Lecture 8: Asymmetric Information

Econ 502: Advanced Microeconomics

Overview

The First Welfare Theorem tells us that competitive equilibria are Pareto efficient, but only under demanding conditions, including that all parties share the same information. In reality, one side of a transaction often knows more than the other.

This lecture studies two fundamental forms of **asymmetric information**:

1. **Adverse selection (hidden types):** One party has private information about a fixed characteristic *before* the transaction. The uninformed party, unable to distinguish types, must deal with the average, which drives out the best types and can cause market failure.

2. **Moral hazard (hidden actions):** After a contract is in place, one party takes an action that the other cannot observe. Because the action is unobservable, incentives are distorted relative to the efficient outcome.

Both problems are studied through specific, tractable models. Understanding them is essential for health policy, financial regulation, labor economics, and public economics more broadly.

Part I: Adverse Selection and the Market for Lemons

1.1 Akerlof's "Market for Lemons"

George Akerlof's 1970 paper, "The Market for Lemons," studied used-car markets where sellers know the quality of their car but buyers do not. He showed that this information asymmetry can cause the high-quality segment of the market to collapse entirely.

The simplified setup. There are two types of used cars:

- **Good cars (G):** worth \$16,000 to sellers and \$20,000 to buyers
- **Lemons (L):** worth \$8,000 to sellers and \$10,000 to buyers

Half of the cars on the lot are good cars and half are lemons. Buyers cannot distinguish quality before purchase, but they know this 50/50 distribution.

With symmetric information (buyers can observe quality), both types trade efficiently: good cars at some price between \$16,000 and \$20,000, lemons at some price between \$8,000 and \$10,000. Total gains from trade are realized.

With asymmetric information, buyers must form an expectation about quality. Since they cannot tell types apart, the most they are willing to pay for any car is the *expected* value:

$$\bar{V}^B = 0.5 \times \$20{,}000 + 0.5 \times \$10{,}000 = \$15{,}000$$

Now ask: will sellers of good cars sell at \$15,000? A good car is worth \$16,000 to its seller. Since \$15,000 < \$16,000, the answer is **no**: good-car sellers will not part with their cars at the going price. Only lemon sellers (who value their cars at \$8,000) are willing to sell at \$15,000.

But if buyers anticipate that only lemons are on the market, they revise their valuation downward to \$10,000, and the price falls accordingly. The market for good cars **collapses**, and only lemons trade.

1.2 The Unraveling Logic

The used-car example captures the beginning of a general unraveling process. When quality is continuous rather than binary, the logic applies step by step:

1. Buyers cannot observe quality → they offer a price based on **average quality**
2. The average price is below the value of the best sellers → the **best types exit**
3. With the best types gone, average quality falls → buyers **lower their price**
4. More sellers exit → quality falls further
5. This process continues until **only the worst types remain** (or the market collapses entirely)

Akerlof's key insight. Quality is *endogenous to price*. Unlike standard supply and demand, where the composition of sellers is fixed and price simply clears the market, here the price itself determines who is willing to sell. A higher price draws in better quality; a lower price drives quality down. This feedback is the source of the adverse selection problem.

The welfare cost. In the symmetric information world, all cars trade (good cars, lemons, and everything in between) at prices reflecting their true quality. Under asymmetric information, high-quality goods may not trade at all even though buyers value them more than sellers do. The surplus that would have been generated by those trades is lost. This is a deadweight loss created purely by the information asymmetry.

1.3 Solutions to Adverse Selection

Since the problem arises from asymmetric information, solutions work by reducing or eliminating the asymmetry.

Signaling. The *informed party* takes a costly, observable action to credibly reveal their type. For a signal to work, it must be **costly to fake**: high-quality sellers must find it worth taking the action, while low-quality sellers must not.

- *Warranties and guarantees:* A seller of a good car is willing to offer a full warranty because the car rarely needs repairs. A seller of a lemon cannot afford to offer the same warranty. The warranty credibly reveals quality.
- *The signal must be costly to fake:* If a lemon seller could offer a warranty as cheaply as a good-car seller, the warranty would convey no information. The signal works only when its cost is inversely related to quality.

Reputation and repeated interaction. In markets where the same parties transact repeatedly, sellers have an incentive to maintain a reputation for quality: the value of future business is at stake.

