



Ciências
ULisboa

Faculdade
de Ciências
da Universidade
de Lisboa

Principles of Economics and Management I

Session 6 · Uncertainty, Risk and Asymmetric Information

Introduction



Where We Came From

So far everything has been **certain**: the consumer knew the prices, the firm knew its costs.

But almost no financial decision works like that. Today **uncertainty** and hidden information enter the picture. 🎲

This is the session closest to what you will actually do in banking, insurance and risk management. 🛡️



Today's Map

1. 🎲 Expected utility and risk aversion
2. 💶 Certainty equivalent and risk premium
3. ☁️ Asymmetric information: adverse selection and moral hazard
4. 🛡️ How insurance works
5. 🧳 How insurers respond to asymmetric information

Part 1 · Choice Under Uncertainty




A Gamble

I offer you a choice:

(A) €52 thousand for **certain**.

(B) A coin toss: €4 thousand on tails, €100 thousand on heads.

The two are worth **the same** on average. Most people choose (A). Why? 🤔



Expected Utility

The answer: people do not maximize expected **value**, they maximize expected **utility**.

$$E[u(w)] = \sum_i p_i u(w_i)$$

Each outcome w_i occurs with probability p_i , and what counts is the utility it delivers, not the euro itself.



A Tool First: Jensen's Inequality

We need one result from probability before going on.

Jensen's inequality: if u is concave and w is a random variable, then $E[u(w)] \leq u(E[w])$. If u is strictly concave and w is not constant, the inequality is strict.

With two outcomes, w_1 with probability p and w_2 with $1 - p$, it reads

$$p u(w_1) + (1 - p) u(w_2) \leq u(p w_1 + (1 - p) w_2),$$

which is the definition of concavity you know from calculus. Jensen extends it to any distribution.

For a **convex** function the inequality flips. You already know one case: x^2 is convex, so $E[w^2] \geq (E[w])^2$, which is $\text{Var}(w) \geq 0$.



Jensen's Inequality: Proof

Let $\mu = E[w]$, and assume u is differentiable.

A concave function lies below each of its tangent lines. Take the one at μ :

$$u(w) \leq u(\mu) + u'(\mu)(w - \mu) \quad \text{for every } w.$$

Take expectations of both sides. The expectation keeps the inequality, and $E[w - \mu] = 0$:

$$E[u(w)] \leq u(\mu) + u'(\mu)(E[w] - \mu) = u(\mu) = u(E[w]). \quad \blacksquare$$

If u is strictly concave the tangent inequality is strict for every $w \neq \mu$, so the result is strict as soon as $w \neq \mu$ with positive probability.



Risk Aversion Is Concavity

If u is **concave**, each additional euro is worth less than the one before ($u'' < 0$).


By Jensen's inequality, for any nondegenerate gamble:

$$E[u(w)] < u(E[w]).$$

Risk aversion: preferring the expected value for certain over the gamble. It is exactly the concavity of u .



Risk Aversion Is Concavity



The chord (the gamble) lies **below** the curve (certainty). That is concavity.



Certainty Equivalent and Risk Premium

Certainty equivalent (CE): the sure amount that gives the same utility as the gamble, $u(CE) = E[u(w)]$.

Risk premium: $E[w] - CE$. It is what you are willing to pay **not** to bear the risk.

And this is exactly where willingness to pay for **insurance** comes from. 



How Risk Averse, Exactly?

“Concave” is qualitative. Ana and Bruno both refuse the coin toss, but Ana would pay more than Bruno to avoid it. How do we measure that?

First idea: use u'' , how fast the marginal utility falls. The problem is that u is not unique.

Take $v(w) = a u(w) + b$ with $a > 0$. Then $E[v(w)] = a E[u(w)] + b$, so v ranks every gamble exactly as u does, and both describe the same preferences. But $v'' = a u''$, so u'' would give a different number for the same person.

Dividing by u' removes the a , since $v''/v' = a u''/(a u') = u''/u'$. This gives the **Arrow-Pratt** coefficient of absolute risk aversion:

$$A(w) = -\frac{u''(w)}{u'(w)}$$

The minus sign makes it positive for a risk averse agent.



Arrow-Pratt: Some Examples

$u(w)$	$A(w)$	Reading
w	0	risk neutral, only $E[w]$ matters
$\ln w$	$1/w$	falls with wealth
$\frac{w^{1-\rho}}{1-\rho}, \rho > 0, \rho \neq 1$	ρ/w	falls with wealth
$-e^{-aw}$	a	constant

When $A(w)$ falls with w , a richer person dislikes the **same** gamble in euros less. For a bet of 1000 euros that sounds right: it hurts a student much more than a bank.

