Transactions with Hidden Action: Part 1

Dr. Margaret Meyer
Nuffield College

2021

Transactions with hidden action

Risk-neutral principal (P) delegates performance of a task to agent (A)

- Key features of the situation
  1. A takes a privately-observed action, e.g. chooses effort level or level of risk
  2. A's measured performance is informative about his effort but is subject to a random shock
  3. A is risk-averse

- Examples:
  - employment contracts for workers and managers
  - sharecropping - landlords and tenants
  - manufacturing firms and subcontractors
  - litigation - clients and attorneys
  - insurance firms and customers
  - government and defence contractors
  - regulators and regulated firms

- While many of these relationships involve repeated interaction, focus today on modeling a one-shot interaction.

Transaction with hidden action

Full-information benchmark: If A's effort were observable, P could force A to choose any desired effort (by basing A's pay appropriately on his effort choice), and A would bear no risk.

If A's effort is not observable, P faces a tradeoff between provision of effort incentives and provision of insurance:

- The more strongly A's pay varies with his measured performance, the stronger A's incentives, but the more risk he is exposed to.

Sources of randomness in A's measured performance:

- imperfect monitoring of effort or subjective biases of evaluator
- uncertain consequences of effort: randomness in market conditions, equipment failure

Special cases with no loss from hidden action:

1. No random shock to performance: as if effort were observable.
2. A not risk averse: P could "sell the firm to A", i.e. pay A the realized value of output (minus a constant) $\Longrightarrow$ A would choose the socially efficient effort and, since not risk averse, would bear no risk costs.

A general principal-agent model

- Mirrlees (1975, 1979), Holmström (1979), Grossman and Hart (1983)
- Agent's preferences: $U(w, e)$, where e = effort, w = payment; $U_w > 0$ and $U_e \leq 0$; $\bar{U}$ = reservation utility
- Principal's preferences: $V(w, z)$, where $z = f(e, x)$ = output and x = random shock; $V_w < 0$, $V_z > 0$
- Timing:
  1. P offers a legally enforceable contract $w(z)$ to A.
  2. A chooses to accept or reject contract.
  3. If rejects, A gets $\bar{U}$; if accepts, A chooses privately-observed effort e.
  4. Random shock is realized, but not observed by P.
  5. P and A observe z, and P pays A $w(z)$.
- P has complete information about A's preferences; the only asymmetry in info. is "hidden action": A's effort choice e is not observable by P.
- P's optimal contract solves

$$\begin{aligned} \max_{w(z),e} \quad & E_{z|e} V(w(z), z) \\ \text{s.t.} \quad & E_{z|e} U(w(z), e) \geq \bar{U} && \text{(individual rationality (IR))} \\ & e \in \arg\max_{e'} E_{z|e'} U(w(z), e') && \text{(incentive compatibility (IC))} \end{aligned}$$

Analytical difficulties with the general model

- Difficult to know when it is valid to replace A's incentive compatibility constraint with A's FOC for effort:

$$e \text{ solves } \frac{\partial}{\partial e'} E_{z|e'} U(w(z), e') = 0.$$

  When can we be sure that A's objective function, *given* the chosen contract, is strictly concave in effort?

- Difficult to derive general results about the optimal contract $w^*(z)$:
  - *Example:* Let $z = e + x$; $e \in \{0, 1\}$; $x \in \{0, 2\}$. Then if it is socially efficient to induce $e = 1$, the following **non-monotonic** contract is an optimal one: $w^*(z) = L$ for $z \in \{0, 2\}$ and $w^*(z) = H$ for $z \in \{1, 3\}$, where $L < H$.
  - Even when effort shifts up the distribution of output, an optimal contract need not be monotonically increasing in output, because **under hidden action, output plays two distinct roles**:
    1. Output directly benefits P
    2. Output is a signal of A's effort
  - If higher output doesn't always signal higher effort, $w^*(z)$ needn't be increasing.

