543: Game Theory **Week 6**
Robustness and worst-case analysis

Updated July 26, 2021.

Preliminary and incomplete.

Thanks to 534 Fall 2020 students for proofreading! (Remaining errors are mine.)

Source(s): [Bergemann and Schlag, 2008], [Bergemann and Schlag, 2011], [Bulow and Klemperer, 1996], [Dhangwatnotai et al., 2015].

1 Robust monopoly pricing

Based on results from [Bergemann and Schlag, 2008] and [Bergemann and Schlag, 2011].

Let's go to our standard screening setup.

There is a single seller, a single buyer, and single product. Buyer's value $\theta \sim F$ with support in $[0, 1]$. As before, θ is privately known only to the buyer. The seller sells the product by offering possibly randomized prices. I.e., she chooses a distribution G over prices (recall from week 1 that we show that restricting attention to randomized pricing mechanisms is without loss of generality.)

We want to now relax the assumption that the seller knows F. How should we think about the problem?

1.1 Pricing without priors

A first try: The seller wishes to design a mechanism that has the highest possible profit guarantee over all possible distributions, that is,

$$\max_G \min_F E_{p\sim G}[p(1 - F(p))].$$

(Technically we should use sup and inf, that are always well-defined. But I'll be lazy and use max and min throughout.) The solution to this problem is 0. Why?

So how do we make the problem "interesting"?

1.2 Nearby priors

Now let's suppose that the seller has some information about F. In particular, she knows that F is close to some prior F_0.

For a given distribution F_0, consider the set of distributions that are ϵ-close to F_0 under the Prohorov metric (which metricizes weak convergence),

$$\mathcal{F}_\epsilon(F_0) = \{F | F(A) \leq F_0(A^\epsilon) + \epsilon, \forall A \subseteq [0,1]\},$$

(Why is the alternative and seemingly more natural definition $\{F | F(A) \leq F_0(A) + \epsilon, \forall A \subseteq [0,1]\}$ not a good one?)

The seller's problem is to maximize the worst case profit over all distributions that are ϵ-close to F_0,

$$\max_G \min_{F \in \mathcal{F}_\epsilon(F_0)} E_{p\sim G}[p(1 - F(p))].$$

Turns out that there is a single $F \in \mathcal{F}_\epsilon(F_0)$ that minimizes the expected profit $E_{p\sim G}[p(1 - F(p))]$ for any G. That distribution is F_ϵ defined as follows:

$$F_\epsilon(v) = \min(F_0(v + \epsilon) + \epsilon, 1)$$

Why? By definition we know that $F(v) \leq F_0(v + \epsilon) + \epsilon$ and also $F(v) \leq 1$. So $F(v) \leq \min(F_0(v + \epsilon) + \epsilon, 1)$. So the point-wise largest possible F is F_ϵ. This distribution minimizes $p(1 - F(p))$ for any given p over all distributions in $\mathcal{F}_\epsilon(F_0)$.

The problem becomes

$$\max_G E_{p\sim G}[p(1 - F_\epsilon(p))].$$

That is, to maximize profit for distribution F_ϵ. We know from week 1 that the solution is to choose a price deterministically. So we have the following result.

Proposition 1.1 *It is optimal to post a deterministic price.*

How does the optimal price $p(\epsilon)$ change in ϵ? It decreases, at least for small enough ϵ.

Proposition 1.2 $p'(0) < 0$.

Proof: See [Bergemann and Schlag, 2011]. **QED**

1.3 Regret

Let's go back to the case where the seller knows nothing about F. But instead consider another objective: regret.

Regret is the difference between the revenue that the seller obtains, and the highest possible revenue in hindsight. So if the value is θ and the price is p, then regret is

$$r(p, \theta) = \theta - p\mathbb{1}_{\theta \geq p}$$

So the problem is

$$\max_G \min_F -r(G, F).$$

where

$$r(G, F) = E_{p \sim G, \theta \sim F}[r(p, \theta)]$$

So this is a two-player zero sum game where the seller chooses p, the nature chooses θ, the payoff to seller is $-r(p, \theta)$, and the payoff to nature is $r(p, \theta)$. The solution is a pair (G, F) satisfying the best-response conditions

