Lectures on Optimal Mechanism Design

Jason Hartline*

Lectures: October 27 – November 10, 2005
Notes: April 28, 2006; June 9, 2006.

These lecture notes cover the second third of the class CS364B, Topics in Algorithmic Game Theory, offered at Stanford University in the Fall 2005 term. They cover the topic of *optimal mechanism design*. As this is a traditional economic objective, we will review the Economics treatment of optimal mechanism design first before moving on to cover recent work on the problem from the theoretical computer science community. Prerequisites for reading these lecture notes are basic understanding of algorithms and complexity as well as elementary calculus and probability theory. I will also assume that the reader has access to the notes on the first third of this course which covers *combinatorial auctions*.

Thanks to the students of CS364B, Ning Chen, and co-instructor Tim Roughgarden. Comments are welcome.

1 Introduction

In the first third of this course we looked at the problem of designing a mechanism to maximize the *social surplus*, the sum of the valuations of the participants. A second, natural, objective to consider is that of maximizing the profit of mechanism. Where as the Vickrey auction maximizes the social surplus by awarding the item to the bidder with the highest valuation, an auctioneer might prefer a mechanism which maximizes their profit instead. Profit maximization is one of fundamental objectives in mechanism design. We adopt terminology from Economics and refer to the mechanism that maximizes the seller's profit as the *optimal mechanism*.[1]

In this course we will cover two approaches to optimal mechanism design. The first is the conventional economics approach of *Bayesian optimal mechanism design* where it is assumed that the valuations of the agents are drawn from a known distribution and the Bayesian optimal mechanism is the one that achieves the largest expected profit for these agents. Here the expectation is taken over the randomness in the agents' valuations. The problem

*Microsoft Research – Silicon Valley, 1065 La Avenida, Mountain View, CA 94043. Email: hartline@microsoft.com.

[1] A mechanism that maximizes the social surplus, we call the *efficient mechanism*.

of optimal mechanism design here is to determine how to take distributional knowledge and use it to construct the optimal mechanism.

The second approach we will consider is reminiscent of the typical Computer Science approach to optimization (and approximation) where we look for a mechanism that is optimal (or near optimal) for any possible valuations the agents may have. This approach is especially appealing as the worst-case guarantees we derive for any agent valuations are comparable to the optimal average-case guarantees for valuations from known distributions. The techniques we discuss for worst-case optimal mechanism design follow from techniques typical for algorithmic design and analysis.

We now briefly motivate the problem of optimal mechanism design by considering the profit of the Vickrey auction for selling a single item. Recall that the Vickrey auction sells the item to the highest bidder at the second highest bid value. Thus, the profit of the Vickrey auction is the second highest bid value. Consider running the Vickrey auction in a setting where there are two bidders who's valuations are known to be drawn independently at random from the uniform distribution on $[0, 1]$.

What follows is a simple argument showing that the expected revenue of the Vickrey auction in this setting is $1/3$. This follows from the calculation of the second order statistic of two uniform random variables on $[0, 1]$. Let v_1 and v_2 denote the valuations of the bidders. By the definition of uniform random variables, $\mathbf{Pr}[v_i < x] = x$ implying that $\mathbf{Pr}[v_i > x] = 1 - x$. Since v_1 and v_2 are independent, $\mathbf{Pr}[\min(v_1, v_2) > x] = \mathbf{Pr}[v_1 > x \wedge v_2 > x] = (1 - x)^2$. Using Fact 1.2, below, about the expectation of non-negative random variables we have

$$\begin{aligned}\mathbf{E}[\min(x_1, x_2)] &= \int_0^\infty \mathbf{Pr}[\min(v_1, v_2) > x]\, dx \\ &= \int_0^1 (1-x)^2 dx = \int_0^1 (1 - 2x + x^2) dx = x - x^2 + x^3/3 \Big|_0^1 = 1/3.\end{aligned}$$

A natural question to ask at this point is whether it is possible to do better and how. Consider the following definition of the Vickrey auction with reservation price r for selling a single item. This auction is of interest, for example, if the seller has valuation r for retaining the item. It is instructive to think of the Vickrey auction with reserve price r as being implemented by running the Vickrey auction with an additional bid placed on behalf of the seller with value r. If this bid wins, then the seller keeps the item. The truthfulness of this auction follows trivially from this viewpoint.

Definition 1.1 (Vickrey Auction with reservation price r) The *Vickrey auction with reservation price* r, VA_r, sells the item if any bidder bids above r. The price the winning bidder pays the maximum of the second highest bid and r.

We now show that the seller can increase their expected profit in our example scenario by pretending to have a valuation of $1/2$ for the item and running the Vickrey auction with reservation price $r = 1/2$. In particular when v_i are independent and identically distributed uniformly on $[0, 1]$, a reservation price of $1/2$ in the Vickrey auction yields an expected profit of 5/12 which is more than 1/3. To show this we consider three cases based on the randomization in the bidders' valuations.

Case 1 (both bids are less than 1/2): This happens with probability 1/4 and in this case $\mathrm{VA}_{1/2}$ does not sell the item and has no profit.

Case 2 (both bids are above 1/2): This happens with probability 1/4 and in this case the expected profit is the expected value of the second highest value. This is the second order statistic of random variables that are uniform on $[1/2, 1]$ which can be calculated, similarly to before, to be 2/3.

Case 3 (otherwise, one bid is less than 1/2 and the other is above 1/2): This happens with probability 1/2 and the expected profit is the reservation price 1/2.

The expected profit of $\mathrm{VA}_{1/2}$ is thus, $1/4 \times 0 + 1/4 \times 2/3 + 1/2 \times 1/2 = 5/12$.

We have thus shown that given prior knowledge of the distribution from which the bidders' valuations are drawn, it is possible to give an auction with higher expected profit than the Vickrey auction. In the next section we will show that for this particular example $\mathrm{VA}_{1/2}$ is the optimal auction and we will give a general technique for finding the optimal auction in more complex scenarios. In the subsequent sections, we discuss doing optimal mechanism design without foreknowledge of the distribution from which the bidders' valuations are drawn.

We finish this section by giving a simple proof of a fact about expectations that will be useful throughout these lectures.

Fact 1.2 *A non-negative random variable X satisfies $\mathbf{E}[X] = \int_0^\infty \mathbf{Pr}[X \geq z]\, dz$.*

Proof: Let $f_X(z)$ be the density function for X. The expectation of X is thus:

$$\mathbf{E}[X] = \int_{y=0}^{\infty} y f_X(y) dy$$

Using the fact that $y = \int_{z=0}^{y} dz$ and then switching the order of integration, we have:

$$\begin{aligned}
\mathbf{E}[X] &= \int_{y=0}^{\infty} \int_{z=0}^{y} f_X(y) dz dy \\
&= \int_{z=0}^{\infty} \int_{y=z}^{\infty} f_X(y) dy dz \\
&= \int_{z=0}^{\infty} \mathbf{Pr}[X \geq z]\, dz.
\end{aligned}$$

■

2 Bayesian Optimal Mechanism Design

In this section we solve the Bayesian optimal mechanism design problem in a very general setting. For this setting we give a simple characterization of truthful mechanisms and then a simple description of the optimal mechanism.

2.1 Single Parameter Agents

Suppose we have n agents. each of whom desires some good or service. We assume that agent i's *valuation* for receiving service is v_i and their valuation for no service we normalize to zero. The mechanism designer, will solicit sealed bids from the agents and compute an outcome which consists of an allocation $\mathbf{x} = (x_1, \ldots, x_n)$ and prices $\mathbf{p} = (p_1, \ldots, p_n)$. Assume that $x_i = 1$ represents agent i being allocated service whereas $x_i = 0$ for no service. In addition to receiving service or not, agent i is required to pay the mechanism p_i. We assume that agents have quasi-linear *utility* expressed by $u_i = v_i x_i - p_i$. Thus, an agent's goal in a mechanism is to maximize the difference between their valuation and their payment.

To make this setting quite general, we assume that there is an inherent cost in producing the outcome $c(\mathbf{x})$ which must be payed by the mechanism. The two most fundamental economic objectives are that of social *surplus*, a.k.a., common welfare, and *profit*. These quantities are defined as:

$$\begin{aligned} \text{Surplus} &= \sum\nolimits_i v_i x_i - c(\mathbf{x}), \text{ and} \\ \text{Profit} &= \sum_i p_i - c(\mathbf{x}). \end{aligned}$$

While in the preceding third of this course when discussing combinatorial auctions we were interested in the *efficient* outcome, i.e, the one that maximizes the social surplus, we now consider the problem of maximizing the mechanism's profit.

Example 2.1 (single item auction)

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

Example 2.2 (single-minded combinatorial auction, known bundles) Let S_i be the desired bundle of bidder i, then

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

Example 2.3 (multicast auctions) Given a graph with costs $c(e)$ on edges, and a designated "root" node. $c(\mathbf{x})$ is the total cost of connecting all of the agents with $x_i = 0$ to the root (I.e., the minimum Steiner tree cost).

We will place two standard assumptions on our mechanisms. The first, that they are *individually rational*, means that no agent has negative expected utility for playing. This condition is alternatively known as *voluntary participation* and essentially it requires that $p_i \leq b_i x_i$. The second condition we require is that of *no positive transfers* which restricts the mechanism to not pay the agents to play. I.e, $p_i \geq 0$.

2.2 Characterization of Truthful Mechanisms

Before presenting this characterization, we extend our analysis to consider randomized mechanisms as well. In a randomized mechanism we view x_i as the probability that agent i is allocated the good, and p_i as their expected payment. Recall that x_i and p_i are results of the mechanism. As such, it will be useful to view them as functions of the input bids as follows. We let $x_i(\mathbf{b})$, $p_i(\mathbf{b})$, and $u_i(\mathbf{b})$ represent agent i's probability of allocation, expected price, and expected utility, respectively. Let $\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, ?, b_{i+1}, \ldots, b_n)$ represent the vector of bids disincluding bid i. Then with $\mathbf{b}_{-i}$ fixed, we let $x_i(b_i)$, $p_i(b_i)$, and $u_i(b_i)$ represent the agent i's probability of allocation, expected price, and expected utility, respectively, as a function of their own bid. We further define the convenient notation $x_i(b_i, \mathbf{b}_{-i}) = x_i(\mathbf{b})$, $p_i(b_i, \mathbf{b}_{-i})$, and $u_i(b_i, \mathbf{b}_{-i})$.

