Lecture 7

Bilateral Private Information

If buyer and seller both have private information,
full efficiency is impossible.

- In the last two classes, we've solved the seller's optimization problem when facing several buyers with private information...
- Under the assumption that the bidders are risk-neutral and have *independent* private values
- We found that...
  - In the symmetric case, the best the seller can do is a "standard" auction with an optimally-set reserve price
  - In the asymmetric case, we can derive a rule for the optimal auction, but it doesn't correspond to anything "obvious"
- Today, we'll do two extensions
- First: what happens when bidders' values are correlated, rather than independent?
- And second: what happens when both buyer and seller have private information?

1 Mechanism Design with Correlated Values

- If bidders' values are correlated, things change significantly
- In fact, if bidders have correlated values, this is good for the seller
- Under the assumptions we've been making – the entire environment (including the joint distribution of bidders' values) is known to both the seller and all the bidders – the seller can use this correlation against the bidders, to extract more of the surplus
- Myerson gives an example to show that the seller may be able to do the best he could *possibly* do – achieve the efficient outcome (maximize total combined payoffs), *and* give every bidder expected payoff 0 regardless of type, so he's maximizing total surplus and capturing all of it

  Given individual rationality, that's the best the seller could possibly hope to do
- The example is discrete, but the intuition extends to continuous cases as well

  In a later paper, Cremer and McLean give a sufficient condition for this type of trick to work in general.
- So, here's Myerson's example.
  - Two bidders
  - The joint distribution of their types (t_1, t_2) is $\Pr(100, 100) = \Pr(10, 10) = \frac{1}{3}$, $\Pr(10, 100) = \Pr(100, 10) = \frac{1}{6}$
  - For simplicity, suppose $t_0 = 0$
- Consider the following direct mechanism:
  - If both bidders report high types, flip a coin, and give one of them the object for 100
  - If one reports high and one reports low, sell the high guy the object for 100, charge the low guy 30 and give him nothing
  - If both report low, flip a coin, give one 15 and give the other 5 plus the object
- First, we'll verify that this is feasible, that is, that truthful revelation is an equilibrium; then we'll consider the logic behind it
  - Suppose your opponent will be reporting his true type
  - If you have a low value, you know, conditional on that, the the other guy is low with probability $\frac{2}{3}$. So if you declare low, you get an expected payoff of

$$\frac{2}{3}(15) + \frac{1}{3}(-30) = 0$$

- On the other hand, if you misreport your type as high, then with probability $\frac{2}{3}$, your opponent reports low and you pay 100 for an object you value at 10; and with probability $\frac{1}{3}$, you both report high, and with probability $\frac{1}{2}$ you pay 100 for the object; so your payoff is

$$\frac{2}{3}(10-100)+\frac{1}{3}\cdot\frac{1}{2}(10-100)=\frac{5}{6}(-90)<0$$

  so your best response is clearly to report truthfully

- Now suppose you have a high value
- If you report high, then either you win the thing and pay 100, or you get nothing, so your payoff is 0
- If you report low, then with probability $\frac{2}{3}$, your opponent reports high and you lose 30. With probability $\frac{1}{3}$, your opponent reports low, and you either get the object plus 5, or you get 15; so your expected payoff is

$$\frac{2}{3}(-30)+\frac{1}{3}\left(\frac{1}{2}(15)+\frac{1}{2}(105)\right)=-20+\frac{1}{6}(120)=0$$

- So it's a weak best-response to report truthfully

• So this mechanism is incentive compatible

• Since both types get expected payoff 0 in equilibrium, it's also individually rational

• So it's feasible

• Also note that under truthful revelation, it's efficient – the object is always allocated to someone, and a low type never gets it if a high type is available

• So the seller is achieving both efficiency and full surplus extraction – his best-case scenario

Now, where the hell did this auction come from?

