Εθνικό και Καποδιστριακό
ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ

Microeconomics II

Spring term 2024

Asymmetric Information

Adverse Selection
Moral Hazard

Theodoros Rapanos

Camilla: You, sir, should unmask.

Stranger: Indeed?

Cassilda: Indeed it's time. We all have laid aside disguise but you.

Stranger: I wear no mask.

Camilla (terrified, aside to Cassilda): No mask? No mask!

–R.W. Chambers, The King in Yellow, Act I, Scene 2

Outline

- Introduction to asymmetric information
- Adverse selection:
  - Main concepts
  - Akerlof's model of the *market for lemons*
  - Spence's model of *job market signalling*
- Moral hazard:
  - Main concepts
  - A simple principal–agent model

Reading material

Recommended textbooks:

- Chapter 10 from **Muñoz-Garcia, F. (2017).** *Advanced Economic Theory: An intuitive approach with examples*, MIT Press: Cambridge, MA.
- Chapters 3 and 4 from **Laffont J.-J. and D. Martimort (2002).** *The theory of incentives: The principal–agent model*, Princeton University Press: Princeton, NJ.

References:

- **Akerlof, G.A. (1970)**. "The market for "lemons": quality uncertainty and the market mechanism", *The Quarterly Journal of Economics,* 84(3): 488–500.
- **Spence, M. (1973)**. "Job market signalling", *The Quarterly Journal of Economics,* 87(3): 355–374.
- **Holmström, B. (1979)**. "Moral hazard and observability", *The Bell Journal of Economics,* 10(1): 74–91.

Information in Competitive Markets

- In perfectly competitive markets, we assume that all agents possess complete information about traded commodities, and all other aspects of the market.
- What about markets for medical services, or insurance, or second-hand items (such as used cars, electric appliances, or tech gadgets)?

Asymmetric Information in Markets

- A doctor (producer) knows more about medical services, and the true needs of his or her patients (consumers), than the patients do themselves.
- An insurance buyer knows more about his or her riskiness than the insurance company (seller).
- The owner of a house knows more about its condition than potential buyers do.
- A prospective borrower has a better knowledge of her financial situation and her risks than the bank she applies to for a loan.
- It is often hard for a car owner to observe the amount of work required to fix his or her car, as well as the quality of the spare parts used by the mechanic.

Asymmetric Information in Markets

- Markets in which at least one of the agents involved is not completely informed about aspects that may affect his or her behavior are markets with **incomplete information**.
- Markets with some agents better informed than others are **markets with asymmetric information**.
- Asymmetric information can cause markets to stray away from efficiency, and lead to suboptimal outcomes.
- The First Theorem of Welfare Economics does not apply in markets with asymmetric information.
- In this class we will discuss in more detail how asymmetric information can affect the functioning and the efficiency of a market.

Adverse Selection and Moral Hazard

- In these two lectures we focus on two important types of asymmetric information: **adverse selection** and **moral hazard**.
- Adverse selection entails *hidden information*: It refers to situations where one side of the market can't observe the quality or "type" of the other side.
- Moral hazard entails *hidden action*: It refers to situations where one side of the market can't observe the actions of the other.

Adverse Selection

- Issues that could potentially arise due to incompleteness of information had been outwardly noted in some papers in the 1950s and 1960s. Yet, before the 1970s there existed no dedicated study on this topic, and its significance had been largely underestimated, even ignored.
- Pathbreaking paper by George Akerlof in 1970. Nobel prize in 2001 (together with M. Spence and J.E. Stglitz).
- Realizing the impact that asymmetric information has on economic outcomes improved greatly our understanding of how markets function, and radically changed our way of thinking in Economics.

Adverse Selection

- Economic models with adverse selection can be used in the study of a wide range of applications such as:
  - principal–agent problem (agent's type unknown to principal; e.g., employer–worker, voters–government, government–contractors, etc.)
  - insurance market analysis (type of insured agents unknown)
  - markets where certain characteristics of the goods or services traded are unknown to the prospective buyers
  - pricing of goods and services when consumers' income and preferences are unknown to the firms
  - regulation of industries with incomplete knowledge of the demand and supply functions (unknown consumer preferences, market size, costs, emissions, etc).

Adverse Selection

Akerlof's model of the *market for lemons* (1970, QJE) :

- Consider a market for used cars of varying quality, denoted with q.
- A car of quality q is valued as such by the buyer, but has only value $d_s q$ for the seller, with $d_s \in (0,1)$.
- Following a car trade at price p, the buyer's utility function is given by

$$u(p, q) = q - p,$$

while the seller's profit function is given by

$$\pi(p, q; d_s) = p - d_s q.$$

Agents' reservation utility (profit) is normalized to $\bar{u} = 0$ ($\bar{\pi} = 0$).

(How can we justify a price showing up explicitly in a utility function?)

Adverse Selection

- Since $d_s < 1$, buyers assign a higher value to a given car than sellers/owners do. (Otherwise, there would exist no scope for trade, and thus no market.)
- Value $v_s(q) = d_s q$ is known as the *reservation price* of the seller for a used car of quality q. It is the lowest price he would accept to sell his car at.
- Therefore, when a car of quality q is traded at some price $p \in (d_s q, q)$, both parties strictly benefit: The buyer enjoys a (positive) surplus, and the seller makes a (positive) profit.

Symmetric Information

Benchmark: Quality observable by the buyers

- If a buyer can perfectly observe car quality q, she accepts to buy the car at price p if and only if

$$q - p \geq 0 \text{ or } p \leq q,$$

that is, if the net utility she gains from such a trade is positive.

- A seller of a car of quality q anticipates such an acceptance rule by the buyer and sets a price p that solves

$$\max_{p} \; p - d_s q$$

$$\text{s.t.} \quad p \leq q$$

where the last inequality is referred to as the buyer's *participation constraint (PC)* or the *individual rationality condition*.

Symmetric Information

- To keep our analysis simple, we shall assume that all bargaining power lies with the seller. This could be for example because there is a large number of buyers, and hence they have no bargaining power.
- The sellers can therefore set any price, and appropriate all trade gains, provided of course that buyers agree to participate in the trade.
- Observe that given car quality q, the seller’s objective function is strictly increasing in p, and hence there exists a corner solution with
$$p^* = q.$$
- Hence the participation constraint will bind (that is, hold with equality).
- Notice that this market outcome is efficient, in the sense that social welfare (which coincides with sellers’ welfare) is maximized.
- *How would our analysis change if buyers had all the market power, or if power was split between the two parties?*

Asymmetric Information

Adverse selection: *Buyers cannot observe car quality*.

- What if buyers are unable to observe a car's true quality, q?
- Assuming that buyers are expected utility maximizers, they will buy a car only if a seller's asking price p satisfies

$$p \leq \mathsf{E}(q)$$

where $\mathrm{E}(q)$ is her expectation of the quality of the car.

