[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: All right, let's get started. So at the end of last week, we did this final calculation, if you remember, on solar and wind. And the idea was to deal with the variability from weather by having storage and continental scale transmission.

And we used this paper by Sheiner to estimate how much we would need. And I forgot to mention that the reason we used this paper and not other papers is because this one uses years of weather data. So you get actual measures of reliability, not, like, one-year snapshot kind of stuff.

So this gets you to 99.9% grid reliability, according to 43 years of fluke weather. And we came up with this number of $131 per megawatt hour. This is not the optimal system. We talked about that.

This is just all renewables. It's not an optimal mix. We, ideally, would have other things, also.

But this is like an upper bound. If we can't do better than $131, we're wasting our time. So we can use this as a basis for evaluating other technologies.

Anyone know, off the top of their head, what the current wholesale price of electricity-- these are wholesale prices-- wholesale price of electricity in Boston?

AUDIENCE: Like $80.

PROFESSOR: It's about $50. Yeah. So this is, like, 2 1/4 times more expensive. Yeah. So that just gives you a sense of how much more this version of carbon-free energy would cost.

So we'll call this wind and solar. This is like a litmus test. You have to do better than that. Does anyone remember what the price of-- this, by the way, is per unit energy, megawatt hours. We never actually talked about price per unit of energy for nuclear, I don't think. I just show you some numbers. But we didn't spend a lot of time talking about it.

The other way that we talk about nuclear, we tend to talk about capital cost, just buying a reactor. Anyone remember what the current EIA, the Department of Energy Information, Energy Information Administration, prediction is for the cost per kilowatt of capacity for [INAUDIBLE] in your head. Nobody remembers. Yes?

AUDIENCE: $7,500.

PROFESSOR: Yes, $7,700 per kilowatt. This is for AP1000. This is just the capital to build it. But because nuclear is mostly the capital cost, that's a reasonable number to have. So we have to do better than this, in terms of capital cost-- or not nuclear-- AP1000.

All right, so these are our thresholds. This is what we have to beat if we're going to bring down the cost of nuclear energy. Sound good? So we're going to try to do that. And we're going to talk about microreactors and their potential to do this.

All right, so microreactors are all the rage. And if you work in fission, there's a good chance you're working on a microreactor. If you want to get a job in a startup, almost certainly, you'll be working on a small reactor or a microreactor.

So what is actually new with these microreactors? Well, it may seem obvious on the surface. But I think we really need to pick it apart to try to understand where we're going to win here. Because just different is not necessarily better.

So here's a history of reactors built in the United States, actually all over the world. Most of these are in the United States. And the first takeaway is that there's nothing magical about the mere fact that they are small or micro.

So we call this a SMR region, which is below a third of a full-scale plant. We call it SMR, a Small Modular Reactor, or a small reactor. And then if you go down here to 150mw thermal, this is electric. This is thermal. So the conversion factor is a factor of 3.

So you get down here. And then we might call these the microreactors. And of course, all the major nuclear vendors, the French, the Koreans, the Japanese, the Russians, they're all selling reactors that are mega reactors.

But all the startups are trying to go down here, where we've already been. And the question is, why did we abandon these small reactors? And the answer was, we already know, because we found out that they weren't capital efficient.

We did this regression with 342 reactors. We got this result that said that, for every doubling of reactor size, you bring down the cost by 20%. And then we did a little calculation on the board and said, well, what if we want to go from AP1000 size reactor at $7,700 per kilowatt of capacity?

And we did this calculation. We said, there's your 20% cost reduction. How many log2s? So you put your power in there. And for new scale, a single module-- new scale is sold as-- and we'll talk about in a moment-- is a multiple module thing.

But single module, we expected to cost 2.35 times more on a per kilowatt of capacity basis. So there's no a priori reason why going small would fix this. Now we're going to talk about other things going on in a moment.

But we should have this scaling in our head. Now, this scaling was devised from historical data for PWRs or LWRs, Light Water Reactors. So if we go to really small sizes or we go to reactors that are not LWRs, this scaling relationship arguably begins to break down.

So you couldn't really use this formula to predict the cost for many of the tiny reactors currently being talked about. However, the fundamental mechanism that drives this scaling, which is a geometric volume surface area kind of relationship, where there are things that don't scale with power, is still going to be present as you go to smaller reactors. So we expect this relationship, this general trend, to hold but maybe not this exact formula to hold for other technologies.

All right, so one thing is they're small. That doesn't seem to really help. It actually hurts.

The second thing is they're modular. Does modular help? Modular in and of itself is not very interesting.

AP1000 is also a modular reactor. These are two modules from AP1000 being installed. They were built offsite in a factory by Chicago Bridge and Iron. They're supposed to help solve the problems of trying to build all these things in the field.

In the end, it didn't really save very much money. It was actually counterproductive, because things didn't fit together, and they had to rework a lot of stuff. But modules just by themselves are not the full story.

So we got to look for something more. And so these are the two things that are left on the table. Either the technology is just so fundamentally different in how it is built that it really is cheaper than LWR, light water reactor, because of some big change. And there are some big changes in the technologies being discussed.

For example, a lot of the current proposals don't use containments. There's a big cost change. So other big changes, we could look at.

The other thing is that small and modularity are actually hinting at something else, some fundamentally different approach to fabricating reactors. And so we're going to look at what it means to actually do factory fabrication, and try to learn what we can. It's a tough problem, because there's only one reactor that's ever been factory fabricated in its totality in history. And those are Naval reactors-- and very little data about those reactors.

So let's just cram through number 1. What is different about the ideas being proposed today? And just see if we can pick out winners.

All right, so let's just start with the things which are not different. Let's start with the NuScale concept. Why am I starting here? It's actually a little PWR.

The reason I'm starting here is because it's had 20 years of R&D. So of all the concepts, this is the most mature concept. The design involves six small 77-megawatt modules. They live under water in this giant bathtub.

And the idea is that, if there's a problem, it's very hard to have a loss of coolant because there's so much coolant everywhere. And that makes them much more safe.

Now, originally, these modules were designed to be 50mw thermal. But in the end, NuScale pushed them up to 77. Who can guess why?

AUDIENCE: Is it the cost?

