Optimal Contracts for Experimentation

Marina Halac   Navin Kartik   Qingmin Liu

October 2013

Introduction

- Principal owns project whose quality (profitability) is unknown
- Agent must work on project, with experimentation or learning
- Beliefs about project depend on agent's effort and ability (and output)
- Expected benefit of effort depends on these beliefs
    - If pessimistic enough, optimal to abandon project
- Agency problem: Agent's ability and effort unobservable to principal

Introduction

- Example: Netflix hires firm to build algorithm to improve movie recommendation accuracy by 10%. Incentive contract must deal with:
  1. Not initially known if 10% target attainable in relevant timeframe
  2. Firm has superior info on its comparative advantage/suitability
  3. How much time/effort firm devotes to task is unobservable
- Features are relevant in many contractual environments
  - Design of incentives for R&D projects
  - Testing of new products
  - Hiring recruiting agency to search for new CEO

Introduction

- Learning about an uncertain state, adverse selection, and dynamic moral hazard are salient features of these agency relationships
- *How well can principal incentivize agent? How do these features affect optimal incentive contracts? What distortions, if any, arise?*
- We provide answers in a simple model of experimentation
    - Each of these features important for dynamic incentive provision
    - It is only their interaction that precludes efficiency
    - Despite intricacy of problem, simple contracts are optimal

Related Literature

- Contracting for experimentation
  - Bergemann and Hege (1998, 2005), Manso (2011), Bonatti and Hörner (2011, 2012), Hörner and Samuleson (2013), Kwon (2013)
  - Gomes, Gottlieb, and Maestri (2013)
- Different environments
  - Gerardi and Maestri (2011)
  - Lewis and Ottaviani (2008), Lewis (2011)
  - Sannikov (2007), Gershkov and Perry (2012)
  - Demarzo and Sannikov (2011), He et al. (2012), Prat and Jovanovic (2012)

Model – Environment (1)

Build on exponential bandit model (Keller, Rady, and Cripps, 2005):

- Project quality or state is either good or bad
  - Prior on good state is $\beta_0 \in (0, 1)$
- In each period $t \in \{1, 2, \ldots\}$, agent covertly chooses to work or shirk
  - Exerting effort in any period costs the agent $c > 0$
- If agent works and state is good, project succeeds with probability λ
- If agent shirks or state is bad, success cannot obtain

$\rightarrow$ *Working is "pulling the risky arm"; shirking is "pulling the safe arm"*

Model – Environment (2)

- Project success yields principal payoff normalized to 1
  - No further effort once success is obtained
- Project success is publicly observable
  - Results also hold if privately observed by agent but verifiable disclosure
- Add adverse selection: Agent privately knows his ability, $\theta \in \{L, H\}$
  - Probability of success in a period (conditional working and good project) is λ^θ, where $1 > \lambda^H > \lambda^L > 0$
  - Prior on ability H is $\mu_0 \in (0, 1)$

First Best

- First best characterized by optimal stopping time t^{θ}:

$$t^{\theta} = \max_{t \geq 0} \left\{ t : \overline{\beta}_t^{\theta} \lambda^{\theta} \geq c \right\},$$

where $\overline{\beta}_t^{\theta}$ is belief on good state at beginning of t given work up to t

Assumption 1. Experimentation is efficient: for $\theta \in \{L, H\}$, $\beta_0 \lambda^{\theta} > c$.

