Lecture 6: The Principal-Agent Problem in Continuous Time

Gonzalo Cisternas
MIT

Contents

Introduction

- Reference: “A Continuous-Time Version of the Principal-Agent Problem" by Y. Sannikov (REStud 2008)
- Focus on the dynamics of a contractual relation between a risk-neutal principal and a risk-averse agent in continuous-time
- Why continuous-time?
  - Tractability
  - New economic insights
  - Computational power

Contents

Model

- A risk-neutral principal and a risk-averse agent sign a contract at time 0 and interact over $[0, \infty)$
- While working at the firm, the agent produces an output $X := (X_t)_{t \geq 0}$ which evolves according to

$$dX_t = A_t dt + \sigma dZ_t,$$

where

  - A_t is the agent's (hidden) effort choice at time $t \geq 0$
  - Z_t is a $\mathbb{F}-$Brownian motion. Integral notation:

$$X_t = X_0 + \int_0^t A_s ds + \int_0^t \sigma dZ_s$$

- The principal observes X only. Let $\mathbb{F}^X := (\mathcal{F}_t^X)_{t \geq 0}$ the associated filtration

Principal's Payoff

- The principal provides the agent with consumption C_t at time $t \geq 0$. The process $C := (C_t)_{t\geq 0}$ must be $\mathbb{F}^X-$progressively measurable (a bit more than adapted)

- The principal's expected discounted payoff at time zero is

$$\Pi_0 = \mathbb{E}_0^A \left[r \int_0^\infty e^{-rs}(dX_s - C_s ds) \right]$$

  when the agent follows the strategy $A := (A_t)_{\geq 0}$

- Substituting $dX_s = A_s ds + \sigma dZ_s$ and using

$$\mathbb{E}_0 \left[\int_0^\infty e^{-rs} \sigma dZ_s \right] = 0$$

  we get

$$\Pi_0 = \mathbb{E}_0 \left[r \int_0^\infty e^{-rs}(A_s - C_s) ds \right]$$

Agent's Payoff

- The agent chooses actions in $\mathcal{A}$, compact and $\min \mathcal{A} = 0$
- His flow utility from consumption C_t and effort A_t at time t

$$u(C_t) - h(A_t)$$

  where
  - $u : [0, \infty) \to [0, \infty)$ is increasing, concave and C^2. $u'(c) \to 0$ as $c \to \infty$
  - $h : \mathcal{A} \to [0, \infty)$ is increasing, convex (disutility of effort).
  - There is $\gamma_0 > 0$ such that $h(a) \geq \gamma_0 a$ for all $a \in \mathcal{A}$.
  - Normalization: $u(0) = h(0) = 0$.
- Agent's expected discounted payoff at $t = 0$ when he follows $A := (A_t)_{t \geq 0}$

$$\mathbb{E}_0^A \left[r \int_0^\infty e^{-rs} [u(C_s) - h(A_s)] ds \right].$$

- A is progressively measurable w.r.t. $\mathbb{F}$

The Principal's Problem

The principal's problem consists of choosing a contract $C := (C_t)_{t\geq 0}$ and a recommendation $A := (A_t)_{t\geq 0}$ such that they solve

$$\begin{aligned}
\max_{(C,A)} \quad & \mathbb{E}_0^A\left[r\int_0^\infty e^{-rs}(A_s - C_s)ds\right] \\
s.t. \quad & (PC_A):\ \mathbb{E}_0^A\left[r\int_0^\infty e^{-rs}[u(C_s) - h(A_s)]ds\right] \geq W_0 \\
& (IC_A):\ A \in \arg\max_{\tilde{A}}\ \mathbb{E}_0^{\tilde{A}}\left[r\int_0^\infty e^{-rs}[u(C_s) - h(\tilde{A}_s)]ds\right]
\end{aligned}$$

Complicated problem: Two optimization problems one embedded into another one. Contracting space is quite large ($\mathbb{F}^X$-prog. measurable processes) and compensation can depend on the whole current path of X's

Contents

Towards a Recursive Formulation of the Problem

- Idea: Derive a version of the **one-shot deviation principle** for this context
- $\Rightarrow$ Reduce the set of incentive constraints to a set of local (one-shot) incentive constraints
- The gains/losses of those one-shot deviations are easily understood by analyzing the continuation value process $W := (W_t)_{t\geq 0}$
- $\Rightarrow$ Continuation value carries all the payoff-relevant information for the principal's problem, so we can solve the problem using stochastic control with state variable $W := (W_t)_{t\geq 0}$

Continuation Value Process

Given a contract C and an effort process A, the agent's continuation value process $W := (W_t)_{t\geq 0}$ is defined as

$$\mathbb{E}^A \left[\int_t^\infty e^{-r(s-t)} [u(C_s) - h(A_s)] ds \Big| \mathcal{F}_t \right], \ t \geq 0.$$

Obs 1: W_t depends on future effort levels $(A_s)_{s\geq t}$ and future consumption levels $(C_s)_{s\geq t}$. Hence $W_t = W_t(C, A)$.