Definition 2.4 A mechanism is *truthful* if and only if for all i, v_i, b_i, and $\mathbf{b}_{-i}$, agent i's utility for bidding their valuation, v_i, is at least their utility for bidding any other bid, b_i. I.e.,

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

The following results were shown by Myerson in his seminal paper on optimal mechanism design [11]. Our proofs will follow those of Archer and Tardos [1] which is a bit more accessible to computer scientists.

Theorem 2.5 *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. *$p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z)dz$.*

Given this theorem, we see that for a fixed allocation rule, $\mathbf{x}(\cdot)$, the payment $\mathbf{p}(\cdot)$ rule is also fixed. Thus, in specifying a mechanism we need only specify a monotone allocation rule and from this the correct payment rule for truthfulness can be derived.

We break the proof of this theorem into two parts, the forward direction "if", and the backwards direction, "only if".

Proof:($\Longleftarrow$) We prove this by picture. Though the formulaic proof is simple, the pictures provide useful intuition. We consider two possible bid values z_1 and z_2 with $z_1 < z_2$. Supposing $v_i = z_2$ we argue that agent i does not benefit by shading their bid down to z_1. We leave the proof of the opposite, that agent i does not benefit by shading their bid up, as an exercise for the reader.

To start with this proof, we assume that $x_i(b_i)$ is monotone and that $p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z)dz$. Recall that the utility of agent i is $u_i(b_i) = v_i x_i(b_i) - p_i(b_i)$.

Consider the diagrams in Figure 1. The first diagram (top, center) shows $x_i(\cdot)$ which is indeed monotone as per our assumption. The column on the left show the bidder's valuation, $v_i x_i(v_i)$; payment, $p_i(v_i)$, and utility, $u_i(v_i) = v_i x_i(v_i) - p_i(v_i)$, assuming that they bid their true valuation, $b_i = v_i = z_2$. The column on the right shows the analogous quantities when the bidder bids $b_i = z_1 < z_2 = v_i$. The final diagram (bottom, center) shows the difference in the bidders utility when bidding $b_i = v_i = z_2$ and bidding $b_i = z_1 \leq z_2 = v_i$. Note that as the picture shows, the monotonicity of the allocation function implies that this difference is always non-negative and that indeed there is no incentive for the bidder to shade their bid to $z_1 < z_2 = v_i$. The reader can draw and verify the analogous pictures to show that there is also no incentive for a bidder with valuation $v_i = z_1$ to shade their bid up to $b_i = z_2$. These two facts combine to show that monotonicity of the allocation rule with the payment rule specified in the lemma statement imply truthfulness.

($\Longrightarrow$) First we show that truthful implies monotone. If a mechanism is truthful then for all valuations, v_i, and possible lies, b_i, $u_i(v_i) \geq u_i(b_i)$. Expanding $u_i(\cdot)$ we require:

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

We now consider z_1 and z_2 and take turns setting $v_i = z_1$, $b_i = z_2$, and $b_i = z_1$, $v_i = z_2$. This yields the following two inequalities:

$$\begin{aligned} v_i = z_2, b_i = z_1 &\Longrightarrow z_2 x_i(z_2) - p_i(z_2) \geq z_2 x_i(z_1) - p_i(z_1), \text{ and} \\ v_i = z_1, b_i = z_2 &\Longrightarrow z_1 x_i(z_1) - p_i(z_1) \geq z_1 x_i(z_2) - p_i(z_2). \end{aligned}$$

Adding these inequalities and canceling the payment terms we have,

$$z_2 x_i(z_2) + z_1 x_i(z_1) \geq z_2 x_i(z_1) + z_1 x_i(z_2).$$

Rearranging,

$$(z_2 - z_1)(x_i(z_2) - x_i(z_1)) \geq 0.$$

Since $z_2 - z_1 > 0$ we see that $x_i(z_2) - x_i(z_1) \geq 0$. This holds for all $z_2 > z_1$ which is precisely the condition that $x_i(\cdot)$ is monotone non-decreasing.

Now we show that the payment rule must satisfy $p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z)dz$. Fix v_i and recall that $u_i(b_i) = v_i x_i(b_i) - p_i(b_i)$. Let $u_i{}'(\cdot)$ be the partial derivative of $u_i(\cdot)$ with respect to b_i (recall that $u_i(\cdot)$ is also a function of v_i). Thus, $u_i{}'(b_i) = v_i x_i{}'(b_i) - p_i{}'(b_i)$, where $x_i{}'(\cdot)$ and $p_i{}'(\cdot)$ are the derivatives of $p_i(\cdot)$ and $x_i(\cdot)$, respectively. Since truthfulness implies that $u_i(b_i)$ is maximized at $b_i = v_i$. It must be that

$$\begin{aligned} 0 &= u_i{}'(v_i) \\ &= v_i x_i{}'(v_i) - p_i{}'(v_i). \end{aligned}$$

This formula must hold true for all values of v_i. For remainder of the proof, we treat this identity formulaicly. To emphasize this, substitute $z = v_i$:

$$0 = z x_i{}'(z) - p_i{}'(z).$$

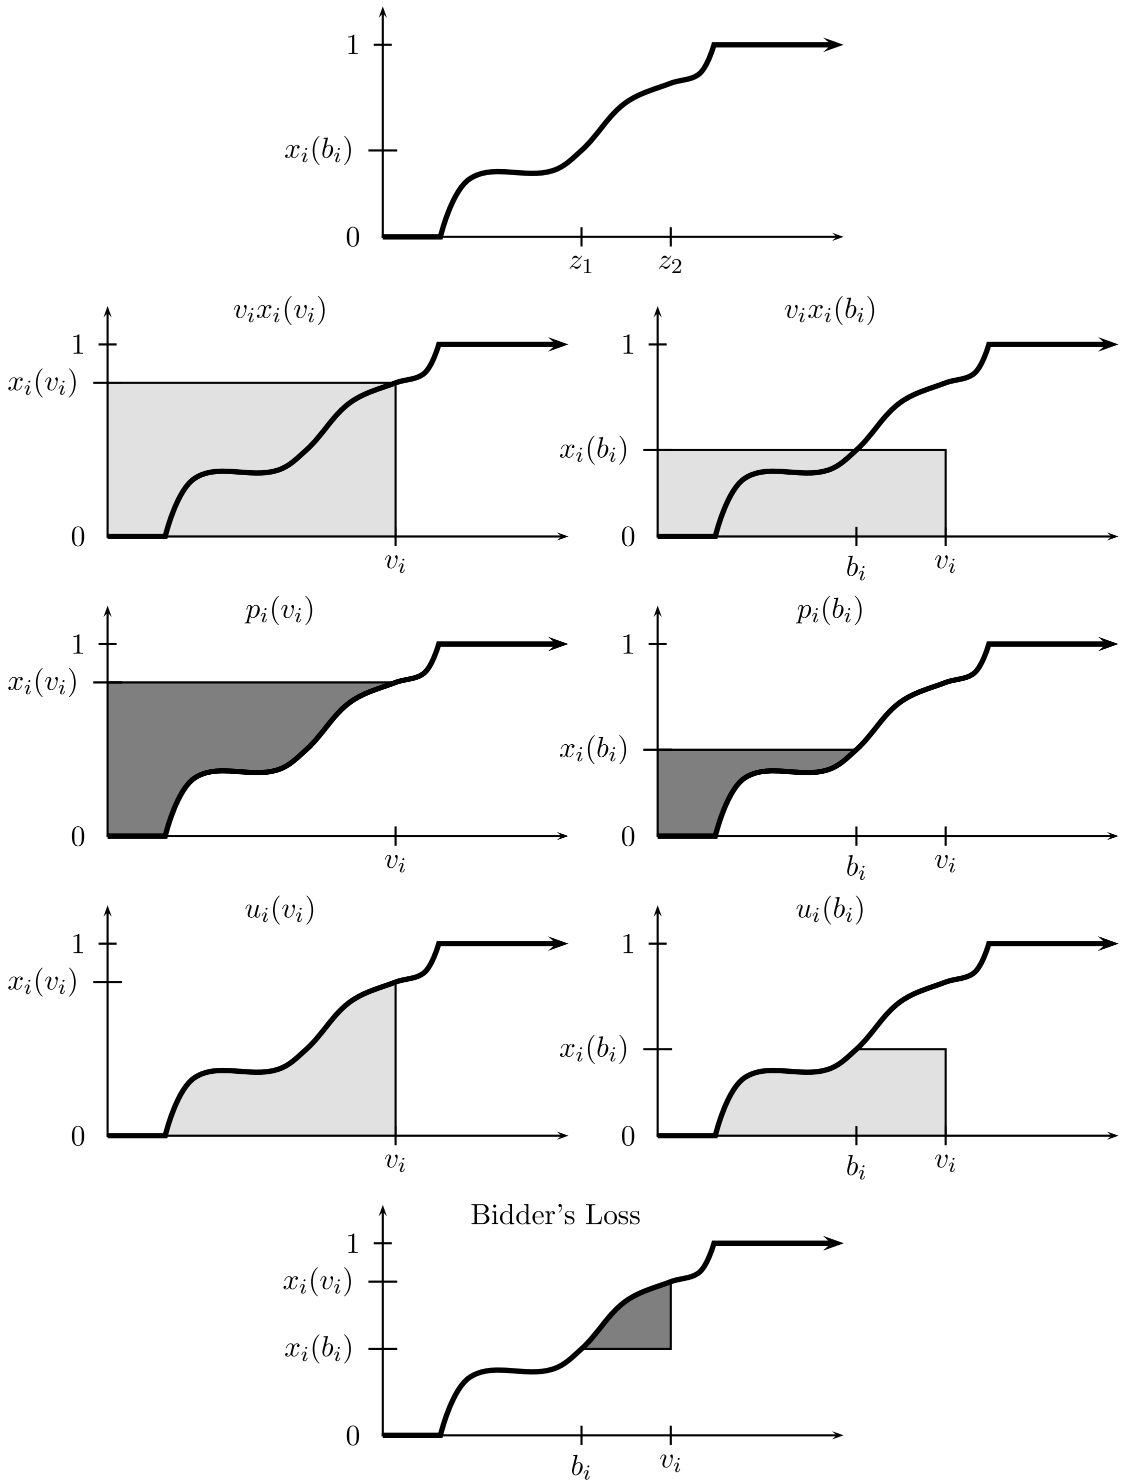

Figure 1: The left column shows the valuation, payment, and utility of the bidder when bidding $b_i = v_i = z_2$. The right column shows the valuation, payment, and utility of the bidder when bidding $b_i z_1 < v_i = z_2$. The final diagram shows the difference between the bidder's utility when bidding their valuation and less than it.

