[SQUEAKING]

[RUSTLING]

[CLICKING]

R SCOTT KEMP: So yesterday, or Monday's class, was probably the most boring class, I hope, of the semester, just reviewing really basic things. This may be the most depressing class of the semester. And then I think things will get more interesting from here. So we're going to just finish out calculating the cost of nuclear and do some of this phenomenology that we talked about to try to understand what drives these costs and what is important.

All right. So I guess it's time. So in the last class, we derived this formula, which is just essentially the investment cost at the end of building your plant. This formula applies to any technology. And we saw the sensitivity to construction timelines. And we mentioned that this is one of the reasons why we're interested in trying to move towards factory fabrication, because these long construction timelines for stick-built technology are driving the costs way up, two and a half to three times what it would be if we were able to just build this very rapidly.

And so that's the motivation. There's a secondary trade-off about factory fabrication that we'll talk about in a couple of classes. But now you know where it comes from. But when we look at the overnight costs, we are ignoring this component. When we look at the investment cost, we're including this component.

And these particular numbers are calculated for essentially a 4% real interest rate, which is a little bit low. 8 minus 4. So 4 is the inflation rate. 8 is the interest rate. Typically, it's-- historically, it's thought to be about 5%. But financing has become a little easier in recent years. There's just so much money floating around that I'm going to put it at 4%.

And then what that does is it biases these numbers. It makes them a little bit smaller at very long timescales because you pay less interest on the loan. So the difference between the overnight and the investment cost is smaller.

So we did this calculation and we put the total investment cost here. But I warned you not to get too worried yet, because these are just before-- these are just the amounts of money the loan is at the beginning of building your plant. But it turns out like not all plants are created equal in terms of how long they run. So nuclear plants traditionally have been scheduled to run for about 40 years. And even today, we still have that as the licensing goal.

But in reality, most nuclear plants run for about 60 years, which really makes them the most enduring kind of plant that we really have out there. So if you're a person whose job is to do this officially for the utility, you're going to select n equals 40, because that's the officially licensed design lifetime. And that's how you're going to run your company. And you're going to set it for 40 years.

But we're not interested in what the accountants do. We are interested in what is the right choice for society. So we need the true number. And the true number is something like 60. So we're going to use 60 to convert this into an annual payment.

So how do we do that? You remember that these are our different types of costs that we have to take into consideration. We're only focusing right now on the capital cost. The capital cost is not the investment cost. This is the monthly capital or annual capital charge that we want. And how do we do that? Someone remember from last class? I explained it.

Nobody remembers? It's this equation. Remember these boring equations I threw up here? Here's why we use them. So here is the investment cost that we pay out at the end of the construction period. And then we want to figure out essentially what is our annual charge over the whole lifetime, over 60 periods, that we're paying to pay off our loan? And this is the formula to do that for n periods at a certain interest rate.

So the question I will present to you is, what value for r and what value for n? We've just talked about what value for n. We're going to use 60 for nuclear. There are other numbers for other technologies. What about the value for r? Thoughts?

AUDIENCE: It means the interest rate, isn't it? So--

R SCOTT KEMP: It is the interest rate. But which interest rate is it? Is it the real interest rate or is it the nominal interest rate?

AUDIENCE: It should be the nominal, right?

R SCOTT KEMP: So if you are doing-- if you're an accountant, you would get a loan that is denominated in nominal dollars. You'd have to pay back that loan in nominal dollars. And you would be interested in figuring out what your payment is. Payment A, not P. And you would do that with a nominal interest rate.

But if we are just trying-- one of the issues with this is if you are thinking, as you go forward, you will get more and more money in revenues because of inflation. The nominal dollar amount will go up and up and up. So we're not-- then you have to do a cash flow analysis. And maybe you have to do amortized loans and other things like this. It depends on the financing structure. Blah, blah, blah, blah.

We're not really interested in that. What we're going to do is just treat everything in real dollars. We just want to just say we'll treat all future dollars as having equal value. And we'll use the real interest rate. I'm not being-- I don't think I'm being adequately clear here. Does that make sense to people?

We're just, by using the real interest rate, we can ignore the effect of inflation on future revenues and just figure out what the real cost is in today's dollars and just be done with it. Yes.

AUDIENCE: It's for loans designed keep everything in nominal dollars. I can see how they incorporate inflation. If it's a loan on your house or something, and they're like, all right, we'll take care of all of that for you. When you're talking about billions, do they still work in nominal?

R SCOTT KEMP: They don't do-- banks don't do loans for nuclear plants. It's like bond financing and all kinds of other instruments. So what I'm doing is like, I'm not really getting into the financing. I'm just trying to give you-- these tools are very simple net present value calculations.

And I'm saying it-- there's a way to still use them without understanding the complexity of actual finance to get an answer that is meaningful. And we can and we can do that if we stick to real dollars. That's the goal. But yeah, it's much more complicated, as you're pointing out.

All right? So we're going to use a smaller value here. We're going to use 4. All right, so there's the lifetime of all the plants. Where do these data come from? You should always ask this when you see it, because these numbers have big impacts on the results.

These data all come from the EIA. It's based on their surveys of historical plant life. So these are what they say plants-- because remember by law, all utilities must constantly report what's going on to the EIA. So these are all good numbers except for these two numbers, which I have unilaterally changed from 40 to 60 to reflect what is happening with life extensions and the fact that now we are pushing nuclear plants through this past the 40-year design threshold.

So what that does is this brings down this number a little bit. But I think this is a reasonable thing to do because basically, we've demonstrated this is feasible. So now we have the capital charge for the different technologies. So I guess you can see that this is where you see there are now some technologies dropping behind nuclear-- solar-thermal, fuel cells.

I guess that's kind of it. But you never see these plants being built. And you can now understand why. They just have very high outgoing capital costs.

Let me also highlight natural gas, which is the most popular thing in the United States, which are these three here. This one with carbon capture and storage, these without, are both under $100 per kilowatt per year of capacity. Solar with storage is almost the same. This is four hours of storage. Four hours of storage is not-- four hours of storage is good now. If you wanted to have a fully solar grid, you would probably need to go a little bit higher than four hours of storage, probably close to 12 hours storage. So that's still pretty good.

Wind is not bad. It's not as cheap. Still 162. This is, remember, only the capital charge. Geothermal is more expensive here at 221. Geothermal is interesting, though, because there's actually quite a lot of it in the South of the United States, the Southwest. And it is dispatchable and almost zero carbon, like nuclear. So it is a competitor with nuclear in that sense.

