Tutorial on Robust Auction Design
Lecture 3

Instructors: Ben Brooks and Songzi Du

Slides @ https://benjaminbrooks.net/
https://econweb.ucsd.edu/~sodu/

EC 2021

RBA with common priors and interdependent values

- Two crucial assumptions were used in the previous lecture: non-common priors and private values
- Each of these is controversial:
  - Arbitrary departures from common priors seem implausible
  - Why should bidders know exactly their own value?
- For the next two lectures, we will impose the common prior assumption, but allow **interdependent values**
- In fact, we will often go to an opposite extreme, and assume that the bidders have **pure common values**, i.e., they have exactly the same value for the good (although many results generalize)
- In addition, whereas we previously assumed that the seller picks the equilibrium, we will instead look for robustness with respect to equilibrium selection... this limits our use of the revelation principle!
- (NB: There's no *a priori* reason why these have to go together, but they lead to a very tractable and fruitful theory)

Robust predictions

- So, the new high-level goal is to characterize auctions that perform well, regardless of the bidders' common-prior beliefs and which equilibrium is played
- Preliminary step: develop a tool for characterizing, for a fixed mechanism, the set of possible outcomes that could arise, consistent with common priors and Bayes Nash equilibrium
- This is **Bayes correlated equilibrium** (BCE) (Bergemann and Morris, 2013, 2016)

Basic environment

- Payoff relevant state $\theta \in \Theta$
- Players $i = 1, \ldots, N$
- A **game form** $\mathcal{G} = (A, u)$:
  - Actions A_i
  - Payoff functions $u_i : A \times \Theta \to \mathbb{R}$
- We can write a common-prior information structures as $\mathcal{I} = (S, \pi)$, where $\pi \in \Delta(S \times \Theta)$

Bayes Nash equilibrium

- Together, $(\mathcal{G}, \mathcal{I})$ constitute a Bayesian game
- Player i's **strategies** are mappings $b_i : S_i \to \Delta(A_i)$
- Under the tuple of strategies $b = (b_1, \ldots, b_N)$, player i's expected payoff is

$$U_i(b) = \sum_{\theta, s, a} \pi(s, \theta) b(a|s) u_i(a, \theta)$$

- A **Bayes Nash equilibrium** is a tuple of strategies b such that for all i and strategies b_i',

$$U_i(b) \geq U_i(b_i', b_{-i})$$

- High-level question: What are the outcomes of $\mathcal{G}$ that are consistent with Bayes Nash equilibrium for some $\mathcal{I}$?

Outcomes and BCE

- An **outcome** of $\mathcal{G}$ is a $\sigma \in \Delta(\Theta \times A)$
- Any $\mathcal{I}$ and BNE b of $(\mathcal{G}, \mathcal{I})$ **induce** the outcome:

$$\sigma(\theta, a) = \sum_{s \in S} \pi(s, \theta) b(a|s)$$

- A **Bayes correlated equilibrium** of $\mathcal{G} = (A, u)$ is an $\sigma \in \Delta(\Theta \times A)$ that satisfies the following **obedience constraints** for all i and a_i,

$$\sum_{\theta, a_{-i}} \sigma(\theta, (a_i, a_{-i}))(u_i(a_i, a_{-i}, \theta) - u_i(a_i', a_{-i}, \theta)) \geq 0$$

- In other words, a BCE is just an information structure of the form (A, σ) for which the identity mapping is a BNE

Epistemic characterization

Theorem

An outcome σ is a BCE of $\mathcal{G}$ iff σ is induced by some information structure $\mathcal{I}$ and a BNE b of $(\mathcal{G}, \mathcal{I})$.

- So, BCE are the outcomes that are consistent with rationality and common priors

Proof

- If: induced by some $(\mathcal{I}, b) \implies$ BCE
- We simply verify the obedience constraints
- If not, there is a profitable deviation for some player i of the form: play a_i' whenever you would have played a_i

Proof

- If: induced by some $(\mathcal{I}, b) \implies$ BCE
- We simply verify the obedience constraints
- If not, there is a profitable deviation for some player i of the form: play a_i' whenever you would have played a_i
- Only if: BCE $\implies$ induced by some $(\mathcal{I}, b)$
- Define $S_i = A_i$ and $\pi = \sigma$
- The strategies are the identity mapping, i.e.,

$$b_i(a_i|a_i) = 1$$

- IC follows from the obedience constraint; if they have a profitable deviation in the proposed information/equilibrium, then an obedience constraint is violated

Comments on BCE

- Relative to Lecture 2, we now impose the common prior
- But still much more general than the classical fixed-information models, reviewed in Lecture 1
- In addition, BCE allows for an arbitrary equilibrium to be played, not just a particular "designer preferred" equilibrium

