

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high-quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: All right, so today we're going to continue our discussion of consumer choice. If you remember the set-up from last time, the main motivation is you're trying to understand what underlies demand curves, how consumers ultimately decide to trade off price and quantity of goods. We said that ultimately that came from the principle of utility maximization, and that utility is maximized when individuals maximize the utility function, which is this mathematical representation of preferences.

And last time we talked about how if individuals were unconstrained how they choose what they want, they would just like more of everything, and their ranking across different bundles would depend on that underlying utility function. Now, of course, what's stopping individuals from consuming everything they want is their budget constraints. And so today we're going to turn to the second part of the problem, which is talking about budget constraints.

Now, we're going to make a very simplifying assumption here for most of the semester, which is we are going to assume that your income equals your budget. That is, you spend your entire income. That is, we're going to ignore the possibility of savings until about the third lecture from the end.

Now, this turns out not to be a terrible assumption for the typical American. The typical American doesn't save. So actually, it's not a terrible assumption for us to work with if we think about typical consumers. In practice, savings is going to turn out to be a very critical part of what we're going to do to think about economics, so we'll come back to that. But we're going to ignore savings for now and assume that your budget equals your income.

So let's say that your parents, probably a good model is you guys, you guys probably aren't in saving mode. You've got some budget saved from your parents. Let's call it y . And let's say that your parents give you some budget at the start of the semester, y , and they say this is your money you have to spend, say each month or for the whole semester. And let's imagine that you have to allocate that budget only across two goods, pizza and movies. So once again, unrealistic, but this is the kind of simplifying assumption that lets us understand how people

make decisions.

So that gives you your budget constraint. You've got some income y that your parents have given you, and you can allocate that across pizza and movies. So how do you allocate that? Well, you can buy movies, the number of movies you can get, plus the number of pizzas. Well, how many of each you can get, that depends on their price. In particular, budget constraint is the number of movies times the price per movie plus the number of pizzas -- plus the number of pizzas times the price for pizza. That's your budget constraint. It's the number of movies times the price per movie or the number of pizzas times the price for pizza.

And this is easiest to see graphically. If you go to figure 5-1, this is a graphical illustration of a budget constraint. Now, let's just carefully talk through this for a moment. You're going to be really good at dealing with budget constraints. You're going to have to be this semester. So let's carefully talk about where this comes from.

OK, the x-axis is going to be how many movies you could have if all you did with your income was consume movies. Well, if all you did with your income was consume movies, you could have y over p sub m movies. If you decided to devote your income solely to movies, then you could have y times y over p sub m movies.

If instead you decided to devote all your income to pizza, then you could have y over p sub p pizza. So the y-axis is going to be the point where you consume zero movies and all pizza. It's going to be where you devote your entire budget to pizzas.

And then there'll be some combination in between, which is our budget line. Which is the combinations of pizzas and movies you can consume given your total income y . And the slope of that line is going to be the price ratio. Or the negative of the price ratio. The slope of that line is going to be minus p_p over p_m .

The slope of that line is going to be the change in the ratio of the price of pizza to the price of movies.

OK, what am I doing wrong here? Negative of the price ratio. Have I got this right? Rise over run. Yeah. Have I got this right? Yeah.

AUDIENCE: I think that the p_p and the p_m 's are in the denominators, because it's y over p_p [INAUDIBLE PHRASE].

PROFESSOR: Right. It's the denominators, that's what I did wrong. Right, so it is pm over pp. Sorry, my bad. OK, right. Because they're the denominators. Because it's the rise over run in terms of the quantity. So that's what I did wrong.

OK, so basically it's the negative of the price ratio, minus the price of movies over the price of pizzas because they're in the denominators as you said, because as the price goes up, the quantity goes down. So the negative of the price ratio of the price of movies to the price of pizzas is the slope.

So let's just do a simple example. Imagine that income equals \$96. Imagine your parents give you \$96, say a month or a semester. Imagine that the price of movies is \$8 and imagine the price of a pizza is \$16. It's a good pizza. So what this means is that with your income of \$96, you could either get eight pizzas or 12 movies.

