5

Optimal Choice

Varian, H. 2010. *Intermediate Microeconomics*, W.W. Norton.

Modeling Choice

Earlier we talked about economists' model of consumer choice:

Consumers choose the *best* bundle they can *afford*.

With the language of budget sets, preferences, and utility functions we can say this more precisely:

1. A consumer chooses the most preferred bundle in her budget set.
2. A consumer chooses a bundle that maximizes utility, subject to the budget constraint.

Conditions for optimal choice

Assume we have well-behaved preferences. We want to pick the bundle in the budget set on the *highest* indifference curve?

This means we want to pick a bundle *on the budget line*. Why?

Answer: With well-behaved preferences, we can rule out bundles inside the budget line because more is always preferred to less!

Optimal choice

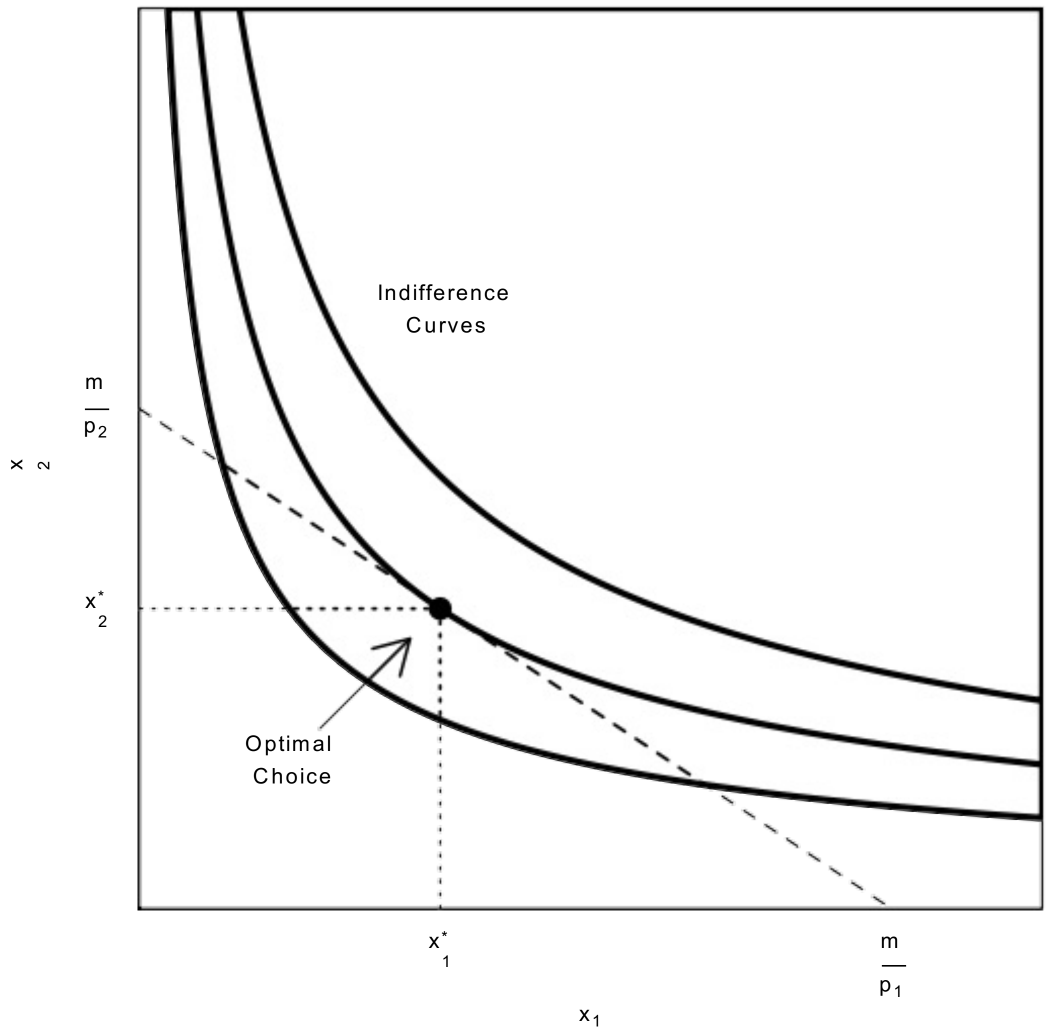

Now start from either corner and ask, as you move towards the center of the budget line, “am I on a higher indifference curve?”

