Moving the Goalposts

Jeffrey Ely   Martin Szydlowski

July 24, 2017

Introduction

- An agent works for a principal
- The agent earns a private reward if he works long enough to complete a task.
- Only the principal knows when the task is complete.
- In the meantime the principal keeps the flow output of the agent.
- The principal provides feedback over time with the goal of maintaining the agent's incentives to work.
- Information is the only incentive device.

Examples

- Worker and mid-level manager.
- Apprentice and expert.
- PhD student and advisor.
- Child and parent.

Ideas

- Feedback is never fully informative: *Leading the agent on*
- Incentive Reversals
  - Early on give the agent hope.
  - Later on tamp down optimism.
- *Moving the Goalposts*
- Information as a carrot.
- The value of sequential feedback.
- The value of uncertainty.

Model

- Continuous time
- The agent decides at each moment whether to continue working or quit.
- There is a task which requires a duration x of effort to complete.
- The agent is uncertain of x and has prior CDF F.
- Quitting at time τ earns expected payoff

$$F(\tau)\, e^{-r\tau} R - c\left(1 - e^{-r\tau}\right).$$

The Principal

- The principal knows the threshold x.
- At each moment he can send a message to the agent conveying information about x.
- If the agent quits at time τ, the principal earns payoff

$$1 - e^{-r_p \tau}$$

Information Policies

The principal designs and commits to an *information policy*: a dynamic, contingent plan of message disclosures.

Binary example

- Two possible thresholds: $\underline{x}, \bar{x}$.

$$0 < \underline{x} < \bar{x}.$$

- "Easy" task versus "difficult" task.
- μ: prior probability that the task is difficult.

Individual Rationality

Define $\bar{\tau}$ by

$$e^{-r\bar{\tau}}R - c\left(1 - e^{-r\bar{\tau}}\right) = 0$$

It is the maximal individually rational effort duration.

Interesting Case

Only the easy task is individually rational:

$$\underline{x} \leq \bar{\tau} < \bar{x}$$

but

$$\bar{x} - \underline{x} \leq \bar{\tau}$$

Plan

We will look at some benchmarks

- No information
- Static optimal mechanism

Then we will see the ways dynamic disclosure allows the principal to improve

Benchmark: No Information

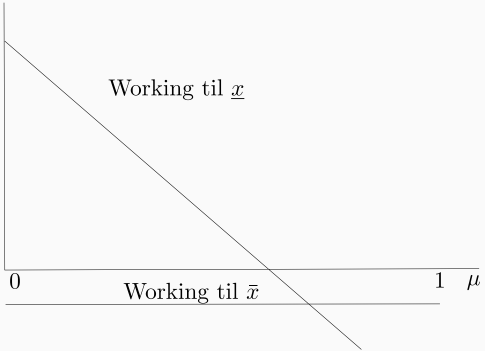

Working til $\bar{x}$ is not individually rational.

Benchmark: No Information

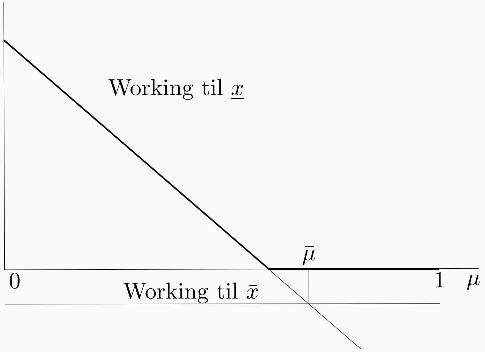

Benchmark: No Information

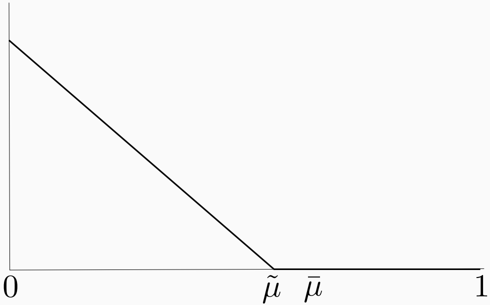

The agent's no-information value function. At $\tilde{\mu}$ the agent is indifferent between $\underline{x}$ and quitting immediately.

