Profit Maximization in Mechanism Design

by

George Pierrakos

scholar of

Bodossaki Foundation

A thesis submitted in partial satisfaction
of the requirements for the degree of

Master of Science

in

Graduate Program in Logic, Algorithms and Computation

in the

Department of Mathematics

of the

UNIVERSITY OF ATHENS

Committee in charge:

Professor Elias Koutsoupias, Chair
Professor Stathis Zachos
Professor Aris Pagourtzis

Spring 2009

Abstract

Profit Maximization in Mechanism Design

by

George Pierrakos

Master of Science

in

Graduate Program in Logic, Algorithms and Computation

The present thesis focuses on a specific area of game theory, known as mechanism design. We review some of the most important results related to revenue maximization in auctions and mechanism design, i.e. we are interested in the case where the mechanism designer aims at maximizing his own profit –defined as the sum of the received payments– rather than the social welfare.

There are two lines of work related to profit maximization in mechanism design. The first and more traditional one, originating by economists and by Myerson's seminal paper, studies the problem when there is some prior knowledge of the distributions from which the bidders' valuations are drawn; it bears the name Bayesian Optimal Mechansim Design. The second one makes no assumption about the distribution of the bids, but rather adopts worst-case analysis –the dominating paradigm in computer science– and employs notions from online algorithms and competitive analysis to approach the problem; this line of work is much more recent and is –in its large part– due to computer scientists.

The structure of this thesis is as follows: in the first chapter we introduce some basic notions about game theory and mechanism design and we rigorously formulate the problem of profit maximization. In chapters two and three we review some of the results related to Bayesian and Worst-Case Optimal Mechanism Design respectively. Finally, in the last chapter we focus our attention to a specific auction that is conjectured to have a good worst-case performance. The problem of finding an elegant (combinatorial) proof of this conjecture remained open by the time this thesis was completed.

Contents

List of Figures

Acknowledgements

First and foremost I want to thank my advisor Elias Koutsoupias. Elias has been no less than a role model for me. As a teacher he gave me the chance of attending some of the most exciting and enjoyable lectures I have ever attended; the way he discusses each topic –ranging from introductory notions in discrete math to very involved topics in online algorithms– there is always something new to learn and take home from his classes. As an advisor, his taste for important (and hard) problems is unquestionable; his great insight and his ability to focus on the essence of a problem, while keeping things crisp and simple keep astonishing me. As a person, his belief in me and his support, in every way possible, has enabled me to realize my potential; for that I am indebted to him. I feel truly privileged to have been able to work with Elias and I look forward to a long lasting collaboration and friendship.

I also want to cordially thank the other two members of my thesis committee. Stathis Zachos, for if it had not been for him and his early classes in NTUA, theoretical computer science would probably have never attired my attention –as is the case with many of Stathis' students, now excelling abroad. Aris Pagourtzis, for mentoring me through my first research steps and for the nice times –usually late at night over leftovers of pizza– we spent discussing about problems. Thanks to both of them for providing CoReLab as a nice and friendly working environment.

Finally I would like to thank all my professors at MPLA for providing me with an excellent opportunity for studies; all the students who have been my friends and collaborators over the one year of my graduate studies spent in Athens; thanks to Christos Papadimitriou and Yaron Singer from Berkeley for stimulating discussions on the RSOP-conjecture; and of course many thanks to all the people in my friendly and family environment, and most of all my mother, for embracing me with their love and support.

Chapter 1

Introduction

1.1 Structure of this thesis

The focus of the present thesis is on mechanism design. Mechanism design is one of the traditional fields of game theory and economists have been studying mechanisms and auctions for the last half century. Only recently though -after the seminal paper of Nisan and Ronen [1] - has the field of mechanism design attracted much attention from the computer science community (along with other fields related to algorithmic game theory, such as computation of equilibria, quantifying the inefficiency of selfish behavior etc). Mechanism design can be viewed as reverse engineering: given a desirable outcome, what is the right way to design a game, so as to reach this outcome. The challenge here is to take into account the selfish behavior of the agents (the players who take part in this game) and try to align their incentives with the desirable outcome.

But what is a desirable outcome? Ideally we would want to be able to design mechanisms for any desirable outcome. Unfortunately very strong impossibility results by Arrow suggest that this is not in general possible, if one is aiming for **truthful** mechanisms, i.e. for mechanisms where each agent is playing honestly; in fact the whole field of Mechanism Design is an attempt to escape from this impossibility result.

Indeed there are some goals one can achieve with truthful mechanisms, by charging the agents for their actions, i.e. by introducing payments as a way to force agents to say the truth. The present thesis aims at presenting some results related to the design of mechanisms (auctions) when the goal is to maximize the profit of the mechanism designer (auctioneer), i.e. to maximize the sum of the agents' payments; this is often referred to as *Optimal Mechanism Design*. Although this may seem as a very natural objective, actually most of the work in the field aims at mechanisms that achieve the maximum "social surplus"; the idea here is that the mechanism designer is not interested in making a profit, rather than achieving some socially desirable outcome. The payments in this model are only a mean to extract the truth out of each agent; in the next section we will briefly discuss the general paradigm in this area, namely the VCG mechanism.

The general outline of this thesis is as follows:

- in the remaining of the introduction we will introduce some basic notions and definitions related to game theory and mechanism design. We will mention some impossibility results for mechanism design and the VCG mechanism.
- in Chapter 2 we will present Myerson's approach to optimal mechanism design. This approach assumes some prior knowledge of the market, in the form of distributions from which the agents' valuations are drawn (hence the name Bayesian Optimal Mechanism Design). The chapter is mostly based on Myerson's seminal paper [2] which dates back in 1981.
- in Chapter 3 we present a much more CS-like approach to the problem of optimal mechanism design, namely worst-case analysis. This is a much more recent line of work, with many results during the last decade (see results by J. Hartline, A. Karlin, A. Goldberg et. al.).
- in Chapter 4 we focus our attention to a specific worst-case competitive auction and we present a proof that it actually achieves a constant fraction of the maximum achievable profit.

Disclaimer: This thesis is not in any way a self contained article for mechanism design and game theory. Most notions are introduced when they are needed, with the exceptions of some fundamental notions which are introduced in the rest of this chapter. For readers without any game-theoretic background a good reference is [3].

1.2 Mechanism Design Basics

We start off by motivating the whole Mechanism Design concept with the simple example of auctioning off a single good.

Assume that we have n bidders (also called players or agents), each of whom desires a single, indivisible good. Each of the bidders has his own private valuation v_i for the good, which is private in the sense that the auctioneer and the other players have no information about it. Each bidder submits a sealed bid b_i to the auctioneer (or mechanism designer), who decides who gets the item and at what price p. Notice that the submitted bid b_i needs not be the same as the bidder's valuation v_i. However our goal is to design truthful mechanisms, i.e. mechanisms such that $b_i = v_i$ for all bidders i. Our only assumption about the bidders is that they are rational and selfish and bid in order to maximize their utility, defined as $v_i - p$ if they get the good and 0 otherwise. Finally our goal is to design a social optimal mechanism, i.e. a mechanism that allocates the good to the person who values it the most.

A first approach would be to give the good to the person who values it the most, i.e. to the highest bidder, and charge him his bid. This would result in 0 utility for all players. Notice now that the highest bidder has an incentive to lie: instead of bidding his true value, he needs only bid slightly more than the second highest bidder. This would result in him still getting the object and paying much less, i.e. having a bigger utility. Since we assume that players bid with the sole goal of maximizing their utility we immediately get that this auction is not truthful.

There is however a very simple way to get past this problem. Allocate the item to the person who values it most but only charge him the second largest bid. This auction does not only achieve the optimal social outcome, but is also truthful since no agent has an incentive to lie about his bid: the highest bidder does not want to raise his bid since he is winning anyway and he also does not want to lower his bid, since this does not affect his payment and may only result into him losing the item (i.e. having zero instead of positive utility). The rest of the bidders on the other hand will always lose the item, as long as they bid below the highest bid, and, if one of them decides to lie and bid more than the highest bidder, then he is going win the item but nonetheless receive negative utility.

What we just saw was *Vickrey's second price auction.* We can formalize all the above in the following simple protocol:

1. Each bidder i submits his sealed bid to the auctioneer.

2. The auctioneer gives the item to the highest bidder.

3. The winner pays the second highest bid.

Notice that the above protocol implements essentially the traditional English (or ascending) auction: the auctioneer announces a price and the bidders declare their willingness to pay this price for the good. The auctioneer then raises the price and some bidders withdraw from the auction until, eventually, only one bidder is left: the highest bidder. So the highest bidder gets the good and has to pay the price at which everyone else dropped out, i.e. the bid of the second-highest bidder (we assume that the increments in the price are relatively small).

We are now ready to give a definition of a mechanism. Since in the present thesis the only mechanisms we will study are auctions we will not use the fully general definition (as it appears in [3] for example) but rather a tailored definition that matches our needs for this thesis. Let us introduce some notation first.

Notation. We shall denote by A the set of possible outcomes of the game; for example, in the case of a single-item auction, A is the set of all bidders, i.e. all possible winners of the item. In the case of a multi-unit auction, A is the set of all possible allocations of items to bidders. Each bidder i has a preference over outcomes in A which is modeled through his valuation $v_i : A \to \mathbb{R}$. We denote the set of possible valuations of bidder i by $V_i \subseteq \mathbb{R}^A$.

Definition 1.1. *A (direct revelation) mechanism is a social choice function $f : V_1 \times \ldots \times V_n \to A$ and a vector of payment functions $p_1, \ldots, p_n$, where $p_i : V_1 \times \ldots \times V_n \to \mathbb{R}$ is the amount that player i pays.*

The term "direct revelation" in the above definition means that we assume that the bidders reveal their true valuations directly to the auctioneer, i.e. they do not lie. This assumption is without loss of generality since it can be shown that for any arbitrary mechanism that implements some social function f in dominant strategies, there exists a truthful one that implements f, where by truthful we mean the following.

Definition 1.2. *A mechanism $(f, p_1, \ldots, p_n)$ is called incentive compatible or truthful if for every player i, every $v_1 \in V_1, \ldots, v_n \in V_n$ and every $v_i' \in V_i$, we have*

$$v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) \geq v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i})$$

One may naturally wonder why truthfulness is so important. After all, in the example we presented above, we had a sealed-bid auction, which means that the bidders do not know what is the profitable deviation, i.e. the right bid (other than their true valuation) that will maximize their profit since this depends on the (private) valuations of the other players. However, one can invoke some game-theoretic jargon to argue for the opposite: thinking in terms of mixed strategies, a player that takes part in a non-trutfful mechanism may have better expected payoff by randomizing over his strategies, rather than declaring his true valuation.

There are however more obvious reasons, why one needs to consider truthful mechanisms. Here are two good ones:

1. In a non-truthful mechanism, knowledge of other players' bids can be very useful. Hence agents are motivated to expend resources to gain information about other players' bids and then more (probably computational) resources to compute their optimal strategy. Truthfulness takes away such considerations: all bidders have to do is submit their true valuation and they will get the best they can out of the mechanism. In some sense truthful mechanisms are a form of guarantee for the agents that no one is going to take away from them what is rightfully theirs; and all they have to do for that is be honest.

2. From the auctioneer's point of view, running an untruthful mechanism is like trying to solve an optimization problem without knowing the inputs. In fact, in some sense this is the biggest challenge in mechanism design and one may as well be surprised by the mere fact that there exist truthful mechanisms: these mechanisms can be seen as functions that compute the right output, without having *a priori* any guarantees on the quality of the input. It is part of the mechanism to enforce truthful behavior and extract the right input out of all possible ones.

Having pinned down the importance of truthful mechanisms, we end this section by giving an impossibility result.

Theorem 1.3. [Gibbard-Satterthwaite] *Let f be an incentive compatible social choice function onto A, where $|A| \geq 3$; then f is a dictatorship. By dictatorship we mean that there exists some agent i such that the outcome is always the one maximizing his own valuation, for any combination of other agents' valuations.*

The interested reader may find a proof of the above theorem in [3].

The above result is devastating: it looks like it eliminates all hopes of designing truthful mechanisms since it is really universal. Nonetheless its generality is also its main weakness. One may hope to design truthful mechanisms for some specific social choice functions; this is indeed the case for -perhaps- the most natural social choice function: optimizing the social welfare.

Definition 1.4. *The social welfare of an alternative $a \in A$ is the sum of the valuations of all players for this alternative $\sum_i v_i(a)$.*

In the next section we present the VCG-mechanism which manages to accomplish the goal of maximum social welfare.

1.3 Vickrey-Clarke-Groves Mechanisms

The VCG-mechanism is the classical mechanism used for combinatorial auctions, i.e. auctions where there are many items and each user may have arbitrarily complex preferences over all possible subsets of items. Nonetheless one can define the VCG-mechanism to apply to the more general mechanism design setting discussed above as follows:

Definition 1.5. *A mechanism $(f, p_1, \ldots, p_n)$ is called a VCG-mechanism if:*

- $f(v_1, \ldots, v_n) \in \arg\max_{a \in A} \sum_i v_i(a)$*, i.e. f maximizes the social welfare*
- *there exist functions $h_1, \ldots, h_n$, where $h_i : V_{-i} \to \mathbb{R}$ such that for all $v_1 \in V_1, \ldots, v_n \in V_n$:* $p_i(v_1, \ldots, v_n) = h_i(v_{-i}) - \sum_{i \neq j} v_i(f(v_1, \ldots, v_n))$

It is obvious that VCG-mechanisms indeed reach the outcome of maximum social welfare. It is also easy to prove that they are incentive compatible. The intuition is as follows: consider the utility $u_i(v_1, \ldots, v_n)$ of player i. This is $v_i(v_1, \ldots, v_n) - p_i(v_1, \ldots, v_n) = -h_i(v_{-i}) + \sum_i v_i(f(v_1, \ldots, v_n))$. Notice how the social welfare appears in the utility of every player; this has the direct implication that a player's own profit is aligned with the socially optimal outcome. The additional term h_i does not depend on an agent's own valuation and has no strategic implications; it could thus be set to zero and we would have a valid VCG-mechanism. Nonetheless, usually we want our mechanisms to have the following additional two properties.

