

Revenue Maximization with Identical Distributions

Instructor: Thomas Kesselheim

Last time, we introduced the problem of maximizing revenue in an auction. We saw a reasonably simple mechanism for the case of a single item. The results also hold for other single-parameter settings—for example, if we have multiple copies of the same items.

We assume that the bidders' valuations v_1, \dots, v_n are drawn from known distributions. Today, we will consider the case that these are identical. That is, v_1, \dots, v_n are independent and identically distributed according to \mathcal{D} . The distribution is continuous; it has density function f and cumulative distribution function F . We define a function φ by $\varphi(t) = t - \frac{1-F(t)}{f(t)}$. This function is called the *virtual-value function*. Today, we will keep the assumption that the distribution is *regular*, which means that φ is increasing.

Our main result last time was that the expected revenue of a truthful mechanism equals its expected virtual welfare:

$$\mathbf{E}_{v_1, \dots, v_n \sim \mathcal{D}} \left[\sum_{i \in \mathcal{N}} p_i(v) \right] = \mathbf{E}_{v_1, \dots, v_n \sim \mathcal{D}} \left[\sum_{i \in \mathcal{N}} \varphi(v_i) x_i(v) \right] .$$

From this, it is easy to conclude that a mechanism that chooses $x(b)$ so as to maximize $\sum_{i \in \mathcal{N}} \varphi(b_i) x_i(b)$ maximizes the expected revenue. Note that this allocation rule is monotone and this way gives a truthful mechanism if and only if the distribution is regular.

For a single item, this means to select the bidder with the highest virtual bid $\varphi(b_i)$ if this value is positive. Otherwise, the item is not allocated.

1 Virtual-Welfare Maximization

We have seen that choosing the allocation with maximum virtual welfare maximizes the expected revenue. In the case of a single item, this means to give the item the bidder i who maximizes $\varphi(b_i)$ if $\varphi(b_i) \geq 0$. In the case of identical, regular distributions, because of monotonicity and because the functions are identical, this is the bidder with maximal b_i . Call this bidder i^* .

The payment is again the smallest bid t that would make him a winner. Two possible cases can happen: If $\varphi(\max_{i \neq i^*} b_i) \geq 0$, then some other bidder would have won in the absence of i^* . So, he has to pay $t = \max_{i \neq i^*} b_i$. If $\varphi(\max_{i \neq i^*} b_i) < 0$, then nobody would have won the item. However, i^* still would have to bid so that $\varphi(t) \geq 0$. In other words, he has to pay $\varphi^{-1}(0)$.

Summarizing, the payment of bidder i^* is

$$\max \left\{ \varphi^{-1}(0), \max_{i \neq i^*} b_i \right\} .$$

That is, we have a second-price auction with a reserve price of $\varphi^{-1}(0)$. Just add a bidder bidding $\varphi^{-1}(0)$ to the second-price auction and if this bidder wins, nobody gets the item.

2 Comparison to Second-Price Auction

We have seen a revenue-optimal mechanism for identically distributed bidders. In comparison to the general case, the mechanism is fairly easy. But what if we want an easier mechanism? For example, if we didn't know about the results on revenue maximization, we would probably use the second-price auction. How well does it perform?

Lemma 22.1. *If all valuations are drawn from the same regular distribution, the second-price auction maximizes the expected revenue among all truthful mechanisms that always allocate the item.*

Proof. Let $\mathcal{M}^* = (x^*, p^*)$ be the second-price auction and let $\mathcal{M} = (x, p)$ be any other truthful mechanism that always allocates the item.

Both mechanisms always allocate the item. That is, given any valuation/bid profile v , there are $i^*, i' \in \mathcal{N}$ such that the second-price auction allocates to i^* and \mathcal{M} allocates to i' . The second-price auction gives the item to the bidder of highest value/bid, so $v_{i^*} \geq v_{i'}$. Furthermore, φ is monotone, that is,

$$\sum_{i \in \mathcal{N}} \varphi(v_i) x^*(v) = \varphi(v_{i^*}) \geq \varphi(v_{i'}) = \sum_{i \in \mathcal{N}} \varphi(v_i) x(v) .$$

By taking expectations on both sides, we get

$$\mathbf{E}_v \left[\sum_{i \in \mathcal{N}} p_i^*(v) \right] = \mathbf{E}_v \left[\sum_{i \in \mathcal{N}} \varphi(v_i) x^*(v) \right] \geq \mathbf{E}_v \left[\sum_{i \in \mathcal{N}} \varphi(v_i) x(v) \right] = \mathbf{E}_v \left[\sum_{i \in \mathcal{N}} p_i(v) \right] . \quad \square$$

This observation already nicely proves some intuition to be correct: The additional revenue that can be obtained compared to a second-price auction is due to the fact that one sometimes does not allocate the item.