PROFESSOR: Economics, yeah. They couldn't make it work with 50s. So even here, that cost relationship is holding. So a key part of making NuScale economically viable is to put six of these reactors in a single building and share a single containment, a single swimming pool, a single steam turbine, a single site, single licensing.

So actually, at the end of the day, this reactor is a full-sized reactor. It just happens to have the nuclear part broken into six smaller pieces. And that's the only way they can get the economics to really work. It's clever. But to call it as SMR is a bit of a deception.

There are other versions that have not done this. So this is the Holtec. This is a one-third to one-sixth scale. It's just a traditional PWR, just built smaller. They haven't really done anything with it.

This is as far as it's gotten as far as I can tell. It's a nice graphic-- not even a really nice graphic. Here's a much better graphic. If you are a startup, you have to get your graphics right.

This is what we learned from Oklo. You know Oklo Graphics. They are built in an A frame in the middle of the floors with no parking lot, no security whatsoever, much smaller than any building that would actually be needed, no administrative building, no power lines, no transformers. But they look beautiful.

So this is Last Energy. It's a new-scale-like concept. Instead of 77mw, they're pushing the size down even smaller to 20mw. So we expect the cost to go up. It's is about 3.6 times more expensive, instead of 2.3 times more expensive than AP1000.

They scrapped the concrete containment, which is expensive, and replaced it with a 12-inch-thick solid steel vault, which is also expensive. And so who knows which is cheaper? There's no water cooling. The thing is buried underground. It's cooled entirely by air.

So on a per-kilowatt basis, to have all these fans cooling the reactor, maybe they tried to price up more. This is a concept that moves towards more safety, smaller reactor, better containment, fewer water issues and at higher costs. So the customer here is one who doesn't care about cost. They really only care about safety.

Whether such customers exist remains to be seen. We can talk about that. Because there are a lot of data center customers who are signing contracts with some of these startups. Though, if you look into the details of those contracts, those contracts basically say, sell us electricity at regular prices. And then those companies can go buy electricity off the grid and then resell it to Amazon or whoever and fulfill their contract.

But in the meantime, they have a deal with Amazon, and they can get better funding. So we have to be careful about whether people are really interested in paying more. But maybe there will be applications or people will pay more than the grid. OK, so beautiful renderings, probably very expensive, not [INAUDIBLE].

Here's another concept, if you can see it-- not really. This is from Deep Fission. They've gotten even smaller, from 77 for NuScale to 20 for the previous one. This is 15. They have thin lines here that you can't see. They're buried, these reactors, one mile underground.

So there's another 50% cost increase because of the smaller size. Here, the hope is they can get rid of the containment by burying it underground. Some problems with this-- one is that it's not actually isolated unless you backfill this whole surface and trap your reactor at the bottom of the hole, which they don't plan to do.

Second is that drilling holes in the ground is much more complicated than you might know. So here's a photograph of a deep borehole. And you can see these weird ovals. And this happens because there's residual stress in the rock.

And as you relieve the stress by creating a hole, the hole distorts and makes these shapes in the hole that's not a straight hole. And it is not round. And if you think you can drop a nice reactor down this without problems, you've got another thing coming.

Let's also hope the reactors never need servicing or anything. Because getting them back up-- anyway, it's kind of a crazy idea. But you can see how it appeals to the better safety crowd. But it's probably not going in the cheap direction.

All right, let's get to some really different technologies, instead. So this is maybe more promising. This was UltraSafe. They were an interesting company. They wanted to do high-temperature gas graphite reactors.

So these are reactors where the fuel is enclosed in, TRISO fuel, little grains of sand enclosed in some kind of graphite matrix that does the moderation. And that is, then, cooled with helium, which is nonreactive gas. In their concept, each reactor produced 15mw.

They claimed, it was obviously a size of a shipping container. But you'll see in a moment, that's not really realistic. They had a great team of engineers doing some good research, some good R&D. But they couldn't manage their financing, and they went bankrupt.

The company went and started another company called Valor. That company is doing exactly the same thing. The upside of these reactors is they're relatively straightforward and simple to design. You have hot graphite. You have gas flowing over them.

There's no phase change. There's reduced complexity. Because there's no fuel to melt, you potentially don't need containment. So there's all these ideas of cost savings.

And the downside is that the power density is very low. So that means you've got a lot of reactor stuff to build for very little power output. It has that effect on the economy economics.

But you can turn that lemon into lemonade by pointing out, well, if the power density is low, then, if you design this right, you can make this passively cool by radiative if you have an event. Yeah?

AUDIENCE: So there is no fuel to melt, you mean?

PROFESSOR: Not really. So the fuel is silicon carbide coated in a graphite matrix. It can withstand very, very high temperatures. It's not a metal.

So it just turns into glowing embers. And then if you design it right, you can make it radiatively cool so that you don't have a puddle of molten fuel on the bottom of your reactor. So that's kind of the whole idea behind the TRISO fuel variables.

There's another variant of this called X-Energy. Actually, they're one of the better-funded entities out there. This is a pebble bed concept. So the fuel is in these individual pebbles.

They put them in the top reactor. They fall down, and they burn up. And every pebble comes out and has to be evaluated to figure out how much is burned. And then they got to put it back in. And they toss the pebble, and they cool it with helium-- same concept.

This one is prismatic. So the blocks are in these large prisms. So they don't recycle the pebbles.

Historically, the pebble concept is technically more complicated and got problems because it has abrasion. These pebbles rub against each other. And there's a whole lot of issues associated with sorting through 30,000 pebbles.

But why did I say historically? I said historically, because we've done this before. These are not new concepts, these high-temperature gas reactors.

Here are the new ideas. But actually, we've done all of these already. And we have some operating experience.

The first thing to note by looking at these historical reactors is that most of them have operated kind of on the scale of one to two decades. You need to get these things up to four to six decades to make nuclear economics. So the lifetime of these reactors have not been very good.

They also have really high cost per kilowatt hour if you adjust them on a capacity factor basis. So Fort Saint Vrain is the largest of these. It's at the top of what we would call a small reactor, a third of a full size.

Prismatic, HTGR, so no pebble pedal problems. Cost was effectively $12,000 per kilowatt instead of-- this is all inflation adjusted-- instead of $7,700 per kilowatt. And the reason for that was because this reactor had a lifetime capacity factor of 15%. 85% of the time, this reactor was being repaired-- just not great.

