JONATHAN GRUBER: Let's get started. I'm a little under the weather today. But don't take my slower style today for lack of enthusiasm about the topic, because today, we're going to cover one of my very favorite topics in economics, and the one I've devoted a lot of my life to studying, which is social insurance.

So today, we're going to start back where we ended. We're going to start back where we ended the last lecture. Remember the last lecture, we talked about information imperfections. We talked about uncertainty. And then we talked about why people would want insurance and why people do spend trillions of dollars on insurance.

But we also mentioned the government also spends trillions of dollars on insurance. And we asked why. And the answer lies in the problem of asymmetric information.

Now, in the used car market example I covered last time, buyers know more about the car they're selling-- I'm sorry, sellers know more about the car they're selling than the buyers know. And that asymmetric information led the market to potentially collapse.

Insurance market is the opposite. Here, the buyers know more than the seller does. But the problem is the same. So for example, suppose you want to start a new health insurance company for MIT graduates. MIT graduates need health insurance.

So you want to start a new company. You figure these are pretty healthy people, health insurance will be cheap. Indeed, you have data that says that out of a class of 100 graduates, let's say you're going to sell to 100 graduates, 90 of them are very careful people. And the only risk, once again, like last time, is getting hit by a car.

And for those 90 people, there's a 10% chance that they'll have a $10,000 bill and a 90% chance that they'll have no cost. That's the expected cost for those people, for those careful people. Now, for the less careful people, or we often refer to as the sicker people-- let's do some sickness. Let's not do this with car accidents.

Let's just say there's illness you can get. We'll call these healthy people. The healthy types, H, have a 10% chance of getting this sickness and 90% chance of nothing happening. The sick people, and then the remaining 10 people out of the 100 have a 50% chance of getting sick and incurring a 10K bill and a 50% chance of nothing. And that's the sick people.

This is all data you know. This is data given to you that is statistically accurate. It's a separate issue about when you don't know the information? We'll come back to that. But for now, imagine you know this information.

So what you do is you say, look, I then know the expected value of buying insurance to MIT students. When I sell to those 100 people, my expected cost-- if I sell this 100 people, my expected cost is what? Well, my expected cost for the healthy people is $1,000, a 10% chance of $10,000.

My expected cost is $1,000 for the healthy people. And they are 90% of the population. So I have $1,000 expected cost of healthy people. They're 90% of the population.

And for the sick people, I'm expected cost of $5,000 because they might get sick 50% of the time. And that happens 10% of the time. So by simple statistics, my expected cost is $1,400. So I expect that when I sell insurance, I will pay out $1,400.

So I say, well, if I sell insurance for $1,500, as long as people are a little risk averse, they'll want to buy it, right? $1,400, everyone wants to buy as long as they're not risk loving. If I charge $1,500, as long as people are little risk averse, they'll buy it, and I can make a profit.

So I roll in, and I sell insurance for $1,500. And I make a tidy profit, I think, of $100 on 100 people, which is $10,000. Why am I wrong? Why do I not, in fact, make that profit? Why, in fact, might I lose money?

Why is that math wrong? Why, if I offer $1,500, might I, in fact, lose money? Yeah?

AUDIENCE: People aren't likely to purchase.

JONATHAN GRUBER: Yeah, because I assumed everyone buys it. But I don't have the power to make everyone buy it. So who's actually going to buy it? Well, let's think about that.

All the sick guys, if I charge $1,500, this is an incredible deal for the sick guys, right? The expected value of them of the insurance is $5,000. And I'm charging $1,500. And they're risk-averse, so they're like, this is a great deal.

What about the healthy guys? Well, the expected value of insurance to them is $1,000. And I'm charging the $1,500. So unless their risk premium is $500 or more, they won't buy.

So let's imagine there's some distribution of risk premiums out there in the population. And let's imagine that it's such that only half the healthy guys decide to buy when I charge $1,500. That half number just I made up.

But there's some distribution of risk aversion. Some healthy guys are super risk-averse. They'll happily buy at $1,500. Some aren't very risk-averse. They won't buy at $1,500.

Let's say half buy. So now what do I pay out? At the end of the year, what do I pay out? Well, to my sick people, to my 90 sick people, I pay out-- to my 90 sick people, I pay out, on average, $45,000.

