Lecture Slides - Part 3

Bengt Holmstrom

MIT

February 2, 2016.

Adverse Selection

- We will solve a procurement problem using a screening mechanism
- Idea: buyer wants to buy from seller, but doesn't know seller's cost

Setup

- Two players, buyer B and seller S
- $v(x)$: value of x units to B
- $c(x, \theta)$: cost of producing x by S depending on his type θ
- Payoffs: $u_B(x, t) = v(x) - t$, $u_S(x, t, \theta) = t - c(x, \theta)$, where t is payment from B to S
- Assumptions: $v' > 0$, $v'' \leq 0$, $v(0) = 0$
- $c_{x\theta} < 0$ (higher types have lower marginal cost), $c(0, \theta) = 0\ \forall\theta$, $c_x > 0$ (positive MC)

- B designs $t(x)$, a nonlinear price schedule specifying a payoff for each quantity
- Given $t(x)$, under some conditions, a seller of type θ will choose a quantity $x(\theta)$ such that marginal cost equals marginal payoff from one more unit: $c_x(x(\theta), \theta) = t'(x)$
- Note: no matter how B designs $t(x)$, *lower cost sellers always produce more*
- Easiest to prove using increasing differences

- Note: if there are k (finitely many) types, I only need t to specify payoffs for k product amounts to implement any outcome
- In equilibrium, given some t, types $\theta_1, \ldots, \theta_k$ choose amounts $x_1, \ldots, x_k$ respectively, so we can design t_2 that pays $t_2(x_i) = t(x_i)$ and $t_2(x) = 0$ otherwise: t_2 implements the same outcome
- So in the 2 type case, we only need to choose two pairs (x_1, t_1), (x_2, t_2) such that type 1 wants to choose x_1 and 2 chooses x_2
- Another of those reformulations that are mathematically equivalent but make the problem more tractable

- Types θ_1, θ_2: $Pr(\theta_1) = p$, $Pr(\theta_2) = 1 - p$
- Cost functions $c_1(x)$, $c_2(x)$
- B chooses $\{(x_1, t_1), (x_2, t_2)\}$ to solve:

$$
\begin{aligned}
&\max p\,(v(x_1) - t_1) + (1-p)\,(v(x_2) - t_2) \\
&\text{s.t. } t_1 - c_1(x_1) \geq t_2 - c_1(x_2) && \text{(IC1)} \\
&\qquad t_2 - c_2(x_2) \geq t_1 - c_2(x_1) && \text{(IC2)} \\
&\qquad t_1 - c_1(x_1) \geq 0 && \text{(IR1)} \\
&\qquad t_2 - c_2(x_2) \geq 0 && \text{(IR2)}
\end{aligned}
$$

- Note: one weird thing about this setup is both types have the same outside option
- Rarely true in reality
- Note 2: the IC conditions are analogous to requiring tangency in the continuous case
- But here "tangency" is not meaningful because there are only 2 options
- Note 3: there may be solutions where we decide to exclude the low type altogether and just offer one pair (x_2, t_2), but we will come back to that later

- General intuition: in the optimal solution, 1's IR constraint will bind but not his IC, and 2's IC constraint will bind but not his IR
- Why?
- Since 2 has lower cost for any x, if 1's IR constraint holds, 2's must hold with slack (could at worst produce x_1 and make positive profit)
- $t_2 - c_2(x_2) \geq t_1 - c_2(x_1) > t_1 - c_1(x_1) \geq 0$
- Hence 2's IR never binds
- If 1's IR did not bind, B could lower both t_1 and t_2 by the same amount and make more money
- Hence 1's IR binds

- Since 2 has lower marginal cost and $x_2 > x_1$, it can't be that IC1 and IC2 both bind
- If IC1 binds, 1 is indifferent between x_1 and x_2, but then 2 strictly prefers x_2, hence IC2 does not bind
- If IC2 binds, 2 is indifferent, hence 1 strictly prefers x_1, hence IC1 does not bind
- Whenever IC2 does not bind, B can improve by lowering t_2 a little:
  - 2 still chooses x_2
  - 1 chooses x_1 even more strongly and his IR is unaffected
  - 2's IR is not violated if change is small enough since it wasn't binding
- Hence in optimal solution IC2 must bind, hence IC1 does not bind