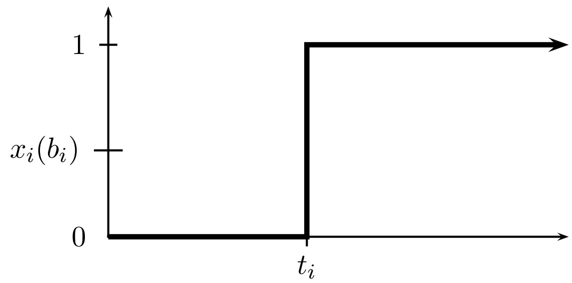

Figure 2: The deterministic allocation with minimum winning bid t_i.

Solving for ${p_i}'(z)$ and then integrating both sides of the equality from 0 to b_i we have,

$$\begin{aligned} {p_i}'(z) &= z{x_i}'(z), \text{ so} \\ \int_0^{b_i} {p_i}'(z)dz &= \int_0^{b_i} z{x_i}'(z)dz, \\ p_i(b_i) - p_i(0) &= zx_i(z)|_0^{b_i} - \int_0^{b_i} x_i(z)dz \\ &= b_i x_i(b_i) - \int_0^{b_i} x_i(z)dz. \end{aligned}$$

Adding $p_i(0)$ from both sides of the equality, we arrive at a valid equation for the payment rule $p_i(\cdot)$ given the allocation rule $x_i(\cdot)$. We take this one step further by assuming voluntary participation and no positive transfers. The former implies that $p_i(b_i) \leq b_i$ (so $p_i(0) \leq 0$) and the latter implies that $p_i(b_i) \geq 0$. Thus, $p_i(0) = 0$ and we have our desired theorem. ■

It is useful to reinterpret this theorem for deterministic mechanisms. For deterministic mechanisms $x_i(\cdot) \in \{0, 1\}$. Therefore, monotonicity implies that $x_i(\cdot)$ is a step function. Furthermore, if an agent wins, then their payment is exactly their "minimum winning bid", i.e, $t_i = \inf_z\{z : x_i(z) = 1\}$ (Figure 2). As an exercise reprove Theorem 2.5 (and derive the analogous diagrams to Figure 1) for the deterministic special case.

As an example, let us consider the goal of maximizing the social surplus. Recall that the surplus is given by the formula $\sum_i v_i x_i - c(\mathbf{x})$. The Vickrey-Clarke-Groves [12, 5, 6] (VCG) mechanism that we considered previously for combinatorial auctions also applies in general single-parameter agent settings. As before, the VCG mechanism outputs the allocation that maximizes the surplus (assuming truthful bidding). To see that this is truthful, we must only verify that it results in a monotone allocation rule. We can then use this monotone allocation rule to calculate payments using the formula for payments: $p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z)dz$.

Lemma 2.6 *The surplus maximizing allocation is monotone.*

Proof: Define $S(\mathbf{v}, \mathbf{x}) = \sum_i v_i x_i - c(\mathbf{x})$. Let $\mathbf{v}'$ be the vector of valuations where agent i's valuation is increased by $\delta \geq 0$ but valuations of other agents remain unchanged. I.e.,

${v_i}' = v_i + \delta$ and for $j \neq i$, ${v_j}' = v_j$. Clearly then,

$$\begin{aligned} S(\mathbf{v}', \mathbf{x}) &= \sum_i {v_i}' x_i - c(\mathbf{x}) \\ &= \delta x_i + \sum_i v_i x_i - c(\mathbf{x}) \\ &= \delta x_i + S(\mathbf{v}, \mathbf{x}). \end{aligned}$$

From this equation it is easy to see that if $\mathbf{x}$ maximizes $S(\mathbf{v}, \mathbf{x})$ with $x_i = 1$, then all $\mathbf{x}'$ that maximize $S(\mathbf{v}', \mathbf{x})$ also have $x_i = 1$. ■

As an exercise, give a simple expression for the truthful payment rule of the VCG mechanism for any single-parameter agent problem.

2.3 Myerson's (Simple) Optimal Mechanism

Now assume that the valuations of the agents, $v_1, \ldots, v_n$, are drawn independently (but not necessarily identically) at random from some distribution. We let F_i denote the distribution from which bidder i's valuation, v_i, is drawn and $\mathbf{F} = F_1 \times \ldots \times F_n$, the distribution from which $\mathbf{v}$ is drawn, is just the product distribution. We adopt the following notation:

Definition 2.7 The *cumulative distribution function* for v_i is $F_i(z) = \mathbf{Pr}[v_i < z]$.

Definition 2.8 The *probability density function* for v_i is $f_i(z) = \frac{d}{dz} F_i(z)$.

We now define the notion of *virtual valuations* and *virtual surplus*. These definitions follow Myerson [11].

Definition 2.9 The *virtual valuation* of agent i with valuation v_i is

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

Definition 2.10 Given valuations v_i, virtual valuations $\phi_i(v_i)$, and allocation $\mathbf{x}$, the *virtual surplus* is

$$\sum_i \phi_i(v_i) x_i - c(\mathbf{x}).$$

We now state the main result of this section. Note that this result implies that maximizing profit is equivalent to maximizing virtual surplus.

Theorem 2.11 *[11] The expected profit of a mechanism is equal to its expected virtual surplus.*

Definition 2.12 *Myerson's optimal mechanism* outputs $\mathbf{x}$ to maximize the virtual surplus.

One intuitive way to view this optimal mechanism is via the following algorithm:

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 $\mathbf{p}$ with $p_i = {\phi_i}^{-1}({p_i}')$.

Of course, we still have to ask ourselves whether Myerson's optimal mechanism is truthful. This is equivalent to the question of whether the allocation rule is monotone: if a winner increases their bid do they still win? Since surplus maximization is monotone in agent valuations (Lemma 2.6), virtual surplus maximization is monotone in agent virtual valuations. Thus, we can conclude that if virtual valuations are monotone in valuations, then virtual surplus maximization is monotone in valuations. It is simple to verify that examples exist such that non-monotone virtual valuation functions result in a non-monotone allocation rule which implies non-truthfulness.

Lemma 2.13 *Myerson's optimal mechanism is truthful if and only if $\phi_i(v_i)$ is monotone in v_i for all i.*

The condition that virtual valuations be monotone is consistent with a common assumption made in Economics: the *monotone hazard rate assumption*. The hazard rate is defined as $f(z)/(1 - F(z))$. Myerson discusses an "ironing" procedure for handling the case where the virtual valuations are not monotone. Interested readers should consult his original paper [11] or alternative presentations (E.g., [2]).

Lemma 2.14 *The expected payment of a bidder satisfies (as a function of their bid b_i and with all other bids fixed $\mathbf{b}_{-i}$):*

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

Proof: (drop the subscript i) The bid b is random from distribution F with density function f. From the definition of expectation and truthfulness we have:

$$\begin{aligned}
\mathbf{E}_b[p(b)] &= \int_0^h p(b)f(b)db \\
&= \int_{b=0}^h bx(b)f(b)db - \int_{b=0}^h \int_{z=0}^b x(z)f(b)dzdb.
\end{aligned}$$

focusing on the second term and switching the order of integration, we have

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

Now, rename z to b and factor out a factor of $x(b)f(b)$.

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

■

Myerson's theorem (Theorem 2.11) follows from linearity of expectation and the fact that the bidders' valuations are independently distributed (i.e., the joint distribution is a product distribution).

2.4 Interpretation of Myerson's Optimal Mechanism

We now consider a few examples which allow us to obtain a precise interpretation of Myerson's optimal mechanism in a few important special cases.

Example 2.15 (single item auction)

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

In the single item auction the surplus maximizing allocation allocates the item to the bidder with the highest valuation, unless the highest valuation is less than zero in which case the auctioneer keeps the item. This point of not allocating the item if all bidders values are less than zero is usually not relevant because it is usually assumed that the bidders valuations are all at least zero. For deriving the Bayesian optimal single item auction Myerson says to find the virtual surplus maximizing allocation and virtual valuations can be negative (see the example below). The optimal auction, therefore, allocates the item to the bidder with the largest positive virtual valuation.

For the case that there are only two bidders, bidder 1 wins precisely when $\phi_1(b_1) \geq \max\{\phi_2(b_2, 0\}$. This is a deterministic allocation rule, and thus the payment that a winning bidder 1 must make is the $p_1 = \inf\{b : \phi_1(b) > \phi_2(b_2) \wedge \phi_1(b) > 0\}$. Suppose that $F_1 = F_2 = F$ which implies that $\phi_1(z) = \phi_2(z) = \phi(z)$. Then we can simplify the price bidder 1 must make upon winning to $p_1 = \inf\{b : b > b_2 \wedge b > \phi^{-1}(0)\}$. Bidder 2's payment upon winning is $p_2 = \inf\{b : b > b_1 \wedge b > \phi^{-1}(0)\}$. Thus we arrive at Myerson's main observation.