$$r(G, F') \leq r(G, F) \leq r(G', F), \forall G', F'.$$

The game has no equilibrium in which the seller chooses a deterministic price p. Suppose the seller chooses p. Then if $\theta \geq p$, regret is $\theta - p$, in which case highest possible regret is $1 - p$ ("upside exposure"). And if $\theta < p$, regret is θ, in which case highest possible regret is p ("downside exposure"). Now if $p > 1/2$, nature's 'doesn't have a best response, since it would like to choose θ strictly less than p, but not equal to it. If $p \leq 1/2$, nature's best unique response is to choose $\theta = 1$, but then the firm's best reply is $p = 1$. The paper claims a bit sloppily that if the firm is forced to choose deterministic prices, then there is an equilibrium with $p = 1/2$. (I think one way to fix this is to consider the non-zero-sum game where nature's payoff if $\theta = p$ is θ, but seller's is still zero. Then $p = 1/2$ and θ drawn uniformly from $\{1/2, 1\}$ is an equilibrium.) Intuitively this makes sense as the seller balances the upside and the downside exposures. Similarly there is no equilibrium where nature chooses θ deterministically.

But this is not an equilibrium of our original game, where regret if $\theta = p$ is 0, since if $p = 1/2$, then nature's unique best response is $\theta = 1$, but $p = 1/2$ is not a best response to $\theta = 1$.

Our original game will have a mixed equilibrium.

Proposition 1.3 *There is a unique equilibrium* (G, F) *where*

$$G(p) = \begin{cases} 1 + \ln p & \textit{if } p \geq 1/e. \\ 0 & \textit{if } p < 1/e. \end{cases}$$

Why? For the seller to randomize over some prices, she must be indifferent between them. For a given F, regret is

$$r(G, F) = E_{p \sim G, \theta \sim F}[\theta - p\mathbb{1}_{\theta \geq p}] = E_\theta[\theta] - E_p[p(1 - F(p))]$$

The first term is independent of G. So for a give F, the seller minimizes regret by maximizing revenue. So if she randomizes over an interval, the revenue curve should be flat over that interval.

$$\frac{d}{dp}(p(1 - F(p))) = 0$$

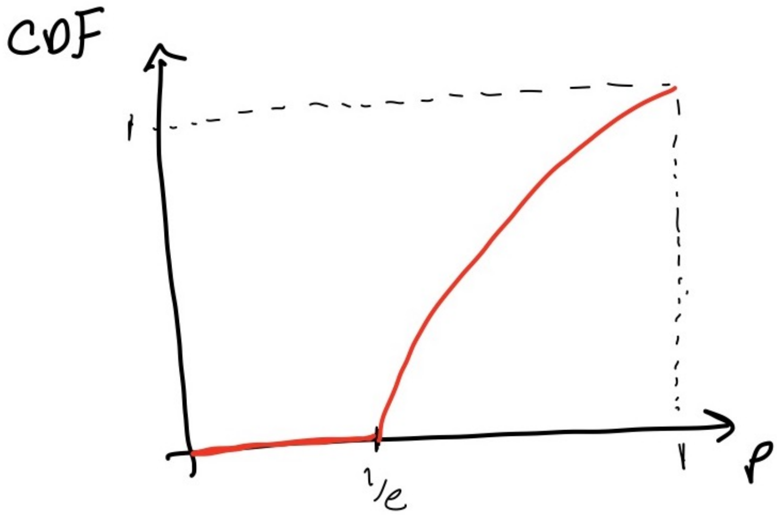

The revenue curve associated with the distribution identified below is flat over interval $[1/e, 1]$. And prices below $1/e$ are not optimal.

$$F(\theta) = \begin{cases} 1 & \text{if } \theta = 1 \\ 1 - 1/(e\theta) & \text{if } 1/e \leq \theta < 1. \\ 0 & \text{if } \theta < 1/e. \end{cases}$$

This is because $p(1 - F(p)) = 1/e$ over the interval. So in fact any distribution over prices in $[1/e, 1]$ is optimal for the seller.

Also for nature to be indifferent over some values in an interval, he must be indifferent over those values:

$$\frac{d}{d\theta}(r(G, \theta)) = 0$$

That is, using g for the density function,

$$\frac{d}{d\theta}(\theta - \int_{p \leq \theta} pdG) = 0$$

$$1 - \theta g(\theta) = 0$$

$$g(\theta) = 1/\theta.$$

Which is the property satisfied by the distribution G identified in Proposition 1.3 above on interval $[1/e, 1]$.

To formally prove the proposition, one needs to show that the G and F above are in fact best responses to each other. We already argued that given F above, any price in $[1/e, 1]$ maximizes revenue (and therefore minimizes regret). Takes a bit more work, but is straightforward, to show F maximizes regret given G.

