

[SQUEAKING]

[RUSTLING]

[CLICKING]

JON GRUBER: Today, we continue our discussion of social insurance by turning to health care. Now, health care is the largest single expenditure of the US federal government. It's the largest single expenditure of most governments around the world.

And it really is an area where-- when we talked about our social insurance programs, I talked about how you need to know some institutional details to understand the area. And I gave you 15 or 20 minutes. This is you need to know an entire sixth of the economy to really understand it, which is a lot tougher.

So we're going to just spend a lot of time in the next few lectures. We'll spend three lectures on this. We'll spend a lot of time on the background on health economics, even to the extent it isn't directly about the public policy because you really need to understand health economics to then think appropriately about the public policy issues.

So what we do in chapter 15 is really cover the basics of health economics. It's a semester of health economics in one chapter covering the things you need to know to be an educated consumer of health care policy and producer of health care policy.

Chapter 16 is then the applications to the biggest government expenditure policies, Medicaid and Medicare. So that's how we're going to lay out. And we'll cover this over three lectures. There's a lot to go through. But the introduction, chapter 15 is my favorite introduction in the whole book. So I'm just going to read it rather than trying to summarize it.

On March 30th, 1981, only two months after entering office, President Ronald Reagan was shot by John Hinckley, a deranged fan of Jodie Foster, who thought the actress would pay more attention to him if he killed the president.

Reagan was rushed to George Washington Hospital, where he underwent three hours of surgery to remove a bullet that entered his lung. Though the injury was potentially fatal, Reagan made a full recovery and went on to serve a full eight years in office.

100 years earlier, President James Garfield was not so lucky. He, too, was shot two months after entering office on July 2, 1881 by Charles Guiteau, who thought God had ordered him to kill the president. Garfield was brought back to the White House, where for 80 days, a dozen doctors attempted to find one bullet that, unbeknownst to them, had lodged itself near Garfield's spine.

The doctors probed the president's wound with their unsterilized fingers and metal rods, succeeding only in widening the wound, infecting it, and puncturing his liver. Alexander Graham Bell ran a metal detector over Garfield's body and soon announced that he found the bullet.

Surgeons went to work, but still failed to locate it, not realizing that bell had mistakenly detected one of the bedsprings underneath the president. The infection resulting from his poor medical care soon caused Garfield to have a heart attack. And when he eventually died, his coroner declared that Garfield would have survived if only his doctors left him alone.

So that's 100 years. That's a big difference over 100 years. And why I love that story, it summarizes the unbelievable change in health care in the US over a hundred-year period. And it's continued, of course, since 1980 over the next 45 years to continue to improve.

But the bottom line is, at the turn of the 20th century, health services were basically a hospital someplace you went to die. Health care was something where essentially it was a 50/50 proposition, whether it help or not. By 1980, we had someone shot in their lung who was surviving with the best health in the country.

Basically, health care has changed enormously. And there's a couple ways to look at it. One way to look at is how much we spend on it. So in 1950, when we started really tracking this stuff, health care was about 5% of the US economy.

So in 1950, it was 5% of this economy. That was less than we spent on cars, fuel, or clothing. Today, health care is 18% of the US economy, which is more than we spend on housing or food. And it's not slowing down. It's been growing historically at a rate well above the growth rate of GDP, so as GDP share continues to rise.

Now we might say, gee, that's terrible. It's going up so fast. But it's not necessarily bad because health care has gotten really good in America. Health care in 1950 sucked. If you had a baby that was born low birth weight, it was four times as likely to die as one born today. If you had a heart attack, you were four times as likely to die as you are today or more relevant for your age group.

If you hurt your knee-- Enoch just hurt his knee. If you hurt your knee, more serious than you, if you tore your ACL in 1950, you were in the hospital for six weeks, or you're in the hospital for a week, you were on crutches for six weeks, and you limp for the rest of your life. Today, the outpatient surgery, you're back on the slopes a week later. Health care is just much, much better than it was in 1950. It's also much, much more expensive.

So at the same time, we have a country with enormous health care disparities. So a common fact you'll hear is that, well, US health care is really expensive, but our outcomes are terrible. We're like 19th in the world infant mortality. That is an incredibly misleading statistic.

A white baby born in America today has as good outcomes as anywhere else in the world. A Black baby born today has worse outcomes than one born in Barbados. So the problem in the US is not that it's bad health care for everyone. The problem is the disparities that lead them to be truly horrific health care outcomes for some groups and pretty good health outcomes for others.

Now, basically, we want to talk about is how do we-- and moreover, the last fact and the most important fact for our course is the US health care system is almost a perfectly split private-public system. Almost exactly 50% of the spending is done by the public sector and 50% by the private sector, within a few percentage points on either side.

So we have a situation where we're starting 50/50. And the question in front of us is should the government doing more, less or different? And that's what we really want to think about as we talk about health care.

To do that, we need to really understand some health economics. We need to understand how health care works before we can talk intelligently about the government's role. So like I said, we'll start with some of health economics basics, and then we'll turn to talking about what the government does.

So let's start with the US health care system. Figure 15-1 shows health spending as a percent of GDP as of 2019 across developed nations. As you can see, the US really stands out as an outlier. There's a pretty smooth rise and then this jump from the second most expensive country to the most expensive country, which is the US.

Where does our money go? Well, then you can go to figure 15-2. About a third of the money goes to hospitals and about a third goes to doctors. A little more to hospitals than doctors, but roughly speaking, it's a third hospitals, a third doctors, a third everything else. Of the everything else, prescription drugs is about 12%. That is a rapidly rising share over time, but still a minority of health care spending. So that's where our dollars go.

The question is, why do we spend so many dollars? Why do we have health care spending that's so much more than the rest of the world? And to be fair, even for people like myself who get great health care, our outcomes still aren't better than the rest of the world. They're comparable. They're terrible for people who don't have good health care. But even for those who have good health care, outcomes are comparable, not better.

So what's the inefficiency? We are spending more. We're getting the same outcomes at best. What's going on? Well, basically, essentially, we can break this down into several different things. One way to start is-- I want to look at figure 15-3. This gives us an international comparison. It just picks five countries, five developed countries.

What we see is the US has the highest health spending per person by a large amount. We also have the worst death rate, although, as I said, that's really driven by disparities. We have the highest preventable death rate. We have the largest number of people that went without health care because of cost of last year and the second longest wait times. Even though we have this free market system, we actually have the second longest wait times.

So why is health care so expensive? Well, basically, if you look at figure 15-4, here is one estimate. This is an estimate of the extra essentially billion dollars. This is billion dollars as of 2018, 2019-- It's more now-- that it's viewed as excess medical spending.

Now, obviously, this is a rough accounting exercise. You can't take every dollar and decide whether it's excess or not. But roughly speaking, this is trying to get at why are health so much more expensive than the rest of the world.

So what's going on? Well, the first category in green is excess administrative costs. That's the easiest one to understand. The fact that it's easy to understand is basically our private health insurance system has much higher administrative costs than do public health care systems.

Our health care system in the US, private health insurance has an administrative cost of about 15%. That is the difference between what-- remember, we solved our model of optimal insurance. We assumed a 0 profit insurer.

That insurer paid out in benefits what they collected in premiums. In reality, in America, insurers collect about 15% more in premiums than they pay out in benefits. That's much, much higher than public insurance systems in other countries.

Moreover, that's only one piece of our higher cost. That's the higher cost of health care insurance. There's also a higher cost of administering health care because we have this multi-payer system. And each payer has a different form. You've got to fill out all these forms and track things.