Another theorem takes this observation to the next level.

Theorem 22.2 (Bulow/Klemperer, 1996). *Consider a regular distribution \mathcal{D} . For any $n \in \mathbb{N}$, the revenue of a second-price auction with $n + 1$ bidders drawn from \mathcal{D} is at least as high as the optimal revenue with n bidders drawn from \mathcal{D} .*

Proof. What makes the proof complicated is the fact that we have to deal with different numbers of bidders. We can make our life easier by assuming that there are $n + 1$ bidders overall. The revenue-optimal mechanism for n bidders can be interpreted in this setting as a mechanism that simply ignores bidder $n + 1$. Let \mathcal{M} be this mechanism.

We can turn \mathcal{M} into a mechanism \mathcal{M}' that always allocates the item by assigning it to bidder $n + 1$ for free if \mathcal{M} does not allocate it. Note that the expected revenue of \mathcal{M} and \mathcal{M}' are identical.

Now, let \mathcal{M}^* be the second-price auction. By Lemma 22.1, we know that the revenue of \mathcal{M}^* is at least as high as the revenue of \mathcal{M}' . \square

3 Posting One Price

There is an even simpler mechanism than a second-price auction: Simply fix a price p^* and ask the bidders one after the other if they want to buy the item at price p^* . To set this price, we will make use of a result in optimal stopping theory.

Theorem 22.3 (i.i.d. Prophet Inequality). *Let Y_1, \dots, Y_n be independent, identically distributed, non-negative random variables. Let τ be a threshold such that $\Pr[Y_i > \tau] = \frac{1}{n}$ and let Y_{select} be the value of the first Y_i variable above this threshold, 0 if there is none. Then*

$$\mathbf{E}[Y_{\text{select}}] \geq \left(1 - \left(1 - \frac{1}{n}\right)^n\right) \mathbf{E}\left[\max_i Y_i\right] .$$

Note that $1 - \left(1 - \frac{1}{n}\right)^n \geq 1 - \frac{1}{e} \approx 0.63$ for all n .

Proof. Note that to bound the expectation it is sufficient to show that for all $z \geq 0$

$$\Pr[Y_{\text{select}} \geq z] \geq \Pr\left[\max_i Y_i \geq z\right] .$$

We distinguish two cases. If $z \leq \tau$, then we use that Y_{select} automatically has value τ if any Y_i exceeds the threshold. Therefore

$$\Pr [Y_{\text{select}} \geq z] = \Pr [\exists i : Y_i > \tau] = 1 - \left(1 - \frac{1}{n}\right)^n .$$

As $\Pr [\max_i Y_i \geq z] \leq 1$, this implies the claim.

Now we consider $z > \tau$. We can decompose the event $Y_{\text{select}} \geq z$ into disjoint events $Y_1, \dots, Y_{i-1} \leq \tau, Y_i \geq z$. Because the Y_i are independent and identically distributed, this is

$$\begin{aligned} \Pr [Y_{\text{select}} \geq z] &= \sum_{i=1}^n \Pr [Y_1, \dots, Y_{i-1} \leq \tau, Y_i \geq z] \\ &= \sum_{i=1}^n \Pr [Y_1 \leq \tau] \dots \Pr [Y_{i-1} \leq \tau] \Pr [Y_i \geq z] \\ &= \sum_{i=1}^n \left(1 - \frac{1}{n}\right)^{i-1} \Pr [Y_1 \geq z] \\ &= \frac{1 - \left(1 - \frac{1}{n}\right)^n}{1 - \left(1 - \frac{1}{n}\right)} \Pr [Y_1 \geq z] \\ &= \left(1 - \left(1 - \frac{1}{n}\right)^n\right) n \Pr [Y_1 \geq z] \\ &= \left(1 - \left(1 - \frac{1}{n}\right)^n\right) \sum_{i=1}^n \Pr [Y_i \geq z] . \end{aligned}$$