Comments on BCE

- Relative to Lecture 2, we now impose the common prior
- But still much more general than the classical fixed-information models, reviewed in Lecture 1
- In addition, BCE allows for an arbitrary equilibrium to be played, not just a particular "designer preferred" equilibrium
- Analytically very tractable: BCE are the intersection of a family of linear incentive constraints

Comments on BCE

- Relative to Lecture 2, we now impose the common prior
- But still much more general than the classical fixed-information models, reviewed in Lecture 1
- In addition, BCE allows for an arbitrary equilibrium to be played, not just a particular "designer preferred" equilibrium
- Analytically very tractable: BCE are the intersection of a family of linear incentive constraints
- Often useful to further discipline the set of BCE by restricting the marginal of σ on Θ
- That way, differences in outcomes across BCE are just due to differences in **beliefs**, rather than differences in **fundamentals**

Application: Common-value first-price auctions

- We apply the BCE methodology to first-price auctions, following Bergemann, Brooks, and Morris (2017)
- N bidders
- **Pure common value** $v \sim F$ with support $[\underline{v}, \overline{v}]$, strictly positive density $f(v)$

Application: Common-value first-price auctions

- We apply the BCE methodology to first-price auctions, following Bergemann, Brooks, and Morris (2017)
- N bidders
- **Pure common value** $v \sim F$ with support $[\underline{v}, \overline{v}]$, strictly positive density $f(v)$
- The bidders compete in a first-price auction:
  - Bids $a_i \in \mathbb{R}_+$
  - High bidder wins and pays their bid

Application: Common-value first-price auctions

- We apply the BCE methodology to first-price auctions, following Bergemann, Brooks, and Morris (2017)
- N bidders
- **Pure common value** $v \sim F$ with support $[\underline{v}, \overline{v}]$, strictly positive density $f(v)$
- The bidders compete in a first-price auction:
  - Bids $a_i \in \mathbb{R}_+$
  - High bidder wins and pays their bid
- No reserve price $\implies$ total surplus is always

$$\hat{v} = \int_{v=0}^{1} vf(v)dv$$

- The split between seller and bidders depends on information
- One possibility: All information is public, so that bidders compete away their rents
- What is minimum revenue across all BCE with the prior f?

Winning bid distributions

- Turns out that we only need to look at certain marginals/conditionals of the BCE
- $H(a|v)$ is the CDF of the winning bid conditional v
- $H(a) = \int_{v=\underline{v}}^{\overline{v}} H(a|v) f(v) dv$ is the unconditional CDF
- $H_i(a|v)$ is the prob that i wins **and** the winning bid is less than a, so that

$$H(a|v) = \sum_{i=1}^{n} H_i(a|v)$$

Winning bid distributions

- Turns out that we only need to look at certain marginals/conditionals of the BCE
- $H(a|v)$ is the CDF of the winning bid conditional v
- $H(a) = \int_{v=\underline{v}}^{\overline{v}} H(a|v)f(v)dv$ is the unconditional CDF
- $H_i(a|v)$ is the prob that i wins **and** the winning bid is less than a, so that

$$H(a|v) = \sum_{i=1}^{n} H_i(a|v)$$

- Bidder i's surplus is

$$U_i = \int_{v=\underline{v}}^{\overline{v}} \int_{x=0}^{\infty} (v-x) H_i(dx|v) f(v) dv$$

Uniform upward deviations

- Consider the **uniform upward deviation** (up to a):
  - Bid a whenever you would have bid $x \leq a$ in equilibrium;
  - if you would have bid $x > a$, do not change your action
- If H does not have an atom at a, deviator's surplus is

$$\int_{v=\underline{v}}^{\overline{v}} \left((v-a)H(a|v) + \int_{x=a}^{\infty} (v-x)H_i(dx|v) \right) f(v)dv$$

Uniform upward deviations

- Consider the **uniform upward deviation** (up to a):
  - Bid a whenever you would have bid $x \leq a$ in equilibrium;
  - if you would have bid $x > a$, do not change your action
- If H does not have an atom at a, deviator's surplus is

$$\int_{v=\underline{v}}^{\overline{v}} \left((v-a)H(a|v) + \int_{x=a}^{\infty} (v-x)H_i(dx|v) \right) f(v)dv$$

- Why? You never bid less than a, so clearly you should win whenever the equilibrium winning bid is less than a

