Microeconomics 2
Lecture notes

Winter 2012

Notes from Lectures by P.Ray (TSE)

Contents

Chapter 1

Moral Hazard

1.1 Introduction

Asymmetric information

1. Adverse selection: Mechanism designer seeks to have agents report certain information.

2. Moral hazard: Mechanism designer seeks to have agents take certain actions.

Examples

- salesman's effort

- managers's decision

These situations involve a trade-off. In adverse selection, the trade off is between efficiency and cost (information rent). In moral hazard situations, we have a similar trade-off, this time between efficiency and the cost of incentive compatibility. This cost takes two forms:

1. Risk-sharing, and

2. Rents.

We will look at two simple illustrations of these two costs.

1.1.1 Efficiency versus Risk-Sharing

Consider an interaction between two parties. The principal, a firm, seeks to maximize $q - w$, where q is the output and w is the wage; the price of the output is normalized to one. The agent, a worker, has utility $u(w) - e$, where e is the effort the agent exerts. We assume that the agent has a reservation level of utility $\hat{u} = u(\hat{w})$, where $u' > 0 > u''$: u is concave, and thus the agent is risk-averse. The output can take one of two values:

$$q = \begin{cases} Q & \text{with probability} \quad p \\ 0 & \text{with probability} \quad 1-p \end{cases}$$

The agent can choose either to work with effort level $e = E$, in which case the probability that the project is successful p is p_E; or to shirk and exert effort $e = 0$, in which case the probability of success p is $p_0 < p_E$. We will assume that $p_E Q - E \geq p_0 Q, \hat{u}$, so that $e = E$ is efficient.

The timing of the interaction is as follows:

- First, the principal offers a contract to the agent.
- The agent then accepts or refuses the contract.
- If the agent refuses the contract he gets a reservation utility $\hat{u}$. If the contract is accepted, the agent then chooses the level of effort $e \in \{0, E\}$, which is unobservable by the principal.
- Finally, as a result of the agent's choice, a quantity q is produced.

Complete Information

Under complete information in which all is observable, a contract can specify the effort level $e = E$ and a wage schedule of the form $\bar{w}$ if $q = Q$ and $\underline{w}$ if $q = 0$. The principal's profit-maximization problem is

$$\begin{aligned} &\max_{\bar{w}, \underline{w}} p_E Q - (p_E \bar{w} + (1 - p_E)\underline{w}) \\ &\text{s.t.} \quad p_E u(\bar{w}) + (1 - p_E) u(\underline{w}) - E \geq \hat{u} \quad (IR) \end{aligned}$$

Since the agent is risk-averse ($u'' < 0$) whereas the principal is risk-neutral, the optimal contract is such that $\bar{w} = \underline{w}$ (perfect insurance); the wage level is then set so as to meet the agent's participation constraint: and $u(\underline{w}) = u(\bar{w}) = \hat{u} + E$. The condition $p_E Q - E \geq p_0 Q$ then ensures that this contract is better than inducing the agent to exert no effort ($e = 0$), whereas $p_E Q - E \geq \hat{u}$ ensures that contracting is better than no contracting.

Incomplete Information

Now suppose that the principal does not observe the agent's effort level, but only the output. The parties can then only contract on a wage schedule $(\bar{w}, \underline{w})$. The principal must still meet the participation constraint of the agent, but must now moreover take into account incentive-compatibility considerations. Assuming that the principal would like to induce the effort level $e = E$, the principal's profit-maximization problem thus becomes

$$\max_{\bar{w},\underline{w}} p_E Q - [p_E \bar{w} + (1-p_E)\underline{w}]$$

$$\text{s.t.} \quad \begin{array}{rr} p_E u(\bar{w}) + (1-p_E)u(\underline{w}) - E \geq \hat{u}, & (IR) \\ p_E u(\bar{w}) + (1-p_E)u(\underline{w}) - E \geq p_0 u(\bar{w}) + (1-p_0)u(\underline{w}). & (IC) \end{array}$$

We can rewrite (IC) as

$$\underbrace{(p_E - p_0)}_{>0} \underbrace{(u(\bar{w}) - u(\underline{w}))}_{\text{bonus}} \geq E.$$

Since $p_E > p_0$, it follows that the bonus must be positive, and thus $\bar{w} \geq \underline{w}$: the agent is no longer perfectly insured by the principal. Since it is efficient to insure the agent as much as possible, this constraint is binding:

$$u(\bar{w}) - u(\underline{w}) = \frac{E}{p_E - p_0}.$$

Using this to express $u(\bar{w})$ as a function of $u(\underline{w})$, the participation constraint (IR) can be rewritten as

$$u(\underline{w}) + \frac{p_E E}{p_E - p_0} - E \geq \hat{u}.$$

Since the principal wants to minimize wages, this participation constraint is binding, which leads to:

$$u(\underline{w}) = \hat{u} + E - \frac{p_E E}{p_E - p_0} < \hat{u} + E,$$
$$u(\bar{w}) = \hat{u} + E + \frac{(1-p_E)E}{p_E - p_0} > \hat{u} + E.$$

Using Jensen's inequality,[1] it can be checked that the expected wage in the second-best world is higher then the first-best wage, to compensate the risk-averse agent for the risk taken (see Figure 1.1). The risk premium is $r = \mathbb{E}[w^{SB}] - w^{FB}$.

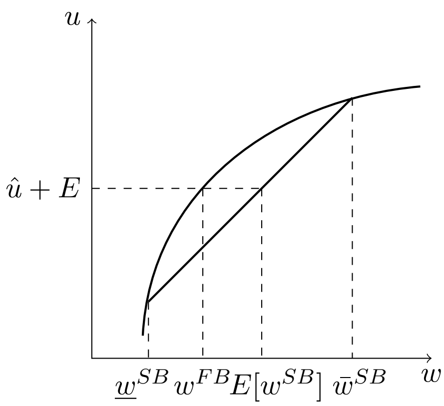

Figure 1.1:

That effort provision is efficient in the first-best world does not guarantee that it remains efficient in the second-best world; this also depends on the size of the risk premium. Indeed, if:

$$p_E Q - E - \hat{u} > 0 > p_E Q - E - \hat{u} - r,$$

then it is no longer optimal to induce the agent to exert effort; instead:

- If $p_0 Q - \widehat{u} > 0$, the second-best is to induce the agent to exert no effort: $e = 0, w = \hat{w}$ (distortion of the level of effort);
- If $p_0 Q - \widehat{u} < 0$, it is better not to contract (distortion of the level of trade).

1.1.2 Efficiency versus Rent

Assume now that the worker, too, is risk-neutral: $u(w) = w$, but has limited liability:

$$w(q) \geq \hat{w}. \quad (LL)$$

That is, the agent is always free to walk away from the job, should the contracted wage fall below his reservation wage.

Complete Information

Under complete information, the parties can contract on the effort level (the principal will pay a non-zero wage only if $e = E$) and the output $(w(q) = \underline{w}$

[1]For any concave function $f(x)$, $\mathbb{E}[f(x)] \leq f(\mathbb{E}[x])$.

if the output q is low, and $w(q) = \bar{w}$ if the output q is high). The principal's profit-maximization problem is thus

$$\max_{\bar{w},\underline{w}} p_E Q - [p_E \bar{w} + (1-p_E)\underline{w}]$$
$$\text{s.t.} \quad \begin{array}{cc} p_E \bar{w} + (1-p_E)\underline{w} - E \geq \hat{w}, & (IR) \\ \underline{w}, \bar{w} \geq \hat{w}. & (LL) \end{array}$$

Ignoring (LL), the principal and the agent only care about the expected wage $w^e = p_E \bar{w} + (1-p_E)\underline{w}$ and, as the principal wants to minimize it, the participation constraint is binding and yields $w^e = \hat{w} + E$. Conversely, setting $\underline{w} = \bar{w} = \hat{w} + E$ yields the desired expected wage and satisfies (LL); this is thus an optimal contract in this first-best. That is, limited liability has no effect on the first-best solution.