- Brand names and franchise systems
- Rating platforms (eBay, Carfax, Yelp, Airbnb)
- Buying from the same dealer over many transactions

Certification and screening. A *third party* observes quality and certifies it to buyers who cannot.

- Third-party inspections (JD Power, Certified Pre-Owned programs)
- Diplomas and professional credentials signal worker quality to employers

Non-market (government) solutions. When private solutions are insufficient:

- *Government certification:* FDA approval, USDA organic labeling, occupational licensing. The government acts as a certifier, replacing the private information asymmetry with a public standard.
- *Legal liability (lemon laws):* Sellers must disclose known defects, on pain of legal penalties. This raises the cost of concealing bad quality.
- *Mandated information provision:* Disclosure requirements (nutritional labels, car safety ratings, financial prospectuses) force the informed party to share information.

All of these work by the same mechanism: reducing the information asymmetry and making it harder for low-quality sellers to hide among high-quality ones.

Part II: Adverse Selection in Health Insurance: A Quantitative Model

The used-car example shows *why* adverse selection causes market failure. The following model, adapted from David Autor's lecture notes, shows *how much* damage it does and evaluates policy responses.

2.1 Setup

Consumers. There is a continuum of consumers indexed by $i \in [0, 1]$. Consumer i:

- Has utility $U(w) = \ln(w)$, which is strictly concave (so consumers are risk-averse)
- Has initial wealth $w_0 = 150$
- Faces a 50% probability of a health loss of size $L_i = 100i$

The type index i captures risk: $i = 0$ is perfectly healthy (zero potential loss), while $i = 1$ faces a potential loss of \$100. Types are uniformly distributed on $[0, 1]$.

Information. Each consumer knows their own type i. Insurers observe only that types are drawn from $U[0, 1]$; they cannot identify any individual's i.

2.2 Willingness to Pay for Insurance: Step-by-Step Derivation

Before studying market outcomes, we need to know how much each consumer type is willing to pay for full insurance (a policy that compensates for the entire loss).

Step 1: Expected wealth without insurance.

Consumer i has wealth $w_0 = 150$ if healthy (probability 0.5) and $w_0 - L_i = 150 - 100i$ if sick (probability 0.5):

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

Step 2: Expected utility without insurance.

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

Step 3: Certainty equivalent.

The *certainty equivalent* CE_i is the guaranteed wealth level that gives the same utility as the risky prospect. It solves:

$$\ln(CE_i) = E[U(w_i)]$$

$$CE_i = e^{E[U(w_i)]} = e^{0.5\ln(150)+0.5\ln(150-100i)} = \sqrt{150 \cdot (150-100i)}$$

The certainty equivalent is below expected wealth, $CE_i < E[w_i]$, because ln is concave. This gap is the **risk premium** $RP_i = E[w_i] - CE_i > 0$.

Step 4: Willingness to pay.

With full insurance at premium P, the consumer receives $w_0 - P$ with certainty. The consumer buys if:

$$\ln(w_0 - P) \geq E[U(w_i)]$$

The maximum premium they are willing to pay solves this with equality:

$$w_0 - WTP_i = CE_i \implies WTP_i = w_0 - CE_i$$

Expanding:

$$\boxed{WTP_i = 150 - \sqrt{150(150-100i)}}$$

Note that $WTP_i = (w_0 - E[w_i]) + RP_i = 50i + RP_i$. The WTP has two components:

- The **actuarially fair premium** $50i$: exactly what the consumer expects to lose, so a risk-neutral person would pay no more
- The **risk premium** $RP_i > 0$: the extra amount the risk-averse consumer is willing to pay on top of the actuarially fair price to eliminate risk

Worked example: Consumer $i = 0.60$.

- Expected wealth: $150 - 50 \times 0.60 = \$120$
- Expected utility: $0.5\ln(150) + 0.5\ln(90) = 0.5(5.011) + 0.5(4.500) = 4.756$
- Certainty equivalent: $e^{4.756} = \$116.19$
- Risk premium: $\$120 - \$116.19 = \$3.81$
- Actuarially fair premium: $0.5 \times 60 = \$30$
- **WTP:** $\$30 + \$3.81 = \$33.81$, or equivalently $WTP = 150 - 116.19 = \$33.81$

The key observation is that WTP *exceeds* the actuarially fair cost for every consumer $i > 0$. Every risk-averse consumer is willing to pay a premium above the expected loss to get insurance.