Obs 2: The continuation value is completely determined by the history $(X_s)_{s\leq t}$ if C and A are $\mathbb{F}^X-$measurable.

W: Representation as a SDE

Theorem

*(**Sannikov, 2008**) For any pair $(C_t, A_t)_{t\geq 0}$ that gives finite utility to the agent:*

(a) *The agent's continuation payoff $W_t(C, A)$ satisfies the SDE*

$$dW_t = r(W_t - u(C_t) + h(A_t))dt + rY_t[dX_t - A_t dt]$$

for some process Y_t s.t. $\mathbb{E}[\int_0^t Y_s^2 ds] < \infty$ for all $t \geq 0$ and

(b) *W_t satisfies the transversality condition $\lim_{t\to\infty} \mathbb{E}_t^A[e^{-rs}W_{t+s}] = 0$.*

Conversely, a process that satisfies the previous SDE and the TVC, then it is the agent's continuation value.

Martingale Representation Theorem

Before proving the theorem, we need an important result

Theorem

*(**Martingale Representation Theorem**) Let* $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ *and* Z *be a* $\mathbb{F}-$ *B.M. Suppose that* $M := (M_t)_{t\geq 0}$ *is a martingale that is square integrable. Then, there exists* θ *with* $\mathbb{E}[\int_0^t \theta_s^2 ds] < \infty$ *for all* $t \geq 0$*, such that*

$$M_t = M_0 + \int_0^t \theta_s dZ_s,\ t \geq 0.$$

Proof ($\Rightarrow$)

- Fix a strategy A and let $\mathcal{Q}^A$ denote the probability distribution over paths of X induced by A. Let $\Omega = C([0,\infty))$ The filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathcal{Q}^A)$ is fixed throughout this proof
- Fix a consumption process C and consider the process

$$V_t := \mathbb{E}^A \left[\int_0^\infty e^{-rs} r(u(C_s) - h(A_s))ds \middle| \mathcal{F}_t \right]$$

  Then, $V := (V_t)_{t\geq 0}$ is a martingale under $\mathcal{Q}^A$
- Notice that when the agent follows A, the process

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

  is a Brownian motion under $\mathcal{Q}^A$

Proof ($\Rightarrow$)

- By the martingale representation theorem, there exists a process $Y := (Y_t)_{t\geq 0}$ such that

$$V_t = V_0 + \int_0^t e^{-rs}\sigma Y_s dZ_t^A$$

- Now, since $V_t = \int_0^t e^{-rs} r(u(C_s) - h(A_s))ds + e^{-rt}W_t(C,A)$ (*)

$$dV_t = re^{-rt}Y_t\sigma dZ_t^A \quad = \quad re^{-rt}(u(C_t) - h(A_t))dt \underbrace{-re^{-rt}W_t(C,A)dt + e^{-rt}dW_t(C,A)}_{d(e^{-rt}W_t(C,A)}$$

  From where we conclude that

$$dW_t(C,A) = r(W_t(C,A) - u(C_t) + h(A_t)) + rY_t\sigma dZ_t^A.$$

- The transversality condition follows from (*) and the dominated convergence theorem

Proof ($\Leftarrow$)

- Conversely, suppose that process W satisfies the SDE

$$dW_t = r(W_t - u(C_t) + h(A_t))dt + rY_t[dX_t - A_t dt]$$

and the TVC. Then,

$$V_t := \mathbb{E}^A \left[\int_0^\infty e^{-rs}(u(C_s) - h(A_s))ds \middle| \mathcal{F}_t \right]$$

is a martingale if the agent follows A, as $dV_t = rY_t \sigma dZ_t^A$ and $Y \in \mathcal{H}^2$ and Z^A is a B.M. under $\mathcal{Q}^A$.

