

[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: So what we're going to do today, chapter 21, is talking about how taxation affects labor supply. So we're going to talk in the next two lectures in particular about how taxation affects two behaviors. One is labor supply, and one is savings.

So we'll just start. The base of taxation labor supply we've kind of done before. It's really just another version of what we did for welfare. But let's just remind ourselves of how it works.

You've got a standard consumption leisure trade off. Here's leisure. Here's consumption. We've got a standard consumption leisure trade off.

And imagine you have an individual whose wage rate is \$15 an hour, during the Massachusetts minimum wage of \$15 an hour. Then basically they have some trade off between consumption and leisure, where the slope of that line is minus 15. That is the rate at which they can trade off consumption and leisure. Alternatively, importantly, it is the price of leisure.

Here's the key thing. We talked about this in 14.01. In economics, we don't know how to measure bads. We only know how to measure-- we don't know how to do optimization with bads. We have optimization with goods.

So if you're choosing between two goods, I know how to do the optimization. If it's a good and a bad, it gets harder. So we do to come back to our basic framework as we turn all the bads into goods.

Labor supply is a bad. Remember, regular people in the real world don't like working. It's something MIT students have a hard time understanding.

But actually, leisure is a normal good. People would rather not work than work. That's our standard assumption. As a result, labor is a bad, and leisure is good.

So what we do is we model consumption and leisure, and then we just create labor as the converse of leisure. So whenever we're talking about anything affecting labor supply, remember, the way we think about it is, first, how does it affect leisure? And then labor supply you solve as a residual. We're not modeling labor, modeling leisure.

So essentially now we have a well-defined trade off of the kind we've been doing since 14.01, lecture 2, which is the trade off between two goods, consumption and leisure, and that the price at which you trade those off is the price of leisure. The price of leisure is the opportunity cost of-- the opportunity cost of leisure is basically that you could be working, earning \$15 an hour.

So you have some indifference curve. And there's some tangency at point A. In the book, we say, let's say, point A happens at 900 hours of leisure. And at 900 hours of leisure, that means you're working-- and let's say there's a total of 2,000 hours possible.

Valerie, once again, can you send me an email that's not labeled in this graph, 21-1? I've got to add the x-- I got to add the x-axis labels.

So you have 2,000 hours possible of leisure. You're going to take 900 hours, which means you're going to work 1,100 hours. If you work 1,100 hours, you make 16,500. So that is your initial allocation. You take 900 hours of leisure, which means you work 1,100 hours, and you consume \$16,500. That's your initial choice.

Now, let's say the government rolls in and taxes your wages. Let's say the government particular rolls in and says, we are going to have a 33% tax on earnings. What does that mean? That means that the price of leisure has now fallen. The price of leisure has now fallen.

If you take 2,000 hours of leisure, nothing's changed because you're not earning anything. But once you take less than 2,000 hours of leisure, it costs less to take that hour of leisure. It used to be an opportunity cost of \$15 an hour to take that hour of leisure. Now, the opportunity cost is only \$10.50.

The slope of this was 15, minus 15 was the slope, minus 15. The slope of this is minus 10.50 because the 33% tax. So we've basically lowered the opportunity cost of leisure. What does that do to your decision of how much leisure to take? Yeah?

AUDIENCE: [INAUDIBLE]

PROFESSOR: What?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Take more leisure. Why?

AUDIENCE: It's cheaper.

PROFESSOR: It's cheaper. And we call that what?

AUDIENCE: The sunk cost fallacy.

PROFESSOR: No. No. Good try, though. What do we call that? That's-- yeah. That's the substitution effect. But remember, there's two effects. What's the other one?

So the substitution effect says that we should take more leisure. The price has fallen, so you should do more of it. That's unambiguous.

But there's also an income effect. You are now effectively poorer. For any amount of work you do, you have less money. You're now effectively poorer.

So that income effect says, when you're poor, you want to buy less of everything. One of the things you want to buy less of is leisure. So that says take less leisure. So substitution and income effects, while they work together in consumption, go against each other with factor supply, like labor supply.

So actually it's ambiguous what a cut in wages will do. Think of it this way. Here's the way I like to-- think of an extreme example. Imagine you have a fixed income target. You want to work enough to earn x.

Well, in that case, it's unambiguous that cutting your wages will cause you to work more, not less. You got to work more to get to your target. That's an extreme case.

But in general, you can see it's ambiguous which way this will go. You could end up taking more leisure or less leisure, and it depends on whether income or substitution effects dominate. Valerie, is [? Ahmed ?] on the way?

VALERIE: [INAUDIBLE] I think five minutes later.

PROFESSOR: OK. So basically, the bottom line is, in this simple model, it is ambiguous whether taxing wages will lead to more leisure or less leisure and, therefore, more labor supply or less labor supply.

Now, of course, this is a very simple model of the labor market. There's a lot of complications we don't have here. For example, overtime creates a kink in the budget constraint. When you work more than a certain hours, your wage changes.

There's the fact that you don't necessarily get to choose your number of hours freely, which means you don't necessarily get to pick the optimal point on this distribution. There's complications like that. But for now, in this simple model, we've got just the trade off, the income and substitution effect trade off.

So which dominates? How do we decide whether income or substitution effects dominate? Well, we go to the data. And we go to the data, and we go to the data. And there's essentially been four different approaches to trying to measure, how do wages affect labor supply?

The first approach is a cross-sectional regression. So this is the old-school approach, which literally just regress. If my wages go up, do I work more? Or if my wages are higher, do I work more?

That is the standard chapter 3 non-comparability problem. High-wage people probably may work harder. Some with a high wage may work harder because they're intense. That's why they have a high wage, not because it actually costs them to work harder. So there's a standard unobserved variable problem, so the cross-sectional thing doesn't really work.

There was actually a real randomized controlled trial of this in the 1970s, something called the negative income tax experiment. What they did is they actually literally ran an experiment. They literally assigned different tax rates to people. And they looked at what happened to their labor supply. So that's the second approach.

There's quasi-experimental approaches, quasi experimental. And here, there's two famous approaches. One is a diff-in-diff approach of the type we discussed in chapter 3. So in particular, imagine you have a tax reform that changes tax rates for one group and not another group. You can look what happens to labor supply of the group whose tax rates went up versus the group whose tax rates didn't, so a standard difference-in-differences analysis, the kind you should be familiar with now.

Another approach would be to do a regression-- would be to do what's called a bunching approach. So the bunching approach, so DD, and then the other, it's a new method we haven't talked much about. But it's gaining a lot of favor. It's been really popularized by the economist Raj Chetty at Harvard.

And the idea of the bunching approach is the following. Imagine there's a kink in a budget constraint. Basically, if people are-- so think of the extreme case. Think of the window tax case.

