

[SQUEAKING]

[RUSTLING]

[CLICKING]

JON GRUBER: We are now going to shift gears to the next section of the course. So really, as I said, the course divided fundamentally into three sections. We have externalities and public goods. That's what we've covered so far. We then talk about social insurance for the next probably eight lectures. And then we turn to taxation.

Now, what is social insurance? Well, let's start by just looking at figure 12-1, and reminding-- it's the figure we saw earlier, just recast a little bit. Reminding ourselves of how the nature of what the US government does has changed.

In 1953, the government was mostly defense. There was a very small Social Security program that took up 3.5%, 3.6% of our government spending, and the government virtually did nothing in the realm of health, 0.5% of government spending.

By 2019, defense had shrunk to 15%, Social Security to 18%, and health to 24%. This has been a massive shift in what the US government does. As Paul Krugman wrote, "The US government is essentially now a giant pension fund with a small army attached."

Basically, we have gone from a government that's fundamental responsibilities defending us to fundamental responsibility of insuring us. Why? What does that mean? And what are all these big programs that are so expensive?

Well, these are a class of what we're going to call social insurance programs. Social insurance programs are going to be programs-- and that's what we're going to cover in the next few lectures, are social insurance programs. These are a class of programs that share several key features.

The first feature is workers-- they're restricted to workers who participate by essentially making involuntary contributions. They can be taxes, they can be insurance premiums, but the point is, they're not voluntary. So basically, involuntary worker contributions is one feature.

The second feature is these programs are not means-tested. What does that mean? That means that they're universal. Everyone gets them. This is not something where only poor people get them.

Now the benefits-- the nature of the program can differ by income, but everyone's entitled. These are universal programs. And the third point is that benefit's payments are triggered by an adverse event. So for example, a classic example of a social insurance program is unemployment insurance. We'll talk about that a lot. Some today, a lot in a couple lectures.

Unemployment insurance is a program where everybody-- actually, their employers on their behalf, pays a tax to the government, everyone who loses their job qualifies, and you qualify by losing your job. That's the adverse event. So everyone's entitled to qualify by losing your job.

Another example is Social Security, the nation's largest single social insurance program. That's a program where we pay-- you know that line FICA on your pay stub? That's the Social Security contribution. Everyone gets Social Security, and it's triggered by retirement or disability.

These are the kinds of programs that are under the social insurance umbrella that we'll cover over the next several lectures, but that's the general idea of social insurance. This would be distinguished from things which are redistribution programs, which explicitly are restricted to some people and not others, and which are triggered by what your income is, not what happens to you. These are really like insurance programs.

Now, to understand why the government spends so much of its revenue on these programs, we have to step back and ask, what is insurance, and why do we want it in the first place? So let's step back and do a little bit of basic insurance theory, this is some review of 14.01. Let's step back and do some basic insurance theory.

Now, first of all, let's stipulate, insurance is big business in America. We have a private health insurance industry, which is currently collects about \$1.2 trillion in premiums. We have an insurance industry, which collects about \$250 billion in premiums. We have a life insurance program which collects about \$160 billion in premiums. And then we have property and casualty insurance, like if something catches on fire or gets stolen, that's about \$600 billion a year.

So we're, as a nation, spending on the order of \$3 trillion-ish a year to insure ourselves. The nature of insurance is essentially you pay a month-- the nature of private insurance is you pay a regular payment, called a premium, usually monthly; in return, that insurance covers you against some adverse event.

So health insurance that covers if we're sick. Property and casualty insurance to cover us if our house burns down. And that's big dollars. It's trillions of dollars in the US.

Now why do people want insurance that badly? Why are we, as the economy, is spending so much of our GDP on insurance? Why do individuals value it? And I'll start with the general intuition. In the book, there's a numerical example, but I'll cover the math.

The general intuition is quite simple. Think about-- remember our typical intertemporal choice problem, which is you can eat today or tomorrow. Instead, we're going to do what we call an uncertainty problem, which is there's multiple states of the world. So it's like a science fiction thing where basically in the future, one of two things can happen. You can think of yourself in either state of the world.

One state of the world is the bad thing doesn't happen, one state of the world is a bad thing does happen. And from today's perspective, you can't know with certainty if it's going to happen. You can know the probability-- or you may know the probability of the bad things can happen, but you don't know, from your current perspective, if it's going to happen or not.

So you have this uncertain risk you face with some probability. And then if that risk happens, if something bad happens to you, you end up in a worse state than if something good happens to you.

And the basic principle that we learned in 14.01, we talked about uncertainty, is that individuals want consumption smoothing. That because our utility functions feature diminishing marginal utility, we're always happier having the same consumption in different states in the world rather than different consumption.

Because remember, with diminishing marginal utility, the extra dollar makes us less happy than one less dollar makes us sad because it's diminishing. So if you're at a certain point and your income goes up a lot, you're happier, but you're not nearly as happy as you are sad if your income goes down a lot because of diminishing marginal utility.

Think of the limit, if everyone goes down all the way to 0, that's much worse, whereas how happy are the billion dollars? It's still limited. So basically, essentially the idea is-- that's about it, forget that intuition, I'll get to math in a second.

The bottom line is, just focus on the curvature, which is, remember, that because of the curvature of the utility function, because diminishing marginal utility, you would want to make sure consumption's smooth across states of the world. That's the basic principle that drives our demand for insurance.

So to understand that, there's a numerical example in the chapter, but let's go to the appendix and actually do the math. So, let's imagine the following model. Let's imagine you are hit by a car with some probability-- here, let's set it up here. There's some probability P that an adverse event happens. Say you're hit by a car.

Regardless of whether you get hit by or not, hit or not, your income is W . So P is the risk getting hit by a car, W is your income regardless, you still go to work either way.

But if you get hit by the car, you incur medical costs δ . You can buy insurance for a premium of M per dollar. So M is the number of cents it costs you to buy a dollar's worth of insurance. That insurance you buy pays out a benefit b . Pays out a benefit b if you get hit. b is the amount of benefit it pays out.

So once again, in this world, you only risk hit by a car. You're a young, healthy man-- and I choose man because there's no risk of pregnancy. You're a young, healthy man. The only thing that can go wrong is you getting hit by a car. This is pre-COVID. Things go wrong, you get hit by a car.

There's a P probability hit by a car. Your income's W regardless of whether you get hit. There's a d chance you get hit. You can buy insurance at a premium of M per dollar of insurance and it pays you a benefit b .

Now, we write down your expected utility. What is your expected utility? Well, your expected utility is 1 minus P , the chance that you don't hit by a car, times the utility of W minus the insurance-- what you pay for insurance, Mb . That's what you pay for insurance, is Mb . So if you want to benefit \$1,000, you pay 1,000 in because M is per dollar.