Incomplete Information

Under incomplete information, the principal cannot observe the agent's effort level and can only contract on a wage schedule $(\bar{w}, \underline{w})$. Assuming that the principal would like to induce the effort level $e = E$, the principal's problem becomes:

$$\max_{\bar{w},\underline{w}} p_E Q - [p_E \bar{w} + (1-p_E)\underline{w}]$$
$$\text{s.t.} \quad \begin{array}{cc} p_E \bar{w} + (1-p_E)\underline{w} - E \geq \hat{w}, & (IR) \\ p_E \bar{w} + (1-p_E)\underline{w} - E \geq p_0 \bar{w} + (1-p_0)\underline{w}, & (IC) \\ \underline{w}, \bar{w} \geq \hat{w}. & (LL) \end{array}$$

Using as variable the base wage $w = \underline{w}$ and the bonus $b = \bar{w} - \underline{w}$, the participation constraint boils down to

$$w + p_E b \geq \hat{w} + E,$$

whereas the incentive constraint can be rewritten as

$$w + p_E b - E \geq w + p_0 b \iff b \geq \frac{E}{p_E - p_0}.$$

The bonus must therefore be positive, which in turn implies that, or the two limited liability constraints $(w, w + b \geq \hat{w})$, only the former one (for the base wage w) matters. The principal problem can thus be expressed as:

$$\max_{\bar{w},\underline{w}} p_E Q - (w + p_E b)$$
$$\text{s.t.} \quad \begin{array}{cc} w + p_E b \geq \hat{w} + E, & (IR) \\ b \geq \frac{E}{p_E - p_0}, & (IC) \\ w \geq \hat{w}. & (LL) \end{array}$$

$$\max_{\underline{w},\bar{w}} p_E Q - [p_E \bar{w} + (1-p_E)\underline{w}]$$

Which of the three constraints are binding? If we ignore the participation constraint (IR), the other two must be binding and determine the wage schedule: $w = \hat{w}$ from (LL) and $b = \frac{E}{p_E - p_0}$ from (IC); but then, the participation constraint is indeed satisfied:

$$w + p_E b = \hat{w} + \frac{p_E E}{p_E - p_0} = \hat{w} + E + R > \hat{w} + E,$$

where

$$R \equiv \frac{p_0 E}{p_E - p_0} \tag{1.1}$$

is the rent that the principal has to leave to the agent on top what is needed to accommodate the limited liability and incentive constraints. Thus, in the second-best world, the participation constraint is not an issue in presence of limited liability: it is instead the limited liability constraint that binds.

The cost of inducing the first-best level of effort is not just the cost of effort E, but $E + R$; the principal must pay *more* than the cost of effort. Therefore, if:

$$p_E Q - E - \hat{w} > 0 > p_E Q - E - \hat{w} - R,$$

then:

- If $p_0 Q - \hat{w} > 0$, the second-best is to induce the agent to exert no effort: $e = 0, w = \hat{w}$;
- If $p_0 Q - \hat{w} < 0$, it is better not to contract.

Variant

Consider the following variant of the above model:

- $p_E = 1$ and $p_0 = 0$: the project succeeds for sure if the agent chooses to work $(e = E)$, and fails for sure otherwise $(e = 0)$;
- there is no cost to effort, but instead the agent enjoys a private benefit B if he chooses to shirk.

In the first-best, $e = E$ is efficient whenever $Q > B$, in which case the firm offers a wage $w = \hat{w}$ for $e = E$. In the second-best, limited liability $\underline{w} \geq \hat{w}$, and incentive compatibility requires $\bar{w} - \underline{w} \geq B$. At the (profit-maximizing) optimum, both constraints are binding and thus $\bar{w} = \hat{w} + B$ (which in turn implies that the participation constraint is satisfied: $\bar{w} > \hat{w}$). If

$$Q - B - \hat{w} < 0 < Q - \hat{w},$$

it can be the case that efficient trade is not taking place.

1.2 The Role of Statistical Inference

The principal's goal is to detect what the agent has done by observing related variables, i.e. variables related to those that are relevant but not observable. In general, the principal will observe imperfect signals of the agent's choice.

1.2.1 The Inference Problem

In the above examples, the principal seeks to infer the agent's unobserved effort e from the observed output q. Suppose more generally that there are n output levels $q_1 < q_2 < \ldots < q_n$, and the firm offers wages based on the realization of the output, $w_1 < w_2 < \ldots < w_n$. A natural question that arises is, should the wage increase with the observed output level? The answer is, "Not necessarily".

To see this, consider the following example:

- the agent is risk-averse (but no limited liability)
- three output levels, $\underline{q} < \hat{q} < \bar{q}$
- two levels of effort, $e \in \{0, E\}$
- the probability distribution of the various outputs, given the agent's effort level, is as follows:

p	$\underline{q}$	$\hat{q}$	$\bar{q}$
$e = 0$	$9/10$	$1/10$	0
$e = E$	$1/10$	0	$9/10$

Assuming that the principal wants to induce the agent to exert effort ($e = R$), what is the second-best optimal wage schedule, $w(q) = (\underline{w}, , \hat{w}, \bar{w})$? In the absence of limited liability constraints it is optimal to punish the agent "infinitely" if $\hat{q}$ is observed: $\hat{w} = -\infty$, as observing $\hat{q}$ would reveal that the agent shirked. This, in turn, would suffice to provide the effort incentives, and thus full insurance is possible: $\underline{w} = \bar{w}$, where the wage level is set so as to meet the agent's participation constraint: $u(\underline{w}) = u(\bar{w}) = \hat{u} + E$. Thus, the optimal wage schedule is not monotonic, but instead goes down steeply when the output increases from $\underline{q}$ to $\hat{q}$, before increasing as steeply when the output increases from $\hat{q}$ to $\bar{q}$.

1.2.2 Full Inference

Suppose that the output level is given by $q = q(e)+\varepsilon$, where $q(\cdot)$ is one-to-one and where the noise ε is distributed over the support $[-\Delta, \Delta]$ according to the c.d.f $F(\cdot)$. The effort level chosen by the agent thus affects the support. Thus, whenever the agent deviates from a prescribed effort level, there is a positive probability that the deviation will be detected. For example, if the agent is asked to choose e^{FB}, and chooses instead $\hat{e} < e^{FB}$, then the support of the output moves from $\left[e^{FB} - \Delta, e^{FB} + \Delta\right]$ to $[\hat{e} - \Delta, \hat{e} + \Delta]$, and thus the deviation is detected whenever $q \in [\hat{e} - \Delta, e^{FB} - \Delta)$, or $\varepsilon \in \left[-\Delta, e^{FB} - \hat{e}\right]$, that is, it is detected with probability $F\left(e^{FB} - \hat{e}\right)$. It follows that the first-best is implementable at no additional cost if you can sufficiently punish the agent.