Theorem 2.16 *The optimal single-item auction for bidders with valuations drawn i.i.d. from distribution F is precisely the Vickrey auction with reservation price $\phi^{-1}(0)$.*

In our example above when F is uniform on $[0, 1]$, we can plug the equations $F(z) = z$ and $f(z) = 1$ into the formula for the virtual valuation function (Definition 2.9) to conclude that $\phi(z) = 2z - 1$. Thus, the virtual valuations are uniformly distributed on $[-1, 1]$. We can easily solve for $\phi^{-1}(0) = 1/2$. Thus, we can conclude that the optimal auction for our example: two bidders with valuations uniform on $[0, 1]$ is the Vickrey auction with reservation price $1/2$.

The reader should now be prepared to answer the following questions:

1. What is the optimal single-item auction when the seller values the item at $v_0 > 0$?

2. What is the optimal digital good auction for bidders with i.i.d. valuations drawn uniformly from $[0, 1]$? from distribution F?

3. What is the optimal auction for a seller with k identical items and $n > k$ bidders with i.i.d. valuations drawn uniformly from $[0, 1]$?

In this section we have seen how to design the optimal truthful auction. The reader might be wondering whether this auction is optimal amongst all auctions including possibly non-truthful auctions. In order to consider this question we must fix an equilibrium concept describing how bidders will behave in a non-truthful mechanism (recall that this issue is precisely the issue we were trying to avoid when we adopted the solution concept of truthfulness). It turns out that for natural equilibrium concepts and given the assumption that the bidders are *risk neutral* there is a revenue equivalence theorem that states that all auctions which yield the same outcome in equilibrium also yield the same expected profit. In particular, the first price auction, the Vickrey auction, and even the *all-pay auction*[2] Interested readers are referred to Krishna's book [9], Klemperer's survey [8], or Myerson's original paper [11].

3 Worst Case Optimal Mechanism Design

We start of with a motivating problem that we will use as a running example throughout this section.

Example 3.1 (Digital Good Auction) In the *digital good auction* there are,

1. n bidders with valuations $v_1, \ldots, v_n$.

2. n units of an item for sale.

We would like to design a truthful auction that maximizes our profit. However, first let us consider how we might design an algorithm (not truthful) for maximizing our profit.

[2] In the all-pay auction, all bidders must pay their bids, yet the item is only allocated to the bidder with the highest bid. Thus, the all-pay auction is not ex post individually rational. It is, however, *ex interim* individually rational (i.e., when bidders know the distribution of other bidders valuations, and their own true valuation, they have no incentive not to participate in the auction).

3.1 Digital Good Algorithms

When designing an algorithm, we make the classical algorithmic assumption that the algorithm is given the true input. In this case, the true input is the bidders' true valuations. The first obvious algorithm for this problem is the following.

Algorithm 3.1 ($\mathcal{T}$) *"Sell to all bidders at their valuation."*

$$\textit{Profit} = \mathcal{T}(\mathbf{v}) = \sum\nolimits_i v_i.$$

We might object to this algorithm because it is unfair in that winning bidders are required to pay different prices. This might cause us to prefer the second most obvious algorithm for this problem. First a definition that will be useful throughout our discussion of digital good auctions.

Definition 3.2 Let the notation $v_{(i)}$ represent the ith largest valuation, i.e., $v_{(1)} \geq v_{(2)} \geq \cdots \geq v_{(n)}$.

Algorithm 3.2 ($\mathcal{F}$) *"Sell at optimal single sale price."*

$$\textit{Profit} = \mathcal{F}(\mathbf{v}) = \max_p p \times \text{"number of bidders with } v_i \geq p.\text{"}$$

1. *Sort valuations:* $v_{(1)} \geq \ldots \geq v_{(n)}$.
2. *Compute* $i^* = \operatorname{argmax}_i iv_{(i)}$.
3. *Output sale price* $p = v_{(i^*)}$.

We now consider the relationship between $\mathcal{T}$ and $\mathcal{F}$. Obviously $\mathcal{T}(\mathbf{v}) \geq \mathcal{F}(\mathbf{v})$. For a bound in the other direction, we show the following lemma where $\mathcal{H}_n = \sum_{i=1}^{n} 1/i$ denotes the nth harmonic number.

Lemma 3.3 *$\mathcal{F}(\mathbf{v}) \geq \mathcal{T}(\mathbf{v})/\mathcal{H}_n$ and this is tight.*

Proof: Recall that $v_{(i)}$ is the ith highest bid value and that $\mathcal{F}(\mathbf{v}) = \max_i iv_{(i)}$. We have:

$$\mathcal{T}(\mathbf{v}) = \sum_i v_i = \sum_i v_{(i)} = \sum_i \frac{iv_{(i)}}{i} \leq \sum_i \frac{\mathcal{F}(\mathbf{v})}{i} = \mathcal{F}(\mathbf{v})\mathcal{H}_n.$$

For tightness, consider the valuations, $v_i = 1/i$. This input is known as the *equal revenue* input because as $iv_{(i)} = 1 = \mathcal{F}(\mathbf{v})$ for all i. Naturally, $\mathcal{T}(\mathbf{v}) = \mathcal{H}_n$. ■

Since $\mathcal{H}_n \approx \ln n$ we see that $\mathcal{F}(\mathbf{v})$ is $\Theta(\mathcal{T}(\mathbf{v})/\log n)$ in worst case.

3.2 Truthfulness in Auctions

In the previous section on Bayesian optimal auctions we saw that given the allocation rule $\mathbf{x}(\mathbf{b})$ we could determine the payment rule $\mathbf{p}(\mathbf{b})$. Of course, the opposite is true as well. In any deterministic auction, the outcome for a bidder is fully specified by an offer price t_i where the bidder wins if they bid at least t_i and they lose if they bid below t_i. Implicit in the definition of this offer price is that all other bids $\mathbf{b}_{-i}$ are fixed. It is useful to write this offer price as a function of the other bids $t_i(\mathbf{b}_{-i})$. This view point gives us the following observation about auctions for digital goods.

Observation 3.1 *Any truthful auction should determine an offer price to bidder i that is a function of all other bids,* $\mathbf{b}_{-i}$*.*

3.3 Worst-case Optimality in Auctions

Recall that when the bidders were drawn from a known distribution there was a very natural notion of what it meant to be the optimal auction. For a given distribution there exists a truthful auction that obtains the highest expected profit. This auction is optimal! The same is not true in worst case settings.

Claim 3.4 *There does not exist an auction that it best on every input.*

Proof: We will show this proof by example. Consider one bidder in an auction. In a truthful auction the offer price t_1 (for bidder 1) must not be a function of her bid. However, there are no other bids, therefore, the offer price must just be a constant. Consider two possible inputs, Case 1 with $b_1 = 1$ and Case 2 with $b_1 = 2$; and two possible auctions, $\mathcal{A}_1$ with $t_1 = 1$ and $\mathcal{A}_2$ with $t_1 = 2$. Consider the following table with the profit of the auction each auction on each input.

	$b_1 = 1$	$b_1 = 2$
$\mathcal{A}_1$	1	1
$\mathcal{A}_2$	0	2

It is easy to see that $\mathcal{A}_1$ is the unique optimal auction in Case 1 and suboptimal in Case 2; whereas, $\mathcal{A}_2$ is the unique optimal auction in case 2 and suboptimal in Case 1. ■

This impossibility leaves us with the question of how we can arrive at a rigorous theoretical framework in which we can determine some auction to be optimal. The key to resolving this issue is in moving from absolute optimality to relative optimality. Indeed, whenever there is an information theoretic obstacle or computational intractability preventing us from obtaining an absolute optimal solution to a problem we can try to approximate. For example, in the design of polynomial time algorithms for NP-hard problems, the goal is to design an algorithm that computes a solution that approximates objective value of the true optimal solution. Likewise, in the design of online algorithms the objective is to find an online algorithm that performs comparably to an optimal offline algorithm. The notable analogy here is between the game theoretic constraint that a mechanism does not know the true bid values in advance and must solicit them in a truth-inducing manner, and the online constraint that an online algorithm does not have knowledge of the future.

Thus, the general approach will be to try to design an auction that is always (in worst case) within a small constant factor of some *benchmark profit*. In the analysis online algorithms it is not always best to gauge the performance of an online algorithm with an unconstrained optimal offline algorithm, likewise, we may choose in the analysis of digital good auctions to compare an auction's profit to some *profit benchmark* that is not necessarily the profit of the optimal algorithm that is given the bidders' true valuations in advance.

Definition 3.5 A *profit benchmark* is a function $\mathcal{G} : \mathbb{R}^n \to \mathbb{R}$ which maps a vector of valuations to a target profit.

We have already seen two profit benchmarks, $\mathcal{F}(\mathbf{v}) = \max_i iv_{(i)}$ and $\mathcal{T}(\mathbf{v}) = \sum_i v_i$, that will be of interest.

The following definition captures the intuition that an auction is good if it is always close to a reasonable profit benchmark.

Definition 3.6 The *competitive ratio* of auction $\mathcal{A}$ is $\beta = \max_{\mathbf{v}} \frac{\mathcal{G}(\mathbf{v})}{\mathcal{A}(\mathbf{v})}$.

Given a profit benchmark $\mathcal{G}$ the task of an auction designer is to give the auction that achieves the minimum possible competitive ratio. This auction is the *optimal competitive auction* for $\mathcal{G}$.

