Principal-Agent Models and Moral Hazard

Christopher W. Miller

Department of Mathematics
University of California, Berkeley

April 11, 2014

Background and Motivation

One major goal of the bottom-up approach to economics is to model the decision-making of individuals, as well as collections of individuals in markets.

- In competitive markets with complete information, this is relatively well-understood.
- With asymmetric information, contracts between individuals can go wrong in many ways.

Background and Motivation

- Complete information
  - Efficient equilibria by Welfare Theorems
- Incomplete information
  - Adverse selection (e.g. lemon markets)
  - Moral hazard (e.g. principal-agent problem)

"All happy families are alike; each unhappy family is unhappy in its own way." - Leo Tolstoy

Background and Motivation

Goals of this talk:

- Set up a basic continuous-time principal-agent model,
- Derive conditions under which moral hazard is irrelevant,
- Cast the choice of a contract as a Hamilton-Jacobi equation.

References:

- "A Continuous-Time Version of the Principal-Agent Problem", Sannikov (2007)
- "How to build stable relationships between people who lie and cheat", Ekeland (2013)

Table of contents

A Simple Model

Consider simple contract between two parties. The first (the principal) hires the second (the agent) to work for him. The agent is in charge of a project which generates a revenue stream dX_t for the principal:

$$dX_t = A_t\,dt + \sigma\,dZ_t. \tag{1}$$

Here Z_t is a standard Brownian motion with filtration $\mathcal{F}_t^Z$ and A_t is $\mathcal{F}_t^Z$-adapted and constrained to a compact set $[0, \overline{a}]$. In exchange for work, the principal pays C_t.

A Simple Model

Assume principal is risk-neutral and agent is risk-averse, with respective utilities:

$$\text{(principal)} \qquad r\mathbb{E}\left[\int_0^\infty e^{-rt}\left(dX_t - C_t\,dt\right)\right] \tag{2}$$

$$\text{(agent)} \qquad r\mathbb{E}\left[\int_0^\infty e^{-rt}\left(u(C_t) - h(A_t)\right)\,dt\right]. \tag{3}$$

We assume that r, u, and h are known to both parties.

A Simple Model

In this case, the principal would want the agent to work as hard as possible. They would require $A_t = \overline{a}$ and compensate at $\overline{c}$ such that $u(\overline{c}) = h(\overline{A}) + \epsilon$. This gives all the power to the principal.

Condition (Moral Hazard)

Let $\mathcal{F}_t^X$ be the filtration generated by X_t. We require that the process C_t is $\mathcal{F}_t^X$-adapted.

Intuitively, this means compensation C must be a function of past outputs X, but cannot depend explicitly upon the effort A.

A Simple Model

- The agent can always blame low revenue ($dX_t < 0$) on bad luck ($dZ_t < 0$) instead of low effort ($dA_t < 0$).
- The principal would like to make $C_t = \phi(A_t)$, but cannot. He must use the information of past values of X_t to determine C_t (non-Markovian).
- The principal would like to threaten punishment if the agent is caught lying. We exclude this with a second condition.

Condition (Limited Liability)

The agent's compensation $C_t \geq 0$ for all t.

Some Examples

Consider a principal owns a farm in the remote countryside and hires an agent to tend to the farm. What are some contracts he can offer?

1. $C_t = c$ (fixed salary)
2. $C_t = aX_t^+$ for $0 \leq a \leq 1$ (share-cropping)
3. $C_t = (X_t - c)^+$ for $c \geq 0$ (farming)

Share-cropping and farming are geographically and historically widespread. Fixed salary would lead to the agent having no incentive to work.

