ECO421: Moral hazard

Marcin Pęski

March 9, 2020

Plan

- Moral hazard:
- two agents with misaligned interests,
- one of them chooses an action that matters for both of them,
- the other can offer a contract to change the incentives of the first agent.

Introduction

Moral hazard: Insurance

- "*Circumstance that increases the probability of occurrence of a loss, or a larger than normal loss, because of a change in an insurance policy applicant's behavior after the issuance of policy.* "

Introduction

Moral hazard: Financial markets

- "*Moral hazard is a situation in which one party gets involved in a risky event knowing that it is protected against the risk and the other party will incur the cost. It arises when both the parties have incomplete information about each other.* "
- "*In a financial market, there is a risk that the borrower might engage in activities that are undesirable from the lender's point of view because they make him less likely to pay back a loan.*"
- more specific applications of the general principle from the previous slide

Introduction

Moral hazard: Job

- Two players:
  - Ann: principal, boss
  - Bob: agent, worker
- the principal wants the agent to take certain action
- the interests of the principal and the agent are misaligned
  - so, the agent may not necessarily take the action that is optimal for the principal.
- To align incentives, the principal and the agent may agree on a *binding contract*.
  - the contract has to provide the agent's incentives,
  - and it must be preferable to no contract.

Plan

How to pay a worker

Example

A principal hires a worker.

- The principal designs contract $w(.)$
- The worker decides whether to accept the contract. If not, the worker receives outside option U_0. The principal recives 0.
- After accepting the contract, the worker chooses effort level e. The worker pays the cost of effort $c(e)$.
- **The effort is observable by the principal** $\Longrightarrow$ The worker receives wage $w(e)$.
- The principal's payoff is equal to $u(e) - w(e)$, where $u(e)$ are the profits from the worker with effort e, and $w(e)$ is the paid out wage.

How to pay a worker

Fully observable effort

- We assume that
  - $u()$ is increasing (i.e., $u' > 0$) - principal likes worker's effort,
  - $c()$ is increasing ($c' > 0$) - worker does not like the effort.
- Outside option $U_0 \geq 0$,
  - alternative employment, cost of time, etc.
  - worker will accept the contract only if he can get out of it at least U_0.

How to pay a worker

Fully observable effort

- Any contract $w(.)$ chosen by the principal induces some effort e_0,
  - it is convenient to consider e_0 as chosen by the principal as well and
  - add IC constraints to the problem.
- Contract = (wage schedule $w(.)$ and effort e_0).

How to pay a worker

Fully observable effort

- Principal's problem: choose contract $w(e)$ and effort level $e_0$so to maximize principal's payoffs, st.

$$\max_{w(.),e_0} u(e_0) - w(e_0) \text{ st.}$$

How to pay a worker

Fully observable effort: IR constraint

- Principal's problem: choose contract $w(e)$ and effort level $e_0$so to maximize principal's payoffs, st.

$$\max_{w(.),e_0} u(e_0) - w(e_0) \text{ st.}$$

$$\text{IR:} w(e_0) - c(e_0) \geq U_0$$

  - worker gets at least his outside option,

How to pay a worker

Fully observable effort: IR constraint

- Principal's problem: choose contract $w(e)$ and effort level e_0 so to maximize principal's payoffs, st.

$$\max_{w(.),e_0} u(e_0) - w(e_0) \text{ st.}$$

$$\text{IR:} w(e_0) - c(e_0) \geq U_0,$$

$$\text{IC:} e_0 \in \arg\max_e w(e) - c(e).$$

  - worker gets at least his outside option,
  - given the contract, the worker's payoff is maximized by choosing e_0.

How to pay a worker

Fully observable effort: IR constraint

- Principal's problem: choose contract $w(e)$ and effort level $e_0$so to maximize principal's payoffs, st.

$$\max_{w(.),e_0} u(e_0) - w(e_0) \text{ st.}$$

$$\text{IR:} w(e_0) - c(e_0) \geq U_0,$$

$$\text{IC:} w(e_0) - c(e_0) \geq w(e) - c(e) \text{ for each } e,$$

  - worker gets at least his outside option,
  - given the contract, the worker's payoff is maximized by choosing e_0.

How to pay a worker

Fully observable effort

- Two questions:
- Step 1. How to find the cheapest way to provide the agent incentives to do e_0?
- Step 2. What is the optimal effort e_0, i.e, the effort that maximizes the principal's profits?

How to pay a worker

Fully observable effort: Step 1

- Step 1. How to find the cheapest way to provide the agent incentives to do e_0?
- Easy here. Forcing contract:
  - Take
$$w(e) = \begin{cases} U_0 + c(e_0), & \text{if } e = e_0, \\ 0, & \text{otherwise.} \end{cases}$$
  - IC holds: for each $e \neq e^*$,
$$w(e) - c(e) = -c(e) \leq 0 \leq U_0 = w(e_0) - c(e_0).$$
  - IR holds:
$$w(e_0) - c(e_0) = U_0 + c(e_0) - c(e_0) = U_0 \geq U_0,$$
  - Also, because any IR contract must pay the worker at least $U_0 + c(e_0)$, the forcing contract is the cheapest way to implement effort e_0.