Remember, the window tax case where basically, once you went above a certain number of windows, there's a huge marginal tax rate for going up. And we showed a lot of people had right below that number of windows. That's a bunching estimate, which is, if people are elastic, you'll see them all bunching right before there's a disincentive coming ahead, right at the notch or the kink.

And that's the idea of this bunching estimator is you can infer elasticities from how much people respond to kinks in the budget set. If they don't respond much, then they're not very elastic. If they respond a ton, they are. So basically, you could have imagined the number of windows wouldn't have changed because people are like, I want to breathe. I don't care about the window tax.

That would say it's not very elastic. The fact that it was bunching of the number of windows right at that limit suggests people were willing to sacrifice some breathing to avoid paying the window tax. And indeed, there's-- I can't draw the graphs well enough, but in the book-- yeah, sorry. Go ahead.

AUDIENCE: [INAUDIBLE]

PROFESSOR: OK, describe [? the ?] [? bunching ?] already. Excellent point. Great, great, point. And I almost made the very mistake here.

The [INAUDIBLE] bunching is the following. An RD estimate, a Regression Discontinuity Estimate-- thank you. Let's just pass them, pass each pass each row here. A regression discontinuity estimate is an estimate where essentially you say, on either side of a cutoff, there's different incentives.

And you look how-- and an RD estimate is basically sort of a difference estimate. You look and ask, do people behave differently on either side of a cutoff? Here it's just a cutoff. It's at the cutoff to behave differently than on either side of it.

So basically, so now you have the handout. So if you look at figure 21.3, this shows Chetty's work around a tax increase. So what he did was he looked at Denmark, where there was a big tax increase around 267,000 kroner. There was not an increase. There was a big notch. Basically your tax jumped at that point.

So basically, he's not doing an RD. This is not something we see dramatic things on either side. It's that you see a dramatic thing at one point relative to what's on either side.

It's a break from smoothness. It's not a discontinuity in the series. It's a break from a smooth series, where it kind of gets non-monotonic.

So basically you see here, this shows essentially how many people are at each income level in Denmark. And what you see is, at 267, there's a little bit of a peak. So in other words, it's a smoothly declining function. But then there's a little peak at 267, which suggests that people are responding by bunching there, but it's not very big.

That's for employed workers. Now look at self-employed workers. Self-employed workers, there's a massive peak at exactly that point. Self-employed workers are responding very, very highly to this incentive.

Employed workers are not. What's the difference? Why? Why do you think there's such more elastic behavior among these self-employed workers and the employed workers? Yeah?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah, self-employed workers are unconstrained. They can work as hard as they want. They can also cheat more.

But certainly, whether through cheating or actually changing the labor supply, they can set their own hours.

Employed workers can't. It's hard to target I'm going to go exactly 267,000. So that's why and what Chetty found was much more elastic response among the self-employed than the employed.

So you have all these methods. These methods actually deliver a fairly consistent answer, which is, if you look back in the old days, say, 1980s and before, you could think of the world as defined as primary earners, e.g. men, and secondary earners, e.g. women, especially married couples, especially with kids.

In that world, what they found was primary earners, e.g. men, were not at all elastic-- basically there was very inelastic labor supply-- but that secondary earners, e.g. moms, were fairly elastic. Elasticity is between 0.5 and 1. The argument for why comes back to early substitution effects.

If you're a guy in 1975 and you don't work, there's nothing else to do. None of your friends are home. They're all working.

Where if you're a married woman in 1975, you have a choice. Some women are home. Some women are working. So it's more plausible to be elastic.

You wouldn't know what to do with yourself if you're a guy at home in 1975. There was no internet. None of your friends were home. There's only so much golf you can play, basically, whereas, for women, there was a choice.

They could stay home, raise the kids. Some of their friends were home, et cetera. So it makes sense. But here's what's fascinating.

Over time, we are losing the distinction in primary and secondary workers as women become more integrated in the labor force. We're losing that distinction. What's happening is men's labor supply remains inelastic, and women's labor supply is becoming more inelastic. So overall labor supply is becoming just much more inelastic, where everybody works. OK, yeah?

AUDIENCE: Is there no [? significant ?] [? apart ?] from the men's side as well?

PROFESSOR: No. This is what's fascinating. So what's happening is men now just work and women just work, and kids are in childcare. So the difference is, it's not that the burden of home care has shifted, where men are more flexible and they might stay home.

Now, obviously, I know a stay-at-home dad. There are more stay-at-home dads than there used to. But by and large, it hasn't really changed for men. It's women are becoming more like men in terms of their working behavior, in terms of more like primary earners. So we're getting to a much more inelastic world of labor supply overall. And I would say, roughly speaking, you can view the labor supply margin as a fairly inelastic one in the US economy.

Now, of course, that is if we measure labor supply by things like hours worked. But in fact, that may be the wrong thing to measure it by. In fact, it may be, in the new economy, that what's more important than how many hours you work is like things like how hard you work and what type of job you have.

And you could imagine, taxes, while they might not affect your hours of work, could very much affect things like how hard you work to get a promotion or what type of job you take. If taxes are high, maybe you don't care as much about taking that high-wage job, working hard to get that promotion and get that high-wage job. But if taxes are low, maybe you do.

So that's ultimately what we might care about is not hours of work but total labor earnings. And this is a foreshadowing of chapter 25, when we'll ultimately conclude that, really, in some sense, it's total income that we care about, not these components. So we're building up to that.

But nonetheless, at least on the easily measurable aspects of labor supply, it doesn't look like there's much of an effect of taxes. Questions about that? All right. yeah, Steven?

AUDIENCE: Are there any other countries where the labor supply is noticeably more or less than others? Or is that not [INAUDIBLE]

PROFESSOR: Oh, I don't know. I don't know. So that's the basics of tax and labor supply. Pretty straightforward.

What I want to do for the rest of chapter 1 and 21 is apply that talking about two interesting policy issues. The first is something we mentioned before, which is the earned income tax credit, the EITC. The earned income tax credit, we mentioned this in chapter 17, it's a tax credit that you get that essentially acts as a wage subsidy.

It's a tax credit you get where, over some range, the harder you work, the more you get from the government. It's a negative tax. It's a subsidy.

So let's describe how it works. First of all, the ITC has grown tremendously over time. If you look at figure 21.4, you see that the EITC has gone from being very small to now being about \$70 billion a year in spending. The ITC has grown a lot in time.

How does the ITC work. Well, the best thing to do is look at figure 21.5. On figure 21.5 on the x-axis is how much earned income a family has. It's irrelevant for things like dividends and stuff. It's just about wages, how much earned income you have. On the y-axis, how big your ITC is.

The way it works is, starting from 0, for the first roughly \$15,000 of earnings, for every dollar you earn, you get an extra \$0.40 from the government. So it's a negative tax. Until you reach the maximum amount of \$5,920, then you stop. You just get a check for \$5,920.

