Asymmetric Information

Lecture 8

Div Bhagia

Introduction

Asymmetric Information

So far, we have largely assumed that all parties to a transaction share the same information. In reality, one side often knows more than the other.

- A used-car seller knows whether the car is a lemon; the buyer does not
- A worker knows how hard they are working; the employer cannot perfectly observe effort
- An insurance applicant knows their own risk type better than the insurer

Key question: Are markets still efficient when there is asymmetry in how much parties know?

Two Types of Asymmetric Information

1. **Adverse Selection (hidden types)**
    - Individuals have private information about their characteristics ("types") *before* entering into a contract.
    - Example: An insurance applicant knows their own risk level better than the insurer.
2. **Moral Hazard (hidden actions)**
    - Individuals take actions that are not observable by others *after* entering into a contract.
    - Example: A worker exerts effort that is not perfectly observable by the employer.

Adverse Selection: Market for Lemons

Akerlof's "Market for Lemons"

Akerlof (1970) studied the market for used cars, where sellers know the quality of their car (good or lemon), but buyers do not. He showed that this information asymmetry can lead to market failure.

Simplified Illustration

- Two types of *used* cars at the dealership:
  - Good cars (G)
  - Lemons (L) that breakdown often
- Buyers cannot tell apart good cars from lemons, but know that half of the cars are lemons.

Equilibrium

- Sellers' and buyers' valuations of the cars is as follows:

Type	Seller's Value	Buyer's Value
Good (G)	\$16,000	\$20,000
Lemon (L)	\$8,000	\$10,000

Since buyers cannot tell apart good cars from lemons, they will only be willing to pay a price equal to the expected value of the car:

$$0.5 \cdot 10,000 + 0.5 \cdot 20,000 = \$15,000$$

But at that price, only sellers of lemons will be willing to sell. The market for good cars collapses and price falls to \$10,000.

Why This Happens

Generally with a continuum of quality, the same logic applies:

1. Buyers can't observe quality → they offer an **average** price
2. Average price is too low for the best sellers → they **exit**
3. Average quality falls → buyers **lower** their price
4. More sellers exit → quality falls further
5. **Unraveling** until only the worst remain

Akerlof's key insight: **quality is endogenous to price**. The composition of goods on the market changes when you change the price.

Private Market Solutions

Signaling: The informed party takes a costly action to reveal type.

- Warranties and guarantees
- The signal must be **costly to fake**, otherwise it's meaningless

Reputation and repeated interaction:

- Brand names and chains
- Rating systems (eBay, Carfax, Yelp)
- Relationship: you buy from the same dealer

Certification and screening:

- Third-party inspections (JD Power, Certified Pre-Owned)
- Diplomas and professional credentials

Non-Market Solutions

Government certification: FDA, USDA, occupational licensing

Legal liability: Lemon laws i.e. seller must disclose known defects

Mandated information provision: Disclosure requirements

All of these work by **reducing the information asymmetry** and making it harder for lemons to hide among good cars.

Adverse Selection in Health Insurance

Insurance Setup

Adapted from David Autor's lecture notes.

- Continuum of consumers $i \in [0, 1]$
- Each has utility $U(w) = \ln(w)$ (risk-averse)
- Initial wealth: $w_0 = 150$
- Probability of loss: 50%
- Loss amount: $L_i = 100 \times i$

Consumer $i = 0$ has zero risk. Consumer $i = 1$ faces a potential loss of $100.

Each consumer knows their own type i. Insurers cannot tell.

Consumer $i = 0.60$

- Expected wealth: $150 - 0.5 \times 60 = \$120$
- Expected utility: $0.5\ln(150) + 0.5\ln(90) = 4.76$
- Certainty equivalent: $e^{4.76} = \$116.19$

Note: Certainty equivalent is the amount of wealth that gives the same utility as the risky prospect.