Education

Separating equilibrium

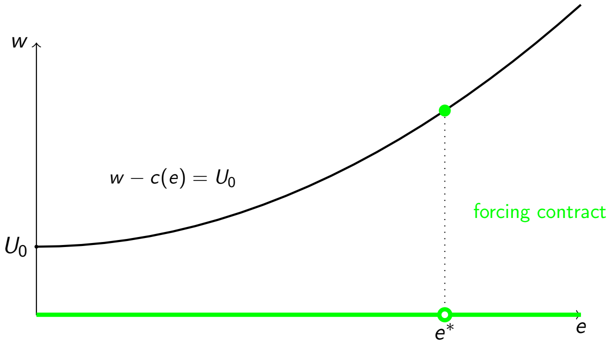

How to pay a worker

Fully observable effort

- Step 2: What is the optimal effort e_0, i.e, the effort that maximizes the principal's profits?
- Define the *social surplus* as

$$\Pi(e) = \underbrace{(u(e) - w(e))}_{\text{principal's payoff}} + \underbrace{(w(e) - c(e) - U_0)}_{\text{agent's payoff}}$$
$$= u(e) - c(e) - U_0$$

  - agent's payoff takes into account the opportunity costs U_0.

How to pay a worker

Fully observable effort

- It the contract $\{w(.), e_0\}$ is IR, $w(e_0) - c(e_0) - U_0 \geq 0$, and

$$\begin{aligned} u(e_0) - w(e_0) &\leq u(e_0) - w(e_0) + [w(e_0) - c(e_0) - U_0]. \\ &= u(e_0) - c(e_0) - U_0 \\ &=: \Pi^*(e_0). \end{aligned}$$

- Thus, the principal's payoff in any IR contract is not higher than the social surplus at the effort level.
- In other words, $\Pi(e_0)$ is an upper bound on the principal's payoff.

How to pay a worker

Fully observable effort

- Let e^* be the socially optimal (*first-best*) choice of effort:

$$e^* \in \arg\max_e \Pi(e) = \arg\max_e (u(e) - c(e) - U_0).$$

  - $\Pi^*(e^*)$ is an upper bound on the principal's profits.
- The guess is if e^* maximizes the upper bound on profits, maybe it also maximizes the proftis.
- Indeed, we will show that it is the case and the upper bound profits can be achieved:

How to pay a worker

Fully observable effort

- Consider a forcing contract to implement e^*:

$$w(e) = \begin{cases} U_0 + c(e^*), & \text{if } e = e^*, \\ 0, & \text{otherwise.} \end{cases}$$

  - IR and IC holds (we checked it already),
  - The principal's profits are equal to the upper bound:

$$u(e^*) - w(e^*) = u(e^*) - c(e^*) - U_0 = \Pi^*(e^*) \geq \max_e \Pi^*(e).$$

  - Hence, the forcing contract with effort level e^* achieves the highest possible profits!
- We found an *optimal contract*.

How to pay a worker

Fully observable effort

- Optimal contract -> general procedure (the other way):
  - Step 1: Find the cheapest way of making sure that the agent chooses any effort e_0.
  - Step 2: Find (principal)-optimal level of effort e_0.
- Here, we found that the principal optimal level of effort is also socially optimal.
  - Not always the case.

How to pay a worker

Fully observable effort

- Step 1: There are many other contracts that achieve the same payoff:
  - forcing contract:

$$w(e) = \begin{cases} U_0 + c(e^*), & \text{if } e = e^*, \\ 0, & \text{otherwise,} \end{cases}$$

  - threshold contract:

$$w(e) = \begin{cases} U_0 + c(e^*), & \text{if } e \geq e^*, \\ 0, & \text{otherwise.} \end{cases}$$

  - linear contract (assuming that c is convex, $c'' > 0$):

$$w(e) = U_0 + c(e^*)'(e - e^*).$$

    - Check FOC of the agent problem to verify the IC constraint.
    - Piece rate

Education

Separating equilibrium

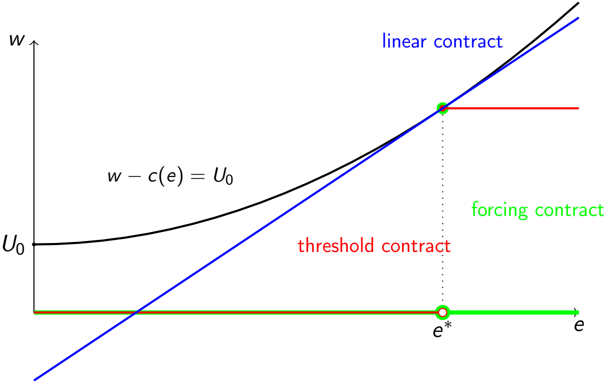

How to pay a worker
Fully observable effort

- Step 2: As long as the cost function is strictly convex and utility is strictly concave, the optimal level of effort is unique.

How to pay a worker

Fully unobservable effort

- Next, suppose the effort is completely unobservable.
- Contract w cannot depend on wage $\Longrightarrow$ flat wage!
- The agent will choose

$$e_0 \in \arg\max w - c(e).$$

  If $c' > 0$, $e_0 = 0$.
  - $e_0 = 0$ is the IC contraint in th unobservable case.
- IR is satisfied if $w \geq U_0 + c(0)$.

How to pay a worker

Fully unobservable effort

- If
$$u(0) > U_0 + c(0),$$
then the Principal offers wage $w_0 = U_0 + c(0)$ if $u(0) > w_0$.
- Otherwise, no contract.

How to pay a worker

Summary

- If the effort is observable, the first-best effort (i.e., the effort that maximizes social surplus) is a solution to the principal's problem.
- If the effort is completely not observable, only flat wage possible, and 0 effort.
- Next, partially observable effort.

