Optimal Long-Term Contracting with Learning

Zhiguo He — Chicago Booth & NBER
Bin Wei — Federal Reserve Board
Jianfeng Yu — U. of Minnesota

January 2014, AFA

Motivation

- Many long-term contractual relationships feature uncertainty and learning: e.g., unknown project's quality or agent's ability
- Dynamic contracting with uncertainty & learning: empirically relevant but theoretically challenging
- Moral hazard interacts with learning — persistent *belief manipulation* effect
  - Agent's shirking lowers today's output
    $\Longrightarrow$ Principal mistakenly thinks the project is worse, and belief distortion persists
    $\Longrightarrow$ Agent keeps getting more compensations (information rent) in the future
- Incentive provisions are intertemporally linked
  - higher future incentives $\Longrightarrow$ more information rents $\Longrightarrow$ more shirking today
  - not in typical dynamic agency literature with repeated moral hazard

This Paper

- We introduce uncertainty into an infinite-horizon Holmstrom-Milgrom setting
  - both transitory "hidden action" and persistent "hidden information"
- We derive the optimal contract, which best trades off working incentive provision against information rent extraction
- Our model is tractable so that the optimal contract is fully characterized by an ordinary differential equation (ODE) w.r.t. information rent only

Main Results

- The optimal effort policy is fully stochastic, surprising in a CARA-Normal framework
  - a result of both long-term contracting and learning
    - w/o learning, Holmstrom-Milgrom (1987): constant optimal effort
    - w. learning but short-term contracts, Holmstrom (1999): deterministic effort
- Under the optimal contract, Agent works harder in earlier periods
  - otherwise stationary framework
  - but later incentive provisions lead to greater information rents
- Option-like feature: incentives increase following good performance
  - mitigating Agent's belief-manipulation incentives
  - new insight on the popularity of option-based compensation

Literature Review

- Dynamic agency models with hidden information and learning
  - DeMarzo and Sannikov (2008) and Prat and Jovanovic (2010)
  - Adrian and Westerfield (2009)

- Short-term contracting with adverse selection and moral hazard
  - Laffont and Tirole (1988), Baron and Besanko (1987), Freixas, Guesnerie, and Tirole (1985), etc.

- Long-term contracting with adverse selection and moral hazard
  - Sannikov (2007), Fong (2009), Gershkov and Perry (2012), Halac, Kartiz, and Liu (2012), etc.

The Model

- An infinite-horizon continuous-time principal-agent model
- Principal hires Agent to manage a project with output

$$dY_t = \Big(\underbrace{\mu_t}_{\text{unobservable effort}} + \underbrace{\theta}_{\text{unknown quality/ability}} \Big) dt + \sigma dB_t$$

- θ is *unknown* quality (or managerial ability), $B_\theta \perp B$:

$$d\theta_t = \sigma \phi dB_{\theta,t}$$

  - common prior $\theta \sim \mathcal{N}(m_0, \Sigma_0)$, posterior

$$m_t^\mu = \mathbb{E}^\mu \left[\theta \,\middle|\, y^t, \mu^t \right]$$

  - stationary Bayesian learning: assume $\Sigma_0 = \sigma^2 \phi$

$$\Sigma_t = Var \left[\theta \,\middle|\, y^t, \mu^t \right] = \Sigma_0 \text{ for all } t$$

The Model (Cont'd)

- Principal is risk-neutral & offers a contract $\{c_t, \mu_t\}_{t=0}^{\infty}$
- Risk averse Agent with **exponential utility** & reservation value v_0 chooses actual effort $\widehat{\mu}_t$ and actual consumption $\widehat{c}_t$

$$\mathbb{E}^{\widehat{\mu}}\left[\int_0^{\infty} e^{-rt} u\left(\widehat{c}_t, \widehat{\mu}_t\right) dt\right]$$

where $$u\left(\widehat{c}_t, \widehat{\mu}_t\right) = -\frac{1}{a}\exp\left[-a\left(\widehat{c}_t - g\left(\widehat{\mu}_t\right)\right)\right]$$