Mirrlees (1975, *RES* 1999) shows that the argument extends to some situations where deviations are never detected for sure. Suppose for example that $\varepsilon \sim F(\cdot)$ over $\mathbb{R}$, where $F(\varepsilon) \to 0$ as $\varepsilon \to -\infty$, and $F(\varepsilon)/f(\varepsilon) \to 0$ as $\varepsilon \to -\infty$. (This last condition simply says that the cdf $F(\cdot)$ goes to zero faster than the density $f(\cdot)$ does; it is satisfied for example by the normal distribution.)

In a first-best world, the principal's profit-maximization problem is

$$\begin{aligned} &\max_{e,w(q)} \mathbb{E}[q - w(q)|e] \\ &\text{s.t.} \quad \mathbb{E}[u(w(q)) - e|e] \geq \hat{u} \end{aligned}$$

Assuming that the agent's utility function is increasing and concave with the wage ($u' > 0 > u''$), it is then optimal for the principal to provide full insurance and to set the wage so as to meet the agent's participation constraint: $w(q) = w = h(e) \equiv u^{-1}(\hat{u} + e)$. The principal's objective function thus becomes (using $\mathbb{E}[q] = \mathbb{E}[e + \varepsilon] = e$)

$$\max_e e - h(e),$$

where $u" < 0 < u'$ implies $h', h'' > 0$. This objective function is thus concave, and the first-order condition yields the first-best level of effort e^{FB}: $h'(e^{FB}) = 1$.

Consider now the second-best setting, and the following schedule (it is more convenient to express it in terms of the agent's utility, $u(w)$, rather than in terms of the wage w):

$$u = \begin{cases} U & \text{if} \quad q \geq Q, \\ U - P & \text{if} \quad q < Q. \end{cases}$$

The parameters of the compensation scheme are U, P and Q: the agent is guaranteed a fixed utility U as long as he "delivers" an output at least equal to Q, and incurs a penalty P otherwise.

Failing to meet the threshold Q amounts to

$$q = e + \varepsilon \leq Q \Longleftrightarrow \varepsilon \leq Q - e,$$

and thus happens with probability $F(Q - e)$. The agent's expected utility is therefore given by

$$F(Q-e)(U-P) + (1 - F(Q-e))U - e = U - F(Q-e)P - e.$$

To induce the agent to choose the first-best level of effort, it then suffices to set P so as to satisfy the first-order condition

$$f(Q - e^{FB})P = 1 \Longleftrightarrow P = \frac{1}{f(Q - e^{FB})}. \tag{1.2}$$

The individual-rationality condition,

$$U - F(Q - e^{FB})P - e^{FB} \geq \hat{u},$$

leading the principal, for any given Q, to set U to:

$$U = \hat{u} + e^{FB} + \frac{F(Q - e^{FB})}{f(Q - e^{FB})}.$$

By assumption, $F(Q - e^{FB})/f(Q - e^{FB}) \to 0$ as $Q \to -\infty$. Therefore, asymptotically, reducing the threshold Q to $-\infty$ (and increasing the penalty P to ∞, so as to maintain incentives) allows the principal to implement the first-best outcome at no additional cost.

Limitations of this approach:

- We needed to make an assumption on the distribution.

- In practice, there is a *de facto* upper bound on the penalty we can impose on the agents, due to limited liability. That is, in real life, agents will not be able to pay a fine of, say, 500 million euros.

- The penalty is chosen relying on the fact that you know exactly the probability of very unlikely events (about the tail of the distribution).

1.2.3 Limited Inference

We just saw several examples in which the inference problem could be fully solved, in which case moral hazard has no bite: incentive constraints are costless. In the remainder of this chapter, we will focus on situations of limited inference, in which incentive constraints come at a cost.

Suppose for example that the output level q is distributed over the support $[0, Q]$, and that the agent can choose between two different levels of effort, 0 and $E > 0$, giving rise to density functions $f_0(q)$ and $f_E(q)$, respectively. Assume that the agent has an outside option that yields utility $\hat{u}$. We assume that the following condition holds:

Assumption (Monotone Likelihood Ratio Property – MLRP): The likelihood ratio

$$l(q) = \frac{f_E(q) - f_0(q)}{f_E(q)}$$

is increasing in q.

The principal wants to induce the high level of effort in the agent, i.e. $e = E$; she must therefore choose a wage schedule $w(q)$ that satisfies individual rationality and incentive compatibility. The principal's profit-maximization problem can thus be stated as:

$$\max_{w(\cdot)} \int_0^Q (q - w(q))\, f_E(q)\, dq$$

$$\text{s.t.} \quad \begin{array}{rr} \int_0^Q u\,(w(q))\, f_E(q)\, dq - E \geq \hat{u}, & (IR) \\ \int_0^Q u\,(w(q))\, f_E(q) - E \geq \int_0^Q u(w(q)) f_0(q)\, dq. & (IC) \end{array}$$

Denote the multipliers for these constraints by λ and μ respectively; the Lagrangian of this problem is

$$L = \int_0^Q [q - w(q) + \lambda u(w(q)) - \lambda(\hat{u} + E) + \mu u(w(q)) \frac{f_E(q) - f_0(q)}{f_E(q)} - \mu E] f_E(q)\, dq.$$

Note that $\lambda, \mu \geq 0$. The first-order condition with respect to $w(q)$, for a given q, yields

$$(\lambda + \mu l(q))\, u'(w(q)) = 1, \tag{1.3}$$

where $l(q)$ is the likelihood ratio defined above. Since $l(q)$ is increasing in q, the first term $\lambda + \mu l(q)$ is increasing in q. Under the assumption of concave utility, the second term $u'(w(q))$ is decreasing in $w(q)$. So as q increases, $w(q)$ increases.

Note: The reasoning relies on $\mu > 0$, which indeed holds at the optimum. To see this, note that if $\mu = 0$, then the first-order condition (1.3) reduces to $\lambda u'(w(q)) = 1$. This would imply full insurance ($w(q) = w$), but then a constant wage does not satisfy the incentive constraint (IC), a contradiction. Hence we must have $\mu > 0$.

1.2.4 Valuable Signals

Intuitively, have more informative signals facilitate inference and thus reduces the cost of providing incentives.

Example

Consider the following example:

- the agent is risk-neutral but has limited liability;
- the agent can choose among two levels of effort: $e \in \{0, E\}$;
- the output is either 0 or Q;
- the principal observes the output and another signal σ which can take two values, σ_0 and σ_E, with probabilities given by the following table:

Effort e	Output q	Signal σ
0	$q = \begin{cases} Q & \text{with probability} \quad p_0 \\ 0 & \text{with probability} \quad 1-p_0 \end{cases}$	$\sigma = \begin{cases} \sigma_E & \text{with probability} \quad \rho_0 \\ \sigma_0 & \text{with probability} \quad 1-\rho_0 \end{cases}$
E	$q = \begin{cases} Q & \text{with probability} \quad p_E \\ 0 & \text{with probability} \quad 1-p_E \end{cases}$	$\sigma = \begin{cases} \sigma_E & \text{with probability} \quad \rho_E \\ \sigma_0 & \text{with probability} \quad 1-\rho_E \end{cases}$

We will assume further that $\rho_E > \rho_0$, so that "right" signal σ_E is more likely when the agent exerts effort.