If you get hit by a car, there's a probability P times the utility of-- you still have income W . You still pay the insurance premium. You pay insurance no matter what. So you still pay Mb , but you also lose δ because you have to pay that medical cost. But you get b , which is the benefit because they pay for your medical costs. They pay some amount b . So, basically that is your that's your setup. Questions about that? Yeah?

AUDIENCE: [INAUDIBLE]

JON GRUBER: Yes. Because M per dollar of insurance, you buy b dollars of insurance. And b is how much you pay for insurance.

Now, we are going to want to optimize this. But if you'll notice, there is two unknowns here. We're going to stipulate W -- we'll tell you what w is, we'll tell you what P is, we'll tell you what δ is, but you have to decide M and b . So we need a second equation to close this system.

And the second equation we're going to get to close the system is the supply side of insurance. And we're going to get the price of insurance by assuming a zero-profit insurer. We're going to assume a zero-profit in expectation. We're going to assume the insurer has zero expected profits.

What does that mean? Well, the insurer's expected profits are what it collects in premiums, which is Mb . That's what it collects in premiums. Minus what it expects to pay out, which is Pb . P percent of the time, they'll pay out b .

Well, what you can see from this is that zero profits implies M equals b . That's the implication of the zero-profit condition. That basically-- I'm sorry, M equals P , my bad. Sorry, M equals P . M equals P . That the zero-profit condition is the rate per dollar charged to you is equal to the probability event of the event.

And it makes sense. If it's a 1 in 100 event, then by charging a cent per dollar of insurance, they break even. So that's how they break even. They basically say a rare event, you pay less; a lesser event, you pay more. So that's what delivers zero profits for the insurer.

Now, with that in mind, let's take a typical assumption. You assume utility equals square root of C , which is our typical risk-averse diminishing marginal utility utility function. And let's plug in and write this down.

So we say expected utility now equals $1 - P$ times the square root of $W - Mb$, which we can rewrite as $W - bP$, because M equals-- I'm sorry. We can rewrite as $W - bP$, right. So we can rewrite that as $W - bP$ because Mb , b is b and M is P . So just rewrote that.

And then for the second part of the equation, we have P plus P times the square root of $W - bP$ minus Δ plus b . So all I did here I did two things. I plugged in from the zero-profit condition, and I assumed a functional form for utility.

We can maximize this with respect to b . So let's maximize this equation with respect to b , you want to choose the optimal benefits. By choosing the optimal benefits, you choose how much you pay in premiums. So basically, you just want to choose the optimal-- you just want to choose the optimal benefits.

And that condition is if you-- I want to ask, what's the optimal benefits? Well basically, you're going to get-- if you take the derivative of this with respect to b , you're going to get $-1 - P$ times P over-- $1 - b$ of P , $1 - P$ times P over square root of $W - bP$. Plus P times $1 - P$ over the square root of $W - bP$ minus Δ plus b . That's what you get when you take the derivative of this equation.

If you set that equal to 0 and solve, you get that b^* equals Δ . That the optimal level of benefits is equal to the medical cost of the adverse event. You can do the math to see that. You can also cheat and say, look, we know that optimal insurance involves consumption smoothing. That involves these two terms being the same.

If these two terms are different, it's not optimal. Why? Because then they have a different utility in different states of the world. They want the same utility in both states of the world. What sets these two equations equal? b^* equals Δ .

So you can take the derivative and do the math, or you can just remember the intuition that the alpha comes consumption smoothing. Consumption smoothing implies the term in this-- here in this term equals this term. And that is only true when b^* equals Δ .

So consumption smoothing requires full insurance. That's the key outcome of insurance theory. That the first best, the optimum, is full insurance. Because if you don't have full insurance, then you're consuming less. If you're less than fully insured, you consume less in the bad state than the good state, you don't like that. You want to consume the same in every state. That's the whole principle of consumption smoothing. Yeah, Alec?

AUDIENCE: This assumes that the insurance firm has no cost other than payouts?

JON GRUBER: I'll get to that. Right now, it's a perfectly zero-profit, perfectly efficient insurer. Yeah, Stephen?

AUDIENCE: So can you just explain specifically what b is and what M --

JON GRUBER: Sure, I understand. M is what you pay per dollar of insurance. b is how many dollars of insurance you get.

AUDIENCE: So it's the--

JON GRUBER: So in other words, it-- it's the cost-- so, the total amount I pay for insurance is Mb . So if I want to buy \$1,000 of insurance, I pay M times 1,000.

AUDIENCE: OK. So b is the amount of coverage that you're receiving?

JON GRUBER: Exactly. The amount of coverage you're receiving.

AUDIENCE: And M is--

JON GRUBER: The cost per dollar of coverage. All right, other questions? Yeah?

AUDIENCE: I'm just a little bit lost how we got the second equation for expected utility where we just plug in the--

JON GRUBER: Where we get this?

AUDIENCE: Yeah.

JON GRUBER: So what I did is two things. One is I replace u with square root. That's the easy part. The second is, wherever I had P -- wherever I had M , I'm sorry, I plugged in P . So I place that M with a P -- I flipped the order just to fuck you up. And then I replaced that M with a P , and once again, I flipped the order. Other questions? Good questions. OK.

So, the bottom line is, people want full insurance. And once again, it's useful to go to the book and look through the-- look through the mathematical example-- look through the numerical example just so you can confirm the intuition.

Now, one important feature is that by assuming this utility function, we've assumed a given level of risk aversion. And remember the key lesson from 14.01. Risk aversion is about the curvature of the utility function. The more rapidly your marginal utility diminishes, the more risk-averse you are, because the more it rapidly diminishes, that's how quickly you get sad when things go south.

So no matter your level of risk aversion, you'll always want full insurance. That's true as long as you're at all risk-averse. In fact, as long as you're not risk-loving.

But how much you care about full insurance will depend on how risk-averse you are. If you're risk-neutral, you won't care a lot. And remember our intuition from 14.01, that the value of insurance depends on the size of the risk relative to your income.

Remember, if your income is-- if it's a \$1 risk and you're very rich, then you're risk-neutral. The utility function is approximately flat at that point, so you don't care-- I mean, you still want full insurance, but you don't really care. But if it's a \$1,000 risk, it becomes \$2,000, you care a lot. And so you really are willing to pay a lot for insurance. So that's the role of risk aversion. Yeah?

AUDIENCE: So which of these statements depend upon the particular form of utility there? Or--

JON GRUBER: None of this is dependent on that. You always want full insurance, unless this is a risk-loving utility function, which violates other norms we have in the-- as long as we have a diminishing marginal utility consumption, this is all true.

AUDIENCE: OK. [INAUDIBLE]

JON GRUBER: Yeah. It's just how much you care depends on how risk-averse you are. We'll come back to that. Now, that's all a setup. That's a review of 14.01, which sets up, well, why do we need the government in this business? We've got a \$3 trillion of private insurance floating around, why do we need, really, almost another \$3 trillion of public insurance added to the mix?