• In a "normal" setting, bidders with high types get positive expected payoffs, because you have to "leave them" some surplus to prevent them from lying and saying they have a low value

• But in this case, a bidder with a high type has information about the other type's likely value

• So what you do here is this: when a bidder claims to have a low type, you also force him to accept a "side-bet" with you about the other bidder's type

• And this bet is rigged to have 0 expected value when his type is actually low, but negative expected value when his type is actually high

- Since bidders are risk-neutral, this doesn't hurt low types, but it lowers the payoff the high type gets from misreporting his type; so it lowers the payoff you have to give him when he reports truthfully

- To see how this mechanism works, think of it this way:

  - Hold a first-price auction, where the only bids allowed are 10 and 100
  - But a bidder who bids 10 also accepts the following bid: "if the other bidder bids 10, I win 15; if the other bidder bids 100, I lose 30"
  - A bidder with low type expects his opponent to have a low type two-thirds of the time, so this bet has 0 expected payoff
  - A bidder with high type expects his opponet to have a high type two-thirds of the time, so this bet has expected value $\frac{1}{3}(15) + \frac{2}{3}(-30) = -15$
  - Since a high bidder who bids low only wins with probability $\frac{1}{6}$, this wipes out the surplus of 90 he would get the times he did win.

- Cremer and McLean (1988, "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions", Econometrica 56.6) generalize this, showing when full surplus extraction is possible with correlated types

  - They assume finite, discrete types
  - Full surplus extraction is possible if for each bidder, the matrix of probabilities of his opponents' types given his own types has full rank – so that his beliefs about his opponents are different enough for the right side bets to be created

- However, this sort of mechanism does not seem very robust

  - more than most auctions, it is extremely sensitive to the seller being right about the true distribution of types
  - it also clearly requires risk-neutrality, since you need the low types to be willing to accept a large zero-expected-value bet
  - finally, it's very sensitive to collusion – both bidders bidding low is profitable for them (it's even an equilibrium!)

- I don't think Myerson is suggesting you would actually run this auction in this way – just making the point that when bidder values are correlated, the optimal auction may be complicated but may outperform anything you would think of as a "regular-looking" auction.

2 Bilateral Private Information

- So far, we've focused on the case of "one-sided" private information
  - buyers have private information, but the seller doesn't
  - or if he does, it's doesn't matter, because he's the one designing the mechanism
- Now we shift to the case where there is a single buyer and a single seller, but both have private information about his own valuation of the good
- Environment: $v_b \sim F_b$, $v_s \sim F_s$
- For simplicity, assume $\text{supp}(F_b) = \text{supp}(F_s) = [0, 1]$
- Think of v_s as the cost to the seller from giving up the good
  - just like the buyer gets payoff $v_b - p$ from buying, the seller gets $p - v_s$ from selling
  - and both get payoff 0 if there is no trade
- The general problem is to maximize ex ante surplus
  - imagine buyer and seller choose a mechanism *before* either one knows their type, to maximize expected payoff once they learn types and trade
  - but individual rationality still has to hold for each type – after they learn their types, either one could still refuse to play
- The revelation principle still holds – we can imagine both buyer and seller report their types to a neutral third party ("mediator"), who then tells them whether to trade and how much to pay each other
- Let $p(v_b, v_s)$ be the probability they trade, given types v_b and v_s
- The question: is there a feasible mechanism where trade happens whenever $v_b > v_s$, i.e.,

$$p(v_b, v_s) = \begin{cases} 1 & if \quad v_b > v_s \\ 0 & if \quad v_b < v_s \end{cases}$$

- The answer turns out to be no. From Myerson and Satterthwaite (1983):

Theorem. *If buyer and seller each have private information about their own private value, and the support of their valuations overlap, there is no feasible mechanism that yields fully efficient trade.*