But starting again at 19,350, they start to reduce that check by a somewhat different slope, by \$0.21, so that by 47,440, you get nothing from the government. So what this is a government subsidy where the level of the subsidy depends on your earnings.

Essentially, you can look at this chart, say, what's my earnings? And that will tell you how big the subsidy is. That's how the ITC works, is basically a subsidy that depends on what you earn.

Now, remember let's step back. What's the idea of the ITC? The idea of the ITC is we're going to distribute to poor people. We're going to do so in a way to try to fix the leaky bucket by paying them to work. That's the idea. I should have said this first.

The idea is we've got this leaky bucket problem, this iron law of redistribution, where we have the G and the T. And we can't solve all our problems. One way to solve it is by actually, instead of taxing, let's promote work.

But actually does the ITC promote work? It's ambiguous. To see this, let's go to figure 21.6, which is a bit hard to read. I couldn't draw this very clearly. But it's actually really interesting in terms of how it helps think about our intuition.

What the ITC does to work varies and depends what kind of work you are. What this does is it takes a standard labor supply, standard budget constraint-- assume a wage of \$20 an hour. It's a standard consumption leisure diagram with a wage \$20 an hour-- and overlays an ITC on top of it.

So the new budget constraint, which is the blue plus the green, is the new budget constraint with the ITC in place. And we have four different types of people. And we can ask, what does the ITC do to the labor supply?

Well, for person A, it's unambiguous. There is only a substitution effect. Why? Because they have no income, so there can't be an income effect.

I'm just saying, I'll pay if you go to work. If you don't go to work, you don't get anything. So it's unambiguous that it will induce increased labor supply among those not working.

Person B is someone who's working and at \$5,000 is below the first notch. So if you flip back to the previous diagram, they're on the upward sloping part of the diagram. What's the effect on them? Well, actually it's ambiguous.

There's income substitution effects. On the one hand, by providing a subsidy to work, I'm raising the price of leisure, which means I'll take less of it. On the other hand, I'm richer, which means I'll take more leisure. So the effect is ambiguous.

What about person C? Person C earns \$15,000. They're on the flat part, if you go to the previous diagram. How does it affect their labor supply, and why? How does the ITC affect someone like person C, and why?

See the answer to the diagram. So you've just got to explain why. Come on guys, you got this. There's nothing I've been talking about. Yeah?

It's a purely income effect. There's no substitution. If you look at that segment, if you look at that segment, it's parallel to the original budget constraint. It's a little hard to see.

There's no substitution effect. They're just richer. If you make someone richer, they work less, so the income effect. So a person C unambiguously they'll work less. I've just given them a check for \$5,000. They're going to work less.

What about person D? Well, person D, there's now both an income and a substitution effect causing them to work less. The substitution effect is now they're on the downward sloping parts, so they're taxed at \$0.21 per dollar they earn. The income effect is they're still getting a check, so still richer than they would be if they weren't getting the check. So that will also cause them to work less.

So look at this diagram. First of all, questions about what's happening with A, B, C and D? Look at this diagram. It doesn't necessarily seem clear that ITC's encouraging labor supply.

Well, it does for A. But it's ambiguous for B, and it ambiguously deters labor supply for C and D. So it's a very open empirical question about whether the ITC will actually-- how it will change labor supply.

Fortunately, the ITC has been fucked around with a lot over the last few decades, so you have lots quasi experiments to look at to understand it. So for example, if you go to Figure 21.7, this shows, for example, how the ITC changed after Tax Reform Act of 1986. And the Tax Act, this was when it was first introduced. The slope used to be minus 0.11 slope.

On the upwards part, you see 0.11. Now it's 0.4. It went up. The slope went up. They raised the maximum. They change the phase out.

So you could set up a standard difference in difference by looking at groups that were more or less affected by this change. Moreover, the ITC depends on the number of children you have. I wrote this-- let's see. What number of kids did I wrote this for?

Single earner with two children. Well, if you have one child, your ITC is smaller than if you have two children, so that sets up another difference-in-difference approach. So there's lots quasi experiments you can do here, looking at, what happens to the ITC changes for women with one kid versus two kids, women with no kids versus some kids?

Women, no kids get a very tiny ITC. By the way, it's generally women because generally married couples, men and married couples earn higher amounts. Well, actually, let me back up. Strike that last statement.

I refer to women because that is the population that folks are thinking about because the big thing of the ITC was, can it get people off welfare? So the initial focus was, what effect does it have on women? And how does it compare with the disincentive effects of welfare?

So that's where most of the literature is, is looking at single moms compared to, say, single women without kids who used to get no EITC and now get a tiny one or, say, single men. So they look at the effect of the ITC on those groups using changes in ITC generosity.

And what did they find? Well, the effects were striking. Huge increases in working. So people type A responded dramatically, big increase in working. And

No apparent other effects on groups B, C, and D. It didn't seem to deter work among groups B, C, and D. So in other words, the group where it unambiguously increases work, it did, large effects. But you didn't really see the offsetting reduced effects among those elsewhere in the budget constraint.

Now, there's two possible explanations for why. One is differential elasticities. Maybe the elasticity of participation that is going 0 to 1 is very elastic, but the elasticity of how many hours you work isn't because you can't really control your hours. You work for the guy. He tells you how hard to work.

So one explanation is just literally the one thing you clearly have a choice over is whether to work or not. How hard you work you don't have as much control. That's one explanation.

The other explanation is that it's really pretty fucking hard to figure out these tax rates, to figure out this system. If I say to you, hey, if you go to work, you get a check. If you don't, you don't. That's easy.

If I say, by the way, the marginal hour you will have [INAUDIBLE] marginal tax rate of minus \$0.21 if you work on this kink and not on this kink, it's like, I don't understand how that works. It's pretty hard to understand.

So the other explanation for what's going on is that basically people recognize the incentive effect of getting a check if they go to work now but don't recognize the marginal tax rate. And indeed, there's a lot of evidence that people aren't responsive to marginal tax rates, at least in much of the range. And remember, as I talked before, I talk about people misunderstand the tax system and this may be an example of that.

So not sure why it is. But the bottom line is that essentially the EITC seems to be, on net, create a large increase in labor supply. This is fascinating because this is unambiguously good news. It is a patch to Okun's bucket because we're clearly redistributing to low income people.

No one with income more than \$50,000 gets the ITC. That's pretty low income. And we're doing so in a way that raises, not lowering labor supply.

So this is wonderful news, which is that basically we've discovered a tool for fighting poverty that actually does so without causing reduction in the size of social pie. Nonetheless, the ITC is still controversial, and that's for two reasons.

One reason is, as I said, it can cause people to have negative tax burdens. And some people, as I said, irrationally feel the tax system should not be paying people money. Tax burden should never be negative. As a result, the ITC's a bad idea.