Plan

Example

Two agents: principal and agent (inventor).

- Inventor may achieve breakthrough. The probability of the breakthrough $q(e)$ depends on the inventor's effort.
- The effort is costly $c(e)$, and not observable by the principal.
- The principal observes whether there is a breakthrough $i = 0, 1$, receives payoff u_i and pays out wages w_i that depend on the existence of breakthrough $i = 0, 1$.

- We assume $u_1 > u_0$.
- Also, $q' > 0, q'' \leq 0$ and $c', c'' > 0$.
- Special case $c(e) = \frac{1}{2}e^2, q(e) = e$.

How to pay an inventor
Partially observable effort

Example

Contract (w_0, w_1).

- w_0 is the *salary*,
- $\Delta w = w_1 - w_0$ is a "*breakthrough bonus*".

- Question: how to pay the inventor? Salary vs bonus?

How to pay an inventor

Inventor's (agent's) problem

- Agent's problem:

$$\max_{e} q(e) w_1 + (1 - q(e)) w_0 - c(e).$$

- FOC:

$$q'(e)(w_1 - w_0) - c'(e) = 0.$$

- Let $\Delta w = w_1 - w_0$ denote *breakthrough bonus.* Then,

$$q' \Delta w = c'.$$

  - let $e_A(\Delta w)$ be the solution to the above equation,
  - optimal effort depends only on the bonus.

How to pay an inventor

Principal's problem

- Principal's problem:

$$\max_{w_0, w_1, e} q(e)(u_1 - w_1) + (1 - q(e))(u_0 - w_0) \text{ st.}$$

- Principal's problem:

$$\max_{w_0, w_1, e} q(e)(u_1 - w_1) + (1 - q(e))(u_0 - w_0) \text{ st.}$$

$$IR : q(e) w_1 + (1 - q(e)) w_0 - c(e) \geq U_0,$$

How to pay an inventor

Principal's problem

- Principal's problem:Also, $q' > 0$ and $c', c'' > 0$.

$$\max_{w_0, w_1, e} q(e)(u_1 - w_1) + (1 - q(e))(u_0 - w_0) \text{ st.}$$

$$IR : q(e) w_1 + (1 - q(e)) w_0 - c(e) \geq U_0,$$

$$IC : e \in \arg\max_{e'} q(e') w_1 + (1 - q(e')) w_0 - c(e').$$

How to pay an inventor

Principal's problem

- Using the agent's solution, we can restate the principal's problem as

$$\max_{w_0, \Delta w} q\left(e_A\left(\Delta w\right)\right)\left(u_1 - w_1\right) + \left(1 - q\left(e_A\left(\Delta w\right)\right)\right)\left(u_0 - w_0\right) \text{ st.}$$

$$IR: q\left(e_A\left(\Delta w\right)\right) w_1 + \left(1 - q\left(e_A\left(\Delta w\right)\right)\right) w_0 - c\left(e_A\left(\Delta w\right)\right) \geq U_0,$$

  - IC condition disappears because we use the solution to the agent's problem,
  - similarly, e disappears from the optimization.

How to pay an inventor

Principal's problem

- Let's take $\Delta u = u_1 - u_0$.
- Using the fact that

$$qA + (1-q)B = q(A-B) + B,$$

we can restate the principal's problem further as

$$\max_{w_0, \Delta w} q(e_A(\Delta w))(\Delta u - \Delta w) + (u_0 - w_0) \text{ st.}$$

$$IR : q(e_A(\Delta w))\Delta w + w_0 - c(e_A(\Delta w)) \geq U_0.$$

How to pay an inventor

Principal's problem

$$\max_{\Delta w} q\left(e_A\left(\Delta w\right)\right)\left(\Delta u - \Delta w\right) + \left(u_0 - w_0\right) \text{ st.}$$
$$IR : q\left(e_A\left(\Delta w\right)\right)\Delta w + w_0 - c\left(e_A\left(\Delta w\right)\right) \geq U_0.$$

- Because lowering w_0 increases principal's payoff (and it does not affect the optimal effort), we can assume that IR binds, or:

$$q\left(e_A\left(\Delta w\right)\right)\Delta w + w_0 - c\left(e_A\left(\Delta w\right)\right) = U_0, \text{ or}$$
$$U_0 + c\left(e_A\left(\Delta w\right)\right) - q\left(e_A\left(\Delta w\right)\right)\Delta w = w_0$$

How to pay an inventor

Principal's problem

- When we substitue w_0 into the principal's objective, we obtain

$$\max_{\Delta w} q\left(e_A\left(\Delta w\right)\right)\left(\Delta u - \Delta w\right) + u_0 + q\left(e_A\left(\Delta w\right)\right)\Delta w - c\left(e_A\left(\Delta w\right.\right.$$

$$= \max_{\Delta w} q\left(e_A\left(\Delta w\right)\right)\Delta u + u_0 - c\left(e_A\left(\Delta w\right)\right) - U_0.$$

and principal's problem has no more constraints.

How to pay an inventor

Optimal contract

$$\max_{\Delta w} q\left(e_A\left(\Delta w\right)\right)\Delta u + u_0 - c\left(e_A\left(\Delta w\right)\right) - U_0$$

- FOC:

$$q' e_A' \Delta u - c' e_A' = 0.$$

We can divide by the derivative e_A', and recall that $1 = \left(q' \Delta w\right)/c'$, to get

$$\frac{\Delta u}{\Delta w} = 1,$$

which implies that

$$\Delta w = \Delta u, \text{ and}$$
$$e^* = e_A\left(\Delta u\right).$$

- Optimal contract!

