Day 1. BNE & the IPV Environment

1 Review – (Static) Nash Equilibrium

- The standard setup for a static game of complete information requires three elements:
  - A set of players $N = \{1, 2, \ldots, n\}$
  - A set of available actions A_i for each player
  - A set of payoff functions
    $$U_i : A_1 \times A_2 \times \cdots \times A_n \to \Re$$
    giving the payoff to each player given each profile of strategies
- A (pure strategy) Nash equilibrium is an action profile $a = (a_1, a_2, \ldots, a_n)$ such that each player is best-responding to the remaining players' actions – that is, such that for all i,
  $$a_i \quad \in \quad \arg\max_{a_i' \in A_i} U_i(a_i', a_{-i})$$
  where a_{-i} is a vector of the other $n-1$ players' actions.
- A key maintained assumption is that the entire environment is *common knowledge*
  - all the players know the set of players, actions and payoffs,
    everyone knows everyone knows it,
    and so on.

- So what happens if we *don't* have common knowledge about the entire environment?
- For example, suppose you know your own payoff function, but you're not sure of mine.
- That's a game of *incomplete information*

2 Games of Incomplete Information

- The way we deal with incomplete information is to assume that there are different possible *types* of each player
- For example, suppose the incomplete information is that you don't know how willing I am to fight
- We assume this is due to you not knowing which type of player I am – a tough player, or a weak player
- I know which one I am, but you don't – so you just assign some probability to me being strong (having one payoff function), and some probability to me being weak (having a different payoff function)
- And of course, my equilibrium strategy would allow the "tough me" to play differently from the "weak me"

- Formally, we assume each player's type is his own private information – I know which type I am, but not which types my opponents turned out to be,

  but we all agree on the ex ante probability of each type of each player
- So a static game of incomplete information consists of...
  1. A set of players $N = \{1, 2, \dots, n\}$, same as before
  2. A set of actions for each player A_i, same as before
  3. A set of possible types T_i for each player
  4. A probability distribution p over the set of type profiles $T_1 \times T_2 \times \cdots \times T_n$
  5. A set of payoff functions, that may depend on every player's type:

$$U_i : A_1 \times \cdots \times A_n \times T_1 \times \cdots \times T_n \to \Re$$

- And *all of that* is common knowledge – if there are lots of players in the game, everyone agrees on the probability that I'm tough, and I know what they think the probability is, and so on – I just also know what my type "turned out" to be

- Several things to note about this setup:
  - It doesn't matter whether we allow the set of available actions A_i to depend on player i's type; if we wanted to "disallow" a particular action a_i for a particular type t_i of player i, we could keep a_i in A_i and just set $U_i(a_i, \cdot, t_i, \cdot) = -\infty$. So for simplicity, we assume A_i is fixed across types.
  - For now, information is defined by the type space – player i is assumed to know the exact value of T_i, but know nothing about T_{-i} beyond what he infers from T_i. (In some papers this semester, we'll relax this assumption.)

    We very often, for simplicity, assume that different players' types are independent, but we don't have to. (If they're not independent, then learning my own type t_i changes my beliefs about the probabilities of your different types.)
  - In many models, we assume that player i's payoff does not depend on other players' types, only his own (and the action profile). That is, we often (but not always) assume that I care about other players' types *only because it may influence their actions*, not because it directly influences my payoffs. But there are exceptions – in an adverse selection model, for example.
  - Like before, we assume that the entire setup above is common knowledge – everyone agrees on the basic universe we live in, but only player i knows which type he turned out to be.
  - It's customary to think of the game happening in two stages: in the first stage, "nature moves" by randomly assigning a type to each player; in the second stage, players play the game given their realized types
  - (This is based on Harsanyi's insight that we can think of a game of incomplete information – a game where, say, your payoff function is not common knowledge – as instead being a game of imperfect information – a game where your payoff function is common knowledge, it just depends on a piece of information (your type) which you have and I don't have.)

3 Bayesian Nash Equilibrium

The solution concept for a static game of incomplete information is Bayesian Nash equilibrium, which is really just a generalization of Nash equilibrium to accommodate types. Specifically...