The other reason it's controversial is because there's a decent amount of cheating. Basically, like with any part of the tax system, there are people who will lie about how much they make if it's not really recorded. It's lying to say you made more than what you made before. It's an odd kind of lying.

But once again, people will do that. They might misreport other features of their taxes, and they might say kids live with them who don't to try to claim the credit, et cetera. Now, we're going to come back to this in chapter 25.

The amount of cheating on the ITC is like a tiny fraction of all cheating done on taxes. Fundamentally, only the rich guys really do the big amount of cheating on taxes. But nonetheless, so I view these as thinly veiled excuses for people who really don't like redistributing to the poor, to be honest, but they are excuses that play in Washington, and that's one reason the ITC is probably even bigger than it is.

There is a third weakness, which is important, and this comes back to chapter 17. You might say, well, gee, this is awesome. Let's get rid of welfare and replace-- let's forget these other problems. Let's imagine a rational system that honest debates.

Let's say, let's get rid of welfare and just replace the ITC. What would the problem would that be? Just say, let's get rid of this whole labor distorting welfare, just having a big EITC system. What's the problem with that? Yeah?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Some people legitimately can't work. You still need-- there's a distortion from subsidizing non-work, but there's also a consumption-smoothing benefit. So an optimal system would never feature nothing for those who can't work, as long as there is some exogenous reason people can't work.

You never want a system that leaves people that work with nothing. So the optimal has to be some combination of money to those who don't work but then incentives to work for those who can. What that combination looks like is complicated. All right.

So that's the EITC. And it's a great example of how we think about applying our logic of income substitution effects to making policy. Yeah?

AUDIENCE: How do people study things like cheating that go, by definition, unrecorded?

PROFESSOR: Awesome. Great question. I'll talk about this more in chapter 25. But basically IRS does audits, and with the audits, they can actually compare people report to what you see in the audit data.

AUDIENCE: Based on that sampling?

PROFESSOR: Yeah. That's your random sample. They literally go and make your life miserable and collect all the data, and they compare it to what you actually declared.

So that's one application. The other application I want to talk about is child care. We talked a little about child care in chapter 17. I want to return to that.

Remember, that basically, today, fewer than of preschool aged children are cared for by a parent. 12.5% are cared for by their relatives, and 62% are in some form of child care. So almost 2/3 of kids today are preschool age are in child care. Child care is a very big deal, and it makes it possible for women to actually finally realize their potential and join the labor force in a way they couldn't in decades past.

The interesting question is, how should we treat childcare for tax purposes? And here we're going to go back to chapter 18 and thinking about Haig-Simons, think about Haig-Simons principles

So basically if you think about it, when you deliver-- when my wife stayed home and took care of our kids, she was delivering something which expanded our resources. We were better off because she stayed home and took care of my kids.

But if she'd gone to work, she would have been taxed. But the work she did staying home was not taxed. And that's a violation of Haig-Simons.

Haig-Simons says you should tax everybody's potential to pay. Well, think about it. Let's say my wife stays home, and your wife goes to work. And they're both delivering the same value of goods.

It's just my wife's still in the home, and yours is in the market. You get taxed. I don't. That's unfair.

That violates the Haig-Simons principle. We should be taxed on our underlying resources, which are the same, yet we pay very different taxes. It's a violation of some concept of horizontal equity.

So what do you do about it? Well, basically there's three approaches. So let's go to an example. Let's go to table 21.1 Go to table 21.1

So imagine you have a couple. Imagine you have a couple where the man's going to go to work, and the woman is considering whether to go to work. The wife is considering whether to go to work. And imagine that, if she went to work, she'd earn \$1,000, but she would pay \$600 in child care costs.

And let's assume markets are perfect. So \$600 in child care costs is literally the inherent social value of caring for the kids. And the \$1,000 is the inherent social value of her work. So what that means is she should work, full stop.

The society for working is \$1,000. The society of her kids being cared for-- ignoring externalities of different people caring for kids, put all that away. Simple model, where basically the market works well [INAUDIBLE] caring for those kids is \$600. So the size social pie is higher if she works.

However, if she goes to work in the 1 to 3 or 5th column, if she goes to work at a 50% tax rate-- because remember, her husband's already earning money, so the marginal tax rate on her, if you've got a high earning spouse, your next dollar is already in a top tax bracket. You don't get to benefit.

Remember, we talked about marriage tax. You don't get to benefit from that low tax rate initially. You're already in the high tax bracket when you go to work. So the marginal dollar, let's say, is taxed at 50%.

What that means is, if she goes to work and earns \$1,000, she'll go \$500 in taxes. So her the after tax value of her work will be \$500. But if she stays home, she'll effectively make \$600. So she won't go to work.

So in other words, we have distorted the market. We have not made a social welfare-improving trade. The social welfare-improving trade is for her to work, but she doesn't. This is the distortion that arises from not taxing the true Haig-Simons base.

We call this a tax wedge. A tax wedge exists whenever taxes impact the incentives to engage in two, if you will, economically similar activities or activities where you basically would like to compare them-- basically, if you have a comparison where there's a relative advantage to one-- let me think about this-- actually, I think I have the definition in the book. Let me get this right.

Basically, tax wedges are basically the difference between the returns to an activity in different sectors. That's the way to think about it. You can deliver labor supply in the market sector or the home sector. If taxes change the relative returns to those activities, that creates a tax wedge. And here that's happened.

Now, what can we do about it? Well, there's two answers. One is what Haig-Simons would have us do, which is say, look, let's measure her full income. Her full income is she makes \$1,000 if she works. She makes \$600 if she stays home.

Why does she make \$600 if she stays home? Because she's delivering \$600 worth of child care. So in the Haig-Simons world, we would impute earnings to her if she stayed home. We'd say, if you go to work, you're wearing 1,000. If you stay home, you earn 600.

And in either case, you'd then pay taxes. You'd pay 500 taxes if you worked, 300 tax if you stayed home. So the after-tax value of work would be 500. The after-tax value of home would be 300, and she'd go to work.

She'd get the right answer because you've gotten rid of the tax wedge. You've gotten rid of the differential effect of taxes on that labor supply delivered in two different ways. It used to be you're taxed differently if you delivered it in market at home. Now you're taxed the same.

That would be one way to do it. The other way to do it would be to allow her to deduct the cost of child care from her income. So she'd earn \$1,000. She'd spend 600 in child care. But she would then get to deduct that 600 so her net taxable income would only be 400.

Therefore she would owe 200 in taxes, and she would make more money at work, \$800, than staying home, which is \$600. So she'd, once again, go to work. So two different approaches that end the distortion caused by this tax wedge.

Here's where public finance gets really fun. How do I think about which one is better? First of all, let's forget feasibility. Let's forget feasibility for a second. Obviously, imputing income would be a nightmare politically, et cetera. Let's forget that.

