

Bayes-Nash Equilibria

Thomas Kesselheim

Last Update: December 10, 2020

We have spent the past weeks discussing dominant-strategy incentive compatible (truthful) mechanisms. In these mechanisms, for every agent it is always a dominant strategy to report the true value. A classic example is the second-price auction. Today, we will broaden our perspective: What statements can we make if the mechanism is not truthful? For example, if it is a first-price auction?

A natural approach would be to consider Nash equilibria. For example, given tie breaking in our favor, the first-price auction has a pure Nash equilibrium, in which everybody bids their value except for the bidder of highest value. She bids the second-highest value. The weakness of this approach is that it requires full information: Essentially, the bidders have to know the other values.

Today, we will get to know an equilibrium concept for *incomplete information*. The players know their own values but only have a *prior belief* about the other players' values.

1 Bayes-Nash Equilibria

We will assume that bidder i 's value $v_i \in V_i$ is drawn independently from some distribution \mathcal{D}_i . These distributions are known to all bidders. A bidder chooses a bid b_i depending on the own valuation v_i , not knowing v_{-i} but only the distributions. We model this by saying that bidder i chooses a *bidding function* $\beta_i: V_i \rightarrow B_i$, mapping valuations to bids. Whenever the valuation is v_i , the bidder bids $\beta_i(v_i)$. For example, truthful bidding is represented by $\beta_i(v_i) = v_i$.

Definition 14.1 (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 $\beta(v) = (\beta_1(v_1), \dots, \beta_n(v_n))$.

So, we take the perspective of a single bidder. She knows her own v_i . The other values $v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n$ are drawn from $\mathcal{D}_1, \dots, \mathcal{D}_{i-1}, \mathcal{D}_{i+1}, \dots, \mathcal{D}_n$ respectively. The bidding function now tells her to bid $\beta_i(v_i)$. In an equilibrium, no other bid should give a higher utility. The other bidders keep playing according to the respective bidding functions. This, in particular, means that no other bidding function yields a higher expected utility when also taking the expectation over v_i .

Example 14.2. *In a truthful mechanism, $(\beta_i)_{i \in N}$ with $\beta_i(v_i) = v_i$ for all $i \in N$ and all $v_i \in V_i$ is a Bayes-Nash equilibrium. It is not necessarily the only one.*

Example 14.3. *Consider a first-price auction with two bidders, in which \mathcal{D}_i is the uniform distribution on $[0, 1]$ for $i = 1, 2$. Let us show that $(\beta_i)_{i \in N}$ with $\beta_i(v_i) = \frac{1}{2}v_i$ for all $i \in N$ is a Bayes-Nash equilibrium.*

Observe that for symmetry reasons, it is enough to only consider bidder 1. Fix any $v_1 \in V_1$ and let us write out the expected utility when bidding some arbitrary $b'_1 \in B_1$. The expectation is over bidder 2's value, respectively the bid.

$$\mathbf{E}_{v_2 \sim \mathcal{D}_2} [u_1((b'_1, \beta_2(v_2)), v_1)] = \int_0^1 u_1((b'_1, \beta_2(v_2)), v_1) dv_2 = \int_0^1 u_1\left(\left(b'_1, \frac{v_2}{2}\right), v_1\right) dv_2 \quad .$$

Here, we used that $\beta_2(v_2) = \frac{v_2}{2}$. Now, what is the value of $u_1((b'_1, \frac{v_2}{2}), v_1)$? If $b'_1 < \frac{v_2}{2}$, then it is 0, if $b'_1 > \frac{v_2}{2}$, then it is $v_1 - b'_1$. Therefore if $b'_1 \leq \frac{1}{2}$ then

$$\mathbf{E}_{v_2 \sim \mathcal{D}_2} [u_1((b'_1, \beta_2(v_2)), v_1)] = \int_0^{2b'_1} (v_1 - b'_1) dv_2 + \int_{2b'_1}^1 0 dv_2 = 2b'_1(v_1 - b'_1) = \frac{v_1^2}{2} - 2\left(b'_1 - \frac{v_1}{2}\right)^2 .$$

