[SQUEAKING] [RUSTLING] [CLICKING]

JONATHAN GRUBER: Today we're going to start talking about what's underneath the demand curve. So basically, what we did last time, and what you did in section on Friday is talk about sort of the workhorse model of economics, which is supply and demand model. And we always start the class with that, because that's the model in the course.

But I think as any good sort of scientists and inquisitive minds, you're probably immediately asking, well, where do these supply and demand curves come from? They don't just come out of thin air. How do we think about them? Where do they come from? And that's what we'll spend basically the first 1/2 of the course going through.

And so we're going to start today with the demand curve, and the demand curve is going to come from how consumers make choices, OK? And that will help us drive the demand curve. Then we'll turn next to supply curve, which will come from how firms make production decisions.

But let's start with the demand curve, and we're going to start by talking about people's preferences, and then the utility functions, OK?

So our model of consumer decision making is going to be a model of utility maximization. That's going to be our fundamental-- remember, this course is all about constrain maximization. Our model today is going to be a model of utility maximization.

And this model's going to have two components. There's going to be consumer preferences, which is what people want, and there's going to be a budget constraint, which is what they can afford. And we're going to put these two things together. We're going to maximize people's happiness, or their choice-- or their happiness given their preferences, subject to the budget constraint they face.

And that's going to be the constraint maximization exercise that actually, through the magic of economics, is going to yield the demand curve, and yield a very sensible demand curve that you'll understand intuitively.

Now, so what we're going to do is do this in three steps. Step one-- over the next two lectures. Step one is we'll talk about preferences, how do we model people's tastes. We'll do that today.

Step two is we'll talk about how we translate this to utility function, how we mathematically represent people's preferences in utility function. We'll do that today as well.

And then next time, we'll talk about the budget constraints that people face. So today, we're going to talk about the max demand. Next time we'll talk about the budget constraint.

That means today's lecture is quite fun. Today's lecture is about unconstrained choice. We're not going to worry at all about what you can afford, what anything costs. We're not going to worry about what things cost. We're not going to worry about what you can afford, OK?

Today's the lecture where you won the lottery, OK? You won the lottery. Money is no object. How do you think about what you want, OK?

Next time, we'll say, well, you didn't win the lottery. In fact, as we learn later in the semester, no one wins the lottery. It's an incredibly bad deal. But next time, we'll impose the budget constraints. But for today, we're just going to ignore that and talk about what do you want, OK?

And to start this, we're going to start with a series of preference assumptions. A series-- remember, as I talked about last time, models rely on simplifying assumptions. Otherwise, we could never write down a model. It'll go on forever, OK?

And the key question is, are those simplifying assumptions sensible? Do they do violence to reality in a way which makes you not believe the model, or are they roughly consistent with reality in a way that allows you to go on with the model? OK?

And we're going to pose three preference assumptions, which I hope will not violate your sense of reasonableness. The first is completeness. What I mean by that is you have preferences over any set of goods you might choose from.

You might be indifferent. You might say, "I like A as much as B," but you can't say, "I don't care," or, "I don't know." You can say, "I don't care." That's indifference. You can't say, "I don't know." You can't literally say, "I don't know how I feel about this."

You might say you're indifferent to two things, but you won't say, "I don't know how I feel about something." That's completeness, OK?

The second is the assumption we've all become familiar with since kindergarten math, which is transitivity. If you prefer A to B and B to C, you prefer A to C, OK? That's kind of-- I'm sure that's pretty clear. You've done this a lot in other classes.

So these two are sort of standard assumptions you might make in any math class. The third assumption is the one where the economics comes in, which is the assumption of nonsatiation or the assumption of more is better. In this class, we will assume more is always better than less, OK? We'll assume more is better than less.

Now, to be clear, we're not going to say that the next unit makes you equally happy as the last unit. In fact, I'll talk about that in a few minutes. Well, in fact, the next unit makes you less happy. But we will say you always want more, that faced with the chance of more or less, you'll always be happier with more, OK? And that's the nonsatiation assumption, OK?