I've pointed out geothermal, and I pointed out wind and solar. And you could argue that this is unfair to point out comparisons at this stage, because we have not taken into account the variable costs yet, or the fixed O&M costs. However, all of these things are basically fuel-free and, therefore, have very low variable costs. And so comparing them at this stage is actually kind of fair.

But yeah, solar is amazingly cheap. As a carbon-free source of energy, it's the cheapest one we have right now, $110. And that explains why China is installing, if you do the math, 1.3 gigawatts of capacity per day. This is a 1 and 1/3 nuclear plants of capacity per day. Capacity does not mean generation. But even if the generation were only, say, 10% or 20%, you're still talking about a couple of nuclear power plants a week being installed by China equivalent.

So this is driven by the insanely low cost of Chinese solar panels and, of course, by their very low cost of labor. But if you're worried about climate change, this gets carbon-free power on the grid fast for cheap. And that's what really matters. Over a period of 10 years, one nuclear construction. But these people can do it in one day. It is not a solution for climate change. And that's really the key takeaway here.

So how is it that any of these plants get built if this is all true? If nuclear is so expensive, why do any nuclear plants get built? Well, it has to do with how reactors are financed. Many, many utility projects are financed through bonds, which are issued in a variety of rates. And they have to pay the money back to the people who gave them, who loaned them money.

But nuclear plants have always been a little different. And that's partly because of their size. It's partly because of the history of where they come from. And then, partly because of a lot of very effective lobbying that goes on from organizations like NEI. So in the United States, there are-- and in South Korea-- there are bond offerings. But they get a lot of risk-free rate loans from the federal government. They call these loan guarantees.

Basically, the federal government says, we will guarantee that, say, $4 billion of your plant, you can take a loan, and we guarantee to whoever loans you that money that you will not default. And so you can charge the risk-free rate, which means a very, very low interest rate. And then as part of that process, if you look at the details, the bank that issues that loan is actually a government-owned. So it's essentially money from the government for free.

I mean, they have to pay it back. But they get to borrow it for free. So that delay in construction that leads to that inflation of cost, that doesn't matter anymore because the multiplier on the overnight cost essentially goes to unity. The other thing that we did for Vogtle is we just sued Westinghouse for $3.7 million. So that was a good chunk of change. And for Vogtle, this is the two new reactors. If you're not familiar, these are the two most recent reactors built in the United States.

They got authorization from the Public Utility Commission to begin billing their customers for nuclear power before 10 years before they started making the power. So they just basically, through ratepayers, billed them for the cost of the plant. So all of this reduced their finance costs dramatically. And that combination of things is allowed Vogtle to be built in the US. But it is not a generalizable strategy. Yeah?

AUDIENCE: You said "we" sued Westinghouse. Who is "we"?

R SCOTT KEMP: The people who built Vogtle, which I can't remember the name of the company. Yeah?

AUDIENCE: Can you talk a little bit more about risk-free rate loans? How often are they given out, and maybe what portion of Vogtle is financed with them?

R SCOTT KEMP: So there've been three or four of these offerings in loss, going back to the Energy Policy Act of 2005. And the rule is that up to 50% of the plant can be financed through these federal loan guarantee programs. So that's kind of the rule.

Initially-- I have another slide, but I can't really remember off the top of my head. But initially, when that policy was announced, there were something like 29 license applications to the NRC. And at the time, we had not built any nuclear plants since the '70s or ordered any plants since the '70s. And no one really knew what it would cost to build a plant.

So these vendor forecasts that I showed you in a previous slide were taken to be somewhat serious. And the size of the loan guarantee program was sized so that it could fully fund something like 15, 14 or 15 nuclear plants, which when you consider that they only pay for half, gives you the 29 licensed applications. So it all may work out. And that's where both the Sumner, VC, Sumner, and the Vogtle plant got started.

Almost everyone-- 20, I guess. The 28 is reactors, not sites. So 24 of the-- or 25 of the 29 plants for which they did a license application to the NRC, went to the point of doing that license application. And then they realized, actually, we're not going to be able to afford this. And the two that went forward were Sumner and Vogtle. And then Sumner also abandoned the project halfway through. I think they spent $7 billion on Sumner before the Public Utility Commission said, despite the fact that we spent $7 billion so far, we still determined that the residual costs or the remaining costs would be so high that it's not in the interest of our customers to continue with this project. So they closed the project.

Only Vogtle survived, of the 29. So that's two of the 29. And it happened because of all of these shenanigans to basically get free money from the public, sue Westinghouse. And they got a second round of loan guarantees. And oh, one thing I didn't mention. This was truly brilliant for Vogtle. They started the project in 2009. Does anyone remember what happened in the economy in 2008?

AUDIENCE: Crash.

R SCOTT KEMP: There was a huge crash, right? So interest rates were really low. And they went on to the bond market, and they issued all these bonds at something like 1%. And then the inflation rate was much higher than 1% So in fact, they made money by issuing bonds, by taking loans, because they had to pay back at 1% but inflation was 3%.

So sometimes you will see a calculation of the overnight cost for Vogtle, and it will be smaller than the number that I show you at $19,000 per kilowatt of capacity. And the reason it's smaller is because they fail to take into account the details of that funding. So Professor [INAUDIBLE] has his own number. And I talked to him about it. He's like, oh, really? Like they actually made money off the bond offering? I go, yeah. You have to add that back in to find the true technical cost.

So that's why that plant is so expensive, but also why it was able to be built. It's a remarkable moment in history where they could do that. OK, other countries, they don't work this way. China, it just comes out every year out of the national budget. That used to be the case for India. India a while ago-- I can't remember how long ago, maybe 10 years ago-- switched so that it's now 33% comes out of the annual budget. And then 60% they take from foreign loans, principally, so issuing bonds on the US bond market.

Japan has these government-backed risk-free loans, essentially loan guarantees but for the full amount of their reactors. And France has this. It's essentially a government-owned utility, right? So most of these people don't have to compete. They don't have to engage in financing of a normal construction project because they're government-owned operations. They just come out of the national budget. Yeah?

AUDIENCE: This is kind of a tangent, but how much does like a fully operational coal plant cost today? Because that should be-- you would expect a nuclear power plant to be that plus, I don't know, a little bit more, because you--

R SCOTT KEMP: Well, OK.

