Information Economics

The Moral Hazard Theory

Ling-Chieh Kung

Department of Information Management
National Taiwan University

Road map

- **Introduction**.
- The KKT condition.
- Deterministic outcome.
- Binary outcome.
- The LEN model.

Moral hazard

- There are two types of private information.
  - Hidden information, which causes the adverse selection problem.
  - **Hidden actions**, which cause the **moral hazard** problem.
- Consider a car insurance company and a driver.
  - The driver's after-purchase driving behavior determines the probability of a car accident.
  - The driving behavior is **hidden** to the company.
  - Once the driver gets an insurance, he will drive less carefully.
  - That is why the company may ask for a **deductible**.
- Consider a sales manager and a salesperson.
  - The salesperson's sales effort determines the sales outcome.
  - The sales effort is **hidden** to the company.
  - Once the salesperson gets a fixed salary, he will work less diligently.
  - That is why the manager may offers a **commission**.

Moral hazard

- Moral hazard is an issue when an agent has a hidden action.
  - Some people call this the **agency problem**: The principal **delegates** an action to the agent.
  - Some people call the theory of moral hazard the **agency theory**.
- In general, the agent takes an action, which affects the realization of an **outcome** that is cared by the principal.
  - The driver's driving behavior affects the realization of a car accident.
  - The salesperson's effort affects the realization of the sales outcome.
- The agent pays the **cost** of taking the action. Therefore, the principal should **pay** the agent to induce a desired action.
- The principal faces a **contract design** problem:
  - If the action is observable, the principal may compensate the agent based on his **action** (and the realized outcome).
  - When the action is unobservable, the principal may compensate the agent based on the realized **outcome** only.

Elements resulting moral hazard

- Delegation (i.e., decentralization) does not necessarily hurts efficiency.
- It will be shown that delegating the action to the agent is a problem **only if** all the following are true:
  - The action is **hidden**.
  - The outcome is **random**.
  - The agent is **risk-averse**.
- We will start from a model with deterministic outcomes to show that delegation does not create moral hazard.
- We then introduce two models with random outcomes.
  - The binary outcome model.
  - The LEN model.
- Before that, we need to talk about **risk attitudes**.

Risk attitudes

- Consider two random payoffs A and B:
  - $\Pr(A = 1) = 1$.
  - $\Pr(B = 0) = \Pr(B = 2) = \frac{1}{2}$.
  - Note that $\mathbb{E}[A] = \mathbb{E}[B]$, but $\text{Var}(A) < \text{Var}(B)$.
- People have different preferences due to different **risk attitudes**.
  - If one prefers A, she is typically believed to be **risk-averse**.
  - If one prefers B, she is said to be **risk-seeking** (or risk-loving).
  - If one feels indifferent, she tends to be **risk-neutral**.
- One's risk attitude is governed by the **shape** of her utility function.
- Consider two utility functions $u_1(z) = z$ and $u_2(z) = \begin{cases} z & \text{if } z \leq 1 \\ 1 & \text{if } z > 1 \end{cases}$.
  - Player 1 is risk-neutral.
  - Player 2 is risk-averse.

Risk attitudes vs. utility functions

- Though in practice it is hard to fully describe one's risk attitude, we adopt the conventional assumption:

Assumption 1

The shape of one's utility function $u(\cdot)$ decides her risk attitude:

- *One is risk-averse if and only if $u(\cdot)$ is concave.*
- *One is risk-seeking if and only if $u(\cdot)$ is convex.*
- *One is risk-neutral if and only if $u(\cdot)$ is linear.*

- We said that player 1 is risk-neutral and player 2 is risk-averse. Are their utility functions really linear and concave?

- But this example is restricted. Is the assumption reasonable in general?

General random payoffs

- Consider a random payoff X and a **concave** utility function $u(\cdot)$:
  - Jensen's inequality: $\mathbb{E}\Big[u(X)\Big] \leq u\Big(\mathbb{E}[X]\Big)$.
  - No matter what the original random payoff is, I always prefer to be offered the expected payoff.
  - A high payoff creates a "not-so-high" utility.
- What if $u(\cdot)$ is convex?
  - $\mathbb{E}[u(X)]$ and $u(\mathbb{E}[X])$, which is higher?
  - A high payoff creates a "very high" utility.
- What if $u(\cdot)$ is linear?
  - Maximizing the expected utility is the same as maximizing the expected payoff.