And I'll talk about that some during the lecture, but that's sort of what's going to give our models their power. That's a sort of new economics assumption. That's going to give-- beyond your typical math assumptions-- this is going to give our models their power, OK? So that's our assumptions.

So armed with those, I want to start with a graphical representation of preferences. I want to graphically represent people's preferences, and I'll do so through something we call indifference curves. Indifference curves, OK?

These are-- indifference curves are basically preference maps. Essentially, indifference curves are graphical maps of preferences, OK?

So for example, suppose your parents gave you some money to begin the semester, and you spent that money on two things. Your parents are rich. They gave you tons of money. You spent your money on two things, buying pizza or eating cookies, OK?

So consider preferences between pizza and cookies. That's your two things you can do. Once again, this is a constrained model. Obviously, in life, you can do a million things with your money.

But it turns out, if we consider the contrast between doing two different things with your money, you get a rich set of intuition that you can apply to a much more multi-dimensional decision case. So let's start with a two dimensional decision case. You've got your money. Either you can have pizza or you can have cookies, OK?

Now, consider three choices, OK? Choice A is two pizzas and one cookie. Choice B is one pizza and two cookies, and choice C is two pizzas, two cookies. OK, that's the three packages I want to compare.

And I am going to assume-- and I'll mathematically rationalize in a few minutes-- but for now, I'm going to assume you are indifferent between these two packages. I'm going to assume you're equally happy with two slices of pizza and one cookie or two cookies and one slice of pizza, OK? I'm going to assume that.

But I'm also going to assume you prefer option C to both of these. In fact, I'm going to assume that, because that is what more is better gives you, OK? So you're indifferent between this.

This indifference doesn't come from any property I wrote up. That's an assumption. That's just-- for this case, I'm assuming that. This comes to the third property I wrote up there. You prefer package C because more is always better than less, OK?

So now, let's graph your preferences, and we do so in figure 2-1, OK, in the handout. OK, so here's your indifference curve. So we've graphed on the x-axis your number of cookies, on the y-axis slices of pizza, OK?

Now, you have-- we've graphed the three choices I laid here, choice A, which is two slices of pizza and one cookie, choice B, which is two cookies and one slice of pizza, and choice C, which is two of both. And I've drawn on this graph your indifference curves. The way your indifference curves looks is there's one indifference curve between A and B, because those are the points among which you're indifferent.

So what an indifference curve represents is all combinations of consumption among which you are indifferent. That's why we call it indifference curve. So an indifference curve, which will be sort of one of the big workhorses of this course, an indifference curve represents all combinations along which you are in different.

You're indifferent between A and B. Therefore, they lie on the same curve, OK? So that's sort of our preference map, our indifference curves.

And these indifference curves are going to have four properties, four properties that you have to-- that follow naturally from this-- it's really three and 1/2. The third and fourth are really pretty much the same, but I like to write them out as four. Four properties that follow from these underlying assumptions--

Property one is, consumers prefer higher indifference curves. Consumers prefer higher indifference curves, OK? And that's just all from more is better. That is, an indifference curve that's higher goes through package that has at least as much of one thing and more of the other thing. Therefore, you prefer it, OK?

So as indifference curve shifts out, people are happier, OK? So on that higher indifference curve, point C, you are happier than points A and B, because more is better, OK?

The second is that indifference curves never cross. Indifference curves never cross, OK? Actually, that's third, actually. I want to come to that in order. Second-- third is the indifference curves never--

Second is indifference curves are downward sloping. Second is indifference curves are downward sloping. Indifference curves are downward sloping. Let's talk about that first, OK?

That simply comes from the principle of nonsatiation. So look at figure 2-2. Here's an upward sloping indifference curve, OK? Why does that violate the principle of nonsatiation? Why does that violate that? Yeah.

AUDIENCE: Either, if you're-- either you're less happy with you have more cookies, or you're less happy if you have more pizza. And like there's-- and that violates nonsatiation.

JONATHAN GRUBER: Exactly. So basically, you're indifferent-- on this curve, you're indifferent with one of each and two of each. You can't be indifferent. Two of each has got to be better than one of each. So an upward sloping indifference curve would violate nonsatiation. So that's the second property of indifference curve.

