

[SQUEAKING]

[RUSTLING]

[CLICKING]

JON GRUBER: The course is essentially divided into three sections. OK, we've got the first section on externalities of public goods, sort of the government regulatory approach. We've got a second section devoted to social insurance and all the things the government spends money on to protect us against various risks in society, and the third section on taxation.

So today we'll start with chapter 5 and talk about externalities. Now, we'll go into more detail next time. But the classic example of what we call an externality is global warming and climate change. I don't need to tell you about the horrors of climate change and what the future looks like under a world of a rapidly-warming climate.

You can see if you look at figure 5.1 on your handout-- you can see the rapid rise in global temperatures that we've seen in recent decades. It's only spiked massively in the years since this textbook was published. There's actually been an upward spike.

As you know, the global temperatures have risen rapidly over the past century and will continue to rise rapidly in the near future. That is unambiguously due to human activity, like driving-- that basically, the more we drive, the more carbon we put in the air, the more the world is warming.

Yet when you last drove, did you think about that? Did you think when you got behind the wheel of a car, that you were microscopically but measurably putting Bangladesh this close to being underwater? No, you did not. You did not think about that when you drove your car. And that's because global warming and emissions from driving are a classic example of what we call a negative externality. Negative externality.

A negative externality-- an externality in general arises whenever the actions of one party impact another party, and the first party does not bear the consequences of doing so. Whenever the actions of one party impact another party, and that first party does not bear the consequence of doing so, there is an externality.

You do not bear the consequence of putting Bangladesh microscopically, but measurably more underwater by driving. Therefore, there is a negative externality in this case. So externalities are the classic-- if you think about the prototypical-- if think about the term market failure in economics, the classic example economists turn to are externalities.

This is sort of the most market-failure-y of the market failures. And that's why we learn about it first. So today we're going to talk about externality theory. It'll bleed over a little in the next lecture. And then we'll spend the rest of next lecture and lecture after talking about examples of externalities in the real world, first in the environmental context and then in the health context.

So we're going to start with the basic theory of externalities. Now, in this case, I'm going to do this primarily graphically. We'll do the mathematics of externalities in section on Friday. But basically, your goal in this class, as I said in 14.01 for those who took it 14.01 with me, is to understand every topic intuitively, graphically, and mathematically.

That's our goal here. I will provide you the intuition and the graphics here in lecture, sometimes the math. Where I don't do the math, section will add the math. But between lecture and section, you should understand every one of these topics at all three of those levels. And it's important that you understand all three of them to really get the full benefit of this course and to do well in the class.

So let's talk about four types of externalities. Let's start by type 1. Let's talk about negative production externalities. It's the first type of externality I want to talk about, negative production externalities.

So basically, let's talk about an example-- a running example used through this lecture-- which is the example of a steel plant that produces steel and sits on a river. And to make life easy, we're going to assume that each unit of steel sells for \$1. So the firm revenues and quantities are the same.

This is a trick in economics we call making a numeraire good. If you make the price of something \$1, then quantities and expenditures are the same. So a unit of steel costs \$1. So the amount of steel produced is the amount of money made by the company.

We're also going to assume that for every unit of steel produced, there is one unit of sludge produced as a byproduct of that steel. And that gets dumped into a river. So mechanically, when you make a unit of steel, you make a unit of sludge that falls into the river.

That floats down the river. Downstream from the plant is a set of fishermen. Those fishermen are catching fish in the river. But when the sludge floats downstream, it kills some of their fish. And the fishermen, as a result, make less money.

This is a classic example of a negative production externality-- negative because it's having a negative effect on the fishermen, production because it comes from the production of steel, and externality because the steel plant owner doesn't give a shit. He's just making his steel and some fish are dying.

OK, so how do we think about that graphically? Well, we think about that in the context of the market for steel. OK? So let's imagine the market for steel. You have some quantity of steel produced on this axis and the price and the steel up here.

Now, people get a benefit from using steel. There's some marginal benefit curve-- marginal benefit, which we also call the demand curve. Remember when you reviewed in chapter 2, the demand curve represents the marginal benefit of using a unit of steel. And then we have some marginal cost curve, which is the marginal cost of producing a unit of steel.

And the private equilibrium, as we know, are where marginal benefits equal marginal costs. And that will be at a unit of production. We'll call it Q_2 At a price and a unit of production-- I'm sorry, P_1 . There's a typo in the book. We'll call this P_1 and Q_1 .

Sometimes when I'm teaching I'll catch typos in the book. Paul, can you send me an email telling me there's a typo with figure 5-2? OK. Figure 5-2 has Q_1 and Q_2 mislabeled. So basically what you see is that's the private equilibrium. That's 14.01.

Now, what we don't do in 14.01-- I did in one lecture, but what you generally don't do in 14.01-- is realize that we've made an implicit assumption here, that these are actually not the marginal benefits. These are the private marginal benefits and private marginal costs.

But in reality, society may value things differently than the private producer does. And in particular, every unit of steel that's being produced imposes damage on the fishermen, damage of an amount D . So we're going to draw a social marginal cost curve, which is the private marginal cost shifted down by the amount D .

I'm assuming that damage is linear. It doesn't have to be a parallel shift. In fact, we'll show examples later where it's not a parallel shift. But for now, assume it's a linear amount of damage. So it's a downward parallel shift of the cost curve.

The social marginal cost is what it actually costs society. What does it cost society to produce a unit of steel? Well, it costs the cost of producing that steel plus society loses by the amount of damage done to the fish.

So from society's perspective, the optimal level of production is here. I'm sorry, I did something wrong. Yes, that's what I did wrong. OK, my bad. I'm sorry. Let's go back. Erase that figure. This is what happens when I try to do my own figures. I didn't make a typo. You can get rid of that email, Paul. I should trust my textbook.

OK, let's do this again, shall we? We have a demand, which is the private marginal benefit. We have a supply, which is the private marginal cost. We have a private market equilibrium at Q_1 , P_1 .

Now, I'm sorry, the costs are of course higher due to the externality. The amount of damage D makes a social marginal cost that is higher than the private marginal cost. I of course should have seen that right away. The social marginal cost is higher than the private marginal cost.

Therefore, the optimal social level of production is Q_2 , a lower level of production. Why? Because private producers here overproduce because they do not account for the damage they're doing to society. They stop at the point where their private costs equal the private benefits, but they should be stopping sooner.

They should be stopping with the social costs equal to marginal benefits. As a result, they are creating a deadweight loss equal to this area. That deadweight loss is units that are produced where the social marginal cost is-- I'm sorry, where the social marginal cost-- once again I did this wrong. I'm sorry. I got to be looking at my notes more.

This has got a deadweight loss of-- this is Q_2 . OK, got to look at my graph more. This is Q_2 . So it is creating a deadweight loss of this amount. And how do I know that's wrong? Because what do deadweight loss triangles do? Deadweight loss triangles always point to the optimum and widen outwards from there.

So the deadweight loss here is the triangle that starts here at the intersection of social marginal cost and private marginal benefits and expands outwards to here and ends at the private equilibrium. And that's the deadweight loss triangle.

Here's the trick with drawing deadweight loss triangles, why I immediately saw I was doing it wrong. Remember, deadweight loss infinitely close to the optimum is zero. OK, if you deviate from an optimum by ϵ , it's a tiny deadweight loss, because you're close enough.

