

[SQUEAKING]

[RUSTLING]

[CLICKING]

JON GRUBER: So today, we're going to continue our discussion of chapter 8. Remember, last time, we went through the costs of we have a highway project. We talked about the costs in terms of labor, in terms of asphalt, and in terms of maintenance. We talked about the issues involved in measuring those costs and came up with the total cost.

Then we turned to the benefits side. And there's two types of benefits. There's valuing time savings and valuing lives saved. For each of these benefits, there's fundamentally three methods that are used to value the benefits, three methods one can use.

The first method is what we call the market value method. And that's basically using a market-revealed price to measure the value of something. So for time, the market-revealed price is wages.

Remember, I talked about the fact that in a standard, no friction, standard economic model of the labor market, the wage is the appropriate measure of people's time. So in that world, we can use the wage to measure what time gained is worth. We then talked about why that may or may not be a good measure. So wage has pros and cons. Remember, it might be too big or too little, depending on various imperfections in the labor market. So in a perfect labor market with no imperfections, the wage would be the appropriate measure of time. In a labor market with imperfections, that may not be.

The second approach we use is an approach we call contingent valuation. Now, this is-- once again, I've mentioned this before, but part of what you get from this course is a few trade secrets. And one of the trade secrets in economics is we make a lot of money by making up fancy terms for things which are really not very hard.

Contingent valuation means asking people, asking people what stuff is worth. And that is basically saying, look, what's it worth to you to save an hour, and doing a survey and ask people what it's worth to save an hour. That seems like the standard place one would start. If I want to know an hour of time is worth, you'd ask people.

The problem is people have no idea, that it's a very hard thing to know how to answer. What's an hour of my time worth? And it gets even more complicated when you ask them harder questions, like, how do we value various other potential benefits of regulation or damage for environmental considerations?

There's a great article which talks about all the problems you have when you ask people their preferences. So for example, one problem is, whether a question is asked isolated or with other questions matter. So if you ask people how much they'd be willing to pay to improve the visibility of the Grand Canyon, if you ask them solely-- if you only ask that question and then ask them as part of a series of questions or randomly asked some people only the question-- other people ask five questions. That's the third. The second group values it one fifth as much as the first group, even though it's random. So when I'm asking the list, it makes the Grand Canyon one fifth as valuable, which is not plausible.

There's also the order of issues matter. So if you ask people if they'd rather save whales or seals, if you ask about seals and then whales, you find that seals are worth about the same as whales, about 142 to save the seals and 195 to save the whales. So whales are worth about 25% more than seals. If you ask about the whales first, now seals are worth twice as much as-- now whales are worth twice as much as seals-- 172 to \$85. So the order you ask matters.

And then finally, there's something called the embedding effect, which is people have a hard time thinking about magnitudes. So if you ask people randomly, what is it worth to save 2,000 birds, 20,000 birds, 200,000 birds, or 2 million birds, they give the same answer. It's because they don't know how to-- and it's not surprising. It's hard to think about what these things are worth.

As a result, it's hard to use contingent valuation methods to value things. And that's why economists typically prefer a third approach, which we call the revealed preference approach, revealed preference, which is basically saying, people might not know how to answer a question, but their actions betray their true valuation. The fundamental rule of economics-- you want to know something about people, look at what they do, not what they say.

And revealed preference methods are ways of using people's actions to value how much they value things. So for example, someone tell me. Raise your hand and tell me. If what you think about valuing how much an hour of your time is worth, how could I use a revealed behavior to tell me how much an hour of your time is worth? What could I do? I want to know what an hour of your time is worth. I see you do lots of things in life. I track you all day. I see everything you do. How could I use your decisions every day to tell me what an hour of your time is worth? Yeah?

AUDIENCE: Without an additional introduction, you could see whether they use something like a grocery delivery app or a--

JON GRUBER: Great. That's a great idea. If you use a grocery delivery app, you pay \$5 more, and you save X minutes. So you could look at, how much more are you willing to pay to save X minutes? You could look at two different ways people can drive home, one which might have a toll and one way which wouldn't have a toll but takes longer. You can actually set ways to avoid tolls. You could look at people who set the Waze in different ways-- Waze, different spellings-- and see how much they'll pay to avoid tolls.

So you can look at people's decisions to try to reveal it. A classic way of doing this is looking at the most important single financial decision every American makes, which is to buy a house. Houses are worth different amounts depending on the commute. The value of the commute is embedded in the value of a house. The further you have to commute, the less a house is worth.

Now, the problem, of course, is house value's determined by lots of things. So while it's true that in general, long commutes are associated with lower valued houses, it's also true that lots of other things matter. So for example, if you look at Everett, which is about-- which is about three miles from downtown Boston, most of the jobs in Boston, the average house price is \$300,000.

If you look at Lexington, where I live, which is three times as far away, the average house price is \$1 million. So that would seem to be backwards, where three times farther away, our house prices are worth more. That's because we've got good schools and nice houses and all that stuff.

So it's very hard. And likewise, all these decisions we've talked about, it's hard to use them. So the fact you use grocery delivery or not, maybe that's not about value of time. Maybe it's just you're socially awkward. You don't want to deal with people. Or the fact you drive one way or another, maybe it's just habit.

So the problem with a revealed preference approach is that decisions are endogenous. If I want to try to value your time by comparing two things, I've got the problem we discussed in chapter 3 of the bias problem, that people who use Waze, who get groceries delivered and don't get delivered, aren't the same. So you can't just compare them. And people live in Lexington and Everett aren't just the same. You can't just compare them.

So what we need is some kind of experimental or quasi-experimental method to look at how people value time. Probably my favorite example, in this case, for valuing time is this super cool paper that was done many years ago, actually in the mid '80s, but it's a really cool paper.

What they said is, look, during the gas crisis of the 1970s-- many of you may know, in the 1970s, OPEC was formed, and gas prices went through the roof. As a result, gas prices got very high in the US. Politicians, responding to this, put a cap on gas prices and said that you could not-- gas stations could not charge beyond a certain amount.

Well, we know from standard economics what that did. It created giant lines at gas stations because basically, gas was now artificially cheap because the government had capped it. So people waited in long lines to get gas.

The problem with that is-- so the-- it was not the problem. The interesting fact that I didn't know when read this paper, it turned out those gas caps only applied to national chains. So your Mobils and Exxons had that cap, but Joe's Gas did not.

So it turned out you had side by side, for random reasons, gas stations with cheap gas and long lines and expensive gas and short lines. And you could use that difference to value time, because essentially, you could choose. You could decide you want to go to Exxon, save money, and wait online. Or do I want to go to Joe's Gas, spend more, not wait online? And that's a true random choice. I mean, there might be some underlying preference for Exxon, but that's probably likely to be not a big deal. So that's kind of a cool quasi-experiment.

So what did he find? He found that-- in today's dollars-- that you saved about \$0.50 a gallon. Actually, precisely \$0.54 a gallon savings-- \$0.54 a gallon savings by going to Chevron, going to a chain.

You waited, on average, 14.6 more minutes. So you saved 14.6 minutes to save \$0.54 a gallon. And the typical purchase was 10.5 gallons.

So what that meant was the typical person was waiting 14.6 minutes to save \$22.50. Was saving 14.6 minutes. Oh, I'm sorry. No, that's wrong. I don't mean that. What I mean is a typical person saved \$27.50 per hour they waited. If you do that math, per hour someone waited, they saved \$27.50.