The estimates on this are varied a lot. But some estimates suggest that as much as 1% of GDP in the US or 1.5% of GDP may just be wasted provider administrative costs. Now that's, I think probably an upper bound.

But the bottom line is there's a lot of excess administrative costs, both in terms of what insurers spend and in terms of what providers have to spend to deal with our multi-payer system.

Now, the next factor is we pay prices that are much higher. This says too high. But let's just start for now saying much higher than the rest of the world. We can see this with international comparison data again. Go to figure 15-5.

So what we do here is we take data from a whole set of developed countries. And for a set of services, we graph the cost of the United States, the cost and the average of these developed countries, and the cost in the lowest cost country. What you see is in every case, the US is multiples of the average cost and many multiples of the low cost.

So if you look at, for example, something simple like a hip replacement, which they do millions of all over the world every year, in the US, that costs on average about \$32,000. This is all in comparable in US dollars. In the average developed country, it's about half that. In the lowest cost country, it's about a seventh or a sixth. You see the hip replacements, angioplasty, C-sections, delivery, appendectomies.

And then you all see it for specific goods. What the second table does is show the cost difference we pay for prescription drugs in the US. Here are common brands of drugs, and we see is once again much higher prices paid in the US than in other countries. So the second reason we spend so much is that we pay a lot higher prices.

Now, why are those prices higher? Well, let's ask ourselves what disciplines prices in a market. We have a market for health care. Unlike other countries, we don't have a single payer health care system. We have a free market for health for many of our payers. What disciplines prices in a free market? What keeps prices closer to marginal cost? Yeah?

AUDIENCE: Competition.

JON GRUBER: Competition and shopping. Well, both of those are problematic in health care. The competition part is problematic because there's an obvious feature of natural monopoly to many features of health care.

If you have a heart attack on Nantucket, you have to go to the one hospital on Nantucket. You can't say, I'm going to go get me on a plane, put me over here a few hours, I'll be OK. You're not going to do that.

Likewise, it's pretty hard to shop. If you're in the back of an ambulance with a heart attack, you're not going to be comparing prices. Moreover, if you wanted to, prices are incredibly opaque and hard to even figure out.

Moreover, the market power here comes not just from geographic monopolies. It can come from also reputational monopolies. Hospitals can just be viewed as the best, regardless of whether they deserve it.

When David Ortiz, beloved Red Sox player, was shot in the Dominican Republic and was flown back to the US for care, did he go to the official hospital of the Red Sox, one of the top 10 hospitals in the country, Beth Israel? No, he went to Mass General Hospital, which is rated somewhat higher. And they decide, well, Big Poppy's got to go to Mass General because that's better.

Now, the truth is I can't say for his case whether it made a difference, but for most people, it doesn't make a difference if you go to MGH or BI, but MGH charges a lot more. That's a reputational monopoly, natural monopoly.

There's basically imperfect competition is rife in the health care provider space. Indeed, the foundational article on health economics is written by Kenneth Arrow, a Nobel Prize-winning economist. Remember, Arrow's impossibility theorem we did earlier this semester.

He wrote in 1963 an article which basically could be summarized as there's no market as fucked up as health care, basically said everything fails in health care in 1963. So we've known for a long time that basically market failures are rife here.

Recognizing that, we find that-- recognizing that, what other countries do is regulate the prices that are charged for health care. Other countries say, well, we've learned in 14.01 that if the market doesn't work, then there could be a welfare loss the government can correct, not necessarily will correct, but can correct. And every other country in the world regulates their prices. And that's why the prices are so much lower in other countries.

If you look at figure 15-6, this shows how much variation we have in prices in the US. This is a pretty cool figure. This is basically across hospitals. Even for the same service, there is a massive variation in what's charged.

Now some of this is geographic. This is for the whole nation. But even if you look within markets, there is things like 2 to 5 to 1 variation. Literally, the cost of basic service in the Boston medical market can vary 5 to 1 across hospitals.

That is not the hallmark of a competitive market, That is not the hallmark competitive market. That basically even within the same market where an ambulance can pretty [INAUDIBLE] take you lots of places, there's enormous variation in prices. It's clearly shopping and competition are not working.

And for example, in Tampa, Florida, if you look at the price of blood tests, there's a 40 to 1 variation in the one city of Tampa in what it costs to get a blood test, 40 to 1. This is clearly competition and shopping are not working. The rest of the world is dealing with this through regulation. How can the US deal with that? And that's something we'll talk a lot about.

Then if you go back to figure 15-4, they have inefficiently delivered and unnecessary services. Let's lump those together What does that mean? That means that we get a lot of health care in the US we do not need.

And here the most famous example, and probably one of the most important health economics pieces ever written, was in 2009 article in *The New Yorker* magazine by Atul Gawande. He's both a physician and writer, brilliant guy, and huge live music fan.

And basically, what he did is he compared two cities, McAllen, Texas and El Paso, Texas. These are two cities very near each other with very comparable population in terms of the age distribution, the income distribution, et cetera. And he looked at the Medicare program, which is I'll talk about next time, pays a fixed price for every service.

So you got the same price. If you do an appendectomy in McAllen and El Paso, you get the same price. And yet Medicare spending per capita was twice as high in McAllen as El Paso. Ignoring the price variation I just talked about, this is purely intensity variation.

So he dug in, and what he found is in McAllen, patients got 50% more specialist visits. They were 2/3 more likely to see 10 or more specialists. They received 20% more abdominal ultrasounds, 60% more stress tests, 200% more nerve conduction studies to diagnose carpal tunnel, and 550% more urine flow studies to diagnose prostate troubles, 2 to 3 times as many pacemakers and 5 times as many home nurse visits.

So what's going on? Well, basically what Gawande argued is that essentially we got into an equilibrium where the doctors in McAllen were greedy and the doctors in El Paso were less so. Doctors in McAllen realized that the more they did, the more money they made, and they just went off the reservation, just started doing a ton of stuff, became the culture of McAllen was to overtreat people.

And he said overtreating because there's a fascinating fact. Despite spending twice as much per capita, people live no longer in McAllen. There was no extra length of life, no obvious improvement in health, despite spending twice as much. And that's what he meant by overtreatment, that basically we're spending doing all this extra stuff and people aren't living any longer.

And indeed, Gawande gets all the credit for this, but, actually, what he did is popularize a resource that economists had known about called the Dartmouth Health Atlas, which was created in the 2000s, essentially took Medicare data and showed for every area in the country what were outcomes and what did health care cost. And there's no correlation.

Health care costs vary enormously around the country, with no correlation with outcomes, suggesting-- and that's for Medicare where prices are fixed. So suggesting that in addition to excess prices, we're paying for a lot of inefficient care.

And that's not just that evidence. There's lots of evidence from things like chart reviews where doctors go back and look at charts of patients treatment and show how they were excessively treated.

For example, by the early 2000s, there were four different clinical trials showing that showing spinal fusion surgery was no better than non-surgical treatment. So one of the most overused procedures in medicine, spinal fusion surgery-- when you have a bad back, they fuse your spine. Despite this evidence, back surgery rates are high and growing in the US, even though medical evidence consistently says we shouldn't be doing it.

Here's another great example, oral intravenous acetaminophen, so Tylenol through an IV, costs hundreds of times as much as oral acetaminophen. There's no evidence it works any better. Obviously, if you're unconscious, there's a difference. But if you're conscious, there's no evidence it works any better. And yet intravenous acetaminophen is used all over the country by many hospitals at a cost hundreds of times as high.