Definition 1.6. *We say that*

- *a mechanism is individually rational if players always receive non-negative utilities, namely if* $p_i(v_1, \ldots, v_n) \leq v_i(v_1, \ldots, v_n)$ *for all* $v_1, \ldots, v_n$. *These mechanisms are also said to have voluntary participation, since in some sense the players are not charged for taking part in the mechanism.*

- *a mechanism has no positive transfers if no player is ever paid money, namely if* $p_i(v_1, \ldots, v_n) \geq 0$ *for all* $v_1, \ldots, v_n$.

Imposing the above two natural constraints to our mechanisms we get that a good choice of h_i is $h_i(v_{-i}) = \max_{b \in A} \sum_{j \neq i} v_j(b)$. The reader may verify that the above two properties hold for this choice of h_i.

Notice finally that the payment of agent i in the VCG-mechanism is $\max_{b \in A} \sum_{j \neq i} v_j(b) - \sum_{j \neq i} v_j(f(v_1, \ldots, v_n))$, which is like charging the i-th player according to the marginal cost of his presence to the happiness of the rest of the players.

1.4 What other social choice functions can we implement?

So far in this section we described a mechanism that implements the social function that wants to maximize social welfare; we also saw an impossibility result. The next natural question is, what else can we implement?

Maybe the next most natural candidate for a social choice function, would be the allocation that maximizes the revenue of the auctioneer. In the mechanisms we have seen so far, the payments were only a way to elicit the truth out of the bidders. On the contrary, in this thesis we are interested in mechanisms that maximize the profit of the auctioneer, i.e. the sum of the payments of all agents.

Definition 1.7. *The profit of an alternative $a \in A$ is the sum of the payments of all players for this alternative $\sum_i p_i(a)$.*

We refer to the design of such profit-maximizing mechanisms as optimal mechanism design and to the corresponding mechanisms as optimal mechanisms. Chapters 2-4 cover this kind of mechanisms.

Another natural objective for example is maximizing the sum of the players' utilities or, as it is called, the residual surplus.

Definition 1.8. *The residual surplus of an alternative $a \in A$ is the sum of the valuations of all players for this alternative minus the payments $\sum_i v_i(a) - p_i(a)$.*

The motivation for such an objective is the following: again we consider payments as a way to enforce truthfulness but, instead of having the money transfered to the auctioneer (and thus in some sense remain in the system) we consider this money to be "burnt". For example, the payments in the context of networks could be some sort of service degradation, or some computational payment. The interested reader may refer to [4] for more information regarding this objective.

Chapter 2

Bayesian Optimal Mechanism Design

In the present chapter we review Myerson's optimal auction. We first give a characterization of truthful mechanisms for single parameter agents, we describe Myerson's optimal mechanism and then give a couple of examples that clarify the form of this mechanism for some particular settings.

Let us start off by a motivating example. Assume we want to maximize profit in the simple setting where we have one good for sale and two bidders. One approach is to use VCG, which for this case is simply Vickrey's second price auction. The first thing to notice is that if we assume nothing about the bids then this mechanism can be arbitrarily bad, in the sense that it guarantees a profit equal to $\min(b_1, b_2)$, which can be arbitrarily smaller than $\max(b_1, b_2)$.

This leads us to the conclusion that we need to have some sort of assumption about the distribution of the bids. So for the rest of this section we will assume that the bids are drawn independently at random from some distribution each. We let $F_i(z)$ be the cumulative distribution function and $f_i(z)$ be the probability density function for each bid b_i. Our goal now is to design a mechanism that gets the maximum expected profit, given the prior (distribution) of the bids.

Example 2.1. Consider the setting where we have one item to sell and two bidders. The bidders' valuations are drawn independently and uniformly at random from $[0,1]$. It is easy to see that if we run VCG then our expected profit is $\mathbb{E}[\min(b_1,b_2)]$. This can be easily computed to be $\mathbb{E}[\min(b_1,b_2)] = \int_0^\infty \Pr[\min(b_1,b_2) > x]dx = \int_0^1 (1-x)^2 dx = 1/3$, where I used the fact that $\mathbb{E}[X] = \int_0^\infty \Pr[X \geq z]dz$ and the independence of the bids.

It turns out that running the VCG auction does not yield the maximum expected profit, even for this simple example. Indeed let us define a new auction, called the Vickrey auction with reservation price, and see how this outperforms VCG. We shall see later that this is in fact the optimal auction for this setting.

Definition 2.2. [Vickrey auction with reservation price r**]** *The Vickrey auction with reservation price VA_r sells the item if any bidder bids above r. It charges the winner the maximum of the second highest bid and r.*

Intuitively VA_r is like Vikcrey's second price auction, only with one extra bid submitted on behalf of the auctioneer. If no one bids higher than the auctioneer then he gets to keep the item.

Consider now again the case of two bidders and a single item, where the bids are i.i.d. random variables with a uniform distribution in $[0,1]$. Then we claim that the optimal auction is $VA_{1/2}$, i.e. the Vickrey auction with reservation price $1/2$. Indeed, one can verify -by analyzing separately the cases where both bids are below or above $1/2$ and the case where $1/2$ is the second largest bid- that this auction yields an expected profit of $5/12$ and thus outperforms VCG. The reason why this is the optimal auction will become clear later.

2.1 Characterizing Truthful Mechanisms for Single Parameter Agents

In the present section we shall prove a theorem that allows us to fully characterize truthful mechanisms for single parameter agents. By single parameter agents we refer to settings where each agent's preferences can be summarized into a single parameter, say the valuation of each bidder for the item being auctioned. Again we will denote the bidders' valuations by v_i and their bids by b_i; in truthful mechanisms we expect to have $b_i = v_i$. Our mechanisms compute an outcome which consists of an allocation vector $\mathbf{x} = (x_1, \ldots, x_n)$ with $x_i \in \{0, 1\}$and a price vector $\mathbf{p} = (p_1, \ldots, p_n)$. As a result each bidder receives a utility of $u_i = v_i x_i - p_i$. We will also consider randomized mechanisms where $x_i \in [0, 1]$ is the probability of bidder i being allocated the good and the p_i and u_i are expected prices and utilities respectively.

In order to be able to model more general settings than just a single item auction we introduce an additional cost term $c(\mathbf{x})$ which must be paid by the mechanism. Thus the quantities of social welfare and profit that we introduced before, become now $\sum_i v_i x_i - c(\mathbf{x})$ and $\sum_i p_i - c(\mathbf{x})$ respectively.

Example 2.3. In a single item auction we have

$$c(\mathbf{x}) = \begin{cases} 0 & : \sum_i x_i \leq 1 \\ \infty & : \text{otherwise} \end{cases}$$

whereas in a single-minded combinatorial auction, i.e. an auction where we have m items for sale but each bidder is interested in only one specific bundle of them, we have

$$c(\mathbf{x}) = \begin{cases} 0 & : \forall i, j, S_i \cap S_j \neq 0 \rightarrow x_i x_j = 0 \\ \infty & : \text{otherwise} \end{cases}$$

We now state and prove the main theorem of this section which consists of an exact characterization of all truthful mechanisms for single parameter agents. Recall that in our current notation a mechanism is truthful if and only if for all i, v_i, b_i and $\mathbf{b}_{-i}$ we have

$$u_i(v_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i}),$$

i.e. all bidders are better off by bidding their true valuation v_i.

Theorem 2.4. *A mechanism is truthful if and only if, for any agent i and bids of other agents $\mathbf{b}_{-i}$ fixed,* *1.* $x_i(b_i)$ *is monotone non-decreasing.*

2. $p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z)dz$.

Proof. First notice that what the above theorem tells us is that for a fixed allocation rule $\mathbf{x}(\cdot)$ the prices are uniquely determined. So in order to specify a truthful mechanism all we have to do is give a monotone allocation rule.

Let us first prove that if a mechanism satisfies the properties of the theorem then it is truthful. The best way to do that is by picture (see Figure 2.1). Let z_1 and z_2, with $z_1 < z_2$ be two possible bids. Without loss of generality assume that $b_i = z_1$ and $v_i = z_2$. The proof for the other case is similar. The top diagram shows the graph of the allocation probability $x_i(\cdot)$ which is indeed monotone non-decreasing. Recall that the utility of player i is $u_i(b_i) = v_i x_i(b_i) - p_i(b_i)$ as a function of his own bid. The three diagrams on the left analyze each term of the utility and the overall sum for the case where the bidder reports his true value $v_i = z_2$, while the three right diagrams analyze the same terms for the case where the bidder reports the false value $b_i = z_1$. These terms can be seen as the shaded areas in the diagrams, thanks to our second condition that determines payments to be $p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z)dz$; it is easy to verify now that the area for bidding v_i is larger than for bidding b_i, because of the monotonicity of x_i, which proves the claim (see bottom figure too).

To prove the inverse statement we start from the definition of truthfulness: a mechanism is truthful if a bidder maximizes his utility by bidding his true valuation, namely if for all i, v_i, b_i and $\mathbf{b}_{-i}$ we have $u_i(v_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i})$, which means that:

$$\forall v_i, b_i \ : \ v_i x_i(v_i) - p_i(v_i) \geq v_i x_i(b_i) - p_i(b_i)$$

Consider now again z_1, z_2 with $z_1 < z_2$ and do the following: first set $v_i = z_1$ and $b_i = z_2$ and then $v_i = z_2$ and $b_i = z_1$. We then get the following two inequalities:

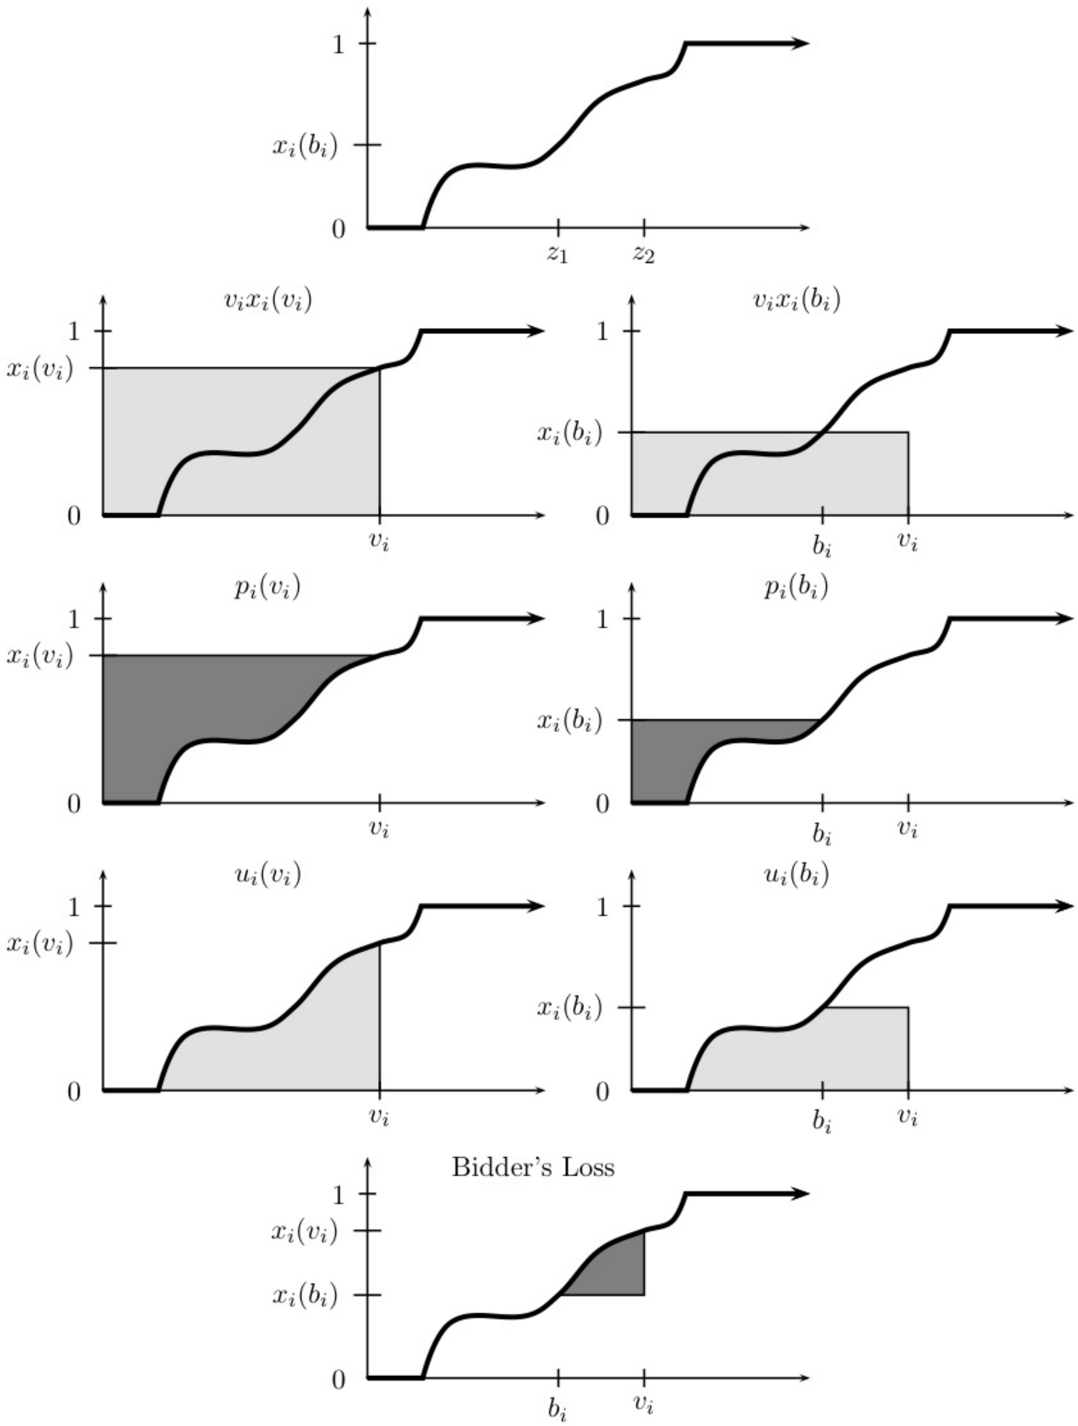

Figure 2.1. Graphical representation of payments and utilities.

$$v_i = z_1, b_i = z_2 \ : \ z_1 x_i(z_1) - p_i(z_1) \geq z_1 x_i(z_2) - p_i(z_2)$$

$$v_i = z_2, b_i = z_1 \ : \ z_2 x_i(z_2) - p_i(z_2) \geq z_2 x_i(z_1) - p_i(z_1)$$

Adding these inequalities and doing some canceling and rearranging, we get that $(z_2 - z_1)(x_i(z_2) - x_i(z_1)) \geq 0$, which implies that the allocation rule has to be monotone non-decreasing.