- To prove it, we will suppose such a mechanism did exist, and then show it would have to violate individual rationality for either low-value buyers or high-cost sellers.
- Remember, anything we can accomplish with any mechanism, we can do with a direct mechanism, so we'll do that
- Let $p_b(v_b) = E_{v_s} p(v_b, v_s)$ be a buyer's expected probability of trade given reported type v_b, and let $x_b(v_b)$ be his expected payment
- Given the seller is reporting truthfully, a buyer with type v_b gets expected payoff

$$U_b(v_b) = \max_{v_b'} v_b p_b(v_b') - x_b(v_b')$$

- The envelope theorem says that $U_b'(v_b) = p_b(v_b)$, so

$$U_b(v_b) = U_b(0) + \int_0^{v_b} p_b(v) dv$$

- Letting $U_b = E_{v_b} U_b(v_b)$ be the bidder's ex-ante expected payoff,

$$\begin{aligned} U_b &= \int_0^1 \left[U_b(0) + \int_0^{v_b} p_b(v) dv \right] f_b(v_b) dv_b \\ &= U_b(0) + \int_0^1 \left[\int_v^1 f_b(v_b) dv_b \right] p_b(v) dv \\ &= U_b(0) + \int_0^1 (1 - F_b(v)) p_b(v) dv \end{aligned}$$

- Similarly, if we let $p_s(v_s)$ and $x_s(v_s)$ be the seller's probability of trade and expected payment received,

$$U_s(v_s) = \max_{v_s'} x_s(v_s') - v_s p_s(v_s')$$

- So by the envelope theorem, $U_s'(v_s) = -p_s(v_s)$, and therefore

$$U_s(v_s) = U_s(1) + \int_{v_s}^1 p_s(v) dv$$

- Taking the expected value over v_s, then,

$$\begin{aligned} U_s &= \int_0^1 \left[U_s(1) + \int_{v_s}^1 p_s(v) dv \right] f_s(v_s) dv_s \\ &= U_s(1) + \int_0^1 \left[\int_0^v f_s(v_s) dv_s \right] p_s(v) dv \\ &= U_s(1) + \int_0^1 F_s(v) p_s(v) dv \end{aligned}$$

- But we can also calculate combined payoffs as expected gains from trade, which are

$$W = \int_0^1 \int_0^1 (v_b - v_s) p(v_b, v_s) f_b(v_b) f_s(v_s) dv_s dv_b$$

- Since $U_b + U_s = W$, this means

$$U_b(0) + \int_0^1 (1 - F_b(v)) p_b(v) dv + U_s(1) + \int_0^1 F_s(v) p_s(v) dv$$

$$= \int_0^1 \int_0^1 (v_b - v_s) p(v_b, v_s) f_b(v_b) f_s(v_s) dv_s dv_b$$

- Now, $p_b(v_b) = \int_0^1 p(v_b, v_s) f_s(v_s) dv_s$, and $p_s(v) = \int_0^1 p(v_b, v_s) f_b(v_b) dv_b$, so we can write these as

$$U_b(0) + \int_0^1 \int_0^1 \frac{1 - F_b(v_b)}{f_b(v_b)} p(v_b, v_s) f_b(v_b) dv_b f_s(v_s) dv_s$$

$$+ \quad U_s(1) + \int_0^1 \int_0^1 \frac{F_s(v_s)}{f_s(v_s)} p(v_b, v_s) f_b(v_b) f_s(v_s) dv_b dv_s$$

$$= \int_0^1 \int_0^1 (v_b - v_s) p(v_b, v_s) f_b(v_b) f_s(v_s) dv_b dv_s$$

- If we move all the integrals to the right hand side, we get

$$U_b(0) + U_s(1) = \int_0^1 \int_0^1 \left[\left(v_b - \frac{1 - F_b(v_b)}{f_b(v_b)} \right) - \left(v_s + \frac{F_s(v_s)}{f_s(v_s)} \right) \right] p(v_b, v_s) f_b(v_b) f_s(v_s) dv_b dv_s$$

  for any feasible mechanism, which is basically Theorem 1 of Myerson and Satterthwaite.