- As is intuitive, t^{θ} is increasing in β_0 and decreasing in c
- But t^{θ} is non-monotonic in λ^{θ}: productivity vs. learning effects

First Best

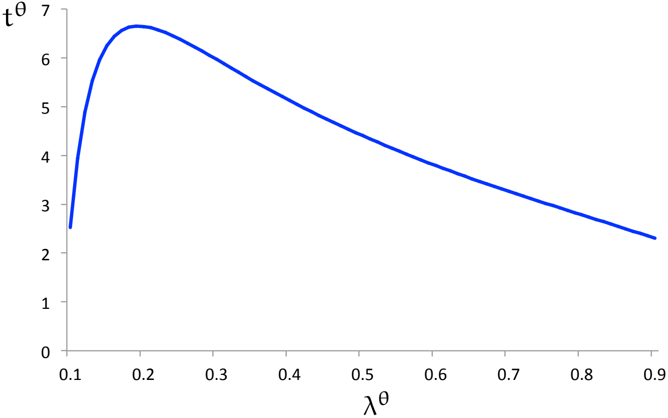

- Both $t^H > t^L$ and $t^H < t^L$ are robust possibilities

Model – Contracts

- Contract at $t = 0$ with full commitment power from the principal
- Hidden info $\Rightarrow$ wlog principal offers *menu* of dynamic contracts
- Contract specifies transfers in each period as function of publicly observable history, i.e. whether or not success has obtained to date
- A contract is $\boldsymbol{C} = (T, W_0, \boldsymbol{b}, \boldsymbol{l})$, where $\boldsymbol{b} = (b_1, \ldots, b_T)$ and $\boldsymbol{l} = (l_1, \ldots, l_T)$
- Agent's actions are $\boldsymbol{a} = (a_1, \ldots, a_T)$, with $a_t \in \{0, 1\}$

Model – Payoffs

- Given agent's type θ, $\boldsymbol{C} = (T, W_0, \boldsymbol{b}, \boldsymbol{l})$, $\boldsymbol{a} = (a_1, \ldots, a_T)$, and $\delta \in (0, 1]$, principal's expected discounted payoff at $t = 0$ is

$$\Pi_0^\theta (\boldsymbol{C}, \boldsymbol{a}) = \beta_0 \sum_{t=1}^{T} \delta^t \left[\prod_{s<t} \left(1 - a_s \lambda^\theta\right) \right] \left[a_t \lambda^\theta (1 - b_t) - \left(1 - a_t \lambda^\theta\right) l_t \right]$$
$$- (1 - \beta_0) \sum_{t=1}^{T} \delta^t l_t - W_0$$

- Agent's expected discounted payoff at $t = 0$ is

$$U_0^\theta (\boldsymbol{C}, \boldsymbol{a}) = \beta_0 \sum_{t=1}^{T} \delta^t \left[\prod_{s<t} \left(1 - a_s \lambda^\theta\right) \right] \left[a_t (\lambda^\theta b_t - c) + \left(1 - a_t \lambda^\theta\right) l_t \right]$$
$$+ (1 - \beta_0) \sum_{t=1}^{T} \delta^t (l_t - a_t c) + W_0$$

Bonus and Penalty Contracts

Definition

A bonus contract is $\boldsymbol{C} = (T, W_0, \boldsymbol{b}, \boldsymbol{l})$ *s.t.* $l_t = 0$ *for all* $t = 1, \ldots, T$.

A bonus contract is constant-bonus if $b_t = b$ *for all* $t = 1, \ldots, T$.

Definition

A penalty contract is $\boldsymbol{C} = (T, W_0, \boldsymbol{b}, \boldsymbol{l})$ *s.t.* $b_t = 0$ *for all* $t = 1, \ldots, T$.

A penalty contract is onetime-penalty if $l_t = 0$ *for all* $t = 1, \ldots, T - 1$.

Proposition

For any contract, $\mathbf{C} = (T, W_0, \boldsymbol{b}, \boldsymbol{l})$*, there exist an equivalent penalty contract* $\widehat{\mathbf{C}} = (T, \widehat{W_0}, \widehat{\boldsymbol{l}})$ *and an equivalent bonus contract* $\widetilde{\mathbf{C}} = (T, \widetilde{W_0}, \widetilde{\boldsymbol{b}})$.