Road map

- Introduction.
- **The KKT condition**.
- Deterministic outcome.
- Binary outcome.
- The LEN model.

Constrains and Lagrange relaxation

- Consider a constrained nonlinear program

$$\begin{aligned} \max_{x \in \mathbb{R}^n} \quad & f(x) \\ \text{s.t.} \quad & g_i(x) \leq 0 \quad \forall i = 1, ..., m. \end{aligned}$$

- We apply **Lagrange relaxation** to the constraints. Given $\lambda = (\lambda_1, ..., \lambda_m) \leq 0$ as the **Lagrange multipliers**, we relax the constraints and move them to the objective function:

$$\max_{x \in \mathbb{R}^n} f(x) + \sum_{i=1}^{m} \lambda_i g_i(x).$$

  - We want the objective value to be large and $g_i(x) \leq 0$.
  - $\lambda_i \leq 0$ is the **penalty** of $g_i(x)$ to be positive.

Constrains and Lagrange relaxation

- The relaxed program is much easier to solve.
- We define the relaxed objective function as the **Lagrangian**:

$$\mathcal{L}(x|\lambda) = f(x) + \sum_{i=1}^{m} \lambda_i g_i(x).$$

  The relaxed problem is to maximize $\mathcal{L}(x|\lambda)$ over x when λ is given.
- If $\bar{x}$ is a local maximizer, it satisfy the **FOC for the Lagrangian**

$$\nabla\left\{ f(\bar{x}) + \sum_{i=1}^{m} \lambda_i g_i(\bar{x}) \right\} = 0 \quad \Leftrightarrow \quad \nabla f(\bar{x}) + \sum_{i=1}^{m} \lambda_i \nabla g_i(\bar{x}) = 0$$

  for some $\lambda \leq 0$.
- Interestingly, if $\bar{x}$ is a local maximizer to the constrained program, it must also be a local maximizer to the relaxed unconstrained program!

The KKT condition

- A very useful constrained optimality condition is the **KKT condition**.

Proposition 1 (KKT condition)

For a "regular" nonlinear program

$$\begin{aligned} \max_{x \in \mathbb{R}^n} \quad & f(x) \\ s.t. \quad & g_i(x) \leq 0 \quad \forall i = 1, ..., m. \end{aligned}$$

If $\bar{x}$ is a local max, then there exists $\lambda \in \mathbb{R}^m$ such that

- $g_i(\bar{x}) \leq 0$ *for all* $i = 1, ..., m$,
- $\lambda \leq 0$ *and* $\nabla f(\bar{x}) + \sum_{i=1}^{m} \lambda_i \nabla g_i(\bar{x}) = 0$, *and*
- $\lambda_i g_i(\bar{x}) = 0$ *for all* $i = 1, ..., m$.

- Most problems in the field of economics are "regular".
- This is only a **necessary** condition in general.
- Note the link between the second part and Lagrange relaxation.

Example

- For a constrained program, the KKT condition may be applied to find candidate optimal solutions.
  - An optimal solution x^* must satisfy all the three parts.
  - x^* must satisfy the second part, which is sometimes useful enough.
- Consider the problem of minimizing $x_1^2 + x_2^2$ subject to $4 - x_1 - x_2 \leq 0$.
  - The Lagrangian is

$$\mathcal{L}(x_1, x_2 | \lambda) = x_1^2 + x_2^2 + \lambda(4 - x_1 - x_2).$$

  - The FOC of the Lagrangian is

$$\frac{\partial}{\partial x_1^*}\mathcal{L} = 2x_1^* - \lambda = 0 \quad \text{and} \quad \frac{\partial}{\partial x_2^*}\mathcal{L} = 2x_2^* - \lambda = 0,$$

    which implies that $x_1 = x_2$.
  - Knowing that $4 - x_1 - x_2 \leq 0$ must be binding at an optimal solution, the only candidate solution is $(x_1^*, x_2^*) = (2, 2)$.

Road map

- Introduction.
- The KKT condition.
- **Deterministic outcome**.
- Binary outcome.
- The LEN model.

