## MITOCW | 14.01-Lecture20-2023Nov15.mp4

[SQUEAKING] [RUSTLING] [CLICKING]

JONATHANSo the new topic we'll focus on today is uncertainty and how you make decisions under uncertainty. Now, theGRUBER:reason it's embarrassing we've taken this long to get to it is most decisions in life are made under uncertainty.

So consider, for example, your decision to study for the final in this course. If you knew exactly which proportion of the final would come from which lecture, you could study them in those proportions, or figure out what you knew best and not and you'd have a clear studying plan. But you don't know that. I do, but you don't. So basically, you've got to figure out in an uncertain world where you're not sure what's going to be on the final, how should you allocate your time to study? That is a decision under uncertainty.

If the forecast said a 50% chance of rain today, you have to decide whether to carry an umbrella. The benefit is you're protected if it's wet. The disadvantage is you might lose it and it's a pain to carry it around. That's a decision under uncertainty.

As you grow up, uncertainties multiply and multiply. Whether to buy a house or rent. Whether to get health insurance. What kind of life insurance to buy. Everything becomes larger and more uncertain as you move into adult life. And so we need, if we're going to have a realistic model of people making decisions, we need a model that allows for decision making under uncertainty. And that's what we'll spend today's lecture on.

Now, once again, while this is late in the course and we're compartmentalizing, you'll be able to see how you can take this model and encode it into the other things we're doing. Indeed, one thing we strive for in the final is to get you to apply concepts across different lectures. So a lot of things we'll talk about here about decision making under uncertainty will immediately tie back to some of the things we learned in consumer theory. And you should be able to see how those things fit together.

So to understand this, let's start by me offering you a bet. And the bet I'm going to offer you is heads, you get \$125. Tails, you pay me \$100. That's the bet. Now, without thinking what the professor wants to hear, just gut instinct, how many of you take that bet? OK, so that's about the proportion I always get. That's about maybe 20% of the class takes the bet. So let's think about that decision, about whether to take the bet or not.

Well, the natural place to start is to compute the bet's expected value. Expected value is simply the probability that you win times the payout. So the probability you win times the value if you win plus the probability that you lose times the value if you lose. Well, for this bet, the expected value is 50%. I mean, I guess I should have specified I'm an honest guy. It's an unweighted coin.

But I presume you guys assume that. 50% times winning \$125 plus 50% times losing \$100 or an expected value of 12.50. It's a positive expected value bet, or what we call in economics a more than fair bet. A fair bet is one with an expected value of 0. This is a more than fair bet. So the definition of a fair bet is 0 expected value. So this bet by definition is a more than fair bet.

But when I offered you this bet, most of you didn't want to take it. Now, I don't think it's that you can't do this math in your head. I think it's that you're human. And humans have a feature of their decision making that we call risk aversion. Risk aversion.

Risk aversion means that individuals value each dollar of winning less than they devalue each dollar of losing. They're made less happy by each dollar of winning than they're made sad by each dollar of losing. And the reason is actually pretty clear when we go back to consumer theory.

And to do this, we're going to turn to a new tool, which is going to be expected utility theory. So it's a new tool which is expected utility theory. The way expected utility theory is going to work is we're going to rewrite this equation, but not in terms of dollar values, but rather in terms of what matters, which is utility. Utility is not about dollars. It's about utility.

So how do we rewrite this expected utility? Well, we'd say, look, the expected utility of a gamble is the probability that you win times your utility if you win plus the probability that you lose times your utility if you lose. Now, if utility functions were linear, I hope you can see this would be the same as expected value. If utility function was just a linear function of your money, it's gonna be expected value.

But in fact, utility functions aren't linear. We typically consider utility functions being concave. That is, we typically assume diminishing marginal product of consumption. We typically assume diminishing marginal utility of consumption or diminishing marginal rates of substitution, going back to consumer theory.