Optimal choice

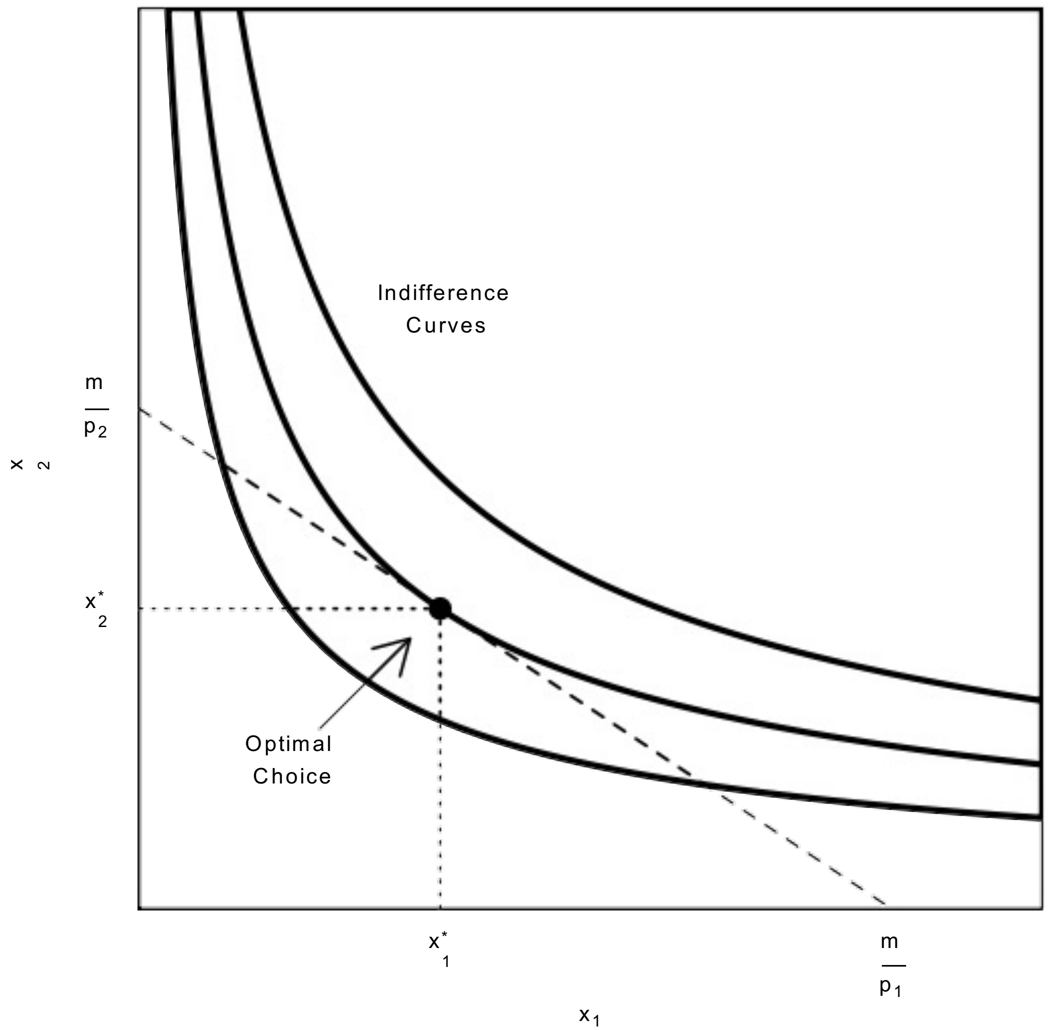

The bundle (x_1*, x_2*) is an **optimal choice** because the set of more preferred bundles does not intersect the budget set.

Optimal choice

Notice that with well-behaved preferences, the optimal bundle puts you on an indifference curve that is *tangent* to the budget line.

Why?

Answer: If the indifference curve weren't tangent, it would either:

1. Cross the budget line, in which case some more preferred bundles would still be affordable.
2. Be above the budget line, in which case, the bundle wouldn't be affordable.

Is tangency necessary?

Why might the tangency condition not be necessary in all cases? **kinky preferences**

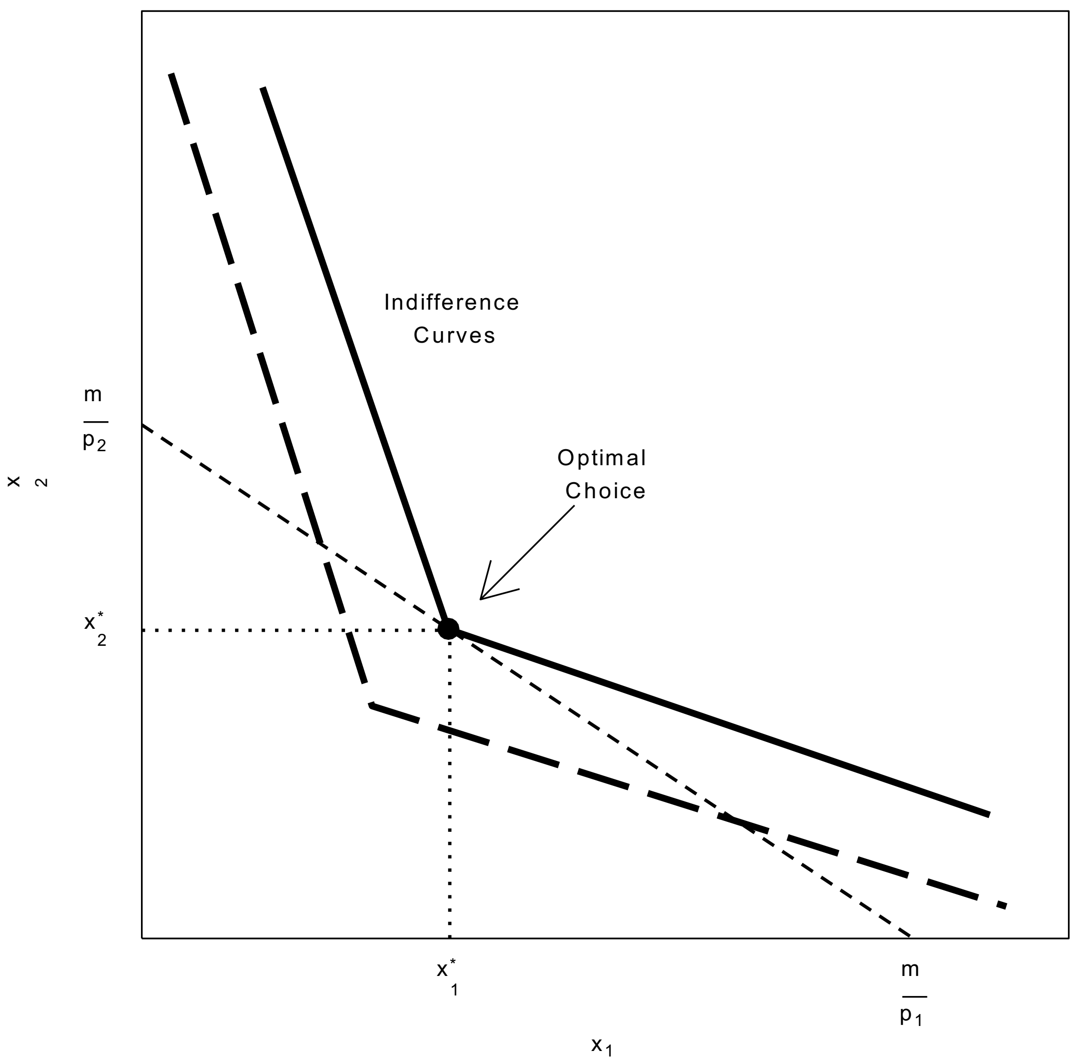

Is tangency necessary?

