CS880: Algorithmic Mechanism Design
Lecture 4: Myerson's mech.; Ironing; Simplicity vs Optimality **Scribe:** Evangelia Gergatsouli

4.1 Recap from previous lecture

Recall from the previous lecture, that in order to approach the revenue maximization problem we needed to assume some previous knowledge for the values of the agents i.e. the *Bayesian Assumption*. More formally the setting for revenue maximization is the following

Setting:

- 1 item
- n single parameter[1] linear[2] agents
- Values v_i for every agent where $\boldsymbol{v} = (v_1, v_2, \ldots, v_n) \sim \mathcal{F} = \mathcal{F}_1 \times \mathcal{F}_2 \times \ldots \times \mathcal{F}_n$, known product distribution

Goal: Find a mechanism $\mathcal{M} = (x, p)$, where $x : \mathbb{R}^n \to \{0, 1\}^n$ is the allocation function and $p : \mathbb{R}^n \to \mathbb{R}^n$ is the payment function, such that

- $\mathcal{M}$ satisfies some form of incentive compatibility (either DSIC or BIC)
- Revenue $\sum_i \mathbb{E}_{\boldsymbol{v}\sim F}[p_i(\boldsymbol{v})]$ is maximized

For this optimization problem, in order to achieve the incentive guarantees, we already saw that Myerson's lemma[Mye81] characterizes the mechanisms that have the DSIC/BIC property.

4.1.1 Myerson's Lemma[Mye81]

Theorem 4.1.1 *A mechanism $\mathcal{M} = (\boldsymbol{x}, \boldsymbol{p})$ for the setting described above is **BIC** if and only if for all agents i we have that*

- *The allocation $x_i(v_i)$ is monotone non-decreasing in v_i*
- *The payment $p(v_i) = v_i x(v_i) - \int_0^{v_i} x_i(t)dt + p_i(0)$*

The payment function, even though it may seem complicated, if we drew it in terms of the allocation and the values it would look like the shaded gray area in figure 4.1.1; essentially the area above the allocation curve until the value v_i.

[1]The agent has a single value for his utility of getting or not getting the item
[2]The utility of agent i is $v_i x_i - p_i$

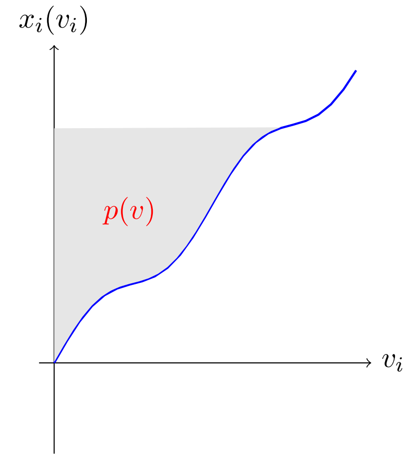

Figure 4.1.1: Payment amount displayed on the $x(v)$ plot

Theorem 4.1.2 *A mechanism $\mathcal{M} = (\boldsymbol{x}, \boldsymbol{p})$ for the setting described above is* ***DSIC*** *if and only if for all agents i and all values v_{-i} we have that*

- *The allocation $x_i(v_i, v_{-i})$ is monotone non-decreasing in v_i*
- *The payment $p(v_i, v_{-i}) = v_i x(v_i, v_{-i}) - \int_0^{v_i} x_i(t, v_{-i})dt + p_i(0)$*

Observe that from the lemma above we get that the revenue is linear in the allocation function $x(\boldsymbol{v})$.

