Understanding Economics · Chapter 3

Market Failures & the Labor Market — narration script

1 · The puzzle

Welcome to Chapter 3 of Understanding Economics. In this chapter, we'll look at the times markets get it wrong: pollution, shared resources, and deals where one side knows more than the other. Then we'll turn to the job market and what really sets your wage.

[SLIDE: Two identical cars side by side. Left: "New, on the lot". Right: "Same model, 3 weeks old, 500 miles". A big question mark between the price tags.]

Here's a puzzle. Picture two cars, the same make, the same model, the same color. One is brand new on the dealer's lot. The other was bought three weeks ago, has five hundred miles on it, and is up for sale. It runs perfectly. It still smells new. And yet the three-week-old car usually sells for noticeably less. Sometimes a lot less.

Why? The car didn't get three weeks worse. Part of the answer is ordinary car-business stuff. But part of it is one of the most useful ideas in this whole course. By the end of this chapter you'll know what it is, and you'll start seeing it in insurance, in hiring, and in the online reviews you read.

2 · When markets fail

[SLIDE: Title "When Markets Fail". A small supply-and-demand cross on the left with a check mark: "Markets usually get it right". On the right, four icons: a smokestack, a fireworks burst, a car with a question mark, a paycheck.]

In Chapters 1 and 2, markets looked pretty impressive. Buyers and sellers, each looking out for themselves, land on a price where the amount people want to buy matches the amount for sale. Nobody plans it. And in Chapter 2 we saw that this point squeezes the most total benefit out of the market. Push the price away from it, with a cap, a floor, or a tax, and you get a deadweight loss: deals worth doing that don't happen.

But, as Chapter 2 flagged, that argument rests on a few assumptions. Three of them matter most here.

[SLIDE: Three boxes. 1: "The price includes all the costs and benefits." 2: "If you don't pay, you can't use it." 3: "Both sides know what they're getting."]

First, the price includes all the costs and all the benefits. Whoever makes the deal feels everything the deal causes.

Second, if you don't pay, you don't get the thing. The coffee shop doesn't hand you a latte unless you pay for it.

Third, both sides know what they're buying. You can see the apples, you can taste the coffee.

When those assumptions hold, markets do a remarkable job. When one of them breaks, the market can get the answer wrong, even with everyone acting sensibly. Economists call that a market failure. It doesn't mean the market collapses. It means the market's result is not the best one available. Too much of something, too little of something else, or deals that never happen at all.

This chapter tours the three big ways that happens, then finishes with the market almost everyone takes part in: the market for jobs.

One quick warning before we start. "The market failed" does not automatically mean "the government should step in and it will fix everything". Governments can get things wrong too, and we'll come back to that. Think of market failure as a diagnosis, not a prescription.

3 · Pollution and externalities

[SLIDE: A factory with a smokestack beside a neighborhood. Label on the factory: "Pays for: workers, materials, power". Label on the neighborhood: "Pays for: the smoke".]

Let's start with the biggest one. Picture a factory that makes plastic packaging. To make each batch, it pays for workers, plastic pellets, electricity, and machines. All of those costs show up on its books, so all of them show up in the price it needs to charge.

But the factory also puts smoke into the air. The people living downwind breathe it. Maybe they get more coughs, maybe they repaint their houses more often, maybe they just can't open their windows in summer. Those are real costs. But the factory doesn't pay them. Neither do the people who buy the packaging. The cost lands on someone who isn't part of the deal at all.

That's an externality. An externality is a cost or a benefit that falls on people outside a transaction. "External" just means outside. The buyer and the seller make their deal, and somebody standing nearby gets part of the bill, or part of the reward.

When the side effect is harmful, like smoke, noise, or traffic, it's a negative externality. When it's helpful, it's a positive externality. We'll do the harmful kind first, because it's the clearest.

You run into negative externalities all day. The neighbor's leaf blower at seven in the morning. The car ahead of you idling and pumping out exhaust. The person on the train playing music without headphones. Every one of those people is making a choice that has a cost, and part of the cost lands on you.

