Lecture 5: The Principal-Agent Problem

Gonzalo Cisternas
MIT

Contents

Static Version: Ingredients

- Risk neutral principal and a risk averse agent
- $q_1 < q_2 < ... < q_n$ outcomes (i.e., output)
- Agent takes an action (i.e., effort) $a \in A$, where $A = [\underline{a}, \overline{a}]$ or $A = \{a_1, ..., a_J\}$. Action is **hidden** to the principal
- Each action induces a probability distribution π over outcomes:

$$\pi(a) := (\pi_1(a), ...\pi(n)),$$

where $\pi_i(a) := \text{Prob}(q_i|a) > 0,\ i = 1, ..., n$, and $\sum_{i=1}^{n} \pi_i(a) = 1$

- The agent's utility from income I and disutility from effort a:

$$U(I, a) = u(I) - h(a)$$

where $u' > 0,\ u'' < 0$. When $A = [\underline{a}, \overline{a}]$ $h' > 0,\ h'' > 0,\ h'(0) = 0$. When A is finite, $U(I, a) = u(I) - a$

The Problem of the Principal

- The principal chooses a **contract** $(I_1, ..., I_n)$, where I_i is the wage paid if q_i is observed, and an **effort recommendation** $a \in A$, such that they maximize her expected utility

$$\mathbb{E}^a[q - I] = \sum_{i=1}^{n} \pi_i(a)[q_i - I_i]$$

subject to:

  - The agent attains a minimum level of utility W:

$$\sum_{i=1}^{n} \pi_i(a)u(I_i) - h(a) \geq W$$

  called *participation constraint* (PC)

  - The agent finds it optimal to actually choose a:

$$a \in \arg\max_{\tilde{a} \in A} \sum_{i=1}^{n} \pi_i(\tilde{a})u(I_i) - h(\tilde{a})$$

  called *incentive-compatibility* constraint (IC)

First-Best Solution

- The **first-best solution** corresponds to the optimal contract for the principal when she can observe the agent's action
  - It is also called the full-information benchmark
- In finding this benchmark, the IC constraint is neglected, as the contract can explicitly depend on the effort chosen by the agent
  - Hence, any undesired effort can be punished so heavily that it is never chosen by the agent
- We start analyzing the binary effort case: $\{a_L, a_H\}$ with $a_L < a_H$

First Best: Cost Minimization

- Divide the principal's problem into two sub-problems:

(1) Minimize the cost of implementing any specific effort level

(2) Choose the effort level that maximizes the principal's utility given the cost function found in (1)

- **Cost minimization**: For each $a \in A$ solve

$$\min_{I_1,\ldots,I_n} \sum_{i=1}^{n} \pi_i(a) I_i$$

$$s.t. \qquad \sum_{i=1}^{n} \pi_i(a) u(I_i) - a \geq W \qquad (PC)$$

- Denote $C_{FB}(a)$ the cost of this problem

First Best: Maximize Benefits

- **Solution:** Choose constant income I such that (PC) holds with equality. Why?
  - For a_L, choose $I_L := u^{-1}(W + a_L) \Rightarrow C_{FB}(a_L) = I_L := u^{-1}(W + a_L)$
  - For a_H, choose $I_H := u^{-1}(W + a_H) \Rightarrow C_{FB}(a_H) = I_H := u^{-1}(W + a_H)$
- **Maximize benefits:** Choose a that solves

$$\max_{a \in \{a_L, a_H\}} \sum_{i=1}^{n} q_i \pi_i(a) - C_{FB}(a)$$

- Depending on the specification (π, q, A), the principal may choose a_L or a_H
- Now we move on to find the **second-best solution**, that is, the optimal contract when effort is not observed, and thus the IC constraint is taken into account. We follow the same two-step method

The Second-Best Solution

- The principal's problem

$$\max_{(I_1,\ldots,I_n),a} \quad \mathbb{E}^a[q - I] = \sum_{i=1}^{n} \pi_i(a)[q_i - I_i]$$

$$s.t. \qquad (PC_a): \; \sum_{i=1}^{n} \pi_i(a)u(I_i) - a \geq W$$

$$(IC_a): \; a \in \arg\max_{\tilde{a}\in\{a_L,a_H\}} \sum_{i=1}^{n} \pi_i(\tilde{a})u(I_i) - \tilde{a}$$