4.2 Solving the Optimization Problem

We start by rewriting the optimal mechanism design problem as a "linear program", as shown in table 1 and try to simplify it. The first set of constraints, *feasibility*, can be anything that the problem requires, for example it could be that we only have 2 items to allocate, so we cannot allocate more than 2. For the final set of constrains we assume that $p(0) = 0$ i.e. we do not charge an agent that gets 0 items.

maximize	$\sum_i \mathbb{E}_{v_i \sim \mathcal{F}_i} [p_i(\boldsymbol{v})]$
subject to	$x_1, \ldots, x_n$ are feasible
	$x_i(v)$ is weakly increasing for all i
	$p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x(t)dt$, for all i

Table 1:

Denote by F_i and f_i the cdf and pdf of agent i respectively, where $F_i(v_i) = \mathbf{Pr}_{t \sim \mathcal{F}_i} [t \leq v_i]$ and $f_i(v_i) = F_i'(v_i)$. Additionally, in the calculations below we drop the i subscripts for convenience. For one agent we get that the payment is

$$
\begin{aligned}
\mathbb{E}_{v_i \sim \mathcal{F}_i}\left[p_i(v_i)\right] &= \int_0^{v_i} p_i(v_i) f_i(v_i) dv_i \\
&= \int_0^{v_i} v_i x_i(v_i) f_i(v_i) dv_i - \int_0^{\infty} \left(\int_0^{v_i} x(t)dt \right) f_i(v_i) dv_i && \text{Definition of payment} \\
&= \int_0^{v} v x(v) f(v) dv - \int_0^{\infty} \left(\int_0^{v} x(t)dt \right) f(v) dv && \text{Drop subscripts} \\
&= \int_0^{v} v x(v) f(v) dv - \int_{t=0}^{\infty} \left(\int_{v=t}^{\infty} f(v)dv \right) x(t) dt && \text{Change order of integration[3]} \\
&= \int_0^{v} v x(v) f(v) dv - \int_{t=0}^{\infty} (1 - F(t)) x(t) dt && \text{Definition of the pdf} \\
&= \int_0^{v} v x(v) f(v) dv - \int_{v=0}^{\infty} (1 - F(v)) x(v) dv && \text{Change variable } t \text{ to } v \\
&= \int_0^{v} x(v) (v f(v) dv - (1 - F(v))) dv && \text{Merge integrals} \\
&= \int_0^{v} x(v) \left(v - \frac{1 - F(v)}{f(v)} \right) f(v) dv \\
&= \int_0^{v} x(v) \Phi(v) f(v) dv && \text{Renaming} \\
&= \mathbb{E}_{v \sim \mathcal{F}}\left[x(v)\Phi(v)\right]
\end{aligned}
$$

Putting all the agents together, in the optimization problem we defined in table 1 we get the optimization problem of table 2. Observe that the third constraint was no longer needed since we used it in the calculations shown before, and is in a sense "embedded" in our objective.

$$
\begin{array}{rl}
\text{maximize} & \sum_i \mathbb{E}_{v_i \sim \mathcal{F}_i}\left[x_i(v_i)\Phi_i(v_i)\right] = \mathbb{E}\left[\sum_i x_i(v_i)\Phi_i(v_i)\right] \\
\text{subject to} & x_1, \ldots, x_n \text{ are feasible} \\
 & x_i(v) \text{ is weakly increasing for all } i
\end{array}
$$

Table 2: Rewritten optimization problem, using virtual values

Let $\Phi_i(v_i)$ be the virtual value for agent i. If we fix some vector $\boldsymbol{v} = (v_1, v_2, \ldots, v_n)$ and an allocation $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$ then the social welfare is $\sum_i x_i v_i$ while the virtual social welfare is $\sum_i x_i \Phi_i(v_i)$. Therefore, what we showed is that

Revenue Maximization $\equiv$ **Virtual Social Welfare Maximization s.t. monotonicity of** x

Now we can think of our problem in terms of *virtual welfare* with an additional monotonicity constraint. The process we use in these problems is the following

1. Ignore the monotonicity constraint

2. Point-wise maximize objective
3. Go back to the constraint, what conditions we need for the constraint to be satisfied?

We show how this is applied in the following example.