Why might the tangency condition not be necessary in all cases?
boundary optima

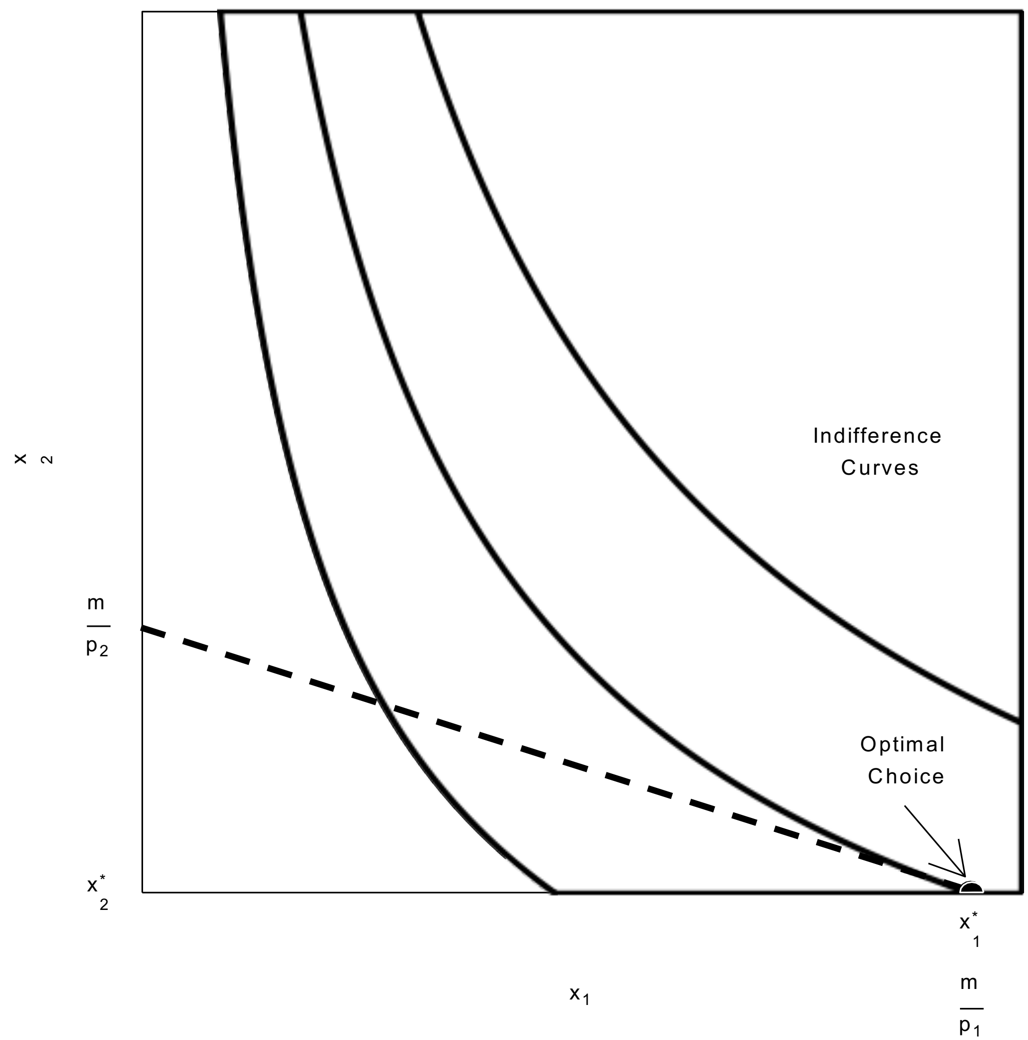

Restrictions

Once again, we're going to appeal to some restrictions to make the math simpler.

If we rule out **kinky preferences** and **boundary optima** then we will always have tangency between the indifference curve and the budget line at the optimal bundle.

Under these restrictions, the tangency condition is called a *necessary* condition for optimality.

The condition is logically *necessary* because the bundle couldn't be optimal if the condition weren't fulfilled.

Is tangengy sufficient?

Why might the tangency condition not be sufficient? **multiple optima**

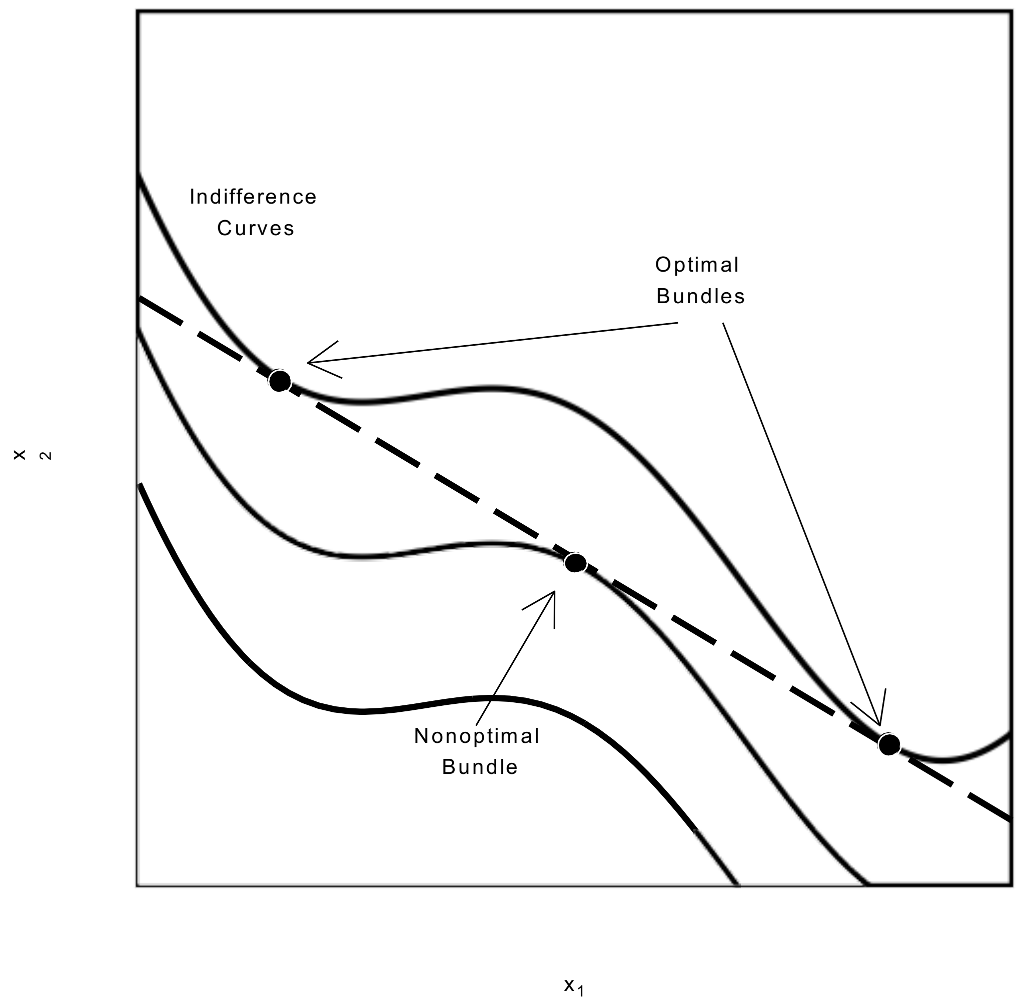

Well-behaved preferences

When preferences are *convex*, any tangency point between the indifference curve and the budget line must also be an optimum.

But this still doesn’t rule out multiple optima.

Why?