So that's the value of time. We can value time. People randomly had to choose. And their preference reveals that they would pay up to \$27.50 to save an hour. But beyond that, they'd rather-- they wouldn't rather save the hour.

What's very cool about that is in today's dollars, that's almost exactly the average wage in America, that actually the estimate you get from that is very similar to the estimate you get from the first approach. That's pretty cool.

These approaches are very different. And these two align. That's very good news. So it seems like we have a pretty good estimate of the value of time. The average wage in the US was about \$23. So it was a little bit higher but close enough for government work. So basically, it looks like you got a pretty similar estimate of what time is worth-- 23, \$27 an hour.

So that is an example of how we use these methods. Now let's come to the more interesting case, which is valuing lives. Let's ask the first question, which is-- God, that sounds awfully morally difficult and ugly to value lives. I mean, if you go to your cocktail party, and they're like, oh my god, we got to save this kid. He's down the well. And you're like, well, it depends how much it's going to cost to save him. Your friends look at you like you're an asshole. And you're just an economist. And that's why they call us the dismal science, because we make statements like that.

But in fact, society does not leave us a choice. We cannot choose whether or not to value lives. We have to choose whether to value lives, because all the time, policymakers face decisions that turn on the value of a life.

So some fun examples in the book-- so in 1993, GM-- swear to God, it's a true example-- GM made a new truck that was exploding. It turned out they'd put the gas tank on the outside of the truck. I'm not kidding you. Literally, they built the gas tank on the outside. It turns out when the gas tank is on the outside the truck, it's more likely to explode. GM probably should have known that.

And so there was a recall. There's recall. They said, look-- so the government said, GM, you've got to recall all these trucks, getting them off the roads. So the government estimated that the recall would cost \$1 billion and would save at most 32 more lives. 32 lives. So that's a value of life of \$1 billion divided by 32, or \$31 million.

So GM could save those lives, \$31 million. GM went back and said, wait a second, that's hugely inefficient. What if instead we said we are going to buy 200,000 child safety seats and provide all this money to educate people about drunk driving and wearing seat belts? And they estimated that that would save at least 50 lives at a cost of at least-- save at least 50 lives at a cost of \$50 million, or only \$1 million per life saved.

So GM said, look, you can make us do this recall, but that's going to be \$31 million per life saved. Or we can save more lives for way less money. And as you can imagine, their-- GM even saying that caused a massive uproar. And they had to recall the cars, but they were right. That's a rational calculation.

Let's look at another example. In 1999, there was a commuter train crash at London's Paddington Station, killing 31 people because basically, the brakes failed on the train. Basically, the safety system failed. So basically, London wanted to introduce a new safety system for their trains. And they had a choice.

They could put in a system called the Train Protection Warning System, which would stop any train traveling under 75 miles an hour, for \$700 million. For \$700 million, they could put in this one system. Or for \$3 billion, 3 to \$9 billion, they could stop trains traveling at any speed. They could have a better system for 3 to \$9 billion that would save one to three more lives per year.

So if we think this system was going to cost-- last 30 to 50 years, that would mean it would-- the cost would be between 20 and \$300 million per life saved. So by going from one system to the better system, we could save as many as three lives a year at a cost of 20 to \$300 million per life saved. Worth it?

That's pretty expensive. They decide not to do it. Were they wrong? Were they right? That's what we need to think about. But the point is, they couldn't avoid that decision.

And let's talk about the most pertinent example to us, which is COVID. This decision was pertinent during COVID. Now, first of all, if-- sometimes it's a no-brainer. So for example, aggressive testing for COVID basically cost a tiny amount of money per life saved. Or vaccines cost a tiny amount of money per life saved. Those are no-brainers.

But what about all the shutdowns? That's a harder case. People have actually valued that and tried to value how much we lost in economic activity versus how many lives we saved by the shutdowns. And they estimate that basically, the shutdowns cost the US about \$7.2 trillion in foregone economic activity and saved 1.24 million lives, falsely precise but roughly speaking, for a cost of about \$6 million per life saved. That's a big number. How big a number? Should we pay \$6 million to save a life?

So you cannot avoid this. And certainly in your lifetimes, you've seen it in bright lights in COVID, which is, how do we think about-- let's take \$6 million, if that was worth it or not.

So once again, we can go to the same three methods to try to decide if \$6 million a life worth it. The first is market value. What's the market value of a life? Well, once again, we have wages. You could say basically, the value of your time is wages. People work on average about 20% of their life. So we can take the average wage, take a 40-hour week-- or take the average wage applied to every hour you're alive. And what you get from that is a value of life about \$4 and 1/2 million. So that's below the \$6 million.

Now, clearly, that has some flaws. What might be wrong with that? Yeah?

AUDIENCE: The wage disparities among industries, positions.

JON GRUBER: Well, that-- this is a very important question. This is why we're unpopular. That is not a problem. Why is that not a problem, that there's wage disparities? Why is that not necessarily a problem? When would it not be a problem that there's wage disparities?

Let's make it simple. Imagine workers in-- garbage men make less than finance bros. Why is that not necessarily a problem? Garbage men lives may be worth less than finance bro lives.

AUDIENCE: That's not a problem if the deaths you're preventing are randomly distributed.

JON GRUBER: No.

AUDIENCE: It's not a problem if it's a competitive market.

JON GRUBER: If the marginal product of finance bros is higher, then they're worth more. That's the kind of shit we got to deal with as economists. So basically, essentially, the point is, in our economic model, the value of life is a function of many things, one of which is the marginal product of what you produce in society.

Now, you go further and argue, as is probably right, that what finance bros makes does not necessarily reflect their underlying social value. No offense to all you future finance bros and sisters. And then maybe we start to have an issue with wages. But let's be very clear. Just like Exxon, we have to start with what the basic economic model says. The basic economic model says that essentially, in a perfect labor market, your wage represents your marginal product, and the value of your life is your marginal product applied to all your hours.

So one problem with that, of course, is we think wages are a bad indicator of marginal product. This isn't a problem of valuing time, but it is a problem of valuing life. We have all the other problems we talked about, a reason why people can't set the value of wage equal to leisure.

We also have the problem that your wage may not reflect how much people like you, once again coming back to the finance bros. So basically, different people will be differentially sad if you die. And that should presumably reflect the value of your life too. Yeah.

AUDIENCE: Yeah. I got a corollary to that. A person's value is not the same as the value that they are able to produce.

JON GRUBER: Yeah, exactly. So market value is not great for lives. That's not really going to fly, although in legal circumstances, typically what's used. If you look at typically-- if you look at damages case, people are killed by accident. They'll often use forgone wages as a measure.

Second would be contingent valuation. Ask people. Now, you can't ask people what their life's worth. That's dumb. But you can ask people questions like, how much would you pay to reduce the risk of your life?

So for example, you could say, look, you're about to take a flight in an airline. You're about to take a flight from A to B. You have two choices. One is an airline where it crashes once every 50,000 times, and one's an airline where it crashes every once every 500,000 times. How much more would you pay for the second airline than the first?

And then you could say, well, that gives you an incremental value of a life, because basically, you're getting a probability of dying. They'll tell you-- they say if it's \$100 more, well, then that gives you \$100 is-- you divide that by the incremental change in the value of dying. That gives you a value of life.