Benchmark: No Adverse Selection

- If θ observable, principal implements first-best and extracts all surplus with simple contracts
  - E.g., constant-bonus contract $\boldsymbol{C}^\theta = \left(t^\theta, W_0^\theta, 1\right)$ where W_0^θ s.t. for $\boldsymbol{a} = \mathbf{1}$, θ's IR constraint at $t = 0$ binds

Benchmark: No Moral Hazard

- If effort observable and contractible, principal implements first-best and extracts all surplus with simple contracts
- Can ignore IC constraints for effort and exploit the two types' differing probabilities of success
    - E.g., $\boldsymbol{C}^H = (t^H, W_0^H, b^H)$ and $\boldsymbol{C}^L = (t^L, W_0^L, b^L)$ with $b^H > 0 > b^L$ s.t. $\boldsymbol{C}^H$ is "too risky" for type L and $\boldsymbol{C}^L$ is "too risky" for type H
- Similar to FSE in mechanism design w/correlated info (Cremer-McLean 1985) or w/ex-post public info (Riordan-Sappington 1988)

Optimal Contracts w/Adverse Selection and Moral Hazard

- Given any contract $\boldsymbol{C} = (T, W_0, \boldsymbol{b}, \boldsymbol{l})$, define

$$\boldsymbol{\alpha}^{\theta}(\mathbf{C}) := \underset{\mathbf{a}}{\operatorname{argmax}}\ U_0^{\theta}(\mathbf{C}, \mathbf{a})$$

- Principal's problem is

$$\max_{(\boldsymbol{C}^H, \boldsymbol{C}^L, \mathbf{a}^H, \mathbf{a}^L)} \mu_0 \Pi_0^H\left(\boldsymbol{C}^H, \boldsymbol{a}^H\right) + (1-\mu_0)\,\Pi_0^L\left(\boldsymbol{C}^L, \boldsymbol{a}^L\right)$$

subject to, for all $\theta, \theta' \in \{L, H\}$,

$$\boldsymbol{a}^{\theta} \in \boldsymbol{\alpha}^{\theta}\left(\boldsymbol{C}^{\theta}\right) \qquad (\mathsf{IC}_a^{\theta})$$

$$U_0^{\theta}\left(\boldsymbol{C}^{\theta}, \boldsymbol{a}^{\theta}\right) \geq 0 \qquad (\mathsf{IR}^{\theta})$$

$$U_0^{\theta}\left(\boldsymbol{C}^{\theta}, \boldsymbol{a}^{\theta}\right) \geq U_0^{\theta}\left(\boldsymbol{C}^{\theta'}, \boldsymbol{\alpha}^{\theta}\left(\boldsymbol{C}^{\theta'}\right)\right) \qquad (\mathsf{IC}^{\theta,\theta'})$$

Optimal Contracts: Contrast with Static Problem

- Standard buyer-seller adverse selection problem: if type deviates to another type's contract, consumes quantity specified by that contract
- Our setting: not a priori clear what "consumption bundle", i.e. effort profile, each type will choose after such a deviation
- No difficulty if systematic relationship between the two types' effort profiles in arbitrary contract, e.g., "single-crossing condition"
- But no analog of single-crossing in general in our dynamic setting

Optimal Contracts: Agency Issues

- Incentives at t shaped by current transfers *and* by subsequent transfers through effect on continuation values $\rightarrow$ dynamic agency effects
  - Increasing b_{t+1} (reducing l_{t+1}) increases incentive to shirk (work) at t
- Continuation value depends on agent's type, future effort profile, and his private belief about the state
- Agent's type also affects current incentive through mg benefit of effort

$\Rightarrow$ For arbitrary contract $\boldsymbol{C}$, we may have $\boldsymbol{\alpha}^H(\boldsymbol{C}) \cap \boldsymbol{\alpha}^L(\boldsymbol{C}) = \emptyset$

  - H may not want to experiment as long as L in arbitrary contract
  - Leads to fixed point problem

