Exam 3 Review

Salome Aguilar Llanes

14.03 Fall 2025

Exam Topics

Market failures

- Externalities
- Private information markets
  - Adverse selection
  - Moral hazard
  - Signaling
  - Statistical discrimination
- Empirical methods used throughout the course special emphasis on IV

1 Coase Theorem

Key: Know the requirements of the Coase Theorem (clearly defined property rights and costless bargaining/transaction costs). When the requirements are satisfied, understand how trading allows externalities to be internalized.

The Coase theorem says that externalities get resolved iff:

1. Property rights are complete
2. Negotiating is costless

In some cases negotiation is infeasible

Airlines cannot realistically negotiate with individual homeowners for overflight rights to their houses

Recall: The problem of remedying externalities can be thought of as two separate problems

1. *What* should be done (sound insulation, quiet machines)?
2. *Who* should pay for it (doctor, baker)?

Other examples in class:

- LA mass transit
- Congestion

Problem: Sean is trying to write 14.03 exam questions, but his roommate wants to have a loud party in the common area - Sean finds the noise bothersome. Suppose that Sean's utility from not having to deal with his roommate's noise is $S > 0$ and his roommate's utility from hosting the party is $R > 0$, such that $S > R$. (Note that sometimes utilities can be expressed as a "disutility" where one party has negative utility for having to suffer an externality rather than positive utility for not having to suffer an externality.)

Suppose that MIT gives property rights to Sean. What Pareto optimum will they reach? If Sean and his roommate trade, at what price (or range of prices) will they trade?

No trade will occur. Sean's utility is S and his roommate's utility is R.

Suppose that MIT gives property rights to Sean's roommate. What Pareto optimum will they reach? If Sean and his roommate trade, at what price (or range of prices) will they trade?

Sean's roommate will sell the property right at a price in the range $R \leq p \leq S$. Sean's utility is $S - p$ and his roommate's utility is p.

Note again that the Coase theorem requires well-defined property rights and zero transaction costs.

2 Remedying pollution

- Consider two oil refineries that both produce fuel, which has a market price of \$3 per gallon. Assume that demand is infinitely elastic so that this price is fixed regardless of the quantity produced.

- Assume that each refinery uses \$2 in raw inputs (crude oil, electricity, labor) to produce 1 gallon of fuel.
- In addition, each plant produces smog, which creates \$0.01 of environmental damage per cubic foot.
- The amount of smog *per gallon of fuel* produced differs at the two plants:

$$\begin{aligned} s_1 &= y_1^2, \\ s_2 &= \frac{1}{2}y_2^2, \end{aligned}$$

  where y_1, y_2 denote the number of gallons of fuel produced at each plant. Plant 2 pollutes only $\frac{1}{2}$ as much as plant 1 for given production.
- Assuming initially that there are no pollution regulations. In this case, each plant will produce as many gallons as possible until it runs out of capacity (since it makes \$1 profit per gallon). Assume each plant can produce 200 gallons.

2.1 Competitive outcome

- What will firms choose to produce in the absence of any regulation of carbon output?

$$\begin{aligned} \max_{y_1} \pi_1 &= y_1 \cdot (3-2) \ s.t. \ y_1 \leq 200, \\ \max_{y_2} \pi_2 &= y_2 \cdot (3-2) \ s.t. \ y_2 \leq 200, \\ y_1^* &= y_2^* = 200. \end{aligned}$$

- Each firm ignores the social damage from its smog production (notice that s_1, s_2 do not enter into the firms' profit maximization problems). Hence, pollution is $s_1 = 40,000$, $s_2 = 20,000$. The negative pollution externality is \$400 and \$200 from plants 1 and 2 respectively.
- What is social surplus in this case? It is consumers' willingness to pay for output (\$3 per gallon) minus the resource costs of production minus the social costs of smog output:

$$(200 + 200) \times (3 - 2) - 0.01 \times (40,000 + 20,000) = -200$$

  Thus, in this example, the social damage from the externality swamps the benefits of consumption: we'd be better off producing no fuel at all than producing 200 gallons at each plant.
- Does this mean that we should produce no fuel? *Of course not.*

2.2 Welfare maximizing case

- Not all activities that generate externalities should be stopped. But if these activities generate negative (positive) externalities, then social efficiency generally suggests that we want to do less (more) of them than would occur in the free market equilibrium. Let's determine the optimal level of pollution.