But how do I think about, between these two approaches, which one is better, and why? In a world where I could do either, there's an obvious answer. Which one would I prefer, and why? Yeah? You got to speak up.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Because of what?

AUDIENCE: [INAUDIBLE]

PROFESSOR: No, I'm saying imagine where we could do. Both I'm getting rid of the administrative problem. Just on pure economics, which would you choose, and why? Appealing to the lessons from last lecture, from chapter 20, which would you choose, and why? Yeah?

AUDIENCE: It seems like [INAUDIBLE] is overall cheaper to do from the government's perspective and does the same--

PROFESSOR: Well, remember, what's the key source of distortion? It's two. It's elasticities and tau squared. The smaller the tax base, the higher the tax rate has to be. So by shrinking the tax base in the second approach to raise a given amount of revenue, we're going to have to tax more highly, so it's less efficient.

So it's unambiguous, from the tau squared perspective, you'd like the first approach because you get the same-- you remove the wedge, but you collect more revenues. And therefore total tax rates can be lower. People understand that?

All else equal, we want taxes with broad bases and lower rates. That's how we maximize efficiency. Yeah?

AUDIENCE: Why is it for [INAUDIBLE], the value [INAUDIBLE]?

PROFESSOR: Why is it not what?

AUDIENCE: Why is it not 900 [INAUDIBLE]?

PROFESSOR: Why is it not 900? No, you're delivering \$600 worth of child care.

AUDIENCE: But with the earnings on top, or is that--

PROFESSOR: No, no. If you stay home, you're delivering 600 worth of child care. What's the extra 300? I don't get it.

AUDIENCE: Because if you're getting earnings from the government in that case but also--

PROFESSOR: You're not getting earnings from the government. You mean the imputed case in the second row? You're not getting earnings from the government. Yeah, you're just getting-- it's you're earning \$600. They're going to tax \$300 away from you.

Now, let's make it more interesting. Let's get rid of the second row. We're never going to do the second row. First row versus third row, Tell me actually why this might-- now here's a really hard one. But once again, it all comes out of the formula for deadweight loss.

Tell me why the first versus the third row is actually ambiguous which is a better one to do? It's actually ambiguous, and the answer comes directly from the formula for deadweight loss. Where's the ambiguity come from?

We know the third row raises less money, so we know which way the tau squared effect goes. But what's the other effect that drives deadweight loss? What's the other thing that drives deadweight loss besides the tau squared term? Think about your formula.

Elasticity. So why might the elasticity go the other way? Well, think about it. If you look at the last-- if you look at the first row, the relevant elasticity is the elasticity of women substituting between work and home.

If you look at the last row, the relevant of elasticity, elasticity of the whole tax base, we're basically collecting less tax revenues. So let's simplify the world. Imagine the world consists of men working and women working. That's it.

Well, in the first row, the thing you're distorting is a woman's decision to work. In the second row, you have to collect more money from both, so you're distorting both their decisions to work. Imagine, as is still true, less true than it used to be but is still true, women's labor supply is more elastic than men's labor supply, then the first row implies a more elastic decision.

The decision for a woman to work or not is more elastic, the decision for all society to work or not because men have a lower elasticity than women. So actually the first versus third row is ambiguous. On the one hand, you raise less revenue than the third row. On the other, the first row, you distort a more elastic margin, which is women's decision to work or not work, and which dominates is unclear.

Pretty mind blowing, I know. Take it home and think about it. Described clearly in the book. But this is how I want to think about this trade off.

What do we do in the US? What do we do in the US is we compromise. We basically say that you get a tax credit for childcare, up to \$3,000 for one child and \$6,000 for two or more. And you get it as a percent of your income.

So you can't like claim \$6,000 and you only make \$5,000. They say it's a percent of income. You can get a credit.

They basically say, you can deduct a percent of your income up to the maximum of \$6,000 for two kids, \$3,000 for one kid. So we have a limited deduction is how we resolve it in reality. So that's how we think about-- how we operationalize the abstract concepts of chapter 20, which is thinking practice about thinking about what's more elastic.

You want to avoid distorting things that are more elastic, like women's labor supply. But you also want to try to have as big a tax base as is possible, and trading those two off is how we think about optimal tax design. Yeah?

AUDIENCE: How elastic-- are there any policies to which people's decision to have kids or not is elastic?

PROFESSOR: So the decision to have kids is just like-- I think you asked the same question about marriage. Decision to have kids is totally inelastic to everything. The timing is elastic. I told you my story of my kid being born in January 2, my sad tale.

But actually, this is a big issue because fertility is falling around the world. A number of countries are really worried about falling fertility. And there's been nothing they figure out how to do.

Now, there's been tools which lower fertility. China's One Child Policy did lower fertility. But you can't really have-- - mandate people have kids. You can mandate they can't have kids. It's hard to mandate they do have kids.

And other than that, there's really been-- there's no evidence of any financial tool that basically promotes people to have kids. It really is a very inelastic margin with respect to incentives.

AUDIENCE: Which, for a lot of economies, is a goal that they would like to achieve?

PROFESSOR: Yeah, they'd like to achieve and they get big child bonuses and things like that, and none of it works. Now, once again, if we think about a tax break, we have the marginal for marginal parts. Marginally, it's a bad idea to give a child tax break because it does induce childbearing.

But inframarginally, it may be you're attributing to exactly the people we want to have more money in society. So I talked about, from an efficiency perspective, the inframarginal piece is a waste. But from an equity perspective, maybe we just equitably want to reward people who have kids, just out of some notion of what's socially just. But it's not inducing more children.

That's chapter 21. Now we're going to move on to chapter 22, which is a parallel analysis. We go through the basic theory, and then we'll talk about applications.

The basic theory for the tax on-- now, instead of looking at the effect of taxes on labor supply, we're going to let the effect of taxes on savings. This is a little more complicated, but the logic is similar. The idea is when you save-- we're going to start in a simple world where the only way you can save is put money in the bank, simple world, money in the bank, and earn some interest rate R [INAUDIBLE] I'm sorry, I. What do you have here? R or I ? R .

We're going to start with a simple world of two periods. There's a working period and a retirement period. This is the classic overlapping generations model we talked about. We talked about Social Security.

There's a working period and retirement period. And people are going to optimize their consumption bundle across time. We call this a model of intertemporal substitution. You're going to substitute your consumption across different periods, intertemporal substitution.

In this simple case, there's two periods. Obviously, in reality, there's infinite periods. There's two periods. There's working life. That's the x-axis. And there's retirement. That's the y-axis.

So on the x-axis, consumption while working, C^W , on the y-axis consumption while retired, C^R . And your question is simply, how am I going to allocate my consumption across these two?