  - moral hazard: effort is unobservable and costly

$$g\left(\widehat{\mu}_t\right) = \frac{1}{2}\widehat{\mu}_t^2$$

  - private savings account with balance S_t:

$$dS_t = rS_t dt + c_t dt - \widehat{c}_t dt \text{ with } S_0 = 0$$

    - private saving in CARA framework helps tractability

The Principal-Agent Problem

- Agent. Given contract with wages and recommended effort $\{c_t, \mu_t\}$

$$\max_{\{\widehat{c},\widehat{\mu}\}} \mathbb{E}^{\widehat{\mu}}\left[\int_0^\infty e^{-rt} u\left(\widehat{c}_t, \widehat{\mu}_t\right) dt\right]$$

$$\text{s.t.} \quad dY_t = \left(\widehat{\mu}_t + m_t^{\widehat{\mu}}\right) dt + \sigma dB_t^{\widehat{\mu}}$$

$$dS_t = rS_t dt + c_t dt - \widehat{c}_t dt \text{ with } S_0 = 0$$

  - w.l.o.g, focus on incentive-compatible & no-savings contracts
- Principal.

$$\max_{\{c,\mu\}} \quad \mathbb{E}_0^{\mu}\left[\int_0^\infty e^{-rt} \left(dY_t - c_t dt\right)\right]$$

$$\text{s.t.} \quad dY_t = \left(\mu_t + m_t^{\mu}\right) dt + \sigma dB_t^{\mu}$$

$$\mathbb{E}_0^{\mu}\left[\int_0^\infty e^{-rt} u\left(c_t, \mu_t\right) dt\right] = v_0$$

$\{c_t, \mu_t\}$ satisfy incentive-compatible & no-savings constraints

Belief Manipulation

- If Agent follows $\{\mu_t\}$ recommended by Principal (on-equilibrium), both share same posterior $m_t^\mu = \mathbb{E}^\mu \left[\theta \,|y^t, \mu^t\right]$:

$$dm_t^\mu = \phi \left[dY_t - \left(\mu_t + m_t^\mu \right) dt \right]$$

Belief Manipulation

- If Agent follows $\{\mu_t\}$ recommended by Principal (on-equilibrium), both share same posterior $m_t^{\mu} = \mathbb{E}^{\mu}\left[\theta \,|y^t, \mu^t\right]$:

$$dm_t^{\mu} = \phi \left[dY_t - \left(\mu_t + m_t^{\mu} \right) dt \right]$$

- Otherwise, if Agent shirked $\widehat{\mu}_t = \mu_t - \epsilon$ only at $[t, t + dt]$ (off-equilibrium), then Principal's belief would be distorted

$$\text{Agent knows the truth: } dm_t^{\widehat{\mu}} = \phi \left[dY_t - \left(\widehat{\mu}_t + m_t^{\widehat{\mu}} \right) dt \right]$$

  - *persistent* belief-manipulation effect for $s \geq t$— hidden information

$$m_s^{\mu} - m_s^{\widehat{\mu}} = -\phi e^{-\phi(s-t)} \epsilon \cdot dt < 0 \text{ for any } s \geq t$$

Belief Manipulation (Cont'd)

One time shirking at $[t, t + dt]$ leads to persistent belief distortion

$$m_s^{\mu} - m_s^{\widehat{\mu}} = -\phi e^{-\phi(s-t)} \epsilon \cdot dt < 0 \text{ for any } s \geq t$$