For the second condition, consider the utility as a function of the bid b_i: $u_i(b_i) = v_i x_i(b_i) - p_i(b_i)$. Having the utility being maximized $b_i = v_i$ means that the derivative $u_i'(z)$ becomes zero for $b_i = v_i$, i.e. that $v_i x_i'(v_i) - p_i'(v_i) = 0$; to simplify the expression, we replace v_i with z and we get $z x_i'(z) = p_i'(z)$. Integrating both sides from 0 to b_i and doing some rearranging we get:

$$p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z) dz + p_i(0)$$

Remember now that we want our mechanisms to have the individual-rationality and the non-positive transfers property (see Chapter 1); the first implies that $p_i(b_i) \leq b_i \Rightarrow p_i(0) \leq 0$ and the second one implies $p_i(0) \geq 0$; together they yield that $p_i(0) = 0$ and the proof is complete. □

The above theorem has a very interesting interpretation for deterministic mechanisms. In deterministic mechanisms, the allocation vector takes values in $\{0, 1\}$, so any monotone non-decreasing function has to look like the one in Figure 2.2. It is easy to see now that the second part of the characterization can be rephrased to say that, "in any deterministic, truthful mechanism, any winning agent has to pay his minimum winning bid, i.e. $\inf_z\{z : \ x_i(z) = 1\}$".

The above theorem has another interesting consequence: in order to define a truthful mechanism, all we have to do is define a monotone non-decreasing allocation rule and the payments are immediately dictated by the second part of the theorem.

Example 2.5. Consider the goal of maximizing the social surplus. We already saw a mechanism that achieves that, namely VCG. It is easy to see that the surplus maximizing allocation rule is

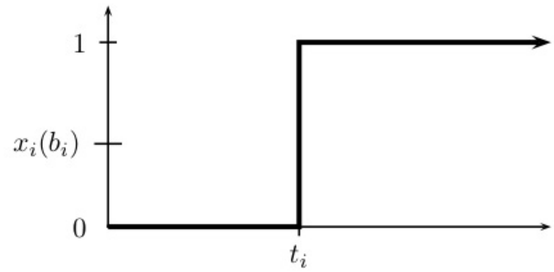

Figure 2.2. The allocation function of any deterministic truthful mechanism.

indeed monotone, since it is linear in all valuations, and so it is truthful. Furthermore the payments are exactly the ones dictated by Theorem 2.4.

2.2 Myerson's Optimal Mechanism

Having Theorem 2.4 in mind, we are ready to describe Myerson's (simple) Optimal Mechanism.

Let 's start by introducing the notion of virtual valuations and virtual surplus.

Definition 2.6. *The virtual valuation of agent i with valuation v_i is*

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

where $F_i(z)$ is the cumulative distribution function and $f_i(z)$ is the probability density function. The virtual surplus is defined to be $\sum_i \phi_i(v_i)x_i - c(\mathbf{x})$.

The basic theorem here is the following:

Theorem 2.7. *The expected profit of a mechanism is equal to its expected virtual surplus.*

which follows easily from the following lemma, using linearity of expectation and the independence of the agents' valuations.

Lemma 2.8. *The expected payment of a bidder satisfies:*

$$\mathbb{E}_{b_i}[p_i(b_i)] = \mathbb{E}_{b_i}[\phi_i(b_i)x_i(b_i)]$$

Proof. The proof of this Lemma exploits Theorem 2.4 in a critical way and leads to the expressions defined as virtual valuations above.

First notice that the bid b (we dropped subscript i) is a random variable drawn from the distribution with cumulative distribution function F and density function f. So, using the result of Theorem 2.4 we have:

$$\begin{aligned}
\mathbb{E}_b[p(b)] &= \int_{b=0}^{h} p(b)f(b)db \\
&= \int_{b=0}^{h} bx(b)f(b)db - \int_{b=0}^{h}\int_{z=0}^{b} x(z)f(b)dzdb \\
&= \int_{b=0}^{h} bx(b)f(b)db - \int_{z=0}^{h} x(z)\int_{b=z}^{h} f(b)dbdz \\
&= \int_{b=0}^{h} bx(b)f(b)db - \int_{z=0}^{h} x(z)[1-F(z)]dz \\
&= \int_{b=0}^{h} bx(b)f(b)db - \int_{b=0}^{h} x(b)[1-F(b)]db \\
&= \int_{b=0}^{h} \left[b - \frac{1-F(b)}{f(b)}\right] x(b)f(b)db \\
&= \mathbb{E}[\phi(b)x(b)]
\end{aligned}$$

□

Having defined all that, we simply need to notice that we already know a mechanism that maximizes social surplus (namely VCG); and since maximizing profit amounts to maximizing a quantity very similar to social surplus, the natural mechanism is the following:

1. Given the bids $\mathbf{b}$, compute "virtual bids": $b_i' = \phi_i(b_i)$.
2. Run VCG on the virtual bids $\mathbf{b}'$ to get $\mathbf{x}'$ and $\mathbf{p}'$.
3. Output $\mathbf{x} = \mathbf{x}'$ and $p_i = \phi_i^{-1}(p_i')$.

One thing that is not obvious is why this mechanism is truthful: VCG is truthful because the allocation rule is monotone in the valuations of the bidders (see Example 2.5). Obviously in the above mechanism the allocation rule is monotone in the virtual valuations. Hence it is easy to show that the above mechanism is truthful if and only if virtual valuations are monotone in real valuations. This depends on the kind of distribution assumed by the bids and it amounts to the *monotone hazard rate assumption*, i.e. that $f(z)/(1-F(z))$ is monotone non-decreasing; in his original paper [2], Myerson describes an optimal mechanism for the general case where $f(z)/(1-F(z))$ does not satisfy the monotone hazard rate assumption. His technique is known as "ironing" and is beyond the scope of this thesis.

We end up this section by giving some specific examples for the use of Myerson's optimal mechanism.

Example 2.9.

single item auction: 1 item, n bidders, i.i.d. valuations from $[0, 1]$

In the case where there are n bidders and a single item for sale, we have that

$$c(\mathbf{x}) = \begin{cases} 0 & : \sum_i x_i \leq 1 \\ \infty & : \text{otherwise} \end{cases}$$

In the single item auction, the allocation that maximizes social surplus (which is also the one resulting from a run of VCG) has to allocate the item to the bidder who values it the most, **unless** this bidder has a negative valuation, in which case the social surplus is maximized by not allocating the item and getting a surplus of 0. However in normal auctions bidders are assumed to have positive valuations, so the issue of not allocating the item does not arise. In our case though, virtual valuations can be negative and therefore the optimal auction has to allocate the item to the bidder with the largest **positive** virtual valuation.

So agent i gets allocated the item precisely when $\phi(b_i) \geq \max\{\max_j\{\phi(b_j)\}, 0\}$, where we took into account that $\phi_i = \phi_j, \forall i, j$ because of the i.i.d. assumption. The payment is the minimum winning bid, i.e. $p = \phi^{-1}(\inf\{b : \forall i : \phi(b) > \phi(b_i) \wedge \phi(b) > 0\}) = \inf\{b : \forall i : b > b_i \wedge b > \phi^{-1}(0)\}$. Hence the resulting auction is the Vickrey auction with reservation price $\phi^{-1}(0)$. If the bids are drawn uniformly from $[0, 1]$ then the reservation price is $\frac{1}{2}$, which is exactly what we claimed in the beginning of this chapter.

Example 2.10.

1 **item,** n **bidders, i.i.d. valuations from** $[0, 1]$**, auctioneer has value** v_0 **for item**

In the case where there are n bidders and a single item for sale, we have that

$$c(\mathbf{x}) = \begin{cases} v_0 & : \sum_i x_i \leq 1 \\ \infty & : \text{otherwise} \end{cases}$$

Then the optimal Bayesian auction, i.e. the virtual-surplus maximizing auction, sells the item to bidder i if he has the highest virtual valuation among all other agents; in addition to that his virtual valuation has to be bigger than v_0, otherwise we get a higher social surplus by having the auctioneer retaining the item and getting happiness v_0. Hence agent i gets allocated the item precisely when $\phi(b_i) \geq \max\{\max_j\{\phi(b_j)\}, v_0\}$, where we took into account that $\phi_i = \phi_j, \forall i, j$ because of the i.i.d. assumption. The payment is the minimum winning bid, i.e. $p = \phi^{-1}(\inf\{b : \forall i : \phi(b) > \phi(b_i) \wedge \phi(b) > v_0\}) = \inf\{b : \forall i : b > b_i \wedge b > \phi^{-1}(v_0)\}$. Hence the resulting auction is the Vickrey auction with reservation price $\phi^{-1}(v_0)$. If the bids are drawn uniformly from $[0, 1]$ then the reservation price is $\frac{v_0+1}{2}$.

Example 2.11.

digital goods auction: n **identical items,** n **bidders, i.i.d. valuations from** F

We assume that we are dealing with bidders with i.i.d. valuations drawn from a distribution F. Myerson suggests we do the following:

1. Replace each valuation v with the virtual valuation $\phi(v) = v - \frac{1-F(v)}{f(v)}$ (where F and f are the cumulative distribution function and the probability density function respectively).

2. In order to maximize the expected profit, we simply maximize the expected virtual surplus $\sum_i \phi_i(v_i)x_i$. This can be achieved by allocating the good to any bidder with a positive virtual valuation, i.e. to any bidder with $v_i - \frac{1-F(v_i)}{f(v_i)} \geq 0$.

3. Charge each bidder his minimum winning bid; under the assumption that the bids are i.i.d. random variables with distribution F, every bidder is charged $\phi^{-1}(0)$ which is the solution to $v - \frac{1-F(v)}{f(v)} = 0$.

Having defined the optimal Bayesian mechanism a-la Myerson, let us try to give an ex-post interpretation of it (more about this approach can be found in [4]). It is easy to verify that the solution to $v - \frac{1-F(v)}{f(v)} = 0$ is the same as $\arg\max_p p(1 - F(p))$ (just differentiate the second expression). This leads us to the following conclusion:

Proposition 2.12. *The optimal Bayesian digital good auction for n bidders with valuations drawn i.i.d. from some distribution F is to make each bidder a take-it-or-leave-it offer for a price of* $\arg\max_p p(1 - F(p))$.

Note how intuitive the above auction is: the offered price is simply the price that maximizes our expected profit. Indeed $p(1 - F(p))$ is our expected profit, in the sense that we expect $(1 - F(p))$ of the bidders to have a valuation that is more than the offered price of p and accept our offer, thus giving us a profit of p each.

Example 2.13.

k **identical items,** $n > k$ **bidders, i.i.d. valuations from** $[0, 1]$

In this case we have that

$$c(\mathbf{x}) = \begin{cases} 0 & : \sum_i x_i \leq k \\ \infty & : \text{otherwise} \end{cases}$$

We remember from the single-item auction for bidders with i.i.d valuations drawn uniformly from $[0, 1]$ with seller's value 0, that the optimal Bayesian auction is the Vickrey auction with reservation price $1/2$. It is easy to verify that for this case as well the social surplus is maximized by allocating the k-items to the k-highest bidders, provided that they have positive virtual valuations, which corresponds to real valuations greater than $1/2$. The price charged is again the minimum winning bid at each case.

More formally: if there are $m \geq k + 1$ bidders with valuations $\geq 1/2$ we allocate the k items to the k highest bidders and charge them the value of the $k + 1$ highest bidder, whereas if there are $m \leq k$ bidders with valuations $\geq 1/2$ we allocate the item only to these m bidders and charge them $1/2$.

Chapter 3

Worst Case Optimal Mechanism Design

Our goal in this chapter is to design a truthful mechanism that achieves the optimal profit for any input bid sequence. Namely, we will henceforth be making no assumptions about the bids coming from a specific (known) distribution, as was the case in the Bayesian setting; rather we will assume that some (cruel) adversary, who knows how our mechanism is working, is picking the bids with the sole objective to keep our profit as low as possible. We aim at designing auctions that achieve a decent profit, no matter how bad the input bid sequence is. We shall allow our auctions to use randomness (which enables us in some way to conceal information from the omniscient adversary) and we will try to analyze the worst-case behavior of our mechanisms, a ubiquitous approach in computer science; in the context of auctions, worst-case analysis is usually referred to as worst case or prior-free optimal mechanism design.

3.1 Another (?) characterization of truthful mechanisms

We already saw in the previous chapter a very crisp characterization of truthful mechanisms for single parameter settings. The punchline was that any truthful mechanism has to be equivalent to a bid-independent mechanism defined as follows:

Definition 3.1. (Bid-independent Mechanism, $\mathcal{A}_f(\mathbf{b})$)

Given a function $f : \mathbb{R}^{n-1} \to \mathbb{R}$, henceforth called the threshold function, do, for each bidder i:

1. $t_i \leftarrow f(\mathbf{b}_{-i})$
2. *If $t_i \leq b_i$, set $x_i \leftarrow 1$ and $p_i \leftarrow t_i$ (bidder i wins)*
3. *Otherwise, set $x_i = p_i = 0$ (bidder i is rejected)*

This auction can be summarized in the following sentence: "the auctioneer makes a take-it-or-leave-it offer to bidder i for the price $f_i(\mathbf{b}_{-i})$". The above definition can be easily generalized to randomized mechanisms as well.

The characterization of truthful auctions for single parameter agents of the previous chapter makes the equivalence of truthful and bid-independent auctions quite obvious. For the sake of completeness we also provide a proof of this fact below.

Theorem 3.2. *A deterministic auction is truthful if and only if it is equivalent to a deterministic bid-independent auction.*

Proof. We first prove that any bid-independent auction is truthful. Let i be a bidder and v_i his true valuation. If $v_i < t_i$ and bidder i declares $b_i = v_i$ then he loses the item and has profit 0. If he lies, he can either bid some other value below t_i, in which case he would still not win the item and have a profit of 0, or he can choose to bid higher than t_i, in which case his profit would be negative. In either case, bidder i wins nothing by lying, hence bidding v_i maximizes his profit. In the case where $v_i \geq t_i$, bidder i pays t_i and wins the item, having a positive profit. If he chooses to bid below t_i he would lose the item and have a profit of 0, and if he chooses some other value greater than t_i he would win the item but his profit would remain the same (as the price paid does not depend on his own bid). Hence, once again truth telling is the best strategy.