Answer: There could be a flat spot on the indifference curve.

But, if preferences are *strictly convex*, the there will be a **unique optimum** for each budget line!

Return of the MRS

Our graphical analysis implies the following:

At an interior optimum, the marginal rate of substitution (MRS) must equal the slope of the budget line.

MRS and optimality

Recall that we can think of the MRS as specifying the rate of exchange between goods 1 and 2 at which a consumer is willing to keep their current bundle.

When the market offers an exchange rate of $-p_1/p_2$, the consumer can give up one unit of good 1 to buy p_1/p_2 units of good 2.

If the consumer has a bundle that they would prefer not to change, then it must be the one where:

$$MRS = \frac{-p_1}{p_2}$$

MRS and optimality

Suppose you had a bundle where your MRS was different from the price ratio.

Say, your $MRS = \Delta x_2/\Delta x_1 = -3/2$ but the price ratio is 1/1.
What would you want to do?

Answer: At your current MRS, you would give up 3 units of good 2 to get an additional 2 units of good 1, but at market rates you only need to give up 2 units of good 2 to get 2 units of good 1!

So you couldn't be at an optimal bundle! You'd want to trade to get a better bundle, and you'd continue to trade until you got to a bundle where MRS $= \frac{-p_1}{p_2}$.

Optimality and demand

We call the optimal bundle, given prices and a budget, the consumer's **demanded bundle**.

Typically, when prices and income change, the **demanded bundle** will also change.

For that reason, we can define a **demand function** that explains how the demanded bundle (the optimal quantities of x_1 and x_2) changes as we change prices and income.

$$x_1(p_1, p_2, m) \ and \ x_2(p_1, p_2, m)$$

Note that the demand function will depend on the consumer's preferences, and this will provide an even more convenient way to represent consumer preferences than the utility function!

Perfect substitutes

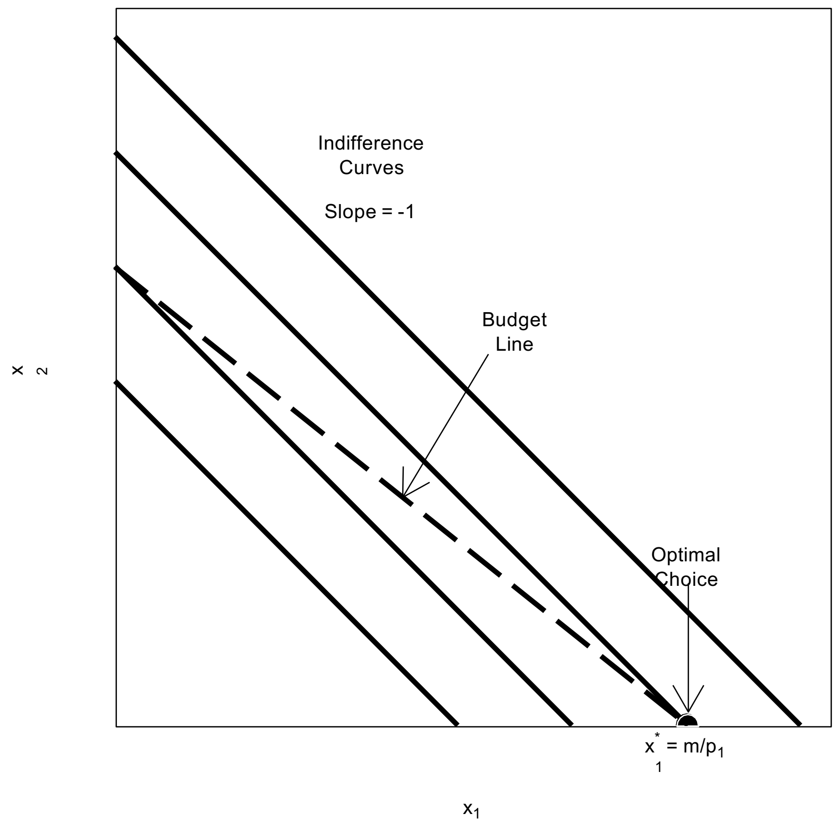

$$x_1 = \begin{cases} \dfrac{m}{p_1} & when\ p_1 < p_2 \\ \left[0, \dfrac{m}{p_1}\right] & when\ p_1 = p_2 \\ 0 & when\ p_1 > p_2 \end{cases}$$

Perfect substitutes

If goods are perfect substitutes, the consumer exclusively buys the cheaper the two goods.

If both prices are the same, it doesn't matter which combination of goods the consumer buys.

Think about what happens to the previous figure when the price of good 1 increases.

Perfect substitutes

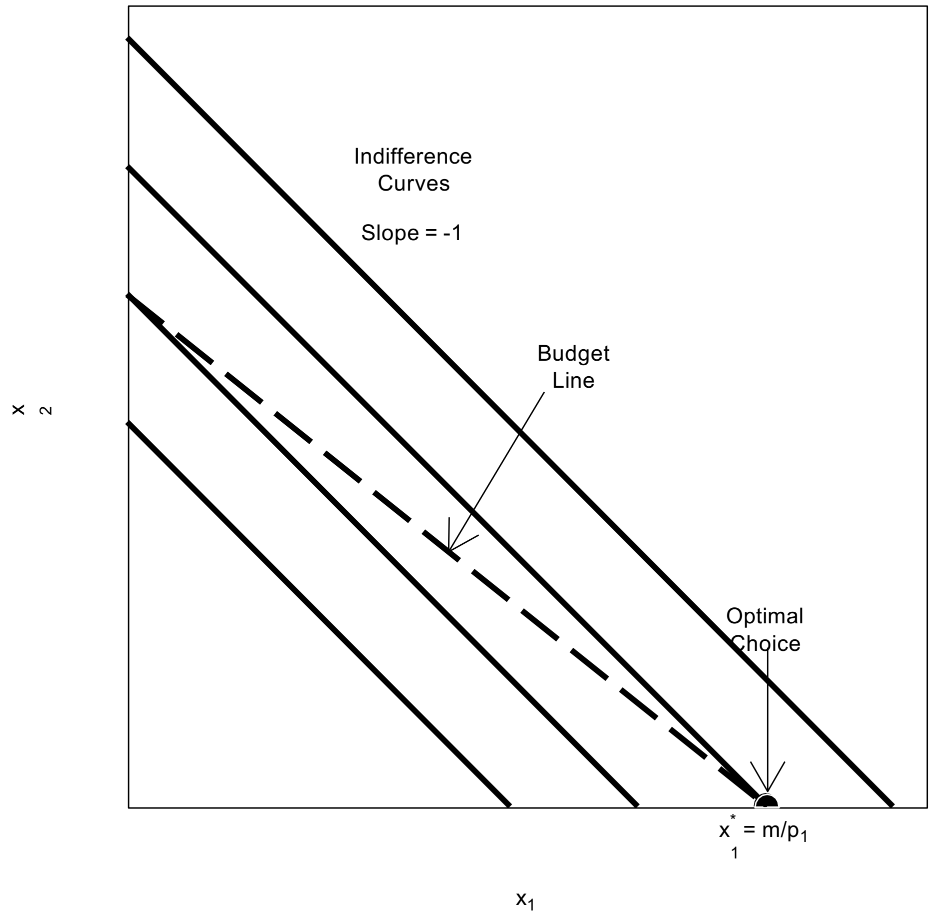

When the price of good 1 rises, the budget line rotates inward, and if it rises high enough, consumers will substitute good 2 for good 1.