If the principal chooses to ignore the signal σ when designing the contract, then we are back to the example studied above, in which she offers the agent a base wage $w = \hat{w}$ and a bonus b designed to meet the incentive-compatibility condition:

$$b = \frac{E}{p_E - p_0},$$

which leaves a rent to the agent, equal to

$$r = p_E b - E = \frac{p_0 E}{p_E - p_0}.$$

Suppose now that the principal includes the signal in the contract design, and offers the agent a base wage w, with a bonus B only if she observes *both* the high-level output Q *and* the "right" signal σ_E. The incentive-compatibility condition is then

$$w + \rho_E p_E B - E \geq w + \rho_0 p_0 B,$$

which we can rearrange as

$$(\rho_E p_E - \rho_0 p_0)\, B \geq E.$$

The minimum bonus b satisfying incentive compatibility in this contract is thus

$$B = \frac{E}{\rho_E p_E - \rho_0 p_0},$$

and the minimum informational rent the principal would need to give is

$$\begin{aligned} R = \rho_E p_E B - E &= \frac{\rho_E p_E E}{\rho_E p_E - \rho_0 p_0} - E \\ &= \frac{\rho_0 p_0 E}{\rho_E p_E - \rho_0 p_0} \\ &= \frac{p_0 E}{\frac{\rho_E}{\rho_0} p_E - p_0} \\ &< r = \frac{p_0 E}{p_E - p_0}, \end{aligned}$$

where the last inequality stems from the assumption that $\rho_E > \rho_0$ (if $\rho_E < \rho_0$, it suffices to swap the roles of σ_E and σ_0; thus, what matters if the signal is "informative", in the sense that $\rho_E \neq \rho_0$). This rent is therefore strictly less than the rent paid in the contract when the principal ignored the signal; that is, the signal is indeed valuable to the principal.

Holmstrom (*Bell* 1979) developed the idea and showed that the principal should base the contract on a sufficient statistic of the signals available. That is, any *informative* signal should be included in the contract. But if one of the signals is perfectly colinear with a linear combination of the other signals, then it need not be included in the contract.

1.3 Comments

1.3.1 Simple Case

Assume limited liability and risk neutrality. There are two possible outcomes, 0 or Q. The agent can choose any effort $e \in \mathbb{R}_+$, in which case the probability of realizing the high output is $p(e)$, where $p(0) = 0$ and $p' > 0 > p''$, and the cost to the agent is $c(e)$, where $c(0) = 0$ and $c', c'' > 0$.

In a complete-information setting, the agent seeks to maximize

$$\max_e p(e)Q - c(e),$$

which is concave in e; the first-best level of effort e^{FB} thus solves the first-order condition:

$$p'(e)Q = c'(e).$$

In a second-best world, principal seeks to offer the bonus b that maximizes

$$\max_e p(e)b - c(e).$$

Define $e(b)$ to be the solution to this maximization problem. By a revealed preference argument, $b(e)$ is increasing: for any b and b', letting $e = e(b)$and $e' = e(b')$ denote the corresponding effort levels, we have: $pb - c \geq p'b - c'$ and $p'b' - c' \geq pb' - c$, which implies: $(p - p')(b - b') \geq 0$. The function $e(b)$ can therefore be inverted; let $b(e)$ denote the bonus needed to induce a given level of effort e. The agent's objective is strictly concave under our convexity conditions, and thus $b(e)$ is given by the first-order condition:

$$p'(e)b = c'(e) \Leftrightarrow b(e) = \frac{c'(e)}{p'(e)}.$$

This bonus satisfies $b(0) = 0$ (maximizing $p(e)b - c(e) = -c(e)$ indeed leads the agent to exert no effort, $e = 0$) and increases in e. The cost of inducing the effort level e is then

$$p(e)b(e) = c(e) + r(e),$$

where the rent

$$r(e) \equiv p(e)b(e) - c(e)$$

is such that $r(0) = 0$ and, using the envelope theorem:

$$r'(e) = p(e)b'(e) > 0.$$

This implies that, the rent $r(e)$ is positive for any $e > 0$; it also implies that the second-best effort level e^{SB}, maximizing the principal net payoff

$$\max_e p(e)Q - c(e) - r(e),$$

satisfies

$$p'(e)Q = c'(e) + r'(e) > c'(e)$$

and is thus strictly lower than the first-best

First-best may no longer be implementable

Suppose now that the agent can choose from three effort levels, $e \in \{0, \underline{E}, \bar{E}\}$, with associated costs $0, \underline{c}, \bar{c}$, respectively. The output is either 0 or Q, where the probabilities of realizing the high level of output Q under each effort level are given by $0, \underline{p}, \bar{p}$, respectively. The agent's utility is such that the agent's marginal utility is equal to $\alpha > 1$ as long as the utility does not exceed $\underline{c}$, and drops to 1 afterwards (see Figure 1.2):

$$u(x) = \begin{cases} \alpha x & \text{for} \quad x \leq \dfrac{\underline{c}}{\alpha}, \\ 1 - \dfrac{1}{\alpha} + x & \text{for} \quad x > \dfrac{\underline{c}}{\alpha}. \end{cases}$$

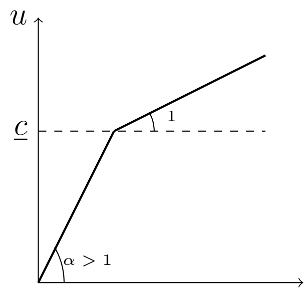

Figure 1.2:

Given this utility function, in a first-best world it would cost the principal an amount $\dfrac{\underline{c}}{\alpha}$ to compensate the agent for exerting $e = \underline{E}$, and $\dfrac{\underline{c}}{\alpha} + \bar{c} - \underline{c}$ to compensate exerting $e = \bar{E}$; we will assume that the first-best effort level is the middle one ($\underline{E}$):

$$\bar{p}Q - \left(\frac{\underline{c}}{\alpha} + \bar{c} - \underline{c}\right) \geq \underline{p}Q - \frac{\underline{c}}{\alpha} \geq 0,$$

or:

$$\frac{\bar{c} - \underline{c}}{\bar{p} - \underline{p}} \geq Q \geq \frac{\underline{c}}{\alpha \underline{p}}.$$

Now consider a second-best setting. The principal offers a base wage (utility) u, plus a bonus Δu if the high-level output Q is realized. Inducing the agent to exert $\underline{E}$ rather than 0 or $\bar{E}$ requires:

$$\underline{u} + \underline{p}\Delta u - \underline{c} \geq \underline{u}, \underline{u} + \bar{p}\Delta u - \bar{c},$$

or:

$$\frac{\bar{c} - \underline{c}}{\bar{p} - \underline{p}} \geq \Delta u \geq \frac{\underline{c}}{\underline{p}}. \tag{1.4}$$

Therefore, if:

$$\frac{\underline{c}}{\underline{p}} \geq \frac{\bar{c}-\underline{c}}{\bar{p}-\underline{p}} \geq Q \geq \frac{\underline{c}}{\alpha \underline{p}},$$

then:

- the first-best effort level is $\underline{E}$;
- and yet inducing this first-best effort level is infeasible in a second-best setting, since the this range characterized by (1.4) is empty (since the lower bound, $\frac{\bar{c}-\underline{c}}{\bar{p}-\underline{p}}$, exceeds the upper one, $\frac{\underline{c}}{\underline{p}}$).

