [13, 46, 26] for example, motivated by Myerson’s astonishing result that the optimal auction for certain simple settings is simply a second-price auction with reserve prices, the authors examine the extent to which simple auctions can achieve good approximations in more general 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 [40]. This direction became especially interesting after [45] developed a framework 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 [26].

Because of the difficulty of addressing the general problem, most of the work for revenue-maximizing auction design has focused on auctions where there is only a single item for sale (we mention some notable exceptions in Section 1.4). In the next section we formally define this setting, which is going to be the focus of this dissertation as well.

1.3 Single item auctions and Myerson’s theorem

The setting. 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 distribution which is public knowledge.

In Chapter 2 we will consider joint distributions whose density function we denote by $f(\mathbf{v})$. We will consider both discrete and continuous f . 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 assume that the support is $[0, 1]^n$ and we follow Myerson in making the analytically convenient assumption that $f(x) > 0$ 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 f very little. In stating an algorithm for the two-dimensional continuous case (Section 2.4), we shall also assume that f is Lipschitz-continuous¹ 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 $Sup(f)$ denote the finite support of the joint discrete distribution f . Then f is presented as a finite set of $|Sup(f)|$ $(n + 1)$ -tuples of the form $(v_1^i, v_2^i, \dots, v_n^i, f^i)$, one for each point $v^i \in Sup(f)$, where f^i is the probability mass concentrated at the point (v^i) of the support.

In Chapters 3 and 4, we will assume that the distributions of the bidders are independent, and we will use F_i to denote the distribution of bidder i . The distributions $\{F_i\}_i$ are not necessarily identical. For simplicity we assume that all F_i 's are differentiable, so we can define the corresponding probability density functions as $f_i(x) = F_i'(x)$. In the discrete case, we will again use $f_i(\cdot)$ to denote the probability density function of bidder i over the discrete support Sup_i . We will also use v_i^k and f_i^k , $k = 1, \dots, |Sup_i|$, to denote the k -th smallest element in the support of bidder i and its probability mass respectively.

Definition of an auction. A single item auction \mathcal{A} 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 . For the most part of this dissertation we focus our attention on deterministic auctions so that $x_i(\cdot) \in \{0, 1\}$. We demand from our auctions to satisfy the standard constraints of incentive compatibility (IC), individual rationality (IR) and no positive transfers (NPT) in the ex-post sense, namely the following:

- **IC:** $\forall i, v_i, v'_i, v_{-i} : v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i}) \geq v_i x_i(v'_i, v_{-i}) - p_i(v'_i, v_{-i})$
- **IR:** $\forall i, v_i, v_{-i} : v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i}) \geq 0$
- **NPT:** $\forall i, v_i, v_{-i} : p_i(v_i, v_{-i}) \geq 0$

Finally, we are interested in the objectives of revenue and social welfare, defined as follows:

$$\text{Rev}(\mathcal{A}) = \mathbb{E} \left[\sum_{i=1}^n p_i(v_1, \dots, v_n) \right] \quad \text{and} \quad \text{SW}(\mathcal{A}) = \mathbb{E} \left[\sum_{i=1}^n v_i \cdot x_i(v_1, \dots, v_n) \right],$$

¹A function f is Lipschitz-continuous if there exists a constant λ such that $|f(x) - f(y)| \leq \lambda|x - y|$ for all x, y .

where the above expectations are with respect to value vectors $v = (v_1, \dots, v_n)$ drawn from the joint distribution with probability density function f or $\times_i f_i$.

A characterization. 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 auctions that we are chiefly 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 bidder given the valuations of the other bidders: $t_i(v_{-i}) = \inf\{v_i \in [0, 1] \mid x_i(v_i, v_{-i}) = 1\}$. In complete analogy, in the discrete case $t_i(v_{-i}) = \min\{v_i \in \text{Sup}_i \mid x_i(v_i, v_{-i}) = 1\}$, where Sup_i is the support of bidder i . The following theorem provides a characterization of auctions for the setting we are interested in; a proof can be found in [60].

Theorem 1. *A deterministic auction satisfies IC, IR and NPT if and only if the following conditions hold.*

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

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

Moreover, 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 optimal auction (with respect to either objective) will always be on the support of bidder i .

Relying on the above characterization, we will describe our auctions (in the discrete case) using the concept of an *allocation matrix* A : a $|\text{Sup}_1| \times \dots \times |\text{Sup}_n|$ matrix where entry (i_1, \dots, i_n) corresponds to the tuple $(v_1^{i_1}, \dots, v_n^{i_n})$ of bidders' 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 her support for which she 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.

Myerson's theorem. In [57] Myerson introduced the notion of a bidder's *virtual valuation function* ϕ_i , defined as follows:

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}.$$

In terms of this notion, we say a distribution F_i is *regular* if the virtual valuation function ϕ_i is non-decreasing, and that it satisfies the *monotone hazard rate condition* if the ratio $\frac{1-F_i(x)}{f_i(x)}$ is non-increasing. For distributions that are non-regular, Myerson's ironing technique [57] can be used to get the corresponding *ironed* virtual valuation function $\hat{\phi}_i(v_i)$. The following result is central to Myerson's analysis, and we will also use it in this dissertation.

Proposition 1 (Myerson's Lemma). *For any deterministic auction $\{(x_i, p_i)\}_i$, satisfying IC and IR, we can express the expected payment of bidder i as follows (with the expectation taken over the bidders' values):*

$$\mathbb{E}[p_i(v_1, \dots, v_n)] = \mathbb{E}[\phi_i(v_i) \cdot x_i(v_1, \dots, v_n)].$$

Proof sketch. For the sake of simplicity assume that each bidder's valuation is supported on some continuous interval $T_i = [a_i, b_i]$. We will show how the above formula for the virtual valuation functions $\phi_i(v_i)$ is derived as a direct consequence of wanting to transform the following expression for revenue:

$$\int_{a_i}^{b_i} p_i(v) f_i(v) dv_i,$$

into the following expression for social welfare, where the valuations v_i are replaced with virtual valuations $\phi_i(v_i)$:

$$\int_{a_i}^{b_i} \phi_i(v_i) x_i(v) f_i(v) dv_i.$$

Using the characterization theorem (Theorem 1) this boils down to solving the following functional equation for $\phi_i(v_i)$, for every bidder i :

$$\begin{aligned} & \int_{a_i}^{b_i} p_i(v) f_i(v) dv_i = \int_{a_i}^{b_i} \phi_i(v_i) x_i(v) f_i(v) dv_i, \text{ for all } i, v_{-i} \\ \stackrel{\text{Theorem 1}}{\iff} & \int_{t_i(v_{-i})}^{b_i} t_i(v_{-i}) f_i(v) dv_i = \int_{t_i(v_{-i})}^{b_i} \phi_i(v_i) f_i(v) dv_i, \text{ for all } i, v_{-i} \\ \iff & \int_{t_i}^{b_i} t_i f_i(v) dv_i = \int_{t_i}^{b_i} \phi_i(v_i) f_i(v) dv_i, \text{ for all } i, t_i \end{aligned}$$

Using elementary calculus we can solve the above functional equation as follows:

$$\begin{aligned} & \int_{t_i}^{b_i} t_i f_i(v) dv_i = \int_{t_i}^{b_i} \phi_i(v_i) f_i(v) dv_i \\ \iff & t_i(1 - F_i(t_i)) = \int_{t_i}^{b_i} \phi_i(v_i) f_i(v) dv_i \\ \iff & \frac{\partial}{\partial t_i} (1 - F_i(t_i)) - t_i f_i(t_i) = -\phi_i(t_i) f_i(t_i) \end{aligned}$$

$$\iff \phi_i(t_i) = t_i - \frac{1 - F_i(t_i)}{f_i(t_i)}$$

□

What the proposition above implies is that, in order to run the revenue-maximizing auction, it suffices to transform the valuations into virtual valuations, and then run the social welfare-maximizing auction, namely the Vickrey or second-price auction, on those virtual valuations. (One caveat is that, for this to work, the resulting virtual valuations have to be increasing in the actual valuations. If this is not the case, the extra step of ironing is needed; since ironing is not central to this dissertation, we refer the bidder to [57, 44] for further discussion.)

In informal, computer science parlance, what Myerson established is the following:

Theorem 2 (Myerson’s Theorem–informal). *The problem of designing the revenue maximizing-auction for a single item and n bidders, whose valuations are distributed according to independent (but not necessarily identical) distributions, reduces (in polynomial time) to that of social welfare-maximization.*

1.4 Beyond Myerson’s theorem: our contribution

Myerson’s result is remarkable in several ways. While it is not the first important paper on auctions of course [75], 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. One of the most interesting aspects of his result is that the auctioneer’s and the bidders’ actions in a given auction situation are easy to compute: *the problem of designing the revenue-maximizing auction for a single item and independent valuations is in P*. On top of that, the resulting auction has a very simple and intuitive format: it is a Vickrey auction in a modified domain, that of virtual valuations. In fact, for the special setting where all bidders have valuations drawn from the same distribution, the resulting auction is nothing more than a Vickrey auction with a reserve price². Even though Myerson did not dwell on this aspect of his auction, i.e. its computational efficiency, it is clearly a sine qua non: no good auction *design* should involve solving intractable problems. Keeping this issue of computational tractability in mind, and looking at Theorem 2 from a computer scientist’s point of view, a couple of very natural questions arise.

Question 1: *Is it possible to obtain a similar result, namely a characterization of the optimal auction and an efficient way of computing it, for the case of more than one items?*

This question –arguably one of the most important questions in mechanism design– is still for the most part open, however a recent surge in the computer science and economics

²The reader can think of the reserve as an extra bidder with a constant valuation that depends on the common distribution of valuations – see Chapter 4 for a formal definition.

literature has provided a lot of new insights into this problem. We can broadly divide the papers into three categories. Most of the papers in the economics literature identify special cases of this problem that are more amenable to mathematical analysis, either by restricting the setting [4], or the class of auctions allowed [48], and then attempt to design the optimal auction for those cases. The other two trends follow the computer science approach of settling with a suboptimal, but computationally tractable solution, that provides a good approximation to the optimal one. One line of work (see for example [5, 14, 13, 3, 6, 49]) aims at designing simple (for the most part posted-price) auctions for the problem, and achieves constant factor approximations for a variety of multi-parameter settings. Another line of work studies the underlying LP implied by the IC and IR constraints, and provides solutions that can approximate the optimal auction within arbitrary precision ϵ , in time polynomial both in the size of the input and $1/\epsilon$ (see for example [9, 25, 11, 10]). Even though most of these papers do not enforce any particular structure on the desired solution à-priori, a very recent paper ([11]) showed that there is in fact structure in the optimal solution of the multi-parameter setting: in particular they prove that the optimal solution is a distribution over virtual VCGs, and they provide an efficient algorithm for computing this distribution. From a lower bound perspective, a recent pre-print shows that the problem of designing the optimal (randomized) auction for general multi-parameter settings is #P-hard [21].

Question 2: *Is it possible to obtain a similar result, namely a characterization of the optimal auction and an efficient way of computing it, for the case where the bidders' valuations are allowed to be arbitrarily correlated (i.e. not necessarily independent)?*

In Chapter 2 we provide a negative answer to this question: we revisit the problem of designing the revenue-maximizing single item auction, focusing on general joint distributions, either discrete or Lipschitz-continuous, seeking the optimal *deterministic*, ex-post incentive compatible and individually rational auction. We give a geometric characterization 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. From a philosophical standpoint -if one is willing to interpret NP-completeness as a sign of mathematical poverty and lack of sufficient structure- this result is a strong indication that a characterization theorem as clean and elegant as Myerson's is unlikely to exist for general joint distributions. Chapter 2 is based on the work in [63].

Question 3: *The reduction of Theorem 2 reveals a lot of similarities and implies a non-trivial interplay between the revenue- and the social welfare-maximizing auctions. Does this mean that these two objectives of revenue and social welfare are compatible with each other, and, if not, what is the trade-off between them?*

It is not hard to notice that the whole point of the Vickrey auction is to deliberately sacrifice auctioneer revenue in order to achieve efficiency and truthfulness (charging the second highest instead of the highest price). Myerson's auction on the other hand, in its very simple format of a Vickrey auction with a reserve price, may sacrifice efficiency by

never allocating the item to a bidder who has non-zero value for it, in order to achieve better revenue guarantees. In Chapter 3 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 auctions, it is easy to see (and we will argue about this later in more detail) that, given any point in the trade-off (Pareto) curve between revenue and welfare, we can efficiently compute an auction that achieves those revenue and welfare guarantees. We investigate whether one can achieve the same guarantees using *deterministic* auctions. 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 auctions 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. Chapter 3 is based on the work in [30].

Finally, in Chapter 4 we study the extent to which simple auctions can simultaneously achieve good revenue and efficiency guarantees in single item settings. Motivated by the optimality of the second price auction with monopoly reserves when the bidders' values are drawn i.i.d. from regular distributions [57], and its approximate optimality when they are drawn from independent regular distributions [46], we focus our attention to the second price auction with general (not necessarily monopoly) reserve prices, arguably one of the simplest and most intuitive auction formats. As our main result, we show that, for a carefully chosen set of reserve prices, this auction guarantees at least 20% of both the optimal welfare and the optimal revenue, when the bidders' values are distributed according to independent, not necessarily identical, regular distributions. We also prove a similar guarantee, when the values are drawn i.i.d. from a –possibly irregular– distribution. Chapter 4 is based on the work in [24].

Chapter 2

Optimal Deterministic Auctions with Correlated Priors

2.1 Introduction

As we mentioned in the Introduction, Myerson left open the case in which the valuations are correlated; in subsequent work, Crémer and McLean [19, 20] consider correlated valuations and solve the problem for the case where auctions are only required to be *interim* individually rational (i.e. individually rational only in expectation over the other bidders' valuations). 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 [53], 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.

Surprisingly, despite the recent surge in the computer science literature on mechanism design, there has been little progress in looking at Bayesian auctions à-la Myerson from this point of view. In particular, Ronen [67] came up with an auction for the correlated case that achieves half of the optimal revenue, while Ronen and Saberi [68] showed that no “ascending auction” can do better than $7/8$, and they conjectured that all relevant auctions are ascending (which we disprove 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.

In this chapter 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 [37], and from this a fully polynomial approximation scheme for two bidders when the distribution is continuous enough and accessible through an oracle. As an aside, we also sketch in the last section a $2/3$ approximation for three bidders, improving the previously best known approximation of [67].

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 [57], the most important of them being that they both admit a marginal revenue interpretation in the spirit of [8]. 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 auctions explored; we achieve that by imposing an additional technical condition which limits the design space into a subset of all auctions, but one which still contains all of the optimal ones.

The work most related to ours is that of Dobzinski, Fu and Kleinberg [32], who also study the problem of designing the optimal auction for the correlated setting. They obtain a collection of interesting results, which however are quite complementary to ours: their work focuses on randomized auctions, a large number of bidders and approximation, while ours focuses on deterministic auctions, a small number of bidders and computational complexity. Based, among others, on insights from [67], they arrive at efficient algorithms for computing the optimal randomized auction that is truthful in expectation, and a constant factor approximation of the optimal deterministic auction for any number of bidders.

2.2 The geometry of optimal auctions

Here we focus on the two-bidder case, and provide an alternative geometric interpretation of the auction 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 2.5, where we will use our geometric characterization to establish the inapproximability of the problem for three 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 2.1): 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 $\beta(x)$.

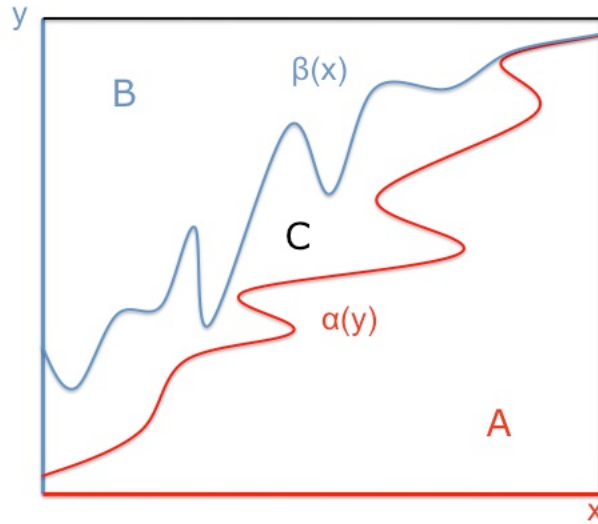


Figure 2.1: A pair of valid allocation rules.

Definition 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 by 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 valuation pair (x, y) at most one bidder gets the item.

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

Definition 2. Let m_1 (resp. m_2) be the marginal profit contribution of a valuation pair (x, y) for bidder 1 (resp. 2), defined as:

$$m_1(x, y) = -\frac{\partial}{\partial x} \left[\max_{x' \geq x} x' \cdot \int_{x'}^1 f(t, y) dt \right], m_2(x, y) = -\frac{\partial}{\partial y} \left[\max_{y' \geq y} y' \cdot \int_{y'}^1 f(x, t) dt \right]$$

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

Intuitively, $m_1(x, y) dx dy$ 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) . Near the end of this section we discuss the intuition behind these functions and their relation to Myerson's virtual valuation functions in more detail.

Definition 3. Call a valid allocation pair **proper** if it satisfies the following condition:

$$\left. \begin{aligned} \alpha(y) \cdot \int_{\alpha(y)}^1 f(t, y) dt &= \max_{x' \geq \alpha(y)} x' \cdot \int_{x'}^1 f(t, y) dt, \quad \text{for all } y \\ \beta(x) \cdot \int_{\beta(x)}^1 f(x, t) dt &= \max_{y' \geq \beta(x)} y' \cdot \int_{y'}^1 f(x, t) dt, \quad \text{for all } x \end{aligned} \right\} \quad (2.1)$$

Intuitively, in a proper valid pair (α, β) , the curve α (resp. β) never goes through points that have zero marginal profit contribution $m_1(x, y)$ (resp. $m_2(x, y)$), as ensured by the first (resp. second) of the two equalities above.

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

Lemma 1. Let (α, β) be a proper valid allocation pair. Then the expected profit of any auction with payments defined as in Theorem 1 is:

$$\int_0^1 \int_{\alpha(y)}^1 m_1(x, y) dx dy + \int_0^1 \int_{\beta(x)}^1 m_2(x, y) dy dx$$

Proof. Let $p_1(x, y), p_2(x, y)$ be the payment functions induced by the allocation rule (α, β) according to Theorem 1 of the previous chapter. Then the expected profit of our auction is:

$$\begin{aligned} & \int_0^1 \int_0^1 p_1(x, y) f(x, y) dx dy + \int_0^1 \int_0^1 p_2(x, y) f(x, y) dx dy \\ &= \int_0^1 \left[\alpha(y) \cdot \int_{\alpha(y)}^1 f(t, y) dt \right] dy + \int_0^1 \left[\beta(x) \cdot \int_{\beta(x)}^1 f(x, t) dt \right] dx \\ &= \int_0^1 \left[\max_{x' \geq \alpha(y)} x' \cdot \int_{x'}^1 f(t, y) dt \right] dy + \int_0^1 \left[\max_{y' \geq \beta(x)} y' \cdot \int_{y'}^1 f(x, t) dt \right] dx \\ &= \int_0^1 \int_{\alpha(y)}^1 m_1(x, y) dx dy + \int_0^1 \int_{\beta(x)}^1 m_2(x, y) dy dx \end{aligned}$$

where in the first equality we used the characterization of truthful payments as every bidder's critical value, in the second equality we made use of condition (2.1) and in the last equality we made use of the definition of marginal profit contribution functions. \square

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

Lemma 2. *For any f and for any valid allocation pair (α, β) there is a proper valid pair (α', β') such that $Profit(\alpha', \beta', f) \geq Profit(\alpha, \beta, f)$.*

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

We start by defining the following sets of points:

$$\mathcal{Y} = \left\{ y \in [0, 1] \mid \alpha(y) \cdot \int_{\alpha(y)}^1 f(t, y) dt \neq \max_{x' \geq \alpha(y)} x' \cdot \int_{x'}^1 f(t, y) dt \right\}$$

$$\mathcal{X} = \left\{ x \in [0, 1] \mid \beta(x) \cdot \int_{\beta(x)}^1 f(x, t) dt \neq \max_{y' \geq \beta(x)} y' \cdot \int_{y'}^1 f(x, t) dt \right\}$$

as the set of all coordinates y (resp. x) where condition (2.1) is violated by function α (resp. β). Consider now the auction defined by the following allocation curves:

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

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

where—in the case of ties— $\arg\{\cdot\}$ returns the largest y or x respectively. By construction, the new pair (α', β') satisfies condition (2.1). 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:

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

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 be restated as the following variational calculus-type problem:

Definition 4 (Problem A).

$$\sup_{(\alpha, \beta) \in AB} \left\{ \int_0^1 \int_{\alpha(y)}^1 m_1(x, y) dx dy + \int_0^1 \int_{\beta(x)}^1 m_2(x, y) dy dx \right\}$$

In this section we have established the following theorem:

Theorem 3. *Finding the optimal auction for two bidders is equivalent to solving Problem A.*

Marginal profit contributions and virtual valuations

There is an intuitive connection between the marginal profit contribution and Myerson’s (ironed) virtual valuations [57]. 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 $x \cdot \int_x^1 f(t, y) dt$ 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 bidder fixed at y ; equivalently, one can write the expected revenue as a function $R(q)$ of the probability of sale $q = 1 - \int_x^1 f(t, y) dt$ (for a thorough discussion of this maneuver the reader is referred to [44]). 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 2.2 we show an example of a revenue curve $R(q)$ and its corresponding $\hat{R}(q)$ and $R'(q)$ curves.




Figure 2.2: The revenue curves R , \hat{R} and R' .

It has been already noted by Myerson [57] 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 2.1. Since the invariance of the allocation rule across ironed regions follows from Myerson’s analysis as well, Condition 2.1 is indeed not a loss of generality in the sense that all optimal auctions satisfy it. This is exactly what we proved (formally) in Lemma 2.

The geometry of optimal auctions for product distributions

We conclude this section by using the machinery developed here to study product distributions and to sketch an alternative “proof by picture” of (a weaker form of) Myerson’s result. In doing so, we will make use of our geometric interpretation of the allocation space, and of the notion of marginal profit contribution functions. Our result is weaker in that it focuses on two bidders and the space of deterministic, ex-post IC and IR auctions. Generalizing to more bidders is relatively straightforward, but we do not know how to extend the analysis to encompass more general auction formats. In particular, note that Myerson proves that his auction is optimal among the larger space of all Bayesian truthful auctions.