- A *strategy* for player i is now a type-dependent choice of action, that is, a mapping

  $$s_i : T_i \to A_i$$

  specifying an action $s_i(t_i) \in A_i$ that I plan to take for each type I might turn out to be

- A Bayesian Nash equilibrium is a profile of strategies $(s_1, s_2, \ldots, s_n)$ such that for every player i and every type t_i,

  $$s_i(t_i) \quad \in \quad \arg\max_{a_i \in A_i} E_{t_{-i}|t_i} U_i(a_i, s_{-i}(t_{-i}), t_i, t_{-i})$$

  That is, each type of each player is maximizing his expected payoff, given his correct beliefs about the probabilities of different opponent types (given his own) and given his correct beliefs about his opponents' strategies

4 An Example of BNE: First Price Auctions

- Let's do a simple example: a private-values, sealed-bid first-price auction

  There is a single object for sale

  Each player's type is the value he would get from winning it; everyone simultaneously submits a bid in writing, and the player with the highest bid pays his bid and receives the object

  Let's suppose types are independent and uniform over the interval $[0, 1]$

  So formally:

  - The set of players is $1, 2, \ldots, n$
  - The set of available actions is $A_i = \Re^+$ – any positive bid is allowed
  - The set of possible types is $T_i = [0, 1]$, and the probability distribution over type profiles is uniform over $[0, 1]^n$ (meaning it has density 1 everywhere)
  - The payoff from winning the auction is your type minus your bid, and the payoff from losing is 0, so if we ignore ties, we can write

    $$U_i(a, t) = \begin{cases} t_i - a_i & if \quad a_i > \max_{j \neq i} a_j \\ 0 & if \quad a_i < \max_{j \neq i} a_j \end{cases}$$

  - To complete the model, let's assume they're broken randomly; so if $a_i = \max_{j \neq i} a_j$, the payoff to player i is

    $$(t_i - a_i) \frac{1}{1 + \#\{j \neq i : a_j = a_i\}}$$

- Note that we've assumed *private values* – other players' types don't enter directly into my payoff function, they affect me only through their effect on my opponents' bids.

- We're also implicitly assuming risk neutrality – bidders maximize the expected value of their payoff

- What does the Bayesian Nash equilibrium of this game look like? It turns out, it's an equilibrium for everyone to bid $\frac{n-1}{n}$ times their type.

- Why?

- To show this, we need to show that if all my opponents are playing this strategy, then this strategy maximizes my expected payoff.

- So suppose my type is t_i, my opponents are all bidding $\frac{n-1}{n}$ times their values, and I bid b.

- First of all, note that if I bid more than $\frac{n-1}{n}$, I'll win with probability 1 – I outbid any opponent with type $t_j < 1$, and tie opponents with type $t_j = 1$, which occurs with prob 0

  So bidding more than $\frac{n-1}{n}$ is a bad idea – it drives up the price I pay, without making me more likely to win. So the only strategies that might be best-responses are in the range $[0, \frac{n-1}{n}]$.

- Now, for b within that range, my expected payoff is

$$\begin{aligned} E_{T_{-i}} U_i &= (t_i - b) \cdot \Pr(win|b) + 0 \cdot \Pr(lose|b) \\ &= (t_i - b) \Pr(\max_{j \neq i} s_j(t_j) < b) \\ &= (t_i - b) \Pr(\max_{j \neq i} \tfrac{n-1}{n} t_j < b) \\ &= (t_i - b) \Pr(\max_{j \neq i} t_j < \tfrac{n}{n-1} b) \\ &= (t_i - b)(\tfrac{n}{n-1} b)^{n-1} \\ &= (\tfrac{n}{n-1})^{n-1} (t_i - b) b^{n-1} \end{aligned}$$

- So now let's maximize this thing: $(t_i - b)b^{n-1} = t_i b^{n-1} - b^n$ has derivative

$$(n-1)t_i b^{n-2} - n b^{n-1} \quad = \quad n b^{n-2} \left[\frac{n-1}{n} t_i - b \right]$$

  This is positive on $b < \frac{n-1}{n} t_i$ and negative on $b > \frac{n-1}{n} t_i$ – so expected payoff is maximized at $b = \frac{n-1}{n} t_i$.