The next best, most successful was AVR. This was a pebble version. It cost more money. AVR had a capacity factor of 62%. It's the highest of any HTGR reactor concepts-- so, still being repaired 40% of the time.

It is, today, the most strontium-contaminated facility in the world, because there were temperature excursions associated with the unpredictable cooling of the pebbles. And so we would expect that problem is still not really a solved problem. There are some weird issues. But helium is not a great conductor.

Density also goes down as it gets hotter. So its cooling capacity goes down and down as it gets hotter, which means, once the hot spot forms, it becomes problematic. So this issue led to fission products diffusing out of the fuel and contaminating-- coating the exterior fuel.

Pebbles rubbed against each other, caused abrasion, caused very fine soot, dust. And the fission products adsorbed on the dust and then coated the entire reactor, all the cooling fans, all the heat exchangers. This thing is just a mess. It's just sitting there, because it's too radioactive to disassemble-- price, not great.

China has a new one, HT-10. And they're bigger. This is the prototype 1. Look at the prototype. Price was very, very, very expensive.

They have a one that they're marketing now called HTRPM. The capacity factor for that reactor, I believe, is now 16%. They're having the same problems. Nothing has changed.

So when Valor and UltraSafe and X-Energy come to the scene, I think we have to view these claims with some skepticism. We've been here before in the 1960s. Smart people worked on these things and abandoned them and went back to PWRs for a reason.

And lots of claims will be made. But many of the people who knew what was really going on with these reactors back in the '60s, they're gone now. They're not around to criticize these proposals. So we'll see if these things actually work as promised-- a lot of techno optimism but some naivete.

So the next idea is, say, well, let's get rid of the helium and use something to improve the power density. And the idea of this is to replace that helium coolant with molten salt. So the traditional molten salt is called FLiBe. It stands for fluorine lithium beryllium.

Now lithium is a great neutron absorber. And it makes tritium, which is hydrogen, which then permeates into metals and creates problems. So you don't want that hydrogen production.

So to avoid it, you have isotopically modify this lithium enrichment to 99.998%, which has never been done. You have to purify it to this level. So this is one problem.

Another problem with FLiBe is it turns out that any other chemical contaminants in the salt make the salt super corrosive. So it just eats away your metal. If you remember the Rickover memo, there are paper reactors, and there are real reactors. And he specifically says in his memo, in paper reactors, there's no corrosion. Yeah, this is a major corrosion issue that has yet to be solved.

There are other issues with this. They plan to build the vessel for the Kairos reactor out of 316H. It's a kind of stainless steel. Given the temperatures at which this reactor is supposed to operate, this vessel will last about eight years.

So we were hoping for a 40-year half lifetime, maybe a 60-year lifetime on a reactor. Eight years is just terrible. That means this thing better be damn cheap.

Now the upside is that you get four times the power density. So you go from 5mw/m, our HTGR, into 10 to 20mw/m by using this better coolant. So that does shrink the reactor.

If you assume just the economics of scale as, say, the radius squared, then you can take the cube root of this, a factor of 4, and then square it. And you get an estimate of something like 65% of the hardware cost of an HTGR for this reactor. But that's immediately offset by the fact that this thing has to be replaced so frequently. And you have to deal with all the corrosion issues which are not solved.

This project has been funded by the father of the founder CEO who's running out of patience. We'll see where this goes. And these ideas have been around for a long time and haven't gone very far.

But I think, all of this aside, the real problem is this. Compare the power density to that of a water reactor. A water reactor is a factor of 5 to 10 better.

And water is cheap. And we know how to deal with water corrosion issues. So why do we really want to do this?

These ideas, these salt-cooled reactors have never been built. So I can't show you actual numbers of what they have cost. But the idea has been around for 23 years. And so far, there are many, many unsolved problems.

So I would say, TBD but not promising. Yeah?

AUDIENCE: Why do you think it's so much bigger, the power density? It is because the overall power is so much bigger, and then the amount of heat transfer area you need is just--

PROFESSOR: What?

AUDIENCE: [INAUDIBLE] power density.

PROFESSOR: So why-- oh, yeah, because well, just because the viscosity. You can run water. It has excellent. You can boil, huge heat transfer from phase change. There's a lot-- water is fantastic [INAUDIBLE] salt.

So if this idea of dealing with molten salt reactors is not crazy enough for you, you can get rid of the fuel altogether. So in this concept, the fuel is still pebbles and instituted at the bottom. And it flowed up in through the dense molten salt. And you hope this thing doesn't freeze. You have to cool off the reaction.

You could just dissolve the uranium in the fuel. This is called a molten salt reactor, where the fuel is liquid. There are lots of these ideas. So one of these was actually built. There's lots of these ideas now-- Terrapower, Terrestrial, Copenhagen, Naarea, Elysium, Natura. A bunch of companies are interested in doing this.

It seems to me a little bit like, the crazier and less developed the idea, the more companies are interested in it. It may have had something to do with the state of funding. It's hard to do due diligence on things that haven't been done.

The basic problem is that, now, all your fission products are not potentially leaching out of the fuel. They're in the fuel. And they're coating the inside of the reactor.

The noble metals plate out against all surfaces and pumps and so on. So if you want to replace a pump or do some maintenance, you basically need to go in there with full radiation protection, because the whole reactor is radioactive. All the fission products that don't dissolve in the salt, things like iodine and radio krypton, they bubble out. And then you have to have separate systems that process all of these gases, which have, it turns out, to be quite a lot of heat in the out gas that has to be managed.

So your safety actually becomes dominated by your ability to control these gases that come off the reactor, cool them and store them safely, as opposed to the reactor itself. That tends to be the problem. Maintenance is hard. These things are a mess.

Just to give you an example of what's been done, MSRE is the most successful. It was a tiny 7 megawatt reactor. It operated 30% capacity. Here's the price. Here's the price with its capacity correction, $72,000.

It's got to beat $7,000. These things are not cheap. Here's a small mini one, 2.5 megawatts thermal. It operated for 38 hours over nine days.

And here's one that is supposedly built in China but has never run. And they claim it only costs $1,000 to build. We'll see. It's in the middle of the desert. I guess they're not hoping for success.