Uniform upward deviations

- Consider the **uniform upward deviation** (up to a):
  - Bid a whenever you would have bid $x \leq a$ in equilibrium;
  - if you would have bid $x > a$, do not change your action
- If H does not have an atom at a, deviator's surplus is

$$\int_{v=\underline{v}}^{\overline{v}} \left((v-a)H(a|v) + \int_{x=a}^{\infty} (v-x)H_i(dx|v) \right) f(v)dv$$

- Why? You never bid less than a, so clearly you should win whenever the equilibrium winning bid is less than a
- But if the equilibrium winning bid is above a, the outcome is not changed by this deviation
- Either the deviator would have won, in which case the deviation does not affect the bid, or someone else would have won, in which case they still outbid the deviator

Uniform upward incentive constraints

- These deviations must not be attractive,

$$\int_{v=\underline{v}}^{\overline{v}} \int_{x=0}^{\infty} (v-x) H_i(dx|v) f(v) dv$$
$$\geq \int_{v=\underline{v}}^{\overline{v}} \left((v-a) H(a|v) + \int_{x=a}^{\infty} (v-x) H_i(dx|v) \right) f(v) dv$$

Uniform upward incentive constraints

- These deviations must not be attractive,

$$\int_{v=\underline{v}}^{\overline{v}} \int_{x=0}^{\infty} (v-x) H_i(dx|v) f(v) dv$$
$$\geq \int_{v=\underline{v}}^{\overline{v}} \left((v-a) H(a|v) + \int_{x=a}^{\infty} (v-x) H_i(dx|v) \right) f(v) dv$$

- Rearranging and integrating by parts:

$$\int_{v=\underline{v}}^{\overline{v}} (v-a)(H(a|v) - H_i(a|v)) f(v) dv \leq \int_{v=\underline{v}}^{\overline{v}} \int_{x=0}^{a} H_i(x|v) dx f(v) dv$$

Stochastic dominance ordering

- A CDF F **first-order stochastically dominates** a CDF F' if $F(x) \leq F'(x)$ for all $x \in \mathbb{R}$
- Equivalent characterization: for any monotonic function, the expectation under F is greater than or equal to the expectation under F'
- NB a partial order on probability distributions

Main theorem

- Define

$$b^*(v) = \frac{1}{(F(v))^{\frac{N-1}{N}}} \int_{w=\underline{v}}^{v} w \frac{N-1}{N} \frac{f(w)dw}{(F(w))^{\frac{1}{n}}}$$

$$H^*(a) = F((b^*)^{-1}(a))$$

Theorem

If H is induced by some $(\mathcal{I}, b)$, then H first-order stochastically dominates H^*. Moreover, there exists $\mathcal{I} = (S, \pi)$ and a BNE b such that $(\mathcal{I}, b)$ induce H^* and the marginal of π on V is f.

- Minimum revenue across BCE is therefore

$$\Pi^{FPA} = \int_{v=\underline{v}}^{\overline{v}} b^*(v)dF(v)$$

The argument

Step 0: We can lower bound equilibrium H by showing that there is a minimal H that satisfies UUIC

The argument

Step 0: We can lower bound equilibrium H by showing that there is a minimal H that satisfies UUIC

Step 1: Linearity of UUIC and expected revenue imply WLOG to look at symmetric solutions

The argument

Step 0: We can lower bound equilibrium H by showing that there is a minimal H that satisfies UUIC

Step 1: Linearity of UUIC and expected revenue imply WLOG to look at symmetric solutions

Step 2: WLOG to restrict attention to solutions in which v and the winning are perfectly positively correlated (lowers the LHS of the UUIC as much as possible)

The argument

Step 0: We can lower bound equilibrium H by showing that there is a minimal H that satisfies UUIC

Step 1: Linearity of UUIC and expected revenue imply WLOG to look at symmetric solutions

Step 2: WLOG to restrict attention to solutions in which v and the winning are perfectly positively correlated (lowers the LHS of the UUIC as much as possible)

Step 3: Monotonic and symmetric solutions are described by **winning bid functions**; can show, using a contraction mapping argument, that b^* is the minimal winning bid function that satisfies UUIC

Step 4: Equilibrium construction

- Last step: Construct $(\mathcal{I}, b)$ that induce H^*
- Worst-case information structure $\mathcal{I}^*$:
  - Signals $S_i = [\underline{v}, \overline{v}]$
  - $s_i \in S_i$ are iid draws from $F^{1/N}$
  - Value is equal to $\max_i s_i$
- Equivalently, a randomly selected bidder observes the true v, and the others see iid draws on $[\underline{v}, v]$ from $(F(s)/F(v))^{1/N}$