So that means that the price ratio of the slope of your budget constraint is minus 1/2. The price ratio is minus 1/2. So the slope of that budget constraint is minus 1/2. Now, we have a name for this slope. We're going to call this the marginal rate of transformation. The marginal rate of transformation is our label for this slope.

Now, why do we use that name? Well, it means that's the marginal rate at which you can transform pizzas into movies. The rate at which you can turn pizzas into movies. Now, once again, like I talked about last time, you're not an alchemist. You're not actually turning pizzas into movies. But the market essentially is giving you a rate at which you can do that given a budget, given that you have a certain amount of money.

Given that you have a certain amount of money, \$96, and given the prices that you face in the market, you could transform pizzas into movies by trading one pizza for 1/2 a movie. Now, once again, you're not actually doing the physical transformation, but that's the trade-off that you face when you're trying to transform one to the other.

So effectively, it's the same as if you're trading them for each other. As I talked about last time. It's essentially the same as you're trading, and that's because of the key economic concept we'll come back to over and over again in this course-- the concept of opportunity cost.

The opportunity cost is the value of the forgone alternative. The value of the forgone alternative is the opportunity cost. So basically what that means is if you decide to forgo a pizza, that's the same as forgoing two movies. Likewise, if you decide to forgo a movie, it's the

same as forgoing half a pizza.

So the opportunity cost of a movie, what essentially the movie is costing you, is $1/2$ a pizza. Now, really it's costing you \$8 and a pizza costs you \$16. But when we think about trading off goods, the opportunity cost of that movie is that you've forgone the ability to eat half a pizza. That's the opportunity cost of the situation.

So that's basically how we're going to think about this trade-off. We're going to think about trading off goods as the opportunity cost of consuming one good instead of another. The opportunity cost of that movie is that you haven't gotten to eat $1/2$ a pizza. The opportunity cost of the pizza is that you've forgone seeing two movies.

And the reason is because you have a fixed budget. If you had an infinite budget, there'd be no opportunity cost. But because you have a fixed budget and you have to allocate that budget, there's an opportunity cost. If you choose not to decide, you've still made a choice. I don't know whether that quote's due to Shakespeare or Rush, I'm not sure. I have to look that up. But basically, if you choose to have one thing, then by definition you're forgoing another.

Now, to understand the budget constraint, let's talk about what happens when we shock the budget constraint. Let's imagine the price of pizzas rose from \$16 to \$24. Pizzas got really expensive. We decided we only want gourmet pizzas or something. The price of pizzas went from \$16 to \$24.

What does this do? Well, let's look at figure 5-2. It'll show what that does. What that does is it says our new budget constraint, instead of being $16p + 8m = 96$, which is what the budget constraint was in our example, it's now $24p + 8m = 96$. That's the new equation for the budget constraint.

Or, more relevantly, the slope of the budget constraint has flattened from minus $1/2$ to minus $1/3$. The slope has fallen from minus $1/2$ to minus $1/3$. The price ratio has been reduced from $1/2$ to $1/3$.

Now, forget utility for a second. Forget the fact that we thought about utility. Just looking at this, can you tell whether you are better or worse off from this price change? You shook your head no, why not?

AUDIENCE: Because you don't know if you like pizzas enough to get any.

PROFESSOR: OK, well in particular, you can almost tell. Who's the only person who doesn't care about this price change?

AUDIENCE: [INAUDIBLE]

PROFESSOR: No, no no. Think about consumers, people with different preferences for pizzas and movies. What would your preference for pizzas and movies have to be for you not to care about this price change? All movies. So long as you care about pizza at all, you're worse off. So in fact, the answer is your opportunity set has been restricted. So we can think about the opportunity set. Your opportunity set is the set of choices you can make given your budget.

Before, you could make choices all the way up to the upper line. Now your set of choices that are available have just fallen. Now, you're no poorer-- it's not like your parents have cut you off. They still give you the \$96. But you're effectively poorer. You're effectively worse off, and why is that? Because the set of things you could afford with that \$96 has just been restricted.