Optimal Contracts when $t^H > t^L$

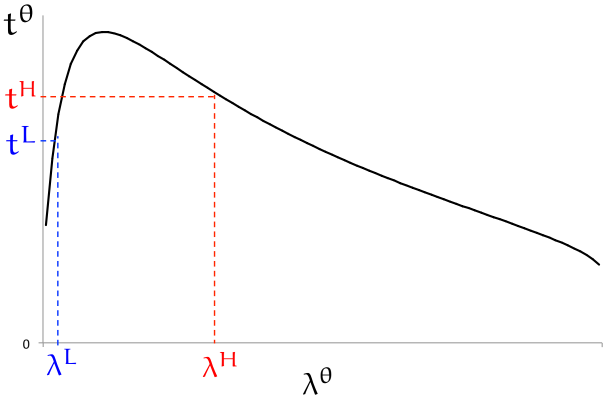

Optimal Contracts: Second-Best Efficiency

Theorem

Assume $t^H > t^L$. In any optimal menu of contracts, each type $\theta \in \{H, L\}$ is induced to work for some number of periods, $\bar{t}^\theta$; if $\delta < 1$, the periods are $1, \ldots, \bar{t}^\theta$. Relative to first-best, second-best has

$$\bar{t}^H = t^H \text{ and } \bar{t}^L \leq t^L.$$

- At the limit when length of time intervals vanishes, $\bar{t}^L < t^L$
- Result is consequence of our characterization of optimal menus

Optimal Contracts: Penalty

Theorem

Assume $t^H > t^L$. There is an optimal menu of penalty contracts:

1. *Onetime-penalty for H, $\boldsymbol{C}^H = \left(t^H, W_0^H, l_{t^H}^H\right)$, with $l_{t^H}^H < 0 < W_0^H$,*
2. *Penalty contract for L, $\boldsymbol{C}^L = \left(\bar{t}^L, W_0^L, \boldsymbol{l}^L\right)$,*

such that

- *For all $t \in \{1, \ldots, \bar{t}^L\}$,* $$l_t^L = \begin{cases} -(1-\delta)\frac{c}{\overline{\beta}_t^L \lambda^L} & \text{if } t < \bar{t}^L, \\ -\frac{c}{\overline{\beta}_{\bar{t}^L}^L \lambda^L} & \text{if } t = \bar{t}^L; \end{cases}$$
- *$W_0^L > 0$ is such that (IR^L) binds;*
- *Type H gets an information rent: $U_0^H(\boldsymbol{C}^H, \boldsymbol{\alpha}^H(\boldsymbol{C}^H)) > 0$;*
- *$\mathbf{1} \in \boldsymbol{\alpha}^H(\boldsymbol{C}^H)$, $\mathbf{1} \in \boldsymbol{\alpha}^L(\boldsymbol{C}^L)$, and $\mathbf{1} = \boldsymbol{\alpha}^H(\boldsymbol{C}^L)$.*

Generically, $\boldsymbol{C}^L$ is unique within penalty contracts.

Optimal Contracts: Penalty

- L's contract has increasing penalty in each period $t < \bar{t}^L$ at which no success, followed by larger penalty that "jumps" in $\bar{t}^L$
- As $\delta \to 1$, L's contract reduces to a onetime-penalty contract
- Among penalty contracts, $\boldsymbol{C}^L$ is generically unique optimal contract
  - Incentivizes L's effort while minimizing H's information rent
- Characterization yields comparative statics: $\bar{t}^L$ is weakly increasing in β_0, and weakly decreasing in c and μ_0

