

Smooth Mechanisms

Thomas Kesselheim

Last Update: December 17, 2020

In this lecture we will continue our exploration of non-truthful mechanisms. We have already got to know an equilibrium concept for settings of incomplete information last time. Today, we will translate the smoothness concept that we have seen a few lectures ago from games to mechanisms. Next time, we will use it to design simple, near-optimal mechanisms.

1 Basic Definitions

Recall our definition of a mechanism-design problem. There is a set \mathcal{N} of n players and a set of feasible outcomes X . Every player $i \in \mathcal{N}$ has a (private) valuation $v_i: X \rightarrow \mathbb{R}_{\geq 0}$ from a set of possible valuations V_i . A *mechanism* $\mathcal{M} = (f, p)$ defines a set of bids B_i for each player $i \in \mathcal{N}$ and consists of

- an *outcome rule* $f: B \rightarrow X$, where $B = B_1 \times B_2 \times \dots \times B_n$, and
- a *payment rule* $p: B \rightarrow \mathbb{R}_{\geq 0}^n$. So far, we assumed that payments could be arbitrary real numbers. Today, they have to be non-negative.

We say that the mechanism is *direct* if $B_i = V_i$ for all $i \in \mathcal{N}$, otherwise we say it is *indirect*. The utility of bidder i on bid profile $b \in B$ is given as $u_i(b, v_i) = v_i(f(b)) - p_i(b)$.

1.1 Complete Information

For a fixed choice of v , these utilities define a normal-form maximization game. If we assume *complete information*, we study the equilibria of this game. For example, a pure Nash equilibrium is a vector of strategies – in this case bids – such that no player wants to unilaterally deviate.

Definition 15.1 (Pure Nash Equilibrium). *Given a fixed valuation profile $v \in V$, a profile of bids $b = (b_1, \dots, b_n) \in B$ is a pure Nash equilibrium (PNE) if for every player $i \in \mathcal{N}$ and every deviation $b'_i \in B_i$,*

$$u_i((b_i, b_{-i}), v_i) \geq u_i((b'_i, b_{-i}), v_i) .$$

Also the concepts of mixed Nash and (coarse) correlated equilibria still make sense here.

The goal is to choose an outcome $x \in X$ that maximizes *social welfare* $\sum_{i \in \mathcal{N}} v_i(x)$. We use $OPT(v) = \max_{x \in X} \sum_{i \in \mathcal{N}} v_i(x)$ to denote the optimal social welfare. For a fixed bid vector b , the mechanism achieves welfare $SW_v(b) = \sum_{i \in \mathcal{N}} v_i(f(b)) = \sum_{i \in \mathcal{N}} u_i(b, v_i) + \sum_{i \in \mathcal{N}} p_i(b)$.

We define the *Price of Anarchy* for any given equilibrium concept as the worst possible ratio between the optimal social welfare and the (expected) social welfare at equilibrium, that is

$$PoA_{\text{Eq}} = \max_{v \in V} \max_{\mathcal{B} \in \text{Eq}(v)} \frac{OPT(v)}{\mathbb{E}_{b \sim \mathcal{B}}[SW_v(b)]},$$

where $\text{Eq}(v)$ denotes the set of equilibria for the game induced by valuations v .

This ratio is always at least 1 because the optimal social welfare can never be smaller than the social welfare in equilibrium. Ratios closer to 1 are better. Furthermore, we have again

$$1 \leq PoA_{\text{PNE}} \leq PoA_{\text{MNE}} \leq PoA_{\text{CCE}} .$$

1.2 Incomplete Information

Last time, we introduced the concept of games with incomplete information. Now, bidder i 's valuation v_i is drawn from a publicly known distribution \mathcal{D}_i . In a Bayes-Nash equilibrium, players choose their strategies depending on their own valuation but not on the other players' valuations.