The problem with that is when you use that method, you get the tight estimate of 1 million to \$32 million per life saved. People just don't know. I mean, how can you know. And that's why economists once again turn to our favorite method, which is a revealed preference. How do you get people to reveal their preference for their life? How do you do that? Yeah.

AUDIENCE: Very similar to the second but slightly different, we can say, if you-- I don't know-- consume normal foods versus organic foods and there's this probability.

JON GRUBER: Exactly. You look at choices that are made that reveal that differ in their risk. And in particular, one of the most significant choices people make is working conditions.

So if you are someone who is an unskilled worker 50 years ago in West Virginia, you could become an administrative assistant, or you could become a coal worker. Those had very different risk profiles. And you could ask-- but they had basically similar skill background requirements. You work at a grocery store, be a coal miner. Let's make that better. Similar requirements, but very, very different risk profiles.

And you can look at, how much higher are the wages in jobs that are riskier? So if a job that has a 1% chance of killing you-- a 1% higher chance of killing you every year pays \$1,000 per year more, then that says your life is worth a thousand over 0.01, or \$1 million.

So you basically look at how much more you have to pay to get people to take jobs that risk their-- that change their life risk by a certain amount. And there's a huge literature doing this. This is the economists' go-to method for valuing life. And the best estimates from this literature are that a life is worth about \$10.5 million, which is a big number, bigger than the COVID number we talked about before.

However, while I agree with most economists, this best method to value life, and it's the number I use in my work, in references, there are a number of problems with this. What are some of the problems with this? What are some of the problems of using some real preference method? Yeah.

AUDIENCE: Determining the risk of someone's occupation.

JON GRUBER: Yeah. Well, not just determining it but the person knowing it. I mean, I can measure it statistically, but the person has to know it. So there's an information problem. Yeah.

AUDIENCE: People don't necessarily have-- like coal miners don't have free choice in--

JON GRUBER: They don't say a free choice in what they do. Great point. Other points? Other reasons why this might be problematic? Yeah.

AUDIENCE: The difference between coal mining versus working in a grocery store might be significant.

JON GRUBER: Great. The only difference is not just that it kills you. There's lots of things. The working conditions suck. You might get injured. This compensating differential, which by the way, this is a concept-- the concept of compensating differentials, I should have mentioned, goes back to Adam Smith. He actually talked about this in 1776. It's *The Wealth of Nations*, the first important economics text. He talks about this idea.

But the problem is that compensating differentials with two jobs is the differential for all the amenities between the two jobs, not just death. And it's hard to find two jobs identical except for the risk of death. Typically, the risk of death comes with other health risks as well.

The other problem is, in some sense, you're overstating the value of life if you use this differential, because you're actually valuing the value of everything that's crappy about this job. Other problems? Yeah.

AUDIENCE: Even within jobs, risks for individuals may differ, and individuals may know that.

JON GRUBER: So this is very interesting. This is a very important point I want to come to that's come home very clearly to me with my daughter's work, which is when you use this method, you only get-- and this comes to the point also about West Virginia workers having choice. When you use this method, you only get the value for the workers who are on the margin of deciding between the two jobs.

So I don't get your guys' value of life by comparing coal miners and grocery workers in West Virginia. You guys are nothing like that. So if there's a difference in value of life, that give a very different answer. My daughter's an EMT. It's an incredibly hard, stressful, and skillful job. She had to study six months, pass an exam, very hard job. She took a pay cut, moving from Starbucks to EMT. Why? Why? Yeah.

AUDIENCE: She was [INAUDIBLE] job.

JON GRUBER: Yeah, she's an adrenaline junkie. She wants that job. She'd do it for almost free. She's in school now. I told her, you don't have to work. I'll pay your living expenses. She said no, I have to work.

So basically, the point is if there's differential tastes for risk-- so the simple example-- imagine a simple example. Imagine the workers are deciding, and there's 900 coal miners jobs, 900 coal mining jobs to be filled. And there's 2,000 people. And 1,000 of them are risk lovers, and 1,000 are risk avoiders. Who's going to take the coal miner job? The 900 risk lovers.

So they're not going to get paid a whole lot more, because they're already risk lovers. So you're getting an underestimate of the value of life if you look at what actually takes. If you took my daughter's job, you'd say, gee, you don't have to-- people don't really care. They're happy to put in putting stressful, life-threatening situations. Well, they're not. It's just people that take those jobs are. It's not an appropriate estimate for the average person.

And that's an important bias here as well, that you're not getting the overall estimate. This is fundamentally at the econometric level what Josh Angrist, my colleague here at MIT, won the Nobel Prize for, was pointing out, basically, that whenever we do economic estimates, we get the estimate not necessarily for the average person but for the marginal person whose behavior has changed.

If you look back at chapter 3, we don't talk about this in chapter 3. Look at chapter 3, all the examples. We're looking at deviations, like a welfare benefit changes by \$100. We're looking at people who get-- whose behavior's changed. That's a small subset-- slice of the population. If they are not representative of the whole population, we're not getting the right average answer. We're just getting the right marginal answer.

So this is an example of that here, which is-- and it's a particularly worrisome example because there's a set of people who may love risk or be less risk averse, and they're the ones who you're getting an estimate for. Questions about that? Yeah.

AUDIENCE: So based on this argument, would you say that 10.5 million is an overestimate?

JON GRUBER: Well, no. So for that argument, I'd say 10.5 million is an underestimate. Other arguments might say it's an overestimate because it accounts for injury and all the stuff that's crappy about the jobs. So that's why I end up using it, because I don't know which way the bias goes.

There's also the fact. There's the information fact. There's another fact, which is we talked about behavioral economics and people are not good at optimizing over these things. So those people may not be making this decision appropriately, even if they have perfect information. They may not make the decision appropriately.

So there's a lot of biases to this. I think there's less bias to this than these other methods. And I think there's no-- there's been no evidence of the bias systematically goes one way or another. So that's why I tend to use that estimate.

Another fun thing about value of life is, should the value of life-- how should the value of life vary with age? So you're a policymaker. You're a king. You don't worry about political correctness. You get to decide how much to spend on saving people of different ages. How would you change the value of-- how do you think the value of life-- what should be the age profile of the value of life? Yeah.

AUDIENCE: [INAUDIBLE] should be.

JON GRUBER: Should what?

AUDIENCE: It should peak around 20, 30.

JON GRUBER: Yeah. So basically, it should be low when you're a baby, peak around-- not actually 20 or 30, because your marginal product is always going. It should really peak-- because remember, it's not just marginal product. It's also attachments you have to people. It's all the things that are value having you alive in society. Those probably peak-- I don't know-- 50, 60, something like that, when you're at the peak of your earnings power and when you maximize your attachments to the world. And then it should fall again, because you're not earning. People are dying. You're not as attached to them, whatever.

If you actually do estimates by age of the value of a life, it actually looks like this and then like that. It doesn't drop. That is, people who are 90 think they're worth just as much as people who are 50, even though our models would say they don't, which is once again says that's a problem with the market value approach. So that's another-- so this is a fascinating literature. Here's another really cool-- yeah.

AUDIENCE: That line comes from self-assessments or other people?

JON GRUBER: This is from compensating differential studies, because when people are older now, you might say they're not working. We have also consumption-based studies of the kind that you suggested, like how much more we pay for an airbag in a car.

