Advanced Microeconomics III

Moral Hazard

Francisco Poggi

Introduction

- So far we focused on the **outside** of trading relationships.
  - Agent's problem was to math with a high-quality trading partner.
  - Once partnership is formed, tasks were trivial.

- Now we consider the **inside** of a trading relationship.
  - Transactions that are too complex to be completely specified.
  - How to write a contract that structures the relationship in the best possible way?

Principal-agent models

- To focus on the inside, we assume away adverse selection.

- We assume that the relationship faces a **moral hazard problem**:
  - One party ("agent") may take actions that are in her own interest rather than in the interest of the other party ("principal")
  - This actions are not observable to the principal (or at least not verifiable by courts)

Contracts

- Parties will try to write a contract that gives the agent the incentives to take the "correct" action.
- **Key idea**: rewards can be conditioned on variables that depend (maybe stochastically) on the agent's action.

- Examples:
  - Firm owner - manager: firm's profit.
  - Insurance firm - insurance taker: whether damage occurs.

Principal-Agent Relationship

- A firm owner (principal) wishes to hire a manager (agent) for a project.
- The manager (if hired) chooses some action $a \in A$ that is not observable to the owner.
  - Effort level.
  - Choice of risky project.
  - Level of care.
- The project yields a stochastic profit $\pi \in [\underline{\pi}, \bar{\pi}]$ that is verifiable.
- Conditional on the action, the distribution of profits has cdf F and density function:

$$f(\pi|a) > 0 \qquad \text{for all } \pi \in [\underline{\pi}, \bar{\pi}].$$

Example: Stochastic dominance

- Example $a \in \{e_L, e_H\}$.
- *$F(\cdot|e_H)$ strictly first-order stochastically dominates $F(\cdot|e_L)$.*

$$F(\pi|e_H) \leq F(\pi|e_L) \qquad \text{for all } \pi \in [\underline{\pi}, \bar{\pi}]$$

and strictly for some π.

- Notice that F strictly FOSD $G \Rightarrow E_F[\pi] > E_G[\pi]$.

$$\begin{aligned} E_F[\pi] &= \int_{\underline{\pi}}^{\bar{\pi}} \pi \cdot f(\pi)\ d\pi \\ &= \pi \cdot F(\pi)\Big|_{\underline{\pi}}^{\bar{\pi}} - \int_{\underline{\pi}}^{\bar{\pi}} 1 \cdot F(\pi)\ d\pi \\ &= \bar{\pi} - \int_{\underline{\pi}}^{\bar{\pi}} F(\pi)\ d\pi \end{aligned}$$

Preferences

- **Agent's preferences**:
  - Utility $u(w, a)$ depends on wage and action.
  - We assume that u is *additively separable*, i.e. there exist functions v and c such that:
  $$u(w, a) = v(w) - c(a)$$
  - $v' > 0$ and $v'' < 0$ guarantee risk aversion.
  - Reservation utility $\bar{u}$.

- **Principal's objective function**:
  - Risk-neutral: $\pi - w$.
  - Reservation utility $\bar{U}$.

Risk-aversion

- Why assume that agent is risk averse and principal risk neutral?
  1. If both are risk neutral with no limits on wealth, the problem becomes trivial.
  2. If both are risk averse the analysis is more complicated, but same general issues and results.
  3. A rationale is that the principal is welthy and is more diversified than the agent.

Overview

1. Verifiable action

2. Non-verifiable actions

Benchmark: Verifiable action

- **No moral hazard**: a can be stipulated in a contract.
  - The principal can basically "choose" the action.
  - action is not only observable, but also verifiable in court.

- A *contract* specifies:
  - an action $a \in \mathcal{A}$, and
  - a wage scheme $w : [\underline{\pi}, \bar{\pi}] \to \mathbb{R}$.

Benchmark: Verifiable action

- Suppose that a contract $(a, w(\cdot))$ is signed.
  - Principal's expected utility:

$$U = \int_{\underline{\pi}}^{\bar{\pi}} (\pi - w(\pi)) \cdot f(\pi|a)\ d\pi$$

  - Agent's expected utility:

$$u = \int_{\underline{\pi}}^{\bar{\pi}} v(w(\pi)) \cdot f(\pi|a)\ d\pi - c(a)$$

- A *feasible utility pair* (u, U) is a pair of expected utilities that can be obtained.
- The *Pareto frontier* is the set of feasible utility pairs that are not Pareto dominated by any other feasible utility pair.