  - How to interpret this? Think of a mechanism designer setting up a mechanism to simultaneously sell to the buyer and buy from the seller
  - $v_b - \frac{1-F_b(v_b)}{f_b(v_b)}$ is the old Myerson term – the "marginal revenue" from selling to a buyer of type v_b, given incentive compatibility
  - $v_s + \frac{F_s(v_s)}{f_s(v_s)}$ is the analogous term for the seller – the "marginal cost" of buying the good from a seller of type v_s
  - So this is the expected value of, whenever we want the two parties to trade, simultaneously buying from the seller and reselling to the buyer
  - If we were a middle man setting up a mechanism to maximize our own profit, the RHS is the thing that we would maximize
  - Here, there's no middleman, just the actual buyer and seller, so that's the "excess" surplus generated by the mechanism (beyond the payoffs needed to satisfy incentive compatibility) – which is the extra payoff one or both of them get

- For our purposes, we want to show there is no way to get trade whenever it's efficient – that is, there's no feasible mechanism such that

$$p(v_b, v_s) = \begin{cases} 1 & if \quad v_b > v_s \\ 0 & if \quad v_b < v_s \end{cases}$$

- To prove this, suppose there was such a mechanism
- In that case, we would get

$$U_b(0) + U_s(1) = \int_0^1 \int_0^{v_b} \left[\left(v_b - \frac{1 - F_b(v_b)}{f_b(v_b)} \right) - \left(v_s + \frac{F_s(v_s)}{f_s(v_s)} \right) \right] f_b(v_b) f_s(v_s) dv_s dv_b$$

- Split that into two integrals,

$$\begin{aligned} U_b(0) + U_s(1) &= \int_0^1 \int_0^{v_b} \left(v_b - \frac{1 - F_b(v_b)}{f_b(v_b)} \right) f_b(v_b) f_s(v_s) dv_s dv_b \\ &\quad - \int_0^1 \int_0^{v_b} \left(v_s + \frac{F_s(v_s)}{f_s(v_s)} \right) f_b(v_b) f_s(v_s) dv_s dv_b \end{aligned}$$

- Simplifying the first one gives

$$\begin{aligned} INT1 &= \int_0^1 \int_0^{v_b} \left(v_b - \frac{1 - F_b(v_b)}{f_b(v_b)} \right) f_b(v_b) f_s(v_s) dv_s dv_b \\ &= \int_0^1 \left(\int_0^{v_b} f_s(v_s) dv_s \right) \left(v_b - \frac{1 - F_b(v_b)}{f_b(v_b)} \right) f_b(v_b) dv_b \\ &= \int_0^1 F_s(v_b) \left(v_b f_b(v_b) - (1 - F_b(v_b)) \right) dv_b \end{aligned}$$

- Simplifying the second one gives

$$\begin{aligned} INT2 &= \int_0^1 \int_0^{v_b} \left(v_s + \frac{F_s(v_s)}{f_s(v_s)} \right) f_b(v_b) f_s(v_s) dv_s dv_b \\ &= \int_0^1 \left(\int_0^{v_b} \left(v_s f_s(v_s) + F_s(v_s) \right) dv_s \right) f_b(v_b) dv_b \\ &= \int_0^1 \left(v_s F_s(v_s) |_{v_s=0}^{v_b} \right) f_b(v_b) dv_b \\ &= \int_0^1 v_b F_s(v_b) f_b(v_b) dv_b \end{aligned}$$

- Recombining the two integrals, then, gives

$$
\begin{aligned}
U_b(0) + U_s(1) &= INT1 - INT2 \\
&= \int_0^1 F_s(v_b)\left(v_b f_b(v_b) - 1 + F_b(v_b)\right) dv_b - \int_0^1 v_b F_s(v_b) f_b(v_b) dv_b \\
&= -\int_0^1 F_s(v_b)\left(1 - F_b(v_b)\right) dv_b
\end{aligned}
$$

- As long as the supports of the two distributions overlap, the right side is strictly negative
- Which means any mechanism that gave efficient trade would have to give $U_b(0) + U_s(1) < 0$
- Which would violate individual rationality for either the highest seller type or the lowest buyer type, which proves the theorem
- (Myerson and Satterthwaite then go on to solve the optimization problem

$$\max_p \{U_b + U_s\}$$

to find the mechanism that maximizes combined payoffs, knowing it won't be efficient; they solve the problem, but the solution is pretty complicated – read the paper if you're interested!)

References

- Jacques Cremer and Richard McLean (1988), "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica* 56
- Roger Myerson and Mark Satterthwaite (1983), "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory* 29