The participation constraint may not be binding

This is clearly the case when the agent is subject to limited liability, as the simple example studied in the introduction shows: incentive-compatibility and limited liability may then be the only relevant constraints in that case, and require the principal to leave a rent to the agent, in addition to what would be needed to meet the agent's participation constraint.

When the agent is not subject to limited liability but is risk-averse, wealth effects may play a role, in such a way that participation may not be binding. That is, if the agent's utility is of the form $u(w, e)$, increasing the effort e may affect the agent's risk aversion, and increasing the wealth w may also lower the agent's disutility of effort. In addition, merely replacing the agent's incentive constraint by the associated first-order condition is not necessarily valid; as shown by Mirrlees (1975, *RES* 1999), the corresponding first-order condition of the principal's optimization problem may then be neither necessary nor sufficient. To avoid these issues, Grossman & Hart (*Econometrica* 1983) consider utility functions that are (additively and/or multiplicatively) separable in w and e, in which case they are able to develop an alternative approach that does not rely on the agent's first-order condition, and which consists in first characterizing the cost of inducing a particular effort level (implementation stage), before studying the optimal choice of effort (optimization stage).

Note: We can also interpret limited liability as an extreme form of risk-aversion; for example, a utility of the form $U = u(w) - c(e)$, where (see Figure 1.3):

$$u(w) = \begin{cases} w & \text{for} \quad w \geq \hat{w} \\ -\infty & \text{for} \quad w < \hat{w} \end{cases},$$

leads to an analysis that is formally identical to the case of a risk-neutral agent facing a limited liability $w \geq \hat{w}$.

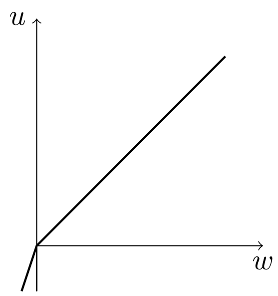

Figure 1.3:

Multi-tasking

Let there be two tasks to be assigned to either one or two agents, both with limited liability ($w \geq \hat{w}$). For each task:

- the agent in charge can choose from two levels of effort, $e \in \{0, E\}$;
- the output level is also binary, $q \in \{0, Q\}$, with independent probabilities of success;
- the probability of realizing the high output Q given by $p_0 > 0$ if the agent chooses $e = 0$ and by $p_E > p_0$ if the agent chooses $e = E$.

If an agent is assigned a single task, then from the analysis of 1.1.2 (see (1.1)), the principal must give the agent a rent

$$R_1 = \frac{p_0 E}{p_E - p_0}.$$

Thus, if the principal assigns the two tasks to different agents, the total information rent he will have to pay is $2R_1$.

The principal can however do better by giving both tasks to the same agent; in that case, instead of rewarding "in cash" the agent for the success of one task, the principal can grant the agent a stake (i.e., a fraction of the output produced) in the other task, which contributes to foster the agent's incentive to behave in the management of that other task. To see this, suppose that the principal gives a bonus B to the agent only if *both* projects succeed. To induce the agent to exert effort in both projects, the principal needs to ensure that the agent's expected payoff when the agent chooses $(e_1, e_2) = (E, E)$ is weakly greater than his expected payoff when (i) he chooses to work on one project but not the other ($(e_1, e_2) = (E, 0)$ or $(0, E)$); and (ii) he chooses not to work on either project ($(e_1, e_2) = (0, 0)$):

$$w + p_E^2 B - 2E \geq \begin{cases} w + p_E p_0 B - E & (ii) \\ w + p_0^2 B & (iii) \end{cases}$$

From condition (i) amounts to

$$p_E B \geq \frac{E}{p_E - p_0},$$

whereas condition (ii) is equivalent to

$$B \geq \frac{2E}{p_E^2 - p_0^2} = \frac{2}{p_E + p_0} \frac{E}{p_E - p_0}.$$

Since $p_0 < p_E \leq 1$, $\frac{2}{p_E + p_0} > 1$ and thus the last constraint is more demanding than the previous one, which we can therefore ignore. Ignoring also the participation constraint $w + p_E^2 B - 2E \geq \hat{w}$, the optimal contract subject to the limited liability constraint $w \geq \hat{w}$ and condition (iii) yields $w = \hat{w}$ and $B = 2E / \left(p_E^2 - p_0^2\right)$, which (more than) satisfies the agent's participation constraint: this contract gives the agent a rent

$$R_2 = p_E^2 B - 2E = \frac{p_E^2 2E}{p_E^2 - p_0^2} - 2E = \frac{p_0^2 2E}{p_E^2 - p_0^2} = \frac{p_0}{p_E + p_0} 2R_1 < 2R_1.$$

Thus, there is a benefit from assigning both tasks to the same agent: by doing so, the principal can reduce the information rent by offering the agent a bonus only if both projects are successful, rather than having to pay a bonus whenever a project is successful, as she would have to do in the case of two agents.

Commitment

Bilateral renegotiation. We begin by noting the discrepancy between what is desirable *ex ante* and what is desirable *ex post*. Consider the following timeline:

Initial contract	$\rightarrow$	Agent decides on effort	$\rightarrow$	Final outcome realized

Ex ante, there is a trade-off between the provision of incentives, according to which the agent should be made the residual claimant and therefore bear all the risks, and efficient risk-sharing. But ex post, once the agent has made his decision, incentives are no longer an issue and the parties would therefore benefit from sharing the risk in an efficient way. Thus, once the effort decision is made, the principal and agent have an incentive to renegotiate the original contract in order to optimize risk-sharing. But then, anticipating that he will be fully insured, the agent will no longer has an incentive to exert any effort. That is, while ex post the parties can benefit from renegotiating the original contract, anticipating this renegotiation will backfire: the agent will not base his effort decision on the initial contract, but rather on the anticipated final contract.

In the same vein, in a multi-period setting the principal would ideally wish to contract on the agent's consumption plan. However, in practice, the parties contract on the agent's compensation, and the agent remains free to borrow or save. Thus, the contract should not only provide incentives with respect to the choice of effort, but also with respect to savings; in particular, Rogerson (*Econometrica* 1985) shows that, the contract that would be optimal if the parties could directly contract on the agent's savings is such that, if the agent would like to save more if he could do so secretly. However, in later periods, there is no reason anymore to account for the agent's incentives about past savings decisions, and thus the parties would benefit from replacing the (continuation of) the original contract with another one that is more efficient for the remaining periods.

One way to alleviate this problem is to make it unclear whether the agent has actually chosen the effort before the renegotiation (that is, induce the agent to randomize). See e.g., Fudenberg & Tirole (*Econometrica* 1990) and Ma (*QJE* 1991) for within-period contracting, and Chiappori *et al.* (*EER* 1994) for multi-period contracting.

Unilateral renegotiation. The above analysis also relies on the assumption that the principal can commit to pay performance-based bonuses; however, ex post, the principal may be tempted to claim that performance was poor, even if it was actually good, in order to avoid paying the bonus. Commitment may thus not be credible if the agent's performance is not readily observed by third parties such as judges or courts. One way consists in introducing tournaments and prize systems; the principal can then commit herself to give a prize to *someone*, and has no reason not to give it the best performer.

1.4 Applications

1.4.1 Partial Insurance

Consider the example of car insurance. The agent is a driver, with wealth w and utility $u(w)$. With probability Π_0 there is no accident; with probability Π_i there is an accident generating a loss L_i, where $i = 1, \ldots, n$. The principal is an insurance company, offering a policy that involves a fee β and reimburses the driver an amount α_i (on top the fee β) in the case of an accident with loss L_i.