I pay out, on average-- I'm sorry for my 10 sick people, I pay out, on average, $5,000. OK? That's what I pay my 10 sick people. That's my expected payout.

What about for my healthy people? I said only 45 of the healthy people choose to buy, right? So for those people, I have a 45 times a $1,000 payout in expectation. So my payout is $95,000 for 45 people.

But what's my revenues? Well, how many people have I sold to? 60 people, 10 sick and 50 healthy. I'm sorry, I've sold to 55 people, 10 sick, and 45 healthy.

What's 55 times 150? Well, 50 times 150 is only $82,250. So I have paid out $95,000. But I've only collected a little more than $82,000, so I lose money. Therefore, I don't offer the insurance.

This is the flip of the car example. In the car example, because I don't know what's a lemon and what's not a lemon, I don't buy the car. Here, because I'm not sure who's sick and who's not, if I offer at that price, I'm going to lose money, so I won't offer it.

Now, you might say to me, well, that's because you were dumb. You chose too low a price. What if I just raised the price to $2,000? Maybe that'll solve my problem.

Well, at $2,000, I make a lot more money per person. But what else happens at $2,000? Even fewer healthy people buy. So if at $2,000, only 10 healthy people buy instead of 45, I'll still lose money.

The point is, this problem of what we call adverse selection, which is only the people on whom you're going to lose money want to buy, screws up insurance markets. Once again, let me be clear on this, since only the people for whom it's a good deal want to buy, and you lose money on the people for whom it's a good deal. You make money for people, it's a bad deal.

It's a bad deal for the healthy. If you charge more than $1,000, it's a bad deal for them. You're making money on them. So essentially, what's insurance doing? It's taking money from the healthy, extraordinarily healthy to cover your costs on the sick.

That only works if the healthy will play. If the healthy won't play, it doesn't work. OK, questions about that? So what happens in the real world? Well, there's five possible outcomes here, five possible outcomes.

Outcome one. Outcome one is the one good outcome, which is the healthy people are so risk-averse, enough of them buy. So if there's enough risk-averse healthy, then you can make this work, right? If something like 70 of the 90 healthy people buy, I forget where the crossover point is. We should add that to the notes.

By the way, there's an error in the notes. It should say 45, not 50. And we should add what the crossover point is. I can't do it in my head. But there's some crossover point where if enough healthy people buy, I will make money at $1,500 or $2000.

So the first outcome is it all works out. Yes, there's adverse selection. But it's not so bad because the healthy are very risk averse. That's outcome one.

Outcome two is the market collapses, the market fails. Why does the market fail? Because at any price, the adverse selection is so bad I can't make money. I never get enough profitable healthy to cover my unprofitable sick so the market fails. That's the second possible outcome.

There's a third outcome. I could just sell to sick. Let's say I rolled in and I said, I've got insurance for MIT graduates. It's going to be $5,500. Well, at that point, as long as the sick are somewhat risk-averse, they'll buy it.

Or at $5,100, if they're all risk-averse, they'll buy it, and I'll make some money. So I could just sell the sick people. At that high price, no healthy people want it. If the healthy people want it, great. Go for it. I'm happy to have them. But that's just icing.

I then make money on the sick as long as I set the price above $5,000. So that's the only thing I could do. I could just have a health insurance policy for the sick priced at the price for the sick. And I know the healthy just won't take it. Or if they do, it's only icing.

A fourth thing I could do is I could just discriminate. I could say, I'm going to sell insurance at $1,500. I'm going to test when you come in. If I find you're sick, I'm not going to sell to you. If I find you're healthy, I'm going to sell to you. Turns out, totally legal until the Affordable Care Act passed.

So for most of my life, insurers were totally legal to say, if you're sick, I'm not going to sell you insurance. Or if you get sick, I'm going to drop you from insurance. As I hope you can see, that undercuts the whole idea of insurance. The point of your insurance is to get healthy people to cross-subsidize sick people, and you drop people as soon as they're sick, then it doesn't really provide insurance.

So insurers were totally free to do that before the Affordable Care Act passed. And I'll talk more about that in the last lecture.

Finally, there's a clever solution, which the math is hard, so I'll just give you the intuition, which is you can sell different products. That is what I can do is I can sell two products, comprehensive insurance for $5,500 and crappy insurance for $1,500. As long as the sick people don't want the crappy insurance, then what have I done? I've solved the asymmetric information problem by getting people to reveal their underlying health.