Why does $1/e$ appear?

1.4 Regret with nearby priors

One could combine the above two sections by seeking to solve

$$\max_G \min_{F \in \mathcal{F}_\epsilon(F_0)} -r(G, F).$$

Turns out the answer looks very similar to the case of no information above, and is provided in [Bergemann and Schlag, 2011].

1.5 Other restrictions on the class of priors

Going back to maximizing worst-case revenue, another way to ask the question is what if the seller knows certain properties of the distribution, such as its mean or variance. [Carrasco et al., 2018] studies this. We are not going to cover it.

2 Revenue Maximization from Samples

Suppose now that the seller doesn't know the distribution, but can engage in "market research". How should the seller choose prices given the results of market research?

First, a very cool digression.

2.1 Auctions vs. Negotiations

This is from [Bulow and Klemperer, 1996].

Setting:

- There is a single seller, n buyers, single product.
- Each buyer i's value θ_i is drawn independently from an identical distribution F with density f
- Suppose that F is regular: $\phi(\theta) = \theta - \frac{1-F(\theta)}{f(\theta)}$ is non-decreasing

Recall Myerson. With non-identical distributions, each bid is mapped to a virtual value, and then the product is given to the highest positive virtual value. But this specification heavily depends on the details of the distribution. Even with identical distributions, the seller needs to know F to calculate the reserve price.

Proposition 2.1 *The second price auction with reserve price $\phi^{-1}(0)$ is optimal.*

What if the seller doesn't know F, or only approximately knows F? Is there a "detail-fee" auction that is approximately optimal?

Proposition 2.2

$$\textit{Revenue of the second price auction} \geq \frac{n-1}{n} \cdot \textit{revenue of the optimal auction}$$

Proof: What auction maximizes revenue among all auctions that are constrained to always sell the product? Second price: you want to give the product to the highest virtual value, regardless of its sign. Since the distributions are identical and regular, this means giving the product to the highest bid. This is what the second price auction does. Let Π be the revenue.

Now consider auction A: it runs the optimal auction for players $1, \ldots, n-1$. If it the product is not sold to any of them, it is given for free to buyer n. Revenue of this auction is $OPT(n-1)$, where $OPT(n)$ denotes the optimal revenue from n buyers. But this auction always sells the product, so its revenue is at most the revenue of the second price auction Π.

$$\Pi \geq OPT(n-1) \geq \frac{n-1}{n} OPT(n).$$

Prove the last inequality at home. **QED**

Bulow and Klemberer originally did not interpret this as an approximation result. Instead, they asked the following question: is it more valuable to "recruit" an extra buyer, or to "learn" the distribution of values. Answer: recruit the buyer.

Yet another interpretation is that the second price auction resembles the case where the seller has not bargaining power. Instead, the buyers offer prices, the highest of which the seller accepts. Then in equilibrium the "competitive price" is equal to the second highest value. Which is what the second price auction achieves. So in this sense, the result says that having extra competition is more important than having bargaining power.

2.2 Revenue maximization from a single sample

This is from [Dhangwatnotai et al., 2015].

Setting:

- Single seller, single buyer, single product
- Value θ drawn from regular distribution F, known to buyer
- Seller doesn't know θ nor F
- Seller observes a draw of a random variable v from F (independent form θ)

- Seller chooses a direct IC and IR mechanism (q^v, t^v)

Proposition 2.3 *The revenue of the mechanism that offers the product at price v is at least half of the optimal revenue if the distribution was known* $\max_p p(1-F(p))$.

Proof: Bulow-Klemperer tells us that $p(1-F(p))$ is at most the revenue of the second price auction with two players. In the second price auction, each buyer faces a random reserve price (the bid of the other player) drawn from F. So the revenue of the second price auction is twice the revenue of the single-sample mechanism that uses v as a posted price. **QED**

A geometric proof.

Consider the revenue in "quantile space", that is, for each $q \in [0,1]$, $R(q)$ is the revenue of the price p that is accepted with probability q, $q = 1-F(p)$. The optimal revenue is $\max_q R(q)$, which is the area of the rectangle below. What is the revenue of posting v as a price? What is the probability with which v is accepted? It is q, where q is the quantile of v. The revenue is $R(q)$. So it is as if we draw q uniformly from $[0,1]$, and then revenue is $R(q)$. The expected revenue is $\int R(q)dq$, which is the area below R. This area is at least half of $\max_q R(q)$ since R is concave.

