RICHARD SCHMALENSEE: Today, we're talking about energy demand. I do owe you 2 pictures however, which I will now supply. The first is-- that's Bangladesh as it is today with 112 million people, and as it would be with a meter and a half sea level rise, about 17 million people or 15% of the population live in the areas that would be flooded.

A meter and a half may be a little aggressive over the century, but it's within the bounds that people talk about. They talk about a meter, meter and a 1/2, sometimes 2. So that's Bangladesh.

This if it comes out, this is a picture that I promised you to talk about-- radiative forcing and methane here versus carbon dioxide. So this is a 2007 study that looks at total contribution to increased radiative impact between 1750 and 2005.

You will see up here, CO2 has a much bigger estimated effect than methane. And the molecules last a lot longer in the atmosphere. So each methane molecule, as Susan Solomon said, is more powerful. But half of these will be gone in 10 years. Whereas ballpark, half of these will still be around in a century.

So there are clear reasons for focusing on carbon dioxide. The halo carbons are fluorocarbons, chlorofluorocarbons and the like. They are very powerful, and they last a long time. But they're not greatly abundant. And people are working to phase them out.

Somebody mentioned black carbon. There's a black carbon impact via the reflectivity, the albedo. It does have a positive impact, land use, negative impact and so forth. Aerosols are interesting. This is like sulfur dioxide put in the air by coal-fired power plants. That reflects, so that tends to cool.

When Mount Pinatubo blew in the Philippines, there was I think about a two degree impact on climate from basically all that dust reflecting sunlight. So that's the best estimate. It appears lots of places of the cumulative impact of human activity.

1750 is usually taken as pre-industrial times, kind of the benchmark. Does this make some sense? Are there questions about it? OK, hearing none.

Today, I want to talk about energy demand. We'll also talk about it Wednesday but from a rather different point of view. I want to talk about where it comes from, the sort of short run, long run difference. How you might estimate characteristics of energy demand and then stability and some puzzles. So this is basically going to be a walk through that teaching note. So it's fairly basic, fairly basic stuff.

With very few exceptions, there is no demand for energy. There's a demand for what energy does. I mean you don't buy gasoline to have it around the house. Gasoline is nasty. Vapors are explosive. You buy it for what you can do with it.

The demand for gasoline is derived from what it does, what it does for you. And if you think about all the things you use energy for, I mean, you use it to heat food. You use it to heat the house. You use it to cool. You use it to make loud music. Energy does a lot of things.

You don't actually want to run electricity through your body, you want to use it to drive speakers and so forth. So the question is, if it's that kind of a commodity, how do we think about it's demand? Where does it come from?

It's clear that to produce any energy service, heat, cooling, transportation, it comes with a bunch of inputs, like any other product. Transportation will involve capital in the form of some kind of vehicle, driving energy, in this case, not materials other than gasoline. But well, maybe spare parts, oil, and so forth.

So you start with that. And I just make this interesting point that you normally think of inputs into production as substitutes. So if capital gets more expensive, you use more labor. You switch from capital to labor.

Well, capital and energy has turned out in a number of studies to be complements, which is pretty rare. If you raise the price of energy, you reduce the demand for energy. You reduce the demand for capital used in production as well, OK. I will talk about this last point in a minute, the progress in these production functions. Clear so far? Yeah.

AUDIENCE: Sorry, what example?

RICHARD SCHMALENSEE: Materials. That's just a general form of a production function. If you want to produce cooked food, you need some piece of capital to contain the heat, some source of heat, your cooking time, and of course, the ingredients. So that would be an example where it would be important, OK?

This is just some indication. There's a Nordhaus paper I think I cited in the note that looks at the cost of lighting. Cents per 1,000 lumens from 1800 to today. And it's dramatic. You can actually do light because you can look at the output of various lighting devices. You're doing whale oil up here and so forth.

And it does get a little flat. Once you get the fluorescent bulbs, it gets flat. And we're down here beginning to think about LEDs. But it's also true in automobiles. I don't have the slide, but 2.6% of the year advance in automotive technology estimated by my colleague, Chris Knittel, means that-- well, here's one example of what it means.

If you kept all the other characteristics of an automobile at 1980 levels, we'd have a 50% increase in miles per gallon, or decrease in gallons per mile between 1980 and 2006. In fact, we didn't. Mileage was pretty stable and all these things increased. But 2.6% a year is a pretty dramatic increase in the advance of a very mature technology.

And you see this in a lot of energy using technology. Power plants have this lighting, motors, lots of energy using technologies, air conditioners have this kind of progress that energy per unit of output has fallen dramatically historically.

All right, so let me do a little bit of characteristics of derived demand. Because when you think about energy, energy demand, you do need to kind think about it as derived demand. So let's suppose capital is fixed. It's the short run. I'm going to just lay this out for the moment.