The third property of indifference curve is the indifference curves never cross, OK? We could see that in figure 2-3, OK? Someone else tell me why this violates the properties I wrote up there, indifference curves crossing. Yeah.

AUDIENCE: Because B and C [is strictly better.

JONATHAN GRUBER: What's that?

AUDIENCE: Because B and C, B is strictly better.

JONATHAN GRUBER: Because the B and C, B is strictly better. That's right.

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: But they're also both on the same curve as A. So you're saying they're both-- you're indifferent with A for both B and C, but you can't be, because B is strictly better than C. So it violates transitivity, OK? So the problem with crossing indifference curves is they violate transitivity.

And then finally, the fourth is sort of a cute extra assumption, but I think it's important to clarify, which is that there is only one indifference curve through every possible consumption bundle, only one IC through every bundle.

OK, you can't have two indifference curves going through the same bundle, OK? And that's because of completeness. If you have two indifference curves going through the same bundle, you wouldn't know how you felt, OK? So there can only be one going through every bundle, because you know how you feel.

You may feel indifferent, but you know how you feel. You can't say I don't know, OK? So that's sort of a extra assumption that sort of completes the link to the properties, OK?

So that's basically how indifference curves work. Now, I find-- when I took this course, before you were-- god, maybe before your parents were born, I don't know, certainly before you guys were born-- when I took this course, I found this course full of a lot of light bulb moments, that is, stuff was just sort of confusing, and then boom, an example really made it work for me.

And the example that made indifference curves work to me was actually doing my first UROP. When my UROP was with a grad student, and that grad student had to decide whether he was going to accept a job. He had a series of job offers, so he had to decide.

And basically, he said, "Here's the way I'm thinking about it. I am indifferent-- I have an indifference map and I care about two things. I care about school location and I care about economics department quality. I care about the quality of my colleagues, and the research it's done there, and the location."

And basically, he had two offers. One was from Princeton, which he put up here. No offense to New Jerseyans, but Princeton as a young single person sucks. OK, fine when you're married and have kids, but deadly as a young single person.

And the other-- so that's Princeton. Down here was Santa Cruz. OK, awesome. [INAUDIBLE] is the most beautiful university in America, OK? But not as good an economics department. And he decided he was roughly indifferent between the two.

But he had a third offer from the IMF, which is a research institution in DC, which has-- he had a lot of good colleagues, and DC is way better than Princeton, New Jersey, even though it's not as good as Santa Cruz. So he decided he would take the offer at the IMF, OK?

Even though the IMF had worse colleagues than Princeton and worse location than Santa Cruz, it was still better in combination of the two of them, given his preferences. And that's how he used indifference curves to make his decision, OK? So that's sort of an example of applying it. Once again, no offense to the New Jerseyans in the room, of which I am one, but believe me, you'd rather be in Santa Cruz.

OK, so now, let's go from preferences to utility functions. OK, so now, we're going to move from preferences, which we've represented graphically, to utility functions, which we're going to represent mathematically.

Remember, I want you understand, everything this course at three levels, graphically, mathematically, and most importantly of all, intuitively, OK? So graphic is indifference curves. Now we come to the mathematical representation, which is utility function, OK?

And the idea is that every individual, all of you in this room, have a stable, well behaved, underlying mathematical representation of your preferences, which we call utility function. Now, once again, that's going to be very complicated, your preference over lots of different things.

We're going to make things simple by writing out a two dimensional representation for now of your indifference curve. We're going to say, how do we act mathematically represent your feelings about pizza versus cookies? OK? Imagine that's all you care about in the world, is pizza and cookies. How do we mathematically represent that?

So for example, we could write down that your utility function is equal to the square root of the number of slices of pizza times the number of cookies. We could write that down. I'm not saying that's right. I'm not saying it works for anyone in this room or even everyone this room, but that is a possible way to represent utility, OK?

What this would say-- this is convenient. We will use-- we'll end up using square root form a lot for utility functions and a lot of convenient mathematical properties. And it happens to jive with our example, right?