AUDIENCE: Would we consider it like your potential to create value? So as a child, you have your whole life ahead of you.

JON GRUBER: Well, that would be-- that's a great point. Many people would say, well, it should be potential. This is a replacement value model. I mean, this says, the first day you're born, you're worth the least, because you're not earning anything, and you're replaceable nine days-- month.

So that's awful. As a parent, I can't imagine anything worse than saying that. But that is what the economic model would be saying. But if you're saying potential, but the point is you could replace that potential. So that's kind of the idea of why it's low. Once again, why we're not popular at parties.

So this is actually a really cool study. In February 2003, the state of Oregon had massive budgetary problems and laid off a bunch of police. So they couldn't do traffic stops anymore. And people drove faster and died more on the road.

So this study looked at how many more deaths were caused by driving faster versus how much the state saved by laying off the police. What they found was there would have been 2,167 fewer fatalities, and the state saved \$670 million, which implies a value of life of \$309,000. That is if life is worth more than \$309,000, they should not have laid off the troopers. Life is definitely worth more than \$309,000.

So by that metric, it was a huge mistake to lay off these troopers, which once again speaks to the fact that society is not always pulling the same direction. They had a budget crisis. They had to do something. There's no social planner saying, oh, don't do it. Life's worth more than that. I'll give you money. So that's an example of the problems that we face in society, making these hard decisions. Questions about that?

What I really like about this lecture is cost-benefit analysis seems like it'd be almost the most-- almost certainly the most boring thing in the world, but it turns out to be really interesting and fun because there's all these really interesting moral questions and ethical issues that come up. Yeah?

AUDIENCE: So my question now is, given all these problems we run into, societally, are we accepting that there is one cost of life for every person or that they differ?

JON GRUBER: Well, not just one cost life for every person. There's a great-- there's age. There's income. There's countries. How do they value in life in Bangladesh versus valuing a life in the US? There's all these difficult issues. And basically, we don't have very clear answers.

So basically, what we do is we have to pick a number. The government picks a number. But that number even varies across government agencies. So it's really, really a challenging issue.

But once again, because it's challenging, we can't throw up our hands and say, we're not going to deal with it. We have to deal with it. It's not a choice. The COVID example shows that. So we just have to keep working on trying to find it and do better and better job.

Some things will never be resolved. We'll never resolve how to value a life in Bangladesh versus the US, because fundamentally, economics would say the life in Bangladesh is worth less. Fundamentally, morality would say it's not. And that's an unresolvable gap. That we're never going to resolve. Yeah, go ahead.

AUDIENCE: Is there a difference between how someone values their own life and people value that person's life?

JON GRUBER: Great question. That's another problem with this measure, which is it's what I think about my life. I may think nobody likes me. And a bunch of people might like me. Or I think people like me, and they don't. So that's another problem with this measure. But this is really interesting.

So this at this point, you can be pretty frustrated, like, well, I don't know what to do. I don't know what's life worth. But we can realize there's an important thing we can obviously do, which is the fact is the government every day makes decisions that save lives. And we can ask, are they making them in the right priority order?

So if you look at table 8-3, this looks at the social cost of a variety of regulations the government puts in place and how many lives and what it cost per life saved. So for example, enforcing that lighters are childproof does not cost the lighter producers much money but saves a ton of lives. So you say, basically, it costs \$140,000 per life saved by childproofing lighters.

On the other hand, regulating cattle feed so we don't get mad cow disease has a tiny chance of affecting our health but costs a huge amount. So that's \$250 million per life saved. So what you can do is you can say, look, even if I don't know what the number is, I know we should put more money towards things on the bottom of this list and less money towards things-- more effort towards things at the bottom of this list and less effort towards things at the top of the list.

So at least it can help you with relative answers or what we call cost effectiveness analysis. Cost effectiveness analysis is not necessarily, what's the value? But what should we do relatively? What action should we take relative to others?

So at least putting valuation on lives can at least help us rank different interventions, even if the absolute level is really hard to pin down. So there is value in that. So to the extent we're a little depressed about the fact it's hard to know how to value this, at least there is a useful piece, which by doing these exercises, we can at least rank things and say, hey, among decisions, this would be the most efficient decision. Yeah.

AUDIENCE: It probably varies from between each of these. How linear are these numbers? If you were to pour in a certain amount of money--

JON GRUBER: Yeah, no, that's great. It may be like we can't do more to childproof lighters. They're childproofed. I agree. But think of this more along the decisions we'll talk about, like, for example, expanding health insurance programs. You could expand health insurance programs, or you could subsidize better highways. They save different things. You could use the valuation to try to get at that. Yeah.

AUDIENCE: Does this include marginal benefits of scaling, because once you have every car with a seat belt, presumably producing--

JON GRUBER: That's a similar question, which is, it's not linear, and there could be cliffs at which beyond it, which it doesn't help. Yeah.

AUDIENCE: How much does altruism matter in this? Like, how do we [? define if ?] I'm a vegan or vegetarian? I don't necessarily-- I mean, I wouldn't care for myself [? or anything like that. ?]

JON GRUBER: No, I mean-- well, actually, that's a great question, which is this whole discussion is taking this theoretical perspective of the social planner, that I'm one benevolent person trying to do what's best for society. As we'll discuss more next chapter, we'll talk about political economics. Right now, people don't seem to feel that way about our society. People aren't benevolent towards their fellow citizens, so that's another factor to put in.

But I think that's the last issue I want to worry about. Let's assume this benevolent social planner. It makes life a lot easier. Other questions, thoughts?

OK, so let me talk about one more issue, which is, what about the fact that a lot of these benefits happen in the future? Now, we talked last time, when costs happen in the future, we discount it. We put it in present discounted value. And you can do a similar thing with benefits.

But it's a little, once again, ethically difficult, especially something like global warming. Even with a very low discount rate, lives in 100 years aren't worth a whole lot relative to lives today. And lives of 500 years are worth almost nothing, by the time you discount them. Do we really want to feel that way?

So imagine global warming worked as no one died for 200 years and then everyone died, or not everyone died, but then a huge number of people died. Well, if you said, gee, I could save 1/10 of those people today. I might rather do that, because by the time I discount, those people in 200 years aren't worth that much. Well, that sounds kind of weird, doesn't it? People aren't worth something in the future?

So that's another tricky ethical issue, which is, how do you discount the future? And once again, the standard model has got a clear answer. You should discount it because we'd save that money today of all that money in the future. But money does no good when the Earth is on fire.

Our standard model is assuming it's a substitution of money, that money can solve the problem, but money can't. 200 years from now, it's too late. Money can't solve that problem. So for global warming, many people say we should use a lower or no discount rate, because you can't easily trade off the money between now and the future.

So for example, think about a different example. Imagine investing in food today versus investing in food in the future or something like that. Then you would want to use discount rate. You want to say, well, if we invest in food today-- if we invest in food in the future, that's money we could put aside today. We'll have more money in the future. We want to account for that. But if it's stopping global warming in 500 years, you can't do that, because the money in the future is no good if we're all underwater. So that's another difficult-- once again, not to make this even more difficult. It's another tricky issue that we have there.

Let's put it all together. We put it all together in table 8-4. We already talked about the cost of this project being \$253 million in present value. We're going to then have time saved and lives saved, but those are going to go off into the future. So we're going to discount them. Let's use the same 7% discount rate. Let's ignore the global warming issue I just talked about. And you end up with benefits that are much, much larger than costs.