Yep, so we want to drive the demand for gasoline. We want to think about where it comes from, or the demand for gas to dry clothes, or to cook food. There is sort of two functions involved.

The first is some demand function for the service. What am I willing to pay to have my condo cool? That's going to depend on the cost of cooling it, other prices, how wealthy I am, capital stock tastes, all kinds of things. But I'm going to focus on the relation between how cool I keep it, my demand for cooling, and what does it cost me to cool?

What it costs me to cool comes out of how I produce cooling. There's some production function determined by my very old air conditioner that I'll get cooling services. It will depend on, well, in this case not particularly materials or labor. It's just going to depend on how much electricity I use.

So in the case of cooling, this function gets simple. The cooling, however we measure it is going to be a function of how much energy, how much electricity I use. We could get fancy. It will depend on when I use it and so forth.

But as a simple story, that's a pretty good story. So if what I'm interested in is the demand for electricity as a function of its price, then somehow I don't have a price for cooling, right? I don't pay. I don't get a bill to cool my apartment. I get a bill for electricity.

So if I want to get a demand for electricity as a function of its price, I need these. I need this function. How much do I care about cooling? And this function, how do I get cooling? And then obviously, something about the price of electricity.

So I'll tell you how you do it in general, and then we'll walk through an example. And there's some stuff in the notes. So if I've got all the prices, conceptually I say, well, what's the best way to produce cooling? That is a function of prices, prices of the way down here rather.

If I know the prices of energy, the price of anything else I use, I can say, what's the cost of cooling as a function of various prices? Then I can substitute the cost of cooling, the price of the service up here and get the demand. That's impossible to follow. Am I right? OK.

You'll come back to it after the example. So let's do lighting with two very simple functions. Suppose up here is my demand for lighting. So lumen hours, which I'm just going to call lumens to keep my slides a little shorter, my bullet's a little shorter.

Lumen hours, this is my demand. If this is the cost per lumen hour, then with A and B as constants, then this is my demand per lumen hour. It's a straightforward demand curve.

B is the price elasticity. You may remember that from first economics course. It's the ratio of percentage changes. Limiting ratio is the changes go to 0, the percentage change in demand induced by a 1% change in cost, or an epsilon change in cost, the ratio.

We use elasticities a lot in economics just as a reminder because they are unit free. There's no natural units here. So you could measure price in pennies or dollars or euros or whatever else. So you don't want to say the notion of change per dollar doesn't have any particular intrinsic meaning. The notion of percentage change from a 1% increase in cost is usually a little more convenient.

So let me suppose that's my demand curve for lighting. And let me suppose it's the short run. So again, the more electricity I use, the more lighting I get. And e is just some constant.

So now what I want to know, I want to get my demand for electricity for lighting, right? It's a derived demand. I don't care about electricity. I care about lighting, about L. I use electricity to produce it. The demand for lighting is derived from the demand for electricity.

So how do I do it? OK, in this simple case, if I've got a price of electricity, say dollars per kilowatt hour, I could easily solve with this thing for a price per lumen. I know what electricity costs. I know the relation between electricity and lumens, it's linear. So I know the price, lumen hours, but lumens.

So just to get the unit straight, the units of this little constant e has to be lumen hours per kilowatt hour, right? And I want a price per lumen. So I get $1 per lumen price by dividing the per kilowatt hour price by lumens per kilowatt hour.

So I've now got the cost per lumen hour, OK? With me so far? Scream and shout if not. Well, don't scream, but a shout would be good. I can substitute then.

I know that the amount of lumens I get is this little constant e times the amount of electricity. The demand curve then is A times the dollars per lumen to the minus B. That's dollars per lumen, PE over little e.

Solve for electricity, and I have the demand for electricity as a function of its price, OK? Walk one more time. I care about lighting as a function of the cost of lighting. I don't care about electricity. I care about lighting as a function of its cost.

What I want to do first is get the cost. And in this case, because it's the short run and very simple, all the bulbs are fixed. It's just proportional to how much electricity I use.

So the cost-- you can think of e as relating to efficiency. You increase little e, basically, I've got more efficient bulbs, which is another way to get quickly. I do it with units because otherwise I get backwards. But the larger is this efficiency parameter little e given the electricity price, the cheaper is lighting, right? So it's another way to think about it.

Once I have the cost of lighting, I can substitute into the demand function. Note that here is the production function. Take this second to quality and solve. And I get the demand for electricity is a function of the price of electricity and the efficiency of lighting.

Question, is this easy? Hard? Confusing? Backwards? Upside down? Will it be on the quiz? Yeah, probably. Yeah, something like it. Yeah. Julian, Julian, you had a question?

AUDIENCE: Oh, no, I was giving you a thumbs up.

RICHARD SCHMALENSEE: You give me a thumbs up, OK. This is a rare and wonderful experience. Let me just treasure this for a moment. Does anybody have an alternative reaction of any kind. OK, all right, I move on.