AUDIENCE: Architecture is all going to be the same. And like obviously, you're going to have your containment. But the differences between the boiler construction and whatnot are not terribly different.

R SCOTT KEMP: There's your answer. Top line, 209, $300 to build a coal plant. And for the nuclear plant, I mean, this is the loan payment. So this is the-- if you just want the total investment cost, $11,000 per kilowatt of capacity versus $5,000. It's about twice.

AUDIENCE: OK.

R SCOTT KEMP: Now remember, this is assuming these EIA numbers, which we've seen in the spectrum for the overnight costs. There's a whole range. These numbers are smaller-- I'm sorry, bigger than the vendor estimates, but quite a bit smaller than what we saw for Vogtle. So these are vendor estimates plus the contingency. And it assumes also that, if you're looking at this number, it assumes that it gets built in six years, which hasn't been done in a while.

So part of it is that this number is high because of safety considerations, because of delays, so many things that make it more complex. Part of it is that this number sometimes gets inflated and also drives up the cost. So it's coming from multiple places. And you've got to deal with all of them if you want to figure out a way to make it cheaper. That's kind of the story I'm trying to weave for you today.

All right. So now you can understand why other countries are able to do this. It basically becomes a matter of national policy. We're going to build these plants. And it doesn't really matter that they're expensive, because they don't have to compete in a competitive market. It's a nationalized electric system, and it just costs whatever costs. And it's just a decision to do it.

All right. So this is just a quick-- yes?

AUDIENCE: Can you repeat again, besides lobbying, what are the reasons why we want to build nuclear plants?

R SCOTT KEMP: Why are the reasons who wants to build? We in the United States, or we the people? We in the US?

AUDIENCE: Yeah.

R SCOTT KEMP: It is principally the vendors who sell the plants, the lobbyists who work for them, and some, I would go out on a limb, and say somewhat corrupt public utility commission officials who get paid off, some of who have gone to jail, who get incentivized to do this. But if you are looking at it from a social benefit perspective, if you're the technology-neutral philosopher king, as it were, you would not choose it at these prices. Today, right?

So we'll look a little bit about-- we're going to keep going. There's going to be some dynamics associated with trying to decarbonize the grid that push you towards technologies that are dispatchable and always available, in a way, for wind and solar. And so the next half of this class, what we're going to do is we're going to study that effect and how much that comes back to rescue nuclear to see if-- because you really want dispatchable technology, and wind and solar are not that, right?

So we're going to keep going. There's a lot here. But we have to construct it piece by piece. So let's keep going. All right. I just want you to remember these terms. You'll come across them in your career.

Investment cost is that number at the end. Future value or rate base cost is basically the same thing. Overnight cost is the technology cost without the part caused by the construction delay, the interest paid during construction, IDC-- Interest During Construction. And levelized cost, we haven't talked about yet.

And I think I have a slide, but I'm going to wait a little bit later to get to levelized cost. Levelized cost is kind of, if you want a single number for the cost of electricity, it's way to do that. But there's a lot of embedded assumptions. So you have to be careful about using levelized cost.

OK. So that gets us into part two. How does nuclear compete in the market given this cost structure? All right, so let me just go back here. This is an equation in Q that gives us a straight line, a family of lines. It intersects the axis at the capital plus fixed costs, the y-axis. And then it has a slope related to the variable cost.

All right. So for every technology, we should have some family of lines. And this is such a family. And using real data, this data is a few years old. But I can't regenerate all the plots every year. So we're just going to use it to get a feel for what's happening.

And this is a four-technology system. And the x-axis is in hours per year. And there's 8,760 hours in a year. So this is, if you will, a measure of capacity factor. Does that make sense? What fraction of the time the plant operates. So if it operates for 2,000 hours a year, it's roughly 20% capacity factor, something like that.

So depending on how often we run the plant, we'll see that the lowest cost option technology will change. So it's looking strictly at the investment cost. And it's just to say if we're only looking at where it intersects here on the y-axis, that's not a sufficient explanation. We also need to look at this slope. So in this system, we see that there are-- well, let's just look at the technologies.

Geothermal is purple. I'm sorry. They're very close, but geothermal is purple. And it has no slope, essentially flat. Almost-- it's not actually technically flat. There's some maintenance costs, but it's almost flat. And that's because there's no fuel cost.

Coal and natural gas have fuel costs. Let's just look at this one. This one is natural gas simple, which means you take the natural gas and you put it in and you burn it. And then you run it through your turbine.

This one is combined cycle. So this has a second heat recovery step which takes the steam that comes off, the cooler steam, and then uses it to drive another turbine. So this has two turbines in it and makes more use of the heat coming out of the natural gas. And so its per-unit energy cost is lower, so the slope is shallower. Does that makes sense?

So what is this telling us? Well, it's saying that if we need to build a plant that's going to run most of the time, we probably want to use-- despite the fact that it costs more to build this plant initially, we want to use this technology if we're using it most of the time. But if we have to build some capacity that just comes on during peak hours, like around 5:00 PM in the afternoon as the sun is setting and everyone is going home and turning on the air conditioners in the summer, then the capital investment for this technology is too high. We'll use the cheaper technology that has more expensive fuel costs and we'll have this.

So the optimal mix will always be at least a couple of technologies. You can also have a three-technology system. So imagine that the cost of building a geothermal plant had come down a little bit from R&D. Then what we would see is that would, if we could bring it down to this level, we would have a three-technology system. And for some small amount of capacity that is running all of the time, which we call base load, we would be using geothermal. That would be the optimal base load generator.

But this is not telling us anything about how much of each of these technologies we should build. This is just telling us which technology to build given a certain capacity factor. How can we convert this into a measure of how many plants we need? So there is a way to do this. So the way we do it is we look at a demand curve. So here is the actual demand. We tend to use ERCOT in class, because ERCOT, which is the Texas electrical grid, publishes all their data on the web. And as a result, academics use ERCOT data all the time, and it's very nice. And ERCOT gets a lot of free research as a result of this.

But you would look at your local demand. And this is an actual demand curve for one year going for every hour, L 8,760 hours. We've plotted the actual total demand. And you can see here in the middle of the year where the summer happens in Texas, the demand for electricity goes up. And if you look really carefully, you can see this periodicity. And there's 52 peaks. So that's the weekday-weekend change with the weekday usage being higher and the weekend usage being lower. So that probably weekend weekend weekend weekend and so on. You can see this in the data.