Furthermore, we're going to assume you only earn income in period W. So you're going to need some income Y that you've got to allocate across the period when you work and the period when you don't work. We're also going to say that you can choose to save some of your money.

When you work, you can choose to save. So savings is Y minus C^W , $Y - C^W$, whatever you earn minus what you consume while working is savings. The rate of return to savings is the interest rate R .

So what we have here is the blue line in figure 22.1. You can either consume Y while you're working. That's the x-intercept of the blue line. Or you could save it all, consume nothing while you're working, and have Y times $1 + R$ when you're retired.

So either consume it all when you're working [INAUDIBLE] while you're working, and nothing while you're retired or save it all, in which case you get Y times $1 + R$ when you're retired and nothing when you're working or some combination in between. As a result, the slope of the budget constraint is minus $1 + R$.

In other words, the opportunity cost of consuming in period 1-- by the way, I'm going to refer to period 1 in working-- and I'm referring interchangeably. When I say period 1 or 2, I mean working and retirement I just slip back and forth.

The bottom line is there's a two-period model. We've got before and after. The opportunity cost of consuming while you're working is that, for every dollar you consume while you're working, you're giving up $1 + R$ dollars consumption when you're retired because you could have put that away and saved and earned interest.

This is the key concept here with savings. It's a little easier with labor. Either you sit at the couch, or you go to work and earn W . It's easy to see W . It's the opportunity cost.

But it's the same logic. The opportunity cost, if I don't eat it today, I put it in the bank, and it becomes $1 + R$ dollars tomorrow. Or when I'm working [INAUDIBLE] dollars when I'm retired. So that's the opportunity cost of consumption today is the savings I could have had when I was retired.

Now, imagine the government comes along and taxes interest. This is what we call capital taxation. It taxes interest.

What does that do? That lowers the effective rate of return to saving.

Is your sweatshirt on backwards? Cool.

So that lowers the effective rate of return to savings. So basically what that does, it moves from the blue line to the red line. You now see the opportunity cost is no longer $1 + R$. It's now $1 + R$ times $1 - \tau$.

It's a little confusing. Think of it this way. Imagine the tax rate was 100%. There'd be no slope to this line. It'd be a flat line.

So basically it's reducing the rate of return by how much that is taxed. So the slope pivots downwards. Once again, if you consume everything while you're working, there's no change. You still have $C^W = Y$.

But now if you wait until-- if you consume nothing while you work and consume it all when you retire, you end up with less when you retire with $Y \text{ times } 1 \text{ plus } R \text{ times } 1 \text{ minus } \tau$. Questions about that? Yeah?

AUDIENCE: [INAUDIBLE] a tax on savings?

PROFESSOR: Tax on interest.

AUDIENCE: [INAUDIBLE]

PROFESSOR: It's not on your stock of wealth. I'll come back to that next. It's a tax on the-- it's what's called a capital income tax, a tax on the flow of income you earn from your capital. Very important distinction.

I call them savings taxes. But technically we're talking about capital income taxes, taxes on the flow of income from your capital holdings. Other questions about how this works? Yeah?

AUDIENCE: That would be experienced as when you sell this, the difference?

PROFESSOR: No, not when you sell it. This is a bank account. Every year you earn interest from the bank, and it's taxed by the government. Just like labor income, this is capital income.

We'll come next time to assets you sell. That's capital gains. We'll come back to that next time. Other questions?

This is hard, so don't be afraid to ask questions. This isn't clear. Now, what happens to your savings decisions when we tax interest? We know the answer is going to be ambiguous. I've given you a hint, so let's go and explore it in figure 22-1.

And this is basically the same as for labor. It's just more complicated. Here is the key insight. Remember this. We are not modeling savings. We're modeling consumption in period 1.

Remember, savings is a bad. Savings is something you have to do to make sure you can survive when you're retired. The good is consumption while working. We model this. This is fixed. And we compute S as a residual.

So just like we don't model labor, we monitor leisure, here we don't model savings. We model consumption while working. So in the first case, we have a substitution effect and income effect.

The substitution effect is the price of first period consumption has fallen. What do I mean by that? The opportunity cost of eating now has fallen because you don't get as much later if you don't eat. Once you think of τ as 100%, why not eat it all now?

The opportunity cost of eating now has fallen. So as a result, you consume more today. That is a substitution effect. Yeah?

AUDIENCE: [INAUDIBLE] that means you gained no interest but you still have--

PROFESSOR: You'd still want to split your income because you don't want to starve.

AUDIENCE: So it's not like you have nothing in the future.

PROFESSOR: No. No if τ is-- it's not negative 100. Yeah, you've gained nothing by saving. So it's not like you're necessarily going to consume it all today. It's just the incentive to wait has lowered.

Good point. I wasn't clear on that. Thank you.

So now on the other hand, you have an income effect. You are now poorer because every dollar you save yields less lifetime income. Since you're poorer, you want less of everything, including first period consumption.

So that says you consume less in the first period. In the first case, the substitution effect dominates. In the second case, the income effect dominates. Once again, not implausible, in fact. Here I find it even more plausible than labor supply case.

Why do we save? We save because we want to buy something. We have some goal. I want a car. I want to have this much when I retire.

Well, in that model, income effects would dominate. In a model where I'm saving to meet a goal, then if the interest rate gets taxed, I'll have to save more. So it's not at all implausible income effects could dominate here.

In fact, from my intuition, it might be more likely, which is, in most parts of the world, we think substitution effects dominate here. I think that presumption is more difficult, although we'll still assume it because we always do. But it's a little more challenging.

But now the diagrams are complicated. Let's go back to the diagrams. They're complicated because let's look at literally what happens.

In the first case, you were consuming C^1 and saving S^1 . Now you're consuming C^2 and saving S^2 . Now, we're not modeling period 2 consumption period-- period R consumption. But it has implications. Look at what happens period R consumption.

It falls from $S^1 + R$ to $S^2 + R$. That is C^2 -- I'm sorry, C^2 falls for two reasons. First of all, it falls because you're saving less. Second of all, it falls because even the savings you have earns less interest.

So that's why, look at the second panel. In the second panel, you save more because income effects dominate. You go from saving S^1 to S^2 .

But look what happens to retirement consumption. It still falls because there's two effects. You save more, but it earns a lower interest rate.

This is why model CW -- and let the rest come out of the equations. Don't be fucking around with C^2 because you're going to mess yourself up because here we have a case where, in one case, savings falls. In other case, it goes up. But in both cases, second period consumption falls because there's two effects.

There's the effect of saving, and there's effect that literally the interest rate's gone down. So you get less for every dollar you save. So once again, make your life easy. This is the choice variable we want to model. Everything else just falls out of the budget constraint. Yeah?

AUDIENCE: [INAUDIBLE] effect.