Lemma 3.7 *For* $\mathbf{v}$ *with* $v_i \in [1, h]$*, no auction is* $o(\log h)$*-competitive with* $\mathcal{T}(\mathbf{v})$ *for any* n*.*

Proof: Exercise. ■

Corollary 3.8 *For* $\mathbf{v}$ *with* $v_i \in [1, h]$*, no auction is* $o(\log h)$*-competitive with* $\mathcal{F}(\mathbf{v})$ *for* $n = 2$*.*

This follows from Lemma 3.7 because for $n = 2$, $\mathcal{F}(\mathbf{v}) \geq \mathcal{T}(\mathbf{v})/2$. Basically, this impossibility is a result of the fact that we cannot get one high bidder to pay a constant fraction of their true valuation. There are two solutions to this problem that have been considered in the literature. The first, which we will consider now, is to compare an auction's profit to the profit of the optimal single sale price that sells at least two items. The second approach is to try to perform well compared to $\mathcal{F}$ less an additive loss term that is $O(h)$ for bids in the range $[1, h]$. We consider this second approach when discussing online auctions in Section 4.

Definition 3.9 ($\mathcal{F}^{(2)}$) The *optimal single priced profit with at least two winners* is

$$\mathcal{F}^{(2)}(\mathbf{v}) = \max_{i \geq 2} iv_{(i)}.$$

For an auction $\mathcal{A}$ that is β-competitive with $\mathcal{F}^{(2)}$ for some small constant β, we have:

$$\mathcal{T}(\mathbf{v}) \geq \mathcal{F}(\mathbf{v}) \approx \mathcal{F}^{(2)}(\mathbf{v}) \geq \mathcal{A}(\mathbf{v}) \geq \mathcal{F}^{(2)}(\mathbf{v})/\beta \geq \mathcal{T}(\mathbf{v})/\beta \log n.$$

3.4 Competing with $\mathcal{F}^{(2)}$

For the remainder of this section we take up the goal of designing an auction that is β-competitive with $\mathcal{F}^{(2)}$ for some constant β. We proceed with a digression into the structure of $\mathcal{F}^{(2)}$. In the following examples we see that the optimal sale price depends on the distribution of input valuations. In both cases, we would like to design an auction $\mathcal{A}$ with revenue $\mathcal{A}(\mathbf{b}) \approx \mathcal{F}^{(2)}(\mathbf{b})$.

Example 3.10 With

- 50 bids \$10 and
- 50 bids at \$1

consider the revenue R_1 and R_{10} from an offer price of \$1 and \$10, respectively:

- $R_{10} = 500$ and
- $R_1 = 100$ (because all 100 bidders win and pay \$1).

Clearly, $\mathcal{F}^{(2)}(\mathbf{b}) = R_{10} = \$500$.

Example 3.11 With

- 5 bids at \$10 and
- 95 bids at \$1

consider the revenue R_1 and R_{10} from an offer price of \$1 and \$10, respectively:

- $R_{10} = 50$ and
- $R_1 = 100$.

Clearly, $\mathcal{F}^{(2)}(\mathbf{b}) = R_1 = \$100$.

It us useful here to consider the profit of the single-priced mechanism on a given set of bidders, $\mathbf{b}$, as a function of its offer price. There is some price at which this function has an optimal value, this value is $\mathcal{F}(\mathbf{b})$ (and if there are at least two winners then this value is $\mathcal{F}^{(2)}(\mathbf{b})$ as well).

Definition 3.12 (opt**)** The *optimal price* for input valuations $\mathbf{v}$ is

$$\text{opt}(\mathbf{v}) = \text{argmax}_p \, p \times \text{“number of bidders with } v_i \geq p.\text{”}$$

As described in the beginning of this section, any truthful auction can be described by an offer price $t_i(\mathbf{b}_{-i})$ for each bidder that is not a function of their bid value (and possibly a function of the other bidders bid values). One very natural auction to consider is the *Deterministic Optimal Price* auction (DOP).

Definition 3.13 (DOP) The *Deterministic Optimal Price* auction offers each bidder

$$t_i(\mathbf{b}_{-i}) = \text{opt}(\mathbf{b}_{-i}).$$

Clearly, by definition DOP is truthful. How about its performance? Is DOP β-competitive for any constant β? The answer to this question is no as the following example illustrates.

Example 3.14 With

- 10 bids at \$10 and
- 90 bids at \$1

consider the prices $t_i(\mathbf{b}_{-1})$ and $t_i(\mathbf{b}_{-10})$ that DOP offers bidders bidding \$1 and \$10 respectively. \$10, respectively:

- $\mathbf{b}_{-1}$ is 89 bids at \$1 and 10 bids at \$10, so $\text{opt}(\mathbf{b}_{-1}) = \$10$, and
- $\mathbf{b}_{-10}$ is 90 bids at \$1 and 9 bids at \$10, so $\text{opt}(\mathbf{b}_{-10}) = \$1$.

Thus, bids at \$10 are accepted, but offered price \$1, while bids at \$1 are rejected. The total profit is \$10 when $\mathcal{F}^{(2)} = 100$. This example can be made arbitrarily bad.

3.5 Competitive Auctions

We start this section by extending the bad example for DOT, above, to show that no deterministic symmetric auction is constant-competitive. An auction is symmetric if its outcome is not a function of the order in which the bids are input.

Theorem 3.15 *No symmetric deterministic truthful auction is competitive with* $\mathcal{F}^{(2)}$.

Proof: We consider only bid vectors with bids $b_i \in \{1, h\}$. Let $n_h(\mathbf{b})$ and $n_1(\mathbf{b})$ represent the number of h bids and 1 bids in $\mathbf{b}$, respectively. That an auction $\mathcal{A}$ is symmetric implies that $t_i(\mathbf{b}_{-i})$ is independent of i and only a function of $n_h(\mathbf{b}_{-i})$ and $n_1(\mathbf{b}_{-i})$. Thus, we can let $t(n_h, n_1)$ represent the offer price of $\mathcal{A}$ for any bidder i when we plug in $n_h = n_h(\mathbf{b}_{-i})$ and $n_1 = n_1(\mathbf{b}_{-i})$. Finally we assume that $t(n_h, n_1) \in \{1, h\}$ as this restriction cannot hurt the auction profit on the bid vectors we are considering.

We assume for a contradiction that the auction is constant competitive and proceed in three steps.

1. Observe that for any auction that is competitive, it must be that for all m, $t(m, 0) = h$. Otherwise, on the all h's input, the auction only achieves profit n while the optimal profit is $\mathcal{F}^{(2)} = hn$. Thus, the auction would be at most h-competitive which is not constant.

2. Likewise, observe that for any auction that is competitive, it must be that for all m, $t(0, m) = 1$. Otherwise, on the all 1's input, the auction achieves no profit and is clearly not competitive with $\mathcal{F}^{(2)} = n$.

3. For the final argument, consider taking m sufficiently large and looking at $t(k, m-k)$. As we have argued for $k = 0$, $t(k, m-k) = 0$. Consider increasing k until $t(k, m-k) = h$. This must occur since $t(k, m-k) = h$ when $k = m$. Let $k^* = \min\{k : t(k, m-k) = h\}$ be this transition point. Now consider the $m+1$ bid input to $\mathcal{A}$ with $n_h(\mathbf{b}) = k^*$ and $n_1(\mathbf{b}) = m - k^* + 1$. Consider separately the offer prices to 1-bidders and h-bidders:

   - For 1-bidders: $t(n_h(\mathbf{b}_{-1}), n_1(\mathbf{b}_{-1})) = t(k^*, m - k^*) = h$. Thus, all 1-bidders are rejected and contribute nothing to our auction profit.
   - For h-bidders: $t(n_h(\mathbf{b}_{-h}), n_1(\mathbf{b}_{-h})) = t(k^* - 1, m - k^* + 1) = 1$. Thus, all h-bidders are are offered a price of one which they accept. Thus, the contribution to the auction profit from the h-bidders is $1 \times n_h(\mathbf{b}) = k^*$.

Now, assume we had chosen $h = n$. If this were the case then $\mathcal{F}^{(2)}(\mathbf{b}) = nk^*$ while our auction profit is only k^*. Thus, auction is at best n-competitive. ∎

We can conclude from this theorem that if we wish to design a competitive auction we must either design one that is asymmetric or randomized. Only very recently has it been shown that there exist asymmetric deterministic auctions that are competitive and these auctions are based on derandomizing randomized auction. Thus, we now consider the design of (symmetric) randomized auctions.

Theorem 3.16 *There exist (symmetric) randomized auctions that are constant-competitive against $\mathcal{F}^{(2)}$.*

Consider the following auction:

Definition 3.17 (RSOP) The *Random Sampling Optimal Price* auction (RSOP) works as follows:

- Randomly partition the bids $\mathbf{b}$ into two parts by flipping a fair coin for each bidder and assigning her to $\mathbf{b}'$ or $\mathbf{b}''$.
- Compute $t' = \text{opt}(\mathbf{b}')$ and $t'' = \text{opt}(\mathbf{b}'')$, the optimal sale prices for each part.
- Offer t' to all bidders in $\mathbf{b}''$ and t'' to all bidders in $\mathbf{b}'$.

It is clear as each bidder is offered a price that is not a function of their own bid, and only a function of the other bids, that RSOP is truthful. Further, via a complicated analysis it is possible to prove that RSOP is no worse than 15-competitive. We omit this analysis here. We instead prove that a lower bound on the competitive ratio of RSOP is four. Thus we have:

Lemma 3.18 *RSOP is truthful.*

Theorem 3.19 *The competitive ratio of RSOP is between 4 and 15.*

Proof:(of lower bound only) Consider the two-bid vector $\mathbf{b} = (\$1, \$2)$. With probability 1/2 both bids end up in the same part and the RSOP profit is zero. Otherwise, with probability 1/2 one bid is in each part. Without loss of generality, $\mathbf{b}' = \{\$1\}$ and $\mathbf{b}'' = \{\$2\}$, then $t' = \$1$ and $t'' = \$2$. Thus, the \$1-bid is rejected (because she cannot pay \$2) and the \$2-bid is offered a price of \$1 which she accepts. The RSOP profit in this case is \$1. The expected profit of RSOP is therefore \$0.50 while $\mathcal{F}^{(2)}(\mathbf{b}) = \$2$ which shows that RSOP is at best 4-competitive. ■

We conclude our discussion of RSOP with the conjecture that this two bid bid vector is actually the worst case and that RSOP is indeed 4-competitive.

