

Lecture Slides - Part 2

Bengt Holmstrom

MIT

February 2, 2016.

Moral Hazard

- Related to adverse selection, but simpler
- A basic problem in principal-agent relationships: how does the principal incentivize the agent to take the right action, given noisy information?

- Setup: two players, P and A
- Technology: $x(e, \theta) = e + \theta$
- x is the outcome, e is the agent's effort, θ is the “state of nature” or measurement error
- Information: P observes x , but not e or θ . A observes e and x (and hence can infer θ)
- x is “verifiable/contractible”: this means it's mutually observed, and moreover the players could show the result to a court, hence can write a contract based on it
- e is private to A
- (Note: you could have a variable that is mutually observed but nonverifiable!)

- Preferences: P is risk neutral: $u_P(x, s) = x - s$, where s is payment to the agent
- A is risk-averse: $u_A(s, e) = u(s) - c(e)$, where u is concave and c is convex
- (Note: could also have $u_A = u[s - c(e)]$ if the cost of effort was monetary)
- The solution to the problem is a *contract* $s(x)$, which specifies payment based on outcome

Timing:

- First, P offers a contract $s(x)$ to A ; A can accept or reject, leading to outside option payoffs (note P has all the bargaining power)
- Second, if A accepts, A chooses e
- Third, Nature chooses θ
- Fourth, x is revealed to both agents and P pays $s(x)$ to A (note there is no commitment problem)

- Note on moral hazard and adverse selection:
- Old paradigm was: MH is the case with hidden action, no hidden info; AS is the opposite
- We now understand it better
- The crucial distinction: MH arises when info is *symmetric at the time of contracting*, AS arises when info is *asymmetric at the time of contracting*
- E.g.: if A had a choice to exert effort before meeting P , and this is private and affects our problem, it is AS with hidden action

Possible formulations:

- State space (as we have done it): think of the outcome as $x(e, \theta)$, where $\theta \sim G$ for some distribution. This is explicit about there being a state
- Conditional distribution (pioneered by Mirrlees): think of the outcome as having a conditional distribution $F(x|e)$. This is equivalent to the first version, if we take
$$F(x_0|e) = P(x(e, \theta) \leq x_0|e) = P(\theta \leq x_e(x_0)^{-1}|e) = G(x_e(x_0)^{-1})$$
- Equivalently, think that the agent is just directly choosing a distribution

- Again, although mathematically equivalent, the second formulation makes you more naturally think of enlightening examples
- E.g., this case: two actions (two distributions), $e_L < e_H$
- Costs $c_L = 0 < c_H$
- $F_H \succ F_L$ in the FOSD sense

- Note: the way we have framed it, principal offers contract $s(\cdot)$, then A chooses e_s , generating F_e and some expected utilities
- However, more natural to solve it this way: imagine the principal chooses a preferred action e^* by A , then designs a contract that guarantees A will choose e^* (i.e., e^* is *incentive-compatible* (IC) given $s(\cdot)$)
- Formally, P solves:

$$\max_{s(\cdot), e} \int_x (x - s(x)) dF(x|e)$$

$$\text{s.t. } \int_x u(s(x)) dF(x|e) - c(e) \geq \int_x u(s(x)) dF(x|e') - c(e') \quad \forall e' \quad (\text{IC})$$

$$\int_x u(s(x)) dF(x|e) - c(e) \geq \bar{u}_A \quad (\text{IR})$$

- The second constraint (individual rationality, IR) assumes A has some outside option paying \bar{u}_A , so P 's contract must pay at least that much
- We can solve this with a two-step approach:
 - First, for a given e , what $s(x)$ is optimal to implement it? Let $B(e)$ be P 's utility under the best possible contract that implements e
 - (Note: an optimal contract never has randomized s because P is risk-neutral and A is risk-averse)
 - Second: what e is optimal? Find $\max_e B(e)$.