Optimal Contracts: Penalty

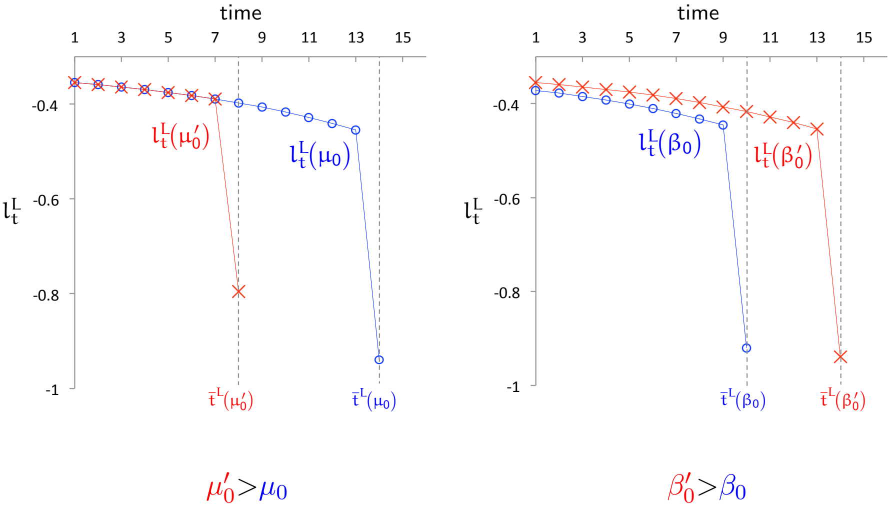

Optimal Contracts: Sketch of Proof

- Without loss, focus on penalty contracts
- **Step 1**: It is without loss to focus on $\boldsymbol{C}^L$ s.t. $\mathbf{1} \in \boldsymbol{\alpha}^L(\boldsymbol{C}^L)$
- **Step 2**: Relax the problem
  - Strategy: Conjecture that "single-crossing" holds in optimal menu

Optimal Contracts: Sketch of Proof

$$\max_{\left(\boldsymbol{C}^H\in\mathcal{C},\boldsymbol{C}^L\in\mathcal{C},\boldsymbol{a}^H\right)} \mu_0\Pi_0^H\left(\boldsymbol{C}^H,\boldsymbol{a}^H\right)+\left(1-\mu_0\right)\Pi_0^L\left(\boldsymbol{C}^L,\mathbf{1}\right) \qquad \textbf{(P)}$$

subject to

$$\mathbf{1}\in\boldsymbol{\alpha}^L\left(\boldsymbol{C}^L\right) \qquad (\text{IC}_a^L)$$

$$\boldsymbol{a}^H\in\boldsymbol{\alpha}^H\left(\boldsymbol{C}^H\right) \qquad (\text{IC}_a^H)$$

$$U_0^L\left(\boldsymbol{C}^L,\mathbf{1}\right)\geq 0 \qquad (\text{IR}^L)$$

$$U_0^H\left(\boldsymbol{C}^H,\boldsymbol{a}^H\right)\geq 0 \qquad (\text{IR}^H)$$

$$U_0^L\left(\boldsymbol{C}^L,\mathbf{1}\right)\geq U_0^L\left(\boldsymbol{C}^H,\boldsymbol{\alpha}^L\left(\boldsymbol{C}^H\right)\right) \qquad (\text{IC}^{LH})$$

$$U_0^H\left(\boldsymbol{C}^H,\boldsymbol{a}^H\right)\geq U_0^H\left(\boldsymbol{C}^L,\boldsymbol{\alpha}^H\left(\boldsymbol{C}^L\right)\right) \qquad U_0^H\left(\boldsymbol{C}^H,\boldsymbol{a}^H\right)\geq U_0^H\left(\boldsymbol{C}^L,\mathbf{1}\right) \qquad (\text{IC}^{HL})$$

Optimal Contracts: Sketch of Proof

- In (**RP1**), (IR^L) and (Weak-IC^{HL}) must bind $\Rightarrow$ write (**RP2**):

$$\max_{(\boldsymbol{C}^H \in \mathcal{C}, \boldsymbol{C}^L \in \mathcal{C}, \boldsymbol{a}^H)} \left\{ \text{ Expected total surplus } - \text{ Information rent of } H \right\}$$

subject to (IC_a^L), (IC_a^H)