- **Step 1:** For each $a \in A$, solve

$$C(a) := \min_{(I_1,\ldots,I_n)} \quad \sum_{i=1}^{n} \pi_i(a) I_i$$

$$s.t. \qquad (PC_a) \text{ and } (IC_a) \tag{1}$$

- When a contract $I := (I_1, ..., I_n)$ satisfies (IC_a), we say that I **implements** a

Implementing Low Effort

- Observe that a flat contract $I_1 = ... = I_n = \bar{I}$ can only implement a_L, as $a_L < a_H$ and

$$\sum_{i=1}^{n} \pi_i(a) u(I_i) - a = u(\bar{I}) - a$$

so only (IC_{a_L}) can be satisfied

- The cheapest way to implement a_L is thus to choose

$$\bar{I}_L := u^{-1}(W + a_L)$$

- This contract yields a total profit of

$$\underbrace{\sum_{i=1}^{n} q_i \pi_i(a_L)}_{B(a_L)} - \underbrace{u^{-1}(W + a_L)}_{C_{FB}(a_L)}$$

to the principal

Implementing High Effort

$$C(a_H) := \min_{(I_1,\ldots,I_n)} \sum_{i=1}^{n} \pi_i(a_H) I_i$$

$$s.t. \qquad (PC_{a_H}): \sum_{i=1}^{n} \pi_i(a_H) u(I_i) - a_H \geq W$$

$$(IC_{a_H}): \sum_{i=1}^{n} \pi_i(a_H) u(I_i) - a_H \geq \sum_{i=1}^{n} \pi_i(a_L) u(I_i) - a_L$$

- Convex problem (principal chooses $u_i := u(I_i)$
- If $(I_1, ..., I_n)$ implements a_H at a minimum cost, then PC$_{a_H}$ holds with equality

$$\sum_{i=1}^{n} \pi_i(a_H) u(I_i) - a_H = W$$

  Why?
- Let $\lambda > 0$ ($\mu > 0$) lagrange multiplier on PC (IC)

Monotone Likelihood Ratio Property and Monotone Contracts

- FOC (from KKT):

$$\frac{1}{u'(I_i)} = \lambda + \mu \left[1 - \frac{\pi_i(a_L)}{\pi_i(a_H)}\right] \;\; i = 1, ..., n$$

- $\ell_i := \pi_i(a_L)/\pi_i(a_H)$ is called the **likelihood ratio**: how likely is to observe outcome i when choosing a_L versus a_H

- **Definition:** In this context, the set of outcomes $q_1 < ... < q_n$ satisfies the **monotone likelihood ratio property** (MLRP) if for all $i > j$

$$\ell_i := \frac{\pi_i(a_L)}{\pi_i(a_H)} < \ell_j := \frac{\pi_j(a_L)}{\pi_j(a_H)}$$

- From the FOC: $\ell_i < \ell_j \Rightarrow I_i > I_j$
- MLRP $\Rightarrow I_1 < I_2 < ... < I_n$, i.e. the contract is **monotone**

Cost of Incentive Provision

- Let, $B(a) := \sum_{i=1}^{n} q_i \pi_i(a)$ and define

$$L := \max_{a \in \{a_L, a_H\}} [B(a) - C_{FB}(a)] - \max_{a \in \{a_L, a_H\}} [B(a) - C(a)].$$