So basically, we have this expensive system because we have a lot of administrative waste. We have prices that are much higher, and we do a lot of extra care. Now there's also another bullet here. If you go back 15-4 for fraud, care coordination failures probably could be lumped in also. That's people not getting passed correctly through their doctors that. Could also be lumped in with unnecessary service, inefficient services.

Fraud is interesting. Once again, we have a lot of health care fraud. So we have lots of situations where people are literally fraudulently billing for services they didn't perform, et cetera, once again, a problem that appears to be worse in the US, not because we're necessarily more dishonest.

It's not the US is necessarily more dishonest. It's that there's so much more money to be made through it. If prices are low, it's not as valuable to commit fraud. When you can make a ton of money, fraud becomes more valuable. And that's why we have a lot more of it. So that's where we are in health care costs. Yeah, Steven?

AUDIENCE: [INAUDIBLE] for example, [INAUDIBLE] Tylenol instead of taking it orally, is that the health care providers are getting some bonus based on what is it they give, or is it just strictly they're not paying the bill because of insurance so they just don't--

JON GRUBER: Great question. And we'll cover that. We'll spend a lot of time on that this lecture. Yeah?

AUDIENCE: We mentioned like reputational monopolies. I also know that a lot of times like, I don't know, insurance companies, and hospitals, and I guess care providers are all either owned by the same company or basically [INAUDIBLE] company? Would that also be [INAUDIBLE]?

JON GRUBER: Yeah, I mean, basically this is a big-- and we'll come and talk about this some. But basically, consolidation in the health care industry has been massive. I think the view is it's been a bad thing for consumers. Yeah?

AUDIENCE: Would you say that regarding drugs, the FDA's rigid patent and approval system either helps the market by like addressing research positive externalities or hurts it by pushing prices really, really high?

JON GRUBER: Great question. And this is the kind of stuff I want to be-- you guys have a lot of questions about health care in your head. I know you do. Let's discuss them. So let's talk about the FDA. This is one I wouldn't come back to normally. So I'll spend a few minutes on this one.

So the FDA, as we learned in chapter 1, regulates all the food we eat and all the drugs we take, et cetera. There's a deep question about how high should the hurdle be to approve a drug. The FDA has what's considered internationally a fairly high hurdle, usually, although sometimes lower, but traditionally higher.

Here's what's interesting about that. I'm not qualified to say if it's too high. But here's what I do know, which is a fascinating disparity in the US, which is to get a drug approved, there's this very high barrier. Once it's approved, go to town. We don't even watch you anymore. And only when someone dies and if the people start to sue you do we pay attention. That's nuts.

OK, it's crazy that it's incredibly hard to get it on the market. Then once it's on the market, we don't do much to keep it-- I'm overstating it. We do keep tabs on you. And if there's a lot of bad incidents, the FDA collects the data and investigates. But the truth is, we're a lot tougher about approving the [INAUDIBLE] that are on the market. That disparity doesn't seem to make a lot of sense to me. Yeah, Steven?

AUDIENCE: Spinal fusion. Isn't there also an issue of-- like that surgery makes people-- like it reduces people's mobility.

JON GRUBER: Well, I mean, that's a separate thing of whether you're actually even worse off getting that surgery than not. I'm not even going there. But clearly, it's not recommended medically. And yet it's done excessively. Yeah?

AUDIENCE: What about generic versus brand name like medications?

JON GRUBER: Yeah.

AUDIENCE: Is that another example of [INAUDIBLE]?

JON GRUBER: Yeah, I mean, look, we talked in 14.01 about the role of patents. And basically, you remember 14.01, we talked about where monopolies come from. And one place monopolies come from is patents. And there's a trade-off with patents.

The trade-off is, on the one hand, you want to create rents to incentivize innovation. On the other hand, once the good's innovated and produced, there's a monopoly. And the optimal patent structure would trade that off.

In a perfect world, we would do what the Nobel Prize winning economist Michael Kremer suggested, which is society would decide what an innovation is worth and would pay you a price for inventing it and then sell it at marginal cost. Then that would be the most efficient thing to do. We don't do that.

So what we do is we're going to set up this patent system. And are patents too long or too short or too tough? That is a fascinating and difficult area. Let me mention one area which illustrates the challenge with patents. This is by a famous article by Heidi Williams, who's at Dartmouth College.

Heidi showed that if you look at drugs and the amount of money that's invested by the private sector, the shorter the trial, the more the private sector will invest. The private sector massively underinvests in drugs where it takes like 10 years to figure out if they work. Why? Because the patent clock starts ticking the minute you file the patent.

So it takes 10 years to figure out if it works, you only have a few years of patent left before you lose your monopoly rights. So that's the kind of thing where patents clearly affect incentives of companies. And there's a whole separate, fascinating field of patent economics and how do we think about the optimal patent, how it should be structured. It's a great question. Yeah?

AUDIENCE: [INAUDIBLE] but that was the one that you mentioned at the beginning of the course of the COVID vaccine.

JON GRUBER: Yeah, I mean, basically, COVID vaccine is one example where we basically put in a prize. The prize in that case was a government guaranteed purchase.

AUDIENCE: [INAUDIBLE] if this was a subsidy [INAUDIBLE].

JON GRUBER: No. Basically, it was a subsidy. It was essentially what we call a pull incentive, saying basically, look, if you invent a good vaccine, the prize is we'll guarantee we'll buy it.

It wasn't just a pure prize for inventing it. You had to then deliver it. But since marginal cost was really low compared to fixed costs, we're essentially going to pay you, guarantee you'll cover your fixed costs if you invent it. Yeah?

AUDIENCE: [INAUDIBLE] we already have [INAUDIBLE] Medicare. Would it be a helpful framework to have more medical innovations sponsored in the way that the COVID vaccine was where you have the prize system that you were mentioning, but the prize is like Medicare is already [INAUDIBLE].

JON GRUBER: It's a great question. I mean, this is something that worked very well in COVID. I think the bottom line on that is you'd want to essentially do that where you think the private market is least likely to deliver. So, for example, classic case is vaccines.

The private market-- I think I covered this-- underinvest in vaccines because they're not sure if the thing will get cured, if it'll go away or not, vaccines or another one is antimicrobial resistance that basically there's all these antibiotics that we're getting [INAUDIBLE] and no one's invested in a new cure to that.

You wouldn't want to say the government does it all, but I think targeting things like that, where clearly there's the private incentives are too low, would be the natural place to go.

So that's the basic of health care spending. Let's now talk about health insurance. Much of our discussion in the next few lectures will focus on health insurance. So let's talk about how health insurance works.

Fundamentally, health insurance is a deal where you pay a regular payment, typically monthly or bi-monthly, which is called the premium, in return for which the health insurer covers the cost of your medical care minus what's called some cost sharing.

So at MIT, I pay my I pay my bi-monthly premium. The health insurer pays my health care costs, but I have to pay some of it too. It's called cost sharing. That cost sharing can come in several forms.

The most common form is a deductible. So at MIT, I'm in a high deductible plan. What that means is the first \$2,500 of my family's health care spending every year, I pay. Then the insurer kicks in.

An alternative is what some of you may have dealt with a lot the doctor, which is a co-payment, which is you pay \$20 every time you go to the doctor. Another alternative is what's called co-insurance, which is you pay X percent of the bill when you go to the doctor. All of these are different forms of cost sharing. So the key feature of an insurance contract is what does it cover?

Under the ACA, insurance contracts cover pretty comparable set of things because we mandated a set of benefits they have to cover. So probably the two big distinctions across insurance products in the US are, A, how much cost sharing you have and, B, which doctors they include in their network. And that comes to the point your team was making earlier. We'll come back to that in a second.