PROFESSOR: So substitution effect is the interest rate, the after-tax interest, rate R times 1 minus τ is the price of working period consumption. Why is it the price of working consumption? Because the higher it is, the more costly it is to consume today.

So when we raise taxes, that is a reduction in price. When you reduce the price of a good, the substitution effects consume more of it. So that would say more period 1 consumption. The income effect if it's-- so it says more of it.

The income effect is, unlike with goods, which I'm buying, this is a factor I'm earning. Now I'm effectively poorer. Why? Because if R minus 1 minus T falls, I have less money in my lifetime budget constraint. We'll do the math in section.

My lifetime budget constraint is basically is essentially going-- my lifetime total consumption is going to turn by how much I save. The more I save, the more my lifetime total consumption is. And so my eventually my resources fall. Yeah?

AUDIENCE: When you're mentioning how people respond to this based on the [? income ?] distribution effects, how do we [INAUDIBLE] so I'm imagining a person is facing this decision, and the ways that they can adjust their income are-- or the ways that they can adjust are maybe extending the amount of time for which they are working--

PROFESSOR: No, no, but I'm taking Y fixed here. This is just a savings decision. Y is fixed. It's not about income. That was the last chapter. This is just about savings, just about how much you consume with income fixed.

Now, you could imagine a world, say, on an exam, where we combine these two insights and allow both income and savings to change. But for now, we're holding it fixed. The last chapter is what drives this. This chapter is about what drives this.

Now, that's one reason this is more complicated than labor supply, why I teach it second, is that graphs are a lot more complicated. It's also a lot more complicated because the intertemporal choice model is only one of a number of models that could plausibly describe savings behavior. Indeed, many people think it might be the third most important determinant of savings, that things like how you respond to the interest rate might be the third most important, because there's two other very valuable model of savings.

So there's three models. There's the intertemporal substitution model, intertemp substitution model. That's the classic model we're working with based on interest rates.

There's also a precautionary savings model. If you ask people why they save, the number one answer they give is they're worried about stuff going wrong, not the [? respond ?] to the interest rate [? and worry ?] about allocating the consumption over working retirement.

The main reason is they want to have money in case something goes wrong. That's not in this model. That's what's called the precautionary savings model.

The precautionary savings model would be basically-- the precautionary savings model would be one where basically-- think of it this way. Here's how I think about-- we're not going to make you write down the math, this, but here's how to operationalize it.

Imagine a world where people face liquidity constraints. They can't borrow. If you could borrow, there's no need for precautionary savings. If you can costlessly borrow, then you wouldn't need to-- this is only a world with liquidity constraints. Otherwise, there's no need for precautionary savings.

But imagine you can't borrow. Then you're going to want to make sure you have some savings in case stuff goes wrong. And so in a world with liquidity constraints, there will naturally be precautionary savings as a form of self-insurance.

So that's an alternative model. And indeed, there's lots of evidence for this. First of all, there's lots of evidence that people who are in riskier jobs save more. Now, that's not super convincing because maybe they're different.

Better evidence is there's evidence that when the generosity of social insurance goes up, people save less. Remember, we talked about this with Social Security. We see in other contexts too. For instance, in Taiwan, they put in national health insurance for some groups and not others, and the ones that got national health insurers suddenly started saving a lot less.

Now, that's because they were saving against getting sick. I found a similar result here in America. When Medicaid expanded, people saved less and consume more because they didn't have to save as much against getting sick.

So clearly the precautionary model has some validity. We'll come back next time to how this changes our conclusions. But that's one alternative model.

The other alternative model that's very relevant is the self-control model, the self-control model, this model which says, it's not really an affirmative model of savings. It's kind of a Homer Simpson model of the world, where basically we just want to consume. We want to eat what's in front of us.

And the only thing that allows us to save is imposing self-control to stop ourselves and that savings is really whatever the motivation. The driver of savings is how much self control do I have to make myself save for either precaution or retirement.

Now, this is very much related to the behavioral economics insights we talked about in chapter 6 when we talked about smoking, which is the notion of self-control problems. And there's lots of evidence that, in savings, people face self-control problems.

The classic evidence? Existence of what we call Christmas clubs. What a Christmas club was the thing banks used to offer, where you could set aside money every week in a noninterest-bearing account or a tiny interest-bearing account to have money at Christmas time.

Well, why the hell do people need to do that? They could just save it somewhere else, that money, Christmas time? Because they're afraid they wouldn't be able to.

More modern evidence comes to the following fascinating confluence of facts, which is that most individuals have virtually nothing in the bank. Or many individuals have virtually nothing in the bank but lots of money in illiquid forms, like housing and pensions, forms they can't get at.

Moreover, it goes further than this. The modal individual in America or many-- I don't know if it's the modal, but almost the majority of people in America have the following [? fee ?] combination of assets-- little money in the bank, money in their pensions and housing, and huge credit card debt on which they're paying 18% interest. Think about that for a second. That makes no sense.

They've got money in the bank, first of all. They should at least spend the money in the bank. Moreover, they've got money in a pension, which is earning 6%, and they have credit card debt at 18%. Why are they doing that? Because it's what we call a commitment device, like we talked about with cigarettes.

They know that they can't control themselves from hitting their credit card maximum. That's just a state of being. So to set a commitment device, they put their money where they can't get it, in illiquid assets, even though they're sacrificing a massive rate of return differential. Essentially the price they're willing to pay for commitment is to hide their money away so they can't get it because they know, if they took it out, yeah, they pay their credit card bill off today. But then they go back up tomorrow, and they would have just blown through their money.

They're still able to keep some money in the bank because they have-- they're able to keep some money in the bank because basically because of another model, self-control models, we like to call self-control/general behavioral models. Another important behavioral model that can explain this is a model of what we call mental accounting.

The mental accounting model is fascinating. This says that people don't think about dollars fungibly, that dollars in different places mean different things. There's a famous story I told in 14.01 about the actors Dustin Hoffman and Gene Hackman. Famous old actors were roommates together when they're acting students and poor.

And Gene Hackman knows Dustin Hoffman and all these bottles labeled like rent, money, et cetera-- rent, food, et cetera. And the food bottle was empty. So he went to Gene Hackman and said, can I borrow some money?

And Gene Hackman said, yeah, but your rent bottle is totally full, and your movie bottle is full. He says, well, I can't touch those. Those are for renting movies.

And that's a way people think about it. And that's how people are able to leave some money in the bank because at least they're like, I won't touch that for my credit card. But once it gets too big, they'll dip into it.

And then maybe the best piece of evidence I find for these models is a fascinating experiment or product introduced by the Nobel Prize-winning economist Richard Thaler. He went to workers and said, I want to offer you the foreign savings plan. It's called the save more tomorrow plan. Here's how it works.