The further you move from the optimum, the worse the deadweight loss is. And that's an important trick to remember, that deadweight loss triangle is going to be smallest right by the optimum, and they're going to get bigger as you move further away. Yeah.

AUDIENCE: Would the deadweight loss be the loss you get because the fish are dying?

JON GRUBER: It's the loss to society. Let's be very clear. It's the loss to society because the steel producer overproduces because he doesn't recognize the fish are dying. So it's a somewhat more complicated version of your answer.

You're right in shorthand, but it's important to understand the full economics here. The deadweight loss is the loss to society, because the steel plant does not recognize the damage he's doing, and therefore overproduces steel. And that creates this loss to society. Good question. OK. Other questions about that? I'm sorry about the graph screw-ups, but other questions about that?

OK, so that is how we think about a negative production externality. Because the private actor does not account for the damage they're doing, they produce too much. And that creates a deadweight loss for the reasons that was just expressed in the answer to this question.

Now, on the other hand, externalities do not always have to come from production. We can also have what we call negative consumption externalities. Negative consumption externalities.

So for example, here we want to think about things where your consumption of the good affects me negatively and you don't account for it. The classic example, of course, would be cigarette smoking. So let's think about the market for cigarettes. Here we have the quantity for cigarettes and the price of cigarettes.

We have some private demand for cigarettes, which is the marginal benefit. That's how much people value the next cigarette. There's some cost of producing cigarettes, which is the private marginal cost, the private marginal benefit. And there's some private market equilibrium with Q_1 cigarettes being smoked at a price P_1 . So that is the private market equilibrium.

Now, when you smoke a cigarette, what you do is you impose damage on me. We'll talk about this more in chapter 6. You pose damage on me by making me cough. You pose damage on me by raising my medical bills. There's a lot of negative external effects you impose on me by smoking.

What that means is that the social marginal benefit of the next cigarette is lower than the private marginal benefit. So in other words, what happens is the social marginal benefit is here. The social marginal benefit is lower than the private marginal benefit by the damage you do to me by smoking.

You smoke, so you impose costs on me that you don't account for. Now, this has nothing to do with production. There could be externalities in how we produce tobacco for tobacco leaf. But I'm not talking about that. This is just what arises from your consumption of the cigarette, so it's on the benefit side.

So by consuming those cigarettes, you have a social marginal benefit that's lower. The social benefit is lower than the private marginal benefit. So once again, the optimal amount of cigarettes is lower. And once again, we have a deadweight loss. This time, the deadweight loss is this area.

We have a deadweight loss because you in your decision to smoke a cigarette are not accounting for the fact that it's costly to me. So you overconsume-- cigarettes. So just like the steel plant overproduces steel because it doesn't account for the damage done to the fishermen, the smoker overconsumes cigarettes because they don't account for the damage they're doing to the people from smoking. OK? Yeah.

AUDIENCE: Why is it that deadweight loss [INAUDIBLE]?

JON GRUBER: Did I get it wrong again? It's where private marginal benefits-- yep, you're right. Absolutely. Got that wrong again. Got that wrong and you're right. It's here. It's pointing to the new optimum, which is here. Pointing to the new optimum, which is here. You're right. Good catch.

See, I'm just tricking you all, quizzing you all. Good job. OK? So once again, how do we know it's there? Because once again, it points to the new optimum. Points to the new optimum and it's the gap between where the cost curve is above the benefit curve. OK? Other questions about that? OK, Yeah.

AUDIENCE: Would you also want a new price? So would you want to lower the price? [INAUDIBLE].

JON GRUBER: No, there's no new price. There's no market here. There's no market. This is an imaginary curve. Nobody sees this. So there's no new market here. There's no new price arising. OK? I'm going to come in a little bit to how a market can lead rise to that price. But for now, there's no new price. So there's just the deadweight loss arising.

AUDIENCE: So wouldn't this, I guess, suggest as a consumer that I should actually be paying less for cigarettes?

JON GRUBER: No, once again, this has nothing to do with prices. This is just saying that you are imposing-- basically, your social marginal benefit, you're imposing a cost on others. The only reason you think you pay less is because you're looking at this intersection. But that intersection is irrelevant for prices. It's only for quantities. This isn't a market transaction, so don't worry about affecting prices, just quantities. Yeah.

AUDIENCE: Can a deadweight loss be quantified to the part that's to the firm as well as the one to the consumers?

JON GRUBER: Well, once again, this is a great question. There's no deadweight loss to the firm here because there's nothing in production. Now, there may be production externalities from tobacco leaf, but I'm ignoring them now.

I'm just talking about the damage that you are doing by smoking. Cigarettes appear magically. The production is efficient. But by using them, you are creating a deadweight loss for society. Yeah.

AUDIENCE: How do you draw the D and the arrow on this graph?

JON GRUBER: Well basically, this is the damage. So basically, this D-- sorry, I covered it up here. This D is the damage that I do to you by smoking. So it's just like this D is the value of lost fish. This D is the value of lost health, say, to you from secondhand smoke or the value of your higher tax bills because you pay for more health care.

AUDIENCE: Is it correct to interpret the D on the left like that is the price of the fish that are now gone? And then if so--

JON GRUBER: This is the value of the fish that are now dead. This is the value of the lungs that are now hurt because you're coughing from secondhand smoke, is another way to think of it. Yeah.

AUDIENCE: Why doesn't the damage-- or maybe this is just [INAUDIBLE]. Why doesn't the damage depend on the marginal benefit?

JON GRUBER: I just said we're assuming it's linear. We're assuming it's linear. More generally a little bit later, we'll talk about a case where it's not linear. Right now, we're assuming it's linear. Great questions. Yeah.

AUDIENCE: So in both cases, the deadweight loss is below the social marginal cost and above the social marginal benefit.

JON GRUBER: Basically, it's the gap between the private-- the it's the gap between the private marginal cost and the social marginal benefit. Or in this case, it's the gap between the social marginal cost and the private marginal benefit. The bottom line is there's no rule, except that you're going to have the deadweight loss that's going to radiate outwards from the social optimum and end at the private equilibrium.

Because we're going to see an opposite case in a minute. OK? Other questions? Great questions. I'm glad you guys are engaged. OK, maybe I need to make more mistakes, because this gets people-- you guys are way more engaged than folks usually are here. Maybe it's just a great class. I love it.

OK, now, externalities do not have to be negative. They can be positive. Let's talk about a positive production externality. And an example of positive production externality, remember, is when I produce something that gives positive benefits to you, but I don't get those benefits.

And the classic example of that is research and development. When firms do R&D, the lessons they learn spill over to other firms that can draw on that R&D. OK, maybe not at first while it's secretive, but eventually when it becomes public knowledge.

But the first firm doesn't benefit from that. So here I have a firm. I have a firm that's making whatever it is, widgets. OK, here's the quantity. Here's the price. There is some demand for widgets. That's the private marginal benefit. There are some supply of widgets. That's the private marginal cost. And then some private amount Q_1 P_1 of widgets.

Now, when I produce a widget-- when I produce a widget-- I'm sorry. This is not production widgets. This is production of R&D. This is the market for R&D. So this is the quantity of R&D in widgets. This is the price of R&D in widgets.