- Summary: How did we solve the principal's problem.
- Step 1a: Solve the agent's problem $e_A(.)$
  - substitute the solution into the principal's problem and IR constraints,
  - takes care of the IC.
- Step 1b: show that the IR constraint is binding
  - substitute w_0 into the principal's problem,
  - takes care of the IR constraint.
- Step 2: Solve unconstrained problem.

How to pay an inventor

Optimal contract

- Let's look at the optimal contract more carefully:

$$\begin{aligned} e^* &= e_A(\Delta u). \\ w_0 &= U_0 + c(e^*) - q(e^*)\Delta u, \\ w_1 &= w_0 + \Delta u. \end{aligned}$$

- The decision is optimal and it maximizes the social welfare:

$$\begin{aligned} W(e) =& q(e)u_1 + (1 - q(e))u_0 - c(e) \\ =& q(e)\Delta u + u_0 - c(e). \end{aligned}$$

  - Indeed, $e_A(\Delta u) = \arg\max W(e)$.

How to pay an inventor

Optimal contract

- Worker's utility is U_0.
- Principal's utility:
  - if there is no breakthrough: $u_0 - w_0$, if there is a breakthrough

$$u_1 - w_1 = u_1 - w_0 - (u_1 - u_0) = u_0 - w_0.$$

  - The same!
- Also, notice that

$$\begin{aligned} u_0 - w_0 &= u_1 - w_1 \\ &= q(e^*)\Delta u + u_0 - c(e^*) - U_0 \\ &= W(e^*) - U_0. \end{aligned}$$

  - Hence, principal gets the value of the social welfare minus the inventor's outside option.

How to pay an inventor

Optimal contract

- The contract is equivalent to *selling the firm*:
  - principal sells the firm to the inventor for $W(e^*) - U_0$.
  - the worker receives payoff u_i depending on the breakthrough,
  - chooses socially optimal action,
- Why selling the firm solves the moral hazard problem!
  - agent (inventor) works for himself,
  - no incentive problem,
- Magical solution.

How to pay an inventor

Optimal contract

- Can we always sell the firm to solve the moral hazard problem?
- Unfortunately no.
- Agent is not always able to realize the same benefit from the breakthrough as the principal (not able to market the invention, etc)
  - maybe both agents and principal are necessary to generate the payoffs
  - which one of the should be the principal?
- Agent does not always have enough money to buy the firm
  - cash- or credit-constraints,
  - in our problem, i may mean that $w_0 \geq 0$, or that the worker cannot have negative wages,
  - added constraint.

Plan

How to pay CEO

- Alternative to selling the firm is to allign the agent incentives with the principal by transferring ownership of the share of the firm.
    - stocks,
    - high-powered incentives
- For instance, the ownership of $x\%$ of the firm.
    - Is $x\%$ enough?
    - It depends on the CEO incentives (effort, empire building, fame) vs the size of the firm
    - If the value of the firm is really large to swamp anything else that the CEO cares about (power, cost of effort, etc), yes.

Plan

How to give grades

- Grading is difficult.
- In the first-best world, the teacher could trust that students would learn for the love of learning, for the usefulness of the knowledge, and it would choose the effort level rationally.
- We are not in the first best world.
- The teacher uses grading policy to design incentives for students to put effort.
- The problem is that grades mix effort and ability.
- Simple model

How to give grades

Example

Each student has ability θ. The ability has cdf $\theta \sim F(.)$.

- A student chooses effort level $a \in \{a_l, a_h\}$. The effort is unobservable.
- The effort a leads to an exam score $q(\theta) = a\theta$, where θ is the student's ability. Grades are observable.
- The high effort costs $c > 0$.
  Student receives utility $u > c$ if she passes the course.
- Teacher designs pass/fail policy.
  Teacher wants to maximize the number of students who choose high effort.

How to give grades

- We consider a various combinations:
  - both student and teacher observe θ,
  - neither teacher nor student observes θ,
  - only the student observes θ (mix of adverse selection and moral hazard).

How to give grades

Example

Version I: Ability θ is observed by both student and teacher.

- If θ is jointly observed, the pass/fail policy may depend on ability and the grade.
- The teacher can pass student θ if her grade is higher than $a_h\theta$ and fail otherwise.
  - because $u > c$, each student prefers to exert high effort and pass.

How to give grades

Example

Version II. Ability is not observed by anybody.

The teacher designs the pass/fail policy based on the grades only

$$c\left(q|q^*\right) = \begin{cases} 0, & \text{if } q < q^*, \\ 1, & \text{if } q \geq q^*. \end{cases}$$

How to give grades

Ability is not observed by anybody

- Student with ability θ and effort a passes if if $a\theta \geq q^*$, or if $\theta \geq \frac{q^*}{a}$.
- Suppose that neither teacher nor student observe θ. The probability that the student with effort a passes is equal to

$$1 - F\left(\frac{q^*}{a}\right).$$

- The expected utility from high effort is

$$\left(1 - F\left(\frac{q^*}{a_h}\right)\right) u - c$$

and low effort is

$$\left(1 - F\left(\frac{q^*}{a_l}\right)\right) u.$$

How to give grades

Ability is not observed by anybody

- Student will choose high effort if

$$\left(1 - F\left(\frac{q^*}{a_h}\right)\right) u - c > \left(1 - F\left(\frac{q^*}{a_l}\right)\right) u,$$

  or

$$F\left(q^* \frac{1}{a_l}\right) - F\left(q^* \frac{1}{a_h}\right) > \frac{c}{u}. \qquad (1)$$

  - notice that $\frac{c}{u} \in (0, 1)$ because $c < u$,
  - maybe everybody, maybe nobody chooses high effort.
- Any q^* such that (1) holds is optimal,
  - it depends on the distribution F whether such q^* exists.