Here's the key point. The person making the choice weighs their own costs and benefits, not yours. That's not because they're bad people. It's because nothing in the price tells them about your cost. And when the price leaves out part of the cost, people do more of the thing than they would if they faced the whole cost.

Let's see it on the graph. This is the most important picture in this chapter, so we'll build it one piece at a time.

[GRAPH: Supply and demand for packaging. Vertical axis "Price ($ per unit)". Horizontal axis "Quantity (units per week)". Demand slopes down; a supply curve labeled "Supply = private cost" slopes up. They cross at price $10, quantity 1,000. All numbers illustrative.]

Here's the market for packaging. The numbers are made up, just to keep the arithmetic simple. Demand slopes down, the way it normally does. Supply slopes up. And remember from Chapter 2 what the supply curve really is: at each quantity, it shows the cost of making one more unit. Here, that's only the factory's own cost, its private cost. So I've labeled it "Supply equals private cost".

Left alone, the market settles where the two curves cross: a price of ten dollars and a thousand units a week. So far, this is just Chapter 1.

Now let's add the smoke. Suppose every unit of packaging puts out enough smoke to do about two dollars of harm to the neighbors. Again, an illustrative number.

[GRAPH: Same axes. A second upward-sloping curve appears exactly $2 above the supply curve, labeled "Social cost = private cost + $2 harm". A small vertical arrow between the curves is labeled "$2 per unit".]

So the true cost to society of making each unit is the factory's cost plus the two dollars of harm. That's this new curve, two dollars above the old one. Economists call it the social cost: everybody's cost, not just the seller's.

Now, here's your first check.

[SLIDE: "Pause and predict #1". Text: "The market ignores the smoke. Compared with what's best for everyone, does the market make too much packaging or too little? Is the price too high or too low?"]

Pause and predict. The market only looks at the private cost curve. Compared with what's best for everyone, does the market make too much packaging or too little? And is the price too high or too low? Pause the video, take a guess, then come back.

Here's the answer.

[GRAPH: The demand curve now also crosses the social cost curve, at price $11 and quantity 900. A dot there is labeled "Best for everyone: 900 units, $11". The original crossing is labeled "Market result: 1,000 units, $10".]

If you count the whole cost, the curves cross at a different point. Demand meets the social cost curve at nine hundred units and a price of eleven dollars. That's the quantity that's best for everyone, once the neighbors count too.

So the market makes too much: a thousand units instead of nine hundred. And the price is too low: ten dollars instead of eleven. If you said "too much, too cheap", you nailed it. The price is low because it leaves out the smoke. And because the price is low, people buy more than they would if they paid the full cost.

Let's slow down on why those extra hundred units are a problem, because this is the part that trips people up.

[GRAPH: A thin shaded triangle between the social cost curve and the demand curve, from 900 to 1,000 units. Label: "Deadweight loss ≈ ½ × 100 × $2 = $100 per week (illustrative)".]

Look at the units between nine hundred and a thousand. For each of those units, the value to the buyer is shown by the demand curve. The true cost to society is shown by the social cost curve. And in this whole stretch, the social cost curve is above the demand curve. That means each of those units costs society more than it's worth to the person buying it. They get made anyway, because the buyer and the seller don't feel the whole cost.

That gap is the shaded triangle. It's a deadweight loss, the same idea from Chapter 2. Last time, a deadweight loss came from deals that didn't happen. This time it comes from deals that did happen and shouldn't have. With these numbers, the triangle is one half, times a hundred units, times two dollars, or about a hundred dollars of lost value every week.

So notice what "too much" means here. It doesn't mean packaging is bad, or that the factory should shut down. Nine hundred units a week are worth more than their full cost, smoke included. The problem is only the last hundred. Most economists don't think the right amount of pollution is zero. They think the right amount is the amount where the benefit of one more unit just covers its full cost.

The same picture works in reverse for positive externalities.

[GRAPH: A new market, "Flu shots". Demand curve labeled "Private benefit". A second demand curve above it labeled "Social benefit = private benefit + protection for others". The market crossing sits at a lower quantity than the social crossing. Label: "Market makes too few". Illustrative.]