Benchmark: No Information

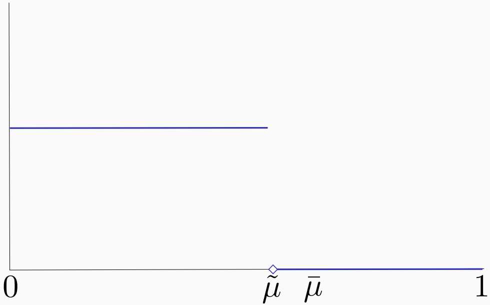

The principal's no-information value function.

Benchmark: Static Mechanism

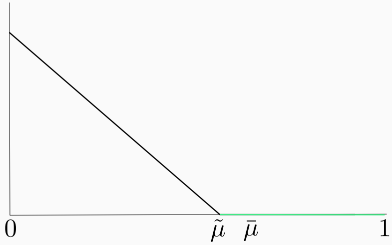

(a) Agent

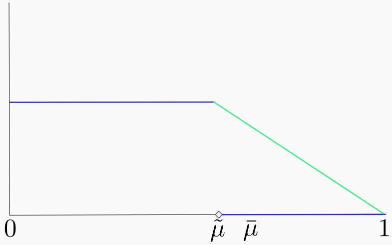

(b) Principal

Concavification

Summary

- The principal would like to persuade the agent that the threshold is low
- Static policy never gets agent to work until $\bar{x}$.
- However, incentives change as output accumulates.
- We can track *continuation values* over time.

As output accumulates

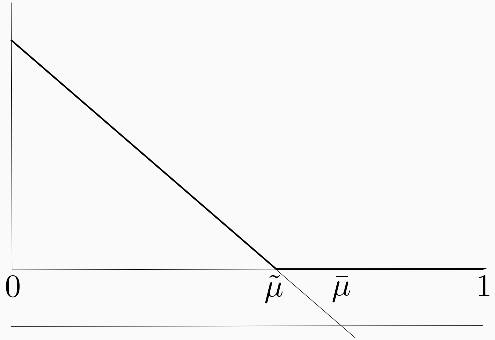

Recall the original payoffs

As output accumulates

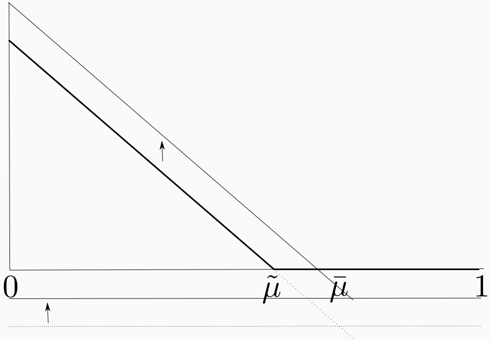

If the agent were to work for some time, the *continuation values* rise because past effort costs are sunk.

As output accumulates

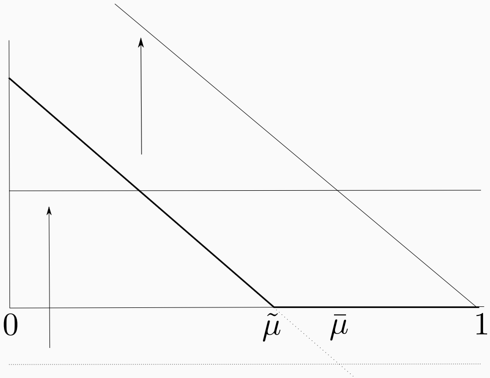

Once his accumulated output reaches $\underline{x}$, his continuation value for continuing on to $\bar{x}$ is now positive.

Meanwhile, the principal

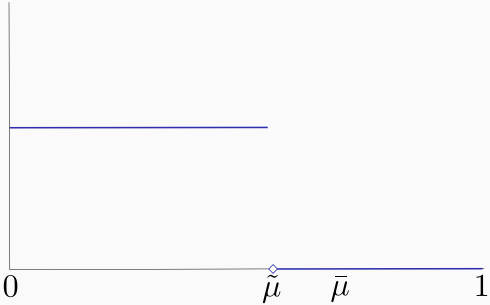

At the beginning

Meanwhile, the principal

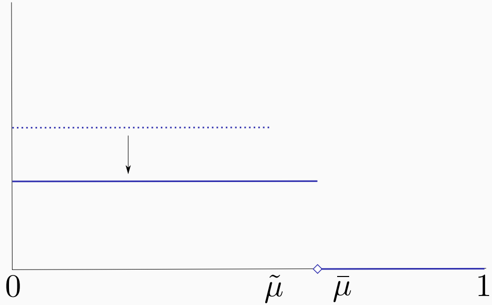

If the agent were to work for some time, the principal's *continuation values* decline because there is less effort remaining to reach the threshold.