- A rational seller anticipates such an acceptance rule by the buyer, and sets a price p^{AI} that solves

$$\max_{p} \; p - d_s q$$
$$\text{s.t. } \; p \leq \mathsf{E}[q]$$

where $p \leq \mathsf{E}[q]$ is the buyer's PC under asymmetric information.

Asymmetric Information

- Since all market power lies with the seller, the participation constraint must again bind at the optimum, so that $p = \mathrm{E}(q)$.

- A seller however will set his asking price at

$$p^{AI} = \mathrm{E}(q)$$

if and only if

$$p^{AI} \geq d_s q.$$

- This will be true for sellers owning cars with quality *q* such that

$$q \leq \frac{\mathrm{E}(q)}{d_s}. \qquad \text{(A.1)}$$

- The rest of the sellers will stay out of the market. *(Why?)*

Asymmetric Information

- Assume that prospective buyers' prior beliefs about car quality q are captured by a distribution with pdf f and support $S \subseteq \left[\underline{q}, \overline{q}\right]$.

<u>Example</u>: *standard uniform distribution,* $q \sim U(0,1)$

- In this case it holds $S = [0,1]$, $f(q) = 1$ for $q \in S$, $f(q) = 0$ for $q \notin S$, and $\mathrm{E}(q) = {}^{1}/_{2}$.

- This implies that the maximum price that a buyer would pay for a car is $p^{max} = \mathrm{E}(q) = {}^{1}/_{2}$.

- Assume moreover that $d_s = {}^{3}/_{4}$. Then according to expression (A.1), participating in the market would be profitable only for those sellers with car quality

$$q \leq \frac{2}{3}.$$

Asymmetric Information

- As a result, cars with relatively high quality (namely, $q > {}^{2}/_{3}$) are *crowded out*: The maximum price that an uninformed buyer would pay, p^{max}, is not high enough to compensate the sellers of such cars for their foregone value.

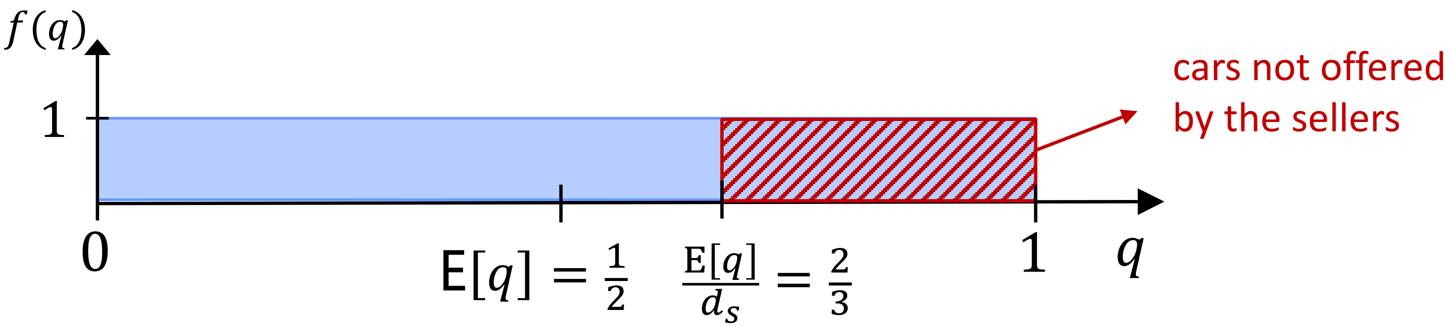

- Buyers' inability to observe q causes the market for good cars ($q > {}^{2}/_{3}$, "peaches") to collapse. Only bad cars ($q \leq {}^{2}/_{3}$, "lemons") are offered in the market.

- This is a *market failure*: Both buyers and sellers of peaches would be better off if peaches were traded in the market, and social welfare would be therefore higher.

Asymmetric Information

- Interestingly though, the story does not end here.
- If buyers are fully rational, they should be able to infer that no cars of quality higher than ${}^{2}/_{3}$ are offered in that market.
- Their beliefs will no longer be represented by the unconditional distribution $f(q)$, but rather by a conditional distribution with pdf $f(q|q \leq {}^{2}/_{3})$.
- The expected quality of those cars will therefore be

$$\mathrm{E}\left[q\middle|q \leq \frac{2}{3}\right] = \frac{\frac{2}{3} - 0}{2} = \frac{1}{3}.$$

Asymmetric Information

- Hence buyers will only buy cars whose price satisfies $p \leq {}^{1}/_{3}$.
- As a result, it will not be profitable for sellers with cars of quality

$$q > \frac{{}^{1}/_{3}}{{}^{3}/_{4}} = \frac{4}{9}$$

to stay in the market.

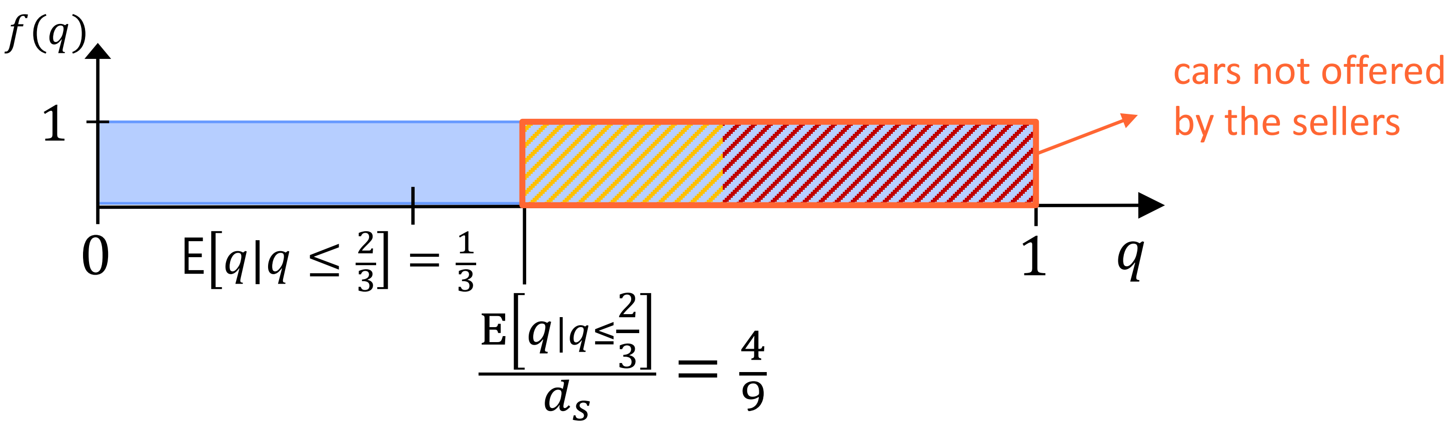

Asymmetric Information

- Anticipating this, buyers will update their beliefs about the quality of the cars offered further downwards.
- Repeating the same argument over again, we find that the market "unravels", and even low-quality cars will be progressively crowded out.
- As a result, only cars of the worst possible quality, $q = 0$, will be available for sale.
- Consequently, buyers will be unwilling to pay a price higher than $p = 0$ for a second-hand car.
- No cars of quality higher than $q = 0$ will be traded in the market in the equilibrium.
- In this extreme example, adverse selection causes the market to collapse completely, leading to a highly inefficient outcome.