Perfect complements

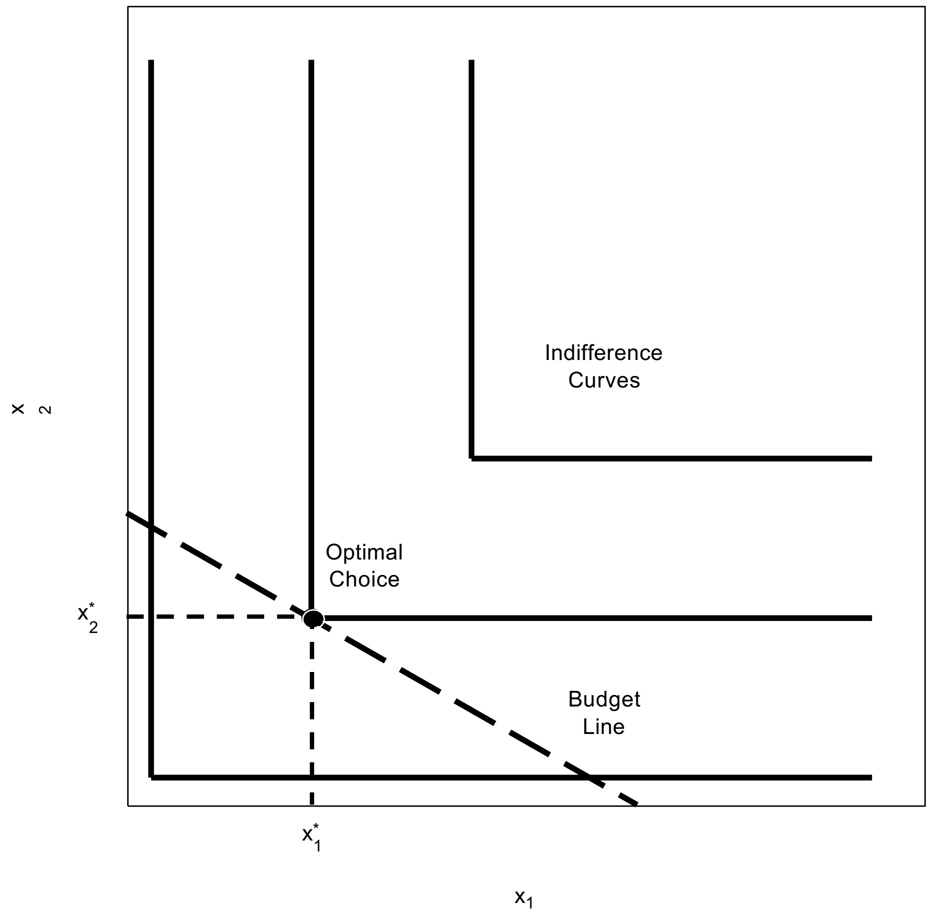

$$x_1 = x_2 = x = \frac{m}{p_1 + p_2}$$

Perfect complements

Where does the demand function come from?

Since the goods must be purchased together to have any value, we know the consumer buys the same quantity of each good, call it x.

Then, we must have:

$$p_1 x + p_2 x = m$$

Then we can solve for x to get:

$$x_1 = x_2 = x = \frac{m}{p_1 + p_2}$$

Neutral and Bads

If either good is a **neutral** or a **bad**, the consumer will spend all of her money on the good she likes and none of it on the neutral or bad.

So the demand functions, supposing good 2 is the neutral (bad) will be:

$$x_1 = \frac{m}{p_1}$$

$$x_2 = 0$$

Discrete goods

When a consumer chooses between discrete goods and spending money on everything else, the bundles can be written as $(1, m - p_1), (2, m - 2p_1), (3, m - 3p_1) \ldots$

Then we just want to find out which bundle has the highest utility.

We can also do this graphically...

