ECON 289, Lecture 6: Maxmin mechanisms

Instructor: Ben Brooks

Spring 2023

Pushing beyond ex post implementation

- Previous unit explored whether ex post implementation addresses our concern about misspecification of beliefs
- Answer: only under fairly restrictive assumptions (KPT type spaces, BM separability, Chung-Ely regularity)
- A case in point is common value auctions, where KPT type spaces are highly restrictive, and rules out the kind of flexible information we allowed on robust predictions for the FPA
- In this unit, we push beyond ex post implementation, and allow the implemented outcome to vary with beliefs
- To do so, we will combine two methodologies:
  Bayesian mechanism design and robust predictions
- After general remarks, we'll study revenue maximization in common-value auctions, following Brooks and Du (2021)

Primitives

- Agents $i = 1, \ldots, N$
- Payoff state Θ
- Known distribution $\mu \in \Delta(\Theta)$
- Finitely many outcomes X
- Agents' utilities $u_i(x, \theta)$, outside option normalized to 0
- Designer's payoff $w(x, \theta)$
- Agents beliefs are described by a (finite) Harsanyi type space (T, π), where $\pi \in \Delta(T \times \Theta)$ with marginal μ on Θ
- Designer chooses a (finite) mechanism (A, g)
- Agents play a Bayes Nash equilibrium
- For a given mechanism (A, g) and type space (T, π), $E(A, g, T, \pi)$ is the set of equilibria

Guarantee

- In our unit on robust predictions, we argued that there is a revelation principle for information design:
  Any outcome that can be implemented in some (common prior) type space and equilibrium can also be implemented in a **direct recommendation type space** and **obedient equilibrium**, i.e., a BCE
- We used this to compute min revenue in a common value FPA
- We can call this the FPA's **guarantee**:
  Minimum designer welfare across all type spaces/equilibria
- Question: What mechanisms maximize the guarantee?
  E.g., in common value auctions, can we improve on the FPA?

Participation security

- In thinking about maximizing the guarantee, obviously we want to impose some kind of limitation on the designer to respect participation constraints
- A mechanism (A, g) is **participation secure** if for all i, there exists a message $0 \in A_i$ such that for all θ and a_{-i},

$$\sum_x u_i(x, \theta) g(x|0, a_{-i}) \geq 0,$$

  i.e., $a_i = 0$ guarantees player i a non-negative payoff ex post
- Henceforth, we restrict the designer to participation secure mechanisms
- We propose to solve:

$$\sup_{(A,g) \text{ p.s.}} \inf_{(T,\pi)} \inf_{\sigma \in E(A,g,T,\pi)} \sum_{\theta,t,a,x} \pi(t, \theta) \sigma(a|t) g(x|a) w(x, \theta)$$

Potential

- In our unit on Bayesian mechanism design, we argued that there is a revelation principle for mechanism design: Any outcome that can be implemented in some mechanism and equilibrium can also be implemented in a **direct revelation mechanism** and **truthful equilibrium**
- BBM (2020) computed max revenue in the maximum type space (worst case for the FPA with common values)
- We may call this the **potential** of the type space: Maximum designer welfare across all mechanisms/equilibria
- Question: Which type spaces minimize the potential?
- In common value auctions, is something worse than the maximum type space?
- Formally, minimizing the potential:

$$\inf_{(T,\pi)} \sup_{(A,g)} \sup_{\sigma \in E(A,g,T,\pi)} \sum_{\theta,t,a,x} \pi(t,\theta)\sigma(a|t)g(x|a)w(x,\theta)$$

Two programs

- Maximize the guarantee:

$$\sup_{(A,g)} \inf_{(T,\pi)} \inf_{\sigma \in E(A,g,T,\pi)} \sum_{\theta,t,a,x} \pi(t,\theta)\sigma(a|t)g(x|a)w(x,\theta)$$

- Minimizing the potential

$$\inf_{(T,\pi)} \sup_{(A,g)} \sup_{\sigma \in E(A,g,T,\pi)} \sum_{\theta,t,a,x} \pi(t,\theta)\sigma(a|t)g(x|a)w(x,\theta)$$

- Questions:
  - When do these coincide?
  - What are the optimal "robust" mechanisms?
  - What are worst-case type spaces?

Looking for saddle points

- This is a saddle point problem, analogous to the LP, where the max guarantee and min potential are dual to one another
- In particular, we have an analogue of weak duality: Max guarantee $\leq$ min potential
- Why? Suppose (A, g) (approximately) achieves the max guarantee of $\underline{W}$; then for any type space, the designer could run (A, g), and in every equilibrium, welfare will be at least $\underline{W}$, so every type space has a potential greater than $\underline{W}$
- Because of weak duality, a sufficient condition for (A, g) and (T, π) to be solutions is that they are a **saddle point**: Each is feasible for their respective programs, and the guarantee of (A, g) is equal to the potential of (T, π)
- Hence, we will investigate the existence of saddle points

Upper bounding the potential

- Suppose we fix (T, π)
- Clearly, we can upper bound the potential by maximizing designer welfare across interim IC/IR direct mechanisms
- The program is

$$\max_{g:T\to\Delta(X)} \sum_{t,\theta,x} w(x,\theta)g(x|t)\pi(\theta,t)$$

$$\text{s.t.} \quad \sum_{x,\theta,t_{-i}} u_i(x,\theta)(g(x|t_i,t_{-i}) - g(x|t_i',t_{-i}))\pi(\theta,t_i,t_{-i}) \geq 0 \;\forall i, t_i, t_i' \;[\alpha_i(t_i,t_i')]$$

$$\sum_{x,\theta,t_{-i}} u_i(x,\theta)g(x|t_i,t_{-i})\pi(\theta,t_i,t_{-i}) \geq 0 \;\forall i, t_i, t_i' \;[\beta_i(t_i)]$$

Upper bounding the potential

- For any choice of Lagrange multipliers $\alpha_i(t_i, t_i')$ and $\beta_i(t_i)$, we have an upper bound on the potential:

$$\max_{g:T\to\Delta(X)} \sum_{t,\theta,x} \pi(\theta,t)\Big[w(x,\theta)g(x|t) + \sum_{i,t_i'} \alpha_i(t_i,t_i')u_i(x,\theta)\left[g(x|t) - g(x|t_i',t_{-i})\right] + \sum_i \beta_i(t_i)u_i(x,\theta)g(x|t)\Big]$$

- The optimal IC/IR multipliers would make this bound tight, but we have an upper bound for any choice of multipliers
- We could minimize this bound across all (T,π) to get an upper bound on the min potential...