Remember, my problem is I don't know who's sick and who's healthy. I know the aggregates. But I don't know who's sick and who's healthy. If people had sick and healthy stamped on their forehead, I'd be done. But I don't know that.

But if I offer a very good product that's expensive and a shitty product that's cheap, people will tell me their type by what they buy. The healthy people will buy the crappy product. Sick people will buy the other product. And I'll make money on both.

So that's another thing I do. Have you ever heard of an HMO? That's what an HMO is, OK? Basically, that's what we've done with health insurance in America to try to deal with this problem is created less high-quality insurance for the healthiest people that they can buy cheaply. OK, so that's another thing we can do.

OK, here's the key point. All of these are market failures, all of them. You might say, well, I know that's a market failure. But why is it a market failure if you just sell to the sick? Why is that a market failure? Yeah?

AUDIENCE: Potentially, we lost some healthy people who would want to buy it.

JONATHAN GRUBER: Yeah, there are traits that, with perfect information, I would be happy to make, and they would be happy to make and not being made. So as long as imperfect information, information asymmetry, leads to trades not being made that make both sides better off, then that's a market failure. So only selling to the sick is a market failure.

Likewise, if I discriminate, that's a market failure because there are sick people who would happily pay me $5,200 for insurance. But I don't know who they are, so I can't sell them at $5,200. Finally, even selling different products is still a market failure. The reason is because what you can see from the insurance model we set up last time is a typical risk-averse consumer wants complete insurance.

They don't want partial insurance. They want to basically make sure that their consumption is the same, no matter if they get hit by a car or not. And partial insurance isn't as good because, once again, the extra they keep doesn't compensate for how much they lose if they get hit by the car.

So incomplete insurance is also a market failure. So the point is, unless you have enough risk averse healthy, the market will not deliver the welfare maximizing outcome, OK? Questions about that?

So now, once again, as I said, the principle that I teach in my classes is the market knows best until the market fails. And then there's a potential role for the government. But it's only a potential role for the government.

And the reason it's especially confusing here is there's many different ways the government could come in and solve this problem. And in fact, governments around the world do this in three different ways. They solve a problem three different ways. So what about government solutions?

The first way governments solve this problem is with subsidies. Let's go back to our example. What if I said to everyone who's an MIT graduate, I will give you $500 if you buy insurance. Then the original product would work because the $1,500 I would charge them, let's say, I, the government, say I'll give you $500.

Well, the insurance I'm offering is $1,500. So it's then a good deal for both sides. So they both buy it. So in some sense, one way to combat adverse selection is through price. If you make it cheap enough, the healthy will buy.

So subsidies just naturally offset adverse selection. If I give a $500 subsidy, in my previous example, it's a good deal for everyone. And everyone buys, and we get solution one. We get this one. Everyone buys, OK? That's one option.

Indeed, we use that option a lot in the US. In the US, we massively subsidize private health insurance. How do we do that? Well, one way we do it is through our tax code. Here's an interesting feature of our tax code.

If MIT pays me in wages, I get taxed. I'm not giving away too much private information to say, if you add up the tax that I pay to Massachusetts and to the federal government, my tax rate approaches 50%. So I get to keep about half of what I make.

But if MIT gives me health insurance, it's not taxed. So if MIT says to me, John, do you want $1,000 in health insurance or $1,000 in wages? I say, wait a second. You give me $1,000 in wages, I get to buy $500 worth of goods. You give me $1,000 worth of health insurance, I get $1,000 worth of health insurance. I'll take the health insurance.

And that is partly why most people in America get their health insurance through their employer. We bribe them to. We bribe them to by taxing their wages but not taxing their health insurance. That's a massive subsidy. That costs the US government $300 billion every single year, big money, that we're paying to bribe people to buy private insurance. So that's one example of how a government can come in and try to solve this problem.

A second example is through a mandate. And if you followed health care policy debates when you were younger, remember this term of the mandate, it was a big centerpiece to the debate over the Affordable Care Act, which is what if I just say everyone has to buy? I say, healthy, sick alike, you have to buy.