And unless you truly have no value on pizza, unless all you care about's the movie-- you're gluten and cheese allergic or something, you have no value on pizza-- then you're worse off. Your opportunity set's restricted. And that's the key insight here, is that you are worse off because the price has increased. A price increase makes you worse off. It restricts your opportunity set, because with the same amount of income, you can now afford fewer goods. Your opportunity set has been restricted.

Likewise, now let's talk about what happens when your income falls. That's the next figure. Now, let's suppose your parents are pissed at you and they cut you down to \$80. Because you didn't do something. You don't write enough or call enough, so they cut you down to \$80. Well, here the slope of the budget constraint has not changed. Because what determines the slope of the budget constraint? It's prices, and no prices have changed.

The slope of the budget constraint is unchanged because prices haven't changed, but your opportunity set is once again restricted because you now have lower income. So you can now afford fewer pizzas and movies. So now, instead of being able to afford up to six pizzas and up to 12 movies, you can now only afford up to five pizzas and up to 10 movies because your income has fallen. So once again, you're unambiguously worse off. Your opportunity set has contracted.

So your opportunity set will contract whenever income falls or whenever price increases. And

how it affects the graph will depend on whether it affects prices, which affects the slope, or just income, which affects the intercepts.

Now, questions about the budget constraints and opportunity sets? Armed with that-- Yeah, I'm sorry, go ahead.

AUDIENCE: Is the area under the curve at all indicative of utility?

PROFESSOR: No, it's not, because that's going to be determined by your preferences. It's indicative of, if you will, potential utility, because that's your opportunity set. But as the example here points out, if you don't like pizzas at all, you'll feel very differently than if you like pizzas a lot. So it's indicative of sort of your potential utility, but not your actual well-being.

But now that's a great segue to the next step. Let's put that together and talk about constrained choice. Which is now, let's put together-- we know what your preferences are, we've mathematically represented those by utility function. We know what your budget set is, we've mathematically represented that based on your income and prices. Now let's put them together and talk about how you make choices.

And the basic question you want to ask is, what's the highest utility you can achieve given the constraints your budget constraints put on you? Or, graphically-- so let's say you wanted to understand this intuitively, graphically, and mathematically.

Intuitively the idea is quite simple, I think, which is just, what's the most you can have given the constraints that are placed on you? Graphically, we represent that as asking, what is the furthest out indifference curve you can achieve? Because remember, more is better. Indifference curves that are further out make you happier. So what's the furthest out indifference curve that you can achieve given your budget constraint?

So to do that, let's actually do an example. Let's imagine, as last time, your utility is the square root of pizza times movies. Once again, this has no fundamental meaning, it's just a mathematical representation of your preferences. So your preferences are mathematically represented by utility equals square root of pizza times movies. And let's have the same budget constraint that we have up here. Income is \$96, price of movies is \$8, price of pizza is \$16.

Now let's go to the next graph. Figure 5-4. What this does is put together our indifference curve analysis with our budget constraint analysis. It's a little complicated. The budget

constraint line is the vertical line running from a y-intercept of six to an x-intercept of 12. The-- not the vertical line, the straight line running from a y-intercept of six to an x-intercept of 12. That's your budget constraint. We saw that before.

Then we have here a series of indifference curves. These curves are drawn-- these are a mathematical representative of this utility function. These are points among which you're indifferent if you have that utility function. And what we see is that the best you can do is to choose point D. Point D, with six movies and three pizzas-- OK, that should be p on the y-axis, not C. It should be movies on the x-axis and pizzas on the y-axis. That should be p on the y-axis. The best you can do is to choose a point D.

Now, to see that. And that gives utility. What's the value of your utility at point D? We understand value's meaningless, but just so we can compare, what's the value of your utility at point D? The square root of 18. The value of utility is the square root of 6 times 3, which is the square root of 18. Which is going to be square root of two, or three times square root of two, but we'll just call it square root of 18 for comparison.