Optimal Contracts: Sketch of Proof

- **Step 3**: Construct $\bar{\boldsymbol{l}}(T^L)$ s.t. (IC_a^L) binds in each period $1, \ldots, T^L$
- **Step 4**: Any $\boldsymbol{C}^L$ that solves (**RP2**) must use $\bar{\boldsymbol{l}}(T^L)$
  - If H takes $\boldsymbol{C}^L$, less likely to incur penalties than L
  - Any slack in (IC_a^L) just increases H's rent

Must deal with dynamic agency problem $\rightarrow$ focus on penalty helps

- Penalties have positive feedback — can use "local variation"
- Bonuses have a negative feedback — need "global variation"

Optimal Contracts: Sketch of Proof

- **Step 5**: Show $\bar{t}^L \le t^L$
  - Standard monotone comparative statics argument
- **Step 6**: Find a solution to (**RP2**) that solves original program (**P**)
  - Induction argument shows $\boldsymbol{\alpha}^H(\overline{\boldsymbol{C}}^L) = \mathbf{1} \quad (\because\ t^H > t^L \ge \bar{t}^L)$
  - Then if $\boldsymbol{C}^H$ satisfies (Weak-IC^{HL}), will satisfy (IC^{HL}) and (IR^H)
  - A onetime-penalty $\boldsymbol{C}^H$ with low enough $l^H_{t^H}$ maximizes surplus from H, satisfies (Weak-IC^{HL}), and also (IC^{LH})

Optimal Contracts: Bonus Implementation

Other implementation of the optimum?

Theorem

Assume $t^H > t^L$. There is an optimal menu of bonus contracts:

1. *Constant-bonus contract $\boldsymbol{C}^H = (t^H, W_0^H, b^H)$ with $b^H > 0$;*
2. *Bonus contract $\boldsymbol{C}^L = (\bar{t}^L, W_0^L, \boldsymbol{b}^L)$, where*

$$b_t^L = \sum_{s=t}^{\bar{t}^L} \delta^{s-t}(-l_s^L),$$

and $W_0^L < 0$ is such that (IR^L) binds.

Generically, $\boldsymbol{C}^L$ is unique within bonus contracts. Implementation satisfies interim IR constraints.

Optimal Contracts: Bonus Implementation

- Can verify that

$$b_t^L = \frac{(1-\delta)c}{\overline{\beta}_t^L \lambda^L} + \delta b_{t+1}^L \quad \text{for any } t \in \{1, \ldots, \bar{t}^L - 1\},$$

- Type L's bonus increases over time
- As $\delta \to 1$, L's contract becomes a constant-bonus contract

Optimal Contracts when $t^H \leq t^L$

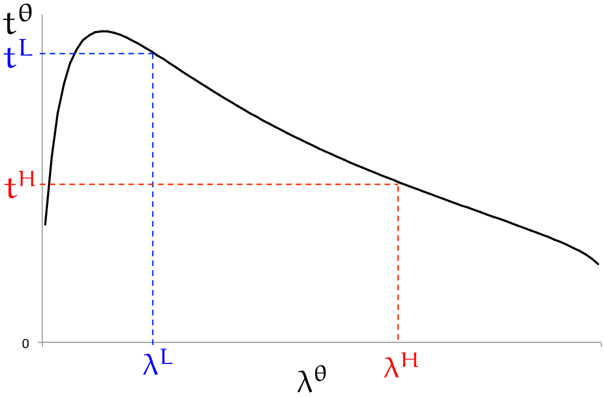

Optimal Contracts when $t^H \leq t^L$

- Major difficulty: $t^H \leq t^L$ compatible with $t^L > t^*$, where t^* given by