The price of R&D, you can think of the cost of a scientist. So it's the quantity of R&D and the price of R&D. Now, the firm privately decides the optimum amount of R&D to do Q_1 . Basically it weighs how much it benefits versus how much it costs to hire a scientist.

What it misses is when it does R&D, it helps other firms. And in particular, the social marginal cost of R&D is actually lower by an amount B than as a private marginal cost. Why is it lower? Because when I do it, I'm also helping other firms.

So the total cost to society-- I think of what does it cost me to hire a scientist? I don't think, but by the way, when I hire a scientist, I'm helping a bunch of other firms too, down the road. I don't think about that. So effectively, the social marginal cost is lower in the private marginal cost.

AUDIENCE: Why would you think of that in terms of cost and not in terms of--

JON GRUBER: Great, great question. This question I do get every year. It's a hard question. Because this is the consumption side of the market. This is the production side of the market. Costs and benefits are confusing.

But think of this as the consumption side of the market, which is how much R&D you're consuming. This is how much R&D you're producing. And you're not producing enough R&D. Now, it's harder, because the firm consumes and produces the R&D together. It's a joint decision. So it's really just in this sense, you got to lean more on the technical thing, which is basically it's through the production of R&D that this externality happens.

As a result, the social optimum of R&D is higher. Let's come to your question of which curve is lower or higher. In this case, the socially optimal level is higher. And as a result, where the marginal cost is above-- where the social marginal cost is below the private marginal benefit, we get this deadweight loss.

We get this deadweight loss because in these units, the private marginal benefit is above the social marginal cost. So firms should be doing that R&D. But they don't because the private marginal benefit is not higher than the private marginal cost.

So you take a unit like here, OK? You take a unit there. The firm says, wait a second. At that unit, my private marginal cost is higher than my private marginal benefit. I'm not going to do R&D.

But what it should say is, my private marginal benefit is above the social marginal cost, so I should do R&D. So it's not doing R&D it should do, creating a deadweight loss from too little R&D. OK? Questions about that?

See? This fits my-- I do empirical modeling. See, I made mistakes, you guys asked a bunch of questions. I didn't, mistake you didn't ask questions. So my theory that my mistakes lead to questions is proved empirically. By the way, that's macroeconomics. Sorry, for my macro friends. All right. Oh wait, question. Go ahead.

AUDIENCE: I have an understanding that when we're talking about whether it's positive or negative production externalities, when we have a production externality we're moving the supply curve. It seems like when we have a consumption externality, we move the demand curve. I can see that, but it doesn't entirely make sense to me why that is.

JON GRUBER: Well, basically because production is about supply and consumption is about demand. And you're right, it's confusing in cases like this because we think about the benefit of R&D, yet we're moving the cost curve. But once again, it's about something you produce, so it's going to be a production externality.

Now, let's come to a final example, which you can guess what it is. It's a positive consumption externality. And here, let's not do a graph, because I won't always do graphs. Sometimes I'll do math and examples. Let's do an example here. OK?

I have a neighbor. Actually, he moved away. But I had a neighbor I was not particularly fond of. I might have, in previous classes, called him my asshole neighbor. But he's moved away, so let's not speak badly of him now.

I had a neighbor who I was not particularly fond of. That neighbor had a habit of getting about halfway through projects and then stopping. And in particular, probably 20 years ago, he got halfway through a major landscape renovation, then stopped.

And he stopped at the point where there was a gigantic dirt pile directly outside our kitchen window, which he then left there for 20 years. Now, this was not great for us. But imagine his thinking. Imagine that to remove that dirt pile would have cost \$1,000.

So that's the private marginal cost of removing the dirt pile. Now, since it was actually an awkward angle for his house, he couldn't see it very well. To him, the benefit of reducing that dirt pile was only \$600. So the private marginal benefit was only \$600.

So he evaluated it and said, I'm not going to remove the dirt pile. It's a shame I didn't finish my project, but what do I care? I'm not going to pay the grand. That's a net loss to me. What he missed was that there's also a \$500 benefit to me, which is the additional benefit.

So that the total social marginal benefit equals \$1,100. So in fact, it would be socially efficient to remove the dirt pile. But he doesn't, because it wasn't privately efficient. Therefore, you get under-consumption of dirt pile removal services. Under-consumption of dirt pile removal because he's not accounting for the social benefits of doing so.

So I realized the graphs were tricky, even for me. The main thing is the intuition. If you got the intuition, everything else will come. The main thing is intuition, which is a negative externality is going to lead to too much of something and therefore deadweight loss from too much production. A positive externality is going to lead to too little of something, and therefore deadweight loss from too little goods being in the market. Yeah.

AUDIENCE: I was wondering if you could clarify why number four is a positive externality if you aren't receiving the extra--

JON GRUBER: Oh, that's a great point. But what am I receiving? A better view. So it's worth-- the point is I'm not getting the \$500. Great question, actually. Remember when we learn about welfare in 14.01. We talk about the fact that we can't really measure utility. So what do we do as a trick?

What we do is we use a compensating variation measure, or we use a revealed preference measure. Which is we say, well, I can't measure how sad I am to have the dirt pile. But you can ask me, how much would you pay to remove the dirt pile?

So this is what I would pay. So if this was over \$1,000, then I could just remove it. Because if it's than \$1,000, I'm not going to remove it. But added together, it'd would be worth it to have that thing removed.

So this is a valuation of my utility. We can't measure it. We don't measure utility. We measure the equivalent variation, the dollar value of that utility. And that's how we do welfare analysis. Good question. Good reminder about a lesson 14.01. Other questions about this?

Now, you may look at this and you may say-- oh, I actually want to talk about one example. One of my favorite examples of negative externalities-- because it's really just something we don't appreciate and incredibly pernicious-- is SUVs.

SUVs are a classic example of a negative externality. Why? Because SUVs-- well, you tell me. What are the ways in which SUVs impose negative externalities on society, above and beyond what people pay? Yeah.

AUDIENCE: They use more gas, they're higher--

JON GRUBER: Well, hold on. They use more gas. And use more gas is bad because it's bad for the environment.

AUDIENCE: They're higher up, so [INAUDIBLE].

JON GRUBER: They are much less safe.

AUDIENCE: They also take up more space.

JON GRUBER: I don't know that space-- I mean, I don't know how big that deal is. But there's one other thing. Yeah.

AUDIENCE: They're heavier, so they damage the road more.

JON GRUBER: Yeah. It turns out-- in a fun MIT-ish fact-- turns out road damage goes up exponentially with weight per axle. So if you have more weight on a given set of axles, road damage goes up exponentially.

And in fact, that's happened. So in 1985, the typical car weighed 3,200-- you don't need to memorize this, just fun facts. Typical car weighed 3,200 pounds and the biggest car on the road weighed 4,600 pounds.

Now the typical car weighs 4,000 pounds and the biggest truck on the road-- this is pre-Cybertrucks, I know what the hell they weigh-- is 6,600 pounds. We have massively- fuel economy, the typical SUV gets about 2/3 the fuel efficiency of a sedan, not even of a compact.