Lower bounding the guarantee

- Similarly, fixing (A, g), we could lower bound the guarantee by minimizing across BCE
- This program is:

$$\min_{\pi \in \Delta(\Theta \times A)} \sum_{a,\theta,x} w(x,\theta) g(x|a) \pi(\theta, a)$$

$$\text{s.t.} \sum_{x,\theta,a_{-i}} u_i(x,\theta) \left[g(x|a_i, a_{-i}) - g(x|a_i', a_{-i}) \right] \pi(\theta, a_i, a_{-i}) \geq 0 \ \forall i, a_i, a_i' \ [\alpha_i(a_i, a_i')]$$

Lower bounding the guarantee

- Again, looking at the Lagrangian, we have multipliers $\alpha(a_i, a_i')$ on obedience
- So, for any choice of multipliers, we have a lower bound

$$\min_{\pi \in \Delta(\Theta \times A)} \sum_{a,\theta,x} \pi(\theta, a) \Big[w(x,\theta) g(x|a) - \sum_{i,a_i'} \alpha_i(a_i, a_i') u_i(x,\theta) \left[g(x|a) - g(x|a_i', a_{-i}) \right] \Big]$$

- The optimal obedience multipliers would make this bound tight, but we have a lower bound for any choice of multipliers
- Again, we could maximize this bound across participation secure (A, g) to get a lower bound on the max guarantee...

But what are the right multipliers?

- In a sense, it should be the multipliers that make these bounds as close together as possible:

$$\min_{(T,\pi)} \max_{g} \sum_{t,\theta,x} \pi(\theta,t)\Big[w(x,\theta)g(x|t) + \sum_{i,t_i'} \alpha_i(t_i,t_i')u_i(x,\theta)\left[g(x|t) - g(x|t_i',t_{-i})\right]$$

$$+ \sum_i \beta_i(t_i)u_i(x,\theta)g(x|t)\Big]$$

$$\max_{(A,g)\text{ p.s.}} \min_{\pi} \sum_{a,\theta,x} \pi(\theta,a)\Big[w(x,\theta)g(x|a) - \sum_{i,a_i'} \alpha_i(a_i,a_i')u_i(x,\theta)(g(x|a) - g(x|a_i',a_{-i}))\Big]$$

- But we have added constraints on the top and subtracted them on the bottom...
- Also the order of moves is different, and we have participation security on the outside in the guarantee lower bound, and IR in the Lagrangian for the potential upper bound

Perhaps the FPA multipliers???

- We only have one example of multipliers for which we've already solved the guarantee for some mechanism: The FPA in the common value case!
- In that model, the actions are ordered (bids) and multipliers correspond to **uniform upward deviations**: $\alpha_i(a_i, a_i') > 0$ iff $a_i' > a_i$, and then they depend only on the deviation a_i', and not the recommendation a_i
- For that choice of multipliers, with respect to an (arbitrary) linear order on actions, we have a non-trivial lower bound when we pick (A, g) to be the first price auction, and we get the "maximum" type space to be the worst case

Second thoughts...

- But even for the common value auction, these don't seem like the right multipliers, for a couple of reasons
- First, BBM (2020) solve for the revenue maximizing mechanism on the maximum type space, and its' value is significantly higher than the guarantee of the FPA
- So, if these were the right multipliers, and the FPA was maxmin optimal, then the minmax type space would have to be different than the type space that's the worst case for the FPA, which doesn't sound quite right
- From another point of view, the uniform upward multipliers would not massage any of the main differences between the two bounding programs, and in particular it doesn't address the different sign on the obedience/truthfulness constraints
- It could be that there is just a duality gap, but let's not give up faith in the universe just yet

Local multipliers for the potential

- Perhaps based on the prominent role played by local constraints in our analyses thus far, we may make an ansatz that local constraints will help us here as well
- Suppose we set

$$A_i = T_i = \{0, \Delta, 2\Delta, \ldots, 1/\Delta\}$$

- Take all multipliers zero except $\alpha_i(t_i, t_i - \Delta)$ and $\beta_i(0)$
- But what exact value for the multipliers?

But which local multipliers?

- In the IPV model, the multiplier has a specific value that is related to the distribution of v_i, but this is just a normalization
- For example, we could have defined $t_i = F_i(v_i)$ (the quantile of v_i), so that $t_i \sim U[0,1]$, the inverse hazard rate of t_i is $1 - t_i$, and the value function is $w_i(t_i) = F_i^{-1}(v_i)$
- Bottom line: we can choose whatever units we want for t_i, which will determine size of a "local deviation" and hence the multipliers, and we might as well choose:

$$\alpha_i(t_i, t_i - \Delta) = \beta_i(0) = 1/\Delta + \mathbb{I}_{t_i = 1/\Delta}\,(1 - 1/\Delta)$$