Because in this example, you're indifferent between two pizza and one cookie or one pizza and two cookie. They're both square root of 2. And you prefer two pizza and two cookies. That's two, OK? So this gives you a high utility for two pizza and two cookies, OK, than one pizza and two cookie, or two pizza and one cookie.

So now, the question is, what does this mean? What is utility? Well, utility doesn't actually mean anything. There's not really a thing out there called utiles OK?

In other words, utility is not a cardinal concept. It is only an ordinal concept. You cannot say your utility, you are-- you cannot literally say, "My utility is x% higher than your utility," but you can rank them.

So we're going to assume that utility can be ranked to allow you to rank choices.

Even if generally, we might slip some and sort of pretend utility is cardinal for some cute examples, but by and large, we're going to think of utility as purely ordinal. It's just a way to rank your choices.

It's just when you have a set of choices out there over many dimensions-- like if your choice in life was always over one dimension and more was better, it would always be easy to rank it, right? You'd never have a problem. Once your choice is over more than one dimension, now if you want to rank them, you need some way to combine them.

That's what this function does. It allows you essentially to weight the different elements of your consumption bundle, so you can rank them when it comes time to choose, OK?

Now, this is obviously incredibly simple, but it turns out to be amazingly powerful in explaining real world behavior, OK? And so what I want to do today is work with the underlying mathematics of utility, and then we'll come back. We'll see in the next few lectures how it could actually be used to explain decisions.

So a key concept we're going to talk about in this class is marginal utility. Marginal utility is just a derivative of the utility function with respect to one of the elements. So the marginal utility for cookies, of cookies, is the utility of the next cookie, given how many cookies you've had.

This class is going to be very focused on marginal decision making. In economics, it's all about how you think about the next unit. Turns out, that makes life a ton easier. Turns out, it's way easier to say, "Do you want the next cookie," than to say, "How many cookies do you want?"

Because if you want the next cookie, that's sort of a very isolated decision. You say, OK, I had this many cookies. Do I want the next cookie? Whereas before you start eating, if you say, how many cookies do you want, that's sort of a harder, more global decision. So we're going to focus on this stepwise decision making process of do you want the next unit, the next cookie, or the next slice of pizza, OK?

And the key feature of utility functions we'll work with throughout the semester is that they will feature diminishing marginal utility. Marginal utility will fall as you have more of a good. The more of a good you've had, the less happiness you'll derive from the next unit, OK?

Now, we can see that graphically in figure 2-4. Figure 2-4 graphs on the x-axis the number of cookies holding constant pizza. So let's say you're having two pizza slices, and you want to say, what's my benefit from the next cookie? And on the left axis, violating what I just said like 15 seconds ago, we graph utility.

Now, once again, the utile numbers don't mean anything. It's just to give you an ordinal sense.

What you see here is that if you have 1 cookie, your utility is 1.4, square root of 2 times 1. If you have 2 cookies, your utility goes up to square root of 4, which is 2. You are happier with 2 cookies, but you are less happy from the second cookie than the first cookie, OK?

And you could see that in figure-- if you flip back and forth between 2-4 and 2-5, you can see that, OK? The first cookie, going from 0 to 1 cookie, gave you one-- so in this case, we're now graphing the marginal utility. So figure 2-4 is the level of utility, which is not really something you can measure, in fact. Figure 2-5 is something you can measure, which is marginal utility, what's your happiness-- and we'll talk about measuring this-- from the next cookie.

You see, the first cookie gives you a utility increment of 1.4, OK? You go from utility of 0 to utility of 1.4. The next cookie gives you utility increment of 0.59. OK, you go from utility of 1.41 to utility of 2.

The next cookie gives utility increment of 0.45, the square root of 3. So now we flip back to the previous page. We're going from the square root of 4, we're going from the square root of 4-- I'm sorry-- to the square root of 6. Square root of 6 is only 0.45 more than the square root of 4, and so on.

So each additional cookie makes you less and less happy. It makes you happier, it has to, because more is better, but it makes you less and less happy, OK?