Meanwhile, the principal

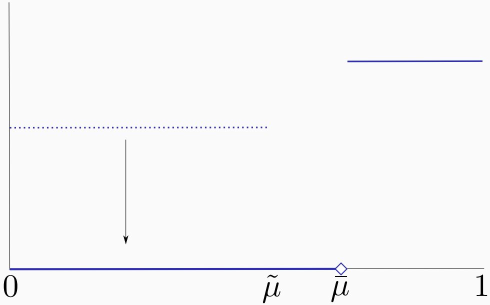

But once the agent reaches $\underline{x}$, the continuation value for $\mu \geq \bar{\mu}$ jumps up.

Meanwhile, the principal

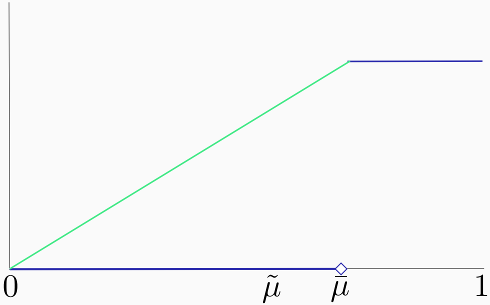

And now the principal wants to persuade the agent that the threshold is high.

This suggests the following mechanism

- Initially ask the agent to work until $\underline{x}$.
- Reveal nothing to the agent.
- Only when he reaches $\underline{x}$ disclose some information.
- As a reward for the initial effort.

Information as a carrot

- The maximum reward the principal can promise is full disclosure.
- The agent's (time-zero) value for full disclosure promised at $\underline{x}$ is

$$V(\mu) = e^{-r\underline{x}}\left\{(1-\mu)\,R + \mu\left[Re^{-r(\bar{x}-\underline{x})} - c\left(1 - e^{-r(\bar{x}-\underline{x})}\right)\right]\right\}.$$

Information as a carrot

- The agent's cost of working til $\underline{x}$ is $c\left(1 - e^{-r\underline{x}}\right)$.
- Suppose that

$$V(\mu) - c\left(1 - e^{-r\underline{x}}\right) \geq 0$$

Leading the agent on

Then the agent can be induced to

- Work to $\underline{x}$ with probability 1
- Continue working on to $\bar{x}$ with positive probability

Leading the agent on

- The agent completes the difficult task even though he would never have knowingly set out to do so.
- We call this *leading the agent on.*

Fine-tuning

- Full-disclosure may be more incentive than necessary (if the inequality is strict.)
- In this case the principal can extract surplus by reducing the informativeness of the disclosure.

Extracting Surplus

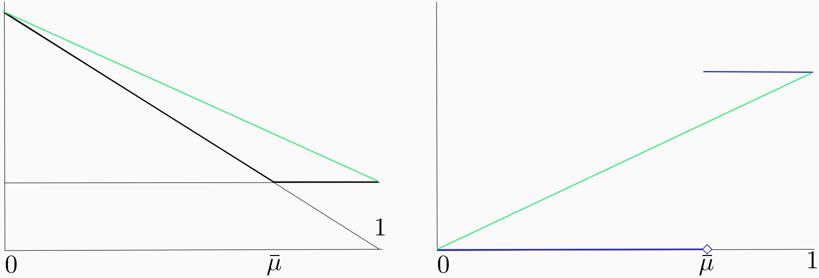

(a) Agent **(b)** Principal

Full disclosure at $\underline{x}$.

Extracting Surplus

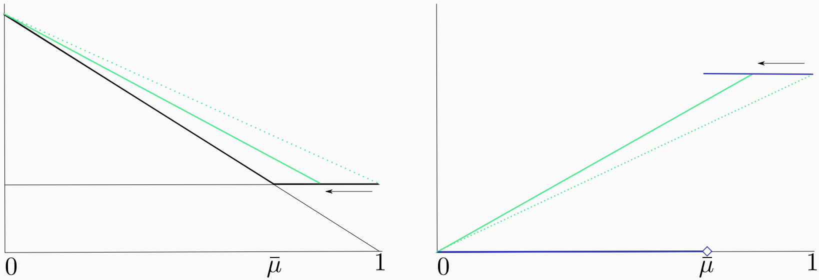

(c) Agent **(d)** Principal

The principal improves by reducing informativeness.