Example In this example, there is 1 item and 2 agents, whose values are distributed as follows: $v_1, v_2 \sim U[0,1]$ therefore $F_1(t) = \mathbf{Pr}[v_i \leq t] = t$, then also $f_1(t) = f_2(t) = 1$

Assume initially that the goal was social welfare maximization, so we wanted to maximize $\mathbb{E}[x_1 v_1 + x_2 v_2]$. In that case, we ask the agents to report v_1, v_2 and we allocate to 1 (so $x_1 = 1, x_2 = 0$) iff $v_1 \geq v_2$ otherwise ($x_1 = 0, x_2 = 1$) agent 2 gets the item and then the agents are charged the critical price.

In the case of revenue maximization, as shown above we calculate the virtual value function Φ as follows

$$\Phi(v) = v - \frac{1 - F(v)}{f(v)} = 2v - 1$$

then the objective becomes

$$\max \ (x_1(2v_1 - 1) + x_2(2v_2 - 1)) = x_1\Phi(v_1) + x_2\Phi(v_2)$$

such that only 1 item is allocated. When we do not have the monotonicity property, we can point-wise maximize the objective, the same way we did in the case of welfare maximization, i.e. if $\Phi_1 \geq \Phi_2$ then $x_1 = 1, x_2 = 0$, and $x_1 = 0, x_2 = 1$ otherwise.

Observe though, that the virtual values, contrary to the actual values v_i, can also be negative. If that is the case, we do not want to allocate the item at all. Therefore, we modify the above rule that decides the allocation to take into account the possible negative virtual values, while still point-wise maximizing the objective:

- If $\Phi_1 \geq \Phi_2$ and $\Phi_1 \geq 0$ then allocate to 1 ($x_1 = 1, x_2 = 0$)
- If $\Phi_2 \geq \Phi_1$ and $\Phi_2 \geq 0$ then allocate to 2 ($x_1 = 0, x_2 = 1$)
- If $\Phi_1 < 0, \Phi_2 < 0$ serve no one ($x_1 = x_2 = 0$)

Now that we made sure the allocation actually maximizes the objective, we need to also make sure that the monotonicity property is also satisfied.

$$x_1(v_1) = \mathbf{Pr}[\Phi_1(v_1) \geq 0, \Phi(v_1) \geq \Phi(v_2)] = \begin{cases} 0, & \text{if } v_1 \leq v_2 \\ \mathbf{Pr}[\Phi_1 \geq \Phi_2], & \text{else} \end{cases} = \begin{cases} 0, & \text{if } v_1 \leq v_2 \\ v_1, & \text{else} \end{cases}$$

where we used that when fixing v_1, then $\mathbf{Pr}[\Phi_2 < \Phi_1] = \mathbf{Pr}[2v_2 - 1 < 2v_1 - 1] = \mathbf{Pr}[v_2 < v_1] = v_1$. The plot or this function is shown in figure 4.2.2. Observe that this function is monotone, therefore putting the constraint back in ,gives us the optimal mechanism.

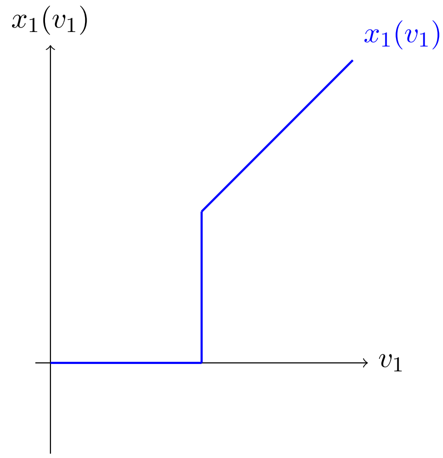

Figure 4.2.2: Allocation function for the example

Observe that int his case, for any fixed v_2 we have that $x(v_1, v_2) = \mathbf{Pr}[v_1 \geq 1/2 \text{ and } v_1 \geq v_2]$ which is a step function, changing from 0 to 1 at $\max(1/2, v_2)$. The function Φ has a nice property in that case, that makes this mechanism satisfy the monotonicity property; as v increases, so does Φ. This condition formally is called *regularity* and we define it below.