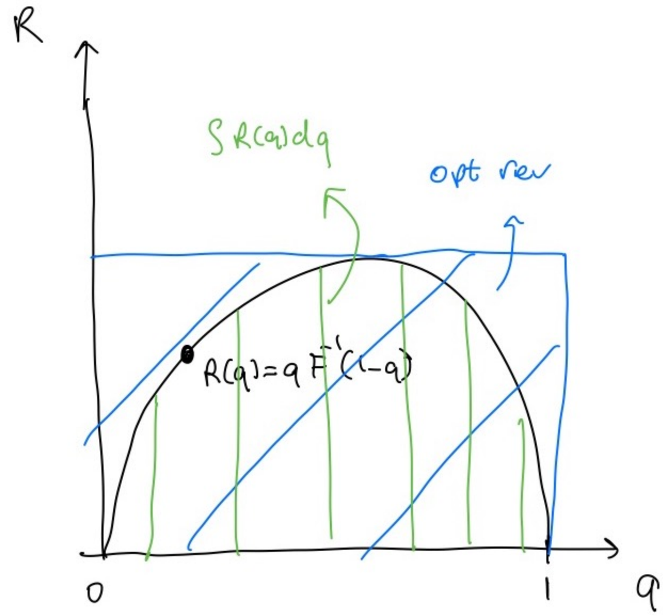

2.3 Multiple buyers, single sample

Turns out that things are significantly different with multiple buyers. [Fu et al., 2021] show that if there are k possible distributions, then the seller can extract the full surplus with at most $k-1$ samples!

3 Multiple products with known marginals

This is from [Carroll, 2017].

Setup:

- Single seller, single buyer
- Products $1, \ldots, n$
- Values $\theta_1, \ldots, \theta_n$, "additive"
- θ drawn from F
- Seller doesn't know F or θ
- Seller knows marginals $F_1, \ldots, F_n$
- $F_i(x) = Pr_{\theta \sim F}[\theta_i \leq x]$
- Mechanism $q : \Theta \to [0,1]^n$, $t : \Theta \to R$

Problem is

$$\max_{\text{IC IR } (q,t)} \min_{F \text{ with marginals } F_1,\ldots,F_n} E_{\theta \sim F}[t(\theta)]$$

Recall: in the standard setup where the seller knows F, selling separately is never optimal.

Proposition 3.1 *Selling separately is optimal.*

Proof: We prove the proposition assuming identical marginals $F_1 = \ldots = F_n = \hat{F}$.

For any distribution F with marginals $\hat{F}$, revenue is $n \times \max_p p(1 - \hat{F}(p))$.

To show that selling separately is optimal, we show that for any mechanism, there exists a distribution F such that the revenue of the mechanism is at most $n \times \max_p p(1 - \hat{F}(p))$. In fact, there is a single distribution that achieves this for all mechanisms: perfect correlation. That is, $\theta_1 = \ldots = \theta_n$. So to draw a type from F, we can draw x from $\hat{F}$, and then set $\theta_1 = \ldots = \theta_n = x$.

What is the optimal revenue for F? It is as if we are selling a single product with value nx, where x is drawn from $\hat{F}$. From week 1 we know that the optimal revenue is $\max_p pPr[nx \geq p] = n \max_p (1 - \hat{F}(p))$. **QED**

References

[Bergemann and Schlag, 2011] Bergemann, D. and Schlag, K. (2011). Robust monopoly pricing. Journal of Economic Theory, 146(6):2527–2543.

[Bergemann and Schlag, 2008] Bergemann, D. and Schlag, K. H. (2008). Pricing without priors. Journal of the European Economic Association, 6(2-3):560–569.

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

[Carrasco et al., 2018] Carrasco, V., Luz, V. F., Kos, N., Messner, M., Monteiro, P., and Moreira, H. (2018). Optimal selling mechanisms under moment conditions. Journal of Economic Theory, 177:245–279.

[Carroll, 2017] Carroll, G. (2017). Robustness and separation in multidimensional screening. Econometrica, 85(2):453–488.

[Dhangwatnotai et al., 2015] Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2015). Revenue maximization with a single sample. Games and Economic Behavior, 91:318–333.

[Fu et al., 2021] Fu, H., Haghpanah, N., Hartline, J., and Kleinberg, R. (2021). Full surplus extraction from samples. Journal of Economic Theory, 193:105230.