We see that that the last term is maximized exactly for $b'_1 = \frac{v_1}{2}$, so for all v_1 and b'_1

$$\mathbf{E}_{v_2 \sim \mathcal{D}_2} \left[u_1 \left(\left(\frac{v_1}{2}, \beta_2(v_2) \right), v_1 \right) \right] \geq \mathbf{E}_{v_2 \sim \mathcal{D}_2} [u_1((b'_1, \beta_2(v_2)), v_1)] ,$$

which is exactly the equilibrium condition.

2 Symmetric Bayes-Nash Equilibria of First-Price Auctions

We will derive a generalization of this equilibrium for arbitrary numbers of players n and arbitrary continuous, identical distributions $\mathcal{D}_1, \dots, \mathcal{D}_n$.

We will assume that for all $i \in N$ and all $x \in \mathbb{R}_{\geq 0}$

$$\Pr [v_i \leq x] = F(x) = \int_0^x f(t) dt .$$

We also write $G(x)$ for $(F(x))^{n-1}$.

Let us assume that there is a Bayes-Nash equilibrium $(\beta_i)_{i \in N}$ in which all functions are identical and differentiable as well as invertible. Then we have for all $y \in \mathbb{R}_{\geq 0}$

$$\mathbf{E}_{v_{-i} \sim \mathcal{D}_{-i}} [u_i((y, \beta_{-i}(v)), v_i)] = (v_i - y) \Pr \left[\bigwedge_{j \neq i} \beta_j(v_j) < y \right] = (v_i - y) \prod_{j \neq i} \Pr [\beta_j(v_j) < y]$$

If we let ϕ denote the inverse of β_i , then, $\Pr [\beta_j(v_j) < y] = \Pr [v_j < \phi(y)] = F(\phi(y))$ as $\beta_j = \beta_i$. So we get

$$\mathbf{E}_{v_{-i} \sim \mathcal{D}_{-i}} [u_i((y, \beta_{-i}(v)), v_i)] = (v_i - y) \prod_{j \neq i} F(\phi(y)) = (v_i - y) G(\phi(y)) .$$

If $\beta_i(v_i) = y$, then y has to be a local maximum of the above function. That is

$$\frac{d}{dy} (v_i - y) G(\phi(y)) = 0 .$$

The derivative can be calculated by standard rules

$$\frac{d}{dy} (v_i - y) G(\phi(y)) = -G(\phi(y)) + (v_i - y) G'(\phi(y)) \phi'(y) .$$

By the inverse function theorem, we have $\phi'(y) = \frac{1}{\beta'_i(\phi(y))}$. That is, if $\beta_i(v_i) = y$ then

$$-G(\phi(y)) + (v_i - y) G'(\phi(y)) \frac{1}{\beta'_i(\phi(y))} = 0 .$$

Replacing all occurrences of y by $\beta_i(v_i)$ (so $\phi(y) = v_i$), we get

$$-G(v_i) + (v_i - \beta_i(v_i)) G'(v_i) \frac{1}{\beta'_i(v_i)} = 0 ,$$

or equivalently

$$\beta'_i(v_i) G(v_i) + \beta_i(v_i) G'(v_i) = v_i G'(v_i) .$$

This has to hold for all $v_i \in \mathbb{R}_{>0}$. Observe that the left-hand side is exactly the derivative of $\beta_i G$. So, all solutions to this equation have the form

$$\beta_i(v_i)G(v_i) = \int v_i G'(v_i)dv_i + \text{constant} .$$

As $\beta_i(0) = 0$, we have

$$\beta_i(v_i) = \frac{1}{G(v_i)} \int_0^{v_i} tG'(t)dt .$$

One can verify that this is indeed an equilibrium the same way we did this in Example 14.3. And, as we have seen, it is necessarily the only symmetric equilibrium.

