Lecture notes on Moral Hazard, i.e. the Hidden Action Principle-Agent Model*

Allan Collard-Wexler

April 19, 2012

*Co-Written with John Asker and Vasiliki Skreta

Reading for next week: "Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information" by Baker and Hubbard, American Economic Review.

Introduction

The issue of moral hazard is key to understanding several topics, most notably how firms are organized, different compensation schemes such as CEO pay, and also many of the monitoring schemes that companies have in place, such as Board of Directors, auditing departments, and more generally accounting and control mechanisms within firms.

The purpose of this lecture is to introduce some of the basic models that revolve around the fact that the principle cannot usually observe exactly what agents are doing, and in particular if agents are doing what would be in the best interests of the firm or just in the best interests of the agents.

The key reason we are interested in moral hazard problems, is that the wrong contractual form can lead to inefficient effort on the part of agents, and outcomes which aren't Pareto efficient.

An individual - the principal - hires another individual - the agent - to perform a task which is determined by an unobservable action chosen by the agent. We study the design of contracts in a setting where actions are not observable: $hidden\ actions$ or $moral\ hazard$ settings.

Examples of moral hazard:

- owner manager relationship
- insurance companies and insured individuals
- manufactures and their distributors
- banks and their borrowers

We analyze the owner-manager case.

The owner hires the manager to perform a project. The manager's actions affect the profitability of the project. However, these actions are not observable by the owner, hence he cannot use a compensation scheme that depends on the actions of the manager. The question we aim to answer is how does an optimal contract look like in this case? Does it always implement an efficient choice of effort level by the manager? In general, the answer to this question is negative. Sometimes providing incentives for that we not be worthwhile - in the sense that it is too costly to induce the correct action.

Simple Model

The agent:

The agent's utility function is:

$$u(w) - a \tag{1}$$

where a is the action taken by the agent, and w is the agent's wealth.

The agent's utility function $u(\cdot)$ is strictly concave, hence the agent is risk averse. This is the key reason for having a principle in the setup, since the principle will be less sensitive to risk than the agent, and hence the optimal solution won't just be to sell the firm to the agent. Another reason for which agents might appear to be risk-averse is the problem of limited liability: agents may have limited capital to start a firm, and if they fail it will be impossible to recuperate a large amount of losses. In general, whenever faced with principal agent problems and your first reaction should be: what's stopping the agents from owning firms, and if the agents did own firms would this eliminates any inefficiency in the model.

Finally, the agent can also choose not to work for the principal, in which case they receive their outside option $\bar{U}$.

Two Action Model

Suppose the agent can take either action $a = 1$ corresponding to high effort, or action $a = 0$ correspondent low effort. If the agent exerts effort, there is a probability P of success, which generates output x^S, and and the probability $1 - P$ of failure which will generate output x^F, with $x^S > x^F$. If the agent does not exert effort, then there is only a probability $p < P$ of success.

The principal:

The principal's utility function is:

$$R = \begin{cases} P(x^S - w^S) + (1 - P)(x^F - w^F) \\ \text{(if the agent exerts effort, i.e. } a = 1) \\ p(x^S - w^S) + (1 - p)(x^F - w^F) \\ \text{(if the agent does not exert effort, i.e. } a = 0) \end{cases}$$

Notice that this implies the principal is risk neutral.

How to make the agent exert effort?

There are two important conditions that need to be satisfied in order for the agent to exert effort.

The first condition is called the incentive compatibility constraint (or IC constraint), which says that the agent must get a higher utility from exerting effort than for not exerting effort:

$$Pu(w^S) + (1-P)u(w^F) - 1 \geq pu(w^S) + (1-p)u(w^F)$$

More generally the IC constraint implies that the agent must get a higher payoff from taking the action that the principal wants her to take that from taking any other action available to him. It appears in virtually all papers of mechanism design.

Individual Rationality:

The second condition is called the individual rationality (or IR constrained) which says that the agent must prefer working for the principle, then not working and receiving his outside option:

$$Pu(w^S) + (1-P)u(w^F) - 1 \geq \bar{U}$$

Note that I have written the IR constraint assuming that the agent chooses to exert effort, but if it is optimal for the agent not to exert effort the IR constraint will have to hold it for the low effort choice.