- Going back to our problem with 2 actions: if we want to implement e_L , just take $s_L(x)$ constant and equal to s_0 , such that $u(s_0) = \bar{u}_A$
- To implement e_H , the contract must satisfy IC:

$$\int u(s(x))dF_H(x) - c_H \geq \int u(s(x))dF_L(x)$$
- Using Lagrange multipliers, we have to solve

$$\begin{aligned} \max_{s(x)} \quad & \int (x - s(x))dF_H \\ & + \mu \left[\int u(s(x))dF_H - c_H - \int u(s(x))dF_L \right] \\ & + \lambda \left[\int u(s(x))dF_H - c_H - \bar{u}_A \right] \end{aligned}$$

- This looks ugly, but since we are maximizing over all contracts $s(x)$, we can effectively maximize point by point (pick the best $s(x)$ for each x)

This gives the FOC:

$$-1 \times f_H(x) + \mu u'(s(x))f_H(x) - \mu u'(s(x))f_L(x) + \lambda u'(s(x))f_H(x) = 0$$

(note we derive with respect to s , not x)

This translates to:

$$\frac{1}{u'(s(x))} = \lambda + \mu \left(1 - \frac{f_L(x)}{f_H(x)} \right)$$

Lecture 4

- We were studying a moral hazard problem with two actions:

$$\max_{s(\cdot)} \int x (x - s(x)) f_H(x) dx$$

$$\text{s.t.} \quad \int x u(s(x)) f_H(x) dx - c_H \geq \int x u(s(x)) f_L(x) dx \quad (\text{IC})$$

$$\int x u(s(x)) f_H(x) dx - c_H \geq \bar{u}_A \quad (\text{IR})$$

- if we wanted to implement e_H
- Note: f_H, f_L don't have to be proper densities for this to work (they can have point masses) but they must have the same support
- Idea: why? If they had different support, you could make perfect inference from some outcomes

- We want to pick $s(x)$, i.e., a value of s for each x
- Can do a change of variables and directly pick $u(s(x))$: let $\phi(x) = u(s(x))$, so that $u^{-1}(\phi(x)) = s(x)$, then we can alternatively solve

$$\max_{\phi(\cdot)} \int_x (x - u^{-1}(\phi(x))) f_H(x) dx$$

$$\text{s.t.} \quad \int_x \phi(x) f_H(x) dx - c_H \geq \int_x \phi(x) f_L(x) dx \quad (\text{IC})$$

$$\int_x \phi(x) f_H(x) dx - c_H \geq \bar{u}_A \quad (\text{IR})$$

- Conceptually, this is a simpler problem because the constraints are now linear in our choice variables

- We obtain a Lagrangian with multipliers μ for the IC constraint and λ for the IR constraint
- Note: if IR is not binding, P can always do better by reducing $\phi(x)$ uniformly for all x (does not affect IR), hence IR is always binding and $\lambda > 0$
- Note: built into our statement that P maximizes his utility subject to IC and IR, is the assumption that if the optimal e is implemented with a program that leaves A indifferent with another action e' , he will pick whichever one is better for P
- But we could frame it the other way, and maximize A 's utility subject to P 's outside option or a market condition, and we would get the same set of results. Both programs return points on the possibility frontier of (EU_A, EU_P)

Let's interpret the resulting FOC:

$$\frac{1}{u'(s(x))} = \lambda + \mu \left(1 - \frac{f_L(x)}{f_H(x)} \right)$$

Some observations:

- $s_H(x)$ is decreasing in $l(x) = \frac{f_L(x)}{f_H(x)}$ (likelihood ratio)
- $s_H(x)$ is increasing in x iff MLRP: if x is higher, $1 - l(x)$ is higher by MLRP, so $u'(s(x))$ is lower, so $s(x)$ is higher by concavity of U
- Note: solution is *as if* P is making inferences about A 's choice (pay more for signals that are more likely under high effort). But paradoxically, in equilibrium, there is actually no inference because A 's action is chosen with certainty, so P knows it
- If P is risk neutral, then the solution is the same if P 's payoff is some $\pi(x) - s(x)$ instead of $x - s(x)$. It just matters that x is a signal of effort, not that it is P 's profits.