- (Why the adjustment for the highest type? This is a subtle issue, that I'll come back to later)
- Then the min potential bound simplifies to:

$$\min_{\pi} \max_{g} \sum_{t,\theta,x} \pi(\theta,t)\left[w(x,\theta)g(x|t) + \sum_i u_i(x,\theta)\nabla_i^- g(x|t)\right]$$

$$\nabla_i^- g(x|t) = \begin{cases} g(x|t) - g(x|t_i - \Delta, t_{-i}) & \text{if } t_i = 1/\Delta \\ \frac{1}{\Delta}\left[g(x|t) - g(x|t_i - \Delta, t_{-i})\right] & \text{if } 0 < t_i < 1/\Delta \\ \frac{1}{\Delta} g(x|t) & \text{if } t_i = 0 \end{cases}$$

Local multipliers for guarantee

- But what about the guarantee? Here we want local multipliers as well, but we need to reverse sign, which means reversing the **direction**
- In particular, take all multipliers zero except **local up**

$$\alpha_i(a_i, a_i + \Delta) = 1/\Delta$$

- Then the max guarantee simplifies to

$$\max_{g \text{ p.s.}} \min_{\pi} \sum_{a,\theta,x} \pi(\theta, a) \left[w(x,\theta) g(x|a) + \sum_i u_i(x,\theta) \nabla_i^+ g(x|a) \right]$$

$$\nabla_i^+ g(x|a) = \mathbb{I}_{a_i < \frac{1}{\Delta}} \frac{1}{\Delta} u_i(x,\theta)(g(x|a_i + \Delta, a_{-i}) - g(x|a))$$

Assessing our progress so far

$$\min_{\pi} \max_{g} \sum_{t,\theta,x} \pi(\theta, t)\Big[w(x,\theta)g(x|t) + \sum_i u_i(x,\theta)\nabla_i^- g(x|t)\Big]$$

$$\max_{g \text{ p.s.}} \min_{\pi} \sum_{a,\theta,x} \pi(\theta, a)\Big[w(x,\theta)g(x|a) + \sum_i u_i(x,\theta)\nabla_i^+ g(x|a)\Big]$$

- Things are starting to look similar! The terms involving ∇_i operators are essentially discrete derivatives
- But there is still an asymmetry in how we handle participation/order of moves/local up versus local down
- The order of moves in particular doesn't matter:
  For fixed Δ, this is a finite zero-sum game, so we could exchange the order of moves without changing the value

Revisiting participation security

- There is still the question of which action(s) should be participation secure
- To make things similar to the min potential upper bound, where IR only binds at the bottom, it's natural to make the lowest action participation secure, i.e., in the max guarantee lower bound, we impose for all i, a_{-i}, θ,

$$\sum_x u_i(0, a_{-i}, \theta) g(x|a) \geq 0$$

- Summing up: We have found particular bounds on the max guarantee/min potential that don't seem so dissimilar, especially in the limit as $\Delta \to 0$

Simplifying a bit further

- Since the marginal on θ is fixed, we can simplify the lower bound program to:

$$\max_{g \text{ p.s.}} \sum_{\theta} \mu(\theta) \min_{a} \sum_{x} \left[w(x,\theta) g(x|a) + \sum_{i} u_i(x,\theta) \nabla_i^+ g(x|a) \right]$$

$$\nabla_i^+ g(x|a) = \mathbb{I}_{a_i < \frac{1}{\Delta}} \frac{1}{\Delta} \left(g(x|a_i + \Delta, a_{-i}) - g(x|a) \right)$$

- This has the form of choosing the mechanism to maximize the expected minimum **strategic virtual objective**
- NB it's a linear program: just introduce variables $\lambda(\theta)$, and maximize $\sum_{\theta} \mu(\theta)\lambda(\theta)$, subject to, for all θ, a,

$$\lambda(\theta) \leq \sum_{x} \left[w(x,\theta) g(x|a) + \sum_{i} \mathbb{I}_{a_i < \frac{1}{\Delta}} u_i(x,\theta) \nabla_i^+ g(x|a) \right]$$

Simplifying a bit further

- For the potential, we can first sum by parts to get:

$$\min_{\pi} \max_{g} \sum_{t,\theta,x} g(x|t) \left[w(x,\theta)\pi(\theta,t) - \sum_i u_i(x,\theta)\tilde{\nabla}_i^+ \pi(\theta,t) \right]$$

$$\tilde{\nabla}_i^+ \pi(\theta,t) = \begin{cases} -\pi(\theta,t) & \text{if } t_i = \frac{1}{\Delta}; \\ \pi(\theta, t_i + \Delta, t_{-i}) - \frac{1}{\Delta}\pi(\theta,t) & \text{if } t_i = \frac{1}{\Delta} - \Delta; \\ \frac{1}{\Delta}[\pi(\theta, t_i + \Delta, t_{-i}) - \pi(\theta,t)] & \text{otherwise.} \end{cases}$$

- And then finally simplify to:

$$\min_{\pi} \sum_{t} \max_{x} \sum_{\theta} \left[w(x,\theta)\pi(\theta,t) - \sum_i u_i(x,\theta)\tilde{\nabla}_i^+ \pi(\theta,t) \right]$$

  i.e., minimizing an expected highest **informational virtual objective**

- Similar as before, this is an LP, and we could introduce auxiliary variables $\gamma(t)$, and minimize $\sum_t \gamma(t)$ subject to for all t, x,

$$\gamma(t) \geq \sum_{\theta} \left[w(x,\theta)\pi(\theta,t) - \sum_i u_i(x,\theta)\tilde{\nabla}_i^+ \pi(\theta,t) \right]$$

"Almost" a dual pair

- We could have switched the order of moves, in which case the min potential upper bound would reduce to:

$$\max_g \sum_\theta \mu(\theta) \min_t \sum_x \left[w(x,\theta) g(x|t) + \sum_i u_i(x,\theta) \nabla_i^- g(x|t) \right]$$

- This also looks like maximizing an expected minimum SVO, but with a slightly different definition, where the type is the action!
- Moreover, it's **exactly** the dual to the min potential upper bound LP from the previous slide
- In particular, $\pi(\theta, t)$ is the Lagrange multiplier on

$$\lambda(\theta) \leq \sum_x \left[w(x,\theta) g(x|t) + \sum_i u_i(x,\theta) \nabla_i^- g(x|t) \right]$$