- So B first chooses a point on 1's zero-profit curve, i.e., B chooses x_1 and $t_1 = c_1(x_1)$
- And then moves up 2's cost curve up to some point, i.e., B chooses x_2 and $t_2 = t_1 - c_2(x_1) + c_2(x_2)$
- So how to choose x_1, x_2?
- x_2 can just be picked as first-best!
- Whatever x_1 is, changing x_2 does not affect 1's incentives, just how much 2 produces and how much B pays 2
- So can just choose x_2 such that $c_2'(x_2) = v'(x_2)$ (first-best)

- What about x_1?
- Picking the first-best x_1 is not good: the more I increase x_1, not only do I have to pay 1 more, but also have to pay 2 more at the same x_2 to satisfy his IC
- For the same reason, x_1 higher than FB is also bad, and optimal x_1 is below FB
- The FOC is: $p = c_1'(x_1) - (1-p)c_2'(x_1) > pc_1'(x_1)$

- If $p < c_1'(x_1) - (1 - p)c_2'(x_1)$ even for small x_1, then may want to choose $x_1 = 0$ (price 1 out of the market)
- p does not affect x_2, but it affects x_1
- The lower p is, the lower x_1 is

- Main tension in this model is between desire to produce at the efficient level (choose x_1, x_2 equal to FB levels) and B's desire to limit type 2's rent
- Have to screw over type 1 to reduce type 2's temptation
- If p is low, lowering x_1 has low efficiency cost (low type is unlikely anyway) but big rent reduction (B pays less to the likely high type)
- Vice versa for high p

- How to derive the FOC: the problem is reduced to

$$\max p\,(v(x_1) - t_1) + (1-p)\,(v(x_2) - t_2)$$
$$\text{s.t. } t_2 - c_2(x_2) = t_1 - c_2(x_1) \qquad \text{(IC2)}$$
$$t_1 - c_1(x_1) = 0 \qquad \text{(IR1)}$$

- Or equivalently

$$\max p\,(v(x_1) - c_1(x_1)) + (1-p)\,(v(x_2) - c_2(x_2) - c_1(x_1) + c_2(x_1))$$
$$\Longrightarrow p(v'(x_1) - c_1'(x_1)) + (1-p)(-c_1'(x_1) + c_2'(x_1)) = 0$$
$$(1-p)(v'(x_2) - c_2'(x_2)) = 0$$

- Reminder: we were solving the screening problem, which we had reduced to:

$$\max p\,(v(x_1) - c_1(x_1)) + (1-p)\,(v(x_2) - c_2(x_2) - c_1(x_1) + c_2(x_1))$$
$$(\text{s.t. } x_2 \geq x_1)$$

- But the condition $x_2 \geq x_1$ does not bind so we can ignore it
- We get the FOCs:

$$p(v'(x_1) - c_1'(x_1)) + (1-p)(-c_1'(x_1) + c_2'(x_1)) = 0$$
$$(1-p)(v'(x_2) - c_2'(x_2)) = 0$$

- From the second FOC, $v'(x_2) = c_2'(x_2)$, so $x_2 = x_2^{FB}$, the first-best value
- Here "first-best" means the value that maximizes the total surplus of the principal and agent
- And also the value that would result from the optimal contract *if the agent were known to be type* 2
- From the first FOC,

$$v'(x_1) - c_1'(x_1) = \frac{1-p}{p}(c_1'(x_1) - c_2'(x_1)) > 0,$$

- so $x_1^* < x_1^{FB}$

- Hence the principal designs the menu $\{(x_1, t_1), (x_2, t_2)\}$ so that type 1 underproduces in equilibrium
- Again, this is to make it cheaper to prevent type 2's temptation to fake being type 1
- In particular, $x_2^* = x_2^{FB} > x_1^{FB} > x_1^*$
- If p is high, there is less distortion in x_1 so x_1^* goes up
- If p is low enough, can go all the way to $x_1 = 0$ (type 1 is shut out of the market)