- To get the socially efficient level of fuel production, we want to equate the marginal social benefit of the last gallon of fuel to the marginal social cost.

  - What is the social benefit? It is \$3. This comes from the infinitely elastic demand curve.
  - The marginal social cost of production is \$2 in raw inputs *plus* whatever pollution is produced.
  - The efficiency condition is $MB_s = MC_s$, marginal social benefit equals marginal social cost.

- We therefore want it to be the case that at the margin, there is no more than \$1 of environmental damage done per gallon of fuel produced. Consequently, no plant should produce more than 100 cubic feet of smog per gallon of fuel.

- Imagine that each plant faced the private *plus* social costs of production. If so, we could rewrite the previous profit maximizing conditions as:

$$\begin{aligned}
\max_{y_1} \pi_1 &= y_1 \cdot (3-2) - 0.01y_1^2 \quad s.t.\ y_1 \leq 200, \\
\max_{y_2} \pi_2 &= y_2 \cdot (3-2) \ - \frac{1}{2} \cdot 0.01y_2^2 \ \ s.t.\ y_2 \leq 200, \\
y_1^{**} &= 50,\ y_2^* = 100.
\end{aligned}$$

- When Plant 1 is producing 50 gallons, the marginal gallon produces 100 cubic feet of smog, which causes \$1.00 in environmental damage. More pollution than this would be socially inefficient since fuel sells for \$3 and uses \$2 in raw inputs to produce. For Plant 2, the corresponding production is 100 gallons because this plant produces less smog per gallon.

- What is the social surplus from output in this case?

$$(100+50) \times (3-2) - 0.01 \times \left(\frac{1}{2} \times 100^2 + 50^2\right) = \$75$$

- We now have an efficient benchmark for welfare maximization.

- How do we get plants to produce the socially efficient level of pollution?

2.3 Regulating quantities

- Permit some positive amount of an activity, but less than a private actor would otherwise undertake.
- How does this apply to the example above? We know the optimal quantity of production for each plant from our calculations above. We could therefore pass a law that says "Plant 1 may produce 50 gallons of fuel and Plant 2 may produce 100 gallons of fuel." This will achieve exactly the desired result.
- But this kind of regulation is clumsy. It's difficult to write laws that regulate the behavior of each plant individually. Once passed, such laws are difficult to modify as technology or pollution costs change.
- Even more importantly, quantity regulation generally provides inefficient incentives (or no incentives) at the margin. For example, let's say the regulation made a mistake and assigned $q_1^* = 100$, and $q_2^* = 50$, that is, the regulator reversed the allocations. Plant 1 would have no incentive to correct the situation, that is, to reduce its pollution, and plant 2 would have no effective means to bargain with plant 1 to increase its allocation since both plants have identical marginal profit from the next gallon of fuel under this regulatory scheme, there are no gains from trade. Notice that because these mistakes are *not* self-correcting, it *must* mean that the incentives provided are inefficient.
- If the law cannot be written to regulate each plant's output differentially, further inefficiencies will result. For example, let's say that the regulator decided to assign each plant the *average* of the efficient output levels (75) since it was not possible to assign them each individual quotas. In this case, total social surplus would be:

$$(75 + 75) \times (3 - 2) - 0.01 \times \left(75^2 + \frac{1}{2}75^2\right) = \$65.63,$$

which is lower than the \$75 in social surplus in the optimal regulation case above.

2.4 Price regulation

- An alternative approach is to use the price system to internalize the externality.
- We know from above that the marginal social cost of pollution is \$0.01 per cubic foot of smog. If we charged firms for polluting, the social cost would be incorporated in the private cost. Done correctly, firms will make optimal choices.
- This type of tax is known as a Pigouvian tax after the economist Pigou who suggested it.