4.2.1 The Road to the Optimal Mechanism

Definition 4.2.1 (Regularity) *A value distribution $\mathcal{F}$ is regular, if Φ is weakly non decreasing.*

We describe now the optimal mechanism, and using the claim below, we only need this regularity to hold for the value distributions.

1)	Agents report $v_1, \ldots, v_n$
2)	Compute $\Phi_1, \ldots, \Phi_n$
3)	Return feasible set $argmax_s \sum_{i \in s} \Phi_i$ (maximizes virtual surplus)

Table 3: Myerson's Optimal Mechanism

Claim 4.2.2 *If all value distributions are regular, then Myerson's mechanism (table 3) is BIC.*

Observe that the mechanism is deterministic, and DSIC[4] Recall that the virtual value is $\Phi(v) = v - \frac{1-F(v)}{f(v)}$.

There is a special class of distributions that satisfy this, that have the monotone hazard rate property. Specifically, we say a distribution is MHR (*Monotone Hazard Rate*) if $h(v) = \frac{f(v)}{1-F(v)}$ is monotone non-decreasing. Observe that for these distributions, we get that the virtual function is non-decreasing.

Examples of distributions to try and find Φ

[4]Since agents will not want to misreport, once the values of the other agents are fixed.

- Uniform in $[a, b]$
- Exponential distribution
- Power law distribution & special case $f(v) = 1/v^2$ for $v \in [1, \infty)$
- Bimodal distributions (for example $1/2U[0, 1] + 1/2U[2, 3]$)

4.3 Change to Quantiles

In this section, we swift gears to a quantile interpretation of the virtual value function. Recall that quantiles are uniform, so, for example, being below the 90th quantile has 90% probability.

Consider for example the function F shown in figure 4.3.3 and described by the cdf: $F = \frac{1}{2}U[0, 1/3] + \frac{1}{2}U[1/3, 1]$. Note that picking a quantile in this plot, we can uniquely map it back to a value, so if we drew the same plot but using quantiles on the x-axis we get figure 4.3.4 shown below, where the low quantiles in this plot correspond to high values and vice versa[5] where the quantile q is $q = 1 - F(v)$ and $v = V(q)$.

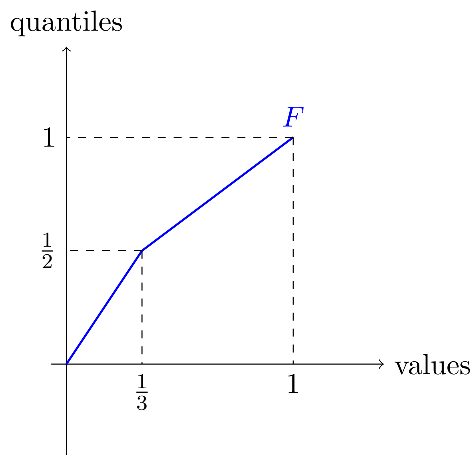

Figure 4.3.3: Cdf for the distribution F

[5] A 42% quantile, means that we are among the best/highest 42% values.

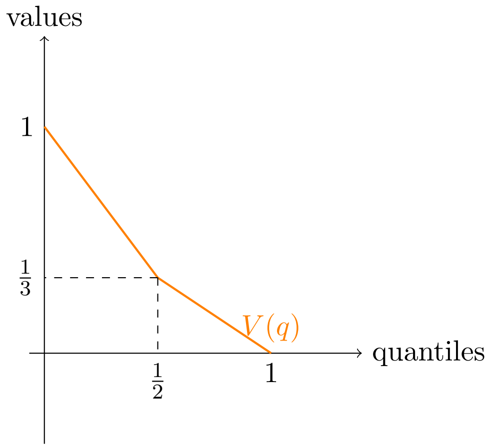

Figure 4.3.4: Quantile plot for the distribution F

4.3.1 Single agent

We focus on a single agent, and let $x(q)$ denote the allocation for the agent with quantile q. The revenue then is the expected revenue from allocation x, which is now a function of the quantile.