Now, let's talk about why that's the best point for you. Let's think about some alternative points. For instance, why is that better than point E? Somebody raise their hand and tell me, why is point D better than point E? Yeah?

AUDIENCE: E is unattainable with your budget.

PROFESSOR: E would be better, that'd be great. We'd love eight movies and four pizzas, but we can't reach it. So E's unattainable. Why is it better than point A? Point A you can afford. So why's point E better than point A? You could afford point A just like you can afford point D. Yeah?

AUDIENCE: Because the utility's only root 10.

PROFESSOR: Because what?

AUDIENCE: It's only root 10.

PROFESSOR: Yes, exactly, because the utility's only the square root of 10. You're on a lower indifference curve at point A. So it's true you could afford point A, but you're on a lower indifference curve. Your utility's a lower value, it's only square root of 10.

So point A is dominated by point D. What about point C? Well, point C you have the same--

point C is just an inward shift from point D, but here that's a dominated choice. Once again your utility's lower. It's the square root of 4.5 times 2.2. And basically that's dominated because you could afford more.

So basically, the point is that the point which will make you happiest is the point at which your indifference curve is tangent to the budget constraint. Because that is the point of the farthest out indifference curve that you can reach given your budget constraint. The tangency of the indifference curve and the budget constraint is the point which makes you best off given your available budget and the available prices. And that's the point where the slope of the indifference curve equals the slope of the budget constraint. The tangency is the point where the slope of the indifference curve equals the slope of the budget constraint.

Or, more relevantly-- OK, let me stop there. That's the graphic intuition. With a sloping indifference curve because the slope of the budget constraint is the optimum, because by definition, that is the point of the furthest out indifference curve you can reach given your budget. The point of tangency is the point of equal slopes. Are there questions about the graphical analysis here? This is very, very important, so. Yeah?

AUDIENCE:

This is sort of about the graphical analysis, but if it only matters in terms of the ordinal values of the utility function and p and M are always positive, does it matter if you got rid of the square root? Would anything change if it was u equals pm ? Because the the marginal rate of substitution, transformation, all that would still be the same, right?

PROFESSOR:

In this particular example, it would not. So actually, you're asking a great question. Because it's ordinal, you can typically do transformations for the ranking of bundles. You will always get the same ranking with a monotone transformation of the utility function. That's exactly right. Later in the course, we'll show different ways why the functional form matters, and I'll show you why I did square root, because it's going to turn out that that's going to matter. But for other things-- but for the ranking bundles, you're right. The ranking of bundles is consistent with the transformation.

So that's the graphical. Now let's come to the mathematical derivation of this. So let's talk about the mathematics of utility maximization. Now what I'm going to do here is, I'm going to do this sort of casually, as is my wont. Friday in section, you're going to work on the underlying calculus that lies behind the mathematics that I'm going to present here.

Now, let's talk about what it means that these slopes are equal. Well, remember, what is--

does anyone remember what the slope of the indifference curve is? What do we call the slope of the indifference curve? Yeah?

AUDIENCE: The MRS.

PROFESSOR: The MRS. The slope of the indifference curve is the marginal rate of substitution. Which is defined as what? What is the marginal rate of substitution? The ratio of what?

AUDIENCE: [INAUDIBLE]

PROFESSOR: No, the marginal rate of substitution. This is just about preferences. What's the marginal rate of substitution defined as? It's the ratio of what to what? Yeah?

AUDIENCE: The amount of one good you have to get to give up a unit of the other good.

PROFESSOR: I'm sorry?

AUDIENCE: The amount of one good you have to get to give up a unit of the other good.

PROFESSOR: So graphically that's what it's defined as, exactly. It's the slope of the indifference curve. Mathematically, what was it? What was it in terms of utility? Does anyone remember? Yeah.

AUDIENCE: The ratio of the partials.