- Hence,

$$W_0 = V_0 = \mathbb{E}^A[V_t] = \mathbb{E}^A \left[\int_0^t e^{-rs} r(u(C_s) - h(A_s))ds \right] + \mathbb{E}^A[e^{-rt} W_t]$$

- The TVC and the dominated convergence theorem yield

$$W_0 = \mathbb{E}^A \left[\int_0^\infty e^{-rs} r(u(C_s) - h(A_s))ds \right] \qquad \square$$

The Agent's Incentives and Sensitivity Process Y

$$dW_t(C, A) = r(W_t(C, A) - u(C_t) + h(A_t)) + rY_t\sigma[dX_t - A_t dt]$$

- Y: **sensitivity of the agent's continuation value** to changes in output. Larger $Y \Rightarrow W$ is more volatile $\Rightarrow$ larger long-term risk
- $Y_t = Y_t(C, A)$; hence, it can be **controlled** by the principal
- Next: W carries all the information that is relevant to understand the agent's incentives via Y. Formally:
  - To implement A, the principal can restrict to constraints that differ from A only over a small interval of time (**one-shot deviation principle**)
  - Such one-shot deviations can be understood studying W
  - The agent's incentives (i.e., which effort he will choose given C) are intimately related to Y. Recall: High effort requires exposure to risk

The Agent's Incentives

Theorem

*(**Sannikov, 2008**) Fix a contract $C := (C_t)_{t\geq 0}$. Then, $A := (A_t)_{t\geq 0}$ is optimal for the agent (i.e., it satisfies (IC)) is and only if*

$$Y_t A_t - h(A_t) \geq Y_t \tilde{A}_t - h(\tilde{A}),\ \forall \tilde{A}, t \geq 0,$$

where Y is the sensitivity process of $W(C, A)$.

Intuition: Infinitesimal version of Spear and Srivastava (1987)

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

If the deviates to $\tilde{A}$ only over $[t, t+dt)$, say $\tilde{A}_t > A_t$, he gains

$$\underbrace{\mathbb{E}_t^{\tilde{A}}[dW_t(C,A)] - \mathbb{E}_t^{A}[dW_t(C,A)]}_{\text{gain from local deviation}} = \underbrace{rY_t(\tilde{A}_t - A_t)dt}_{\text{higher output}} + \underbrace{(h(A_t) - h(\tilde{A}))dt}_{\text{additional effort}}$$

Proof

- Fix C and consider the process

$$\tilde{V}_t = \underbrace{r\int_0^t e^{-rs}(u(C_s) - h(\tilde{A}_s))ds}_{\text{follow } \tilde{A} \text{ in } [0,t]} + \underbrace{e^{-rt}W_t(C,A)}_{\text{switch to } A \text{ thereafter}}$$

$$\begin{aligned}\Rightarrow d\tilde{V}_t &= \underbrace{re^{-rt}(u(C_t) - h(\tilde{A}_t))dt}_{\text{drift of } \tilde{V}} - re^{-rt}W_t \\ &\quad + e^{-rt}\underbrace{\left[r(W_t - u(C_t) - h(A_t))dt + rY_t\sigma dZ_t^A\right]}_{dW_t(C,A)}\end{aligned}$$

- But when the agent deviates to $\tilde{A}$, the process

$$Z_t^{\tilde{A}} = \frac{1}{\sigma}\left(X_t - \int_0^t \tilde{A}_s ds\right)$$

is a $\mathcal{Q}^{\tilde{A}}-$Brownian motion. Furthermore, Z_t^A and $Z_t^{\tilde{A}}$ are related via

$$\sigma Z_t^A = \sigma Z_t^{\tilde{A}} + \int_0^t (\tilde{A}_s - A_s)ds$$

Proof: Only if Part ($\Rightarrow$)

- Hence,

$$d\tilde{V}_t = re^{-rt} \underbrace{[(h(A_t) - h(\tilde{A}_t)) + re^{-rt} Y_t(\tilde{A}_t - A_t)]dt}_{(*)} + re^{-rt} Y_t \sigma dZ_t^{\tilde{A}}$$

- ($\Rightarrow$) Suppose that A_t does not maximize $Y_t a - h(a)$, over A, $t \geq 0$. Choose $\tilde{A}$ that maximizes the previous expression. Then, $(*) \geq 0$ with strict inequality over a set of positive measure under $\mathcal{Q}^{\tilde{A}}$. Hence,

$$\mathbb{E}^{\tilde{A}}[\tilde{V}_t] > \tilde{V}_0 = W_0(C, A), \text{ some } t \geq 0.$$

- Consequently, the strategy "follow $\tilde{A}$ up to time t and A thereafter" yields more utility than A, so the latter is sub-optimal.