We now show the inverse statement: any truthful auction is equivalent to a bid-independent auction. Recall the characterization of truthful auctions of Theorem 2.4 and for simplicity let us focus on the deterministic case; we saw that in any truthful auction there exists a threshold price for each bidder i, which is independent of his own valuation and corresponds to his minimum winning bid, i.e. the smallest bid he could announce and still win the item. In the language of bid-independent auctions, this would be $t_i = f(\mathbf{b}_{-i})$. The payment –in the case where bidder i wins the item– was also set to be t_i, so what we have is exactly a bid-independent mechanism. □

The above theorem is a very useful tool: it allows us to limit our attention only on auctions that can be formulated as bid-independent auctions. In the rest of this chapter we shall try to design optimal truthful (aka bid-independent) auctions for the specific setting of "digital goods", which we already saw in Example 2.11.

3.2 Digital Good Auctions and Profit Benchmarks

We start off by giving the definition of the auction setting which we will study in the rest of this chapter and in the next one.

Definition 3.3. (Digital Good Auction) *In the digital good auction there are,*

- *n bidders with valuations $v_1 \geq \ldots \geq v_n$.*
- *n identical copies of an item for sale (for example copies of an MP3 file).*

Let us now try to design the optimal, prior-free mechanism for the above setting. By optimal mechanism we have in mind something similar to what we did in Chapter 2: there, we designed mechanisms that have the highest (optimal) possible expected profit. Now we do not make any statistical assumptions about the distribution of the bids, so we do not aim at maximizing expected profit, but rather achieving the optimal profit for any input bid sequence, a globally optimal auction in some sense.

However, having the above characterization of truthful auctions in hand it is easy to see that the goal of designing a globally-optimal prior-free mechanism is over-ambitious: indeed, consider the scenario where we have a single bidder in the auction and consider the two cases where this bidder values the good for 1 and 2 respectively. Since any truthful auction has to be a bid-independent auction, the offer made to the bidder has to be independent of his own valuation and, since there are no other bidders, it cannot be but a constant. Obviously though, the optimal take-it-or-leave-it offer is 1 in the first case and 2 in the second, so no auction can be globally optimal, even for this simple case.

To get past this obstacle we shift our focus from absolute to relative optimality. This is a common approach in computer science. For example, to cope with limited computational resources we introduce the notion of approximation algorithms, where the goal is to compare (favorably) against a computationally unbounded algorithm (adversary). Likewise, when we study online algorithms, what we lack is a knowledge of the future; in this case we want to compete (hence the name competitive analysis) against an omniscient adversary who knows all the input in advance.

The situation here has a lot in common with competitive analysis: the obstacle in our case are the incentives of the agents. We need to design a truthful mechanism, i.e. we need to compete against an omniscient adversary who has more information than we do, namely the true valuations of the agents. Our goal will be to achieve a good profit **compared** to the profit achieved by an omniscient algorithm and we shall use the framework of competitive analysis to formalize this idea.

Definition 3.4. *A profit benchmark is a function $\mathcal{G} : \mathbb{R}^n \to \mathbb{R}$ which maps a vector of valuations to a target profit*

Intuitively this is the omniscient algorithm-mechanism we need to compare against. Let us see a couple of potential adversaries.

Benchmark 1. Sell to all bidders at their valuation.

$$Profit = \mathcal{T}(\mathbf{v}) = \sum_i v_i$$

Benchmark 2. Sell at optimal single sale price.

$$Profit = \mathcal{F}(\mathbf{v}) = \max_i i \cdot v_i$$

Though the first candidate seems much more promising (in fact it is the best one can hope for) the next proposition shows that these two algorithms yield profits that are at most $\log n$ apart, where n is the number of the bidders. Formally we have the following theorem:

Theorem 3.5. *There exist bid vectors* $\mathbf{b}$ *for which*

$$\mathcal{F}(\mathbf{b}) = \Theta(\mathcal{T}(\mathbf{b})/\ln n).$$

Moreover, for all bid vectors $\mathbf{b}$

$$\mathcal{F}(\mathbf{b}) \geq (\mathcal{T}(\mathbf{b})/\ln n).$$

Proof. For the first part we just need to consider the bid $\mathbf{b}$ with $b_i = n/i$. We can easily check that $\mathcal{F}(\mathbf{b}) = n$ and $\mathcal{T}(\mathbf{b}) = n(\ln n + \Theta(1))$ (from a standard approximation of the harmonic mean).

For the second part, let $\mathcal{F}(\mathbf{b})= \max_i iz_i = kz_k$. Then for all i, $iz_i \leq kz_k$, hence

$$\mathcal{T}(\mathbf{b}) = \sum_{i=1}^{n} z_i = \sum_{i=1}^{n} \frac{iz_i}{i} \leq \sum_{i=1}^{n} \frac{kz_k}{i} = \mathcal{F}(\mathbf{b}) \sum_{i=1}^{n} \frac{1}{i} = \mathcal{F}(\mathbf{b})(\ln n + \Theta(1))$$

□

By keeping our expectations down to a minimal level we shall try to design auctions that are constant competitive against $\mathcal{F}(\mathbf{v})$; by constant competitive we mean that the competitive ratio of the designed auction $\mathcal{A}$, defined as $\beta = \max_{\mathbf{v}} \frac{\mathcal{F}(\mathbf{v})}{\mathcal{A}(\mathbf{v})}$, is constant.

Nonetheless such an auction does not exist; in fact no (truthful) auction is $o(\log h)$ competitive with $\mathcal{F}(\mathbf{v})$ for bids $v_i \in [1, h]$. The reason for that is that if h gets arbitrarily large (i.e. if we have one very high bidder) then the profit of the benchmark is dictated by this bidder's price. However, when we try to define a truthful auction, we have to make a take-it-or-leave-it offer to the highest bidder and the price offered to him will only depend on the rest (probably much lower) bids. Hence our profit from the high bidder will be negligible, whereas our profit from the rest of the bidders will also be low if not zero, since we will probably make them an offer for a price that is comparable to h, which most bidders will probably turn down. Formally we have the following theorem:

Proposition 3.6. *For any truthful auction $\mathcal{A}_f$ and any $\beta \geq 1$, there exists some bid vector $\mathbf{b}$ such that the expected profit of $\mathcal{A}$ on $\mathbf{b}$ is less than $\mathcal{F}(\mathbf{b})/\beta$.*

Proof. The key idea is to use a bid vector where the largest bid is a lot bigger than all the rest, thus determining the profit of the optimal single price omniscient auction. When this bid is used as input to any truthful (i.e. bid independent) auction, the threshold function f charges to the high bidder a relatively small amount, since it does not take into account his own bid. Since the contribution of the rest of the bidders is negligible the total profit may be quite far from the optimum.

Formally, consider a bid-independent randomized auction on two bids, 1 and $x \geq 1$. Let h be the smallest value greater or equal to 1 such that $Pr[f(1) \geq h] \leq \frac{1}{2\beta}$. Then the profit on input vector $\mathbf{b} = (1, H)$ with $H = 4\beta h$ is at most

$$\frac{H}{2\beta} + h(1 - \frac{1}{2\beta}) + 1 < 4h = \frac{H}{\beta} = \frac{\mathcal{F}(\mathbf{b})}{\beta}$$

□

The above lemma implies that we cannot expect a close matching to the performance of $\mathcal{F}$ by any truthful auction. Hence we have to set our goals even lower: instead of comparing the performance of an auction to the performance of $\mathcal{F}$, we compare it to the performance of $\mathcal{F}^{(2)}$ -the optimal single price omniscient auction that sells at least two units[1].

Benchmark 3. Sell at optimal single sale price to at least two bidders.

$$Profit = \mathcal{F}^{(2)}(\mathbf{v}) = \max_{i \geq 2} i \cdot v_i$$

To clarify things let us see an example.

Example 3.7. Consider the simple setting consisting of 10 high bidders with values \$10 and 90 low bidders with values of \$1. We have two options: either sell at price \$1, in which case all 100 bidders will get the items and we will get a profit of \$100, or sell at price \$10, in which case only the 10 highest bidders will get the items and we will get a profit of \$100 again. So, in this example the optimal single sale price can be either \$1 or \$100.

Consider now having 10 high bidders with values \$5 and 90 low bidders with values of \$1. It is easy to verify that in this case the optimal single sale price is \$1, yielding a profit of \$100.

In the next section we shall see that we can design auctions that compare favorably to $\mathcal{F}^{(2)}$.

3.3 The Competitive Framework and Deterministic Auctions

Having formalized what we need by the term "adversary" we now need to make the notion of optimality rigorous. Our metric of quality for a mechanism $\mathcal{A}$ will be the worst-case (over all input bid sequences) ratio of its profit to the profit of some carefully chosen benchmark $\mathcal{G}$ (in fact from now on we will limit our attention to the benchmark $\mathcal{F}^{(2)}$). This ratio is commonly known as competitive ratio; formally:

Definition 3.8. *We say that a (possibly randomized) auction $\mathcal{A}$ is β-competitive against $\mathcal{F}^{(2)}$ if*

[1]an alternate strategy would be to consider $\mathcal{F}$ as a metric, but limit our attention to input bids where the highest bidder is not much larger than the rest

for all bid vectors $\mathbf{b}$*, the expected profit of* $\mathcal{A}$ *on* $\mathbf{b}$ *satisfies*

$$\mathbf{E}[\mathcal{A}(\mathbf{b})] \geq \frac{\mathcal{F}^{(2)}}{\beta}.$$

We say that an auction is competitive against $\mathcal{F}^{(2)}$ *if the auction is* β*-competitive, where* β *is a constant. We refer to* β *as the competitive ratio of* $\mathcal{A}$*.*
From now on "$\mathcal{A}$ *is competitive" stands for "competitive against* $\mathcal{F}^{(2)}$*", unless otherwise stated.*

We first try to design a deterministic, truthful and competitive auction. Bearing in mind Theorem 3.2, perhaps the most natural candidate is the following:

Definition 3.9. (Deterministic Optimal Price (DOP)) *The Deterministic Optimal Price auction is a bid independent auction defined by the function*

$$f(b_{-i}) = \mathrm{opt}(\mathbf{b}_{-i}) = \arg\max_p (p \times \textit{number of bidders } j \neq i \textit{ with } b_j \geq p)$$

Hence DOP makes each bidder a take-it-or-leave-it-offer for the optimal sale price that is computed for the rest of the bids.

DOP is clearly truthful, because it is bid-independent; however Example 3.7 shows that its performance is not that good: consider again the auction consisting of 10 high bidders (with value \$10) and 90 low bidders (with value \$1). It is easy to verify that DOP will make an offer of \$1 to all bids at 10 and accept them, and will make an offer of \$10 to all bids at 1 and reject them; thus its total profit is \$10, while the optimal single-price mechanism will get a profit of \$100. This example can be made arbitrarily bad and in fact we can generalize it to prove that no deterministic, symmetric and truthful auction is competitive; an auction is symmetric if its outcome is not a function of the order of the bids, but only of their values.

Theorem 3.10. *No symmetric, deterministic, truthful auction is constant-competitive against* $\mathcal{F}^{(2)}$*.*

Proof. Let $\mathcal{A}_f$ be any symmetric deterministic auction, where the threshold prices are given by the function f. We will show that $\mathcal{A}_f$ is not competitive: in fact, we will show that there exists a bid vector $\mathbf{b}$ of arbitrary size n, such that the profit of $\mathcal{A}_f$ on $\mathbf{b}$ is at most $\mathcal{F}^{(2)}/n$.

Consider the set of bid vectors whose bids are all n or 1. For $0 \leq j \leq n-1$ write $f(j)$ for the price the auction assigns to a vector with exactly j bids at n and $n-1-j$ bids at 1 (f is well defined since $\mathcal{A}_f$ is symmetric). We will make the assumption that f takes one of the values in $\{1, n\}$ as this restriction cannot hurt the auction profit.

Note first that we must have $f(0) = 1$; indeed, if $f(0) = n$ then for the bid vector that has all bids equal to 1 $\mathcal{A}_f$ would yield a profit of 0 and the theorem would hold trivially.

Also it has to be that $f(n-1) = n$ otherwise, for the bid vector with all bids equal to n, $\mathcal{A}_f$ will yield a profit of n while the optimal profit is n^2.

Let k now be the largest integer in $\{0, \ldots, n-1\}$ such that $f(k) = 1$; because of the previous discussion we know that $k \in \{0, \ldots, n-1\}$. Let $\mathbf{b}$ now be the vector with $k+1$ bids at n and $n-k-1$ bids at 1. Since $f(k+1) = n$ all bids at 1 are rejected and all high bids win at price $f(k) = 1$ so the total profit of $\mathcal{A}_f$ is $k+1$. If $k = 0$ then $\mathcal{F}^{(2)}$ has profit n, by accepting all bids at price 1 and $\mathcal{A}_f$ has profit 1, so the conclusion holds. If $k \geq 1$ then $\mathcal{F}^{(2)}$ has profit at least $(k+1)n$ by accepting all $k+1$ high bids at price n and $\mathcal{A}_f$ has profit $k+1$, so the theorem holds in this case as well. $\square$

There exist asymmetric deterministic auctions that are truthful and competitive. However these auctions result from derandomization of randomized auctions, so we will focus our attention on randomized auctions for the rest of this chapter.

3.4 Randomized Auctions

In this section we will describe two very natural randomized auctions, based on the method of random sampling. Before moving on to the details we should mention that even randomized auctions have limits in their performance, i.e. they cannot match the performance of $\mathcal{F}^{(2)}$. A lower bound on the performance of any auction (including randomized ones) is $\frac{\mathcal{F}^{(2)}}{2.42}$. The proof of this bound can be found in [5].