- But also, $g(x|t)$ is the multiplier on

$$\gamma(t) \geq \sum_\theta \left[w(x,\theta) \pi(\theta,t) - \sum_i u_i(x,\theta) \tilde{\nabla}_i^+ \pi(\theta,t) \right]$$

Complementary slackness

- Everything that we know about duality in LP applies: In particular, saddle points for the min potential bound and its dual would just be pairs of mechanisms/type spaces that satisfy complementary slackness, i.e.,
  - The type space puts probability one on types/actions that minimize the modified SVO
  - The mechanim puts probability one on types/actions that maximize the informational virtual objective
- When we pass to the continuum limit, if the difference between ∇_i^- and ∇_i^+ washes out, then we might expect the min potenial and max guarantee bounds to be dual to one another,
- In that case, saddle points consisting of min potential type spaces and max guarantee mechanisms would satisfy the same notion of complementary slackness
- This is a useful heuristic for engineering saddle points...

Applying the bounds

- Brooks and Du (2021) applied these bounds to the problem of revenue maximization in auctions when there is a pure common value $v \in V$ (like in our analysis of the FPA):
- This is slightly outside of the setup I just gave you, since transfers are unbounded (and this is kind of important)
- But we can formulate analogous bounding programs
- We will look at the simpler special case of this problem where the good must be allocated
- The mechanism then consists of an allocations $q_i(a) \geq 0$ subject to $\Sigma q(a) = 1$ and payments $p_i(a)$ (we previously called then t but now we need to distinguish from types)
- The designer's payoff is $\Sigma p(a)$ and agent i's payoff is $vq_i(a) - p_i(a)$

Bounds for revenue maximization in auctions

$$\min_{\pi} \sum_{t} \max_{(q,p)} \sum_{v} \left[\left(\sum_{i} p_i \right) \pi(v,t) - \sum_{i} \underbrace{[vq_i - p_i]\, \tilde{\nabla}_i^{+} \pi(v,t)}_{=\sum_x u_i(x,\theta) \tilde{\nabla}_i^{+} g(x|t)} \right]$$

$$\max_{q,p \text{ p.s.}} \sum_{v} \mu(v) \min_{a} \left[\sum_{i} p_i(a) + \sum_{i} \underbrace{(v \nabla_i^{+} q_i(a) - \nabla_i^{+} p_i(a))}_{=\sum_x u_i(x,\theta) \nabla_i^{+} g(x|a)} \right]$$

- These two LPs bound the potential and guarantee
- NB Participation security just means that $vq_i(0, a_{-i}) \geq p_i(0, a_{-i})$
- Also, note that π is a distribution over values and types in $\{0, \Delta, \ldots, 1/\Delta\}$

Key facts from simulations

- On your final problem set, you will solve them for the case where the value is uniform on $[0, 1]$
- You will find:
  1. The programs have approximately the same value when $\Delta \approx 0$
  2. The optimal π has independent signals, but the density is also just a function of the sum of the signals!
  3. $w(t) \equiv \mathbb{E}[v|t]$ is non-decreasing and only a function of the sum of the signals
  4. The optimal q has a **proportional** form:

$$q_i(a) = \frac{a_i}{\Sigma a}$$

  5. Transfers don't (appear) to have a lot of structure... multiplicity?
- Let's use these clues to analyze and solve the bounding programs

Min potential type space

- For the min potential, since p_i is free, clearly it must drop out of the objective, so we need to have

$$\sum_v \pi(v,t) = -\sum_v \tilde{\nabla}_i^+ \pi(v,t)$$

  (Here is where it was important how we define ∇_i^- at the top, so that it is possible to satisfy this equation when $t_i = 1/\Delta$, where it reduces to $\sum_v \pi(v,t) = \sum_v \pi(v,t)$)
- In fact, going to the continuum limit, where) types are in $\mathbb{R}_+$, and writing $f(t)$ for the marginal density of the signals, we would have

$$\begin{aligned} f(t) &= -\frac{\partial}{\partial t_i} f(t) \\ \implies f(t) &= \exp(-t_i) f(0, t_{-i}) \\ \implies f(t) &= \prod_i \exp(-t_i) = \exp(-\Sigma t), \end{aligned}$$

  so the signals are iid exponential with unit arrival rate!
- (Transfer neutrality again!)

Remaining terms in the min potential

- Note that if we write write $w(t) = \mathbb{E}[v|t]$, then

$$\sum_{v,t} v \frac{\partial}{\partial t_i} \pi(v,t) = \frac{\partial}{\partial t_i} \sum_{v,t} v\pi(v,t) = \frac{\partial}{\partial t_i}(w(t)\exp(-\Sigma t))$$

- Plugging this and the exponential distribution into the rest of the objective, we get

$$\begin{aligned} & -\int_t \sum_i q_i(t) \frac{\partial}{\partial t_i}(w(t)\exp(-\Sigma t))dt \\ & = \int_t \sum_i q_i(t) \left(w(t) - \frac{\partial}{\partial t_i} w(t) \right) \exp(-\Sigma t) dt \end{aligned}$$

- So, why should we expect $w(t)$ to be a function only of Σt? Well, this equalizes the "virtual value" $w - w'$ across bidders!