Literally, if you look at total insurance expenditures in America, it's about half-public, half-private-- maybe a little more private, but close to half and half. Why do we need that?

And the main reason we feel like we need a public sector intervention is you have to ask yourself, what's the market failure? And the main market failure here is the market failure of asymmetric information. That's the main market failure we're dealing with here, is asymmetric information.

The classic example that goes back to the Akerlof lemons problem. George Akerlof won the Nobel Prize for-- George Akerlof, Nobel Prize-winning economist and half of the power couple because his wife is Janet Yellen, who's Secretary of Treasury, so that's quite a power couple.

George Akerlof made the following talk in-- and once again, it's a review in 14.01, but he talked in his 1970 article about the used car market, and talked about the fact that in the used car market, there's asymmetric information. The person selling the car knows more about the car than the person buying the car.

As a result, you can't necessarily come to an agreement on what the car is worth. And if you can't agree with what the car is worth, you might not be able to make a mutually beneficial trade.

Buyers are perfectly happy to pay a lot of money for a good car. Sellers may be perfectly happy to accept even less money than that for their car, but as they can't prove it's good, they can't make the deal.

Health insurance is the opposite. In health insurance, it's the buyers who know more and the sellers who know less. So basically, the problem is, the insurance sellers don't know exactly what's wrong with you, so they don't know how to price it. And that can, but not necessarily will, lead to market failure. And you want to remember, it can, but not necessarily will.

So to do that, once again, we'll skip over the easy stuff in the chapter, let's go to the appendix. OK. So basically, imagine we're in the same world, but now there's two types, careful and careless. Careful and careless. So now we're going to say the probability of an accident is, a, we're going to say P_a -- I shouldn't use careful and careless, that's really not helpful.

So P_a is the careless and P_c is the careful. Valerie, can you send me an email, remind me to change that in the next section of the textbook? That's not very helpful. OK. So, we're going to assume P -- the probability of an accident, P_a , is greater than P_c . So now we have the same model as before with two types, with different odds of an accident.

Now, imagine that insurers know your type. When you walk into the office, it's stamped on your forehead, it's irremovable, and it says I'm careful, I'm careless. Then this whole model works, they just set two separate prices. They just say, well, I've got a P_c and a P_a , and everything goes through. Nothing changes. It's just-- they say, oh, I see you're careful, so I'm going to do this whole math, and I'm going to set up this math according to P sub c .

If I see you're reckless, I'll do the whole math, set to P sub a , and I'll get different premiums. Which premium will be higher? Which premium will be higher? P_a . Because remember, the premium that you set here, M , is equal to the probability. So you get a higher M because P is higher. Full stop, done.

But now imagine, as in the real world, we don't know who's careful and careless. Now imagine we don't know who's careful and careless. And imagine the insurance company only knows the proportions of each in society. They know the proportions of each in society, but they don't know who's careful and who's careless.

So what's the insurance company do? Well, one strategy is they could say, look, I'm going to assume people are honest. And I'm just going to charge them according to what they tell me they are.

So what happens is, the careful people come in and say, I'm careful, so they say, fine, I am going to charge you-- I'm going to charge you a premium-- I'm going to-- so for the careful person, my profit condition-- so go back to expression for profit, would be M times b . M -- so my profit condition is M , which is going to be b . Remember, people are still going to choose-- they're still going to choose M equals-- right.

So I'm going to choose M careful-- let's write this out, let's not skip steps. M careful minus b times P_c . I'm going to set that equal to 0, just like here. M minus b times P_c I'm going to set to 0.

That is-- and remember, M equals P , so it's just going to be P minus b times P_c equals 0-- wait, hold on one second, one second. Man, I'm getting the math wrong. That, by definition, equals 0. M minus b times P_c -- M , b times P_c -- right. Exactly. So they break even. They make zero profit.

So if they charge you the right premium, which is M equals P , they make zero profit. Contrast that with what happens with the careless. The careless walks in and they say-- what do they say? I'm a careful guy. I walk in carefully, I'm a careful guy, trust me. I'm a careful guy.

So what happens there? What happens with the careless? Well, the careless, they charge them M_c , they charge them M_c minus-- they charge them M_c , they charge them the careful premium. But their payout is b times P_a . That's what they pay out, because actually, they're careless. Which we can rewrite as b times P_c -- M_c is b times P_c minus b times P_a . And we know P_a is greater than P_c , so we know this is less than 0. So profits fall.

So once again, in a world with careful, I charge them Mc -- I charge them Mc minus b times Pc , I rewrite this, and I see that it equals 0. In a world with careless, I charge them Mc , but I pay out b times Pa . So I earn less than 0. So I lose money, I go out of business. The market fails. The market fails when you don't have supply equaling demand. Well, here, if I'm going to lose money, I'm going to be in the market, therefore, there's no supply to meet the demand.

So you can see in this world where we trust people will be honest, you lose money. If we trust people to be honest, you lose money. Questions about that? OK.

Now, let's say I don't trust people to be honest. Let's say instead, let's say instead of trusting people to be honest, I just say, look, I know the population is half-careful and half-careless. And so what I'll do is I'll set the price to be the average, to be Pv , which is the average of the careful and the careless. And then I'll get it right.

I was stupid to trust people to be honest. Nobody's honest. But I know that half the people are careful and half the people are careless, so I'll just set Pv .

Now, at that price, what happens? Well, what might happen, will the careful buy insurance at a price Pv ? The careful--

AUDIENCE: Depending on what other insurers are offering.

JON GRUBER: Well, no, forget there's other insurers, there's one insurers, but it still depends on something.

AUDIENCE: Probably not, but depends on risk aversion?

JON GRUBER: Depends on risk aversion. But let's-- but will the careless buy insurance? For sure, because it's a great deal for them, because Pv is less. Pv is less than Pa , but more than Pc . So for sure the careless will buy it, it's a great deal for them.

The careful may or may not buy. Remember, whether they buy depends on something we derived in 14.01 called the risk premium. Are they risk-averse enough that they're willing to pay the extra amount?

So in that case, if only-- if, in fact, the rich-- the careful don't buy, then you'll see, you lose money by the same math here. Replace this with Pv . And since Pv is still less than Pa , you still lose money. Yeah?

AUDIENCE: Would you expect to be careful to buy some amount of--

JON GRUBER: Hold on, you're getting ahead, we're getting ahead. Let's imagine you buy it or you don't. Now-- so that's kind of the-- let me back up. Let's say that we're selling a certain-- yeah, we're selling a certain product that's for X dollars. So you'll see that basically, you're going to end up losing money even if you use the average price.