Recall that Myerson’s optimal auction first computes each bidder’s (ironed) virtual valuation $\phi_i(v_i)$ and it then runs Vickrey’s auction on the $\phi_i(v_i)$. One way to interpret the resulting auction is the following. First we set a reserve price for every bidder independently, based on her valuation and her prior distribution; this is the threshold value that a bidder needs to exceed in order to have a positive virtual value and therefore a chance at being allocated the item. Then, for every pair of valuations that are both above their respective reserve prices, the auction carefully optimizes the allocation based on the virtual valuation of each bidder. The allocation space for this auction looks like the one in Figure 2.3.

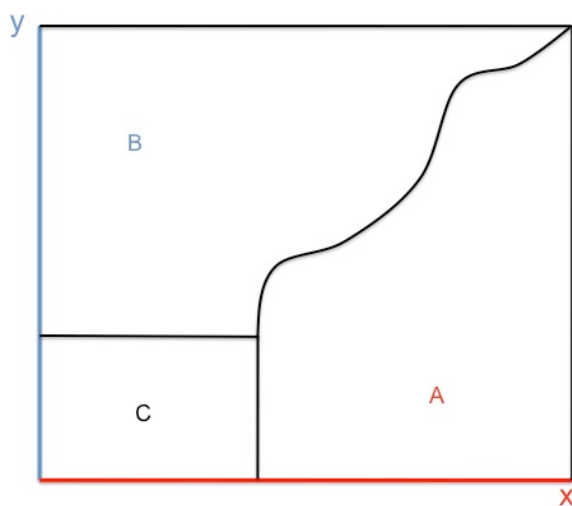


Figure 2.3: The allocation regions of an optimal auction for a product distribution.

In what follows we use the geometric interpretation of optimal auctions introduced earlier in this section to show that the optimal auction for two bidders and product distributions must indeed have the structure of Figure 2.3: namely, that the only region where the auctioneer may keep the item is a rectangular area at the bottom left of the unit square. This is in stark contrast with what happens in the case of correlated distributions, where the allocation region of the auctioneer may be arbitrarily shaped (subject to monotonicity constraints of course). We establish this result in two steps through the following two lemmas:

Lemma 3. *For a product distribution the allocation curves $\alpha(y)$ and $\beta(x)$ of the optimal auction are non-decreasing functions of y and x respectively (see Figure 2.4(a)).*

Proof sketch. Let (α, β) be a proper valid allocation pair with the additional property (which is not a loss of generality) that the area for which the auctioneer keeps the item is maximized; in other words, if for a given y there are two possible points $(\alpha(y), y)$ that achieve the same profit, we pick the rightmost such point (and analogously for points $(x, \beta(x))$). We prove that $\alpha(y)$ is non-decreasing and the proof for $\beta(x)$ is similar.

Suppose that $\alpha(y)$ is not non-decreasing, in particular that there exist y_1, y_2 with $y_1 < y_2$ and $\alpha(y_1) > \alpha(y_2)$. Consider then the allocation pair (α', β) , where $\alpha'(y) = \alpha(y)$ for all $y \neq y_1$ and $\alpha'(y_1) = \alpha(y_2)$, i.e. we “pull” the red curve at y_1 to the left until it takes the same value as in point y_2 . Because y_1 lies “below” y_2 and because of the non-crossing property satisfied by the original allocation pair (α, β) , pair (α', β) is still valid. The crucial observation now is that, because of the product distribution assumption, the marginal profit contribution functions for the values y_1 and y_2 are the same (i.e. $m_1(x, y_1) = m_1(x, y_2)$ for all $x \in [0, 1]$). Using this and the fact that $\alpha'(y_1) = \alpha(y_2)$ we get that

$$\int_{\alpha'(y_1)}^1 m_1(x, y_1) dx = \int_{\alpha(y_2)}^1 m_1(x, y_2) dx > \int_{\alpha(y_1)}^1 m_1(x, y_2) dx = \int_{\alpha(y_1)}^1 m_1(x, y_1) dx$$

where the inequality is strict because of the fact that ties in (α, β) are broken in favor of the rightmost point. Therefore, the profit of (α', β) is **strictly** larger than that of (α, β) , a contradiction. \square

Lemma 4. *For a product distribution the allocation regions look like the ones depicted in Figure 2.3, where the curve separating regions A and B is a non-decreasing function with respect to both x and y .*

Proof sketch. Suppose this is not true and that the allocation rules are only monotone – which must hold as we proved in Lemma 3; in particular suppose that they look like the ones depicted in Figure 2.4(a) (there are other cases, which nonetheless admit a similar analysis). Let $x_1 = \alpha(0)$; as in Lemma 3 we will exclude the possibility of a tie, so that it is always **strictly** better to extend the allocation region A all the way to x_1 for all $y \in [0, y_1]$ where $y_1 = \beta(x_1)$ (see Figure 2.4(b)). Repeating this argument all the way up to $y = 1$ we get the new auction defined by the black allocation rule of Figure 2.4(b). This auction is of the type depicted in Figure 2.3 and has a **strictly** better profit, a contradiction. \square

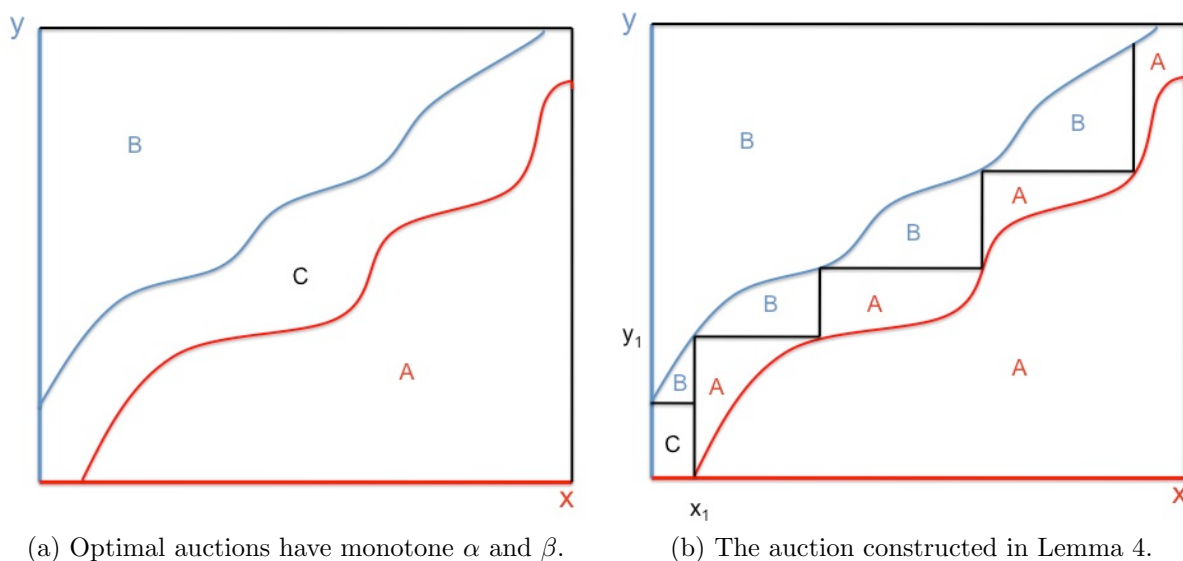


Figure 2.4: The allocation curves discussed in Lemmas 3 and 4.

Finally, for the special case where the distributions of the two agents are identical (and regular), it follows easily that the allocation regions need also be symmetric across the $y = x$ line. In terms of the Figure 2.3, this means that the bottom left rectangular region C where the auctioneer keeps the item is actually a square, and the curve separating regions A and B is actually a straight line. We therefore re-derive Myerson's result, that the optimal auction for bidders with i.i.d. valuations is the second price auction with a reserve price (the same for both bidders); see also Figure 2.5.




Figure 2.5: The allocation regions of an optimal auction for i.i.d. bidders.

2.3 Two bidders: the discrete case

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

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

Definition 5. *The discrete analogues of the marginal profit contribution functions (Definition 2) for each bidder are defined as follows:*

$$m_1(i, j) = \max \left\{ i \cdot \sum_{i' \geq i} f(i', j) - \sum_{i' > i} m_1(i', j), 0 \right\} \quad \text{for bidder 1}$$

$$m_2(i, j) = \max \left\{ j \cdot \sum_{j' \geq j} f(i, j') - \sum_{j' > j} m_2(i, j'), 0 \right\} \quad \text{for bidder 2}$$

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

Definition 6. *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 $\text{Sup}(f)$ into three sets A, B , and C , where $A = \{(i, j) \in \text{Sup}(f) : i \geq \alpha(j)\}$, $B = \{(i, j) \in \text{Sup}(f) : j \geq \beta(i)\}$, and $C = \text{Sup}(f) - A - B$. 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 $m_1(\alpha(j), j) = 0$. 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 $m_1(\alpha(j), j), m_2(i, \beta(i)) > 0$. We call such valid pairs **proper**. Finally, we define $\Pi(f)$ to be the set of all partitions A, B of $\text{Sup}(f)$ induced by proper valid pairs.*

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

$$\max_{(A, B) \in \Pi(f)} \sum_{(i, j) \in A} m_1(i, j) + \sum_{(i, j) \in B} m_2(i, j).$$

All that remains now is to provide a useful characterization of the set $\Pi(f)$. 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 7. *We call a pair of disjoint subsets (A, B) of $\text{Sup}(f)$ **monotone** if the following holds:*

- If $(i, j) \in A$ and $(i', j) \in \text{Sup}(f)$ with $i' > i$, then $(i', j) \in A$.
- If $(i, j) \in B$ and $(i, j') \in \text{Sup}(f)$ 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 $m_1(i, j) > 0$.
- If $(i, j) \in B$ and $(i, j') \notin B$ for all $j' < j$, then $m_2(i, j) > 0$.

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 profit contribution, or there is another optimal auction with this property in the case of a zero marginal profit contribution).

Let us define now a bipartite graph $G^f = (U, V, E)$:

- $U = \{u_{i,j} \mid i, j \in [N]\}$ and $V = \{v_{i',j'} \mid i', j' \in [N]\}$;
- $(u_{i,j}, v_{i',j'}) \in E$ if and only if $i \leq i'$ and $j \geq j'$. In other words, there is an edge between u and v if and only if, informally, (i', j') “lies to the southeast” of (i, j) (see Figure 2.6); (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 $m_1(i, j)$ and the weight of any node $v_{i,j}$ is $m_2(i, j)$.

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




Figure 2.6: The bipartite graph for 2 bidders with $N = 2$.

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

Lemma 5. *Let (A, B) be a pair of subsets of $\text{Sup}(f)$. Then $(A, B) \in \Pi(f)$ if and only if (A, B) is monotone and proper, and $\{u_{i,j} : (i, j) \in A\} \cup \{v_{i,j} : (i, j) \in B\}$ is an independent set of G^f .*

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

(Only if.) If $(A, B) \in \Pi(f)$ then by the definition of $\Pi(f)$ 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 $\{u_{i,j} : (i, j) \in A\} \cup \{v_{i,j} : (i, j) \in B\}$ is an independent set of G^f , assume towards contradiction that there are nodes $u_{i,j}$ and $v_{i',j'}$ such that $(i, j) \in A$ and $(i', j') \in B$, with an edge between $u_{i,j}$ and $v_{i',j'}$; from the construction of G^f , 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 6. *Let (A, B) be a pair of subsets of $\text{Sup}(f)$ such that the set $\{u_{i,j} : (i, j) \in A\} \cup \{v_{i,j} : (i, j) \in B\}$ is a maximum weight independent set of G^f , of minimum cardinality among all independent sets of the same weight. Then $(A, B) \in \Pi(f)$.*

Proof. It suffices to show that the set (A, B) is monotone and proper, and then the lemma follows from Lemma 5. Indeed, the monotonicity of (A, B) follows from the fact that the set $\{u_{i,j} : (i, j) \in A\} \cup \{v_{i,j} : (i, j) \in B\}$ 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 $m_1(i, j) = 0$, unless it also includes a node $u_{i',j}$ corresponding to some valuation (i', j) with $m_1(i', j) > 0$ for some $i' < i$ (and analogously for $v_{i,j}$). Hence, by definition, (A, B) is proper as well. \square

Theorem 4. *Given a discrete joint valuation distribution f for two bidders, the optimal ex-post IC and IR deterministic auction can be computed in time $O(|\text{Sup}(f)|^3)$.*

Proof. It follows from Lemma 6 that computing the optimal auction for two bidders with a joint valuation distribution f reduces to computing a maximum weight independent set on the induced bipartite graph G^f . In particular, the optimal allocation rule corresponds to a partition (A, B) such that $\{u_{i,j} : (i, j) \in A\} \cup \{v_{i,j} : (i, j) \in B\}$ is a maximum weight independent set of G^f , 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

2.4 Two bidders: the general case

In this section we return to the continuous two-bidder problem of Section 2.2. Our main result is an efficient algorithm that approximates the optimal solution of Problem A within an arbitrarily small additive ϵ . Our algorithm is quite natural: we first round the input, namely the bidders' distribution, in multiples of some constant ϵ , and then run the maximum weight independent set algorithm on the resulting $1/\epsilon \times 1/\epsilon$ grid. The analysis however is more elaborate: we use as our main tool a duality theorem, generalizing the duality between the maximum-weight independent set problem in a bipartite graph and a minimum-cost flow in an associated network. In particular, we show that the maximization Problem A defined in Section 2.2 is equivalent to a minimization problem, reminiscent in some aspects of the classic Monge-Kantorovich [37] mass-transfer problem, namely Problem B defined below.

Definition 8 (Problem B).

$$\begin{aligned} \inf_{\gamma} \quad & \left\{ \int_0^1 \int_0^1 \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 \right\} \\ \text{s.t.} \quad & \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 \geq m_1(x_1, y_1), \forall (x_1, y_1) \in [0, 1]^2 \\ & \int_{y_2}^1 \int_0^{x_2} \gamma(x_1, y_1, x_2, y_2) dx_1 dy_1 \geq m_2(x_2, y_2), \forall (x_2, y_2) \in [0, 1]^2 \\ & \gamma(x_1, y_1, x_2, y_2) \geq 0, \forall (x_1, y_1), (x_2, y_2) \in [0, 1]^2 \end{aligned}$$

Problem B is a continuous version of the min-cost flow dual of the maximum weighted independent set problem of the discrete case. “Continuous” here operates at many levels: The nodes of the bipartite graph (both sides thereof) are the points in the unit square; the edges of the bipartite graph are all possible directed edges going from one point in the unit square to another *in the southeast direction*. The capacities of the nodes are the values of the functions m_1 (left-hand side) and m_2 (right-hand side).

Finally, one way of understanding the min-cost flow problem is the following: Suppose that the unit square is a garden of a particular geomorphology (hills, peaks and valleys, all above sea level) captured by the function m_1 . We want to transform this landscape to the one captured by function m_2 , and we want to do this by moving one grain of earth at a time. For each grain of earth in landscape m_1 we have two options: Either (1) we remove it, or (2) we move it in the southeast direction (or we keep it where it is, this counts as moving it a zero distance). We then complete our crafting of landscape m_2 by repeatedly (3) bringing in new grains. 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.

Let us denote by Γ the set of all functions $\gamma : [0, 1]^4 \mapsto \mathfrak{R}$ satisfying the above constraints. We next show that this problem coincides, at optimality, with the optimal auction:

Theorem 5 (Duality Theorem). *For any joint density function f on $[0, 1]^2$:*

$$\begin{aligned} & \sup_{(\alpha, \beta) \in \mathcal{AB}} \left\{ \int_0^1 \int_{\alpha(y)}^1 m_1(x, y) dx dy + \int_0^1 \int_{\beta(x)}^1 m_2(x, y) dy dx \right\} \\ &= \inf_{\gamma \in \Gamma} \left\{ \int_0^1 \int_0^1 \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 \right\} \end{aligned}$$

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 m_1 and m_2 defined on $[0, 1] \times [0, 1]$ by two discrete functions m_1^d and m_2^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.

$$m_1^d(i, j) = \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} m_1(x, y) dy dx, \text{ for } i, j = 0, \dots, n-1$$

and

$$m_2^d(i, j) = \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} m_2(x, y) dy dx, \text{ for } i, j = 0, \dots, n-1.$$




Figure 2.7: The graph consisting of two grids, one for each bidder, 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 $w_u = m_1^d(i, j)$, and similarly vertex $v(i', j')$ of V has a weight $w_v = m_2^d(i', j')$. 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 2.7 for an example with $n = 3$).

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

Definition 9 (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 10 (Problem D: MINIMUM COST TRANSSHIPMENT). *Given the weighted bipartite graph $G = (U, V, E)$ above,*

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

The Inequalities. The crux of the proof is a sequence of inequalities relating the various solutions and optimal solutions of these four problems. In what follows we use $SOL(\cdot)$ to denote the cost of a feasible solution of any of the problems defined above, and $OPT(\cdot)$ to denote the cost of the optimal solution of a problem (sometimes SOL and OPT also denote 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 7. *For any two feasible solutions of A and B we have: $SOL(A) \leq SOL(B)$.*

Proof.

$$SOL(A)$$

$$\begin{aligned}
 &= \int_0^1 \int_{\alpha(y_1)}^1 m_1(x_1, y_1) dx_1 dy_1 + \int_0^1 \int_{\beta(x_2)}^1 m_2(x_2, y_2) dy_2 dx_2 \\
 &\leq \int_0^1 \int_{\alpha(y_1)}^1 \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 \\
 &\quad + \int_0^1 \int_{\beta(x_2)}^1 \int_{y_2}^1 \int_0^{x_2} \gamma(x_1, y_1, x_2, y_2) dx_1 dy_1 dy_2 dx_2
 \end{aligned}$$

where we used the inequality constraints of Problem B to upper bound the values of $m_1(x_1, y_1)$ and $m_2(x_2, y_2)$. We next notice that:

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

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:

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

□

Lemma 8. *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 bidder 1 for all valuations (x, y) such that $x_u(\lfloor \frac{x}{\epsilon} \rfloor, \lfloor \frac{y}{\epsilon} \rfloor) = 1$; we allocate to bidder 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 2.8, 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 bidder 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 bidder 1’s grid, where $i' > i$. Notice that the set of nodes on bidder 2’s grid that node $u(i, j)$ of bidder 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 bidder 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 2.1, 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 2 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 bidder 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 bidder 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:

$$\sum_{u \in U} x_u m_1^d(u) + \sum_{v \in V} x_v m_2^d(v) - \epsilon = OPT(C) - \epsilon$$

To show this we first argue that the number of nodes $v(i, j)$ for which this happens is small, and in particular that 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 bidder 1 gets the item and bidder 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 m_2 at any point (x, y) is at most 1: indeed, $m_2(x, y)$ is defined as $yf(x, y) - \int_y^1 f(x, t)dt$, wherever $\frac{\partial}{\partial y} \left[\max_{y' \geq y} y' \cdot \int_{y'}^1 f(x, t) dt \right]$ is defined, and extended to full range by right continuity. It

follows immediately that $m_2(x, y) \leq 1$ and $w_{v(i,j)} = m_2^d(i, j) \leq \epsilon^2$; therefore the total weight loss is at most $\frac{1}{\epsilon} \cdot \epsilon^2$ and the lemma follows. \square

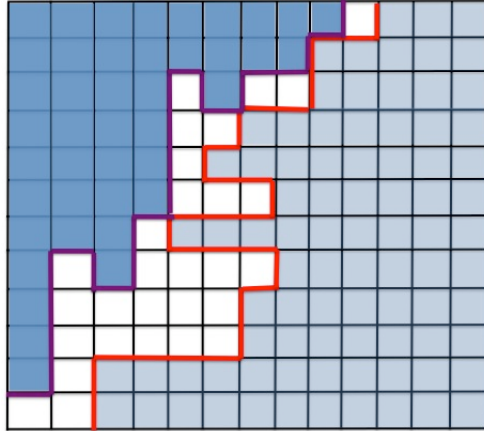


Figure 2.8: The solution of Problem A (red and purple curves) resulting from the solution of Problem C (blue and light-blue shaded regions) as described in Lemma 8.

Lemma 9. *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 optimal solution for Problem B has 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:

$$\gamma(x_1, y_1, x_2, y_2) = m_1(x_1, y_1) \cdot m_2(x_2, y_2) \cdot \frac{y_{uv}}{w_u w_v}, \quad \text{if } \left\lfloor \frac{x_1}{\epsilon} \right\rfloor < \left\lfloor \frac{x_2}{\epsilon} \right\rfloor \quad \text{and} \quad \left\lfloor \frac{y_1}{\epsilon} \right\rfloor > \left\lfloor \frac{y_2}{\epsilon} \right\rfloor$$

where $u \in U$ (resp. $v \in V$) is the grid point $u \left(\left\lfloor \frac{x_1}{\epsilon} \right\rfloor, \left\lfloor \frac{y_1}{\epsilon} \right\rfloor \right)$ (resp. $v \left(\left\lfloor \frac{x_2}{\epsilon} \right\rfloor, \left\lfloor \frac{y_2}{\epsilon} \right\rfloor \right)$) that corresponds to the little ϵ^2 -area square containing point (x_1, y_1) (resp. (x_2, y_2)). Finally, we let:

$$\gamma(x_1, y_1, x_2, y_2) = 0, \quad \text{if } \left\lfloor \frac{x_1}{\epsilon} \right\rfloor \geq \left\lfloor \frac{x_2}{\epsilon} \right\rfloor \quad \text{or} \quad \left\lfloor \frac{y_1}{\epsilon} \right\rfloor \leq \left\lfloor \frac{y_2}{\epsilon} \right\rfloor$$

We next verify that the function γ defined above satisfies the constraints of Problem B. Since the non-negativity constraint is obviously satisfied, we only need to check that γ satisfies the first and second constraints of Problem B. We only provide the proof for the first constraint and the proof for the second constraint follows along the exact same lines:

$$\int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2$$

$$\begin{aligned}
 &= \sum_{i > \lfloor \frac{x_1}{\epsilon} \rfloor, j < \lfloor \frac{y_1}{\epsilon} \rfloor} \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} m_1(x_1, y_1) \cdot m_2(x_2, y_2) \cdot \frac{y_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor), v(\lfloor \frac{x_2}{\epsilon} \rfloor, \lfloor \frac{y_2}{\epsilon} \rfloor)}{w_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor) \cdot w_v(\lfloor \frac{x_2}{\epsilon} \rfloor, \lfloor \frac{y_2}{\epsilon} \rfloor)} dy_2 dx_2 \\
 &= m_1(x_1, y_1) \cdot \sum_{i > \lfloor \frac{x_1}{\epsilon} \rfloor, j < \lfloor \frac{y_1}{\epsilon} \rfloor} \frac{y_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor), v(i, j)}{w_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor) \cdot w_v(i, j)} \cdot \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} m_2(x_2, y_2) dy_2 dx_2 \\
 &= m_1(x_1, y_1) \cdot \sum_{i > \lfloor \frac{x_1}{\epsilon} \rfloor, j < \lfloor \frac{y_1}{\epsilon} \rfloor} \frac{y_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor), v(i, j)}{w_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor) \cdot w_v(i, j)} \cdot w_v(i, j) \\
 &= m_1(x_1, y_1) \cdot \sum_{i > \lfloor \frac{x_1}{\epsilon} \rfloor, j < \lfloor \frac{y_1}{\epsilon} \rfloor} \frac{y_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor), v(i, j)}{w_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor)} \\
 &\geq m_1(x_1, y_1)
 \end{aligned}$$

where in the first equality we split the integration over discretized square regions of area ϵ^2 (the same that are used in the discrete auxiliary Problems *C* and *D*) and in the second equality we 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_2, y_2)). In the third equality we used the definition of the weight w and in the last inequality we used the fact that y is a feasible solution for Problem *D* and therefore $\sum_{v \in E} y_{uv} \geq w_u$.