So let's focus on the first issue, which is essentially how much people have to pay. But first, let's talk about health insurance in the US. Table 15-1 shows the distribution of health coverage, just a little out of date. This is now 2020. It's not that different today.

Essentially, we have, three quarters of the population or 70% of the population covered by private insurance. Almost all of that comes through employers. Most private insurers come through employers. But a substantial minority comes through direct purchase, through what we call the non-group insurance market.

So employer insurance is typically of the form-- so I can just give you the numbers for MIT. MIT spends about 24 grand a year on my health insurance, but they don't spend all of that. They buy a health insurance plan for me from Blue Cross, cost \$24,000 a year. They then asked me to pay some of the premiums.

So the premiums are paid partly by my employer and partly by me. So typical arrangement is your employer will say, we're going to get you health insurance, but you have to pick up some of the premiums. So I probably pay about 25% of the premiums. I probably pay about six grand a year. MIT pays 18 grand.

So that is the structure of a typical employer insurance plan will work is they'll buy the insurance. And then typically for a family plan, the employee pays about a quarter of the cost.

OK, a non-group insurance plan is one where you buy it and you pay it all yourself. The employer is not involved. The non-group insurance market was a market that was fundamentally broken before the Affordable Care Act. And I'll talk about that, how the Affordable Care Act fixed it. But now it's a very functional, comparable market to the employer market.

So the question is if you're going to buy private insurance, why do you buy through the employer market? Why is the employer market four times as big as the-- why is the employer market, yeah, four to five times as big? No, 6 times as big as the direct purchase market. Why is the employer market so much bigger? Yeah?

AUDIENCE: [INAUDIBLE]

JON GRUBER: Well, OK. Go ahead.

AUDIENCE: [INAUDIBLE] income tax?

JON GRUBER: Well, OK, that's subtle. Let's get to that in a second. The first reason-- let's forget taxes. The first reason is because of the benefits of pooling. Remember insurance and adverse selection. Remember an insurer who's worried about adverse selection.

What does that insurer want? To minimize the cost of-- to minimize the risk of adverse selection, they want two things. They want a big pool, and they want a pool that's constructed independent of health status. So you don't want the skydiving club

But MIT, that's pretty good. People don't work at MIT-- they're not working at MIT because they're sick or healthy. We got some people, [INAUDIBLE] get some healthy people. That's fine. They want to live off the law of large numbers. That's basically what insurers want to do is they want to live off the law of large numbers.

They want a big enough group they can rely on the law of large numbers and a symmetric enough group that they can map the distribution. With MIT, they can pretty much guess what our costs are going to be year after year.

They can't predict the cancer, but they can know that X percent of our employees got cancer every year, statistically, and some year might be X plus 3% somebody X minus 3%. But roughly speaking, our costs aren't going to double next year. So insurers want large, predictable risks. And employers provide exactly that, especially large employers.

As a result, large employers-- moreover, on top of that, the administrative costs of insurance have a large fixed component. So for both of those reasons, the bigger the group, the cheaper insurance can be, both because insurers can predict it better.

So they have to charge less of a-- they can be more risk neutral, so they can charge a low price. And because administrative costs are lower per person. For both those reasons. Insurance is much cheaper for large firms than small firms. And for that reason, virtually every large firm in America offers health insurance.

The fact is that 94% of firms with more than 50 employees, not even that big, even medium sized, offer health insurance. 94% with 50 employees and 99% with more than 1,000 employees.

On the other hand, among firms with less than 50 employees, only 50% offer health insurance, and it shrinks. If you're down at firms that are less than 10 employees, only about 35% offer health insurance. And that's largely because you lose the risk pooling if you get too small.

Even if your firm isn't the skydiving club, if it's only nine people and one gets cancer, that can double your costs. And so insurers are more wary of that. Insurers become more risk averse, if you will. They say, look, I'm going to charge you more because I'm worried about what's going on. I'm worried about adverse selection. So that's the first reason why employers are a source of health insurance.

The second reason as well as mentioned, which is the tax subsidy employer health insurance, which is what MIT pays me in wages, I pay a tax rate that after my federal income tax plus Social Security plus state income tax adds up to about 50%. When MIT pays me in health insurance, I'm not taxed. It is a non-taxable benefit. As a result, if I'm going to get health insurance, I really want to get it through MIT.

So to see that go to table 15-2. We have two individuals, Nigel and Khadija. They both have the same marginal product, \$30,000, which is their wage. So it's a perfect labor market. They're paid their wage. But Nigel's employer doesn't provide health insurance. Khadija's does.

The health insurance costs \$5,000. So the employer, since it's a perfectly competitive labor market, just takes it out of Khadija's wages. So the employer says, we're going to pay Khadija \$25,000 in wage and \$5,000 in health insurance.

Now imagine there's a 33% tax rate. So after tax, Nigel takes home \$20,000, and Khadija takes home 16,666. Now imagine insurance is actually cheaper outside the employer setting, which isn't true, but imagine it was. Imagine Nigel could actually get insurance outside employer setting for \$4,000.

So they say, great, he should do that. But look what happens at the end of the day, Nigel still ends up with less money. Why? Because he bought his insurance after paying taxes, whereas Khadija got to buy health insurance before paying taxes, so it became tax deductible. She wasn't taxed on it.

Therefore, there is a bias towards getting your health insurance through your employer. Because if you get health insurance, you might as well take advantage of this tax break.

Now here's a very important point, very important point that people misunderstand. If you understand it, you'll be in a very small, elite group of Americans who understand it. Most people don't know about the tax subsidy, but people who do get it wrong. This is not a tax break to employers. Employers don't care. It's a tax break to you, to the employee.

So it's not that employers are treated differently if they give-- if I pay you wages or pay health insurance, either way, I'm paying you \$30,000. I don't care. But as an employee, if I get paid wages or health insurance, I care. This is the employee tax subsidy. It's the employee tax subsidy to employer insurance, not the employer tax subsidy. It's a subsidy to individuals because the employer's indifferent. Either way, he's paying you 30,000. What does he care?

So that's the major way people get their health insurance. The second major way is the nongroup market. As I said before, the Affordable Care Act, this is a very messed up market. And I'll talk more about that when we talk about the Affordable Care Act, probably in two lectures. But now, it's a fairly successful market. We then have the two big government programs. Yeah?

AUDIENCE: The [INAUDIBLE] has this ripple effect on wages. Where does that subsidy actually really [? land? ?]

JON GRUBER: On the taxpayer. So if you think about this, that basically Khadija is getting the same compensation as Nigel, \$30,000. But the US tax system is giving her a lot more money. So if you add up the loss-- if health insurance was taxed as wages in the US, we would raise almost \$350 billion every single year more in tax revenues. That's a lot. So basically, this is a huge, huge tax break.

Because there's no-- I'll come back to this in chapter 18. It's just compensation. We're just saying one form of compensation gets a tax break. One doesn't. If we treated all forms of compensation the same, we'd raise a lot more revenue.

AUDIENCE: Wages being able to fluctuate, though, don't employers claim some of that, at least?

JON GRUBER: Well, that's the open question. Does it all get shifted to wages or not? And the evidence looks like it all gets shifted to wages. Yeah?

AUDIENCE: What's the rationale for doing this? Was it to induce people to buy--

JON GRUBER: The history of it is that during World War II, the government was worried about prices, so they put in place wage limits. You couldn't raise your employees' wages. So employers got around that by introducing benefits as a way to compensate their employees more. And then since those were new, the IRS said, well, I guess they're not taxable. Because it's some new thing, we're not going to tax it. And that became entrenched. So pensions are the same thing. Many other employee benefits are the same thing.