3 A Welfare Bound for First-Price Auctions

Let us have a closer look at the symmetric equilibrium that we have just derived. We observe that for any distribution the functions β_i are always strictly increasing. This means, whenever a bidder has a higher value, the bid will also be higher. Consequently, always the bidder with the highest value wins.


Observation 14.4. *In the symmetric Bayes-Nash equilibria $(\beta_i)_{i \in N}$ of a first-price auction with identical distributions for all $v \in V$*

$$\sum_{i \in N} v_i(f(\beta(v))) = \max_{i \in N} v_i .$$

If distributions are different, the equilibrium is usually asymmetric and it is not always true that the bidder with the highest value wins the item. For example, consider the case that v_1 is uniformly distributed on $[0, \frac{4}{5}]$ and v_2 is uniformly distributed on $[0, \frac{4}{5}]$. Using similar arguments as before, one can derive that the unique Bayes-Nash equilibrium is

$$\beta_1(v_1) = -\frac{1 - \sqrt{1 + v_1^2}}{v_1} \quad \text{and} \quad \beta_2(v_2) = \frac{1 - \sqrt{1 - v_2^2}}{v_2} .$$

With constant probability, it happens that $v_1 \in (\frac{4}{5}, 1]$ but $v_2 \in (\frac{\sqrt{2}}{2}, \frac{4}{5}]$. Whenever this is true, $v_1 > v_2$ but $\beta_1(v_1) \leq \beta_1(1) = \sqrt{2} - 1 = \beta_2(\frac{\sqrt{2}}{2}) < \beta_2(v_2)$. So, bidder 2 wins despite having the smaller value.



However, we can still derive a guarantee. This is in the spirit of a Price-of-Anarchy bound.

Theorem 14.5. *In any Bayes-Nash equilibrium $(\beta_i)_{i \in N}$ of a first-price auction*

$$\mathbf{E}_{v \sim \mathcal{D}} \left[\sum_{i \in N} v_i(f(\beta(v))) \right] \geq \frac{1}{2} \mathbf{E}_{v \sim \mathcal{D}} \left[\max_{i \in N} v_i \right] .$$

Before we come to the proof for Bayes-Nash equilibria, let us first see the argument in the full-information setting for pure Nash equilibria. That is, the valuations v and the bids b are fixed now.

It is important to observe that we can write the social welfare $\sum_{i \in N} v_i(f(b))$ also as the sum of utilities and payments: $\sum_{i \in N} v_i(f(b)) = \sum_{i \in N} u_i(b, v_i) + \sum_{i \in N} p_i(b)$.

Let i^* be a player of maximum value. If this bidder now bids $\frac{1}{2}v_{i^*}$, then her utility is $\frac{1}{2}v_{i^*}$ if she wins the item with this bid, meaning that $\max_{i \neq i^*} b_i < \frac{1}{2}v_{i^*}$. Otherwise it is 0. So, always the utility is at least $\frac{1}{2}v_{i^*} - \max_{i \neq i^*} b_i$

As we are in an equilibrium, $u_{i^*}(b, v_{i^*}) \geq u_{i^*}(\left(\frac{1}{2}v_{i^*}, b_{-i^*}\right), v_{i^*}) \geq \frac{1}{2}v_{i^*} - \max_i b_i$. Also, $u_i(b, v_i) \geq 0$ for all $i \in N$ because one option would be $b_i = 0$. Therefore

$$\sum_{i \in N} u_i(b, v_i) + \sum_{i \in N} p_i(b) \geq \frac{1}{2}v_{i^*} - \max_i b_i + \sum_{i \in N} p_i(b) = \frac{1}{2}v_{i^*} .$$