And what that yields naturally without any new math is risk aversion. Someone who has a concave utility or diminishing marginal utility of consumption will naturally be risk averse. Why is that? Well, intuitively, it's because if you have a concave utility function, each additional dollar makes you less happy than losing \$1 makes you sad. And that yields naturally risk aversion.

So let's go through-- let's do an example. Let's use our typical utility function form, which is utility is square root of consumption. A one argument utility function. Make life super easy. All you care about is consumption. Utility square root of consumption. And let's say that C0 is 100. So U0 equals 10. You start with 100 and you end up with 10.

Now, let's calculate the expected utility of the gamble I offered you. The expected utility of the [INAUDIBLE] would be a 50% chance that you win. Well, if you win, someone raise their hand and tell me, what's your utility if you win? Yeah? 15. How did you get that?

## **STUDENT:** [INAUDIBLE]

JONATHANYou won 100. You started with 100. You won 125. So it's the square root of 225 or 15. That's exactly right. WhatGRUBER:is your utility if-- what is your utility if you lose? What is utility if you lose? Yeah. 0. It's 0.5 times the square rootof 0. You started with 100. You lose 100. So it's half of 15 plus 0 or 7.5.

So your expected utility is lower. You started with 10. It fell to 7.5. I offered you a more than fair bet and you didn't take it simply because you have diminishing marginal utility of consumption. I'll come back to why some people might have taken it. But the reason most people don't take that bet is simply because the property that we all derived as natural feature of utility back in lecture two, which is diminishing marginal utility of consumption.

I find a great way to see this as graphical. So let's go to Figure 20-1, which I find a little confusing but ultimately helpful. So let's take our time on this. Figure 20-1 is a different figure than we've ever seen before. On the x-axis is your wealth or your consumption. In this case, you consume your wealth. On the y-axis, we're just graphing your utils And once again, we don't care about the level of utils. We only care about the relationships between different decisions and how they affect your utils.

So for example, you originally start at point A. You have \$100. That translates to utility of 10, 10 utils. Now, I've given you a choice between a gamble with a 50% chance of it ending up at 0 and a 50% chance of ending up at point B.

Well, if you value that, you just say, well, a 50% chance of each means that I end up at point C. I take 50% of point A. I say 50% of 0 or 50% of point B and point C, which is below point A because of the concavity of the utility function. Because by going up towards 225, it gets flatter. So I don't get nearly as much. But when I go down towards 0, it gets steeper. So I'm really sad about losing. 0 sucks. 225 is good, that's great, but I'm not nearly as happy with 125 as I am having nothing to eat this week.

So because of the concavity utility function, you don't want to bet even though it's more than fair. Even though it's the bet that's more than fair, you still don't want it. Once again, no new assumptions here. No new math. This is just using the property of diminishing marginal utility of consumption. Questions about that?

OK, what if utility function is in a different form? What if your utility function instead of the form of utility of square root of C, what if utility function of the form-- actually, no. Let me cover something else first.

So what you see, like I said, you're better off at point A than point C. That's why you choose not to take the gamble. Indeed, here's a fascinating fact. Imagine I came to you and I said, look, I'm your boss. You're going to take this gamble unless you pay me not to. So change the framing. I say to you, you're going to take this gamble unless you pay me not to. Can anyone do the math and tell me how much they'd be willing to pay me not to take this gamble? Can anyone figure out how much they'd be willing to pay me not to take this gamble? Yeah.

**STUDENT:** [INAUDIBLE]

**JONATHAN** What's that?

**GRUBER:** 

**STUDENT:** [INAUDIBLE]

JONATHAN	No, not quite. But let's keep going. Think about it. So let's work out the math of this. You basically are saying I'm
GRUBER:	asking you what payment to me plus the gamble leaves you back at this original point. So if you take the
	gamble, that's as if your consumption if you take that gamble, if you look at the answers here on the graph, you
	look at point C and go to the left. That gamble leaves you at the same level of happiness as having \$56.25.
	Yeah?