- Note: if there was no incentive problem, the optimal solution would just involve $\frac{1}{u'(s(x))} = \lambda$, as in a risk sharing problem (since P is risk-neutral)
- This formula tells us to what extent the risk-sharing incentive is distorted by the need to incentivize A
- Moral: tension in this model is between incentives and mitigating the cost of A 's risk aversion

- When is additional information valuable? E.g., suppose we also observe y . When can we design a contract $s(x, y) \succ s(x)$?
- The solution for info (x, y) is given by the FOC

$$\frac{1}{u'(s(x, y))} = \lambda + \mu \left(1 - \frac{f_L(x, y)}{f_H(x, y)} \right)$$

- We can show that $s(x, y) \succ s(x)$ iff $\frac{f_L(x, y)}{f_H(x, y)} \neq I(x)$: y should always be included in the contract in some form unless it adds no info about e , given x
- Equivalently, y is not useful when $f_i(x, y) = g(x)h(x, y)$, so that $h(x, y) = P(y|x)$ is independent of i
- Note that if y is very risky, the contract probably won't use it much, but some positive use is still optimal (because small changes at the margin add little risk)
- Randomization is always suboptimal, since for each x there is a unique optimum (due to the concavity of U)

Now look at the case with continuous action:

$$\max_{s(\cdot), e} \int_x (x - s(x))f(x|e)dx$$

$$\text{s.t.} \quad \int_x u(s(x))f(x|e)dx - c(e) \geq \int_x u(s(x))f(x|e')dx - c(e') \quad \forall e' \quad (\text{IC})$$

$$\int_x u(s(x))f(x|e)dx - c(e) \geq \bar{u}_A \quad (\text{IR})$$

Here, there are infinitely many constraints so it's a harder problem

- However, if A 's problem is smooth and concave, we can perhaps replace all the IC's with a single FOC:

$$u(s(x))f_e(x|e)dx - c'(e) = 0$$

- And from there, we get the FOC for P 's optimal contract problem:

$$\frac{1}{u'(s(x))} = \lambda + \mu \frac{f_e(x|e)}{f(x|e)}$$

- Note that this is a local condition (based only on comparison with e 's very close to the chosen one)

- This is the right solution *if* the local approach is valid
- But what if it's not? Suppose that $f(x|e) = N(e, \sigma^2)$
- Then $\frac{f_e(x|e)}{f(x|e)} \propto \frac{x-e}{\sigma^2}$
- Plugging that into our FOC, we get a contradiction: low enough x 's would get negative marginal utility
- What's going on? There is no solution satisfying the FOC!

- What's the true solution? Since the normal distribution offers potentially so much info (there are x 's with very extreme likelihood ratios), P can incentivize A by simply punishing very hard for very bad (but unlikely) outcomes
- This can be done at vanishingly low cost, so we can approach the first best, but not reach it

Lecture 5

- Reminder: we were studying the Mirrlees example
- In general, for a continuous action problem we had the FOC

$$\frac{1}{u'(s(x))} = \lambda + \mu \frac{f_e}{f}$$

- For a normal f , this becomes

$$\frac{1}{u'(s(x))} = \lambda + \mu \frac{(x - e)}{\sigma^2}$$

- It can be shown that $\mu > 0$, so the expression becomes negative for x small enough: a contradiction!

- This implies there is no proper solution (any solution must satisfy the FOC)
- Why? The informativeness of normal signals allows us to get arbitrarily close to costless punishments (but not reach it)

- For comparison: compare a case where $x = e + \epsilon$, $e = e_L$ or e_H , and $\epsilon \in U[0, 1]$
- Here we can implement the first best at no cost: if $x \in [e_L, e_H)$, then agent definitely chose e_L
- So we can design a contract where $x \in [e_L, e_H)$ is punished very hard, otherwise agent gets constant income
- On the equilibrium path, perfect insurance (because the agent can avoid the bad outcome with prob. 1)
- The key to this example: moving support allows infinitely informative signals (infinite likelihood ratio)

- In the normal case, by punishing only very low x 's very hard, we get a similar result
- The likelihood ratio for very low x 's becomes so extreme that it's almost like the case with moving support
- But can't fully reach first best (the limit of these contracts approaching the first best is degenerate, and has no punishment)

First-best cases

When can we implement the efficient e without suffering any cost due to A 's risk aversion?

- If A is risk neutral (then we can just choose $s(x) = x - \beta$ to implement the optimum)
- If e is verifiable (then we can choose $s = c_0$ low unless $e = e^*$)
- If $x(\theta, e)$ and θ is verifiable (then we can back out e and we are back in the previous case)
- With moving support

What's wrong?