But if the careful do buy, then Pv is calculated correctly and you'll be back to making zero profits. So the bottom line is, the market can fail if you don't have perfect information. If you don't have perfect information, you assume honesty, the market will fail. If you don't have perfect information, you use an average price, the market may fail.

So basically, that is kind of-- and whether it fails or not depends on whether we end up what we call a pooling equilibrium. A pooling equilibrium is an equilibrium where everyone does buy even though you can't have full information. As long as everyone buys, you're all in one big pool, so the average premium gets it right. But if you don't have that pooling equilibrium, the market will fail.

Now, that is not the only thing that can happen. Stephen's question comes another point. There's another way the market cannot collapse here. There's another way where the market might exist. And that's if you say, well, you know what? I'm going to offer two separate products. I'm going to offer a crappy product for cheap money and a good product for expensive money.

It is possible-- there are certain conditions, you may or may not have to solve in a problem set or exam. It's possible that careful people are OK with the crappy product at low money because the risk is low. Careless people only want the good product at high money and they're happy to buy it.

So what you do is you offer two products. One is the actual careless product. So remember, we said up here, you can literally just offer two products. If you had full information-- well, if you had information you offer two products. OK, you redo this math, once with P_a and once with P_c .

In that world, you do it once with P_a , you offer that product, and remember, the careless people will buy it because they want full insurance. Then you do it once at P_c , but at a much-- you do it once at P_c , but at a smaller amount. A small enough amount that the careless people aren't interested anymore buying it, but the careful people might want to buy it. And you can get what's called a separating equilibrium.

We have another word for this. We call it HMOs. When I was young, the way health insurance worked is they covered you for everything very generously. And basically what happened was they charged a high price. And since it was cheap, a lot of people bought it, but as the price of medical care started going up, the careful-- i.e., the healthy, started dropping out because they didn't want to pay that price.

So companies came up with this, they came up with HMOs, Health Maintenance Organizations, we'll talk about these a lot more in a few lectures. What are they? They're kind of like mediocre insurance at low prices. And what that did was it offered a cheaper option so the healthy people could still buy something. They created a separating equilibrium. And in chapter 16, we'll talk about the welfare effects of separating equilibrium.

But they created a separating equilibrium where you had really good insurance for the sick and mediocre insurance for the healthy, and both bought.

Now, here's the real tricky insight, which is separating market failure from market collapse. In the world where you counted on people to be honest, the market collapsed because insurers couldn't make money. In the world where you charge average prices and you didn't get a pooling equilibrium, the market collapsed.

In the separating equilibrium, the market doesn't collapse, but it still fails. Why? Why does the market still exist but is no longer efficient? In a separating equilibrium-- yeah?

AUDIENCE: Because the optimal is fully insured.

JON GRUBER: Yes, because the healthy people are getting less than full insurance, so it's not optimal. So market collapse equals market failure. Market existing doesn't equal market efficiency. It's an inefficient outcome because healthy people are getting less than full insurance, which is the first best, which is the efficient outcome. Important distinction.

So the only case in which the market actually doesn't fail is the pooling equilibrium, where everyone still buys full insurance because the rich are just risk-averse-- because the careful are just risk-averse enough. Then there's markets-- the market doesn't collapse and it doesn't fail.

But in every other example, it either collapses, in which case it fails, or it exists, but still fails. It still fails to achieve-- by market failure, we mean not that the market doesn't exist. Plenty of markets that exist have still failed. By market failure, we mean it does not deliver the efficient outcome, the welfare-maximizing outcome. And that's true with the separating equilibrium. Yeah?

AUDIENCE: So just to clarify, complete insurance means that in both the good and bad state of the world, utilities are equal?

JON GRUBER: Yes.

AUDIENCE: So in this case, it's incomplete because healthy people are better off in the good--

JON GRUBER: Yes, because they're not fully insuring the risk. Well put. Well put. So it's not-- we're not getting full consumption smoothing because the healthy people are ending up worse off if they actually get sick.

OK. Let's talk about how this works in practice with a particularly pernicious example, down the street at Harvard University. All the same thing sort of happened at MIT. And the consequence of what can happen, we call this problem-- the problem that arises only because people who are healthy might not want to buy is called the problem of adverse selection. Adverse selection is a problem that arises because only the sick people want to buy, and they might skip the market and it might collapse.

Adverse selection can have very pernicious examples, as we can see-- you don't have to write this down, it's in the book, but let's talk about what happened at Harvard University.

Harvard University had a health insurance plan which was what was called experience-rated. Experience-rated. What that means is Harvard effectively was its own insurer. The insurers administered the program, at the end of the year, they charged Harvard whatever people spent in medical care. They were essentially what's called self-insured. That basically, the more people spent, the more Harvard paid. The insurer effectively didn't bear the risk, Harvard did.

Now, Harvard had a set of plans, some which are more generous and some which are less generous. And Harvard had a system where basically they subsidized the more generous plans, so the gap between the two wasn't very large. The more generous plans didn't really cost much more than the less generous plans. So everyone just was in the more generous plans because why the hell not? Everyone likes more generous insurance.

Then Harvard changed its system where now they allowed there to be-- they allowed there to be a price gap between the two plans that reflected their own experience. So each plan now reflected the actual experience of people in it.

What that meant was, the really generous plans that attracted the very sickest people was expensive. The other plan, which didn't attract as sick people because it wasn't very generous, was less expensive. So the healthy people went to the less expensive plan.

What happens, then? Let's just think about the math for a second. You've got sick people here, healthy people here. Healthy people leave here and move to here. What happens to the price of this plan? Goes up. What happens to the price of this plan? Goes down.

Well, now what's going to happen? Healthy people are going to leave here. Now it's like, well shit, none of the healthy people want to be here, so they keep leaving, and they keep leaving and leaving. And we call that a death spiral. It got to the point where just like three super sick guys left and Harvard just closed down the plan. Same thing happened at MIT.

So adverse selection can actually destroy markets. That's how adverse selection destroys a market, is through that kind of mechanism, where basically-- because each-- because basically, healthy people want to be with healthy people because the premiums are low. Sick people don't want to be with healthy people because they don't want cheap insurance, and the thing drifts further and further apart until it collapses. Yeah?

AUDIENCE: So once the more expensive plan collapses, and those people who were covered under it, I guess, move into the--

JON GRUBER: They end up-- yeah.

AUDIENCE: It just ends up always returning to one--

JON GRUBER: It ends up with one shitty plan. It's a race to the bottom. That basically you end up in one mediocre plan because that plan is the one that everyone can-- that can exist in equilibrium. That's where the new pooling equilibrium is.

The old pooling equilibrium worked because the price gap wasn't very large. Now you get this new pooling equilibrium in a mediocre plan. And reacting to this, Harvard actually then reintroduced the subsidies to bring them back closer together.