## **STUDENT:** [INAUDIBLE]

JONATHANExactly. That gamble leaves you at the same level of have 56. So therefore, you're willing to pay \$43.75 to avoidGRUBER:that gamble. Because by paying that, you end up-- by paying that-- because if you pay me, because if you pay<br/>me \$43.75 and take the gamble, it leaves you indifferent to not taking the gamble. Can people see that? Think<br/>about that for a second. It's stunning.

I've offered you a more than fair gamble, and yet you're willing to pay me almost half your wealth to get out of it. It's pretty amazing, right? That is the power of concavity of utility functions. You are so sad at 0 that you'll give up half your wealth to avoid a 50% chance of ending up there.

Another way to see this. How much would the winning payoff have to be to get you to take the gamble? Someone tell me that. How much would the win-- lose 100, win x? How big is x before you take the gamble? Yeah.

**STUDENT:** 400.

## JONATHAN How did you get 400?

**GRUBER:** 

**STUDENT:** Because the expected value is-- you're trying to get-- you're trying to have your utility be 10 before and after. So if it's half of 4, the square root of 400 is 20 and then half of 20 is 10. So then if you add that to the expected value, that's going to remain as 10.

**JONATHAN** That's exactly right with one wrinkle. It's 300, not 400. Why? Because you start with 100.

**GRUBER:** 

**STUDENT:** Oh yeah.

JONATHANSo you're right. You want 400 in that square root, but you start with 100. So the point is, so another way to thinkGRUBER:of this is literally this is how crazy I have to offer you to take it. Literally tails you lose 100, heads you win 300. So<br/>that is a really massive fair bet I have to give you to take it.

And that's the key lesson of this lecture. Everything we talk about in this lecture and next lecture is all driven by one simple thing. It sucks to go to 0. And the closer you get to 0, that makes you enormously sadder than getting above where you are today makes you happier. And I don't think that's unreasonable as a way to think about the world.

Now, let's do some extensions and applications of this model before we get to do some extensions and applications. To talk about why this model can explain lots of things in the world. Why adding this wrinkle, essentially, remember what I said about modeling in the first lecture? We start with the simplest possible model and then add wrinkles to allow it to explain more things. Here's a wrinkle. That's making our modeling harder. But it's going to explain a lot more things by doing so. So for example, let's change this to instead of utility being the square root of C, let's say utility is 0.1 times C. That's utility. So once again, you start with C0 equals 100, U0 equals 10. But now if I offer you that gamble, do you take it? What's your expected utility? Well, there's a 50% chance you win times 0.1 times 225.

And there's a 50% chance you lose. 0.5 times 0.1 times 0. What is that? It's 12 and 1/2. I'm sorry. It's not 12 and 1/2. If you do that, it's 11 and 1/4. So this equals-- which is greater than 10. So you do take the bet. Why do you take the bet in this case and not in the previous case? Yeah.

**STUDENT:** The function isn't concave anymore. It's linear.

JONATHANExactly. It's linear. So all you care about is expected value. The expected value is positive. You're done. Now,GRUBER:what if we had a different utility function? What if we had a utility function? And we see that in figure 20-2. I'm<br/>sorry. This shows we call this risk neutrality. So you're no longer risk averse. With linear functions, you are risk<br/>neutral. You just do expected value. You don't care about the risk. You're just taking expected value, and that's<br/>shown here. That basically, look, I start at 10. This bet leaves me at point C, which is 11.25, so I'm better off.

So with concave functions, you get risk aversion and unwillingness to take fair bets. With risk neutrality, I'm sorry, with linear utility, you get risk neutrality and an indifference towards risk. But of course, there's also a third case, which is what if your utility function was of the form-- what if utility function was of the form-- oh yeah. What if the function was the form utility equals C squared over-- utility equals C squared over 1,000? Utility equals C squared over 1,000.