So I can take this demand data. And I could take out the most-- probably this peak here in the winter was probably some kind of weather anomaly, one that does not appear in general. You can go and see. In 2022 in December, was there some kind of weird weather event? I can take this and I can start sorting these by how high they are. So I'll these bars from highest to lowest.

And here, I'll just do the sort. There they are. So this is the same, all the same bars. And I've just sorted them on the graph. And now it's telling me how many hours of the year I'm using 80,000 megawatts and how many hours I'm using 40,000 megawatts. And I can basically, because I resorted the axis, I can now see how often I need a certain amount of capacity. Does that make sense?

All right. So I can now put this plot, which is also a fraction of the year that I need a certain capacity. And I can draw a line's going straight down. And I can see that here is the switch. If I'm using a plant for less than whatever this is, 3,000-ish hours-- I guess I didn't write it down-- then I want the yellow technology. And how much capacity do I need? I need this much capacity to do it. Does that make sense? Which is 30,000 megawatts.

So that's 30,000 megawatts of yellow. And then for the rest, I'm going to have green. And I'll built 50,000 megawatts of green. Does that make sense?

So this is called a screening curve. And this is a way that you can determine how much of each technology you can build. And of course, it depends entirely on the shape of these lines. And I will point out that if I generate this plot and I omit a line, I will get the wrong answer. Like if I forget to put in a certain technology, that matters. I won't get the right answer.

This is a modeling problem that happens all the time. So you'll see people do studies, and they'll choose their technologies. They'll say natural gas. And they'll choose nuclear, and they'll choose something else. And then they omit geothermal. And then they get an answer, and it's not the right answer because if you had included geothermal, the answer would be different, because you have these substitution effects.

So whenever you look at a study about what is the optimal mix, you need to keep this in your mind and saying, OK, what were their modeling assumptions? Is this a complete set of assumptions? And therefore, can I believe these conclusions?

Or have they-- so you can be sophisticated. Like well, they've omitted this particular technology. Let's say it has a very high fuel cost, but a very low capital cost. So it would actually come down here. And we can say, OK, so they've overestimated the amount of yellow. And there's another technology that would actually substitute in there. You can be smart about it if you understand the dynamics of each technology.

Let's see. So I think I have this? Yes. So I've added this to the graph. What is this number, 32,000?

AUDIENCE: It's always used all the time?

R SCOTT KEMP: Oh, it never turns off. And that is what we call base load. You may hear phrases like oh, we need nuclear for base load. Who's heard this? Yep. Do you see any problem with natural gas providing this base load? Is there any reason I can't run my natural gas plant all the time or build 32,000 megawtts of natural gas?

AUDIENCE: Depends what you consider. If it's just a matter of price or cost, you get an answer. You consider, as you said before, you're carbonizing the grid. Or maybe the answer [INAUDIBLE]. It depends.

R SCOTT KEMP: Yeah. But in terms of the question, we need something for base load, base load will simply be the technology which is the lowest on this end of the graph. That's what will be dispatched as base load. There's no special properties of a generator that make it suitable for base load.

However, the inverse is not true. There are properties of generators that make them unsuitable for non-base load applications. So the correct statement is not, we need nuclear for base load. The correct statement is, nuclear is really only suitable for base load. And the reason is that it has a very high x-axis.

Hold on. There's nuclear. The reason is it has a very high intercept over here. Now this is nuclear with a 10-year construction. This is not the EIA-predicted nuclear. But you can see how it kind of operates. It operates similarly to geothermal.

So if we could bring the cost of nuclear down, the capital costs, which is to bring it down with the intersect on this axis, at some point, it would come across and it would pop out underneath the green line here. And then it would be the base load technology. So that's what base load means.

Best hope of financial success is to be operated as a base load generator. But we don't-- there's no requirement that base load must be serviced by something like nuclear. It can be serviced by solar. You just have to have the reliability to be on all the time. All right?

So I guess I'm going to-- I think I'm going to skip a discussion of bidding. Yeah?

AUDIENCE: What does this cost look like for nuclear that's already on the grid? So Texas has, I know, four large-scale nuclear power plants right now. What does this look like?

R SCOTT KEMP: So yeah. So this is-- the question is, how much do we build? And we're taking into account all of the costs. However, as just pointed out, if the plant is already built, the annual payment is a sunk cost. So all that really matters is the slope.

So what happens is the plants look at what they call their long-run marginal cost, and just say like not the instantaneous amount of cost it takes me to burn the next unit of gas or use the next gram of uranium, but that plus the additional maintenance cost that comes from operating the plants. And then they say, OK, it's going to cost me $0.05 to make a kilowatt hour of electricity.

And so then in a market-based system-- and places like Vogtle are not a market-based system. But most of the United States is a market based system. Each power plant will make a bid. And there's a day-ahead market. And there's an hour-ahead market. There's various markets that people can bid in.

And they will say, OK, I will make electricity at, and they quote their long-run marginal cost in principle, $0.05 a kilowatt hour. And then what the grid operator does is they say, OK, well, we need-- they have all these models. Let me just show you. So they have all these models for what they think the amount of electricity demanded will be. So maybe they're here. And OK, we're going to need to have a little bit more power. And we think it's going to be this. And we'll have-- there's also a market for generators that are on standby, and all this to deal with anomalies.

But basically, they say, we can think we're going to need this much power. And they start going through all the bids, and they take the cheapest bid, and then the next most expensive bid, and the next more expensive bid, until they get enough power. And then they pay everyone the price of the highest, most expensive bid that was necessary to get enough electricity. And that's how it works.

So when the demand is high, the generators that can make that demand might be peakers. They may be very expensive, maybe a couple of dollars. I have whole charts on this, but I put them out. They may be several dollars per megawatt, something really, really expensive. And for an hour, everyone will get this amazing amount of revenues of a dollar. And then if the demand goes down at the end of the day and we're going into the weekend, now everyone is getting like pittance, like $0.01 per kilowatt hour.

So the nuclear plants, they never want to shut down, because shutting down a nuclear plant is a very time-consuming process. So they will basically bid 0. They will say, we'll make electricity at 0. The actual cost is something close to 0. And they'll make sure they get the batch. And they'll take whatever price they get. So if there's a lot of demand, they'll make $1.00 per kilowatt one day. And then when there's no demand, they'll make no money.