So that would say you should do this project, that the value of time saved and lives saved is massive relative to the costs. It turns out that lives are worth so much that when you're saving lives, most projects want to get done. So it turns out that once a project involves saving lives, it's going to be a pretty inefficient project not to want to do it, because the number is so massive. \$10.5 million a life is a lot. So that's where we end up. Questions about that?

Let me talk about a couple more things before we stop. The first is how to-- if you ever find yourself in this situation, how to avoid the mistakes noneconomists make. And a common mistake noneconomists make is double-counting benefits.

So for example, if you do this analysis, some would say, well, that's good, but you've understated the benefits because you've also created all these jobs. What about that? What are the value of creating the jobs? What would your answer be? Why would you explain that's wrong? Yeah.

AUDIENCE: We already accounted for the opportunity cost.

JON GRUBER: OK, that's kind of a nerdy answer. Like, they'll be like, OK, I don't understand you. Well, how would you answer them in more intuitive terms? Yeah.

AUDIENCE: They would have found a job elsewhere.

JON GRUBER: Yeah, they would have found a job elsewhere, not creating a job or taking a job from somewhere else. Or, for example, if someone said, yeah, well, not only did you save time. You raised the value of houses, because now it's quicker to get back and forth from work. Why shouldn't that be counted? Why shouldn't you count raising the value of time? I'm sorry. Why shouldn't you count raising the value of houses along the highway?

And the answer is you shouldn't. Why not? Think about why those houses are worth more. Why are those houses worth more? Yeah, Alec.

AUDIENCE: Because of driving time saved, which is already included.

JON GRUBER: It would be double counting. You've already counted that. They're worth more because you saved time. You already counted saving time. These are the kinds of mistakes that people commonly make in this that we need to use our economic intuition to avoid.

Now, the last point I want to come to-- I'll discuss this early in my office hours-- is distributional concerns. This course so far has been efficiency, efficiency, efficiency. And we don't care rich versus poor. We discussed a little bit here, but largely, we haven't talked about that.

But distributional concerns can be really critical when doing cost-benefit analysis, because the costs and benefits can accrue to very different types of people. For example, a lot of the most serious environmental damage in our country has been done to some of the most disadvantaged communities.

So if we said, Should we clean up this environmental damage? we might say, well, overall, it's not worth it. Or we might even say more. Well, it's certainly not worth it, because they're poor people, so their lives are worth less. But that's worrisome.

How do we deal with that in a rigorous way? What economists do is they say we have distributional weights or social welfare weights. Remember, let's go back to our standard utilitarian social welfare model. In our standard utilitarian social welfare model, we want to redistribute income until everyone's marginal utility of income is the same.

Well, poor people have a higher marginal utility of income than rich people. So that would say we want to redistribute from rich to poor. That would say that a project which taxes the rich and distributes the poor should get a special plus up relative to a project which is just tax the rich and benefits the rich, or even worse, tax the poor and benefits the rich. That's what we call a distributional weight.

Distributional weight, W-E-I-G-H-T, would be that basically, you're going to put extra value based on who's helped and who pays for this project. And that value would be weights which are higher on the poorer people because in a utilitarian social welfare function, they have a higher marginal utility of consumption, so valuing resources then more highly. Yeah, Steven.

AUDIENCE: [INAUDIBLE] highways [INAUDIBLE]. I remember a city where the highway was built over a historically marginalized community. And crime there is an issue there that-- there's crime in that area of the city that may have been exacerbated by having that project over them, for example. When you have lives saved numbers like this, I imagine that there's not-- it's hard for them to capture that information. [INAUDIBLE]

JON GRUBER: Yeah, no, I mean, this is once again-- once again, you can look at any problem like this and view it as an opportunity or challenge. You can say, God, this is really fucking hard and throw your hands up and walk away. Or you could say, that's really cool. Let me try to add that into my analysis and figure that out. That's not been done well, but it should be. We should look at all the benefits and all the costs and to whom they're attributed. And that would be exactly the kind of interesting thing to build in these kind of cost-benefit analyses. Yeah.

AUDIENCE: [INAUDIBLE] if you assume that there is some marginal damage to society by having income inequality?

JON GRUBER: We're going to talk in chapter 17 more about theories of income inequality. But essentially, there are many theories of income inequality which lead you to the same spot, which is from today's perspective, we'd be better off taking a dollar from rich guy and giving it to a poor guy, and there's many ways to get there. But at the end of the day, you get to that same spot, which fundamentally most Americans-- not all, far from all-- but most Americans would agree that all else equal, taking a dollar from a rich guy, giving it to a poor guy would make society better off. And we'll talk at length about that in chapter 17.

Now, speaking of difficulty of getting Americans to agree on things, let's now turn to chapter 9 and political economy. As I said, the fourth fundamental question-- in the first lecture, the fourth fundamental question about economics is, why do governments do what they do?

At one level, this is the ultimate overarching question of public finance, which is I spend this whole lecture assuming some benevolent social planner. And it's hard to teach this chapter without revealing one's preferences. But I certainly think it's hard to imagine Donald Trump as a benevolent social planner. You may like or dislike him, but certainly, that would not be your description of him.

So basically, the question is, how do we think about a world where decisions are not made by some theoretical social planner but by real people? And that's the field of political economy, which builds in government decision-making.

Now I'll only do a lecture and a half on this, which is very unfortunate. Fortunately, our department has the best political economists in the world teaching here. Daron Acemoglu, who wrote the book *Why Nations Fail*, probably the most important political economy book ever written, or at least in the last century. Ben Olken, who does amazing work on how we think about governance in developing countries. Abhijit Banerjee, Nobel Prize winner who's done incredible work about the theory of political economy.

We have incredible resource in our department, and there's multiple courses you can take. So I urge you, if you want to learn more, to take those classes. Unfortunately, here, I only have about a lecture and a half.

So here's how we're going to think about political economy. Next time, we'll do the fun gossipy stuff. Now we're going to do the more rigorous economic stuff. Here's how we're going to think about political economy.

How we're going to think about it is we're going to proceed in several steps. The first step is going to be, imagine a theoretical world where we could implement the perfect form of political economy. Then we'll say, OK, that doesn't exist.

A second alternative would be, imagine a world where literally everybody, every American, voted to choose every item, no politicians, or what we call direct democracy, which occurs in some cases but not in most. And then third, and most interesting, we'll talk next time about representative democracy, which is what we have and most countries have, where we don't literally make every decision. We elect people who make the decisions. So that's the three steps we're going to proceed.

And we're going to start with an example of what we call a model of unanimous consent. So this first model is a model of unanimous consent, which is, imagine a world where we can only do things if everyone agrees. And imagine a world where we could actually make that happen. We get everyone to agree.

That is a world of what we use-- what we call Lindahl pricing, named after a Swedish economist who developed this model in the early 20th century, where he basically said, look, as long as we charge every individual their marginal willingness to pay for a public good, we can get unanimous consent. To see that, let's go to figure 9-1, which is a replication back from an earlier-- this replication example from an earlier chapter, chapter 7, which was demand for fireworks.