- Should we be happy with this model?
- Agent has a fairly simple, restricted choice: just choose one-dimensional level of effort
- If effort is made over many days, A just chooses the mean
- P has infinite-dimensional control

- We might think that giving A a small choice space simplifies the problem, but that may not be true
- Giving A more options can force P to design a contract that is less manipulable, and hence simpler
- (For an extreme example of that, see Carroll (2015))
- E.g.: maximizing a smooth function over an interval is easier than over a large finite set
- Maximizing a function over a plane is often easier than over some curve embedded in the plane
- One-dimensional choice for A constrains him to a small family of distributions (e.g. can't take a convex combination of available distributions)

- Imagine that $x = e + \theta$ where $\theta = \epsilon + \gamma$
- A chooses e , then observes ϵ before choosing γ at some cost (last-minute gaming of the outcome)
- Even if cost of γ is high, so A constrained to very small manipulation, this breaks contracts with discontinuities (agent will game near the discontinuities)

- This happens in real life with target-based bonuses:
- If A just chooses one e , a contract that pays out iff $x \geq x^*$ (you get a bonus if you meet the quota) may be optimal, for the right distribution f
- But if A is making sales every day over a month, he will want to work hard as the end of the month approaches if he is close to the target, but shirk once he reaches it (or give up if too far)
- Clearly suboptimal now, and the culprit is A 's richer choice set
- Intuitively, we expect linear contracts to avoid this problem

- Holmstrom and Milgrom (1987) formalizes this idea for CARA utility: $u(m) = -e^{-r(m-c(e))}$, and normal noise
- With this utility, there is no income effect (agent's marginal incentive to work does not depend on accumulated wealth)
- Then, in a problem where agent chooses effort N times and sees previous outcomes before next choice, linear contract gives the right incentives
- Can also do in continuous time (accumulated product is Brownian motion, A chooses drift at every t)

LEN Model

- Let's study the Linear Exponential Normal model
- (Holmstrom and Milgrom (1987) show that linear contracts are optimal in this case; this is hard, but finding the best linear contract is easy)
- $x = e + \epsilon$, $\epsilon \sim N(0, \sigma^2)$ so $x \sim N(e, \sigma^2)$
- $u(s(x) - c(e)) = -e^{-r(s(x) - c(e))}$
- $s(x) = \alpha x + \beta$
- Can consider a single outcome (given linear contract, the model is effectively separable across outcomes)

Certainty Equivalent

- Given a random variable X (e.g. the money payout of a lottery), the certainty equivalent $CE(X)$ is a certain payoff that would leave the agent indifferent compared to getting X
- Formally, $CE(X)$ is such that $u(CE(X)) = E(u(X))$
- This depends on A 's attitude towards risk (more risk averse means lower CE for same lottery)
- For a normal distribution and CARA utility, get mean-variance decomposition: $CE_A(s) = E(s(x)) - \frac{1}{2}r\text{Var}(s(x)) - c(e)$

- Meanwhile $CE_P(s) = E(x - s(x)) = (1 - \alpha)e - \beta$
- $CE_A = \alpha e + \beta - \frac{1}{2}r\alpha^2\sigma^2 - c(e)$
- Hence total surplus is $TS = e - \frac{1}{2}r\alpha^2\sigma^2 - c(e)$
- First best effort maximizes TS , satisfies $1 = c'(e)$

- In practice, for a given α , A maximizes CE_A and chooses e such that $c'(e) = \alpha$
- If c is convex, this gives $e_\alpha < e_{FB}$
- Then we can choose α to maximize TS given e_α , i.e.,

$$TS(\alpha) = e_\alpha - c(e_\alpha) - \frac{1}{2}r\alpha^2\sigma^2$$
- Find solution: $\alpha^* = \frac{1}{1+r\sigma^2 c''}$

- Corollary: $\alpha^* < 1$, decreasing in r (risk aversion) and σ^2 (noise of signal)
- This model offers much more natural predictions; we would trust it more to answer new questions
- But note we can only do this because we have a proof that, under some conditions (richer A choices), Mirrlees-style contracts are bad and linear contracts are optimal, and we understand the difference between the settings
- Just saying “I don’t like the optimal solution to my original problem, so let’s just assume linear contracts” would not be kosher

- Aside remark: these models assume increasing cost of effort for simplicity
- But we can obtain the same results in models where cost of effort is U-shaped (agents intrinsically want to work up to some point): if we need them to work more than that, at the margin it is the same problem
- Too many papers claiming these models are irrelevant because they assume agents don't like working

