

Lecture 21

Lecturer: Aaron Roth

Scribe: Aaron Roth

Mechanism Design via Differential Privacy

In this class, we continue using digital goods auctions with valuations $v_i \in [0, 1]$ as a case study for another technique in mechanism design. Recall that two lectures ago, we designed an auction for the digital goods setting that was dominant strategy truthful, and for any vector of valuations v , obtained revenue:

$$\text{Rev}(v) \geq \text{OPT}(v) - O(\sqrt{n})$$

Where here $\text{OPT}_v = \max_{p \in [0, 1]} p \cdot |\{i : v_i \geq p\}|$ denotes the revenue of the best fixed price.

In this class, we will aim for another point on the tradeoff space – we will relax our solution concept to *asymptotic* dominant strategy truthfulness, and in exchange, will try and obtain a better revenue guarantees. The tool we will use to do this is *differential privacy*, an algorithmic stability constraint that is primarily studied as a means of doing data analysis while protecting the privacy of individuals in the data set.

Recall: Why can we not simply compute the price $p^* = \arg \max_{p \in [0, 1]} p \cdot |\{i : v_i \geq p\}|$ that optimizes revenue, and use that price? The reason is that this price will be highly manipulable by one of the bidders – it will be v_{i^*} for some bidder i^* , who will have strong incentive to lower his bid. But what if we can compute a price p that is *almost* independent of the reported valuation v_i for every buyer i ? Will this yield in some sense an approximate truthfulness guarantee? This will be the idea behind our approach.

Definition 1 Two bid vectors $v, v' \in [0, 1]^n$ are neighbors if they differ in just a single agent's bid: i.e. if there exists an index i such that $v_j = v'_j$ for every index $j \neq i$.

We can now define differential privacy:

Definition 2 A mechanism $M : [0, 1]^n \rightarrow \mathcal{O}$ is ϵ -differentially private if for every pair of neighboring bid vectors $v, v' \in [0, 1]^n$, and for every outcome $x \in \mathcal{O}$:

$$\Pr[M(v) = x] \leq \exp(\epsilon) \Pr[M(v') = x]$$

. Here you should think of $\epsilon < 1$ as a small constant, and think of $\exp(\epsilon) \approx (1 + \epsilon)$. Using the exponential form will simply be more convenient mathematically. (And for $\epsilon \leq 1$ we have $1 + \epsilon \leq \exp(\epsilon) \leq 1 + 2\epsilon$)

We can also define what we mean by *approximate* dominant strategy truthfulness:

Definition 3 A mechanism $M : [0, 1]^n \rightarrow \mathcal{O}$ is ϵ -approximately dominant strategy truthful if for every bidder i , every utility function $u_i : [0, 1] \times \mathcal{O} \rightarrow [0, 1]$, every vector of valuations $v \in [0, 1]^n$, and every deviation $v'_i \in [0, 1]$, if we write $v' = (v_{-i}, v'_i)$, then:

$$\mathbb{E}_{o \sim M(v)}[u_i(v_i, o)] \geq \mathbb{E}_{o \sim M(v')} [u_i(v_i, o)] - \epsilon$$

In other words, we require that no bidder can substantially (by more than ϵ) increase his utility by mis-reporting his valuation.

Differential privacy will be useful for us because differentially private mechanisms are automatically ϵ -approximately dominant strategy truthful.

Theorem 4 If a mechanism $M : [0, 1]^n \rightarrow \mathcal{O}$ is ϵ -differentially private, then M is also ϵ -approximately dominant strategy truthful.

Proof Fix any buyer i , valuation vector v , and utility function $u_i : [0, 1] \times \mathcal{O} \rightarrow [0, 1]$.

$$\begin{aligned}
 \mathbb{E}_{o \sim M(v)}[u_i(v_i, o)] &= \sum_{o \in \mathcal{O}} u_i(v_i, o) \cdot \Pr[M(v) = o] \\
 &\geq \sum_{o \in \mathcal{O}} u_i(v_i, o) \cdot \exp(-\epsilon) \Pr[M(v') = o] \\
 &= \exp(-\epsilon) \mathbb{E}_{o \sim M(v')} [u_i(v_i, o)] \\
 &\geq \mathbb{E}_{o \sim M(v')} [u_i(v_i, o)] - \epsilon
 \end{aligned}$$

where the last inequality follows because for $\epsilon < 1$, $\exp(-\epsilon) \geq 1 - \epsilon$, and $u_i(v_i, o) \leq 1$. ■

Great! So to design an approximately truthful mechanism that guarantees high revenue, it is sufficient to design a differentially private mechanism with high revenue. Lets see if we can do so in a straightforward way: directly picking a price that approximately maximizes revenue for the particular bidder valuations that have been reported.

As we have done in the last two lectures, lets pick a finite subset of prices $P \subset [0, 1]$ to select from. Now consider the following mechanism (an instantiation of what is called “the exponential mechanism” in its more general form):

ExpMech(v, ϵ, P):

Define $\text{Rev}(p, v) = p \cdot |\{i : v_i \geq p\}|$.