Optimal value function

- But what is the exact value function?
- Well, the expectation of w is fixed; this is just the expected value
- So we should pick w to maximize the expectation of w'
- This is achieved by a **fully revealing value function** that matches Σt and v comonotonically
- Let G_N be the distribution of the sum of N iid standard exponential distributions
- Then if F is the CDF of the value, we define $\hat{w}$ according to

$$F(\hat{w}(x)) = G_N(x)$$

- Then the critical value function is $w(t) = \hat{w}(\Sigma t)$

The optimal potential

- The potential is

$$\int_t (\hat{w}(\Sigma t) - \hat{w}'(\Sigma t)) \exp(-\Sigma t) dt = \int_{x=0}^{\infty} (\hat{w}(x) - \hat{w}'(x)) g_N(x) dx$$

where g_N is the PDF of the Erlang distribution:

$$g_N(x) = \frac{x^{N-1}}{(N-1)!} \exp(-x)$$

- Not hard to show that the potential is also

$$\int_{x=0}^{\infty} \hat{w}(x) g_{N-1}(x) dx$$

What's going on here?

- We might have guessed that the worst case would have to have independent signals, to avoid Myerson-Crémer-McLean style bets that just give extra rents to the designer
- Given independent signals, the marginals are just a normalization, so we might as well make them iid exponential
- In this setting, generalization of the "local" virtual value (e.g., Bulow and Klemperer 1996) is

$$w_i(t) - \frac{1 - F_i(t_i)}{f_i(t_i)} \frac{\partial w_i(t)}{\partial t_i}$$

  where $w_i(t) = \mathbb{E}[v_i|t]$
- (We could have derived this using the same envelope theorem-logic and it is intuitive that the sensitivity term should appear at the end)
- With the standard exponential normalization, the inverse hazard rate is just $(1 - (1 - \exp(-t_i)))/\exp(-t_i) = 1$, so it drops out
- In fact, this is related to our choice of Lagrange multipliers to be constant (since we know the multiplier is related to the inverse hazard rate term)
- And with common values, $w_i(t) = w(t) = \mathbb{E}[v|t]$

Max guarantee mechanism

- On the mechanism side, in the continuum limit, the strategic virtual objective reduces to

$$\Sigma p(a) - \nabla \cdot p(a) + v\nabla \cdot q(a),$$

where $\nabla\cdot$ is the divergence

- Now, what is special about the proportional form?

$$\nabla \cdot q(a) = \sum_i \frac{\partial}{\partial a_i}\left(\frac{a_i}{\Sigma a}\right) = \frac{N-1}{\Sigma a}$$

- What's important here? Independent of a, conditional on Σa! So it equalizes the strategic virtual objective across lots of action profiles

What about transfers?

- Recall the constraint in the max guarantee program, which in the continuum limit with a continuous allocation/payment becomes:

$$\lambda(v) \leq \Sigma p(a) + v\nabla \cdot q(a) - \nabla \cdot p(a) \; \forall a, v$$

  where $\nabla\cdot$ is the divergence
- Of course, since we are maximizing λ, the optimal multipliers should make this constraint hold with equality:

$$\lambda(v) = \min_a \left[\Sigma p(a) + v\nabla \cdot q(a) - \nabla \cdot p(a)\right]$$

- But what action profile should be the minimizer? $v = \hat{w}(\Sigma a)$!
- Why? This is what makes Nature willing to put probability mass on (v, a) (complementary slackness)

The optimal λ

- λ is a concave and non-decreasing, so differentiable a.e.
- In fact, the envelope theorem implies that

$$\lambda'(\hat{w}(\Sigma a)) = \nabla \cdot q(a) = \frac{N-1}{\Sigma a}$$

- This differential equation determines λ up to a constant
- But what constant? We will come back to that in a few slides
- Denote the optimal λ by $\hat{\lambda}$

The aggregate excess growth equation

- But since all a are in the support, evidently, we must have

$$\begin{aligned}\nabla \cdot p(a) - \Sigma p(a) &= \hat{w}(\Sigma a)\nabla \cdot q(a) - \hat{\lambda}(\hat{w}(\Sigma a)) \\ &= \hat{w}(\Sigma a)\frac{N-1}{\Sigma a} - \hat{\lambda}(\hat{w}(\Sigma a))\end{aligned}$$

- Brooks and Du (2021) call the LHS of this equation the **aggregate excess growth**
- Evidently, any transfer rule that satisfies the AEG equation and $p_i(0, a_{-i}) = 0$ will work (to satisfy participation security)
- In general, there are many solutions to this equation, some of which are quite messy

Proportional transfers

- But a bold guess is that there is also a transfer rule of the proportional form:

$$p_i(a) = \frac{a_i}{\Sigma a} P(\Sigma a),$$

  where $P \equiv \sum_i p_i$ is the aggregate payment
- The transfer equation reduces to

$$P'(x) + \left(\frac{N-1}{x} - 1 \right) P(x) = \hat{w}(x) \frac{N-1}{x} - \hat{\lambda}(\hat{w}(x))$$

- This ODE is easily solved, subject to the boundary condition $P(0) = 0$, to obtain the transfer rule, and in fact

$$P(x) = \frac{1}{g_N(x)} \int_{y=0}^{x} \left[\hat{w}(y) \frac{N-1}{y} - \hat{\lambda}(\hat{w}(y)) \right] g_N(y) dy$$

  where $g_N(x)$ is the PDF of the Erlang distribution

The revenue guarantee

- NB the integral must converge to zero as $x \to \infty$... this pins down the constant in $\hat{\lambda}$
- The resulting guarantee is therefore

$$
\begin{aligned}
0 &= \int_{y=0}^{\infty} \left[\hat{w}(y) \frac{N-1}{y} - \hat{\lambda}(\hat{w}(y)) \right] g_N(y) dy \\
\iff \int_{y=0}^{\infty} \hat{\lambda}(\hat{w}(y)) g_N(y) dy \\
&= \int_{y=0}^{\infty} \hat{w}(y) \frac{N-1}{y} g_N(y) dy \\
&= \int_{y=0}^{\infty} \hat{w}(y) g_{N-1}(y) dy,
\end{aligned}
$$

as desired

What does it all mean?