The main idea behind the design of both randomized auctions we shall present here is the following: we partition the set of bids in two sets. The partitioning is performed at random, by flipping a fair coin for each bid, in order to decide which side to assign it to. We then use one partition for market analysis and plug what we learn to the other side of the partition. At this point it is instructive to think about Example 2.11 again; we saw there that the optimal Bayesian auction makes a take-it-or-leave-it offer to all bidders at an "appropriate" price, namely at the price that maximizes the expected revenue. Since in prior-free mechanism design we have no statistical information about the distribution of the bids, the best we can hope for is to try to elicit some sort of information, by performing some sort of market analysis. Intuitively we try to build an empirical distribution for the bids which we will use to pick the optimal single sale price offer we will make to the bidders.

Our first auction example, the Random Sampling Optimal Price Auction (RSOP), defined in [5], is arguably the most natural randomized auction. RSOP basically runs DOP in each side of the partition and uses the resulting threshold for the other side of the partition; this is essentially a dual-priced auction. Nonetheless we can easily modify RSOP so as to reject all bids of one partition (by skipping step 4 below), hence making it single-priced. Nonetheless, as we shall see in the next Chapter, RSOP practically (i.e. in the worst case) only gets profit from one side of the partition.

Definition 3.11. (RSOP)

1. *Randomly partition the bids* $\mathbf{b}$ *into two parts* $\mathbf{b}'$ *and* $\mathbf{b}''$ *by flipping a fair coin for each bidder.*

2. *Compute* $t' = opt(\mathbf{b}')$ *and* $t'' = opt(\mathbf{b}'')$*, the optimal sale prices for each part.*

3. *Make a take-it-or-leave-it offer of* t' *to all bidders in* $\mathbf{b}''$ *and a take-it-or-leave-it offer of* t'' *to all bidders in* $\mathbf{b}'$*.*

Since the offered price to each bidder is not a function of his own bid the auction is truthful (bid-independent). It is also competitive as the following theorem suggests.

Theorem 3.12. *RSOP is 15-competitive.*

We defer the proof of this theorem until the next section, which will be entirely devoted to the study of RSOP.

The best lower bound we have for RSOP is 4 and it comes from the following example: consider a bid vector $\mathbf{b}$ consisting of only two very high bids $h+\epsilon$ and h and all other bids negligibly small. The profit of RSOP in this example is not zero only if both of the high bids fall on different sides of the partition (which happens with probability 1/2); if this is indeed the case, then RSOP has non zero profit only from the side of the partition where $h+\epsilon$ lies – since the price offered to bidder h is $h+\epsilon$ and is rejected. Hence the expected profit of RSOP in this example is $h/2$, while $\mathcal{F}^{(2)} \approx 2h$, so we get a competitive ratio of 4.

Another auction considered in [5] is the Sampling Cost-Sharing Auction, which makes use of a well known approach for the design of cost sharing mechanisms, called the Shapley Value mechanism (due to Moulin and Shenker). For our purposes we define the following profit extractor:

ProfitExtract$_C$: Given bids $\mathbf{b}$, find the largest k such that the highest k bidders' values are at least C/k. Charge each C/k.

ProfitExtract$_C$ can be implemented by the following algorithm which is given n, C and $\mathbf{b}$:

1. Offer C/n to all bidders.

2. If $b_i \leq C/n$ reject bidder i.

3. Let $\mathbf{b}$ (resp. n) be the bidders (resp. number of bidders) remaining.

4. Repeat until nobody is rejected in Step 2.

Example 3.13. Consider $\mathbf{b} = \{1, 2, 3, 4, 5\}, R = \$9$ and $n = 5$ bidders. The first iteration offers 9/5 to all bidders and bidder 1 is rejected. In the second iteration every bidder is offered 9/4 and bidder 2 is rejected. In the third iteration every bidder is offered 9/3 and everybody accepts the offer and pays \$3.

It is easy to verify that ProfitExact$_C$ is truthful; in fact it is group strategyproof (truthful), meaning that no coalition of deviating players can achieve an outcome that is at least as good for all deviating agents and strictly better for at least one. Furthermore, if $C \leq \mathcal{F}(\mathbf{b})$ then ProfitExtract$_C$ has revenue C, otherwise it has no profit; it is hence a proper profit extractor.

It is obvious that if we knew the right value for C, which is $\mathcal{F}(\mathbf{b})$ in our case, we could run ProfitExtract$_C$ for this C and get the desired profit. However, as usual we are limited by the incentives of the players and we do not know a priori their true valuations and $\mathcal{F}$.

To get past this difficulty, we once again employ random partitioning of the bids: we split the bidders in half at random and –based on the assumption that the two partitions must look quite similar– we try to extract the optimal profit of each side of the partition from the other side. Using ProfitExtract$_C$ we can write down the Sampling Cost-Sharing Auction as follows:

Definition 3.14. (SCS)

1. *Partition bids* $\mathbf{b}$ *uniformly at random into two sets: for each bid with probability 1/2 put the bid in* $\mathbf{b}'$ *and otherwise in* $\mathbf{b}''$.

2. *Compute* $\mathcal{F}' = \mathcal{F}(\mathbf{b}')$ *and* $\mathcal{F}'' = \mathcal{F}(\mathbf{b}'')$.

3. *Compute the auction results by running ProfitExtract*$_{\mathcal{F}'}$ *on* $\mathbf{b}''$ *and ProfitExtract*$_{\mathcal{F}''}$ *on* $\mathbf{b}'$.

SCS is a bid-independent auction and as such it is truthful. Furthermore it is 4-competitive and this is tight; in fact SCS is tight for the exact same bid sequence as RSOP. We end this chapter with a proof of 4-competitiveness for SCS. We first state (without proof) a probabilistic lemma that will be useful.

Lemma 3.15. *If we flip $k \geq 2$ fair coins, then*

$$\mathbb{E}[\min(\#heads, \#tails)] \geq k/4$$

Theorem 3.16. *SCS is 4-competitive*

Proof. To gain some useful intuition to SCS notice that ProfitExtrac_C allows us to treat a set of bidders $\mathbf{b}$ as one bidder with bid value $\mathcal{F}(\mathbf{b})$. Recall also from Theorem 3.2 that a truthful auction must just make a take-it-or-leave-it offer at price t_i to bidder i who accepts it if his value is at least t_i; similarly ProfitExtrac_C extracts profit C when $\mathcal{F}(\mathbf{b}) \geq C$. The SCS auction can then be viewed as randomly partitioning the bidders into two parts, treating each part as a single bidder and performing the second-price Vickrey auction on these two "bids".

With this intuition it is obvious that the profit of SCS is $\min(\mathcal{F}', \mathcal{F}'')$, for $\mathcal{F}', \mathcal{F}''$ as in Definition 3.14. Hence, it suffices to analyze $\mathbb{E}[\min(\mathcal{F}', \mathcal{F}'')]$.

Assume that $\mathcal{F}^{(2)}(\mathbf{b}) = kp$ has $k \geq 2$ winners at price p. Of those k winners, let k' be the number of winners that are in $\mathbf{b}'$ and k'' the number that are in $\mathbf{b}''$. Since there are k' bidders in $\mathbf{b}'$ at price p, then $\mathcal{F}' \geq k'p$. Likewise, $\mathcal{F}' \geq k''p$. We now have:

$$\begin{aligned}
\frac{\mathbb{E}[SCS(\mathbf{b})]}{\mathcal{F}^{(2)}(\mathbf{b})} &= \frac{\mathbb{E}[\min(\mathcal{F}', \mathcal{F}'')]}{kp} \\
&\geq \frac{\mathbb{E}[\min(k'p, k''p)]}{kp} \\
&= \frac{\mathbb{E}[\min(k', k'')]}{k} \\
&\geq \frac{1}{4}
\end{aligned}$$

where in the last inequality we employed Lemma 3.15. □

Chapter 4

Analysis of the Random Sampling Optimal Price Auction

The subject of this chapter is the Random Sampling Optimal Price Auction. We already briefly discussed the auction in Chapter 3 in the context of randomized, competitive auctions (see Definition 3.11). We then stated Theorem 3.12 claiming that RSOP is constant-competitive with a competitive ratio of 15. It is in fact conjectured that the competitive ratio of RSOP is 4 and that the lower bound example of the previous chapter is the worst possible.

Conjecture 4.1. *RSOP is 4-competitive.*

In the next section we provide a sketch of the proof of Theorem 3.12 and we conclude with a section containing some results aiming to prove Conjecture 4.1; these results are mostly meant to be kept as a log of various attempts to prove Conjecture 4.1. I hope people who try to read them do not keep any hard feelings...

4.1 Proof of Theorem 3.12

The first proof of RSOP having a constant competitive ratio was given in [5], where they proved a competitive ratio of 7600. Nonetheless this result was far from being tight. Indeed in [6] a tighter analysis led to an upper bound of 15 on the competitive ratio of RSOP. Their proof is based on a lemma about random walks. Moreover, in this paper they provide further support to the conjecture that RSOP is 4-competitive, by exhibiting a family of inputs for which RSOP has expected profit at least $\mathcal{F}^{(2)}/4$; this is the equal revenue input vector $\mathbf{b}$ with $b_i = 1/i$, for which random sampling at first glance does not seem to be an appropriate technique. In this section we sketch the proof technique of [6] for the upper bound of 15.

Notation 4.2. Let us first introduce some terminology and notation; we shall not be using the one of [6], mostly because we want to use the same notation as in the next section.

- Let $B = \{b_1, \ldots, b_n\}$ be the set of all bids and let $B_2 = \{b_2, \ldots, b_n\}$. We henceforth assume that $b_1 \geq \ldots \geq b_n$. We refer to subsets $S \subseteq B$ as *sequences* of bids and we consider them to be sorted in descending order as well.
- Given a specific partition of bids $b_1, \ldots, b_n$ in two, we use $b_{j_1}, \ldots, b_{j_k}$ to denote the side of the partition that does not contain the highest bid b_1, i.e. by writing $b_{j_1}, \ldots, b_{j_k}$ we assume implicitly that $j_1 \geq 2$. As in [6] we call this the "good" side of the partition and the other side the "bad" side.
- We define $y(b_{j_1}, \ldots, b_{j_k}) = \max\{b_{j_1}, 2b_{j_2}, \ldots, kb_{j_k}\}$, where $b_{j_1} \geq \ldots \geq b_{j_k}$.
- We denote by $z(b_{j_1}, \ldots, b_{j_k})$ the resulting profit from offering the optimal price of the "good" side $(b_{j_1}, \ldots, b_{j_k})$ to the "bad" side. This can be formally written as

$$z(b_{j_1}, \ldots, b_{j_k}) = \max_{t=1\ldots k}\{t \cdot b_{j_t}\} \cdot \left(\text{number of bids } \geq \arg \max_{t=1\ldots k}\{t \cdot b_{j_t}\} \text{ on the "bad" side}\right)$$

$$= y(b_{j_1}, \ldots, b_{j_k}) \cdot \left(\text{number of bids } \geq \arg \max_{t=1\ldots k} \{t \cdot b_{j_t}\} \text{ on the "bad" side} \right)$$

The first observation made in [6] is that in our analysis we can ignore the profit from the "good" side; this is because, in the worst case, the highest bid will be so large that the profit from the "good" side of the partition will be zero. In fact, we can assume without loss of generality that the highest bid is larger than $\mathcal{F}^{(2)}$.

The goal now is to prove the following two things:

- First we want to prove that the best single-price profit from the good side of the partition will be at least half of the single-price profit for the whole set of bids with probability 1/2. More formally, we want to argue that $y(b_{j_1}, \ldots, b_{j_k}) \geq \mathcal{F}^{(2)}/2$ with probability at least 1/2.
- Then we want to show that the profit we get by offering the optimal price for the "good" side of the partition to the "bad" side, is at least $\frac{1}{3}y(b_{j_1}, \ldots, b_{j_k})$, with probability at least 0.9.

By combining the above we can easily prove an upper bound of 15 on the competitive ratio of RSOP. Let us briefly explain how the above conclusions can be drawn.

For the first one we merely need to prove that there are "enough" "high" bidders in the good side; by "high" bidders we mean bidders who have valuations larger than the optimal sale price for the full set of bids, i.e. if the optimal price for the full set of bids is $b_{i'}$ with $i' = \arg\max ib_i$, we mean all bidders with $b_j \geq b_{i'}$; by "enough" we demand at least half of them. It is straightforward then that indeed, at least half of the "high" bidders will lie on the same side of the partition as the highest bidder, with probability 1/2; then with probability at least 1/2,

$$y(b_{j_1}, \ldots, b_{j_k}) \geq b_{i'} \cdot [\text{ number of bids } \geq b_{i'} \text{ on good side of the partition}] \geq b_{i'} \cdot \frac{i'}{2} = \mathcal{F}^{(2)}/2$$

For the second observation things are slightly more complicated. Intuitively, what we need to show is that the two sides of the partition are not very different in the following sense: a price that is good for one side should not be too bad for the other side either. Let b_{j_t} be the optimal price for the "good" side $(b_{j_1}, \ldots, b_{j_k})$, i.e. $t \cdot b_{j_t} = \max\{b_{j_1}, 2b_{j_2}, \ldots, kb_{j_k}\}$; obviously then $z(b_{j_1}, \ldots, b_{j_k}) = (j_t - t)b_{j_t}$. We already know that $tb_{j_t} = y(b_{j_1}, \ldots, b_{j_k})$ is $\geq \mathcal{F}^{(2)}/2$ with probability $1/2$. So it suffices to show that $(j_t - t)b_{j_t} \geq \alpha t b_{j_t}$, for some constant α. This boils down to proving that –with some probability– $(j_t - t) \geq \alpha t$ for some constant α.

To show that consider a discrete random walk on a line such that in each time step the walk takes one step forward or stays put independently with probability $1/2$. Notice that the partition of the bids in two sides is very much like this random walk: at time 1 we are at the origin (the "bad" side in our case). Then, when we consider a bid b_{j_i} of the good side, it is like looking at the j_i-th step of this walk; the index i of this bid is exactly the number of bids that have picked the good side of the partition so far or equivalently the distance of the random walk from the origin at time j_i (one would expect the notation to be the other way here, but I kept this notation as I will need it for the next section). Hence, for a given time j_i we would like to know what is the maximum distance i of the random walk from the origin. All we need now is the following lemma, which is proved in [6].