The first example

- An agent takes an **action** $a \geq 0$ (as some kind of effort) by paying $c(a)$ as his **cost**. For simplicity, let $c(a) = a$.
- The **outcome** $q(a)$ depends on a in a deterministic way. We have $q(\cdot)$ strictly increasing and strictly concave.
- The principal **compensates** the agent for his action by paying w.
  - If a is observable, w can be $w(q, a)$, i.e., contingent on q and a.
  - If a is unobservable, w will be $w(q)$, i.e., contingent only on q.
- The principal's payoff is $q(a) - w$.
- The agent may be risk neutral or risk averse.
  - If he is **risk neutral**, his payoff is $w - a$.
  - If he is **risk averse**, his payoff is $u(w) - a$, where $u(\cdot)$ is strictly increasing and strictly concave.

Risk-neutral agent: first best

- Consider the first-best scenario with a risk-neutral agent.
  - The risk-neutral agent's utility is $w - a$.
  - First best: The action is observable.
- The principal's problem:

$$\begin{aligned} \max_{w(\cdot,\cdot),a} \quad & q(a) - w(q(a), a) \\ \text{s.t.} \quad & w(q(a), a) - a \geq 0. \end{aligned}$$

- The constraint must be binding at an optimal solution. The problem reduces to $\max_a q(a) - a$. The optimal a^* satisfies $q'(a^*) = 1$.
- The compensation plan $w(\cdot, \cdot)$ satisfies $w(q, a^*) = a^*$ for any q.
  - Simply compensate the agent the **cost of the efficient action**.
  - The **input-based** compensation is not contingent on the outcome.

Risk-neutral agent: second best

- Consider the second-best scenario with a risk-neutral agent.
- The principal's problem:

$$\begin{aligned}\max_{w(\cdot)} \quad & q(a) - w(q(a)) \\ \text{s.t.} \quad & w(q(a)) - a \geq 0 \\ & a \in \operatorname*{argmax}_{\hat{a}} \{w(q(\hat{a})) - \hat{a}\}.\end{aligned}$$

- May the principal induce the first-best a^*, which satisfies $q'(a^*) = 1$?
  - Let $q^* = q(a^*)$.
  - Because the outcome is deterministic, only a^* can result in q^*.
  - The principal can "shoot" the agent as long as the outcome is not q^*.
  - The **output-based** compensation plan is efficient and optimal:

$$w(q) = \begin{cases} a^* & \text{if } q = q^* \\ -\infty & \text{otherwise} \end{cases}.$$

Risk-averse agent: first best

- Consider the first-best scenario with a risk-averse agent.
  - The risk-averse agent's utility is $u(w) - a$.
  - First best: The action is observable.
- The principal's problem:

$$\begin{aligned} \max_{w(\cdot,\cdot),a} \quad & q(a) - w(q(a), a) \\ \text{s.t.} \quad & u(w(q(a), a)) - a \geq 0. \end{aligned}$$

- Let a^* be an optimal action chosen by the principal.
- $w(q, a)$ can be designed so that $u(w(q, a^*)) = a^*$ for any q:

$$w(q, a^*) = u^{-1}(a^*).$$

Risk-averse agent: second best

- Consider the second-best scenario with a risk-averse agent.
- The principal's problem:

$$\begin{aligned} \max_{w(\cdot)} \quad & q(a) - w(q(a)) \\ \text{s.t.} \quad & u(w(q(a))) - a \geq 0 \\ & a \in \underset{\hat{a}}{\operatorname{argmax}} \{u(w(q(\hat{a}))) - \hat{a}\}. \end{aligned}$$

- May the principal induce the first-best a^*?
  - Let $q^* = q(a^*)$. Only a^* can result in q^*.
  - The principal can still "shoot" the agent if the outcome is not good:

$$w(q) = \begin{cases} u^{-1}(a^*) & \text{if } q = q^* \\ -\infty & \text{otherwise} \end{cases}.$$

  - The **output-based** compensation plan is still efficient and optimal.