2.3 The Adverse Selection Problem: The Naive Policy

An insurer who does not know individual types might naively offer insurance at the **population-average** expected loss:

$$P^{\text{naive}} = 0.5 \times E[L_i] = 0.5 \times E[100i] = 0.5 \times 50 = \$25$$

where $E[i] = 0.5$ since $i \sim U[0,1]$.

Which consumers buy at $P = \$25$?

Consumer i buys if $WTP_i \geq \$25$, i.e., if $150 - CE_i \geq 25$, i.e., if $CE_i \leq 125$.

The cutoff type i_0 is indifferent: their expected utility without insurance equals the utility of paying \$25 and having \$125 for certain:

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

Taking exponentials of both sides:

$$\sqrt{150(150 - 100i_0)} = 125$$

$$150(150 - 100i_0) = 15{,}625$$

$$150 - 100i_0 = 104.17$$

$$i_0 \approx 0.46$$

Only consumers with $i \geq 0.46$ (the sicker 54%) buy insurance at \$25. The healthiest 46% opt out because the \$25 premium exceeds their WTP.

Does the insurer break even?

The insurer's cost per policy is the expected payout to the insured pool:

$$\text{Cost} = 0.5 \times E[L_i \mid i \geq 0.46] = 0.5 \times 100 \times E[i \mid i \geq 0.46]$$

Since $i \mid i \geq 0.46$ is uniform on $[0.46, 1]$, its mean is $(0.46 + 1)/2 = 0.73$:

$$\text{Cost} = 50 \times 0.73 = \$36.50$$

The insurer collects \$25 but pays \$36.50 on average, a **loss of \$11.50 per policy**. The naive policy is unsustainable.

Why? The insurer priced the policy for the average consumer, but only the sicker-than-average consumers found it worthwhile to buy. This is adverse selection: the insured pool is "adversely selected" toward high-cost types, relative to the general population.

2.4 The Break-Even Policy

A profit-maximizing insurer must raise the premium until it breaks even on its actual customer pool. Let i_0 be the cutoff type (who is indifferent between buying and not buying) and P be the premium. Two conditions must hold simultaneously:

Break-even condition (insurer covers its costs):

$$P = 0.5 \times E[L_i \mid i \geq i_0] = 0.5 \times 100 \times \frac{1 + i_0}{2} = 25 \cdot \frac{1 + i_0}{2}$$

Indifference condition (the cutoff consumer is just willing to buy):

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

which is equivalent to $CE_{i_0} = 150 - P$, or $P = 150 - CE_{i_0} = WTP_{i_0}$.

Substituting the break-even premium into the indifference condition and solving numerically:

$$i_0 = 0.75, \qquad P = 25 \times \frac{1 + 0.75}{2} = 25 \times 0.875 = \$43.75$$

Verification. At $i_0 = 0.75$:

$$CE_{0.75} = \sqrt{150(150 - 75)} = \sqrt{150 \times 75} = \sqrt{11{,}250} \approx 106.07$$

$$WTP_{0.75} = 150 - 106.07 \approx \$43.93 \approx \$43.75\checkmark$$

In the break-even equilibrium, only the sickest 25% of consumers ($i \geq 0.75$) are insured. The other 75% (who would genuinely benefit from insurance at their actuarially fair price) are priced out.

2.5 Why the Insurance Market Doesn't Collapse (Unlike Used Cars)

Unlike the used-car market, the health insurance market doesn't necessarily collapse to zero coverage. Why?

Risk aversion saves the market, partially. The sickest consumers are willing to pay *more* than their actuarially fair cost precisely because they are risk-averse. Even if the premium is very high (well above the expected loss), consumers with severe illness risk prefer insurance to no insurance. This keeps at least the high-risk segment of the market alive.

But the outcome is still highly inefficient. High-risk types impose a negative externality on low-risk types: because the insured pool is dominated by expensive cases, the premium must be high, which prices out the low-risk consumers who might have bought at a lower premium. Most consumers who *could* benefit from insurance end up uninsured.