AUDIENCE: And in that final state where everyone is in a mediocre plan, presumably people would all be better off if they were fully insured, but there's no way of getting--

JON GRUBER: Well, OK, so let's talk about that. What can the government do about adverse selection? Yeah?

AUDIENCE: [INAUDIBLE]

JON GRUBER: Yeah. The government could enforce a pooling equilibrium. The government could just say, look, everyone has to buy insurance at the average price. Let's go back to our simple example of one plan, one product. The government could enforce a pooling equilibrium so everyone has to buy insurance at the pooling equilibrium price.

So, the government could say, everybody has to pay P_v , and you have to be an individual mandate. That would be efficient in the sense that everyone will be fully insured. But now let's think about it politically. Who does not like that plan?

AUDIENCE: People who don't want to pay--

JON GRUBER: People who are healthy and aren't risk-averse are basically forced into something they don't want. So foreshadowing what happened with the ACA, individual mandate is economically efficient, but politically challenging because you're forced into by people to buy something they don't want. Yeah?

AUDIENCE: Is forcing a minimum standard of coverage, does that also have the same--

JON GRUBER: No. Forcing a minimum standard of coverage would be different because you don't force people to buy a minimum standard, you just make the adverse selection spiral worse because people just drop out.

AUDIENCE: [INAUDIBLE]

JON GRUBER: Yeah, exactly. As long as you have an outside operative with no insurance, the more expensive you make the insurance-- this is what I struggle with, I had to set insurance for the state of Massachusetts under our health insurance plan, and the thing we struggled with is, well, we want people to have really good insurance, but our mandate wasn't that strong. So unless you have a super strong mandate, if you force a minimum standard, you're just going to chase people out of the insurance market.

AUDIENCE: Can you force a full equilibrium, but only if you mandate people have to have insurance, insurance has to be this good?

JON GRUBER: Yeah, exactly. OK, so that's one thing you could do. What's the other thing you could do? How else could you get to a pooling equilibrium? What's one way you could absolutely get to pooling equilibrium if you didn't care about the government budget? Yeah? What? Make insurance free. At a price of zero, you get a pooling equilibrium.

By that same math, subsidizing insurance offsets the adverse selection problem. The more you subsidize insurance with the limit being zero, the less adverse selection there is.

Now by the way, it turns out, in practice-- this is the theory chapter, in practice with health insurance, that's still not good enough. Because even at free, a lot of healthy people just don't bother to sign up. It's never really free. You've got to go sign up, stuff like that. So you still, even with free insurance, get adverse selection, but it's not-- but it's minimized by the amount of subsidy.

So the more-- so there's an interesting trade-off here. The more you make insurance-- subsidize insurance, make it cheaper, the less adverse selection problem you have. So you can mandate or you can subsidize, and that's how you can fight the adverse selection problem. Questions about that?

Now, I started this whole thing by saying, why does the government provide social insurance? Well, adverse selection is one a set of reasons. There's a number of other reasons that matter as well, we just don't spend as much time on them because they're not as fun.

One is externalities, what we started the very lecture with. Remember, I was sick, I had the flu, I coughed, I got you sick. So we started the course with-- I wasn't actually sick, it was the example. In that case, there's another argument for government intervening insurance markets, externalities.

Another argument was mentioned in a question by Stephen, which is administrative costs. This whole model only works, only delivers this efficient outcome under the $M = P$ assumption. So only if insurers-- if we have zero profit for insurers, then we get the efficient outcome.

What if insurers make profit? Well, if insurers make profit, then M will be higher than P . And if M is higher than P , people will not choose full insurance. You can solve that in your own free time, but you can see that if M is greater than P , full insurance is no longer optimal for individuals, which means you will not be the efficient outcome.

So-- and it turns out, administrative costs matter a lot, and they're much higher for private insurance. The administrative costs for private insurance are on the order of 12% to 15%. So every dollar you spend on private insurance, \$0.12 to \$0.15 goes to paying bills and profit and things like that.

The administrative cost for government insurance is on the order of 2%. Why is it so much lower? Well, there's two reasons. First of all, there's no profits being made. But second of all, many administrative costs are fixed, so the larger the pool, the smaller they are as a percentage, which is what we care about.

So an argument for social insurance, you can create giant pools by having the government do it, and that can lower administrative costs. Not in dollars, but as a percent of the premium. So that's another reason why we might want social insurance. We've got adverse selection, we've got externalities, administrative costs.

Redistribution, we might just want social insurance because the private market might not deliver a fair outcome. It might be the rich can afford really good insurance and we don't like that. And then finally, we have the reason that economists hate, but it's probably the major reason governments do things, which is paternalism.

The government may just think, look, the private market is not getting it right, people just aren't figuring this out, people aren't buying insurance when they should. There's all sorts of behavioral problems, psychological issues that people have in buying insurance, so as a result, we need to force them to buy it or subsidize it because people aren't doing correctly.

Now, a version of all this is something I discussed in the book, we don't have time to go through. There's an application about what we call the Samaritan's dilemma. The Samaritan's dilemma is the problem the government faces, which is that because I know, I as a government, can't commit to let you die in the street, I'm going to force you to buy insurance in the first place. How does this work?

That is the following. If I don't force you to buy insurance, you won't buy it. If you don't buy insurance, you get sick. You stumble into the hospital and I treat you. So the Samaritan's dilemma is, because I'm a good Samaritan, because I can't not treat you, I end up spending money on you anyway, so I might as well force you to insure.

The classic example of this is disaster insurance. If disaster happens, FEMA is going to come and rescue you and fix your house and stuff. And as a result, the government has set up what's called the National Flood Insurance Program to make sure that-- to basically subsidize and try to mandate, although it's not very effective, try to mandate people to have flood insurance so that when the time comes, the government doesn't end up picking the cost.

The idea is the government get paid for at the beginning or the end. And by mandate and by mandating insurance, we make you pay for it rather than the government picking it up at the end. So you can actually argue that it's efficient to subsidize insurance upfront to avoid the spending you get in the back. And that is an argument for subsidized flood insurance.

The problem, once again, we run into is the politics, which is that flood insurance, if properly priced, would make houses unaffordable in a number of places where houses exist. And so flood insurance is still not properly priced. People don't-- people pay too little for their flood insurance.

As a result, even though we have a system, and it functions, it constantly runs out of money because it basically aren't charging people the fair price for insurance because that would mean houses on Long Island and on parts of Florida become incredibly expensive, and politicians in those areas don't want that to happen, so the flood insurance prices actually end up where they should be.

So as a result, we end up-- we don't end up in a perfect system where basically people pay x instead of the government paying x post. We have the government subsidize x and the government pays some x post. So because we can't politically set flood insurance prices properly, we end up with this mixed system where we people do buy flood insurance, but FEMA still comes and bails them out at the end. Questions about that?