So the point is that, all of what they call gen IV ideas move away from the economics of the PWR. And there's a reason why these ideas have been around for decades and decades and decades and have never really been built. They're being pursued today because there's a lot of money being dumped into the nuclear field. And people are looking for ideas that are interesting to work on.

But when we start looking from the perspective of reliability, of cost efficiency, they all kind move in the wrong direction. Now there are some really novel ideas, like the nuclear battery world. So this is something called the eVinci reactor that has been proposed by Westinghouse.

It is supposed to be a TRISO fuel reactor. That was the little, tiny grains of sand reactor that don't melt. That's good.

Graphite moderated-- just like the HTGRs. But it's supposed to get the heat out, instead of by flowing helium through it, instead, it's supposed to have heat pipes. There are lots of ideas in this direction, where you have really small, totally different approach to-- instead of a pressure vessel with fuel in it and coolant, there are lots of other ideas of heat pipes and so on, batteries. And it's difficult to know what these things will cost.

So now I want to look at some of these smaller battery ideas and try to estimate what have been said about these things. So this is just to remind you, that's our goal, $131 per megawatt hour at the top. And here's the cost predictions that have been published.

So I think there might be one more now since I made this chart. But at the time I made this chart, these were all the cost predictions. So these are for abstract, battery-style reactors that have yet to be fully designed.

They are cost per unit energy, dollars per megawatt hour. And the yellow is-- they're first of a kind. Cost is what they predict it would cost to build a single reactor. And the purple ones are n-th of a kind, which say, once they build a lot of reactors. But they never specify how many is a lot.

And we'll figure out how some of these numbers are being generated. All of these are assuming the reactor is running basically 100% of the time, 95% of the time. The reactor is assumed not to be load following.

So if I asked you, what would you do to these prices if these reactors had to load follow, off the top of your head, what would you say?

AUDIENCE: [INAUDIBLE]. It would go up.

PROFESSOR: Go up. But can you estimate how much?

AUDIENCE: Oof. I do not feel qualified to estimate that-- 3x?

PROFESSOR: Not 3x. Does anyone remember what a demand curve looks like? Here's the demand curve for you.

Oh, here's how I would do it. I would say, here's the demand. Peak demand is 80,000. About half of it is on all the time. And the other half is on something like 40% of the time.

So I would say this would be, like, 0.2 of the total. And this would be 0.5-- so 0.7-- so cost divided by 0.7, 70% capacity factor, something like that. 70% capacity factor would be 1.4 times. 1 divided by 0.71-- 1.4 times.

So I would look at these numbers. And I would multiply by 1.4 my head if I wanted to do a full grid of small nuclear batteries that had to do all the load following. This is a full grid of wind and solar that has to do all the load following.

So now we're doing an apples to apples comparison if we compare this number to this times 1.4. Does that make sense to everyone? Yes? All right, I just want to make sure we're making a fair comparison here.

All right, so obviously, there's a lot of variability here. So let's just go through these really quickly so you can see what has been said about these batteries. So first of all, CAREM is not actually a battery. This is a 25 megawatt reactor.

It is not factory made. It's a traditional PWR. And it's just very, very small. And it's currently being built. And it's the only of these ideas being built in Brazil.

Its design goes back to 1980. So it's been being built for 45 years. That's partly or largely because Brazil is not a very rich country. And they might have money for it. And then they don't have money for it and so on.

According to the published historical cost data for what it actually cost, I've calculated the price of electricity from this reactor-- and you just have to trust me-- is $370 per megawatt hour based on the actual cost this reactor has cost them. But the reason I put it up here is it gives us a baseline for what a very tiny reactor has actually caused.

All right, so now let's go through some of these more battery concepts. Jacopo Buongiorno in our department published the most optimistic estimate. He does not give a first of a kind cost. He gives only an n-th of a kind cost in his paper.

And to design this reactor, which had never actually been designed yet at the time he published this, he decided to take the internals of the reactor and say it was basically the equivalent of a GE 90 jet airplane engine, which you could go either way. At one level, that's a high precision machine with real reliability constraints, a lot of complicated parts to make. It sounds kind of like a nuclear reactor. So you could say that it's a good concept.

On the other hand, it's not a reactor. It doesn't have to deal with radiation. It doesn't have to deal with as much corrosion. So maybe it's-- who knows?

Anyway, he comes up with an n-th of a kind cost that goes between-- you can see these bars are not actually error bars. They're ranges. He says the most probable is about 85, but his range goes all the way up to 355. So if we multiply by 1.4 to deal with the capacity factor correction, that goes from $120 per megawatt hour to $500 per megawatt hour. And so at the low end, you just beat out the non-optimal wind and solar grid price.

NEI's numbers, Nuclear Energy Institute, they've also published estimates. This is a fantastic paper. They're not based on any design. In fact, they give absolutely no basis for their cost estimates whatsoever, except to say that they Surveyed microreactor companies.

And then in appendix, they list all the microreactor companies. And they say none of these companies were consulted for this paper. So I don't know where they get their number.

It seems to come out of thin air. But there's the number. They have $270 and $148.

Idaho National Lab, which is our government lab that actually builds reactors, has a real design. It's not a completed design, but it's a real design. And they went and priced out their design, which they called Design A. And this is what it costs, $2,174 per megawatt hour if operated at 95% capacity.

So then they say, well, what if we build a lot of these? So they came up with an n-th of a kind cost. And if you read very carefully in the paper, how did they come up with the n-th of a kind cost?

They run around to their colleagues and asked them, what do you think it should cost? And they took the average of that number. And that's that. So you can believe it or not.

And finally, there's the Department of Defense called Project Pele. And there's actually only a single reactor here. It's really hard to know very much.

Part of this reactor's uncertainty is that they say it's somewhere between 1 and 5mw, which, given the current budget, means a factor of 5 variation in a per unit energy cost. And that's why this has a huge error bar on it. Very little is known about the details of this reactor concept.

So I think the key takeaway here is that no one is actually claiming that these microreactors are cheaper than conventional nuclear power, these battery-size microreactors are cheaper than-- we're putting small modular aside-- these tiny reactors, which is what these are. And that's interesting. So it really does depend on there being a market for these things that is not the grid.

And those markets may exist. They may exist in heat markets. They may exist in remote locations where the grid doesn't exist, things like that. But this is they're not arguing this is the path forward to making nuclear cheap to solve the general climate change problem.