Think about a flu shot. When you get one, you protect yourself. That's your private benefit, and it's why you'd pay for it. But you also make it less likely that you'll pass the flu to a coworker, a grandparent, or the person next to you on the bus. They benefit too, and they didn't pay you anything.

So the true benefit to society sits above the private benefit. The market only counts the private part, so it lands on too few flu shots. Same with education, which tends to make a whole community more productive, not just the person who studied. Same with a neighbor who keeps a beautiful garden that everyone on the street gets to enjoy.

So here's the rule of thumb. Harmful side effects mean the market makes too much. Helpful side effects mean the market makes too little. In both cases, the price is sending the wrong signal, because it leaves something out.

4 · Fixing externalities

[SLIDE: Four fixes, as four tiles: "Rules", "A tax on the harm", "Permits you can trade", "Bargaining". Each with a small icon.]

So what can a society do about it? There are four main tools. All four are trying to do the same thing: make the people making the choice feel the cost they're putting on others. Economists call that internalizing the externality. Bringing the outside cost inside the decision.

Tool one: rules. The government simply sets a limit. Cars must have a certain kind of exhaust system. Factories can only release so much of a chemical. Construction noise stops at ten at night. Rules are clear and easy to understand. Their weakness is that they treat everyone the same, even when cutting pollution is cheap for one firm and very expensive for another.

Tool two: a tax on the harm. This one's elegant, so let's go back to our graph.

[GRAPH: The packaging market again. A $2 per-unit tax shifts the supply curve up by $2, so it lies exactly on the social cost curve. New result: 900 units. Labels: "Buyers pay $11", "Factory keeps $9", "Tax: $2 × 900 = $1,800 per week". Illustrative.]

Suppose the government charges the factory two dollars for every unit, exactly the harm each unit causes. You learned in Chapter 2 that a per-unit tax shifts the supply curve up by the amount of the tax. So the supply curve moves up by two dollars and lands right on top of the social cost curve.

Now the market crosses at nine hundred units, the best amount. Buyers pay eleven dollars. The factory keeps nine after the tax. And the government collects two dollars times nine hundred units, eighteen hundred dollars a week, which it could use to help the neighbors or cut some other tax.

In Chapter 2, a tax created a deadweight loss, because it pushed a market away from its best point. Here, the tax does the opposite. It pushes the market toward its best point. That's why economists like this kind of tax so much. It's often called a Pigouvian tax, named after the economist Arthur Pigou, who's usually credited with the idea. A tax on cigarettes or on gasoline is sometimes defended this way.

And notice something else. With a tax, the government doesn't have to know how each factory should cut its smoke. Each factory figures out for itself whether it's cheaper to pay the tax or to clean up. The ones that can clean up cheaply do. That's the big advantage over a one-size-fits-all rule.

The catch: you have to put a number on the harm. How many dollars of damage does a ton of smoke really do? That's hard to measure, and people argue about it.

Tool three: permits you can trade. This is often called cap-and-trade. The government sets a total cap on pollution, say a fixed number of tons a year, and hands out or sells permits that add up to exactly that cap. To pollute one ton, you need one permit. And firms can buy and sell permits with each other.

Here's why trading helps. Say one firm can cut a ton of pollution for twenty dollars, and another would have to spend two hundred. The cheap cutter cleans up and sells its spare permit to the expensive one. Same total pollution, much lower total cost. A tax fixes the price of pollution and lets the amount adjust. Cap-and-trade fixes the amount and lets the price adjust.

Tool four: bargaining. Sometimes the people involved can just work it out. Your neighbor's dog barks all night, and you offer to pay for a better fence. The idea, associated with the economist Ronald Coase, is that if property rights are clear and bargaining is cheap, people can often reach a good outcome on their own. The catch is the "if". When the smoke hits a hundred thousand people, they can't all sit down at one table with the factory. So bargaining works for small, local problems, and the other tools handle the big ones.