Definition 15.2 (Bayes-Nash equilibrium). *A (pure) Bayes-Nash equilibrium (BNE) is a profile of bidding functions $(\beta_i)_{i \in N}$, $\beta_i: V_i \rightarrow B_i$ such that for all $i \in N$, all $v_i \in V_i$, and all $b'_i \in B_i$*

$$\mathbf{E}_{v_{-i} \sim \mathcal{D}_{-i}} [u_i(\beta(v), v_i)] \geq \mathbf{E}_{v_{-i} \sim \mathcal{D}_{-i}} [u_i((b'_i, \beta_{-i}(v)), v_i)] \quad ,$$

where $\beta(v) = (\beta_1(v_1), \dots, \beta_n(v_n))$.

Also for this setting, we can define the Price of Anarchy

$$PoA_{\text{BNE}} = \max_{\text{distribution } \mathcal{D}} \max_{\beta \text{ is BNE for } \mathcal{D}} \frac{\mathbf{E}_{v \sim \mathcal{D}} [OPT(v)]}{\mathbf{E}_{v \sim \mathcal{D}} [SW_v(\beta(v))]} .$$

So, we now consider the worst choice of distributions, and, again, the worst equilibrium. The value $OPT(v)$ is now a random variable, therefore we take its expectation.

2 Recap: First-Price Auction

Last time, we showed that $PoA_{\text{PNE}}, PoA_{\text{BNE}} \leq 2$ in a first-price auction. Let us recap the argument for pure Nash equilibria. The valuations v are fixed. We showed that for any b , we have for all $i \in N$

$$u_i \left(\left(\frac{v_i}{2}, b_{-i} \right), v_i \right) \geq \frac{v_i}{2} - \max_{i'} b_{i'} \quad \text{and} \quad u_i \left(\left(\frac{v_i}{2}, b_{-i} \right), v_i \right) \geq 0 \quad .$$

The first inequality follows by a simple case distinction: Either i wins the item with bid $\frac{v_i}{2}$, then the utility is $v_i - \frac{v_i}{2} = \frac{v_i}{2}$. Or i loses, so then $\max_{i'} b_{i'} \geq \frac{v_i}{2}$. The second inequality follows because in any case the utility is non-negative. In combination, this gives us for any v and any b

$$\sum_{i \in N} u_i \left(\left(\frac{v_i}{2}, b_{-i} \right), v_i \right) \geq \max_{i \in N} u_i \left(\left(\frac{v_i}{2}, b_{-i} \right), v_i \right) \geq \max_{i \in N} \frac{v_i}{2} - \max_{i \in N} b_i = \frac{1}{2} OPT(v) - \sum_{i \in N} p_i(b) \quad . \tag{1}$$

We get that if b is a pure Nash equilibrium $SW_v(b) = \sum_{i \in N} u_i(b, v_i) + \sum_{i \in N} p_i(b) \geq \sum_{i \in N} u_i \left(\left(\frac{v_i}{2}, b_{-i} \right), v_i \right) + \sum_{i \in N} p_i(b) \geq \frac{1}{2} OPT(v)$.

3 The Smoothness Framework

We define smooth mechanisms and show how smoothness implies that all equilibria of a mechanism are close to optimal.

Definition 15.3 (Smooth Mechanism, simplified version). *Let $\lambda, \mu \geq 0$. A mechanism \mathcal{M} is (λ, μ) -smooth if for any valuation profile $v \in V$ for each player $i \in N$ there exists a bid b_i^* such that for any profile of bids $b \in B$ we have*

$$\sum_{i \in N} u_i((b_i^*, b_{-i}), v_i) \geq \lambda \cdot OPT(v) - \mu \sum_{i \in N} p_i(b) \quad .$$

Note that, by the order of the quantifiers, b_i^* may depend on the profile of valuations but not on the bids. Equation (1) already shows that the single-item first-price auction is $(1/2, 1)$ smooth by setting $b_i^* = \frac{v_i}{2}$.