Conjecture 3.20 *RSOP is 4-competitive.*

3.6 Lower bounds

Now that we have seen that there exists an auction that has constant competitive ratio to $\mathcal{F}^{(2)}$ it is interesting to consider the following questions. What is the *optimal auction* in terms of worst case competitive ratio to $\mathcal{F}^{(2)}$? What is the competitive ratio of this optimal auction? One way to try to answer this question is to look for good auctions; indeed, we have already seen one auction, RSOP, that is 15-competitive. Thus, we can conclude that the competitive ratio of the optimal auction is at most 15. The other direction is to look for lower-bounds on the competitive ratio. In this section we discuss a proof that shows that no auction is better that 2.42-competitive.

Theorem 3.21 *No auction has competitive ratio less than 2.42.*

Instead of proving this theorem which involved a complicated analysis of the expected value of $\mathcal{F}^{(2)}(\mathbf{b})$ when $\mathbf{b}$ generated from a particular distribution, we will show a simpler result which highlights the main ideas of the main theorem.

Lemma 3.22 *No 2-bidder auction is has competitive ratio less than 2.*

Proof: The proof follows a simple structure that is useful for proving lower bounds for this type of problem. First, we consider bids drawn from a random distribution. Second, we argue that $\mathbf{E_b}[\mathcal{A}(\mathbf{b})] \leq \mathbf{E_b}\left[\mathcal{F}^{(2)}(\mathbf{b})\right]/2$. By the definition of expectation this implies that there exists a bid vector $\mathbf{b}^*$ such that $\mathcal{A}(\mathbf{b}^*) \leq \mathcal{F}^{(2)}(\mathbf{b}^*)/2$ (as otherwise the expected values could not satisfy this condition).

We choose a distribution to make the analysis of $\mathbf{E_b}[\mathcal{A}(\mathbf{b})]$ simple. This is important because we have to analyze it for all auctions $\mathcal{A}$. The idea is to choose the distribution for $\mathbf{b}$ such that all auctions obtain the same expected profit. Consider $\mathbf{b}$ with b_i satisfying $\mathbf{Pr}[b_i > z] = 1/z$. Note that whatever price t_i that $\mathcal{A}$ offers bidder i, the expected payment made by bidder i is $t_i \times \mathbf{Pr}[b_i \geq t_i] = 1$. Thus, for $n = 2$ bidders the expected profit of the auction $\mathbf{E_b}[\mathcal{A}(\mathbf{b})] = n = 2$.

We must now calculate $\mathbf{E_b}\left[\mathcal{F}^{(2)}(\mathbf{b})\right]$. $\mathcal{F}^{(2)}(\mathbf{b}) = \max_{i\geq 2} ib_{(i)}$ where $b_{(i)}$ is the ith highest bid value. In the case that $n = 2$, this simplifies to $\mathcal{F}^{(2)}(\mathbf{b}) = 2b_{(2)} = 2\min(b_1, b_2)$. We recall (Fact 1.2) that a non-negative random variable X has $\mathbf{E}[X] = \int_0^\infty \mathbf{Pr}[X > z]\, dz$ and calculate $\mathbf{Pr}\left[\mathcal{F}^{(2)}(\mathbf{b}) > z\right]$.

$$\begin{aligned}
\mathbf{Pr_b}\left[\mathcal{F}^{(2)}(\mathbf{b}) > z\right] &= \mathbf{Pr_b}[b_1 \geq z/2 \wedge b_2 \geq z/2] \\
&= \mathbf{Pr_b}[b_1 \geq z/2]\, \mathbf{Pr_b}[b_2 \geq z/2] \\
&= 4/z^2.
\end{aligned}$$

Note that this equation is only valid for $z \geq 2$. Of course for $z < 2$, $\mathbf{Pr}\left[\mathcal{F}^{(2)}(\mathbf{b}) \geq z\right] = 1$.

$$\begin{aligned}
\mathbf{E_b}\left[\mathcal{F}^{(2)}(\mathbf{b})\right] &= \int_0^\infty \mathbf{Pr}\left[\mathcal{F}^{(2)}\mathbf{b} \geq z\right] dz \\
&= 2 + \int_2^\infty \frac{4}{z^2} dz \\
&= 4.
\end{aligned}$$

Thus we see that for this distribution and any auction $\mathcal{A}$, $\mathbf{E_b}[\mathcal{A}(\mathbf{b})] = 2$ and $\mathbf{E_b}\left[\mathcal{F}^{(2)}(\mathbf{b})\right] = 4$. Thus, the inequality $\mathbf{E_b}[\mathcal{A}(\mathbf{b})] \leq \mathbf{E_b}\left[\mathcal{F}^{(2)}(\mathbf{b})\right]/2$ holds and there must exist some input $\mathbf{b}$ such that $\mathcal{A}(\mathbf{b}) \leq \mathcal{F}^{(2)}(\mathbf{b})/2$. ■

The above proof is based on analysis of $\mathbf{E_b}\left[\mathcal{F}^{(2)}(\mathbf{b})\right]$ for $\mathbf{b}$ of length $n = 2$. This bound can be generalized by a much more complicated analysis for larger n. Such an analysis leads to bounds of 13/6 for $n = 3$ and eventually to a bound of 2.42 for general n. It is conjectured that these bounds are tight. Indeed they are tight for $n \leq 3$.

Exercise 3.1 *Give an auction that is 2-competitive for* $\mathbf{b}$ *of length* $n = 2$*?*

Next we see how to design an auction that is 4-competitive (and is thus a 4/2.42-approximation to the optimal competitive auction).

3.7 The Digital Good Auction Decision Problem

A standard approach to solving problems is by reduction to an "easier" problem or one with a known solution. In what follows we define the notion of a decision problem for mechanism design and reduce the problem of designing a good competitive auction to it. Since we will show that the decision problem is solvable, this will give a solution to the competitive auction problem. This approach is reminiscent of a typical approach for solving optimization problem via reduction to the decision problem.

Consider a classical optimization problem which can be phrased as "find the feasible solution maximizes some objective function." The decision problem version of this is "is there a feasible solution for which the objective function has value at least V." Notice that we can use a solution to the decision problem to solve the optimization problem by doing a binary search over values V. Unfortunately, such an approach will not work for mechanism design as it is not truthful to run several truthful mechanisms and then only take the output of the one that is the most desirable.

Definition 3.23 The *digital good auction decision problem* is given:

- n bidders,
- n copies of an item, and
- a target profit R;

design a mechanism that obtains profit R if possible, i.e., if $R \leq \mathcal{F}(\mathbf{b})$.

This digital good auction decision problem is also known as the *profit extraction* problem as its goal is to extract a profit R from a set of bidders. It turns out that this problem is solved by a special case of a general *cost sharing* mechanism due to Moulin and Shenker [10].

Definition 3.24 (ProfitExtract$_R$) The *basic auction profit extractor* with target profit R sells to the largest group of k bidders that can equally share R and charge each R/k.

ProfitExtract$_R$ can be implemented be the following algorithm which is given n, R, and $\mathbf{b}$.

1. Offer R/n to all bidders.
2. If $b_i \leq R/n$ 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 3.25 Consider:

- $\mathbf{b} = (1, 2, 3, 4, 5)$,
- $R = \$9$, and
- $n = 5$ bidders.

The first iteration offers 9/5 to all bidders. This causes bidder 1 to be rejected since $9/5 > 1$. The second iteration of the algorithm (with $n = 4$) offers $9/4 > 2$ to each bidder and bidder 2 is rejected. The final iteration offers $9/3 = 3$ to each bidder. Since all bidders remaining (which are just $(3, 4, 5)$ can afford to pay \$3, none are rejected and the auction terminates with an offer price of \$3 to each of the three highest bidders.

Exercise 3.2 *Prove that* ProfitExtract$_R$ *is a truthful profit extractor:*

1. *Prove it is truthful.*
2. *Prove it is a profit extractor (that it always obtains a profit of at least $\mathcal{F}(\mathbf{b})$ when $\mathcal{F}(\mathbf{b}) \geq R$).*

3.8 Reduction to the Decision Problem

Definition 3.26 (RSPE) The *Random Sampling Profit Extraction* auction (RSPE) works as follows:

- Randomly partition the bids $\mathbf{b}$ into two by flipping a fair coin for each bidder and assigning her to $\mathbf{b}'$ or $\mathbf{b}''$.
- Compute $F' = \mathcal{F}(\mathbf{b}')$ and $F'' = \mathcal{F}(\mathbf{b}'')$, the optimal sale prices for each part.
- Run ProfitExtract$_{F'}$ on $\mathbf{b}''$ and ProfitExtract$_{F''}$ on $\mathbf{b}'$.

The intuition for this auction is that ProfitExtract$_R$ allows us treat a set of bidders, $\mathbf{b}$, as one bidder with bid value $\mathcal{F}(\mathbf{b})$. Recall that a truthful auction must just offer a price t_i to bidder i who accepts if her value is at least t_i. This is analogous to trying to extract a profit R and actually getting R in profit when $\mathcal{F}(\mathbf{b}) \geq R$. The RSPE auction can then be viewed as randomly partitioning the bidders into two parts, treating each part as a single bid and performing the Vickrey auction on these two "bids". This intuition is crucial for the proof that follows as it implies that the profit of RSPE is the minimum of F' and F''.

Lemma 3.27 *RSPE is truthful.*

Before we prove that RSPE is 4-competitive, we give a simple proof of the following lemma that will be important in the analysis.

Lemma 3.28 *If we flipped $k \geq 2$ fair coins, then*

$$\mathbf{E}[\min\{\#\mathit{heads}, \#\mathit{tails}\}] \geq k/4.$$

Proof: Let M_i be a random variable for the $\min\{\#\text{heads}, \#\text{tails}\}$ after only i coin flips. We make the following basic calculations (verify these as an exercise):

- $\mathbf{E}[M_1] = 0$.
- $\mathbf{E}[M_2] = 1/2$.
- $\mathbf{E}[M_3] = 3/4$.