Theorem 1 *The IR constraint will bind, i.e. will be satisfied with equality, for the the optimal contract.*

To show that the IR constraint will bind, I will do a proof by contradiction. If $Pu(w^S) + (1-P)u(w^F) - 1 > \bar{U}$, then there exists an ϵ amounts by which I can lower the wages in the successful and unsuccessful state without violating the IR constraint, i.e. $Pu(w^S - \epsilon) + (1-P)u(w^F - \epsilon) - 1 > \bar{U}$. Now let's check the IC constraint, which we can rearrange as:

$$(P-p)u(w^S - \epsilon) - 1 \geq (P-p)u(w^F - \epsilon)$$

However since $w^S > w^F$, then by concavity of utility function the new IC will hold as well. Since both the IR and IC constraint still hold, this means that the principal's payoff can be raised by ϵ, which violates the assumption that the contract was optimal in the first place.

Theorem 2 *The IC constraint will bind, i.e. will be satisfied with equality, for the the optimal contract.*

To see why the IC constraint needs to bind, I can do a proof by contradiction and suppose suppose that the IC constraint does not bind. Then I can lower the gap between w^S and w^F. Since $u(\cdot)$ is concave, this makes it easier to satisfy the IR constraint, i.e.

$$Pu(w^S - \delta) + (1 - P)u(w^F + \delta) - 1 \geq \bar{U}$$

so I can lower the wages here by ϵ

$$Pu(w^S - \delta + \epsilon) + (1 - P)u(w^F + \delta - \epsilon) - 1 \geq \bar{U}$$

but then the original contract wasn't profit maximizing.

Multiple Outcomes Version

$e-$ the action chosen by the agent: we will call it effort. For simplicity we will assume that effort can just take two values and it is either low or high:

$e \in \{e_L, e_H\}$

Manager's utility: $u(w, e) = v(w) - g(e)$, where w stands for wage and e stands for effort.

$v' > 0$ and $v'' < 0$

$g(e_H) > g(e_L)$

$\bar{u}$ is the managers reservation utility.

The profit of the project π is random and can take values in $\{\pi_1, ..., \pi_n\}$, where $\pi_1 < .... < \pi_n$. The distribution of profits - how likely each profit level is depend on the effort exerted by the manager. Let $p(\pi_k |e)$ denote the probability that π_k is realized when the manager chooses effort e.

Exerting high effort versus low effort is more costly for the manager in terms of disutility but it makes higher profits more likely. In order to formalize this, that is, that higher effort induces higher profits we use the notion of first order stochastic dominance. That is we assume that $p(. |e_H)$ first order stochastically dominates $p(. |e_L)$, which means that for any $m \in \{1, ..., n\}$

$\Sigma_{k=1}^{m} p(\pi_k |e_H) \leq \Sigma_{k=1}^{m} p(\pi_k |e_L)$.

Where the CDF looks like:

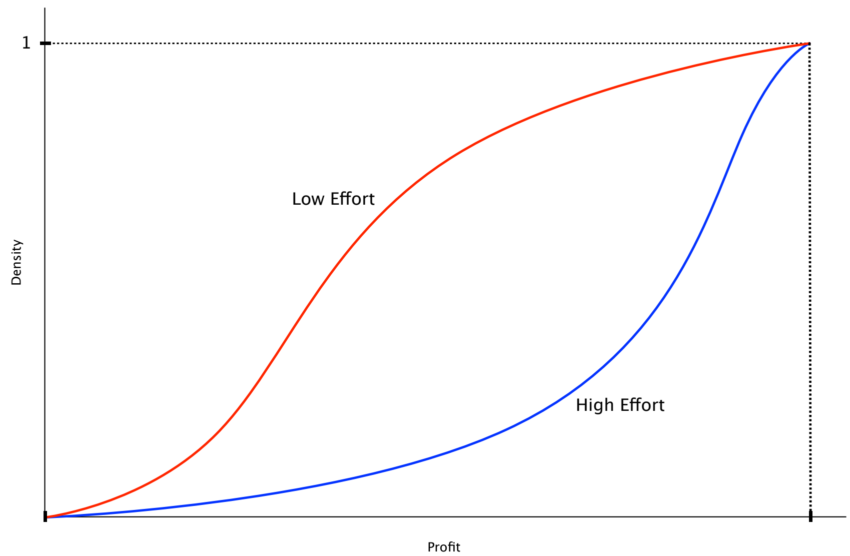

The firm owner's problem is as follows: Choose e and the wage scheme w that maximizes profits: $\pi - w$.