We conclude our proof by showing that the cost of the feasible solution we produced is exactly $OPT(D)$:

$$\begin{aligned}
 &\int_0^1 \int_0^1 \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 \\
 &= \sum_{i, j} \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} \int_0^{y_1} \int_{x_1}^1 \gamma(x_1, y_1, x_2, y_2) dx_2 dy_2 dx_1 dy_1 \\
 &= \sum_{i, j} \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} m_1(x_1, y_1) \cdot \sum_{i' > \lfloor \frac{x_1}{\epsilon} \rfloor, j' < \lfloor \frac{y_1}{\epsilon} \rfloor} \frac{y_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor), v(i', j')}{w_u(\lfloor \frac{x_1}{\epsilon} \rfloor, \lfloor \frac{y_1}{\epsilon} \rfloor)} dx_1 dy_1 \\
 &= \sum_{i, j} \sum_{i' > i, j' < j} \frac{y_u(i, j), v(i', j')}{w_u(i, j)} \cdot \int_{\epsilon i}^{\epsilon(i+1)} \int_{\epsilon j}^{\epsilon(j+1)} m_1(x_1, y_1) dx_1 dy_1 \\
 &= \sum_{i, j} \sum_{i' > i, j' < j} \frac{y_u(i, j), v(i', j')}{w_u(i, j)} \cdot w_u(i, j) \\
 &= \sum_{i, j} \sum_{i' > i, j' < j} y_u(i, j), v(i', j') \\
 &= OPT(D)
 \end{aligned}$$

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 8 and 9 we get $OPT(A) \geq OPT(B) - \epsilon$. From Lemma 7 we get $OPT(A) \leq OPT(B)$. By having $\epsilon \rightarrow 0$ we get the result. \square

The Algorithm

The proof of the Main Theorem suggests a fully polynomial-time approximation scheme (FPTAS) for the continuous case, that is, an auction 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 f 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 f 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 $f \in [0, 1]^2$
 - 2: Compute $m_1(\cdot, \cdot), m_2(\cdot, \cdot)$
 - 3: Discretize the unit plane and construct the bipartite graph G as described in Section 2.4
 - 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 2.8
-

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

Theorem 6. *Algorithm 1 returns a truthful auction that approximates the optimal profit within ϵ additive error; moreover the algorithm runs in time polynomial in both the support size and $1/\epsilon$.*

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 (2.1) and may consequently not be a feasible solution to Problem A . This is problematic since it does not allow us to use Lemma 1 to compute the profit of the auction returned. To that end we need to establish that the violation of condition (2.1) is –in some sense– negligible; we do that next.

Suppose that curve $\alpha(y)$ violates condition (2.1) for some $y^* \in [0, 1]$ and let

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

be the minimum x for which we could create a new solution α' by setting $\alpha'(y^*) = x^*$ and have condition (2.1) restored, while not altering the profit of our auction.¹ We will argue that

$$\alpha(y^*) \cdot \int_{\alpha(y^*)}^1 f(t, y^*) dt \geq x^* \cdot \int_{x^*}^1 f(t, y^*) dt - \Theta(\epsilon). \quad (2.2)$$

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 bidder 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 $m_1(x_1, y_1) > 0$. By the definition of m_1 this immediately implies that

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

and therefore in particular that, if $x^* \geq x_1$ then

$$x_1 \cdot \int_{x_1}^1 f(t, y_1) dt \geq x^* \cdot \int_{x^*}^1 f(t, y_1) dt. \quad (2.3)$$

Hence, if $x^* \geq x_1$, then, 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:

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

where in the first and third inequalities we used the fact that $x \cdot \int_x^1 f(t, y) dt$ is Lipschitz-continuous for some constant λ and in the second inequality we used inequality (2.3).

If on the other hand $x^* \leq x_1$, then the points $(\alpha(y^*), y^*)$ and (x^*, y^*) already have l_1 -distance at most 2ϵ and it therefore follows immediately from Lipschitz continuity that:

$$\alpha(y^*) \cdot \int_{\alpha(y^*)}^1 f(t, y^*) dt \geq x^* \cdot \int_{x^*}^1 f(t, y^*) dt - 2\lambda\epsilon.$$

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

$$\alpha'(y) \cdot \int_{\alpha'(y)}^1 f(x, y) dx = \int_{\alpha'(y)}^1 m_1(x, y) dx \quad \text{and} \quad \beta'(x) \cdot \int_{\beta'(x)}^1 f(x, y) dy = \int_{\beta'(x)}^1 m_2(x, y) dy \quad (2.4)$$

¹The reader is referred to Lemma 2 for further discussion on this point.

so the profit of the algorithm is:

$$\begin{aligned}
 & \int_0^1 \left[\alpha(y) \cdot \int_{\alpha(y)}^1 f(x, y) dx \right] dy + \int_0^1 \left[\beta(x) \cdot \int_{\beta(x)}^1 f(x, y) dy \right] dx \\
 \geq & \int_0^1 \left[\alpha'(y) \cdot \int_{\alpha'(y)}^1 f(x, y) dx - \Theta(\epsilon) \right] dy + \int_0^1 \left[\beta'(x) \cdot \int_{\beta'(x)}^1 f(x, y) dy - \Theta(\epsilon) \right] dx \\
 = & \int_0^1 \int_{\alpha'(y)}^1 m_1(x, y) dx dy + \int_0^1 \int_{\beta'(x)}^1 m_2(x, y) dy dx - \Theta(\epsilon) \\
 = & OPT(C) - \Theta(\epsilon) \\
 \geq & OPT(A) - \Theta(\epsilon)
 \end{aligned}$$

where in the first inequality we used inequality (2.2), in the first equality we used (2.4) and in the last inequality we used our Main Theorem from the previous section.

In terms of running time, the discretized approximations of m_1 and m_2 are trivial (because of Lipschitz continuity, we can take $m_1^d(i, j) = m_1(i\epsilon, j\epsilon)$, and similarly for m_2). Solving the MAXIMUM WEIGHT INDEPENDENT SET problem is done exactly as in the previous section. \square

2.5 Three bidders: NP-completeness

In this section we show that, for three bidders, the problem of designing an approximately optimal (deterministic) auction becomes NP-hard, and in fact from 3SAT via approximation-preserving reductions. This implies that not only is the computational problem of designing the optimal deterministic auction intractable, but also that so is the problem of *approximating* this optimal auction within some finite relative error ϵ . Finding a reasonably high value of the ϵ for which this statement holds is an interesting open problem.

A geometric reformulation

We are interested in the complexity of the following problem:

Definition 11 (3OPTIMALAUCTIONDESIGN). *Given a joint discrete probability distribution f supported on $\mathcal{G} = [S] \times [S] \times [S]$ for some integer $S > 0$, find the ex-post IC and IR deterministic auction –that is, the 3-dimensional allocation matrix A , where $A[x, y, z] \in \{0, 1, 2, 3\}$, and i is the index of the bidder who gets the item when the valuation vector is (x, y, z) , or 0 if the auctioneer keeps the item– which maximizes revenue.*

To calculate the revenue (last phrase of the previous definition), the following notion of marginal profit contribution, appropriately modified for the three bidder discrete case, will be useful to our proof.

Definition 12. *The discrete analogues of the marginal profit contribution functions (Definition 2) for each bidder are the following:*

$$m_1(x, y, z) = \max \left\{ x \cdot \sum_{x' \geq x} f(x', y, z) - \sum_{x' > x} m_1(x', y, z), 0 \right\} \quad \text{for bidder 1}$$

$$m_2(x, y, z) = \max \left\{ y \cdot \sum_{y' \geq y} f(x, y', z) - \sum_{y' > y} m_2(x, y', z), 0 \right\} \quad \text{for bidder 2}$$

$$m_3(x, y, z) = \max \left\{ z \cdot \sum_{z' \geq z} f(x, y, z') - \sum_{z' > z} m_3(x, y, z'), 0 \right\} \quad \text{for bidder 3}$$

The revenue corresponding to allocation A is therefore

$$\sum_{A(x,y,z)=1} m_1(x, y, z) + \sum_{A(x,y,z)=2} m_2(x, y, z) + \sum_{A(x,y,z)=3} m_3(x, y, z)$$

The segments. Given a distribution $f(x, y, z)$ supported on \mathcal{G} , any point (x, y, z) of \mathcal{G} with $m_1(x, y, z) > 0$ is the apex of what we shall henceforth be calling an x -segment: a sequence of points starting at point (x, y, z) and including all points (x', y, z) with $x' \geq x$. The *weight* of this segment is defined to be $\sum_{x' \geq x} m_1(x', y, z)$, which, according to Definition 12, is equal to $x \cdot \sum_{x' \geq x} f(x', y, z)$. We define segments across the other dimensions analogously. The following problem is equivalent to the auction design problem:

Definition 13 (3SEGMENTS). *Given a joint discrete probability distribution f 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 10. *The problem 3SEGMENTS(f) is equivalent to 3OPTIMALAUCTIONDESIGN(f) (via approximation-preserving reductions).*

Proof sketch. The correspondence between solutions of the two problems is straightforward, illustrated in Figure 2.9. □




Figure 2.9: The reduction of Lemma 10, shown here for 2 bidders for ease of depiction.

The reduction

We must therefore show that 3SEGMENTS is APX-hard; we prove that by reducing from the problem TYPED3SAT, a special case of MAX3SAT:

Definition 14. *Let TYPED3SAT be the MAX3SAT problem with the input formula restricted to be of the following form. We have three types 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 $(\bar{x}_5 \vee z_7)$). Moreover, every variable appears at most 7 times in the formula.*

Lemma 11. *TYPED3SAT is APX-complete.*

Proof. We reduce from the 3SAT-5 problem, which is 3SAT with the additional constraint that every variable appears at most 5 times (this was shown in [66] to be APX-complete). In order to turn an instance ϕ of 3SAT-5 to an equivalent instance Ψ of TYPED3SAT, for every variable x we create three copies x_1, x_2, x_3 , and add the consistency clauses $(\bar{x}_1 \vee x_2), (\bar{x}_2 \vee x_3), (\bar{x}_3 \vee x_1)$ to guarantee that all copies of the variable x will have the same value. We then consider an arbitrary ordering of the variables within each clause, and replace any occurrence of variable x at position $i \in \{1, 2, 3\}$ within a clause with x_i . Since a variable x can appear at most 5 times in the original formula, each one of the copies x_1, x_2, x_3 can appear at most 7 times: 5 times in the original formula, if all occurrences of x happen to be at the same position, and 2 times in the consistency clauses.

NP-hardness is immediate: Any satisfying truth assignment of the original formula yields immediately a satisfying truth assignment of the resulting formula, and vice-versa. To show that the reduction is approximation-preserving, suppose that we can satisfy a fraction of $(1 - \epsilon)$ of Ψ . We claim that we can satisfy a fraction of $(1 - 50 \cdot \epsilon)$ of the original ϕ . To see this, notice that Ψ has $m + 3n \leq 10m$ clauses, and if an ϵ fraction of them is unsatisfied, or at most $10\epsilon m$, this can affect at most $50\epsilon m$ of the m clauses of ϕ . This is because each clause of Ψ either is itself a clause of ϕ , or it is a consistency clause for some variable, which therefore affects the at most five clauses in ϕ in which the variable appears. Therefore, the remaining clauses of ϕ are satisfied. \square

Overview of the construction

The instance of 3SEGMENTS we create has three types of segments:

- *Literal segments* capture the truth assignment to 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* capture the truth assignment to literals of a particular clause; we create one such clause segment for every literal that appears in a clause. These segments intersect with each other and with some of the literal segments corresponding to those

literals. As a result, if the clause is satisfied with the truth assignment implied by the literal segments, we will be able to pick at least one clause segment per clause (because the literal segment that intersects it will not be picked). Moreover, we cannot pick two or more clause segments per clause, since they will all intersect with each other, so we will end up picking *exactly* one such clause segment per satisfied clause.

- The *scaffolding segments* complete the construction in a way outlined later.

Suppose that the instance of TYPED3SAT contains the clause $C = (x_1 \vee y_1 \vee \bar{z}_1)$. The corresponding 3SEGMENTS gadget is presented in Figure 2.10, where literal segments are depicted in red and clause segments are depicted in blue. There are two intersecting literal segments α and $\bar{\alpha}$ for each variable α (the intersections of segments are denoted for clarity as small circles); as only one of these can be selected in the solution (and it will be clear that one *will* be selected in the optimal solution), the solution will imply a truth assignment.

Conversely, for any truth assignment we pick a set of literal segments for our 3SEGMENTS solution as follows: if the truth assignment sets x_1 to true we are going to pick the literal segment labeled x_1 in Figure 2.10, otherwise, if x_1 is false, we are going to pick the literal segment labeled \bar{x}_1 . Notice each type of variables is assigned its own axis in 3-dimensional space. Notice also how the fact that these two segments intersect ensures the consistency of the assignment. The idea now is that any truth assignment that satisfies the clause, i.e. it sets at least one of the literals x_1, y_1 or \bar{z}_1 to true, results in a solution of 3SEGMENTS that includes at least one of the literal segments labeled x_1, y_1, \bar{z}_1 . Therefore, the solution *does not* include at least one of the complementary literal segments labeled $\bar{x}_1, \bar{y}_1, z_1$; this allows us to include one of the blue clause segments to our 3SEGMENTS solution. In fact, we can include *at most one* blue segment corresponding to clause C , because the three blue segments associated with it intersect.

There are many details missing, of course. For example, recall that an instance of 3SEGMENTS is a probability distribution f , which induces the aforementioned segments as explained in the paragraph preceding Definition 13. Intuitively, this distribution should assign non-zero probability mass *only* at the points at which we want our segments to begin.

This leads to a technical complication, which necessitates the third kind of segments, the scaffolding segments: consider for simplicity the two-dimensional example of Figure 2.11, where the distribution is supported on the set $\{1, 2, 3\}^2$ and our goal is to include the blue segment depicted. We can achieve this by assigning some non-zero probability mass at point $(1, 1)$, say $f(1, 1) = 1$, which will make $m_2(1, 1) = 1 > 0$ and will indeed induce the blue segment; however, since $m_1(1, 1) = 1 > 0$ as well, the instance will also induce the undesired red segment. To resolve this issue we introduce a scaffolding segment: we split the probability mass between points $(1, 1)$ and $(3, 1)$, for example $f(1, 1) = f(3, 1) = 0.5$, so that $m_1(1, 1) = 0$ and the undesired red segment disappears. The new green segments starting at $(3, 1)$ because of $m_1(3, 1) = 1.5, m_2(3, 1) = 0.5$, are what we refer to as scaffolding segments.




Figure 2.10: The gadget of 3SEGMENTS resulting from a single clause of TYPED3SAT. This picture is not accurate –especially with respect to the location of the segments’ apices– and is solely meant to convey the basic idea of the reduction. We give a precise example in the next section.

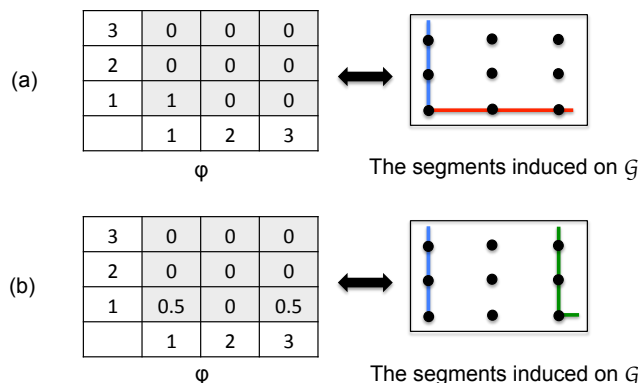


Figure 2.11: The usage of scaffolding segments.

The detailed construction

In this section we give the full details of the reduction, showing how to construct a probability distribution f that serves as the input of 3SEGMENTS, given an instance of TYPED3SAT. We conclude with a worked out example.

Let \hat{n} denote $\max\{n_x, n_y, n_z\}$. The support² of f is: $\{s(i) | i = 1, \dots, \hat{n} + 2m + 4\}^3$, for an appropriate choice of the values $s(i)$ which we will fix later; for now all we assume is that s is an increasing function of i . The size of the support is at most $(\hat{n} + 2m + 4)^3$, so this is

²The support is actually a subset of this; these are all the points (bidder valuations) with *potentially* non-zero probabilities. This will become clear in the actual construction.

clearly a polynomial time construction.

In what follows we shall abuse notation and write $f(x, y, z)$ instead of $f(s(x), s(y), s(z))$ when there is no ambiguity. Also, in terms of phrasing, we shall sometimes refer to the sub-matrices $f(i, \cdot, \cdot)$, $f(\cdot, j, \cdot)$, $f(\cdot, \cdot, k)$ as the “planes” $x = i$, $y = j$ and $z = k$ respectively; finally, remember that a segment that runs parallel to the x (resp. y or z) axis is an x -segment (resp. y -segment or z -segment).

The construction proceeds as follows: Consider an arbitrary ordering of the clauses 1 through m and 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) = $1, \dots, n_x$ (resp. n_y, n_z) (the same construction also works for clauses with less than 3 variables). For this clause we are going to introduce the set of literal segments and clause segments described below.

Literal segments. To formally define the literal segments we first need to set the probability mass of the point that is the apex of each segment, and then use scaffolding segments to ensure that there is exactly one segment starting at each such point, towards the appropriate direction. We defer the discussion of scaffolding segments to a subsequent section, and here we only show which points are picked as apices.

In particular, the probability mass that is assigned to each such apex is c_1 (to be determined later); for notational convenience, and in order to give an idea of what kind of segments we are expecting to get per type of literal, in what follows we write x -segment(i, j, k) to denote that $f(i, j, k) = c_1$ and that we are later going to use scaffolding segments to ensure the existence of an x -segment *only* (and analogously for y and z -segments).

- $pos(x_i) \Rightarrow y$ -segment($i + 1, \hat{n} + 2, \hat{n} + 2 + l$) and z -segment($i + 1, \hat{n} + 2 + m + l, \hat{n} + 2$)
 $neg(x_i) \Rightarrow z$ -segment($i + 1, \hat{n} + 2 + l, \hat{n} + 2$) and y -segment($i + 1, \hat{n} + 2, \hat{n} + 2 + m + l$)
- $pos(y_j) \Rightarrow z$ -segment($\hat{n} + 2 + l, j + 1, \hat{n} + 2$) and x -segment($\hat{n} + 2, j + 1, \hat{n} + 2 + m + l$)
 $neg(y_j) \Rightarrow x$ -segment($\hat{n} + 2, j + 1, \hat{n} + 2 + l$) and z -segment($\hat{n} + 2 + m + l, j + 1, \hat{n} + 2$)
- $pos(z_k) \Rightarrow x$ -segment($\hat{n} + 2, \hat{n} + 2 + l, k + 1$) and y -segment($\hat{n} + 2 + m + l, \hat{n} + 2, k + 1$)
 $neg(z_k) \Rightarrow y$ -segment($\hat{n} + 2 + l, \hat{n} + 2, k + 1$) and x -segment($\hat{n} + 2, \hat{n} + 2 + m + l, k + 1$)

Note that every positive occurrence of a variable of type, say, x results in the following two segments: a *negative* literal segment starting at $(i + 1, \hat{n} + 2, \hat{n} + 2 + l)$ that intersects with the corresponding clause segment starting at $(1, \hat{n} + 2 + l, \hat{n} + 2 + l)$ (see next paragraph), 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-positive* literal segment. Negative occurrences of variables analogously result in *positive* and *dummy-negative* literal segments; Figure 2.10 provided an illustration, minus the dummy segments. Dummy literal segments are introduced because (for reasons that will become apparent in the proof Lemma 12) we want to ensure that we have an equal number of positive and negative literal segments (when dummies are included).

The point masses described above are depicted in Figure 2.12 for variables of type z . (The reader should ignore for now the c_4 entries of Figure 2.12 which correspond to scaffolding segments.)