Next time you get a raise, you will commit to saving 50% of that raise in a separate account. It doesn't cost you anything today. You don't have to give up your consumption today. It's just money you're going to get in the future anyway. Just decide to set half of it apart.

Now, in a rational, non self-controlled model, I would say no. Why should I commit myself? What good does [INAUDIBLE] I want to save in the future? I'll save in the future.

But the model where I know I have trouble saving money, I might find that attractive because it commits me at no cost today to having savings in the future. Well, he offered that plan. 80% of employees signed up for it, and their savings rate went from 4% to 14%, which is exactly consistent with the fact that people had self-control problems and they couldn't save.

So clearly self-control is a major determinant of savings. And once again, we'll come back to what that means next time for policy.

Now, like last time, I want to have two applications for thinking about how tax affects savings. Like in chapter 21, we have two applications. I'll cover one today and one next time.

So basically step back. Just like the other slide, we have a basic theory, a basic traditional theory. Before, it was the model of allocation across consumption, leisure. It's the model of substitution.

Unlike with labor supply, we have other very plausible models that may, in fact, be more important, those other models. And that's the theory setup. Now let's go to two empirical applications.

The first empirical application is talking about inflation. Now, in the chart I drew, I used R , the real interest rate. But in fact, banks don't pay a real interest rate. They pay a nominal interest rate.

And we typically assume that the nominal interest rate is the real interest rate plus expected inflation. So basically if inflation is stable at 3%, then nominal interest rate will just be real interest rate plus 3%. If we expect inflation to go up, the nominal rate will go up. So that's our model of interest rate determination.

It turns out that, in inflationary environments, capital taxes have a particularly bad feature. They have a particularly bad feature. And that's because, when we tax interest, we tax nominal interest, not real interest. And this creates an additional wedge between consumption and savings.

To see that, let's go to table 22.1. Imagine we're all Marshawn Lynch and all we eat is Skittles. Nothing? Nobody? No Marshawn Lynch fans here? OK.

So basically, I want to consider three worlds. In the first world, there's no inflation, And there's no taxation of interest earnings. I'm going to save \$100. I want to save \$100. And the interest rate is 10%

So in this world, in the first row, I saved \$100 at 10%. I get \$10. After tax of \$110, skittles are \$1, so I buy 110 bags of Skittles. People understand the setup?

Now imagine the government comes in and says, no, I'm going to tax your interest at 50%. Well, now that \$10 of savings only becomes \$5 after tax return. That \$10 of interest only becomes \$5. So now I have \$105 after tax, and I only buy 105 bags of Skittles.

So what taxation has done is cause me to-- has distorted my consumption so that I can only consume 105 out of 110 bags of Skittles if I save, which may make me want to cause to save less. So it's distorted my consumption. That could have income and a substitution effect. We know which way it goes. But that illustrates the distortion of capital taxation.

Now, imagine a world where inflation comes along, and the interest rate doesn't adjust. Imagine there's now 10% inflation, and interest rate doesn't adjust. It stays at 10%. This is an implausible world, but it's a useful middle step.

Well, now I saved my 100. I earn \$10. In row 3, I keep the 110. But there's 10% inflation. Skittles have gone to \$1.10. So you only get 100 bags of Skittles.

And with interest taxation, that 10, I only get \$105, so only get 95.5 bags of Skittles. So basically in the first row, I end up with fewer bags of Skittles because inflation basically offset my interest. In the second row, I ended up with even fewer bags of Skittles. People clear on that?

Now let's go to the realistic case. The realistic case is one where the real interest rate adjusts so that now the real interest rate is 20%. There's a 10% inflation. Interest rate's 20%.

So look at the next to last row. Inflation is 10%. I'm sorry, it's 21% because it's multiplicative.

I save \$100. I end up with a nominal rate of 21% because it's not really plus pi. It's times 1 plus pi. So that's why it's 21%. But we can always round and just say plus.

So I end up with \$121. The price of Skittles is \$1.10. So guess what? I'm the same as the first row, which is what should happen.

Absent taxation, inflation should not affect how many Skittles I can buy. Why? Because the interest rate goes up to offset, goes up to offset. The price change, I'm indifferent.

But now look what happens when there's taxation. There's a 50% tax rate on interest. My savings is 100. I earn \$21, but half of that is taxed away.

So after tax, I have 110.5. Skittles are \$1.10, so I can only afford 100 bags of Skittles, plus half a bag. So look at the difference between rows 1 and 2 and rows 5 and 6.

In rows 1 and 2 without inflation, capital taxation has reduced the number of Skittles I can eat by 5%. In rows 5 and 6 with inflation, capital taxation has reduced the number of Skittles I can eat by almost 10%. What's going on?

What's going on is there's another wedge. And the wedge here is that we tax nominal interest rates, not real interest rates. So when inflation comes along-- and think of it as an extreme case. If there's massive inflation, then even if real interest rates adjust, I'm still not going to save because basically I'm going to-- when interest rates go up, I'm going to lose a lot of that through taxation, even though, in fact, it should be like row 5.

I should be indifferent to inflation. But I'm not because it's a wedge created by the fact that we only tax nominal interest, not real interest. Now, this point was made in the 1970s by the famous economist Martin Feldstein, really, if you will, the father of empirical public finance.

This was a very important point. And then we forgot about it starting about the mid-1980s. Why? Because inflation went away. So every year I teach this and be like, yeah, whatever. There's no inflation.

Well, guess what? There was inflation. And now you understand why something like this matters. Now, fortunately, we're back to low inflation.

But you could see we were in a world with almost 10% inflation. And you could see in that world the fact that capital taxes are more distortionary in a high inflation environment because we tax nominal interest, not real interest.

Now, in fact, we could fix this by we could just simply index capital taxation. Indeed, we've done this otherwise. When I was a kid, back in the day, in the 1970s, the whole tax system used to have this problem, which is that basically the tax brackets were fixed, but your income was not.

So let's say inflation went up, and my boss gave me a raise. That could move me to a higher tax bracket. And I could pay more taxes, even though I didn't have any more resources. So imagine where prices went up 20%, as they were going up in the 1970s, or 15%.

I get a 15% wage increase. I can't buy any more goods. But my taxes have gone up because the tax bracket was fixed, so I'm bumped into higher tax bracket. That was called bracket creep.

We fixed that in 1979. We indexed the tax brackets. So the tax brackets themselves go up with inflation.

We could do the same thing with capital taxation, which is we could index what you pay to what happens to inflation. We don't. Part of the reason we don't is that, who would that help? The wealthiest, the ones who save the most.

So that'd be no. We do lots of things that help the wealthy, and maybe this Congress will do that. But traditionally, that's been viewed as a wonky fix that plays bad politically. So it's not really been done.

But in principle, you could solve this by indexing taxation to account for inflation, just like we did with brackets. Questions about that? All right. If not, let's stop here.