Benchmark Model: The optimal contract when effort is **observable**.

Wage payment can depend both on e and π : $w(e, \pi_k) = w_k(e)$

$w_1(e_L), ..., w_n(e_L); w_1(e_H), ..., w_n(e_H)$

The principle's problem is:

$\max_{\substack{e \in \{e_L, e_H\} \\ w_1(e_L), ..., w_n(e_L); w_1(e_H), ..., w_n(e_H)}} \Sigma_{k=1}^{n} \{\pi_k - w_k(e)\} p(\pi_k \,|e)$

such that $\Sigma_{k=1}^{n} v(w_k(e)) p(\pi_k \,|e) - g(e) \geq \bar{u}$. (P)

where (P) is called the participation constraint

This problem can be solved in two steps:

1) For each choice of e, what is the best compensation scheme?

2) What is the optimal level of effort?

Step 1: For each choice of e, what is the optimal compensation scheme?

Suppose that the owner wants to implement effort level e : Then the problem reduces to finding the cheapest way to do so. That is the minimum level of way that will induce the manager to choose effort level e :

$Min_{w_1(\widehat{e}),.....,w_n(\widehat{e})} \Sigma_{k=1}^{n} w_k(\widehat{e}) p(\pi_k \,|\widehat{e})$

such that $\Sigma_{k=1}^{n} v(w_k(\widehat{e})) p(\pi_k \,|\widehat{e}) - g(\widehat{e}) \geq \bar{u}$. (P)

It is easy to see that the participation constraint will be binding at the optimum.

The Lagrangian for this problem is given by

$\mathcal{L}(\widehat{w}_1, ..., \widehat{w}_n, \lambda) = \Sigma_{k=1}^{n} \widehat{w}_k p(\pi_k \,|\widehat{e}) - \lambda \{\Sigma_{k=1}^{n} v(w_k(\widehat{e})) p(\pi_k \,|\widehat{e}) - g(\widehat{e}) - \bar{u}\}$

where $\lambda \geq 0$.

FOC's

$\frac{\partial \mathcal{L}}{\partial \widehat{w}_k} = p(\pi_k \,|\widehat{e}) - \lambda v'(\widehat{w}_k)_k p(\pi_k \,|\widehat{e}) = 0$, $k = 1, 2, ..., n$

Assume that $p > 0$ then we get $\frac{1}{\lambda} = v'(\widehat{w}_k) > 0$

hence the participation constraint is binding

This condition implies that is the case that the manager's effort is observable - the firm owner will pay the manager the same wage offer all the time:

$v'(\widehat{w}_1) = ... = v'(\widehat{w}_n) = \frac{1}{\lambda}$; $v'' < 0$, v' is weakly decreasing and hence it is a $1-1$ function.

$\widehat{w}_1 = ... = \widehat{w}_n$.

Now we will use the fact that the participation constraint is binding to get the optimal level of w :

$v(w_e^*) - g(e) = \bar{u}$.

$w_e^* = v^{-1}(\bar{u} + g(e^*))$: This is the optimal wage conditional on e^* being implemented.

Step 2: What is the optimal level of effort being implemented?

$e^* = e_L$ or $e^* = e_H$.