Furthermore, we can re-write the event that the maximum Y_i is at least z as a union of disjoint events of the form $Y_1 < z, \dots, Y_{i-1} < z, Y_i \geq z$

$$\begin{aligned} \Pr \left[\max_i Y_i \geq z \right] &= \Pr \left[\bigcup_{i=1}^n Y_1 < z, \dots, Y_{i-1} < z, Y_i \geq z \right] \\ &= \sum_{i=1}^n \Pr [Y_1 < z, \dots, Y_{i-1} < z, Y_i \geq z] \leq \sum_{i=1}^n \Pr [Y_i \geq z] . \end{aligned}$$

In combination, this gives us $\Pr [Y_{\text{select}} \geq z] \geq \left(1 - \left(1 - \frac{1}{n}\right)^n\right) \Pr [\max_i Y_i \geq z]$, which is exactly what we claimed. \square

Coming now back to our problem of setting a price of maximize revenue, we choose it as

$$p^* = \max \left\{ \varphi^{-1}(0), F^{-1} \left(1 - \frac{1}{n}\right) \right\} .$$

The first term ensures that we only select bidders of non-negative virtual value. The second term makes sure that the probability that a bidder is selected is at most $\frac{1}{n}$.

Theorem 22.4. *The mechanism that uses price p^* and approaches bidders in order $1, \dots, n$ gives a $1 - (1 - \frac{1}{n})^n \geq 1 - \frac{1}{e}$ -approximation to the optimal revenue.*

Proof. We use Theorem 22.3 to prove this theorem. Note that it is enough to bound the virtual welfare rather than the revenue. So, let $Y_i = \max\{0, \varphi(v_i)\}$ be the virtual value of bidder i but at least 0. The optimal virtual welfare is exactly $\max_{i \in \mathcal{N}} Y_i$. Our mechanism's virtual welfare is the first Y_i such that $v_i > p^*$, which is nothing but the first Y_i for which $Y_i > \varphi(p^*)$.

If $\varphi^{-1}(0) \leq F^{-1} \left(1 - \frac{1}{n}\right)$, Theorem 22.3 gives us exactly what we want. In this case,

$$\Pr [Y_i > \varphi(p^*)] = \Pr \left[v_i > F^{-1} \left(1 - \frac{1}{n}\right) \right] = \frac{1}{n} .$$

So,

$$\mathbf{E} [Y_{\text{select}}] \geq \left(1 - \left(1 - \frac{1}{n}\right)^n\right) \mathbf{E} \left[\max_i Y_i\right] .$$

If $\varphi^{-1}(0) > F^{-1}\left(1 - \frac{1}{n}\right)$, then

$$\Pr [Y_i > \varphi(p^*)] = \Pr [Y_i > 0] < \frac{1}{n} .$$

So, we choose a higher threshold than Theorem 22.3 tells us to. (It is simply not possible to set the threshold exactly as the theorem requires because there is a discontinuity of $\Pr [Y_i > \tau]$ at $\tau = 0$. However, there is no actual difference because we only might continue the sequence although Theorem 22.3 would tell us to stop at 0. The value of Y_{select} can therefore only be higher. \square

4 Bonus Content: Beyond Truthfulness

Our insights so far only concern truthful direct mechanisms. One would be tempted to think that non-truthful mechanisms might yield a higher revenue. For example, given the same bids, in a first-price auction the revenue is higher than in a second-price auction, which is truthful. However, this is not true if one takes into consideration the strategic behavior of the bidders. We will now sketch a proof that considering truthful mechanisms is not actually a restriction.

Theorem 22.5. *If all bidder's distributions are regular, there is no mechanism \mathcal{M}' that has a Bayes-Nash equilibrium $(\beta_i)_{i \in \mathcal{N}}$ in which the revenue is higher than in the dominant-strategy equilibrium of the virtual-welfare maximizer.*

Proof Sketch. Assume that there is a mechanism \mathcal{M}' with a Bayes-Nash equilibrium $(\beta_i)_{i \in \mathcal{N}}$ of higher revenue. Our *first observation* is that we can build a direct mechanism \mathcal{M}'' that given v first computes $\beta(v)$ and then runs \mathcal{M}' . For any valuation profile v , Mechanism \mathcal{M}'' gives the same outcome and payments on truthful reports that \mathcal{M}' gives on $(\beta_i)_{i \in \mathcal{N}}$. (This technique is called the *revelation principle*.)