Discrete goods

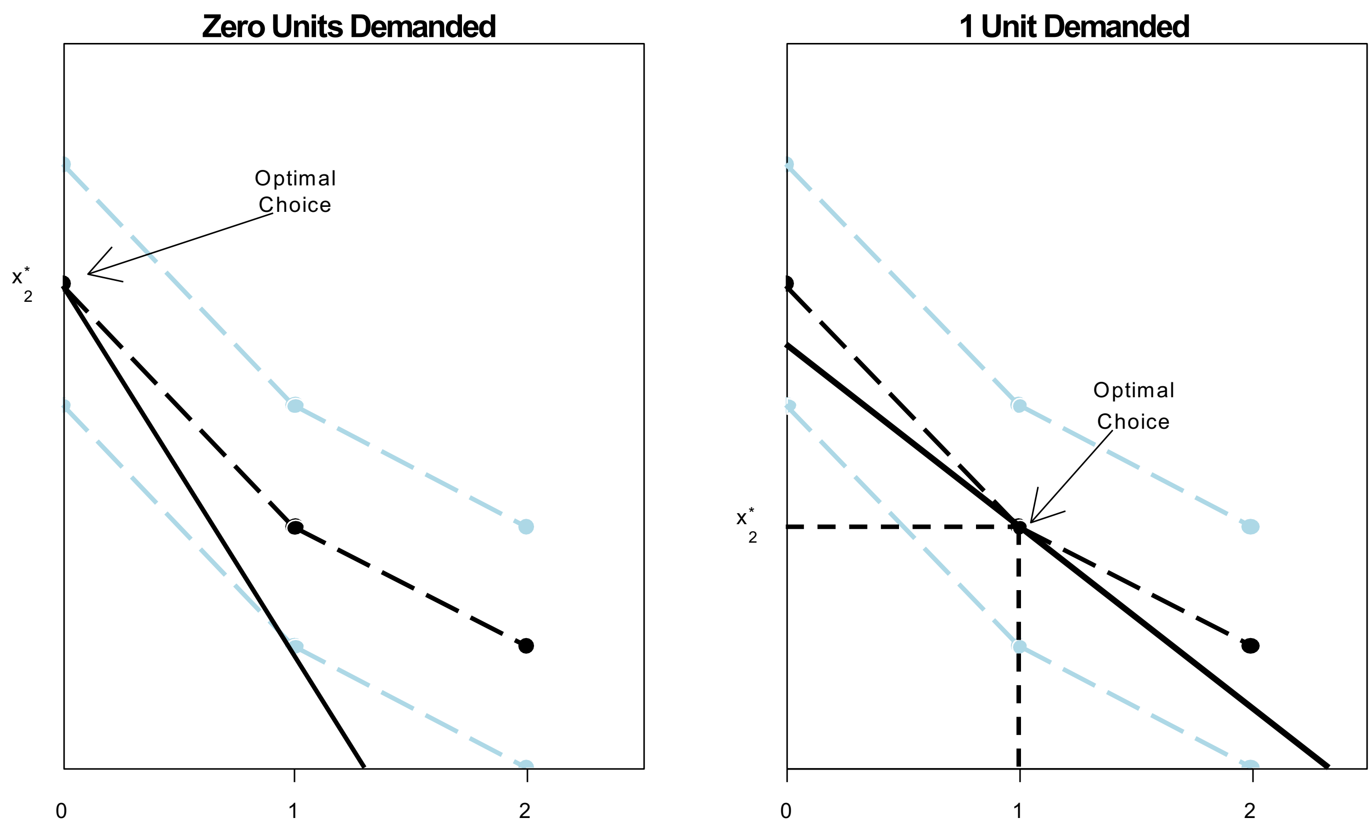

The optimal bundle is the one on the highest indifference curve. As the price of good 1 falls, the consumer will continue *not* to purchase it until some threshold.

Concave preferences

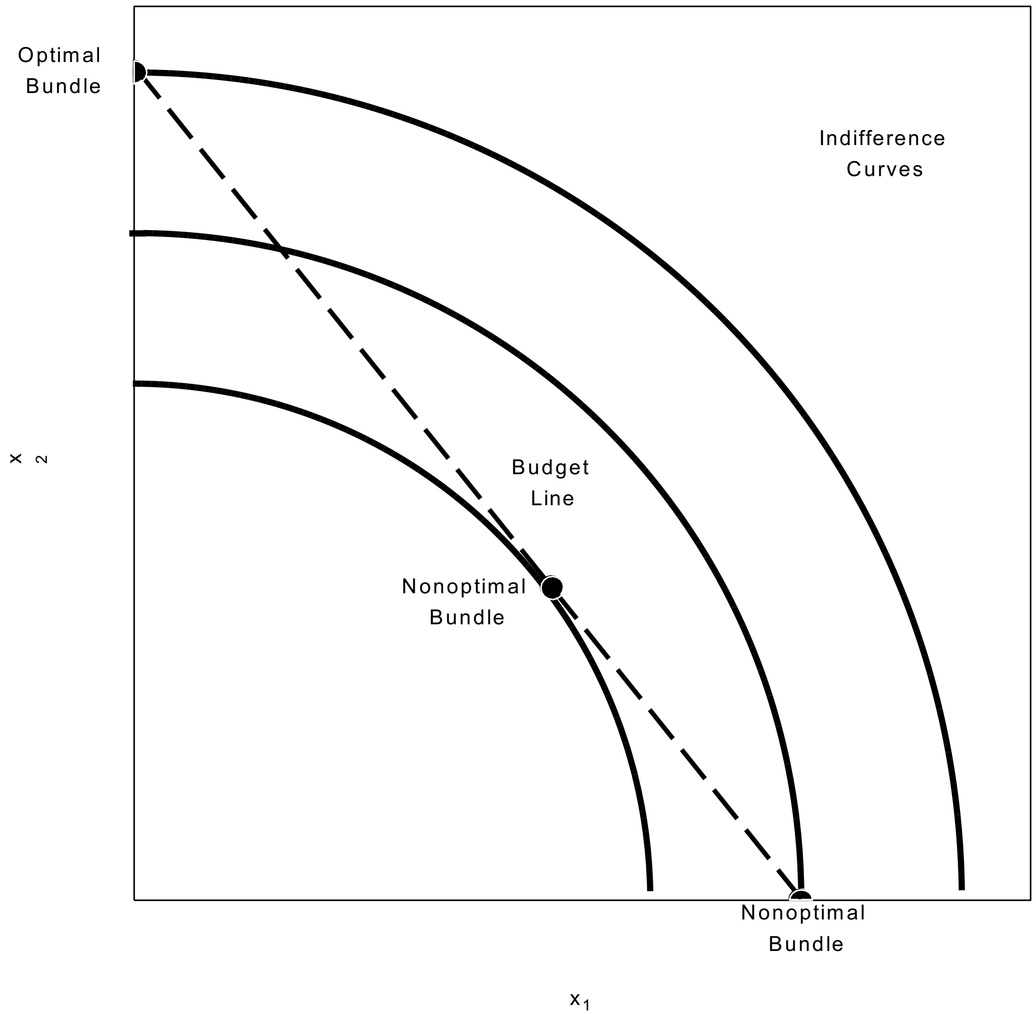

These will *always* yield **boundary optima**.

Cobb-Douglas preferences

Recall the Cobb-Douglas utility function $u(x_1, x_2) = x^c x^d$.

We want to find the **demand function** that gives the optimal choices (x_1, x_2) for all prices (p_1, p_2) and m.

Before we solve this particular case, let's go back and look at the utility maximization problem more generally.

Utility (preference) maximization

The reason we developed the **utility function** representation of preferences was so that we can treat the consumer's choice problem as a simple maximization problem (hopefully you remember these from calculus).

The goal is to solve this maximization problem.

The good news is that we already have the tools to do it!

Optimal choice and MRS

Recall that an optimal choice, that is, a bundle (x_1, x_2) that solves the utility maximization problem, must satisfy the following condition:

$$MRS(x_1, x_2) = -\frac{p_1}{p_2}$$

And recall from the Utility lecture (Chapter 4) that the MRS can be expressed as:

$$MRS(x_1, x_2) = -\frac{\partial u(x_1, x_2)/\partial x_1}{\partial u(x_1, x_2)/\partial x_2} = -\frac{p_1}{p_2}$$

And the negative signs cancel out.