If $e^* = e_L$ then $w_L^* = v^{-1}(\bar{u} + g(e_L))$ and the owner's expected payoff is given by

$\Sigma_{k=1}^{n}(\pi_k - v^{-1}(\bar{u} + g(e_L)))p(\pi_k \,|e_L) = \Sigma_{k=1}^{n}\pi_k p(\pi_k \,|e_L) - v^{-1}(\bar{u} + g(e_L))$

If $e^* = e_H$ then $w_H^* = v^{-1}(\bar{u} + g(e_H))$and the owner's expected payoff is given by

$\Sigma_{k=1}^{n}(\pi_k - v^{-1}(\bar{u} + g(e_H)))p(\pi_k \,|e_H) = \Sigma_{k=1}^{n}\pi_k p(\pi_k \,|e_H) - v^{-1}(\bar{u} + g(e_H))$

Proposition. In the principal-agent model with observable managerial effort, an optimal contract specifies that the manager choose the effort e^* that maximizes $\Sigma_{k=1}^{n}\pi_k p(\pi_k \,|e) - v^{-1}(\bar{u} + g(e))$

and pays the manager a fixed wage $w^* = v^{-1}(\bar{u} + g(e^*))$. This is the uniquely optimal contract if $v''(w) < 0$ at all w.

The optimal contract when effort is **not observable**.

The optimal contract in the case of observable achieved two goals: It specifies the optimal level of effort by the manager and it insures him against risk. When effort is not observable the two goals may be in conflict. To disentangle the roles of the choice of the optimal level of effort and risk-sharing we will first study the case of a risk neutral manager.

Before doing this let us note that the set of compensation schemes available to the owner in the case that effort is observable is strictly larger than the one available to the owner in the case of unobservable effort. With observable effort the principal can make compensation scheme proposed to the agent contingent on the effort level he chooses. This kind of compensation schemes are not available to the owner when effort is not observable. Hence the choice set in the case of observable effort is strictly larger then the one with unobservable effort, which implies that the value of the owner's program is (weakly) higher then the value of owners program with unobservable effort.

We will now demonstrate that the case that the manager is risk neutral the maximized values of these programs is the same

In order to do so, we consider first the manager's problem in the case that effort is observable. The owners problem is given by

$$\max_{e \in \{e_L, e_H\}} \Sigma_{k=1}^{n} \pi_k p(\pi_k \,|e) - \bar{u} - g(e)$$

The owner's profit in this case is the value of the program and the manager receives utility at least $\bar{u}$.

The following proposition establishes this point.

Proposition. In the principal-agent model with unobservable managerial effort and a risk neutral manager, an optimal contract generates the same effort choice and expected utilities for the manager and the owner as when effort is observable.

Proof:

The owner offers the manager a compensation schedule $w_k = \pi_k - \alpha$, where α is some constant. If the manager accepts this contract, he chooses e to maximize his expected utility

$\Sigma_{k=1}^{n} w_k p(\pi_k \,|e) - g(e) = \Sigma_{k=1}^{n} \pi_k p(\pi_k \,|e) - g(e) - \alpha.$

The solution to manager's problem solves also the owner's problem in the case of non-observable effort. Let e^* denote the optimal effort level chosen by the manager given this sell-out contract. The manager is willing to accept this contract so long as

$\Sigma_{k=1}^{n} \pi_k p(\pi_k \,|e^*) - \alpha - g(e^*) \geq \bar{u}$. Let α^* denote the α such that $\Sigma_{k=1}^{n} \pi_k p(\pi_k \,|e^*) - g(e^*) - \bar{u} = \alpha^*$. Given a compensation scheme of $w_k = \pi_k - \alpha^*$ the owner and the manager get exactly the same payoff as with observable effort.

The case of risk-averse manager

Since the optimal compensation scheme cannot be conditioned on the optimal effort level, it has to be designed in such a way that it induces it. Suppose that the owner wants the manager to choose effort level e. Then in the case that effort is observable the owner seeks for the cheapest way that gets the manager to agree to do the project - that is the owner chooses w to minimize payments over the set of wages that guarantee the manager expected utility level of $\bar{u}$ when he chooses e. When effort e is not observable the owner faces an additional constraint: in order that w implements e it must belong in the set of compensation schemes such that the agent prefers to choose e versus choosing e'. This kind of constraint is called incentive constraint.