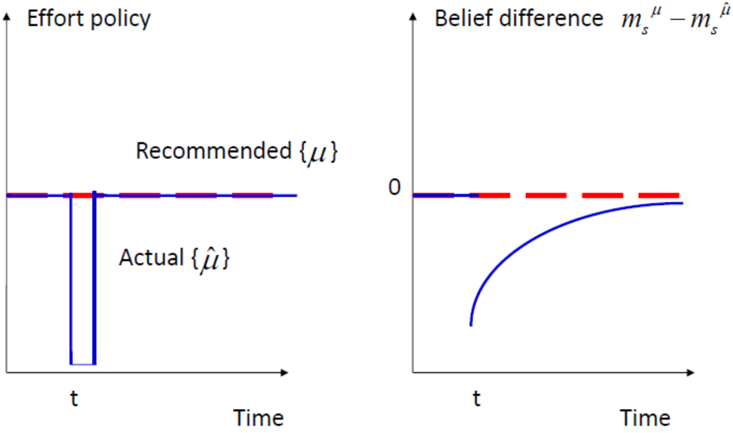

The Agent's Problem: Continuation Value

- Given the contract $\{c_t, \mu_t\}$, continuation value, no-savings

$$v_t \equiv \mathbb{E}_t^{\mu}\left[\int_t^{\infty} e^{-r(s-t)} u\left(c_s, \mu_s\right) ds\right]$$

The Agent's Problem: Continuation Value

- Given the contract $\{c_t, \mu_t\}$, continuation value, no-savings

$$v_t \equiv \mathbb{E}_t^{\mu} \left[\int_t^{\infty} e^{-r(s-t)} u\left(c_s, \mu_s\right) ds \right]$$

- Martingale Representation Theorem (Sannikov, 2008):

$$dv_t = rv_t dt - u\left(c_t, \mu_t\right) dt + \underbrace{\beta_t}_{\text{\$ incentive}} \underbrace{\left(-arv_t\right)}_{\text{Marginal u, \$}\rightarrow\text{utils}} \underbrace{\left(dY_t - \left(\mu_t + m_t^{\mu}\right) dt\right)}_{\text{performance } \sigma dB_t^{\mu}}$$

- Private savings + CARA imply

$$rv_t = u\left(c_t, \mu_t\right) \text{ and } u_c = -arv_t$$

  - CARA utility $\Rightarrow$ u_c is proportional to u

- v_t is an exponential Martingale, with loading being incentives $\{\beta_u\}$

$$v_s = v_t \exp\left(-\int_t^s ar\beta_u \sigma dB_u^{\mu} - \frac{1}{2}\int_t^s a^2 r^2 \beta_u^2 \sigma^2 du\right), \text{ for } s > t$$

The Agent's Problem: Incentive-Compatibility Condition

- Shirking $\widehat{\mu}_t = \mu_t - \epsilon$ at $[t, t+dt]$ affects the agent's total payoff

$$u\left(c_t, \mu_t - \epsilon\right) dt + v_t + \mathbb{E}_t^{\mu}\left[\int_t^{\infty} e^{-r(s-t)} dv_s\right]$$

$$= \underbrace{u\left(c_t, \mu_t - \epsilon\right) dt}_{\text{Saving effort cost } (+)} + \mathbb{E}_t^{\mu}\left[\underbrace{\beta_t\left(-arv_t\right)\left(dY_t\left(\mu_t - \epsilon\right) - \mu_t dt - m_t^{\mu} dt\right)}_{\text{Hurting performance \textbf{today} (-)}}\right]$$

$$+\mathbb{E}_t^{\mu}\left[\underbrace{\int_{t+dt}^{\infty} e^{-r(s-t)} \beta_s\left(-arv_s\right)\left(dY_s\left(\mu_s\right) - \mu_s dt - m_s^{\mu^{\epsilon}} dt\right)}_{\text{Creating \textbf{future} belief divergence } (+)}\right] + v_t$$