How to give grades

Example

Version III. Ability is observed by student, but not the teacher.

The teacher designs the pass/fail policy based on the grades only

$$c\left(q|q^{*}\right)=\begin{cases}0, & \text{if } q<q^{*},\\ 1, & \text{if } q\geq q^{*}.\end{cases}$$

How to give grades

Ability is observed by student, but not by the teacher

- The student will choose high effort if
  - it will allow him to pass, and
  - low effort will lead to the failure.

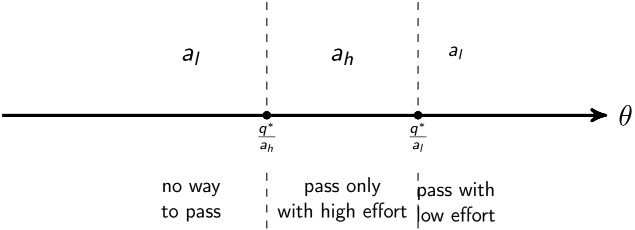

How to give grades

Ability is observed by student, but not by the teacher

- Grading policy q^* implies that fraction

$$F\left(\frac{1}{a_l}q^*\right) - F\left(\frac{1}{a_h}q^*\right)$$

  of students puts an effort.
- For any grading policy, there are abilities
  - who can never pass, even with high effort, and
  - who can pass even with low effort,
  - Any such abilities will choose low effort.
- In general, there is no way to ensure that all students put an effort.

How to give grades

Ability is observed by student, but not by the teacher

Example

Uniform distribution $F(\theta) = \theta$ for $\theta \in [0, 1]$.

- In such a case,

$$F\left(\frac{1}{a_l}q^*\right) - F\left(\frac{1}{a_h}q^*\right)$$

$$= \begin{cases} 0, & q^* < 0 \\ \left(\frac{1}{a_l} - \frac{1}{a_h}\right) q^*, & q^* \in [0, a_L] \\ 1 - \frac{1}{a_h}q^* & q^* \in [a_l, a_h], \\ 0, & q^* > a_h. \end{cases}$$

How to give grades

Ability is observed by student, but not by the teacher

- Optimal choice

$$q^*_{opt} = a_L.$$

- All the students who pass the test choose the high effort
  - more precisely, all but $\theta = 1$ who is the only ability type who can pass with low effort.
- This property depend son the distribution.
- If the ability density is decreasing, it might be optimal to choose low threshold that motivates lower ability student, but

How to give grades

Ability is observed by student, but not by the teacher

Example

Consider density

$$f(x) = c - \frac{c^2}{2}x \text{ for } x \in \left[0, \frac{2}{c}\right].$$

- Let $\gamma = \frac{a_l}{a_h} < 1$.
- Find the optimal grading policy as a function of γ and c.

How to give grades

Ability is observed by student, but not by the teacher

- TBA

Plan

How to pay a VP

- Next, a small variation on the above problem.

Example

A VP in a large company responsible for a part of the market chooses effort level $a \in \{a_l, a_h\}$. The high effort costs $c > 0$. The effort is unobservable by the CEO.

- The effort a leads to the output equal to $q(\theta) = a\theta$, where θ is the "market condition".
- θ is observed by the VP, but not by the CEO.

- The previous discussion suggests that there is no way to design the wages as functions of q to ensure that the department chief always chooses high effort.
- For very bad or very good market conditions, the low effort is optimal.
- But, sometimes we can do better.

Example

There are two departments headed by two VPs. The CEO has a budget u to motivate each VP.
The outputs in the two departments are affected by *the same market condition* θ.

- Instead of working with two VPs separately, we combine the problems together.
- *Relative performance*.
- Why would it help?
- We are learning something from the behavior of one agent, that may help us to provide better incentives for the other one.

How to pay a VP

Relative performance

- Consider the following scheme.
- If one VP has a higher output,
    - she gets bonus $2u$,
    - and the other VP gets nothing.
- If the two VPs have the same output, both of them get bonus u.
- In this scheme, the value of the bonus depends on the performance of the other VP
    - that makes it a game between VPs!

How to pay a VP

Relative performance

- We are going to show that high effort is strictly dominant for each θ!
- Indeed, if the other VP chooses low effort, then
  - $a_i = a_h$ leads to payoff of $2u - c$, vs
  - $a_i = a_l$ leads to payoff of u.
  - because $u > c$, choosing a_h is better.
- If the other VP chooses high effort, then
  - $a_i = a_h$ leads to payoff of $u - c$, vs
  - $a_i = a_l$ leads to payoff of 0.
  - because $u > c$, choosing a_h is, again, better.

How to pay a VP

Relative performance

- High effort is chosen for any θ!
- Thus, a big improvement on any scheme that treated the VP problems separately.
- Thus, using the information about the other VP is beneficial.
- We assumed that the VPs can observe θ.
  - but the argument actually does not depend on it.
  - High effort is dominant even if the market condition cannot be observed by the VPs.