And for positive externalities, the same logic runs in reverse. Instead of taxing the activity, you subsidize it. Free or cheap flu shots. Public funding for schools. Patents, which give inventors a temporary right to profit from an idea, are another way to reward something that helps lots of people beyond the inventor.

One honest note before we move on. These fixes are run by governments, and governments are made of people with their own incentives. A tax can be set at the wrong level. A rule can be written to favor a well-connected industry. A subsidy can outlive its purpose by decades because the people receiving it lobby hard to keep it. Economists call this government failure. So the real question usually isn't "market or government?" It's "which imperfect option does better here?"

5 · Public goods and the commons

[SLIDE: A fireworks show over a crowd. Some people in a paid area; many more watching free from a hill and from their apartment windows.]

Now the second broken assumption: if you don't pay, you can't use it.

Picture a town that wants a fireworks show for a holiday. It costs, say, fifty thousand dollars. An illustrative number. Lots of people would love it. But think about how a private company would sell tickets. It can't. Once the fireworks are in the sky, everyone within a few miles can see them, ticket or not.

So if you're a typical resident, why would you pay? If enough other people chip in, you get the show for free. If they don't, your twenty dollars won't save it anyway. That's the free rider problem: when people can enjoy something without paying, many of them won't pay. And if everyone reasons that way, the show never happens, even though the town as a whole would gladly have paid for it.

To see which goods have this problem, economists ask two questions about any good.

[SLIDE: A 2×2 table. Columns: "Can you keep people out? Yes / No". Rows: "Does one person's use leave less for others? Yes / No". Four boxes, filled in as the narration goes.]

Question one: can you keep people out? If you can stop people who don't pay from using it, the good is called excludable. A coffee is excludable. A fireworks show isn't.

Question two: does one person's use leave less for others? If my using it means there's less for you, the good is called rival. If I drink a coffee, you can't drink that same coffee. Rival. If I watch the fireworks, there's just as much sky for you. Not rival.

Put those two questions together and you get four kinds of goods.

[SLIDE: Table fills in. Top-left "Private goods: coffee, phones, groceries". Top-right "Common resources: ocean fish, a free beach, a busy road". Bottom-left "Club goods: streaming services, gyms, a toll bridge". Bottom-right "Public goods: fireworks, streetlights, national defense".]

Private goods are excludable and rival. Coffee, phones, a sandwich. This is where markets work best, and it's everything we studied in Chapters 1 and 2.

Club goods are excludable but not rival, at least until they get crowded. A streaming service is the classic case. The company can lock you out if you don't pay, but you watching a show doesn't use it up for anyone else. Markets handle these fine, because the seller can charge.

Public goods are neither. You can't keep people out, and one person's use doesn't reduce anyone else's. Fireworks. Streetlights. A flood wall that protects a whole neighborhood. National defense, which protects everyone in the country whether they pay or not. These are where free riding hits hardest. A private market tends to provide too little of them, or none at all.

That's the main reason governments pay for public goods out of taxes. Taxes are basically a way to make everyone chip in, so the free rider problem goes away. There are other routes too. Public radio

runs on donations. Some open-source software gets built by volunteers and by companies that need it themselves. But for the big ones, like defense, roads, and basic scientific research, governments do most of the paying.

Now the last box, the trickiest one. Common resources: you can't keep people out, but one person's use does leave less for others. Fish in the open ocean. A free beach on a hot day. A road at rush hour.

These lead to a famous problem called the tragedy of the commons. The phrase was made popular by an essay from the late 1960s, usually credited to the ecologist Garrett Hardin, but the idea is much older.

[SLIDE: A shared lake with ten fishing boats. Caption: "Each boat thinks: if I don't catch it, someone else will."]

Here's how it works. Picture a lake that ten fishing families all share. Nobody owns the fish. Each family knows that if everyone takes too much, the fish won't have time to breed, and the lake will run dry. But each family also knows this: any fish it doesn't catch today will probably just get caught by another boat. So holding back doesn't save the lake. It just means someone else gets your fish.

So every family fishes hard, and the lake gets emptied, even though every single family would be better off if they all held back. Nobody's being foolish. Each choice is reasonable on its own, and together they destroy the thing everyone depends on.