The Pareto frontier

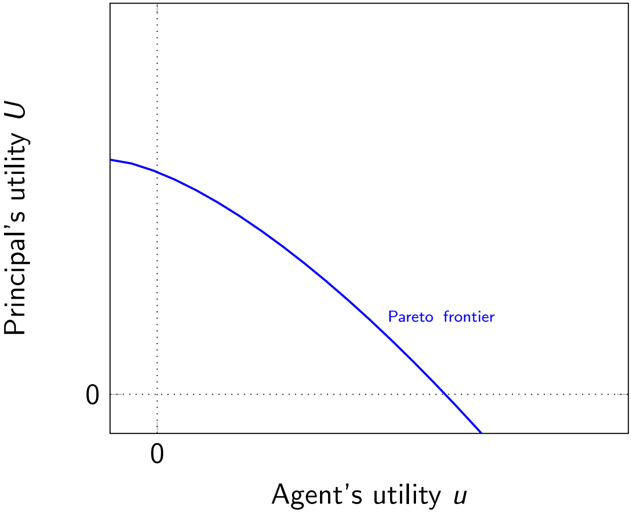

Reservation Utilities

- We assumed that participants had reservation utilities $\underline{u}$ and $\underline{U}$.
- This is 'wlog': Any surrunding market can be summarized by some *reservation utility* of the participants.
  - What do the participants expect to obtain if they don't sign the contract?

The contract curve

- The *contract curve* is the section of the Pareto frontier that is above the reservation utilities.
- Which point on the contract curve is chosen depends on the relative *bargaining power* of the participants.

- Note: The contract curve might be empty.

The contract curve

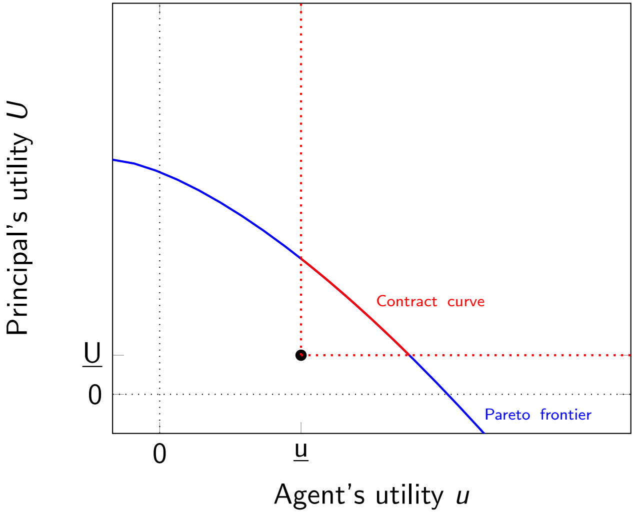

Characterizing the Pareto frontier

- Fix any level $\bar{u}$ of utility for the agent.
- Any point on the Pareto frontier is found by maximizing the principal's utility subject to leaving at least $\bar{u}$ utility to the agent.

$$\max_{a,w(\cdot)} \int (\pi - w(\pi)) f(\pi|a)\, d\pi \qquad (**)$$

$$\text{s.t.} \int v(w(\pi)) f(\pi|a)\, d\pi - c(a) \geq \bar{u}$$

- We solve the problem in two steps:
  1. Fix $a \in \mathcal{A}$ and maximize principal's utility over all wage schemes.
  2. find the maximizing action a^*.

Characterizing the Pareto frontier: Step 1.