Denote by $\mathcal{T}_t(q) = \mathbb{I}_{q \leq t}$ the step function that is 1 before t and 0 after. Then any function $x(q)$ can be rewritten as an integral over these step functions as shown below[6].

$$x(q) = -\int x'(t)\mathcal{T}_t(q)dt$$

Having defined x in this quantile space, we write revenue as follows

$$\begin{aligned}
Rev(x) &= -\int_t x'(t)R(t)dt \\
&= \int_0^1 x(t)R'(t)dt - R(1)x(1) + R(0)x(0) && \text{integration by parts} \\
&= \int_0^1 x(t)R'(t)dt && \text{Since } R(0) = 0 \text{ and } R(1) = 0
\end{aligned}$$

where we used that the expected revenue of $\mathcal{T}_t$ is $R(t) = tV(t)$ which is the probability of serving the agent (t) times the value it gives $(V(t))$, and for the first equality integration by parts.

Therefore, we have shown that for any **weakly monotone** allocation function x then $Rev(x) = \int_q x(q)R'(q)dq$. If we call $R'(q)$ the *virtual value* we get that

Expected revenue $=$ **Expected Virtual Surplus**

[6]The negative sign is because x is decreasing.

We show now that this is the same virtual surplus we calculated before. Recall that $R(q) = qV(q)$ and that $V(q)$ is the inverse of $1 - F(v)$.

$$\begin{aligned} R'(q) &= V(q) + qV'(q) \\ &= V(q) + q\frac{-1}{f(v)} \\ &= v + \frac{1 - F(v)}{-f(v)} \\ &= v - \frac{1 - F(v)}{f(v)} \\ &= \Phi(v) \end{aligned}$$

where we used that $V(q) = dv/dq = 1/(dq/dv) = 1/(1 - F(v))' = -1/f(v))$.

Recall that the regularity condition from before required for $\Phi(v)$ to be non-decreasing in v. Using the rewritten version of Φ we have that we want $\Phi(q) = \Phi((V(q))$ is non increasing in q[7] and therefore $R(q)$ is concave in q (since we require the derivative to be non-increasing).

Example: apply this to the example distribution F The cdf of the distribution in shown in figure 4.3.3 and the plot of $V(q)$ in figure 4.3.4. We calculate $V(q)$ and $R(q)$ as follows

$$V(q) = \begin{cases} 1 - \frac{4}{3}q, & \text{For } q \leq \frac{1}{2} \\ \frac{2}{3}(1 - q) & \text{For } q > \frac{1}{2} \end{cases}$$

and

$$R(q) = \begin{cases} q - \frac{4}{3}q^2, & \text{For } q \leq \frac{1}{2} \\ \frac{2}{3}\left(q - q^2\right) & \text{For } q > \frac{1}{2} \end{cases}$$

If we plot the revenue curve above we get the plot in figure 4.3.1, and as we can see this is not concave. Ideally we would want to allocate the item only when $\Phi > 0$, but this will lead to a non-monotone allocation function. In order to avoid this problem, we use a technique called *ironing*, which basically corresponds to "drawing" the purple lines in the plot (see figure 4.3.1), and essentially removing the non-monotonicities in Φ and the non-concavities in $R(q)$.

[7] This comes form the interpretation of quantiles, and the fact that low quantile means high value.