The Agent's Problem: Incentive-Compatibility Condition

- Shirking $\widehat{\mu}_t = \mu_t - \epsilon$ at $[t, t+dt]$ affects the agent's total payoff

$$u\left(c_t, \mu_t - \epsilon\right) dt + v_t + \mathbb{E}_t^{\mu}\left[\int_t^{\infty} e^{-r(s-t)} dv_s\right]$$

$$= \underbrace{u\left(c_t, \mu_t - \epsilon\right) dt}_{\text{Saving effort cost } (+)} + \mathbb{E}_t^{\mu}\left[\beta_t\left(-arv_t\right)\underbrace{\left(dY_t\left(\mu_t - \epsilon\right) - \mu_t dt - m_t^{\mu} dt\right)}_{\text{Hurting performance \textbf{today} (-)}}\right]$$

$$+\mathbb{E}_t^{\mu}\left[\underbrace{\int_{t+dt}^{\infty} e^{-r(s-t)} \beta_s\left(-arv_s\right)\left(dY_s\left(\mu_s\right) - \mu_s dt - m_s^{\mu^{\epsilon}} dt\right)}_{\text{Creating \textbf{future} belief divergence } (+)}\right] + v_t$$

- The incentive-compatibility (IC) condition

$$u_{\mu} + \left(-arv_t\right)\beta_t - \mathbb{E}_t^{\mu}\left[\int_t^{\infty} \phi e^{-(\phi+r)(s-t)} \beta_s\left(-arv_s\right) ds\right] = 0$$

  - future incentives $\beta_s \uparrow$, shirking to manipulate belief $\uparrow$, offsetting today's incentive β_t

The Incentive-Compatibility Condition — Information Rent

- Recall $u_\mu = -u_c \mu_t = (arv_t)\,\mu_t$,
  $v_s = v_t \exp\left(-\int_t^s ar\beta_u \sigma dB_u^\mu - \frac{1}{2}\int_t^s a^2 r^2 \beta_u^2 \sigma^2 du\right)$
- Cancel v_t on both sides, IC condition becomes

$$\mu_t = \beta_t - \underbrace{\mathbb{E}_t^\mu \left[\int_t^\infty \underbrace{\phi e^{-(\phi+r)(s-t)}}_{\text{discounting}} \underbrace{\beta_s}_{\text{belief manipulation}} \underbrace{e^{-\int_t^s ar\beta_u \sigma dB_u^\mu - \frac{1}{2}\int_t^s a^2 r^2 \beta_u^2 \sigma^2 du}}_{\text{future marginal utility effect}} ds \right]}_{p_t \text{ information rent}}$$

  - $\mu_t = \beta_t - p_t$, only $\{\beta\}$ enters in the IC condition!
- Intuitively, CARA implies level invariance, the level of today's v does not matter
- We use Pontryagin's Maximum Principle to rigorously derive it and prove its sufficiency

The Principal's Problem

- The Principal's Problem:

$$\max_{\{\beta_t\}} \mathbb{E}\left[\int_0^\infty e^{-rt}\left(dY_t - c_t dt\right)\right]$$

$$\begin{aligned} s.t \qquad & dY_t = (\mu_t + m_t)dt + \sigma dB_t \text{ and } dm_t = \phi\sigma dB_t \\ & c_t = \frac{1}{2}\mu_t^2 - \frac{\ln(-arv_t)}{a} \\ & dv_t = \beta_t(-arv_t)\sigma dB_t \text{ given } v_0 \\ & \mu_t = \beta_t - p_t(\{\beta_s\}) \end{aligned}$$

- Value function is separable w.r.t. two state variables: (v_t, p_t)

$$J(v,p) = \underbrace{V(p)}_{\text{Project Value}} - \underbrace{\left[-\frac{\ln(-arv)}{ar}\right]}_{\text{Certainty Equivalent to Agent}}$$