So they're inefficient. They consume more gas. That's an externality through causing global warming. They're Bangladesh killers. OK? Second, they do more damage to the road because road damage goes up with weight per axle exponentially. And third, and actually most interestingly, they're incredibly unsafe. And why are they especially unsafe? Because they're designed where their grills at exactly the height of people's heads.

So actually, if you look at traffic deaths in the US, we see a dramatic reduction-- the entire world has seen a dramatic reduction of driver deaths from traffic accidents-- driver and passenger deaths. Dramatic reduction, because cars have gotten safer. But only in the US that has been offset by a dramatic increase in pedestrian deaths. Only in the US.

So while the rest of the world is seeing traffic-- among developed countries-- while the rest of the world is seeing traffic deaths fall in general, the US, they've fallen a little bit, but not nearly as much as the rest of the world because our savings in driver fatalities are being offset by pedestrian fatalities because we have these grills at exactly head height and they're killing people.

Moreover, what's fascinating about this is it's what we call a dynamic externality, because since you're driving a big car, I got to drive a big car. So as this dynamic effect of reinforcing itself that when somebody's driving an SUV, you want to drive an SUV because you don't want to get hit by this SUV. So that's basically a perfect example of a negative externality.

Now, if we go back, let's go back to my neighbor's dirt pile. Now, you might say-- you hear that example and you're like, well Jon, this is a course about the government. But that seems like you could solve it privately.

Just go to your neighbor. It's like, look, it's \$1,000 for this. I know it's worth \$600 to you. I know it's worth \$500 to me. Let's just pool together and let's just get rid of this dirt pile.

And indeed, there's a famous example of applying this, which is what's called the Coase theorem, named after the economist Ronald Coase from the University of Chicago. He said, look, let's go back to our first example.

Let's go back to the steel plant. And let's look at that example. And he says, what if the fishermen owned the river? What if fishermen owned the river? They would march up to the steel plant owner and they'd say, look, steel plant owner, you need to produce less steel, because you are-- and by the way, we'll come later to pollution reduction technologies.

Imagine there's no pollution reduction technology. The only way to reduce the sludge is to produce less. We'll come back to pollution reduction technology. But for now, imagine the only way to reduce the pollution is to produce less.

The fishermen would come and say, look, we own the river. We will not let you dump unless you reduce the amount or unless you pay us the damage that you're doing. So they say look, steel plant owner, we don't want to mess with your production process. You know what's best. But we want you to pay us D for every unit of steel you produce to compensate us for the fish that are dying.

Well, think about what that does. Think about what that does to the steel plant owner. So we've got, once again, the market for steel-- quantity, price. We've got our private marginal benefit. We've got our private marginal cost. We've got our social marginal cost. We've got the private market quantity and the optimal public quantity. OK?

And we have some damage D that the production of steel is doing. We actually have our handout now. Thank you, Valerie. Valerie biked all the way back to the department. So now we can go to figure 5-5. Here I have a marginal damage of, say, 100.

Now, so what if the marginal damage was 100? And what if the fishermen said to the steel plant owner, well, you got to pay us 100 for every unit of steel you produce. What does that do? Well, that internalizes the externality.

What mean by that is the steel plant owner now has a higher private marginal cost. It's the original marginal cost plus the \$100 they have to pay. And guess what, that new private marginal cost, which in figure 5-5 we call PMC_2 , is equal to the social marginal cost.

So the new private marginal cost is equal to the social marginal cost. And there's no more externality because the steel plant now produces optimally at Q_2 . OK, once again, the steel plant has to pay the fishermen \$100 for every unit of steel they produce.

That equates their new private marginal cost with the social marginal cost, and there's no more overproduction of steel. This indeed is the proof-- is a graphical proof of what we call the Coase theorem.

The Coase theorem states that as long as property rights are well-defined-- as long as you have well-defined property rights-- property rights are well-defined, and as long as you have costless and efficient bargaining, if those conditions are met, then there will be no externalities.

Externalities will all be internalized through the market mechanism. Unsurprisingly, this guy was in Chicago. If you know the history of economic thought, there's a big-- Chicago's very pro-market, Cambridge is sort of more pro-government. That's kind of the debate.

So this is very much within the Chicago school of thought that, in fact, you don't need the government here, because in fact, there's no externality. They just internalized it and we're done.

OK, so that is Coase. Actually, Coase theorem has a part two that's kind of cool. Coase theorem part two-- that's part one. Part two says it actually doesn't matter who has the property rights.

Note, I didn't say who has the property rights. I just said they have to be defined. So imagine the steel plant owner owned the river. Well, then the fishermen could come up to him and pay him \$100 not to produce the unit of steel. And we get the same equilibrium.

Now, of course, they wouldn't pay him 100. They'd want to pay him 99.9. Of course, he wouldn't pay them 100. You'd pay them 99.9. So I'm using 100 to round. But roughly speaking, the point is, as long as property rights are well-defined, whoever owns the property rights will get the benefits.

It will determine distribution. In the first case, the fishermen are better off. In the second case, the steel plant is better off. But in both cases, you get efficiency. And that's, once again, an example where you have to keep your mind separated between efficiency and redistribution. We're not talking about winners and losers here. We're talking about efficiency. Yeah.

AUDIENCE: What specifically does the word well-defined-- or what does it mean--

JON GRUBER: Somebody owns the river. There's an enforceable contract that you own the river. That if I do not get permission to use the steel plant to dump in the river, and you do, you will be shut down or you'll be subject to some extreme fine.

AUDIENCE: Sure, so the examples of some sort of shared ownership or government ownership, do they fall under that? No. Well, I mean, that's a complicated case. But the bottom line is you need property rights to be well-defined.

JON GRUBER: So Coase's point is, just because there's this externality doesn't mean you need the government. As long as you just define property rights and you allow people to bargain, they'll figure it out.

Now, this is an example of a trick we'll use several times this semester, which is to lay out a theory which has a lot of good insights but is clearly wrong in application. The Coase theorem is very valuable for thinking about lots of things. It also almost never actually works.

And why does it not work? Well, there's a number of problems with the Coase theorem. The first problem is what we call the assignment problem. The assignment problem is that basically when you have many bad actors and many people being hurt, you don't know how to assign blame or how to assign benefits.

So let's say there's six steel plants along the river, and you don't know which steel plant sludge is killing your fish. How do you know who to charge? Or let's say there's three pods of fishermen. How does each pod of fishermen get paid?

So whenever there's multiple players, you have an assignment problem. There's also the problem that we don't know how much the damage is. The fishermen could go to the steel plant owner and say you're doing \$100 of damage to fish. How do we know? How is the steel plant owner supposed to verify that and know?

So basically, the first problem is we don't know how to assign the benefits and the costs in a way that will make it seem easily theoretically here. Think about how this gets worse with climate change.

Remember, every single unit of carbon emitted anywhere in the world has an equal effect on climate change. So how could you possibly negotiate with every single person emitting carbon at every point and every single person being damaged, and some being benefited? Right?

Duluth, Minnesota advertising itself as the climate refuge of the world. We're now are growing grains in Siberia that were never possible before. Some folks are benefiting from global warming. Most are losing. How would you possibly arrange that negotiation? It's implausible. So that's one problem.

The second problem is what we call the holdout problem. OK, imagine that we've solved this assignment problem. We figured it all out. We know what's going on. And let's say that we've decided that there are 100 fishermen, and each unit of sludge kills \$1's worth of fish for each of the 100 fishermen. Thus, the \$100 damage.