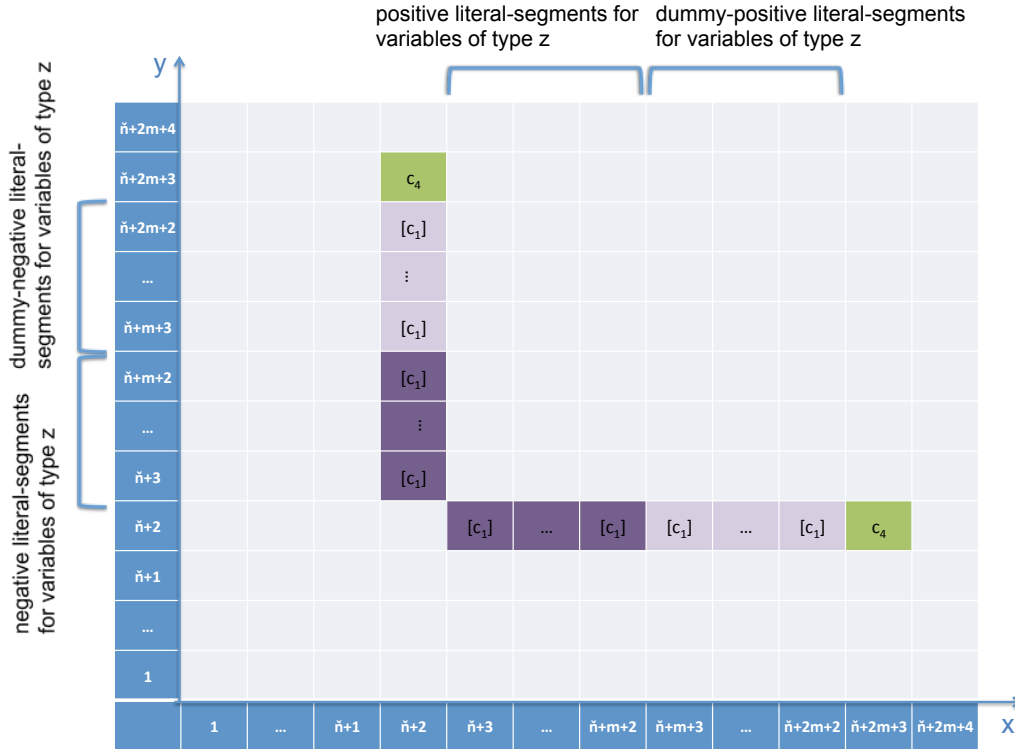


Figure 2.12: The plane $z = k + 1$, $k = 1, \dots, n_z$, contains the literal segments of variable z_k . By $[\cdot]$ we mean that this point does not appear at all levels z . Every such level must have an equal number of c_1 -entries in line $\hat{n} + 2$ and column $\hat{n} + 2$; also, since every variable appears at most 7 times, there can be at most 7 entries in line and column $\hat{n} + 2$ for each such level.

Clause segments. The point masses below (with the appropriate usage of scaffolding segments to be specified later) will give rise to an x -segment, a y -segment and a z -segment, which are the clause segments of the variable of type x , y and z respectively, for clause l .

$$f(1, \hat{n} + 2 + l, \hat{n} + 2 + l) = f(\hat{n} + 2 + l, 1, \hat{n} + 2 + l) = f(\hat{n} + 2 + l, \hat{n} + 2 + l, 1) = c_2$$

The point masses for clause segments of variables of type z are depicted in Figure 2.13. (The reader should ignore for now the c_3 entries of Figure 2.13 which correspond to scaffolding segments.)




Figure 2.13: The plane $z = 1$. 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 at $(l, l, 1)$, so that it intersects with the non-dummy literal segment. The only reason why a c_2 point might be missing is because of a clause of the form $(x_i \vee y_j)$.


An example. To better understand the translation of a TYPED3SAT formula into a 3SEGMENTS instance, we consider the following example:

$$(x_1 \vee y_1 \vee \bar{z}_1) \wedge (\bar{x}_1 \vee \bar{y}_1) \wedge (x_2 \vee y_1 \vee z_1)$$


According to the construction above, the probability masses of the regular and dummy literal segments (c_1 probability mass) and of the clause segments (c_2 probability mass) are assigned to the coordinates displayed in the following table.

	$x_1 \vee y_1 \vee \bar{z}_1$			$\bar{x}_1 \vee \bar{y}_1$		$x_2 \vee y_1 \vee z_1$		
	x_1	y_1	\bar{z}_1	\bar{x}_1	\bar{y}_1	x_2	y_1	z_1
regular	(2,4,5)	(5,2,4)	(5,4,2)	(2,6,4)	(4,2,6)	(3,4,7)	(7,2,4)	(4,7,2)
dummy	(2,8,4)	(4,2,8)	(4,8,2)	(2,4,9)	(9,2,4)	(3,10,4)	(4,2,10)	(10,4,2)
clause	(1,5,5),(5,1,5),(5,5,1)			(1,6,6),(6,1,6),(6,6,1)		(1,7,7),(7,1,7),(7,7,1)		


The resulting set of segments per clause (after the application of appropriate scaffolding segments which will be explained in the next paragraph) is depicted in Figure 2.14. Figure 2.15 displays the full set of segments (minus the scaffolding segments) for this formula.



(a) The gadget for clause $x_1 \vee y_1 \vee \bar{z}_1$.



(b) The gadget for clause $\bar{x}_1 \vee \bar{y}_1$.



(c) The gadget for clause $x_2 \vee y_1 \vee z_1$.

Figure 2.14: The gadgets resulting from $(x_1 \vee y_1 \vee \bar{z}_1) \wedge (\bar{x}_1 \vee \bar{y}_1) \wedge (x_2 \vee y_1 \vee z_1)$.




Figure 2.15: The full set of segments for $(x_1 \vee y_1 \vee \bar{z}_1) \wedge (\bar{x}_1 \vee \bar{y}_1) \wedge (x_2 \vee y_1 \vee z_1)$.

Scaffolding segments. 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 this, we need the aforementioned scaffolding segments; these will ensure that the only segments starting from the c_1 and c_2 points specified in the previous paragraphs will indeed be the desired literal and clause segments mentioned above.

c_3 : 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: $f(1, \hat{n} + l + 2, \hat{n} + 2m + 3) = f(1, \hat{n} + 2m + 3, \hat{n} + l + 2) = f(\hat{n} + l + 2, 1, \hat{n} + 2m + 3) = f(\hat{n} + 2m + 3, 1, \hat{n} + l + 2) = f(\hat{n} + l + 2, \hat{n} + 2m + 3, 1) = f(\hat{n} + 2m + 3, \hat{n} + l + 2, 1) = c_3$, for all $l = 1, \dots, m$; see Figure 2.13 for an illustration. In order to force the marginal profit contribution function to be zero at a c_2 point for some appropriately chosen bidder/direction, we impose the requirement that:

$$s(\hat{n} + l + 2) \cdot (c_2 + c_3) \leq s(\hat{n} + 2m + 3) \cdot c_3, \quad \forall l = 1, \dots, m \quad (2.5)$$

For example, suppose that we want point $(1, \hat{n} + 2 + l, \hat{n} + 2 + l)$ to be the apex of an x -segment only, and that no y or z -segment should start at this point. We achieve this by including the points $f(1, \hat{n} + l + 2, \hat{n} + 2m + 3) = f(1, \hat{n} + 2m + 3, \hat{n} + l + 2) = c_3$; the above inequality then ensures that

$$m_2(1, \hat{n} + 2 + l, \hat{n} + 2 + l) = h(1, \hat{n} + 2 + l, \hat{n} + 2 + l) = 0$$

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

c_4 : We next ensure that there are no segments starting from c_1 points that go along row or column $\hat{n} + 2$, by introducing: $f(i + 1, \hat{n} + 2, \hat{n} + 2m + 3) = f(i + 1, \hat{n} + 2m + 3, \hat{n} + 2) = f(\hat{n} + 2, j + 1, \hat{n} + 2m + 3) = f(\hat{n} + 2m + 3, j + 1, \hat{n} + 2) = f(\hat{n} + 2, \hat{n} + 2m + 3, k + 1) = f(\hat{n} + 2m + 3, \hat{n} + 2, k + 1) = c_4$, for $i = 1, \dots, n_x, j = 1, \dots, n_y, k = 1, \dots, n_z$; see Figure 2.12 for an illustration. Since we have at most 7 occurrences of any variable, and hence at most 7 c_1 -entries on any row or column $\hat{n} + 2$ at a given level, the requirement we impose (along the same lines as above) is that:

$$s(\hat{n} + l + 2) \cdot (7c_1 + c_4) \leq s(\hat{n} + 2m + 3) \cdot c_4, \quad \forall l = 1, \dots, 2m \quad (2.6)$$

c_5 : Finally, we ensure that there are no x -segments (resp. y -segments and z -segments) starting from a point c_1 on plane $x = i + 1$ (resp. $y = i + 1, z = i + 1$) for $i = 1, \dots, \hat{n}$, by introducing: $f(\hat{n} + 2m + 4, l, \hat{n} + 2) = f(\hat{n} + 2m + 4, \hat{n} + 2, l) = c_5$ (resp. $f(\hat{n} + 2, \hat{n} + 2m + 4, l) = f(l, \hat{n} + 2m + 4, \hat{n} + 2) = c_5$ and $f(l, \hat{n} + 2, \hat{n} + 2m + 4) = f(\hat{n} + 2, l, \hat{n} + 2m + 4) = c_5$) for all $l = \hat{n} + 3, \dots, \hat{n} + 2m + 2$; see Figure 2.16 for an illustration. Noticing that c_1 can appear at position $(i + 1, l, \hat{n} + 2)$ only for one $i \in \{1, 2, \dots, \hat{n}\}$ (and analogously for the other positions of c_5 points) it suffices to impose the requirement that:

$$s(i + 1) \cdot (c_1 + c_5) \leq s(\hat{n} + 2m + 4) \cdot c_5, \quad \forall i = 1, \dots, \hat{n} \quad (2.7)$$

There is one final issue to address: from each of the scaffolding points defined above, there is not one, but three segments starting from it. Since the scaffolding segments do not intersect with any of the literal and clause segments but only with each other, for the purposes of our reduction it suffices to consider the most profitable set of scaffolding segments we can include in our solution. This way, when we prove our reduction's guarantees, we can make sure to tune the target profit in a way that will ensure the inclusion in the optimal solution of this most profitable set, so that we won't have to worry about the other scaffolding segments. In the remaining of this paragraph we will point out the most profitable set of segments in each of the three cases above. To facilitate the exposition we will move across planes $z = 1, \dots, \hat{n} + 2m + 4$, so that the reader can consult Figures 2.12, 2.13 and 2.16 along the way.

Starting with plane $z = 1$ and the c_3 scaffolding points, we will focus wlog on the points $(\hat{n} + 2 + l, \hat{n} + 2m + 3, 1), l = 1, \dots, m$, the apices of an x -segment of weight $s(\hat{n} + 2 + l) \cdot (m - l + 1)c_3$, a y -segment of weight $s(\hat{n} + 2m + 3) \cdot c_3$ and a z -segment of weight $s(1) \cdot c_3$. The y -segment of every point clearly dominates the z -segment of the same point, while the whole set of all m y -segments from all such points (with total weight $m \cdot s(\hat{n} + 2m + 3) \cdot c_3$) is better than any combination of l y -segments followed by a (long) x -segment (with total weight $l \cdot s(\hat{n} + 2m + 3) \cdot c_3 + s(\hat{n} + 2 + l) \cdot (m - l)c_3$). Therefore, the most profitable set of segments consists of m y -segments, each with a weight of $s(\hat{n} + 2m + 3) \cdot c_3$.

Moving on to planes $z = k + 1, k = 1 \dots n_z$ and the c_4 scaffolding points, we will focus wlog on the points $(\hat{n} + 2, \hat{n} + 2m + 3, k + 1)$, the apices of an x -segment of weight $s(\hat{n} + 2) \cdot c_4$, a y -segment of weight $s(\hat{n} + 2m + 3) \cdot c_4$ and a z -segment with weight $s(k + 1) \cdot (n_z - k + 1)c_4$. The y -segment of every point clearly dominates the x -segment of the same point, while the whole set of all n_z y -segments from all such points (with total weight $n_z \cdot s(\hat{n} + 2m + 3) \cdot c_4$) is better than any combination of k y -segments followed by a (long) z -segment (with total weight $k \cdot s(\hat{n} + 2m + 3) \cdot c_4 + s(k + 1) \cdot (n_z - k)c_4$). Therefore, the most profitable set of segments consists of n_z y -segments, each with a weight of $s(\hat{n} + 2m + 3) \cdot c_4$.

Finally, for plane $z = \hat{n} + 2m + 4$ and the c_5 scaffolding points, we will focus wlog on the points $(\hat{n} + 2, \hat{n} + 2 + l, \hat{n} + 2m + 4), l = 1, \dots, 2m$, the apices of an x -segment of weight $s(\hat{n} + 2) \cdot c_5$, a y -segment of weight $s(\hat{n} + 2 + l) \cdot (2m - l + 1)c_5$ and a z -segment with weight $s(\hat{n} + 2m + 4) \cdot c_5$. The z -segment of every point clearly dominates the x -segment of the same point, while the whole set of all $2m$ z -segments from all such points (with total weight $2m \cdot s(\hat{n} + 2m + 4) \cdot c_5$) is better than any combination of l z -segments followed by a (long) y -segment (with total weight $l \cdot s(\hat{n} + 2m + 4) \cdot c_5 + s(\hat{n} + 2 + l) \cdot (2m - l)c_5$). Therefore, the most profitable set of segments consists of $2m$ z -segments, each with a weight of $s(\hat{n} + 2m + 4) \cdot c_5$.

The support. We are now ready to set the values of the support $s(\cdot)$ and of the variables c_1, \dots, c_5 . Those values will need to satisfy the aforementioned constraints (2.5,2.6,2.7), and since $s(\cdot)$ is increasing, it actually suffices to satisfy the following stronger constraints:

$$(2.5) \quad \Leftarrow \quad s(\hat{n} + 2m + 2) \cdot (c_2 + c_3) \leq s(\hat{n} + 2m + 3) \cdot c_3 \quad (2.8)$$

$$(2.6) \quad \Leftarrow \quad s(\hat{n} + 2m + 2) \cdot (7c_1 + c_4) \leq s(\hat{n} + 2m + 3) \cdot c_4 \quad (2.9)$$


 Figure 2.16: The plane $z = \hat{n} + 2m + 4$.

$$(2.7) \Leftrightarrow s(\hat{n} + 2m + 2) \cdot (c_1 + c_5) \leq s(\hat{n} + 2m + 4) \cdot c_5 \quad (2.10)$$

Moreover, the values chosen need to satisfy the following inequality, which we will use later in the proof of soundness of our reduction:

$$m \cdot s(1)c_2 \leq \bar{n} \cdot s(\hat{n} + 2)c_1 \quad (2.11)$$

where \bar{n} denotes the total number of literal occurrences in the formula.

It is easy to check that if we pick the values for $s(\cdot)$ as follows:

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

then constraints (2.8,2.9,2.10,2.11) above are all satisfied with equality, as long as the variables c_1, \dots, c_5 satisfy the following system (where we can choose the value c_1 so that all of the entries across $f(\cdot, \cdot, \cdot)$ sum up to 1 and we get a proper probability distribution):

$$\begin{aligned} c_2 &= \frac{\bar{n}}{m} \left(1 + \frac{\hat{n} + 1}{\hat{n} + 2m + 1} \right) c_1 \\ c_3 &= c_2 \\ c_4 &= 7c_1 \end{aligned}$$

$$c_5 = \frac{2}{3}c_1$$

We are now ready to provide completeness and soundness guarantees for our reduction.

Lemma 12. *The construction described above is an approximation-preserving reduction from TYPED3SAT to 3SEGMENTS.*

Proof. We first note that regardless of the instance of TYPED3SAT we are reducing from, we can always obtain a fixed profit for 3SEGMENTS from the scaffolding segments; by picking the most profitable set of segments, as discussed above, we get a total profit of:

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

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

- If the TYPED3SAT formula is satisfiable then the optimal profit of 3SEGMENTS is exactly:

$$m \cdot s(1)c_2 + \bar{n} \cdot s(\hat{n} + 2)c_1 + F \tag{2.12}$$

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**³ we include its **positive** literal segments (both regular and dummies). Notice that –thanks to the dummy literal segments– there is an equal number of positive and negative literal segments overall (dummies included), with a total of $2\bar{n}$ 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 $\bar{n} \cdot s(\hat{n} + 2)c_1$.

Moreover, since the formula is satisfiable, at least one literal per clause is satisfied. If this is a positive³ literal, then the variable has been set to true and the literal segments for this variable included in our 3SEGMENTS solution will be the positive ones. Furthermore, since the variable appears as a positive literal at this clause, our construction ensures that the corresponding clause segment intersects only with the negative literal segment of this variable. Hence, since we have already included the positive literal segment in our solution, the negative literal segment is not included, and we are therefore free to include the clause segment in our solution as well. Each one of these clause segments contributes $s(1)c_2$; noticing that we cannot include more than one clause segment from each clause (because they intersect) we get that their total contribution is exactly $m \cdot s(1)c_2$.

- If the optimal assignment for TYPED3SAT satisfies at most a $(1 - \epsilon)$ fraction of the clauses, then the optimal profit for 3SEGMENTS is at most:

$$(1 - \epsilon)m \cdot s(1)c_2 + \bar{n} \cdot s(\hat{n} + 2)c_1 + F \tag{2.13}$$

³The other case is completely symmetrical.

Notice that, since the quantities F , n , \bar{n} and \hat{n} are all linear in m , this claim immediately establishes that the reduction is approximation-preserving.

To prove this we show how to transform a solution of profit $\geq (1 - \epsilon)m \cdot s(1)c_2 + \bar{n} \cdot s(\hat{n} + 2)c_1 + F$ to a truth assignment that satisfies more than a $(1 - \epsilon)$ fraction of the clauses. First notice that, wlog, any solution of 3SEGMENTS that achieves this profit will always include the most profitable set of scaffolding segments discussed above: these segments do not intersect with any literal or clause segments, so it makes no sense not to include them. Clause and literal segments do intersect with each other, however in order to achieve the additional profit of $\bar{n} \cdot s(\hat{n} + 2)c_1$, we need to include exactly \bar{n} literal segments: Obviously we cannot include any more literal segments without having intersections between them. More importantly though, we cannot substitute any literal segments for clause segments. Indeed, even if we included one clause segment per clause (remember that all clause segments of a given clause intersect with each other), and skipped some of the literal segments, we would only get a profit of $m \cdot s(1)c_2$ which, as (2.11) guarantees, is less than the required profit $\bar{n} \cdot s(\hat{n} + 2)c_1$. Finally, we include exactly one clause segment per clause for a certain fraction of the clauses: we cannot include more than one clause segment per clause without having intersecting clause segments.

The \bar{n} literal segments that are included in our solution correspond to a truth assignment to the variables of the formula as described above: we set every variable whose positive (resp. negative) literal segments are included to true (resp. false). The claim follows by noticing that this truth assignment satisfies every clause for which a clause segment was included in the 3SEGMENTS solution. If the fraction of the clauses for which we included a clause segment is more than $(1 - \epsilon)$, this means that this particular truth assignment satisfies more than a $(1 - \epsilon)$ fraction of the clauses of the original formula, a contradiction.

□

Combining Lemmas 10 and 12 immediately yields our main theorem for this section:

Theorem 7. 3OPTIMALAUCTIONDESIGN is APX-complete.

2.6 Discussion and open problems

Even though in this chapter we focused on deterministic auctions, our geometric characterization has interesting consequences for randomized auctions as well. Remember that for the discrete case the optimal (deterministic) auction immediately follows from solving the integer program of Problem *C* in Section 2.4, i.e. computing a maximum weight independent set in the corresponding graph. Our first observation is that the linear programming relaxation of this integer program corresponds to computing the optimal randomized auction. For two bidders, where the graph is bipartite and the integer program is totally unimodular, the

optimal integer solution is also the optimal of the relaxed linear program. Therefore, for two bidders, the program of Problem C computes a deterministic auction that is optimal among all randomized auctions: this is reminiscent of Myerson's original result, where the deterministic auction obtained is optimal for the (larger) class of *Bayesian truthful* randomized auctions [57]. For a constant number of three or more bidders, the appropriate generalization of our geometric characterization (which was not discussed at length in this chapter) gives polynomial-time algorithms for computing the optimal randomized auction. This is in sharp contrast with the intractability of computing the optimal deterministic auction, already from three bidders. Of course, for a large number of bidders the size of this linear program becomes 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 [32].

The main insight in deriving the auction for two bidders is the connection of the discrete case with the weighted independent set problem in a bipartite graph. To derive the solution to the continuous case, we resorted to a long-winded proof of a duality theorem. We believe the theorem is quite interesting in its own right, but the question remains: is there a direct argument, through a simple quantization of the distribution? There is a technical problem which has thwarted our attempts at such a direct proof, relating to the fact that the curves in Figure 2.1 may not be Lipschitz continuous. We do not know whether this obstacle is real (that is, whether there are Lipschitz continuous distributions that yield optimal auctions that are not Lipschitz continuous), or whether such a more direct proof is ultimately possible.

Our APX-completeness proof establishes that there is some constant lower bound on the approximability of the optimal auction design problem. An important open problem of this work is to close the gap between the best approximation algorithm known for the optimal auction problem (which was improved from the 0.5 of [67] to 0.6 in [32]) and this tiny constant. We believe that progress there is attainable, by more sophisticated reductions. Indeed, two different papers have recently brought the two bounds closer to each other: In [17] the authors bring the best known approximation ratio up to 0.622 through a tighter analysis of the k -lookahead auction of [67], while [12] provides the first explicit inapproximability threshold for the three bidder optimal auction design problem, which is around 0.983.

Our work also implies an approximation of $2/n$ for n bidders: 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 bidders. 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 [67], but for $n = 4$ already, the approximation ratio drops below $1/2$.

Chapter 3

Efficiency-Revenue Trade-offs in Auctions

3.1 Introduction

The objectives of 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 chapter. See Figure 3.1 (a) for a graphical illustration.