  - CARA setting, v_t only matters to the extent of certainty equivalent

The Principal's Problem: Dynamic Programming

- Hamilton-Jacobi-Bellman equation for $V(p)$, with only state variable p_t

$$V(p) \equiv \max_{\{\beta_t, \sigma_t^P\}} \mathbb{E}\left[\int_0^\infty e^{-rt} \left(\begin{array}{cc} \underbrace{(\beta_t - p_t)}_{\text{effort benefit}} & \underbrace{-\frac{1}{2}(\beta_t - p_t)^2}_{\text{effort cost}} \\ \multicolumn{2}{c}{\underbrace{-\frac{1}{2} ar\sigma^2 \beta_t^2}_{\text{incentive (risk) compensation}}} \end{array} \right) dt \right]$$

$$s.t. \quad dp_t = \left[(\phi + r) p_t + \beta_t \left(ar\sigma\sigma_t^P - \phi \right)\right] dt + \sigma_t^P dB_t, \; p_0 = p$$

- Optimal control is on β_t (drift of dp_t) and σ_t^P (volatility of dp_t)
- Mapping to consumption/portfolio problem: β_t is like consumption affecting the drift of wealth, σ_t^P is like portfolio choice affecting the volatility of wealth

Benchmark Case: Deterministic Contracts

- Under the restriction $\sigma^p = 0$,
  - $\{\beta_t\}$ is deterministic and

$$(IC)\ \mu_t = \beta_t - \int_t^\infty \phi \beta_s e^{-(\phi+r)(s-t)} ds$$

  - The optimal deterministic contract is solved in closed-form with quadratic value function

$$V^d(p) = -\frac{1}{2} A^d p^2 + B^d p$$

- Result: p^d, β^d, and μ^d declines exponentially over time toward zero
  - more costly to provide incentives later, because later incentives give agent information rents early on
  - remains to hold in the stochastic optimal contract

The Optimal Contract: Characterization

- Value function $V(p)$ satisfies ODE: for $p \in [0, \overline{p}]$

$$rV = \frac{1}{2}\frac{\left(1+p-\phi V_p\right)^2}{1+ar\sigma^2+a^2r^2\sigma^2\frac{V_p^2}{V_{pp}}} - p - \frac{1}{2}p^2 + V_p\left(\phi+r\right)p$$

  where p_t is bounded between an endogenous entry-no-exit upper boundary $\{\overline{p}\}$ and an absorbing lower boundary $\{0\}$

- Under the optimal contract, effort $\mu_t = \beta_t - p_t$ is stochastic!

Optimal Contracting: Option-Like Feature

- Besides time-decreasing effort, the optimal contract exhibits an **option-like** feature
  - incentives increase after good performance, i.e., diffusion of β is positive
- Interestingly, it is due to an asset pricing intuition
- To reduce the agent's information rent

$$p_t = \frac{1}{u_c(c_t)} \mathbb{E}_t \left[\int_t^{\infty} \phi e^{-(\phi+r)(s-t)} \beta_s u_c(c_s)\, ds \right].$$

  - intuition: (i) following good performance, $u_c(c_s, \mu_s) \downarrow$
    (ii) increasing incentives, $\beta_s \uparrow$, induces negative covariance

$$Cov_t\left[\beta_s, u_c(c_s, \mu_s)\right] < 0$$

$\Rightarrow$ as a result, information rent p_t is lower
  - option: give belief manipulation rewards (high incentives) in states where Agent does not care much (low marginal utility)

Concluding Remarks

- We study optimal dynamic contracting in a context of uncertainty and learning
- Our continuous time agency model features both hidden action and hidden information
  - the optimal effort policy is stochastic and front-loaded
- The optimal contract has an option-like feature, which is used to mitigate the agent's belief-manipulation motives
  - it is typically hard to have option in optimal contracting standard models (Dittmann and Maug, 2007)
- Empirical predictions/evidence
  - Core and Guay (1999) find that the annual grant of options and stocks to CEO is increasing in past stock returns
  - the industry/firm with higher uncertainty should have more option-based contracts for managerial compensation
  - Ittner, Lambert, and Larcker (2003) and Murphy (2003) who document new-economy firms (e.g., computer) grant more stock options to managers than old-economy firms