- If we look for robustness to beliefs, but allow the implemented outcome to vary with beliefs, then we get a very different kind of "robust" mechanism
- Robustness in this sense means controlling the strategic virtual objective:

$$\sum_x \left[w(x,\theta)g(x|a) + \sum_i u_i(x,\theta)\nabla_i g(x|a) \right]$$

- Robust mechanisms maximize the expected minimum SVO
- The proportional auction in particular does so by equalizing the SVO across all actions with the same sum
- Analogously, worst case information minimizes the expected highest informational virtual objective, which in the context of auctions means independent signals and minimizing an expected highest local virtual value, in the sense of Myerson (1981) and Bulow and Klemperer (1996)

Revenue guarantee when N is large

- Recall that the optimal guarantee is

$$\int_{x=0}^{\infty} \hat{w}(x) g_{N-1}(x) dx.$$

- We can change units to $y = (x - N)/\sqrt{N}$, with corresponding distribution $G_N^c(y) = G_N(\sqrt{N}y + N)$
- In these "central limit" units, $\hat{w}^c(y) = F^{-1}(G_N^c(y))$
- Then G_N^c converges to a standard normal Φ, $\hat{w}_N^c$ converges to $F^{-1}(\Phi(y))$, and maxmin revenue converges to

$$\int_{y=-\infty}^{\infty} F^{-1}(\Phi(x))\phi(x) dx = \mathbb{E}_F[v]$$

- So, full surplus extraction in the limit! NB contrast with what we found for the FPA
- Generalizes and refines a result of Du (2018)
- In fact, can even show that $\hat{\lambda}(v) \to v$ pointwise, so that expected revenue converges to $\mathbb{E}_F[v]$ even if the prior is misspecified

Can-keep case

- What about when the seller can keep the good?
- Well, the problem with our construction is that in general, $\hat{w}(x) - \hat{w}'(x)$ (the virtual value when the sum of signals is x) can be negative, in which case the seller would prefer to keep the good
- For example, this always happens when $\underline{v} = 0$, or it will happen whenever there are gaps in the support of F
- As a result, optimal revenue on the type space just constructed among all mechanisms is greater than the optimal guarantee if the good must be sold
- If the fully revealing virtual value is negative, it is necessary to "grade" the value function so that it's not too steep, and the virtual value stays non-negative
- Equivalently, we grade the value function so that it always grows slower than exponential

Grading

- We define

$$\hat{W}(x) = \int_{y=0}^{x} \hat{w}(y) g_N(y) dy$$
$$E(x) = \int_{y=0}^{x} \exp(y) g_N(y) dy$$

- $\hat{w}$ is sub-exponential if and only if $\hat{W} \circ E^{-1}$ is sub-linear
- Let $\mathrm{cav}(\hat{W} \circ E^{-1})$ denote the smallest concave function that is everywhere above $\hat{W} \circ E^{-1}$, and define

$$\overline{W} = \mathrm{cav}(\hat{W} \circ E^{-1}) \circ E$$

- The graded value function is

$$\overline{w}(x) = \frac{1}{g_N(x)} \frac{d}{dx} \overline{W}(x)$$

What is grading?

- In effect, we change units so that linear growth corresponds to exponential growth in the original units
- By taking the concavification in this space, we make sure that growth is sub-exponential, but we preserve the expectation of the value over any graded interval
- In particular, in a graded interval $[a, b]$ where $\bar{W}(x) > \hat{W}(x)$, we have $\overline{w}(x) = k \exp(x)$ for some constant k, but

$$\int_{x=a}^{b} \hat{w}(x) g_N(x) dx = \int_{x=a}^{b} \overline{w}(x) g_N(x) dx$$

- In non-graded intervals, $\overline{w}(x) = \hat{w}(x)$
- In the uniform example, there is a single graded interval around zero

The can-keep potential minimizing type space

- Again, we have iid exponential signals but now the value function is $\overline{w}$
- The seller is now always willing to allocate the good, so optimal revenue is just

$$\int_{x=0}^{\infty} \overline{w}(x) g_{n-1}(x) dx$$

- This also turns out to be maxmin revenue

The can-keep mechanism

- The guarantee maximizing mechanism retains the proportional form, but now there is an aggregate supply that depends on Σa:

$$q_i(a) = \frac{a_i}{\Sigma a} Q(\Sigma a)$$

- On non-graded intervals, the divergence of q is $(N-1)/x$, as before, and $Q(x) = 1$
- On graded intervals, the divergence turns out to be constant, which leads to rationing
- The aggregate allocation must be continuous and solve

$$\frac{N-1}{x} Q(x) + Q'(x) = C$$

- If there is a graded interval $[0, b]$ at the bottom, then $Q(0) = 0$ and $Q(b) = 1$, so $Q(x) = x/b$
- Otherwise, there is a more complicated polynomial solution to the boundary conditions $Q(a) = Q(b) = 1$
- The construction of the transfers is the same as before, but with a different multiplier $\overline{\lambda}$ so that the SVO is still minimized at the fully revealing value

Will these bounds coincide in general?

- No. A key issue is that participation security, while it works well in certain settings, may be too demanding in general
- For example, consider the following public goods problem:
  - Society wants to produce a bridge that costs c
  - Each member of society values the bridge at $c > v > c/N$
  - All funds must be raised from society (budget balance)
  - The designer's goal is to maximize social welfare
- There is complete information, so information plays no role: Min potential is just $Nv - c$, since sharing the cost equally would satisfy participation constraints
- But in a participation secure mechanism, each agent has the ability to "opt out" and not pay anything, and if others are opting out, then an agent who opts in and pays for the bridge would receive a payoff $v - c < 0$
- As such, any participation secure mechanism will have an opt out equilibrium, in which case welfare is zero
- As such, the duality gap in this problem is $Nv - c$
- Nonetheless, there are other specifications of the public goods problem where the bounds do coincide...