Proof: If Part ($\Leftarrow$)

- If A maximizes $Y_t a - h(a)$, drift of $\tilde{V} \leq 0$. Thus, $\tilde{V}$ is a $\mathcal{Q}^{\tilde{A}}-$supermartingale
- Since $W_t(C, A)$ is bounded from below, we can add

$$\tilde{V}_\infty := r \int_0^\infty e^{-rs}(u(C_s) - h(\tilde{A}_s))ds$$

as its last element. Therefore,

$$W_0(C, A) = \tilde{V}_0 \geq \mathbb{E}^{\tilde{A}}[\tilde{V}_\infty] = W_0(C, \tilde{A}),$$

which implies that A is indeed optimal. $\Box$

Summary: Controlled Dynamics

Theorem

There is a one-to-one correspondence between pairs of contracts and incentive-compatible effort strategies pairs (C, A) *which give the agent finite utility, and controlled processes*

$$dW_t = r(W_t - u(C_t) + h(a(Y_t)))dt + rY_t \sigma dZ_t^A$$

with controls C, Y *that satisfy the transversality condition* $\lim_{s \to \infty} \mathbb{E}^A[e^{-rsW_{t+s}}] = 0$*, where* $A_t = a(Y_t)$*,* $t \geq 0$*, and*

$$a(Y_t) = \arg\max_{a \in A} Y_t a - h(a).$$

The Principal's Optimal Control Problem

- Recall that exposing the agent to risk $(Y \neq 0)$ is costly to the principal. Let

$$\gamma(a) = \inf\{Y \geq 0 | aY - h(a) \geq \tilde{a}Y - h(\tilde{a}),\ \forall \tilde{a} \in A\}$$

  i.e., $\gamma(a)$ is the minimum sensitivity that implements a. Observe that γ is increasing: higher effort requires more powerful incentives

- The principal's problem becomes

$$\begin{aligned} \max_{C,A} \quad & \mathbb{E}^A \left[\int_0^\infty e^{-rt}(A_s - C_s) ds \right] \\ s.t. \quad & dW_t = r(W_t - u(C_t) + h(A_t))dt + r\gamma(A_t)\sigma dZ_t^A, \\ & W_0 = \bar{W} \end{aligned}$$

  where Z^A is an exogenous B.M. from the principal's perspective.

- Let $F(\overline{W})$ denote the value of this problem

HJB Equation and Boundary Conditions

$$rF(W) = \sup_{a,c}\{a - c + r(W - u(c) + h(a))F'(W) + \frac{1}{2}r^2\gamma(a)^2\sigma^2 F''(W)\}$$

- Domain of W? $W \in \mathcal{W} := [0, u(\infty))$, where $u(\infty) := \lim_{c\to\infty} u(c)$, as the agent can always guarantee himself $W \geq 0$ by putting zero effort.
- Observe that the principal can always **retire** the agent, i.e., she can offer a constant consumption level c such that

$$u(c) = W$$

  and the agent exerts zero effort. Let $F_0(u(c)) = -c$ denote the profit under that contract
- Thus, since the only way in which the principal can deliver $W = 0$ is by retiring the agent at $W = 0$,

$$F(0) = F_0(0) = 0 \ \rightarrow \text{first boundary condition}$$

HJB Equation and Boundary Conditions

$$rF(W) = \sup_{a,c}\{a - c + r(W - u(c) + h(a))F'(W) + \frac{1}{2}r^2\gamma(a)^2\sigma^2 F''(W)\}$$

- Observe that for any $W \geq 0$,

$$F(W) \geq \underbrace{F_0(W) = -u^{-1}(W)}_{\text{profit when agent is retired}}$$

- Hence, Sannikov looks for the a solution F to the previous ODE that is defined over the largest interval $[0, W_{gp}]$ with the property that $F \geq F_0$ and such that

$$(BC):\ F(0) = 0,\ \underbrace{F(W_{gp}) = F_0(W_{gp})}_{\text{value matching}} \text{ and } \underbrace{F'(W_{gp}) = F_0'(W_{gp})}_{\text{smooth pasting}}$$

Big Result: Sannikov's Verification Theorem

Theorem

*(**Sannikov**, 2008) There exists a unique solution $F \geq F_0$ to the HJB equation satisfying the boundary conditions (BC) for some $W_{gp} \geq 0$. Such F is concave and, furthermore, $F(\cdot)$ corresponds to the principal's value function. For $\overline{W} \in (0, W_{gp})$, and as long as $W_t \in (0, W_{gp})$, the optimal contract is characterized by payments $C_t = c(W_t)$ and recommended effort $A_t = a(W_t)$, with $a(\cdot)$ and $c(\cdot)$ the maximizers in the HJB, and W is the controlled process governed by*

$$dW_t = r(W_t - u(C_t) + h(A_t))dt + r\gamma(A_t)[dX_t - A_t dt],\ t > 0,\ W_0 = \overline{W}$$

with $\gamma(A_t) > 0$, until the time τ at which W_τ hits either 0 or W_{gp}. For $W > W_{gp}$, F_0 is an upper bound to the principal's profit and hence it is optimal for the principal to retire the agent. In that case, the agent gets constant consumption equal to $-F_0(W_\tau)$ thereafter.