And we talked about marginal willingness to pay fireworks. Back then, it was Ben and Jerry. Now it's Ava and Jack. So the funny thing is when I wrote this textbook 20 years ago, 22 years ago, every example involved the names of one of my kids. I've now had to change it because that's not really fair. And I was very uptight about my kids being mad at me, Ava, Jack, and Sam. And Sam is now Rachel. But Tom was Sam, and I was very upset about my kids being mad at me, so I carefully went through and made sure I used every kid's names exactly the same number of times. And the ironic part is, none of my kids have ever opened the book. So I spent all that time for nothing. Anyway, Ava and Jack or two of my kids.

And let's imagine that they have a willingness to pay for fireworks. And let's say Ava's willingness to pay is depicted in the first-- in panel a. Jack's in panel b. And the total social willingness to pay, in panel C. Just like we do with Ben and Jerry, it's their willingness to pay for fireworks.

OK, so Ava's willing to pay \$1 for the first firework. And by the time you get to the 100th firework, it's worth nothing to her. Jack's willing to pay \$3 for the first firework, but also the time you get to 100 fireworks, worth nothing to him.

And so society as a whole is willing to pay \$4 for the first firework, with its diminishing demand. And if the marginal cost of fireworks is \$1, then the optimal number of fireworks to provide is provide is 75. That's the graphical implementation of the Samuelson rule, the sum of the marginal willingness to pay.

So here's what Lindahl said. What if we set up the following system? It's in your textbook. So just talk it through. It's hard to write it all down.

Imagine the government says, we're going to set a set of tax prices for fireworks. For every firework we choose, we're going to charge each person a certain amount of money. Fireworks cost \$1 each. It could be \$0.50 each, 90/10, whatever it is. We're going to charge each person a tax price for each firework. It's going to add up to \$1. That's the only constraint. Fireworks are \$1. We're not going to make money. It's going to add to \$1. But we set these tax prices.

Each individual would then say, well, at that price, this is how many fireworks I want. Well, the government, by doing that over and over again can essentially trace out the demand curve. So I go to people and say over and over again, how many fireworks do you want at \$0.10, at \$0.20, at \$0.30? I trace out the demand curve. Once I have the demand curve and know marginal cost, I can figure out the optimal number of fireworks. And then I can charge people what they said it was worth to them.

So in this case, if the optimal number of fireworks is 75, I say, look, it's 75 fireworks. Fireworks cost \$1. Fireworks cost \$1. So that's \$75. How am I going to pay for that? Well, Eva said, for 75 fireworks, she's willing to pay \$0.25 each. Jack said, for 75 fireworks, willing to pay \$0.75 each. So I charge Jack \$0.75 on 75 fireworks, so like \$49. I charge Eva \$0.25 for 75 fireworks. It's like \$21. And it adds up-- or \$26. It adds up to the \$75, whatever the numbers are. I forget what I have here.

I charged Jack-- so I charge-- whatever, I just charge the amount. So basically, this delivers the Samuelson equilibrium. Each person you are adding up to the proper number of fireworks. You're covering your costs, and everybody is paying their true marginal value. You're done.

And there's no reason to deviate. Ava and Jack are both perfectly happy. They'd be happier paying less, but they're on the demand curve, so they're happy. So you've achieved the first best.

So that would be the ideal political economy system, would be we could get the first best Samuelson equilibrium. Now, in fact, we can't do that. And we can't do that for the three reasons I laid out. Was it last lecture or a couple lectures ago? I think it was last lecture.

The first reason is the preference revelation problem, which is if I set up this system, people have an incentive to lie. Let's work through the math. Imagine that Jack lied and said his preferences were the same as Ava's. In that case, we would say, well, if Jack and Ava were the same, we would end up producing 50 fireworks at \$0.50 each. If there were two Avas, the optimal number of fireworks is 50 fireworks, \$0.50 each.

What would that mean? That would mean that what Ava pays-- Ava used to pay \$0.25 for 75 fireworks, or \$18.75. She used to pay \$18.75 and get the same 75 fireworks Jack got. Now she pays more. She pays \$25 and gets fewer fireworks. She gets 50.

Now, Jack gets fewer too, but he doesn't pay that much less. Jack gets fewer too, but he pays a ton less. I'm sorry. Jack used to pay \$56.25 for 75 fireworks. Now he pays \$25 for 50 fireworks, so he gets fewer fireworks, but he pays a ton less. What's happened here? Jack has become a free rider on Ava. By lying about his preferences, he's made-- forced Ava to contribute more.

He's hurt himself. He gets fewer fireworks, but he's paying less than half as much to get 3/4-- to get 2/3 of many fireworks. Ava is paying more to get fewer fireworks. So by lying, he's managed to make himself better off at Ava's expense. So there's a huge incentive to lie. That's the first problem. People get that? Huge incentive to lie.

Second problem is people may not know. We don't know what things are worth. People may not know what fireworks are worth.

And the third problem is aggregation, which is this is all well and good in a two-person society. But if we had to collect every person in America's preference for every single public good we provide at every possible price, the entire population would work for the government collecting that data from each other. It wouldn't work.

So this system is a great theoretical benchmark of how you achieve the unanimous first best, but it's irrelevant in practice. What it does highlight-- and we'll come back to this in later lectures-- is that decisions will be more popular, the extent to which they can link the taxes people pay and the benefits they get. One way to make a decision popular-- and this will come to explain a lot of decisions governments make-- is try to link to get closer to Lindahl equilibrium by linking what people actually pay and the benefits they actually get. Yeah.

AUDIENCE: I think that makes sense in theory. But in practice, wouldn't most people say that they would want lower taxes if that meant not spending more money on roads, even though we all benefit. But people behaviorally are going to say, I want lower taxes.

JON GRUBER: Right. So basically, first of all, it's a couple of things. Maybe they don't realize what it's going to be like to have shitty roads. Second of all, they may be lying. And third of all, how do you put that all together? Once again, the tax benefit linkage would be, how do you raise money in a way that more directly links the benefits of better roads to the taxes people pay?

Well, let me ask you a question? At what level of government do we fund road repairs? The federal government, the state government, or the city government?

AUDIENCE: Here, it's going to be city.

JON GRUBER: City. Why does that help answer your next question? Alec, yeah.

AUDIENCE: Their taxpayers are going to be the ones directly benefiting from the roads.

JON GRUBER: Yeah, if I'm taxing you to pay for roads in Alabama, you're like, screw that. So basically, that's exactly-- that's why this model can actually help explain decisions governments make. And we'll come back to that in chapter 10.

All right. So that's the first method. Well, it has some interesting directional predictions, obviously relevant in reality. Second method says, OK, let's move from unanimous decision-making, which never works, to the common decision-making mode used in the world, which is majority decision-making.

And the best way to tell if a majority of people want something is with direct democracy. Let's just have everybody vote over everything. If you vote, you solve the aggregation problem because you're voting. And you potentially-- and we'll come back to this-- solve the revelation problem, because there's no reason to lie when you're voting. But we'll come back to why that may not be true.

So direct democracy-- and direct democracy has a long history in the US. The first town meeting in the US that's recorded was held in 1657, where the-- in 1657, the residents of Huntington, New York, held a meeting to vote to hire the first schoolmaster of their town. There's still town meetings held all over New England. Many decisions in New England are still made by town meetings. People get together. And they vote on what they're going to have. Literally, there's no politician making the decision.