- fix $a \in \mathcal{A}$.

$$\max_{w(\cdot)} \int (\pi - w(\pi)) f(\pi|a)\ d\pi \tag{*}$$

$$\text{s.t.} \quad \int v(w(\pi)) f(\pi|a)\ d\pi - c(a) \geq \bar{u}$$

- Equivalently,

$$\min_{w(\cdot)} \int w(\pi) f(\pi|a)\ d\pi$$

$$\text{s.t.} \quad \int v(w(\pi)) f(\pi|a)\ d\pi - c(a) \geq \bar{u}$$

- The feasible set of the problem is non-empty if and only if

$$\lim_{w \to \infty} v(w) > \bar{u} + c(a).$$

Illustration with two profit levels

- For illustration, suppose (contrary to our earlier assumptions) that only two profits can occur π_A and π_B.
- (Remember that we are fixing an action $a \in \mathcal{A}$.)
- Let w_A and w_B denote the wages that the agent receives when profits are π_A and π_B respectively.
- Let $p_A(a)$ and $p_B(a) := 1 - p_A(a)$ be the probabilities of outcomes π_A and π_B respectively.

Illustration with two profit levels

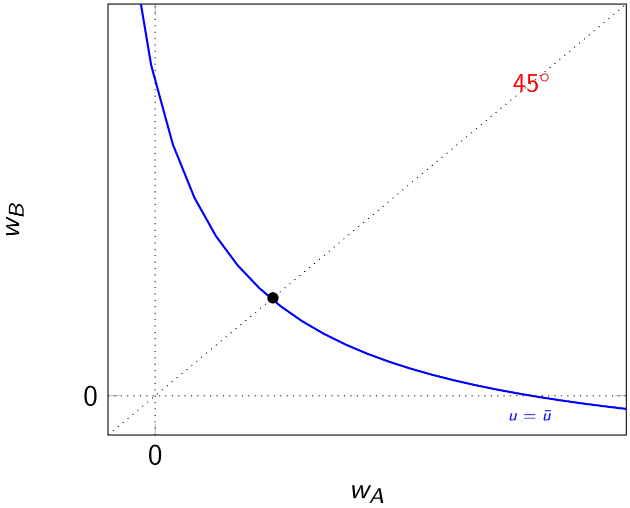

Illustration with two profit levels

- Indifference curve for the principal:

$$p_A(a)(\pi_A - w_A) + p_B(a)(\pi_B - w_B) = \tilde{U}$$

  - Differentiating.

$$p_A(a) \cdot (-dw_A) + p_B(a) \cdot (-dw_B) = 0$$

  - Rearranging:

$$\frac{dw_B}{dw_A} = -\frac{p_A(a)}{p_B(a)}$$

Illustration with two profit levels

- Indifference curve for the agent:

$$p_A(a)v(w_A) + p_B(a)v(w_B) - c(a) = \tilde{u}$$

  - Differentiating.

$$p_A(a) \cdot v'(w_A) \cdot \; dw_A + p_B(a) \cdot v'(w_B) \cdot dw_B = 0$$

  - Rearranging:

$$\frac{dw_B}{dw_A} = -\frac{v'(w_A) \cdot p_A(a)}{v'(w_B) \cdot p_B(a)}$$

Illustration with two profit levels

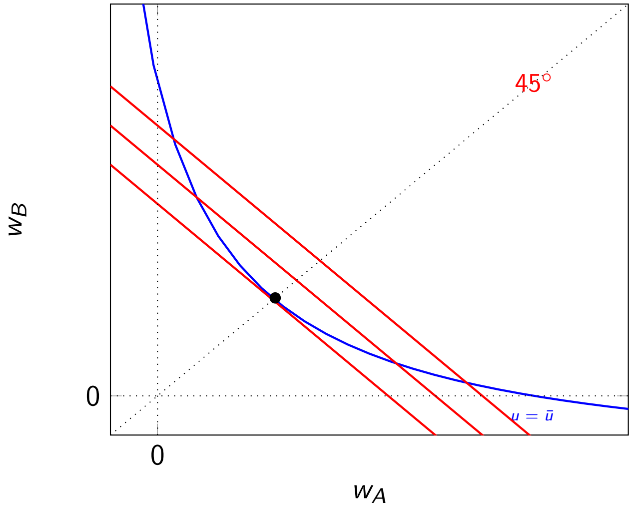

Illustration with two profit levels

- **Full insurance**: The optimal wage scheme satisfies $w_A = w_B$.
- This logic extends to the original setup with a continuum of outcomes.
  - Intuition: Any random payments can be replaced by the certainty equivalent, which is less costly to the principal.