Lemma 4.3. *Consider a discrete random walk on a line such that in each time step the walk takes one step forward or stays put, independently with probability* $1/2$*. If we start at the origin at time* $t = 1$*, then let* $\mathcal{E}_\alpha$ *be the event that at no time* $t \geq 1$ *is the random walk further than* αt *from the origin. It holds that* $\Pr[\mathcal{E}_{3/4}] \geq 0.9$*.*

In terms of our bids, this lemma tells us that $\Pr[\forall i : \ i \leq \frac{3}{4} j_i] \geq 0.9$ which implies that $\Pr[\forall i : \ j_i - i \geq \frac{1}{3} i] \geq 0.9$.

To prove Theorem 3.12 just note that the probability of the two events described above happening together is at least 0.4, so

$$RSOP = (j_t - t)b_{j_t} \geq \frac{tb_{j_t}}{3} = \frac{1}{3} y(b_{j_1}, \ldots, b_{j_k}) \geq \mathcal{F}^{(2)}/6$$

with probability at least 0.4, so $\mathbb{E}[RSOP] \geq 0.4 \cdot \mathcal{F}^{(2)}/6 = \mathcal{F}^{(2)}/15$.

4.2 Aiming for a competitive ratio of 4

This section is based on joint work with Elias Koutsoupias.

We first note that we can easily express RSOP 's expected profit as an expression involving only the z values:

$$\text{RSOP} = \sum_{S \subseteq B_2} z(S) 2^{-n+1}$$

To be more precise, the above expression gives us the worst case expected profit of RSOP in the following sense: the adversary will always pick a bids' sequence $b_1, \ldots, b_n$ with sufficiently large b_1, so that RSOP gets no profit from the "good" side (see the proof in Section 2 too). So it suffices to examine the profit collected from the "bad" side which is exactly the above expression.

It is easy to see now that the following two claims prove Conjecture 4.1.

Claim 4.4.

$$\sum_{S \subseteq B_2} z(S) \geq \sum_{S \subseteq B_2:\ b_2 \in S} y(S)$$

Proof. We do not have the proof for this Claim. We discuss some ideas at the end of this section. □

Claim 4.5.

$$\sum_{S \subseteq B_2:\ b_2 \in S} y(S) \geq 2^{n-3}\ ib_i$$

Proof. Let b_i be the optimum single sale price for the whole set of bids, i.e. $\mathcal{F}^{(2)} = ib_i$ (although our result holds for any bid b_i).

We will introduce a mapping between the set of sequences $X = \{S \subseteq B_2 |\ b_2 \in S \ \&\ b_i \notin S\}$ and the set $Y = \{S \subseteq B_2 |\ b_2 \notin S \ \&\ b_i \in S\}$. Given a sequence of bids $S \in X$ let $t = \max\{j : j < i, b_j \in S\}$. We then define the following mapping for each bid $b_j \in S$:

$$f(b_j) = \begin{cases} b_{j+i-t} & : \text{if } j < i \\ b_j & : \text{if } j > i \end{cases}$$

It is easy to see that the mapping $g : X \longrightarrow Y$ defined as $g(b_{j_1}, \ldots, b_{j_k}) = (f(b_{j_1}), \ldots, f(b_{j_k}))$ is in fact a bijection. For example, if the optimal price is b_5 then $i = 5$ and we have $g(b_2, b_3, b_8) = (b_4, b_5, b_8)$, $g(b_2, b_4, b_8) = (b_3, b_5, b_8)$ and so on. Moreover, it is easy to see that $b_1 \geq \ldots \geq b_n$ implies that $y(S) \geq y(g(S))$.

Hence we have that:

$$\begin{aligned}
\sum_{S \subseteq B_2:\ b_2 \in S} y(S) &= \sum_{S \subseteq B_2:\ b_2 \in S, b_i \in S} y(S) + \sum_{S \subseteq B_2:\ b_2 \in S, b_i \notin S} y(S) \\
&\geq \sum_{S \subseteq B_2:\ b_2 \in S, b_i \in S} y(S) + \sum_{S \subseteq B_2:\ b_2 \in S, b_i \notin S} y(g(S)) \\
&= \sum_{S \subseteq B_2:\ b_2 \in S, b_i \in S} y(S) + \sum_{S \subseteq B_2:\ b_2 \notin S, b_i \in S} y(S) \\
&= \sum_{S \subseteq B_2:\ b_i \in S} y(S) \\
&= 2^{n-i} \sum_{j=0}^{i-2} \binom{i-2}{j} (j+1) \cdot b_i
\end{aligned}$$

where the last equality follows from the following simple counting argument: consider all possible positions of b_i in a sequence of B_2. There can be j bids larger than b_i where j ranges from 0 to $i-2$; there are $\binom{i-2}{j}$ ways to pick these bids and 2^{n-i} ways to pick the bids that are smaller than b_i and for this specific position the coefficient of b_i is $(j+1)$.

A straightforward calculation shows that $\sum_{j=0}^{i-2} \binom{i-2}{j} (j+1) = i2^{i-3}$, and the claim follows.

□

In order to prove Claim 1 we need the following simple Lemma:

Lemma 4.6.

$$z(b_{j_1},\ldots,b_{j_k}) \geq y(b_{j_1},\ldots,b_{j_k}) - \max\left(0, \max_{t=2,\ldots,k}\left\{\frac{2t-j_t}{t-1}\right\}\right)\cdot y(b_{j_2},\ldots,b_{j_k})$$

Proof. Let b_{j_t} be the optimal price for $(b_{j_1},\ldots,b_{j_k})$, i.e. $t\cdot b_{j_t} = \max\{b_{j_1}, 2b_{j_2},\ldots,kb_{j_k})$. Then

$$\begin{aligned}
z(b_{j_1},\ldots,b_{j_k}) &= (j_t - t)b_{j_t} \\
&= tb_{j_t} - (2t-j_t)b_{j_t} \\
&= tb_{j_t} - \frac{2t-j_t}{t-1}(t-1)b_{j_t} \\
&= y(b_{j_1},\ldots,b_{j_k}) - \frac{2t-j_t}{t-1}(t-1)b_{j_t} \\
&\geq y(b_{j_1},\ldots,b_{j_k}) - \max\left(0, \max_{t=2,\ldots,k}\left\{\frac{2t-j_t}{t-1}\right\}\right)\cdot y(b_{j_2},\ldots,b_{j_k})
\end{aligned}$$

where we need $\frac{2t-j_t}{t-1}$ to be positive for the inequalities to work correctly, which is why we take $\max\left(0, \max_{t=2,\ldots,k}\left\{\frac{2t-j_t}{t-1}\right\}\right)$. □

Ideally we would like to sum all the inequalities resulting from Lemma 2 for all $S \subseteq B_2$ and get Claim 1 right away –the inequalities are clearly of the right form. However such a result is not that straightforward. For example, for $n = 7$ bids if we sum the above inequality for all $S \subseteq B_2$ we are left with the following expression:

$b_2 + b_4 + b_5 + b_6 + b_7 + \max(b_2, 2b_3) + \max(b_2, 2b_4) + \max(b_2, 2b_5) + \max(b_2, 2b_6) + \max(b_2, 2b_7) + \max(b_4, 2b_6) + \max(b_4, 2b_7) + \max(b_5, 2b_6) + \max(b_5, 2b_7) + \max(b_6, 2b_7) + \max(b_2, 2b_3, 3b_4) + \max(b_2, 2b_3, 3b_5) + \max(b_2, 2b_3, 3b_6) + \max(b_2, 2b_3, 3b_7) + \max(b_2, 2b_4, 3b_5) + \max(b_2, 2b_4, 3b_6) + \max(b_2, 2b_4, 3b_7) + \max(b_2, 2b_5, 3b_6) + \max(b_2, 2b_5, 3b_7) + \max(b_2, 2b_6, 3b_7) - 1/3\max(b_4, 2b_5, 3b_6) + 1/3\max(b_4, 2b_6, 3b_7) + \max(b_2, 2b_3, 3b_4, 4b_5) + \max(b_2, 2b_3, 3b_4, 4b_6) + \max(b_2, 2b_3, 3b_4, 4b_7) + \max(b_2, 2b_3, 3b_5, 4b_6) + \max(b_2, 2b_3, 3b_5, 4b_7) + \max(b_2, 2b_3, 3b_6, 4b_7) + \max(b_2, 2b_4, 3b_5, 4b_6) + \max(b_2, 2b_4, 3b_5, 4b_7) + \max(b_2, 2b_4, 3b_6, 4b_7) + \max(b_2, 2b_5, 3b_6, 4b_7) - 1/2\max(b_4, 2b_5, 3b_6, 4b_7) + \max(b_2, 2b_3, 3b_4, 4b_5, 5b_6) + \max(b_2, 2b_3, 3b_4, 4b_5, 5b_7) + \max(b_2, 2b_3, 3b_4, 4b_6, 5b_7) + \max(b_2, 2b_3, 3b_5, 4b_6, 5b_7) + \max(b_2, 2b_4, 3b_5, 4b_6, 5b_7) +$

$\max(b_2, 2\,b_3, 3\,b_4, 4\,b_5, 5\,b_6, 6\,b_7)$

Although this expression includes all $y(S)$ for S that contain b_2 (since these terms cannot be canceled in the telescoping summation) it contains some negative terms. The number of these terms increases for larger n but so do the remaining positive terms like $\max(b_4, 2b_6)$ etc. In fact it can be verified -by exhaustively trying all combinations- that the sum of the positive terms is large enough to cover the negative terms, for $n \leq 12$, for all possible orderings of $b_2, 2b_3, 3b_4$....

In addition to that, numerical simulations suggest that the total number of $-b_i, i = 4, \ldots$ approximately doubles as we go from n to $n+1$ bids; indeed, if the coefficient $\mu = \max_{t=2,\ldots,k} \frac{2t-j_t}{t-1}$ were to stay the same, then we would expect this number to double. However, the addition of another bid leads to a small increase on μ (which gets smaller and smaller as we move to higher bids). Nonetheless the rate of increase seems to go down as n grows and to approach 2.

Moreover, as we said, by moving to larger n's we have more remaining positive terms as well. In fact the number of positive terms with coefficient 1, i.e. the sequences for which $\mu = 0$ has a nice property, that is very easy to prove: the number of positive terms with coefficient 1, starting with b_i when we have n bids is at least F_{n-i+1}, the $n-i+1$ Fibonacci number. To see that look at the numerator of μ; this is $2t - j_t$ and it suggests that, in order to have $\mu \leq 0$ (so that we end up with $0 \cdot y(b_{j_2}, \ldots, b_{j_k})$) it must be that $\forall t : 2t \leq j_t$. Consider now the following example: say we start with b_4 and we have 4 bids, then the only appropriate sequence is (b_4); for 5 bids we still have only (b_4); for 6 bids we get (b_4) and (b_4, b_6); for 7 bids $(b_4), (b_4, b_6), (b_4, b_7)$ and so on. It is easy to see that the allowable sequences at step t include all the allowable sequences at step $t-1$ (which end with b_{t-1}) and all the sequences that result from the allowable sequences at step $t-2$ (ending with b_{t-2}) if we add a b_t at the end. A straightforward induction shows that indeed the number of allowable sequences increases at least as fast as the Fibonacci sequence (in fact faster, since there are even more allowable sequences which we did not take into account here).

This observation suggests that it is quite realistic to expect the number of positive terms to be approximately doubling each time we increase n by 1: we have an increase of almost ϕ even when we ignore the additional terms. This in turn suggests that we may be able to match the increase in the negative terms, which -as we said- also seem to double each time we increase n by 1.

One may notice that there are no sequences starting with b_3. This is due to the fact that μ is exactly 1 for sequences $b_3, \ldots$ and there is exactly one way we can get a $-y(b_3, \ldots)$, namely by a $z(b_2, b_3, \ldots)$, so these terms all cancel out. As we move to higher bids, there are $j_2 - 2$ ways of getting a $-\mu y(b_{j_2}, \ldots, b_{j_k})$, while there is only one appearance of $+y(b_{j_2}, \ldots, b_{j_k})$. So one may lay his hopes on μ that is strictly less than 1 and gets smaller as j_2 increases, to make all $j_2 - 2$ appearances of $-\mu y(b_{j_2}, \ldots, b_{j_k})$ sum up to 1. As the previous example suggests, this is not the case and we are left with some negative terms. So, the next hope is -as we said- to have these terms canceled out with the remaining positive terms. However, even if we manage to pin down the exact rate of increase of both the positive and the negative terms, we are left with one more problem. The problem is that though the remaining terms are not identical, there may be significant canceling out, e.g. $\max(b_2, 2b_4)$ may cancel out with $\max(b_2, 2b_5, 3b_6)$, depending on the values of the bids. Capturing this in an elegant combinatorial proof is not trivial, but seems necessary.

Another issue is the behavior of μ for growing values of n and for different sequences. Notice that μ does not depend on the values of the bids but only on the current partition. One may also notice that μ is very closely connected to the behavior of the random walk defined in the proof of Theorem 3.12. However, we are not sure to what extent a probabilistic argument is useful when trying to prove 4-competitiveness; it seems that there is a subtle balance of the positive and negative terms that needs careful handling. The goal is to bound μ in a way that allows to argue that the sum of the negative terms does not exceed the sum of the corresponding positive ones.

One final observation is that Claim 2 and Lemma 2, become tight for the case of input bids $(h+\epsilon, h, 0, \ldots, 0)$ as they ought to.

4.2.1 The fight against the remaining negative terms: May 2009

We started trying to find a way to handle the negative coefficients. As mentioned above, maybe the first problem we must handle is the heterogeneity of the various negative terms; what we we need is to find an expression that bounds the total number of negative terms from above, i.e. for the example with $n = 7$ bids above, to find a function of the bids $f(b_4, \ldots, b_7)$ such that $f(b_4, \ldots, b_7) \geq -1/3 \max(b_4, 2\,b_5, 3\,b_6) - 1/2 \max(b_4, 2\,b_5, 3\,b_6, 4\,b_7)$, **for all values** of $b_4, \ldots, b_n$. One obvious choice for f is to pick

$$f(b_4, \ldots, b_n) = \sum_{neg.\ terms\ r} \mu_r \cdot \max\{b_4, 2b_5, \ldots, (n-3)b_n\} \ , \tag{4.1}$$

where μ_r is the ratio $\max_{t=2,\ldots,k} \frac{2t-j_t}{t-1}$ for the specific sequence of bids r.