- $L \geq 0$: loss in profits to the principal due to her informational disadvantage $\Leftrightarrow$ cost of incentive provision
- $L = 0$ if and only if a_L is optimal in the first-best contract.
  - ($\Leftarrow$) Since, $C_{FB}(a_L) = C(a_L)$, we have that $L = 0$
  - ($\Rightarrow$) Suppose a_H in the first-best contract. Then a_L can't be optimal in the second-best contract (otherwise, a_L is optimal in the first-best contract!). Since a_H is optimal, there must be $I_i \neq I_j$. Hence,

$$u\left(\sum_{i=1}^{n} \pi_i(a_H) I_i\right) \underbrace{>}_{\text{u concave}} \sum_{i=1}^{n} \pi_i(a_H) u(I_i) = W + a_H = u(C_{FB}(a_H))$$

$$\Rightarrow C(a_H) := \sum_{i=1}^{n} \pi_i(a_H) I_i > C_{FB}(a_H)$$

Important Takeaway

- When actions cannot be monitored, inducing high effort is costly to the principal
- This is because in order to provide such incentives, the contract must reward/punish performance, i.e., $I_i \neq I_j$ some $i \neq j$
- Hence, the contract must expose the agent to risk
- But since the agent is risk averse, providing any level of utility with volatile income is always more expensive that doing so without any risk.

Tension between incentive provision and risk exposure

Extension 1: Finite Number of Actions

- Suppose that the principal wants to implement $a_j \in \{a_1, ..., a_J\}$, where $a_1 < a_2 < ... < a_J$. FOC for cost minimization:

$$\frac{1}{u'(I_i)} = \lambda + \sum_{j'=1}^{J} \mu_{j'} \left[1 - \frac{\pi_i(a_{j'})}{\pi_i(a_j)} \right]$$

  where λ is the lagrange multiplier of the PC, and $\mu_{j'}$ is the multiplier of the IC constraint that prevents the agent from deviating to $a_{j'}$

- The contract will be monotone if it is optimal to implement a_J and MLRP holds:

$$\frac{\pi_i(a_{j'})}{\pi_i(a_J)} < \frac{\pi_{i'}(a_{j'})}{\pi_{i'}(a_J)}, \text{ for } \; i > i'$$

- However, if a_J is not optimal, the contract may not be monotone. This is because the contract will reward the signal that is most likely to arise under the optimal effort ▸ Continuum of Actions

Contents

Repeated Moral Hazard (Rogerson 1985)

- Ingredients:
  - $T = 2$; A, W, $(q_1, ..., q_n)$ and $(\pi_1(a), ..., \pi_n(a))$ as before
  - Actions only affect outcomes in the period in which they are taken only
  - Outcomes are (conditionally) independent across time
  - Agent cannot save nor borrow
  - $\delta =$ discount factor
- A **contract** is a tuple

$$\mathcal{I} := ((I_1, ..., I_n), (I_{11}, ...I_{1n}), ..., (I_{n1}, ..., I_{nn})) \text{ where}$$

  where I_i is the wage paid at the end of $t = 1$ if q_i was observed and I_{ij} is the wage paid at the end of $t = 2$ if q_i was observed at $t = 1$ and q_j was observed at $t = 2$
- The principal chooses $\mathcal{I}$ and recommends an effort policy $a := (a_0, (a_1, ..., a_n))$ with $a_0 =$ effort at time 1, and $a_j =$ effort at time 2 after observing q_i, i$>$0

The Problem

$$\max_{\mathcal{I},a} \quad \sum_{i=1}^{n} \pi_i(a_0)\left[q_i - I_i + \delta \sum_{j=1}^{n} \pi_j(a_i)[q_j - I_{ij}]\right]$$

$$s.t. \quad \sum_{i=1}^{n} \pi_i(a_0)\left[u(I_i) - a_0 + \delta \sum_{j=1}^{n} \pi_j(a_i)[u(I_{ij}) - a_i]\right] \geq W$$

$$a \in \arg\max_{\tilde{a}} \sum_{i=1}^{n} \pi_i(\tilde{a}_0)\left[u(I_i) - \tilde{a}_0 + \delta \sum_{j=1}^{n} \pi_j(\tilde{a}_i)[u(I_{ij}) - \tilde{a}_i]\right]$$