And this makes sense. Just think about any decision life starting with nothing of something and having the first one, slice of pizza, a cookie, deciding on which movie to go to. The first movie, the one you want to see the most, is going to make you happier than the one you want to see not quite as much. The first cookie when you're hungry will make you happier than the second cookie. The first slice of pizza make you happier--

Now, you may be close to indifferent. Maybe the second slice of pizza makes you almost as happy as the first. But the first will make you happier, OK? If you think about-- that's really sort of that first step. You were hungry, and that first one makes you feel happier.

Now, but you got to remember, you always want more cookies. Now, you might say, "Wait a second. This is stupid. Once I've had 10 cookies, I'm going to barf. The 11th cookie can actually make me worse off, because I don't like barfing."

But in economics, we have to remember, you don't have to eat the 11th cookie. You can give it away. So if like say, you don't want the 11th cookie, you can save it for later. You can give it to a friend.

So you always want it. In the worst case, you throw it out. It can't make you worse off, it can only make you better off. And that's what our sort of more is better assumption comes from.

Obviously, the limit-- you know, if you get a million cookies, your garbage can gets full. You have no friends to give them to. I understand at the limit, these things fall apart, OK? But that's the basic idea of more is better and the basic idea of diminishing marginal utility. OK, any questions about that? Yeah.

AUDIENCE: Can the utility function ever be negative?

JONATHAN GRUBER: Utility function can never be negative because we have-- well, utility-- once again, utility is not an ordinal concept. You can set up utility functions such that the number is negative. You can set that up. OK, the marginal utility is always positive. You always get some benefit from the next unit.

Utility, once again, the measurement's relevant. So it could be negative. You could set it up-- I could write my utility function like this, you know, something like that. So it could be negative. That's just a sort of scaling factor.

But marginal utility is always positive. You're always happier, or it's non-negative. You're always happier or at least indifferent to getting the next unit. Yeah.

AUDIENCE: So when you're looking at 2-5, if you get like a fraction of a cookie, is the marginal utility still going to go up?

JONATHAN GRUBER: I'm sorry, you look-- figure 2-5-- no, the marginal is going to go down. Each fraction of a cookie, the marginal utility-- marginal utility is always diminishing.

AUDIENCE: So if you start with zero, and you get 1/2 a cookie based on this graph--

JONATHAN GRUBER: Well, it's really hard to do it from zero. That's really tricky. It's sort of much easier to start from one.

So corner solutions, we'll talk about corner solutions in this class, they get ugly. Think of it starting from one. Starting with that first cookie, every fraction of a cookie makes you happier, but less and less happy with each fraction. Good question. All right, good questions.

All right, so now, let's talk about-- let's flip back from the math to the graphics, and talk about where indifference curves come from. I just drew them out. But in fact, indifference curves are the graphical representation of what comes out of utility function, OK?

And indeed, the slope of the indifference curve, we're going to call the marginal rate of substitution, the rate essentially at which you're willing to substitute one good for the other. The rate at which you're willing to substitute cookies for pizza is your marginal rate of substitution.

And we'll define that as the slope of the indifference curve, delta p over delta c. That is your marginal rate of substitution. Literally, the indifference curve tells you the rate at which you're willing to substitute. You just follow along and say, "Look, I'm willing to give up--"

So in other words, if you look at figure 2-6, you say, "Look, I'm indifferent between point A to point B. One slice of pizza-- I'm sorry-- one cookie and four slices of pizza is the same to me as two cookies and two slices of pizza." Why is it the same? Because they both give me utility square root of four, right?

So given this mathematical-- I'm not saying you are. I'm saying, given this mathematical representation, OK, you are indifferent between point A and point B. So what that says-- and what's the slope with the indifference curve? What's the arc slope between point A and point B? The slope is negative 2.

So your marginal rate of substitution is negative 2. You are indifferent, OK? You are indifferent between 1, 4 and 2, 2. Therefore, you're willing to substitute or give away two slices of pizza to get one cookie. Delta p delta c is negative 2, OK?

Now, it turns out you can define the marginal rate of substitution over any segment of indifference curve, and what's interesting is it changes. It diminishes.