Some Examples

Moral hazard and limited liability are fundamental to modern finance industry:

1. Fund managers effort is not easily observable by clients; poor returns may be blamed on the market (moral hazard)
2. Losses come out of the pocket of the client, not the manager (limited liability)

Problem Statement

Definition

A contract is a pair (C_t, A_t), *with* $C_t \geq 0$ *and* $\mathcal{F}_t^X$*-adapted, while* A_t *is* $\mathcal{F}_t^Z$*-adapted.*

Definition

A contract (C_t, A_t) *is incentive-compatible (IC) if the effort process* A_t *maximizes the agent's total expected utility.*

Problem Statement

Definition

A contract (C_t, A_t) is individually-rational (IR) if both the principal and agent have:

$$\text{(principal)} \qquad r\mathbb{E}\left[\int_0^\infty e^{-rt}\left(dX_t - C_t\,dt\right)\right] \geq 0 \tag{4}$$

$$\text{(agent)} \qquad r\mathbb{E}\left[\int_0^\infty e^{-rt}\left(u(C_t) - h(A_t)\right)\,dt\right] \geq 0. \tag{5}$$

Goal

We want to find a contract (C_t, A_t), which is both incentive-compatible and individually-rational, that maximizes the principal's expected utility.

Continuation Value

Let (C_t, A_t) be a contract. We want a characterization of when this contract is incentive-compatible. Consider the following natural process:

Definition

We define the continuation value the agent derives from following the contract (C_t, A_t) at time t as:

$$W_t = r\mathbb{E}\left[\int_t^\infty e^{-r(s-t)}\left(u(C_s) - h(A_s)\right) ds \mid \mathcal{F}_t^Z\right]. \tag{6}$$

It is useful to derive an SDE governing this process.

Continuation Value

Lemma

There exists a unique $\mathcal{F}_t^Z$-adapted process Y_t such that:

$$dW_t = r\,(W_t - u(C_t) + h(A_t))\,dt + r\sigma Y_t\,dZ_t \tag{7}$$

$$= r\,(W_t - u(C_t) + h(A_t) - Y_t A_t)\,dt + rY_t\,dX_t. \tag{8}$$

One key observation is that both the agent and the principal can evolve W_t using (7) and (8) respectively.

Continuation Value

Proof

The idea is to re-write W_t such that a martingale shows up, then use the Martingale Representation Theorem. In particular:

$$\begin{aligned} W_t &= r\mathbb{E}\left[\int_t^\infty e^{-r(s-t)}\left(u(C_s) - h(A_s)\right) ds \mid \mathcal{F}_t^Z\right] \\ &= re^{rt}\mathbb{E}\left[\int_0^\infty e^{-rs}\left(u(C_s) - h(A_s)\right) ds \mid \mathcal{F}_t^Z\right] \\ &\quad - re^{rt}\int_0^t e^{-rs}\left(u(C_s) - h(A_s)\right) ds. \end{aligned}$$

Continuation Value

Under some growth assumptions, the expectation is a continuous $\mathcal{F}_t^Z$-martingale, so there exists unique $\mathcal{F}_t^Z$-adapted process Y_t such that:

$$e^{-rt} W_t = rW_0 + r \int_0^t \sigma e^{-rs} Y_s \, dZ_s - r \int_0^t e^{-rs} \left(u(C_s) - h(A_s) \right) \, ds.$$

By Ito's lemma, we see:

$$-re^{-rt} W_t \, dt + e^{-rt} \, dW_t = r\sigma e^{-rt} Y_t \, dZ_t - re^{-rt} \left(u(C_t) - h(A_t) \right) \, dt$$

$$\implies dW_t = r \left(W_t + u(C_t) - h(A_t) \right) \, dt + r\sigma Y_t \, dZ_t.$$

Continuation Value

Next, note that because the agent is assumed to follow the contract, we have:

$$Z_t = \frac{1}{\sigma}\left(X_t - \int_0^t A_s\, ds\right),$$

so we can replace the dZ_t in the previous SDE to get:

$$dW_t = r\left(W_t + u(C_t) - h(A_t) - Y_t A_t\right) dt + rY_t\, dX_t.$$

□

Girsanov Theorem

In the proceeding, we need to consider an alternative effort A_t^*. Formally, A_t and A_t^* generate different probability laws for X_t. We call them $\mathbb{P}^A$ and $\mathbb{P}^{A^*}$ respectively. Now we have:

$$\begin{cases} Z_t^A = \frac{1}{\sigma}\left(X_t - \int_0^t A_s\, ds\right) & \mathbb{P}^A\text{-Brownian motion} \\ Z_t^{A^*} = \frac{1}{\sigma}\left(X_t - \int_0^t A_s^*\, ds\right) & \mathbb{P}^{A^*}\text{-Brownian motion} \end{cases}$$

related by:

$$Z_t^{A^*} = Z_t^A + \int_0^t \frac{A_s - A_s^*}{\sigma}\, ds. \tag{9}$$

Girsanov Theorem

Theorem (Girsanov)