Now, the fishermen go up to the steel plant, and they say you have to give us each \$1, or you can't use the river. We have property rights on the river. You can't use the river.

Well, let's say Enoch here is the 100th fisherman. Now, the other 99 have gotten their dollar. Now he's the last fisherman. The steel plant can't produce unless he also says yes. Is he going to ask for a dollar? No, he's going to ask for \$50. He'd be like, look, you've already given away \$99. I have all the power now. You have all the negotiating power. You're the holdout.

Well, working backwards, the first fishermen should see that coming and should also not settle for \$1, and the whole thing falls apart. Classic problem with negotiation is this holdout problem. OK?

The flip side of that is what we call the free rider problem. And this is a term we'll come back to a lot this semester, the free rider problem. Which is that, what if we said, OK, fine, that holdout problem is going to be bad. So let's say the steel plant owner had the problem.

OK, let's forget Coase theorem part two. What about Coase theorem part one? What if the steel plant holds the property rights? Would that solve the problem? Now, the problem if the steel plant holds the property rights is that basically-- so let's say they have the rights and they agree to reduce production by one unit for every \$100 they get from the fishermen.

Now basically, suppose that they agree that-- let's say the optimal reduction is 100 units, that the social optimum is they produce 100 units. So they're going to get \$100 from the fishermen for each unit. They're going to get \$10,000 from the fishermen.

So each fisherman goes and pays his \$100. OK? Well, now you're the last fisherman. Enoch again, he's the last fisherman. All the other fishermen have paid \$100. He comes up and he says, well look, if I don't pay, you still reduce 99 units. And I'll get 99/100 of the benefit. So why should I pay you?

I'm a free rider, because it's a private cost and a social benefit. When the steel plant reduces dumping sludge, all the fishermen benefit. As a result, why should any one fisherman pony up? They should try to get all the other fishermen to pony up. So that same logic again unravels and no fisherman will pony up. OK? Yeah.

AUDIENCE: So would the last fisherman would say?

JON GRUBER: So the way I set this up, the output needs to reduce 100 units of steel. So each fisherman should pay you \$100 to reduce the unit. So they all get in a line. They each give you \$100.

The last fisherman says, well, 99 people have given you that money. I know you're reduced by 99 units. If I give you my 100, you reduce by the 100th unit. But I only get 1/100 of the benefit of that, whereas I bear all the cost. So I'm not going to give you the money. And that's the free rider problem.

And then finally, we have the problems with this other assumption, which is bargaining is not costless and efficient. So we have bargaining problems. I already talked about one kind of bargaining problem, which is, how you get all the people affected by climate change together?

But let's just think about my neighbor and I. In fact, it'd be super fucking weird if I went to him and said, well, it's worth \$600 to you and I'll give you \$500, it's \$1,000-- I mean, he'd punch me in the face.

Think about it. Think about a classic externality, which is the person next door to you plays their music too loud. You do one of two things. You either ignore it or you complain and make them shut it off-- tell them to shut it off.

You don't calculate the loss in lifetime earnings and say, well look, if you want to pay it, you've got to pay me \$6.63. Basically, that's just not how the world works. There's a famous apocalyptic tale of an economist on a plane-- who shall remain nameless-- who offered the guy next to him \$10 to shut up so he could get some sleep.

That's apocalyptic, because it probably didn't happen. So the point is, the Coase theorem is great in theory, and it's a valuable insight in theory. In practice, that's just not how things work. OK? Questions about that. Yeah. Oh, go ahead.

AUDIENCE: Why is the holdout different from the free rider?

JON GRUBER: It's basically which side of the problem you look at. I mean, Enoch's the bad guy in either example. But the holdout problem is when the property right owner can hold out and basically say, I'm not going to give you this property right unless you pay me the most.

The free rider problem is when the other party has the ownership, and you say, well, if most of my friends pay up and they do the right thing, I'm going to benefit from that. So why should I pay up? So it's really which side owns it.

So it's really fighting with part two of the Coase theorem, saying that the side of the market matters, but it screws up no matter which-- it matters who owns it. It screws up either way. There's a question here. Yeah.

AUDIENCE: You were saying that actually they wouldn't pay 10, but 9.9999. Why is that?

JON GRUBER: So this is a great question. So basically, I assumed that it worked out that you paid exactly 100. But how does that work out? I mean, if the fishermen own the river, why don't they charge a steel plant 150? Why should they settle for 100?

And in fact, I don't have a good answer to that, other than if there were multiple groups of fishermen, they'd bid against each other and drive it down to 100. So in fact, in this example, I don't have a good answer for that.

In fact, they probably would ask for 150. So basically the point is, I'm assuming we end up at where efficient negotiation would end. That's what I'm covering there. Other questions? Yeah.

AUDIENCE: So I understand both the assignment problem and the bargaining problem-- well, part of it comes down to that it's hard to quantify damage. Could you a little bit elaborate on how they both differently?

JON GRUBER: Yeah. So the assignment problem is a technical problem, which is A, basically, it's hard to measure damage. And B, it's hard to compensate the losers, practically, when there's global warming. The holdout problem-- wait, which is the other one?

AUDIENCE: Bargaining.

JON GRUBER: Oh, the bargaining problem is that basically, it's related to the measurement. But it's also just that literally the conditions of bargaining get very tense and people don't act rationally.

My \$100 example, even if the fishermen-- even if it's \$100, the steel plant may say to the fisherman, I'll give you \$110. The fisherman might say yes to that. They may say, screw you, I hate you. You've poisoned my water forever. I want \$150. I'm not doing it. OK? That's an example of bargaining problems. People are not Homo economicus. They don't behave hyper-rationally. Yeah.

AUDIENCE: For the example for the holdout, that was when we had several steel plants along the river?

JON GRUBER: No, the holdout example was basically when there are multiple fishermen, still one steel plant. There's multiple fishermen, and the holdout problem is the steel plant needs permission of every fisherman to pollute, and mean Enoch here won't give him permission. Yeah.

AUDIENCE: Are there cases in the real world or markets, I guess, where the Coase theorem is like pseudo correct?

JON GRUBER: Oh, yeah. No, this is a great question. When we teach these theorems, the Coase theorem, others we'll talk about, they definitely have directional effects. There are certainly times where there's private negotiations that can happen and can move you part of the way there.

But the right way to think about the Coase theorem is, when's it going to work the best? When there's few parties, when information is transparent, and when they have a good relationship, then it can work. Yeah. Other questions? OK.

So we've just answered the when question. When does the government intervene? When the private market fails. I've just shown you the private market fails. So now we have turned to the next question, which is, how might the government intervene? OK?

We've answered the first question in public finance, when does the government intervene? Now we have to turn to the second question in public finance, which is, how might the government intervene? So now we're going to talk about public sector solutions to externalities.

All right. Now, the first type of public sector solution we'll talk about is-- so now we're going to talk about public sector solution. So we're on the second question in public finance, which is the how.

The first one we'll talk about is corrective taxation-- corrective taxation as a solution to negative externalities. How does corrective taxation work? Well, let's go to figure 5-6.