- Solve using backward induction using the **promised utility** approach
  - Any contract $\mathcal{I}$ can be seen as a tuple $(I_1, ..., I_n, W_1, ..., W_n)$ where W_i is the expected utility that the principal offers the agent in the next period after observing q_i (i.e., $(I_{i1}, ...I_{in})$ must deliver on average W_i)
  - Advantage: solve two identical problems at $t = 1$ and $t = 2$
- Restrict to binary effort case: $h(a_L) = 0$ and $h(a_H) = \psi > 0$

$T = 2$: Static Contract

- After observing q_i and promising the agent W_i, the principal solves a static problem. W is a state variable
- First step: cost minimization. Suppose a_H is to be implemented:

$$
\begin{aligned}
C_2(W_i) := \quad \min_{(I_{i1},\dots,I_{in})} \quad & \sum_{j=1}^{n} \pi_j(a_H) I_{ij} \\
s.t. \quad & \sum_{j=1}^{n} \pi_j(a_H) u(I_{ij}) - \psi \geq W_i \\
& \sum_{j=1}^{n} \pi_j(a_H) u(I_{ij}) - \psi \geq \sum_{j=1}^{n} \pi_j(a_L) u(I_{ij})
\end{aligned}
$$

- $C_2(W_i) =$ cost of delivering W_i utils to the agent at $t = 2$
- FOC: $\frac{1}{u'(I_{ij})} = \lambda_2 + \mu_2 \left[1 - \frac{\pi_j(a_L)}{\pi_j(a_H)}\right]$; λ_2, μ_2 lagrange multipliers
- Envelope: $C_2'(W_i) = \lambda_2 = \sum_{j=1}^{n} \frac{\pi_j(a_H)}{u'(I_{ij})}$

$T = 1$: Going Backwards

- The principal solves

$$C_1(\overline{W}) := \min_{(\vec{I},\vec{W})} \sum_{i=1}^{n} \pi_i(a_H)[I_i + \delta C_2(W_i)]$$

$$s.t. \quad \sum_{i=1}^{n} \pi_i(a_H)[u(I_i) + \delta W_i] - \psi \geq \overline{W}$$

$$\sum_{i=1}^{n} \pi_i(a_H)[u(I_i) + \delta W_i] - \psi \geq \sum_{i=1}^{n} \pi_i(a_L)[u(I_i) + \delta W_i]$$

- FOC:

$$(u_i) \quad : \quad -\pi_i(a_H)/u'(I_i) + \lambda_1 \pi_i(a_H) + \mu_1(\pi_i(a_H) - \pi_i(a_L)) = 0$$

$$(W_i) \quad : \quad -\pi_i(a_H)\delta C_2'(W_i) + \lambda_1 \delta \pi_i(a_H) + \mu_1 \delta(\pi_i(a_H) - \pi_i(a_L)) = 0$$

$$\Rightarrow \frac{1}{u'(I_i)} = C_2'(W_i) = \sum_{j=1}^{n} \frac{\pi_j(a_H)}{u'(I_{ij})}$$

Inter-temporal Cost Smoothing

Theorem

Consider an optimal contract that implements $a > a_L$. Then, the inverse of marginal utility is a martingale:

$$\frac{1}{u'(I_i)} = \mathbb{E}^{a_H}\left[\frac{1}{u'(I_{i\cdot})}\right] := \sum_{j=1}^{n} \frac{\pi_j(a_H)}{u'(I_{ij})}$$

where $I_{i\cdot}$ is the $t = 2$ income seen as a random variable from $t = 1$.

Intuition: Let $J(v) := u^{-1}(v)$: income level that yields utility v. Then,

$$J'(v_i) = \mathbb{E}^{a_H}[J'(v_{i\cdot})]$$

$J'(v)$: marginal cost of an extra util at $t = 1$; $\mathbb{E}^{a_H}[J'(v_{i\cdot})]$: marginal cost of an extra util at $t = 2$.

The optimal contract optimally smooths out costs of incentive provision across time

History Dependence: Memory and Promise Keeping

Proposition

*(**Memory**) Consider an optimal contract that implements $a \in A$. Suppose that $I_i \neq I_j$. Then, there exists $i' \neq j'$ such that $I_{ii'} \neq I_{jj'}$.*

Proof: From the martingale property of $1/u'$.