You've probably seen a small version of this at work: the shared office fridge, or a shared kitchen that nobody cleans. Everyone uses it, nobody owns it, and it slowly falls apart.

And notice, this is really an externality in disguise. Each family's extra catch puts a cost on all the others, a cost it doesn't pay.

The fixes look familiar. Rules, like fishing seasons and catch limits. Prices, like a toll on a crowded road at rush hour. Or ownership: give each family a share of the catch it can keep, sell, or save. Once you own something, you have a reason to look after it. And sometimes communities manage a shared resource well on their own, with local customs and rules everyone agrees to.

6 · Hidden information

[SLIDE: A buyer and a seller across a table. The seller has a thought bubble showing the car's full history. The buyer's thought bubble shows only a question mark.]

Now the third broken assumption: both sides know what they're getting. And this is where we come back to the opening puzzle.

When one side of a deal knows something important the other side doesn't, economists call it asymmetric information. "Asymmetric" just means uneven. Sellers often know more than buyers. Sometimes buyers know more than sellers. And this uneven knowledge creates two separate problems. One happens before a deal. The other happens after. Let's take them one at a time, slowly, because people often mix these up.

The first problem is called adverse selection. It's about who shows up to the deal.

The classic example is used cars, and it's usually traced to a famous paper by the economist George Akerlof. Let's build it with simple, made-up numbers.

[SLIDE: Two piles of identical-looking used cars. Left pile labeled "Peaches: good cars, worth $12,000 to buyers". Right pile labeled "Lemons: hidden problems, worth $6,000 to buyers". Numbers illustrative.]

Picture a used car market where half the cars are good and half have hidden problems. People call the bad ones lemons and the good ones peaches. To a buyer, a good car is worth twelve thousand dollars, and a lemon is worth six thousand.

The sellers know which kind they have. They've been driving it. The buyers can't tell, because from the outside a lemon and a peach look exactly the same.

So what's a careful buyer willing to pay? If it's a coin flip which one you get, a reasonable buyer might offer something like the average: nine thousand dollars.

Now look at it from the sellers' side. If you own a lemon, nine thousand dollars is a great deal. Your car is only worth six. You'll happily sell. But if you own a peach, nine thousand is too low. Your car is worth twelve. Many owners of good cars will say no thanks and keep driving them.

[SLIDE: The peach pile shrinks as owners walk away. The lemon pile stays. Caption: "At $9,000, good cars leave. The cars still for sale are mostly lemons."]

So which cars are left on the market? Mostly lemons. And buyers aren't stupid. They figure out that the cars for sale are worse than average. So they lower their offer. Which pushes even more good cars off the market. In the extreme, only the lemons are left.

That's adverse selection. The people who choose to show up for a deal are not a random sample. They're tilted toward the ones the other side least wants to deal with. Bad cars push good cars out of the market, even though there are buyers who'd gladly pay full price for a good car if only they could tell which was which.

And now the puzzle.

[SLIDE: Back to the opening picture: the new car and the three-week-old car. Left, a short list "Ordinary reasons: dealer incentives and margins, warranty and first-owner perks, buyers who want to be first". Right, the question mark is replaced by a caption: "Plus: why is the owner selling?"]

Why does a car that's three weeks old sell for noticeably less than the identical new one? Start with the ordinary reasons, because they're real. New cars often come with dealer incentives, like rebates or cheap financing, that a used-car seller can't match. A dealer buying a nearly-new car to resell needs room for its own margin. The second owner may get less of the warranty, or miss perks meant for the first owner. And plenty of buyers simply like being first, and will pay for it. Together, those can explain a lot of the gap.

But there's one more piece, and it's the interesting economic part, because it can push the price down even when everything else is equal. When a buyer sees a car that's nearly new for sale, the first question in their head is: why is the owner getting rid of it already? Maybe the owner's plans changed. Or maybe the owner found a problem. The buyer can't tell which, and the seller knows. So the buyer knocks the price down for the risk of a lemon. And that lower price makes owners with perfectly good nearly-new cars less likely to sell at all, which tilts the pool even further. The car didn't get three weeks worse. Partly, it got a question mark attached to it. That's hidden information putting a price on itself.