Remarks

- When the outcome is **deterministic**, delegation does not create the moral hazard problem.
  - It does not matter whether the agent is risk-averse or not.
- The optimal contract is a "do-it-or-I-shoot-you" contract.
  - The agent gets a payment that is **just enough to cover** his cost for taking the first-best action.
  - The agent gets a **huge penalty** otherwise.
  - The agent in equilibrium earns nothing (no information rent).
  - The principal can **implement the first best** with an output-based compensation plan.
- This is all because the deterministic outcome can be used to accurately infer the agent's action.
- This is no longer the case if the outcome is **random**.

Road map

- Introduction.
- The KKT condition.
- Deterministic outcome.
- **Binary outcome**.
- The LEN model.

Binary outcome

- Let the outcome $q \in \{0, 1\}$ follows a Bernoulli distribution where

$$\Pr(q = 1|a) = p(a) = 1 - \Pr(q = 0|a).$$

  Let $p(\cdot)$ be strictly increasing, strictly concave, and no greater than 1.
- We should still discuss four cases:
  - The action is observable or unobservable.
  - The agent is risk-neutral or risk-averse.
- In each case, the principal should design a **compensation plan**.
  - Because the outcome is **binary**, the plan contains only two numbers w_0 and w_1, the payments for the agent when $q = 0$ and $q = 1$, respectively.
  - If the action is observable, we can have $w_0(a)$ and $w_1(a)$. However, this is not needed because in equilibrium the agent will be assigned a value of a.
- The shape of $u(\cdot)$ determine the agent's risk attitude.
  - Let's work with the risk-averse agent directly.
  - The case with the risk-neutral agent will be a special case with $u(w) = w$.

Risk-averse agent: first best

- If the action is observable, the principal's problem is

$$\begin{aligned} \max_{w_0, w_1, a} \quad & p(a)(1 - w_1) + (1 - p(a))(-w_0) \\ \text{s.t.} \quad & p(a)u(w_1) + (1 - p(a))u(w_0) - a \geq 0. \end{aligned} \tag{1}$$

- The constraint is binding at any optimal solution. However, it does not help a lot (due to the nonlinearity of $u(\cdot)$).
- We rely on the KKT condition to reduce the problem.
  - Because the constraint is a greater-than-or-equal-to one, we have the Lagrange multiplier $\lambda \geq 0$.

Proposition 2

An optimal contract to the problem in (1) satisfies $w_0 = w_1$.

- Because the agent is risk-averse, he prefers a **fixed payment**.

Proof of the proposition

- Given $\lambda \geq 0$, the Lagrangian is

$$\mathcal{L}(w_0, w_1, a|\lambda) = p(a)(1 - w_1) + (1 - p(a))(-w_0) + \lambda\Big[p(a)u(w_1) + (1 - p(a))u(w_0) - a\Big].$$

- The FOC requires

$$\frac{\partial}{\partial w_0}\mathcal{L} = -(1 - p(a)) + \lambda(1 - p(a))u'(w_0) = 0 \quad \Leftrightarrow \quad \lambda = \frac{1}{u'(w_0)}$$

$$\frac{\partial}{\partial w_1}\mathcal{L} = -p(a) + \lambda p(a)u'(w_1) = 0 \quad \Leftrightarrow \quad \lambda = \frac{1}{u'(w_1)}.$$

  As $\lambda \geq 0$ and $u'(\cdot) > 0$, this is possible.

- In any optimal contract, $w_0 = w_1$!

Risk-averse agent: second best

- If the action is unobservable, the agent choose a to maximize her expected utility $p(a)u(w_1) + (1 - p(a))u(w_0) - a$. An optimal a satisfies $p'(a)[u(w_1) - u(w_0)] = 1$.
- The principal's problem is

$$\begin{aligned} \max_{w_0, w_1} \quad & p(a)(1 - w_1) + (1 - p(a))(-w_0) \\ \text{s.t.} \quad & p(a)u(w_1) + (1 - p(a))u(w_0) - a \geq 0 \\ & p'(a)[u(w_1) - u(w_0)] = 1. \end{aligned} \tag{2}$$

- To solve this problem, again we rely on the KKT condition.

Proposition 3

An optimal contract to (2) satisfies $w_1 > w_0$.

- To induce the agent to "work," a **bonus** for a good outcome is needed.