Back to steel market, back to our sludge-producing steel guy. Once again, he wants to produce Q1. Society wants to produce Q2. He's going to overproduce. Coasian bargaining breaks down so he can't get it that way.

What if we imposed a tax on him, which was exactly equal to the marginal damage he's doing, or \$100 in the example I gave a few minutes ago? What if we imposed \$100 tax per unit produced? Well, that would raise his marginal cost by \$100 and have the same effect of internalizing the externality.

Not through the private market, which involves negotiation, but just through government fiat. We would say, you have to pay-- and once again, I'm assuming there's this linear damage. More complicated than the real world, and we'll get to that.

We would say, look, you're doing this linear damage of \$100 of a unit you produce, so I'm going to tax you \$100 for every unit you produce. So that's going to create a new private marginal cost, which is equal to the social marginal cost.

So PMC_2 is equal to PMC_1 plus the marginal damage. Essentially, if I want to produce a unit of steel, I got to pay all my workers, plus I now got to pay a tax to the government. And that government tax is set such that we end up at the optimal equilibrium. Done.

It's like Coase, but the government is doing it, so you don't have to worry about assignment problems. All you have to worry about is measuring the externality and you're done. You measure the externality, boom. You set the tax. Yeah.

AUDIENCE: So in the case of pollution, how is marginal damage even calculated? Is it based on how much everyone else emits too, or is it just based on [INAUDIBLE]?

JON GRUBER: Awesome question. Calculating marginal damage is really-- that's very hard by itself. Now, of course, that doesn't make Coase any easier, because then you've got to put it in a private negotiating framework. But even this is not so easy. What's marginal damage? Is it linear? How do you measure it? I'm going to come to that in a little bit.

Now, what about positive externalities? Well, there's an equivalent for positive externalities, which is a subsidy. Which is, you could just say-- in figure 5-7, we've got our R&D example. We can just say, look, firms are producing Q_1 R&D. That's too little, because they're not accounting for the positive benefit on other firms for producing R&D.

So we give them a subsidy, which is equal to the marginal benefit, which effectively lowers their private marginal cost, causes them to produce more. So it's like a tax raises their private marginal cost, causing them to produce less. A subsidy lowers their private marginal cost, causes them to produce more, and move into the optimal level of production of Q_2 .

So this is using what we call the price mechanism. OK? A great example of the price mechanism in the real world that's been in the news lately is what we call congestion pricing. And you might have read-- this has been in the news a lot because of a debate in New York State.

Economists since the dawn of time have said that, look, traffic is a classic negative externality. When you drive, you make my drive a little bit longer by clogging the road. The estimates are that the cost of traffic to the US, if you look at people's wasted time-- we'll talk about how we value time later-- is well over \$100 billion a year.

The idea is, well, that's an externality. By you driving, you're not only putting Bangladesh underwater. You're making my commute longer. So as a result, you should pay for that externality. You should pay for driving, in particular in areas where the traffic's bad or at times when the traffic's bad.

Thus, the idea of congestion pricing. Congestion pricing has been used in Europe successfully in London and in Stockholm. In London, if you go into an inner city, into parts of London during certain hours, there is a fee that's imposed on you. You have your little electronic thing and they ping you.

And they charge you a fee for driving in parts of London at certain hours. Same in Stockholm. These have been shown to be incredibly effective. They've lowered traffic, but not only that, they've improved health. Why?

Well, first of all, there's fewer accidents. Second of all, there's less asthma among the kids because there's less car pollution in the cities. Pretty awesome indirect effect. Literally they're showing that kids' health improves with congestion pricing because there's less pollutants in the air and the kids are less likely to have asthma. Pretty awesome example of an unintended benefit. And of course, there's less greenhouse gas being emitted, which benefits all of us in terms of slowing global warming.

Now, New York City had a mayor called Bloomberg. You might have heard of Bloomberg. Bloomberg is like Mr. Externalities. He totally loves this stuff. And he actually proposed congestion pricing. And after many decades, it actually made it through the debate.

For years they wanted to pass it, and they couldn't get the support, because who wants a new tax? What they did is they did a very clever thing. They said, look, our subways are terrible. We'll take all the money from the congestion tax and put it into improving our subways. And it's a win-win.

We get less driving and better subways. And it was almost going to be signed. And then the, I'm going to say it, cowardly governor of New York bagged it at the last minute because she was worried about the blowback, which is really, really unfortunate. This is a clear example of a win-win.

But we see examples of this where there's policies that really only an economist could love don't often happen. And this is a classic thing where you've got a lot of opponents, people from my home state of New Jersey, who don't want to pay to commute into New York make a stink. And this gets canceled. So that's unfortunate. And hopefully New York will revert-- will see the error of its ways and reverse in putting congestion pricing. Yeah.

AUDIENCE: Why did it need to go through the state government and not just city level?

JON GRUBER: I don't know the logistics of the New York law, but for some reason the governor had a veto. The legislature passed it. Somebody had to go to the state level. And I don't quite know. Maybe because it's inter-state, because people from New Jersey would pay it. I don't know. Good question.

All right. So now, you might be thinking to yourself, well, this is all interesting, kind of cool. But if we want Q2, why don't we just say, produce Q2. Why don't we just regulate?

That is, instead of using a corrective tax or the price mechanism, why don't we just regulate? That is, if you go back to figure 5-6, we could put in this convoluted tax and measure. Someone said figure out what the marginal damage is.

We say look, Q2 is what you should produce. Produce Q2, done. None of this fancy economic stuff, none of this politics over taxes, just say produce Q2 and we're done. And in fact, that would have a very similar impact.

If you just did Q2, you'd be done with it. And in fact, that is the approach we've typically taken in environmental regulation in the US historically, has been quantity regulation, what's often called command and control regulation. Which you just said, look, you're producing too many pollutants. Produce less, boom. Full stop.

So what we want to turn to next after one detour is talking about why would you use price versus quantity regulation? When would they make sense? Now, to do so, we need a new framework. Before we do so, I want to talk about one other thing that is very important as a historical lesson, which is we talked about corrective subsidies.

We talked about corrective subsidies when you have underproduction, say underproduction of R&D. And there's no better example of the power of corrective subsidies than Operation Warp Speed. Now, when COVID came along, we realized we needed a vaccine.

Vaccines are a classic example of a positive production externality. And why is that? That's because they're often not very profitable, or the profits are highly uncertain. Take, for example, the vaccine that was going to be developed for Ebola.

Ebola struck West Africa in 2013. Immediately began to work on a vaccine for Ebola. It took them six years to make it, which is actually very fast in vaccine time. And by the time they did, Ebola was gone. No one wanted it.

They'd invested billions into making this vaccine and it just got wasted. No one wanted the vaccine. So companies don't invest in vaccines. You invent a cancer drug, you know cancer is still going to be there in six years. Vaccines, people aren't sure if the demand is going to be there, if people are going to want them. So there's massive under-investment in vaccines.

So what did the government do when COVID came along? They set up Operation Warp Speed. The Trump administration, to their credit, set up Operation Warp Speed. What this did was two kinds of subsidies, we call them push and pull.

The push incentive was the kind that we have in figure 5-7, which is literally we subsidized R&D. We gave billions of dollars to companies to do R&D into vaccines. The second, which is not really depicted on this graph, is what we call a pull incentive.