Check the third line yourself: $u' = w^{-\rho}$ and $u'' = -\rho w^{-\rho-1}$. 



The Risk Premium from Arrow-Pratt

Write the gamble as $w = \mu + \varepsilon$, with $E[\varepsilon] = 0$ and $\text{Var}(\varepsilon) = \sigma^2$. The risk premium π is defined by

$$u(\mu - \pi) = E[u(\mu + \varepsilon)].$$

Right side, a Taylor expansion of order two around μ , then expectations:

$$E[u(\mu + \varepsilon)] \approx E\left[u(\mu) + u'(\mu)\varepsilon + \frac{1}{2}u''(\mu)\varepsilon^2\right] = u(\mu) + \frac{1}{2}u''(\mu)\sigma^2.$$

Left side, order one is enough: $u(\mu - \pi) \approx u(\mu) - u'(\mu)\pi$.

Equate the two and solve for π :

$$\pi \approx \frac{1}{2} \left(-\frac{u''(\mu)}{u'(\mu)} \right) \sigma^2 = \frac{1}{2} A(\mu) \sigma^2.$$



The Risk Premium from Arrow-Pratt

Why order two on the right and only order one on the left?

ε is of the size of σ , and its first order term has mean zero. So the first term that survives the expectation is the one in ε^2 , of size σ^2 .

The answer π is itself of size σ^2 . On the left a π^2 term would be of size σ^4 , much smaller when the risk is small. So we drop it.

The formula is for **small** risks. Example with $u = \sqrt{w}$ and a 50/50 gamble between 64 and 144: $\mu = 104$, $\sigma = 40$. Exact: $E[u] = (8 + 12)/2 = 10$, so $CE = 100$ and $\pi = 4$.

Approximation: $\frac{1}{2} \cdot \frac{1}{208} \cdot 1600 \approx 3.85$. Close.

So the risk premium is proportional to the **variance**, and A is the factor. This is one reason variance is the risk measure you will meet everywhere in finance. 



Diversification

If risk aversion is expensive, how can we reduce risk without giving up return?

Two assets with the same mean and variance σ^2 , held half and half. If $\text{Cov}(X_1, X_2) = 0$:

$$\text{Var}\left(\frac{1}{2}X_1 + \frac{1}{2}X_2\right) = \frac{1}{4}\sigma^2 + \frac{1}{4}\sigma^2 + 2 \cdot \frac{1}{4}\underbrace{\text{Cov}}_{=0} = \frac{\sigma^2}{2}.$$

The same expected return, **half** the variance. With n independent assets, $\sigma^2/n \rightarrow 0$. 

 Source: OpenStax *Micro* 3e, §17.2 “How Households Supply Financial Capital”, “The Tradeoffs between Return and Risk”, which makes the argument for diversification without the algebra.



Diversification

Independence was doing all the work above. In general, for n equally weighted assets:

$$\text{Var}(\bar{X}) = \frac{1}{n} \overline{\sigma^2} + \frac{n-1}{n} \overline{\text{Cov}}$$

As $n \rightarrow \infty$ the first term vanishes and the second converges to $\overline{\text{Cov}}$. Individual variance can be diversified away; **average covariance cannot**.

That residual is **systematic risk**: what remains when everything moves together. It is why a portfolio of 500 European bank stocks is still a bet on European banking.

? Multiple Choice 1

An agent with $u(w) = \sqrt{w}$ is offered a fair gamble. They:

A. Accept it, because the expected value is unchanged.

B. Are indifferent, because $E[w]$ is the same either way.

C. Refuse it, because $\sqrt{\cdot}$ is concave and Jensen gives $E[u(w)] < u(E[w])$.

D. Cannot decide without knowing their initial wealth.

✓ **C.** A fair gamble leaves the mean unchanged and adds variance, and a concave u dislikes variance. Note that D is wrong here for a specific reason: the *sign* of the answer does not depend on wealth, only its size does.

? Multiple Choice 2

The risk premium is:


A. The difference between the expected value and the certainty equivalent.

B. The price charged by the insurance company.

C. Always equal to the expected loss.

D. Zero for any risk averse agent.

✓ **A.** It is $E[w] - CE$: how much the agent will give up, on average, to be rid of the risk. It is what makes an insurer's margin possible. It is not the premium itself.



Numerical Question

Utility $u(w) = \sqrt{w}$, and a 50/50 gamble between 4 and 100 thousand euros.