Lagrange conditions

$$L(\gamma, w(\cdot)) = \int w(\pi) f(\pi|a)\, d\pi - \gamma \left[\int v(w(\pi)) \cdot f(\pi|a)\, d\pi - c(a) - \bar{u} \right]$$

There is a wage scheme $w^*(\cdot)$ that solves (*) if and only if there exists $\gamma \geq 0$ such that

$$w^* \quad \text{solves} \quad \min_{w(\cdot)} \; L(\gamma, w(\cdot))$$

And

$$\int v(w^*(\pi)) \cdot f(\pi|a)\, d\pi - c(a) - \bar{u} \geq 0$$

With equality if $\gamma > 0$.

- Reference: *Luenberger (1969), "Optimization by Vector Space Methods"*

Step 1: Minimizing Lagrange function

- We can rewrite the Lagrange function as:

$$L(\gamma, w(\cdot)) = \int [w(\pi) - \gamma v(w(\pi))] f(\pi|a)\, d\pi - \gamma c(a) - \gamma \bar{u}.$$

- The problem of minimization is equivalent to minimize

$$w(\pi) - \gamma v(w(\pi)) \qquad \text{for almost all } \pi.$$

- A solution can be chosen such that w is independent of π.
- Because $w - \gamma v(w)$ is convex in w, the first-order condition is sufficient for a minimum.
- Hence, $1 - \gamma v'(w^*) = 0$ implies that w^* is a minimum.

Step 1: The optimal wage scheme

$w^*(\pi) = \hat{w} := v^{-1}(\bar{u} + c(a))$ for all π is an optimal wage scheme.

- First, observe that $E[v(\hat{w}) - c(a)] = \bar{u}$.
- Define $\gamma = 1/v'(\hat{w})$.
- Then the Lagrange conditions are satisfied.
- Hence $w^*(\cdot)$ solves the problem (*).

- $v^{-1}(\bar{u} + c(a))$ can be thought as the *cost of implementing action a*.

Step 2: The optimal action

A contract $(a^*, w^*(\cdot))$ that solves the problem (**) is given by $w^*(\pi) = \hat{w} := v^{-1}(\bar{u} + c(a^*))$ for all π and such that

$$a^* \in \arg \max_{a \in \mathcal{A}} \int \pi \cdot f(\pi|a) \, d\pi - v^{-1}(\bar{u} + c(a))$$

- The principal provides full insurance to the agent.
- Principal stipulates an action that optimally trades off the expected profit against her cost of implementing the action.

Overview

1. Verifiable action
2. Non-verifiable actions

Non-verifiable actions: outlook

- Now we assume that actions are non-verifiable to courts.
  - May or may not be observable by the principal.
- Let $A_\circ \subset \mathcal{A}$ be the set of actions that minimizes $c(a)$.

If the agent is strictly risk-averse, then no point on the Pareto frontier that is only feasible with an action in $\mathcal{A} \setminus A_\circ$ can be achieved.

Intuition

- In every point on the Pareto frontier the agent is *fully insured*.
- But any fully insured agent will choose the least costly action.

- However, if the agent is risk-neutral, every point on the Pareto frontier can still be achieved.

Incentive compatible contracts

- **Question:** what can the principal do to implement an action $a \in \mathcal{A}$?
  - Align incentives via $w(\cdot)$.

- As before, a *contract* is a pair $(a, w(\cdot))$.
- Now, however, a is interpreted as a recommendation that the agent may or may not follow.
- A contract is *incentive compatible* if the agent has no incentive to deviate from the recommendation.

$$a \in \arg\max_{a \in \mathcal{A}} \int v(w(\pi)) f(\pi|a)\, d\pi - c(a) \tag{IC}$$

Constrained feasibility

- An expected utility pair (u, U) is *constrained feasible* if it can be obtained via some incentive compatible contract.
- The *constrained Pareto frontier* is the constrained feasible utility pairs that are not Pareto dominated by any other constrained feasible pair.

Observation

Any point on the constrained Pareto frontier is either on the Pareto frontier, or is (unconstrained) Pareto dominated.

Risk-neutral agent

- Suppose $v(w) = w$ for all $w \in \mathbb{R}$.

If the agent is risk-neutral, then the constrained Pareto frontier is identical to the Pareto frontier.