However this bound is too crude: we tried to see if the positive terms are enough to cancel out $(i-3)b_i \sum \mu_r$ for all $i = 3 \ldots n$ and this seemed to be the case for up to $n = 16$ bids. By positive terms we mean all terms that appear with a positive sign and do not contain a b_2 in the $\max(\cdot)$; i.e. we even included terms with fractional coefficients (whose exact contribution is much harder to compute than for the ones with coefficient 1). Furthermore, in our computation we also use the fact that $b_4 \geq \ldots \geq b_n$, to bound b_i with b_j-terms, $j < i$, if needed.

In the next figure we show the difference between the total positive and the negative coefficients of each b_i for $n = 14, \ldots, 17$ bids. This difference is increasing for larger bids; however for $n = 17$ bids, for the first time, the negative coefficient for b_5 is larger than the corresponding positive one.

In the next figures we show the behavior of the actual total positive vs negative coefficients for all bids, for the case of $n = 14$ and 17 bids.

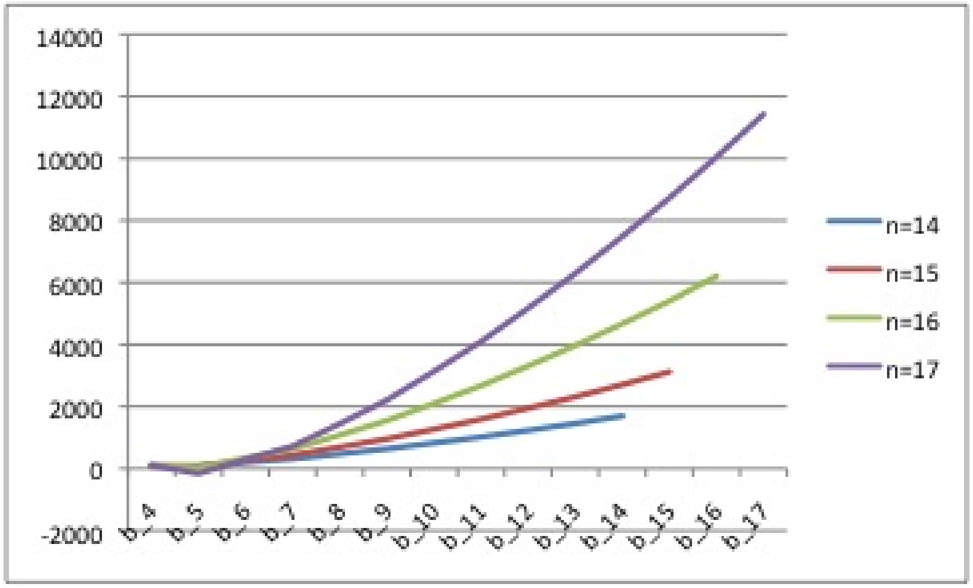

Figure 4.1. Difference between positive and negative coefficients, for $n = 14, 15, 16, 17$ bids

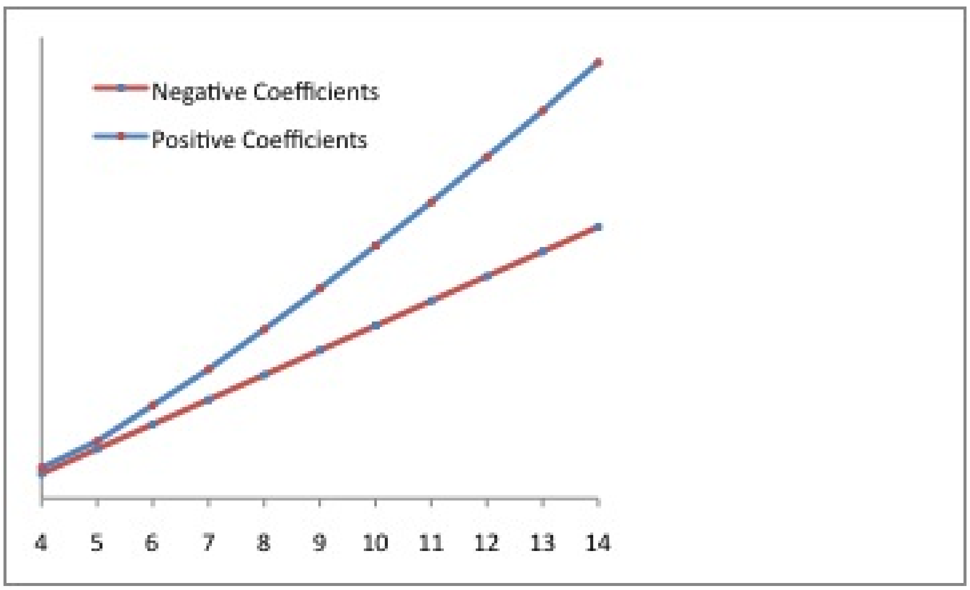

Figure 4.2. Total number of positive vs negative coefficients for $n = 14$

These simulations suggest that the upper bound of 4.1 should not work. However it gives us some good information on the behavior of the negative coefficients. Currently we are trying to exploit this information and use some recursive argument (some sort of induction on the number of bids n) to bound the negative coefficients; we want to base this on the observation that by increasing

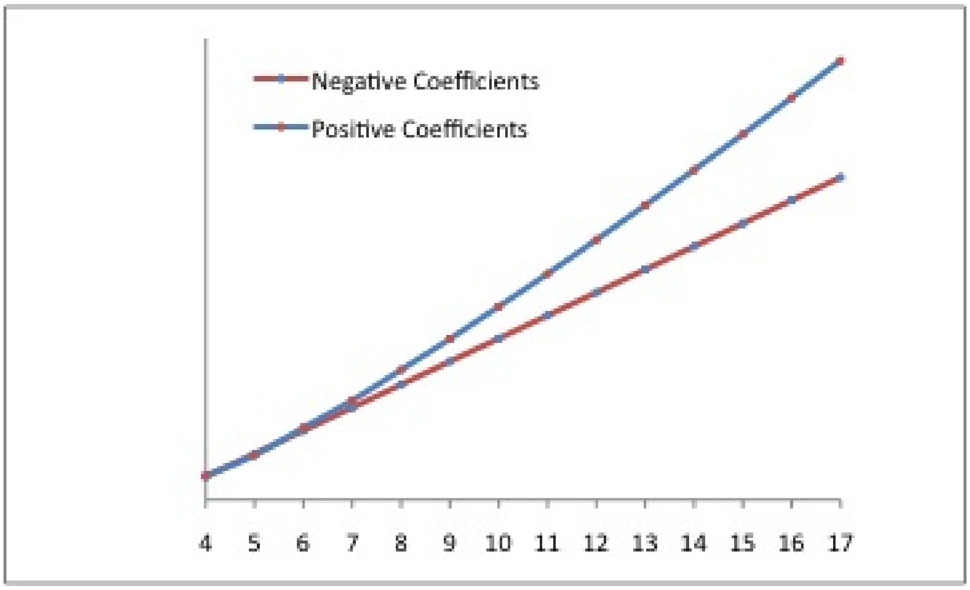

Figure 4.3. Total number of positive vs negative coefficients for $n = 17$

the number of bids by 1, we get new negative terms of two kinds: one that results by just adding the new bid b_n to the end of all previous negative sequences and another that corresponds to sequences with consecutive terms in the end (eg for $n = 9$ we get $(b_4, b_6, b_7, b_8, b_9)$, while for $n = 8$ there was no (b_4, b_6, b_7, b_8)).

4.2.2 A more involved bound - the "lambdamu" bound: June 2009

As trying to bound the total number of negative sequences was a much harder task than expected, we started thinking of ways to improve our bound in a way that would not yield any negative sequences at all in the end. Here is a motivating example: consider the case of $n = 7$ bids of the previous section; notice that we have a $-1/3\max(b_4, 2b_5, 3b_6)$ term. This term could well be bounded with a better (?) use of $z(b_4, b_5, b_6)$: what we get now is $z(b_4, b_5, b_6) \geq y(b_4, b_5, b_6)$; we could also get $z(b_4, b_5, b_6) \geq \frac{4}{3}y(b_4, b_5, b_6) - \frac{1}{2}y(b_5, b_6)$, with $-\frac{1}{2}y(b_5, b_6)$ canceling out with the remaining $y(b_5, b_6)$. In general the idea would be to use a bound of the form $z(b_{j_1}, \ldots, b_{j_k}) \geq \lambda y(b_{j_1}, \ldots, b_{j_k}) + \mu y(b_{j_2}, \ldots, b_{j_k})$, with λ chosen carefully, so as to cancel out all existing negative terms of the form $-cy(b_{j_1}, \ldots, b_{j_k})$ that will occur (these are in fact easy to be counted).

The problem with this approach is that we cannot have an arbitrarily large λ; it holds that $\lambda \leq j_1 - 1$ for any sequence $(b_{j_1}, \ldots, b_{j_k})$. The following lemma gives a "lambdamu" type of bound that is as close to optimal as possible (or so we think).

Lemma 4.7. *Let $(b_{j_1}, \ldots, b_{j_k})$ be a set of at least 2 bids and λ a real in $[0, j_1 - 1]$. We can bound $z(b_{j_1}, \ldots, b_{j_k})$ with*

$$z(b_{j_1}, \ldots, b_{j_k}) \geq \lambda y(b_{j_1}, \ldots, b_{j_k}) + \mu y(b_{j_2}, \ldots, b_{j_k}), \tag{4.2}$$

where μ is defined by

$$\mu = \begin{cases} \frac{k}{k-1} \min_{t=1,\ldots,k}\{\frac{j_t - t - \lambda t}{t}\}\}, & \textit{when } \min_{t=2,\ldots,k}\{j_t - t - \lambda t\} \geq 0 \\ \min_{t=2,\ldots,k}\{\frac{j_t - t - \lambda t}{t-1}\}, & \textit{otherwise} \end{cases}$$

Proof. Assume that

$$y(b_{j_1}, \ldots, b_{j_k}) = t \cdot b_{j_t}$$
$$y(b_{j_2}, \ldots, b_{j_k}) = (s-1) \cdot b_{j_s}$$

From these we get that $tb_{j_t} \geq sb_{j_s}$ and $(s-1)b_{j_s} \geq (t-1)b_{j_t}$. Notice that the latter holds even when $t = 1$.

Assume that $\min_{r=2,\ldots,k}\{j_r - r - \lambda r\} \geq 0$. We will show that inequality (4.2) is satisfied for $\mu = \frac{k}{k-1} \min_{r=1,\ldots,k}\{\frac{j_r - r - \lambda r}{r}\}\}$. We will use the fact that μ is nonnegative and the inequality $tb_{j_t} \geq sb_{j_s}$. Indeed we have,

$$\begin{aligned} \lambda y(b_{j_1}, \ldots, b_{j_k}) + \mu y(b_{j_2}, \ldots, b_{j_k}) &= \lambda t b_{j_t} + \mu(s-1) b_{j_s} \\ &\leq \lambda t b_{j_t} + \mu(s-1)\frac{t}{s} b_{j_t} \\ &\leq \lambda t b_{j_t} + \mu(k-1)\frac{t}{k} b_{j_t} \\ &\leq \lambda t b_{j_t} + \frac{k}{k-1}\frac{j_t - t - \lambda t}{t}(k-1)\frac{t}{k} b_{j_t} \end{aligned}$$

$$= (j_t - t)b_{j_t}$$

$$= z(b_{j_1}, \ldots, b_{j_k})$$

Now we consider the case of $\min_{r=2,\ldots,k}\{j_r - r - \lambda r\} < 0$. Assume first that $t \geq 2$. We will now show that inequality (4.2) is satisfied for $\mu = \min_{r=2,\ldots,k}\{\frac{j_r - r - \lambda r}{r-1}\}\}$. We will use the fact that μ is now negative and the inequality $(t-1)b_{j_t} \leq (s-1)b_{j_s}$. Indeed we have,

$$\begin{aligned}
\lambda y(b_{j_1}, \ldots, b_{j_k}) + \mu y(b_{j_2}, \ldots, b_{j_k}) &= \lambda t b_{j_t} + \mu(s-1)b_{j_s} \\
&\leq \lambda t b_{j_t} + \mu(t-1)b_{j_t} \\
&\leq \lambda t b_{j_t} + \frac{j_t - t - \lambda t}{t-1}(t-1)b_{j_t} \\
&= (j_t - t)b_{j_t} \\
&= z(b_{j_1}, \ldots, b_{j_k})
\end{aligned}$$

The case $t = 1$ must be handled separately because $t - 1$ appears in the denominator in the above. When $t = 1$ we have that

$$\begin{aligned}
z(b_{j_1}, \ldots, b_{j_k}) &= (j_1 - 1)b_{j_1} \\
&\geq \lambda b_{j_1} + (j_1 - 1 - \lambda)b_{j_1} \\
&\geq \lambda y(b_{j_1}, \ldots, b_{j_k}) \\
&\geq \lambda y(b_{j_1}, \ldots, b_{j_k}) + \mu y(b_{j_2}, \ldots, b_{j_k}).
\end{aligned}$$

Notice that we used the fact that $\lambda \leq j_1 - 1$ and that $\mu \leq 0$. $\square$

By using this bound for $n = 15$ bids we get the following (scary) bound:

$y([5]) + 9/2\,y([6]) + 17/2\,y([7]) + 13\,y([8]) + 18\,y([9]) + 24\,y([10]) + \frac{61}{2}\,y([11]) + 38\,y([12]) + 2/3\,y([5,7]) + y([5,8]) + y([5,9]) + y([5,10]) + y([5,11]) + y([5,12]) + 1/2\,y([6,7]) + 7/2\,y([6,8]) + 9/2\,y([6,9]) + 9/2\,y([6,10]) + 9/2\,y([6,11]) + 9/2\,y([6,12]) + 4\,y([7,8]) + 7\,y([7,9]) + 8\,y([7,10]) + 17/2\,y([7,11]) + 17/2\,y([7,12]) + \frac{91}{12}\,y([8,9]) + \frac{31}{3}\,y([8,10]) + \frac{35}{3}\,y([8,11]) + 13\,y([8,12]) + \frac{38}{3}\,y([9,10]) + \frac{43}{3}\,y([9,11]) + 16\,y([9,12]) + 17\,y([10,11]) + 19\,y([10,12]) +$