Open questions

- So, an important open question is under what conditions do these bounds coincide and are non-trivial
- And when they do coincide, what is the form of saddle points
- In such cases, the solutions to the max guarantee problem fulfill a lot of desiderata:
  - Bounds on welfare for all type spaces/equilibria
  - Actions are not dependent on a particular choice of interim types or ex post types
  - Also, the mechanisms allow the outcome to vary with beliefs, and escape the ex post implementation trap
- Nonetheless, this model demands an extreme form of robustness (all type spaces and all equilibria!)
- Would be nice to be able to restrict the set of type spaces...

Restricting mechanisms?

- Alternative approach: restrict the class of mechanisms that we allow the designer to consider
- Perhaps we do not like the randomization inherent in the proportional auction... What if we look at mechanisms that have a more "conventional" auction structure, in which a high bidder wins?
- This question was investigated by BBM 2019, who find that the FPA maximizes the guarantee among a class of auctions of this form

First-price auction recap

- First-price auction, denoted $\mathcal{M}^{FPA}$:
  - $A_i = \mathbb{R}$
  - q allocates to a high bidder (breaking ties randomly)
  - $p_i(a) = q_i(a)a_i$
- Recall our analysis of revenue-minimizing BCE of the FPA
- Worst-case type space $\mathcal{T}^*$:
  - Types $T_i = [\underline{v}, \overline{v}]$
  - $t_i \in T_i$ are iid draws from $F^{1/n}$
  - Value is equal to $\max_i t_i$
- We showed there is an equilibrium of the FPA in $\mathcal{T}^*$ in which a type t_i bids

$$\beta^{FPA}(t_i) = \frac{1}{(F(t_i))^{\frac{n-1}{n}}} \int_{x=\underline{v}}^{t_i} x d(F(x))^{\frac{n-1}{n}}$$

- Minimum revenue is precisely revenue in this equilibrium:

$$R^{FPA} = \int_{v=\underline{v}}^{\overline{v}} \beta^{FPA}(v) dF(v)$$

Other mechanisms in the worst-case type space

- Recall that $\beta(v_i)$ is the equilibrium strategy in the FPA in the independent private value model where $v_i \sim F^{1/n}$
- We denote this type space by $\mathcal{T}^{IPV}$
- This connection played a role in our proof that these strategies are an equilibrium in $\mathcal{T}^*$
- In fact, we can use this connection, and the revenue equivalence theorem, to show that lots of other mechanisms have an equilibrium on $\mathcal{T}^*$ that induces a revenue of R^{FPA}

Standard mechanisms

- We say that a mechanism $\mathcal{M} = (A, q, p)$ is **standard** if
  - $A_i = \mathbb{R}$
  - A high bidder is allocated the good
  - There is a symmetric and monotonic pure strategy equilibrium in symmetric IPV type spaces in which equilibrium bidder surplus is non-negative
- This is in a sense an endogenous condition on the mechanism, but it boils down to just choosing a different ex post implementation of the interim allocation for the efficient allocation rule
- First-price, second-price, all-pay, and combinations thereof are standard auctions

Strategic equivalence

Proposition (Bergemann, Brooks, and Morris (2019))

Suppose that $\mathcal{M}$ is a standard mechanism and β is a symmetric and monotonic pure-strategy equilibrium of $(\mathcal{M}, \mathcal{T}^{IPV})$. Then β is also an equilibrium of $(\mathcal{M}, \mathcal{T}^*)$.

- Proof: The allocation induced by β is precisely that induced by the monotonic pure-strategy equilibrium of the FPA
- Revenue equivalence then implies that the interim expected transfer is $P_i(t_i) = P^{FPA}(t_i) + c_i$, where T^{FPA} is the interim transfer in the FPA in the equilibrium β^{FPA}, and c_i is a constant

Proof, continued

- Thus, if we let $U(t_i, t_i')$ denote the payoff from type t_i bidding $\beta(t_i')$ in $(\mathcal{M}, \mathcal{T}^*)$, and U^{FPA} denotes the corresponding payoff in $(\mathcal{M}^{FPA}, \mathcal{T}^*)$, then we have

$$U_i(t_i, t_i') = U_i^{FPA}(t_i, t_i') - c_i$$

- But $U_i^{FPA}(t_i, t_i) \geq U_i^{FPA}(t_i, t_i')$ for any t_i', so $U_i(t_i, t_i) \geq U_i(t_i, t_i')$ as well
- Finally, just have to check that bidders don't want to deviate to a report b_i that is not in the range of β
- But for any such bid, there must be an equilibrium bid with the same winning probability and a weakly lower transfer, so such deviations cannot be attractive $\square$

Robust optimality of the FPA

- This result immediately yields the following characterization:

Theorem

If $\mathcal{M}$ is standard, then

$$R^{FPA} \geq \inf_{\sigma \in BCE(\mathcal{M})} \int_{M \times [\underline{v}, \overline{v}]} \sum_{i=1}^{N} t_i(m) \sigma(dm, dv).$$

Thus, the first-price auction maximizes revenue among standard mechanisms the minimum revenue in BCE.

Proof of Theorem

- For any standard mechanism $\mathcal{M}$, there is an equilibrium β on $(\mathcal{M}, \mathcal{T}^{IPV})$, which by the proposition is also an equilibrium on $(\mathcal{M}, \mathcal{T}^{*})$
- In each of these games, the strategies β induce revenue R
- But R^{FPA} is revenue from the equilibrium β^{FPA} of $(\mathcal{M}^{FPA}, \mathcal{T}^{IPV})$, which is maximum revenue among all efficient mechanisms and equilibria on $\mathcal{T}^{IPV}$, subject to bidder utilities being non-negative
- Hence, $R \leq R^{FPA}$
- Since we have a type space and equilibrium of $\mathcal{M}$ in which revenue is less than R^{FPA}, we conclue that infimum revenue across BCE is also less than R^{FPA} $\square$