PROFESSOR: Ratio of the marginal utilities. In particular, it's the negative of the marginal utility of movies over the marginal utility of pizza. Remember, it's the negative of the marginal utility of the x-axis over the marginal utility of the y-axis. So marginal rate of substitution is the rate at which you're willing to substitute between movies and pizza, which is a function of your marginal utilities.

If your marginal utility for movies is very high, then you need a lot of pizzas. Then you wouldn't trade a movie unless you get a lot of pizza for it. If your marginal utility of movies is very low, you'd be happy to give up a movie even for a small fraction of a pizza. So that's the marginal rate of substitution. That's about preferences only.

At the same time, we're saying that that marginal rate of substitution is equal to-- this slope is equal to the slope of the budget constraint. Well, the slope of the budget constraint we call the marginal rate of transformation, which is the price ratio.

The slope of the budget constraint is the negative of the price ratio. That's where you were

sort of one step ahead of us here. So preferences give us this, the marginal rate of substitution. The mechanics of the market give us the marginal rate of transformation. And utility maximization gives us that those are equal, because they're equal at that tangency.

At that tangency is where you get to the highest possible indifference curve. So at the optimum, you get that the ratio of marginal utility equals the ratio of prices. Now, I want to try to see you understand this a bit more intuitively, given this mathematics.

The way I like to think about this is, think about the ratio of the marginal utilities as the marginal benefit. So it's the benefit of another movie in terms of pizza. The marginal rate of substitution is the benefit of another movie in terms of pizza. It's how much you like that next movie relative to how much you like that next pizza.

The marginal rate of transformation, the cost, is the price of that next movie relative to the price of that next pizza. So what we're saying here is we're setting benefits equal to the costs. In particular, we're setting marginal benefits equal to marginal cost.

The marginal benefit, the benefit to that next movie in terms of pizza, has got to be equal to the marginal cost, the cost to you in terms of that next movie in terms of pizza. And this notion that the optimum will be where marginal benefit equals marginal cost will pervade through the whole course. When we do firm maximization, it'll be the same thing. Any maximization we'll do in this course, any optimization, will be about equating these margins. Setting the marginal benefits equal to marginal costs.

Now, this is different than benefits equals cost, because it's about the next unit. It's saying, how do we feel about that next movie compared to the price of that next movie. Now, prices here we have being constant. You could imagine prices of movies changing as you see more, but that gets complicated. We'll worry about that later. For now, the price is constant. But the marginal utilities are not constant. Marginal utilities are obviously changing the more moves you see.

Once again, with intuition, you have to develop your own intuition. The way that I like to think about this intuitively is to actually rewrite this a little bit, and rewrite it as saying that the marginal utility of movies over the price of movies equals the marginal utility of pizza over the price of pizza. At the optimum, this'll be true. I like this because to me this term sort of says, look, the bang for the buck has to be the same across all goods.

For each dollar of movie expenditure, what's it buying? What's that next dollar of movie expenditure buying you? This is saying, what's that next dollar of pizza expenditure buying you? And they've got to be equal. If the next dollar of movie expenditure buys you a lot more happiness than the next dollar of pizza expenditure, then you're not at the right place. You should shift your money and spend more on movies and less on pizza.

If the next dollar of pizza expenditure buys you a lot more happiness than the next dollar of movie expenditure, then you're not in the right place either. You should see fewer movies and buy more pizza. So basically, it's where the marginal benefit to you-- the bang for the buck of that next movie is the same as the bang for the buck of that next pizza.

So to see this, let's go back to figure 5-4 and let's actually think through the mathematics of a couple of these points. Let's take point A. At point A, you have two pizzas and five movies. So pizza equals two, movies equals five at point A. So utility is equal to square root of 10 at point A.

Now, in particular, your marginal utility for pizzas at that point-- what's the marginal utility for pizzas? Well, we can differentiate. That's the derivative of the utility function with respect to prices, du/dp . Which is going to be 0.5 times movies over the square root of p times m , which at these values is going to be one over the square root of 10. That's your marginal utility of pizzas.