Then basically, I get this outcome again because I force everyone to buy so I know what the price is. Remember, I know, in aggregate, what the price should be. I just don't know who's going to buy and who's not. So I force everyone to buy. Everyone comes in, and I know what price to set.

Now, you might think, well, gee, that seems like a way better option than option one. Option one costs the government a ton of money. Option two costs the government nothing other than to just say, buy insurance.

But in fact, who gets upset when there's a mandate? Which particular group is most upset when you mandate health insurance? The healthy who are not super risk-averse.

They are getting hit. They didn't want to buy insurance at 1,500, you're forced to do something they didn't want to do. That's politically very challenging, as we learned with the Affordable Care Act, which is that it's politically very challenging to force people to do something they don't want to do in the American political system.

So in some sense, you can think of a mandate as implicitly forcing the healthy to subsidize the sick. That's great because it gets the government out of the way. But it's essentially pissing off the healthy, non-risk-averse people.

And believe me, during the Affordable Care Act debate, every single one of them sent me an email. People did not like that feature of the mandate in the Affordable Care Act. And eventually, the mandate went away for that reason.

But we have mandates. They're not new. For example, the oldest program of this type in the country is called Workers' Compensation Insurance. That's a program that insures you if you get hurt at work. That's been a mandated benefit of employers since the early 20th century.

No one's protesting Workers' Comp mandates. In most states, you're mandated to buy auto insurance. People don't protest that. They bitch about how expensive insurance is. People don't protest it.

So interestingly, mandates are not a new concept brought up by the Affordable Care Act. It was just in this application that got people very upset.

The final option for the government is you can just provide the insurance. Cut out the middleman. Say, sorry, John, we're just going to have gov care.

It's going to be whatever it's going to be. And if we lose money, we lose money because we're the government. We don't need to break even. We're just going to provide it.

Obviously, you can combine it with subsidies. You say provide it, we're going to subsidize so everyone's in, or you just provide it and make it free. This is what most countries do. Most countries just say, we have a national health insurance system. It's free.

Take Canada. When you're born, you're entitled to national health insurance. It's free. You just go. There's no adverse selection because it's free. When stuff is free, there's no adverse selection.

So that's another way you could do it is have the government just provide it. Indeed, that's what most of the world does. OK. Questions about any of that? Yeah?

AUDIENCE: Do mandates hurt private insurers at all?

JONATHAN GRUBER: No, mandates help private insurers because when you mandate, you're giving them more customers.

AUDIENCE: During the whole, I don't know, government health care thing back then, some private insurance weren't happy about it.

JONATHAN GRUBER: Because what the government did is they wanted to make a trade. They said, we'll mandate insurance, but you need to stop discriminating. That was the trade, which they're saying you're discriminating because of adverse selection.

Here's the deal we're going to offer you. Stop discriminating, and we'll end your adverse selection. Some people would rather discriminate. And they're upset about it. OK, good question.

Now the bottom line is this all sounds pretty good. We've got a problem. We've identified potential solutions. Most governments in the world are using some combination of solutions. What's the problem?

Well, the problem is that there's another asymmetric information problem, which poses a fundamental trade-off. And that's what we call the fundamental social insurance trade-off. And that other problem, that other asymmetric information problem is what we call moral hazard.

The other asymmetric information problem is what we call moral hazard. That's basically the adverse behavior that is encouraged by insuring against an adverse event. When you insure against an adverse event, you can cause adverse behavior.

So basically, for example, if families buy fire insurance for their homes, they may be less likely to have a fire extinguisher in their house. They don't need it because their house is insured if it burns down. If people have health insurance, maybe they're less worried about taking precautions.

OK, I'm not feeling so well today. I decided I'd come in with a mask. I'd be OK. But I don't have health insurance, maybe I wouldn't. So basically, maybe there's risks people take.

And then also, take another example. One of the biggest insurance programs in America is unemployment insurance, insurance you get if you lose your job. Well, if we insure you for the money you lose when you lose your job, what's the incentive to search for a new job? Remember, leisure is a normal good.

These are all forms of moral hazard. And the best and most fun examples come in the context of Workers' Compensation, which I already mentioned already. That's a $60 billion a year program that insures people for injuries that are job related.

The problem is, now, you might think, well, that sounds really sensible insurance. The problem is a lot of these problems are actually hard to diagnose and hard to figure if they're job-related. If someone is at a machine, you watch them saw their arm off, that's pretty easy.