So we could kind of stop there. But why? Because there's so much more to talk about.

One of the issues is, I think this n-th of a kind cost methodology is kind of bad. Going around and asking your colleagues for what they think of your cost is not a good method. And I also think they're leaving some potential wins on the table when they go to n-th of a kind cost.

So I want to talk a little bit more about, first, how realistic are these first-of-a-kind cost? And then let's look a little bit more at this issue of factory fabrication just to see if we can do better at figuring out what we can get for n-th of a kind. Sound good?

All right, so let's just start with the beginning, which is do we believe these first-of-a-kind numbers to begin with? So generally, as nuclear reactor designs have become more mature-- remember, the ones I just showed you are as immature as they could possibly be. They're basically paper concepts. They generally, tend to get more expensive as they go up, as they get more mature.

So here's NuScale. These are all in 2023 dollars. So they're all inflation adjusted. NuScale in 2007, in 2020, 2021, 2023-- and you see how the cost estimate just keeps growing as the design is filled out.

They think, oh, we need this other thing, or we need that thing. Oh, this costs more than we thought. Oh, this is going to be hard to manufacture.

And this is the same for Fusion. This is the exact same thing is going to happen in Fusion. So NuScale started at $5,800. And in the last three years, it has doubled and doubled, doubled and doubled.

X-Energy here-- there's the initial cost. It's gone up by a factor of 4. I really wish this thing wouldn't vibrate.

GE-Hitachi, this is a small BWR, same story, up by a factor of 4 1/4. So the takeaway from this is you really need to know something about how mature the design concept is when you look at these cost predictions. If you just take these numbers that are on these websites and they haven't actually done the design work, you can't take that as the final cost.

But I want to give you a rule of thumb that you can use to apply so that you can be smart about this. And the rule of thumb is going to be based on the AP1000. So the original design was set to cost around $3,300 per kilowatt of capacity.

And the completed one cost $20,000 per kilowatt of capacity. Remember, they now claim they can build it for $7,700. But the one we actually built costs that much.

So the lesson is that, going from the early days to the start of construction, we saw the price triple. And then construction-related issues doubled the price. And then whether or not this is something we have to live with forever or not is what people argue about. So the DOE number is basically this number with some cost reduction from learning in it-- basically, what the $7,700 per kilowatt hour is.

So you might say, well, maybe AP1000, maybe this cost escalation is a one-off. Probably not. And I'll tell you why.

The Congressional Budget Office has done an evaluation of 75 nuclear plants in the US. And on average, they find that the finished reactor costs 207% more than the original reactor, according to their survey of 75 reactor builds. So I would argue that this construction-based cost escalation is not a fluke. But many people argue that we can somehow get past it if we just start building a lot of reactors.

If we were to apply the same rule to NuScale, then we would expect this $21,000 number to eventually become $64,000, which would just be crazy. So it's just something to bear in mind. I guess I would say, bottom line, expect a factor of 2 to 3 between the finished design and the finished construction and a factor of 3 between early-stage design and finished design, something like that.

This is a factor of 3 1/2, a factor of 4. These escalations are basically factors of 3. So that's kind of a good rule of thumb. So overall, factor of 3 times a factor of something like 2 or factor of 3 for construction gives you something like 6 to 9 times the original cost estimate.

AUDIENCE: So is this just, like, a lack of foresight? Or is it just they need to price it this way in order to get the initial funding? Why are we so bad at estimating the costs from the get go?

PROFESSOR: It's a good question. I sometimes wonder if it's malicious. It certainly has repeated itself in every single reactor build since Shippingport, basically. But it's also true that people are just notoriously bad at estimating the cost of mega projects. And you find these cost escalations everywhere in other projects.

So it's some combination of these things. And I can't say what part is malicious and what part isn't. I certainly know, when I see the cost estimates for Spark, and they're like, well, it cost this today, but we're just going to divide this by 4, because we think we can get the price down. I'm like, well, that's unsubstantiated. You shouldn't really be doing that.

And so that I call the malicious pricing. And I think there's a lot of that in the startup world. If we go back here, this company had pretty good engineering. And they said it was going to cost $10,000 per kilowatt hour.

This company is building a reactor that's more complicated, because it's pebble instead of prismatic. And they're like, oh, $2,000. And you start rolling your eyes at that point. Yeah?

AUDIENCE: In low cost, it's expressly encouraged, at least in the US VC culture. That's what, in the whatever workshops I attended, they would say, if you approach an investor, take your most optimistic case to the point that you're just not lying, and do that. Because they expect you to do it. So whatever you tell them, they're going to internally multiply by 4 or 5.

But essentially, the one liner they say, is if you don't provide me a step change, I'm not interested. So if you say I'm going to be 10% cheaper, they don't want to hear it. They want to hear it, you'd be 2 times cheaper. So at least in the startup world, it just seems to be the game that is played.

PROFESSOR: It seems to be true. That seems to be what I get from other people. If you just look at the history of what gets funded, the more audacious the claims, the more likely they are to get funded. I don't disagree. Yeah?

AUDIENCE: Do we have number? I'm just curious what the comparison is for other mega projects. Because bridges, for instance, notoriously do the same thing, where they end up massively--

PROFESSOR: Dams, bridges, tons of stuff do this.

AUDIENCE: I'm curious as to if it's the same magnitude, roughly, if it's also 6x?

PROFESSOR: That would be a really great paper.

AUDIENCE: That's good to know.

PROFESSOR: Yeah, lots of great paper topics all over this thing. Yeah?

AUDIENCE: So do we expect that there's hopefully [INAUDIBLE] to your price [INAUDIBLE] idea, even though we think microreactors maybe-- we're not expecting them to be cheaper, necessarily, but that they will maybe avoid just escalation of price from design to construction?

PROFESSOR: Yeah, I think that's part of the hope-- is that maybe part of it is the mega project phenomenon. Yeah. But then it's counteracted by what you just said, this crazy VC nonsense and people pursuing ideas that are really not very mature.

AUDIENCE: Is the same true in other countries? Or is it just the US that's particularly bad at pricing mega projects?

PROFESSOR: Oh, no, no.

AUDIENCE: This is across the board. It's more of a human thing.