The insurance company's expected profit is

$$V = \Pi_0 \beta - \sum_{i>0} \Pi_i \left(c + (1+\gamma)\alpha_i\right),$$

where c denotes a fixed transaction cost per accident, and γ denotes an additional transaction cost that varies proportionally with the amount at stake, α_i. The insurance company seeks to maximize its expected profit, subject to the driver's participation constraint:

$$U = \Pi_0 u(w-\beta) + \sum_{i=1}^{n} \Pi_i u(w - L_i + \alpha_i) \geq \hat{U} = \Pi_0 u(w) + \sum_{i=1}^{n} \Pi_i u(w - L_i).$$

Let $\lambda \geq 0$ denote the Lagrangian multiplier of the participation constraint. In a first-best world, this constraint is necessarily binding (since the principal wishes for example to maximize the fee β) and thus $\lambda > 0$; the corresponding first-order conditions are

$$\begin{aligned} \text{w.r.t. } \beta: \quad & \Pi_0 - \lambda \Pi_0 u'(w-\beta) = 0, \\ \text{w.r.t. } \alpha_i: \quad & -\Pi_i(1+\gamma) + \lambda \Pi_i u'(w - L_i + \alpha_i) = 0, \end{aligned}$$

which we can simplify to

$$\begin{aligned} \text{w.r.t. } \beta: \quad & w - \beta = u'^{-1}\left(\frac{1}{\lambda}\right), \\ \text{w.r.t. } \alpha_i: \quad & w - L_i + \alpha_i = u'^{-1}\left(\frac{1+\gamma}{\lambda}\right). \end{aligned}$$

It follows that, for $i = 1, \ldots, n$, the driver's net wealth $w - L_i + \alpha_i$ remains constant and equal to $u'^{-1}\left((1+\gamma)/\lambda\right)$. That is, conditional on having an accident, the driver ends up with the same net wealth: the driver is thus fully insured against the *gravity* of the accident. Furthermore, if $\gamma = 0$, then $w - L_i + \alpha_i = w - \beta \left(= u'^{-1}(1/\lambda)\right)$: we have full insurance. If instead

$\gamma > 0$, the net wealth is lower in case of an accident: the driver then pays a set amount out-of-pocket, while the insurance company covers the rest. We therefore have a franchise contract: $w - L_i + \alpha_i = w_- < w_+ = w - \beta$.

Let us introduce moral hazard: the driver can now choose to exert an effort to be cautious; denote this level by e, which can take one of two values 0 or E. If the driver exerts effort ($e = E$: driving carefully, locking the doors of the car when parking, ...), exerts the low-effort level 0, the probabilities of accidents are given by Π_0, Π_i; otherwise, they are given by π_0, π_i. The insurance company's maximization problem in this second-best setting is

$$\begin{aligned}
&\max \Pi_0\beta - \sum_{i>0} \Pi_i \left(c_0 + (1+\gamma)\alpha_i\right) \\
&\text{s.t.} \quad \Pi_0 u(w-\beta) + \sum_{i>0} \Pi_i u(w - L_i + \alpha_i) - E_i \geq \hat{U} \\
&\qquad \Pi_0 u(w-\beta) + \sum_{i>0} \Pi_i u(w - L_i + \alpha_i) - E_i \\
&\qquad\qquad \geq \pi_0 u(w-\beta) + \sum_{i>0} \pi_i u(w - L_i + \alpha_i) - 0.
\end{aligned}$$

Denote the multipliers of the two constraints by λ and μ respectively. The first-order conditions are

$$\begin{aligned}
&\text{w.r.t. } \beta : & \Pi_0 &= (\lambda + \mu l_0)\, \Pi_0 u'(w-\beta), \\
&\text{w.r.t. } \alpha_i : & (1+\gamma)\Pi_i &= (\lambda + \mu l_i)\, \Pi_i u'(w - L_i + \alpha_i),
\end{aligned}$$

where, for $h = 0, 1, ..., n$,

$$l_h = \frac{\Pi_h - \pi_h}{\Pi_h}$$

denotes the likelihood ratio. If l_i is independent of i, then we have the same pattern as before: the agent's marginal utility is constant in case of accident. In other words, when exerting effort only affects the probability of having an accident, but not its seriousness, than it remains optimal to opt for a fixed franchise policy. Otherwise (i.e., if the likelihood ratio l_i varies across accidents), the franchise should depend not only on whether an accident occurred, but also on its gravity.

Remark: Should we expect any correlation between the reimbursement paid by the insurance agency and the gravity of the accident? Not necessarily; it depends on which type of accident (the very serious ones or the less serious ones) is mostly affected by the agent's effort.

1.4.2 Efficiency Wage

Consider the following interaction between a firm (the principal) and a worker (the agent). The agent can exert effort $e = E$; this costs the agent E and results in output Q with probability 1. Alternatively, the agent can exert effort $e = 0$, which costs the agent nothing but yields the output Q only with probability $p < 1$, and zero output otherwise. The principal offers a contract with a wage w if the observed output is Q. There is limited liability ("no-slavery") in that the agent is always free to walk away from the job and start contracting with a new principal. If the worker quits, he obtains utility $\hat{U}$ (which is endogenous; more on this below).

The limited liability condition implies that the principal must offer a contract satisfying $w \geq \hat{U}$. The principal must moreover satisfy the incentive-compatibility condition, that is, the worker must prefer exerting effort, in which case he gets the wage w with probability 1 but incurs the cost E, to shirking, in which case:

- With probability p, the project is successful; the agent then gets the wage w.
- With probability $1 - p$, the project is unsuccessful; the worst the principal can do in that case is to fire the agent, who then obtains $\hat{U}$ on the labour market.

The incentive-compatibility condition can thus be expressed as

$$w - E \geq pw + (1 - p)\hat{U},$$

which can be rearranged to yield

$$(1 - p)(w - \hat{U}) \geq E.$$

Under "full employment", the agent can walk away and starts immediately a new relationship with another firm; therefore: $\hat{U} = w - E$. But the, the incentive-compatibility condition becomes:

$$\begin{aligned} w - E \geq pw + (1 - p)\hat{U} &= pw + (1 - p)(w - E) \\ \iff p(w - E) &\geq pw, \end{aligned}$$

a contradiction. In other words, full employment is incompatible with the provision of effort incentives.

If instead there is some unemployment, then an agent who walks away does not find another job immediately; the value of this outside option can then be expressed as

$$\hat{U} = (1 - \delta(u))(w - E),$$

where u denotes the unemployment rate, and the discount rate δ increases with unemployment. Each firm, maximizing its expected profit subject to the above incentive constraint, will make this constraint binding; in equilibrium, we thus have:

$$w - E = pw + (1-p)\hat{U} = pw + (1-p)(1-\delta(u))(w-E),$$

or

$$w = w^*(u) \equiv E + \frac{pE}{(1-p)\delta(u)}.$$

Notice that the equilibrium wage w^* decreases as the rate of unemployment u increases. For further analysis, see Shapiro & Stiglitz (*AER* 1984).