If someone says, I've been driving all day and I have back pain and I can't stand up, that's harder. Most of claims for worker's compensation today are mostly pain and mental health. Those are hard things to evaluate as a result.

And then on top of that-- so you've got two problems with Worker's Compensation. One is it's hard to evaluate if people are really injured. The other is if they are injured, they get a pretty attractive benefit. If you miss work and get Workers' Comp, you get 75% of the wages you would have gotten if you worked.

So you can sit at home and get 75. Remember, leisure is normal. People would rather not work. So you can work or sit at home with 75% of wages. But it's better than that because your wages are taxed and Workers' Comp isn't.

So for a typical worker after tax, their Workers' Comp amounts to about 90% of what they take home working. So you've got a program which says, look, leisure is a normal good. I'm going to pay you roughly 90% of what you'd earn at work. All you have to do is pretend you're hurt. That's a clear incentive for moral hazard.

So basically, the idea is that there might be moral hazard in something like Workers' Comp. And there's been a large, really fascinating literature that's tried to measure that. Sometimes there's fun examples.

My favorite example was the woman who-- in what state? I'm talking about bad behavior. What state am I going to talk about? Florida. Come on.

So there was a woman in Florida who was at her desk, while she was at her desk, a sprinkler head fell out of the ceiling and landed on her desk and shocked her. She looked around, bonked herself on the head with the sprinkler piece, and claimed Workers' Comp. But, of course, the reason I'm telling this story is it's all on camera.

Have you guys ever heard of show *Wicked Tuna?* It's about tuna fishing out of Gloucester, about a guy who wrestles these 1,000, two ton tunas onto his boat. Turns out, the whole time he was on *Wicked Tuna,* he was collecting Workers' Compensation because he couldn't work. There are lots of really fun examples.

But what we care about in economics is not stories. We care about facts. And how do we test, in empirical economics, whether moral hazard is a real problem? The way we do it is we say, it turns out that what you get from Workers' Comp varies across different states and over time.

So let's say I'm in one state, Andrew is in another state. Andrew's state randomly raises the benefits. Now, there's no reason that Andrew should be more injury prone just because the state raised the benefits.

But if Andrew becomes more likely to be injured when the state raises its benefits, clearly that's moral hazard because there's no other reason why, just because the state raises benefits, Andrew should be more injury prone. And the evidence is that it's incredibly powerful. We talked about elasticities. The evidence is that for every 10% that benefits go up, people miss 10% to 20% more work, an elasticity of 1 to 2.

People are very responsive to these benefits. And that's clear evidence of moral hazard, OK? So Moral hazard exists. But we might ask, look, so what's the problem?

OK, adverse selection, I explained the problem. The market collapses. That's pretty bad. Collapse, we know, is a bad word. That's bad. What about moral hazard? Why is it so bad?

Well, it's bad for two reasons. The first is you remove productive labor supply. Let's just do some quick math to show why this is.

Imagine that an individual is choosing between staying at home-- is choosing for the marginal hour, Do I stay at home or do I work? Well, we know the way they'll make that decision is they will maximize their utility over consumption and leisure, subject to some budget constraint. And we know that this will yield the first order condition that du dl, or the marginal utility of leisure, equals the wage.

That will be the equilibrium condition. Why? Because it's where marginal benefits equal marginal costs. What's the marginal benefit of staying home? How much happier it makes you, the marginal utility of leisure.

What's the marginal cost of staying at home? You give up your wage, W. So an efficient market-- this is a very important insight, by the way. Many people talk about work as a virtue. That is wrong. Work is not a virtue.

People should only work until this condition is true. If you're someone who really hates work because you're really disabled, it's awful, you shouldn't work. It's not like God has said we must work.

The optimal social condition is that people should work until this condition is true, which means no one should work all the time. That can't be true 24 hours a day for everyone. At some point, you got to sleep. And some people should work less or more than others. OK? That's what's optimal.

And we know that the wage in a perfectly competitive labor market equals the marginal revenue product of labor. We know that's where the wage comes from in a perfectly competitive labor market. So basically, think of this from the employer's perspective.

From the employer perspective, I take leisure until the value of leisure equals the value of what I produce at work. So it's based on what we want society. We want people trading off the value of what they produce at work versus how much they value they get from sitting at home.