Look what happens when we move from two pizzas and two cookies, from point B to point C. Now the marginal rate of substitution is only negative of 1/2. Now I'm only willing to give up one slice of pizza to get two cookies. What's happening?

First, I was willing give up two slices of pizza to get one cookie. Now I'm only willing to give up one slice of pizza to get two cookies. What's happening? Yeah.

AUDIENCE: You don't want a cookie as much?

JONATHAN GRUBER: Because of?

AUDIENCE: Diminishing marginal utility.

JONATHAN GRUBER: Exactly. Diminishing margin utility has caused the marginal rate of substitution itself to diminish. For those who are really kind of better at math than I am, it turns out technically, mathematically, marginal utility isn't always diminishing. You can draw up cases.

MRS is always diminishing. So you can think of marginal as always diminishing. It's fine for this class. When you get to higher level math and economics, you'll see marginal utility doesn't have to diminish. MRS has to diminish, OK? MRS is always diminishing. As you go along the indifference curve, that slope is always falling, OK?

So basically, what we can right now is how the MRS relates to utility function. Our first sort of mind-blowing result is that the MRS is equal to the negative of the marginal utility of cookies over the marginal utility of pizza. That's our first key definition. It's equal to the negative of the marginal utility of the good on the x-axis over the marginal utility of the good on the y-axis, OK?

Essentially, the marginal rate of substitution tells you how your relative marginal utilities evolve as you move down the indifference curve. When you start at point A, you have lots of pizza and not a lot of cookies. When you have lots of pizza, your marginal utility is small.

Here's the key insight. This is the thing which, once again, it's a light bulb thing. If you get this, it'll make your life so much easier. Marginal utilities are negative functions of quantity. The more you have of a thing, the less you want the next unit of it.

That's why, for example, cookies is now in the numerator and pizza is in the denominator, flipping from this side, OK? The more you have a good, the less you want it.

So start at point A. You have lots of pizza and not a lot of cookies. You don't really want more pizza. You want more cookies.

That means the denominator is small. The marginal utility of pizza is small. You don't really want it. But the marginal utility of cookies is high. You want many of them. So this is a big number.

Now let's move to point B. Think about your next decision. Well, now, your marginal utility of pizza, if you were going to go from two to one slice of pizza, now pizza is worth a lot more than cookies. So now it gets smaller.

So essentially, as you move along that indifference curve, because of this, you want-- because of diminishing marginal utility, it leads this issue of a diminishing marginal rate substitution, OK?

So basically, as you move along the indifference curve, you're more and more willing to give up the good on the x-axis to get the good on the y-axis. As you move from the upper left to the lower right on that indifference map, figure 2-6, you're more you're more willing to give up the good on the x-axis to get the good on the y-axis.

And what this implies is that indifference curves are-- indifference curves are convex to the origin. Indifference curves are convex to the origin. That's very important.

OK, let's see, they are not concave. They're either convex or straight. Let's say they're not concave to the origin, to be technical. Indifference curves can be linear. We'll come to that. But they can't be concave to the origin.

Why? Well, let's look at the next figure, the last figure, figure 2-7. What would happen if indifference curves were concave to the origin?

Then that would say, moving from one pizza-- so now I've drawn a concave indifference curve. And with this indifference curve, moving from point A to point B leaves you indifferent. So you're happy to give up one slice of pizza to get one cookie. Starting with four slices of pizza and one cookie, you were happy to give up one slice of pizza to get one cookie.

Now, starting from two and three, you're now willing to give up two slices of pizza to get one cookie. What does that violate? Why does that not make sense? Yeah.

AUDIENCE: Law of diminishing marginal returns?

JONATHAN GRUBER: Yeah, law of diminishing marginal utility. Here, you were happy to have one slice of pizza to get one cookie. Now you are willing to have two slices of pizza to get one cookie, even though you have less pizza and more cookies.

That can't be right. As you have less pizza and more cookies, cookies-- pizza should become more valuable, not less valuable, and cookies should become less valuable, not more valuable. So a concave to the origin indifference curve would violate the principle of diminishing marginal utility and diminishing marginal rate of substitution, OK? Yeah.