Linear-exponential-normal (LEN) principal-agent model

Following Holmstrom and Milgrom (1991) and Milgrom and Roberts (1992, Ch.7), we'll work with the "linear-exponential-normal" model:

- Agent's utility function: $U(w, e) = u(w - C(e))$, where
  - w = payment to agent and e = A's effort
  - $u(m) = -e^{-rm}$ (constant absolute risk aversion (CARA) = r)
  - $C(e)$ = cost of effort (including direct and opportunity cost); $C'(e) \geq 0$, $C''(e) \geq 0$, and for e s.t. $C(e) > C(0)$, $C''(e) > 0$.
  - e is chosen once and for all by A and *cannot be observed* by P
  - A's reservation utility = $\bar{U}$
- Risk-neutral principal's utility function $V(w, e) = B(e) - w$
  - $B(e)$ = expected value of P's revenue when A exerts e
  - $B(e)$ not observable in short run, so can't be contracted on
- Contract bases A's payment w on measured performance: $z = e + x$
  - random shock $x \sim N(0, \sigma_x^2)$; realized after e chosen; not observed by P
  - effort affects the mean of z, but not the variance of z
- Additional contractible signal $y \sim N(0, \sigma_y^2)$
  - y is informative about x, but unaffected by e

Linear-exponential-normal (LEN) principal-agent model

Full-info. benchmark: P rewards A if and only if A chooses $e \geq e^*$, where $B'(e^*) = C'(e^*)$ (socially efficient e); reward chosen so A's IR constraint binds.

Under hidden action:

- Assume contract takes *linear* form:
  $w = \alpha + \beta(z + \gamma y) = \alpha + \beta(e + x + \gamma y)$
  - β measures intensity of incentives for effort
  - γ determines relative weight on y and z
- Linear contracts
  1. simple to understand and administer and commonly used (e.g. sales commissions; contingency fees to attorneys; piece rates for workers)
  2. not in general optimal in a static model
  3. have the benefit of providing uniform "incentive pressure" (in contrast, with contracts paying a bonus if A surpasses an output target, A's incentives vary according to how "far" from target he is)
- Holmström and Milgrom (1987) formally analyzed benefit 3.
  - In a model where A chooses efforts continuously over time and can observe his cumulative performance, the optimal contract induces A to choose a constant effort over time and makes A's wage a linear function of the final outcome.

Linear-exponential-normal (LEN) model: optimal contract

The optimal contract, described by $(\alpha, \beta, \gamma, e)$, solves

$$\begin{aligned}
\max_{\alpha,\beta,\gamma,e} \quad & E_{x,y}\left[B(e) - (\alpha + \beta(e + x + \gamma y))\right] \\
\text{s.t.} \quad & E_{x,y}\left[u\left(\alpha + \beta(e + x + \gamma y) - C(e)\right)\right] \geq \bar{U} && (IR) \\
& e \in \arg\max_{e'} E_{x,y}\left[u\left(\alpha + \beta(e' + x + \gamma y) - C(e')\right)\right] && (IC)
\end{aligned}$$

Simplifying the optimal contracting problem

Definition: Given utility function $u(m)$, the *certainty equivalent* (CE) of the random income m is defined by $u(CE) = E_m u(m)$.

Fact: For an individual with (CARA) utility function $u(m) = -e^{-rm}$ and random income $m \sim N(\bar{m}, \sigma_m^2)$, the CE of m is $\bar{m} - \frac{1}{2} r \sigma_m^2$.

- $\frac{1}{2} r \sigma_m^2$ is the *risk premium* associated with random income m; here, it does not depend on $\bar{m}$.

Under the contract $w = \alpha + \beta(e + x + \gamma y)$,

- CE of A's income $= \alpha + \beta e - C(e) - \frac{1}{2} r \beta^2 \text{Var}(x + \gamma y)$
  - Maximizing expected utility is equivalent to maximizing CE.
  - A's CE is strictly concave in effort, so A's optimal e solves FOC $\beta = C'(e)$.
- CE of P's income $= B(e) - (\alpha + \beta e)$
- Total CE $=$ (A's CE) $+$ (P's CE) $= B(e) - C(e) - \frac{1}{2} r \beta^2 \text{Var}(x + \gamma y)$

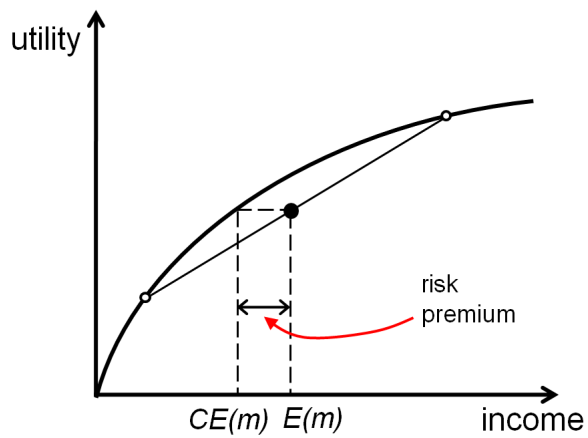

Simplifying the optimal contracting problem

- At the optimum, A's IR constraint will bind. Rewrite it in CE terms, then substitute result into P's objective function. This converts the objective into the sum of P's and A's CE's.
- Replace IC constraint with FOC for e: $\beta = C'(e)$.
- These steps imply that optimal β, γ and e solve