And the hope is that they just make as much money as they possibly can to pay their bills. They may not be able to cover their bills. And so eventually the plant may shut down. But that's the goal, is to try to keep the plant ticking along by being in the market. So that's why these plants-- I won't say they're able to compete, but that's why they can sustain themselves, because their marginal cost, which goes back all the way to that slide, the NEI slide I showed you, is actually very low.

It's certainly in a couple of cents per kilowatt hour, depending on the age of the plant. So all the plants have higher maintenance costs. And so what you'll see is older plants starting to shut down early because their maintenance costs are creeping up, and they basically can't cover their long-run marginal costs. Did that make sense to people? Does anyone have questions about bidding? Yeah.

AUDIENCE: Would you not need to include the interest payment or the investment costs on top of the long-run marginal cost in order to break even?

R SCOTT KEMP: So it depends if the utility wants to go bankrupt.

[LAUGHTER]

Yeah. If they're willing to go bankrupt, then they would include that and say, OK, it's not interesting, and they go bankrupt. But otherwise, they're going to eat that cost and spread it around their other operations. And then they would just say, to run this plant, which I already own, what is the incremental cost of running it versus incremental amount of money I can make?

And you might ask, why don't people bid higher? They say, well, instead of bidding $0, I'll bid $0.50, even though my true cost is $0.30. And the answer is, well, then you might not be dispatched. It turns out that the best solution is always to bid your true marginal cost. And that's the dynamic of margin. This is the miracle of the Adam Smith's invisible hand in the market, that drives the cost to the true marginal cost. Yeah?

AUDIENCE: How does nuclear ever break even? Or they just--

R SCOTT KEMP: How does what?

AUDIENCE: Nuclear ever break even? Or they just don't?

R SCOTT KEMP: Well, I think we'll actually kind of answer that in a moment. But here's the preview. It wasn't always so expensive.

So here's the cost in real dollars. And you can see early on the costs were low. And it would just be slowly rising. So historical nuclear is both so old that they've paid off their loans, and the plants are basically just now being run. Remember, the loans are based on 40 years, and the plants go past 40 years, or sometimes even less than that. And so their costs are now just maintenance and operating costs.

So it's a very different situation. An existing nuclear plant is very different economics than the choice to build a new nuclear plant. Anyway, we'll get to that in a moment.

All right. So now you know how to dispatch different amounts of material or different numbers of plants. Let me just-- well, since we're talking about bidding, let me just mention one thing. Sometimes prices will for electricity will go negative because a bunch of generators on the grid, and demand drops. And it's a market signal to get your generators off the grid now or we're charging you. Otherwise, you're destabilizing this grid.

And sometimes you hear people talk about negative electricity prices like it's a problem. And it's not. It's a market signal for the next hour that you better get your generator off the grid. It's not negative all the time. It's negative for an hour or two hours, and then it becomes positive again. And on average, they make money.

But a lot of made of this because solar has created a situation where their marginal cost is $0. It cost them nothing to make electricity. If there's power coming out of the solar cells, it costs them nothing. So they bid $0. They always get dispatched. And in order to shut down those solar farms, the only way you get them to curtail is you start billing them so the price goes negative. So then you make less money.

And this reduces the profits for everyone, right? And so some people who are very expensive generators don't like it because the average cost has gone down, because occasionally there are hours of negative pricing. And then they complain, oh, the electricity is negative. Well, the zero point is kind of arbitrary. All that really matters is that at the end of the year, you're still making some profits. It's not negative on average.

So you're making less profit. And because you were an expensive technology and a cheaper technology has come around, yeah, you go out of business. And that's good for society because we want to make the most electricity we can for the cheapest amount of money. So don't get caught up in the whole negative electricity price thing. You see it occasionally from people who don't understand how markets work. Yeah?

AUDIENCE: Is it negative for big industrial users of power, too? As in if you're operating a cement plant or something and you need a lot of heating power and the price goes negative, can you make a lot of money by suddenly shooting up demand?

R SCOTT KEMP: You could. Yeah, if you were a consumer, you could-- yeah, so this is the business model for batteries. You build a battery plant. And then you wait for the price to go really low or negative. And then you make money taking power off the grid, which helps the grid operator stabilize the grid. And then when the price is positive, you sell it back.

And that's a good thing. And that market incentive to create that business will both stabilize the grid and make you money. And that's what it's supposed to do. So yeah, there's indeed a way to make money there. Yeah?

AUDIENCE: Isn't the main thing to be said about negative pricing, just the price volatility? Because equally well, in this whole picture in general, marginal price setting only works if you let the market fully do its thing. But we don't do that in the positive sense either. If there's a price peak that's massive, it gets capped. So peaker plants get in trouble. So you need capacity markets. Blah, blah, blah, blah, blah.

So isn't that the whole argument, that you just increase volatility, and that negative prices in and of themselves, as you mentioned, are not bad? I feel like typically the argument is made like, look how much volatility is increasing and the negative prices are just a symptom of that?

R SCOTT KEMP: Correct. Yeah. So that's correct. So volatility is an issue if the generators can't respond to volatility. So yeah, that's absolutely correct that for traditional thermal generators, a highly volatile market is problematic. For inverter-based generators, a highly volatile market is not problematic because they can instantaneously curtail within a millisecond.

So people who are thermal generators or inertial generators complain about volatility. And those who are not inertial generators-- but there's also benefit to having inertial generators, because sometimes you have unpredictable fluctuations in demand, and inertial generators help prevent, essentially, voltage drops.

So there's a lot-- as you pointed out, there's a lot of devil in the details. And there's a lot of additional what they call ancillary services that, in addition to just providing kilowatts on the power grid, you also want to stabilize the grid. You want to have high responsivity. All of these things are called ancillary services. And there are many markets emerging for all of them as the types and diversity of generators go up.

In principle, these things are small perturbations as to the question that we're addressing in this class, which is how do we decarbonize society? Those are fixable through various small tweaks.

OK, I mentioned levelized cost. And I didn't-- so now I'm going to finally tell you what it is. It's a reasonable cost to just talk about where things are. It's a reasonable metric to talk about where technologies are on the margin. Like today, and if we're talking about adding the next plant so that we're not fully reconfiguring the generation system, and we're just looking at, what is the-- this plant is shutting down because of end of life. I need to look at a new plant. What are my options? LCOE giving you a good measure of that.