Adverse Selection

- The above problem can be studied as game of incomplete information.
- Since in the version discussed above the seller would move first and propose a price, it is a dynamic game with the appropriate equilibrium concept being the *perfect Bayes–Nash* (or *perfect Bayesian*) equilibrium.
- Formally, an equilibrium in such games comprises a system of consistent beliefs and a vector of sequentially rational strategies given that belief system (i.e., one strategy for each player).
- Different versions of the game can be either dynamic or static (e.g., an auctioneer proposes a price and the seller and the buyer simultaneously choose whether to accept or reject the deal).

Adverse Selection

- An equilibrium in which different types of car are traded and cannot be distinguished by the buyers is called a **pooling equilibrium.**
- An equilibrium in which only one type of car is traded, or several types are traded but can be distinguished by the buyers (e.g., no asymmetric info), is called a **separating** (or alternatively **sorting**) **equilibrium**.

Some Solutions to Adverse Selection

- The market failure described above can be (partially) overcome by a number of methods and tools, including:
  - **Information acquisition**: Uninformed agents have an incentive to improve their information. Acquiring information however may be costly, or even impossible in practice.
  - **Signaling**: The informed party undertakes an action to communicate its type to the uninformed party. Such a signal is credible only if it would not be optimal for a different type to undertake this action (for example, sellers can offer warranties for their cars in order to signal their quality).
  - **Screening**: The principal (buyer) offers a menu of contracts to the agent (seller) that induce each type of agent to voluntarily select only one contract, whereby the contracts induce self-selection.
  - **Reputation** (dynamic setup): Informed agents who engage repeatedly into an activity or market have the opportunity to establish a reputation that serves as a credible signal of their type (e.g., public ratings, brand name). This however requires that they intend to continue engaging into that activity or market in the future in order to be credible.

Asymmetric Information

- Adverse selection is an outcome of an informational deficiency.
- High-quality sellers have an incentive to improve information in the market by credibly signaling their type.
- As discussed above, such examples are warranties, certifications, professional credentials, references from previous clients, etc.
- Next, we study such an example from the labour market.

Signaling

Michael Spence's model of job market signaling (1973, QJE)

- Consider a competitive labour market with many firms (buyers of labour services) seeking to hire workers (sellers of labour services) for a specific position.
- Workers may differ in terms of productivity (how "hard-working" or "efficient" they are); this is captured by parameter θ, which denotes their marginal product.
- Suppose there exist two types of workers: high-productivity and low-productivity. The marginal products of each type are θ_H and θ_L respectively, with $\theta_L < \theta_H$.

Signaling

- It is common knowledge that a share s_H of all workers are of high productivity, while the remaining, $s_L = 1 - s_H$, is the share of low- productivity workers.

- Assume that the utility a worker of type θ enjoys from offering a unit of labour is given by

$$u(w, \theta) = w.$$

- If she chooses not to work, her outside option (e.g., her current wage or the unemployment benefit she receives) is assumed to be equal to $\overline{u}$, with $\overline{u} < \theta_L$.

- The produced output sells at price $p = 1$ per unit.

Signaling

- We start by considering the **complete information benchmark**, where the firms can observe each worker's type.
- In a perfectly competitive market with risk-neutral firms, each worker is paid the (expected) value of her marginal product:

$$w_i^* = p\theta_i.$$

- Hence a *separating equilibrium* emerges. Equilibrium wages will be
  - $w_H^* = \theta_H$ for high-productivity workers,
  - $w_L^* = \theta_L$ for low-productivity workers.

Signaling

- Under **asymmetric information** though, the workers *privately* observe their own productivity θ; firms don't observe this and hence cannot tell apart workers of different types. As a result, a pooling equilibrium emerges.

- Assuming that firms are risk-neutral, all workers are paid the (pooling) wage rate, that is, the value of their *expected* marginal product

$$w^p = (1 - s_H)\theta_L + s_H\,\theta_H.$$

- Notice that $w^p < \theta_H$, that is, the pooling wage rate is lower than the wage rate that firms would pay high-productivity workers under complete information.

- High-productivity workers have thus an incentive to find a credible way to signal their productivity to the firms.

Signaling

- Workers can acquire some form of "education". Suppose that this "education" has no effect on workers' productivity, and its cost is a deadweight loss (you can think it of as taking a test).

- Education costs a high-productivity worker c_H per unit, and a low-productivity worker c_L per unit, with $c_L > c_H > 0$.

- Assume that the utility of a worker of type i, with $i \in \{H, L\}$ is given by

$$u_i(e) = w_i - c_i e$$

  if she chooses to acquire level e of education and work for wage w_i.

- Recall that if she chooses not to work, she will receive her outside option, $\overline{u} < \theta_L$.

Signaling

- Is there a separating equilibrium, in which high- productivity workers choose education level e_H, and low- productivity workers choose education level e_L, with $e_L \neq e_H$?
- A necessary condition for such an equilibrium to exist is that high-productivity workers –but not low-productivity ones– have indeed an incentive to acquire e_H units of education.
- In order for this to be the case, the following two conditions need to be satisfied:

$$w_H - w_L = \theta_H - \theta_L \geq c_H e_H - c_H e_L \quad \text{(S.1)}$$

$$w_H - w_L = \theta_H - \theta_L \leq c_L e_H - c_L e_L. \quad \text{(S.2)}$$

- *Given* that workers decide to acquire education and apply for work at the firm, conditions (S.1) and (S.2) guarantee that each type of worker will select the intended level of education.

Signaling

- Expression (S.1) states that acquiring e_H units of education would benefit high-productivity workers more than acquiring e_L units. Otherwise, it wouldn't be worth for them to acquire level of education e_H.
- Expression (S.2) stipulates that acquiring e_H education units would not be optimal for low-productivity workers. Otherwise, it would be optimal for them to "disguise" themselves as high-productivity workers by acquiring education level e_H.
- Expressions (S.1) and (S.2) combined imply that

$$\frac{\theta_H - \theta_L}{c_L} \leq e_H - e_L \leq \frac{\theta_H - \theta_L}{c_H}. \tag{S.3}$$

- Provided that condition (S.3) holds, acquiring an education level equal to e_H credibly signals a high productivity, allowing high-productivity workers to separate themselves from low-productivity workers.