A couple of comments about this. Think about that function. You know, I really should have put it on this slide. You notice a couple of things.

You'll notice that the elasticity of demand for electricity turns out to be the same as the elasticity of demand for lighting. That is not a general property. It can be greater or less. It's not a general property.

You will also notice that what happens to the demand for electricity when efficiency increases depend on that parameter B. So that if B is greater than 1, then making my light bulbs more efficient will increase my use of electricity, right?

Because B greater than 1 means my demand for lighting is really price sensitive. And if you've made lights more efficient, you've reduced the cost of lighting. And since my demand for lighting is price sensitive, you make lights more efficient. I use more lighting, OK.

Those are basically the two comments. The first is that relationship, it's kind of a function of this example. The second comes up in debates about auto efficiency standards. Because it makes sense, if I impose a miles per gallon standard on cars, I have made driving less expensive.

And if driving is really price sensitive, that could on balance increase the demand for gasoline. Now, this gets debated. This is called the rebound effect in talking about CAFE standards, auto fuel efficiency standards.

Most estimates say that it happens. But as far as CAFE standards are concerned, in effect, B is less than 1. The rebound effect is small. But it's there. Estimates show it to be there because driving is price sensitive. Price of gasoline goes up, people drive less. Cars get cheaper to drive, people drive more.

Not a big effect, not a huge effect, but it's there. OK so far? Very quiet today, OK. All the enthusiasm went into negotiations then, OK.

Obviously price isn't the only thing that matters. I find a very interesting graph. This is passenger kilometers traveled per capita on the vertical axis. It's a log, log scale. And GDP per capita on the horizontal axis, you will notice that the relation looks linear. What you can't see is that this is a scatter that includes all kinds of different countries, sub-Saharan Africa, Latin America, North America.

The universe is somewhere in this scatter over time. So this is probably North America. It points to a relatively constant elasticity. That's not driven, that's traveled. So you get airplanes at the top, and so forth.

There is this law, and I don't remember the parameter that says you spend x% your time traveling, whether you're a poor person walking or a rich person flying, the travel percentage turns out to not vary much with income. That's a big income effect. Let me just say that.

That's an income elasticity floating around 1, double income, double vehicle miles traveled. It doesn't necessarily mean double gasoline, doubling anything else, people travel other ways. But income affects energy demand, obviously, a lot.

What about the long run? Well, when I talked about the demand for lighting or the demand for driving or anything else, in that simple example, or in the example in the note, you hold the capital stock fixed. In the longer run, and here gasoline is the thing to think about, I'll change my car.

So in the short run, price of gasoline moves. Maybe I drive a little more. Maybe I drive a little less. Price of gasoline goes up by a multiple of 5 and stays there, let's say. I probably would change my car. I might want to consider changing where I live. I might want to rethink my use of public transportation.

I can do a lot of things in the longer run. I can certainly change my car and think about where I live. Do I really want to live in Weston as opposed to-- oh, you're not from here. Do I really want to live way out in the suburbs if gasoline is really expensive. Or do I want to think about moving into town. And if I'm locating a factory, and I'm thinking about, do workers like to drive very far, or an office building, I have to rethink where to put it.

So in the longer run, you expect more change. You expect whatever the response to an energy price changes in the short run, bip, bip, bip, you expect if it's maintained, the response will be longer. That will be larger. And I will come back to this. This is sort of a basic principle. Paul Samuelson was the first to really lay it out.

You have more flexibility in the long run, more things, more degrees of freedom. You expect a greater response to almost anything, income change, price change, whatever, if it's maintained over time. You expect the response to grow to a change that's maintained over time.

Part of that is because of the long-lived capital. But you will also recall that we talked about all kinds of path-dependent reasons. So let me just slip outside economics for just a slight moment and say that for the pure economic side, one reason why the long run response is greater than the short run response is because you change capital stock. And you don't change capital stock quickly.

At the societal level, the things that affect response may be policies, may be habits, may be culture. And they change very quickly. For Boston to look like Copenhagen as far as bicycle usage is concerned would not be easy or quick. They're having ferocious debates about how to set up the Longfellow bridge for bicycles versus cars. And that bridge hasn't been reconfigured in 100 years.

It will now have, I think, more bicycle space going forward. And if you've biked across that bridge, you will think that's a good thing because it's a little tricky-- so long run, short run. So let me now talk about how you estimate these things.

So that's theory. I'm a rational calculator, I'm Walmart, and I want to know my demand for electricity for heating and cooling. I don't care about electricity. I care about heating and cooling. I want to think through how my demand for electricity will be reflected by changes in the price.

And I will look at my production function, the HVAC system I've got. And I'll look at my demand function, how much do I care about store temperature and figure out what I want to do in terms of electricity. And then in the longer run, depending on the cost of electricity, I will insulate more or less. I will design stores in different ways. I have all kinds of response options.