- This is achieved with a contract that "sells the firm to the manager".
- Sufficient to show that, for any $\bar{u}$, the problem (**) has the same solution value as the problem:

$$\max_{a, w(\cdot)} \int (\pi - w(\pi)) \cdot f(\pi|a)\, d\pi$$

$$\text{s.t} \int v(w(\pi)) \cdot f(\pi|a)\, d\pi - c(a) \geq 0$$

and Incentive Compatibility.

Risk-neutral agent

- Using risk-neutrality and participation binding, previous problem can be written as:

$$\max_{a,w(\cdot)} \int (\pi - w(\pi)) \cdot f(\pi|a)\, d\pi$$

$$\text{s.t} \int w(\pi) \cdot f(\pi|a)\, d\pi - c(a) = \bar{u}$$

and Incentive Compatibility.

- This can be rewritten as:

$$\max_{a,w(\cdot)} \quad \int \pi \cdot f(\pi|a)\, d\pi - c(a) - \bar{u}$$

$$\text{s.t} \int w(\pi) \cdot f(\pi|a)\, d\pi - c(a) = \bar{u}$$

and Incentive Compatibility.

Risk-neutral agent

$$\max_{a,w(\cdot)} \quad \int \pi \cdot f(\pi|a)\ d\pi - c(a) - \bar{u}$$

$$\text{s.t} \int w(\pi) \cdot f(\pi|a)\ d\pi - c(a) = \bar{u}$$

and Incentive Compatibility.

- This problem has a simple solution:
  - Choose a to $\max_{a \in \mathcal{A}} \int \pi \cdot f(\pi|a)\ d\pi - c(a) - \bar{u}$.
  - Let $w(\pi) = \pi - \alpha$ for all π for some constant α. The agent becomes the *residual claimant* to the profit.
  - IC constraint holds: agent's incentives is aligned with the principal's.
  - Choose α such that $\int w(\pi) \cdot f(\pi|a)\ d\pi - c(a) = \bar{u}$.

Formulating the problem

- Back to the risk-averse agent.
  - v' strictly decreasing,
  - For simplicity, v unbounded above.

- Fix a utility for the agent $\bar{u} \in \mathbb{R}$.
- As before, we can split the problem in two:
  1. For any action $a \in \mathcal{A}$, we look at the lowest cost to implement it, i.e. find $w(\cdot)$ which is the lowest cost incentive scheme that *implements a*.
  2. Given the costs to implement each action, choose a^*_{SB} that maximizes profits, given the utility that the agent must obtain.

Step 1

- Fix an action $a \in \mathcal{A}$.

$$\min_{w(\cdot)} \int w(\pi) \cdot f(\pi|a)\, d\pi \qquad \text{s.t.} \tag{*'}$$

$$\int v(w(\pi)) f(\pi|a)\, d\pi - c(a) \geq \bar{u}$$

$$a \in \arg \max_{a' \in \mathcal{A}} \int v(w(\pi)) \cdot f(\pi|a')\, d\pi - c(a') \tag{IC}$$

- Only difference with (*) is the presence of the IC constraint.

Implementing actions

- Suppose that principal wants to implement an action $a_\circ \in \mathcal{A}_\circ$.
  - Trick often useful in solving optimization problems:
    - Look at a *relaxed problem* where certain constraints are ingnored.
    - Then check that the solution to the relaxed problem in fact satisfies the ignored constraints.
  - Ignoring IC, the solution to the problem is the one with verifiable actions.

$$w^*(\pi) = v^{-1}(\bar{u} + c(a_\circ))$$

  - Because the wage is constant, the IC constraint is satisfied.
  - Hence, w^* solves the problem with the IC constraints.

Implementing actions

- Suppose that we want to implement an action $a \notin \mathcal{A}_\circ$.

- **Observation**: the feasible set of problem (*') may not be convex.
  - IC constraint is a collection of constraints:

$$\int v(w(\pi)) \cdot f(\pi|a)\ d\pi - c(a) \geq \int v(w(\pi)) \cdot f(\pi|a')\ d\pi - c(a')$$

- Therefore, not clear whether Lagrange conditions are sufficient for optimum.