Optimality

- The agent completes the task with probability 1.
- The agent earns his no-information value.
- Any policy which improves the principal's payoff must involve more total effort
- This must be worse for the agent, and hence not individually rational.

Hardest

- Recall the inequality:

$$V(\mu) - c\left(1 - e^{-r\underline{x}}\right) \geq 0$$

- For initial beliefs μ close enough to 1 this inequality fails.
- Not even full-disclosure promised at $\underline{x}$ can incentivize the agent to start working.

Moving the goalposts

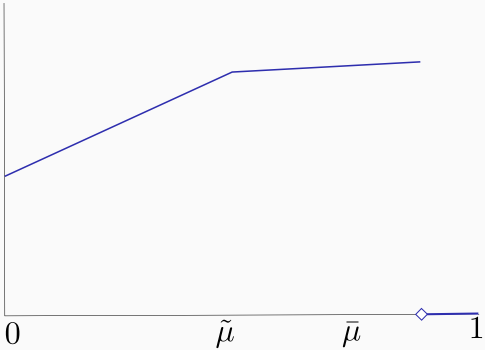

Principal's value function

Moving the goalposts

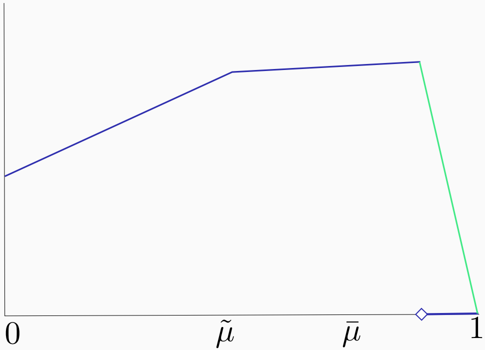

Concavification: Persuade agent that threshold is low, then lead him on

Moving the goalposts

- First, persuade agent that it's likely the threshold is low
- Stay silent while agent works
- Once $\underline{x}$ reached, persuade agent that it's likely the threshold is high
- Even though the agent anticipates all of this.

Optimality

	$x = \underline{x}$	$x = \bar{x}$
$\tau = 0$	0	α
$\tau = \underline{x}$	$1 - \mu$	0
$\tau = \bar{x}$	0	β

Designing the Threshold Distribution

Information disclosure is useful because the agent is uncertain *ex ante* about the threshold. Does the principal benefit from the uncertainty?

Deterministic Threshold

The optimal deterministic threshold is $\bar{\tau}$.

Random Threshold

Now suppose the principal can randomize the threshold (and commit to it.)

Define $\tilde{\tau}$ by

$$R/2 \leq -c(1 - e^{-r\tilde{\tau}}) + e^{-r\tilde{\tau}}R.$$

and choose a binary threshold distribution with equal probability on

$$\bar{x} = \bar{\tau} + \tilde{\tau}/2, \quad \text{and}$$
$$\underline{x} = \bar{\tau} - \tilde{\tau}/2.$$

No Information

The agent's no information value from this distribution is zero.
(He would quit immediately.)

Full Delayed Disclosure

Full delayed disclosure would induce the agent to start working and (just) complete the task with probability 1.

The agent prefers the randomization

The expected discounted payoff from full delayed disclosure is strictly higher than the deterministic $\bar{\tau}$.

The agent is risk loving

Indeed, the expected effort costs are smaller,

$$c\left[\frac{1}{2}(1-e^{-r\underline{x}})+\frac{1}{2}(1-e^{-r\bar{x}})\right]$$
$$=c\left[1-\left(\frac{e^{-r\underline{x}}+e^{-r\bar{x}}}{2}\right)\right]$$
$$<c\left[1-e^{-r\bar{\tau}}\right]$$

and the expected reward value is higher

$$R\left[\frac{1}{2}e^{-r\underline{x}}+\frac{1}{2}e^{-r\bar{x}}\right]$$
$$>Re^{-r\bar{\tau}}.$$

But so is the principal

For the same reason, the principal is worse off.

$$\frac{1}{2}(1 - e^{-r_p \underline{x}}) + \frac{1}{2}(1 - e^{-r_p \bar{x}}) < 1 - e^{-r_p \bar{\tau}}$$