- a) Find $E[w]$ and $E[u(w)]$.
- b) Find the certainty equivalent.
- c) Find the risk premium.
- d) Compute $A(w)$ at $w = E[w]$ and compare $\frac{1}{2}A\sigma^2$ with your answer to (c).

Solution

- a)** $E[w] = 0.5(4) + 0.5(100) = 52$, and $E[u] = 0.5(2) + 0.5(10) = 6$ thousand euros.
- b)** $\sqrt{CE} = 6$, so $CE = 36$ thousand euros.
- c)** Risk premium $= 52 - 36 = 16$ thousand euros. Almost a third of the expected value: this is a very large gamble.
- d)** $u' = \frac{1}{2}w^{-1/2}$ and $u'' = -\frac{1}{4}w^{-3/2}$, so $A(w) = 1/(2w) = 1/104$. With $\sigma^2 = 48^2 = 2304$: $\frac{1}{2}(2304)/104 \approx 11$. The right order of magnitude, but short of 16: the approximation is for small risks, and this gamble is large. 

Part 2 · Asymmetric Information



Imperfect Information

We rarely decide with perfect information.

You buy a used phone without knowing the real state of the battery.

Asymmetric information: when one side of the transaction knows **more** than the other.

The Used Car Market

A classic example (Akerlof). There are good cars and “lemons” (bad ones). Only the **seller** knows which is which.

The buyer, not knowing, will only pay an **average** price.

At that average price, whoever has a good car will not sell (it is worth more). Only the **lemons** remain. 

 Source: OpenStax *Micro* 3e, [§16.1](#) “The Problem of Imperfect Information and Asymmetric Information”, “‘Lemons’ and Other Examples of Imperfect Information”, where Marvin shops for a used car. The original paper is Akerlof (1970), “The Market for Lemons”.



Adverse Selection

Adverse selection: when hidden information causes the “worst” types to dominate the market.

The average price drives out the good ones, which lowers quality and price further still. The market can **collapse**.

Hold on to this idea: it is exactly the central problem of **insurance**. 

Moral Hazard

A different problem, arising **after** the contract is signed.

Moral hazard: when someone shielded from a consequence changes their behavior and takes on more risk.

With your car fully insured against theft, you might leave it unlocked more often. 

 Source: OpenStax *Micro* 3e, [§16.2](#), “The Moral Hazard Problem”; adverse selection follows immediately after it, in “The Adverse Selection Problem”.



Adverse Selection vs Moral Hazard

	When it occurs	The problem
Adverse selection	before the contract	you hide your type
Moral hazard	after the contract	you change your behavior

Both arise from asymmetric information, but at different moments.



Signaling

How do you solve hidden information? One side can **signal** its quality.

A degree signals competence. A warranty signals that the car is good.

The signal only works if it is **expensive to imitate** for whoever lacks the quality.

 Source: OpenStax *Micro* 3e, [§16.1](#), “Mechanisms to Reduce the Risk of Imperfect Information”: warranties, guarantees, reputation, references and the rest.

? Multiple Choice 3


An insurer offers one policy at a single price to everyone. Low risk clients leave, so the insurer raises the price, so more clients leave. This is:

- A. Moral hazard, since behavior changed after the contract.
 - B. Risk pooling working exactly as intended.
 - C. A signaling equilibrium.
 - D. Adverse selection, because the price cannot separate hidden types.
- ✓ **D.** A single price averages over types the insurer cannot observe. It is too expensive for the good risks and a bargain for the bad ones, so the pool deteriorates at every round. Insurers call it the **death spiral**.

? Multiple Choice 4

A risk averse agent is offered insurance at an **actuarially fair** premium (premium equals expected loss). They will buy:

- A. No insurance, since the premium equals the expected loss and there is no gain.
 - B. Full insurance, eliminating all the risk.
 - C. Partial cover, to balance premium against risk.
 - D. An amount depending on how large the loss is.
- ✓ **B.** At a fair price, buying cover has zero expected cost and strictly cuts variance, so a concave u takes all of it. Partial cover (C) is what happens once the premium carries a **loading** above fair, which is the real world case.



Numerical Question

At a dealership, half the cars are worth 10,000 euros (good) and half 4,000 euros (lemons). Buyers cannot tell them apart, and are risk neutral.

- a) What is the most a buyer will pay?
- b) What happens to the owners of good cars?
- c) What is the new market price once they have left?
- d) Name one institution that stops this happening in practice.