Signaling

- Yet another set of constraints need to hold in order for the above separating equilibrium to exist:

$$u_H(e_H) = \theta_H - c_H e_H \geq \bar{u} \quad \text{(S.4)}$$

$$u_L(e_L) = \theta_L - c_L e_L \geq \bar{u}. \quad \text{(S.5)}$$

- Conditions (S.4) and (S.5) state that it is preferable for workers of the respective type to acquire education levels e_H and e_L and work, compared to opting for their outside option. [*Notice that if (S.1), (S.2), and (S.5) hold, then (S.4) will also hold.*]

- *Given* that high-productivity and low-productivity workers choose to acquire education levels e_H and e_L respectively, conditions (S.4) and (S.5) guarantee that both types of workers will be better off working at that firm than opting for their outside option.

Signaling

- **Q**: Given that high-productivity workers acquire e_H units of education, how much education should low-productivity workers acquire?
- **A**: Zero. Low-productivity workers will be paid $w_L = \theta_L$ as long as they do not have e_H units of education, and they are better off with zero education and w_L than with e_H units of education and salary w_H.
- Then condition (S.3) collapses to

$$\frac{\theta_H - \theta_L}{c_L} \leq e_H \leq \frac{\theta_H - \theta_L}{c_H}, \qquad \text{(S.3')}$$

and condition (S.5) holds since $\theta_L > \overline{u}$.

- It follows then that condition (S.4) will be also satisfied.

Signaling

- **Q**: How much education e_H will high-productivity workers acquire in the most efficient separating equilibrium?
- **A**: By (S.3'), the minimum amount of education required to receive wage $w_H^* = \theta_H$ is:

$$e_H = \frac{\theta_H - \theta_L}{c_L}.$$

- If therefore condition (S.3') is satisfied, there will exist a *separating equilibrium* in which high-productivity workers acquire level of education e_H and receive wage w_H^*, and low-productivity workers acquire no education, choosing instead to work for wage w_L^*.
- Hence, if such an "education" option exists, an equilibrium similar to the one under complete information is achieved: workers are separated according to their productivity. Yet even the least wasteful outcome $(e_L = 0)$ is inefficient from a social welfare perspective. *(Why?)*

Signaling

- Signaling can transmit information and improve knowledge in a market.
- Yet, as shown in the job market signaling model, not only it may not increase total output, but it may even be costly to those who opt for it.
- As a result, signaling leads to decreased aggregate welfare in the model we discussed.
- We see thus that better information may not improve gains-to-trade if it is costly (and in some cases, even if it's free).

Moral Hazard

- The term *moral hazard* is used to describe settings in which an agent cannot observe the actions of the other agent(s). It is often referred to as “hidden action”.
- Moral hazard may cause agents’ behaviour to change after they enter an agreement or sign a contract.

Moral Hazard

- Economic models of moral hazard can be used in the study of a wide range of applications such as:
  - principal–agent problem (unobserved behaviour; voters–government, shareholders–managers/board members, employer–worker, government–contractors, etc)
  - insurance market analysis (behaviour of insured agents)
  - design of incentive schemes for bank managers in the presence of a bailout mechanism
  - optimizing bankruptcy protection legislation
  - design of social welfare policy (e.g., unemployment insurance schemes)

Moral Hazard

- Examples of efforts to avoid moral hazard are:
  - productivity- or output-bonuses for managers
  - higher life and medical insurance premiums for smokers or heavy drinkers of alcohol
  - lower car insurance premiums for contracts with higher deductibles for drivers with a record of safe driving
  - harsh terms for bailout beneficiaries

Moral Hazard: An Agency Problem

Holmström's model of *moral hazard and observability (1979, Bell JE)*

- Bengt Holmström and Oliver Hart received the Nobel prize in 2016 for their work on contract theory.
- Example: The principal–agent problem
  - The owner of a firm ("principal") cannot observe the effort of the manager ("agent") she hires to run the firm. This may be true even if the owner is perfectly informed about the manager's skills, ability and productivity.
  - As a result, the manager has an incentive to *shirk* instead of exerting costly effort, giving thus rise to moral hazard problems.

Moral Hazard: An Agency Problem

- The owner (“principal”) can offer a contract that provides incentives to the manager (“agent”) to work hard. For example, pay a higher salary (“bonus”) if the firm’s output is high, and a low salary otherwise.
- Incentivizing the manager to work hard is costly for the owner.
- The owner will choose to induce high effort from the manager only if the firm’s expected profits are sufficiently higher than they would be if the manager exerted low effort.

Moral Hazard: An Agency Problem

- Consider a principal with benefit function
$$B(\pi - w)$$
where π is the (gross) profit that arises from the agent's effort, and w is the salary that the principal pays to the agent.

- The benefit function satisfies $B' \geq 0$ and $B'' \leq 0$.

- The agent has a quasi-linear utility function given by
$$U(w, e) = u(w) - \psi(e)$$
where $u(w)$ is the utility the agent derives due to his salary (i.e., consumption), with $u'(w) > 0$ and $u''(w) \leq 0$, e is the agent's effort level, and $\psi(e)$ is the agent's disutility from exerting effort level $e \in S_e$, with $\psi'(e) > 0$ and $\psi''(e) \geq 0$.

- The agent has an *outside option* which would provide him utility equal to $\overline{u}$. This is his *reservation utility*.

Moral Hazard: An Agency Problem

- Apart from effort, the gross profit generated by the agent depends on factors that are beyond his control, and it will be therefore treated as a *random variable*.
- Of course, the agent's effort level e affects the probability that a certain level of gross profit will be realized.
- For a given effort level e, the conditional probability that gross profit will be equal to π_i is given by pmf

$$f(\pi_i|e) = \Pr(\pi = \pi_i|e) \geq 0$$

where π_i, with $i = \{1,2,\ldots,n\}$, are the various profit levels that can emerge for effort level e.

- Hence, it will be possible in general for a high profit to arise even if the agent shirks, or a low profit level to arise even if the agent exerts effort.

Moral hazard: An Agency Problem

- An important assumption affecting the conclusions drawn from this model has to do with the risk attitude of each party.
- Three cases are of special interest:

  Case 1 (C1): The principal is risk-neutral but the agent is risk-averse

  Case 2 (C2): The principal is risk-averse but the agent is risk-neutral

  Case 3 (C3): Both the principal and the agent are risk-averse

- We discuss only Cases 1 and 2 in this lecture. Students interested in Case 3 are encouraged to go through the relevant part in Chapter 10.1 in Muñoz-Garcia (2017).

Symmetric Information

- **Benchmark case:** Assume that the agent's effort level e is *observable and verifiable* by the principal.

- The general principal's maximization problem is then

$$\max_{\{e, w(\pi_i)\}_{i=1}^{n}} \mathrm{E}[B(\pi_i - w)|e] = \sum_{i=1}^{n} f(\pi_i|e)\, B(\pi_i - w)$$

$$\text{s.t. } \mathrm{E}[u(w) - \psi(e)|e] = \sum_{i=1}^{n} f(\pi_i|e)\, [u(w) - \psi(e)] \geq \bar{u}.$$

- The principal seeks to maximize expected profits, subject to the agent participating in the contract. Hence the agent's *participation constraint* must be satisfied.