- So if everyone else bids $\frac{n-1}{n}$ times their type, my best-response is to bid $\frac{n-1}{n}$ times my type; so everyone bidding $\frac{n-1}{n}$ times their type is an equilibrium.

5 Second-price auctions

- I want to talk about one other common auction format
- Description of open oral ascending, or English, auctions
- A modeling convenience: for the private values case, think of a sealed-bid *second*-price auction
  - Bidders simultaneously submit written bids, as in a first-price
  - High bidder wins, but his payment is the *second*-highest bid, not his own
- With private values, turns out it's a dominant strategy to bid your valuation!
- Think of y a random variable being the highest opponent bid; you have a choice between bidding above y (for payoff $t_i - y$) and bidding below y (and getting 0)
- So when $t_i > y$, you want to overbid, and when $t_i < y$, you want to underbid – which you accomplish by bidding t_i
- So it's a Bayesian Nash equilibrium for everyone to bid their valuation

6 Revenue Equivalence

- We determined that in a first-price auction, it's an equilibrium for each bidder to bid $\frac{n-1}{n}$ times his or her valuation
- So in a first-price auction, the seller's expected revenue is the expected value of the highest bid, which is $\frac{n-1}{n}$ times the expected value of the highest valuation
- Turns out, with n independent uniform random variables, the expected value of the highest is $\frac{n}{n+1}$
- So expected revenue in the first-price auction is $\frac{n-1}{n} \times \frac{n}{n+1} = \frac{n-1}{n+1}$

  (Expected value of the highest of n independent $U[0,1]$ random variables is calculated by taking the CDF, x^n; calculating the PDF, nx^{n-1}, and taking the integral $\int_0^1 x \cdot nx^{n-1} dx = \int_0^1 nx^n dx = \frac{n}{n+1}$.)
- What about the second-price auction?
- In the second-price auction, everyone bids his own type, and the payment to the seller is the second-highest bid
- So revenue is the second-highest valuation
- Which has expected value $\frac{n-1}{n+1}$

  (The CDF of the second-highest of n is $nx^{n-1}(1-x) + x^n = nx^{n-1} - (n-1)x^n$, so the PDF is $n(n-1)(x^{n-2} - x^{n-1})$, and our expected value is $\int_0^1 n(n-1)(x^{n-1} - x^n)dx = (n-1)x^n - \frac{n(n-1)}{n+1}x^{n+1}\Big|_0^1 = n - 1 - \frac{n^2-n}{n+1} = \frac{n^2-1-n^2+n}{n+1} = \frac{n-1}{n+1}$.)
- So the two auction formats give the same expected payoff to the seller
- This turns out to be a much more general result, which we'll prove (in passing) next week

7 Anyone Remember the Envelope Theorem?

- Put aside strategic concerns, and think of a one-agent decision problem

  $$\max_{a \in A} h(a, \theta)$$

  where a is the agent's chosen action and θ an exogenous parameter

- In an auction, θ could be your valuation, and a your choice of bid

- A could be either discrete or continuous, but let θ be continuous

- Let $a^*(\theta)$ be the set of optimal choices, and $V(\theta)$ the value function;

  and let h_a and h_θ denote partial derivatives of h

Theorem (The Envelope Theorem)**.** *Suppose $\forall \theta$, $a^*(\theta)$ is nonempty, and $\forall (a, \theta)$, h_θ exists. Let $a(\theta)$ be any selection from $a^*(\theta)$.*

1. *If V is differentiable at θ, then*

   $$V'(\theta) = h_\theta(a(\theta), \theta)$$

2. *If V is absolutely continuous, then for any $\theta' > \theta$,*

   $$V(\theta') - V(\theta) = \int_\theta^{\theta'} h_\theta(a(t), t) dt$$

- This says: the derivative of the *value function* (or maximum) is the derivative of the *objective function*, evaluated *at the maximizer*

- Absolute continuity says that $\forall \epsilon > 0$, $\exists \delta > 0$ such that for any finite, disjoint set of intervals $\{[x_k, y_k]\}_{k=1,2,\ldots,M}$ with $\sum_k |y_k - x_k| < \delta$, $\sum_k |V(y_k) - V(x_k)| < \epsilon$.