But now suppose I tell you that for every hour you sit at home, you also get an amount, G. Well, now if I sit at home, what do I get? I get du dl plus G. And I set that equal to the marginal revenue product of labor. So du dl plus G, I set equal to marginal product of labor.

Let's think for a second about what this does to my leisure decision. This means I set my marginal utility of leisure above-- I'm sorry, I set marginal utility of leisure below my marginal product because G is positive. I set my marginal utility of leisure equal to MRPL minus G.

Does that imply more leisure or less leisure than what we know is efficient? We know efficient is du dl equals MRPL. If du dl plus G equals MRPL, am I supplying more leisure or less leisure than the efficient outcome? We know this-- yeah, go ahead.

AUDIENCE: There's diminishing marginal returns there.

JONATHAN GRUBER: More, because remember, as leisure goes up, marginal utility goes down. We know we want more. Since G is positive, to keep MRPL, du dl has to go down. du dl goes down by taking more leisure.

So by offering this benefit, I've created inefficiency. I'm causing people to take more leisure than is optimal. The optimal leisure would be du dl equals MRPL. But I'm making du dl equal to MRPL minus G. So I'm causing them to take too much leisure.

In other words, I'm getting in the way of productive traits. There's a productive trade to be made here, which is, I work, you pay me. But now I'm like, fuck it, I'm not to work because I get G.

So that's a productive trade that goes away because of moral hazard. And that's why we care. That's a deadweight loss. OK?

There's a second problem, which is minor but important, remember, which is that, how do we pay for social insurance programs? By taxing workers. The less people work, the more taxes we have to raise on everybody else.

So it means higher tax rates. Because if we're shrinking the base of workers, but we've got to pay for all these workers sitting at home, that means higher tax rates. And we'll talk later about why higher tax rates cause inefficiency.

So this leads us to the fundamental trade-off with social insurance. On the one hand, unless the government intervenes, there will be a market collapse, unless you get condition one. And we have lots of evidence condition one isn't met in most markets.

The market will collapse. And that is bad. That's a welfare loss.

On the other hand, by providing that insurance, we create an offsetting welfare loss, which people don't work as hard. What does that mean? Well, that has a simple implication, which is we should have social insurance, but it shouldn't be too generous.

Clearly, the right answer is not zero social insurance unless moral hazard is deadly. I mean, this is just you're losing some product. This one, people starve if they don't have insurance. So clearly, you want some insurance.

But if we have too much, we might encourage people to sit at home. So the idea is we want in between. We want a trade-off. We want enough insurance that people don't starve, but not so much that we cause a lot of moral hazard. And to see that, let's talk about the largest social insurance program in America, which is the Social Security program.

We have a number of social insurance programs in America. I mentioned Workers' Compensation if you get hurt at work. I mentioned unemployment insurance if you lose your job. Our biggest aggregate social insurance programs are for health care. I'll talk about them in the last lecture. They're called Medicare and Medicaid.

But the single biggest program of social insurance, at about $800 billion a year, is Social Security. Social Security is insurance against the income loss that comes from retirement. Social Security is insurance against the income loss of retirement. It's insuring you to make sure that once you're retired, you have enough money and don't starve.

Let's talk about how it works. Here's how Social Security works. First, when you retire, you get a check from the government. You get a check from the government, OK?

That check is a function of how much you made before you retired. That check is progressive. What does that mean? That means that if you're very poor, you get a check that's almost as much as you made before you retired. If you're rich, you get a check that's about 30% of what you made before you retired.

So you get a check that basically depends on how much you made before you retire but in a progressive way. The more you made before you retired, the less you get back from the government. Everybody gets something. Everyone gets at least 30%. But there's a range, 30 to 90.

Actually, I'm wrong. It's at least 15%. There's a range, 15 to 90, over which it varies. So Bill Gates will get Social Security. Sorry to keep picking on Bill. He'll get Social Security, but he'll only get 15% of his previous wages, or subject to a cap, so even less.

You get this check until you die. It's what's called an annuity. This is an important term, not a lot of people know, even financially savvy people. An annuity is the opposite of life insurance.

We all know life insurance. Life insurance is I pay money while I'm alive. And when I die, my descendants get a check. Annuity is the opposite.