Symmetric Information

- The principal could
  - present the agent with a contract linking the latter's salary directly on his level of effort e, or
  - make the agent's salary contingent on the realized profit π.
- In either case, intuition suggests that the participation constraint must be binding (i.e., hold with equality) at the optimal contract for the principal. Otherwise, since B is decreasing in w, and u is a continuous function, the principal could increase her benefit by decreasing w, and the agent would still agree to participate.

Symmetric Information

- To solve for $w(\pi_i)$ we can set up the Lagrangian function

$$\mathcal{L}(w(\pi_i),\lambda) = \sum_{i=1}^{N} f(\pi_i|e)\, B(\pi_i - w(\pi_i)) + \lambda\left[\sum_{i=1}^{N} f(\pi_i|e)\, u(w(\pi_i)) - \psi(e) - \bar{u}\right]$$

- The first-order condition (FOC) with respect to $w(\pi_i)$ is

$$-f(\pi_i|e)\, B'(\pi_i - w(\pi_i)) + \lambda f(\pi_i|e)\, u'(w(\pi_i)) = 0$$

for each $w(\pi_i)$, where B' and u' are the derivatives of the functions B and u respectively.

Symmetric Information

- Rearranging yields

$$\lambda u'\big(w(\pi_i)\big) = B'\big(\pi_i - w(\pi_i)\big),$$

and solving for λ

$$\lambda = \frac{B'\big(\pi_i - w(\pi_i)\big)}{u'\big(w(\pi_i)\big)}$$

which is positive since $B'(\cdot) > 0$ and $u'(\cdot) > 0$.

- $\lambda > 0$ suggests that the agent's participation constraint must indeed bind at the optimum:

$$\sum_{i=1}^{N} f(\pi_i|e)u\big(w(\pi_i)\big) - \psi(e) = \bar{u}.$$

Moral Hazard: A Simple Model

- To simplify our analysis, suppose that the agent's effort space S_e contains only two elements, namely $S_e = \{e_0, e_1\}$.
- That is, the agent can choose only between two different actions: shirk (e_0) or work (e_1), and it will be assumed that $e_0 = 0$ and $e_1 = 1$.
- Suppose, moreover, that the generated profit can assume only two values: $\pi_1 = \pi_L$ and $\pi_2 = \pi_H$, with $\pi_H > \pi_L$.
- The agent's reservation utility is set to $\overline{u} = 1$, and the principal's outside option is normalized to $\overline{B} = 0$.
- Social welfare is simply defined as the sum of the principal's benefit and the agent's utility (the so-called *Walrasian utility function*: a common, but highly normative, non-innocuous assumption).
- To simplify notation, denote

$$p_0 := f(\pi_H | e = 0), \text{ and}$$
$$p_1 := f(\pi_H | e = 1).$$

Moral Hazard: A Simple Model

- The probability mass function of the random variable π_i representing gross profit conditional on the effort level e exerted by the agent, is given in Table 1 below.

π_i	$f(\pi_i \vert e)$	
	$e = 0$	$e = 1$
8	$\frac{3}{4}$	$\frac{1}{4}$
20	$\frac{1}{4}$	$\frac{3}{4}$

Table 1: Conditional pmf of gross profit given agent's effort

- We can now calculate expected gross profits for each effort level:

$$\mathrm{E}[\pi \,|e = 0] = p_0\pi_H + (1 - p_0)\pi_L = \frac{1}{4} \cdot 20 + \frac{3}{4} \cdot 8 = 11$$

$$\mathrm{E}[\pi \,|e = 1] = p_1\pi_H + (1 - p_1)\pi_L = \frac{3}{4} \cdot 20 + \frac{1}{4} \cdot 8 = 17$$

Case 1: The Setup

Case 1: Risk-neutral principal, risk-averse agent

- Assume that the agent is risk-averse, with his utility from wealth (consumption) given by $u(w) = \sqrt{w}$, and his disutility of effort given by $\psi(e) = e$.

- The principal shall be assumed to be risk-neutral, and hence her benefit coincides with her net profits:

$$B(\pi - w) = \pi - w.$$

- If effort is observable and verifiable, then the principal could present to the agent a contract linking his salary directly to the agent's realized level of effort, e.

- In particular, if the agent exerts effort, he is to receive a salary equal to w_1; if he doesn't, he is to receive w_0.

C1: Symmetric Information

- The principal's maximization problem can be written as

$$\max_{w_e} \mathsf{E}[B(\pi - w_e)|e] = p_e\pi_H + (1-p_e)\pi_L - w_e$$
$$\text{s.t.} \quad \mathsf{E}[u(w_e) - e|e] = \sqrt{w_e} - e \geq \bar{u}$$

- One approach would be to find the benefit-maximizing salary w_e^{SI}, and calculate the corresponding expected benefit for each effort level; then compare the two benefit levels and choose the highest one.
- Notice that the principal's objective function is decreasing in w. Hence the participation constraint must be binding at the optimal contract for the principal. Otherwise (since u is a continuous function) the principal could increase her expected benefit by decreasing w_e.

C1: Symmetric Information

- So the principal can induce the agent to work (exert effort level $e = 1$) by paying him a salary w_1 that satisfies his participation constraint with equality:

$$u\left(w_1^{SI}\right) - \psi(1) = \overline{u}$$
$$\sqrt{w_1^{SI}} - 1 = 1$$
$$w_1^{SI} = 4\,.$$

- Similarly, the principal can induce the agent to shirk (effort level $e = 0$) by offering salary w_0 such that

$$u\left(w_0^{SI}\right) - \psi(0) = \overline{u}$$
$$\sqrt{w_0^{SI}} - 0 = 1$$
$$w_0^{SI} = 1.$$

C1: Symmetric Information

- The principal's expected benefit (net profit) is

$$\mathsf{E}[B(\pi - w_e)|e] = \mathsf{E}[\pi|e] - w_e$$

and hence we can calculate that

$$\mathsf{E}[B(\pi - w_e)|e] = \begin{cases} 13, & \text{if } e = 1 \\ 10, & \text{if } e = 0. \end{cases}$$

- If effort is therefore observable and verifiable, it would be optimal for the principal to induce the agent to work by offering him a contract such as the following:

$$w_e^{SI} = \begin{cases} 4, & \text{if } e = 1 \\ a, & \text{if } e = 0, \end{cases}$$

for some $a < 1$.

- It would be optimal for the agent to accept that contract, and then choose to work. Social surplus would be maximized (*first best* solution).