Proof of the proposition

- Given $\lambda \geq 0$ and μ urs.,[1] the Lagrangian of the reduced problem is

$$\begin{aligned}\mathcal{L}(w_0, w_1 | \lambda, \mu) = {} & p(a)(1 - w_1) + (1 - p(a))(-w_0) \\ & + \lambda\Big[p(a)u(w_1) + (1 - p(a))u(w_0) - a\Big] \\ & + \mu\Big[p'(a)[u(w_1) - u(w_0)] - 1\Big].\end{aligned}$$

- The FOC requires

$$\frac{\partial}{\partial w_0}\mathcal{L} = -(1 - p(a)) + \lambda(1 - p(a))u'(w_0) - \mu p'(a)u'(w_0) = 0 \text{ and}$$
$$\frac{\partial}{\partial w_1}\mathcal{L} = -p(a) + \lambda p(a)u'(w_1) + \mu p'(a)u'(w_1) = 0.$$

[1] The Lagrange multiplier for an equality should be "unrestricted in sign."

Proof of the proposition

- The FOC implies

$$\frac{1}{u'(w_0)} = \lambda - \mu \frac{p'(a)}{1 - p(a)} \text{ and}$$
$$\frac{1}{u'(w_1)} = \lambda + \mu \frac{p'(a)}{p(a)}.$$

- If $\mu = 0$, we go back to the first-best contract (and $w_0 = w_1$).
- The principal now may alter μ to improve her expected profit.
- It can be shown that an optimal contract satisfies $\mu > 0$ (how?).
- As $u'(w)$ decreases in w, $\frac{1}{u'(w)}$ increases in w.
- Therefore, if $\mu > 0$, we have $w_1 > w_0$.

Summary

- When the agent is risk-averse and outcome is random:
  - If the effort is observable: $w_0 = w_1$ to **remove risks** from the agent.
  - If the effort is unobservable: $w_0 < w_1$ to **incentivize** the agent.
- Information asymmetry (more precisely, hidden actions) results in efficiency loss.
- It can be shown that if the agent becomes risk-neutral, the second-best contract will also be efficient (how?).
  - **Risk aversion** is necessary for moral hazard.

Road map

- Introduction.
- The KKT condition.
- Deterministic outcome.
- Binary outcome.
- **The LEN model**.

The LEN model

- Sometimes we want to allow the random outcome to be **continuous**.
- A moral hazard model with a random outcome that has a general distribution can be easily intractible.
- A tractible model with a continuous outcome is **the LEN model**.
  - The compensation plan is **linear**.
  - The utility function is a negative **exponential** function.
  - The random outcome is **normally** distributed.
- More precisely:
  - Let the outcome $q = a + \epsilon$, where a is the action and $\epsilon \sim \text{ND}(0, \sigma^2)$.
  - Let the agent's utility function be $u(z) = -e^{-\eta z}$, where $\eta > 0$ is his coefficient of absolute risk aversion and z is the payoff.
  - Let the compensation plan be $t + sq$, where t is the fixed payment and s is the commission rate.
- The agent's cost of taking action a is $c(a) = \frac{1}{2}a^2$.

The agent's expected utility

- Given an offer (t, s), the agent chooses a to maximize

$$\mathbb{E}[u(z)] = \mathbb{E}[-e^{-\eta z}] = \mathbb{E}\Big[-e^{-\eta(t+sq-\frac{1}{2}a^2)}\Big] = \mathbb{E}\Big[-e^{-\eta(t+s(a+\epsilon)-\frac{1}{2}a^2)}\Big].$$

As only ϵ is random, we may simplify the expected utility to

$$\mathbb{E}\Big[-e^{-\eta(t+sa-\frac{1}{2}a^2)} \cdot e^{-\eta s\epsilon}\Big] = -e^{-\eta(t+sa-\frac{1}{2}a^2)}\mathbb{E}\Big[e^{-\eta s\epsilon}\Big],$$

where the expectation is the bilateral **Laplace transformation** of ϵ:

Proposition 4

Given $\epsilon \sim \text{ND}(0, \sigma^2)$ *and* $r \in \mathbb{R}$, *we have*

$$\mathbb{E}[e^{r\epsilon}] = e^{r^2\sigma^2/2}.$$

Proof of the proposition

$$
\begin{aligned}
\mathbb{E}[e^{r\epsilon}] &= \int_{-\infty}^{\infty} e^{rx} \overbrace{\frac{1}{\sqrt{2\pi\sigma}} e^{-x^2/(2\sigma^2)}}^{\text{pdf of ND}(0,\sigma^2)} dx \\
&= \frac{1}{\sqrt{2\pi\sigma}} \int_{-\infty}^{\infty} e^{-(x^2-2rx\sigma^2)/(2\sigma^2)} dx \\
&= \frac{1}{\sqrt{2\pi\sigma}} \int_{-\infty}^{\infty} e^{-((x-r\sigma^2)^2-r^2\sigma^4)/(2\sigma^2)} dx \\
&= \frac{1}{\sqrt{2\pi\sigma}} \int_{-\infty}^{\infty} e^{-((x-r\sigma^2)^2)/(2\sigma^2)} \cdot e^{r^2\sigma^2/2} dx \\
&= e^{r^2\sigma^2/2} \int_{-\infty}^{\infty} \underbrace{\frac{1}{\sqrt{2\pi\sigma}} e^{-((x-r\sigma^2)^2)/(2\sigma^2)}}_{\text{pdf of ND}(r\sigma^2,\sigma^2)} dx = e^{r^2\sigma^2/2}.
\end{aligned}
$$

Certainty equivalents

- Now the agent's expected utility is simplified to

$$\mathbb{E}[u(z)] = -e^{-\eta(t+sa-\frac{1}{2}a^2)} \cdot e^{\eta^2 s^2 \sigma^2/2} = -e^{-\eta(t+sa-\frac{1}{2}a^2-\frac{1}{2}\eta s^2 \sigma^2)}.$$

- We define the **certainty equivalent** of the agent's utility function as

$$CE(a) = t + sa - \frac{1}{2}a^2 - \frac{1}{2}\eta s^2 \sigma^2.$$

  - $t + sa - \frac{1}{2}a^2$ measures the expected return.
  - $\frac{1}{2}\eta s^2 \sigma^2$ measures the risk due to the uncertainty.
- Because $-e^{-\eta z}$ increases in z, maximizing the expected utility is equivalent to maximizing the certainty equivalent.
- The agent's optimal action is $a^* = s$.
  - A higher commission rate **induces** a higher effort level.

The contract design problem

- The principal's expected profit in equilibrium is

$$\mathbb{E}[(1-s)q-t] = (1-s)s - t + (1-s)\mathbb{E}[\epsilon] = (1-s)s - t.$$

- The agent's certainty equivalent in equilibrium is

$$CE(s) = t + \frac{1}{2}s^2 - \frac{1}{2}\eta s^2\sigma^2 = t + \frac{1}{2}s^2(1-\eta\sigma^2).$$

- The principal's problem is

$$\begin{aligned} \max_{t,s} \quad & (1-s)s - t \\ \text{s.t.} \quad & t + \frac{1}{2}s^2(1-\eta\sigma^2) \geq 0. \end{aligned}$$

The contract design problem

- As the constraint is binding at any optimal solution, the principal's problem reduces to

$$\max_{s} \, (1-s)s + \frac{1}{2}s^2(1-\eta\sigma^2).$$

The FOC gives the optimal commission rate

$$s^* = \frac{1}{1+\eta\sigma^2}.$$

- Economic interpretations:
  - s^* decreases in η: When the agent becomes **more risk-averse**, he prefers a lower commission rate (and a higher fixed payment).
  - s^* decreases in σ^2: When the outcome becomes **more unpredictable**, the agent prefers a lower commission rate (and a higher fixed payment).
- Remark: A linear contract is suboptimal.

Summary

- Hidden actions create the moral hazard problem.
  - The agent must be incentivized (compensated) for his action.
  - Compensation may or may not be inefficient.
- This is really a problem when all the following elements exist:
  - Unobservability of the action.
  - Uncertainty of the outcome.
  - Risk aversion of the agent.
- Information asymmetry:
  - Adverse selection: screening and signaling.
  - Moral hazard.
- The world is decentralized.
  - **Decentralization** brings in the **incentive** issue.
  - **Information asymmetry** aggravates the issue.