So how do markets fight back? They find ways to make the hidden information visible.

Sellers can send a signal: something that's cheap for an honest seller and costly for a dishonest one. A warranty is a good example. If you know your car is good, offering a one-year warranty costs you very little, because it probably won't break. If you know it's a lemon, a warranty could cost you a fortune. So a warranty tells the buyer something.

Buyers can screen: they gather information themselves. Paying a mechanic to inspect the car, or checking a vehicle history report. Online reviews, star ratings, and seller scores do the same job across the whole internet. They exist largely because buyers can't see inside the product before they pay.

Now let's go to the market where adverse selection matters most: insurance. Here, the hidden information flips. The buyer knows more than the seller. You know far more about your own health, or your own driving, than an insurance company does.

[SLIDE: "Pause and predict #2". Text: "A health insurer charges everyone the same price, set for an average customer. Who is most eager to sign up? What happens to the price next year?"]

Pause and predict. Imagine a health insurer that charges every customer the same price, set for the average person. Who is most eager to sign up? And what's likely to happen to the price next year?

Pause, think it through, then come back.

Here's how it tends to play out. The people most eager to buy are the ones who expect big medical bills. People who expect to stay healthy may decide the price isn't worth it and skip it. So the insurer's customers end up sicker than average, and it pays out more than it planned. Next year, it has to raise the price. That pushes out the healthiest of the remaining customers. Costs per customer go up again, and so does the price. In the worst case it spirals, and the insurance gets more and more expensive for a smaller and sicker group.

If you predicted the price goes up and the healthy people leave, that's adverse selection at work. It's why insurers ask questions, look at records, and charge different prices for different risks. It's why many employers offer insurance to all their workers as a group, which brings in healthy and sick people alike. And it's a big part of why some governments require people to have health coverage, or provide it directly: it keeps the healthy people in the pool.

Now the second problem, the one that happens after the deal. It's called moral hazard.

Moral hazard is when someone changes their behavior after a deal, because someone else now carries part of the risk. And the other side can't fully watch what they do.

[SLIDE: Two scenes. Left: a rental car being driven gently, labeled "My car". Right: the same car bouncing over a pothole, labeled "Rental car, full coverage". Caption: "The deal changed the behavior."]

Think about how people treat a rental car compared with their own car. Or think about what happens after you buy full insurance on a phone. Some people get a little less careful about where they set it down. Nobody's planning to break it. But when the cost of a mistake falls on someone else, people take a bit less care.

It shows up everywhere. A worker who's paid by the hour and isn't supervised might not hurry. A contractor paid for every hour might find the job takes longer. And at the scale of a whole economy, a bank that expects to be rescued if things go wrong may take bigger risks than one that knows it's on its own.

So here's the clean way to tell them apart. Adverse selection is about hidden information before the deal: who chooses to show up. Moral hazard is about hidden actions after the deal: how people behave once they're covered.

And again, markets have tools. For moral hazard, the main trick is to leave some of the risk with the person who controls it. That's what a deductible does. If you pay the first few hundred dollars of any claim yourself, you still have a reason to be careful. Same with co-pays at the doctor, security

deposits on an apartment, or paying salespeople partly on commission. Each one keeps a bit of the cost of a bad outcome on the person whose choices cause it.

7 · Wages and the job market

[SLIDE: A supply-and-demand graph for labor. Vertical axis "Wage ($ per hour)". Horizontal axis "Number of workers". Labor demand slopes down (employers). Labor supply slopes up (workers).]

Let's finish with the market almost every adult takes part in: the job market. We'll keep this light, because the main idea is one you already know.

A wage is a price. It's the price of an hour of work. And the job market works like any other market, just with the roles switched from what you might expect. Workers are the sellers: they supply their time and skills. Employers are the buyers: they demand labor.