Virtual Cost Function

- An alternative way to think about the problem of choosing x_1
- We can define

$$\tilde{c}(x_1) \equiv c_1(x_1) + \frac{1-p}{p} \Delta c(x_1)$$

- Then the choice of x_1 made in the screening mechanism is actually the FB choice, for a hypothetical agent that had this (higher) cost function
- The cost function captures both the real cost of 1 producing more x, and the cost of having to pay type 2 more as a result of increasing x_1

m-state case

- Suppose I have types $\theta_1, \ldots, \theta_m$
- Cost functions $c_1, \ldots, c_m$ such that $c_i'(x) > c_j'(x)$ for all $i < j$ and any x (higher types have lower marginal cost)
- Probabilities $p_1, \ldots, p_m$ adding up to 1
- How to design the mechanism?

- As before, we need to define at most m points: $(t_1, x_1), \dots, (t_m, x_m)$
- Could be fewer if I want to shut out some types, but not more (can just drop options from the contract which no one picks in equilibrium anyway)
- Now there are m IR constraints: $\mathrm{IR}_1, \dots, \mathrm{IR}_m$
- How many IC constraints? For each type k, need one IC constraint for each $i \neq k$, saying k prefers picking (t_k, x_k) to (t_i, x_i)
- So $k(k-1)$ IC constraints: $\mathrm{IC}_{k1}, \dots, \mathrm{IC}_{k(k-1)}, \mathrm{IC}_{k(k+1)}, \dots, \mathrm{IC}_{kn}$

- Which ones bind?
- We can show (with similar arguments to the 2-state case) that:
  - Only IR_1 binds (higher types have lower cost so necessarily positive profits)
  - Only $IC_{k(k-1)}$ binds for each $k = 2, \ldots, n$
- Lowest type who is not priced out is left indifferent about entering
- Each type is indifferent about not mimicking the next type with higher cost
- (But strictly does not want to mimic others)

- This gives the right amount of conditions: *given* some values of $x_1, \ldots, x_m$, the conditions uniquely pin down $t_1, \ldots, t_m$
- From IR_1, we know $t_1 = c_1(x_1)$: pins down t_1
- From IC_{21}, we know that $t_2 - c_2(x_2) = t_1 - c_2(x_1)$: pins down t_2
- And so on
- Finding the optimal $x_1, \ldots, x_m$ still requires solving for some FOCs
- (Side note: choosing t_i's with this algorithm allows us to implement *any* sequence $x_1, \ldots, x_m$ we want, as long as it's non-decreasing, but some are better for the principal than others)

- $x_m^* = x_m^{FB}$, but for $i < m$ we will have $x_i^* < x_i^{FB}$
- As before, increasing x for low types forces principal to pay all higher types more (by the same amount)
- Hence distortion is worst for the lowest i's (highest cost types)

Continuous Case

- Suppose now we have a continuum of types $\theta \in [0, 1]$
- θ distributed according to a continuous cdf F, with density f
- (Could deal with atoms in distribution; holes in the support are more annoying)
- Suppose $c_{x\theta} < 0$, $c(0, \theta) = 0$ for all θ, and (hence) $c_\theta < 0$
- Higher types have lower marginal cost, hence lower cost

Now principal solves:

$$\max_{x(\cdot),t(\cdot)} \int_0^1 (x(\theta) - t(\theta))\, dF(\theta)$$

$$\text{s.t. } t(\theta) - c(x(\theta), \theta) \geq t(\theta') - c(x(\theta'), \theta) \ \ \forall\, \theta,\ \theta' \qquad (\mathrm{IC}_{\theta,\theta'})$$

$$t(\theta) - c(x(\theta), \theta) \geq 0 \ \ \forall\, \theta \qquad (\mathrm{IR}_{\theta})$$

- Define $\Pi(\tilde{\theta}, \theta) \equiv t(\tilde{\theta}) - c(x(\tilde{\theta}), \theta))$
- This is the profit θ gets from pretending to be $\tilde{\theta}$
- Define $V(\theta) \equiv \Pi(\theta, \theta)$
- This is type θ's equilibrium payoff
- Then the IC conditions can be rewritten as $V(\theta) \geq \Pi(\tilde{\theta}, \theta)$ for all $\theta, \tilde{\theta}$