I chose this particular function because, once again, C0 equals 100. U0 equals 10. That's why I chose that function. Same starting point. It makes it easier to start from the same point. But now let's ask what happens with this function. What's the expected utility to gamble? Well, the expected utility is 0.5 times if you win. That's 225 squared over 1,000 plus 0.5 times if you lose, which is 0.

What is the value of that? It's 25.3. You are really happy to take this gamble. You've increased utility. Why? Figure 20-3, utility is now convex. This is what we call a risk loving individual. They're no longer risk averse. Not even risk neutral. They're risk loving. They are made happier by taking risk. That yields a convex utility function.

In that case, you are much better off at point C than point A. Because in this case, you don't have diminishing marginal consumption. You have increasing marginal utility of consumption. So going and getting \$1 richer makes you happier than getting \$1 poorer makes you sad.

Now, this is really unreasonable as a description of people in general, but it's not an unreasonable description of some people in some places maybe risk loving. And if you're risk loving, you will take an unfair gamble. In this case, imagine I said, heads you win 75, tails you lose 100. That is an unfair bet. That's a negative expected value bet. Yet this person would still take it. They would get an expected utility from that of 15.3.

So if I said heads 75, you get 75, tails, you pay me 100. So it's a negative expected value of 12.5. Yet with this utility function, that person will take that bet because it yields an expected utility of 15.3. So if you're risk loving, you'll even take unfair gambles.

Indeed, it goes further than that. You will pay me for the opportunity to take an unfair gamble. Risk loving implies that literally, if I'm walking along and say, well, I might offer you this deal if you're really nice. I might offer this unfair bet. You'll pay me to take the unfair bet, because it makes you happy to take that risk. Now, you prefer a fair bet, obviously. More is better. But you prefer an unfair bet to not taking a bet. OK, questions about that? So that's the basics.

Let me do one more extension, which is very interesting. Let me ask another question. Once again, gut reaction. Just give me the answer of your gut. Let's say that I change the example, which is if you win, you get \$12.50. If you lose, you pay me \$10. How many people take that bet? This always happens to. Now you're all being a little inconsistent, because many more of you should take that bet. Why is that? Why? Yeah.

**STUDENT:** [INAUDIBLE]

JONATHANIt's the same bet divided by 10. But why does being worth less money make you more want to take it as a riskGRUBER:averse individual?

**STUDENT:** Less worse off [INAUDIBLE].

JONATHANExactly. Think about this concave utility function back in Figure 20-1. Think about what happens as you shrink theGRUBER:gamble towards the initial point. What happens? The function becomes locally linear. So in fact, the smaller the<br/>gamble, the less concave effectively is your decision making. And once you're linear, you only care about<br/>expected value.

So indeed, in what I just said, if I just did that gamble I gave you guys and put it on an original risk averse utility function, the expected utility of that gamble was 0.5 times the square root of 112.50 plus 0.5 times the square root of 100 minus 10 or 90. And if you add that, you actually get 10.05, which is higher than your initial expected utility.

This is very important. What matters is not the size of the gamble. What matters is the size of the gamble relative to your resources. Risk aversion is relative. What matters is not the size of the gamble. It's a gamble relative to your resources. Bill Gates and you feel differently about \$100 gamble. So basically, and that's because as the gamble gets smaller relative to your initial wealth, utility function becomes locally linear and you become risk neutral. It's a very, very important insight.

You should be able to see this, actually, I can you a couple of things. You should be able to see the same insight answers a couple of other questions. So what this says is as the gamble-- basically the bottom line is as the gamble gets small, or as your initial wealth gets bigger, you become more risk neutral. That's the key lesson. Any questions about that? Yeah.

**STUDENT:** [INAUDIBLE]

JONATHANAh, because we never care about the values. We only care about comparing decisions. Note I always said, well,GRUBER:it's 10.05 bigger than 10. That's all I care about. I never said anything about interpreting 10.05. I just said, is it<br/>bigger than 10 or smaller than 10?