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 3.1 (b)). A second common-sense approach is the so-called *slope search*: To explore the trade-off space, just optimize the objective “revenue + λ · 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 [51]. The problem is that the trade-off curve may not be convex (Figure 3.1 (c)), and hence the algorithm of [51] can miss vast areas of trade-offs:

Proposition 2. *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.*

Proof. We start with a simple example with 2 bidders (presented in Figure 3.2), for which

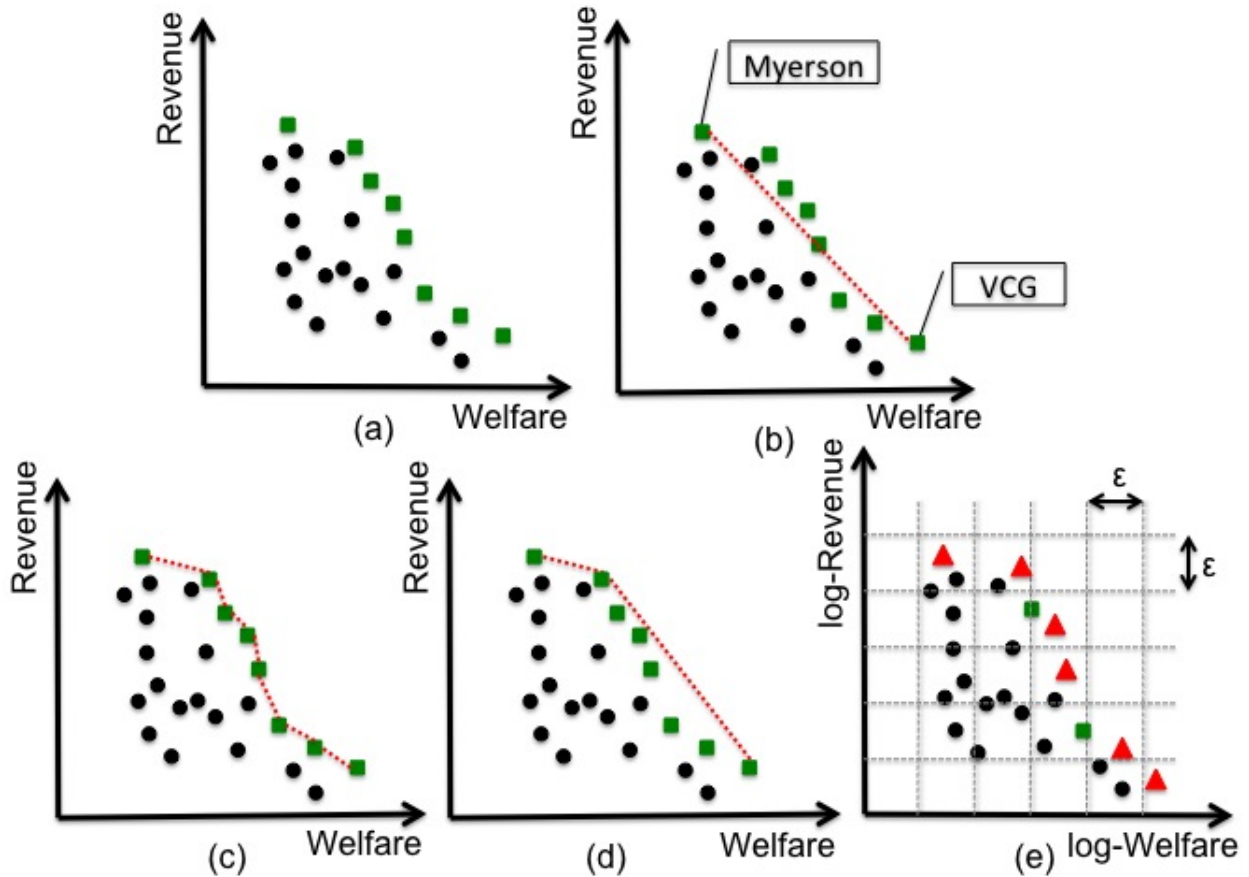


Figure 3.1: The Pareto points of the bi-criterion auction problem are shown as squares (a); the Pareto points may be far off the line connecting the two extremes (b), and may be non-convex (c). The Pareto points of randomized auctions comprise the upper boundary of the convex closure of the Pareto points (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 Pareto curve is not convex, while the bidders' valuations are drawn independently from two *non-identical* distributions of support 2; since any binary-valued distribution satisfies the monotone hazard rate condition (for the discrete case) the first part of the claim follows.

On the positive side we can show that the Pareto curve is convex for a single bidder with valuation drawn from a monotone hazard rate distribution. Since in the discrete case one can simply enumerate the set of all feasible auctions in linear time anyway, this result is of interest only in the continuous case.

Let $M(r)$ be the single-bidder auction (pricing) that makes a take-it-or-leave-it offer of




Figure 3.2: The objective space of all two-bidders auctions when $(v_1^1, v_1^2) = (11, 20)$, $(v_2^1, v_2^2) = (2, 5)$ and $(f_i^1, f_i^2) = (1/3, 2/3)$ for $i = 1, 2$.

r to the bidder; notice that in this case the Pareto curve is a mono-parametric curve on the plane where $x = \text{SW}(M(r))$ and $y = \text{Rev}(M(r))$. We next show that this monoparametric curve is in fact convex; the following is a necessary and sufficient condition for (local) convexity of mono-parametric (continuous) curves [52]:

$$\begin{vmatrix} x'(r) & x''(r) \\ y'(r) & y''(r) \end{vmatrix} = x'(r)y''(r) - y'(r)x''(r) \geq 0$$

Substituting $x = \text{SW}(M(r)) = \int_r^1 v f(v) dv$ and $y = \text{Rev}(M(r)) = r \int_r^1 f(v) dv$ and doing the algebra we get the following necessary and sufficient condition:

$$r(f(r))^2 + \int_r^1 f(v) dv \cdot [f(r) + r f'(r)] \geq 0 \tag{3.1}$$

By definition, for monotone hazard rate distributions the ratio $\frac{1-F(r)}{f(r)}$ is a non-increasing function of r ; taking derivatives¹ we get that for these distributions it must hold that:

$$(1 - F(r))f'(r) + (f(r))^2 \geq 0 \tag{3.2}$$

By substituting $f'(r)$ from (3.2) into the LHS of (3.1) we get that it is indeed ≥ 0 . \square

An interesting open question is whether the property of convex Pareto curves extends to 2 or more bidders with valuations distributed *identically* according to some monotone hazard rate distribution.

¹We make the analytically convenient assumption that f is differentiable here.

It follows that the slope search approach of [51] is incorrect. However, the correctness of the slope search approach is restored if one is willing to settle for randomized auctions: The trade-off space of randomized auctions is always convex (in particular, it is the convex hull of the deterministic auctions (Figure 3.1 (d)). It is easy to see (and it had been actually worked out for different purposes already in [56]) that the optimum randomized auction with respect to the metric “revenue + $\lambda \cdot$ welfare” is then easy to calculate:

Proposition 3. *The optimum randomized auction for the objective “revenue + $\lambda \cdot$ welfare” can be computed in polynomial time. Hence, any point of the revenue/welfare trade-off for randomized auctions can be computed in polynomial time.*

Our results

In this chapter we consider the problem of exploring the revenue/welfare trade-off for *deterministic* auctions, and show that it is an intractable problem in general, even for two bidders (Theorem 9). Comparing with the tractability of the corresponding randomized problem (as pointed out in the previous section), this result adds to the recent surge in literature pointing out complexity gaps between randomized and deterministic auctions [63, 32, 35, 31]. Randomized auctions are of course a powerful and useful analytical concept, but it is deterministic 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 auctions has exponential size.

On the positive side, we show that the problem can be solved for two bidders, even for correlated valuations (Theorem 12). 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 [65]) that an approximate summary of the trade-off of polynomial size (the ϵ -Pareto curve – see Figure 1(e)) can be computed in polynomial time. The derivation of the two-bidders auction (see Section 3.4) is quite involved. We first find a pseudo-polynomial dynamic programming algorithm for the problem of finding an auction with welfare (resp. revenue) *exactly* a given number. This algorithm is very different from the one in [63] 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 [65] (Section 3.2) 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 we discussed in the previous chapter, 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 [4]. We show that the trade-off problem is in PSPACE and (weakly) NP-hard (Theorem 13).

Related work

Efficiency-revenue trade-offs. Although [51] 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 [7] who show that the revenue benefits of adding one extra bidder and running the efficiency-maximizing auction surpass 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 [69] for keyword auctions (multi-item auctions), and in [58] 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 auctions, which depend on the number of bidders and the size of the support.

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 [47, 36, 54]. 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 [65] 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 [74, 28, 29, 41, 22, 15, 27], and much methodology has been developed, some of which has been applied to

mechanism design before [42]. In this chapter we use this methodology for studying Bayesian auctions under the two criteria of expected revenue and social welfare.

3.2 Preliminaries

The Bi-Criterion Auction problem. We want to design deterministic auctions that perform favorably with respect to (expected) social welfare, defined as $\text{SW} = \mathbb{E}[\sum_i x_i v_i]$ and (expected) revenue, defined as $\text{Rev} = \mathbb{E}[\sum_i p_i]$. Based on the characterization with allocation matrices (see Section 1.3), 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 (auctions) 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 points in A . Given two auctions 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 [65], 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 an auction 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 auction 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 [65].

Remark 1. *Even though our exposition focuses on discrete distributions, our results easily extend to continuous distributions as well. As in [63], given a sufficiently smooth continuous density (say Lipschitz-continuous), whose support lies in a finite interval $[\underline{v}, \bar{v}]$,² we can appropriately discretize (while preserving the optimal values within $O(\epsilon)$) and run our algorithms on the discrete approximations.*

²This is the standard approach in economics, see for example [57].

From exact to bi-criterion. We will make essential use of a result from [65] which reduces 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 implies an FPTAS for the multi-objective version:

Theorem 8 ([65]). *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 12), we design a pseudo-polynomial algorithm for the exact version of the BI-CRITERION AUCTION problem and apply Theorem 8 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 negative coefficients. We show however (Lemma 13) that there exists an alternate representation with monotonic linear functions.

3.3 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 bidders³ 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 9. *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 .*

Before proving Theorem 9, we will first provide a 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. In particular we will prove the following theorem.

Theorem 10. *For two bidders with independent regular distributions, it is NP-hard to decide whether there exists an auction with welfare exactly equal to W .*

Proof. 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

³Note that for a single bidder one can enumerate all feasible auctions in linear time.

$b'_i := b_i/(10k \cdot T)$, 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 relatively straightforward to verify that both distributions are indeed regular; we do that in the next claim:

Claim 1: The distributions defined above are regular.

Proof of Claim 1. We define the virtual valuation of a bidder with valuation v_i , taking values from $\{v_i^1, \dots, v_i^k\}$ with probabilities $\{f_i^1, \dots, f_i^k\}$, as follows:

$$\phi_i^j = v_i^j - (v_i^{j+1} - v_i^j) \frac{f_i^{j+1} + \dots + f_i^k}{f_i^j}$$

Substituting for the setting in hand, we get that the virtual valuation of the row bidder is $\phi_r^j = 2j - k$, while the virtual valuation of the column bidder is:

$$\phi_c^j = j + b'_j - (b'_{j+1} - b'_j + 1)(k - j)$$

A distribution is called regular iff $\phi_i^j \leq \phi_i^{j+1}$; it follows immediately that the row bidder’s distribution is regular, while for the column bidder we need that

$$j + 1 + b'_{j+1} - (b'_{j+2} - b'_{j+1} + 1)(k - j - 1) \geq j + b'_j - (b'_{j+1} - b'_j + 1)(k - j)$$

and since $b'_{j+1} \geq b'_{j+2}$ it suffices that

$$j + 1 + b'_{j+1} - (k - j - 1) \geq j + b'_j - (b'_{j+1} - b'_j + 1)(k - j)$$

Rearranging terms and doing some calculations we get that it suffices to have $b'_{j+1} \geq b'_j - 2/k$, which follows from the fact that $\sum_{i=1}^k b'_i = 1/(10k)$. \square

The main idea of the proof is this: appropriately *isolate* a subset of 2^k feasible auctions whose welfare values encode the sum of values $\sum_{i \in S} b'_i$ for all possible subsets $S \subseteq [k]$. The existence of an auction 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 auction M^* with (expected) welfare

$$\text{SW}(M^*) = (2/3) \cdot (k - 1)k(k + 1) + (1/2) \cdot k(k + 1) + \sum_{i=2}^k (i - 1)b'_i + 1/(20k) \quad (3.3)$$

Consider the allocation matrix of a feasible auction. Recall that an auction is feasible iff its allocation matrix satisfies the monotonicity constraint. The main claim is that *all auctions that could potentially satisfy (3.3) 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 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 auctions. We prove the claim above next, showing that auctions in \mathcal{R} maximize welfare:

Claim 2: We have $\max_{M \notin \mathcal{R}} \text{SW}(M) < \min_{M \in \mathcal{R}} \text{SW}(M) < \text{SW}(M^*)$.

Proof of Claim 2. Recall that this lemma implies that only auctions in \mathcal{R} can potentially satisfy (3.3). To prove it we proceed as follows: Consider a partition of the allocation matrix A into three subsets: (i) the subset of the matrix above the secondary diagonal, (ii) the subset below the diagonal and (iii) the diagonal itself. The gist of the proof is this: The contribution to the welfare for subsets (i) and (ii) is maximized for auctions in \mathcal{R} . The welfare contribution from (i) and (ii) for *any* other auction (i.e. not in \mathcal{R}) is strictly smaller by a quantity sufficiently large that outweighs any effects on the welfare from subset (iii).

Let us first compute $\min_{M \in \mathcal{R}} \text{SW}(M)$, the minimum welfare of an auction in \mathcal{R} . It is easy to see that the welfare minimizing auction is the one that assigns the item to the row bidder for all entries in the diagonal. Hence, we have:

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

So, we obtain the second inequality of the lemma. To bound from above $\max_{M \notin \mathcal{R}} \text{SW}(M)$ we consider three cases: Consider first the subset of the allocation matrix above the diagonal. If any entry of this subset is allocated to the column bidder, then it is not hard to see that this would lower the welfare value by at least $1 - \max_i b'_i \geq 1 - 1/(10k) \geq 0.9$. Similarly, if any entry is not allocated at all (i.e. the auctioneer keeps the item), this would cost us at least 1. For the subset below the diagonal the situation is analogous; if an entry is allocated to the row bidder, this costs us at least 1, same if an entry is not allocated at all. This decrease in the value of the welfare cannot be compensated by the diagonal entries; indeed, if all such entries are allocated to either bidder, contribution to the welfare lies in $[k(k+1)/2, k(k+1)/2 + 1/(10k)]$ (an interval of length $1/(10k) \leq 1/10$). As a consequence, any auction that disagrees with \mathcal{R} either below or above the diagonal has welfare strictly smaller than $\min_{M \in \mathcal{R}} \text{SW}(M)$. Now consider an auction that agrees with \mathcal{R} except potentially at the diagonal. Note that a non-allocated entry of the diagonal costs at least 1, and again this cannot be compensated by the $1/10$ potential contribution of the column bidder. \square

To complete the proof, observe that all 2^k auctions in \mathcal{R} satisfy monotonicity, hence are feasible. Also note that there is a natural bijection between subsets $S \subseteq [k]$ and these auctions: we include b_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 auction 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 below 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 above 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:

$$\text{SW}(M(S)) = (2/3) \cdot (k-1)k(k+1) + (1/2) \cdot k(k+1) + \sum_{i=2}^k (i-1)b'_i + \sum_{i \in S} b'_i \quad (3.4)$$

Recalling that $\sum_{i=1}^k b'_i = 1/(10k)$, (3.3) and (3.4) imply that there exists a partition of B' iff there exists a feasible auction satisfying (3.3). This completes the proof. \square

We are now ready to provide the proof of Theorem 9.

Proof of Theorem 9. The idea is similar to that of the previous proof for the exact version of the welfare objective but the details are more elaborate. At a high-level, the difficulty is that the two objective functions (welfare, revenue) depend on each other in a subtle way. Thus, a more complicated construction is required to “decouple” these two criteria. (It is not hard to see that the construction presented for the exact version fails for the bi-objective problem.) Very roughly, the reduction ends up using non-uniform distributions on larger and carefully selected supports.

As before, our reduction is from PARTITION. We start with a set $A = \{a_1, \dots, a_k\}$ of positive numbers (rescaled so that they sum to a sufficiently small positive constant) and we want to decide whether there exists a partition of this set. We will construct an instance of the auction problem with 2 bidders and distributions of support size $2k+1$. Before presenting the actual instance we first give some intuition behind the construction.

Similarly, our goal is to establish a bijection between an appropriate subset of feasible auctions and subsets S of $[k]$; since the number of feasible auctions greatly exceeds that of subsets of A , we have to limit our attention to a subset of feasible auctions. To that end, we are going to appropriately pick the target values for welfare and revenue, so that the only relevant auctions in our reduction will be those that allocate the item to bidder 2 (column bidder) for entries above the secondary diagonal (i.e. (v_1^i, v_2^j) with $j > i$) and to bidder 1 (row bidder) for entries below the secondary diagonal (i.e. (v_1^i, v_2^j) with $j < i$). We will also rule out the possibility of not allocating the item across the diagonal entries of the allocation matrix. As a result, the only relevant auctions will be the 2^{2k+1} different auctions that allocate to either bidder 1 or 2 across the diagonal, all of which respect monotonicity and are therefore feasible. Call this subset of auctions \mathcal{R} . We will then use the i -th *odd*

entry of the diagonal to encode the decision of including or not the i -th element of A in the set S : we shall include element a_i iff bidder 1 gets allocated for entry $(2i - 1, 2i - 1)$ of the allocation matrix.

The first step is to ensure that the only relevant auctions are the ones with the above property. To this end we ask that the following relation between the bidders' valuations holds:

$$v_1^i < v_2^i < v_1^{i+1} < v_2^{i+1} < v_1^{i+2}, \text{ for } i = 1 \dots 2k - 1 \quad (3.5)$$

As a result of relation (3.5), the social welfare from entries other than those on the diagonal is maximized by an auction that allocates to bidder 2 on top of the diagonal and to bidder 1 below the diagonal. Therefore, by setting a sufficiently high welfare target W in our reduction we will be able to guarantee that the only relevant auctions will be of this format.

More specifically, the distributions are defined as follows (where $\epsilon > 0$ a sufficiently small parameter that we will fix later):

The (unnormalized) probabilities of the two bidders are:

$$f_j^i = \begin{cases} 1 & \text{if } i \text{ is odd;} \\ \epsilon & \text{if } i \text{ is even.} \end{cases}$$

for both bidders $j = 1, 2$. (The point of the small probability elements is to achieve the desired decoupling between welfare and revenue; it may be convenient for the reader to think of ϵ as if it was 0. In the course of the proof we will provide a sufficient upper bound on its magnitude.)

The values of bidder 1 are:

$$v_1^i = \begin{cases} i + a_{\frac{i+1}{2}} & \text{for } i \in \{1, 3, \dots, 2k - 1\}; \\ i + a_{\frac{i}{2}} \left(1 + \frac{4}{(2k-i+2)(1+\epsilon)}\right) & \text{for } i \in \{2, 4, \dots, 2k\}; \\ 2k + 1 & \text{for } i = 2k + 1. \end{cases}$$

The values of bidder 2 are:

$$v_2^i = \begin{cases} i & \text{for } i \in \{1, 3, \dots, 2k - 1\}; \\ i & \text{for } i \in \{2, 4, \dots, 2k\}; \\ 2k + 1 & \text{for } i = 2k + 1. \end{cases}$$

We note that there are 3 different scales of numbers in the reduction. The values of the elements in the support (big scale), the magnitudes of the elements of A (medium scale), and the magnitude of ϵ (small scale).

What we would like to do next is to make the sum of welfare and revenue remain constant across all auctions in \mathcal{R} . By doing so we can ensure that, whenever an auction achieves the target welfare and revenue values, the relations will in fact hold with equality, allowing us to encode an instance of PARTITION. To achieve that, we impose an even stronger requirement: In particular, consider the following entries of the allocation matrix: $(v_1^i, v_2^j), j = i \dots 2k + 1$ and $(v_1^j, v_2^i), j = i \dots 2k + 1$, where i is an odd number. Assuming our auction has the

format discussed above, entries $(v_1^i, v_2^j), j = i + 1 \dots 2k + 1$ are allocated to bidder 2, entries $(v_1^j, v_2^i), j = i + 1 \dots 2k + 1$ are allocated to bidder 1, and we are left to decide which bidder to allocate entry (v_1^i, v_2^i) to. Now let SW_j^i (resp. Rev_j^i), where i is odd and $j \in \{1, 2\}$, denote the welfare (resp. revenue) that results from the aforementioned entries if we allocate entry (v_1^i, v_2^i) to bidder j . The stronger requirement that we impose is that $SW_1^i + Rev_1^i = SW_2^i + Rev_2^i$ for all odd i . To see what this entails we next write the expressions for SW_j^i and Rev_j^i :

$$\begin{aligned} SW_1^i &= v_1^i + \sum_{j=\frac{i+1}{2}}^k v_1^{2j+1} + \sum_{j=\frac{i+1}{2}}^k v_2^{2j+1} + \epsilon \cdot \left(\sum_{j=\frac{i+1}{2}}^k v_1^{2j} + \sum_{j=\frac{i+1}{2}}^k v_2^{2j} \right) \\ Rev_1^i &= v_1^i \left(\frac{2k-i+1}{2}(1+\epsilon) + 1 \right) + v_2^{i+1} \left(\frac{2k-i-1}{2} + \frac{2k-i+1}{2}\epsilon + 1 \right) \\ SW_2^i &= \sum_{j=\frac{i+1}{2}}^k v_1^{2j+1} + v_2^i + \sum_{j=\frac{i+1}{2}}^k v_2^{2j+1} + \epsilon \cdot \left(\sum_{j=\frac{i+1}{2}}^k v_1^{2j} + \sum_{j=\frac{i+1}{2}}^k v_2^{2j} \right) \\ Rev_2^i &= v_1^{i+1} \left(\frac{2k-i-1}{2} + \frac{2k-i+1}{2}\epsilon + 1 \right) + v_2^i \left(\frac{2k-i+1}{2}(1+\epsilon) + 1 \right) \end{aligned}$$

Notice that $SW_1^i - SW_2^i = v_1^i - v_2^i$. In order to have $SW_1^i + Rev_1^i = SW_2^i + Rev_2^i$ we ask that:

$$Rev_1^i - Rev_2^i = v_2^i - v_1^i \quad (3.6)$$

The only difficulty in satisfying the relation above, is that equation (3.6) necessarily imposes some additional constraints on the values v_1^{i+1}, v_2^{i+1} . We get around this by using a support of roughly twice the size of A , and using only half of the points in the support to encode the elements of A ; the remaining points are assigned a very small probability, so that they have a negligible effect on the overall welfare and revenue. It is now easy to verify that the aforementioned choice of distributions for the two bidders satisfies properties (3.5) (since the a_i are assumed to be sufficiently small) and (3.6) above.