Observation 15.4. *A single-item first-price auction is $(1/2, 1)$ -smooth.*

To get a little more intuition, let us consider the single-item all-pay auction. Everything is the same as in the first-price auction except for the payments: Everybody pays his bid, no matter if he wins or loses.

Theorem 15.5. *A single-item all-pay auction is $(\frac{1}{2}, 2)$ -smooth.*

Proof. Let j be a bidder with highest value v_j . Set $b_j^* = v_j/2$ and $b_i^* = 0$ for $i \neq j$. Consider an arbitrary bid profile $b \in B$.

We show that always $u_j(b_j^*, b_{-j}) \geq \frac{1}{2}v_j - 2 \max_{i \neq j} b_i$. We distinguish two cases. If bidder j wins the item in (b_j^*, b_{-j}) , then his utility is $v_j - b_j^* = \frac{1}{2}v_j$. So, the bound is fulfilled because bids are non-negative. If he does not win the item, then his utility is $-\frac{1}{2}v_j$. As he loses, somebody must outbid him, meaning that $\max_{i \neq j} b_i \geq \frac{1}{2}v_j$. So, the bound holds as well.

Furthermore, by non-negativity of bids, $\sum_i p_i(b) = \sum_i b_i \geq \max_{i \neq j} b_i$.

Finally, for all $i \neq j$, we have $u_i(b_i^*, b_{-i}) \geq 0$ because $b_i^* = 0$ and therefore regardless of b_{-i} the bidder does not have to pay anything.

In combination, this gives us

$$\sum_i u_i(b_i^*, b_{-i}) \geq u_j(b_j^*, b_{-j}) \geq \frac{1}{2}v_j - 2 \max_i b_i \geq \frac{1}{2}v_j - 2 \sum_i p_i(b) = \frac{1}{2}OPT(v) - 2 \sum_i p_i(b) .$$

Therefore, this auction is $(\frac{1}{2}, 2)$ -smooth. □

4 Price-of-Anarchy Bound for Full Information

To see how smoothness bounds the price of anarchy, we first consider the full-information case. Here, the proof works just like for smooth games. To keep things simple, we consider pure Nash equilibria only.

Theorem 15.6 (Syrgkanis and Tardos, 2013). *If a mechanism \mathcal{M} is (λ, μ) -smooth and players have the possibility to withdraw from the mechanism then*

$$PoA_{\text{PNE}} \leq \frac{\max\{\mu, 1\}}{\lambda} .$$

Proof. Suppose bid profile b is a pure Nash equilibrium. This means that no player wants to unilaterally deviate from the equilibrium bid to some other bid. That is,

$$u_i((b_i, b_{-i}), v_i) \geq u_i((b'_i, b_{-i}), v_i) ,$$

for all players $i \in \mathcal{N}$ and bids $b'_i \in B_i$.

Now in particular players do not want to deviate to the bid b_i^* whose existence is guaranteed by smoothness. Considering, for each player $i \in \mathcal{N}$ the deviation to b_i^* and summing over all players,

$$\sum_{i \in \mathcal{N}} u_i((b_i, b_{-i}), v_i) \geq \sum_{i \in \mathcal{N}} u_i((b_i^*, b_{-i}), v_i) \geq \lambda \cdot OPT(v) - \mu \cdot \sum_{i \in \mathcal{N}} p_i(b) .$$

Since players have quasi-linear utilities $u_i(b, v_i) = v_i(f(b)) - p_i(b)$ or $v_i(f(b)) = u_i(b, v_i) + p_i(b)$. Using this we obtain

$$\sum_{i \in \mathcal{N}} v_i(f(b)) \geq \lambda \cdot OPT(v) + (1 - \mu) \cdot \sum_{i \in \mathcal{N}} p_i(b) .$$

Notice that the left-hand side is precisely the social welfare at equilibrium. So if $\mu \leq 1$ we can bound $(1 - \mu) \cdot \sum_{i \in \mathcal{N}} p_i(b) \geq 0$ and obtain

$$\sum_{i \in \mathcal{N}} v_i(f(b)) \geq \lambda \cdot OPT(v) ,$$

which shows a Price of Anarchy of $1/\lambda = \max\{1, \mu\}/\lambda$.