Optimal choice and MRS

From the Budget Set lecture (Chapter 2) we also know that optimal choice must satisfy:

$$p_1 x_1 + p_2 x_2 = m$$

Therefore, we're left with two equations and two unknowns, and we just need to solve for the optimal choices of x_1 and x_2 as a function of prices and income.

Solution concept #1: substitution

We can solve the budget constraint for one of the choices and then substitute that choice into the MRS condition.

E.g. we can solve the budget constraint for x_2:

$$x_2 = \frac{m}{p_2} - \frac{p_1}{p_2} x_1$$

When we substitute that into the MRS condition, we get:

$$\frac{\partial u(x_1, \frac{m}{p_2} - \frac{p_1}{p_2} x_1)/\partial x_1}{\partial u(x_1, \frac{m}{p_2} - \frac{p_1}{p_2} x_1)/\partial x_2} = \frac{p_1}{p_2}$$

Since this equation only has one unknown (x_1), it can then be solved for x_1 in terms of (p_1, p_2, m), and substitution back in the budget constraint will yield a solution for x_2.

Back to Cobb-Douglas

How can we solve for the general form of the demand function with Cobb-Douglas utility?

$$u(x_1, x_2) = x_1^c x_2^d$$

First, to make life easier, let's take a monotonic transformation:

$$\ln u(x_1, x_2) = c \ln x_1 + d \ln x_2$$

Then, we can write the maximization problem:

$$\max_{(x_1, x_2)} c \ln x_1 + d \ln x_2$$
$$s.t. \, p_1 x_1 + p_2 x_2 = m$$

Solution #1

Using the expression for the MRS from Chapter 4:

$$\frac{cx_2}{dx_1} = \frac{p_1}{p_2}$$

$$p_1 x_1 + p_2 x_2 = m$$

These are two equations in two unknowns, so we can solve them by substituting the second equation into the first:

$$\frac{c(m/p_2 - x_1 p_1/p_2)}{dx_1} = \frac{p_1}{p_2}$$

Solution #1

Then, cross multiplying the last expression gives:

$$c(m - x_1 p_1) = d p_1 x_1$$

We can rearrange this to get:

$$cm = (c + d) p_1 x_1$$

Which is equivalent to:

$$x_1 = \frac{c}{c + d} \frac{m}{p_1}$$

Solution #1

To get the demand function for x , we can substitute back into the budget constraint to get:

$$x_2 = \frac{m}{p_2} - \frac{p_1}{p_2}\frac{c}{c+d}\frac{m}{p_1}$$

$$= \frac{d}{c+d}\frac{m}{p_2}$$

Notice, these demand functions imply that with Cobb-Douglas preferences, the consumer will spend a *fixed percentage* of her income on each good!

Constrained maximization

We can also use the calculus conditions for maximization to solve the following *constrained maximization* problem:

$$\max_{(x_1,x_2)} u(x_1, x_2)$$

$$s.t.\, p_1x_1 + p_2x_2 = \mathrm{m}$$

This says that we want to choose x_1 and x_2 to maximize the *objective (utility) function*, subject to the (budget) *constraint* that the chosen bundle must be affordable.

There are two ways to solve this problem.

Solution concept #2: convert to uncostrained maximization

The first is to solve the budget constraint for one of the variables (say x_2) and then to substitute it into the objective function.

E.g. given x_1, the amount of x_2 needed to ensure we're on the budget line, that is x_2 as a function of x_1, is:

$$x_2(x_1) = \frac{m}{p_2} - \frac{p_1}{p_2} x_1$$

Then, substitute this back into the objective function to get:

$$\max_{x_1} u(x_1, \frac{m}{p_2} - \frac{p_1}{p_2} x_1)$$

Which is an *unconstrained* maximization problem.

Solution concept #2

Then, all we need to do is differentiate (take the partial derivative of) the new objective function with respect to x_1 and set the result equal to 0 to get the *first-order condition*:

$$\frac{\partial u(x_1, x_2(x_1))}{\partial x_1} + \frac{\partial u(x_1, x_2(x_1))}{\partial x_2}\frac{dx_2}{dx_1} = 0$$

The first term shows how utility changes as x_1 increases.

The second term shows:

1. The rate of increase of utility as x_2 increases
2. The rate of decrease of x_2 as x_1 increases in order to continue to satisfy the budget constraint

Solution concept #2

Differentiate the equation for $x_2(x_1)$ in order to get the second part of the second term:

$$\frac{dx_2}{dx_1} = -\frac{p_1}{p_2}$$

Then, substituting that back into our equation from the last slide gives:

$$\frac{\partial u(x_1^*, x_2^*)/\partial x_1}{\partial u(x_1^*, x_2^*)/\partial x_2} = \frac{p_1}{p_2}$$

Which just says (again) that the MRS between x_1 and x_2 must equal the price ratio at the optimal choice (x_1^*, x_2^*), which must also satisfy the budget constraint $p_1\, x_1^* + p_2\, x_2^* = m$, giving us two equations and two unknowns.

Solution #2

Let's try this concept out with the Cobb-Douglas Utility function.

We can also just substitute the budget constraint into the maximization problem:

$$\max_{x_1} c \ln x_1 + d \ln(m/p_2 - x_1\, p_1/p_2)$$

Then take the derivative w.r.t. x_1 to get the first-order condition:

$$\frac{c}{x_1} - d\,\frac{p_2}{m - p_1 x_1}\frac{p_1}{p_2} = 0$$

Solution #2

Then, p_2 cancels out to give:

$$\frac{c}{x_1} - \frac{dp_1}{m - p_1x_1} = 0$$

Adding the second term to both sides and multiplying by x_1 gives:

$$c = \frac{dp_1x_1}{m - p_1x_1}$$

Multiplying both sides by the denominator and adding cp_1x_1 to both sides leaves:

$$cm = dp_1x_1 + cp_1x_1 = (c + d)p_1x_1$$

Solution #2

Finally, divide by $(c + d)p_1$ to get:

$$x_1 = \frac{c}{c + d}\frac{m}{p_1}$$

Again, by substituting into the budget constraint, we can solve for x_2 :

$$x_2 = \frac{d}{c + d}\frac{m}{p_2}$$

Which is the same solution we got before!

Solution concept #3: Lagrange multipliers

We can also solve the maximization problem using **Lagrange Multipliers**.