AUDIENCE: What if it's like something like trading cards?

JONATHAN GRUBER: OK.

AUDIENCE: I mean, I mean, as you get more trading cards, you have-- you're already made a complete set.

JONATHAN GRUBER: That's very interesting. So in some sense, what that is saying is that your utility function is really over sets. You're saying your utility functions isn't over trading cards. It's over sets.

So basically, that's what's sort of a bit-- you know, our models are flexible. One way is to say they're loose. Another way is to say they're flexible.

But one of the challenges you'll face on this course is thinking about what is the decision set over which I'm writing my utility function? You're saying it's sets, not trading cards. So that's why it happens. Other questions? Good question. Yeah, at the back.

AUDIENCE: What about like addictive things, where like, the more you have it, the more you want to buy?

JONATHAN GRUBER: Yeah, that's a really relishing question. I spent a lot of my research life, actually-- I did a lot of research for a number of years on thinking about how you properly model addictive decisions like smoking.

Addictive decisions like smoking, essentially, it really is that your utility function itself shifts as you get more addictive. It's not that your marginal utility-- the next cigarette is still worth less than the first cigarette. It's just that as you get more addicted, that first cigarette gets worth more and more to you.

So when you wake up in the morning feeling crappy, that first cigarette still does more for you than the second cigarette. It's just, the next day you wake up feeling crappier, OK?

So we model addiction as something where essentially, each day, cigarettes do less and less for you. You get essentially adjusted to new-- you habituate to higher levels.

And this is why I do a lot of work-- you know, this is why, unfortunately, we saw last year, the number-- the highest number of deaths from accidental overdose in US history. 72,000 people died from drug overdoses last year, more than ever died in traffic accidents in our nation's history, OK?

Why? Because people get habituated to certain levels, and they get habituated to certain levels. So people get hooked on Oxycontin. They get habituated to a certain level. They maybe switch to heroin, and they habituate to a certain level.

And now there's this thing called fentanyl, which is a synthetic opioid brought over from China, which is incredibly powerful. And dealers are mixing the fentanyl in with the heroin. And the people shoot up, not realizing-- at their habituated level-- not realizing they have this dangerous substance, and they overdose and die.

And that's because they've got habituated to high levels. They don't realize they're getting a different product. So it's not about not diminishing marginal utility. It's about different-- underlying different products. All right? Other questions?

Sorry for that depressing note, but it's important to be thinking about that. That's why, once again, we're the dismal science. We have to think about these things.

OK, now, let's come to a great example that I hope you've wondered about, and maybe you've already figured out in your life, but I hope you've at least stopped and wondered about, which is the prices of different sizes of goods, in a convenience store, say.

OK, take Starbucks. You can get a tall iced coffee for 2.25, or the next size, whatever the hell they call it, bigger, OK? You can get, for 70 more cents-- so 2.25, and you can double it for 70 more cents. Or take McDonald's. A small drink is $1.22 at the local McDonald's, but for 50 more cents, you can double the size, OK?

What's going on here? Why did they give you twice as much liquid, or if you go for ice cream, it's the same thing. Why do they give you twice as much for much less than twice as much money? What's going on? Yeah.

AUDIENCE: Since your marginal utility is diminishing as you have more coffee available to you, you're willing to pay less for it, so they make the additional coffee cheaper.

JONATHAN GRUBER: Exactly. That's a great way to explain it. The point is it's all about diminishing marginal utility.

OK, when you come in to McDonald's on a hot day, you are desperate for that soda, but you're not as desperate have twice as much soda. You'd like it. You probably want to pay more for it, but you don't like it nearly as much as that first bit of soda.

So those prices simply reflects the market's reaction to understanding diminishing marginal utility. Now, we haven't even talked about the supply side of the market yet. I'm not getting to how providers make decisions. That's a much deeper issue. I'm just saying that this is diminishing marginal utility in action, how it works in the market, and that's why you see this, OK?