Output each $p \in P$ according to the following probability distribution:

$$\Pr[p] = \frac{1}{\phi(v)} \exp\left(\frac{\epsilon \cdot \text{Rev}(p, v)}{2}\right)$$

where

$$\phi(v) = \sum_{p \in P} \exp\left(\frac{\epsilon \cdot \text{Rev}(p, v)}{2}\right)$$

First, we prove that the exponential mechanism is ϵ -differentially private (and hence ϵ -approximately truthful):

Theorem 5 For any ϵ, P : $\text{ExpMech}(\cdot, \epsilon, P)$ is ϵ -differentially private.

Proof Fix any pair of neighboring bid vectors v, v' and any output p . We have:

$$\begin{aligned}
 \Pr[\text{ExpMech}(v, \epsilon, P) = p] &= \frac{1}{\phi(v)} \exp\left(\frac{\epsilon \cdot \text{Rev}(p, v)}{2}\right) \\
 &\leq \frac{1}{\phi(v)} \exp\left(\frac{\epsilon \cdot (\text{Rev}(p, v') + 1)}{2}\right) \\
 &= \frac{1}{\phi(v)} \exp\left(\frac{\epsilon}{2}\right) \exp\left(\frac{\epsilon \cdot \text{Rev}(p, v')}{2}\right) \\
 &\leq \exp\left(\frac{\epsilon}{2}\right) \frac{1}{\phi(v')} \exp\left(\frac{\epsilon}{2}\right) \exp\left(\frac{\epsilon \cdot \text{Rev}(p, v')}{2}\right) \\
 &= \exp(\epsilon) \Pr[\text{ExpMech}(v', \epsilon, P) = p]
 \end{aligned}$$

■

It remains to bound the revenue that the exponential mechanism obtains.

Theorem 6 For any P, v, ϵ, δ , with probability $1 - \delta$, $\text{ExpMech}(v, \epsilon, P)$ outputs a price p such that:

$$\text{Rev}(p, v) \geq \max_{p^* \in P} \text{Rev}(p^*, v) - \frac{2}{\epsilon} \cdot \ln \left(\frac{|P|}{\delta} \right)$$

Proof Let $p^* = \max_{p^* \in P} \text{Rev}(p^*, v)$. For any value x , we have:

$$\begin{aligned} \Pr_p[\text{Rev}(p, v) \leq x] &\leq \frac{\Pr_p[\text{Rev}(p, v) \leq x]}{\Pr_p[\text{Rev}(p, v) = \text{Rev}(p^*, v)]} \\ &\leq \frac{|P| \cdot \exp(\epsilon x / 2)}{\exp(\epsilon \text{Rev}(p^*, v) / 2)} \\ &= |P| \cdot \exp \left(\frac{\epsilon \cdot (x - \text{Rev}(p^*, v))}{2} \right) \end{aligned}$$

Now choose $x = \text{Rev}(p^*, v) - \frac{2}{\epsilon} \cdot \ln \left(\frac{|P|}{\delta} \right)$. Plugging that in above, we get:

$$\begin{aligned} \Pr_p[\text{Rev}(p, v) \leq x] &\leq |P| \cdot \exp \left(-\ln \left(\frac{|P|}{\delta} \right) \right) \\ &= |P| \cdot \frac{\delta}{|P|} \\ &= \delta \end{aligned}$$

■

Now we can put all the pieces together. We have an approximately truthful way to select a revenue maximizing price from a finite set of prices P , where the revenue guarantees with respect to the best price in P degrade as a function of $|P|$. This is a familiar tradeoff – the larger P is, the closer it is to containing an optimal price, but the worse we will be able to approximate it. However, in this case, the dependence we see on $|P|$ is only logarithmic... Lets again see what happens when we take the natural discretization:

$$P = \{\alpha, 2\alpha, 3\alpha, \dots, 1\}$$

Note that just as in the last lecture, $|P| = 1/\alpha$, and that we have the guarantee that for all v :

$$\max_{p \in P} \text{Rev}(p, v) \geq \max_{p \in [0, 1]} \text{Rev}(p, v) - \alpha n$$

Combining this bound with the guarantee of the exponential mechanism, we see that if we discretize the price space by α , with probability 0.99, we obtain revenue:

$$\text{Rev}(p, v) \geq \text{OPT} - \alpha \cdot n - O \left(\frac{1}{\epsilon} \ln \left(\frac{1}{\alpha} \right) \right)$$

Choosing $\alpha = 1/n$, we find that for any ϵ , we can obtain an ϵ -approximately dominant strategy truthful mechanism which obtains revenue:

$$\text{Rev}(p, v) \geq \text{OPT} - O \left(\frac{\log n}{\epsilon} \right)$$

If we take e.g. $\epsilon = O(1/\log(n))$, then we have an *asymptotically truthful* mechanism (in the sense that it becomes exactly truthful in the limit as $n \rightarrow \infty$, that improves by an exponential factor on the revenue guarantee that we were able to obtain with an exactly truthful mechanism for the same problem.