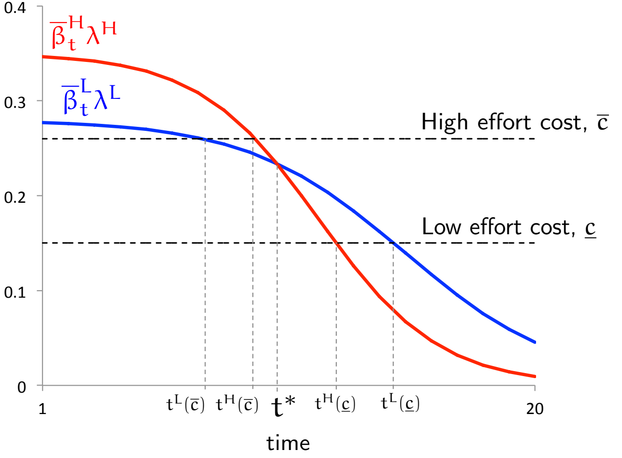

Optimal Contracts when $t^H \leq t^L$

- Second-best stopping time for L can also satisfy $\bar{t}^L > t^*$ (e.g., μ low)
- Cannot conjecture $\mathbf{1} \in \boldsymbol{\alpha}^H(\boldsymbol{C}^L) \rightarrow$ No "single-crossing" at optimum
- For arbitrary δ, difficult to find valid restriction on $\boldsymbol{\alpha}^H(\boldsymbol{C}^L)$; example:

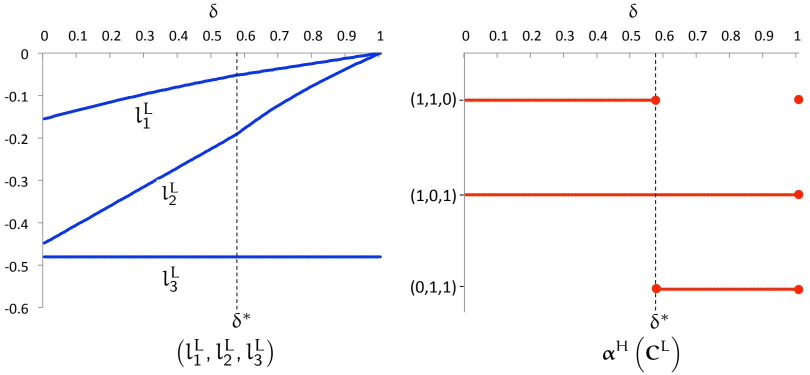

- However, we are able to solve the problem when $\delta = 1$

Optimal Contracts: Penalty

Theorem

Assume $\delta = 1$, $t^H \leq t^L$. There is an optimal menu of onetime-penalty contracts:

1. $\boldsymbol{C}^H = (t^H, W_0^H, l_{t^H}^H)$ *for* H*, with* $l_{t^H}^H < 0 < W_0^H$*,*
2. $\boldsymbol{C}^L = (\bar{t}^L, W_0^L, l_{\bar{t}^L}^L)$ *for* L*,*

such that

- $\bar{t}^L \leq t^L$*;*
- $l_{\bar{t}^L}^L = \min\left\{-\frac{c}{\overline{\beta}_{\bar{t}^L}^L \lambda^L}, -\frac{c}{\overline{\beta}_{t^{HL}}^H \lambda^H}\right\}$*, for* $t^{HL} := \max\limits_{\mathbf{a} \in \boldsymbol{\alpha}^H(\mathbf{C}^L)} \#\{n : a_n = 1\}$*;*
- $W_0^L > 0$ *is such that (IR*^{L}*) binds;*
- *Type* H *gets an information rent:* $U_0^H(\boldsymbol{C}^H, \boldsymbol{\alpha}^H(\boldsymbol{C}^H)) > 0$*;*
- $\mathbf{1} \in \boldsymbol{\alpha}^H(\boldsymbol{C}^H)$, $\mathbf{1} \in \boldsymbol{\alpha}^L(\boldsymbol{C}^L)$.