So that's terrific for the decision maker. Suppose you want to estimate the demand for electricity, the demand for gasoline or anything else by looking at the market. You're a seller. You're a policymaker. You care about what all of that stuff ends up looking like at the end of the day when all the decisions shake out.

So a natural thing to do is to look at data. The first place this was done was in agriculture. People said well, you know, we'd like to know what the demand curve for wheat looks like, what the demand curve for corn looks like for planning purposes.

If you've seen the Department of Agriculture building in Washington, it was for a long time the largest building in the world. We spend a lot of effort studying agricultural markets. And when they first did this they found demand curves that sloped up. They graphed price, and they graphed quantity, and they got demand curves sloping up. This was disturbing. How could that happen. Yes, Marie.

AUDIENCE: Well, the demand curve could shift with the supply curve. And if they shift more easily, you might a bad-- your data is not going to follow the trend of specifically the demand curve.

RICHARD SCHMALENSEE: OK, close. You're saying that the data are going to be affected by the supply curve as well as the demand curve. In fact, suppose the supply curve doesn't shift at all. What are you going to get?

AUDIENCE: If the supply curve doesn't shift at all then you're going to get every change in-- like the equilibrium point is going to be a point along the demand curve.

RICHARD SCHMALENSEE: Think again. The supply curve doesn't shift, the demand curve does shift. What do you get?

AUDIENCE: You are going to trace out the supply curve.

RICHARD SCHMALENSEE: You're going to trace out the supply curve, which is what David was dying to say, right? Am I right? OK. We're going to move to thumbs, OK?

The classic picture, the classic picture, so you have what's come to be called an Identification problem. How do you figure out, how do you get to demand? I care about demand. I'm not getting demand. I'm getting supply.

And in fact, if that's the world, if all I know were price and quantity, and all I know is the demand curve shifts and the supply curve doesn't, I can't estimate demand. That's to say, it's not identified. I could gather all the data in the world from that setting. And I could not learn about demand from the data. Demand curve is not identified.

So what you need-- and Marie was on to it. What you need is you've got to move the supply curve. There has to be something that affects supply. If you've got something that affects supply, and you've got something that affects demand and they both move around, then if you work through the example in the teaching note, you can, in principle estimate both of them. There are nice techniques. This is a lot of fun.

But I mention this not because this course is about learning these techniques, which it's not, but it is about being skeptical when people tell you, well, we plotted price and quantity and we got-- here's our demand curve. Whenever you're dealing with a demand supply situation, you have to ask, how do you know that you're actually getting demand not supply not some mixture? Right? You can have both demand and supply shifting, and you could get a scatter of points that slopes down or is flat. You could say, well, you know, demand is very price elastic because the scatter shows this. Well, it doesn't.

This is just a caution. Because you will see in your life, I expect, with reasonable probability, estimates of the demand for this and the demand for that. And question one, is how do you know you've got demand? What did you do to deal with this identification problem? What did you do to deal with the fact that demand and supply are simultaneously determined? OK, clear, simple, classic stuff, OK.

Now, I talked about the long run and the short run. So how do we do that in practice? Well, there's a simple way that's pretty commonly used. And it's called a partial adjustment model.

And the equation let's say, E is the demand for energy. And here all the things that determine it. E star is a function. E star is considered to be the long run demand function. And lambda is a constant between 0 and 1, OK?

So this is a common formulation. If you think about it and you fiddled with difference equations, or you can just substitute, suppose E star is constant. Suppose the demand for energy is constant because the prices and other things have changed, and they've changed permanently.

Then E is going to converge to E star. Again, you just start out at some E0 and just substitute. And you will see that if lambda is less than 1, the initial condition goes away and it converges to E star. So this is a model.

But in the short run, in the short run, the change in demand will be reduced by this multiple lambda, right? So suppose doubling income doubles E star. OK, that's the long run response. If lambda is 0.5, then doubling income today will by itself only increase energy by 50% of E star, right?

So it's a smoothing. You could think about this as smoothing as signal detection is whatever you like. But the effective lambda less than 1 is that the response in the short run is less than the response in the long run.

So you estimate that. You assume a function that involves everything else you can think of. You use some technique to deal with the fact that price is determined along with quantity, along with E and estimate a function of that form, pretty commonly done, pretty commonly done.

I'm going to give you some actual estimates in a minute. They tend to come from a framework like that where you say, it's a simple way of saying if a change is maintained over time, the response will be greater than if the change happens today and gets reversed tomorrow. It has that property for lambda between 0 and 1, OK?

Yeah, that's all straightforward. You can recover so to speak. You estimate a function like this that has lagged demand on the right as well. Its coefficient is an estimate of 1 minus lambda. That lets you recover lambda. That lets you get the long run function, right?