Annuity is I pay money now. They send me a check. As long as I'm alive. And then it ends when I die. So if I want to annuitize my wealth, what I do is I'd walk into a broker's office. I'd say, here's my money. I'll give it all to you now, and you're going to send me a check until I die.

Why is this appealing? It's appealing because we don't know when we'll die. And as a result, savings for retirement is very imperfect. Definitionally, you will always save if you just use saving for retirement, you always save either too much or too little.

If you live longer than you thought, you'll save too little. If you die sooner than you thought, you'll save too much. Now, if you care about your kids, that's a different thing. But let's talk about a childless person, OK?

So basically, the point is, savings isn't perfect because you don't know when you're going to die. Annuities are insurance against that uncertainty. Think about it. Just like health insurance insures you against the medical bill, annuities insure against how long you're going to live.

You turn out to live a long time, then you really will turn out to be a good deal. You've paid an amount, and you get a check forever. If you turn out to live only a short time, annuity would be a bad deal.

Just like with health insurance, the health insurer, if you turn out to be healthy, health insurance is a bad deal. If you turn out to be sick, it's a good deal. But you still want it because you want to smooth your consumption. And that's what annuities do.

Now, that is the benefit of Social Security, that we make sure that people have enough money till they die. They don't under save, that they don't under save and end up without enough money because they don't think they're going to live that long. They live a long time, and suddenly they're impoverished.

OK, and indeed, one of the reasons we set up Social Security was because there'd be all these impoverished people, especially after the Great Depression. The main motivation Social Security is the Great Depression wiped out a bunch of people's savings. And we were worried they wouldn't have enough money to live to till their death, so we set up this program.

On the other hand, there's a moral hazard element of Social Security, which is, remember, how do you qualify for Social Security? By retiring. Retirement is a choice. Leisure is normal. People would like to be retired.

So the more generous we make Social Security, the earlier you'll retire. Because basically, we're saying go back to the math over here. Think about this not as a period by period decision, but a year by year decision.

This year, do I want to work? Well, this year, I think, gee, how much do I value leisure if I sit at home versus how much will I make this year, and decide, should I work or not? Well, now I'm going to be more likely to quit because if I quit, I get money I get if I don't work.

So I've essentially put a thumb on the scale for quitting. I've put a thumb on the scale for retirement. So the moral hazard aspect of Social Security is it causes increased retirement.

Why is that bad? Well, because older workers are productive. Not that they shouldn't retire. But they should only retire when this condition is true, not when this condition is true, OK?

They're retiring too early. And in fact, there's tons of evidence from very well-crafted studies that Social Security has a huge impact on retirement. And basically, the way they study this often is by looking at the fact that Social Security has a minimum age, and when you lower that minimum age, all of a sudden, everyone retires earlier. When you age, people retire later.

In Germany, the median age was 60 and the average retirement age was 58. Some people retire early because they have their own money. They raised to 65. Within five years, the average retirement age in the whole country of Germany was 63.

Think about that. Within five years, an entire country changed when retired by five years. Pretty phenomenal, OK? That illustrates the power of moral hazard. And that says there's a trade-off.

So basically, the question is, how do you think about this trade-off? Well, you write down an economic model and you talk about the benefits of insuring against longevity and the costs of early retirement. And you come out with some optimal Social Security number.

This isn't what we do here. We'll do that in 1441 if you want to follow up. That's 1441. We'll write down a model where we put both these things in one model and optimize. And you get some optimal level of Social Security.

It turns out, it's probably not that far from where we are. Social Security, on average, replaces about 45% of people's wages across everyone. That's probably not a terrible number given this trade-off.

But the problem is, this is a much bigger issue in Europe than it is in the US because of a seemingly harmless little tweak in the way we do Social Security versus the way Europe does Social Security. Social Security has an age of early retirement, the earliest at which you can get your benefits.

So the way it works in Social Security in the US is as soon as you're 62, you can start getting a check. But if you wait till 63, we'll give you a bigger check. Why? Because you're going to get it for fewer years.

Your death date is the same. So at 62, you're going to die at 60 to 100, and we'll give you x. If you get 63 to 100, we're going to give you x plus a little bit to give you the same total by age 100. So we basically have what we call, in Social Security, an actuarial adjustment. We basically, look, the longer you work, the more money we'll let you keep because you're going to get fewer years of benefits. It's a trade-off.