$$\begin{aligned} \max_{\beta,\gamma,e} \quad & B(e) - C(e) - \frac{1}{2} r\beta^2 \text{Var}(x + \gamma y) - u^{-1}(\bar{U}) \\ \text{s.t.} \quad & \beta = C'(e), \qquad (IC) \end{aligned}$$

  and optimal α makes IR constraint bind (given β, γ, e).
- NB: CARA utility for $A \Longrightarrow$ optimal β, γ and e are independent of $\bar{U}$.
- We will solve this problem in 2 steps:
  1. Given any e, what is optimal way to induce A to choose e, that is, what (β, γ) are optimal?
  2. What is the optimal level of e to induce A to choose?

Step 1: Optimal contract for given effort

- Given e, optimal β determined by IC constraint: $\beta = C'(e)$.
- Given e and β, P's maximization problem is equivalent to choosing γ to minimize $\text{Var}(x + \gamma y)$ (i.e. to minimize A's risk premium):

$$\min_{\gamma} \text{Var}(x + \gamma y) \Longleftrightarrow \min_{\gamma} \text{Var}(x) + 2\gamma \text{Cov}(x, y) + \gamma^2 \text{Var}(y)$$

- Therefore $\gamma^* = -\text{Cov}(x, y)/\text{Var}(y)$
  - If $\text{Cov}(x, y) = 0$, then $\gamma^* = 0$, so w should be independent of y: there is no benefit from making risk-averse A bear risk associated with variability of y, when y is uninformative about x and hence about e.
  - If $\text{Cov}(x, y) \neq 0$, then $\gamma^* \neq 0$.
  - If $\text{Cov}(x, y) > 0$, then $\gamma^* < 0$: a high value of y indicates that a high z is more likely to be due to a high x and not a high e, so A should get a lower w.
  - Ceteris paribus, as $\text{Var}(y) \uparrow$, $|\gamma^*| \downarrow$: place less weight on y, the more noisy it is.

Step 1: Optimal contract for given effort

Claim: With γ set optimally (i.e. $\gamma^* = -\text{Cov}(x,y)/\text{Var}(y)$),

$$x + \gamma^* y = x - E(x|y) \qquad \text{and} \qquad \text{Var}(x + \gamma^* y) = \text{Var}(x|y).$$

Derivation: For jointly normal r.v.'s x and y with correlation ρ:

$$E[x|y] = E[x] + \rho \frac{\sqrt{\text{Var}(x)}}{\sqrt{\text{Var}(y)}} (y - E[y]) \quad \text{and} \quad \text{Var}(x|y) = (1 - \rho^2)\text{Var}(x).$$

Informativeness Principle: It is optimal to base A's pay on signal y if and only if doing so allows A's effort e and random shock x to be estimated with lower error. The optimal weight minimizes the variance of the estimation error.

- P can perfectly predict A's choice of e in response to any contract.
- But P cannot prove that A chose the e that P anticipated.
- Optimal contract uses signal y to estimate x and hence e, and pays A acc. to how the actual output z differs from best prediction given y:

$$\begin{aligned} w = \alpha + \beta(e + x - E(x|y)) &= \alpha + \beta(z - E(x|y)) \\ &= \textit{constant} + \beta(z - E(z|y)), \end{aligned}$$

where $E(z|y)$ denotes the predicted value of output given observed y and anticipated choice of e. (NB: e is not a random variable.)