For the aforementioned choice of values v_j^i the social welfare contributions now become:

$$SW_1^i = v_1^i + X_i = i + a_{\frac{i+1}{2}} + X_i \text{ and } SW_2^i = v_2^i + X_i = i + X_i,$$

for some X_i whose exact value is irrelevant (and can be derived from the expressions above). Analogously for revenue we have:

$$Rev_1^i = v_2^i + Y_i = i + Y_i \text{ and } Rev_2^i = v_1^i + Y_i = i + a_{\frac{i+1}{2}} + Y_i,$$

for some Y_i . We therefore have $SW_1^i + Rev_1^i = SW_2^i + Rev_2^i = 2i + X_i + Y_i + a_{\frac{i+1}{2}}$ and we have thus ensured that all auctions with the property of allocating to bidder 2 on top of the

diagonal and to bidder 1 below the diagonal have a sum of (total) revenue and welfare that can be upper-bounded as follows:

$$SW + \text{Rev} \leq \sum_{\text{odd } i} (2i + X_i + Y_i + a_{\frac{i+1}{2}}) + v_1^{2k+1} + v_2^{2k+1} + \epsilon \cdot 2n^2(2k+1),$$

where the last term is an upper bound on the contribution in revenue and welfare of the even rows and columns (where we took into account that the maximum contribution of any entry is at most the maximum value appearing in the support of any bidder, namely $2k+1$). We next fix the value of ϵ so that the quantity

$$\epsilon \cdot 2n^2(2k+1)$$

is smaller than the accuracy used in the rational numbers a_i . Note that this can always be done with an ϵ that has polynomially many bits – since the a_i 's are by assumption rational numbers with polynomially many bits.

We are now ready to argue that there exists a partition of A iff there exists an auction with:

$$SW \geq \sum_{\text{odd } i} (i + X_i) + 2k + 1 + \frac{1}{2} \sum_{i=1}^k a_i \quad \text{and} \quad \text{Rev} \geq \sum_{\text{odd } i} (i + Y_i) + 2k + 1 + \frac{1}{2} \sum_{i=1}^k a_i \quad (3.7)$$

Given any partition S of A , we can turn it into an auction with the above welfare and revenue guarantees by allocating to bidder 2 on top of the diagonal, bidder 1 below the diagonal and allocating to bidder 1 for entries $(2i-1, 2i-1)$, $i = 1 \dots k$, for all i s.t. $a_i \in S$; the even entries on the diagonal, as well as the entry $(2k+1, 2k+1)$ can be allocated to either bidder.

Conversely, given an auction with welfare and revenue as above we can get a partition of A . To see how, first notice that because of property (3.5) above (and because the a_i are much smaller) the only auctions that can achieve a social welfare of at least $\sum_{\text{odd } i} i + X_i$ and revenue of at least $\sum_{\text{odd } i} i + Y_i$ are those that allocate to bidder 2 above the diagonal, bidder 1 below the diagonal, and always allocate to either bidder 1 or bidder 2 on the diagonal. In the discussion above we established that for those auctions it holds that:

$$SW + \text{Rev} \leq \left(\sum_{\text{odd } i} (i + X_i) + 2k + 1 + \frac{1}{2} \sum_{i=1}^k a_i \right) + \left(\sum_{\text{odd } i} (i + Y_i) + 2k + 1 + \frac{1}{2} \sum_{i=1}^k a_i \right) + \epsilon \cdot 2n^2(2k+1) \quad (3.8)$$

By our choice of ϵ and inequalities (3.7) and (3.8) it follows that the inequalities in (3.7) must hold with equality. We then get a partition by including in S all elements a_i for which $(2i-1, 2i-1)$ is allocated to bidder 1. This completes the proof. \square

Finally, we 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 $P = NP$).

Theorem 11. *There exists a family of two-bidders instances for which the size of the Pareto curve for BI-CRITERION AUCTION grows exponentially.*

Proof. The construction is similar to the reduction for the exact problem in Theorem 9. We will construct a two-bidder auction and we will argue that there exists an appropriate *subset* of the Pareto curve with exponential size.

We describe an instance with two bidders, both with uniform distributions over the following supports of size k : Bidder 1 has values in $\{1 + a_1, 2 + a_2, \dots, k + a_k\}$, and bidder 2 has values in $\{1, 2, \dots, k\}$, with all $a_i \ll 1$; the exact value of a_i will be determined later. Assuming as usual that bidder 1 is the row bidder, consider an auction that allocates the item to bidder 2 for all entries above the diagonal (i.e. (v_1^i, v_2^j) with $j > i$), to bidder 1 for all entries below the diagonal (i.e. (v_1^i, v_2^j) with $j < i$), and to either bidder on the diagonal (see Table 3.1). Such an auction can be concisely described through the diagonal entries of its allocation matrix. In what follows we write $\text{SW}(A[v_1^1, v_2^1], \dots, A[v_1^k, v_2^k])$ and $\text{Rev}(A[v_1^1, v_2^1], \dots, A[v_1^k, v_2^k])$ to denote the welfare and revenue respectively of this auction. We note that this subset of feasible auctions maximizes the welfare over all feasible auctions; hence, it suffices to show that the Pareto set of this subset of auctions is exponential. In fact, we will choose the a_i 's appropriately so that *all* these auctions are undominated.

Our goal is to pick values a_1, \dots, a_k such that all 2^k different auctions of the above type will be Pareto optimal. To do that we observe that, under some mild conditions on the a_i , satisfied by picking for example $a_i = 3^{i-1}$ (and normalizing so that the normalized sum is small, e.g. $< 1/1000$), we can impose orderings on the welfares and revenues of those 2^k auctions that go in opposite directions, i.e. one auction has larger revenue than another iff it has smaller welfare. To this end we make the following two claims, which can be verified by explicitly writing down the expressions for revenue and welfare and doing some elementary calculations:

Claim 1: If $a_i > 0$ and $a_i < \frac{n-i}{n-i+1}a_{i+1}$ for all i , it holds that:

1. $\text{SW}(\xi_1, \dots, \xi_{i-1}, 1, \xi_{i+1}, \dots, \xi_k) > \text{SW}(\xi_1, \dots, \xi_{i-1}, 2, \xi_{i+1}, \dots, \xi_k)$
2. $\text{Rev}(\xi_1, \dots, \xi_{i-1}, 1, \xi_{i+1}, \dots, \xi_k) < \text{Rev}(\xi_1, \dots, \xi_{i-1}, 2, \xi_{i+1}, \dots, \xi_k)$

for any values of the $\xi_i \in \{1, 2\}, i = 1 \dots k$.

Claim 2: If $\sum_{j=1}^{i-1} a_j < a_i$ and $a_i < \frac{n-i}{2(n-i+1)}a_{i+1}$ for all i , it holds that:

1. $\text{SW}(1, \dots, 1, 2, \xi_{i+1}, \dots, \xi_k) < \text{SW}(2, \dots, 2, 1, \psi_{i+1}, \dots, \psi_k)$
2. $\text{Rev}(1, \dots, 1, 2, \xi_{i+1}, \dots, \xi_k) > \text{Rev}(2, \dots, 2, 1, \psi_{i+1}, \dots, \psi_k)$

for any values of the $\xi_i, \psi_i \in \{1, 2\}, i = 1 \dots k$; note that in general we may have $\psi_i \neq \xi_i$.

Intuitively, Claim 1 says that switching any 1 into a 2 on any entry of the diagonal has the effect of decreasing the welfare while increasing the revenue. Claim 2 on the other hand

says that the (negative) effect that a 2 on the diagonal has on the welfare, is bigger for 2's that are placed in higher positions –and the opposite is true for revenue.

Using the above two claims one can now prove that for any two auctions $M_1 = (\xi_1, \dots, \xi_k)$ and $M_2 = (\psi_1, \dots, \psi_k)$, it holds that $\text{SW}(M_1) > \text{SW}(M_2)$ iff $\text{Rev}(M_1) < \text{Rev}(M_2)$, and therefore all the 2^k auctions are Pareto optimal. We convey the idea by means of the following example, for $k = 5$. Consider the two auctions $M_1 = (1, 2, 1, 1, 1)$ and $M_2 = (1, 1, 2, 2, 1)$. We then have:

$$\text{SW}(M_1) = \text{SW}(1, 2, 1, 1, 1) > \text{SW}(2, 2, 2, 1, 1) > \text{SW}(1, 1, 1, 2, 1) > \text{SW}(1, 1, 2, 2, 1) = \text{SW}(M_2)$$

with the inequalities above following from Claims 1.1, 2.1 and 1.1 respectively. In complete analogy we can show that $\text{Rev}(M_1) < \text{Rev}(M_2)$. \square

$v_2^k = k$	2	2	2	...	1 or 2
$v_2^{k-1} = k - 1$	2	2	...	1 or 2	1
\vdots	\vdots	\vdots	\dots	\vdots	\vdots
$v_2^2 = 2$	2	1 or 2	...	1	1
$v_2^1 = 1$	1 or 2	1	...	1	1
	$v_1^1 = 1 + a_1$	$v_1^2 = 2 + a_2$...	$v_1^{k-1} = k - 1 + a_{k-1}$	$v_1^k = k + a_k$

Table 3.1: An instance with an exponential size Pareto set

3.4 An FPTAS for 2 bidders

In this section we give our main algorithmic result:

Theorem 12. *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 8. There is a difficulty, however, in showing that the problem satisfies the assumptions of Theorem 8: 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 8.

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

Definition 15. *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 $v_2^{i_2}$:*

$r_1^{i_1, i_2} = \sum_{j \geq i_1} v_1^{i_1} \cdot f(v_1^j, v_2^{i_2})$ and $w_1^{i_1, i_2} = \sum_{j \geq i_1} v_1^j \cdot f(v_1^j, v_2^{i_2})$, 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 13. *The BI-CRITERION AUCTION problem can be expressed in a way that satisfies the conditions of Theorem 8.*

Proof. We consider variables x_{ij}, y_{ij} , $i \in [|Sup_1|]$, $j \in [|Sup_2|]$. 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:

$$\begin{aligned} \text{Rev}(x, y) &= \sum_{i=1}^{|Sup_1|} \sum_{j=1}^{|Sup_2|} x_{ij} r_1^{i,j} + \sum_{i=1}^{|Sup_1|} \sum_{j=1}^{|Sup_2|} y_{ij} r_2^{i,j} \\ \text{SW}(x, y) &= \sum_{i=1}^{|Sup_1|} \sum_{j=1}^{|Sup_2|} x_{ij} w_1^{i,j} + \sum_{i=1}^{|Sup_1|} \sum_{j=1}^{|Sup_2|} y_{ij} w_2^{i,j} \end{aligned}$$

□

An algorithm for the exact version of Bi-Criterion Auction

The main idea behind our algorithm, inspired by the characterization of Lemma 13, is to consider the contribution from each bidder (fixing the value of the other) independently, by going over all (linearly many) single-bidder auctions for both bidders. The challenging part is to combine the individual single-bidder auctions into a single two-bidders auction, 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 her 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 14. *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. We start by fixing a pair (i, j) and considering what $A[i, j]$ can be, given the entries of A for either larger i , or larger j , or both. First note that any allocation matrix A can have one of the following four formats:

- 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 format F1 can be valid, while the other three formats correspond to the three terms of the recurrence. Finally, note that for format, say F2, the second 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 $w_j^{i_1, i_2}$.

3.5 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 an auction. 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 13. *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.

Lower bound. The NP-hardness reduction is from PARTITION. Let $B = \{b_1, \dots, b_n\}$ be a set of positive rationals; we can assume wlog (because of rescaling) that $\sum_{i=1}^n b_i = 1/100$. We construct an instance of the auction problem as follows: there are n 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 $\sum_{i=1}^n h_i + (1/2) \sum_{i=1}^n b_i$. To see this notice that – since the sequence $\{h_i\}_{i=1 \dots n}$ is super-increasing – any auction

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 auctions, 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 $v_i = l_i$. The first part of the theorem now follows.

Upper bound. 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 [32]) with a number of nodes that is exponential in the number of bidders: Assuming that each bidder has two values $\{v_i^1, v_i^2\}$, with $v_i^2 > v_i^1$, we create a node labeled by $(v_i^{k_i})_{i=1\dots n}$, where $k_i \in \{1, 2\}$, for each one of those 2^n valuation tuples. We then connect two such nodes if their labels differ in exactly one coordinate, say the i -th one; the weight of this edge is $\prod_{j \neq i} f_j^{k_j} \cdot (f_i^1 v_i^1 + f_i^2 v_i^2)$ for welfare and $\prod_{j \neq i} f_j^{k_j} \cdot v_i^1 (f_i^1 + f_i^2)$ for revenue. We also introduce a set of dummy nodes: for every node with label $(v_i^{k_i})_{i=1\dots n}$, where $|\{i \mid k_i = 2\}| = r$, we introduce r dummy nodes and associate each one of them with the bidder i for whom $v_i^{k_i} = v_i^2$. We then add an edge between this node and all its dummy nodes; the weight of the edge connecting to the dummy node of the i -th bidder is $\prod_{j \neq i} f_j^{k_j} \cdot f_i^2 v_i^2$ both for welfare and revenue.

It is easy to verify that every matching in the above graph corresponds to a deterministic truthful auction as follows. For each valuation tuple consider the corresponding node in the graph. If it is not matched to any other node, then allocate nothing; if it is matched to a dummy node, then allocate the item to the bidder that is associated with this dummy node; otherwise, it is matched to another non-dummy node and these two nodes differ in exactly one coordinate, say the i -th, in which case we allocate to the i -th bidder. It is easy to check that the resulting auction is both feasible and monotone (IC and IR). Moreover, the welfare (or revenue depending on the kind of weights used) is equal to the weight of the matching.

The EXACT MATCHING problem is known to be solvable in RNC [55]; since our input provides an exponentially succinct representation of the constructed graph, we are interested in the so-called *succinct version* of the problem [38, 61]. 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.

3.6 Open questions

Is there is an FPTAS for three bidders with independently distributed valuations? We conjecture that there is, and in fact for any constant number of bidders. Of course, the approach of our FPTAS for two bidders cannot be generalized, since it works for the correlated case, which is APX-complete for three 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 three 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 Section 3.1. Are there improved approximation guarantees for more general kinds of distributions and n bidders?

Chapter 4

Simple, Optimal and Efficient Auctions

4.1 Introduction

Somewhat surprisingly, as we have already discussed, when the bidders' values are independently and identically distributed according to some regular¹ distribution, the Vickrey and Myerson auctions behave very much alike: Myerson's auction is just Vickrey's auction with an additional reserve price. Motivated by this astonishing similarity (and the somewhat peculiar format of Myerson's auction in more general settings), [46] showed that a Vickrey auction with appropriately chosen reserve prices can approximate the revenue of the optimal auction in more general settings. Inspired by their result, and the fact that the auction of Myerson and Satterthwaite [56] is at least as complicated as Myerson's auction and potentially randomized, in this chapter we ask the question of whether one can design *simple and deterministic* auctions that achieve *approximately-optimal* guarantees for both objectives *simultaneously*.

For the remaining of this chapter we will use the term *simple auction* to refer to a Vickrey auction with appropriately chosen reserve prices, formally defined as follows:

Definition 16. *The Vickrey auction with reserve prices $\mathbf{r} = (r_1, \dots, r_n)$, denoted $\text{Vic}_{\mathbf{r}}$, is the following auction:*

1. *Reject all bidders whose values are $v_i < r_i$.*
2. *Allocate the item to the remaining bidder with the highest valuation (or to none if no one clears their reserve in Step 1). Tie-break lexicographically if there are multiple highest bidders.*
3. *Charge the winner the maximum of the second highest bidder (among those who were not eliminated in Step 1) and her reserve price.*

¹See Section 1.3 for a definition.

Two cases of particular interest are the Vickrey auction with an anonymous reserve, where a common reserve r is used for all bidders, and the Vickrey auction with monopoly reserves, denoted by Vic_m , where $m_i = \phi_i^{-1}(0)$, the monopoly reserve of bidder i .

At first glance it is not obvious why such simple auctions that achieve approximately-optimal guarantees for both objectives simultaneously should even exist. Indeed, despite the fact that Vickrey's auction achieves at least half of the optimal revenue, when the values are drawn i.i.d. from regular distributions (see e.g. [34]), this is no longer the case when the values are independent but drawn from different regular distributions. In particular, it is easy to see that the revenue of Vickrey's auction can be arbitrarily far from the optimal revenue: just consider $n - 1$ bidders distributed independently and uniformly in $[0, 1]$, and a single bidder distributed uniformly in $[h, h + 1]$, for some large $h > 1$. We can try to fix that by using the Vickrey auction with monopoly (player-specific) reserves. However, despite the revenue guarantees provided in [46], the auction now can be arbitrarily inefficient even for a single bidder whose value is distributed according to a regular distribution. Indeed, consider the (almost) equal revenue distribution, where the bidder's value is supported on $\{1 - \epsilon, 2 - \epsilon, \dots, h - \epsilon\}$, for some $\epsilon \in (0, 1)$ and $h > 1$, and the probability that it is larger than or equal to $i - \epsilon$ is exactly $1/i$, for $i = 1 \dots h$. In this chapter, we show that, by appropriately tweaking the reserve price of each bidder, we can fix this inefficiency:

Main Result (Theorem 14 of Section 4.2): *In every single item setting with n bidders whose values are distributed according to independent (possibly non-identical) regular distributions, and for every $p \in [0, 1]$, there exists a Vickrey auction with (generally non-anonymous) reserve prices that simultaneously achieves a p -fraction of the optimal social welfare and a $(\frac{1-p}{4})$ -fraction of the optimal revenue. In particular, there exists a Vickrey auction with reserve prices that achieves at least a 20% of the optimal social welfare and revenue.*

We can use our techniques to prove a similar approximation guarantee for non-identical distributions satisfying the monotone hazard rate condition (which has already been obtained by [26]). We also show that a Vickrey auction with an anonymous reserve simultaneously approximates both objectives for general (possibly non-regular) distributions, as long as all values are i.i.d (Theorem 16). We summarize our results together with already known welfare and revenue guarantees for various settings in Table 4.1.

	i.i.d.	independent
mhr	$(1, \frac{1}{e})$ and $(\frac{1}{e}, 1)$ [2]	$(\frac{1}{e}, \frac{1}{2})$ [26]
regular	$(1, \frac{1}{2})$ [34]	$(\frac{1}{5}, \frac{1}{5})$ and $(p, \frac{1-p}{4})$, for all $p \in [0, 1]$ [24]
non-regular	$(\frac{1}{2}, \frac{1}{2})$ [24]	?

Table 4.1: (α, β) stands for α -approximation for welfare and β -approximation for revenue. Notice that our result for regular distributions gets a handle on the whole Pareto boundary achieved by the Vickrey auction with non-anonymous reserve prices.

Two questions left open are whether one can extend our results to the setting of n bidders distributed according to independent, but not necessarily identical, and possibly non-regular distributions, and to general single-dimensional settings (such as matroids and general downward-closed environments).

Related Work

The work closer in spirit to ours is that of [46], where the authors show that, for a variety of single-dimensional settings, Vickrey auctions with carefully chosen reserve prices are approximately revenue-optimal. In particular, when the bidders' values are independently drawn from (possibly different) regular distributions, they show that Vickrey's auction with monopoly reserve prices achieves at least half of the optimal revenue. Moreover, they show that, for single item settings, Vickrey's auction with an anonymous reserve achieves a factor 4 approximation to the optimal revenue. Finally, [46] also extends Bulow and Klemperer's result (which we discussed in the previous chapter) to more general single-dimensional settings, as follows: They show that by duplicating all bidders (whose values are drawn independently from not necessarily identical, regular distributions), and then running the VCG auction, one can guarantee at least half of the optimal revenue (while being optimal with respect to welfare). Our result shows that in single item settings with independent (but not necessarily i.i.d.) bidders, one can simultaneously achieve constant factor approximations to both the optimal revenue and welfare *without adding any extra bidders*, via the use of a Vickrey auction with appropriate (non-anonymous) reserve prices.

The work presented in the previous chapter is obviously very relevant to the problem at hand as well. The two-bidder auction we presented there however, despite being deterministic, is far from simple; this chapter complements the previous one by showing that, if one is willing to settle for less than an arbitrarily small loss in performance, simple auctions are possible, even when the number of bidders is large. Moreover, the existence of an auction that simultaneously achieves a constant factor approximation to both objectives, characterizes the "knee" of the Pareto curve, a structural result which is of independent interest.

The literature mentioned in the related work section of the previous chapter, which studies the revenue and welfare guarantees of welfare-maximizing and revenue-maximizing auctions respectively, is also very relevant. In particular, in [2] the authors show that, for values drawn independently from the same monotone hazard rate distribution, both the welfare and revenue ratios of Vickrey and Myerson's auctions are bounded by $1/e$ (see the top-left square of Table 4.1). Furthermore, in [34] and [26] the authors present simple auctions that simultaneously achieve constant factor approximations to both objectives in single item settings where bidders' values are either i.i.d. from a regular distribution (see the middle-left square of Table 4.1), or independently (but not necessarily identically) distributed according to a monotone hazard rate distribution (see the top-right square of Table 4.1). Some of their results also hold for more general single-dimensional settings, namely when the feasibility constraints form a matroid.

Finally, despite our different motivation, methodologically our work is somewhat related to [3]: in that paper the goal is to provide a general reduction from the mechanism design problem for many bidders to that of a single bidder, while preserving the value of a separable objective (such as welfare or revenue) within a constant factor. In Lemmas 17 and 19 we establish analogous many-to-one reductions; however, our goal is not only to preserve the approximation factor, but also for the resulting auction to be of a specific simple format, in contrast to the much more generic reduction of [3].

Notation

We say that an auction \mathcal{A} is an α -approximation for welfare (resp. revenue) if $\text{SW}(\mathcal{A}) \geq \alpha \cdot \text{SW}(\text{Vic})$ (resp. $\text{Rev}(\mathcal{A}) \geq \alpha \cdot \text{Rev}(\text{Mye})$), where Vic denotes Vickrey's auction and Mye denotes Myerson's auction. We say that an auction is an (α, β) -approximation if it is simultaneously an α -approximation for welfare and a β -approximation for revenue. Also, given an auction \mathcal{A} , and a set of bidders $B \subseteq \{1, \dots, n\}$, we may write $\mathcal{A}(B)$ to denote the auction \mathcal{A} run only on the subset B of bidders. When we use this notation it will be clear from the context how the “projected” auction operates. Finally, for convenience, we sometimes write $\mathcal{R}_{\mathcal{A}} = \sum_{i=1}^n p_i(v_1, \dots, v_n)$, so that $\text{Rev}(\mathcal{A}) = \mathbb{E}[\mathcal{R}_{\mathcal{A}}]$.

4.2 The regular, independent case

In this section we focus on the setting of n bidders whose values are distributed according to regular, but not necessarily identical, distributions. We start with a couple of probabilistic lemmas –not requiring regularity– which have easy proofs. Since those proofs make use of the notion of *stochastic dominance of measures* we remind this to the reader first.