Lecture 6

- Reminder: we provided a justification for looking at linear contracts (Holmstrom and Milgrom 1987)
- We then found the optimal linear contract, characterized by $\alpha = c'(e(\alpha))$ and $\alpha = \frac{1}{1+r\sigma^2 c''}$

- Where does this come from? $\alpha = c'(e(\alpha))$ just comes from A 's IC condition
- Then, deriving with respect to α , $1 = c''(e(\alpha))e'(\alpha)$
- $\frac{de}{da} = \frac{1}{c''}$ is how much more A works if I increase the commission a little
- Substituting these into the derivative of TS , we find α^*

Multi-tasking

- So far we studied a one-activity model where the cost of providing A with incentives is burdening A with risk
- In multi-tasking models, A has several activities
- P may want to give more incentives for activities that he can monitor well (less noise means A suffers less from risk)
- Incentives, even for perfectly monitored activities, may be distorted if cost function is not separable
- Idea: if I can monitor job 1 well and 2 badly, I want to give more incentives for 1
- But not to the efficient level: else 1 will crowd out too much 2

- Suppose A can invest effort into increasing quality and quantity
- $x_1 = e_1 + \epsilon_1$ is quality, $x_2 = e_2$ is quantity
- $B(e_1, e_2) = p_1 e_1 + p_2 e_2$ is P 's payoff from (e_1, e_2)
- $C(e_1, e_2)$ is A 's cost of (e_1, e_2)
- (e_1, e_2) may interact in the cost function, e.g., if it is $C(e_1 + e_2)$ with C convex, doing more e_1 increases the marginal cost of e_2 and vice versa)

- P designs a contract $s(x_1, x_2) = \alpha x_1 + \alpha x_2 + \beta$
- (Again, if we assume exponential utility, ..., then linear contracts are optimal)
- So P solves:

$$\max_{\alpha_1, \alpha_2, \beta} B(e_1, e_2) - C(e_1, e_2) - \frac{1}{2} r \alpha_1^2 \sigma_1^2$$

$$\text{s.t. } \alpha_1 = \frac{\partial c}{\partial e_1} \quad (\text{IC1})$$

$$\alpha_2 = \frac{\partial c}{\partial e_2} \quad (\text{IC2})$$

- Note: we are not exactly solving P 's problem, but instead maximizing total surplus
- This is OK because the two problems are equivalent
- We can also drop the IR condition because optimal α 's are independent of the preferred pie distribution, and adjusting β is how we divide the pie (thanks to exponential utility)

With more detail: P solves

$$\max_{\alpha_1, \alpha_2, \beta} B(\mathbf{e}_1, \mathbf{e}_2) - E(s(x_1, x_2)) \equiv (p_1 - \alpha_1)\mathbf{e}_1 + (p_2 - \alpha_2)\mathbf{e}_2 - \beta$$

$$\text{s.t. } \max_{\mathbf{e}_1, \mathbf{e}_2} \{E(s(x_1, x_2)) - \frac{1}{2}r\text{Var}(s) - C(\mathbf{e}_1, \mathbf{e}_2)\}$$

$$\equiv \{\alpha_1\mathbf{e}_1 + \alpha_2\mathbf{e}_2 - C(\mathbf{e}_1, \mathbf{e}_2) + \beta - \frac{1}{2}r\alpha_1^2\sigma_1^2\}$$

- To solve, go back to the total surplus problem and derive with respect to α_1 and α_2 , getting FOCs:

$$\frac{\partial TS}{\alpha_1} = (p_1 - C_1) \frac{\partial e_1}{\alpha_1} + (p_2 - C_2) \frac{\partial e_2}{\alpha_1} - r\alpha_1\sigma_1^2 = 0$$

$$\frac{\partial TS}{\alpha_2} = (p_1 - C_1) \frac{\partial e_1}{\alpha_2} + (p_2 - C_2) \frac{\partial e_2}{\alpha_2} = 0$$

- Then plug in the IC conditions and their versions derived with respect to α_j

- From here, we get the optimal α 's:

$$\alpha_1^* = \frac{p_1}{1 + r\sigma_1^2 C_{11}}$$

$$\alpha_2^* = p_2 - r\sigma_1^2 C_{12}\alpha_1^2$$

- What does this mean?
- *Both* α_j are lower than their efficient effort-inducing levels, p_1 and p_2
- But $\alpha_2 < p_2$ only when $C_{12} > 0$, i.e., when doing 2 increases the cost of 1
- Idea: as in the single activity case, you want to choose $\alpha_1 < p_1$ to reduce risk
- The choice of α_1 is unaffected by α_2 except for the fact that e_2 may affect C_{11} (make the cost of 1 steeper)
- But, because e_2 crowds out e_1 through the cost function (if $C_{12} > 0$), want to choose lower α_2 when this interaction is strong