PROFESSOR: Yeah. Yeah. All right, so where were we? All right, so I guess the point is, if you take this and you apply our factors of 2 and factors of 3 multipliers to these things, you see that the prices are just really, really high. They're all well, well, well above the $131 per kilowatt hour. But we haven't done mass production yet.

So let's see what we can do here. So the thing is, we're starting with the fact that no credible SMR vendor is claiming that these things are cheaper to begin with. But they all say it will get cheaper if we build lots. And the question is, do we have to build 10, or do we have to build 10,000?

And so, generally, when people have looked at this-- they call it learning. People just look at the cumulative number of reactors built. And that this is how I've taught this class in the past.

But it dawned on me recently, this is probably not the right way to look at the problem. Because it certainly matters whether you're building one reactor a year for 100 years or you're building 100 reactors a year. If you have a plan that's producing 100 reactors a year, you're going to have much better economies of scale than if you're producing one reactor a year over 100 years.

So I think we need to look at both total number of reactors built and the rate of production. Yeah?

AUDIENCE: OK, so this is a question that I haven't seen posed that often. But how much does standardized components versus standardized reactors make a difference? So for the Apollo program, by the end of it, we were getting very, very good at having standardized components for engines and all that stuff. And despite them being extraordinarily complex, they were fairly mass producible.

For nuclear stuff, I haven't seen a lot of people considering, can I standardize all the components of my shield wall or whatever? And I don't know if anyone has actually ran the cost estimates on that, versus, if I have hundreds of the same component in a reactor, versus, I need 100 different components but I'm making 100 reactors--

PROFESSOR: Yeah, I think what you're basically hinting at is the right kind of thinking. You want to basically have some internal component that is replicated that you produce a lot of. Obviously, you have a lot of spacers in an LWR. But they're still kind of boutique-ly made.

If we could really get to the point where we made tons and just tons and tons of these things, I think that would be interesting. So I don't have an answer for you. I think there's another good paper topic. Yeah?

AUDIENCE: And they can actually--

PROFESSOR: Oh, good.

AUDIENCE: [INAUDIBLE] to a certain degree. In the most recent AP1000 Bill, mass-producible parts within the US was actually a core consideration. The main issue between, say, like a nuclear steam supply system and something you use for gas is the quality of the parts. So you would need a slightly different grade of steel, and you need a slightly different grade of concrete.

But they were actually able to set up these vendor relationships with specialized manufacturers, who were able to create these steam supply systems. So there is some of that. But nuclear still would be-- they tend to be categorically different equipment requirements than any other--

PROFESSOR: Yeah, the N-stamp is expensive. And that's why the falsified safety certificates in the Korean reactors were falsified. Getting to that nuclear-grade construction is hard. And you see companies like X-Energy, which are building prototype systems. None of those systems are N-stamps.

They're all U-stamped. So they need pressure vessel requirements. But we haven't gotten to the point where they're actually doing the welds in a way that will stand up to radiation and things like that.

So yeah, I think a lot of it is this N-stamped construction is very boutique. It's very expensive. We'll see to it. But potentially, all that could be scaled up.

Anyway, I wanted to start by talking about the first of these two things. So the one is, how many did we build? And the other one is, how fast are we building them, the rate of production? So let's just start with the traditional way of looking at it, which is, how many have we built?

And the question is, if we factory fabricate a reactor-- well, we know what the number is if we don't factory fabricate the reactor. We did the regression. I showed you the regression result. And the learning was 1% cost reduction for every doubling of the number of reactors, which was not great.

But those weren't factory fabricated. So then, what about factory fabrication? So we have Naval reactors to go on. Unfortunately, I have yet to find US Naval reactor construction data, cost data to really derive anything on this front.

However, there is one claim from the Soviets, from a manager of a Soviet Naval reactor factory in the Soviet Union. And he claims in a paper that he wrote in some obscure journal that their factory saw a 5% cost reduction for every doubling of the number of reactors they produced in their factory, which is 5 times better than what we are observing for sig build. So let's just see what that means.

So here for comparison, are the observed learning rates for solar cells, which are phenomenally good-- and yeah, just solar cells. And you can see how you came from $40 per watt down to down here under $1 per watt. And that's because of this amazing power of learning.

Now obviously, these are also being made in factories that are churning them out in large numbers. I've put the 5% learning rate on the plot arbitrarily. What matters is the slope, not the number. Because I don't have the number, the cost per kilowatt of the reactor on this chart.

But you can just see how much slower that is. These faster learning rates make a big difference. So let's see.

The DOE prediction for traditional nuclear was $80 per megawatt-- so twice-- OK, that's per watt capacity. But you can just multiply it by 2. And you can see that, if we built 100,000 reactors, we would be something like 20. So we would get the price down by a factor of 2 by building 100,000 factory reactors.

It's not probably realistic that the market is that big, that we can build 100,000 of these things. We could work out on the board how many we would have to build up. It turns out, you need to build-- I forget. I have it somewhere in my notes-- lots and lots of reactors, more than 100,000 reactors.

Oh, you need to build about a trillion reactors to get cost competitive if this is the only thing going on, this kind of learning. So I think this is probably not the whole story. So then I was thinking, OK, well, what is it about the factory fabrication that makes things cheaper?

And at first, I thought, is it replacing manual labor with machine labor and better quality machine labor and going from manual machines to, say, computer-operated machines and so on and so forth. So I looked into this a little bit. And I was trying to find a comparable technology with some comparable level of complexity as a reactor.

And I thought, oh, Jacopo Buongiorno decided to use jet engines. So why don't I use jet engines? So here is the history of the CFM56 engine, which underwent a bunch of changes over time going back to the late 1970s, all the way up to today. Mid 2015s was the last order for this engine.

And during this time, we have seen a lot of changes. We saw the introduction of advanced CNC, the introduction of 5-axis machining. We saw computer-controlled machining go for something that ran on tape to something that ran off of a CAD file.

So there's been considerable advances in machining during this time. But the cost has basically been exactly the same. And so at first I was scratching my head. But when you think about it, it kind of makes sense.

Basically, it appears that manufacturers are trading labor for machine for capital. And the price of the capital that they're buying, which is highly specialized, is going to be priced such that it's basically the same cost as their neighbor, maybe a little bit cheaper. So it kind of makes sense from a market dynamics perspective. As long as these machines are highly customized for the factory, for the application in hand, we don't really expect to see a lot of cost reduction just from the technology change.