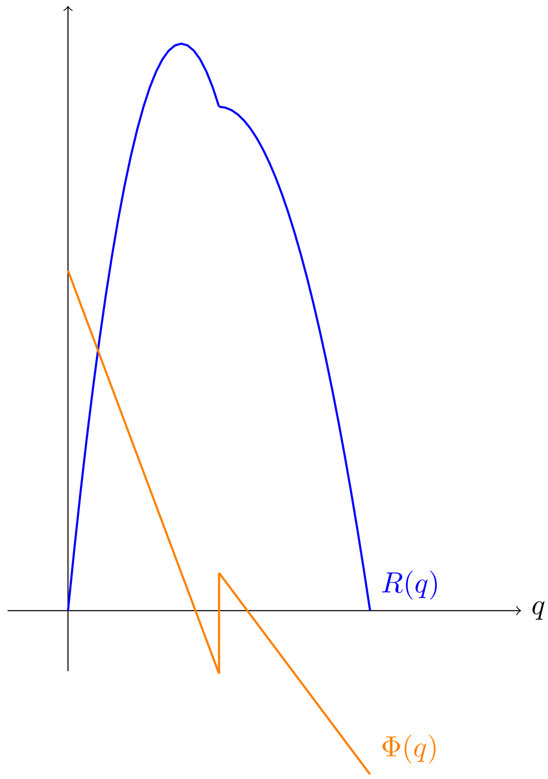

Figure 4.3.5: $R(q)$ plot

This can be achieved by taking the concave upper envelope $\overline{R}(q)$ for R, which formally is the smallest concave function that lies above $R(q)$. This is shown in figure 4.3.1 below; the envelope is the same as the function R with the difference that it follows the purple part to avoid the non-concavity of R.

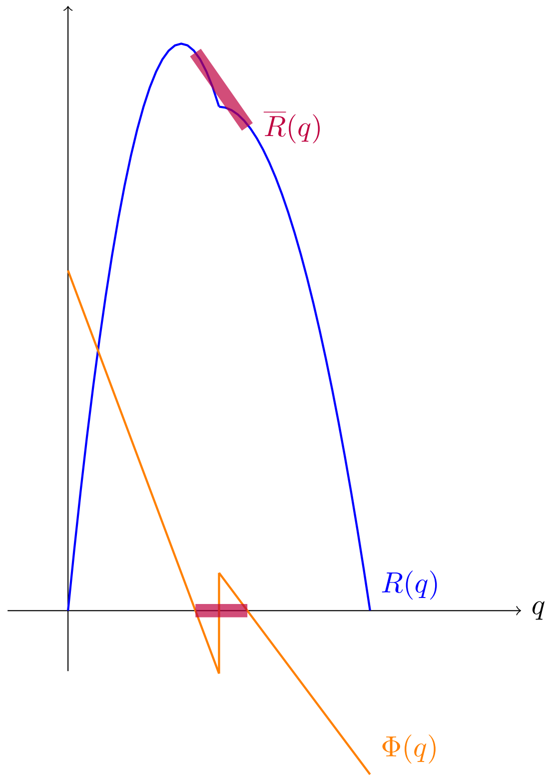

Figure 4.3.6: $R(q)$ ironed plot

Observe the following properties satisfied by the concave envelope:

1. $\overline{R}(q)$ is concave and so $\overline{\Phi}(q)$ is non increasing

2. $Rev(x) = \mathbb{E}_q\left[x(q)\Phi(q)\right] = -\mathbb{E}_q\left[x'(q)R(q)\right]$
   If x is monotone (i.e. $x'(q) \leq 0$) then $Rev(x) = -\mathbb{E}\left[x'(q)R(q)\right] \leq \mathbb{E}\left[x'(q)\overline{R}(q)\right] = \mathbb{E}\left[x(q)\overline{\Phi}(q)\right]$

   Therefore, we have that **Expected Revenue of x $\leq$ Ironed Virtual Surplus**. This happens essentially because the ironed surplus "pretends" that the mechanism gets more revenue than it actually does.