- What do our conditions imply about $V(\theta)$?
- Since it's the value function of an optimization problem, we can use the envelope theorem:

$$V'(\theta) = \frac{d\Pi(\theta,\theta)}{d\theta} = \frac{\partial\Pi(\theta_1,\theta_2)}{\partial\theta_2}|_{(\theta,\theta)} = -c_\theta(x(\theta),\theta) > 0$$

- Note: $V(\theta)$ a priori doesn't have to be differentiable, as it is endogenous: the principal could pick a non-smooth x or t
- But we know c_θ is well-defined by assumption
- There are versions of the envelope theorem for non-differentiable functions, which guarantee we can use it without knowing ex ante that V is differentiable
- But too complicated for this class, so just assume functions are differentiable

- Now we can integrate $V'(\theta)$:

$$V(\theta) = \Pi(0,0) - \int_0^\theta c_\theta(x(\tilde{\theta}),\tilde{\theta})d\tilde{\theta}$$

- Since $V(\theta) = t(\theta) - c(x(\theta),\theta)$,

$$t(\theta) = \Pi(0,0) + c(x(\theta),\theta) - \int_0^\theta c_\theta(x(\tilde{\theta}),\tilde{\theta})d\tilde{\theta}$$

- This has a similar flavor to the finite types case: given some $x(\theta)$, we can pin down $t(\theta)$
- But it is *not* logically equivalent!
- In the finite case, given $x_1, \ldots, x_m$, there were many $t_1, \ldots, t_m$ that could be paired with them that would implement production x_1 for θ_1, ..., x_m for θ_m
- The uniqueness of the t_i followed from making some IR and IC conditions bind, to achieve optimality for the principal
- (You could design other t_i schedules such that no IR or ICs would bind, and which would also implement the same x_i's, but they would give some agent types free money)

- On the other hand, in the continuous case, the conditions which uniquely pin down $V(\theta)$ and $t(\theta)$ (up to $\Pi(0,0)$) follow *exclusively* from the assumption that picking the schedule $x(\theta)$ is optimal (i.e., incentive compatible) *for the agent*
- We have not yet exploited in any way the assumption that we're trying to achieve optimality for the principal!
- The only way optimality for the principal will show up, in terms of conditions on t, is that we should set $\Pi(0,0) = 0$ (no free money for lowest type)
- But we still have to find the optimal schedule $x(\theta)$

- The problem

$$\max_{x(\cdot),t(\cdot)} \int_0^1 (x(\theta) - t(\theta))\, dF(\theta)$$

- now becomes

$$\max_{x(\cdot)} \int_0^1 \left(x(\theta) - c(x(\theta), \theta) + \int_0^\theta c_\theta(x(\tilde{\theta}), \tilde{\theta}) \right) dF(\theta)$$

- Subject only to the condition that $x(\theta)$ is non-decreasing

- Changing the order of integration, we can rewrite this as

$$\max_{x(\cdot)} \int_0^1 (x(\theta) - \tilde{c}(x(\theta),\theta)) f(\theta) d\theta$$

- where

$$\tilde{c}(x(\theta),\theta) \equiv c(x,\theta) - c_\theta(x,\theta)\frac{1-F(\theta)}{f(\theta)}$$

- Deriving with respect to each $x(\theta)$, we get the FOC:

$$c_x(x,\theta) - c_{x\theta}(x,\theta)\frac{1-F(\theta)}{f(\theta)} = 1 \;\; \forall \; \theta$$

- This gives us an equation in $x(\theta)$ which generally pins down $x(\theta)$
- As before, the solution satisfies that $x^*(\theta) < x^{FB}(\theta)$ for $\theta < 1$, and $x^*(1) = x^{FB}(1)$

- One question left: is the solution $x^*(\theta)$ pinned down by this condition necessarily non-decreasing?
- Not always!
- It turns out that, when the solution to this system of FOCs is non-monotonic, you can find the “real” solution by smoothing out the decreasing parts
- Surprisingly, this does not affect the optimal value of $x(\theta)$ outside of the regions we’re smoothing out
- This is because of the agent’s quasilinear utilities: changing x, and t, for some θ affects required payoffs for all θ’s equally, so does not affect local incentives