The owner's problem in the case of unobservable effort and risk averse manager is given by

$Min_{w_1,\ldots,w_n} \Sigma_{k=1}^{n} w_k p(\pi_k \,|e)$

such that $(i) \Sigma_{k=1}^{n} v(w_k) p(\pi_k \,|e) - g(e) \geq \bar{u}$ (P)

participation constraint

(ii) e solves $\max_{\tilde{e}} \sum_{k=1}^{n} v(w(\pi_k)) p(\pi_k \,|\tilde{e}) - g(\tilde{e})$ (IC)

incentive constraint

Implementing e_L : in this case the optimal compensation scheme is for the owner to offer a fixed wage payment $w_e^* = v^{-1}(\bar{u} + g(e_L))$, the payment he would offer if contractually specifying effort e_L when effort is observable. This compensation scheme is not affected by the effort exerted by the manager - hence he will choose the lowest possible level of effort, which is indeed e_L.

Implementing e_H. First note that the incentive constraint can be rewritten as:

$\Sigma_{k=1}^n v(w_k) p(\pi_k | e_H) - g(e_H) \geq \Sigma_{k=1}^n v(w_k) p(\pi_k | e_L) - g(e_L)$ (IC)

that is the owner chooses among the compensation schemes that satisfy this condition.

Let $\lambda \geq 0$ and $\mu \geq 0$ denote the Lagrange multipliers associated with the constraints (P) and (IC).

Lemma. When the owner wants to induce high effort both constraints (P) and (IC) are binding that is $\lambda > 0$ and $\mu > 0$.

The Lagrangian for this problem is given by

$L(w_1, ..., w_n, \lambda, \mu) = \Sigma_{k=1}^{n} w_k p(\pi_k \,|e_H)$

$-\lambda\{\Sigma_{k=1}^{n} v(w_k) p(\pi_k \,|e_H) - g(e_H) - \bar{u}\}$

$-\mu\{\Sigma_{k=1}^{n} v(w_k) p(\pi_k \,|e_H) - g(e_H)$

$-\Sigma_{k=1}^{n} v(w_k) p(\pi_k \,|e_H) + g(e_H)\}$

FOC's

$\frac{\partial L}{\partial w_k} = p(\pi_k \,|e_H) - \lambda v'(w_k) p(\pi_k \,|e_H)$

$-\mu\{v'(w_k) p(\pi_k \,|e_H) - v'(w_k) p(\pi_k \,|e_L)\} = 0$

$\lambda \geq 0, \mu \geq 0$

Complementary Slackness

Divide by $v'(w_k) p(\pi_k \,|e_H) > 0$

$\frac{1}{v'(w_k)} - \lambda - \mu \left[1 - \frac{p(\pi_k|e_L)}{p(\pi_k|e_H)}\right] = 0,\ k = 1, ..., n$

Suppose that $\lambda = 0$ then

$\frac{1}{v'(w_k)} = \mu \left[1 - \frac{p(\pi_k|e_L)}{p(\pi_k|e_H)}\right]$

Recall that $p(.|e_H)$ first order stochastically dominates $p(.|e_H)$

For any m :$\Sigma_{k=1}^{m} p(\pi_k |e_H) \leq \Sigma_{k=1}^{m} p(\pi_k |e_L)$

for $m = 1$ $p(\pi_1 |e_H) \leq p(\pi_1 |e_L)$

$1 \leq \frac{p(\pi_k|e_L)}{p(\pi_k|e_H)}$ contradiction. Hence $\lambda > 0$ which implies that the participation constraint is binding.

We now proceed to show that (IC) is binding as well. Suppose that $\mu = 0$, then

$\frac{1}{v'(w_1)} = ... = \frac{1}{v'(w_n)} = \lambda$ is and only if $w_1 = ... = w_n$, but then this violates the incentive constraint since the manager has no incentive to exert high effort. Contradiction. Hence the wage payment will depend on k.