The pull incentive was we guaranteed, no matter what happened, we would buy them. The government said, even if COVID goes away, we guarantee we'll buy a billion doses. So we've taken the uncertainty out of your decision.

Due to those two incentives, the COVID vaccine, which all experts predicted would take 10 years, took 10 months. It's unbelievable. It's truly a modern miracle how fast we got the COVID vaccine. It's all a blur in the rear view mirror.

But if you were talking to people like me early in the COVID, I would have said, well look, if we're lucky, we'll get it in five or six years. The notion of doing it in 10 months is just beyond belief.

And the vaccines probably have saved-- the best estimates are two million lives have been saved worldwide by the COVID vaccine. It's amazing. And if we didn't have anti-vax movements of the kind we discussed a bit last time, it would be even higher.

So basically, this is a classic example of the power of corrective subsidies. Now, that detour aside, let's go back and talk about how do we think about quantity versus price approaches to correcting externalities? And to do so, I'm going to have to freak you out a little bit and switch the framework.

So we're now going to move to a framework like figure 5-8. In figure 5-8, we're going to change the example. In the previous example, the only way that I could pollute less was by producing less steel. But of course that's wrong.

I can do other things, right? I could put a grill over-- just picture the Lorax. How many of you guys have read *The Lorax*? You can picture Lorax speak for the trees. You could see the plant and stuff dumping in the river.

I could put a grill over it. Or I could do something which actually reduces the pollution. I wouldn't have to actually slow my production of steel. So that's a more realistic case. And that more realistic case is depicted in figure 5-8.

In figure 5-8, we have the market for pollution reduction. It's a weird market, because usually markets are for positive things. There's a market for pollution reduction. What do you mean by that? On the x-axis is the amount of pollution reduction a firm does.

So as you move to the right, the firm is doing more pollution reduction. As you move to the left, the firm has more pollution. So to the left is more pollution, to the right is more pollution reduction. OK?

On the y-axis it's just dollars, costs and benefits. We'll come back to that. So in other words, if you look at-- we talked about an example where there was a constant marginal damage done by pollution. Well, in this framework, that becomes the marginal benefit of pollution reduction.

The marginal damage done by pollution is the marginal benefit of pollution reduction. For every unit of pollution you reduce, you save \$100 worth of fish. So that red line is the benefit of pollution reduction, the marginal benefit, every unit of pollution. And once again, I've assumed it's flat. It doesn't have to be. That's the benefit.

What is the marginal cost of pollution reduction? OK, well, here I'm going to make an assumption. I'm going to assume that the more pollution reduction you do, the more expensive it is per unit to do it. This is sort of intuitive.

When you're doing sludge removal, it's pretty inexpensive to just put a grate over the pipe and you maybe cut your pollution by 20%. That's cheap. But if you want the pollution all the way to nothing, you've got to change the entire power plant and production process. That's expensive.

So what we assume is that the marginal cost of pollution reduction is rising. An assumption, but a sensible one. So at zero pollution reduction at point A, there's zero marginal cost. As you reduce pollution more, the blue line-- which is pollution reduction. And that's the private marginal cost of pollution reduction. We'll also assume it's the social marginal cost, that there's no externalities in the market for pollution reduction.

So we're going to call this the social marginal cost of pollution reduction is the blue curve and the social marginal benefit of pollution reduction is the red curve. So now let's ask, what is the privately optimal amount of pollution reduction?

Well, the privately optimal pollution reduction is where the private marginal benefit of pollution reduction equals the private marginal cost. Where is that? What is the private marginal benefit of pollution reduction? Someone raise their hand and tell me. Yeah.

AUDIENCE: Zero.

JON GRUBER: Zero. So the private optimum is at A. That's where the private marginal benefit equals the private marginal cost. They don't do any pollution reduction. Why should they? They don't care about the fishermen.

But the social optimum is where the social marginal benefit equals the private and social cost, which is at point B. So the social optimum is R star of reduction or P star of pollution. That's the social optimum.

So it's the same logic as before. We've just changed what we're looking at to have the more realistic case of pollution reduction. We don't really talk about we talk about pollution. You want to be looking to steal or you want to look at pollution reduction market. That's what we're looking at here. OK, questions about that? I realize I've changed things up on you. Yeah.

AUDIENCE: Is the social marginal benefit flat just because it's in that specific example?

JON GRUBER: In that specific example, I've assumed it's flat. I guess next time I'll get to examples where it's not flat. Now, in this case-- yeah, I'm sorry.

AUDIENCE: Yeah, why is the social marginal benefit-- why does it not intersect the marginal cost curve as high as possible?

JON GRUBER: Oh, that's a great question. So in other words, why is the red line not higher? Well, the red line is not higher because that literally is the social marginal benefit. So in other words-- actually, that's a great segue to what I want to say next. I'll answer your question in a second.

I want you to contrast two approaches to government policy in this diagram. One approach is regulation. Regulation would say, look, tell them they got to produce R star. Done, easy. The other is a tax.

Imagine I set the tax at \$50. Imagine I said that firms had to pay a tax of \$50 for every unit of pollution that made. They would choose a point between A and B, because the marginal cost-- they would choose a point where the marginal cost of pollution reduction is 50.

If pollution reduction was less than 50, they'd do it. But if it's more than 50, why do it? They'll just pay the tax. Now let's come to this question. The actual cost is 100. Why do I not set the tax at 500?

Why am I only setting the tax at 100? Look, why not just say, look, pollution is bad. Set the tax of 500. Why is that not optimal? That's a sort of version of your question. Why is that not optimal?

AUDIENCE: Because then they'll produce less, and it will be inefficient.

JON GRUBER: Because you're missing the fact that steel is valuable. And pollution should almost never be zero. This is the thing that has been a slowly-learned lesson by environmentalists over time.

I'll talk more about this next time. But pollution should almost never be zero. And we could make pollution zero today. OK? No lights, no heat, no cars, no production. Pollution zero. That isn't optimal, because we derive value from lights and heat and cars.

You only want to tax them to the extent it's causing external damage. And that's why the red line-- that's why you want it fair for the external damage. And that sets the optimal tax. If you set the tax at 100-- one second-- if you set the tax at 100, what do you get? You get the same thing as you get from the quantity regulation, reduction of R^* . Question.

AUDIENCE: So if you get into [INAUDIBLE] quantities, which is a decision between going from another thing-- deciding whether you want to--

JON GRUBER: OK, great. So that's the question I want to ask. That's the question-- I've just showed you first how in this simple example they're the same. Now, let's talk about why they might not be the same. Let's move to an example with two firms.

Figure 5-9 is complicated. So let's walk through it together and ask me questions OK. We have two firms, A and B. The difference is that firm A has-- the difference is that they have different cost curves.

OK, firm A has a higher marginal cost of reduction than firm B. So firm B, to reduce by 50 units, costs them 100. That's a point X. Firm A is a lower marginal cost of reduction. Lower marginal cost of reduction. Firm A is more efficient at reducing.

So they, for, \$100 can get rid of 150 units. Firm B at a cost of \$100 can only get rid of 50 units. Do you see that? So look at the marginal cost curves. Firm A's marginal cost curve is to the right.