- At the margin, if α_2 is close to (close to p_2), reducing it a little reduces e_2
- Impact on payoffs generated by e_2 ($p_2 - \frac{\partial C}{\partial e_2}$) is second-order since we are close to the optimum
- But reduction in e_2 increases e_1 , which is a first-order benefit

- Moral of the story: *low-powered* incentives are good when activities are badly monitored
- When A has multiple activities that vary in the quality of monitoring, P should make incentives weakest for the poorly monitored activities, but also make *all* incentives low so poorly monitored jobs don't get crowded out by the well monitored
- In some examples, could even want *no* incentives (fixed wage)
- This idea has big real-world implications

Multi-task Lab

- Up to now, we were in a model where P has as many levers in the contract as A has jobs
- So P can choose how to incentivize each activity (choose both α_1 and α_2), and the tension is between incentives and insurance against risk
- But in many real-life jobs, A has a lot of activities: vector (e_1, \dots, e_n)
- And P only has access to a few performance measures
- So A always has opportunities to game the measures using certain e_i 's that are well-rewarded

- E.g.: A is a teacher, can teach 100 different topics
- P has two measures: class grades and standardized exam
- The moral will be: in this world, we want low-powered incentives *even if A is risk-neutral*
- Beyond some point, increasing incentives will just lead to A doing too much of tasks undesired by P
- E.g. “teaching to the test”

- E.g.: imagine that $B(e) = e$ is P 's activity (e.g. coding)
- $(e_k)_{k=1,\dots,K}$ are private activities (e.g. using the computer to chat, watch videos)
- A 's cost is $C(e + \sum_k e_k) - \sum_k v_k(e_k)$ (A enjoys private activities, but they may increase the cost of coding by distracting A)
- P observes $x = e + \epsilon$ and pays $s(x) = \alpha x + \beta$
- Suppose also that P can *exclude* some tasks (e.g. block Youtube on the company network)
- How to design the optimal contract $(\alpha, \beta, \text{exclusions})$?

- FOCs:

$$\alpha = C'(e + \sum_k e_k)$$

$$v'_k(e_k) = C' = \alpha$$

- Note: this means that, given α , total effort $e + \sum_k e_k$ is constant!
- Exclusions simply transfer effort from an excluded task e_k to e

- So should P just exclude all private activities? No
- Exclude e_k iff it generates less total surplus than transferring to the job
- Exclude k iff $v_k(e_k) < e_k$
- (Remember P can always adjust pie through β , so efficient to make A happy if at low cost: pay you less and let you use Youtube)
- Then we still have to choose the optimal α , but note that choice of what tasks to exclude is conditional on α
- If the v_k are concave, then e_k will be declining in α , so $\frac{v(e_k)}{e_k}$ increasing in α , so task exclusions will decrease in α
- (If incentives are strong, A will mostly ignore Youtube because reward for work is high; if they are weak, A will shirk a lot unless Youtube is blocked)

- Before moving on to the next topic: a reminder of why we can study moral hazard problems as problems of maximizing total surplus
- P is solving: $\max EU_P$ subject to (IC) and $EU_A \geq \bar{u}_A$ (IR) or equivalently $CE_A \geq u^{-1}(\bar{u}_A)$
- Under exponential utility, if a point (CE_P, CE_A) is achievable, then $(CE_P + \beta, CE_A - \beta)$ is achievable too: just transfer β no matter what the state, or in other words make a new contract $s_2(x) = s(x) - \beta$ for all x
- Hence, if there is a contract maximizing $CE_P + CE_A$, then it is optimal to implement essentially that contract no matter the desired distribution of the pie, and then just change the β to achieve different distributions

- Without exponential utility, the idea is less clear because certainty equivalents are not as handy, and the optimal contract changes depending on desired distribution due to income effects
- But it is still true that we can essentially solve by maximizing TS subject to IC and IR

MIT OpenCourseWare
<https://ocw.mit.edu>

14.124 Microeconomic Theory IV
Spring 2017

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>.