1.4.3 Credit Markets

Consider a firm with an initial asset A which has a project costing $I > A$; to finance the project, the firm must therefore seek external investors. The firm manager (the agent) can either "behave", in which case the project produces output Q with probability 1; or he can "shirk", in which case the project produces no output Q but the manager then enjoys a private benefit B. It is efficient to finance the project when the manager behave if

$$Q > I.$$

The project is efficient, and Let R denote the amount the firm must reimburse the lender. The incentive-compatibility constraint in this problem is

$$Q - R \geq B,$$

or:

$$R \leq \hat{R} \equiv Q - B,$$

where $\hat{R}$ represents the *pledgeable income*, that is, the maximal amount that the firm can credibly repay – any higher amount $R > \hat{R}$ would violate the incentive-compatibility, implying that the firm manager will choose to shirk – and thus no repayment would ever be made.

Note that this pledgeable income is negative when $B > Q$, that is, when shirking would actually be efficient; in that case, while it would still be efficient to undertake the project if $B > I$, it cannot be financed. Even if $Q > B$, so that inducing the manager to behave is more efficient, the most a firm can raise via external financing is

$$I - A \leq \hat{R},$$

or

$$A \geq \hat{A} \equiv I - \hat{R} = I - (Q - B),$$

where $\hat{A} > 0$ as long as $Q < B + I$. In that case, it is only the firms with initial assets at least as large as $\hat{A}$ who get financed – in other words, "the rich get richer".

1.4.4 Group Lending

Consider a variant of the above model in which “shirking” still produces the output Q with some probability $p < 1$ (and no output otherwise). Going through the same steps as above, the pledgeable incomes becomes

$$\hat{R}_1 \equiv Q - \frac{B}{1-p},$$

and the associated threshold level for the initial assets becomes

$$\hat{A}_1 \equiv I - \hat{R}_1 = I - (Q - \frac{B}{1-p}).$$

In what follows, we will suppose that this threshold is positive $\left(\hat{A}_1 > 0\right)$, so that an entrepreneur with initial assets $A < \hat{A}_1$ cannot finance his project. Suppose now that:

- there are n such entrepreneurs, each with a project similar as the first one (and independent realizations, in case of shirking);
- each entrepreneur’s initial asset is lower than $\hat{A}_1$, which prevents the entrepreneur from seeking investors on an individual basis.

We now show that, by grouping their projects, the entrepreneurs may be able to secure financing, provided they can coordinate their effort decisions.

To see this, suppose first that their regroup their projects, so that the reimbursement for one project may now depend on the outcomes of all the projects, but keep choosing their efforts independently from each other.

An entrepreneur’s individual incentive-compatibility constraint becomes

$$Q - \mathbb{E}[R] \geq p(Q - \mathbb{E}[R]) + B,$$

where $\mathbb{E}[R]$ denotes the expected reimbursement, given the distribution of the outcomes of the other projects. But then, $\mathbb{E}[R]$ cannot exceed $\hat{R}_1$, which in turn implies that the entrepreneur would need at least $\hat{A}_1$ to secure financing.

Suppose now that the n entrepreneurs take their effort decisions jointly, so as to maximize their joint payoff. Building on the insights from section 1.3.1 on multi-tasking, the best contract then consist in leaving the entrepreneurs a return only if all projects are successful – that is, successful projects should pay back the loan to unsuccessful ones. The contract must induce the entrepreneurs to behave rather than to shirk on any $k \leq n$ projects; therefore, the following incentive-compatibility conditions must be satisfied for $k = 1, ..., n$:

$$n\left(Q-R\right) \geq np^k(Q-R)+kB,$$

or

$$Q-R \geq \frac{kB}{n\left(1-p^k\right)}.$$

It can be checked that the right-hand side of RHS this condition increases in k:

$$\frac{d}{dk}\left(\frac{kB}{n\left(1-p^k\right)}\right) = \frac{B\left(1-p^k+k\left(\log p\right)p^k\right)}{n\left(1-p^k\right)^2},$$

and thus has the same sign as $f\left(x\right) = 1 - x + x\log x$, where $x = p^k \in [0,1]$. Since $f'\left(x\right) = \log x < 0$ on this range, $f\left(x\right) \geq f\left(1\right) = 0$.

Therefore, the most demanding constraint is the one for $k = n$, which amounts to

$$R \leq \hat{R}_k \equiv Q - \frac{B}{1-p^n},$$

which in turn leads to

$$A \geq \hat{A}_k \equiv I - (Q - \frac{B}{1-p^n}),$$

where the threshold level $\hat{A}_k$ decreases as n increases. It is thus easier to finance the projects by regrouping them.

1.4.5 Moral Hazard in Teams

Consider a team of two agents, where the output Q produced by the team is equal to the sum of the agents' contributions: $Q = e_1 + e_2$. Assume that the cost of effort to Agent i is given by $c_i(e_i)$, where $c_i', c_i'' \geq 0$ for $i = 1, 2$.

In a first-best world, the agents solve

$$\max_{e_1,e_2} e_1 + e_2 - \left(c_1(e_1) + c_2(e_2)\right).$$

This program is concave, and the solution e_i^{FB} thus satisfies $c_i'(e_i^{FB}) = 1$ for all i.

Let us now turn to a second-best world, in which the agents cannot contract directly on their effort levels, but can only share the output produced by the team: each agent then solves

$$\max_{e_i} s_i(Q) - c_i(e_i),$$

where by constructions the agents' shares, $s_1(Q)$ and $s_2(Q)$, must satisfy, for any Q:

$$s_1(Q) + s_2(Q) = Q.$$

Assuming interior solutions, each agent $i = 1, 2$ will choose an effort e_i^{SB} characterized by the first-order condition

$$s_i'(Q) = c_i'(e_i^{SB}).$$

It follows that the first-best cannot be achieved: this would require $s_1'(Q^{FB}) = s_2'(Q^{FB}) = 1$, but by construction $s_1'(Q) + s_2'(Q) = 1$ for any Q, a contradiction.

A way out consists in introducing a "budget breaker", so as to allow $s_1(Q) + s_2(Q) \neq Q$ (at least for output levels other than the equilibrium one). For example:

- A smooth sharing rule such as $s_i(Q) = Q - \alpha_j Q^{FB}$, where $\alpha_1 + \alpha_2 = 1$, leads to $e_i = e_i^{FB}$ and gives each agent i a share $\alpha_i Q^{FB}$.

- Alternatively, the agents could adopt a discontinuous scheme such as

$$s_i(Q) = \begin{cases} \alpha_i Q^{FB} & \text{if } Q \geq Q^{FB}, \\ 0 & \text{otherwise.} \end{cases}$$

1.4.6 Career Concerns

Suppose that there are two periods, and the output produced in period t depend on an implicit productivity parameter θ, which is constant over time and *unknown to everyone*, as well as on the effort e_t provided by the agent in that period: $q_t = \theta + e_t$, where exerting effort e_t costs the agent $c(e_t)$ in that period.

Suppose further that, in each period, the wage obtained by the agent is equal to the expected output produced in that period.

In period 2, the agent is not incentivized to exert effort; therefore,

$$w_2(q_1) = \mathbb{E}[q_2|q_1] = \mathbb{E}[\theta|q_1] = q_1 - e_1^*,$$

where e_1^* denotes the effort expected from the agent in period 1. It follows that, in period 1, the agent will seek to maxzimize

$$\max_{e_1} \mathbb{E}[w_1 - c(e_1) + \delta w_2(q_1) = w_1 - c(e_1) + \delta(\theta + e_1 - e_1^*),$$

where δ denotes the discount factor attached to period 2. This leads to $e_1 = e_1^*$, characterized by:

$$c'(e_1^*) = \delta.$$

It follows that the agent will exert too much effort when $\delta > 1$ (that is, at the beginning of the career, interpreting "period 2" as representing all the future periods of activity), and will instead exert too little effort when $\delta < 1$ (towards the end of the career).