Let Z_t be a $\mathbb{P}$-Brownian motion and a_t an adapted process. Define a new process by:

$$M_t = \exp\left(\int_0^t a_s\, dZ_s - \frac{1}{2}\int_0^t a_s^2\, ds\right).$$

If M is uniformly integrable, then a new measure $\mathbb{Q}$, equivalent to $\mathbb{P}$, may be defined by:

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = M_\infty.$$

Furthermore, we have:

$$Z_t - \int_0^t a_s\, ds \text{ is a } \mathbb{Q}\text{-Brownian motion.}$$

Key Theorem

Theorem (Characterization of IC Contracts)

A contract (C_t, A_t) is IC if and only if

$$Y_t A_t - h(A_t) = \max_a \{aY_t - h(a)\} \ a.e. \tag{10}$$

Proof

First, consider an arbitrary alternative strategy A_t^* and the process corresponding to the return from following A_t^* then switching to A_t:

$$V_t = r \int_0^t e^{-rs} \left(u(C_s) - h(A_s^*)\right) ds + e^{-rt} W_t.$$

Key Theorem

Our goal is to show V_t is a $\mathbb{P}^{A^*}$-submartingale. By Ito's lemma and (7):

$$\begin{aligned} dV_t &= re^{-rt}\left(u(C_t) - h(A_t^*) - W_t\right) dt + e^{-rt}\, dW_t \\ &= re^{-rt}\left(u(C_t) - h(A_t^*) - W_t\right) dt \\ &\quad + re^{-rt}\left(W_t - u(C_t) + h(A_t)\right) dt \\ &\quad + r\sigma e^{-rt} Y_t\, dZ_t^A \\ &= re^{-rt}\left(h(A_t) - h(A_t^*)\right) dt + r\sigma e^{-rt} Y_t\, dZ_t^A. \end{aligned}$$

Key Theorem

Next, by (9), we rewrite this as:

$$dV_t = re^{-rt}\left(h(A_t) - h(A_t^*) + A_t^* - A_t\right) dt + r\sigma e^{-rt} Y_t \, dZ_t^{A^*}. \quad (11)$$

Then if A_t does not satisfy (10) almost everywhere, choose A_t^* that does. Then the $\mathbb{P}^{A^*}$-drift of V_t is positive and non-zero on a set of positive measure, so it is a strict submartingale. In particular, there exists t large enough that:

$$\mathbb{E}^{A^*}[V_t] > V_0 = W_0.$$

Therefore, we conclude if A_t does not satisfy (10), then the contract is not IC.

Key Theorem

Alternatively, suppose A_t does satisfy (10). Then for any alternative strategy A_t^*, the expected gain process V_t is a supermartingale by (11). It is easy to see it is bounded below, so there exists V_∞ such that:

$$W_0 = V_0 \geq \mathbb{E}^{A^*}[V_\infty].$$

So the expected payoff of A is at least as good as that of A^*. Therefore, the contract is IC.

□

Control Formulation

Now if we assume for simplicity that the maximum in (10) occurs in the interior, we conclude:

$$Y_t = h'(A_t).$$

The maximization of the principal's expected profit over IC contracts then becomes:

$$\begin{aligned} \max_{(C_t, A_t)} & \left\{ r\mathbb{E}\left[\int_0^\infty e^{-rt}\left(dX_t - C_t\,dt\right)\right]\right\} \\ \text{s.t. } & dW_t = r\left(W_t - u(C_t) + h(A_t)\right) + \sigma h'(A_t)\,dZ_t \\ & dX_t = A_t\,dt + \sigma\,dZ_t \\ & W_0 = w_0,\ X_0 = x_0. \end{aligned}$$

Hamilton-Jacobi Equation

These notes will be completed for next week...