3. For any mechanism $\mathcal{M} = (x, p)$ such that $x'(q) = 0$ whenever $\overline{R}(q) \neq R(q)$ we have that **Expected Revenue of x = Ironed Virtual Surplus**.

4. There is a monotone mechanism maximizing virtual surplus that satisfies 3 (above)

4.4 Putting it all together: Myerson's Mechanism

1. Pointwise maximize the **ironed** virtual surplus

2. If there are any ties in the ironed virtual values→ break ties consistently. In any interval where $\overline{R}(q) > R(q)$ we will get the ironed version of the virtual surplus to be constant: $\overline{\Phi}(q)$ =constant, and therefore the allocation is the same over this region[8]. This implies that $x'(q) = 0$ in this region, and therefore that **ironed virtual surplus = virtual surplus**.

Example for Myerson's Mechanism We have 1 item, 2 buyers where $F_1 = U[0,2]$ and $F_2 = U[0,3]$, therefore $F_1(t) = t/2$ and $F_2(t) = t/3$. We find the virtual value functions:

- $\Phi_1(t) = t - \frac{1-F_1(t)}{f_1(t)} = t - \frac{1-t/2}{1/2} = 2t - 2$
- $\Phi_2(t) = t - \frac{1-F_2(t)}{f_2(t)} = t - \frac{1-t/3}{1/3} = 2t - 3$

The plot of the virtual functions is shown in figure 4.4 below. It is clear now, when the optimal mechanism should allocate to each agent:

- If $v_1 \leq 1$ then reject agent 1
- If $v_2 \leq 1.5$ then reject agent 2
- Else, agent 2 needs to outbid 1 by 1/5 to win

There are some useful observations to make for this mechanism, which we describe here.

1. The mechanism is discriminating, since if we know which agent has more value, we can set different "entry" fees, depending on their values.
2. If all values are iid and regular, then all the value functions are identical. This implies that the agent with the higher Φ is the agent with the highest v, then we allocate to the agent with the highest value as long as it is more than 0. This is called the monopoly reserve price $\Phi^{-1}(0) = r^*$. The mechanism for iid regular agents, described concretely is
   - remove everyone below r^*
   - allocate to the highest surviving agent

   which is essentially the Vickrey auction with monopoly reserve.

[8] This is the purple region in figure 4.3.1.

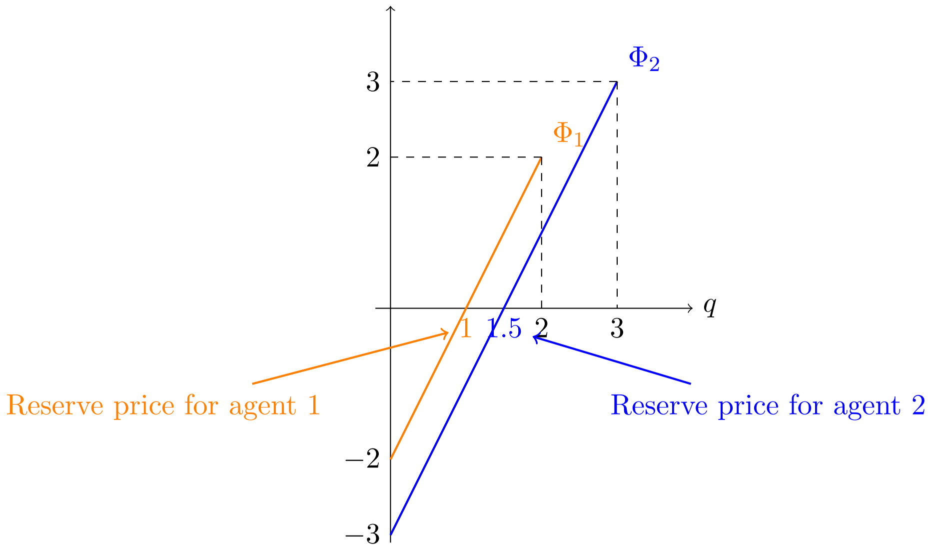

Figure 4.4.7: Virtual Values

4.5 Simplicity vs Optimality

In this section, we redirect our focus to designing simple mechanisms, that also are good approximations to the optimal, since even from the example before, we saw that the rules can get quite complicated. The way we try to approximate the optimal mechanism is through Vickrey Auction with (different) reserve prices, for non iid settings. The auction is described below:

1. Compute $r_i^* = \Phi_i^{-1}(0)$
2. Reject all agents with $v_i \leq r_i^*$
3. Run Vickrey auction over the values reported
4. Charge critical prices $= \max(r_i^*, \max_{i \neq j,\ v_j \geq r_j^*} v_j)$

Theorem 4.5.1 *In the non-iid regular setting, Vickrey with monopoly reserves sets a 2-approximation to the optimal expected revenue. This factor is also tight.*

Proof: Initially note that the optimal mechanism also rejects the agents that have a negative virtual value (steps 1+2 of the algorithm above). In the 3rd step though the optimal will run the Vickrey auction over the virtual values in stead of the reported values.

The difference the optimal mechanism and the Vickrey mechanism might have, comes from the fact that the order in virtual values is not the same as the order in the values v_i[9], so for example the ordering of the values might be $v_1 \geq v_2 \geq v_3 \ldots$ but the ordering of the virtual values is

[9] Since the agents do not have the same distribution

$\Phi_2 \geq \Phi_1 \geq \Phi_4 \geq \ldots$. Note that Vickrey has a monotone allocation rule, so we can use the characterization of the expected revenue; $Rev(Vickrey) = \mathbb{E}\left[Virtual\ Surplus\right]$, and in this case the revenue of Vickrey is Φ_1 and the revenue of the optimal is Φ_2.

Denote by i the winning agent for optimal and j the winning agent for Vickrey. Observe that for every agent we have that $\Phi_i(v_i) \leq v_i$. The optimal revenue is

$$\mathbb{E}\left[\Phi_i\right] = \mathbb{E}\left[\Phi_i | i = j\right] \mathbf{Pr}\left[i = j\right] + \mathbb{E}\left[\Phi_i | i \neq j\right] \mathbf{Pr}\left[i \neq j\right]$$

we bound each of the two terms by the revenue of the Vickrey auction: for the fist term we get

$$\begin{aligned}
\mathbb{E}\left[\Phi_i | i = j\right] \mathbf{Pr}\left[i = j\right] &= \mathbb{E}\left[\Phi_j | i = j\right] \mathbf{Pr}\left[i = j\right] \\
&\leq \mathbb{E}\left[\Phi_j | i = j\right] \mathbf{Pr}\left[i = j\right] + \mathbb{E}\left[\Phi_j | i \neq j\right] \mathbf{Pr}\left[i \neq j\right] \\
&= Rev(Vickrey)
\end{aligned}$$

where for the inequality we used that $\Phi_j \geq 0$, from the specification of Vickrey. Similarly for the second term (where the optimal and Vickrey do not pick the same agent)

$$\begin{aligned}
Rev(Vickrey) &= \mathbb{E}\left[p_j | i = j\right] \mathbf{Pr}\left[i = j\right] + \mathbb{E}\left[p_j | i \neq j\right] \mathbf{Pr}\left[i \neq j\right] \\
&\geq \mathbb{E}\left[v_i | i \neq j\right] \mathbf{Pr}\left[i \neq j\right] \\
&\geq \mathbb{E}\left[\Phi_i | i \neq j\right] \mathbf{Pr}\left[i \neq j\right]
\end{aligned}$$

where the first inequality comes from the fact that j needs to beat at least the value v_i to win.v Putting all of this together, we get that $Rev(OPT) \leq Rev(Vickrey)$. ■

References

[Mye81] Roger B. Myerson. Optimal auction design. *Math. Oper. Res.*, 6(1):58–73, 1981.