Observations:

1)Recall that $\pi_1 < .... < \pi_n$. Does this imply that at the optimum we have $w_1 \leq .... \leq w_n$? No. In order that the optimal compensation scheme to be increasing it must be the case that $\frac{p(\pi_k|e_L)}{p(\pi_k|e_H)}$ is decreasing π_k,

$$\frac{p(\pi_1|e_L)}{p(\pi_1|e_H)} \geq \frac{p(\pi_2|e_L)}{p(\pi_2|e_H)} \geq .... \geq \frac{p(\pi_k|e_L)}{p(\pi_k|e_H)}.$$

which implies that as π increases the likelihood of getting profit level π if effort is e_H relative to the likelihood if effort is e_L must increase. This is known as the monotone likelihood ratio property. MLRP implies first order stochastic dominance but not the opposite.

2) From the FOC it follows that the optimal compensation scheme is likely to be non-linear

3) The fact that effort is non-observable increases the owner's compensation cost of implementing high effort. First note that we have $E[v(w(\pi)\,|e_H] = \bar{u} + g(e_H)$ and $v''() < 0$. From Jensen's inequality it follows that $v(E[w(\pi)\,|e_H]) > \bar{u} + g(e_H)$ - since v is concave. But from the previous analysis we have that $v(w^*_{e_H}) = \bar{u} + g(e_H)$, and so $E[w(\pi)\,|e_H] > w^*_{e_H}$.

Hence unobservability increases the cost of implementing e_H but does not change the cost of implementing e_L- this means that sometimes an inefficiently low level of effort will be implemented. As usual, in order for the owner to choose the optimal effort level he must compare the benefits, with the cost

benefits: $\Sigma_{k=1}^{m} p(\pi_k | e_H)\pi_k - \Sigma_{k=1}^{m} p(\pi_k | e_L)\pi_k$

costs: the difference in expected wage

Proposition. In the principal-agent model with unobservable manager effort, a risk-averse manager, and two possible effort choices, the optimal compensation scheme for implementing e_H satisfies the FOC given above, given the manager expected utility level of $\bar{u}$ and involves a larger expected wage payment than is required when effort is observable. The optimal compensation scheme for implementing e_L involves the same fixed payment as if effort were observable. Whenever the optimal effort level with observable effort would be e_H, nonobservability causes a welfare loss.

The First-Order Approach:

The first-order approach deals with the problem when there are an infinite number of actions that can be chosen, i.e. a continuous action space. What's nice about the first-order approaches that you can use derivatives, however there are some serious technical conditions that need to be applied in order for this derivative-based approach to be valid such as Rogerson(1985)'s Econometrica.

Suppose that the probability of success is given by $p(a)$ and $p(a)$ is increasing and concave.

Let's first showed that $p(a)$ satisfies two important conditions for the the first-order approach to work:

1. Convexity of the distribution function condition (CDFC): This just means that:

$$a_j = \lambda a_i + (1-\lambda)a_k$$

   where $\lambda \in [0,1]$ then by CDFC:

$$P_{jl} \leq \lambda P_{il} + (1-\lambda)P_{kl}$$

   for all outcomes l. For the continuous problem this implies that:

$$p(\lambda\bar{a} + (1-\lambda)\underline{a}) \leq \lambda p(\bar{a}) + (1-\lambda)p(\underline{a})$$

2. Monotone Likelihood Ratio Property (MLRP)

This is equivalent to saying that $\frac{1-p(a)}{p(a)}$ is decreasing in a. So

$$\frac{d}{da}\left[\frac{1-p(a)}{p(a)}\right] = -p'(a)\left(\frac{1-p(a)}{p(a)^2} + \frac{1}{p(a)}\right)$$

which will be negative since $p(a)$ is increasing.

Let's suppose that the IR and IC constraints both blind. First let's take a look at the IR constraint:

$$p(a)u(w^s) + (1-p(a))u(w^f) - a = \bar{U}$$

Next the IC constraint can be derived from the fact that the individuals payoff is just:

$$p(a)u(w^s) + (1-p(a))u(w^f) - a$$

So if we take the first-order condition with respect to a, i.e. I choose the action that maximizes my payoff then we obtain:

$$\frac{\partial U}{\partial a} = p'(a)u(w^s) - p'(a)u(w^f) - 1 = 0$$

Okay now let's derive the wages in the case of success and failure given that the IC in the IR constraint will bind. From the IC we get:

$$u(w^s) = u(w^f) + \frac{1}{p'(a)}$$

Then plugging this into the IR constraint:

$$p(a)\left(u(w^f) + \frac{1}{p'(a)}\right) + (1 - p(a))u(w^f) = \bar{U} + a$$

$$u(w^f) = \bar{U} + a - \frac{p(a)}{p'(a)}$$

And likewise for the wages of success we get:

$$u(w^s) = \bar{U} + a - \frac{1 - p(a)}{p'(a)}$$

Notice that if there wasn't an IC constraint, then I would would pay the agents just $\bar{U} + a$ a flat wage. I need to get a higher wage for success in order to induce effort, and this magnitude of this inducement depends on the marginal effect of efforts on the probability of success. So if effort translates into a very large change in the probability of success, then after will be easy to detect and I get closer to the case where effort is fully observable.

Repeated Moral Hazard: Holmstrom 1982 First, some notation:

Agent's ability: $\eta_{t+1} = \eta_t + \delta_t$, where δ is a random i.i.d. normal shock.

Action: a_t

Output: $y_t = \eta_t + a_t + \epsilon_t$, where ϵ_t is a random i.i.d. normal shock, and a_t is the agent's effort.

History: The agent's history is $y^{ht} = (y_1, \cdots, y_{t-1})$.

Wages are set competitively, but aren't dependent on current output, just your employment record:

$$w_t(y^{ht-1}) = E[y_t | y^{ht-1}]$$

Agent's Utility:

$$\sum_{t=1}^{\infty} \beta^{t-1}[c_t - g(a_t)] = \sum_{t=1}^{\infty} \beta^{t-1}[w_t - g(a_t)]$$

essentially, risk neutral agent, with no motive to save here, so $c_t = w_t$.

Efficiency:

$$\underbrace{g'(a_t)}_{\text{MC}} = \underbrace{\frac{\partial[\eta_t + a_t + \epsilon_t]}{\partial a_t}}_{\text{MB}} = 1$$

Two Period Model, $t = 1, 2$

Let's start with a two-period model to get some intuition. We can work by backward induction.

- $t = 2$

In period 2 the agent does not work, so $a_2 = 0$. Thus the wage to the agent is just

$$y_2 = E[\eta_2 + a_2 + \epsilon_2] = E[\eta_2]$$

The expectation of $\eta_2 = E[y_1 - a_1 - \epsilon_1] = y_1 - a_1^*$.

- $t = 1$

In period 1 the agent has to choose wether to work. The agent's wage is

$$y_1 = E[\eta_1 + a_1^* + \epsilon_1] = E[\eta_1] + a_1^*$$

- Finding a_1, the first-order condition is:

$$-g'(a_1) + \beta \frac{\partial (y_1 - a_1^*)}{\partial a_1} = -g'(a_1) + \beta = 0$$

So the agent works a bit less that would be implied by the first-best where $g'(a_1) = 1$ instead of $g'(a_1) = \beta$.

Stuff on Linear Filtering Problems:

Suppose that X and Y are independent normal variables. $X \sim \mathcal{N}(\mu, \sigma_X)$ and $Y \sim \mathcal{N}(0, \sigma_Y)$.

Define $h_X = 1/\sigma_X^2$ and $h_Y = 1/\sigma_Y^2$

We want to know $X|X+Y=q$, i.e. we observe a noisy signal q on X. It can be shown that:

$$X|(X+Y=q) \sim \mathcal{N}(\frac{h_X}{h_X+h_Y}\mu + \frac{h_Y}{h_X+h_Y}q, \frac{1}{h_X+h_Y})$$

Say I don't know anything about X, i.e. $X \sim \mathcal{N}(\mu, \infty)$ then $h_X = 0$ and:

$$X|(X+Y=q) \sim \mathcal{N}(q, \frac{1}{h_Y})$$

End of $t = 1$

$$\eta_1 | (\eta_1 + \epsilon_1 = y_1 - a_1^*) \sim \mathcal{N}(\frac{h_1}{h_1 + h_\epsilon} 0 + \frac{h_\epsilon}{h_1 + h_\epsilon}(y_1 - a_1), \frac{1}{h_1 + h_\epsilon})$$
$$= \mathcal{N}(\frac{h_\epsilon}{h_1 + h_\epsilon}(y_1 - a_1), \frac{1}{h_1 + h_\epsilon})$$