On the other hand, if $\mu > 1$, we can use that players have the right to withdraw from the mechanism and obtain a utility of zero to argue that $u_i(b) = v_i(f(b)) - p_i(b) \geq 0$ and so $p_i(b) \leq v_i(f(b))$. Since $(1 - \mu) < 0$ we obtain

$$\sum_{i \in \mathcal{N}} v_i(f(b)) \geq \lambda \cdot OPT(v) + (1 - \mu) \cdot \sum_{i \in \mathcal{N}} v_i(f(b)) .$$

Subtracting $(1 - \mu) \cdot \sum_{i \in \mathcal{N}} v_i(f(b))$ and dividing by $\mu > 1$ we obtain

$$\sum_{i \in \mathcal{N}} v_i(f(b)) \geq \lambda/\mu \cdot OPT(v) ,$$

which again shows a Price of Anarchy bound of $\mu/\lambda = \max\{1, \mu\}/\lambda$. \square

This argument extends to more general equilibrium concepts such as coarse correlated equilibria. The only point where we used the equilibrium condition is when we argued that players do not want to deviate from the equilibrium bid b_i to some other bid b'_i . In fact, the specific deviations that we considered only depended on the valuation profile v and did not depend on the bids b . Hence the exact same argument applies to coarse correlated equilibria and shows an upper bound on the Price of Anarchy of $\max\{1, \mu\}/\lambda$.

5 Price-of-Anarchy Bound for Incomplete Information

Next, we turn to Bayes-Nash equilibria. Last time, we saw how to bound the Price of Anarchy for Bayes-Nash equilibria in the case of a first-price auction. We used that bidder i would not prefer bidding $\frac{v_i}{2}$ instead of $\beta_i(v_i)$ for any v_i . This is exactly in the spirit of a smoothness-based proof. The difficulty is that the deviation bid b_i^* does not only depend on v_i but also on the other bidders' valuations v_{-i} . When we show smoothness of the all-pay auction, this is indeed crucial. Interestingly, using a very smart argument, this is not a problem and we can still derive the same bound.

Theorem 15.7. *If a mechanism \mathcal{M} is (λ, μ) -smooth and players have the possibility to withdraw from the mechanism then*

$$PoA_{\text{BNE}} \leq \frac{\max\{\mu, 1\}}{\lambda} .$$

Proof. We write out the dependence of b_i^* on v explicitly as $b_i^*(v)$ for this proof.

Let \tilde{v} be any valuation profile. Because $(\beta_i)_{i \in \mathcal{N}}$ is a Bayes-Nash equilibrium, bidder i would not prefer to unilaterally switch to strategy $b_i^*(v_i, \tilde{v}_{-i})$, that is

$$\mathbf{E}_{v_{-i}} [u_i(\beta(v), v_i)] \geq \mathbf{E}_{v_{-i}} [u_i((b_i^*(v_i, \tilde{v}_{-i}), \beta_{-i}(v)), v_i)]$$

for every v_i and every \tilde{v} . This is, in particular, true if \tilde{v} is a random valuation profile, drawn independently from the distributions that v is drawn from. We also take the expectation over v_i to get

$$\mathbf{E}_v [u_i(\beta(v), v_i)] \geq \mathbf{E}_{v, \tilde{v}} [u_i((b_i^*(v_i, \tilde{v}_{-i}), \beta_{-i}(v)), v_i)] . \quad (2)$$

This inequality is not surprising at all: Bidder i would not prefer to bid some other bid, which is only based on some valuation for the other bidders that he imagined.

