DS 574 **Prior Independence**
Fall 2026 Prof. Kira Goldner

Robustness: Prior-Independence

Previously, we knew the distribution F and designed an auction for it. What if we do not know F? Values are still drawn from a distribution, and we still compare our revenue with the optimal auction that knows it. The mechanism we run must work without knowing F.

We consider one item and independent values drawn from the same continuous distribution F, with nonnegative values, positive density f, and finite expected value. Assume F is regular: $\varphi(v) = v - (1 - F(v))/f(v)$ is nondecreasing in value. Write $\textsc{opt}(n, F)$ for the optimal expected revenue with n buyers, among truthful, individually rational auctions.

The Bulow–Klemperer Result

One famous result takes the form of resource augmentation.

Theorem 1 (Bulow–Klemperer)**.** *For i.i.d. regular single-item environments,*

$$\textsc{Rev}_{\text{VCG}}(n+1, F) \geq \textsc{opt}(n, F).$$

Vickrey is the second-price auction with no reserve: allocate to a highest bidder and charge the second-highest bid.

What extra resource replaces knowledge of F? Which assumptions make the comparison possible?

The optimal auction can leave the item unsold. Why might that help its revenue?

The Bulow–Klemperer proof

First, require allocation. Among truthful, individually rational auctions that must allocate, which allocation maximizes expected virtual welfare? Use i.i.d. regularity to identify the corresponding auction.

Now compare with the optimal auction. Run the optimal auction on the first n buyers. If it does not sell, how can we use buyer $n+1$ to obtain an always-allocate auction M with the same revenue? Check truthfulness and individual rationality.

Finish the revenue comparison:

$$\textsc{opt}(n, F) \qquad \textsc{Rev}_M(n+1, F) \qquad \textsc{Rev}_{\text{VCG}}(n+1, F).$$

Quantile Space

Instead of describing a price by its dollar value, describe how often it sells. Assume F is continuous and strictly increasing on its support. A value v has *tail quantile* $q = 1 - F(v)$: the fraction of the distribution with higher values. Write $v(q) = F^{-1}(1 - q)$ for the corresponding value.

For $U[0, 10]$, a value of 9 is in the top 10%: $q = 0.1$ and $v(0.1) = 9$. What is the quantile of value 2? Does a smaller quantile mean a higher or lower value?

Revenue curve. A price $v(q)$ sells with probability q, so its expected revenue is $R(q) = qv(q)$. For $U[0, 10]$, express $R(q)$ and sketch it. Label the axes and the best posted price's expected revenue.

A random quantile. Draw $S \sim F$ and let $Q = 1 - F(S)$. The event $Q \leq a$ means the sample lies in the top fraction a of values. What is its probability? What distribution does Q have?

The Single Sample Mechanism

Let $S \sim F$ be one sample, independent of the buyer's value $V \sim F$. Post price S: the buyer pays S if $V \geq S$, and zero otherwise. The sample is not an eligible buyer.

Theorem 2 (Single-sample pricing)**.** *For a continuous regular F as above, posting an independent sample as a price earns at least $\frac{1}{2}\textsc{opt}(1, F)$ in expectation over both draws. Assume the optimal posted price is attained.*

Proof: in quantile space. First fix the sample at s. What is the expected revenue over the buyer? Then average over the sample's quantile to express revenue as an area.

Let q^* maximize $R(q)$. Regularity makes R concave, and its endpoint revenues are nonnegative. Draw the curve and represent $\textsc{opt}(1, F) = R(q^*)$ by a rectangle of width 1. Find an area beneath the curve that proves the guarantee.

What information does the seller actually need to run this mechanism?