So consumer $i = 0.60$ is willing to pay a **$3.81 risk premium** ($120 − $116.19) for insurance.

In total, the consumer is **willing to pay $33.81** for insurance that covers the $30 expected loss (the $3.81 risk premium plus the $30 actuarially fair cost).

Willingness to Pay for Insurance

Expected wealth for i:

$$E[w_i] = 0.5 \cdot (150) + 0.5 \cdot (150 - 100i) = 150 - 50i$$

Expected utility for i:

$$E[U(w_i)] = 0.5 \ln(150) + 0.5 \ln(150 - 100i)$$

Certainty equivalent and risk premium for i:

$$CE_i = e^{E[U(w_i)]}, \qquad RP_i = E[w_i] - CE_i$$

Consumer i's willingness to pay for insurance:

$$WTP_i = \underbrace{w_0 - E[w_i]}_{\text{Expected Loss}} + RP_i = w_0 - CE_i$$

Willingness to Pay vs. Risk Type

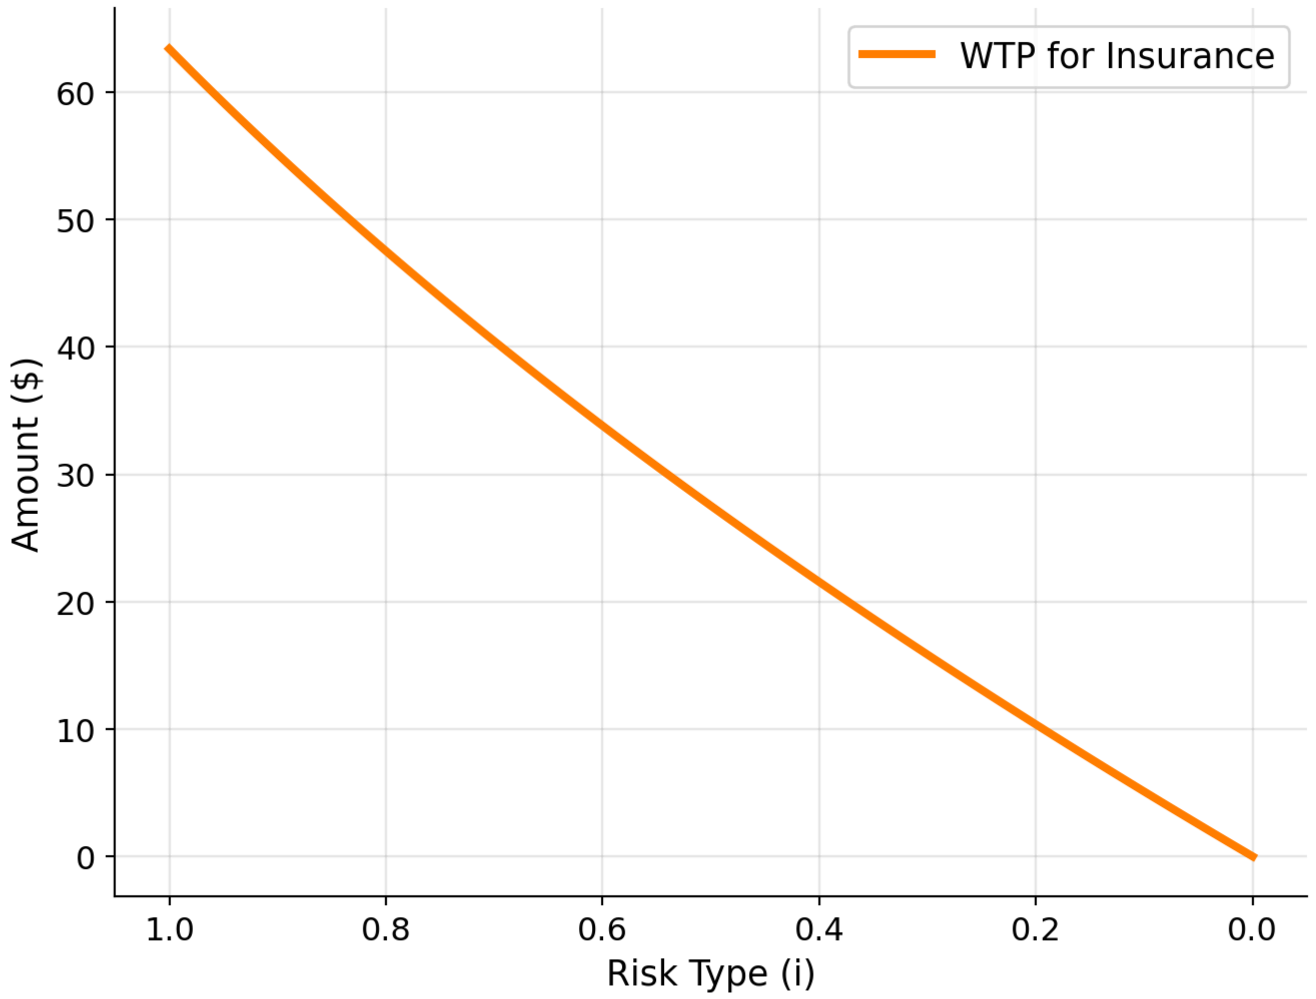

The Naive Policy

An insurer offers **full insurance** at the **population-average** expected loss:

$$\text{Premium} = 0.5 \times E[L_i] = \$25$$

Note that $L_i \sim U(0, 100)$, so $E[L_i] = 50$, and each consumer faces a 50% probability of loss.

Which consumers will buy at this price?

The consumer i_0 who is indifferent solves:

$$0.5 \ln(150) + 0.5 \ln(150 - 100 i_0) = \ln(125)$$

$$i_0 = 0.46$$

Only consumers with $i \geq 0.46$ buy i.e. the **sicker** 54%.

Who buys insurance at $25?