2.6 Adverse Selection Through the Demand/Supply Lens

The market for insurance has an unusual feature: the demand curve and the cost curves all run in the same direction (the diagram uses risk type i on the horizontal axis, with high-risk consumers on the right).

- **Demand (WTP) curve:** Downward sloping from right to left: sicker consumers (i near 1) are willing to pay more.
- **Marginal cost (MC):** Also downward sloping: the cost of insuring consumer i is $50i$, which is higher for sicker consumers.
- **Average cost (AC):** The cost of insuring *all consumers with type* $\geq i_0$, which lies above MC because the insured pool always includes some high-cost types above i_0.

The competitive equilibrium is where **demand = average cost**, not where **demand = marginal cost**. This is the source of the inefficiency: Einav and Finkelstein (2011) summarize it precisely:

> "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."

At the efficient allocation, every consumer for whom WTP exceeds marginal cost ($50i$) should be insured. At the market equilibrium, only consumers for whom WTP exceeds average cost are insured. Since AC > MC (due to adverse selection), the equilibrium pool is smaller than the efficient pool: too few consumers are insured.

2.7 Policy Responses

Option 1: Mandatory insurance at $25.

The government requires *everyone* to buy insurance at the average-cost premium of $25.

- The insurer breaks even: all consumers are pooled, and the premium exactly equals the average expected loss.
- Every consumer is insured. For all $i > 0$, the actuarially fair premium $50i < WTP_i$, so every consumer has positive net benefit from insurance *at their fair price*. But at $25, consumers with $i < 0.46$ are paying above their fair price — they are **cross-subsidizing** the sicker consumers.
- From a social welfare standpoint, the mandate provides two benefits:
  - **Risk pooling:** All consumers are insured, eliminating wealth risk for all types.
  - **Income spreading (redistribution):** Low-risk consumers transfer to high-risk consumers. Under concave utility (diminishing marginal utility of wealth), this transfer raises aggregate welfare: a dollar transferred from the low-risk (who would have healthy wealth even without insurance) to the high-risk (who would face large wealth losses when sick) generates a net welfare gain.

Option 2: Free screening.

A free genetic test or health examination reveals each consumer's type i to insurers. Insurers then offer actuarially fair individual premiums $P_i = 50i$.

- Every consumer is fully insured at their fair price. Adverse selection is completely eliminated.
- **Why would everyone take the test?** The "full disclosure principle" explains: the unraveling logic works in reverse. The healthiest consumers benefit most from revealing their type (they get a low premium). As they reveal, the pool of untested consumers skews sicker, raising the premium for the untested. The next-healthiest group now also benefits from disclosure. This continues until everyone has tested, a complete "unraveling" toward full disclosure.

Which policy is better?

Policy	Insured?	Redistribution?	Avg. welfare
No insurance	No	No	Lowest
Free market (break-even)	25%	No	Low
Individual pricing (screening)	100%	No	Medium
Mandatory pool ($25)	100%	Yes	**Highest**

Surprisingly, the **mandatory pool dominates screening** on average welfare, even though screening eliminates adverse selection. The reason: the mandate both insures everyone *and* redistributes from low-risk to high-risk. Because utility is concave, this redistribution is welfare-enhancing in aggregate: the marginal utility of wealth is higher for the sick (who would otherwise face large losses) than for the healthy. The dollar-for-dollar transfer raises total utility.

Screening achieves the insurance benefit but foregoes redistribution. The mandated pool achieves both.

Part III: Moral Hazard

3.1 The Basic Idea

Adverse selection is a problem of *hidden types*: characteristics that are fixed before the contract. Moral hazard is a problem of *hidden actions*: choices the insured party makes *after* the contract is in place.

Core idea: When you do not bear the full cost of a bad outcome, you have less incentive to prevent it.

- *Fully insured car:* You are less likely to lock it, less likely to park carefully, more willing to park in unsafe areas.
- *Health insurance:* Insured patients visit the doctor more frequently, invest less in preventive care, and choose riskier activities.
- *Bank bailout guarantee:* A bank that knows it will be rescued takes on riskier investments. If the bet pays off, it keeps the profit. If it fails, taxpayers bear the loss.