And it turns out, this is a cost per kilogram of engine manufactured, cost per unit thrust of engine manufactured. It's all the same. It doesn't go down.

So I thought, OK, maybe the answer is something else. Maybe it's really the rate of production. And I would say this is the operative question.

It turns out, there's not a whole lot on this in the literature about the rate of production. But here's the best thing I was able to find. This is a study done by NREL for these stationary fuel cells.

So inside this fuel cell are these very carefully manufactured membranes and electrolytes. And they move protons or ions across these membranes. And they have to be very carefully assembled. And then there's some computer and control equipment and other things.

So this is some complex process that is, as earlier mentioned, eventually built with out of many identical parts and mass produced. And then these things are dropped in to places that need backup generation. So these things are basically a kind of giant battery.

And NREL has been a very extensive study done on producing these things. And this is their takeaway. If you produce 10 systems a year, 100 systems a year, you're here, 1,000 systems a year, you're here, 10,000 systems a year, you're here and so on. You do this math.

And what they find is basically there's a 30% cost reduction for every tenfold increase in the production rate. Does this apply to nuclear? I have no idea.

This would be a really good question for someone to study. There is a lot of work in the automotive world, which I'll show you in a moment. Yeah?

AUDIENCE: So the INL published a paper for the MARVEL microreactor, where they worked together with some manufacturing company to try and estimate, based on their workflow, what learning rates could be for microreactors. So it's obviously not real-world data. But they seem to go pretty in detail.

PROFESSOR: I haven't seen this. Yeah. So do they do rate of prediction? Or they do per unit prediction?

AUDIENCE: So I vaguely remember, they looked at producing 10 per year or producing 100 per year and then kind of splitting up the tasks. And then based on what kind of scenario you have, they used different multipliers. It seems to be pretty detailed. And they work together with a--

PROFESSOR: A real company.

AUDIENCE: Yeah, that actually knows how this works. But it's all still obviously on paper.

PROFESSOR: Yeah. Do you remember, what was their conclusion?

AUDIENCE: Their conclusion was that they could get pretty high learning rates in the highest scenario, I think, like, 15% or something. So it wouldn't be an entire moving assembly line. It would be kind of like a rotating system, where it kind of moves between different workshops. And they seem to think that you could get decent learning rates.

PROFESSOR: This is very interesting. We'll go look at it. 50% is the high end of what people quote for modular reactors. So did they have a certain production rate used to sustain that?

AUDIENCE: It's been, I guess, a year and a half or more since I've read it. But they looked at, I think, 1, then 10, and 100 per year as their scenarios.

PROFESSOR: And they think 15%--

AUDIENCE: In the highest. So in the 100th, 100x.

PROFESSOR: And that's for every doubling produced? Or is that 15% off of the single unit cost?

AUDIENCE: I think that would have been per doubling. I think that's-- but, yeah. I don't trust my recollection. I can look it up again.

PROFESSOR: Yeah, definitely want to look at this paper. Because I'm surprised I didn't find it. OK, super.

Well, anyway, I took this. And I just thought, let's run with this, and see what we can do-- a 30% cost for every tenfold increase in the production rate. So here's kind of my thinking. All non-renewables generation in the US planned for the next five years is about 4 gigawatts of additional capacity per year, which is really remarkable, for the next five years.

And this is really being driven by data centers. So while there's a lot of deals about data centers and nuclear out there, what it's actually being built is natural gas. So what if what if an SMR were to replace all of that natural gas that is coming online in the next five years here in the United States?

So X-Energy says their reactor is 80mw. So that would be 50 units of capacity. And they say it costs $18,000 per kilowatt. That's their most recent price assumption. So if we assume a factor of 2 to 3, because that's an early price and we want to get to a real price, maybe we're something like $45,000. Maybe their price is correct. But historically, no one has been correct.

And then we take the fuel cell. And we say, OK, if we could produce all of these at the same rate for gigawatts per year, then, instead of one per year coming out, coming out, we would have 50 per year coming out. And we would have the price drop down to this number, about $20,000, which is about the same price as AP1000 turned out to be.

So you could potentially see how this would be competitive with traditional nuclear. If we could produce 500 reactors per year instead of 50, this comes down another 30% to $14,000. But you quickly run out of market, even with such a small reaction. So it's still too expensive. But who knows?

So this is, I think, the real unknown, which is alluding to what has been said. What is the minimum efficient scales? And it turns out that, for automobiles, which are other complex objects, we have some real data. And so economies of scale for the industry suggests an automobile plant has to produce 200,000 to 300,000 units to be cost efficient. Engines should produce 400,000 units and transmissions 500,000 units.

And that just gives you a sense of why super cars cost so much more than your Ford Focus. So something like this is needed. Some conclusion like this is needed for the nuclear field in order for us to really know what we can do.

But I think there is some work to be done here-- maybe some of it has been done-- to try to understand, where is the optimum size? Obviously, the smaller the reactor, the more you produce per year, the better your rate of production per year. So somewhere in there, there is an optimal. And it would be really neat to see someone try to pursue that.

All right, that brings me to my last topic, which is, let's imagine there is a future here. I don't know exactly what it is. It does not really emerge from the review. But let's imagine there's a future. When will we see these reactors coming online?

Well, the NuScale reactor has been, I mentioned already, under development for 20 years. The project actually has just been halted. But even before it was halted, they said they had planned to have the first unit in service no earlier than 2030, which would have been 25 years from conceptual inception to first unit, a quarter century.

How about the AP1000? AP1000 started its life as the AP600 in 1984. The design was submitted to the NRC 10 years later in 1994. NRC approved it in 1999. So that's 15 years from conception to approval.

Then they said, well, this reactor is not big enough. We need to redo the economics. They stretched it vertically to make it the AP1000 from the AP600. That took one year, not that hard.

And then they resubmitted for another NRC license approval. That took four more years. So the total time from design conception to final license approval was 20 years, same as NuScale.

So this just illustrates how long it takes to conceive of, engineer and license a reactor, 20 years, basically. And at five to 10 years of construction time on that, you get a quarter century. So I think the idea that a bunch of these reactors are going to come online and supply electricity for data centers is crazy.