Going further

- The theorem is a kind of robust revenue ranking: the FPA has a higher "revenue guarantee" than any other standard auction
- An important standard mechanism is the second-price auction
- Of course, we didn't need the theorem to tell us that minimum revenue across BCE in the SPA is less than R^{FPA}
- Indeed, in any type space, the SPA has equilibria with zero revenue, in which one bidder always bids $\overline{v}$ and all others bid zero
- What is perhaps a bit more surprising is that minor perturbations of the mechanism that might kill off that equilbrium (e.g., placing a small probability on pay-as-bid) cannot lead to greater minimum revenue than R^{FPA}

Affiliated type spaces

- This result is in stark contrast to a finding of Milgrom and Weber (1982):
- A common value type space is **affiliated** if
  - $T_i \subseteq \mathbb{R}$, $T = \times_{i \in N} T_i$
  - Bidder i's interim value $w_i(t)$ is non-decreasing in $t \in T$
  - The distribution of t is affiliated; if there is a density f, this is equivalent to $f(t)$ being log-supermodular
- Independence is a special case of affiliation
- For symmetric and affiliated type spaces, MW '82 construct symmetric monotonic pure-strategy equilibria of the FPA, SPA, and English auction
- These equilibria are "natural" extensions of the Vickrey equilibria to correlated and interdependent value type spaces

Linkage principle

- MW '82 show that revenue in the English auction is greater than that of the SPA, which is greater than that of the FPA
- Here is some intuition for the SPA/FPA comparison:
  - In the FPA, the winner's payment just depends on their own bid; so if a type t_i bids as $t_i' < t_i$, they expect to pay $\beta^{FPA}(t_i')$ conditional on winning, the same amount as t_i' would expect to pay conditional on bidding t_i' and winning
  - In the SPA, a type t_i expects to pay more than t_i' expects to pay conditional on winning with a bid of $\beta^{SPA}(t')$, because with strict affiliation, the conditional distribution of t_{-i} is "higher" from the perspective of t_i than t_i'
- More broadly, "linking" the payment to variables that are jointly affiliated with signals leads to higher revenue
- In a sense, this is a limited form of "separation with beliefs" a la Crémer and McLean, although if one could design any mechanism, then full surplus extraction is possible when there is strict affiliation (McAfee, McMillan, and Reny 1989)

Generalized revenue equivalence

- Importantly, this result depends on there being correlation
- When signals are independent, revenue equivalence applies
- In particular, in any direct mechanism, the indirect utility is

$$U_i(t_i) = \max_{t_i'} \mathbb{E}_{t_{-i}} \left[w_i(t_i, t_{-i}) q_i(t_i', t_{-i}) - t_i(t_i', t_{-i}) \right]$$

- If w, q, and t are sufficiently well behaved, then the envelope theorem implies that

$$U_i'(t_i) = \mathbb{E}_{t_{-i}} \left[\frac{\partial}{\partial t_i} w_i(t_i, t_{-i}) q_i(t_i, t_{-i}) \right],$$

  which reduces to the interim allocation in the special case where $w_i(s) = t_i$

Generalized revenue equivalence, cont'd

- This leads to a generalized revenue equivalence formula, where the virtual value is

$$\phi_i(t) = w_i(s) - \frac{\partial}{\partial t_i} w_i(t) \frac{1 - F_i(t_i)}{f_i(t_i)}$$

- The size of the information rent depends on how informative is bidder i's signal
- Expected revenue is the expected virtual value of the bidder who is allocated the good
- This is another way to see why revenue equivalence holds on the type space $\mathcal{T}^*$
- Indeed, the English auction also has an equilibrium on $\mathcal{T}^*$ in which expected revenue is R^{FPA}
- Bulow and Klemperer (1996) make this point, and also give conditions with interdependent values and independent signals under which a generalization of the English auction maximizes expected revenue

Revenue guarantee equivalence

- Thus, one way to rescue the SPA and English auctions from the "bad" equilibria is to focus on symmetric affiliated type spaces and the MW equilibria
- Within this class of type spaces and equilibria, SPA and English auction weakly outperform the FPA
- But because $\mathcal{T}^*$ is affiliated, we know that SPA and English auction cannot be strictly better than the FPA in the worst case
- This is reported as Theorem 2 in Bergemann, Brooks, and Morris (2019)

A saddle point

- Returning to general type spaces, the FPA achieves maxmin revenue (with the min over BCE) when we restrict attention to standard mechanisms
- Moreover, $(\mathcal{M}^{FPA}, \mathcal{T}^*)$ are a **saddle point** in the following sense:
  - For the mechanism $\mathcal{M}^{FPA}$, revenue is at least R^{FPA} in all type spaces and equilibria
  - On the type space $\mathcal{T}^*$, no standard mechanism can achieve more revenue than R^{FPA} in all equilibria

Posted prices on $\mathcal{T}^*$

- But as we showed, not a saddle point for the unrestricted problem!
- At $\mathcal{T}^*$, the seller could just offer to sell the good at a posted price of

$$p = \int_{x=\underline{v}}^{\overline{v}} x d(F(x)^{(n-1)/n}),$$

i.e., the expected highest value conditional on $t_i = \underline{v}$

- All bidders would want to purchase at this price
- Moreover, the expected highest of $n-1$ draws is always greater than the expected second-highest of n draws (because sometimes the former would be highest among n draws)
- So $p \geq R^{FPA}$
- BBM (2020) solve for the optimal auction on $\mathcal{T}^*$, which is an interesting subject, but not directly related to our discussion of robust mechanisms

Final thoughts

- The analysis of "guarantees" and "potentials" provide a different take on robust mechanism design
- Pros: The model gives us new ideas for mechanisms (proportional auctions), or new insights about the "robustness" of standard mechanisms (FPA)
- Cons: Still relies heavily on the common priors, BNE, no restrictions on information
- Mechanism design still needs a more flexible model of "robust" design, that can incorporate some restrictions on beliefs, while still not assuming a particular type space, and possibly also relaxing the CPA