– Specifically, it we set the pollution tax at $t = \$0.01$ per cubic foot of smog, then each plant would choose the optimal quantities as a result of profit maximization:

$$\begin{aligned}\max_{y_1} \pi &= y_1(3-2) - t \cdot y_1^2, \text{ where } t = 0.01 \rightarrow y_1^p = 50 \\ \max_{y_2} \pi &= y_2(3-2) - t \cdot \frac{1}{2} y_1^2, \text{ where } t = 0.01 \rightarrow y_2^p = 100\end{aligned}$$

– This solution achieves the desired result with arguably less complexity. Facing this tax, plants will choose the efficient amount of production. *We do not have to write a separate law for each plant. In fact, we don't even need to know firms' production functions to write this regulation correctly. All we need to do is calculate and price the marginal social damage of pollution (of course, we also need to enforce these regulations—a separate though important issue).*

– Note that this problem is made especially simple by the assumption that the marginal damage of pollution is always \$0.01 per cubic foot. If the marginal damage varied with the amount of pollution (plausible), then setting the right tax schedule would be much harder.

– For example, if pollution above a certain threshold caused mass extinction but pollution below this level did little harm, this Pigouvian taxation scheme would be *quite* risky. Setting the tax slightly too low would result in calamity.

3 Adverse Selection vs Moral Hazard

Key: Understand the similarities and differences between adverse selection and moral hazard. For adverse selection, quantitatively understand the conditions that need to be met in order for someone to purchase insurance.

Where there is private information, there is an incentive for agents to engage in strategic behavior. For example, if you are selling a product, and your buyer knows the distribution of product quality but not the quality of the individual product that you posses, how much should the buyer be willing to pay? The intuitive answer might be the *expected value* of the product, or perhaps the certainty equivalent of this lottery.

But this answer ignores an important consideration: the choice of what product you sell may depend on what price the buyer offers. And the price that the buyer offers may depend on what product she thinks you'll sell at that price. The equilibrium outcome in which buyer and seller expectations are aligned, that is, the buyer gets what she wants at the price she offers, may be far from efficient.

Examples in class

- The market for lemons
- Health insurance
- Credits and loans. Financial crisis both adverse selection and moral hazard.

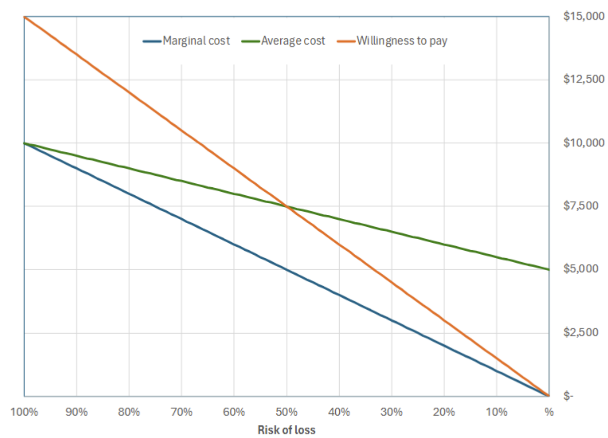

Figure 1: Insurance market with risk aversion

Problem: Suppose that buyers of Apple Care are more likely to damage their Apple devices. Is this observation due to adverse selection? Moral hazard? Both?

Adverse selection occurs if the people who buy Apple Care are high-risk individuals (i.e. those who don't take care of their devices).

Moral hazard occurs if the people who buy Apple Care become less careful as a result of having Apple Care (because they don't bear the full cost of damaging their devices).

Problem: Suppose I conduct an experiment in which I randomly assign iPhone users a free insurance plan, in which people receive either $50 or $500 if they break their phone. I find that people assigned to receive $500 are more likely to break their phone than those assigned to receive $50. Is this adverse selection, moral hazard, or both?

Since this is a randomized experiment, there is no opportunity for adverse selection. The difference in probabilities is then due to moral hazard since people assigned to the $500 group face less cost if they break their phone.