These data centers are not going to be here in 20 years. This is all just [INAUDIBLE]. And there's a lot of history from reactor to reactor to reactor that will support these timescales.

If we are serious, we ought to be thinking about these reactors as something that is coming into a future world that has seen a quarter century of wind and solar installations. Because that stuff, as I already showed you, is happening at a gigawatt a day in China today. So in 20 years, when your first design comes online, that's the world in which it needs to compete.

And that means these reactors need to be load-following reactors that are going to deal with the variability of a lot of wind and solar. And so this is now shaping your design space and how these reactors are going to operate. That, or they're going to be doing things that are not electrical grid, like maybe providing processes for manufacturing, this second-tier decarbonization efforts, making paper, making iron, things like that.

So that's my bottom line conclusion for microreactors. Let's think far in the future. Let's think factory fabrication in large numbers at large production rates. We need a single standardized design, not 10 different competing designs.

And let's just forget about electricity for now, unless you're really going to do like peaking or something like that-- probably something else. So that's my takeaway. I thought, since we have a few more minutes, maybe I can blast through this, unless there's any questions about what I've just mentioned.

I wasn't planning to show this. But we do have 10 minutes. [INAUDIBLE]. So here's what people have said for small modular reactors or microreactors-- disaster relief, Arctic and remote communities, mining operations, island communities, remote military installations.

Here's how my brain works. Disaster relief-- every now and then, you need reactors. This is not thousands per year. This is a couple reactors on standby.

Arctic and remote communities-- we'll look at this in a moment. Remote mining operations-- we'll look at this in a moment. Island communities-- OK, how many islands are there?

Remote military installations-- this is a little bit complicated. There are hundreds of military installations. But whether you want to actually put nuclear reactors there and whether it's a good idea is TBD. Because even if the reactor is not going to melt down because it's made of TRISO fuel, if you blow that thing up, those TRISO particles will fly a mile away and contaminate a huge region. And a mile is not unrealistic for a kinetic explosion of a microreactor.

So whether you want to put these things in a war zone, I think is a bit crazy. But let's look at these things. And then let's look at this market, the developing world that currently doesn't have electricity, that doesn't have much of a grid and what we can do.

All right, so let's do this real fast. So here's all the rural communities with installed power production. I forget the threshold requirement, because I wasn't planning to present this. Oh, this is total megawatt capacity.

So it turns out, it's something like four to 20 reactors, depending on their size, in Alaska. And if you expand this to all of Canada, you get another 20 to 30 reactors. So you're talking 40, 60 reactors, something like this. And if you're hoping to get hundreds of produced per year, this is probably not going to give you the economies of scale you need.

It turns out, mining operations are also not great. There's a couple of mining operations that need a lot of electricity-- maybe worldwide, a few dozen. And this is probably also not a sufficiently large market to create the mass production that we need.

So the thing that's really left is actually this at the end. What can be done in the developing world? So here's what we did.

We took satellite data that showed where all the lights are from space during night. And we used that as a proxy for how much electricity was installed. And we looked at GIS data for where people lived. And we had road network and land and topological data and so on and said, OK, let's find all the areas that the land is suitable for installation of a reactor where people are living but there's low nighttime light, indicating poor electrification of those countries.

And we did that. And here's the results. Here's the electricity needed to raise everyone living in the world to Tier 5 electrification standard, according to the World Bank, which is 8.2kw hours per day per household. And you see that, especially here in Africa and here in India, a lot of potential for electrification.

It turns out that these dark blue areas, the power density is too low, less than a megawatt, to really support a reactor. So we're kind of really interested in the green to red zones. Then what we did, we said, you need to be kind of close to a road. We can't be going through a forest or something to try to bring in a reactor. So you need to be close to a road, close to water or a rail networks.

And you have to have physically stable ground. It can't be a wetland or anything like that. And we came up with-- this is the site ability, where you could site reactors of different sizes.

And then we combine these things. And these are places where you could actually build reactors and there's not so much forest that you can't actually do it. So this is the hope. 78% of the 630 million people living in these areas could be potentially theoretically served with small nuclear-- so very exciting, in principle.

And then we said, well, but you need to have a country where there's not a lot of corruption going on. You have tc have laws that work. You have to have some minimum regulatory framework for safety. And you have to have an effective government that can manage this project.

And you could debate how this should be done. But we said basically, Ukraine was like the threshold. Ukraine currently has nuclear that it inherited from the Soviets. It relies on the United States heavily to keep those nuclear plants safe and operating. It's kind of the threshold of where nuclear makes sense. So we just took the Ukraine and said you have to be at least as good as the Ukraine.

And when you do that, you exclude 80% of the population. But that's still 98 million people. And then we said you need to not be in a war zone. And I don't know how this is shown on this plot. We did an exclusion of war zones, which I guess I don't have a slide.

So what we found was that you still have room for about 300,000 reactors, potentially. And that's an interesting number, because 300,000 is where you really would get economies of manufacturing. And moreover, if you look at the cost of African electricity prices, there's something like $80 to $600 per megawatt hour compared to that. Today, Boston is $50 a megawatt hour. So you could convince yourself that this would be affordable.

But the flaw in the logic is that they want to spend that price and that that's scalable. And the real reason why the prices are that way is because it's not scaled. And no one can afford more. So you might look at these numbers and say, oh, $80 to $600, maybe some of these technologies would work. But the problem is, really, that's kind of like taking the poorest people and saying we're going to make you pay for the economies of scale for our reactors.

And it turns out, as we later learned, of the 39 countries that are here in Africa, only two of these countries are actually able to cover their costs making electricity. The other 37 operate their electricity sectors at a loss. And they are basically subsidized by the government.

But it's one thing for the government to subsidize electricity in the capital city so that they can operate a city. It's another thing to say, oh, now we're going to subsidize 100 million additional people. Where is that revenue going to come from?

So it turns out that this hope doesn't really hold water when you consider developmental economics in play. But then this turned out to be the real thing. Where are you going to find all these people to run all these reactors?

I think this is kind of a general problem for an expansion of nuclear in any country. If every reactor requires two engineers, then, for 400,000 reactors for five shifts, you're going to need, like, 4 million nuclear engineers. So either these things really do have to be completely computer controlled, or we have to have some massive expansion in reactor training to actually operate all these reactors that people are talking about.

All right, I will wrap up there. We're out of time.