So if I want to know what the long run response is, I look at the lag dependent variable. And again, if you think about it, suppose this coefficient is 0.9. That means all else equal, E is going to tend to be pretty smooth, right?

The changes in E star are going to be highly damp. This will be 0.1. And E itself is going to be pretty smooth. But I can sort of undo that by saying, all right, I'll divide by 0.1 over here and get an estimate of the function that would govern response to permanent changes, OK?

There are other approaches. There's a note, the note cites this very interesting paper by Hill Huntington that looks at oil price changes. He notes that the price of oil at any time is equal to the sum of the maximum it's ever been. And the difference between what it is today and the maximum it's ever been. He puts both of these into an equation and finds that the maximum has the largest negative effect on demand for gasoline.

In his argument-- so it's not this. It's a different adjustment mechanism. He says people get a sense of what kind of cars they ought to have and how they ought to heat their houses based on how high oil or gasoline has been. So the previous maximum has a very powerful shock impact.

We do remember gasoline prices over $4 a gallon. You may be able to say what it is right now, I can't. So P max the argument is has a large effect. I'm not sure that's right. I haven't seen any other study that tested that effectively or that reproduced it. But it it's an interesting approach. Most of the time, again, the intuition, because all this stuff is durable, all these assets are durable, changes take a long time, OK.

Some results, this is something that people have studied a lot, the demand for various forms of energy. And I will have assigned a paper or two. But they all go properly into the deep weeds, technically, and in terms of data. So I'm just going to give you a summary. The first summary point is the difference is normally huge between the short run and the long run.

So to go back to that other formulation, it's as if lambda is generally quite small in models like that, estimated to be quite small. Second, short run income and price elasticity, say a year, is generally less than 0.5. So if you double price, you will reduce consumption by less than half. I mean generally, you will see a good deal less, double income price consumption will increase by less than half.

The results seem to depend on the economy, rich countries, poor countries. That might make some sense, density. You can think of all kinds of things that might affect how people respond to changes in price and income. A lot of the studies have been of gasoline and electricity. And here's some of the ranges for elasticities.

Income elasticity is the percentage increase in demand by 1%-- small 1%-- increase in income. So income elasticity of 1 means if I double your income, you double your demand. The short run elasticities for gasoline-- yeah, the short run price elasticities for gasoline and electricity tend to be quite low.

Just looking at the literature, I think that the studies tend to cluster more toward the low end of those ranges. So price elasticity of 0.1 means I double the price. Consumption goes down by 10%. 0.2 means 20%. So these are small responses. These are low elasticities.

Notice there's a big difference-- factor of 2, 2 and 1/2 there. Bigger in gasoline than electricity, which makes a little more sense but only a little. The income elasticity of electricity I'm going to return to. In the long run, that's tended to be pretty high. Most estimates say, as people get richer, they use more electricity. And it's plus or minus proportional to income. You add appliances, bigger house. All the things that take electricity tend to go up with income. And that affect is not small.

In the short run, well, that's pretty small, right? I mean. My income goes up significantly in the short run. What do I use electricity for? It's mostly related. These are residential, mostly residential. There's relatively little I can change until I start changing assets.

So it makes sense that again, there's a big long run, short-run gap. What's interesting here is the long-run elasticities are pretty big historically. Yeah, Andrew.

AUDIENCE: Do these numbers-- they seem to take on something from the '70s to now.

RICHARD SCHMALENSEE: No, I'm going to come right to that. The width of those ranges ought to tell you that these are not constants of nature. This is a summary of human behavior, human behavior changes. So they do change over time.

And the problem, there is a piece I could have assigned because it's pretty simple. There's a piece that's argued, there's been a big change in the price elasticity of demand for gasoline in the '80s. And I'm trying to remember whether it goes up or down. So people have looked at this, but not much, not much.

You just see different studies from different periods and in different countries getting different results. And you say, well, maybe that's a change in reality or maybe that's a different method, right? It depends on data. It depends on, can you study households? Is this aggregate data? What are you working with?

But I want to make a couple of points about this stuff being constants of nature. This is one of the interesting puzzles. This is electricity demand. And if you look at electricity use in the United States before about 1970, you get sort of the top two rows.

You get pretty good growth in GDP. And these are not per capita. So this is total electricity use, total GDP. That's pretty good annual growth rates, 3 and 1/2, 4, and very big growth rates in electricity.

So the conventional wisdom was the income elasticity is greater than 1. And the price elasticity is small. Now, there really wasn't much historical experience you could base that on except the economy had grown. And if you go back farther, the figures get even bigger for growth in electricity.

But I didn't have everything back far enough. The experience was GDP grows pretty smoothly. Not much falloff in electricity in recessions. So we're looking here at the long run over decades. And prices do vary from time to time but nobody can see an effect.

So we come now to the '70s. And there's an oil shock. And what happens is a consequence of the oil shock is all kinds of inputs to electricity go up. And instead of declining faster than other prices, the price of electricity is rising faster than other prices.