OK. So then we have the two big government programs. We have Medicare and Medicaid. Now, look, I probably won't run you run into you later in life. But if I do, there's only one thing I'm going to hold you responsible for, which is knowing the difference between Medicare and Medicaid. Medicare is insurance for the elderly. Medicaid is insurance for the poor. Now, there are poor elderly, who get both. But the bottom line is, think of it that way. Medicare is insurance-- actually, for the elderly and disabled. I won't hold you responsible for that.

But if you're over 65 or on disability insurance-- we talked about last class-- you qualify for Medicare, which is a single-payer government program. It's our version of what the rest of the world has for everybody. It's a single-payer government program. Or at least, it was traditionally. It's changing over time-- was traditionally. This program, it's paid for by payroll tax. You'll notice on your pay stub, in addition to the FICA row, there's a row for the Medicare tax, 1.45%. The difference with the Medicare tax, unlike the FICA tax, it's uncapped. You pay that 1.45% all the way on your income.

Indeed, as part of the Affordable Care Act, they actually raised that by another 0.9%. So if your income's above \$250,000, it's actually 2.35%. So basically, the way it works is, if you've worked for 10 years in Medicare-covered employment-- so this is like Social Security. If you work long enough to qualify, then when you turn 65, you get Medicare. It's just an entitlement for turning 65.

And I'm going to come back and talk a lot more about Medicare. I just want to give the basic lay of the land. That's Medicare, so for the elderly. Medicaid is much more complicated because Medicaid is targeted at several different groups. So the groups we'll focus on, and where most of the people are, are low-income children and families. That is traditionally.

So traditionally, when Medicaid was set up, it was the health insurance component that went along with the cash welfare we delivered to very poor people. It was originally for just very poor people. In the 1980s and 1990s, it got expanded not just to poor people but to higher-income people for pregnancy and children. So children became eligible if their incomes were higher and women became eligible for the expenses of pregnancy. So it was very poor families. Then it was expanded more for up to middle class for kids and pregnant women.

And then as part of the Affordable Care Act, we made everyone eligible up to 133% of the poverty line. We haven't talked about the poverty line yet, but 100% of poverty line is about \$16,000 for a person or \$32,000 for a family. So basically, Medicaid today is universal coverage for people who are poor or near poor in the US. That is most of the people. That is 3/4 of the people on Medicaid. It is only about 1/3 of the spending.

Most of the spending is the last group covered by Medicaid, which is the people who are people who are using long-term care services, nursing homes and home care, mostly elderly. That's why I said there's the overlap. The poor elderly get both Medicare and Medicaid, but not exclusively elderly. They include disability. That includes some people who are disabled who, say, are mentally disabled, never worked at all. They'll also be on Medicaid. So essentially, most of Medicaid spending is actually on people who are in nursing homes and other institutions, as well as people getting home care.

But most of the people-- it's sort of a universal entitlement to low-income Americans. So that's Medicaid. So if you look at your table, Medicare and Medicaid are each about 56, 59 million people on those programs. Then the last cartridge category of insurance is Tricare. Tricare is insurance for those in the military and their dependents. It's a small number of people. And then, finally, we have the uninsured. So who are the uninsured?

Well, the uninsured are poor, but they're typically not the poorest because the poorest can get free Medicaid. You could think of a typical uninsured person as falling into two groups. One is a very poor person who simply hasn't taken up their Medicaid. Remember, I talked about the problem of take-up. Some people don't take in things they're entitled to. Probably the majority of the uninsured in America, or a large share of the uninsured, maybe not the majority, but near majority, could get free health insurance if they wanted to. They just don't take it up. So that's one group.

The other group is what we might call the working poor. These are people who work, earn more than 133% of poverty, so can't get Medicaid, but don't work for employers that offer health insurance or can't afford the health insurance their employer offers. That's a big chunk of the uninsured, is that group.

OK. So why are people uninsured? Well, remember, in chapter 12, we said, look, the optimum is full insurance. People should want full insurance. Why aren't they buying insurance? Well, there's a number of reasons. The first reason could be that insurance is not actuarially fair. Remember, we said full insurance was optimal under a situation of zero-profit insurance being sold, where insurers sold insurance at exactly what it cost them to deliver the care.

That's not true. As I said a few minutes ago, we have about an average about a 15% administrative load on insurance in America. Therefore, many people don't find it worth it because it's 15% more than what will be actuarially fair. So that's one reason. The second reason is adverse selection. Remember, we talked about adverse selection in health insurance markets. And we talked about the fact that because only the sick people want to buy insurance, and insurers know that, they may not offer insurance or may not offer complete insurance. That's the second reason they'll be uninsured people.

A third reason is, it could be a completely rational decision to forgo insurance. Once again, insurance is expensive. If you're a 25-year-old male, and once again, I'm saying male because there's no risk of pregnancy-- 25-year-old single male. Basically, the main risk is getting hit by a car or some other catastrophe. And if that happens, you can go to the hospital, and the hospital has to treat you. We have a law in America which says that if someone shows up at the emergency room, you have to treat them, regardless of their insurance status.

So for many people, they might say, look, if my only risk is something send me in the emergency room, then why should I pay for insurance? If I have to go, the emergency room will be free anyway. So basically, there's what we might call rational free riding. Because when people go to emergency room, it may be free to them. It's not free to society. That becomes what we call uncompensated care, uncompensated care, which is basically care that hospitals deliver they're not reimbursed for. And that is an externality on the rest of us because someone has to pay those costs. Yeah?

AUDIENCE: And it's free to them because they don't expect the hospital to sue them or because they could go bankrupt. And then [INAUDIBLE].

JON GRUBER: No. The hospital's bill them, but people just don't pay. Turns out, it's pretty hard to get people to pay. There's a paper recently about someone who went to a hospital and tried to buy the debt that people owed. And the hospital will sell it to them at 5 and 1/2 cents on the dollar. So basically, most people just don't pay.

AUDIENCE: So they'll make it that that individual and see it as like [INAUDIBLE] credit score?

JON GRUBER: Yeah, the credit score or something like that. So that's another reason. Another reason is, people just can't afford it. There's a [INAUDIBLE] liquidity constraints. And you can't borrow on your human capital. So maybe health insurance would be an optimal investment. They can't afford it. And then there's just standard paternalism. People may just be making bad decisions. Maybe they should have health insurance and they're just not appropriately calculating the risk and understanding the risks they face if they're uninsured. Yeah?

AUDIENCE: What about very healthy, very rich people that see themselves not needing health care frequently at all? But in the case where they do, they could dish out the money?

JON GRUBER: It's a great point. Remember, insurance loses its value as the risk becomes small relative to your resources. So if you are a billionaire and the largest possible expense that could ever happen is \$100,000, then you're roughly risk neutral. So why pay money to the insurance company? Why don't you pay out of pocket? The problem is, the largest expense is not \$100,000. You can have a medical bill of \$10 million. Even for a billionaire, that's real money. And we're going to come back to what kind of insurance that suggests they should buy. You can probably figure that out already. But that has features of what kind of insurance they should buy.

So that's why people are uninsured. What's more interesting is, why do we care? Like I said in the beginning of the course, lots of people don't have stuff in America. We don't feel like we have to fix all that stuff. Why do we feel we have to fix the uninsurance problem? And one reason I gave-- so why do we care? Why care?

One reason I gave the first lecture was externalities, which is that basically, people uninsured, they might get you sick. Or alternatively, they might get uncompensated care, which amounts to about \$30 billion a year. And typically, our view is that the rest of us pay that \$30 billion through taxes or through higher insurance prices or whatever. So the first reason might care is externalities, so basically, either physical or financial externalities.