The key thing is we always compare cardinal values, but it's only the relative implication that matters. We're always comparing values. When you calculate utility, you calculate a value. It's just a value in a vacuum is meaningless. It's only meaningful relative to an alternative choice with a different utility value. And that's how you evaluate. Does that make sense?

**STUDENT:** [INAUDIBLE] if you average the 10 and the 20 [INAUDIBLE].

JONATHANExcellent point. That's a great point. Expected utility theory is making a particular assumption about how weGRUBER:combine utility in different states. You could imagine, in fact, when you get into higher order theory, it's a great<br/>question, there's other kinds of models people use for expected utility theory that don't make this simple linear<br/>combination assumption.

And you talked about one reason that assumption may not be great. It's an approximation, because it says that we're combining cardinal things. It's a cardinal sin. We're combining things that we said we really shouldn't be combining. And that's a great point. It works very well in 99% of cases, but there are cases where you get paradoxes with this model. And to explain them, you need more complicated models. Great point. Other questions?

Now, let's go on and talk about a couple of applications. Let's start by talking about insurance. Insurance is big business in the US. US consumers spend about 10% of GDP, or over \$2 trillion a year, on insurance. Health insurance. Life insurance. Accident insurance. Insurance against fire or theft, et cetera. Auto insurance. Insurance is big business, and it's literally 10% of our GDP is insurance.

Why do they do this? I mean, that's a huge amount of money for this product, more than many, many other sectors in the economy. Why is that true? Well, to understand that, let's go through an example. Suppose that you're a 25-year-old, and I'll explain why in a second why I'm being gendered in a second. You're a 25-year-old single male here in Cambridge.

The reason I picked male is because basically, a 25-year-old single male is invincible. The only risk to bad health is getting hit by a car, whereas women can get pregnant, and that's another risk to their health. So let's pick a male who basically the only thing that could cause medical expenditures would be getting hit by a car, which unfortunately is a real risk in Cambridge.

So imagine that you have a situation where your income is \$40,000 a year. That's your initial income. \$40,000 a year. And let's say that there's a 1% chance that you get hit by a car every year. Higher than we'd like, but probably not unreasonable for Cambridge. And if you get hit, you end up with a hospital bill of \$30,000. That's your hospital bill if you get hit.

So how would you evaluate? And I'm going to come along and offer you insurance. And what that insurance is going to do is pay. The way insurance is going to work is you're going to pay me a certain amount, no matter what happens, if you get hit or not. But if you get hit, I pay your medical bill. So let's be clear. Insurance you pay no matter what. If you get hit or don't get hit, you pay me an insurance premium, a certain amount of money. We call that the premium. So you pay me a premium whether you get hit or not. In return, if you get hit, I pay your medical bill.

Well, let's ask, how much should someone be willing to spend for insurance? What is the amount they'd be willing to pay? Well, imagine utility function is square root of C, as we've been working with. Square root of C. So your expected utility is what happens if you do get hit? Well, if you do get hit, there's a 0.01 chance of what? That utility is the square root of what? Well, you still earn your income. I'm assuming you don't lose your job. You still earn your 40,000. And you still pay your insurance premium no matter what. Let's call that x.

But you don't have any medical bill. I'm sorry. This is if you don't have insurance. If you don't have insurance. And then I'm sorry, if you don't have insurance, you pay 40,000 minus the 30,000 medical bill. This is the no insurance case. It's expected utility, no insurance. There's a 0.1 chance that you end up with 40,000 income, 30,000 medical bills. There's a 99% chance that you end up with the square root of 40,000. And that is an expected utility of \$199. That's your expected utility if you don't buy insurance.