Step 4: Equilibrium construction

- Last step: Construct $(\mathcal{I}, b)$ that induce H^*
- Worst-case information structure $\mathcal{I}^*$:
  - Signals $S_i = [\underline{v}, \overline{v}]$
  - $s_i \in S_i$ are iid draws from $F^{1/N}$
  - Value is equal to $\max_i s_i$
- Equivalently, a randomly selected bidder observes the true v, and the others see iid draws on $[\underline{v}, v]$ from $(F(s)/F(v))^{1/N}$
- Can show that b^* is a symmetric pure-strategy equilibrium!
  - In fact, b^* is the equilibrium of the FPA in the IPV model where $v_i \sim F^{1/N}$, denoted $\mathcal{I}^{IPV}$
  - Downward deviations look the same in $(\mathcal{I}^*, b^*)$ and $(\mathcal{I}^{IPV}, b^*)$
  - In $(\mathcal{I}^*, b^*)$, bidders are **indifferent** to all upward deviations

Back to auction design

- We have been analyzing a fixed mechanism, the FPA
- Turns out that the FPA is robustly optimal, among a certain class of mechanisms, when values are common
- We say that a mechanism $\mathcal{M} = (A, q, t)$ 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 information structures in which equilibrium bidder surplus is non-negative
- 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 b is a symmetric and monotonic pure-strategy equilibrium of $(\mathcal{M}, \mathcal{I}^{IPV})$. Then b is also an equilibrium of $(\mathcal{M}, \mathcal{I}^*)$.

- Proof: The allocation induced by b is precisely that induced by the monotonic pure-strategy equilibrium of the FPA
- Revenue equivalence then implies that the interim expected transfer is $T_i(s_i) = T^*(s_i) + c_i$, where T^* is the interim transfer in the FPA in the equilibrium b^*, and c_i is a constant

Proof, continued

- Let $U(s_i, s_i')$ denote the payoff when s_i bids $b(s_i')$ in $(\mathcal{M}, \mathcal{I}^*)$
- U^{FPA} is the corresponding payoff in $(\mathcal{M}^{FPA}, \mathcal{I}^*)$
- Then

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

- But $U_i^{FPA}(s_i, s_i) \geq U_i^{FPA}(s_i, s_i')$ for any s_i', so $U_i(s_i, s_i) \geq U_i(s_i, s_i')$ 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 b
- 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

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

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

Proof

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

Going further

- The theorem is a kind of robust revenue ranking: the FPA has a higher "revenue guarantee" across all BCE than any other standard auction, e.g., 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 Π^{FPA}
- Regardless of beliefs, 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 Π^{FPA}...

Revenue guarantee equivalence with affiliated values

- A famous result of Milgrom and Weber (1982) is that in affiliated-values environments, English and second-price auctions generate more revenue than first-price auctions (in the particular equilibria they construct)
- Interestingly, $\mathcal{I}^*$ is affiliated
- But the strict revenue ranking relies on correlation in signals, which is absent on $\mathcal{I}^*$
- Combining these observations, we get the following result:
- If one selects the MW equilibria of the SPA and English auction, and if we restrict attention to symmetric affiliated-values information structures, then the three mechanisms are **revenue guarantee equivalent**

A saddle point

- The FPA achieves maxmin revenue (with the min over BCE) when we restrict attention to standard mechanisms
- Moreover, $(\mathcal{M}^{FPA}, \mathcal{I}^*)$ are a **saddle point**:
  - For the mechanism $\mathcal{M}^{FPA}$, revenue is at least Π^{FPA} in all information structures and equilibria
  - On $\mathcal{I}^*$, no standard mechanism can achieve more revenue than Π^{FPA} in all information structures and equilibria
- Could this be a saddle point for the unrestricted problem, where the seller can choose any mechanism?

A saddle point

- The FPA achieves maxmin revenue (with the min over BCE) when we restrict attention to standard mechanisms
- Moreover, $(\mathcal{M}^{FPA}, \mathcal{I}^*)$ are a **saddle point**:
  - For the mechanism $\mathcal{M}^{FPA}$, revenue is at least Π^{FPA} in all information structures and equilibria
  - On $\mathcal{I}^*$, no standard mechanism can achieve more revenue than Π^{FPA} in all information structures and equilibria
- Could this be a saddle point for the unrestricted problem, where the seller can choose any mechanism? No!