Definition 17 (Stochastic Dominance). *Let \mathcal{F}_1 and \mathcal{F}_2 be distributions over \mathbb{R} . We say that \mathcal{F}_2 stochastically dominates \mathcal{F}_1 iff there exist random variables X_1 and X_2 that are marginally distributed according to \mathcal{F}_1 and \mathcal{F}_2 respectively, and a coupling of X_1 and X_2 , such that $X_1 \leq X_2$, with probability 1.*

Lemma 15. *Let X and Y be independent random variables and $g : \mathbb{R} \rightarrow \mathbb{R}$ a (weakly) increasing function. Then, for any constant $c \in \mathbb{R}$,*

$$\Pr[X \geq Y \mid g(X) \geq c] \geq \Pr[X \geq Y \mid g(X) \leq c].$$

Proof. Fix c , let $\mathcal{I} := \{x \mid g(x) = c\}$ and, without loss of generality, assume that $\Pr[g(x) \leq c] \leq \Pr[g(x) \geq c]$. Let \mathcal{F}_1 be the distribution of X , conditioning on $g(X) \leq c$, and \mathcal{F}_2 the distribution of X , conditioning on $g(X) \geq c$. We claim that \mathcal{F}_2 stochastically dominates \mathcal{F}_1 . Indeed, let X_1 be a random variable distributed according to \mathcal{F}_1 and X_2 a random variable distributed according to \mathcal{F}_2 . Define any coupling of X_1 and X_2 enforcing that whenever $X_2 \in \mathcal{I}$, $X_1 = X_2$. This is easy to achieve since, at every point $x \in \mathcal{I}$, \mathcal{F}_1 has more

probability mass than \mathcal{F}_2 (using our assumption that $\Pr[g(x) \leq c] \leq \Pr[g(x) \geq c]$). It is now easy to verify that any such coupling satisfies that $X_1 \leq X_2$, with probability 1. (For completeness, we note that, if instead we had $\Pr[g(x) \leq c] \geq \Pr[g(x) \geq c]$, we would pick any coupling satisfying that whenever $X_1 \in \mathcal{I}$, $X_2 = X_1$.)

Suppose that X_1 and X_2 are coupled as above and sample Y independently from X_1 and X_2 . In the joint distribution \mathcal{F} thus defined, whenever $X_1 \geq Y$, it must also be that $X_2 \geq Y$ (since by stochastic domination $\Pr[X_2 \geq X_1] = 1$). Hence under \mathcal{F} : $\Pr[X_2 \geq Y] \geq \Pr[X_1 \geq Y]$. The lemma now follows by simply noticing that the marginal of \mathcal{F} over the pair (X_1, Y) is identical to the distribution of X and Y conditioning on $g(X) \leq c$, and similarly for (X_2, Y) . \square

Lemma 16. *Let X and Y be independent random variables and $g : \mathbb{R} \rightarrow \mathbb{R}$ a (weakly) increasing function. Then*

$$\mathbb{E}[g(X)] \leq \mathbb{E}[g(X) \mid X \geq Y].$$

Proof. For any constant $c \in \mathbb{R}$:

$$\begin{aligned} \mathbb{E}[g(X)] &= \mathbb{E}[g(X) \mid X \geq c] \Pr[X \geq c] + \mathbb{E}[g(X) \mid X \leq c] \Pr[X \leq c] \\ &= \mathbb{E}[g(X) \mid X \geq c] + (\mathbb{E}[g(X) \mid X \leq c] - \mathbb{E}[g(X) \mid X \geq c]) \cdot \Pr[X \leq c] \\ &\leq \mathbb{E}[g(X) \mid X \geq c], \end{aligned} \tag{4.1}$$

where the inequality follows from the fact $\mathbb{E}[g(X) \mid X \leq c] \leq \mathbb{E}[g(X) \mid X \geq c]$. This is true since g is a non-decreasing function, and the conditional distribution of X , conditioning on $X \geq c$, stochastically dominates the conditional distribution of X , conditioning on $X \leq c$ (this is a special case of what we argued in the beginning of the proof of Lemma 15).

To conclude the lemma, let $f(y)$ be the density function of Y . Notice that:

$$\mathbb{E}[g(X)] \equiv \int_y \mathbb{E}[g(X)] f(y) dy \leq \int_y \mathbb{E}[g(X) \mid X \geq y] f(y) dy \equiv \mathbb{E}[g(X) \mid X \geq Y],$$

where the equalities follow from the independence of X and Y and the inequality follows from applying (4.1) pointwise for all y . \square

Our next lemma shows that if we take the Vickrey auction and add a reserve price for each bidder, such that the probability of any single bidder's value exceeding her reserve price is at least p , then the resulting welfare is at least a p fraction of Vickrey's (optimal) social welfare $\mathbb{E}[\max_i \{v_i\}]$. In what follows we use $\mathbb{I}_{(\cdot)}$ to denote the indicator function.

Lemma 17 (Many-to-One Reduction—Welfare). *Suppose that X_1, \dots, X_n are independent, non-negative random variables (possibly non-identically distributed), t_1, \dots, t_n are (possibly different) thresholds, and $p \in [0, 1]$. If it holds that $\Pr[X_i \geq t_i] \geq p$, for all $i = 1 \dots n$, then:*

$$\mathbb{E} \left[\max_i \{X_i \cdot \mathbb{I}_{(X_i \geq t_i)}\} \right] \geq p \cdot \mathbb{E} \left[\max_i \{X_i\} \right].$$

Proof. Let $\mathcal{E}_i = \{X_i = \max_j \{X_j\}\}$. Then $\mathbb{E}[\max_i \{X_i\}]$ is:

$$\sum_{i=1}^n \mathbb{E}[X_i \mid \mathcal{E}_i, X_i \geq t_i] \Pr[\mathcal{E}_i, X_i \geq t_i] + \mathbb{E}[X_i \mid \mathcal{E}_i, X_i \leq t_i] \Pr[\mathcal{E}_i, X_i \leq t_i]$$

To proceed we need the following claims:

Claim 1: $\Pr[\mathcal{E}_i, X_i \geq t_i] \geq \frac{p}{1-p} \cdot \Pr[\mathcal{E}_i, X_i \leq t_i]$.

Proof of Claim 1.

$$\begin{aligned} \Pr[\mathcal{E}_i, X_i \geq t_i] &= \Pr[\mathcal{E}_i \mid X_i \geq t_i] \cdot \Pr[X_i \geq t_i] \\ &\geq \Pr[\mathcal{E}_i \mid X_i \leq t_i] \cdot \frac{p}{1-p} \Pr[X_i \leq t_i] \\ &= \frac{p}{1-p} \cdot \Pr[\mathcal{E}_i, X_i \leq t_i], \end{aligned}$$

where in the inequality above we used the following facts: First, we use that

$$\Pr[\mathcal{E}_i \mid X_i \geq t_i] \geq \Pr[\mathcal{E}_i \mid X_i \leq t_i],$$

which follows from Lemma 15 taking $X = X_i$, $c = t_i$, g the identity function, $Y = \max_{j \neq i} \{X_j\}$, and noticing that the event \mathcal{E}_i is the same as the event $X \geq Y$. The second fact we use is that $\Pr[X_i \geq t_i] \geq \frac{p}{1-p} \Pr[X_i \leq t_i]$, which in turn follows from the fact that $\Pr[X_i \geq t_i] \geq p$ and $\Pr[X_i \leq t_i] \leq 1-p$. \square

Claim 2: $\mathbb{E}[X_i \mid \mathcal{E}_i, X_i \geq t_i] \geq \mathbb{E}[X_i \mid \mathcal{E}_i, X_i \leq t_i]$.

Proof of Claim 2. Just note: $\mathbb{E}[X_i \mid \mathcal{E}_i, X_i \geq t_i] \geq t_i \geq \mathbb{E}[X_i \mid \mathcal{E}_i, X_i \leq t_i]$. \square

From the above claims and the non-negativity of X_i it follows that:

$$\mathbb{E}\left[\max_i \{X_i\}\right] \leq \frac{1}{p} \cdot \sum_{i=1}^n \mathbb{E}[X_i \mid \mathcal{E}_i, X_i \geq t_i] \Pr[\mathcal{E}_i, X_i \geq t_i]. \quad (4.2)$$

Next we write $\mathbb{E}[\max_i \{X_i \cdot \mathbb{1}_{(X_i \geq t_i)}\}]$ as follows:

$$\begin{aligned} &\sum_{i=1}^n \mathbb{E}[X_i \mid \mathcal{E}_i, X_i \geq t_i] \Pr[\mathcal{E}_i, X_i \geq t_i] + \mathbb{E}[X_i \mid \neg \mathcal{E}_i, X_i \geq t_i] \Pr[\neg \mathcal{E}_i, X_i \geq t_i] \\ &\geq \sum_{i=1}^n \mathbb{E}[X_i \mid \mathcal{E}_i, X_i \geq t_i] \Pr[\mathcal{E}_i, X_i \geq t_i], \end{aligned}$$

where in the last inequality we used the non-negativity of X_i . The lemma follows by combining the above lower bound on $\mathbb{E}[\max_i \{X_i \cdot \mathbb{1}_{(X_i \geq t_i)}\}]$ with (4.2). \square

Lemma 17 immediately implies the following corollary, already known from [26].

Corollary 1 (mhr, independent). *In every single item setting with n bidders whose values are distributed according to independent (possibly non-identical) distributions that satisfy the monotone hazard rate condition, the Vickrey auction with monopoly reserves is a $(1/e, 1/2)$ -approximation.*

Proof. It is known from [46] that, if \mathbf{m} is the vector of monopoly reserve prices, then $\text{Vic}_{\mathbf{m}}$ (the Vickrey auction with monopoly reserves) is a $1/2$ -approximation to the optimal revenue. The welfare guarantee follows from Lemma 17 and the following fact from [2]: if v is drawn from a monotone hazard rate distribution, then $\Pr[v \geq \phi^{-1}(0)] \geq 1/e$. \square

Unfortunately, as discussed in Section 4.1, the Vickrey auction with monopoly reserve prices may be arbitrarily inefficient when we allow for regular distributions. In particular, we cannot employ Lemma 17 directly as the probability of any single bidder being above her monopoly reserve may be arbitrarily small. To fix this, we recall a lemma for regular distributions from [9]. For a single bidder setting, this lemma guarantees that there is always a reserve price r (which generally needs to be smaller than the monopoly reserve) that achieves a constant factor of the optimal revenue, while at the same time is smaller than the bidder's value with constant probability.

Lemma 18 ([9]). *Let F be a regular distribution, and let $R_F(x) = x \cdot F^{-1}(1 - x)$, for all $x \in [0, 1]$,² be the revenue curve in quantile space. Then, for all $0 < \tilde{q} \leq q \leq p < 1$,*

$$R_F(\tilde{q}) \leq \frac{1}{1-p} R_F(q).$$

If we try to use Lemma 18 to generalize Corollary 1 to regular distributions, we run into an additional difficulty. Indeed, if we lower the bidders' reserve prices to some vector $\mathbf{r} \leq \mathbf{m}$ below their monopoly reserves and run $\text{Vic}_{\mathbf{r}}$, the bidders will start contributing negative virtual values to the expected virtual welfare of the auction (i.e. its expected revenue). So we need to control the absolute value of the overall negative contribution to the expected virtual social welfare. This is not straightforward and is established in the following lemma, which alongside our main result is one of the main contributions of the work in this chapter.

Before providing its proof, it is worth noting that the obvious approach of decomposing the auction's virtual welfare into every bidder's contribution (using the law of total expectation) and then comparing each bidder's contribution under different reserve prices poses technical challenges. In particular, the terms of the decomposition cannot be directly compared as each of these terms depends on the probabilistic experiment that determines the winner of the auction, and this experiment depends on the reserves in ways that make it hard to find a useful coupling that enables term-by-term comparisons. Our technique tries to disentangle the contribution of each bidder to the virtual welfare of the auction from

²See the discussion in [9] for why F^{-1} is a well-defined function for a differentiable regular distribution.

the competition among the bidders, enabling us to first relate the revenue of $\text{Vic}_{\mathbf{r}}$ with the revenue of a *hybrid* auction, instead of $\text{Vic}_{\mathbf{m}}$ (for which we have good revenue guarantees from [46]). Our hybrid auction uses the tweaked reserves \mathbf{r} to truncate the bidders' values, but only gives the item to the winner of $\text{Vic}_{\mathbf{r}}$ if the winner also meets her monopoly reserve. Next we relate the revenue of our hybrid auction to $\text{Vic}_{\mathbf{m}}$. This is quite more challenging and involves a calculation that matches events where the hybrid auction makes no sale while $\text{Vic}_{\mathbf{m}}$ makes a sale to events where both auctions make a sale, establishing a factor 2 approximation. We expect our technique to find broader use in auction design.

Lemma 19 (Many-to-One Reduction—Revenue). *Consider a single item setting with n bidders whose values are distributed according to independent (possibly non-identical) regular distributions. Let also $\mathbf{r} = (r_1, \dots, r_n)$ be a vector of reserve prices such that, for all $i \in \{1, \dots, n\}$, $r_i \leq \phi_i^{-1}(0)$ (i.e. r_i is no larger than the monopoly reserve for bidder i) and $\text{Rev}(\text{Vic}_{r_i}(\{i\})) \geq (1-p) \cdot \text{Rev}(\text{Mye}(\{i\}))$, for some $p \in (0, 1)$. (That is, if bidder i were considered in isolation, then the Vickrey auction with reserve price r_i would achieve a $(1-p)$ -fraction of the optimal revenue.) Then it holds that $\text{Rev}(\text{Vic}_{\mathbf{r}}) \geq \frac{1-p}{4} \cdot \text{Rev}(\text{Mye})$.*

Proof. Let \mathcal{E}_i denote the event that i is the winner of the Vickrey auction with reserves \mathbf{r} , i.e. $i = \arg \max_j \{v_j \cdot \mathbb{I}_{(v_j \geq r_j)}\}$ ³ and $v_i \geq r_i$. Using Proposition 1 we can write $\text{Rev}(\text{Vic}_{\mathbf{r}})$ in terms of the bidders' virtual values as follows:

$$\begin{aligned} \text{Rev}(\text{Vic}_{\mathbf{r}}) &= \sum_{i=1}^n \mathbb{E}[\phi_i(v_i) \mid \mathcal{E}_i, \phi_i(v_i) \in [\phi_i(r_i), 0]] \Pr[\mathcal{E}_i, \phi_i(v_i) \in [\phi_i(r_i), 0]] \\ &\quad + \mathbb{E}[\phi_i(v_i) \mid \mathcal{E}_i, \phi_i(v_i) \geq 0] \Pr[\mathcal{E}_i, \phi_i(v_i) \geq 0]. \end{aligned} \quad (4.3)$$

In the course of the proof, we use the following inequalities:

$$\mathbb{E}[\phi_i(v_i) \mid \phi_i(v_i) \in [\phi_i(r_i), 0]] \leq \mathbb{E}[\phi_i(v_i) \mid \mathcal{E}_i, \phi_i(v_i) \in [\phi_i(r_i), 0]] \quad (\leq 0) \quad (4.4)$$

$$(0 \leq) \mathbb{E}[\phi_i(v_i) \mid \phi_i(v_i) \geq 0] \leq \mathbb{E}[\phi_i(v_i) \mid \mathcal{E}_i, \phi_i(v_i) \geq 0] \quad (4.5)$$

$$\begin{aligned} &|\mathbb{E}[\phi_i(v_i) \mid \phi_i(v_i) \in [\phi_i(r_i), 0]]| \cdot \Pr[\phi_i(v_i) \in [\phi_i(r_i), 0]] \leq \\ &p \cdot \mathbb{E}[\phi_i(v_i) \mid \phi_i(v_i) \geq 0] \cdot \Pr[\phi_i(v_i) \geq 0] \end{aligned} \quad (4.6)$$

Inequalities (4.4) and (4.5) follow from Lemma 16 when g is ϕ_i and $Y = \max_{j \neq i} \{v_j \cdot \mathbb{I}_{(v_j \geq r_j)}\}$. Inequality (4.6) involves a single bidder, and follows immediately from our assumption that $\text{Rev}(\text{Vic}_{r_i}(\{i\})) \geq (1-p) \cdot \text{Rev}(\text{Mye}(\{i\}))$ and noting that

$$\begin{aligned} \text{Rev}(\text{Vic}_{r_i}(\{i\})) &= \mathbb{E}[\phi_i(v_i) \mid \phi_i(v_i) \in [\phi_i(r_i), 0]] \cdot \Pr[\phi_i(v_i) \in [\phi_i(r_i), 0]] \\ &\quad + \mathbb{E}[\phi_i(v_i) \mid \phi_i(v_i) \geq 0] \cdot \Pr[\phi_i(v_i) \geq 0]; \end{aligned}$$

³Throughout the proof we assume that all maximizations have a unique maximizer. This is ok, since we consider continuous distributions so this happens with probability 1.

$$\text{Rev}(\text{Mye}(\{i\})) = \mathbb{E}[\phi_i(v_i) \mid \phi_i(v_i) \geq 0] \cdot \Pr[\phi_i(v_i) \geq 0].$$

Using (4.4), (4.5) and (4.6), we can bound the terms of the negative contribution to the expected revenue (4.3) as follows:

$$\begin{aligned} & \mathbb{E}[\phi_i(v_i) \mid \mathcal{E}_i, \phi_i(v_i) \in [\phi_i(r_i), 0]] \cdot \Pr[\mathcal{E}_i, \phi_i(v_i) \in [\phi_i(r_i), 0]] \\ \leq & \underbrace{\mathbb{E}[\phi_i(v_i) \mid \phi_i(v_i) \in [\phi_i(r_i), 0]]}_{\leq p \cdot \mathbb{E}[\phi_i(v_i) \mid \phi_i(v_i) \geq 0]} \underbrace{|\Pr[\phi_i(v_i) \in [\phi_i(r_i), 0]] \Pr[\mathcal{E}_i \mid \phi_i(v_i) \in [\phi_i(r_i), 0]]|}_{\leq \Pr[\phi_i(v_i) \geq 0] \cdot \Pr[\mathcal{E}_i \mid \phi_i(v_i) \geq 0]} \\ \leq & p \cdot \mathbb{E}[\phi_i(v_i) \mid \mathcal{E}_i, \phi_i(v_i) \geq 0] \cdot \Pr[\mathcal{E}_i, \phi_i(v_i) \geq 0] \end{aligned}$$

where for the first inequality we used (4.4) (and the fact that both sides of the inequality are non-positive), for the second inequality we used (4.6) and Lemma 15 taking g equal to ϕ_i , $X = v_i$ (conditioned on $X \geq r_i$), $Y = \max_{j \neq i} \{v_j \cdot \mathbb{I}_{(v_j \geq r_j)}\}$ and $c = 0$, and in the third inequality we used (4.5). We can now bound the revenue as follows:

$$\text{Rev}(\text{Vic}_{\mathbf{r}}) \geq (1 - p) \cdot \sum_{i=1}^n \mathbb{E}[\phi_i(v_i) \mid \mathcal{E}_i, \phi_i(v_i) \geq 0] \cdot \Pr[\mathcal{E}_i, \phi_i(v_i) \geq 0]. \quad (4.7)$$

To continue, we observe that the summation on the right-hand-side of (4.7) can be interpreted as the revenue of the following *hybrid* auction, \mathcal{H} , which lies “between” $\text{Vic}_{\mathbf{r}}$ and $\text{Vic}_{\mathbf{m}}$: \mathcal{H} truncates all bidders at their respective reserve prices r_i ; among the surviving bidders it identifies the larger bidder i^* as a potential winner, but only allocates the item to i^* if she clears her *monopoly* reserve m_{i^*} ; if this happens, i^* pays the maximum of her reserve price m_{i^*} and $\max_{j \neq i^*} \{v_j \cdot \mathbb{I}_{(v_j \geq r_j)}\}$. We can clearly lower bound the expected payment of bidder i in the hybrid auction by the following expression:

$$\int_{x=0}^{m_i} \Pr \left[\max_{j \neq i} \{v_j \cdot \mathbb{I}_{(v_j \geq r_j)}\} = x \right] \cdot m_i \cdot \Pr[v_i \geq m_i] \, dx.$$

Hence:

$$\mathbb{E}[\mathcal{R}_{\mathcal{H}}] \geq \sum_{i=1}^n \int_{x=0}^{m_i} \Pr \left[\max_{j \neq i} \{v_j \cdot \mathbb{I}_{(v_j \geq r_j)}\} = x \right] \cdot m_i \cdot \Pr[v_i \geq m_i] \, dx. \quad (4.8)$$

Next we compare the revenue of \mathcal{H} to that of the Vickrey auction with monopoly reserves $\text{Vic}_{\mathbf{m}}$. Our first observation is that whenever (i.e. for any value vector for which) \mathcal{H} sells to some bidder i , $\text{Vic}_{\mathbf{m}}$ also sells to the same bidder i ; moreover, the payment of bidder i in \mathcal{H} is at least as large as her payment in $\text{Vic}_{\mathbf{m}}$.⁴ So the contribution of bidder i to the revenue

⁴The reason for this is that \mathcal{H} uses lower reserves to truncate the bidders' values. So if i wins in \mathcal{H} her value is larger than her monopoly reserve as well as all other bidders' values truncated at the reserves \mathbf{r} . So her value must also be larger than the other bidders' values truncated at the (higher) monopoly reserves \mathbf{m} . By the same token, the second highest truncated value will be higher if truncation happens at \mathbf{r} than if it happens at \mathbf{m} .

from the event where she gets the item in both auctions is larger in the hybrid auction. This implies that the revenue in the event that both \mathcal{H} and $\text{Vic}_{\mathbf{m}}$ sell the item is larger in \mathcal{H} than $\text{Vic}_{\mathbf{m}}$. Let us call this event the *good event* \mathcal{G} . We have just argued that

$$\mathbb{E}[\mathcal{R}_{\mathcal{H}} \mid \mathcal{G}] \cdot \Pr[\mathcal{G}] \geq \mathbb{E}[\mathcal{R}_{\text{Vic}_{\mathbf{m}}} \mid \mathcal{G}] \cdot \Pr[\mathcal{G}]. \quad (4.9)$$

So it suffices to bound the revenue of $\text{Vic}_{\mathbf{m}}$ under the event that $\text{Vic}_{\mathbf{m}}$ sells to some bidder, but the hybrid auction does not sell to any bidder. Let us call this event the *bad event*, \mathcal{B} . We claim that the bad event is contained in the union of the following disjoint events:

$$B_i = \left\{ v_i \cdot \mathbb{I}_{(v_i \geq r_i)} = \max_j v_j \cdot \mathbb{I}_{(v_j \geq r_j)} \text{ and } v_i \leq m_i \right\}, \text{ for all } i.$$

Indeed, if the bad event happens it must be that the winner j^* of $\text{Vic}_{\mathbf{m}}$ does not satisfy $v_{j^*} \cdot \mathbb{I}_{(v_{j^*} \geq r_{j^*})} = \max_j \{v_j \cdot \mathbb{I}_{(v_j \geq r_j)}\}$. Suppose instead that $v_i \cdot \mathbb{I}_{(v_i \geq r_i)} = \max_j \{v_j \cdot \mathbb{I}_{(v_j \geq r_j)}\}$. For i not to be the winner in the hybrid auction it must be that $v_i \leq m_i$. Hence B_i is satisfied.