  In many auction models, it's easy to show that $\|V(y) - V(x)\| \leq \|y - x\|$, which means V is absolutely continuous

  Absolutely continuity is equivalent to V being differentiable almost everywhere and being the integral of its derivative, so the second part of the theorem follows directly from the first.

  (For the types of problems we'll be dealing with, value functions will generally be absolutely continuous.)

8 Using the envelope theorem to solve for equilibrium bids in FPA

- For the setup we just defined (PV *i.i.d.* $\sim U[0,1]$), we can use the envelope theorem to recover equilibrium bid functions

- Suppose a symmetric equilibrium exists, and uses a strictly-increasing bid function, so the bidder with the highest type always wins in equilibrium

- Equilibrium bids solve $\max_b (t_i - b) \Pr(b > \max_{j \neq i} b_i)$, call maximand $h(t_i, b)$ and U its max

- The partial of h with respect to t_i is $\Pr(b > \max_{j \neq i} b_i)$

- If we assume the equilibrium is symmetric, then at the maximizer, this equals the probability my type is the highest, which is t_i^{n-1}

- Since the lowest-type bidder always loses, $U(0) = 0$, so

$$U(t_i) \quad = \quad \int_0^{t_i} s^{n-1} ds \quad = \quad \frac{1}{n} t_i^n$$

- But if we calculate expected payoff based on bid and probability of winning,

$$U(t_i) \quad = \quad (t_i - \beta(t_i))\, t_i^{n-1}$$

- And if we equate these,

$$\frac{1}{n} t_i^n \quad = \quad (t_i - \beta(t_i)) t_i^{n-1}$$

  implies $\beta(t_i) = \frac{n-1}{n} t_i$

- So if a symmetric, strictly-increasing equilibrium exists, bids must be $\frac{n-1}{n} t_i$ – exactly like we found earlier

9 Another example – the all-pay auction

- We can also consider an *all-pay* auction – an auction where the highest bidder wins, but *every* bidder pays their bid, even the losers

- Again, if we suppose a symmetric, strictly-monotonic equilibrium exists, we can use it to calculate the expected payoff to each type of bidder, which is the same $U(t_i) = \frac{1}{n} t_i^n$

- We then write $\frac{1}{n} t_i^n = t_i^{n-1} t_i - b$ and calculate equilibrium bids as $b = \frac{n-1}{n} t_i^n$

- Finally, we can calculate the seller's expected revenue as $n E_{t_i} \frac{n-1}{n} t_i^n$ and see that it's the same $\frac{n-1}{n+1}$ as before – more anecdotal evidence of revenue equivalence!

10 Now on to Myerson's Optimal Auctions

First, we define the environment, and our goal.

- Our environment is as follows.
    - Players $N = \{1, 2, \ldots, n\}$
    - Independent types $T_i \perp T_j$
    - Player i's type T_i has distribution F_i with support $[a_i, b_i]$
    - F_i is strictly increasing on $[a_i, b_i]$, with density f_i
    - One object to be allocated, value to player i is his type t_i, and players value money linearly, so i's payoff is $t_i - x$ if he receives object and pays x, and $-x$ if he doesn't receive object but still pays x
    - (Players are risk-neutral and maximize expected payoff)
    - Seller values keeping the good at t_0
- Note that this is the environment we looked at above for first- and second-price auctions, except that type distributions F_i can be arbitrary (not just uniform) and don't have to be the same
- We're still assuming:
    - Independent private values
    - Risk-neutrality (linear valuation for money)
    - Whole environment is common knowledge
- Our challenge: design an auction (or other) protocol that maximizes the seller's expected revenue, given that the buyers play a Bayesian Nash equilibrium, over all conceivable sales mechanisms.
- That is, we basically want to solve

  $$\max_{all\ \ conceivable\ \ auction\ \ formats} E_{t_1,\ldots,t_n} \{Revenue + t_0 \Pr(seller\ \ keeps\ \ object)\}$$

  subject to the constraints that the buyers have to willingly participate and play an equilibrium
- The surprising thing: this is actually doable.