In each case, the insurer (or government) can observe the **outcome** (car stolen, patient got sick, bank failed) but not the **action** (did you lock the car? exercise? manage risk carefully?). This is the **hidden action** problem.

3.2 Moral Hazard in Unemployment Insurance

The setting. A worker loses their job and qualifies for unemployment insurance (UI), which pays benefit b per period while unemployed. The government cannot observe whether the worker is searching hard or sitting at home. This is the key informational asymmetry.

Without UI: The worker bears the full cost of unemployment. Lost income is the penalty for not finding a job quickly. The worker has a strong incentive to search intensively and accept reasonable offers.

With UI: The cost of remaining unemployed falls from "no income" to "b per period." At the margin, the worker searches less intensively, is pickier about which offers to accept, and stays unemployed longer.

The government observes only whether the worker is employed or not. It **cannot observe**:

- How many applications the worker sent
- How hard the worker prepared for interviews
- Whether the worker turned down a reasonable offer
- Whether the worker is genuinely "looking" for work

Search effort is the hidden action. If effort were observable, the government could simply write a contract: "receive b but search at intensity $\bar{e}$." Because effort is hidden, that contract is unenforceable, and benefits distort incentives.

3.3 A Two-Period Model

Setup.

- Period 1: Worker is unemployed, receives benefit b

- Worker chooses search effort $e \in [0, 1]$, at cost $c(e) = \frac{1}{2}e^2$
- Period 2: With probability e, finds a job earning wage w; with probability $1 - e$, remains unemployed and receives b
- The worker has utility function $u(\cdot)$ with $u' > 0$, $u'' < 0$

Worker's consumption in each state:

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

Worker's expected utility.

The worker maximizes expected utility minus effort cost:

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

Collecting terms:

$$EU = (2 - e)\, u(b) + e\, u(w) - \frac{1}{2}e^2$$

First-order condition. Differentiating with respect to e and setting equal to zero:

$$\frac{\partial EU}{\partial e} = -u(b) + u(w) - e = 0$$

$$\boxed{e^* = u(w) - u(b)}$$

Interpretation. The worker searches until the marginal cost of effort (e, i.e., the e-th unit costs e in utility terms given the quadratic cost function) equals the marginal benefit: the utility gain from being employed ($u(w)$) rather than remaining unemployed ($u(b)$).

3.4 How Benefits Affect Effort

From $e^* = u(w) - u(b)$, we can directly read off how UI benefits affect search effort:

$$\frac{de^*}{db} = -u'(b) < 0$$

Since $u'(b) > 0$ (more wealth is always better), an increase in benefits *unambiguously reduces* search effort. The mechanism:

- Higher b raises $u(b)$: being unemployed is now less painful
- This shrinks the gap $u(w) - u(b)$: the gain from finding a job falls
- So e^* falls: the worker searches less

At the extremes: - If $b = 0$: maximum effort $e^* = u(w) - u(0)$. With no income while unemployed, the worker has maximum incentive to find a job. - If $b = w$ (full insurance): $e^* = u(w) - u(w) = 0$. With the same income regardless of employment status, the worker has zero incentive to search.

The takeaway. Full insurance ($b = w$) completely eliminates the incentive to search. Any UI system that provides meaningful income support must accept some reduction in search effort. This is the core tension.

3.5 The Insurance–Incentives Tradeoff

More generous UI ($\uparrow b$) has two opposing effects:

Insurance value increases (+): - Smooths consumption across employment states: the worker has income in both period 2 states, not just when employed - Protects workers from income loss due to unemployment, which is partly involuntary (layoffs due to business conditions, not just worker behavior) - Particularly valuable for risk-averse workers: the utility gain from reducing income volatility can be large

Moral hazard cost increases (−): - Reduces search effort e^* - Workers stay unemployed longer on average: expected unemployment duration $= 1/e^*$ (in a simple model) - Government pays more in benefits (longer spells, same benefit level) - Output is lost: the job that would have been filled quickly is filled later

The **optimal** b balances these two forces. Neither extreme is optimal:

- $b = 0$: maximizes effort but leaves workers fully exposed to income risk; zero insurance value is socially suboptimal if workers are risk-averse
- $b = w$: provides full insurance but eliminates effort, which is also suboptimal because search output falls to zero

3.6 The Government's Problem

The government chooses b to maximize the worker's expected welfare subject to two constraints:

$$\max_b \quad u(b) + e^*(b)\, u(w) + (1 - e^*(b))\, u(b) - \tfrac{1}{2}(e^*(b))^2$$

Incentive compatibility: The worker's effort responds to b according to $e^* = u(w) - u(b)$. The government takes this as given; it cannot directly dictate effort.