How to pay to a VP

Tournaments

- The above scheme is a special case of a *tournament*.
  - one way to pay VP is to reward the better one with a job of CEO.
  - let the best one win.
- Tournaments are the most common, oldest relative performance incentive schemes.
- Tournaments are great when there is a common shock,
  - the relative performance allows to isolate the effect of the common shock from the incentives.
  - The might be observed by the participants (adverse selection), but it does not have to be.
- Tournaments are great way to provide incentives for scientific research.
  - patent races!

How to pay to a VP

Tournaments

- Examples of famous tournaments:
  - Longitude Act, 1714
    - series of rewards: £10,000 (worth over 1.33 million in 2016[4]) for anyone who could find a practical way of determining longitude at sea to an accuracy of not greater than one degree of longitude (equates to 60 nautical miles (110 km; 69 mi) at the equator). The reward was to be increased to £15,000 if the accuracy was not greater than 40 minutes, and further enhanced to £20,000 if the accuracy was not greater than half a degree.[5]
    - John Harrison, inventor of chronometer
  - Scottish book, 1930ies, Lvov
    - For problem 153, closely related to Stefan Banach's "basis problem", Stanisław Mazur offered the prize of a live goose.
    - This problem was solved only in 1972 by Per Enflo.
    - The winner got the goose.
  - Millennium Prize Problems are seven problems in mathematics that were stated by the Clay Mathematics Institute in 2000,
  - Google Lunar X prize.

Plan

How to pay a teacher

- Huge increase in use of standardized tests to evaluate teaching performance for the last 20 years or so.
  - precise measures,
  - can isolate impact of individual teachers,
  - seems like a perfect tool to measure teaching quality.
- However, such tests do not measure all that is important to teaching:
  - arithmetic ability, but not
  - general problem solving skills, creativity, etc.
- Using the test outcomes to incentivize the teachers
  - may cause them to teach for the test,
  - but not focus on skills that may be much more important in the long-run.

How to pay a teacher

Example

(Multitasking) The teacher can allocate effort into two tasks.

- A teacher picks effort allocation into two tasks $e_1, e_2 \geq 0$. Only the first effort e_1 is observed.
- The teacher (i.e., agent) and the school district (principal) have different preferences over the tasks:

$$\pi_{teacher}(e_1, e_2; w) = w + a(e_1 + e_2) - \frac{1}{2}(e_1 + e_2)^2, \text{where}$$

$$\pi_{district}(e_1, e_2, w) = -w + e_1 + \beta e_2,$$

  - $\beta > 0$ is the relative value of the unobserved activity for the school district,
  - $a > 0$ is the value of "honest work" for the teacher, and
  - $\frac{1}{2}(e_1 + e_2)^2$ is the cost of effort.

- We assume $\beta > 1$ so the unobserved activity is more important.

How to pay a teacher

Linear contracts

- Suppose that the school district offers a linear contract:

$$w\left(e_1\right) = w_0 + \alpha e_1.$$

  - w_0 is a (fixed) salary,
  - $\alpha \geq 0$ is a performance bonus.
- Simple type of wage contract, well-approximates real life,
- It turns out, in this example, without loss of generality.
  - we can replicate any behavior under general contracts by linear contracts.

How to pay a teacher

District's problem

- The district's problem:

$$\max_{w_0,\alpha,e_1^*,e_2^*} -w_0 - \alpha e_1^* + e_1^* + \beta e_2^*$$

$$\text{st.}(IC): (e_1^*, e_2^*) \in \arg\max_{e_1,e_2} \pi_{teacher}(e_1, e_2; w_0, \alpha),$$

$$(IR): \pi(e_1^*, e_2^*; w_0, \alpha) \geq U_0.$$

- Solution procedure:
  - we separately consider $\alpha = 0$ and $\alpha > 0$,
  - we solve the teacher's problem - takes care of IC,
  - we use the IR to figure out w_0,
  - we solve the unconstrained disctrict problem,
  - compare payoffs between $\alpha = 0$ and $\alpha > 0$.

How to pay a teacher

$\alpha = 0$, Teacher's problem

- First, $\alpha = 0$.
- The teacher's problem

$$\max w_0 + a\left(e_1 + e_2\right) - \frac{1}{2}\left(e_1 + e_2\right)^2 .$$

- Optimal choice:

$$e_1^* + e_2^* = a.$$

How to pay a teacher

$\alpha = 0$, IRs constraint,

- IR constraint:

$$w_0 + a\left(e_1^* + e_2^*\right) - \frac{1}{2}\left(e_1^* + e_2^*\right)^2 \geq U_0,$$

or, using the solution to the teacher's problem,

$$w_0 + a^2 - \frac{1}{2}a^2 \geq U_0,$$

- Because the IR constraint is going to bind, we have

$$w_0 = U_0 - \frac{1}{2}a^2.$$

How to pay a teacher

$\alpha = 0$, principal's problem

- School's payoffs
$$-w_0 + e_1 + \beta e_2,$$
where $e_1, e_2 \geq 0$ and $e_1 + e_2 = a$.
- Hence, using the formula for wages,
$$\begin{aligned} & -U_0 + \frac{1}{2}a^2 + (a - e_2) + \beta e_2 \\ = & -U_0 + \frac{1}{2}a^2 + a + (\beta - 1)\, e_2, \end{aligned}$$
subject to $e_2 \in (0, a)$.
- If $\beta > 1$, the payoffs are maximized by $e_2 = a$ and equal to
$$-U_0 + \frac{1}{2}a^2 + a + (\beta - 1)\, a.$$