The second reason is that by people being uninsured, they don't get care efficiently. And ultimately, in some sense-- externality is sort of 1a is inefficiencies, which is, if people are uninsured, they then get sick, and then that imposes big externalities once they're sick. So for example, there was a famous study I loved back from when I was in grad school, where a guy went to an emergency room and talked to people sitting around. Obviously, these aren't the worst cases. They're sitting around.

But most people in the emergency room are sitting around. And he said to them, how many of you would happily trade this visit for a doctor's appointment in three days? And like 70% said yes. Well, what's the actual percentage? I have that here. The actual-- I'm sorry, 40% said yes. Because basically, they didn't need to go to the emergency room. They just didn't have insurance.

And doctors don't have to take you if you don't have insurance. Emergency rooms do. So that's inefficient. They should not be getting their doctor care in the emergency room. But why not for them? So basically, there's inefficiency that arises from people being uninsured, which ultimately leads to these externalities through health or through fiscal externalities. Then there's the reason that we might actually feel altruism, that basically. We feel that people are uninsured are suffering. Maybe we're paternalistic. We feel they're not making the right decision. And we feel badly for them. We want to make sure they have health insurance. So that's another reason we might care.

But maybe be economically-- so there's no altruism slash paternalism. But there's another economic reason we might care, which is unique American problem of job lock, which is that since people get their health insurance through their job, people may be unwilling to leave their job if they risk becoming uninsured.

So I like to write down a simple model that illustrates this that makes it very clear. So imagine that Emma has utility over two goods, consumption, health insurance. She likes health insurance. Just think of health insurance. She's really valuing not having to pay her bills, but just making it easier to say there's two goods of utility function. So it's $U(C, HI)$, where HI is 1 or 0, whether or not she has health insurance.

Suppose we're in a perfectly well-functioning labor market where compensation equals your marginal product. And as I said, in such a market, if your employer gives you health insurance, they'll just take it out of your wages. That's the way a perfectly functioning labor market would work. So suppose that Emma currently works for a-- what are they, Big Three now, accounting firms? They used to be Big Seven. I don't know what they are now-- Big Four accounting firm. It's a good job. She gets health insurance.

Now suppose she has a great idea for a startup, where she could go and make a shitload more money, but she'd lose her health insurance. So what's her decision? Well, her decision is, do I work in job 1, with a marginal product of 1? So our choice is utility of working in job 1, where she has a marginal product of 1, MP_1 . She pays the cost of health insurance. Let's call that C .

But she gets health insurance-- versus going to job 2, versus the utility of MP_2 , which is higher than MP_1 and certainly higher than MP_1 minus C . So this term is clearly higher. But she doesn't get health insurance. Well, if M is someone who's really risk averse or sick, she might not make that move, even though society cares-- what does society care about? Society cares about this. Society wants people in their most productive positions. That's not happening. Emma is not moving to a job where she'd be more productive because of fear of losing health insurance.

This is an economic inefficiency. And it's big. The hallmark of what makes the US so successful is the mobility of our labor market, the fact that people can seek out the most productive opportunities for themselves. When we put sand in the wheels of that system, we lower the size of the pie. Yeah?

AUDIENCE: Are we to assume that she can't just purchase health care?

JON GRUBER: Great question. So if she could purchase health insurance-- so you can put an upper bound on how bad job lock is. What's the upper bound on how bad job lock could be?

AUDIENCE: Like in terms of--

JON GRUBER: Yeah. So you can tell me mathematically-- you can measure the upper bound on the efficiency cost of job lock. What is that? Yeah?

AUDIENCE: The cost of [? private ?] insurance.

JON GRUBER: What's that?

AUDIENCE: The cost of [? private ?] insurance.

JON GRUBER: Minus the cost of employer insurance. So basically, it's because here, she can buy insurance. So if MP_1 is 20,000, MP_2 's a million, then she just move and buy health insurance. So job lock, you might say it can't be that bad. It's just it's the gap between what your employer insurance cost, the cost of buying your own. The problem is, before the Affordable Care Act, that gap could be infinite because nongroup insurers did not have to offer you insurance. If you were sick, you were not insurable.

I said I'd come back to this. But let me-- I forgot I had to cover this first. The way the nongroup insurance worked before the Affordable Care Act is, insurers were free to discriminate against you. They could not offer you insurance. They could charge you a million bucks a month. They could charge women more than men. Or they could say, for example, if you'd ever had cancer in your life and you ever get cancer again, we're not covering it. So it's possible this gap is infinite, or at least, the upper bound would be the maximum medical spending.

So that's why job lock can still be bad. Now, post ACA, we can always go buy in the exchange. [INAUDIBLE] job lock should have gone down because now that upper bound is defined by what it costs you out in the world versus what your employer spends on you. Now, there was a huge literature devoted to studying this back when I was young, and I contributed to it. And the bottom line is, job lock is a real significant phenomenon. And there were studies which did two things to try to test whether job lock was real.

One thing is, they looked at two employees, both at the same firm, of different health statuses, in a firm that offers health insurance and one that doesn't. So a diff in diff-- sick versus healthy in the firm that offers health insurance and sick versus healthy in the firm doesn't offer health insurance and looked at their odds of changing jobs. And what they found was, in the firm that offered health insurance, sick people were much less likely to change jobs than in the firm that didn't offer health insurance, relative to their healthy counterparts.

Another is to say, look at state laws, which are called COBRA laws. You might have heard of COBRA. It's a law which allows you to continue to buy your health insurance, even when you leave your employer. What we found is, when states passed those laws, people were much more willing to leave their jobs, once again, also indicating health insurance was a barrier to job mobility.

So this is suggesting this is a real potential economic inefficiency and maybe a reason why universal coverage could be economically efficient. Questions about that? Yeah?

AUDIENCE: During COVID, a lot of people lost their jobs because COVID. But then if they got sick, they would not have insurance anymore?

JON GRUBER: During COVID, we basically made insurance super available. So what we did is, first of all, you can always buy in the ACA exchanges. And that is dependent on your income. So if you lose your job, it's free. We also made basically anyone who wanted to eligible for Medicaid. So basically, we had a huge expansion in insurance. Insurance coverage didn't fall nearly as much during COVID as we feared. Yeah?

AUDIENCE: So during COVID, we had like a single payer [INAUDIBLE]?

JON GRUBER: Well, not really. And we still had the private system, but we had a lot looser conditions on getting public insurance. Yeah?

AUDIENCE: The state eligibility, is it just determined by income, or is it also like, generally, the wealth that a person has?

JON GRUBER: No, it's just income. Wealth comes into getting longterm care for Medicaid, like a nursing home. But for health insurance, it's just your income.

AUDIENCE: So if someone is unemployed, they're [INAUDIBLE].

JON GRUBER: Yeah, yeah. So if a super rich person is getting no income, then they're eligible. OK. So last big topic I want to cover today, which is-- that's a huge amount of background. Let's now get into more of the economics. And let's start by applying the lesson of chapter 12 in two ways to health care. The first way we're going to apply the lesson of chapter 12 is to ask, how generous should insurance be to individuals? The second way we're going to apply the lesson of chapter 12 is, how generous should insurance be to providers, in terms of reimbursement?

We now have a two-sided thing. There's both, how much money should we give people? And how much money should we pay the providers who deliver the care? We'll cover the first one today. So what we're going to ask is-- let's think about a simple model where what determines the generosity of insurance-- determined by one parameter, which is how much I ask you to pay, your cost sharing. Let's take everything else out of the equation. It's comprehensive coverage. Everything's covered. The only thing that really varies is how much I make you pay.