Numerical example

- Taking $v \sim U[0,1]$, then

$$b^*(v) = \frac{1}{v^{\frac{N-1}{N}}} \int_{w=0}^{v} w \frac{N-1}{N} \frac{dw}{w^{\frac{1}{N}}} = \frac{N-1}{2N-1} v$$

- Thus, the minimum winning bid distribution is $U[0,(N-1)/(2N-1)]$
- As $N \to \infty$, converges to $U[0,1/2]$, i.e., limit revenue and bidder surplus are $1/4$
- NB $1/4$ is less than the total surplus of $1/2$, so bidders still get rents in the limit
- Very different from older positive results of Wilson (1977) and Milgrom (1979) that show asymptotic full surplus extraction of the FPA in the mineral-rights model

Robust surplus extraction with many bidders

- Begs the question, are there other mechanisms that extract more revenue in all BCE?
- Du (2018) constructs a sequence of mechanisms that asymptotically have revenue equal to $\hat{v}$, regardless of F
- *A fortiori*, even with finite N, the revenue guarantee of the FPA is improvable

Du (2018)

- Consider the following $\mathcal{M}^N = ([0,1]^N, q^N, t^N)$
- If $a_{i_1} \geq a_{i_2} \cdots \geq a_{i_N}$, then

$$q^N_{i_k}(a) = \sum_{k'=1}^{k} \frac{a_{k'} - a_{k'+1}}{k'}$$

- Interpretation:
  - The good is composed of a continuum of pieces $x \in [0,1]$
  - Action a_i mans “I want to buy the pieces $x \leq a_i$
  - Each piece is divided equally among the bidders that demand it
- The transfer is

$$t^N_i(a) = \frac{2\overline{v}}{N \log(N)} \left(\exp(a_i \log(N)) - 1\right)$$

- NB $t_i(0, a_{-i}) = 0$, so in any equilibrium, $U_i \geq 0 \; \forall i$

Main result

Theorem (Du, 2018)

In any BCE of $\mathcal{M}^N$, revenue is equal to $\hat{v} + O(\log(N))$.

- Proof: Expected revenue is

$$\Pi = \sum_i \int_{(v,a)} t_i^N(a)\sigma(dv, da)$$

- Note that

$$t_i^N(a) = \frac{1}{\log(N)} \frac{\partial t_i^N(a)}{\partial a_i} - \frac{2\overline{v}}{N\log(N)}$$

- Hence,

$$\Pi = \frac{1}{\log(N)} \sum_i \int_{(v,a)} \frac{\partial t_i^N(a)}{\partial a_i} \sigma(dv, da) - \frac{2\overline{v}}{\log(N)}$$

Proof, cont'd

- Can show that $a_i < 1$ with prob 1 in equilibrium (that's why we have the 2 in the transfer)
- Moreover,

$$\sum_i \int_{(v,a)} \left(v \frac{\partial q_i^N(a)}{\partial a_i} - \frac{\partial t_i^N(a)}{\partial a_i} \right) \leq 0$$

- (Otherwise a positive measure of "local upward" obedience constraints would be violated for some player)
- We conclude that

$$\Pi \geq \frac{1}{\log(N)} \sum_i \int_{(v,a)} v \frac{\partial q_i^N(a)}{\partial a_i} \sigma(dv, da) - \frac{2\overline{v}}{\log(N)}$$

Proof cont'd

- Furthermore,

$$\sum_{i=1}^{N} \frac{\partial q_i(a)}{\partial a_i} = \sum_{i=1}^{N} \frac{1}{|\{j|a_j > a_i\}| + 1} \geq \sum_{i=1}^{N} \frac{1}{i} > \log(N+1)$$

- Putting it all together,

$$\begin{aligned}
\Pi &= \sum_i \int_{(v,a)} t_i^N(a)\sigma(dv, da) \\
&= \frac{1}{\log(N)} \sum_i \int_{(v,a)} \frac{\partial t_i^N(a)}{\partial a_i} \sigma(dv, da) - \frac{2\overline{v}}{\log(N)} \\
&\geq \frac{1}{\log(N)} \sum_i \int_{(v,a)} v \frac{\partial q_i^N(a)}{\partial a_i} \sigma(dv, da) - \frac{2\overline{v}}{\log(N)} \\
&\geq \frac{1}{\log(N)} \hat{v} \log(N+1) - \frac{2\overline{v}}{\log(N)},
\end{aligned}$$

which goes to $\hat{v}$ at a rate of $1/\log(N)$ $\square$

Remarks

- Both BBM (2017, 2019) and Du (2018) use obedience constraints to lower bound revenue
- Du (2018) achieves a better bound (asymptotically)
- The key difference is that revenue is tied to local incentives
- This is done through the particular choice of transfer rule
- One then has to control the sensitivity of the allocation in order to achieve a favorable revenue guarantee
- The next lecture will pursue this analysis to its logical conclusion, and use local incentives/allocation sensitivity to identify maxmin mechanisms for finite N