3.1 Insurance and Adverse Selection (adapted from 2016 exam)

Suppose a high school can be divided into two groups: high-risk (h) and low-risk (l) students. In any given year, 25% of high-risk students will get into a car accident, while 10% of low-risk students will. Car accidents cause $10,000 of damage.

Actuarially Fair Insurance

If the insurance company could not tell high-risk and low-risk students apart, what would they charge as the actuarially fair price of insurance? To find this, we need to calculate the expected losses:

$$E[L] = 0.5 \cdot 0.25 \cdot 10{,}000 + 0.5 \cdot 0.1 \cdot 10{,}000 = \$1750$$

So, the insurance company would need to charge $1750 per student in order to, in expectation, break even on insurance costs.

Willingness-to-Pay for Insurance

Suppose that families' utility functions are given by $u(x) = \ln(x)$ and their initial wealth is $50,000. In order to determine willingness-to-pay for both high-risk and low-risk students, we need to equate families' utility with and without insurance. For low-risk students' families, we get:

$$\begin{aligned} \text{with insurance: } & \ln(50{,}000 - x) \\ \text{without insurance: } & 0.9\ln(50{,}000) + 0.1\ln(40{,}000) \\ & \Rightarrow \ln(50{,}000 - x) = 0.9\ln(50{,}000) + 0.1\ln(40{,}000) \end{aligned}$$

We could then solve this equation for x to determine willingness-to-pay, but in an exam setting, you certainly would not be asked to solve an equation this complicated (because you cannot use a calculator). Doing the exact same thing for high-risk students, we would get:

$$\ln(50{,}000 - x) = 0.75\ln(50{,}000) + 0.25\ln(40{,}000)$$

Competitive Equilibrium

Suppose the willingness-to-pay for low-risk students is \$1200, and the willingness-to-pay for high-risk students is \$2700. Who would buy insurance in equilibrium, and what would the price of insurance be?

Since we calculated earlier that the insurance company would have to charge \$1750 to insure both groups, the low-risk students' families would be unwilling to buy insurance. Once they drop out of the market, the insurance company would need to charge

$$0.25 \cdot 10{,}000 = \$2500$$

for insurance. The high-risk students' families would still be willing to pay at that price, so the high-risk students are the only ones who are insured, at a price of \$2500.

3.2 Some solutions

3.2.1 Insurance Mandate

If the government wanted to mandate that everyone must buy insurance for \$1750, what would be the impact on total surplus? Would this be Pareto-improving compared to the competitive equilibrium?

Since low-risk students' families are only willing to pay \$1200 for insurance, they would lose \$550 in surplus. However, the high-risk students' families would gain \$750 in surplus. Since there are equal numbers of high-risk and low-risk students, this means that total surplus would increase. However, this would certainly *not* be Pareto-improving, since low-risk students are worse off than when they were not required to buy insurance.

3.2.2 Insurance Discount for Good Students

Suppose that the insurance company can charge different prices to students with high GPAs and low GPAs. Further assume that 10% of high-risk students get high GPAs, and 90% of low-risk students get high GPAs. Who will be insured, and what will the insurance company charge?

Within the high-GPA pool, 90% of the students are low-risk (in expectation) and the rest are high-risk. So the insurance company's actuarially fair price would be:

$$0.9 \cdot 0.1 \cdot 10{,}000 + 0.1 \cdot 0.25 \cdot 10{,}000 = \$1150$$

Both low-risk and high-risk students are willing to pay for insurance at this price, so high-GPA low-risk and high-risk students are all insured. Within the low-GPA group, the insurance company's expected losses are:

$$0.9 \cdot 0.25 \cdot 10{,}000 + 0.1 \cdot 0.1 \cdot 10{,}000 = \$2350$$

However, at this price, the low-risk students are unwilling to buy insurance, so they will drop out of the market. This means that low-GPA low-risk students will not be insured at all. **This could be an example of statistical discrimination**. Instead, only low-GPA high-risk students will be insured at a price of $2500.