Labor supply slopes up. At higher wages, more people are willing to work, or to work more hours. Labor demand slopes down. At higher wages, employers hire fewer people, find other ways to get the job done, or use more machines. Where the curves cross, you get the market wage and the number of jobs. You saw in Chapter 2 what happens when a minimum wage is set above that point: the price-floor picture.

One thing is special about labor demand, though. Nobody wants workers for their own sake. A bakery hires a baker because people want bread. So the demand for workers comes from the demand for what they make. Economists call it derived demand. If people suddenly want more bread, bakeries want more bakers, and the bakers' wage tends to rise. If people stop buying something, the jobs that make it are in trouble, however good the workers are.

And how many workers does a business hire? It's thinking at the margin, straight from Chapter 1. A business keeps hiring as long as the next worker adds more to its revenue, after the extra supplies they use, than they cost in wages. When one more worker would bring in less than they cost, it stops.

[SLIDE: Illustrative table for a coffee shop. Columns: "Baristas", "Extra drinks sold per hour", "Extra revenue at $5 per drink", "Minus beans, milk, cups at $1 per drink", "Left over", "Wage $20 per hour". Rows: 1st barista 12 drinks, $60 − $12 = $48; 2nd 8 drinks, $40 − $8 = $32; 3rd 6 drinks, $30 − $6 = $24; 4th 3 drinks, $15 − $3 = $12. The 4th row is crossed out.]

Here's a quick illustrative example. A coffee shop sells drinks for five dollars and pays baristas twenty dollars an hour. Each drink also uses about a dollar of beans, milk and a cup, so each extra drink leaves four dollars to cover the wage. The first barista can make twelve extra drinks an hour: sixty dollars in, twelve for supplies, forty-eight left. Easy hire. The second adds eight drinks, thirty-two dollars left. Still worth it. The third adds six drinks, twenty-four dollars left. Still more than twenty. But the fourth only adds three drinks, because they're bumping into each other behind the

counter. That's twelve dollars left over, for a twenty-dollar wage. So the shop stops at three. And notice: if drinks got more popular and the price went up, that fourth barista might become worth hiring. Derived demand, right there in a table.

So why do wages differ so much from job to job? Here are the main reasons, quickly.

[SLIDE: Four reasons, as a list: "Skills and training", "Unpleasant or risky work", "Scarcity of the skill", "Bargaining power".]

One: skills and training. Economists call the skills you build human capital, because it works a lot like a machine you invest in. It costs time and money up front, and it makes you more productive later. Jobs that need rare, hard-to-learn skills tend to pay more.

Two: some jobs are unpleasant, risky, or at awkward hours. Night shifts, dangerous work, jobs far from home. Employers usually have to pay extra to get people to take them. That extra pay is called a compensating differential. It compensates for the downside.

Three: scarcity. Supply and demand again. If lots of employers want a skill and few people have it, its wage gets bid up. If many people can do a job, the wage stays lower, even if the job is hard.

Four: bargaining power. This one is where the job market can break the simple model. Picture a small town where one big employer hires most of the workers. If you don't like its wage, there's nowhere else nearby to go. When there's only one main buyer of labor, economists call it monopsony. It's the mirror image of a monopoly from Chapter 2. A monopoly is one seller with power over price. A monopsony is one buyer with power over price. And just as a monopoly can charge more, a monopsony can pay less than a competitive market would, and hire fewer people too.

That's one reason labor unions exist. A union bargains for many workers at once, which can balance out an employer's power. Unions can also push wages above the market level, which, like any price floor, may mean fewer jobs. Economists argue a lot about where each effect is bigger, and the honest answer is that it depends on the market.

And of course, wages aren't only shaped by markets. Discrimination, where people are paid or hired differently for reasons that have nothing to do with their work, is a real and well-documented problem, and it's both unfair and a kind of waste, because the economy loses out on talent.

Notice how this connects to the rest of the chapter. Hidden information is all over the job market. An employer can't see how hard you'll work before hiring you. That's adverse selection. And once you're hired, it can't watch everything you do. That's moral hazard. Degrees, references and work samples are signals. Interviews and trial periods are screens. Bonuses and commissions are the job market's version of a deductible: they keep some of the outcome riding on the effort. Same tools, new market.