Mechanism \mathcal{M}'' is what we call *Bayes-Nash incentive compatible*: Truthful bidding is a Bayes-Nash equilibrium. (This does not mean that it is dominant-strategy incentive compatible; we have seen such examples.)

For such Bayes-Nash incentive compatible mechanisms, a variant of Myerson's lemma holds. It tells us that the expected payments are unique. For every $i \in \mathcal{N}$ and every b_i , we have

$$\mathbf{E}_{v_{-i}} [p_i(b_i, v_{-i})] = \mathbf{E}_{v_{-i}} \left[b_i x_i(b_i, v_{-i}) - \int_0^{b_i} f_i(t, v_{-i}) dt \right] .$$

It is derived exactly the same way as Myerson's lemma for dominant-strategy incentive compatible mechanisms: One can use the payment-difference sandwich, which now holds in expectation.

The exact same calculations as above give us

$$\mathbf{E}_v \left[\sum_{i \in \mathcal{N}} p_i(v) \right] = \mathbf{E}_v \left[\sum_{i \in \mathcal{N}} \varphi_i(v_i) x_i(v) \right] ,$$

meaning that still, even under the weaker assumption of Bayes-Nash incentive compatibility, the expected revenue is equal to the expected virtual welfare.

Among all mechanisms (without any further assumption), the virtual-welfare maximizer has the highest expected virtual welfare. \square

5 Bonus Content: Multiple Items

So far, all our results were restricted to single-parameter settings and mostly just to allocating a single item. If we turn to multiple items, the theory suddenly becomes much more complicated. To get an idea what is happening, let us consider a single bidder and two items. Note that a single item you would sell at the price p^* that maximizes $p^* \Pr[v \geq p^*]$.

We assume that the bidders' values are v_1 for item 1 and v_2 for item 2. The valuation is additive, that is, for getting both items the value is $v_1 + v_2$. We did not consider maximizing social welfare for such additive valuations because it is trivial: We can treat the items separately. For example, run a second-price auction for each of them. One might think that this is also true for revenue. But it is not.

5.1 Item Bundling

Assume that v_1 and v_2 are independent and both uniformly drawn from $\{1, 2\}$. Let us first set item prices p_j^* . If we set $p_j^* = 1$ or $p_j^* = 2$, then the expected revenue from item j is 1. For other prices it is even less. So, the maximum revenue from item prices is 2.

An alternative mechanism is to only sell items 1 and 2 as a bundle. If we set the price of bundle $\{1, 2\}$ to 3 (single items cannot be purchased), then the bidder will always buy the bundle unless $v_1 = v_2 = 1$. So, the revenue is $3(1 - \Pr[v_1 = v_2 = 1]) = 3(1 - \frac{1}{4}) = \frac{9}{4} > 2$.

5.2 Lotteries

There is something even stranger that happens. Our results for a single item are still true if we allow our mechanism to sell "lottery tickets". For example, this might mean that with a medium bid one wins the item with probability $\frac{1}{2}$, with a small bid the probability is 1, with a high bid it is 1. Also for those lottery mechanisms, the expected revenue is equal to expected virtual welfare. We can maximize expected virtual welfare without resorting to selling lotteries. So the revenue-optimal mechanism works without lotteries.

In the case of multiple items, this is again different. Consider the case that v_1 and v_2 are independent; v_1 is drawn uniformly from $\{1, 2\}$, v_2 is drawn uniformly from $\{1, 3\}$. The optimal mechanism that does not sell lottery tickets prices item 1 at 1 and item 2 at 3. Its expected revenue is $1 + \frac{3}{2} = \frac{5}{2} = 2.5$. However, there is a mechanism with higher revenue. It offers the buyer the following two options. One choice is to get both items for sure and pay 4. The other one is a lottery ticket for 2.5. This ticket includes item 1 for sure and item 2 with probability $\frac{1}{2}$. One can show that this mechanism gives expected revenue 2.625 and that this is the optimal revenue.