So, those are all the reasons-- that's the first hour of the class, is all the reasons why we need social insurance, from the detailed math to the intuitive.

Now, remember, the next question of our four questions is, what should the government actually do? And we're going to rephrase that here as, well, how much social insurance should the government provide? Should the government-- how much-- remember, essentially we want to ask, if you think about a social insurance program, and the government can insure you fully or less than fully, what should it do?

Now you might say, well, you've already given me that answer, Jon, the government should insure fully, but there's two reasons why that's wrong. There's two reasons why, in fact, optimal government insurance is rarely 100% government insurance.

The first reason is because-- the first reason why that's not right is because individuals have self-insurance. Remember, the goal of social insurance is not to provide you insurance to smooth your consumption. But if you already have the ability to smooth your consumption, the government is not helping you out that much.

Now self-insurance is real. Think about someone who loses their job. How can they self-insure? What could they do? They've lost their wage income, what could they do so they don't starve? What's that? They could have savings. What else could they do? What else could they do? They could have savings. What else could they do? They could find a new-- well, they can find new job. Let's imagine they can't get a job. What else could they do? Yeah?

AUDIENCE: They could sell their house.

JON GRUBER: They could. They have assets they could sell. They could borrow on their credit card. They might get help from family. They might get help from friends. Their spouse might go to work. There's lots of ways-- people lose their job, but not suddenly, by definition, going to starve. There's lots of other avenues they have that they can turn to. We call those self-insurance.

In other words, they are private market-- it's almost a Coasian point, that there are private market responses to this adverse event. And if we provide social insurance-- so let's take a simple example.

Imagine someone is set up such that if they lose their job, they can easily replace their full income. They could easily-- if someone lose their job, they can easily replace their full income. OK, let's take here's a simple example of that. My wife's aunt. My wife's aunt worked for a manufacturing plant in the Midwest. The manufacturing plant announced, every summer we are going to lay you off with perfect certainty.

Well, what did she do? She saved for nine months of the year so that she could consume the same amount in the summer. She could self-smooth her consumption. I'm a professor. I only have-- I get paid annually. Think about school teachers. My kid's a school teacher, you get paid for 10 months a year. Does that mean you're going to starve the other two months? No. You know with certainty. You just save so you have the same consumption every month. You can self-insure.

And indeed, my aunt was a statewide softball pitching champion during the summer, she had a great time, and her consumption didn't suffer at all because she saved all year and knew it.

In that case, what happens when she collects UI? Let's say UI is available. She got laid off, UI is available. What does UI do in that case? Does it smooth consumption? No. What does it do? Yeah?

AUDIENCE: [INAUDIBLE]

JON GRUBER: Well, no, no, no, but you already have-- you're already start consuming less. What does it do to your budget? What does it do to your decisions? How does it affect your behavior? Yeah?

AUDIENCE: It increases your budget in a way that you--

JON GRUBER: Well, it could. But remember the first rule of government intervention. If you can undo a government intervention, you always will. How can you perfectly undo the government intervention here? Thinking back to fireworks.

AUDIENCE: Save less.

JON GRUBER: Save less. So indeed, full self-insurance is the same as full crowdout. That all UI is doing here is crowding out her savings. All UI is doing is crowding out her savings. It is doing no consumption smoothing. Either way, she consume the same amount all year. The difference has to come from her savings or the government. But there's no consumption being done by the UI program.

Contrast that to another case. Imagine I'm someone who's just like, I have no-- I literally cannot afford to save. Literally, I'm at a job where literally I pay my rent and my bills, I literally have to spend my salary every week. I have no friends, I have no family, I have no credit card because I have no good credit, and I have no savings.

Now I lose my job. In that case, I will die. I am unable to smooth my consumption. In that case, the government-- in that case, there is no self-insurance, so as a result, I get no crowdout. There's nothing to crowd out. But I get full consumption smoothing.

Whereas when I have full self-insure, I get full crowdout and no consumption smoothing. So look what we have here, we have a trade-off. The more self-insurance you have, the more government insurance is not smoothing your consumption, it's just crowding out your self-insurance.

And since the goal of social insurance to smooth consumption, it's less valuable the more self-insurance you have. UI is worth nothing to my aunt in terms of smoothing consumption. It makes her wealthier. That's just a transfer from the government to her. But there's nothing in terms of her desired consumption-- in terms of her consumption smoothing. All it does is crowded her savings.

For the poor sap I just described a few minutes ago, this person, UI is vital. It's fully smoothing their consumption, but it's not crowding any self-interests because no self-insurance to crowd out. So it's the same concept as fireworks. But instead of crowding out Ben & Jerry's contributions, it's crowding out what you're doing to protect yourself.

What does that mean? That means that the value of social insurance is determined by how much self-insurance you have. We assumed, in this example, that each dollar of social insurance was providing \$1 of consumption smoothing, but that's not true. Each dollar of social insurance is providing less than a dollar of consumption smoothing because part of what it's doing, it's just crowding out your self-insurance.

So one reason why social insurance is less valuable than you might think is that all it's doing is crowding out your self-insurance. So it's less-- so that's the first point to remember.

And what is going to determine how much crowding out there is, or conversely, how much consumption smoothing is done? What's going to determine that is two things. The first of all is how predictable is the event? And the second is, how severe is the event?

If an event-- if I said to you, for sure, every year I'm going to take \$1 from you, you will need to insure that. You just know that I just basically have to save a penny less than-- \$0.08 a month. At the end of the year, when they take the dollar, I've got a dollar of savings to replace it. Social insurance would be useless.

But if I said to you, randomly, you're going to be hit by a car and it's going to cost you \$50,000, well, that's going to be hard to self-insure unless you're very rich because you need to have an extra 50K of savings sitting around at all times. That's going to make it hard to self-insure.

So basically, the more predictable is the event, the more you can self-insure like my aunt. That was a major loss. She was losing three months of income every year, but it was predictable. And how big it is, because even it's unpredictable, if it's not big, it doesn't really matter, you can just self-insure it.

So my dollar example, imagine I said I'm going to take \$1 for you, but randomly throughout the year. Or I could say it could be \$1 or \$2. You're like, I don't care, I can just make up for that. So self-insurance will work best when events are predictable and not severe. But the less predictable, the more severe, the less effective will self-insurance be.

Now, that is one reason why our model isn't quite so simple. The other reason is more important, which is that insurance also has a cost. And the cost of insurance is what we call moral hazard, a term you guys might have heard bandied about in various contexts.

Moral hazard, to be specific, is the adverse behavior encouraged by insurance. The adverse behavior encouraged by insurance is what we call moral hazard. Basically, the idea is if you're insured against your house burning down, you don't worry about keeping your fire extinguisher up to date. If you're insured against a bike accident, you might bike a little more recklessly.