We now obtain a general bound on $\mathbf{E}[M_i]$ for $i > 3$. Let $X_i = M_i - M_{i-1}$ representing the change to $\min\{\#\text{heads}, \#\text{tails}\}$ after flipping one more coin. Notice that linearity of expectation implies that $\mathbf{E}[M_k] = \sum_{i=1}^{k} \mathbf{E}[X_i]$. Thus, it would be enough to calculate $\mathbf{E}[X_i]$ for all i. We consider this in two cases:
Case 1 (i even): This implies that $i-1$ is odd, and prior to flipping the ith coin it was not the case that there was a tie, i.e, #heads $\neq$ #tails. Assume without loss of generality that #heads ¡ #tails. Now when we flip the ith coin, there is probability $1/2$ that it is heads and we increase the minimum by one; otherwise, we get tails have no increase to the minimum. Thus, $\mathbf{E}[X_i] = 1/2$.

Case 2 (i odd): Here we use the crude bound that $\mathbf{E}[X_i] \geq 0$. Note that this is actually the best we can claim in worst case since $i-1$ is even and it could have been that #heads = #tails in the previous round. If this were the case then regardless of the ith coin flip, $X_i = 0$ and the minimum of #heads and #tails would be unchanged.
Case 3 ($i = 3$): This is a special case of Case 2; however we can get a better bound using the calculations of $\mathbf{E}[M_2] = 1/2$ and $\mathbf{E}[M_3] = 3/4$ above to deduce that $\mathbf{E}[X_3] = \mathbf{E}[M_3] - \mathbf{E}[M_2] = 1/4$.

Finally we are ready to calculate a lower bound on $\mathbf{E}[M_k]$.

$$
\begin{aligned}
M_k &= \sum\nolimits_{i=1}^{k} X_i \\
\mathbf{E}[M_k] &= \sum\nolimits_{i=1}^{k} \mathbf{E}[X_k] \\
&\geq 0 + 1/2 + 1/4 + 1/2 + 0 + 1/2 + 0 + 1/2 \ldots \\
&= \tfrac{1}{4} + \tfrac{\lfloor k/2 \rfloor}{2} \\
&\geq k/4.
\end{aligned}
$$

■

Theorem 3.29 *RSPE is 4-competitive.*

Proof: Recall our metaphor of using ProfitExtract_R to treat one partition of the bids $\mathbf{b}'$ as a single bid with value $F' = \mathcal{F}^{(2)}(\mathbf{b}')$, the other partition $\mathbf{b}''$ as a single bid with value $F'' = \mathcal{F}^{(2)}(\mathbf{b}'')$, and then running the Vickrey auction on these two bids. With this intuition it is obvious that the profit of RSPE is $\min(F', F'')$. In what remains, we analyze $\mathbf{E}[\min(F', F'')]$.

Assume that $\mathcal{F}^{(2)}(\mathbf{b}) = kp$ has with $k \geq 2$ winners at price p. Of the k winners in $\mathcal{F}^{(2)}$, let k' be the number of them 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 $F' \geq k'p$. Likewise, $F'' \geq k''p$.

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

The last inequality follows from applying Lemma 3.28 when we consider $k \geq 2$ coins and heads as putting a bidder in $\mathbf{b}'$ and a tails as putting the bidder in $\mathbf{b}''$.

It is evident that RSPE is no better than 4-competitive via an identical proof to that of RSOP. ■

The conclusion from Exercise 3.1 is that the Vickrey auction in the case that $n = 2$ is 2-competitive. Since this matches the lower bound, the Vickrey auction is the optimal 2-bidder auction. We can then view RSPE as partitioning the bidders into two parts, treating each part like a bid that is equal to its optimal single price revenue, and then running the optimal two bidder auction. The best known competitive auction has competitive ratio 3.25 and is based on randomly partitioning into three parts, treating each part as a single bid with value equal to its optimal single price revenue, and then running the optimal 3-bidder auction on these three "bids".

4 Online Auctions

In this section we consider the problem of selling a digital good *online*. In an online selling setting, we imagine that the bidders arrive one at a time and the auction must decide whether to sell to each bidder and at what price, before the next bidder arrives. The assumption that we are selling a digital good implies that we do not have to worry about running out.

The online auction we consider proceeds in rounds. In round i, bidder i arrives. The auctioneer must decide whether to allocate, x_i, and at what price, p_i. Then bidder i leaves, either with or without a copy of the good, and does not return. Our goal is to obtain a truthful online auction that maximizes the profit of the seller.

As before it will be useful to consider the problem of determining x_i and p_i for a truthful auction from t_i, the minimum bid value that bidder i can bid and still win. As before, $x_i = 1$ and $p_i = t_i$ if $b_i \geq t_i$, otherwise $x_i = p_i = 0$. The following constraints on t_i are imposed by online truthful auctions:

- For truthful auctions the offer price t_i is not a function of bidder i's bid.
- For online auctions the offer price t_i is not a function of future bids.

This second constraint comes from the fact that we cannot tell the future. At time i when bidder i arrives we must decide t_i. Thus, we consider the offer price to bidder i to be $t_i(b_1, \ldots, b_{i-1})$.

What function $t_i(b_1, \ldots, b_{i-1})$ gives a good auction profit? As with the offline digital good auction we would like an auction to obtain profit close to that of the *optimal single price sale*, $\mathcal{F}(\mathbf{b})$. In the remainder of this section we answer this question.

4.1 Learning from Expert Advice

The problem of online auctions is closely related to the machine learning problem of *learning from expert advice*. Compare the following problems statements:
Online Auction:

- bid range: $[1, h]$,
- n bidders, and
- in round i the auction
    1. chooses offer price t_i,

2. learn b_i, and
3. get payoff $p_i = t_i$ if $t_i \leq b_i$, 0 otherwise.

- Goal: profit close to optimal single price sale in hindsight.

Online Learning:

- k experts,
- n iterations, and
- in round i the expert algorithm:
  1. chooses an expert j,
  2. learn payoffs $q_1^{(i)}, \ldots, q_k^{(i)}$, and
  3. get payoff $q_j^{(i)}$.
- Goal: profit close to payoff of best expert in hindsight, $\text{OPT} = \max_j \sum_{i=1}^{n} q_j^{(i)}$.

The prototypical example of a problem of learning from expert advice is that of predicting the weather. In this example we imagine that the experts are weather reporters. A weather reporter receives payoff of one if their prediction is correct and zero otherwise. The problem of learning from expert advice calls for designing a strategy for selectively adopting the predictions of the weather reporters so as to be correct almost as often as the best one.

Is is not too difficult to see how to reduce the Online Auction problem to the problem of Online Learning from Expert Advice. Consider instantiating an expert j for the online learning problem that suggests single sale price of 2^j. The payoff of such an expert when facing bidder i is precisely 2^j when $2^j \leq b_i$ and zero otherwise. In this application k, the number of experts, is $\log h$ since we have assumed that all bid values are in the interval $[1, h]$. In this reduction there is a 1-to-1 correspondence between the best expert and the best single sale price that is a power of two. This implies that an expert algorithm that has performance as good as the best expert in hindsight can be used to derive an online auction that has performance as good as the optimal single sale price that is a power of two. It is clear that the best power of two sale price is within a factor of two of the profit of the optimal sale price, $\mathcal{F}$.

Of course we could also apply the above reduction with expert j advocating a sale price of $(1+\alpha)^j$ which would result the best expert being within a factor of $(1+\epsilon)$ of the optimal sale price. In this lecture we omit the details of optimizing constants in the online learning and auction problems and will use "2" in place of "$1+\epsilon$" and "1/2" in place of "$1/(1+\epsilon)$".

Consider the following algorithm and the example below that shows that it does not solve the online learning problem.

Definition 4.1 (Follow the Leader) The *follow the leader* algorithm:

- Let $s_j^{(i)} = \sum_{i'=1}^{i} q_j^{(i')}$ be the total payoff of expert j up to and including round i.

- In round i choose expert $j = \text{argmax}_j\, s_j^{(i-1)}$.

Example 4.2 Consider input:

	1	2	3	4	5	
Expert 1:	1/2	0	1	0	1	$\cdots$
Expert 2:	0	1	0	1	0	$\cdots$

The standard approach to learning from expert advice is to randomize a little bit so that examples like the one above are unlikely to be a problem. Probably the most well known algorithm for learning from expert advice is the weighted majority algorithm.

Definition 4.3 (Weighted Majority) The *weighted majority* algorithm for expert payoffs in $[0, h]$:

- Let $s_j^{(i)} = \sum_{i'=1}^{i} q_j^{(i')}$ (as before).
- In round i choose expert j with probability proportional to $2^{s_j^{(i-1)}/h}$.

Theorem 4.4 *When expert payoffs are in $[0, h]$, Weighted Majority satisfies $\mathbf{E}[\textit{Payoff}] \geq \text{OPT}/2 - O(\frac{h}{k} \log k)$.*

By plugging in our reduction from Online Auctions to Online Learning we obtain:

Corollary 4.5 *The Online Auction via Weighted Majority has $\mathbf{E}[\textit{Profit}] \geq \mathcal{F}/4 - O(h \log\log h)$.*

4.2 Kalai's Online Learning Algorithm

Definition 4.6 (Perturbed Follow-the-Leader) The perturbed follow-the-leader algorithm when expert payoffs are in $[0, h]$ is:

1. Hallucinate: set $s_j^{(0)} = h \times \#$ heads flipped in a row.
2. Follow the leader: in round i choose expert $j = \text{argmax}_{j'}\, s_{j'}^{(0)} + s_{j'}^{(i-1)}$

Theorem 4.7 *the perturbed follow-the-leader algorithm satisfies $\mathbf{E}[\textit{Payoff}] \geq \text{OPT}/2 - O(h \log k)$.*

Note that this is the same bound given for the Weighted Majority algorithm. Kalai's proof of this theorem follows from two steps. First, he shows that if Step 2 was "be the leader" instead of "follow the leader", i.e., if we choose expert $j = \text{argmax}_{j'}\, s_{j'}^{(0)} + s_{j'}^{(i)}$ in round i then we would obtain profit within $O(h \log k)$ of OPT. Second, he shows that the using the follow-the-leader subroutine is pretty similar to using be-the-leader as about half the time they make the same choices. Therefore, the expected profit of the randomly perturbed follow-the-leader algorithm is within a factor of two of the expected profit of the randomly perturbed be-the-leader algorithm. We make these arguments explicit in the following two lemmas.