So the question is, what is the optimal amount to make someone pay? Well, here, we go back to chapter 12. And we ask, how much consumption smoothing does it do versus how much moral hazard does it cause? Well, a key insight that is developed by the famous economist Martin Feldstein in 1970s was that, in fact, if you look at insurance and ask, should insurance cover all your spending? The answer is almost certainly no because the consumption smoothing's small, and the moral hazard's big.

Let's start with the consumption smoothing. Let's say you have insurance where, as much care as you use in a year, you pay \$100. And I want to lower that to 0. And your job's making \$60,000. That's worth to you almost exactly \$100. There's no insurance value there. There's not a lot of consumption smoothing going on there. You basically pay 100 in premiums to get 100 in protection.

It comes to your earlier question. If the risk is small relative to your income, there's not a lot of utility of consumption smoothing. There's not a lot of benefit of insurance. Insurance benefit comes to protect you against big risks, where you're going to really hit the curvature of that utility function. If it's a tiny risk, you're risk neutral, insurance doesn't do a lot of good for you. So that says that the consumption smoothing will be small-- remember, what determines the size of consumption smoothing?

What determines how much consumption smoothing is delivered, or vice versa, how much crowd out, is if the event is unpredictable and large. Well, if I know that every year, the most I could spend in health care is \$10, it's predictable I'm going to spend \$10 because we all get sick sometime. And it's small. So there's no consumption smoothing. What about moral hazard? Well, Feldstein did this famous little chart we see in figure 15-7.

He said, look, imagine we're going to think about doctor's visits. And imagine doctor's visits have a marginal cost of \$100. That is, the social marginal cost of delivering doctors is \$100. Imagine further that you're in a health insurance plan where they only make you pay \$10 when you go to the doctor. That's your copayment. And then finally, let's imagine the demand for health care is downward sloping, that basically, as the price of health care goes up, you want less of it, like any other good.

In that case, the optimal amount of health care you should get is at point A because the social marginal benefit of that care is dictated by your demand curve. The social marginal cost is dictated by the supply curve. And so the optimal amount of health care you should get is point A, Q1. But in fact, you choose Q2. Why do you choose Q2? You choose Q2 because that's where your private benefits equal your private costs. And that is too much health care.

And indeed, it leads to a somewhat odd-looking deadweight loss triangle. But remember, deadweight loss is where costs are above benefits. And that's everything to the right of Q1 and to the left of Q2. So that becomes your deadweight loss triangle. So Feldstein pointed out that by making health insurance cheap, we induce moral hazard.

And what he pointed out is that basically, what that means is-- we induce moral hazard. So basically, this moral hazard problem has led people to claim that we overuse health care. And it's led people to refer to something that's called the flat of the curve, which you see in 15-8. 15-8's is a weird figure. Let me explain it.

On the y-axis is the dollar of benefit you get from the next dollar of healthcare spending. On the x-axis is how much you're currently spending. So point A says, for someone currently spending \$1,000, the next dollar of spending delivers \$5 of benefits. For someone currently spending \$2,000, the next dollar of spending delivers \$1 of benefits. People understand how to interpret that graph?

We think the graph for medical care production theoretically looks like this, where at first, health care's is incredibly productive, that if you go from getting no health to getting some, it's very valuable. But eventually, you reach the flat of the curve, where additional medical care is doing you no good. Indeed, some would argue it goes below zero, like the spinal fusion example, that the additional health care is actually making you worse off.

But I don't need to go that far. The point is, it just gets flat. And so basically, what that says is that essentially, the situation where-- when you make health insurance free, when you make health care free, you don't do a lot of consumption smoothing, but people get a lot of health care that isn't that valuable.

And that is essentially the driver of how we think about optimal health insurance, is saying that optimal health insurance would almost certainly never be free. Because on the one hand, certainly, optimal health insurance would never have you pay nothing. They'd have you pay something. But you pay nothing-- paying nothing doesn't do a lot of consumption smoothing. When I say pay, I mean what you pay at the doctor, not your premium, what you pay at the doctor, the cost sharing. Optimal cost sharing is almost certainly never zero because it doesn't do a lot of consumption smoothing, but it causes a lot of moral hazard. Yeah, question in the back?

AUDIENCE: Is there a difference [? ranked ?] by the age?

JON GRUBER: By age?

AUDIENCE: Yeah. For example, I don't know if consumption smoothing when insurance is free, it's lower for people that are in the 60s.

JON GRUBER: No, the consumption-smoothing piece will only vary by income. You might ask if the moral hazard piece varies by age because the smoothing piece is about income. If you're a millionaire, at 25 millionaire at 75, either way, \$10 is irrelevant to you. So the consumption-smoothing piece doesn't depend on age. The moral hazard piece might. Yeah?

AUDIENCE: Would the reason that there's an overconsumption of health care in 15-7 be different if this were an individual who was very, very close to poverty, where if they had more money, the intersection of their demand curve with the supply curve wouldn't [INAUDIBLE]?

JON GRUBER: So once again, we're only doing efficiency here, not distribution. But let me ask the question differently, which is, in figure 15-7, what is the key driver of the size of the deadweight loss? There's two key drivers of the size of the deadweight loss. What drives the size of the deadweight loss in figure 15-7? Yeah?

AUDIENCE: [INAUDIBLE] demand.

JON GRUBER: Yeah. Well, it's the cost sharing, the gap between the two lines. But economically, more relevantly, you're exactly right. It's the elasticity of demand for health care. Another version of Steven's question is-- look, maybe healthcare demand is inelastic, especially for a per person. Maybe they only get health care if they really need it. And in some sense, it's not about the money. If you think about it, it's kind of a silly view of the world at one level, which is like, well, gee, people go to the doctor when they're sick. I don't say, that's a good deal. I'll go today. You go when you're sick.

And that was the standard view of the world before the 1970s, was that basically, healthcare demand's pretty inelastic. And as healthcare demand's approximately inelastic, this deadweight loss goes to 0 because the curve gets vertical. So we don't have to worry about this. So the elasticity demand becomes the key determinant.

So there became an enormous field of health economics trying to measure the elasticity of demand of health care. So initially, what studies did is, they said, well, let's compare me to Ahmet. I've got a plan with a \$2,000 deductible. So I pay a lot for my health care. Ahmet's got a plan with no deductible. He doesn't pay much for his health care. Let's compare how much we spend. And what will be wrong with that? Why is that study not causal? What's the potential challenge to that study? I want to compare two people. Yeah?

AUDIENCE: Is there a bias?

JON GRUBER: What?

AUDIENCE: The two groups of you guys will be different.

JON GRUBER: In particular, how might they be different?

AUDIENCE: The person who wants a lower deductible might be more health seeking [INAUDIBLE].

JON GRUBER: Yeah, exactly. So basically, you have to ask yourself-- Ahmet didn't get assigned that policy. He chose. And I didn't get assigned. I chose. Why did I choose a policy where I pay a lot out of pocket, he didn't? Presumably, because he's sicker than me and he didn't want to have to pay out of pocket. So we're going to look at-- we're going to say, wow, check it out. John has a high deductible. He spends less. Therefore, health care must be really elastic. Well, that'd be incorrect. I spend less not because of my high deductible, because people who spend less take high-deductible plans. It's classic bias, as talked about in chapter 3.

So, how do you solve this? Well, as we said, the perfect solution is a randomized controlled trial. And actually, we ran one. We ran perhaps the most famous randomized controlled trial in social policy history in the US, which was the RAND Health Insurance Experiment. And what the Rand HIE did is, they literally randomized-- RAND isn't because of randomization. RAND is the name of the company.