So basically, what you see is that that first bite of ice cream, for example, is worth more, and that's why the ice cream that's twice as big doesn't cost twice as much. Now, so basically, what this means is, if you think about our demand and supply model, on a hot day, or any day, the demand for the first 16 ounces is higher than the demand for the second 16 ounces. But the cost of producing 16 ounces is the same.

So let's think about this. It's always risky when I try to draw a graph on the board, but let's bear with me. OK, so let's say we've got a simple supply and demand model.

You have this supply function for soda, and let's assume it's roughly flat. OK, let's assume sort of the cost the firm proceeds within some range. The firm-- basically, every incremental 16 ounces costs them the same. So that's sort of their supply curve.

And then you have some demand curve, OK? You have some demand curve which is downward sloping, OK, and they set some price. And this is the demand for 16 ounces.

Now, what's the demand for the next 16 ounces, OK? Yeah, this isn't going to work. We have to have an upward-sloping supply curve. Sorry about that. We have a slightly upward sloping supply curve, OK?

Now we have the demand for the next-- so here's your price. Here's your $1.22, OK? Now, you say, "Well, what's my demand when I sell 32 ounces?" Well, it turns out demand doesn't shift out twice as much. It just shifts out a little bit more. So you can only charge $1.72 for the next 16 ounces.

Probably, if you want to go to the big-- if you go to 7-Eleven, where you can get sizes up to, you know, as big as your house, OK-- they keep these curves keep getting closer and closer to each other. So those price increments get smaller and smaller. And that's why you can get the monster, you know, ginormous Gulp at 7-Eleven-- is really just not that different from the price of getting the small little mini size, OK, because of diminishing marginal utility.

All right, and so that's how the market-- that's essentially how we can take this abstract concept, this sort of crazy math, and turn it into literally what you see in the store you walk into, OK? Questions about that? Yeah.

AUDIENCE: So how does this [? place ?] [INAUDIBLE], like if for example, you wanted to buy a snack that you were going to have for breakfast every day--

JONATHAN GRUBER: Awesome. Awesome question.

AUDIENCE: And then every single day, it was going to be your first granola bar, right? So I think that it's going to diminish every single time, but it's still cheaper to buy in bulk than it would be to buy a single granola bar every single time.

JONATHAN GRUBER: Great, great question. Yeah?

AUDIENCE: I think that has more to do with packaging cost than marginal utility.

JONATHAN GRUBER: Well, I mean, the risk of my going to this model is, once we get nonlinear, the order we do things in this class, we have to start talking about supply factors I want to talk to. But there's two answers. One is packaging efficiencies.

But the other is, if you actually go to Costco and look at their prices, for many things, they're not actually better than the supermarket. So actually, the price of buying the giant like, 8,000 bars of granola is actually not that much more-- not that much less than 1,000 time buying eight granola bars.

It's less, but it's not nearly as much less of these examples as sodas in McDonald's, which is exactly your point. Utility diminishes less, so they don't want to charge as much less for multiple packages.

So you can actually-- if you compare the gap in perishable product pricing by size, it's much larger than the gap in nonperishable pricing by size. Great point. Yeah.

AUDIENCE: Is there also just like a different time frame to which the utility starts diminishing for every product? Because you gave the example of soda, but it's like, would that reset later in the day, if we wanted-- were thirsty again, or--

JONATHAN GRUBER: Awesome, and that is why they don't let you walk back in with the same cup and refill it, right? That's exactly right, and that comes to this point. It's sort of like it's nonperishable as you get longer apart.

But you know, it's all just really interesting. So at Fenway, OK, you can get-- you get like a regular sized soda, it's like crazy. It's like $6. Then for like $8, you get a big soda. Then for $10, you get a refillable big soda, OK?

Now, the question is, can you bring that refillable soda back to additional games? Technically not, but I do.

[LAUGHING]

And basically they sort of understand-- so this interesting question of sort of the perishability of things and how that's going to affect things going on. It's a really-- it's an interesting question. Other comments?

OK, I'm going to stop there. Those are great comments. Thanks everyone for participating. And we will come back next time and talk about the sad reality that we haven't won the lottery, and we have limited amounts of money.