This actually led to an interesting case in New Hampshire. The problem with that, of course, is many people don't show up. And what happened in a town in New Hampshire is there was this guy who was a very extreme antigovernment person. And he realized in this town that town meetings were poorly attended. So he gathered a few of his friends, and they went to the town meeting and voted to basically end public financing of all things in the town. It basically turned into a purely capitalistic town with no government, no schools, no fire department, nothing.

People got upset. They had another town meeting where people actually showed up and voted him down. But this is the challenge. Aggregation only works if people show up. So that's an example of direct democracy.

More generally, many of you may be from states where there are ballot referenda or ballot initiatives, which are when you go to vote for your politicians, you also vote on a series of questions. A referendum means that-- is a case where there's been a law that's already passed, and citizens can basically ratify or not ratify it. An initiative is when voters can actually decide an issue.

Now, initiatives are-- the popularity of these methods have waxed and waned over time. 24 states allow initiatives. It started with Oregon in 1904. More than 2,600 initiatives have been in state ballots. And about 41% have passed, although 60% of all initiative activities occurs in only six states-- Arizona, California, Colorado, North Dakota, Oregon, and Washington.

The popularity of initiatives waxes and wanes over time, as do the topics. So early in their lives, initiatives were used for things like taxes. So the most famous, of which we'll discuss in chapter 10, was the passage in California in 1978 of proposition 13, which limit how much revenue towns could collect in property taxes. That was an early initiative.

Then, basically, the earliest-- there were things like election rules, alcohol regulation, labor laws, and government, were first things on. Then there was taxes. Now, the most popular initiative, of course, is marijuana legalization. But there's lots of other things that are on ballots as well. And for Massachusetts, we've had initiatives about charter schools, initiatives about the minimum wage. Lots of things are decided by initiative. Yeah.

AUDIENCE: How up to date is marijuana legalization [INAUDIBLE] most popular, because I feel like abortion may have certain--

JON GRUBER: Abortion may have. Now, by this round, abortion's on a lot of ballots. That's a common example. Good point. Yeah.

AUDIENCE: Are there ever federal referendums or initiatives?

JON GRUBER: No, never.

AUDIENCE: Is there a reason why?

JON GRUBER: It's just not the way the US government is done. Now, there are in other countries. Switzerland makes almost all its decisions by full majority voting on initiatives, things like what the tax rate should be. That is not done in the US. There are state-level initiatives on things like tax levels but never done federally that I know of. I could be misspeaking for the record, but that I know of, there's none.

OK, now let's-- what I want to do in the remaining time is write down the economic model of majority voting. And in particular, I'm going to talk about when majority voting actually works and when it doesn't work. I'll define what working and doesn't working means.

So let's talk about, what principles would we like in a voting system? The fundamental principle we'd like in a voting system is that it consistently aggregates preferences. The fundamental goal of a voting system is that it consistently aggregates preferences.

What do I mean by consistently aggregating preferences? I mean three things. The first is it must achieve dominance, which means that if one choice is preferred by all the voters, by a majority of voters, that choice should be chosen. No, all voters. I'm sorry. If there's one choice preferred by all voters, that should be chosen. If that's not true, the system is not consistently aggregating preferences. If everyone wants something and the voting system doesn't deliver it, then you're clearly not aggregating preferences appropriately.

The second is our old friend from kindergarten-- transitivity. And I say kindergarten because you've probably-- now in all the math classes you guys have been in, you're sick of the term "transitivity." You all know what it means. So it's just transitivity. We want if a choice A is preferred to B and B is preferred to C, we need A to be preferred to C. That's the second condition we'd like.

So those are two I would say noncontroversial decisions. The third decision is a little more controversial but in political economy generally assumed to be important, which is what we call independence of irrelevant alternatives. Independence of irrelevant alternatives. And this system says that the choice between A and B should not be affected by the random introduction of C.

So let's say, ignoring the financing-- ignore the money. Let's say your town's choosing between erecting a statue and building-- basically a statue and a park. You're going to build a park or a big monument. And your town votes and prefers the statue to the park.

Now, let's say a new option is introduced of a new police station. Now, that could win, but it should not affect people's preferences on a statue versus a park. The fact that this third alternative-- I'm not saying the alternative couldn't win. I'm just saying relative ranking should not change when a random independent alternative gets introduced. That is our third condition we impose. And that's what political economists impose, is the conditions of a consistent voting system, consistent aggregation.

Now, what we're going to show you is when majority voting systems do and do not meet this condition. So let's start with an example of a town that's deciding on levels of funding for public education. Let's go to table 9-1.

In table 9-1, there are-- we are voting over whether to have how-- whether we should vote-- how we should fund the level of school. And there's three choices-- high, medium, and low. High means great schools, high taxes. Medium means medium schools, medium taxes. Low means crappy schools, low taxes.

And there's three groups of voters. There's parents. There's elders. And there's young couples who aren't sure if they're going to have kids before they move. And then we have, for each group, the ranking of their choices.

Parents choose-- their first choice is high funding. They want good schools. They have to pay high taxes. Second, medium. Third, low. Elders are the reverse. These are selfish elders, not altruistic elders. Elders are like, I want low taxes. I don't care how crappy the schools are. The worst outcome for me. So I want low taxes, then medium, then high.

The young couple, they're not really sure. So they're like, well, I'm not sure I'm going to be here. So I want to hedge my bets. My best is medium. My worst is low, because I want to be stuck in a town with crappy schools. I'm sorry my middle is low, because if the schools are crappy, I'll just move. The worst for me is high. I pay high taxes in a town that may not even stay in.

So I mean, I'm not saying these are right preferences. But I'm saying these are not crazy preferences for people to have. So imagine those are people's preferences. If you had those preferences use majority voting, you will deliver a consistent outcome.

Now, majority voting-- let's think of majority voting as happening pairwise. Think of pairwise voting. So for example, no matter how you have pairwise voting, you always end up with consistent outcomes.

So for example, imagine we voted on H versus L. Look at this table. If we had H versus L, what would win? Yeah.

AUDIENCE: L.

JON GRUBER: L would win. L is preferred to H by two out of three voters. So L would win. So now what about H versus M? Same person. H versus M?

AUDIENCE: M.

JON GRUBER: M wins. So we got our two semifinal winners. So let's go to the finals, M versus L. What wins?

AUDIENCE: M wins.

JON GRUBER: M wins. And it turns out, no matter what order I did those in, you'd get the same answer. M would win. No matter what order I did those in, you would get the same answer. M will always win in this case. Majority voting works, as opposed to-- you might say, that's-- what do you mean by works?

Let's talk about an example where it doesn't work. Imagine that the elderly person moves out, and they are replaced by someone who either is going to send their kids to the public schools if their good or the private schools if they're not. So we're going to turn, flip over to the next page. Go to table 9-2. Oh, you lost table 9-2. OK, we don't have table 9-2.

Table 9-2 looks like this. And this when I copy from the book so I don't screw this one up, although I will say you guys were most active when I screwed that graph up. So I'm not going to screw stuff up on purpose, but I want to keep that level of activity.

So now let's imagine you've got your public school parents. You've got your private school parents. And you've got your young couples. First choice, second choice, third choice. So we know the first and third. This is H, M, L. This is M, L, H.