Your marginal utility of movies is going to be du/dm , which is going to be 0.5 times p over square root of p times m , which is going to be 2.5 over square root of 10. So the marginal rate of substitution between these two is 2.5. The marginal rate of substitution, 2.5.

What does that mean intuitively? Can someone tell me what that means intuitively, the marginal rate of substitution is 2.5? What does that mean? Someone explain that like you'd explain it to someone who's speaking English. What does it mean that the marginal rate of substitution is 2.5? Yeah?

AUDIENCE: You'd give up one pizza for 2.5 movies.

Yes. No, actually, the opposite. You'd give up two and 1/2 pizzas-- the marginal rate of substitution is 2.5-- no, that's right. You would give up 2.5-- one second, let's make sure I have this right. Right, so you would give up one pizza to see 2 and 1/2-- no, it's the other way around. You're getting a lot of pizza and not enough movies. So you would give up two and

1/2 pizzas to get one more movie.

This is confusing, OK? You'd give up two and 1/2 pizzas to get one more movie. That's what it means. You would give up two and 1/2 pizzas to see one more movie, and why is that? Why at a point like A would you give up two and 1/2 pizzas to see one more movie? Yeah?

AUDIENCE: Well, if you just look at the line tangent to the indifference curve at A, it's really not a downward slope.

PROFESSOR: It's really steep, which means what?

AUDIENCE: Which means you get a lot more benefit--like you don't care if you give up a lot of pizzas to see more movies.

PROFESSOR: Exactly. Actually, that's a great way to bring the graphics and the intuition together. Here's thinking of it intuitively. I'm getting a lot of pizza at point A. I'm not getting a lot of movies. So I would happily give up a lot of pizza to get my next movie.

What you pointed out is the tie to the graphics here. The indifference curve is very steep at that point through A. A steep indifference curve in that way means I don't really care a whole lot at this point about how many pizzas I get, but I care a lot about getting more movies. So at a point like A where it's very steep, you are willing to give up two and 1/2 pizzas to see a movie. But what do you have to give up? What's the market telling you you have to give up to see a movie? How much pizza do you actually have to give up to see a movie in practice? You're willing to give up two and 1/2 pizzas to see a movie, but how many pizzas do you actually have to give up?

AUDIENCE: 1/2.

PROFESSOR: 1/2 a pizza to see a movie. So that can't be the optimum. You're willing to give up two and 1/2 pizzas to see a movie, but you only have to give up 1/2 a pizza to see a movie. So you can't be at the right place. You should be changing your consumption bundle. You should be changing your consumption bundle, because the market is only asking for 1/2 a pizza to see a movie, but you're willing to give up two and 1/2 pizzas to see a movie.

So at a point like A, you're clearly not at the optimum. Because you are willing to make a trade. Remember, we talked about-- we can go all the way back to the first lecture. The key point was-- the first or second lecture-- inefficiency happens when trades aren't made that people

value.

Here's a trade that you value. You're willing to give up two and 1/2 pizzas to see a movie. The market only wants 1/2 a pizza to see a movie. You're not making a trade you value, so that's not the efficient outcome.

Likewise, let's do the same mathematics at point B. Well, at point B, if you do the math and work it out, you see that the marginal utility of pizza is five over square root of ten. The marginal utility of movies is 0.5 over the square root of 10. So the MRS is 0.1. At this point, you'd only be willing to give up 0.1 pizzas to see a movie.

At a point like B, the indifference curve is very flat. You're only willing to give up 0.1 pizzas to see a movie. But remember, the market is willing to say, look-- you can flip it around, the market's willing to say, look, you can get a movie. You're only willing to give up 0.1 pizzas to see a movie. Well, that means clearly that you have too many movies and not enough pizza.

You're clearly at that point happy to say, wow, you mean that I can gain a whole pizza by just giving up two movies? Heck, I'd be willing to give up 10 movies to get a pizza. At the point I'm at right there, I'd be willing to give up 10 movies to get a pizza. You're telling me we only have to give up two movies to get a pizza? Great. I'm going to do that trade. I'm going to move back towards point D. So that's why this idea of what you're willing to do, which is the MRS, and what the market's making you do, you want to equilibrate those to decide how much you want to consume.