Basically, any action you take, anytime insurance changes your behavior, that's the moral hazard response. Anytime it changes your risk-taking behavior or your protective behavior, that's a moral hazard response. In fact, Shakespeare wrote about this. He said nothing emboldens sin so much as is mercy. That basically, the more you're protected, the more you might fuck about.

Now, what that means is that if there is moral hazard, full social insurance is no longer optimal. It's no longer optimal because now there's a trade-off. On the one hand, insurance is smoothing consumption depending on how much self-insurance you have. On the other hand, insurance is inducing moral hazard, which is costly.

And there's no better way to see this than to think about the workers' compensation program, a program we'll talk about in a couple of lectures, probably one of the social insurance programs you're least familiar with, but it's actually the oldest social insurance program in America, started in the beginning of the 20th century, before Social Security and other programs we know about.

This is a program that, if you get hurt on the job, pays for cash benefits. You know your Aflac ads? Aflac is like a top of this. This is the stuff, which is a basic government program, that if you get hurt on the job, they replace some of your income. They also provide you medical benefits, we'll come back to that. But fundamentally, for ease of thinking about it, forget the medical piece. Focus on the cash piece. The idea is if you get hurt on the job, they'll replace some of your income.

Now the Aflac ads will be like, well, they'll replace of your income, so you should buy Aflac, and here's why you should never buy Aflac. Because the truth is, they replace almost all of your income. Worker's compensation insurance replaces about three-quarters of your income that you lose because you're hurt. And it's not taxed, while your income is.

So if you do the math, which we'll do in a few lectures, turns out, you make about almost \$0.90 on the dollar when you're hurt compared to when you're working. Why do you need to buy Aflac at some high premium to top that up? The truth is, workers' comp is already pretty damn good.

Now, the problem with that is, it's also true with workers' comp that it's actually kind of hard sometimes to tell if someone was really hurt on the job. Why? Because most workers' comp is not lacerations-- your arm cut off, it's mostly back pain and mental health issues. Those are very hard to adjudicate, A, if they're even happening; and B, if they were caused at work.

As a result, you have something where it's pretty easy to make it up, and it's very rewarding if you do. And that can cause moral hazard. That can cause people to claim benefits and stop working even when they're not really hurt.

In the next chapter, we're going to talk about Social Security. I'll highlight the end of the lecture that in some European countries, if you work, you make a \$1, and if you don't work, you make \$1.40. What do you think happens? No one works. Moral hazard is a real consequence of these types of programs.

Worker's comp is fun because all sorts of fun examples. You guys ever hear about a show called *Wicked Tuna*? It's a National Geographic about guys who fish for these giant tuna. Turns out, the guy in the show was actually collecting workers' comp at the same time as he was wrestling these 3,000-pound tuna fish and he got busted. There was another guy who was collecting workers' comp, was also a tandem skydiving instructor at the same time.

But my favorite example, which comes from, of course, Florida, where all bad behavior in America happens-- not all bad behavior, but a high percentage.

A woman was working at her desk in Fort Lauderdale when she noticed a ceiling fixture had fallen on her desk. So she looks around, picks it up, bonks herself in the head with it, and says, oh my God, it fell on my head, I need workers' comp, not realizing, of course, the reason I can tell you this story is it was all on film.

So, there is-- now, if it wasn't on film, she would have gotten workers' comp because how are you going to verify the thing bonked her on the head rather than falling on her desk? The bottom line is, moral hazard is real.

Now what determines moral hazard? Moral hazard will be determined by two things. The first is the observability. So what will determine moral hazard? The first is observability. How easy it is to see whether you really need the thing you're qualifying for. Like, if it's your arm cut off, I can see it pretty easily. If it's a back pain, it's hard to verify.

And the second is the financial incentive. Because basically, fundamentally, people are pretty honest. So unless there's a-- makes it worth it, they're not going to try to pretend they're hurt.

Just to emphasize that the second can be important, let's talk about the most disgusting case of moral hazard in American history, which is that in the 1950s and 1960s, the investigators determined that two-thirds of all claims for losing limbs came from one area of Florida. And it turned out, people were cutting off their limbs to get benefits. They even called it Nub City.

So it turns out, even though it was observable, there was some moral hazard because the benefit was big enough relative to people's income. So basically, the point is, moral hazard is real, it can be pernicious, and it's going to depend on these two things. So that's the first thing about moral hazard, is what causes it is going to be observability and financial incentives.

The second thing about moral hazard is that it multidimensional. That's what causes it. The second thing it has is multidimensional. There could be multiple types of moral hazard. The types we typically think about are reduced precaution, or what we might call ex-ante moral hazard.

This is the, I don't have a fire extinguisher in my house, or I don't ride my bike as carefully. Or when I'm climbing the ladder at work, I rush up instead of taking my time. That's the classic example of moral hazard we teach, but it's actually probably by far the least important.

Most people don't want to get hit on their bikes. Most people don't want their house to burn down. Most people don't want to get hurt at work. Except maybe in Nub City, So the bottom line is-- the bottom line is that basically, while this is the easiest case to think about as an example, it's probably the least important case.

The second case is increased odds of entering the adverse state. Odds of adverse event. Increased odds of adverse event. This is people not trying harder in their job. So it's like, this is reduced precaution-- let me separate-- these are hard to separate. Let me be more precise.

Reduce precaution is I don't-- because I have fire insurance, I don't buy a fire extinguisher. Increased odds of any adverse state is, because this is-- I actually misspoke here. This is the, I'm not careful-- this is the, I don't worry about keeping my job because I have unemployment insurance. So clumsy on the ladder is actually in this bucket.

So this is literally, I don't take protective things I could do. So this would be, I don't buy a really good ladder to make sure I don't fall. This would be when I'm climbing the ladder, I'm just figure, what the hell, I'll fall off, I don't care. Or I don't really care about who's in my job. The difference is really in the timing. Do you take precautionary activities or while you're doing the thing, are you careful enough?

The third example that may actually be the most costly is increased spending while under social insurance. This is the idea that, OK, I lost my job. I got hurt. Now I'm not going to try to get better. I'm not going to try to find a new job. I'm just going to hang out.

I lost my job, so basically, I didn't try hard, I entered the state, and now that I'm in the state, I'm getting checks every week, I'm just going to sit around and watch TV, I'm not going to look for a new job. Or I fell off the ladder because I wasn't being careful, now that I'm home collecting 90% of my wages, I'm not going to try to get better. I'm not going to go to PT. I'm not going to do things to get me better because why bother? That can end up being very expensive as well because that increases the duration of your social insurance expenditures.

If everyone fell off the ladder was back at work three days later, we wouldn't really care. We care because they're out of work for six to 10 weeks. And the question is, could they actually be back to work three days later, but they don't bother? That's another type of moral hazard, is extending your stay in the social insurance state.