Solution

a) $0.5(10,000) + 0.5(4,000) = 7,000$ euros.

b) At 7,000 they will not sell: the car is worth more to them than the market offers. They **withdraw**.

c) Only lemons are left, so buyers work that out and pay 4,000. The good half of the market has **disappeared**, and nobody behaved irrationally at any step. 🍋

d) A warranty, an independent inspection, a dealer's reputation, a manufacturer certification. Each one is a **signal** that is costly to fake, which is why it works. 

Part 3 · Insurance



What Insurance Is

Insurance trades an **uncertain, large loss** for a **certain, small cost**: the premium.

You pay €300 a year so as not to risk a €20,000 loss in an accident.

Why are you willing to do this?

Because you are **risk averse**: you prefer certainty.



The Idea of Risk Pooling

The insurer brings many clients together. Individually, the risk is unpredictable.


But across **thousands** of policies, the fraction of claims is **stable** and predictable.

Risk pooling: bringing together many independent risks makes the total predictable.

 Source: OpenStax *Micro* 3e, [§16.2](#) “Insurance and Imperfect Information”, “How Insurance Works”, which opens with a table of who pays for which kind of policy and when it pays out.



The Law of Large Numbers



With more policies, the fraction of claims approaches the true value (10%). The red line is the true risk.

How the Premium Is Set

The premium covers the expected loss plus costs and margin.

If the probability of a claim is 2% and the loss is €20,000, the expected loss is $0.02 \times 20,000 = \text{€}400$.

The premium will sit a little above €400, to cover administrative costs and margin. 

 Source: OpenStax *Micro* 3e, [§16.2](#), “Risk Groups and Actuarial Fairness”, which is where the term **actuarially fair** comes from.



The Two Enemies of Insurance

The problems from Part 2 return, now from the insurer's side:



Adverse selection: whoever knows they are high risk is the one who most wants coverage.



Moral hazard: whoever is insured may become less careful.



How Insurers Respond



Screening: medical questionnaires, driving history, age of the car.



Deductibles: the insured pays part of the loss, which keeps them careful (this fights moral hazard).



Differentiated premiums: prices by risk profile, to fight adverse selection.



Why This Matters for You

Managing asymmetric information is the **heart** of the insurance and banking business.

Assessing credit risk, setting premiums, designing deductibles: all of that is applying what we saw today. 

? Multiple Choice 5


An insurer covers a region against flooding. Risk pooling works **poorly** here because:

- A. The losses are too large to insure at any price.
 - B. Policyholders will behave less carefully once covered.
 - C. The risks are strongly correlated, so the covariance term does not vanish.
 - D. The law of large numbers does not apply to natural events.
- ✓ C. Pooling needs risks that are close to independent. When one flood hits every policy at once, $\overline{Cov}$ stays large no matter how many policies are written. This is exactly the systematic risk term from Part 1, and it is why catastrophe cover is reinsured or state backed.

? Multiple Choice 6

A deductible, where the insured pays the first part of any claim, works mainly by:

- A. Leaving the insured exposed enough to stay careful.
 - B. Raising the insurer's revenue per policy.
 - C. Removing the need to price by risk profile.
 - D. Making the policy compulsory.
- ✓ **A.** It targets **moral hazard**, by keeping some of the loss with the person whose behavior affects it. Screening and risk-based pricing are the tools aimed at adverse selection instead.



Numerical Question

The probability of a claim is 5 percent and the loss is 30,000 euros.

- a) Compute the expected loss.
- b) The insurer loads 20 percent for costs and margin. What is the premium?
- c) A client with $u(w) = \sqrt{w}$ and wealth 100,000 euros faces this risk. What is the **most** they would pay for full cover?
- d) Do they buy the policy at the price in (b)?



Solution

a) $0.05 \times 30,000 = 1,500$ euros.

b) $1,500 \times 1.20 = 1,800$ euros. The 300 euros above fair is the **loading**.

c) $E[u] = 0.95\sqrt{100,000} + 0.05\sqrt{70,000} \approx 313.65$, so $CE \approx 313.65^2 \approx 98,373$. The most they would pay is $100,000 - CE \approx 1,627$ euros.

d) **No.** 1,800 is above their 1,627 ceiling. But at the **fair** premium of 1,500 they would buy, and buy in full, which is question 4 in numbers: the loading is what makes cover refusable. 

Wrap-Up

What to Take From This Session

 **Asymmetric information** generates adverse selection (before) and moral hazard (after the contract).

 **Insurance** trades an uncertain loss for a certain cost, resting on **risk pooling**.

 Insurers fight asymmetric information with screening, deductibles and differentiated premiums.

End of Block 1

Next comes the **Midterm Test**, and after that we move into **macroeconomics**: GDP, inflation, money and economic policy.



Good luck with the studying. 