3.3 Insurance and Moral Hazard

Moral hazard arises when an individual's behavior changes *after* entering into a contract because that behavior is imperfectly observed by the other party. In insurance markets, moral hazard occurs when people take less care to avoid accidents once they are insured, since they no longer bear the full cost of their actions.

Setup

Suppose again that a high school consists of high-risk and low-risk students. A car accident causes $10,000 of damage. Suppose that insurers can perfectly observe students' risk types, but they *cannot* observe how much effort students put into driving safely.

Each student can choose either:

- **High effort (safe driving)**: reduces accident probability but has a cost $1000 > c_h > 0$
- **Low effort (careless driving)**: increases accident probability

For a representative student:

- With **high effort**, the probability of an accident is 10%.
- With **low effort**, the probability of an accident is 25%.

High effort is privately costly (e.g., time, inconvenience, reduced enjoyment), but the insurer cannot observe effort.

Behavior Without Insurance

If a student is uninsured, they bear the full \$10,000 cost of an accident. The expected cost of an accident is:

$$\text{High effort: } 0.1 \cdot 10{,}000 = \$1{,}000 + c_h$$
$$\text{Low effort: } 0.25 \cdot 10{,}000 = \$2{,}500$$

Since low effort leads to much higher expected losses, uninsured students have a strong incentive to choose **high effort** in order to avoid accidents.

Behavior With Full Insurance

Now suppose the student becomes fully insured. The insurance covers the entire \$10,000 loss in the event of an accident. From the student's perspective, the financial cost of an accident is now zero.

Expected private costs become:

$$\text{High effort: } 0 + c_h$$
$$\text{Low effort: } 0$$

Since the student no longer bears any financial cost from accidents, they now have a much weaker incentive to exert high effort. If high effort is even slightly costly, the student will optimally choose **low effort**. This increase in risky behavior caused by insurance is **moral hazard**.

3.3.1 Impact on the Insurance Market

Because insured students now engage in riskier behavior:

- The probability of accidents rises.
- The insurer's expected payouts increase.
- Insurance premiums must increase to cover these higher losses.

Thus, even though insurance provides valuable risk protection, it also distorts behavior by weakening incentives for precaution.

3.3.2 Policy Tools to Reduce Moral Hazard

Insurance contracts are designed specifically to mitigate moral hazard using:

- **Deductibles**: insured individuals pay part of the loss
- **Copayments and coinsurance**: individuals share costs
- **Monitoring and enforcement**
- **Experience rating**: future prices depend on past behavior

These tools restore some private cost of risky behavior and encourage safer choices.

Key Distinction from Adverse Selection

- **Adverse selection** arises from hidden information *before* contracts are signed.
- **Moral hazard** arises from hidden actions *after* contracts are signed.

Both lead to market inefficiencies, but the policy responses are fundamentally different.

Moral hazard illustrates the fundamental tradeoff in insurance markets: insurance provides valuable protection against risk, but by insulating individuals from the consequences of their actions, it can also induce inefficiently risky behavior. Optimal insurance balances risk protection with incentives for precaution.

4 Signaling

Signaling arises when informed agents take observable actions in order to reveal their private information to uninformed parties. In contrast to adverse selection, where private information distorts market outcomes, signaling can *restore* efficiency by allowing agents to credibly communicate their type through costly actions.

Setup

Suppose again that a high school consists of two types of students, which employers cannot directly observe:

- **High-productivity students (H)**
- **Low-productivity students (L)**

Students may choose whether to obtain a college degree. Getting a degree does *not* increase productivity in this example; it only serves as a signal.

The cost of obtaining a degree differs by type:

- Cost for high-productivity students: \$5,000
- Cost for low-productivity students: \$30,000

Employers observe whether a student has a degree but cannot observe productivity directly.

Productivity (annual output to the firm) is:

- High-productivity: \$60,000
- Low-productivity: \$30,000

Wages Without Signaling

If employers cannot condition wages on education, they must pay a pooled wage equal to expected productivity. Suppose half of students are high-productivity and half are low-productivity:

$$w^{\text{pool}} = 0.5 \cdot 60{,}000 + 0.5 \cdot 30{,}000 = \$45{,}000$$

At this wage, high-productivity students are underpaid relative to their productivity, and low-productivity students are overpaid.