Proof of Theorem 14.5. We bound $\mathbf{E}_{v \sim \mathcal{D}} [\sum_{i \in N} u_i(\beta(v), v_i)]$. To this end, we use that for each bidder for each v_i

$$\mathbf{E}_{v_{-i} \sim \mathcal{D}_{-i}} [u_i(\beta(v), v_i)] \geq \mathbf{E}_{v_{-i} \sim \mathcal{D}_{-i}} \left[u_i \left(\left(\frac{v_i}{2}, \beta_{-i}(v) \right), v_i \right) \right] .$$

This holds for every v_i , so it also holds if we draw v_i from \mathcal{D}_i and take this expectation:

$$\mathbf{E}_{v \sim \mathcal{D}} [u_i(\beta(v), v_i)] \geq \mathbf{E}_{v \sim \mathcal{D}} \left[u_i \left(\left(\frac{v_i}{2}, \beta_{-i}(v) \right), v_i \right) \right] .$$

And by linearity of expectation, we also get

$$\begin{aligned} \mathbf{E}_{v \sim \mathcal{D}} \left[\sum_{i \in N} u_i(\beta(v), v_i) \right] &= \sum_{i \in N} \mathbf{E}_{v \sim \mathcal{D}} [u_i(\beta(v), v_i)] \\ &\geq \sum_{i \in N} \mathbf{E}_{v \sim \mathcal{D}} \left[u_i \left(\left(\frac{v_i}{2}, \beta_{-i}(v) \right), v_i \right) \right] \\ &= \mathbf{E}_{v \sim \mathcal{D}} \left[\sum_{i \in N} u_i \left(\left(\frac{v_i}{2}, \beta_{-i}(v) \right), v_i \right) \right] . \end{aligned}$$

For every fixed v , we also have

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

This gives us

$$\sum_{i \in N} u_i \left(\left(\frac{v_i}{2}, \beta_{-i}(v) \right), v_i \right) \geq \max_{i \in N} u_i \left(\left(\frac{v_i}{2}, \beta_{-i}(v) \right), v_i \right) \geq \max_{i \in N} \frac{v_i}{2} - \max_{i \in N} \beta_i(v_i) .$$

As we are in a first-price auction, $\max_{i \in N} \beta_i(v_i) = \sum_{i \in N} p_i(\beta(v))$, so

$$\sum_{i \in N} u_i \left(\left(\frac{v_i}{2}, \beta_{-i}(v) \right), v_i \right) + \sum_{i \in N} p_i(\beta(v)) \geq \max_{i \in N} \frac{v_i}{2} .$$

The rest follows directly by linearity of expectation. □

4 Outlook: Smooth Mechanisms

The last proof followed a very particular template: We use the fact that bidders do not want to deviate from the equilibrium to a fixed other strategy. We do not use further properties of the equilibrium—which is entirely different from the argument for symmetric equilibria. Indeed, there is a formalization of the latter proof pattern. In analogy to smooth games, we also call mechanisms smooth.

Definition 14.6 (Smooth Mechanism, simplified version). *Let $\lambda, \mu \geq 0$. A mechanism $\mathcal{M} = (f, p)$, $f: B \rightarrow X$, $p: B \rightarrow \mathbb{R}^n$, is (λ, μ) -smooth if for any valuation profile $v \in V$ for each player $i \in \mathcal{N}$ there exists a bid b_i^* such that for any profile of bids $b \in B$ we have*

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

In particular, our proof uses that a single-item first-price auction is $(\frac{1}{2}, 1)$ -smooth. It uses $b_i^* = \frac{v_i}{2}$. Next time, we will once again see this definition and how it allows us to bound the welfare in equilibria of other mechanisms.

References and Further Reading

- Vijay Krishna, Auction Theory, Academic Press. (Book on many aspects of auction theory, including the symmetric Bayes-Nash equilibria of the first-price auction.)
- Vasilis Syrgkanis and Éva Tardos. Composable and Efficient Mechanisms. STOC'13. (Smoothness for mechanisms)