Implementing actions: reformulation

- A transformation of variables fixes this problem.
- One-to-one correspondence between wages and utilities from wage.

$$\hat{v}(\pi) = v(w(\pi))$$

- Problem (*') can be reformulated as:

$$\min_{\hat{v}(\cdot)} \int v^{-1}(\hat{v}(\pi)) \cdot f(\pi|a)\, d\pi \qquad \text{s.t.} \qquad (*'')$$

$$\int \hat{v}(\pi) \cdot f(\pi|a)\, d\pi - c(a) \geq \bar{u}$$

$$a \in \arg \max_{a' \in \mathcal{A}} \int \hat{v}(\pi) \cdot f(\pi|a')\, d\pi - c(a') \qquad \text{(IC)}$$

- This problem has a convex objective and a feasible set that is convex as well. K-T applies.

Solution Existence

- **Question**: when will (*") have a solution?
  - If the feasible set is empty there is trivially no solution.
  - We will show that the set is non-empty.
- Consider linear utility schedules:

$$\hat{v}(\pi) = \alpha\pi + \beta.$$

- The constraints of the problem become:

$$\alpha \int \pi \cdot f(\pi|a)\; d\pi + \beta - c(a) \geq \bar{u}$$

$$\alpha \left(\int \pi \cdot f(\pi|a)\; d\pi - \int \pi \cdot f(\pi|a')\; d\pi \right) - [c(a) - c(a')] \geq 0$$

- With **two actions**, by choosing α and β one can guarantee that both constraints are satisfied as long as the expected profits are different for both actions.

Example: Implementing high effort

- From now on, we continue with the example $\mathcal{A} = \{e_L, e_H\}$.

- Define Lagrange function as:

$$
\begin{aligned}
L(\gamma, \mu, \hat{v}(\cdot)) = \int & v^{-1}(\hat{v}(\pi)) \cdot f(\pi|e_H)\ d\pi \\
& - \gamma \left(\int \hat{v}(\pi) \cdot f(\pi|e_H)\ d\pi - c(e_H) \right) \\
& - \mu \left(\int \hat{v}(\pi)[f(\pi|e_H) - f(\pi|e_L)]\ d\pi - [c(e_H) - c(e_L)] \right)
\end{aligned}
$$

Implementing high effort: Lagrange conditions

- A utility scheme $\hat{v}^*(\cdot)$ solves problem (*") if and only if there exist $\gamma \geq 0$ and $\mu \geq 0$ such that

$$\hat{v}^*(\cdot) \quad \text{solves} \min_{\hat{v}(\cdot)} \quad L(\gamma, \mu, \hat{v}(\cdot))$$

$$\int \hat{v}(\pi) \cdot f(\pi | e_H) \, d\pi - c(e_H) \geq \bar{u} \qquad \text{with equality if } \gamma > 0.$$

$$\int \hat{v}(\pi) \cdot f(\pi | e_H) \, d\pi - c(e_H) \geq \int \hat{v} \cdot f(\pi | e_L) \, d\pi - c(e_L)$$

with equality if $\mu > 0$.

- Another application of the convex optimization theorem (Luenberger's book pages 217 and 220).

Implementing high effort

- We can rewrite the Lagrange function as:

$$L(\gamma, \mu, \hat{v}(\cdot)) = \int \left(\left(v^{-1}(\hat{v}(\pi)) - (\gamma + \mu)\hat{v}(\pi) \right) f(\pi|e_H) \right.$$
$$\left. + \mu \hat{v}(\pi) f(\pi|e_L) \right) \; d\pi \; + \; (\gamma + \mu) g(e_H) - \mu g(e_L).$$

- Equivalent to minimizing

$$\left(v^{-1}(\hat{v}(\pi)) - (\gamma + \mu)\hat{v}(\pi) \right) f(\pi|e_H) + \mu \hat{v}(\pi) f(\pi|e_L) \qquad \text{a.e.}$$

Implementing high effort

- The previous function is strictly convex in $\hat{v}(\pi)$, so the FOC is necessary and sufficient for minimization.

- Hence, $\hat{v}(\cdot)$ is a minimizer if and only if

$$\left(v^{-1'}(\hat{v}(\pi)) - (\gamma+\mu)\right) f(\pi|e_H) + \mu f(\pi|e_L) = 0 \qquad \text{a.e.}$$

Implementing high effort

If the Lagrange conditions are satisfied, then $\gamma > 0$. (In other words, the participation constraint is binding.)