Proposition

*(**Promise-keeping**) Consider an optimal contract that implements $a \in A$. Then $u_i > u_j \Leftrightarrow W_i > W_j$.*

Proof: Comes from $\frac{1}{u'(u_i)} = C_2'(W_i)$ and $C_2(\cdot)$ convex. **Intuition:** cost minimization $\Rightarrow$ prizes are spread out across time

Contents

The Problem

The principal must choose a contract $I := (I_t)_{t\geq 0}$ and a recommendation $a := (a_t)_{t\geq 0}$ with $I_t = I(q_1, ..., q_{t-1})$ and $a_t = a(q_1, ..., q_{t-1})$ (i.e. $\mathbb{F}^q$-measurable stochastic processes) such that they solve

$$\begin{aligned}
\max_{(I,a)} \quad & \mathbb{E}^a \left[r \sum_{t=1}^{\infty} \frac{q_t - I(a_t)}{(1+r)^t} \right] \\
s.t. \quad & (PC):\ \mathbb{E}^a \left[r \sum_{t=1}^{\infty} \frac{u(I_t) - h(a_t)}{(1+r)^t} \right] \geq \overline{W} \\
& (IC):\ \mathbb{E}^a \left[\sum_{t=1}^{\infty} \frac{u(I_t) - h(a_t)}{(1+r)^t} \right] \geq \mathbb{E}^{\tilde{a}} \left[\sum_{t=1}^{\infty} \frac{u(I_t) - h(\tilde{a}_t)}{(1+r)^t} \right],\ \forall \tilde{a}
\end{aligned}$$

Spear and Srivastava (Restud, 1987)

- They show that a principal-agent model in infinite horizon can be solved recursively using the **continuation value** (or promised utility) as a state variable

$$W_t := \mathbb{E}\left[r \sum_{s=t}^{\infty} \frac{u(I_s) - h(a_s)}{(1+r)^{s-t+1}} \middle| h^t \right],\ t = 1, 2, ...$$

  where $h^t := (q_1, ..., q_{t-1})$ is the history up to time t

- **Key insight**: the continuation value captures the agent's incentives to exert effort. Hence, all the incentive constraints can be replaced by IC ones that involve W_t only

- $\Rightarrow$ It is optimal for the principal to keep track only of W_t, as it summarizes all the payoff-relevant information

Recursive Formulation: One-Shot Deviations

Theorem

*(**Spear and Srivastava, 1987**) There is a one-to-one mapping between contracts* $(I_t(h^t), a_t(h^t))_{t\geq 0}$ *that are IC, and maps*

$$I(h^t), a(h^t), W(h^t)$$

such that

$$\mathbb{E}[W(h^{t+1})|h^t, a_t] - W(h^t) = r(W(h^t) - u(I(h^t)) + h(a(h^t))), \quad (PK)$$

$$\mathbb{E}[W(h^{t+1})|h^t, a_t] - rh(a(h^t)) \geq \mathbb{E}[W(h^{t+1})|h^t, \tilde{a}_t] - rh(\tilde{a}(h^t)) \quad (IC)$$

for any $\tilde{a} \in A$*, where* $W(h^t)$ *is the agent's continuation value after history* h^t*.*

Intuition: (IC) says that preventing all one-shot deviations is enough to deter any type of deviation; (PC) is just accounting

Recursive Formulation

Let $\vec{u} = (u_1, ..., u_n)$ and $\vec{W}' = (W_1', ... W_n')$. $F(W)$: the maximum benefit to the principal when the agent's reservation utility is W'. Then

$$F(W) := \max_{(\vec{u}, \vec{W}', a)} \quad \frac{1}{1+r} \sum_{i=1}^{n} \pi_i(a) \left[r(q_i - I_i) + F(W_i') \right]$$

$$s.t. \qquad (PK): \sum_{i=1}^{n} \pi_i(a) W_i' - W = r(W - u_i + h(a))$$

$$(IC): \; a \in \arg\max \sum_{i=1}^{n} \pi_i(\tilde{a}) \left[r(u_i - h(\tilde{a})) + W_i' \right]$$