So what do you think happened to the total consumption of electricity? How many think it went up by 1% to 2%? 2% to 3%? Thank you, Julian. 3 to 4? How many think we shut off the electricity? 4 to 5. Oh, come on. You've got to vote. 4 to 5, all right. It went up by 4.

So you immediately say, wow, income growth actually wasn't that bad over that period. There must be more price sensitivity than we thought. We got it wrong. And we were fooled by the fact that we didn't have much price variation. Prices moved and growth slowed sharply from those earlier numbers.

So we come now to the decade of the '80s. Income growth picks up. Prices start going down again. What happens to electricity growth? Is it above 4%? How many think it's above 4%?

How many think it's below 4%? Some of you think it's an imaginary number, I'm missing votes here, OK. Those who voted say it's above 4%. And yet it wasn't.

So now you say, how are people making these decisions? There's clearly some change that's happened on the demand side, right? I mean, something's different here.

Is it that people were shocked by the high prices they encountered and the price increases they encountered in the '70s and for the first time they thought about not buying every electrical device? And so you got slower growth, even though incomes were growing and prices were declining a bit-- reasonable hypothesis. You can test it. Is there any price and income effect?

Growth picks up in the '90s. Prices come down in the '90s for electricity and other things as well. What happens to electricity demand? Now, you've got to vote folks. This is no fun watching people go, oh, geez, I don't know.

How many think its above 3%? OK, how many people think is below 3%? OK, the smart people win. It's below 3%. Why is it below 3%? Income growth picked up. Prices started declining pretty reasonably. No one quite knows. You know.

AUDIENCE: A little earlier on in the years people were buying more appliances like refrigerators and TVs. And then as you get into the later years, everybody already has those appliances. But they were getting more efficient. And so they're replacing their existing appliances with more efficient appliances and so electricity use didn't grow.

RICHARD SCHMALENSEE: Maybe, maybe, yeah, it's not implausible. You could test it, right? You could look at purchase patterns of various appliances. This is national use. You're getting the impact of probably a decline of manufacturing. This isn't all residential.

So you're getting less industrial load. And you may also be getting more efficient new construction in office buildings. People began to think about efficiency. They didn't think much about efficiency. But when you start to think about efficiency, you begin to get gains in efficiency.

So I'm guessing its efficiency. It's not that people stopped wanting energy services, but it's kind of implausible that suddenly you know, oh, we're going to go back to small black and white televisions because the big colors use too much electricity. And I'm willing to bet refrigerators didn't shrink over this period, so probably efficiency. But I can tell you, it's a lot easier to talk about it after the fact than it was to see it, to see it coming on the basis of historical experience.

Oh, I should take this after '07. So this was not a great decade. And I won't trouble you more. This is not a great decade. Income growth slowed. And if I took it to the end of the decade, it's really ugly.

Let me not run into the recession. But even before the recession, income growth slowed. Energy prices began to go up. And as you might expect, electricity use slowed as well.

Forecasts out to about 20, 30 show less than 1% growth a year. So the business has changed from one where-- and even with reasonable economic growth assumptions-- where 10% and bigger than that before 1950, 5%, 10% to maybe 1%. Maybe 1% over the next 20 years, OK?

That should have come up earlier. That's one caution. Demand functions are not constants of nature. And if anybody says the elasticity is x, it probably means the elasticity has been x. Or we've got a pretty good study that estimates that the elasticity is x. These are not constants of nature.

And now let me sort of lead into the intro to Wednesday. The earlier discussion of derived demand and long run, short run and how you estimate sort of assumed in the background the standard economist model of rational behavior. Comes now, McKinsey.

And you've got a very short McKinsey assignment on the readings. McKinsey has made a lot of money, basically country to country saying, how much energy conservation can be free? How can you do it for free? And they do curves like this. I forget what this one is for.

But this is potential gigatons of carbon. But they do it for energy, as well. Gigatons of carbon dioxide reductions per year and the cost. And this is all negative. This is all better than free.

And they have it for energy consumption. And they have it for a bunch of countries. As I say, they've made a lot of money. But they're not the only ones. Let me not pick on McKinsey, as much fun as that is to do.

This is called the efficiency paradox. And it's been around since the '70s that there are lots of studies that say, look, you would save money if. Now, we're going to take a look at reasons why this persists.

But I will give you an example I encountered personally. I came back from Washington, EPA had just started its Energy Star program. I went to people at MIT, and I said, why not sign up for it?

Why doesn't MIT sign up for the Energy Star program because A, you get PR points. And B, all you promised to do is to use more efficient lighting when it saves you money. You promise to use more efficient lighting when it saves you money.

I could not get that sale made. And nobody's behaving irrationally, and nobody's being a jerk. And we will come back to it later when we talk about biodiesel at MIT. Because it's exactly the same phenomenon. Why did it take them three years to sell biodiesel at MIT.