Begining of $t = 2$

$$\eta_2 | (\eta_1 + \epsilon_1 = y_1 - a_1^*) \sim \mathcal{N}(\underbrace{\frac{h_\epsilon}{h_1 + h_\epsilon}(y_1 - a_1)}_{\eta_2}, \underbrace{\frac{1}{h_1 + h_\epsilon} + \frac{1}{h_\epsilon}}_{1/h_2})$$

End of $t = 2$

$$\eta_2 | (\eta_1 + \epsilon_1 = y_1 - a_1^*), (\eta_2 + \epsilon_2 = y_2 - a_2^*) \sim$$
$$\mathcal{N}(\frac{h_2}{h_2 + h_\epsilon}\eta_2 + \frac{h_\epsilon}{h_2 + h_\epsilon}(y_2 - a_2), \frac{1}{h_2 + h_\epsilon})$$

Beginning of $t = 3$

$$\eta_2 | (\eta_1 + \epsilon_1 = y_1 - a_1^*), (\eta_2 + \epsilon_2 = y_2 - a_2^*) \sim$$
$$\mathcal{N}(\underbrace{\frac{h_2}{h_2 + h_\epsilon}\eta_2 + \frac{h_\epsilon}{h_2 + h_\epsilon}(y_2 - a_2)}_{\eta_3}, \underbrace{\frac{1}{h_\epsilon} + \frac{1}{h_2 + h_\epsilon}}_{h_3})$$

So by induction we get:

$$\eta_{t+1} = \frac{h_t}{h_t + h_\epsilon}\eta_t + \frac{h_\epsilon}{h_t + h_\epsilon}y_t$$

By the way this property that I just need to keep track of two numbers to summarize the entire history of the agent's performance, just the mean and the variance of my prediction is a very special property the linear model with normal measurement error. If you try to tweak this model in any direction on it gets very hard to preserve this feature that there's only two summary statistics that you need to keep track of. Dealing with histories of quickly leads you to very intractable models.

Let's get back to:

$$m_{t+1} - m_t = \lambda_t[y_t - a_t^* - m_t]$$

Suppose the agent deviates and expends a tiny bit more effort, which I'll call da. The induced change in the principles estimate of the agents ability is:

$$dm_{t+1} - dm_t = -\lambda_t da$$

Okay, how does the effect total wages over time, and let's assume we've reached a t large enough so that $\lambda_t \to \lambda$. Let's look at the wage change tomorrow:

$$dw_1 = \lambda$$

An likewise in s periods from now:

$$dw_s = \lambda(1 - \lambda)^{s-1} da$$

since the effect of effort today on my reputation for ability decays over time, with parameter $1 - \lambda$.

Okay now let's add up all these effects:

$$\begin{aligned}
dU &= dw_1 - g'(a)da + \sum_{t=2}^{\infty} \beta^t dw_t \\
&= \lambda da - g'(a)da + \sum_{t=1}^{\infty} \beta^t \lambda (1-\lambda)^{t-1} da \\
&= -g'(a)da + \lambda \sum_{t=0}^{\infty} (\beta(1-\lambda))^t da \\
&= -g'(a)da + \frac{\lambda\beta}{1-\beta(1-\lambda)} da = 0
\end{aligned}$$

Remember the efficient level of effort if $g'(a) = 1$. So the whether the agent will provide more or less effort that what is optimal depends on if:

$$\frac{\lambda\beta}{1-\beta(1-\lambda)} \geq\leq^{?} 1$$

If $\beta < 1$ then $\frac{\lambda\beta}{1-\beta(1-\lambda)} \leq 1$, so less effort is provided that what is optimal. Essentially this happens since the rewards of effort are delayed in the future. What happens is the consequences of effort are growing over time. This may be reversed.

Empirical Applications

1. Lafontaine on Franchising.

2. Partnerships versus Corporations.

3. Hubbard (monitoring).

4. Hubbard and Baker (monitoring again).