But to do it, it embeds an assumption about how often that plan will operate. So basically what it is-- oh, come on. I mean push this button. There we go.

It's essentially discounted the net present value of all investment costs, maintenance costs, and fuel costs discounted going into the future, plus-- there's the discount rate-- divided by all the energy being produced by that plant. So it's a net present value of all of, essentially, the electricity over the whole life of the plant. And this energy out term, the denominator, embeds an assumption about how often that plant operates. It's like every period, there's going to be some amount of generation being produced. So it embeds that inside of it.

Sometimes you see this discounted. They actually discount the energy itself. That is a formulation rooted in investment analysis. And it assumes that the value of energy can be moved around in time, which is like other products. You could consume it today, or you could consume it tomorrow. That's not a really good assumption for electricity because electricity demand is highly inelastic.

So people need the power when they need the power. So it's generally better to use the LCO formulation that does not apply this discount factor to the denominator. But you'll see both of them in the literature. So you can just give a sense-- yeah?

AUDIENCE: Wouldn't, then, the assumption be of discounting your energy that essentially your revenue that you are going to get from it in like a project finance model would be discounted itself? So wouldn't that be where it's coming from? Or--

R SCOTT KEMP: I think in a-- I guess it would depend on what you're using for the discount rate also, then. The way I think about it as a policy person is I use a real interest rate here. Revenues are constant. They don't get discounted. And it just-- right. So, yeah. Yeah?

AUDIENCE: So then, the fuel costs are variable costs and difficult to predict. So what is normally used in this case? I mean, unless you hedge and you pursue or agree long-term contracts, you cannot be sure of the natural gas or oil price.

R SCOTT KEMP: Yeah.

AUDIENCE: Even for the next six months.

R SCOTT KEMP: Yep. Yeah. So how do they deal with that?

AUDIENCE: Yeah.

R SCOTT KEMP: I think the answer is I don't have a good answer. Yeah. This is not something I've looked into. Yeah.

OK. There's the other version that comes out in the finance world. But after spending a lot of time trying to understand the relative benefits for most electricity decision making, I would recommend the top version of this equation. So let me just show you what it means.

So here's our more sophisticated screening curves. And LCOE is basically taking this and saying, OK, we're going to assume this plant operates 10% of the time. And we're going to assume this plant operates 24% of the time. We're going to assume this plant operates 28% of the time. And so then they make these assumptions, and then they calculate the LCOE, the single cost for electricity.

So they assume nuclear plants run 90% of the time, and then you get-- which is about what they run. And so you can see these numbers tend not to change very much. It's really where big changes occur is down here in the peaker domain, people who are being dispatched just very occasionally. And they will have bigger fluctuations in their capacity factors.

So LCOE is not a good metric if you're looking at single turbine, low-efficiency, very-cheap-to-build peaker technology. It is a reasonably good thing if you're trying to compare base load natural gas to base load geothermal to base load coal to base load nuclear. It's generally a reasonable thing to compare. Does that make sense? Yeah? All right.

All right. So I promised you we would do a sensitivity study to the social cost of carbon. And so far, we have totally ignored the fact that natural gas emits carbon. So in practice, what does this do? Well, if we added a carbon tax, the natural gas line would go like that, right? It would get steeper. The cost to society is really steeper than the internalized cost.

So let me show you what this looks like. So I wrote up this little code. (HUMS) All right, here it is. So again, these are slightly old inputs for the technologies.

But you can see this is basically what we were looking at. They don't match the one in the previous slide. Here's geothermal, has come down in price. So in this case, geothermal is the base load technology of choice. Then natural gas is-- let's see, natural gas NCCS. So note, this is natural gas without carbon capture and storage, natural gas with carbon capture and storage.

Nuclear empirical is basically the average construction time, not Vogtle, but the average construction time for recent nuclear plants in Western countries. And coal plus carbon capture and storage up there. So the only CO2-emitting option here is this natural gas without carbon capture and storage. So these are things that I think are reasonably in play.

And you see the importance of geothermal is really an amazing technology. The Department of Energy is continuously, every year, issuing reports that say we should build more geothermal. And there's a lot of potential legal issues in the way, and also transmission issues, because you can't build geothermal anywhere. It has to be built in the Southwest. So you would need long-distance transmission to bring it to other parts of the country. But I think building big data centers in the Southwest would be a good idea.

All right. So we see here we have a two-technology system, right? Geothermal is base load, and natural gas without carbon capture and storage as our peaker technology. Does that make sense? Everyone can interpret that?

And let's just increase-- the social cost of carbon is currently zero. So let's just start bumping that up and see what happens. So here we go. You can see what's happened to the gas with no CCS. And you can see actually the geothermal is going up, too, because it turns out there's a little bit of emissions from geothermal. Not a lot, but it's not zero.

All right. So there it is. And so far, the amount you build of each one isn't changing, but the technologies are still the same, too. And here, we have entered into the regime of a three-technology system. And now we have geothermal still as a base load. We have a little bit of natural gas with carbon capture and storage. And then we have-- just for peaker plants, the natural gas without carbon capture and storage. Yep?

AUDIENCE: Where does the carbon emissions from geothermal? What's the mechanism?

R SCOTT KEMP: Yeah. I'm not actually 100% sure about that. This would be a great-- geothermal, I think, would make-- for anyone interested, I think it would make a great final paper topic. Anytime I say I don't know, that's a final paper topic. Yeah, I've never really looked into it in depth. I just trust the data that is published.

It used to be we would do this in a homework assignment. You'd go look up the data. But I took this homework assignment out.

All right. So we got ourselves to a three-technology system and a social cost of carbon of $335. And I believe it's per kilogram of CO_2, I believe. The units, I forgot. And then we can keep going.

And now, nuclear shows up. This is nuclear if you can build it as fast as Westinghouse says it can be built at Westinghouse prices, not their original price prediction back in 2002 of $1,000 [INAUDIBLE], but their current price prediction. And yeah, so we get quite a lot of nuclear as a base load showing up. And so social cost of carbon, to have some amount of nuclear, we see it is somewhere around $500 that it begins to move over.

And then as we keep upping the social cost of carbon-- let's get to $1,000, somewhere like there-- we see quite a lot of nuclear, essentially an entire nuclear, with just some, I think that's natural gas with CCS, for most of the peaker, and then just a tiny amount of carbon-emitting natural gas for like a very short occasional peaking during Super Bowl Sunday, right?