Lemma 4.8 *Perturbed be-the-leader has* $\mathbf{E}[Profit] \geq \mathrm{OPT} - O(h \log k)$.

Proof: Fix the random coins flipped in the algorithm. The score of the best expert with hallucination initially is $H_0 = \max_j s_j^{(0)}$. At the end of the algorithm, the score of the best expert with hallucination is $H_n = \max_j s_j^{(0)} + s_j^{(n)}$. Consider the difference in the score of the best expert with hallucination between two consecutive rounds, $H_i - H_{i-1}$. Notice that the be-the-leader algorithm must obtain payoff of $H_i - H_{i-1}$ in round i. This is simply because the total score of the leading expert in round i is exactly H_i, while this experts score is a most H_{i-1} in the previous round. The change in an expert's score from one round to the next is their payoff. Thus, the leaders payoff is at least $H_i - H_{i-1}$. We can conclude that the be-the-leader algorithm has total payoff at least $H_n - H_0$.

We are interested in the expected payoff of the algorithm. By linearity of expectation this is just $\mathbf{E}[H_n] - \mathbf{E}[H_0]$. Clearly, $H_n \geq \mathrm{OPT}$ as H_n is the maximum of expert scores with added hallucination and OPT is just the maximum of the scores. To get a bound on $\mathbf{E}[H_0]$ we wish to know the expected value of h times the maximum of k geometric random variables. The maximum of k geometric random variables is well known to be $O(\log k)$. Thus, $\mathbf{E}[H_0] = O(h \log k)$.

We conclude that the lemma holds and that the expected profit of the perturbed be-the-leader algorithm is $\mathrm{OPT} - O(h \log k)$. ■

Lemma 4.9 *Perturbed follow-the-leader has expected payoff at least half of the expected payoff of the perturbed be-the-leader algorithm.*

Proof: We show that in each round the probability that the leader in the previous round is the same as the new leader is 1/2. Thus, with this probability follow-the-leader obtains the same payoff as be-the-leader. This, along with the simple observation that when the leaders in these consecutive rounds are not the same then follow-the-leader at least has non-negative payoff, shows that the expected payoff in each round for follow-the-leader is half of that of be-the-leader. Linearity of expectation, then, proves the lemma.

We now focus on a single round and show that the probability that the leader before and afterwards are the same with probability 1/2. Think about time i after tallying the payoffs of the experts in this round. The scores of the experts without hallucinations are $s_0^{(i)}, \ldots, s_k^{(i)}$. Now consider adding in the hallucinations in the following way:

1. Initially for all j, $s_j^{(0)} = 0$, all hallucinations are zero.

2. Pick the expert j with the lowest score plus hallucination, i.e., $j = \mathrm{argmin}_{j'} s_{j'}^{(0)} + s_{j'}^{(i)}$.

3. Flip a coin:

   (a) Heads: add an additional h to expert j's hallucination, $s_j^{(0)} \leftarrow s_j^{(0)} + h$.

   (b) Tails: permanently discard this expert. Their final hallucination is their current value of $s_j^{(0)}$ and they are total score with hallucination is worse than all other remaining experts. They cannot be the best expert at this point.

4. Repeat from Step 2 until there is only one expert remaining.

Notice that in the preceding algorithm the final hallucinations of all discarded experts are precisely those of a geometric distribution. For each of these expert we flipped a coin until we get tails and give them an additional score of h times the number of heads they obtained. The final remaining expert j, on the other hand, is definitely the leader after this round; however, j still has a coin to flip for calculating their total initial hallucination.

Consider the possible outcomes of expert j's next coin flip. We now argue that if this coin comes up heads then both the follow-the-leader and be-the-leader algorithms picked expert j. Since this happens with probability $1/2$, it completes the lemma. Notice that if the coin comes up heads then this expert, who was already leading, is now leading by more than h. Even if we subtracted the score of this expert for this round, $q_j^{(i)}$, which is at most h, the expert would still be leading. Since this expert is leading with their total score from the previous round, they must have also lead the previous round. ■

To conclude, the perturbed follow-the-leader algorithm solved the problem of learning from expert advice and we have given a proof that its performance matches the performance bound of the more standard weighted majority algorithm. It turns out that this proof is more revealing of the problem we are solving than the proof of the analogous bound for the weighted majority algorithm. In what follows, we show how to take advantage of Kalai's proof to derive an online auction with expected profit at least $\mathcal{F}/4 - O(h)$.

4.3 Non-uniform Bounds on Payoff

Notice that it was important that we had a bound of h on the payoff of each expert. It was no coincidence that when we hallucinated, we added multiples of h the experts' scores. In particular, it was important in the proof of Lemma 4.9 that h was bigger than $q_j^{(i)}$, so that one additional h from hallucination insured that the best expert after round i was also best before round i.

In our auction application of expert learning there is a special structure of the payoffs of the experts. In particular the expert that suggest "offer price 2^j" will always either make profit 2^j or zero. Thus, we can establish non-uniform bounds on each expert's payoff, $h_j = 2^j$. Consider the following algorithm.

Definition 4.10 (Non-Uniform Perturbed Follow-the-Leader) The non-uniform perturbed follow-the-leader algorithm when expert j's payoff in $[0, h_j]$ is:

1. Hallucinate: set $s_j^{(0)} = h_j \times \#$ heads flipped in a row.

2. Follow the leader: in round i choose expert $j = \text{argmax}_{j'}\, s_{j'}^{(0)} + s_{j'}^{(i-1)}$

Theorem 4.11 *The non-uniform perturbed follow-the-leader algorithm satisfies* $\mathbf{E}[\mathit{Payoff}] \geq \mathrm{OPT}/2 - \frac{1}{2}\sum_j h_j$.

Notice that when we apply this to our auction example with $h_j = 2^j$ the additive loss term telescopes to become $\frac{1}{2}\sum_j h_j = h$. Thus, we have the corollary:

Corollary 4.12 *The online auction constructed from the non-uniform perturbed follow-the-leader algorithm has* $\mathbf{E}[Profit] \geq \mathcal{F}/4 - h$.

Theorem 4.11 follows from the following lemmas. The non-uniform perturbed be-the-leader algorithm below corresponds to replacing Step 2 in the above non-uniform perturbed follow-the-leader algorithm with the be-the-leader algorithm.

Lemma 4.13 *Non-uniform perturbed be-the-leader has* $\mathbf{E}[Profit] \geq \mathrm{OPT} - \sum_j h_j$.

Proof: The proof of this theorem is identical to the proof of Lemma 4.13 except in the final step when we get a bound on $\mathbf{E}[H_0]$, the expected maximum hallucination. Instead we argue that the expected maximum hallucination is certainly less than the expected sum of the hallucinations. Using linearity of expectation and the fact that the expected number of heads flipped in a row is one, it is easy to see that $\mathbf{E}[H_0] = \sum_j h_j$. This gives the desired bound. ■

Lemma 4.14 *Non-uniform perturbed follow-the-leader has expected payoff at least half of the expected payoff of the non-uniform perturbed be-the-leader algorithm.*

Proof: This mirrors the proof of Lemma 4.9 substituting the fact that $q_j^{(i)} \leq h_j$ in the appropriate place. ■

This concludes our analysis of online auctions. There are two interesting directions this work has taken that we have not discussed here. The first is to *online posted price mechanisms*. In online posted price mechanisms in each round the seller must post an offer price; however, unlike in the auction setting, the seller only learns if the bidder accepted or rejected this price. The seller does not learn the bidder's actual valuation. We view this as the partial information version of the online auction problem and it can be solved with the partial information version of the problem of learning from expert advice, a.k.a., the *multi-armed bandit problem* (See [4, 3]).

The other interesting direction is in *limited supply* online auctions. Here, we have a fixed number of units of the item for sale. The limited supply constraint introduces state into the auction problem. With state in an online optimization problem, a decision in one round affects the payoff (or set of allowable decisions) in subsequent rounds. In particular, if we run out of items to sell, then we are not allowed to sell any more units. This presents a difficulty because the auctioneer may run out of units before encountering a number of really high valued bidders. Solutions to this problem assume that the bidders arrive in a random order which makes a prefix of the bidders appear like a random sample. Analysis similar to that of the random sampling optimal price auction from the preceding section can then be applied (See [7]).

References

[1] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.

[2] S. Baliga and R. Vohra. Market research and market design. *Advances in Theoretical Economics*, 3, 2003.

[3] A. Blum and J. Hartline. Near-Optimal Online Auctions. In *Proc. 16th Symp. on Discrete Alg.* ACM/SIAM, 2005.

[4] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online Learning in Online Auctions. In *Proc. 14th Symp. on Discrete Algorithms*, 2003.

[5] E. H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11:17–33, 1971.

[6] T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.

[7] M. Hajiaghayi, R. Kleinberg, and D. Parkes. Adaptive Limited-Supply Online Auctions. In *Proc. of 5th ACM Conference on Electronic Commerce*. ACM Press, New York, 2004.

[8] P. Klemperer. Auction theory: A guide to the literature. *J. of Economic Surveys*, 13:227–286, 1999.

[9] V. Krishna. *Auction Theory*. Academic Press, 2002.

[10] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Economic Theory*, 18:511–533, 2001.

[11] R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.

[12] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16:8–37, 1961.