Proof.

- Suppose that $\gamma = 0$.
- Then,

$$v^{-1'}(\hat{v}^*(\pi))f(\pi|e_H) - \mu f(\pi|e_H) + \mu f(\pi|e_L) = 0 \qquad \text{a.e.}$$

- First term is strictly positive, hence

$$f(\pi|e_H) - f(\pi|e_L) > 0 \qquad \text{a.e.}$$

- This contradicts the fact that both $f(\cdot|e_H)$ and $f(\cdot|e_L)$ are densities (integrate to 1).

□

Implementing high effort

If the Lagrange conditions are satisfied, then $\mu > 0$. (IC is also binding.)

Proof.

- Suppose that $\mu = 0$.
- Then
$$\left(v^{-1'}(\hat{v}(\pi)) - \gamma\right) f(\pi|e_H) = 0 \qquad \text{a.e.}$$
- Hence, there is no wage uncertainty.
- This contradicts the IC constraint.

□

Implementing high effort: summary

A utility scheme $\hat{v}^*(\cdot)$ solves (*'') if and only if there exist $\gamma > 0$ and $\mu > 0$ such that

$$\begin{aligned}
\left(v^{-1'}(\hat{v}^*(\pi)) - (\gamma+\mu)\right) f(\pi|e_H) + \mu f(\pi|e_L) &= 0 \qquad \text{a.e.} \\
\int \hat{v}^*(\pi) f(\pi|e_H)\, d\pi - c(e_H) &= \bar{u} \\
\int \hat{v}^*(\pi)[f(\pi|e_H) - f(\pi|e_L)]\, d\pi &= c(e_H) - c(e_L)
\end{aligned}$$

Implementing high effort: reformulation

A wage scheme $w^*(\cdot)$ solves (*') if and only if there exist $\gamma > 0$ and $\mu > 0$ such that

$$\begin{aligned} \left(\frac{1}{v'(w^*(\pi))} - (\gamma + \mu)\right) f(\pi|e_H) + \mu f(\pi|e_L) &= 0 \qquad \text{a.e.} \\ \int v(w^*(\pi)) \cdot f(\pi|e_H)\, d\pi - c(e_H) &= \bar{u} \\ \int v(w^*(\pi)) \cdot [f(\pi|e_H) - f(\pi|e_L)]\, d\pi &= c(e_H) - c(e_L) \end{aligned}$$

Implementing high effort: structure

- Solving the first equality for w^* yields

$$w^*(\pi) = v'^{-1}\left(\frac{1}{\gamma + \mu\left(1 - \frac{f(\pi|e_L)}{f(\pi|e_H)}\right)}\right)$$

The optimal wage to implement high effort is increasing in the *likelihood ratio* $f(\pi|e_H)/f(\pi|e_L)$.

MLRP

We say that the monotone likelihood ratio property (MLRP) is satisfied if:

$$\frac{f(\pi|e_H)}{f(\pi|e_L)} \text{ is weakly increasing in } \pi$$

- MLRP is stronger than FOSD.

MLRP $\Rightarrow$ FOSD

- Consider the function

$$\chi(\pi) = F(\pi|e_L) - F(\pi|e_H) = \int_{\underline{\pi}}^{\pi} [f(\pi|e_L) - f(\pi|e_H)] \, d\pi$$

- Clearly, $\chi(\underline{\pi}) = \chi(\bar{\pi}) = 0$.
- The MLRP implies that there exists a profit π° such that

$$\frac{f(\pi|e_H)}{f(\pi|e_L)} \leq 1 \qquad \Leftrightarrow \qquad \pi \leq \pi^\circ$$

- Hence, $\chi(\cdot)$ is weakly increasing on $[\underline{\pi}, \pi^\circ]$ and weakly decreasing on $[\pi^\circ, \bar{\pi}]$.
- Thus, $\chi(\pi)$ is weakly positive.

Corollary

Corollary

if the MLRP is satisfied, then the optimal wage scheme to implement e_H is weakly increasing in the profit π.