So it totally depends. This is the phenomenology that I was talking about two classes ago. We want to do the sensitivity study to the social cost of carbon to understand what's going on. Now the key observation here is nuclear has a role, potentially a very large role. But you have to build it at this price, not this price. That's the key thing.

So we still have to bring down the cost of nuclear in order to compete. Yeah?

AUDIENCE: Just for regions without geothermal potential, what is the minimal SEC for nuclear?

R SCOTT KEMP: It's already-- at this point, it's already more interesting. So let's just go back.

AUDIENCE: Because in the Northeast, for instance, at least unless there's extra resource that I'm not aware of, the there's not a lot of potential for geothermal.

R SCOTT KEMP: So let's see. Here, if we take geothermal out, nuclear still dominant. You have to keep going. So in this model, actually, if you can build nuclear fast, nuclear would be dispatched. Yeah. If you can build it at this price. If you can build it at this price, it would not be. It would still be better to do natural gas.

Someone might have noticed there's a couple of generators that are not on this slide.

AUDIENCE: Wind. Solar.

R SCOTT KEMP: Wind and solar.

AUDIENCE: Hydropower.

R SCOTT KEMP: Hydro. Yeah. Hydro could, in principle, be on the slide. Hydro is like geothermal in that it's kind of regionally limited. But it's important. And there are places that have a lot of access to hydro. And if we had a lot of long-distance transmission, we could deal with hydro. But we'll have a class about grid issues.

Does anyone want to guess why I omitted wind and solar? It's not just to make you feel better about nuclear. Yeah?

AUDIENCE: If you take it off, the demand has met demand because you can't-- they always have zero costs anyways. So you essentially treat the remaining demand as a demand to be covered by these ones.

R SCOTT KEMP: You're close. You're very close. Yeah?

AUDIENCE: They can't generate at 100% capacity?

R SCOTT KEMP: No, you can if you just build enough. Yeah. You can't choose when to have your wind power. This calculation, if you all remember, I went back and I sorted my demand curve. And I said, OK, I'm going to build some plant for this part of the demand, and some plant for this part of demand, some plant for this part of demand. Well, that assumes I can run that plant when that demand happens. And I can't necessarily do that with wind and solar.

It's going to give me energy. And I might have demand. And they may or may not match up. So I can't use this methodology for generators that are not dispatchable. And wind and solar are not dispatchable. So how do we do it?

So we use GenX. Who has heard of GenX? All right. A couple of people. So GenX was a code written as a final project paper for this class. And then it turned into a master's thesis, and then also a PhD of somebody else. And now there's a whole professor who got tenure on the basis of GenX.

Let me just show you what GenX does. GenX is a linear programming optimization code. And what do I mean by that? Well, it means that I can define a functional constraint. So this is a functional constraint that says, for every four units of x, I have to have at least three units of y. That's just a function. I can write all kinds of arbitrary functions. And it gives me of family of lines. All these lines are solutions to that function.

And then I have some feasible region. And I want to optimize or maximize or minimize something in this feasible region to come up with a solution. That's the basic idea. So what we do is we write functional constraints. For every generator, the cost of their electricity, how fast you can turn on the plant, how long the plant has to be off if you turn off the plant, because nuclear plants can't just turn on and off every second.

You write all of these functional constraints and you put it into this model. And then you have this feasible region, which is defined by what is available for fuel, wind, sun, et cetera. And you do a cost minimization procedure to find out what is the minimum cost system. So this is higher dimensional concept. We have these planes. But the real answer is that this thing is optimizing over hundreds of thousands of dimensions. It's a very, very high-dimensional optimizer. And that's what it's doing to try to find the best outcome.

Here are the assumptions that go into GenX. You should always know assumptions before you look at any model. They assume that each technology presents a specific dispatch profile that must be obeyed. That's part of the set of functions. Minimum load, maximum ramp rate, minimum cycle time-- if you turn off the plant, how long does it take to turn back on? That's cycle time.

Minimum load-- you can't run a nuclear plant at 10% capacity. You have to run it at 30% or higher. Non-dispatchable-- and that's true for a lot of technologies. Non-dispatchable technologies are only available when the resources are available. You can only use solar if the sun is shining, and so on.

Storage technologies can be added, but at additional cost. Each utilized component contributes some amount of electricity and CO2, which we track. You minimize the overall annual cost, which is the same thing as maximizing social welfare, by adjusting the generation mix to fulfill the energy demand at all times, subject to an imposed carbon constraint. So this is all kinds of solutions. And as I impose a carbon constraint, the solutions that gets more and more expensive, the cheapest solution is more and more expensive as we do the decarbonization.

We don't have forecast errors, essential reliability services mentioned earlier, dealing with volatility, things like that. They're not incorporated. And we assume that demand essentially is unchanged over time. So the demand that we're going to stick into the model, we're going to stick in one year of demand. We assume that the demand basically looks the same next year and the year after that. And that's obviously not totally true. So you could generate synthetic demand and try to rerun the model.

So these are the inputs. This is the original demand for a certain region over the year. This is just a zoom-up. You can see the hour of the year. You can see the day night day night day night observation. In fact, earlier it was mentioned that volatility is an issue. Specifically, the thing that people are worried about is the steepness of this slope, that as you're running all this demand during the day, and then at the end of the night, you go-- or steepness of these slopes, as you [INAUDIBLE] up demand. So that's where the issues are.

We have wind data, real historical wind data for a given region. There's just a zoom-up just to give an example of what it looks like. We have solar data and the example, what it looks like. So it's pretty predictable. Middle of the day, sun shines.

So this is the model that was used by the Future of Nuclear Energy in a Carbon-Constrained World report. This is a study that was run by Professor Jacopo Buongiorno in our department. And they have a certain set of assumptions. You can decide whether those assumptions are fair assumptions or not. My goal is to teach you about the assumptions. And I will show you their results. These are not models that I ran. These are models that they ran for their report.

I'm also going to show you the complete model results as opposed to just select cuts in the model results, so you can see the phenomenology. That's what I wanted you to take from this class, the phenomenology of what happens as carbon becomes more problematic.

So here are the assumptions that go in. I mentioned earlier that the cost of nuclear has been going up. This is in the United States. This is the cost of building nuclear plants. And there's Vogtle at the very end. Actually, the Vogtle number is a little low by $1,000, because when I generated this plot, the final accounting was not in. So it's a little bit higher than it shows.