Step One

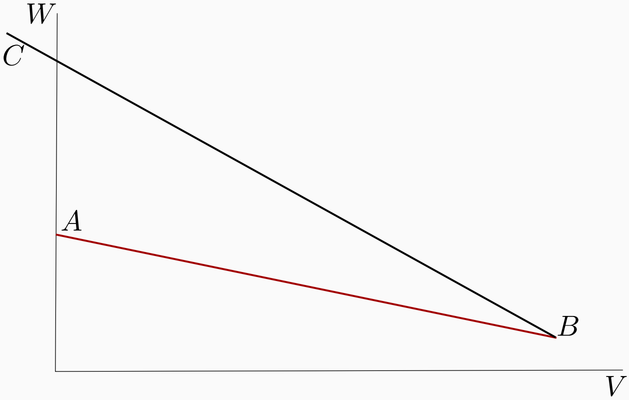

Partial Delayed Disclosure

Now consider transferring surplus from the agent to the principal by reducing information.

Step Two

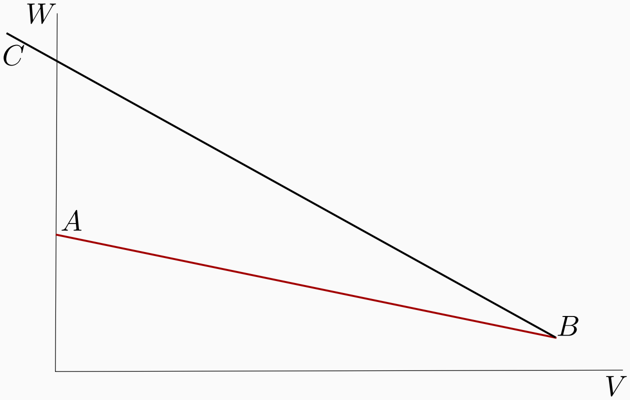

Point C represents no disclosure which would give the agent a negative payoff.

The agent likes risk more than the principal dislikes it.

Because $r_p \leq r$.

Intuition:

- Exponential discounting is like constant absolute risk aversion/affinity.
- The movement to point B is about tolerance to risk (second derivative).
- The movement to point C is about marginal value/cost of effort (first derivative).

Slightly more formally

- We are comparing
$$\frac{\Delta_p^2}{\Delta_a^2} \text{ with } \frac{\Delta_p^1}{\Delta_a^1}$$
- Equivalently Δ_p^1/Δ_p^2 to Δ_a^1/Δ_a^2.
- These are the ratios of risk tolerance to marginal value/cost of effort.
- Essentially the coefficient of absolute risk aversion/affinity, i.e. r, r_p.

General Analysis

We frame the analysis in terms of *effort schedules*: joint distributions over thresholds and effort durations (whose marginal is the exogenous F).

- Implementable schedules
- Efficient schedules
- Optimal implementable schedules

The schedule g^∞

Consider the *pure* schedule g^∞ in which the agent works until x with probability 1.

- It yields expected discounted effort $\mathbf{E}_F\left(1 - e^{-r_p x}\right)$.
- It is the most efficient way to provide the principal that payoff.
- Is it implementable? (Note that it has the agent working arbitrarily long.)

Implementable Schedules

A schedule is implementable if and only if it provides the agent at least his no-information value at all times.

Implementing g^∞

Assume that F has an increasing hazard rate. Then the schedule g^∞ is implementable if and only if it gives the agent a non-negative *ex ante* expected payoff.

Extracting Surplus

If g^∞ gives the agent a strictly positive *ex ante* payoff the principal extract surplus by implementing more effort.

The efficient way to do this is to front-load effort, witholding information as an incentivizing carrot (in addition to the role it is already playing.)

The schedule g^*

There exists a t^* such that the following schedule

$$g^*(x) = \begin{cases} t^*, & \text{if } x \leq t^* \\ x & \text{otherwise.} \end{cases}$$

gives the agent zero *ex ante* expected utility.

Leading the Agent On

The mechanism t^* is efficient.

Under the following condition it is also implementable and therefore optimal.

$$\frac{f'(t)}{f(t)} < r, \qquad \text{(A1)}$$

The implementing mechanism is *leading the agent on.*

Moving the Goalposts

If g^{∞} is not implementable then the following schedule is optimal

$$g^*(x) = \begin{cases} x & \text{if } x \leq t^* \\ 0 & \text{otherwise.} \end{cases}$$

where t^* is chosen so that the agent's *ex ante* expected payoff is zero.

It is implementable by a policy of *moving the goalposts*