Numerical Examples

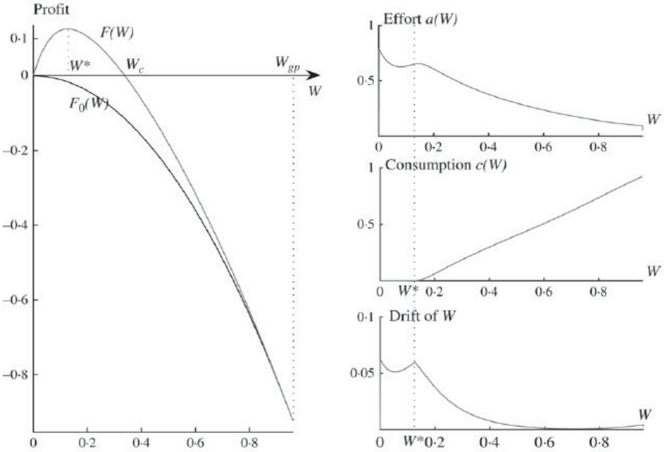

FIGURE 1

Function F for $u(c) = \sqrt{c}$, $h(a) = 0{\cdot}5a^2 + 0.4a$, $r = 0{\cdot}1$ and $\sigma = 1$. Point W^* is the maximum of F

Reference: Sannikov (2008)

Analysis: Policy Functions

- $a(W)$ maximizes

$$\underbrace{a}_{\text{benefit from extra output}} + \underbrace{h(a)F'(W)}_{\text{gain/loss of compensating effort}} + \underbrace{\frac{1}{2}F''(W)r\sigma^2\gamma(a)^2}_{\text{cost of risk exposure}<0}$$

$$\Rightarrow \quad 1 + F'(W)h'(a(W)) + F''(W)r\sigma^2\gamma(a(W))\gamma'(a(W)) = 0.$$

- $c(W)$ maximizes $-c + F'(W)u(c)$

$$\Rightarrow \quad -1 - F'(W)u'(c(W)) = 0, \text{ if } c(W) > 0 \qquad (1)$$

- Note that $c(W) = 0$ if $-F'(W) < 1/u'(0)$: marginal cost of providing incentives through W is smaller than the mg. cost of doing so through flow utility. Let W^{**} such that $-F'(W^{**}) < 1/u'(0)$
  - In $[0, W^{**}]$, the drift of W is maximized in order to make it drift away from the inefficient retirement point $W = 0$

Analysis: Dynamics of W and Retirement

- From the envelope theorem in $(0, W_{gp})$:

$$0 = rF''(W)[W - u(c(W)) + h(a(W))] + \frac{1}{2}r^2\sigma^2\gamma(a(W))^2 F'''(W)$$

  Since $F'' < 0$, sign$(W - u(c(W)) + h(a(W)))$ = sign$(F'''(W))$.

  - Hence, drift(W_t) points in the direction in which $F''(W_t)$ increases, i.e. in the direction in which **it is cheaper to provide incentives!** consumption is chosen so as to minimize the cost of incentive provision

- Why retirement at W_{gp}? **Income effect:** For large flow payments it becomes too expensive to compensate effort:
  - $c(\cdot)$ is increasing in $[0, W_{gp}]$
  - Cost of providing flow utility $1/u'(c)$ grows indefinitely with c, yet cost of output stays bounded above zero $h(a) \geq \gamma_0 a$
  - If $h''(0) = 0$ for instance, F approaches F_0 asymptotically

Analysis: Martingale Property of Inverse Marginal Utility

Since $-1/u(c(W_t)) = F'(W_t)$ when consumption is positive, by Ito's lemma:

$$\begin{aligned} \text{drift}\left(\frac{1}{u'(c(W_t))}\right) &= \text{drift}(-F'(W_t)) \\ &= -rF''(W_t)[W - u(c(W)) + h(a(W))] \\ &\quad -\frac{1}{2}r^2\sigma^2\gamma(a(W))^2F'''(W) \\ &= 0 \text{ (previous slide!)} \end{aligned}$$

Proposition

In the region $[W^{**}, W_{gp}]$ the inverse of the agent's marginal utility $\frac{1}{u'(c_t)}$ follows a martingale.