More generally, learning about the agent's productivity is more progressive (the output also depends on a noise ε_t, say), and this productivity may not be constant over time. See e.g. Dewatripont, Jewitt and Tirole (*RES* 1999) for more discussion.

Chapter 2

References

Akerlof, G. (1970), "The market for Lemons: Quality uncertainty and the market mechanism," *Quarterly Journal of Economics*, 89:488-500.

Attar, A., Th. Mariotti and F. Salanié (2011), "Nonexclusive Competition in the Market for Lemons," *Econometrica*, 79(6):1869-1918.

Baron, D., and R. Myerson (1982), "Regulating a Monopolist with Unknown Costs", *Econometrica*, 50:911-930.

Bester, H., and R. Strausz (2001), "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case," *Econometrica*, 69(4):1077-98.

Bolton, P. and M. Dewatripont (2005), *Contract Theory*, MIT Press.

Chiappori, P.-A., I. Macho, P. Rey and B. Salanié (1994), "Repeated Moral Hazard: The Role of Memory, Commitment, and the Access to Credit Markets," *European Economic Review*, 38(8):1227-1253.

Clarke, E. H. (1971), "Multipart pricing of public goods," *Public Choice*, 17-33.

Dasgupta, P., P. Hammond and E. Maskin (1979), "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility", *Review of Economic Studies*, 46:185-216.

D'Aspremont, C., and L.A. Gérard-Varet (1979), "Incentives and Incomplete Information", *Journal of Public Economics*, 11:24-45.

Dewatripont, M., I. Jewitt and J. Tirole (1999), "The Economics of Career Concerns, I: Comparing Information Structure," *The Review of Economic Studies*, 66(1):183-198.

Fudenberg, D., and J. Tirole (1990), "Moral Hazard and Renegotiation in Agency Contracts," *Econometrica*, 58(6):1279-1319.

Fudenberg, D., and J. Tirole (1991), *Game Theory*, MIT Press, Cambridge.

Gibbard, A. (1973), "Manipulation of Voting Schemes: A General Result", *Econometrica*, 41:587-601.

Gibbons, R. (1992), *A Primer in Game Theory*, Harvester Wheatsheaf, N.Y.

Green, J., and J.J. Laffont (1979), *Incentives in Public Decision Making*, North Holland: Amsterdam.

Grossman, S., and O. Hart (1983), "An Analysis of the Principal-Agent Problem", *Econometrica*, 51:7-45.

Groves, T. (1973), "Incentives in Teams", *Econometrica*, 41:617-631.

Guesnerie, G., and J.-J. Laffont (1984), "A Complete Solution of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm", *Journal of Public Economics*, 25:329-369.

Holmstrom, B. (1979), "Moral Hazard and Observability", *Bell Journal of Economics*, 10:74-91.

Hirshleifer, J. (1971), "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 61:561-574.

Laffont, J.J. (1988), *Fundamentals of Public Economics*, MIT Press, Cambridge.

Laffont, J.J. (1986), *The Economics of Uncertainty and Information*, MIT Press, Cambrigde. Version française: *Economie de l'Incertain et de l'Information*, Economica, Paris.

Laffont, J.J. and D. Martimort (2002), *The Theory of Incentives: The Principal Agent Model*, Princeton University Press.

Laffont, J.J., and E. Maskin (1979), "A Differentiable Approach to Expected Utility Maximizing Mechanisms", Chap. 16 in Laffont, J.J. ed., *Aggregation and Revelation of Preferences*, North-Holland: Amsterdam.

Laffont, J.J., and E. Maskin (1980), "A Differential Approach to Dominant Strategy Mechanisms", *Econometrica*, 48:1507-1520.

Laffont, J.J., and E. Maskin (1982), "The Theory of Incentives: An Overview", in W. Hildenbrand (ed.), *Advances in Economic Theory*, Cambridge University Press.

Laffont, J.J., and J. Tirole (1986), "Using Cost Observation to Regulate Firms", *Journal of Political Economy*, 94(3):614-641

Leininger, W., P. B. Linhart and R. Radner (1989), "Equilibria of the sealed-bid mechanism for bargaining with incomplete information", *Journal of Economic Theory*, 48(1):63-106.

Ma, A. (1991), "Adverse Selection in Dynamic Moral Hazard," *Quarterly Journal of Economics*, 106:255-275.

Mas-Colell, A., M.D. Whinston and J. Green (1995), *Microeconomic Theory*, Oxford University Press, New York and Oxford.

Maskin, E. (1985), "The Theory of Implementation in Nash Equilibrium: A Survey", *Social Goals and Social Organization, Essays in memory of Elisha Pazner*, L. Hurwicz, D. Schmeidler and H. Sonnenschein eds., Cambridge University Press.

Maskin, E. (1999), "Nash Equilibrium and Welfare Optimality", *Review of Economic Studies*, 66:23-38.

Mirrlees, J. (1975), "The Theory of Moral Hazard and Unobservable Behaviour: Part I", reprinted in *Review of Economic Studies,* 66(1):3–21, January 1999.

Miyazaki, H. (1977), "The rat race and internal labor markets," *Bell Journal of Economics*, 8:394-418.

Moore, J. (1992), "Implementation, Contracts, and Renegotiation in Environments with Complete Information", in *Advances in Economic Theory*, Vol. I, Laffont ed., Cambridge University Press.

Myerson, R. B. (1979), "Incentive Compatibility and the Bargaining Problem", *Econometrica*, 47: 61-74.

Myerson, R. B., and M. Satterthwaite (1983), "Efficient Mechanisms for Bilateral Trading", *Journal of Economic Theory*, 28: 265-281.

Palfrey, T. (1992), "Implementation in Bayesian Equilibria: The Multiple Equilibrium Problem in Mechanism Design", in *Advances in Economic Theory*, Vol. I, Laffont ed., Cambridge University Press.

Picard, P. (2009), "Participating insurance contracts and the Rothschild-Stiglitz equilibrium puzzle," Working Papers hal-00413825, HAL.

Rogerson, W. (1985), "Repeated Moral Hazard," *Econometrica*, 53:69-76.

Rothschild, M., and J. E. Stiglitz (1976), "Equilibrium in Competitive Insurance Markets," *Quarterly Journal of Economics*, 90:629-649.

Salanié, B. (1994), *The Economics of Contracts: A Primer*, MIT Press, 1997.

Satterthwaite, M. (1975), "Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions", *Journal of Economic Theory*, 10:187-217.

Satterthwaite, M., and S. R. Williams (1989), "Bilateral trade with the sealed bid k-double auction: Existence and efficiency", *Journal of Economic Theory*,48(1):107-133

Shapiro, C., and J. E. Stiglitz (1984), "Equilibrium Unemployment as a Worker Discipline Device," *The American Economic Review*, 74(3):433-444.

Spence, M. (1973), "Job Market Signalling," *Quarterly Journal of Economics*, 87:355-374.

Vickrey, W. (1961), "Counterspeculation, auctions and competitive sealed tenders," *Journal of Finance*, 8-37.

Wilson, C. (1977), "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory*, 16:167-207.

Wilson, C. (1980), "The Nature of Equilibrium in Markets with Adverse Selection," *Bell Journal of Economics*, 11:108-130.