But in the following step, the magic happens. Note that v_{-i} and \tilde{v}_{-i} are independent and identically distributed. Therefore, the expectation is the same if we swap them

$$\mathbf{E}_{v, \tilde{v}} [u_i((b_i^*(v_i, \tilde{v}_{-i}), \beta_{-i}(v)), v_i)] = \mathbf{E}_{v, \tilde{v}} [u_i((b_i^*(v), \beta_{-i}(\tilde{v})), v_i)] .$$

Let us have a closer look at what this means. On the left-hand side, we assumed that the bidders except i just follow whatever the Bayes-Nash equilibrium tells them to do with respect to their actual valuations. Bidder i just “hallucinates” some valuations \tilde{v}_{-i} and chooses a bid that is good against this hallucination—which does not have any meaning with respect to the real values v . On the right-hand side, things have flipped. Now, bidder i actually does the right thing against v_{-i} but the other bidders potentially do something strange: They are bidding what the Bayes-Nash equilibrium tells them to do for a *different* valuation profile. The reason why this is true is that u_i does not depend on v_i directly.

Fixing v and \tilde{v} , $\beta(\tilde{v})$ is just some bid profile. By the smoothness inequality, we therefore have

$$\sum_{i \in N} u_i((b_i^*(v), \beta_{-i}(\tilde{v})), v_i) \geq \lambda OPT(v) - \mu \sum_{i \in N} p_i(\beta(\tilde{v})) .$$

By linearity of expectation, this implies

$$\sum_{i \in N} \mathbf{E}_{v, \tilde{v}} [u_i((b_i^*(v), \beta_{-i}(\tilde{v})), v_i)] \geq \lambda \mathbf{E}_v [OPT(v)] - \mu \mathbf{E}_{\tilde{v}} \left[\sum_{i \in N} p_i(\beta(\tilde{v})) \right]$$

and in combination with Equation (2)

$$\sum_{i \in N} \mathbf{E}_v [u_i(\beta(v), v_i)] \geq \lambda \mathbf{E}_v [OPT(v)] - \mu \mathbf{E}_{\tilde{v}} \left[\sum_{i \in N} p_i(\beta(\tilde{v})) \right] .$$

Now, we use again that v and \tilde{v} are identically distributed, which means that $\mathbf{E}_{\tilde{v}} [\sum_{i \in N} p_i(\beta(\tilde{v}))] = \mathbf{E}_v [\sum_{i \in N} p_i(\beta(v))]$. So

$$\sum_{i \in N} \mathbf{E}_v [u_i(\beta(v), v_i)] \geq \lambda \mathbf{E}_v [OPT(v)] - \mu \mathbf{E}_v \left[\sum_{i \in N} p_i(\beta(v)) \right] ,$$

which implies because $u_i(b, v_i) = v_i(f(b)) - p_i(b)$

$$\sum_{i \in N} \mathbf{E}_v [v_i(f(\beta(v)))] \geq \lambda \mathbf{E}_v [OPT(v)] - (\mu - 1) \mathbf{E}_v \left[\sum_{i \in N} p_i(\beta(v)) \right] .$$

If $\mu \leq 1$, then we are done. Otherwise, we use again $u_i(\beta(v_i)) = v_i(f(\beta(v))) - p_i(\beta(v)) \geq 0$ because bidders can withdraw from the mechanism and so $p_i(\beta(v)) \leq v_i(f(\beta(v)))$. This again implies

$$\sum_{i \in N} \mathbf{E}_v [v_i(f(\beta(v)))] \geq \frac{\lambda}{\mu} \mathbf{E}_v [OPT(v)] ,$$

which is what we claimed. □

References and Further Reading

- Vasilis Syrgkanis and Éva Tardos. Composable and Efficient Mechanisms. STOC'13. (Smoothness for mechanisms)
- Paul Dütting and Thomas Kesselheim. Algorithms against Anarchy: Understanding Non-Truthful Mechanisms. EC'15. (Characterization of algorithms with small PoA)