But they literally randomized health insurance. They literally took people, flipped a coin, and assigned them to three different health insurance plans, one where it was all free, one where you paid some, and one where you paid a huge amount. Essentially, think of it as comparing free insurance, where you paid nothing, to one where you had a \$5,000 deductible. That was sort of the comparison they did. So everyone was insured. No one had to pay more than \$5,000. But the difference was, how does it matter if you pay the first \$5,000 versus if you don't?

Now, you might ask yourself, little side note, why would you be in this experiment if you could be assigned to a group where you have to pay \$5,000 for your insurance? And how is that even ethical? How do you set that up? How do you do it so it's ethical and you get people to participate? Yeah?

AUDIENCE: If it's random, you also have the same chance of being [INAUDIBLE].

JON GRUBER: Yeah. But that means, on average, it's a \$2,500 loss. It's random. The 0 or 5,000-- on average, you're still screwed. So what do they do? Yeah?

AUDIENCE: [INAUDIBLE]

JON GRUBER: What's that?

AUDIENCE: I think people who [INAUDIBLE].

JON GRUBER: They could do that. But then you're getting a particularly biased answer. What else could you do?

AUDIENCE: [INAUDIBLE]

JON GRUBER: Give everyone \$5,000. There's no bias to the experiment because the treatment and control are both getting it, and no one's worse off. And that's how you run an experiment. That's how you get people to participate in an experiment with some risk. You make sure that no one's worse off. Now, if you get the zero, you're better off. But that's just a transfer. That doesn't affect your decision to use health care.

So, basically, we ran this experiment. They had several interesting findings. The first finding was that the healthcare demand curve is downward sloping, although it's not very elastic. They estimate elasticity at minus 0.2. That's for every 10% increase in healthcare prices, people consume 2% less health care. So there is deadweight loss, it is downward sloping, but it's not super steep. That was the first finding.

The second finding was that basically, we had a situation where healthcare spending, the different categories of healthcare spending lined up like you think they should. What do you think should be most elastic, and what should be most inelastic? What health services should be most elastic, and what should be most inelastic? Let's get some more folks involved. Come on. This is just intuition. Let's get some of you quiet folks involved. Come on. What should be most sensitive to price? What should be least sensitive to price? You don't need to be a rocket scientist to figure this out.

What services are you going to most avoid if the price goes up? And which ones are you going to not avoid if the price goes up? I don't have to cold call. Let's get some folks speaking here. Come on. Yeah?

AUDIENCE: You get, like, heart surgery, [? where ?] people pay [INAUDIBLE].

JON GRUBER: Yeah. Heart surgery is pretty much inelastic. You're not going to say, yeah, \$20, I'd rather die. So that's a great example. What about elastic? What's something where people-- yeah?

AUDIENCE: Plastic surgery.

JON GRUBER: Plastic surgery. The most elastic thing-- dental care, mental health care, things which are necessary medical treatments for many people, but which people pay attention to price and are more price sensitive. And RAND insurance found exactly, that the elasticities line up exactly as you'd think. Then the most striking finding was, OK, people use less health care. They actually use less of all health care, including preventive care.

So one criticism you hear of making people pay for health care is, wait a second. What if they stop using not just the stuff they don't need but the stuff they do need? And maybe their health is worse. So RAND also looked to people's health and found no effect. Then we're on the flat of the curve. People who, with \$5,000 deductibles, who used a lot less health care, end up just as healthy as people who didn't.

So what's going on with that? Well, there's two possibilities. One possibility is that we don't really know what effective and ineffective health is, and people do. And they just use less of it. And it turned out, it didn't matter. They just use less of it. The other possibility is, maybe the people who used less understood when they could afford to use-- the related is, people sort of understood when it was OK to use less and not use less.

But that was a striking finding, that we really look like we're on the flat of the curve, that going from the 0 to \$5,000 deductible-- now, remember, no one's uninsured. We're not saying you pay it all. There was a cap on what you'd pay. We'll come back to the uninsured next time. There was a cap on what you paid. But at least within that range, where you are spending a little versus a lot, it's not mattering for your health, except they did find suggestive evidence of effects for low-income, very sick people, what we call chronically ill people, people with diabetes or high blood pressure or depression and other things, chronic conditions.

For that group, they did find suggestion that these higher prices were bad for them. And that comes-- someone mentioned different incomes. That is a very interesting lesson for how we think about optimal health insurance, which says that basically, what RAND tells us is that if you are in a range where the money is not that meaningful to you, then probably, it doesn't really affect your health about whether you use the health care or not. If it's really going to affect your health, you'll get it anyway.

But if you're poor and sick where you really just can't afford it, you might not have it, even if it could improve your health. And that leads us to a subtle explanation of optimal health insurance, which is, optimal health insurance is not just saying, well, we should have a deductible. So before I made that point-- in fact, let me go back. I did this in the wrong order.

If I'd stop before that last point, if I just said, look, Rand found, \$5,000 deductible, no worse health, you'd say, well, great. We should at least have a decent amount of cost sharing, not a lot of consumption smoothing, and reduce moral hazard without impacting health. But then there's this other finding that for low-income people, that may impact their health. What does that suggest we should do? That suggests that we should have income-related cost sharing, where the amount you have to pay when you go to the doctor should depend on the resource you have to pay it.

In that way, you reduce moral hazard to people where you can really reduce it without sacrificing health. But the ones that will sacrifice health, you avoid them reducing their care. That would suggest maybe an income-based cost sharing. And we actually have that in the Affordable Care Act. We actually have, the higher your income is, the more you pay out of pocket, to try to address that exact concern. An income-based out-of-pocket system is one form of optimal insurance.

The other piece of optimal insurance comes to something we call value-based insurance design, which is saying, look, we do know with certainty some services are high value, and some are low value. For example, if you're diabetic, there's a drug you can take which is \$1 a day, which basically keeps you from getting your foot amputated, keeps you alive. That's effective. You should be doing that. On the other hand, we know this back surgery is ineffective, maybe even bad for you. So what a value-based insurance design would say is, look, we ought to vary what people pay by how effective the service is.

So insurance, it's not just one cost sharing based on your income but actually varies based on how appropriate the service is for you. You could even have personalized cost sharing, where you say, well, John's a diabetic. Therefore, we're going to make this pill free for him. But his back surgery is going to be really expensive, et cetera. And that would be the ultimate way to design insurance. The problem is, it's really hard. So I was appointed-- when we passed healthcare reform in Massachusetts, the law was incredibly vague. The law said, we should have health insurance. There should be good benefits. Go figure that out. And they passed that on to a board of 10 people, of which I was one. We had to figure out what constitutes good health insurance and what people should pay for it.

So I was pretty excited. I called my buddy, who came up with the term value-based insurance design, sort of the king of this. I said, this is so exciting. I'm in charge, man. We can set up value-based insurance design. What should we-- he goes, this is so exciting. He said, let's charge higher copayments for low-value services and lower copayments for high-value services. I'm like, that's awesome, I want to do that. What should we actually do? And he said, let's charge higher copayments for low-value services and lower copayments for high-value services. Because it's actually hard to figure out how you should actually do it.

So the problem is-- this is a standard thing economists run into in policy design, which is, the optimal policy in theory may be actually hard to implement in practice. So where do we come, in terms of optimal health insurance? Where we come down, in terms of optimal health insurance, is that basically, zero cost sharing almost certainly doesn't make sense unless you're very poor, but that the cost sharing you pay should depend on your income and probably should depend on what the service is as well. OK. I'm going to stop there.