Wages With Signaling

Suppose employers believe:

- Students with a degree are high-productivity.
- Students without a degree are low-productivity.

Then wages become:

$$\begin{aligned} w_D &= 60{,}000 \\ w_{ND} &= 30{,}000 \end{aligned}$$

A **separating equilibrium** requires that:

- High-productivity students prefer to get the degree.
- Low-productivity students prefer *not* to get the degree.

High-productivity students compare:

$$\begin{aligned} \text{Degree payoff: } & 60{,}000 - 5{,}000 = 55{,}000 \\ \text{No degree payoff: } & 30{,}000 \end{aligned}$$

So high-productivity students strictly prefer to get the degree.

Low-productivity students compare:

$$\begin{aligned} \text{Degree payoff: } & 60{,}000 - 30{,}000 = 30{,}000 \\ \text{No degree payoff: } & 30{,}000 \end{aligned}$$

So low-productivity students strictly prefer *not* to get the degree.

Thus, education functions as a valid **signal**: only high-productivity students obtain it, and employers can infer productivity from education status.

Efficiency

Even though education does not increase productivity in this example, it allows firms to:

- Pay workers according to their true productivity
- Avoid misallocation caused by pooling

However, signaling is socially costly because resources are spent purely to transmit information rather than to increase output.

Key Distinction from Screening

- **Signaling**: the *informed* side (students) takes actions to reveal hidden information.
- **Screening**: the *uninformed* side (employers or insurers) designs contracts to induce revelation.

Signaling shows how markets can sometimes overcome adverse selection through costly actions that credibly transmit private information. The key requirement is that the signal must be *differentially costly* across types; otherwise, all types would mimic each other and the signal would convey no information.

4.1 Statistical Discrimination

Statistical discrimination arises when decision-makers use observable group-level characteristics as proxies for unobservable individual traits. Unlike adverse selection and moral hazard, statistical discrimination can occur even when there is no strategic behavior by individuals.

Setup

Suppose an employer is hiring workers but cannot directly observe individual productivity at the time of hiring. Instead, the employer observes a worker's group identity (for example, neighborhood, school quality, or previous employment sector).

There are two groups of workers:

- **Group A**
- **Group B**

Within each group, workers can be either:

- **High productivity (H)**
- **Low productivity (L)**

Productivity is not directly observable at the time of hiring.

True Productivity Distributions

Suppose the true productivity distributions differ across groups:

- In Group A, 70% of workers are high productivity and 30% are low productivity.
- In Group B, 30% of workers are high productivity and 70% are low productivity.

Productivity levels are:

$$y_H = 60{,}000,$$
$$y_L = 30{,}000.$$

Wage Setting Under Statistical Discrimination

Because individual productivity is unobservable, the employer pays each worker the *expected productivity of their group*.

For Group A:

$$w_A = 0.7 \cdot 60{,}000 + 0.3 \cdot 30{,}000 = \$51{,}000$$

For Group B:

$$w_B = 0.3 \cdot 60{,}000 + 0.7 \cdot 30{,}000 = \$39{,}000$$

All workers in Group A are paid $51,000, and all workers in Group B are paid $39,000, regardless of their true productivity.

Individual-Level Mispricing

At these wages:

- High-productivity workers in Group B are underpaid relative to their true productivity.
- Low-productivity workers in Group A are overpaid relative to their true productivity.

The wage gap arises *not* from taste-based prejudice, but from rational statistical inference based on imperfect information.

Dynamic Feedback Effects

Statistical discrimination can generate self-reinforcing dynamics. High-productivity workers in the lower-paid group (Group B) may choose to:

- Invest less in skills due to weaker incentives, or
- Exit the market altogether.

This can further reduce the average productivity of Group B over time, making the initial statistical belief self-fulfilling.

Even absent prejudice, statistical discrimination can still generate persistent inequality.