8 · Recap, cheat sheet & self-check

[SLIDE: "Recap". Five short lines appear one at a time.]

Let's lock it in. Here's the whole chapter in five lines.

Markets work well when prices include all the costs and benefits, when people have to pay to use things, and when both sides know what they're getting. When one of those breaks, you get a market failure: the market's result isn't the best one available.

Externalities are costs or benefits that land on people outside the deal. Harmful ones mean the market makes too much at too low a price. Helpful ones mean it makes too little. The fixes, rules, taxes on the harm, tradable permits, bargaining, and subsidies, all try to make people feel the full cost or benefit of their choices.

Public goods, which you can't keep people out of and don't get used up, suffer from free riders, so markets provide too few of them. Common resources, which you can't keep people out of but do get used up, get overused. That's the tragedy of the commons.

Hidden information creates two problems. Adverse selection, before the deal, tilts who shows up. Moral hazard, after the deal, changes how people behave. Signals, screening, and deductibles fight back.

And wages are prices set by supply and demand for labor, shaped by skills, unpleasant work, scarcity, and bargaining power.

[SLIDE: "Cheat sheet". Three columns. Left: the externality graph, with private cost, social cost, the market point and the best point, the shaded deadweight loss triangle, and the arrow for a corrective tax. Middle: the 2×2 table of goods. Right: two boxes, "Adverse selection: hidden info, before the deal, fix with signals and screening" and "Moral hazard: hidden action, after the deal, fix by sharing the risk".]

Here's your cheat sheet. Screenshot this one, or grab it from the course page. On the left, the externality graph: the gap between private and social cost, the market making too much, and the tax that closes the gap. In the middle, the four kinds of goods. On the right, the two information problems and how each one gets fixed. And the one question to ask about any market that seems to be going wrong: which assumption broke? Does the price leave out a cost? Can people use it without paying? Or does one side know something the other doesn't?

Now three questions to check yourself. Pause after each one, answer out loud, then let it play.

Question one. A city has a crowded downtown where cars are stuck in traffic every rush hour, and the exhaust fills the air. Name the market failure, and suggest one fix an economist would like.

Answer. It's a negative externality, and the road is also acting like a common resource: each extra car slows everyone else down and adds pollution it doesn't pay for. So there are too many car trips at rush hour. An economist's favorite fix is a charge for driving downtown at peak times. It's a tax on the harm. People who really need to drive still can, and others shift to other times, other routes, or transit.

Question two. A company starts offering phone insurance with no deductible: any damage at all, a free replacement. Over the next year, more customers' phones get broken than the company expected. Which information problem is this, and what would you change?

Answer. That's moral hazard. Once people were fully covered, some took less care. The fix is to leave some of the cost with the customer, like a deductible or a fee per claim. If you said adverse selection, that would be the story if careless people were the ones most likely to buy the insurance in the first place. In real life, both can happen at once.

Question three. A neighborhood would love a set of new streetlights. Everyone agrees they're worth it. But when a volunteer goes door to door asking for donations, hardly anyone pays. Why, and what usually happens instead?

Answer. Streetlights are a public good. You can't stop non-payers from enjoying the light, and one person's use doesn't reduce anyone else's. So people free ride, hoping others will pay. That's why things like streetlights are usually paid for through taxes, which makes everyone chip in.

[SLIDE: "Where we're going". The 6 chapters listed, with 1, 2 and 3 checked: "1 Foundations & Supply and Demand", "2 How Markets Work", "3 When Markets Fail", "4 The Whole Economy", "5 Money, the Fed & Interest Rates", "6 Trade, Stocks & Bonds, and Real Estate". Chapter 4 highlighted: "Next up".]

That's Chapter 3, and it finishes the zoomed-in half of the course: microeconomics.

In Chapter 4, we zoom out to the whole economy at once. How do you measure the size of a whole country's economy? What does the unemployment rate actually count? Why do prices rise everywhere at the same time? And what are recessions, and what can a government do about them? See you in Chapter 4.