Budget constraint: UI benefits must be funded, e.g., through payroll taxes on employed workers. A higher b raises both the per-period payout and the expected duration of unemployment, increasing the fiscal cost.

The key constraint is that the government **cannot choose** e **directly**: it can only choose b and let e adjust through the worker's optimization. This is why the second-best policy involves incomplete insurance.

3.7 First Best versus Second Best

First best (effort observable).

If the government could observe and contract on effort, it would choose:

- **Full insurance:** $b = w$. This perfectly smooths consumption across employment states, eliminating the income risk from unemployment. A risk-averse worker strictly prefers this.
- **Mandated effort:** The government also mandates the effort level that maximizes total surplus: choose e to maximize $e \cdot w - \frac{1}{2}e^2$, giving $e^{FB} = w$ (in utility units).

Under the first best, the worker is fully insured and faces the correct incentives. Risk and incentives are *separable* because effort is contractible.

Second best (effort hidden).

Now the government can only choose b, and e responds through $e^* = u(w) - u(b)$. Two things are true:

1. Setting $b = w$ (full insurance) implies $e^* = 0$: no search.
2. Setting $b < w$ restores positive effort but exposes the worker to income risk.

The second best requires **incomplete insurance** ($b < w$) to preserve search incentives. The worker must bear some income risk as the price of maintaining incentives. This is the **cost of moral hazard**: the gap between the first-best welfare level (full insurance + optimal effort) and the second-best level (partial insurance + lower-than-optimal effort) is the welfare loss created by the hidden action problem.

3.8 Moral Hazard Beyond UI

The same structure (insurance reduces effort, creating an insurance/incentives tradeoff) appears across many domains.

Health insurance. Insured patients consume more health care (more frequent doctor visits, more procedures, less prevention). Deductibles and copayments are the standard policy response: they make the patient bear part of the marginal cost, preserving at least some incentive to economize on care while still providing most of the insurance value.

Banking. Deposit insurance and implicit "too big to fail" guarantees protect depositors and the financial system. But they give banks an incentive to take on excessive risk: heads, the bank keeps the profit; tails, taxpayers bear the loss. Capital requirements (forcing banks to fund themselves partly with equity) are the policy response, giving banks "skin in the game."

Employment contracts. In the Gibbons principal-agent model, a firm (principal) hires a worker (agent) whose effort is unobservable. The contract pays $s + b \cdot y$, where y is measurable output. Steeper incentive pay (higher b) induces more effort, but also imposes risk on the worker because y is stochastic. The optimal contract trades off incentive provision against risk imposition.

Key Takeaways

	Adverse Selection	Moral Hazard
Type of problem	Hidden type (before contract)	Hidden action (after contract)
Example	Used cars, health insurance	UI, car insurance, banking
Market failure	Bad types drive out good types	Over-consumption of insurance/under-provision of effort

	Adverse Selection	Moral Hazard
Private fix	Signals, screening, reputation	Deductibles, copays, incentive pay
Public fix	Mandatory participation, disclosure	Time limits, monitoring, experience rating

The core lesson of adverse selection. When sellers know quality and buyers do not, the price must reflect average quality, which drives out above-average sellers, lowering average quality further. This feedback can unravel markets. Policy works by reducing the information asymmetry or mandating participation to prevent the most severe unraveling.

The core lesson of moral hazard. Full insurance eliminates risk but also eliminates incentives. Any practical insurance system must leave the insured party bearing some fraction of losses to maintain incentives. The optimal policy balances the welfare gain from consumption smoothing against the welfare loss from distorted behavior.

Together: adverse selection determines who gets coverage; moral hazard determines how the covered party behaves. Both create a wedge between what competitive markets deliver and what a fully-informed planner could achieve.