Now, obviously a point like-- let's talk about point C. Point C is interesting, because at point C, what's true? The marginal rate of substitution is equal to the marginal rate of transformation at point C. The slope of the indifference curve and the budget constraint are equal. That's why you have to check two conditions for optimization.

First of all, those slopes have to be equal. Second of all, you've got to spend all your money. So it's true there's a whole host of points-- in fact, there's a vector running through CDE, all which are points where the margin rate of substitution equals the margin rate of transformation. But only D is optimal, because you also have to remember more is better. You never want to leave money on the table.

So the two conditions you have to meet is that you're at the point where your desired trade-off between pizzas and movies is the same as the market's, and where you're spending all your

budget. And that's the optimum. OK, questions about that? So that's basically how we think about optimization. That's how we think about consumers making their decisions deciding between consuming pizzas and movies.

Now, let's come back-- however, this is a particular case we've looked at. This is a case in particular where we've imposed that there's an interior solution. In fact, in practice you could end up in these kinds of choices with corner solutions. So let's take a look at the last figure, figure 5-5.

We've chosen a case where your optimal bundle includes both pizza and movies. But you could imagine a situation where your optimal bundle-- and once again, that should be a p , not a c in figure 5-5, that should be p on the y -axis-- where your optimal bundle includes only one or the other.

So this is a particular case-- we have the same budget line as before, which is you're trading off pizza and movies. You have an income of 96, the price of pizzas is 16, the price of movies is 8. So same budget line as before. But now your indifference curves look very different.

Now your preferences are such that you've got these linear indifference curves of the form I_1 , I_2 , I_3 . You've got these linear indifference curves. What that means is you've got-- these indifference curves mean that you have a constant rate at which you're willing to trade off pizza for movies. If we go back to figure 5.4, the rate at which you're willing to trade off pizza for movies changes. Your preferences are such that-- because the square root function as pointed out-- your preferences are such that you are willing to make different rates of trade at different amounts.

In figure 5-5, your preferences are constant, no matter how many movies or pizzas you have. You're always willing to make that trade-off at the same rate. Well, in that case you can end up with a corner solution, where in fact, you're going to consume only six pizzas and no movies. And why is that? That's because this is a person who loves pizza relative to movies. It's a very flat indifference curve. They love pizza relative to movies. And they love pizza so much relative to movies that given the prices they face, they'll just go ahead and choose six pizzas and no movies. So that's a corner solution.

So mathematically, as you'll go through a section on Friday, you're going to have to check for corner solutions. You may solve these problems and end up with negative quantities and be befuddled about what happens. Well, if an answer looks wrong, it is wrong. If you solve

problems with negative quantities, that's probably because there's a corner solution to the problem, and actually the optimal quantity is to have zero of one thing and spend your entire budget on something else. Questions about that?

So now let's think about applying this to the kinds of decisions that you all have to make. So I talked about this particular example of pizza and movies. And in fact, you might say, well, that's sort of unrealistic. Gee, I spend my budget on lots of things, and how do I do this? Well, in practice, what you'd have to do is you'd have to draw a multi-dimensional graph and solve a multi-dimensional problem. And that's a bear.

But in practice, in fact, we can often think about breaking down the choices we make into pairs of choices. In practice, you could think about saying, look-- many people do what a lot of psychologists call mental accounting, where basically they say, look, yes, I have a whole budget and lots of things I can buy. But in fact, I like to think of my budget in sort of subcategories. I think of a certain amount I'm willing to spend on entertainment and a certain amount I'm willing to spend on food. And I take my budget and I mentally put it in different buckets. And within each of those buckets, you can then do the same kind of optimization problem that we've done here.

So even though in reality we choose across a whole host of goods, in practice what you're going to see is that people will do this kind of mental accounting, where they sort of divide their goods into different buckets and optimize within each of those buckets. What this means in practice is that in fact, if we now stop for a second and think about the government, and how it affects our consumption decisions, what this means is that in practice, the government-- so, one way we typically of the government affecting consumption decisions is through the power of taxation.