C1: Symmetric Information

- What if the two parties wrote instead a contract contingent on the level of realized gross profit? *(We will see later why we study this case.)*
- According to that contract, the agent would commit to work ($e = 1$).
- His salary then would be a function of the gross profit that his effort would generate for the principal: $w(\pi_H)$ and $w(\pi_L)$.
- What would be the optimal levels of $w(\pi_H)$ and $w(\pi_L)$, to be denoted with w_H^{SI} and w_L^{SI} respectively?
- The principal's programme would be:

$$\max_{\{w(\pi_i)\}_{i\in\{H,L\}}} \mathrm{E}[B(\pi - w(\pi)|e = e_1] = p_1(\pi_H - w(\pi_H)) + (1 - p_1)(\pi_L - w(\pi_L))$$

$$\text{s.t.} \quad \mathrm{E}[u(w(\pi)) - e|e = e_1] = p_1\sqrt{w(\pi_H)} + (1 - p_1)\sqrt{w(\pi_L)} - e_1 \geq \bar{u}.$$

C1: Symmetric Information

- We can set up the Lagrangian function:

$$\mathcal{L}(w(\pi_H), w(\pi_L), \lambda) = p_1\big(\pi_H - w(\pi_H)\big) + (1-p_1)\big(\pi_L - w(\pi_L)\big)$$
$$+\lambda\left[p_1\sqrt{w(\pi_H)} + (1-p_1)\sqrt{w(\pi_L)} - e_1 - \bar{u}\right]$$

- The first-order conditions (FOC) are

$$w(\pi_H): \quad -p_1 + \lambda p_1 \frac{1}{2\sqrt{w_H^{SI}}} = 0 \qquad (\text{SI}.1)$$

$$w(\pi_L): \quad -(1-p_1) + \lambda(1-p_1)\frac{1}{2\sqrt{w_L^{SI}}} = 0 \qquad (\text{SI}.2)$$

$$\lambda: \quad p_1\sqrt{w_H^{SI}} + (1-p_1)\sqrt{w_L^{SI}} - e_1 - \bar{u} \geq 0 \qquad (\text{SI}.3)$$

C1: Symmetric Information

- Rearranging conditions (SI.1) and (SI.2) yields

$$\lambda = 2\sqrt{w_H^{SI}} \text{ and}$$
$$\lambda = 2\sqrt{w_L^{SI}},$$

which implies that the two salary levels will be equal in the optimal contract. In particular, it must hold that

$$w_H^{SI} = w_L^{SI} = \frac{1}{4}\lambda^2. \qquad \text{(SI. 4)}$$

- Since the agent is risk-averse, he would like to be fully insured, and receive the same salary irrespectively of the realized profit.

- This is a *risk sharing* result. In the present case the principal, being risk-neutral, is only interested in maximizing her net profit. She will thus undertake all the risk herself, providing full insurance to the agent.

C1: Symmetric Information

- As discussed above, the participation constraint must bind at the optimum, and thus it follows from (S.4) that

$$p_1\sqrt{w_H^{SI}} + (1 - p_1)\sqrt{w_H^{SI}} - e_1 - \bar{u} = 0.$$

- Using values $e_1 = 1$, and $\bar{u} = 1$, we find that

$$w_H^{SI} = w_L^{SI} = 4.$$

- Hence this profit-contingent contract gives rise to an outcome (effort, expected benefit, and expected utility) identical to the one under the effort-contingent contract analyzed above.

- The first best (maximum social surplus) is achieved.

C1: Asymmetric Information

- **Moral hazard:** Assume that the effort exerted by the agent is either not observable by the principal, or not verifiable (or both).
- As a result, in any contract, the agent's salary can no longer be conditional on his effort level, but only on realized gross profit.
- A contract could thus specify only profit-contingent salaries, $w(\pi_H)$ and $w(\pi_L)$, not effort-contingent ones.
- The corresponding participation constraint of course still needs to be satisfied at the effort level e_j chosen by the principal, $j \in \{0,1\}$:

$$p_j\sqrt{w(\pi_H)} + (1 - p_j)\sqrt{w(\pi_L)} - e_j \geq \bar{u}.$$

C1: Asymmetric Information

- If, however, the principal would like the agent not anoly to accept the contract, but also to work, the following *incentive compatibility constraint* (IC) must hold as well:

$$\mathrm{E}[U(w(\pi), e)|e = e_1] \geq \mathrm{E}[U(w(\pi), e)|e = e_0],$$

because the agent cannot be forced to work via a contract. Expanding this expression gives:

$$p_1 u\big(w(\pi_H)\big) + (1 - p_1) u\big(w(\pi_L)\big) - \psi(e_1) \\ \geq\ p_0 u\big(w(\pi_H)\big) + (1 - p_0) u\big(w(\pi_L)\big) - \psi(e_0)$$

- The principal's problem can be written as

$$\max_{\{e, w(\pi)\}} \quad \mathrm{E}[B(\pi - w(\pi))]$$

$$\text{s.t.} \quad \mathrm{E}[U(w(\pi), e)|e = e_1] \geq \bar{u}. \qquad \text{(PC)}$$

$$\mathrm{E}[U(w(\pi), e)|e = e_1] \geq \mathrm{E}[U(w(\pi), e)|e = e_0] \qquad \text{(IC)}$$

C1: Asymmetric Information

- The Lagragian is

$$\mathcal{L}(w(\pi_H), w(\pi_L), \lambda, \mu) = \quad p_1(\pi_H - w(\pi_H)) + (1 - p_1)(\pi_L - w(\pi_\mathrm{L}))$$

$$+\lambda\left[p_1\sqrt{w(\pi_H)} + (1 - p_1)\sqrt{w(\pi_L)} - e_1 - \bar{u}\right]$$

$$+\mu\left[p_1\sqrt{w(\pi_H)} + (1 - p_1)\sqrt{w(\pi_L)} - e_1\right.$$

$$\left.-p_0\sqrt{w(\pi_H)} - (1 - p_0)\sqrt{w(\pi_L)} + e_0\right].$$

C1: Asymmetric Information

- The corresponding FOC are

$$w(\pi_H): \quad -p_1 + \lambda p_1 \frac{1}{2\sqrt{w_H}} + \mu(p_1 - p_0)\frac{1}{2\sqrt{w_H}} = 0 \qquad (\mathrm{AI.}\,1)$$

$$w(\pi_L): \quad -(1-p_1) + \lambda(1-p_1)\frac{1}{2\sqrt{w_L}} - \mu(p_1 - p_0)\frac{1}{2\sqrt{w_L}} = 0 \qquad (\mathrm{AI.}\,2)$$

$$\lambda: \quad p_1\sqrt{w_H} + (1-p_1)\sqrt{w_L} - e_1 - \bar{u} \geq 0 \qquad (\mathrm{AI.}\,3)$$

$$\mu: \quad (p_1 - p_0)\sqrt{w_H} - (p_1 - p_0)\sqrt{w_L} - e_1 + e_0 \geq 0 \qquad (\mathrm{AI.}\,4)$$

C1: Asymmetric Information

- First, notice that the principal's objective function is decreasing in $w(\pi_H)$ and $w(\pi_L)$. This suggests that both constraints must bind at the optimal contract for the principal. Otherwise—since u is a continuous function—she could increase her expected benefit by decreasing either $w(\pi_H)$ or $w(\pi_L)$, or both. This implies that $\lambda > 0$ and $\mu > 0$.