It is a more efficient-- in other words, the other way to think about this is literally firm A can reduce pollution by 50 for nothing. Their X-intercept is at 50. So firm A is a more efficient pollution reducer.

Now, let's say that we decide the optimal pollution reduction is 200. Let's say we start in a world where there is 400 units of pollution. With no government intervention, you get 400 units of pollution. Let's say that's our starting point.

And let's say that we decide that the optimum-- so that's the private, that's Q^1 . We decide the optimal amount of pollution is 200. So that's P^1 is 400. P^2 or the optimum we decide is 200. So we want to reduce pollution by 200. We want a reduction of 200. Once again, stop me when this isn't clear. OK? Want a reduction of 200. How do we do that? Yeah.

AUDIENCE: Is that point Z on this graph?

JON GRUBER: I haven't gotten to points yet. I'm just giving you the example. OK, how can we do that? Well, one example is we can say, well, look, we know we want a reduction of 200. There's two firms. Let's make them each do 100. That would be the quantity approach. We know we've got 200 to reduce, two firms. Let's make them each do 100.

What would that result in? Well, that would result in firms ending up at the points MCA, 100 and MCB, 100. If firm B has to reduce by 100, they end up at the point MCB, 100, with a very high cost of pollution reduction. Firm A ends up at MCA, 100 with a low cost of pollution reduction. Why is that inefficient? Yeah.

AUDIENCE: Because firm A should be producing more.

JON GRUBER: Yeah, firm A can-- this is inefficient because firm A is better at this. So they should do more of it. So here, the quantity regulation doesn't give you the best outcome.

Compare that to price regulation. Imagine we said that we are going to charge-- we are going to charge firms \$100 for every unit of pollution they produce. OK? We're going to charge them \$100.

Well, in that case, firm B will say, look, at \$100, I'm going to reduce by the amount X, 50 units. Why am I going to do that? This comes to the question here about why it's the red line, not the top.

Because if it's above 50 units, it costs more than \$100. I'll just pay the damn fine. Below 50 units, I want to reduce. Firm A says, I'll reduce 150 units. Why? Because that's where their marginal cost equals the marginal benefit.

Guess what-- that is up to 200 units. You're done. So by setting the tax-- and that is efficient. That's efficient because every firm is spending the same amount to reduce pollution. Once again, we're forgetting redistribution.

And I'm not talking about A better than B. It's all about efficiency. The efficient outcome is where every firm is paying the same marginal cost of reduction. And you get that with the tax. You don't get that with the crude quantity regulation. Unless-- how could I solve the quantity regulation problem? You know the answer.

AUDIENCE: Cap and trade.

JON GRUBER: You could trade. So now let's get really funky. Let's say that I say I want to reduce pollution by 200, but I don't care who does it. I know I want to reduce it by 200, but I don't care who does it.

How do I do that? Well, what I do is I say, look, firm A and B, I'm going to give each of you 100 permits for pollution. At the end of the day, any pollution you do which you cannot hand me a permit for you will pay-- I'm sorry. Yes.

Let me go back. You can only do the amount of pollution you have permits for. So I'm issuing 200 permits because that's the optimal amount of pollution. I can give them away any way. I'm assuming I split them 50/50, but that doesn't have to be true.

We'll talk more about this next time, how this was proposed in practice. So I have 200 pollution permits. I'm going to give them away, say, 50/50, each get 100. And I'm going to say, you guys can trade them if you want. What happens there?

Well, our starting point is they don't trade. If they don't trade, they each can only pollute 100 units. And you end up back at MCB and MCA. Now, at that point, what should happen?

Well, at that point, B should go to and say you know what, I'll buy your permit, because look, it's only costing you like \$50 to reduce at that point. It's costing me like \$200. So as long as I can buy your permit for between 50 and 200, you should be happy to sell it to me and I should be happy to buy it. So I'll buy that permit.

In fact, I'll keep buying permits until we end up at X and Y, because every point between MCA and MCB, every point other than that, there's an efficiency gain from selling permits. And we'll continue until the equilibrium price of permits is \$100.

So you'll replicate the tax outcome through trading. Now, once again, that assumes that the trading doesn't break down for all these various reasons. We're back in the private market problems. Firm A could hate firm B. The government has to assign the value, et cetera. There's lots of problems.

But in principle, you could get to that outcome with a quantity regulation and cap and trade. But we're already starting to see the reason why economists tend to prefer price solutions to quantity solutions. They tend to prefer price solutions for two reasons.

First of all, you need to add an extra step to the quantity solution to be as efficient as the price solution. Second of all, you need a lot more information. Let's go back to figure 5-6.

Now, here I said I could either tax you by the amount MD or I could impose Q2. What do I need to know for the tax? I need to know one thing. What's the marginal damage? What do I need to know for the optimal regulation? I need to know the marginal damage. I need to know demand curves. I need to know supply curves.

I need to know a ton more stuff that's super hard to measure. So taxes basically have a lower information requirement. Now, once again, easy in this case where damage is flat. But as long as damage is locally flat, or you can approximate it, you need to know a lot less for a tax than you do for a regulation.

For regulation, I got to know everything that's involved in finding point B. That means I got to know demand curve, supply curves, and marginal damage. For tax, I just need to know marginal damage and I'm done.

So taxes have two big advantages. There's an informational advantage-- you need to know less. And there is a heterogeneous firms advantage which is that by pricing the externality, you get firms to do what's in their best interest. You get firms to efficiently reduce pollution.

This is very important. And the question that was asked here about the red line not being at the top is so important, because this exactly traces the intellectual history of this type of approach.

Which is when permits were first-- when people first started talking about using taxes or cap and trade, it was dramatically opposed, indeed, in almost religious terms. Many of you that know your religious history and will know what led to the downfall of the Catholic Church was in the Middle Ages, the Catholic Church used to sell absolutions.

So the way it worked was, if you sinned you went to hell, unless you gave the Catholic Church a piece of gold. And they gave you a paper that said you won't go to hell. This was viewed as not super appropriate. And Martin Luther decided to fight against it. And that led to the birth of Protestantism.

Many people referred to cap and trade systems as pollution absolution. They said it is wrong to allow firms to buy their way out of polluting. Look, what is cap and trade? Firm B is buying their way out of polluting. They're saying, I'm still going to pollute and just pay for it.

And people say it's ethically wrong. And they were wrong. It's ethically fine because all we really care about is the damage done to society by the pollution. Now, you may say there are reasons why some firms should pay more and some firms should pay less. And we get into that.

But the bottom line is, if you want to reduce pollution, you want to reduce it in the most efficient way. Let me just leave you with one story to illustrate this, which is-- I'll talk next time about negotiating the Global Warming Treaty.

I was at Kyoto in 1992 when they negotiated the original global warming treaty. And at that meeting in Kyoto, the liberals from the US were very much in favor of these kinds of cap and trade solutions.

And I said, gee, I thought you environmentalists kind of didn't like this pollution abatement stuff. And they said, no, actually, we've realized that it's efficient to regulate in this way. And the more efficient we can regulate, the more we can do.

So their view was everybody wins when you regulate efficiently. You don't just allow efficient production. But by making regulation less painful, you allow the government to do more of it. OK, so let me stop there. We'll come back next time to talk about the last section of chapter 5. And then we'll start on chapter 6 and talk about environmental externalities in practice.