How to pay a teacher

$\alpha > 0$, teacher's problem

- Case $\alpha > 0$. The teacher's problem

$$\max w_0 + \alpha e_1 + a(e_1 + e_2) - \frac{1}{2}(e_1 + e_2)^2 .$$

- Optimal choice:
  - $e_2 = 0$: no effort put into the unobservable activity,
  - to see why, notice that if $e_2 > 0$, then the teacher can increase his payof by taking ${e_1}' = e_1 + e_2$ and ${e_2}' = 0$,
  - if $e_2 = 0$, $e_1$can be derived from the FOCs:

$$e_1^* = \alpha + a.$$

How to pay a teacher

$\alpha > 0$, IR constraint

- IR constraint:

$$w_0 + \alpha e_1^* + a\left(e_1^* + e_2^*\right) - \frac{1}{2}\left(e_1^* + e_2^*\right)^2 \geq U_0,$$

or using the teacher's solution

$$w_0 + \alpha\left(\alpha + a\right) + a\left(\alpha + a\right) - \frac{1}{2}\left(\alpha + a\right)^2 \geq U_0.$$

- Because the IR constraint is going to bind, we have

$$w_0 = U_0 - \frac{1}{2}\left(\alpha + a\right)^2.$$

How to pay a teacher

$\alpha > 0$, principal's problem

- Case $\alpha > 0$. Then, $e_1 = a + \alpha$, the school's payoffs are

$$\begin{aligned}\pi_{school} &= -w_0 - \alpha e_1^* + e_1^* + \beta e_2^* \\ &= -U_0 + \frac{1}{2}(a+\alpha)^2 - \alpha(\alpha + a) + a + \alpha.\end{aligned}$$

- FOC (wrt. α):

$$a + \alpha - \alpha - (a + \alpha) + 1 = 0,$$

or

$$\alpha = 1.$$

  - $\alpha = 1$ shifts the school's incentives onto teacher,
    - this is a familiar result (spelling firm).

How to pay a teacher

$\alpha > 0$, principal's problem

- School's payoff:

$$\pi^*_{school} = -U_0 + \frac{1}{2}(a+1)^2 - (a+1) + a + 1$$
$$= -U_0 + \frac{1}{2}(a+1)^2$$
$$= -U_0 + \frac{1}{2}a^2 + a + \frac{1}{2}.$$

How to pay a teacher

- We can compare

$$\alpha = 0: \quad -U_0 + \frac{1}{2}a^2 + a + (\beta - 1)\,a.$$
$$\alpha > 0: \quad -U_0 + \frac{1}{2}a^2 + a + \frac{1}{2}.$$

- Thus, if

$$\beta a > a + \frac{1}{2},$$

  the school's payoffs are higher if $\alpha = 0$.

How to pay a teacher

Lemma

Suppose that $(\beta - 1)\, a > \frac{1}{2}$. *Then, the optimal contract does not involve any incentives* $\alpha = 0$.

- If β is high, the payoff from the unobserved activity is high relative to the benefit of incentivizing the teacher.
- If a is high, the teacher has lots of internal motivation.
- In both cases, the optimal choice of the remuneration is a flat wage contract, without any incentive components.

How to pay a teacher

- Optimality of flat wage sheds some light on the puzzle:
  - moral hazard is prevalent,
  - but most of the wage contracts have no explicit incentive component.
- Trade-off between incentivizing observable activity, and making sure that unobservable tasks are also executed.
  - monitoring is difficult, in general,
  - there are tasks that are very difficult to monitor,
  - but they are valuable to principal,
  - incentivizing only observable task may result in shifting effort away from unobservable tasks, which is bad for the principal.

How to pay a teacher

- Examples:
  - sales: maximizing the number of sales vs. customer satisfaction and good will,
  - services: time in which service is accomplished vs attention to detail and care,
  - production: speed vs quality,
  - production: the output vs taking care of the machines and/or safety.

Plan

Efficiency wages

- Even if a wages do not explicitly depend on performance, there are other ways of providing incentives:
  - keeping vs firing the worker,
  - promotion.
- If the fired worker can immediately find an equivalent job, he is not scared of being fired,
  - so, no incentives.
- Moral hazard model leads to a simple theory of unemployment.

Efficiency wages

Economy

- 1 competitive firm with production function $f(L)$,
  - L number of hired workers,
  - $f'(L) = w$ is the wage
  - workers and firms separate for random reasons with prob. b,
- N workers in the economy,
  - $u = \frac{N-L}{N}$ unemployment rate.
- Probability of finding a job

$$
\begin{aligned}
p(u) =& \frac{\#\text{job openings}}{\#\text{separations} \ + \ \#\text{unemployed}} \\
=& \frac{bL}{bL + N - L} = \frac{bL/N}{bL/N + (N-L)/N} \\
=& \frac{b(1-u)}{b(1-u)+u} = 1 - \frac{u}{b + u(1-b)}.
\end{aligned}
$$

is decreasing in u.

Efficiency wages

Moral hazard

- Each worker can either shirk or work hard
  - $c > 0$ cost of working hard
- If the worker shirks, it does not generate any output
  - is caught with prob. q and fired,
- Firms will hire the worker only if they expect her to work hard.
- Firms and workers separate for other reasons (randomly) with probability $b \in (0, 1)$.
- Discounting $\beta < 1$.