If you do buy insurance, expected utility yes insurance, there's a 0.01 chance that you get your 40,000 minus the insurance premium x. We're going to come to what x is. But you don't pay your medical bill. And if you don't get hit by the car, you actually get the same thing. 40,000 minus x. So literally, this is just square root of 40,000 minus x.

So now we can ask, what is the price you'd be willing to pay for insurance against getting hit by a car? And the answer is we answer that by setting your expected utility equal with or without insurance. So your expected utility with insurance is the square root of 40,000 times x. 40k minus x. Your expected utility without insurance is \$199. So this is one equation, one unknown. 40, 40k minus x equals 199. And you solve that and you get that x equals 399. You will pay \$399.

Let's be clear. The reason you will is because that's the premium that equates your well-being in the state with or without the accident. Think about that number for a second relative to what the insurance is worth. What's the expected value of this insurance? What's the expected value of the insurance you're buying? What's the expected? Yeah. \$300. There's a 1% chance that they pay you \$30,000. So you are willing to pay 399 for insurance with an expected value of 300.

Why? Because you're risk averse. You are happy to pay more-- just like you're unhappy to take the gamble that's fair, you're happy to buy insurance that's not fair. Because this is the opposite of gambling. Insurance is the opposite of gambling. So just as a risk averse person will take a gamble, won't take a gamble that's fair, they're willing to buy insurance even if it's unfair.

And we call this extra \$99 the risk premium. So the extra 99 we call the risk premium. How much are individuals willing to pay through insurance to avoid bearing risk? This is a striking fact, and it's consistent with the fact that insurance is big business in America. It's consistent with that, because risk averse individuals will pay to have insurance.

Now, I would like you to do, as an exercise to show yourself that the following things are true. A, as the size of the loss rises holding income constant, the risk premium rises. B, holding the loss constant as your income rises, the risk premium falls. Both of those things simply follow from the intuition we've developed in gambling. Basically, the more you become-- the more the gamble is irrelevant, relevant to your income, the less you care and the less you're willing to do unfair things, like take an unfair bet or pay or pay a risk premium. Questions about that?

That's it. Now, you're right, these are making certain mathematical assumptions. But the idea that people want a risk premium just comes-- the magnitude of this risk premium will depend on what you assume about the form of the expected utility function or the utility function itself. But the notion that there's a risk premium will always be true with risk averse individuals. Yeah?

**STUDENT:** Without overcomplicating the model, do people-- well, we won't even talk about people. If you incorporated the fact that if you are paying for insurance, that's preventing you from buying something else, does that change how much?

JONATHANOh no. That's incorporated. It's square root of C. So by definition if I buy insurance, I can't buy something else.GRUBER:That's incorporated here.

Now, that's one application. Let's talk about another application. Let's talk about the lottery. The lottery in the United States is a total rip off. The expected value of a lottery ticket of every dollar you spend on the lottery in the US is \$0.50. It's a massively unfair bet. Yet it's also massively popular, and a major source of revenue for states is state lotteries. Why is that? How can you have a situation where you have a popular unfair bet?

Now, there's four theories for why this is. Theory one is people are risk loving. They play the lottery because they're risk loving. How can you use something I just talked about 37 seconds ago disprove that theory? How can you disprove the fact that people are risk loving? How can you say Americans are clearly not risk loving? They spend 1.5, they spend \$2 trillion on insurance. They're clearly not risk loving. So that theory goes out the window.

The second theory is a little more subtle, which is that people are both risk averse and risk loving. And to see this, we go to Figure 20-4 of what's called Friedman Savage preferences. You don't need to know that. That's just what they're called. This is preferences developed to explain riddles like this. And here are those preferences look. They're preferences where they start concave and then become convex.

So the way this works is, follow along with me, imagine it's a small 50/50 gamble between W1 and W3. For that small gamble, people clearly don't want to take it. They are worse off with the gamble of point B than the utility of B star. So the same logic I talked about before, people won't take that small gamble.