And then here's the construction duration for historical plants built in the US. Here's Vogtle. You can see it's actually in line with traditional durations. So in fact, the cost of Vogtle-- the very, very high cost of Vogtle was not because of construction delays. The construction delays are basically no worse than they've ever been. Yep?

AUDIENCE: What plant was built in like the 1995-ish range that had like below-average costs by the looks of it, or roughly standard with the lower grouping there, but somehow was built over the course of 25 years?

R SCOTT KEMP: That's a very good question. I'll go look it up.

AUDIENCE: OK.

R SCOTT KEMP: I mean, there's [INAUDIBLE] 2, which is like, has a 30-year construction time. And it was-- and that may be [INAUDIBLE] which was a government-owned plant and did not have, was not subject to financing issues.

OK, so this just gives you the historical view of nuclear power in the United States. And here's what MIT did. They said, we're going to pick a high, middle, and low cost. And this is where-- just for your information-- this is where they come out on the dispersion. So you can judge for yourself whether you think the low cost is a reasonable number, if you think the high cost is a reasonable number. If you think the high cost should be closer to where the Vogtle number is, if you think this trend is real or not real, my magic ball is as good as yours.

So you, you just pick a number that you think is the most fair cost and keep that number in your head, all right? Although I would say take a note at where-- how low is low? Just so you know, because we're all in it, that's going to show up in the results. But you pick a number. Here's what MIT picks, or Jacopo picked for his construction time, way down here, lower than essentially every one. We know what that does, right? When you pick a very low construction time, you get very low investment costs and very low capital return, right?

So I would say they're being a little dodgy here. Real numbers are, like, here. But they only ran the model for one. This was a hidden parameter they didn't reveal. But they only chose one. So this is what we've got to go with, right? OK, so you've all chosen your favorite cost number. You ready for the results?

All right. So there's the results for low cost, for mid cost, and for high cost. Across the x-axis is the percent decarbonization, assuming that we start from a period where everything is essentially natural gas. And this shows us-- this dotted line shows us where we are right now in our efforts to decarbonize the electric grid. And we don't know how far to go. We know the optimal amount is not zero. So it's not 100% decarbonization. It's included at earlier in this class. But whether it's like right there, or whether it's right there, or whether it's here, we don't know.

So you want to proceed somehow along this curve, building technologies as you go, until we decide through whatever process, political process, that we've had enough decarbonization. OK, so this shows you the fraction of total generation that should be partitioned by different mixes. So you'll see here natural gas is dominant if you don't care about carbon.

But actually, some wind and solar always gets built because it turns out even though it's not dispatchable, a little bit of solar that turns on in the middle of the day when the peak demand is highest is always going to be cheaper than paying for natural gas to do that. So even in a no-decarbonization scenario, you have some wind and solar.

And then, as you decarbonize, various low-carbon technologies begin to grow into the system. So CCS is natural gas with carbon capture and storage. Pure battery storage on the grid is shown here in gray. And nuclear is shown here in purple. And depending on your cost assumptions about nuclear, it either shows up really early and you just build more and more nuclear, or you only build it after you reach 90% decarbonization. And even then, you may only build like 20% of your generation capacity of nuclear, which is-- we already have 20% of our generation capacity as nuclear today.

So this is the result of the MIT study, assuming these plants can be built really fast, remember? And these are the prices of nuclear. So you pick your price, and this tells you whether nuclear is a good idea and when it's a good idea or not. And you don't have to-- this is not-- I didn't do this. This is Jacopo Buongiorno, so you can hopefully trust him that this is somewhat favorable towards nuclear.

This is in the northern United States, New England, where there are very poor quality or poorer quality renewables. There's not a lot. It's high latitude. It's not a good place to build wind and solar. Here is the same simulation if we do it in the southern half of the United States. You'll see wind and solar just expands, because wind and solar makes a lot more sense there. And nuclear becomes, really, very tiny.

So the story is that nuclear makes sense if it is cheap or if you need 100% decarbonization or something close to that. And we don't know how much we need. But both of these things must be true, basically. Yeah?

AUDIENCE: And then the qualified renewables doesn't include hydro, in this case?

R SCOTT KEMP: It does not include hydro. Yeah. So thank you for saying that. This does not include hydro. This does not include geothermal. So the only dispatchable sources of low-carbon electricity here are CCS and nuclear. And you'll see that-- well, OK. Yeah. So those are the two dispatchable technology.

So you can say, well, what would happen if we put hydro? I mentioned this earlier. You get the wrong answers if you leave out a technology. So that depends on whether you think hydro or geothermal, that depends on whether you think those are reasonable technologies that we could build. I would say we probably have close to as much hydro as we can build, unless we start bringing down hydro from Canada, which is totally feasible, because they have excess hydro. But that requires long-distance grid transmission.

Or we could build a lot of geothermal, which also requires long-distance grid transmission. But those were excluded from the model. [INAUDIBLE]. Yep?

AUDIENCE: Do you know how long the natural gas supply for the world was supposed to last, how many years [INAUDIBLE]?

R SCOTT KEMP: I have not actually looked at that. So we'll talk about all the wind and solar. Could you really have enough wind and solar and geothermal to do it? But I don't have data for natural gas.

AUDIENCE: I believe the estimate in Sustainable Energy was about 50 years. But he talked about how he keeps growing as we keep finding more.

R SCOTT KEMP: Yeah, this is the problem. 50 years ago they said there was only 50 years left.

AUDIENCE: I'd say about 100.

R SCOTT KEMP: Yeah, it's--

AUDIENCE: Seems that actually [INAUDIBLE] the time that would take might make enough gas distribution.

AUDIENCE: But then [INAUDIBLE].

R SCOTT KEMP: Yeah, that is if, in fact, natural gas really is constrained in that way, that would be a particularly interesting thing. So what would you see? We'd have to take it out of the model and rerun it. But you can think about it.

Probably, this is the model saying, I don't care about carbon, I care about price. So then you say, OK, if you don't care about carbon, you care about price, what's the next best thing? Probably wind and solar. But it's a little bit dodgy because it's still subject to availability constraints and things like that. Yeah.

All right. That's the end. So I promised you today would be the most depressing day. This is the depressing slide. But we'll talk more about other things that might change this.