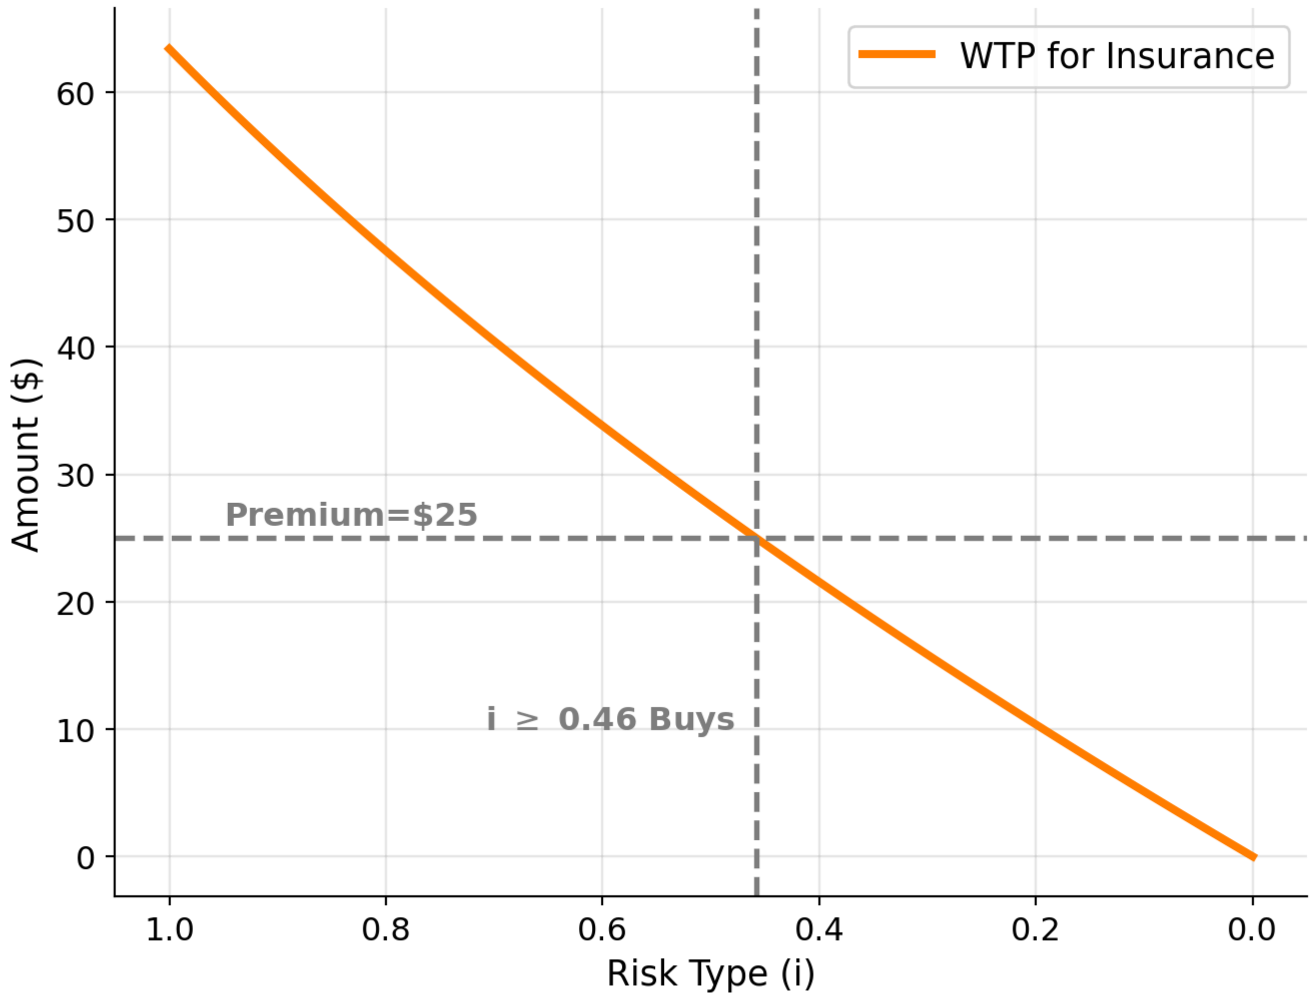

The Naive Policy Loses Money

Expected cost per insured:

$$0.5 \times E[L_i \mid i \geq 0.46] = 0.5 \times 100 \times \frac{1 + 0.46}{2} = \$36.50$$

Premium collected: \$25. Cost per insured: \$36.50.

The insurer loses \$11.50 per policy.

This is adverse selection at work: at the average price, only the high-risk consumers buy. The healthy ones opt out because the premium far exceeds their expected loss.

What should be insurer's optimal policy?

Insurer should set the premium high enough to break even given the composition of buyers it attracts.

The Break-Even Policy

Let i_0 be the cutoff type who enrolls. Then break-even requires that the premium P equals the expected cost of insuring those types:

$$P = 0.5 \times E[L_i \mid i \geq i_0] = 0.5 \times 100 \times \frac{1 + i_0}{2}$$

Since consumer i_0 is indifferent between buying and not buying, we have:

$$0.5 \ln(150) + 0.5 \ln(150 - 100 i_0) = \ln(150 - P)$$

Solution: $i_0 = 0.75$

Premium: $0.5 \times (100 + 75)/2 = \$43.75$

The Break-Even Policy

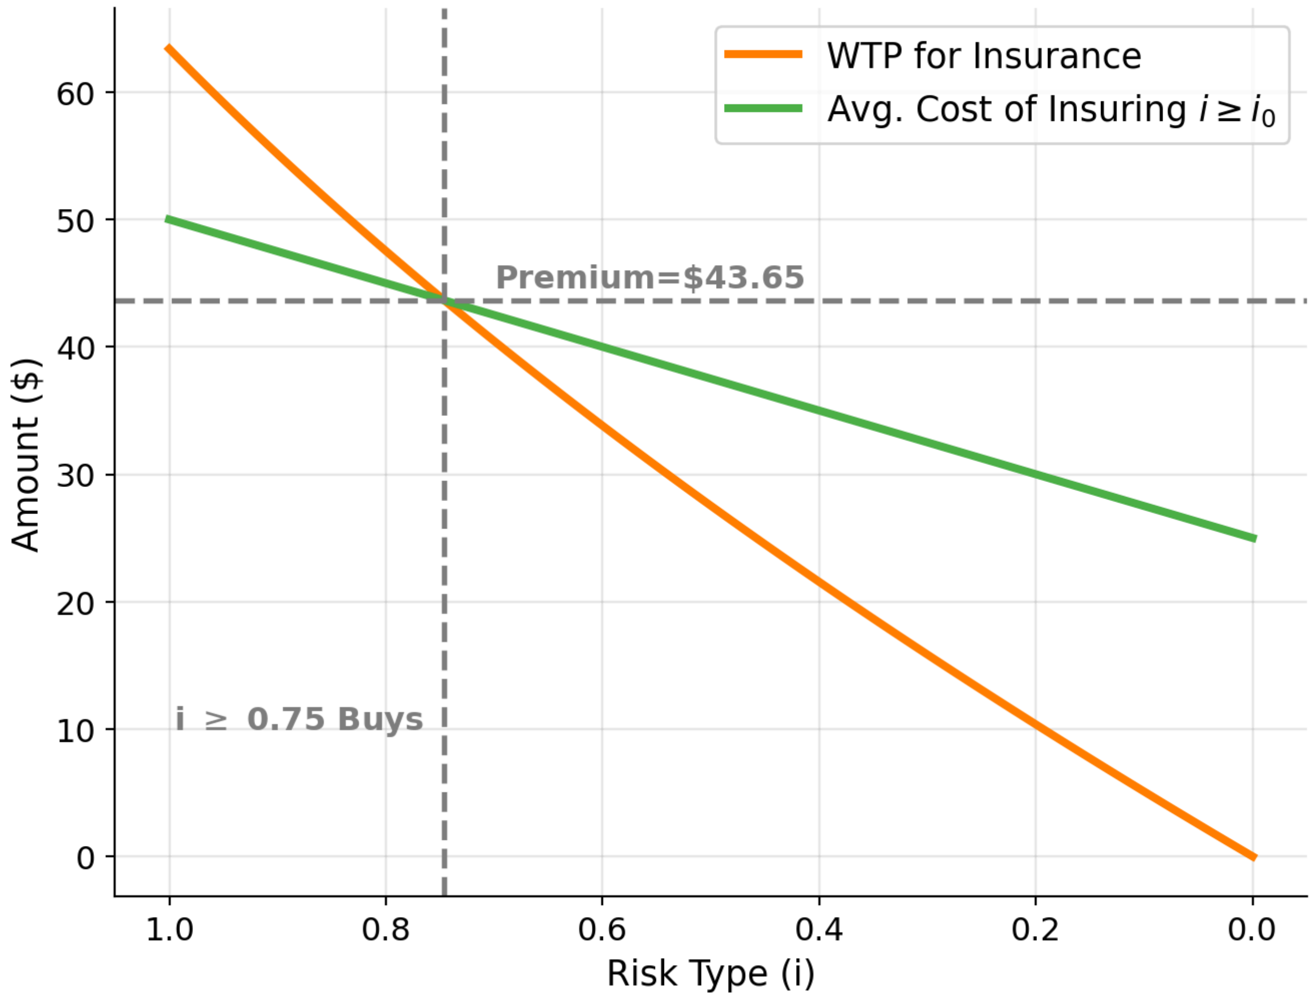

Is the market functioning well?