So let's say for example, the government decided that pizzas were bad. They caused obesity. That there's too much obesity because people are eating too much pizza, and we need to deal with that through a government policy that involves taxation. Somebody talk me through the analysis of how we'd think about analyzing a tax on pizza given these diagrams.

Let's go to figure 5-4. And imagine I said that we're going to place a 50% tax on pizza. So we're going to say that every dollar you pay on pizza, you're going to have to pay \$0.50 to the government. Because we're really worried people are eating too much pizza. What would that do the budget constraint? Yeah.

AUDIENCE: Well, effectively you're increasing the price of pizza.

PROFESSOR: Effectively you're increasing the price of pizza to the consumer.

AUDIENCE: The budget constraint would shift down like this.

PROFESSOR: It's actually going to have the same effect as we saw in figure 5-2. In fact, I've just replicated figure 5-2. Because in figure 5-2, the price of pizza went up by 50%, from \$16 to \$24. That's the same thing that the government's just done. It's raised the price of pizza effectively from \$16 a pizza to \$24 because instead of paying your \$16, you're also going to pay \$8 to the government in tax.

So what's that going to do? That is going to, in general, lower consumption of pizza. So that kind of price increase is going to, in general, lower the consumption of pizza. So the government has tried to accomplish its goal by shifting people away from pizza towards movies, away from pizza toward other things. Yeah?

AUDIENCE: Would it be meaningful to say that it also affects the utility curves?

PROFESSOR: It would not be meaningful, actually-- it will affect your optimal choice. In general, you will choose a different amount of pizza and movies. We'll talk about that next time. But it's very important, it would not be right to say it affects utility. Utility is like what you're born with, it's an innate concept about your underlying preferences. However, you're actually getting to my point, which is in fact, the way economists typically think about this is the government can't affect your preferences. But in fact, if people do mental accounting, the government maybe can affect your preferences.

So let's say that basically the way I think about it is, let's say I think I have a budget for food and I have a budget for entertainment. And let's say I think my budget for entertainment's pretty small because I'm low income. I've gotta have a budget for food. And so let's say I put pizza in my budget for food, so I allocate some of my budget for food to pizza. And the government can then cause me to eat less pizza by taxing it.

But what if somehow the government could get me to think of pizza differently? What if somehow the government could get me to think of pizza as entertainment? And suddenly I put it in that bucket, where I'm trading it off not against other food, but against the fact that I want to see a movie and I want to download stuff from iTunes, et cetera. Maybe then I'd buy less

pizza at the same price because I'm putting it in the bucket where I have less money.

So in other words, we've imagined a world with two goods. And the only way you can affect the choice across those two goods is to lower your income or your price. But in fact, if people have lots of different bundles of goods, somehow I could shift you mentally from considering pizza in a bucket where you have a lot of money to putting pizza in a bucket where you don't have as much money. I can lower your consumption of pizza without affecting prices, without an obtrusive government tax policy.

This is the kind of thing that we call in economic policy a nudge. And there's a new book called *Nudge* by Richard Thaler, who's a famous behavioral economist, basically bringing psychology into economics. There's this very important field now in economics called behavioral economics, which is all about, how can we bring the lessons of psychology into economics?

We don't do that in 14.01. 14.01's all about, we assume everybody's these perfectly rational people who would never really be fooled into thinking about pizza differently just because of what the government told them. But in fact, in reality people think about things differently based on the kinds of information they have. And given that tax policies can feel very intrusive - imagine some of you are like, wow, they're raising my pizza to \$24, that seems very intrusive.

If the government could somehow through nudging you to think about pizza differently change your pizza consumption, that might be a much more acceptable and palatable policy to many people. And that's the kind of role the government can play, or policy-makers can play, is not just by changing your prices or income, but by actually changing how you categorize things mentally, they can change the choices you make.