The Bulow-Klemperer Theorem

CSCI 1440/2440

2025-10-15

We present the Bulow-Klemperer theorem, which argues that a straightforward way to surpass the revenue of Myerson's optimal auction is to run a Vickrey auction with one additional bidder.

1 *The Bulow-Klemperer Theorem*

Myerson's optimal auction relies on knowledge of the distribution of bidders' values, while the Vickrey auction does not. What is the relationship between the revenue earned by the Vickrey auction, as compared to Myerson's optimal auction?[1]

Bulow and Klemperer [2] provide an interesting take on this question. They argue that increasing competition (i.e., attracting more bidders) is at least as important as designing an optimal auction.

[1] In general, the performance of the former can be arbitrarily worse than that of the latter: e.g., assume one select bidder's values are distributed in $[K, K+1]$, for some very large $K \gg 0$, while all others are distributed in $[0, \epsilon]$, for some very small $\epsilon > 0$.

[2] Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *The American Economic Review*, 86(1):180–194, 1996

Theorem 1.1. *Let us denote by $Rev(VA_{n+1})$ the revenue of the Vickrey auction with $n+1$ bidders. Let us further denote by $Rev(OPT_n)$ the revenue of the optimal auction with n bidders whose values are drawn i.i.d. from some regular distribution F. Then*

$$\mathbb{E}_{v_1,\ldots,v_{n+1}\sim F}\left[Rev(VA_{n+1})\right] \geq \mathbb{E}_{v_1,\ldots,v_n\sim F}\left[Rev(OPT_n)\right] \quad . \tag{1}$$

The proof relies on the following simple lemma.

Lemma 1.2. *The Vickrey auction is revenue-optimizing among all auctions that always allocate the good, when the bidders values are drawn i.i.d. from a regular distribution.*

Proof. By Myerson's theorem, expected revenue equals expected virtual value. The optimal auction thus allocates in order of virtual values—although it does not allocate to negative virtual values. Likewise, the optimal auction that *always* allocates also allocates in order of virtual values, even if virtual values are negative.

Assuming the bidders' values are drawn i.i.d. from a regular distribution, all bidders share the same weakly increasing virtual value function. So allocating to the bidder with the highest value, as the Vickrey auction does, is the same as allocating to the bidder with the highest virtual value, as Myerson's optimal auction does.

Moreover, inverting the virtual value function at the second-highest virtual value recovers the second-highest bid/value. □

Proof of Theorem 1.1. Let $\mathscr{A}_{n+1}$ be the following auction, run with $n+1$ bidders.

1. Run the optimal auction OPT_n on the first n bidders.
2. If the good is not allocated via this optimal auction, then allocate it to bidder $n+1$ for free.

Here are two important observations about Auction $\mathscr{A}_{n+1}$:

1. The good is always allocated.
2. $\mathbb{E}_{v_1,\ldots,v_{n+1}\sim F}\left[\mathrm{Rev}(\mathscr{A}_{n+1})\right] = \mathbb{E}_{v_1,\ldots,v_n\sim F}\left[\mathrm{Rev}(\mathrm{OPT}_n)\right]$.

Auction $\mathscr{A}_{n+1}$ thus serves as a bridge between Vickrey's auction and Myerson's, via Lemma 1.2:

$$\mathbb{E}_{v_1,\ldots,v_{n+1}\sim F}\left[\mathrm{Rev}(\mathrm{VA}_{n+1})\right] \geq \mathbb{E}_{v_1,\ldots,v_{n+1}\sim F}\left[\mathrm{Rev}(\mathscr{A}_{n+1})\right] \quad (2)$$

$$= \mathbb{E}_{v_1,\ldots,v_n\sim F}\left[\mathrm{Rev}(\mathrm{OPT}_n)\right] \quad (3)$$

□

References

[1] Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *The American Economic Review*, 86(1):180–194, 1996.