- We can now use equations (AI.3) and (AI.4) with equality (a system of two equations in two unknowns) to characterize (and potentially find explicitly) the salaries under the optimal contract, w_H and w_L:

$$p_1\sqrt{w_H} + (1 - p_1)\sqrt{w_L} - e_1 - \bar{u} = 0 \qquad \text{(AI. 3')}$$

$$(p_1 - p_0)\sqrt{w_H} - (p_1 - p_0)\sqrt{w_L} - e_1 + e_0 = 0. \qquad \text{(AI. 4')}$$

C1: Asymmetric Information

- Yet we can draw some conclusions even if we cannot explicitly solve the above system.
- It follows from conditions (AI.1) and (AI.2) respectively that

$$w_H = \left[\frac{1}{2}\lambda + \frac{1}{2}\mu\left(1 - \frac{p_0}{p_1}\right)\right]^2 \qquad (\mathrm{AI.}\,1')$$

$$w_L = \left[\frac{1}{2}\lambda - \frac{1}{2}\mu\left(1 - \frac{p_0}{p_1}\right)\right]^2. \qquad (\mathrm{AI.}\,2')$$

- Notice that unlike the corresponding expressions (SI.1’) and (SI.2’) in the case of symmetric information, (AI.1’) and (AI.2’) above suggest that the optimal contract under asymmetric information **prescribes different salaries** depending on the realized profits.
- Working is costly for the agent. In the presence of moral hazard therefore, he cannot commit to work if he is promised a fixed salary (due to moral hazard).

C1: Asymmetric Information

- We can now use the values of our example to find a numerical solution to the principal's problem. Substituting into conditions (AI.3') and (AI.4'), and using elementary algebra yields

$$3\sqrt{w_H} + \sqrt{w_L} = 8$$

$$\sqrt{w_H} - \sqrt{w_L} = 2.$$

- Solving this system we find the optimal profit-contingent salaries:

$$w_H = \frac{25}{4} = 6.25$$

$$w_L = \frac{1}{4} = 0.25.$$

- The principal's expected benefit (net profit) if the agent works is now

$$\mathrm{E}[B(\pi - w(\pi))|e = e_1] = p_1(\pi_H - w_H) + (1 - p_1)(\pi_L - w_L)$$

$$= \frac{3}{4}(20 - 6.25) + \frac{1}{4}(8 - 0.25)$$

$$= 12.25.$$

C1: Asymmetric Information

- It can be seen that the principal's expected benefit if she is risk-neutral and the agent is risk-averse is lower compared to the case of complete information (perfect observability/verifiability).
- In order for the principal to give the agent an incentive to work in the presence of moral hazard, she needs to punish him in the state of the world that is more likely to occur if he shirks.
- However, to compensate for this "punishment" and induce the agent to accept the contract, the principal needs to increase the agent's salary in the "good" state of the world, and more than proportionally so, since the agent is risk-averse.
- Is it worth though for the principal to induce the agent to work?

C1: Asymmetric Information

- If the principal would like the agent just to accept the contract, only constraint (PC) needs to be satisfied.
- The principal's problem becomes thus identical to the one under symmetric information, where it was shown that
$$\mathrm{E}[B(\pi - w(\pi))|e = e_0] = 10.$$
- It is hence optimal for the principal to induce the agent to work.
- The agent's maximized expected utility remains at its reservation level $\bar{u}$, yet the principal's expected benefit is now lower.
- Due to the presence of moral hazard, social welfare is now lower, compared to the symmetric info case, and only a *second best* can be achieved.

Case 2: The Setup

Example 2: Risk-averse principal, risk-neutral agent

- Assume now instead that that the principal is risk-averse, with her benefit function being

$$B(\pi - w) = \sqrt{\pi - w},$$

and the agent is risk-neutral, with $u(w) = w$ and $\psi(e) = e$, so that his utility function is given by

$$U(w, e) = w - e.$$

- Observe that the principal's benefit does not coincide with her net profits in this case.

- Parameter values and distributions are as stipulated above.

C2: Symmetric Information

- What would be now the optimal profit-contingent contract, if effort is *observable and verifiable*?
- Let the agent's salary be again a function of the gross profit generated, given by $w(\pi_H)$, and $w(\pi_L)$, and denote the optimal levels with w_H^{SI} and w_L^{SI} respectively.
- The contract may also state whether the agent is required to work.
- The principal's programme will be now:

$$\max_{\{e_i, w(\pi_j)\}_{\substack{i\in\{0,1\} \\ j\in\{H,L\}}}} \mathrm{E}[B(\pi - w(\pi))|e] = p_e\sqrt{\pi_H - w(\pi_H)} + (1-p_e)\sqrt{\pi_L - w(\pi_L)}$$

$$\text{s.t.} \qquad \mathrm{E}[u(w(\pi)) - e|e] = p_e\, w(\pi_H) + (1-p_e)\, w(\pi_L) - e_i \geq \bar{u}$$

C2: Symmetric Information

- The Lagrangian function will be

$$\mathcal{L}(w(\pi_H), w(\pi_H), \lambda) = \quad p_e\sqrt{\pi_H - w(\pi_H)} + (1 - p_e)\sqrt{\pi_L - w(\pi_L)}$$
$$+\lambda\,[p_e\, w(\pi_H) + (1 - p_e)w(\pi_L) - e\ - \bar{u}],$$

and the associated FOC:

$$w(\pi_H): \qquad -p_e \frac{1}{2\sqrt{\pi_H - w_H^{SI}}} + \lambda p_e = 0 \qquad (\mathrm{SI.}\,4)$$

$$w(\pi_L): \qquad -(1 - p_e)\frac{1}{2\sqrt{\pi_L - w_L^{SI}}} + \lambda(1 - p_e) = 0 \qquad (\mathrm{SI.}\,5)$$

$$\lambda: \qquad p_e w_H^{SI} + (1 - p_e)w_L^{SI} - e\ - \bar{u} \geq 0 \qquad (\mathrm{SI.}\,6)$$

C2: Symmetric Information

- Combining conditions (SI.4) and (SI.5) gives

$$\lambda = \frac{1}{2\sqrt{\pi_H - w_H^{SI}}} = \frac{1}{2\sqrt{\pi_L - w_L^{SI}}} \qquad \text{(SI.7)}$$

  which implies that

$$B\left(\pi_H - w_H^{SI}\right) = B\left(\pi_L - w_L^{SI}\right).$$

- This suggests that the principal, who is risk-averse, would like to insure herself against an adverse profit outcome, irrespectively of the effort level to be exerted by the agent.
- The agent is risk-neutral, and hence willing to undertake all risk himself.