Efficiency wages

Moral hazard

- Let
  - V_e - value of being employed,
  - V_u - value of being unemployed,
- Value of shirking

$$w + \beta \left[(1-b) \left(qV_u + (1-q) V_e \right) + bV_u \right],$$

Value of not shirking

$$w - c + \beta \left((1-b) V_e + bV_u \right)$$

Efficiency wages
Moral hazard

- Let
  - V_e - value of being employed,
  - V_u - value of being unemployed.
- Worker won't shirk if

$$\beta\left[(1-b)\left(qV_u+(1-q)V_e\right)+bV_u\right] \leq -c+\beta\left((1-b)V_e+bV_u\right),$$

or

$$\Delta := V_e - V_u \geq \frac{c}{(1-b)\beta q}.$$

  - value of being employed must be higher than the value of being unemployed to compensate for the effort,
  - Δ is the difference between the two values.

- Value of being employed (and working hard)

$$\begin{aligned} V_e &= w - c + \beta (1 - b) V_e + \beta b V_u \\ &= w - c + \beta V_e - \beta b \Delta \end{aligned}$$

or

$$(1 - \beta) V_e = w - c - \beta b \Delta.$$

Efficiency wages

Equilibrium

- Value of being unemployed:
  - the convention is that if a worker finds job, she immediately benefits from being employed,
  - otherwise, she waits for a whole preiod, and becomes unemployed again:

$$\begin{aligned} V_u &= p(u)\,V_e + \beta\left(1 - p(u)\right)V_u \\ &= \beta V_u + \beta p(u)\,\Delta + (1-\beta)\,p(u)\,V_e \\ &= \beta V_u + \beta p(u)\,\Delta + p(u)\left(w - c - \beta b\Delta\right) \\ &= \beta V_u + p(u)\,\beta\,(1-b)\,\Delta + p(u)\,(w-c)\,, \end{aligned}$$

  where we used the above expression for V_e.

Efficiency wages

Equilibrium

- After subtracting the two equations and some algebra, we get

$$\Delta\left(1-\beta+\beta b-p\beta\left(1-b\right)\right)=\left(1-p\left(u\right)\right)\left(w-c\right),$$

  and

$$\Delta=\frac{1-p\left(u\right)}{1-\beta\left(1-p\left(u\right)\right)\left(1-b\right)}\left(w-c\right).$$

- Putting the above into the incentive inequality,

$$\begin{aligned} w &\geq c+\frac{c}{\left(1-b\right)\beta q}\left(\frac{1-\beta\left(1-p\left(u\right)\right)\left(1-b\right)}{1-p\left(u\right)}\right). \\ &= c+\frac{c}{\beta q}\left(\frac{1}{\left(1-b\right)\left(1-p\left(u\right)\right)}-\beta\right). \end{aligned}$$

Efficiency wages

Equilibrium

- Incentive inequality

$$w \geq c + \frac{c}{\beta q}\left(\frac{1}{(1-b)(1-p(u))} - \beta\right).$$

  - wages have to be sufficiently high to stop the worker from shirking and risking being fired
  - if it is easy to find job, $p(u) \approx 1$, and the incentive wages go to ∞
  - hence, unemployment in equilibrium must exist.

Efficiency wages

Equilibrium

- Equilibrium unemployment is at intersection of
  - labor demand:
    $$f'\left(\left(1-u\right)N\right)=w,$$
  - labor supply
    $$w \geq c+\frac{c}{\beta q}\left(\frac{1}{1-p\left(u\right)}-\beta\left(1-b\right)\right).$$
- Full employment is impossible.
- There are workers who would be happy to work for wages below the market level, but they cannot find jobs.

Plan

Career incentives

- Risk of being fired and unemployment is not the only tool for the firm to provide incentives for the worker.
- One argument is that firms have other tools to provide incentives for the workers.
- For example,
  - promotion,
  - or progress-through-the-ranks.
- Increasing wage profiles are common.

Career incentives

- A worker who is exerting full effort, is kept and receives a higher wage in the next period
  - a worker may be paid below their productivity initially to pay up for future raise,
  - because she is paid more than her productivity later,
  - on average, the worker is paid its productivity.
- A worker who is caught shirking is fired and needs to start again.
  - this provides incentives to put an effort,
  - as long as the cost of effort is smaller than the future raise.
- This works without unemployment.
- One of the criticisms of the efficiency wage theory of unemployment.

Plan

Conclusions

What did we learn - concepts

- Principal agent problem
  - contract as a function of observables,
- Agent's problem:
  - IR constraint and incentive conditions,
- Principal's problem
  - maximize objective subject to the agent's problem.

Conclusions

What did we learn - Types of principal's agent problem

- Fully observable effort,
  - first-best allocation
- Observable output (but not effort)
  - selling the firm contract,
  - minimum wage-constraints,
- Moral hazard with adverse selection

Conclusions

What did we learn - Types of principal's agent problem

- Relative performance
  - tournaments,
- multitasking,
  - optimality of flat wages
- Unemployment as an incentive
  - efficiency wages,
- Multi-period moral hazard.

Conclusions

What did we learn - skills

- Write down principal agent problem,
- Show optimality of simple contracts in given situations
  - threshold contract (fully observable effort)
  - selling the firm (observable output)
  - flat wage (multitasking)
  - tournament (common shocks with two agents)
- Explain how unemployment and career prospects can be used to provide incentives.