Now imagine a giant gamble. I'm sorry. Now imagine going from W3 to W5. Well, for that gamble, they become risk loving. They're more than happy to take that gamble even though it's 50-50. It's above the utility.

So now imagine a gamble from W1 to W5. If you do a giant gamble, people might end up looking risk loving. So it could be, yeah, they're risk averse over small things, but they love those big gambles, and that's why they play the lottery.

Why is that empirically false? Why is it an empirically false description of the lottery? Because 99% of the money made in the lottery is not on Mega Millions. It's on scratch tickets, which are you gamble for \$100. So people are in the range where they're at the bottom of this function. If that theory was true, we'd see lotteries, but only Mega Millions. You wouldn't see little lotteries and you do. So that's how we can empirically prove that this theory isn't right.

The third and fourth theories are in some sense the flip sides of each other. And the third theory is that this is entertainment. And this comes to how we go further in Microeconomics 1401. Imagine that I put in a utility function not just consumption, but the thrill of scratching something off, the thrill of seeing, gee, maybe I won money.

And maybe it's worth \$1 to you just for a chance, even though you know you'll probably lose. Maybe utility function is the thrill you get from that. There's nothing wrong or illegitimate about that. It's just the way you write down the utility function. So many people just gamble because they find it entertaining. The thrill is in the utility function.

The problem is there's another theory, which is that people are uninformed or just making mistakes. Which is not that they get utility enjoyment from it. They just don't understand what they're doing. They don't understand what a bad deal. It is, or if they do understand, they just don't think it through properly and they're making a mistake.

Now, we can't really tell. Both these are true to some extent. Why does it matter which one is true? Well, it matters enormously. If we think the truth is one of these, it matters enormously which one it is primarily. And why is that? Because it drives what the government should do.

If this is the correct view, government should support lotteries. It's a great way to make money. It's what we call voluntary taxation. If people like gambling, then why let the private sector make the money? The government can make the money by setting up a lottery. So if that's true, we should be encouraging state lotteries.

If this is true, we should be discouraging them. The government is sanctioning people to make bad decisions. That's not the role of a government, is not to set up more opportunities for people to make bad decisions. So it turns out it matters enormously for government policy which of these are right. It matters enormously which of these are right. Because basically, it depends on is this a voluntary tax or government encouraged mistakes?

And we don't know. But the stakes are huge. In some low income communities, people spend, I kid you not, up to 20% of their income on the lottery. These are huge stakes for many people's lives. This is a very important question to be thinking about. OK. Questions about that?

One last topic I want to cover. And this is going to be a lead into we're going to cover after Thanksgiving. So you have to think about this and put it in your head and chew it over for a week and a half. Which is the problem of information asymmetry. Information asymmetry.

What do I mean by that? Well, let's go back to our insurance fact. I said that insurance is big business, 10% of GDP. You know what else amounts to about 6% or 7% of GDP? Government provided insurance. That is, not only do people spend 2 point something trillion dollars a year in private insurance, the government spends almost another \$2 trillion on public insurance.

Well, you might say, wait a second. How could that make sense? People want insurance. They're buying it. What's the government doing in this business? Why should there be a government role? People are just buying the insurance they should buy it. Some people are risk averse. They buy it. Some people aren't. They don't. What's the government doing with it? And the reason is because private people buy too little private insurance. Not because they're not risk averse, but because of an information asymmetry. Because some parties have information that others do not, and that can cause a failure of insurance markets.

And to see that, let's go to the best example, the classic example, the example that won the Nobel Prize for George Akerlof, which is the lemons problem. Akerlof set up the following problem. And the lemons problem is basically the insight that information asymmetry can cause market failure. Remember, if markets don't fail, the governments just piss off. Government is only useful if markets fail. And Akerlof's insight was that information asymmetry, much like monopoly or other things we'll talk about, can create a market failure. How does it do so? Let's go through a simple example.