Generically, $\boldsymbol{C}^L$ *is essentially-unique within penalty contracts*

Optimal Contracts: Penalty

- Share common properties with optimal contracts for $\delta = 1$, $t^H > t^L$
- However, two differences when $t^H \leq t^L$:
  - In general, optimal $\boldsymbol{C}^L$ such that $\mathbf{1} \notin \boldsymbol{\alpha}^H(\boldsymbol{C}^L)$
  - Can be optimal to satisfy L's IC constraint for effort with slack
- Intuition stems from information-rent considerations:
  - Because H less likely to incur penalties if he mimics L, want to minimize penalties $\implies$ onetime-penalty with $l^L_{\bar{t}^L} = -\frac{c}{\overline{\beta}^L_{\bar{t}^L}\lambda^L}$
  - But when $\bar{t}^L > t^*$, H would then work for some $T < \bar{t}^L$ periods

  $\implies$ Possible that T is such that H more likely to incur $l^L_{\bar{t}^L}$

  $\implies$ Want lower $l^L_{\bar{t}^L}$ to reduce rent $\implies$ $l^L_{\bar{t}^L} = -\frac{c}{\overline{\beta}^H_{T+1}\lambda^H}$

Optimal Contracts: Sketch of Proof

- **Key Step**: When $\delta = 1$, we show that one can focus on penalty contracts for L under which H has an optimal *stopping strategy*:

  $$\mathbf{a} \in \boldsymbol{\alpha}^H(\boldsymbol{C}^L) \text{ s.t. for some } t \geq 1, a_s = 1 \text{ for } s \leq t, a_s = 0 \text{ for } s > t$$

- Given this, we show that *onetime*-penalty contracts are optimal

- Rent-minimization considerations are used to complete argument

- While restriction will not generally be valid for $\delta < 1$, optimal contracts are continuous in δ (recall example)

Optimal Contracts: Bonus Implementation

Other implementation of the optimum?

Theorem

Assume $\delta = 1$, $t^H \leq t^L$. There is an optimal menu of constant-bonus contracts

1. $\boldsymbol{C}^H = (t^H, W_0^H, b^H)$ *for* H*, with* $b^H > 0$*;*
2. $\boldsymbol{C}^L = (\bar{t}^L, W_0^L, b^L)$ *for* L*, where*

$$b^L = -l^L_{\bar{t}^L}$$

and $W_0^L < 0$ *is such that (IR*^{L}*) binds.*

Generically, $\boldsymbol{C}^L$ *is essentially-unique within bonus contracts.*
Implementation satisfies interim IR constraints.

Discussion: The role of learning

- If $\beta_0 = 1$, first-best has both types working until success is obtained
- Suppose exogenous end date T so $t^L = t^H = T$
- If it is not optimal to exclude type L, then no distortion: $\bar{t}^L = t^L$
  - Efficiency loss larger than gain from rent reduction
- When $\beta_0 < 1$, logic fails: social surplus from L vanishes over time
- Learning important for results: whenever with $\beta_0 < 1$ distort t^L without entirely excluding L, with $\beta_0 = 1$ would not distort t^L

Concluding Remarks

- Study dynamic principal-agent contracts for experimentation
- Interaction of private learning, adverse selection, and moral hazard
  - New conceptual issues, each plays role in structuring dynamic incentives
- Optimal menu induces low type to end experimentation too early
  - But first best without either adverse selection or moral hazard
- Derive explicit optimal menus: bonus and penalty contracts

Thank you!

Discussion: Private observability and disclosure

- Suppose project success is privately observed by agent but can be verifiably disclosed

Theorem

Even if project success is privately observed, the menus of contracts identified above remain optimal and implement same outcome as above

- Intuition: $l_t^\theta \leq 0$ and $\delta b_{t+1}^\theta \leq b_t^\theta$ for all t
- Desirable robustness property
  - Not true for every optimal menu under public observability

Discussion: Limited liability

- Restrict transfers to be positive; assume $t^L < t^H$, $\delta = 1$
  - Without loss, principal uses constant-bonus contracts
  - Principal distorts both types: cannot distort t^L without distorting t^H
  - Optimal b^L has same form we characterize
  - Even though both t^L and t^H distorted, order preserved: $\bar{t}^L \leq \bar{t}^H$
- Other less severe forms of limited liability may also be relevant