The issue is getting this actual adjustment right can strongly impact people's work decisions. So let's do an example. Let's say you earn y percent per year. And Social Security replaces k percent of that.

Remember, I said, on average, Social Security replaces about 45%. So k would be 0.45, or 45%, in the US. So Social Security replaced about 45%. And ignore the fact that rich people and poor people are different. Imagine one representative person.

That representative person earns y. And when they retire, they're going to get k times y, where k is a number less than 1. And let's imagine that you're going to live until 80. And you're deciding, should I retire at age 62 or not? Should I have 62 or work till 63?

Well, if I retire 62, what do I get from social security? Well, I get the present discounted value-- remember, I got a PDV it because the money's coming in the future-- of getting 18 years of k times y. I retire at 62. I get it for 18 years. And it's going to be in the future, so I got to PDV it.

If I work till 63, what do I get? If I work till 63, well, then I get y. I get the PDV of y this year plus 17 times k times y, because I get it for one fewer year. So what's the difference between these two? The difference between these two is 1 minus k times y. That's the difference between these two.

If you work one more year, you get 1 minus k times y more. Questions about that? Why is this important? Because this k term effectively is a tax.

Think about it. If I work one more year without Social Security, what do I get? Y. If I work one more year with Social Security, I get 1 minus k times y. Why is that? Because I'm giving up the money I could have had by sitting at home.

And my decision is always sitting at home versus working. So there's a trade-off. The bigger k is, imagine k was 100%. Imagine k was 1 and leisure's normal. Would anyone work past 62? No.

Why would you work past 62? You get the same thing. If k is 1, that means if I work, I get y. And if I'm home, I get y. And leisure is normal. Why would I ever work? So the bigger k is, the bigger the moral hazard is. The less incentive there is to work.

So what do we do in the US to offset that? We have a simple, what I call, actuarial adjustment, which is we say if you work till 63, we're going to give you an extra bonus. We're going to give you y plus 17 times k times y times 1.067. We're going to give you more.

And that 1.067 is calculated so that you're indifferent. So if the average person is indifferent, yes, they give up, k. But they get more per month. And this calculates to the average person is indifferent. So roughly speaking, in the US, when you're deciding to retire at 62 or 63, Social Security shouldn't affect that decision. It's basically pretty neutral.

Now let's go to the Netherlands. In the Netherlands, two things are different. k, instead of being 0.5 and 0.9, B, the second thing different, they don't have this. So in the Netherlands, your decision is, do I work? Age 55 is the early retirement age.

Your decision is, do I work till age 55 and get y? Do I work my 56th year and get y? Or do I stay at home and get k times y, which is 90% of y? OK, people understand that? That's your decision.

But it's even worse than that. Because how does the Netherlands pay for this program? By taxing workers, taxes you don't pay if you don't work. So if you work, you get y, but you lose almost half of it in taxes. So you get half y.

If you don't work, you get 90% of y. You make money by not working. Even MIT students might find that attractive. As a result, literally no one works past age 55 in these countries.

We essentially made it not in their financial interest. Some people go back to work off the books in jobs that aren't recorded, so they can make some money. But that's why you see all these old people sitting around cafes in Europe, having a great old time in the middle of the day, because basically, there's no financial benefit to working. We're essentially making you indifferent if you're working or non-working once you've reached the retirement age.

So that is an example of moral hazard. That's an example of setting the benefits too large so that the moral hazard cost is enormous relative to potential smoothing. If the Netherlands, for example, lowered the 90% or even just introduce a factor like this, that would offset things a lot.

So the point is, this is a wonky discussion. But the point is this shows you how these little wonky numbers that policymakers fight about all the time could have a fundamental impact on the way an entire economy functions. Literally, by setting k to 0.9, and by not having this actuarial adjustment, the Netherlands has essentially taken every old person out of the labor force.

That's a pretty fundamental change. And that's something that policymakers-- that's why I love policymaking. That's what policymakers need to think about, is essentially, what are the unintended consequences of what I started this semester with.

I want to ensure people in Netherlands if they can't work anymore. But I don't want to take every old person out of the labor force. And that's what they end up doing.

OK, that's all I got for today. Hope you guys had a good Thanksgiving. And we'll come back Wednesday.