Application: Comparative Performance Evaluation

- Examples of comparative performance evaluation:
  - paying firm CEO relative to performance of industry or stock mkt.
  - paying division manager relative to performance of other divisions
  - grading students relative to average performance of class
  - yardstick regulation
- *Example*: 2 agents, A and B, choose efforts e_A and e_B
  - A's performance measure $z = e_A + x_A + x_C$
  - B's performance measure $y = e_B + x_B + x_C$
  - Shocks x_A, x_B (idiosyncratic), x_C (common): indep., normally dist.
- Is it better to base A's pay on z or on $z - y = e_A - e_B + x_A - x_B$?
- For a given β, both $w = \beta z$ and $w = \beta(z - y)$ induce same effort: $C'(e) = \beta$. Hence Informativeness Principle implies better to base w on measure which generates smaller risk premium:
  - $z - y$ preferable if $\text{Var}(x_A - x_B) < \text{Var}(x_A + x_C) \Leftrightarrow \text{Var}(x_B) < \text{Var}(x_C)$
  - z preferable if $\text{Var}(x_A - x_B) > \text{Var}(x_A + x_C) \Leftrightarrow \text{Var}(x_B) > \text{Var}(x_C)$

Exercise: What's the optimal statistic $z + \gamma^* y$ to base A's pay on? How does γ^* vary with $\text{Var}(x_C)$?

Application: Auto insurance vs. health insurance

- Insurance companies approximately risk neutral, customers risk averse. Consequently, in absence of informational asymmetries (here, hidden action) customer would be fully insured.
- Auto insurance policies for collisions and theft typically specify a deductible - a fixed amount the insured must pay before insurance company makes any payment. Why?
  - Owner of car can, by taking care, reduce probability that car will be stolen or damaged, but has little control over size of loss, if loss occurs.
  - Hence, size of loss is not very informative about care taken by owner.
  - Hence, by Informativeness Principle, owner's contribution should not depend on size of loss, only on whether or not loss occurs.
- In contrast, health insurance policies often require the insured to pay a fraction of the cost of any medical services used (co-payments). Why?
  - Total costs incurred are informative about how hard insured tried to conserve on health expenses.
  - Hence, by Informativeness Principle, insured's payment should depend on total costs incurred.

Step 2: Optimal level of effort to induce

- What is the optimal level of e to induce? Equivalently, since $\beta = C'(e)$, what is the optimal intensity of incentives β?
- Given a choice of γ, define $V \equiv \text{Var}(x + \gamma y)$. Substitute IC constraint $\beta = C'(e)$ into objective to get following problem:

$$\max_e \; B(e) - C(e) - \frac{1}{2} r (C'(e))^2 V$$

  - $C(e)$ is direct cost of inducing effort
  - $\frac{1}{2} r (C'(e))^2 V$ is indirect cost of inducing effort under moral hazard, arising from imposing more risk on risk-averse A
- First-order condition for optimal e for P to induce:

$$B'(e) - C'(e) - rVC'(e)C''(e) = 0$$

Incentive Intensity Principle: The optimal levels of β and e satisfy

$$\beta^* = \frac{B'(e^*)}{1 + rVC''(e^*)} \qquad \text{and} \qquad \beta^* = C'(e^*).$$

Comparative statics on optimal strength of incentives

$$\beta^* = \frac{B'(e^*)}{1 + rVC''(e^*)} \qquad \text{and} \qquad \beta^* = C'(e^*)$$

Comparative statics on β^*:

- higher $B'(e^*)$ (marginal benefit of effort) $\Longrightarrow$ higher β^*
- higher r (risk aversion) $\Longrightarrow$ lower β^*
- higher V (variance in measurement of effort) $\Longrightarrow$ lower β^*
- higher $C''(e^*)$ (inflexibility of A's effort) $\Longrightarrow$ lower β^*

If $r = 0$ or $V = 0$, then $\beta^* = B'(e^*) = C'(e^*)$, so optimal to induce socially efficient level of effort (as in full-info benchmark).

If $r > 0$ and $V > 0$, then $B'(e^*) > \beta^* = C'(e^*)$, so optimal contract in presence of hidden action induces strictly less effort than socially efficient level.

Application: incentives for Japanese subcontractors

- P: manufacturing firm, A: subcontractor supplying parts
- Contract: subcontractor is paid $x + \beta(\bar{x} - x)$
  - $\bar{x}$ is target level of cost, and x is actual level of cost incurred
  - so A gets x reimbursed and receives incentive payment $\beta(\bar{x} - x)$
- How does β depend on r, V and C'' ?
  - β estimated by comparing variance of A's profits with variance of A's cost
  - r assumed smaller for larger firms
  - V estimated by calculating variation in actual costs around trend level
  - C'' assumed smaller, the larger A's responsibility for supplying technology and designs
- Empirical findings for automobiles (Kawasaki and McMillian, 1987) and electronics (Asanuma and Kikutani, 1991): β varied in predicted direction with measures used for r, V and C''.