And then finally, we have the other interesting case, which is supplier responses. This is the idea that basically, because-- this is the idea that because someone has insurance, I spend more-- I make more money on them. So a physician-- I don't-- you guys probably haven't taken your cars for repairs much.

When you take your cars for repairs, the mechanic will ask you, is this for insurance or not for insurance? So insurance, they'll bill you more than if it's not for insurance because you don't care. So you're not price-sensitive, so they'll just charge more.

Doctors will do more if you're insured than you're uninsured because you don't care. And so that is an example of another way that moral hazard can increase government costs, is through these type of supplier responses.

My job after freshman year of MIT, I worked as a filing clerk in-- actually started IAP of my freshman year, in the summer of my freshman year. I was a filing clerk at a doctor's office, I did workers' comp. And eventually I was good enough at filing, they promoted me to assistant bookkeeper. Assistant bookkeeper, and then I got fired. Only time I've ever been fired.

It turned out, much to my justification, the doctor then went to jail, along with his accountant. The reason they fired me, I think, is because they thought I might catch them-- maybe. Because the doctor was running this incredible scam where basically, people would come in for workers' comp claims, and he would just say yes. He wouldn't even really look at him.

He would see one person every three minutes all day-- it was one doctor with five full-time assistants-- full-time secretaries, that should have given some hint, that literally, all he did was just people come in, he'd say, yes, you're hurt, give me-- I'll bill you now. And just do that over and over and over again. So that was a supplier response that increased the moral hazard, affects this program, and I lost my job. So that's another example.

Now, why do we care about all this? Why do we care about any of this stuff? Why do I care if there's moral hazard? Why do I care if the guy takes a few more weeks at home getting better rather than working hard with PT to get better? Or he stays in relax a bit before he starts his new job? Why do I care?

AUDIENCE: [INAUDIBLE]

JON GRUBER: Yeah. I care for two reasons. First of all, why do I care? I care for two reasons. One is government cost. Cost to the government. As we'll talk about in the tax chapters, taxation causes inefficiency. So the more the government has to raise in taxes to pay its costs, the more inefficiency it induces in society. Did you do chapter 4? You did that this Friday? Valerie, was that this Friday?

AUDIENCE: That was last--

JON GRUBER: Last Friday. If you did chapter 4, you talked about the consequences of deficits. We don't like to have deficits, we don't want to have excess government spending. And as a result, it's costly to have these expenditures.

But that may not be the biggest reason why we don't like moral hazard. This is-- we can call this the deadweight loss of taxation. That we care about government costs rising because we've got to pay for it with taxes and that causes deadweight loss.

But there's an even bigger deadweight loss problem here, which is what? What's even bigger deadweight loss problem than moral hazard causes? Why do I care if you're sitting at home not looking for a job? Why would a social planner care? Yeah?

AUDIENCE: [INAUDIBLE]

JON GRUBER: Well basically, more fundamentally, you're not producing your marginal product. And as a result, the social pie shrinks. Remember, labor supply is optimized at the point where your wage equals your marginal value of leisure.

That is the first best for the labor market. You should work until you're indifferent between the next step work and the next hour of producing stuff. Your wage, which I've assumed, by the way, I've assumed that equals your value marginal product of labor. That's really the technical condition. The wage, I'm just assuming a perfect labor market, but this is really the technical condition.

When your marginal product of labor, what you make with the next hour equals how happy you'd be not making something the next hour. Well, what happens when you get a workers' comp benefit? Well now, if you work, you get W . What do you get if you sit at home? You get the marginal value of leisure plus a benefit, which means that by definition, you will have a wage that's less. You will now choose-- you will now only work at a wage that is greater than your marginal value of leisure because you've got to make these equal.

Here, you will work until wage equals marginal value of leisure. Here, you'll work until wage equals marginal value of leisure plus b , which means you'll need a higher wage. Which means that wages that you should work, that would be socially efficient to work, you won't work.

So if I come in and say to you, here's \$15 an hour, and \$15 an hour is your marginal value of leisure, you should take that job. But if you get a \$5 an hour workers' comp benefit, you won't. You'll say, forget it, I'm sitting at home. Society has lost value. The pie has shrunk because people are not working optimally.

Remember, work is not a virtue. Work is something that we do to try to get the socially efficient outcome. Work should only occur to this point-- if someone has a high marginal value of leisure, they maybe shouldn't work. I'm not saying everyone should work.

What I'm saying is, if you are working less than your marginal value of leisure, if the wage you require is higher than your marginal value of leisure, then you're turning down productive opportunities for society, and the pie is shrinking. And that's an even bigger reason we care. There's an inefficiency for moral hazard.

Where does this leave us? Well, this leaves us-- why I love this topic is that we finally, in this class, actually get to write down a condition for optimal government policy. We don't say, well, on the one hand, on the other hand. We actually say, now we know what the government should do. And what the government should do is quite simple. It should provide you social insurance such that the benefits equal the cost. What's the cost of social insurance? The moral hazard. What's the benefit? The consumption smoothing.

So the government should provide you insurance until-- where's our-- over here. Is this-- yeah. It should provide you insurance until the amount of consumption smoothing you get equals the moral hazard it's causing. So in other words, in cases where there's very little self-insurance, then social insurance is valuable. In case where there's very little moral hazard, social insurance is valuable. Social insurance should be generous.

But if you have a case where there's a lot of self-insurance and a lot of moral hazard, social should not be generous. Imagine my wife's aunt, imagine that having unemployment insurance causes you to stay out of work longer and it was doing nothing to smooth their consumption, then there shouldn't be unemployment insurance for her because it was doing nothing to smooth the consumption and it was shrinking and it was distorting social efficiency. It was causing us to have less than the fully efficient equilibrium.

So the optimal social insurance program will almost never be 100% full insurance. Almost never, despite where we started this lecture. Almost never. And the reason is because it will be, to some extent, also it's crowding out, and to some extent, it's causing moral hazard.

So think about starting from 100. You shrink it because it's crowding out, and you shrink it more because it's causing moral hazard, you get a number well below 100. And that's how we think about optimal social insurance sides. And what we'll do in the next set of lectures is talk about those two sides of that equation.

The one side of the equation is how much consumption smoothing is a program doing. The other side is how much moral hazard it's causing. And then we can put them together, they actually tell you how big social insurance programs should be. which is really cool.

It'll be able to tell you what's the benefit, that's how much consumption smoothing it's doing, or alternatively, how much crowdout it's not doing. What's the cost? How much moral hazard it's causing, we'll put together to tell you what optimal social insurance looks like. And that's what we're doing for the next set of lectures.

So I'll stop there with your mind freshly blown. Go away because this is hard stuff, spend some time with chapter 12 and you'll get it, and we'll come back on Wednesday and talk about Chapter 13.