So we'll talk about that set of reasons why this persists. But let me talk about some others. I mean, does anyone know who Amory Lovins is? OK, Amory Lovins has been an efficiency guru, an energy efficiency guru since the 1970s.

Amory Lovins, he heads the Rocky Mountain Institute. He lives in a house that has quintuple pane glass and highly insulated. And I think uses no outside energy. And Amory, yeah, you know Amory. OK.

AUDIENCE: Yeah, well, He has a super efficient home. But he can't run the computer off solar power so he does actually-- office equipment he has to get outside help for.

RICHARD SCHMALENSEE: I'm really disappointed right there. Yeah, I'll tell you. But I was once lobbied by him in Washington. And he came in and explained how there are all these triple-pane glass and all these wonderful energy saving devices that would more than pay for themselves and reduce US energy consumption.

And a guy who was working for me looked at him and said, you are saying consumers save money doing this stuff. He said, yes, that's right. And my guy's response was, that's fabulous. It means we don't have to have a government program.

There was a pause. We'll come back to again. There ought to be a pause. But it was a great line. It was a great line. He hadn't thought of that. He has by now.

So why? Well, one set of answers in the '70s was these. And I don't mean to demean engineers by saying engineering estimates. They're sort of back of the envelope, typical building, typical household kind of analysis that a lot of the early work was just too optimistic.

It just assumed, you know, you'd throw out the furnace. Well, you don't throw out the furnace. And you don't save money by throwing out a five-year-old furnace and replacing it early. So there were a lot of critiques that said, these estimates are flawed. Well, they might be flawed. But there have been lots of clear examples later. The MIT lighting example was an absolutely dead clear example of money on the sidewalk, money on the sidewalk.

So the note sites, and I think I put on the list the Hausman-Joskow paper. And I say, well, if you look at consumer decision making-- and Hausman did a classic study of air conditioners-- consumers act like they have very high discount rates. So you say, can you rationalize why consumers don't spend a little more for energy efficient air conditioners? This was using data well before they were all really efficient with standards. How do you rationalize that?

Well, you rationalize it by saying, consumers have limited access to capital. They don't want to borrow to buy an air conditioner. So they act as if their discount rate is 20%, 30% a year, which means another way to think about it, I need a really quick payback if I'm going to put money up front in energy efficiency.

Well, two points, and I'll come to the second point in just a minute. The first point is, that doesn't explain business behavior. The 20% to 30% discount rate, why would a business do that? If you can raise capital at let's say a canonical rate in the old days of 10%, and you got an investment in efficiency that has a 20% return, why wouldn't you make it?

You can see the household might have trouble raising capital to insulate the house or to buy air conditioning, to buy a more efficient air conditioner. But Walmart, why wouldn't Walmart do it? Well, Walmart probably does.

But why did we seem to see money lying on the ground for commercial and industrial demand? And we'll see a case later in the course of an energy saving opportunity being confronted by a perfectly rational profit maximizing firm, high rate of return that is a very hard sell. And we'll see biodiesel was a very hard sell. So you got to have other explanations. The high discount rate doesn't make sense for businesses. Yeah.

AUDIENCE: Yes, would it make more sense for the utility firms to kind of take charge of going and trying to make buildings more efficient in that sense? They're going to be more aligned in terms of the profit they can make because they'll be using less fuel for the same sort of benefit they are going to provide for the people.

RICHARD SCHMALENSEE: Well, we're going to come back to that. It depends on, to get to that, you have to have a diagnosis of the problem So if the problem is ignorance, which is a common argument, then you say the utility has better information. And so it could probably do this.

Although, the utilities expertise is generating electricity, not cooling buildings, right? They can acquire the expertise, but so could some third party. And some firms do, right? People do go around. There are a number of firms, I'm involved with one of them that is trying to make a business out of selling energy efficiency to households. It is not an easy business, not an easy business.

But the problem with utilities though-- and we'll come back to this for sure-- they get compensated for selling electricity. So they have an inherent conflict. If I can find a way not to rely on somebody that's conflicted, I will.

If you say, I'm paying you for selling electricity now, I'd like you to go around and tell people how to consume less, I'd begin to wonder how effective you're going to be. I can make you do it. I can make you go around. And states do this, but the effectiveness varies because your motivation isn't necessarily terrific.

So we've got this low hanging fruit problem. And we're going to see this in a number of settings that how come opportunities to increase efficiency that are apparently profitable don't get picked up. But the world is even more complicated. Sometimes people over-invest.

Let me talk about the Prius. How many people or their parents have Prius, or relatives have Prius? Yeah, lots of people have Prius. And I'm not going to criticize your rationality in what I'm saying here, OK? OK.

So three million people have bought them and they sold three million worldwide. They've sold over a million in the US. I'm now relying on others who have looked hard at this.