Obs: If both u and a depend only on the history up to the previous period (i.e., wages in every period are paid before output is realized), (IC) if and only if

$$\mathbb{E}_t^{a_t}[W_{t+1}] - rh(a_t) \geq \mathbb{E}_t^{\tilde{a}_t}[W_{t+1}] - rh(\tilde{a}_t)$$

Computation: Phelan and Townsend (Restud, 1991)

$$F_T(w) = \max_{\Pi(a,q,c,w')} \frac{1}{1+r} \sum_{A\times Q\times C\times W_{T-1}} \Pi(a,q,c,w')(r(q-c) + F_{T-1}(w'))$$

$$\sum_{A\times Q\times C\times W_{T-1}} \Pi(a,q,c,w') = 1$$ The probabilities add up to 1

$$\forall a,q,c,w',\ \Pi(a,q,c,w') \geq 0$$ The probabilities are nonnegative

$$\forall \hat{a},\hat{q},\hat{c} \sum_{W_{T-1}} \Pi(\hat{a},\hat{q},\hat{c},w') = P(\hat{q} \mid \hat{a}) \sum_{Q\times C\times W_{T-1}} \Pi(\hat{a},q,c,w')$$

The joint probability distribution over a, q and c is consistent with the conditional probabilities P(q|a) with which output q is realized given effort a[1]

$$\sum_{A\times Q\times C\times W_{T-1}} w'\,\Pi(a,q,c,w') - w = r \sum_{A\times Q\times C\times W_{T-1}} \left(w - u(c) + h(a)\right)\Pi(a,q,c,w')$$ (PK)

$$\forall a,\hat{a},\ \sum_{Q\times C\times W_{T-1}} \left(w' - rh(a)\right)\Pi(a,q,c,w') \geq \sum_{Q\times C\times W_{T-1}} \left(w' - rh(\hat{a})\right)\frac{P(q \mid \hat{a})}{P(q \mid a)}\Pi(a,q,c,w')$$ (IC)

Bibliography

- Grossman and Hart (1983): "An Analysis of the Principal-Agent Problem," *Econometrica*
- Holmstrom (1979): "Moral Hazard and Observability," *The Bell Journal of Economics*
- Phelan and Townsend (1991): "Computing Multi-period Information-Constrained Optima," *Review of Economic Studies*
- Rogerson (1985): "Repeated Moral Hazard," *Econometrica*
- Spear and Srivastava (1987): "Computing Multiperiod Constrained Optima," *Review of Economic Studies*

Static Model: Continuum of Actions

$A = [\underline{a}, \overline{a}]$; $q \in [\underline{q}, \overline{q}]$; $f(q|a)$; I(q). Problem:

$$\max_{a \in A, I(\cdot)} \int_{\underline{q}}^{\overline{q}} [q - I(q)] f(q|a) dq$$

$$s.t. \qquad (PC): \int_{\underline{q}}^{\overline{q}} [u(I(q)) - h(a)] f(q|a) dq - \psi(a) \geq W$$

$$(IC): \; a \in \arg\max_{\tilde{a} \in A} \int_{\underline{q}}^{\overline{q}} [u(I(q)) - h(\tilde{a})] f(q|\tilde{a}) dq - \psi(\tilde{a})$$

Static Model: Continuum of Actions

- If the agent's problem is strictly concave we can replace (IC) by

$$\int_{\underline{q}}^{\overline{q}} [u(I(q)) - h'(a)] f_a(q|a)\, dq - \psi'(a) = 0.$$

- FOC becomes

$$\frac{1}{u'(I(q))} = \lambda + \mu \left[\frac{f_a(q|a)}{f(q|a)} \right]$$

- Also, f_a/f increasing in q iff $f(q|a')/f(q|a'')$ is decreasing in q if $a'' > a'$ (MLRP analog)

- Hence, $I(\cdot)$is increasing if MLRP holds

▸ Back to Multiple Actions Case