Imagine the world of George Akerlof was writing back in the 1970s, where when you bought a used car, there was no Carfax. There was none of that stuff. You had no idea what had happened to the car. When you bought a used car when I was young, you were basically taking a gamble. How good was the car?

So imagine that I have a 10-year-old car that I keep in pristine shape. And let's say that that is worth \$5,000. So the value to John, the value of the car to John is \$5,000. That's what everyone says a good car in good shape is worth. And let's say I'm going to sell to Andrew. And let's say the value of the car to Andrew is \$6,000. For a pristine car, he's willing to pay \$6,000.

This is a welfare improving transaction. It makes Andrew happier than it makes me sad. I am happy to sell it at 5,000. He's happy to buy it at 6,000. There's \$1,000 of consumer surplus created by this transaction. It should happen.

Now, but imagine that Andrew also knows that most 10-year-old cars are crappy. They're not kept in good condition. And he knows that on average, if you buy a 10-year-old car, you'll have to put in \$2,000 of repairs to make it run. So Andrew doesn't know that my car is not crappy. I can tell him it's pristine. I can show him things. It's like, whatever, you could have gotten an accident. I have no idea.

I don't have Carfax. Carfax [INAUDIBLE] the internet. None of that exists. He said, I don't know. All I know is if I buy a 10-year-old car, my value is 6,000 minus the \$2,000 I expect to put in to fix it. So I'm not buying your car. The market has failed A transaction that makes both parties better off has not happened. Why?

Because information asymmetry. Imagine now, in today's world, imagine we have perfect information about how good a car is. Then there's no more information asymmetry in today's world. I could show them the Carfax and everything and say, look, it's pristine. They'd say, great, I know I want to put \$2,000 in. Let's do it. But without that information, he doesn't do it. The information asymmetry is that I know how good the car is, but he doesn't. And the fact that different parties have different information just by itself can cause an entire market to collapse.

This is very important. This is an incredibly deep insight, one of the deepest insights in all economics. That basically a market where nothing else is wrong, no monopoly, nothing else is wrong. Just the fact that the buyer and seller have different information can cause the market to collapse. Here the market is just one sale, but you can expand it to the market to see how it could collapse. And the intuition is with perfect information, this transaction would happen. But without perfect information, it doesn't. Therefore, the information asymmetry has caused a market failure. Questions about that? Now, let's think for a second about how that applies to insurance. Andrew's got his insurance company. Andrew Co. I walk in and I'm like, hey, I'm John. I want insurance. He's like, why? Are you a skydiver? Do you engage in risky sexual behavior? Why aren't you buying insurance through MIT or someone else? I'm worried. I'm worried that there's a problem that we're going to call adverse selection.

I'm worried that there's adverse selection. That is the set of people, the people who come to me to want to buy insurance are only the people who really need it. And I'm going to lose money if I sell them insurance. The information asymmetry is on the other foot now. Now the buyer knows more than the seller does. In the car example, the seller knew more than the buyer did. Now, with insurance, the buyer knows more than the seller does, but it's the same problem.

If Andrew knew that, in fact, I'm a clean living, non-skydiving guy, then he'd be happy to sell me insurance. Say let's say I come and I say I'm willing to pay 399. He says, great, I make \$99 in profit. But since he doesn't know that, he's worried the fact I'm coming in means I'm someone who runs in the middle of the street and waves my arms, gets hit a lot. So he's going to lose money if he gives me insurance. Therefore, he won't insure me, even though society would be better off-- if he had full information, he'd be happy to insure me. Because partial information, he won't. That's a market failure.

The reason government provides insurance is because that market failure leads to underinsurance. That market failure leads people who could have insurance at a fair price not getting it. What we'll talk about after Thanksgiving is how the government specifically addresses that market failure through things like mandates, taxes, subsidies, single payer coverage. We'll talk about that and how the government addresses those market failures through government policy.