$22\,y\,([11,12])+1/2\,y\,([5,7,9])+2/3\,y\,([5,7,10])+2/3\,y\,([5,7,11])+2/3\,y\,([5,7,12])+1/2\,y\,([5,8,9])+ y\,([5,8,10])+y\,([5,8,11])+y\,([5,8,12])+y\,([5,9,10])+y\,([5,9,11])+y\,([5,9,12])+y\,([5,10,11])+ y\,([5,10,12])+y\,([5,11,12])+\frac{5}{18}\,y\,([6,7,9])+1/2\,y\,([6,7,10])+1/2\,y\,([6,7,11])+1/2\,y\,([6,7,12])+ \frac{13}{12}\,y\,([6,8,9])+3\,y\,([6,8,10])+7/2\,y\,([6,8,11])+7/2\,y\,([6,8,12])+3\,y\,([6,9,10])+\frac{15}{4}\,y\,([6,9,11])+ 9/2\,y\,([6,9,12])+\frac{15}{4}\,y\,([6,10,11])+9/2\,y\,([6,10,12])+9/2\,y\,([6,11,12])+\frac{59}{18}\,y\,([7,8,10])+ 4\,y\,([7,8,11])+4\,y\,([7,8,12])+\frac{21}{5}\,y\,([7,9,10])+6\,y\,([7,9,11])+7\,y\,([7,9,12])+6\,y\,([7,10,11])+ 7\,y\,([7,10,12])+7\,y\,([7,11,12])+\frac{23}{60}\,y\,([8,9,10])+\frac{371}{60}\,y\,([8,9,11])+\frac{91}{12}\,y\,([8,9,12])+\frac{31}{4}\,y\,([8,10,11])+ 9\,y\,([8,10,12])+9\,y\,([8,11,12])+\frac{173}{60}\,y\,([9,10,11])+\frac{659}{60}\,y\,([9,10,12])+11\,y\,([9,11,12])+ \frac{398}{45}\,y\,([10,11,12])+2/5\,y\,([5,7,9,11])+1/2\,y\,([5,7,9,12])+2/5\,y\,([5,7,10,11])+2/3\,y\,([5,7,10,12])+ 2/3\,y\,([5,7,11,12])+2/5\,y\,([5,8,9,11])+1/2\,y\,([5,8,9,12])+2/5\,y\,([5,8,10,11])+4/5\,y\,([5,8,10,12])+ 4/5\,y\,([5,8,11,12])+2/5\,y\,([5,9,10,11])+4/5\,y\,([5,9,10,12])+4/5\,y\,([5,9,11,12])+4/5\,y\,([5,10,11,12])+ \frac{3}{20}\,y\,([6,7,9,11])+\frac{5}{18}\,y\,([6,7,9,12])+\frac{3}{20}\,y\,([6,7,10,11])+1/2\,y\,([6,7,10,12])+1/2\,y\,([6,7,11,12])+ \frac{14}{15}\,y\,([6,8,9,11])+\frac{13}{12}\,y\,([6,8,9,12])+6/5\,y\,([6,8,10,11])+\frac{27}{10}\,y\,([6,8,10,12])+\frac{27}{10}\,y\,([6,8,11,12])+ 6/5\,y\,([6,9,10,11])+\frac{27}{10}\,y\,([6,9,10,12])+\frac{27}{10}\,y\,([6,9,11,12])+\frac{27}{10}\,y\,([6,10,11,12])+\frac{9}{40}\,y\,([7,8,10,11])+ \frac{43}{15}\,y\,([7,8,10,12])+\frac{43}{15}\,y\,([7,8,11,12])+\frac{19}{40}\,y\,([7,9,10,11])+\frac{19}{5}\,y\,([7,9,10,12])+\frac{62}{15}\,y\,([7,9,11,12])+ \frac{62}{15}\,y\,([7,10,11,12])+\frac{307}{120}\,y\,([8,9,11,12])+\frac{191}{60}\,y\,([8,10,11,12])-4/5\,y\,([7,8,9,10,11])- \frac{19}{15}\,y\,([8,9,10,11,12])-2\,y\,([7,8,9,10,11,12])-2/7\,y\,([6,7,8,9,10,11,12])$

Just notice that there are only 4 negative terms and their coefficients are not very large and they can certainly be bounded by some of the resulting positive terms. Further experiments for larger values of n suggested that the negative terms do not grow too fast, and that they could in fact be bounded by the remaining positive ones for all values of n (and all values of bids). However proving something like that did not seem any easier than what we had in the previous section (rather more involved).

4.3 Getting beaten to the result: end of June 2009

At the time of the completion of this thesis, a paper appeared in EC 2009 that partially settled Conjecture 4.1. The paper [7] by Alaei, Malekian and Srinivasan proves that if the optimal single sale price for the full set of bids is at least the sixth largest bid then RSOP is indeed 4-competitive, while in the remaining cases it is 4.68-competitive. This paper thus leaves a very small gap between the lower and the upper bounds.

Their proof is based on the following idea: we need some lower bounding function on the profit of RSOP from each side of the partition (A, B). Wlog of generality we can assume that the profit from side A (where b_1 falls) is zero, that we have an infinite number of bids, with all bids larger than b_n being zero and that the optimal single sale price is $OPT = 1$. A very important parameter is the index of the winning bid in the full set of bids, denoted by λ.

Let us use S_i to denote the number of bidders larger than b_i on side A of the partition, i.e. $S_i = \sharp\{j | j \leq i, b_j \in A\}$. The authors make two main observations:

1. The optimal profit from side A of the partition is $S_{\lambda_A} b_{\lambda_A}$ which is at least $S_\lambda b_\lambda$. Since we assumed that $OPT = 1$, $\lambda b_\lambda = 1$, so $b_\lambda = \frac{1}{\lambda}$, so the profit from side A, when offered its optimal single sale price b_{λ_A}, denoted by $\text{Profit}(A, b_{\lambda_A})$, is at least $\frac{S_\lambda}{\lambda}$:

$$\text{Profit}(A, b_{\lambda_A}) \geq \frac{S_\lambda}{\lambda}$$

2. The profit extracted from side B of the partition when offered the optimal single sale price for side A is $\text{Profit}(B, b_{\lambda_A})$ and we have

$$\text{Profit}(B, b_{\lambda_A}) = \frac{\lambda_A - S_{\lambda_A}}{S_{\lambda_A}} \text{Profit}(A, b_{\lambda_A})$$

   It is hence obvious that the quantity of interest is $z_i = \frac{i - S_i}{S_i}$. Since λ_A is a quantity that depends on the actual values of the bids, the authors, instead of working with z_{λ_A} they work with the quantity $z = \min_i z_i$ which is a lower bound for z_{λ_A} and only depends on the partition of the bids and not their actual values. Hence we have:

$$\frac{\text{Profit}(B, b_{\lambda_A})}{\text{Profit}(A, b_{\lambda_A})} \geq z$$

Using the above two observations, a lower bounding function on the expected profit of RSOP is straightforward:

$$\begin{aligned}
\mathbb{E}[RSOP] &\geq \mathbb{E}[\text{Profit}(B, b_{\lambda_A})] \\
&= \mathbb{E}\left[\text{Profit}(A, b_{\lambda_A}) \frac{\text{Profit}(B, b_{\lambda_A})}{\text{Profit}(A, b_{\lambda_A})}\right] \\
&\geq \mathbb{E}[\frac{S_\lambda}{\lambda} z]
\end{aligned}$$

What is really important about the expression $\mathbb{E}[\frac{S_\lambda}{\lambda} z]$ is that it depends on the actual values of the bids in only one way: the value of λ; other than that it is a function of the partition of the bids only. To achieve that we had to do two things: first, in observation 1 we had to bound $\text{Profit}(A, b_{\lambda_A})$ with $S_\lambda b_\lambda$ rather than $S_{\lambda_A} b_{\lambda_A}$ and in observation 2 we had to take the $\min_i z_i$ to get rid of the dependance on λ_A. Both these assumptions do not seem to hurt the tightness of the analysis by a large factor.

The proof then proceeds as follows: ideally we would like to split $\mathbb{E}[\frac{S_\lambda}{\lambda} z]$ in $\mathbb{E}[\frac{S_\lambda}{\lambda}]\mathbb{E}[z]$ and bound the two terms separately (which is much easier to do). However the two terms are clearly correlated and we cannot do that; the key observation is that the correlation decreases for larger values of λ: remember from Section 4.1 that we have an implicit random walk on a line and we are interested on its relative offset at some time t; as this time grows we expect this ratio to come arbitrarily close to 1/2. Now, if this is the case, say for $\lambda > 5000$, then they can decompose $\mathbb{E}[\frac{S_\lambda}{\lambda} z]$ in $\mathbb{E}[\frac{S_\lambda}{\lambda}]\mathbb{E}[z]$ minus some correlation terms, and show that the competitive ratio is bounded by 4. For the case $10 \leq \lambda \leq 5000$ they employ a combination of probabilistic techniques and dynamic programming to show that for this case too the competitive ratio is 4. Their proof is highly technical and is omitted.

The authors finally address the case where $\lambda < 10$. A similar –but even more technical– approach that involves exhaustive search over all possible **values** of the bids for these values of λ yields a competitive ratio of 4 for $6 \leq \lambda \leq 10$ and 4.68 for $2 \leq \lambda \leq 5$.

Let us conclude this section by trying to briefly compare this proof with the approach in [6] and our approach. The common feature of [6] and [7] is that they both try to bound the profit extracted by one side of the partition when offered the optimal price of the other partition, with the profit extracted from the other side when offered the same price. Moreover, they both bound this profit by some fraction of the optimal profit ($\frac{S_\lambda}{\lambda}$ in [7] and 1/2 in [6]) – note that the simplicity of the [6]-bound is largely due to the fact that the authors there just use the simple fact that with probability 1/2 profit $\text{Profit}(A, b_{\lambda_A})$ is at least $OPT/2$, while in [7] the authors use the most general –and much harder to analyze– $\frac{S_\lambda}{\lambda}$ bound. The second observation is also common: the ratio $z_i = \frac{i-S_i}{S_i}$ is the exact same quantity bounded in [6] for the random walk defined there. In fact in both papers, in the end of the day, the authors need a lower bound on $\mathbb{E}\left[\min_i \frac{i-S_i}{S_i}\right]$. However in [6] they first prove a bound on the probability of $\frac{i-S_i}{S_i}$ getting too large (see Lemma 4.3) and use it to bound the expectation, while in [7] they shoot for a direct bound on $\mathbb{E}[\min_i z_i]$, which is a much harder task. One could say that the ideas behind [7] follow very closely those at [6], with the difference of a much more tight and elaborate analysis of the bounds, which differs at the points mentioned above.

Our approach bares a lot in common with both approaches; as we mentioned already we make an implicit use of the $\frac{i-S_i}{S_i}$-quantity, and in fact of $\min_i \frac{i-S_i}{S_i}$. We believe that taking the min over all i should not hurt us too much. Our main goal is to derive a lower bound on $\text{Profit}(B, b_{\lambda_A})$ (which we call $z(S)$), as in [7]; the main difference is that we make no assumption whatsoever about the actual values of the bids (while in [7] they have λ to address this issue). This is also the reason why our analysis is heavily affected by the total number of bids n. We believe –at this point– that taking into account some property of the winning bid in the full set of bids (such as λ) could be unavoidable. Our final –when this thesis was filed– attempt to address the problem allows for some experimentation towards this direction as well; we briefly discuss that in the next section.

4.4 Future attempts: July 2009-...

We conclude the thesis with yet another bounding approach; this approach is based on the following observation (which in some sense gets us back to the beginning: all we assume for now is that Claim 4.4 holds and we try to prove it).

Ideally what we would like is to have $z(b_{j_1},\ldots,b_{j_k}) \geq y(b_{j_1},\ldots,b_{j_k})$ for all sequences with $b_{j_1} = b_2$. Unfortunately this does not hold, and what we get is

$$\sum_{S\subseteq B_2:\ b_2\in S} z(S) \geq \sum_{S\subseteq B_2:\ b_2\in S} y(S) - \sum_{i=2\ldots k} r_i b_{j_i}$$

where $r_i \geq 0$ are some resulting negative coefficients which we want to keep at a minimum. Our goal then is to bound these coefficients using the remaining

$$\sum_{S\subseteq B_2:\ b_2\notin S} z(S)$$

In order to do that we also need some appropriate bound on $z(S)$ for any S not containing b_2; by appropriate we mean some bound that is linear in the bids, i.e. it does not contain any $\max(\cdot)$, etc terms.

For the $z(S)$, with $b_2 \in S$ we use the following bound:

$$\sum_{S\subseteq B_2:\ b_2\in S} z(S) \geq \sum_{S\subseteq B_2:\ b_2\in S} y(S) - \sum_{i=2\ldots k} (2i - j_i - \phi_{i-1}) b_{j_i}$$

where $\phi_1 = 0$ and $\phi_i = \max(\phi_{i-1}, 2i - j_i)$.

For example we get that $z(b_2, b_3, b_4, b_5) \geq y(b_2, b_3, b_4, b_5) - b_3 - b_4 - b_5$.

It is easy to see that this bound is "good", in the sense that it does not subtract a bid unless it really has to; the way of achieving that is to take into account all previous (larger) bids that have already been subtracted; this is exactly what ϕ_i is taking into account.

For the case of S with $b_2 \notin S$ we believe that we somehow need to take into account the index of the winning bid in the full set of bids in order to derive a good bound. The exact way of doing that is something we still have to think about.

References

[1] Noam Nisan and Amir Ronen. Algorithmic mechanism design (extended abstract). In *STOC*, pages 129–140, 1999.

[2] Roger B. Myerson. Optimal auction design. Discussion Papers 362, Northwestern University, Center for Mathematical Studies in Economics and Management Science, December 1978.

[3] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

[4] Jason D. Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *STOC*, pages 75–84, 2008.

[5] Andrew V. Goldberg, Jason D. Hartline, and Andrew Wright. Competitive auctions and digital goods. In *SODA*, pages 735–744, 2001.

[6] Uriel Feige, Abraham Flaxman, Jason D. Hartline, and Robert D. Kleinberg. On the competitive ratio of the random sampling auction. In *WINE*, pages 878–886, 2005.

[7] S. Alaei, A. Malekian, and Aravind Srinivasan. On random sampling auctions for digital goods. In *EC*, 2009.