For the private school parents, they have an interesting choice. Their first choice is crappy schools. They send the kid to private school. But their second choice is H. And their worst choice is M. Why? Because M is the worst. They pay medium taxes, and they can't send their kids to school. At least with H, they'll then send their kids to local school. But the worst is a town with pretty high taxes and mediocre schools. They'd either rather have a town with low tax and crappy schools-- they'll go to private school-- or a town with high taxes and good schools. Then they'll go to public school. Once again, not crazy preference for someone to have.

But with these preferences, we get a very different outcome. Let's do the same voting as before. H versus L, what wins here? H versus L?

AUDIENCE: L.

JON GRUBER: L wins. L gets two votes. H gets one. OK, L wins. H versus M?

AUDIENCE: H.

JON GRUBER: H wins. So H wins over M. And H wins versus L. So you might think we're done, but what about L versus M? So I'm sorry. L wins over H. H wins over M. So we've got L versus-- but imagine we did-- OK, so now let's do L versus H. Who wins that?

No, we already did that. I'm sorry. L beats H. H beats M. So L versus M, what do we get?

AUDIENCE: M.

JON GRUBER: M, which is inconsistent. It violates transitivity. And in fact, no matter how you do this vote, you're going to violate transitivity. You're never going to get the condition of aggregation. No matter how you do this vote, the order is going to violate transitivity.

And indeed, no matter how you set up the voting mechanism, you can never guarantee a consistent outcome. I say, well, this is one voting system, Jon. What about other voting systems? Let's consider some other systems.

Imagine, let's let everyone vote in their first choice. Well, in the first case, if we let everyone vote on their first choice, we'd get a tie. In the second case, we get a tie. It wouldn't help us. So even first choice doesn't help us.

Let's do another thing like my dorm used to do. We'll do weighted voting. 3 points for first choice. 2 points, second choice. 1 point for third choice. In the first example, M would win with 7 points. In the second example, once again, it's going to be a three-way tie with 6 points each.

So no matter how do you set up this voting system, you're not going to get a consistent outcome. Indeed-- let me just come to this exciting punchline. Indeed, this is the proof, a graphical proof, of a very complicated set of math that I cannot do-- because I'm not going to win the Nobel Prize, but someone who won the Nobel Prize did do it, Kenneth Arrow-- that proves mathematically that you cannot ever have a system which ever and always satisfies these three conditions, except a dictator. Absent a dictator, you can never have a system which satisfies all three of those properties in generality. Alec, question.

AUDIENCE: So how often is this actually a problem? Because to create this, you created a very artificial--

JON GRUBER: Let's talk about that. Great point. Let's go to the figure. Let's go to the figure. The figure shows figure a. So what we're doing here is we're mapping. On the x-axis is the levels of school spending. On the y-axis is preferences-- first, second, third choice.

So look at, for example, the green line. The green line is parents. High spending is the highest utility. That's at point A. Lowest spending is the lowest utility. That's at point B. People see that? Now look at the elders. It's the opposite line. Now look at the young couples. It peaks in the middle.

Now go to figure B. The difference of private school parents is it goes down then back up. The red line is different than any of the lines in panel A, which all have what we call a single peak.

So as long as preferences are single peaked, majority voting works. So the restriction you can impose to get around the impossibility-- the impossibility theorem holds. You have to impose a restriction. The restriction is what we call single-peaked preferences, preferences where the local maximum is the global maximum. That's not true in the second panel.

When is that likely or not likely to be true? And The example, Alec, was not that weird. In fact, when-- the situation was likely to be violated are exactly situations when there's a good private alternative to the public good. That's exactly the conditions in which it is likely to get violated, because those are the conditions which people might want one extreme but not the middle.

Generally, people prefer the middle. Or they have linear or they have monotonic rankings. The fact that this is nonmonotonicity, where they like the middle worst. That's because there's this private option.

So Alec is generally right. Generally, it's a pretty safe assumption that preference is single peaked but not always. And the theoretical proposition is important, which is not always means that it will not always work.

So basically, we restrict preferences. If we restrict preferences to be single peaked, we get that majority voting works, but we get more than that. We actually get an even better, more powerful outcome, which is that basically, if preferences are single peaked, you can prove mathematically what we call the median voter theorem.

The median voter theorem states that if preferences are single peaked, the outcome of majority voting were the outcome preferred by the median voter. So you don't even need to know anybody else's preferences as long as I count people up-- if there's 1,001 people, I don't need to know the 1,000 people preferences. I just need to know the one guy who's the median voter, because they're the swing vote. Their preferences determine. So all you need to know is the median voter's preferences. We call them swing voters now. Back in the old day, we called them median voters.

I don't care if there's a million people, 2 million people. As long as I can decide who the median voter is, I just need to know what they want, which is incredibly powerful in terms of reducing our information that we require to understand decisions. That's an unbelievably powerful theorem. It doesn't matter how big society is. I just need to know the median voter's preferences.

That's the good news. The good news is that in a model with single-peak preferences with majority voting, the median voter will always determine the outcome. Yeah.

AUDIENCE: These conditions-- is the thing that these conditions are necessary for to have a consistent top of the ranking or consistent in aggregate.

JON GRUBER: Consistent voting mechanism, that voting delivers the outcome that meets these three conditions. So the good news is we can meet these three conditions with single-peak preference with majority voting. And even better news is always knows who the median voter is and what their preferences are.

But here's the bad news. Is that outcome efficient? Is the median voter outcome efficient? When would it be? And when would it not be? And here, I'll accept a graphical intuition, not that kind of graph but thinking about statistics. When would the median voter-- remember, what's efficiency? Remember the Samuelson efficiency condition. When will the median voter outcome achieve the Samuelson efficiency condition? Or when will it not? There's actually a mathematical answer to this. Yeah.

AUDIENCE: [INAUDIBLE] preference is the same thing.

JON GRUBER: Yeah. And what would that mean in terms of a distribution? It would mean the median equals the mean, because the Samuelson condition is that the sum of values gets set equal to the cost. That means on average, you want the average person's value to be the average cost. In a median voter system that will only be true if the median equals the mean. But if the median doesn't equal the mean, it will no longer be true. In particular, if there's more intensity of preference on one side than another, you will not get the efficient outcome from the median voter theorem.

So for example, imagine your town is considering building a statue to you. You got in MIT. Life's made. You're all set. They don't care what you're going to do. You already deserve a statue. And let's imagine there's 1,001 voters in your town. And the statue will, conveniently for our mathematical purposes, cost \$40,040. And the town is going to finance it by an equal tax on everyone, of \$40, conveniently enough. So a tax of \$40 will finance building the statue to you.

Now, let's say everyone in this town knows you. And let's say that 500 of the people in this town appropriately recognize you are awesome. And each of those 500 people would happily pay up to \$100 to have the statue, for a value of \$50,000. So the Samuelson rule would say the statue should get built. If I aggregate-- assuming the other people's value isn't negative, the statue should get built. Assuming even if everyone else is zero, the statue should get built, because basically, \$50,000 social value exceeds the \$40,000 in costs.

But let's say you have a vote. And the other 501 don't know you, don't like you, whatever. And they vote no. What's going to happen? Will the statue get built or not? No, because the median voter doesn't like you. The median voter values you at less than \$40. So the statue won't get built. So we do not meet the Samuelson condition through the median voter theorem.

So the median voter theorem will only meet the Samuelson condition if we have equal intensity of preferences, or if the median distribution equals the mean. Only in that case will the convenience of the median voter theorem meet the efficiency of the Samuelson condition, which is going to be pretty rare. So let me stop there.