C2: Symmetric Information

- As in Case 1, the participation constraint must bind at the optimum. Expressions (SI.6) and (SI.7) respectively give thus rise to the following system:

$$p_e w_H^{SI} + (1 - p_e) w_L^{SI} = e + \bar{u}$$

$$\pi_H - w_H^{SI} = \pi_L - w_L^{SI}.$$

- Solving this system of equations gives the optimal salaries:

$$w_H^{SI} = \bar{u} + e + (1 - p_e)(\pi_H - \pi_L)$$
$$w_L^{SI} = \bar{u} + e - p_e(\pi_H - \pi_L).$$

- Notice that in this case, the optimal salary is different in each state of the world. In particular, it is higher than the agent's reservation utility plus his cost of effort in the event of a high-profit outcome, and lower in the event of a low-profit outcome.

C2: Symmetric Information

- Observe that under the contract stipulated above, the principal's net profit is independent of the profit realization, and given by

$$\pi_H - w_H^{SI} = \pi_L - w_L^{SI} = p_e \pi_H + (1 - p_e)\pi_L - \bar{u} - e$$
$$= \mathrm{E}[\pi|e] - (\bar{u} + e).$$

- In either state of the world, the principal receives the expected value of the gross profits generated by the agent's effort, less the minimum compensation required for the agent to participate in the contract.

- This is effectively a *franchise contract*: Notice that it is as if the agent transfers to the principal a fixed payment t_e that depends on the agent's effort,

$$t_e := p_e \pi_H + (1 - p_e)\pi_L - \bar{u} - e,$$

and then keeps for himself all profit generated by his work; see indeed that it holds that $w_H^{SI} = \pi_H - t_e$, and $w_L^{SI} = \pi_L - t_e$.

C2: Symmetric Information

- Plugging in the numerical values used above, the principal's benefit if the agent works ($e = 1$) is given by

$$B(t_1) = \sqrt{p_1\pi_H + (1 - p_1)\pi_L - \bar{u} - e}$$
$$= \sqrt{\tfrac{3}{4} \cdot 20 + \tfrac{1}{4} \cdot 8 - 1 - 1}$$
$$= 3.873\,.$$

- If the agent does not work ($e\ = 0$) the principal's benefit will be

$$B(t_0) = \sqrt{p_0\pi_H + (1 - p_0)\pi_L - \bar{u} - e}$$
$$= \sqrt{\tfrac{1}{4} \cdot 20 + \tfrac{3}{4} \cdot 8 - 1 - 0}$$
$$= 3.162\,.$$

- It will be therefore optimal for the principal to offer a contract stipulating that the agent should work.

C2: Symmetric Information

- The optimal (first best) profit-dependent salaries will be

$$w_H^{SI} = 1 + 1 + \frac{1}{4} \cdot 12 = 5$$
$$w_L^{SI} = 1 + 1 - \frac{3}{4} \cdot 12 = -7,$$

or equivalently, the agent's transfer to acquire the franchise will be

$$t_1 = \frac{3}{4} \cdot 20 + \frac{1}{4} \cdot 8 - 1 - 1 = 15.$$

- It will be optimal then for the agent to accept this contract.

C2: Asymmetric Information

- **Moral hazard:** Let us now examine how asymmetric information affects the first-best contract by assuming that the agent's effort is not observable or verifiable.
- The contract cannot be contingent on the agent's effort level. The principal's only contractual tool is to set differentiated salaries for the agent, w_H and w_L, depending on the realized level of profits.
- It follows that the principal must make sure that the agent's incentive compatibility constraint is satisfied (apart from his participation constraint), that is it must hold that

$$\mathrm{E}[U(w(\pi), e) | e = e_1] \geq \mathrm{E}[U(w(\pi), e) | e = e_0],$$

or

$$p_1 w_H + (1 - p_1) w_L - \psi(e_1) \geq p_0 w_H + (1 - p_0) w_L - \psi(e_0).$$

C2: Asymmetric Information

- Observe however that the salaries stipulated in the first-best contract under symmetric information do in fact satisfy the incentive constraint too:

$$\frac{3}{4} \cdot 5 + \frac{1}{4} \cdot (-7) - 1 \geq \frac{1}{4} \cdot 5 + \frac{3}{4} \cdot (-7)$$
$$1 \geq -4.$$

- It follows then that the optimal contract under moral hazard will be identical with the one under symmetric information, that is,

$$w_H = w_H^{SI} = 5$$
$$w_L = w_L^{SI} = -7,$$

and the agent will choose to work.

- It can be therefore seen that although moral hazard *may* distort behaviour and lead to suboptimal outcomes, this does not *need* to be always the case.

"If this paper was correct,
Economics would be different…"

– *Excerpt from a referee report for the Journal of Political Economy, recommending rejection of the paper "The market for lemons" submitted by George Akerlof*

Practice exercises

1. Consider a market for used cars. The value of a car depends on its quality. Assume that there are only two types of used cars: *lemons*, with quality $q_L = 5$, and *peaches*, with quality $q_P = 10$. It is known that a share s_L of all used cars in the market are lemons. Each buyer's utility function following the purchase of a used car of quality q is given by $u(p, q) = m_b q - p$, and each seller's profit function is given by $\pi(p, q) = p - q$, where p is the price at which the car is traded, and m_b is a "*mark-up*" that represents the factor by which the buyer's valuation of a car of a given quality is higher than that of the seller. If no trade occurs it is assumed that $\overline{u} = \overline{\pi} = 0$.

   a) What condition must hold true for m_b in order for a market for used cars to exist?

   b) Let $m_b = {}^6/_5$. Assuming that buyers can observe car quality q, find the equilibrium price $p^*(q)$ of a car of quality q if sellers have all the bargaining power.

   c) Suppose now that car quality is not observable by the buyers. Find the equilibrium price $p^{AI}(q)$ of a car of quality q if as a function of s_L if both car types are traded in the market.

Practice exercises

d) Assume that $s_L = {}^1/_6$. Find explicitly $p^{AI}(q)$. Are both lemons and peaches traded in the market? Is this equilibrium pooling or separating? Is a first-best achieved?

e) What is the maximum value of s_L for which both types of cars are traded in the market? What happens if s_L is higher than that value?

Practice Exercises

2. Consider the agency problem with the risk-neutral principal and the risk-averse agent examined in Case 1. Assume that the principal's and the agent's objective functions are as above, $\pi_L = 9$, $\psi(e_0) = 1$, $\psi(e_1) = 2$, and the rest of the values are as above.

   a) Set up formally the principal's programme and find the optimal contract in the case of symmetric information (effort observable and verifiable).

   b) Set up formally the principal's programme, and write the first-order conditions for the case of asymmetric information (unobservable/ non-verifiable effort).

   c) Find explicitly the optimal contract under asymmetric information. Is it different from the first-best contract?

   d) Is a first best attained under asymmetric information? Explain intuitively why or why not.