Unlike the used car market, the insurance market doesn't collapse entirely.

The reason is risk aversion. The sickest consumers are willing to pay *more* than their actuarially fair cost because they prefer a bad deal on insurance to no insurance at all.

But the outcome is still highly inefficient. The high-risk types impose a **negative externality** on low-risk types by driving up premiums. Most consumers who would benefit from insurance don't get it.

Adverse Selection on a Plot

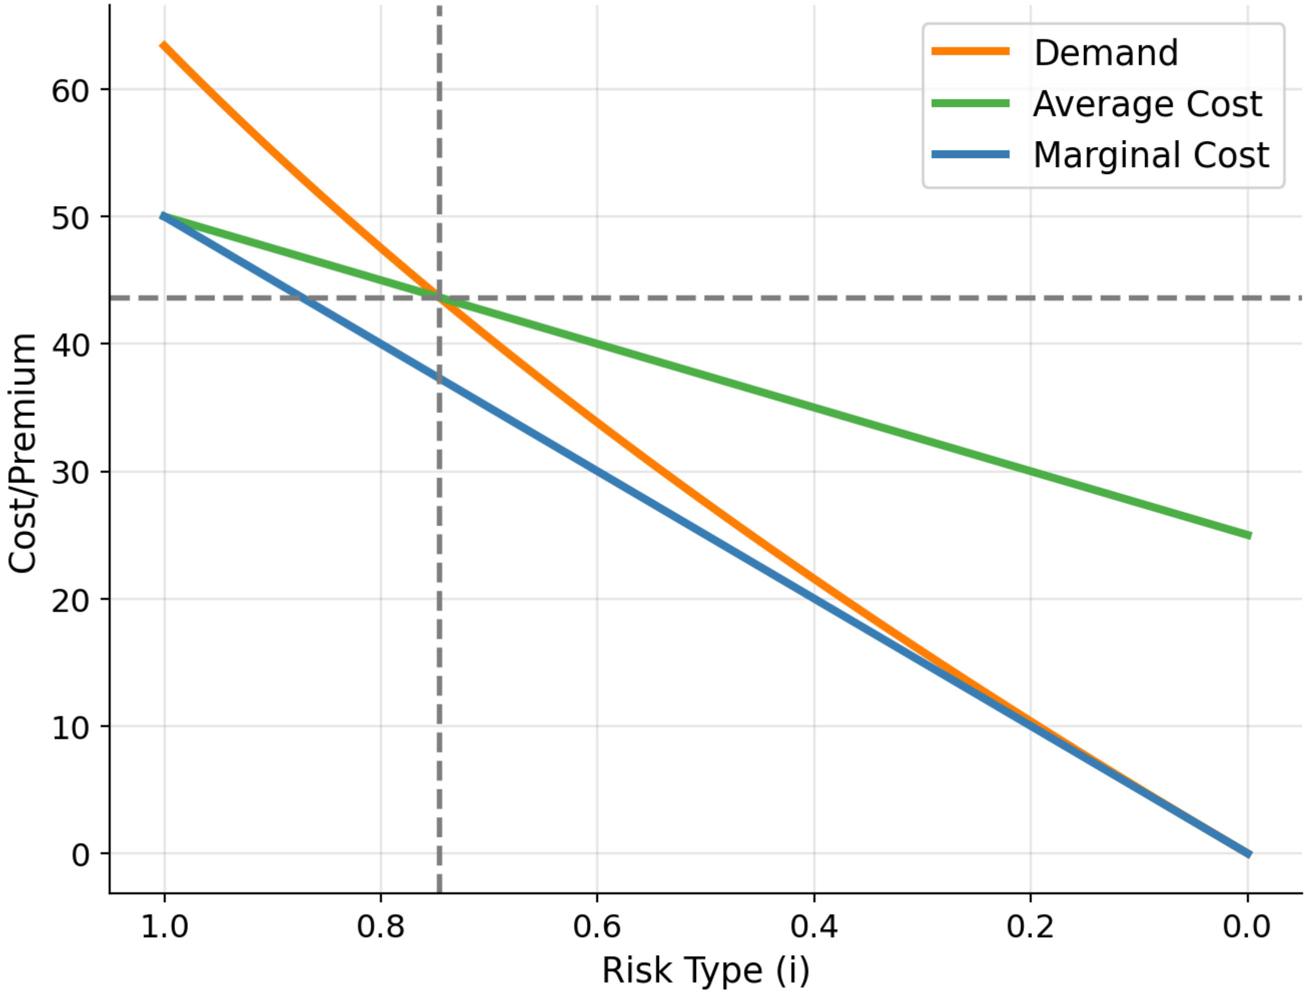

Quoting Einav & Finkelstein (2011)

> “The fundamental inefficiency created by adverse selection arises because the efficient allocation is determined by the relationship between marginal cost and demand, but the equilibrium allocation is determined by the relationship between average cost and demand. Because of adverse selection (downward sloping MC curve), the marginal buyer is always associated with a lower expected cost than that of infra-marginal buyers.”

Policy Options: Mandatory Insurance

What if everyone must buy at $25?

- Marginal cost of insuring each consumer ≤ their willingness to pay (for all $i > 0$)
- Everyone is insured. Premium = population-average cost. The policy **breaks even**.

Not everyone is better off individually as consumers with $i < 0.46$ would prefer no insurance at this price. They are **subsidizing** sicker consumers.

But from a social welfare standpoint, the mandate provides both **risk pooling** and **income spreading** and average welfare is higher than in the free market case.

Policy Options: Free Screening

Suppose a free test reveals each consumer's type i. Insurers then offer actuarially fair individual premiums: $50 \times i$.

Full disclosure principle: Everyone volunteers for the test.

Why? The healthiest half discloses first. Then the healthiest half of the remainder. Then the next. "Turtles all the way down" and eventually everyone is tested.

No more adverse selection. Everyone is insured at an individualized fair price.