Now, in event B_i , the maximum possible revenue that any auction (and hence $\text{Vic}_{\mathbf{m}}$) could be making is $\max_{j \neq i} v_j \cdot \mathbb{I}_{(v_j \geq r_j)}$. Hence, the revenue of $\text{Vic}_{\mathbf{m}}$ from the event B_i can be upper bounded as:

$$\begin{aligned} \mathbb{E}[\mathcal{R}_{\text{Vic}_{\mathbf{m}}} \mid B_i] \cdot \Pr[B_i] &\leq \int_{x=0}^{m_i} \Pr \left[\max_{j \neq i} \{v_j \cdot \mathbb{I}_{(v_j \geq r_j)}\} = x \right] \cdot x \cdot \Pr[x \leq v_i \leq m_i] \, dx \\ &\leq \int_{x=0}^{m_i} \Pr \left[\max_{j \neq i} \{v_j \cdot \mathbb{I}_{(v_j \geq r_j)}\} = x \right] \cdot x \cdot \Pr[v_i \geq x] \, dx \\ &\leq \int_{x=0}^{m_i} \Pr \left[\max_{j \neq i} \{v_j \cdot \mathbb{I}_{(v_j \geq r_j)}\} = x \right] \cdot m_i \cdot \Pr[v_i \geq m_i] \, dx \end{aligned} \quad (4.10)$$

where the last inequality follows from the definition of the monopoly reserve m_i .

Hence, the revenue of $\text{Vic}_{\mathbf{m}}$ from the bad event \mathcal{B} can be upper bounded as:

$$\mathbb{E}[\mathcal{R}_{\text{Vic}_{\mathbf{m}}} \mid \mathcal{B}] \cdot \Pr[\mathcal{B}] \leq \sum_{i=1}^n \mathbb{E}[\mathcal{R}_{\text{Vic}_{\mathbf{m}}} \mid B_i] \cdot \Pr[B_i] \leq \mathbb{E}[\mathcal{R}_{\mathcal{H}}], \quad (4.11)$$

where for the first inequality we used that $\mathcal{B} \subseteq \cup_i B_i$, and for the second inequality we combined (4.10) and (4.8). Combining (4.11) and (4.9) we obtain:

$$\text{Rev}(\mathcal{H}) \geq \frac{1}{2} \cdot \text{Rev}(\text{Vic}_{\mathbf{m}}). \quad (4.12)$$

The lemma follows by combining (4.7), (4.12) and noticing that the revenue of $\text{Vic}_{\mathbf{m}}$ is known by [46] to be a 1/2-approximation to the optimal revenue, i.e. $\text{Rev}(\text{Vic}_{\mathbf{m}}) \geq \frac{1}{2} \cdot \text{Rev}(\text{Mye})$. \square

We are now ready to prove our main theorem:

Theorem 14 (Main). *For every single item setting with n bidders whose values are distributed according to independent (possibly non-identical) regular distributions, and any $p \in [0, 1]$, there is a vector of reserve prices $\mathbf{r} = (r_1, \dots, r_n)$ such that $\text{Vic}_{\mathbf{r}}$ is a $(p, (1-p)/4)$ -approximation.*

Proof. We argue that, for all i , there exists a price r_i such that the following are satisfied:

$$\Pr[v_i \geq r_i] \geq p; \text{ and}$$

$$\text{Rev}(\text{Vic}_{r_i}(\{i\})) \geq (1-p) \cdot \text{Rev}(\text{Mye}(\{i\})).$$

Indeed, we distinguish two cases. If $1 - F(\phi_i^{-1}(0)) \geq p$, we take $r_i = \phi_i^{-1}(0)$ and the above are satisfied automatically. Otherwise, the existence of a reserve with the above properties is implied by Lemma 18. Given reserves r_1, \dots, r_n as above, the theorem follows immediately from Lemmas 17 and 19. \square

Picking $p = 1/5$ we obtain a $(1/5, 1/5)$ -approximate auction for regular distributions.

Corollary 2 (regular, independent). *For every single item setting with n bidders whose values are distributed according to independent (possibly non-identical) regular distributions, there exist reserve prices \mathbf{r} such that $\text{Vic}_{\mathbf{r}}$ achieves a $(1/5, 1/5)$ -approximation.*

4.3 The non-regular, i.i.d. case

In this section we show that the Vickrey auction with an anonymous reserve price achieves a constant factor approximation to both objectives for general distributions, when the bidders' values are distributed independently and identically. We will follow the approach of [14], which makes use of *prophet inequalities* [72] to show that this auction achieves a 1/2-approximation to the optimal revenue.

We first describe prophet inequalities. Imagine a gambler facing a series of n games in a casino, one on each of n days. Game i has a prize associated with it, whose value is distributed according to some distribution F_i . The distributions of the prize values are known to the gambler in advance, but their exact realization is not known in advance, and neither is the order of the games. On day i a game is chosen by an adversary trying to minimize the gambler's profit, and its prize value is drawn from the corresponding distribution; the gambler needs to decide whether to pick the prize and leave the casino, or ignore it and keep playing. Clearly the gambler's optimal strategy can be computed using backwards induction; on the other hand, there exists a simple *threshold* strategy that guarantees the gambler at least half of the expected value of the maximum prize. A threshold strategy is a single value t , such that the gambler accepts the first prize i with $v_i \geq t$; the proof of the following theorem can be found in [72, 44].

Theorem 15. *There exists a threshold t such that, independently of the order the games are played, the expected prize of the gambler is at least half of the expected value of the maximum prize, and the probability that the gambler receives a prize is exactly 1/2.*

In [14] they leverage this theorem to show that the Vickrey auction with an anonymous reserve price achieves at least half of the optimal revenue. We can easily extend this to show a guarantee for both social welfare and revenue.

Theorem 16. *In every single item setting with n bidders whose values are drawn independently from the same (possibly non-regular) distribution, a Vickrey auction with an anonymous reserve price achieves a $1/2$ -approximation to both the optimal revenue and welfare.*

Proof. For the sake of completeness we first sketch the proof for revenue. (For full details we refer the reader to [44].) Observe that the problem a revenue-optimizing auctioneer faces is similar to the gambler's problem described above, if prizes are taken to be the bidders' ironed virtual values (assuming that the gambler's strategy treats all values in every flat region of the ironed virtual valuation functions the same). Indeed, let t be the threshold that is guaranteed by Theorem 15, and pick the reserve price to be $p = \hat{\phi}^{-1}(t)$, where $\hat{\phi}$ denotes the ironed virtual valuation of the bidders. If there are multiple p 's mapped to t by $\hat{\phi}$ pick the smallest such p . Given this tie-breaking, observe that the Vickrey auction with reserve price p treats all flat regions in the ironed virtual valuation function the same; hence its revenue is equal to the expected ironed virtual value of the winner (prize picked), which by Theorem 15 is at least $1/2$ of the optimal expected ironed virtual surplus (expected maximum prize). Since the latter is an upper bound to the optimal revenue, the revenue of the Vickrey auction with reserve p is a $1/2$ -approximation to the optimal revenue. Moreover, Theorem 15 guarantees that a prize will be picked with probability exactly $1/2$, and so

$$\Pr \left[\max_i \{v_i\} \geq p \right] \geq 1/2 \geq \Pr \left[\max_i \{v_i\} \leq p \right]. \quad (4.13)$$

Note that the way we defined our tie-breaking rule is important for this to hold. Next we show that this auction achieves at least half of the optimal social welfare as well:

$$\begin{aligned} \mathbb{E} \left[\max_i \{v_i\} \right] &= \int_0^p x \cdot \Pr \left[\max_i \{v_i\} = x \right] dx + \int_p^\infty x \cdot \Pr \left[\max_i \{v_i\} = x \right] dx \\ &\leq p \cdot \int_0^p \Pr \left[\max_i \{v_i\} = x \right] dx + \int_p^\infty x \cdot \Pr \left[\max_i \{v_i\} = x \right] dx \\ &\stackrel{(4.13)}{\leq} p \cdot \int_p^\infty \Pr \left[\max_i \{v_i\} = x \right] dx + \int_p^\infty x \cdot \Pr \left[\max_i \{v_i\} = x \right] dx \\ &\leq \int_p^\infty x \cdot \Pr \left[\max_i \{v_i\} = x \right] dx + \int_p^\infty x \cdot \Pr \left[\max_i \{v_i\} = x \right] dx \\ &= 2 \cdot \mathbb{E} \left[\max_i \{v_i \cdot \mathbb{I}_{(v_i \geq p)}\} \right] \end{aligned}$$

□

Bibliography

- [1] Vineet Abhishek and Bruce E. Hajek. “Efficiency loss in revenue optimal auctions”. In: *CDC*. 2010, pp. 1082–1087.
- [2] Gagan Aggarwal, Gagan Goel, and Aranyak Mehta. “Efficiency of (revenue-)optimal mechanisms”. In: *ACM Conference on Electronic Commerce*. 2009, pp. 235–242.
- [3] Saeed Alaei. “Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers”. In: *FOCS*. 2011, pp. 512–521.
- [4] Mark Armstrong. “Optimal Multi-object Auctions”. In: *Review of Economic Studies* 67.3 (July 2000), pp. 455–81. URL: <http://ideas.repec.org/a/bla/restud/v67y2000i3p455-81.html>.
- [5] Sayan Bhattacharya et al. “Budget constrained auctions with heterogeneous items”. In: *STOC*. 2010, pp. 379–388.
- [6] Patrick Briest et al. “Pricing Randomized Allocations”. In: *SODA*. 2010, pp. 585–597.
- [7] J. Bulow and P. Klemperer. “Auctions versus Negotiations”. In: *American Economic Review* 86.1 (Mar. 1996), pp. 180–94. URL: <http://ideas.repec.org/a/aea/aecrev/v86y1996i1p180-94.html>.
- [8] Jeremy Bulow and John Roberts. “The Simple Economics of Optimal Auctions”. In: *Journal of Political Economy* 97.5 (Oct. 1989), pp. 1060–90. URL: <http://ideas.repec.org/a/ucp/jpolec/v97y1989i5p1060-90.html>.
- [9] Yang Cai and Constantinos Daskalakis. “Extreme-Value Theorems for Optimal Multi-dimensional Pricing”. In: *FOCS*. 2011, pp. 522–531.
- [10] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. “An algorithmic characterization of multi-dimensional mechanisms”. In: *STOC*. 2012, pp. 459–478.
- [11] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. “Optimal Multi-dimensional Mechanism Design: Reducing Revenue to Welfare Maximization”. In: *FOCS*. 2012, pp. 130–139.
- [12] Ioannis Caragiannis, Christos Kaklamanis, and Maria Kyropoulou. “Limitations of deterministic auction design for correlated bidders”. In: *Manuscript* (2013).
- [13] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. “Algorithmic pricing via virtual valuations”. In: *ACM Conference on Electronic Commerce*. 2007, pp. 243–251.

- [14] Shuchi Chawla et al. “Multi-parameter mechanism design and sequential posted pricing”. In: *STOC*. 2010, pp. 311–320.
- [15] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. “Multi-budgeted Matchings and Matroid Intersection via Dependent Rounding”. In: *SODA*. 2011, pp. 1080–1097.
- [16] Xi Chen and Xiaotie Deng. “Settling the Complexity of Two-Player Nash Equilibrium”. In: *FOCS*. 2006, pp. 261–272.
- [17] Xue Chen et al. “On the Approximation Ratio of k-Lookahead Auction”. In: *WINE*. 2011, pp. 61–71.
- [18] Edward H. Clarke. “Multipart pricing of public goods”. In: *Public Choice* 11 (1 1971). 10.1007/BF01726210, pp. 17–33. ISSN: 0048-5829. URL: <http://dx.doi.org/10.1007/BF01726210>.
- [19] Jacques Cremer and Richard P McLean. “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions”. In: *Econometrica* 56.6 (Nov. 1988), pp. 1247–57. URL: <http://ideas.repec.org/a/ecm/emetrp/v56y1988i6p1247-57.html>.
- [20] Jacques Cremer and Richard P McLean. “Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent”. In: *Econometrica* 53.2 (Mar. 1985), pp. 345–61. URL: <http://ideas.repec.org/a/ecm/emetrp/v53y1985i2p345-61.html>.
- [21] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. “The Complexity of Optimal Mechanism Design”. In: *CoRR* abs/1211.1703 (2012).
- [22] Constantinos Daskalakis, Ilias Diakonikolas, and Mihalis Yannakakis. “How Good is the Chord Algorithm?” In: *SODA*. 2010, pp. 978–991.
- [23] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. “The complexity of computing a Nash equilibrium”. In: *STOC*. 2006, pp. 71–78.
- [24] Constantinos Daskalakis and George Pierrakos. “Simple, Optimal and Efficient Auctions”. In: *WINE*. 2011, pp. 109–121.
- [25] Constantinos Daskalakis and S. Matthew Weinberg. “Symmetries and optimal multi-dimensional mechanism design”. In: *ACM Conference on Electronic Commerce*. 2012, pp. 370–387.
- [26] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. “Revenue maximization with a single sample”. In: *ACM Conference on Electronic Commerce*. 2010, pp. 129–138.
- [27] I. Diakonikolas. “Approximation of Multiobjective Optimization Problems”. PhD thesis. Columbia University, 2010.
- [28] I. Diakonikolas and M. Yannakakis. “Small Approximate Pareto Sets for Biobjective Shortest Paths and Other Problems”. In: *SIAM J. Comput.* 39 (2009), pp. 1340–1371.

- [29] Ilias Diakonikolas and Mihalis Yannakakis. “Succinct approximate convex pareto curves”. In: *SODA*. 2008, pp. 74–83.
- [30] Ilias Diakonikolas et al. “Efficiency-Revenue Trade-Offs in Auctions”. In: *ICALP (2)*. 2012, pp. 488–499.
- [31] Shahar Dobzinski. “An impossibility result for truthful combinatorial auctions with submodular valuations”. In: *STOC*. 2011, pp. 139–148.
- [32] Shahar Dobzinski, Hu Fu, and Robert D. Kleinberg. “Optimal auctions with correlated bidders are easy”. In: *STOC*. 2011, pp. 129–138.
- [33] Shahar Dobzinski and Jan Vondrák. “The computational complexity of truthfulness in combinatorial auctions”. In: *ACM Conference on Electronic Commerce*. 2012, pp. 405–422.
- [34] Shaddin Dughmi, Tim Roughgarden, and Mukund Sundararajan. “Revenue submodularity”. In: *ACM Conference on Electronic Commerce*. 2009, pp. 243–252.
- [35] Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. “From convex optimization to randomized mechanisms: toward optimal combinatorial auctions”. In: *STOC*. 2011, pp. 149–158.
- [36] M. Ehrgott. *Multicriteria optimization*. Springer-Verlag, 2005.
- [37] Lawrence C. Evans. “Partial differential equations and Monge-Kantorovich mass transfer (survey paper)”. In: *Current Developments in Mathematics, 1997, International Press*. 1999.
- [38] H. Galperin and A. Wigderson. “Succinct Representations of Graphs”. In: *Information and Control* 56.3 (1983), pp. 183–198.
- [39] Allan Gibbard. “Manipulation of Voting Schemes: A General Result”. In: *Econometrica* 41.4 (July 1973), pp. 587–601. URL: <http://ideas.repec.org/a/ecm/emetrp/v41y1973i4p587-601.html>.
- [40] Andrew V. Goldberg et al. “Competitive auctions”. In: *Games and Economic Behavior* 55.2 (2006), pp. 242–269.
- [41] Fabrizio Grandoni, R. Ravi, and Mohit Singh. “Iterative Rounding for Multi-Objective Optimization Problems”. In: *ESA*. 2009, pp. 95–106.
- [42] Fabrizio Grandoni et al. “Utilitarian Mechanism Design for Multi-Objective Optimization”. In: *SODA*. 2010, pp. 573–584.
- [43] Theodore Groves. “Incentives in Teams”. In: *Econometrica* 41.4 (July 1973), pp. 617–31. URL: <http://ideas.repec.org/a/ecm/emetrp/v41y1973i4p617-31.html>.
- [44] Jason Hartline. *Lectures on Approximation in Mechanism Design*. Lecture Notes. Northwestern University, 2010. URL: <http://users.eecs.northwestern.edu/~hartline/amd.pdf>.

- [45] Jason D. Hartline and Tim Roughgarden. “Optimal mechanism design and money burning”. In: *STOC*. 2008, pp. 75–84.
- [46] Jason D. Hartline and Tim Roughgarden. “Simple versus optimal mechanisms”. In: *ACM Conference on Electronic Commerce*. 2009, pp. 225–234.
- [47] Ed. J. Climacao. *Multicriteria Analysis*. Springer-Verlag, 1997.
- [48] Philippe Jehiel, Moritz Meyer-ter Vehn, and Benny Moldovanu. “Mixed bundling auctions”. In: *Journal of Economic Theory* 134.1 (May 2007), pp. 494–512. URL: <http://ideas.repec.org/a/eee/jetheo/v134y2007i1p494-512.html>.
- [49] Robert Kleinberg and S. Matthew Weinberg. “Matroid prophet inequalities”. In: *STOC*. 2012, pp. 123–136.
- [50] Elias Koutsoupias and Christos H. Papadimitriou. “Worst-case Equilibria”. In: *STACS*. 1999, pp. 404–413.
- [51] Anton Likhodedov and Tuomas Sandholm. “Mechanism for optimally trading off revenue and efficiency in multi-unit auctions”. In: *ACM Conference on Electronic Commerce*. 2004, pp. 268–269.
- [52] C. Liu and C. R. Traas. “On convexity of planar curves and its application in CAGD”. In: *Comput. Aided Geom. Des.* 14 (7 Sept. 1997).
- [53] R. Preston McAfee and Philip J. Reny. “Correlated Information and Mechanism Design”. In: *Econometrica* 60.2 (Mar. 1992), pp. 395–421. URL: <http://ideas.repec.org/a/ecm/emetrp/v60y1992i2p395-421.html>.
- [54] K. M. Miettinen. *Nonlinear Multiobjective Optimization*. Kluwer, 1999.
- [55] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. “Matching Is as Easy as Matrix Inversion”. In: *STOC*. 1987, pp. 345–354.
- [56] R. B. Myerson and M. A. Satterthwaite. “Efficient mechanisms for bilateral trading”. In: *Journal of Economic Theory* 29.2 (Apr. 1983), pp. 265–281.
- [57] Roger B. Myerson. “Optimal Auction Design”. In: *Mathematics of Operations Research* 6 (1981), pp. 58–73. DOI: 10.1287/moor.6.1.58.
- [58] Z. Neeman. “The effectiveness of English auctions”. In: *Games and Economic Behavior* 43.2 (May 2003), pp. 214–238.
- [59] Noam Nisan and Amir Ronen. “Algorithmic Mechanism Design (Extended Abstract)”. In: *STOC*. 1999, pp. 129–140.
- [60] Noam Nisan et al. *Algorithmic Game Theory*. New York, NY, USA: Cambridge University Press, 2007. ISBN: 0521872820.
- [61] C. H. Papadimitriou and M. Yannakakis. “A Note on Succinct Representations of Graphs”. In: *Information and Control* 71.3 (1986), pp. 181–185.
- [62] Christos H. Papadimitriou. “Algorithms, Games, and the Internet”. In: *ICALP*. 2001, pp. 1–3.

- [63] Christos H. Papadimitriou and George Pierrakos. “On optimal single-item auctions”. In: *STOC*. 2011, pp. 119–128.
- [64] Christos H. Papadimitriou, Michael Schapira, and Yaron Singer. “On the Hardness of Being Truthful”. In: *FOCS*. 2008, pp. 250–259.
- [65] Christos H. Papadimitriou and Mihalis Yannakakis. “On the Approximability of Trade-offs and Optimal Access of Web Sources”. In: *FOCS*. 2000, pp. 86–92.
- [66] Christos H. Papadimitriou and Mihalis Yannakakis. “Optimization, approximation, and complexity classes”. In: *Journal of Computer and System Sciences* 43.3 (1991), pp. 425–440. ISSN: 0022-0000. DOI: 10.1016/0022-0000(91)90023-X. URL: <http://www.sciencedirect.com/science/article/pii/002200009190023X>.
- [67] Amir Ronen. “On approximating optimal auctions”. In: *ACM Conference on Electronic Commerce*. 2001, pp. 11–17.
- [68] Amir Ronen and Amin Saberi. “On the Hardness of Optimal Auctions”. In: *FOCS*. 2002, pp. 396–405.
- [69] T. Roughgarden and M. Sundararajan. “Is Efficiency Expensive?” In: *3rd Workshop on Sponsored Search*. 2007.
- [70] Tim Roughgarden. “The price of anarchy is independent of the network topology”. In: *J. Comput. Syst. Sci.* 67.2 (2003), pp. 341–364.
- [71] Tim Roughgarden and Éva Tardos. “How bad is selfish routing?” In: *J. ACM* 49.2 (2002), pp. 236–259.
- [72] Ester Samuel-Cahn. “Comparison of Threshold Stop Rules and Maximum for Independent Nonnegative Random Variables”. English. In: *The Annals of Probability* 12.4 (1984), pp. 1213–1216. ISSN: 00911798. URL: <http://www.jstor.org/stable/2243359>.
- [73] Mark Allen Satterthwaite. “Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions”. In: *Journal of Economic Theory* 10.2 (Apr. 1975), pp. 187–217. URL: <http://ideas.repec.org/a/eee/jetheo/v10y1975i2p187-217.html>.
- [74] S. Vassilvitskii and M. Yannakakis. “Efficiently computing succinct trade-off curves”. In: *Theoretical Computer Science* 348 (2005), pp. 334–356.
- [75] William Vickrey. “Counterspeculation, Auctions, and Competitive Sealed Tenders”. In: *The Journal of Finance* 16.1 (1961), pp. 8–37. URL: <http://jmvidal.cse.sc.edu/library/vickrey61a.pdf>.