But they tell me that in terms of driving, in terms of interior amenities, the styling is different. But in terms of interior design, it's a Corolla. But it's a Corolla that gets 50 miles a gallon instead of 28 miles a gallon.

So let's ask the question, is it worth the difference in upfront cost? Did these three million people make a good Investment? Remember, these are people who are not buying energy efficient air conditioners. Well, maybe they are, but not that many others are. So let's look at it.

We're going to ask if the cost difference-- so you save money on gasoline. A cost difference is the upfront cost difference. S is the annual savings on gasoline. R is the interest rate. And T is, how long do you have to keep the car until the savings on gasoline justify the purchase price. OK, so that's just a discounting exercise.

If there's no interest rate, you just divide the upfront cost by the annual savings and that gives you the number of years. That's 0 interest rate. If there is a positive interest rate, then you ask, you discount S per year out to T. And you say, how large does T have to be so that discounted savings equals the original cost.

Clearly if the savings are small, there can be no solution. There can be no positive solution to this. But if RC is less than S, it's this simple formula you can get in less than 5 seconds. So let's look at it.

Toyota, these are data a few years old. And they don't count the subsidies that were there for the first-- in the US I think it was 60,000 units sold. But the subsidies have long since run out. So the average premium, 2008, maybe was $7,450. Assume maintenance is the same. It may not be. Prius may be cheaper.

Car's average about 12,000 miles a year. I can work out the annual cost. The difference, the savings is $566. So if the interest rate is 0, you've got to keep the car for 13.2 years to break even.

Now again, if you have a subsidy or if the gasoline is more expensive, this will get smaller. But that's kind of a long time. Cars tend to turn over after about 10, really. And most people don't say, hold it for 10. 10 is when they go out of the fleet.

So if you're going to resell, you've got to sell to somebody who also loves Prius. If the interest rate is 5%, not that high, you've got to hold the car for 21.6 years before those gasoline savings pay for the cost difference. If the interest rate is above 7.6%, there is no break even. You can verify that you don't satisfy that. So you can hold it forever and the discounted gasoline savings don't equal the original purchase.

Now, you can play with it. You can put in a per vehicle subsidy. You can make gasoline more expensive. But these don't seem like crazy numbers in light of how those cars work.

And this is the puzzle. People act when they're buying air conditioners like their discount rate is 20% to 30%. And they act when they're buying a Prius like they have a really low discount rate. Or maybe this is not the right way to look at it.

Tell no other economists I said this. But it may be that this sort of economic approach I was taking here is missing stuff. It's missing a lot. Why do people buy a Prius by the way? Do people buy a Prius to save gasoline? Scot.

AUDIENCE: A lot of people just buy environmental friendly. And they want to do something and something like that. And so even if it might not be saving them that much money or saving them money at all, they still do it.

RICHARD SCHMALENSEE: They want to do the right thing. There's another argument that people make is, there's a reason the Prius doesn't look like the Corolla, so not only do you feel satisfied that you're doing a good thing, people can see you're doing a good thing.

Because I mean, that's a funny looking car to be honest. And people know you're driving a Prius. There are other hybrids you don't people are driving. And they don't sell as well as the Prius. None of that's here. The do the good thing isn't there. The other people see I'm doing a good thing isn't there.

This is not the drive demand for transportation services story that gets the Prius sold. This is other things. And we're going to hear from professor Silbey on Wednesday about some of the other things. Have you questions or comments. Paper going OK. Obena.

AUDIENCE: Did they consider the time saved at the gas pump?

RICHARD SCHMALENSEE: Hmm, time saved at the gas pump?

AUDIENCE: Right.

RICHARD SCHMALENSEE: I don't know the size of the tank on that thing. Does it save you time? It probably does.

AUDIENCE: Full time, so that would be two or three weeks.

RICHARD SCHMALENSEE: These are people who value their time very highly. I have to say, I value mine really highly. And it never occurred to me. But it could be. It could be you people all bill out at $1,000 an hour. And a few minutes is worth saving. Any other questions, comments. David.

AUDIENCE: Yes. You could also regulate a small town than the average driven mile? So somebody drives less in 2003.

RICHARD SCHMALENSEE: That might be low. So maybe it's segments that way. Also testable. Sarah, does your family drive a lot?

AUDIENCE: No, what about risk aversion and price fluctuations in gas?

RICHARD SCHMALENSEE: You figure this is kind of an insurance policy that when gas prices move, we won't be hit too hard. That's another possibility, also testable, yeah.

AUDIENCE: I read an article that production of the Prius is actually more inefficient than the production of the Corolla. And the energy that people claimed they are saving actually because of the inefficiency of production, you know, they established it doesn't actually save energy.

RICHARD SCHMALENSEE: So they think they're doing a good thing but they're not. It's a terrible tragedy. Bring your name cards next time.