Which Is Better?

Surprisingly: **the mandate beats free screening** on average welfare.

Policy	Insurance?	Redistribution?	Avg. welfare
No insurance	No	No	Lowest
Free market (break-even)	Partial (25%)	No	Low
Individual pricing (screening)	Full	No	Medium
Mandatory pool ($25)	Full	Yes	**Highest**

The mandatory policy does two things: it provides **insurance** and it **transfers** from low-risk to high-risk. Under concave utility (diminishing marginal utility of wealth), that transfer raises average welfare.

Moral Hazard

The Basic Idea

You behave differently when you don't bear the full cost of your actions.

- Fully insured car → leave it unlocked, drive less carefully
- Health insurance → visit the doctor more, worry less about prevention
- Bank bailout guarantee → take on riskier investments

The insurer can observe the **outcome** (car stolen, got sick, bank failed) but not the **action** (did you lock the car? exercise? manage risk?).

This is the **hidden action** problem.

Moral Hazard in UI

A worker loses their job and qualifies for unemployment insurance (UI). The government offers unemployment benefits b per period.

Without UI: The worker bears the full cost of unemployment. Strong incentive to search hard and accept offers quickly.

With UI: The cost of remaining unemployed falls from "no income" to "b per period." At the margin, the worker:

- Searches less intensely
- Is pickier about which offers to accept
- Stays unemployed longer

The Hidden Action

The government can observe:

- Whether the worker is **employed or not**

The government **cannot** observe:

- How many applications the worker sends
- How hard they prepare for interviews
- Whether they turned down a reasonable offer
- Whether they're really "looking"

Search effort is the hidden action. If the government could observe it, they'd simply require a minimum effort level. But they can't so benefits distort incentives.

A Simple Model

Two periods. A worker is unemployed in period 1 and chooses search effort $e \in [0, 1]$.