Statistical discrimination arises when market participants rationally use group averages to compensate for missing individual-level information. While it may be privately optimal for firms or insurers, it can generate efficiency losses, misallocation of talent, and persistent inequality at the individual level.

Practice: Identify the Information Problem

For each of the following scenarios, identify whether the main economic issue is:

- **Adverse Selection**
- **Moral Hazard**
- **Statistical Discrimination**
- **Signaling**

Scenarios

1. A health insurer raises premiums because it observes that people who voluntarily buy the most generous plans end up being the sickest on average. Healthy individuals gradually stop buying these plans.
2. After getting full car insurance with zero deductible, a driver begins parking in riskier areas and drives more aggressively.
3. A bank charges very high interest rates to borrowers from a particular low-income neighborhood because, on average, past borrowers from that area defaulted more, even though some current applicants are very safe borrowers.
4. A job applicant obtains a costly professional certification that does not increase productivity but is easier for high-ability workers to complete. Employers pay certified workers higher wages.
5. A lender cannot observe how carefully borrowers manage their businesses after receiving a loan. Once funded, many borrowers take on very risky projects.
6. Only the riskiest drivers choose to buy collision insurance when premiums are high, while safe drivers opt out.
7. Workers from School A earn higher starting wages than workers from School B because employers believe School A produces stronger students on average, even though individual ability is not observed.
8. Entrepreneurs who know their projects are unlikely to succeed are the most eager to accept loans with extremely high interest rates.
9. An employee works much less hard after being given a fixed salary with no performance bonuses.
10. A worker chooses to get a master's degree mainly to convince employers that she is highly motivated and capable, not because the degree itself raises her productivity.

11. An insurance company offers lower premiums to people who install a telematics device that tracks their driving behavior.
12. A firm offers two contracts:
    - a low wage with strong job protection, and
    - a high wage with weak job protection.

    High-productivity workers choose the second contract, while low-productivity workers choose the first.
13. A microfinance institution believes informal workers default more often than formal workers and therefore charges all informal workers higher interest rates, even though some informal workers have very stable incomes.
14. After receiving unemployment benefits, a worker reduces the intensity of their job search.
15. A student from a disadvantaged background works extremely hard to achieve a perfect GPA in order to convince employers that he is just as able as students from elite schools.

Answer Key

1. Adverse Selection
2. Moral Hazard
3. Statistical Discrimination
4. Signaling
5. Moral Hazard
6. Adverse Selection
7. Statistical Discrimination
8. Adverse Selection
9. Moral Hazard
10. Signaling
11. Moral Hazard (mitigation through monitoring)
12. Screening (by the firm; dual to signaling)
13. Statistical Discrimination
14. Moral Hazard
15. Signaling

5 IV recap

In class, we learned that instrumental variables (IV) provide a powerful tool to estimate causal effects when the variable of interest is endogenous. Often, the challenge is that the treatment X is correlated with unobserved characteristics that also affect the outcome Y, making simple regression estimates biased.

An IV design helps isolate variation in X that is plausibly exogenous. We use an instrument Z that affects X but has no direct effect on Y other than through X:

$$Z \to X \to Y$$

The instrument must satisfy two key assumptions:

- **Relevance:** Z must affect the endogenous variable X.

$$E[X_j|Z_j = 1] > E[X_j|Z_j = 0]$$

- **Exclusion restriction:** Z must affect Y only through X, not through any other channel. $E[Y_j|X_j = k, Z_j = 1] = E[Y_j|X_j = k, Z_j = 0]$.
- **Exchangeability** requires that $E[Y_j^1|Z_j = 1] = E[Y_j^1|Z_j = 0]$ and $E[Y_j^0|Z_j = 1] = E[Y_j^0|Z_j = 0]$

Mathematically, the IV estimator can be written as the ratio of the reduced form to the first stage:

$$\hat{\gamma} = \frac{E[Y_j|Z_j = 1] - E[Y_j|Z_j = 0]}{E[X_j|Z_j = 1] - E[X_j|Z_j = 0]}$$