- With probability e: finds a job, earns wage w in period 2
- With probability $1 - e$: stays unemployed in period 2
- Cost of search effort: $c(e) = \frac{1}{2}e^2$

The government pays benefit b while unemployed (both periods). If employed in period 2, benefits stop.

Worker's consumption

In this model, the worker's consumption in each state is given by:

Period	State	Consumption
Period 1	Unemployed	b
Period 2	Employed (prob e)	w
Period 2	Unemployed (prob $1 - e$)	b

Worker's Problem

The worker chooses e to maximize expected utility minus effort cost:

$$\max_{e} \quad u(b) + e \cdot u(w) + (1 - e) \cdot u(b) - \frac{1}{2}e^2$$

First-order condition:

$$u(w) - u(b) = e^*$$

The worker searches until the **marginal cost of effort** (e) equals the **marginal benefit** i.e. the utility gain from being employed vs. staying on UI.

How Benefits Affect Effort

$$e^* = u(w) - u(b)$$

When b increases:

- $u(b)$ rises: unemployment is **less painful**
- $u(w) - u(b)$ **shrinks**: the gain from finding a job falls
- So e^* **falls**: the worker searches less

In the extreme:

- $b = 0$: maximum search effort, $e^* = u(w) - u(0)$
- $b = w$: zero search effort, $e^* = 0$

The Tradeoff

More generous UI ($\uparrow b$) does two things:

Insurance value ↑

- Smooths consumption during a bad shock
- Protects workers from income loss they didn't choose
- Especially valuable if workers are risk-averse

Moral hazard cost ↑

- Reduces search effort
- Workers stay unemployed longer
- Output is lost, and the government pays more in benefits

The optimal b **balances** these two forces. Neither $b = 0$ (no insurance) nor $b = w$ (full insurance) is optimal.

The Government's Problem

The government chooses b to maximize the worker's expected welfare, subject to a budget constraint:

$$\max_{b} \quad u(b) + e^*(b) \cdot u(w) + (1 - e^*(b)) \cdot u(b) - \frac{1}{2}(e^*(b))^2$$

subject to:

- **Incentive compatibility:** $e^* = u(w) - u(b)$ (the worker chooses effort optimally given b)
- **Budget constraint:** benefits must be funded (e.g., through taxes on wages)

The government **cannot** choose e directly. It can only choose b, and e responds.

First Best vs. Second Best

First best (effort observable): The government picks b and e directly.

- Full insurance: $b = w$ (perfect consumption smoothing)
- Mandate high search effort
- Worker bears no risk and works hard

Second best (effort hidden): The government picks b, worker picks e.

- $b < w$ – **incomplete insurance** to preserve search incentives
- Worker bears some risk as the **price** of maintaining incentives
- Fundamental tradeoff: insurance vs. incentives

The gap between first best and second best is the **cost of moral hazard**.

Moral Hazard Beyond UI

The same structure appears whenever insurance or protection changes behavior:

Health insurance: Insured patients consume more health care. Copays and deductibles are the "incomplete insurance" that preserves incentives as you bear some cost at the margin.

Banking: Deposit insurance and "too big to fail" protect depositors and the system. but banks take on more risk knowing they'll be bailed out. Capital requirements are the policy response.

Employment contracts: The Gibbons principal-agent model. The worker's effort is hidden. Steep incentive pay (high b in $w = s + b \cdot y$) is the response, but imposes risk on the worker.

Key Takeaways

1. **Moral hazard** arises when insurance changes behavior because the insured party doesn't bear the full cost of their actions
2. The hidden action (effort, care, risk-taking) cannot be contracted on: only outcomes are observable
3. The fundamental tradeoff is **insurance vs. incentives**: full insurance eliminates risk but also eliminates effort
4. The optimal policy provides **incomplete insurance**: enough to protect against bad outcomes, not so much that it kills incentives