We start by defining a new function called the *Lagrangian*:

$$L = u(x_1, x_2) - \lambda(p_1 x_1 + p_2 x_2 - m)$$

The new variable λ is known as the **Lagrange multiplier**, because it multiplies the constraint. Lagrange's theorem says that an optimal choice (x_1^*, x_2^*) must satisfy:

$$\frac{\partial L}{\partial x_1} = \frac{\partial u(x_1^*, x_2^*)}{\partial x_1} - \lambda p_1 = 0$$

$$\frac{\partial L}{\partial x_2} = \frac{\partial u(x_1^*, x_2^*)}{\partial x_2} - \lambda p_2 = 0$$

$$\frac{\partial L}{\partial \lambda} = p_1 x_1^* + p_2 x_2^* - m = 0$$

Solution concept #3

The three equations are just the partial derivatives of the *Lagrangian* with respect to x_1, x_2 and λ.

Now we have 3 equations and 3 unknowns, and we can solve for x_1 and x_2 in terms of (p_1, p_2, m).

Note that we will not need to prove Lagrange's Theorem (or necessarily to use it), but in some cases you may find it convenient to use.

But note, that by dividing the first condition by the second, we end up with:

$$\frac{\partial u(x_1^*, x_2^*)/\partial x_1}{\partial u(x_1^*, x_2^*)/\partial x_2} = \frac{p_1}{p_2}$$

And the budget constraint (condition 3) as before, and we're back to 2 equations and 2 unknowns.

Solution #3: Cobb-Douglas

As before, we start with $u(x_1, x_2) = c \ln x_1 + d \ln x_2$

Set up the Lagrangian:

$$L = c \ln x_1 + d \ln x_2 - \lambda(p_1 x_1 + p_2 x_2 - m)$$

Differentiate to get the three first-order conditions:

$$\frac{\partial L}{\partial x_1} = \frac{c}{x_1} - \lambda p_1 = 0$$
$$\frac{\partial L}{\partial x_2} = \frac{d}{x_2} - \lambda p_2 = 0$$
$$\frac{\partial L}{\partial \lambda} = p_1 x_1^* + p_2 x_2^* - m = 0$$

Solution #3

To solve them, let's start by solving for λ.

Rearranging and cross multiplying the first two equations leaves us with:

$$c = \lambda p_1 x_1$$

$$d = \lambda p_2 x_2$$

Adding these together gives us:

$$c + d = \lambda(p_2 x_2 + p_2 x_2) = \lambda m$$

Which implies:

$$\lambda = \frac{c + d}{m}$$

Solution #3

Then, substituting back into the first two equations and solving for x_1 and x_2 gives us:

$$x_1 = \frac{c}{c+d}\frac{m}{p_1}$$

$$x_2 = \frac{d}{c+d}\frac{m}{p_2}$$

Which is the same solution as in the other two cases!

Cobb-Douglas utility

To summarize, the Cobb-Douglas utility function:

$$u(x_1, x_2) = x_1^c x_2^d$$

Generates demand functions (found in many ways):

$$x_1 = \frac{c}{c+d}\frac{m}{p_1}$$

$$x_2 = \frac{d}{c+d}\frac{m}{p_2}$$

Which imply that the consumer will always spend a fixed percentage of her income on each good.

This is why it is convenient to normalize the exponents so that they sum to 1!

$$u(x_1, x_2) = x_1^\alpha x_2^{1-\alpha}$$

The utility of utility functions

One situation in which you might expect people to have Cobb-Douglas preferences is where good 1 is something like the percent of income that people spend on housing and good 2 is the percent left over for everything else.

There's some evidence that this is pretty stable within countries (e.g. 15-20% in the UK).

So then the utility function might look like:

$$u(x_1, x_2) = x_1^{\frac{1}{5}} x_2^{\frac{4}{5}}$$

The utility of utility functions

Since x_2 is a **composite good** we can assume p_2 is 1.

Suppose m = 1000, and suppose the population is growing. The government is considering a policy that would restrict the ability of developers to build new units.

With a growing population and a fixed supply, suppose that would raise the price of 1 "unit of housing", say 1000 sq. ft., (p_1) from 100 to 200.

Then demand for x_1 and x_2 at initial prices is:

$$x_1 = \frac{1}{5}\frac{1000}{100} = 2$$

$$x_2 = \frac{4}{5}\frac{1000}{1} = 800$$

And the initial utility is: $u(x_1, x_2) = 2^{\frac{1}{5}} 800^{\frac{4}{5}} = 241.4$

The utility of utility functions

After the building restriction is imposed and $p_1 = 200$, the demand becomes:

$$x_1 = \frac{1}{5}\frac{1000}{200} = 1$$

$$x_2 = \frac{4}{5}\frac{1000}{1} = 800$$

Which generates utility: $u(x_1, x_2) = 1^{\frac{1}{5}}800^{\frac{4}{5}} = 210.1$

Remember, the actual number is meaningless. What is important is that the policy moves the consumer to a *lower* indifference curve (e.g. to a smaller housing unit).

What might be some of the other responses by homeowners?

More on the MRS

Think about it, if:

1. The market generates a single price for each good, say butter and milk, (as is often the case in well-organized markets),
2. Everyone is choosing optimally,
3. Everyone is at an interior solution

Then, everyone must have exactly the same MRS for butter and milk!

The market offers everyone the same rate of exchange for milk and butter to everyone, and everyone is adjusting their consumption until their *own* marginal valuation of the two goods equals the market's marginal valuation.

More on the MRS

Crucially, this is true *regardless of income and preferences*!

Some people will value their total consumption differently, and some will have more milk while others will have more butter.

And wealthier people may have more of both.

But, everyone who has optimized consumption will have exactly the same **marginal rate of substitution**.

Which says that everyone will agree on the rate at which they would exchange one good for the other.

And because everyone has consumed to the point where that rate equals the market exchange rate, no one will want to engage in further trade.

More on the MRS

Now, we might doubt that people actually optimize perfectly. . .

BUT the combination of utility-driven choice and market prices will allow choices to tend in the direction of optimality over time.

Prices contain information, and that information creates incentives.

Just thinking about a) how policies/events change relative prices and b) how changes in relative prices will change choices will help you understand a lot about how the world works.

The role of the entrepreneur

Entrepreneurs get rich by finding a way to transform goods at better than the current market rate of exchange!

If the current exchange rate has 1 quart of milk going for \$2 and 1 pound of butter for \$2, then an entrepreneur can profit by finding a way to convert butter to milk more efficiently.

Suppose she develops a technology that converts a pound of butter into 2 quarts of milk!

She can then buy up pounds of butter for \$2, convert them to 2 quarts of milk and sell the result for a \$4!

The role of prices

This all points to one of the most important insights in economics, the resolution to the water-diamond paradox:

Prices are not arbitrary numbers, they tell us how people value goods at the margin.

Diamonds are more expensive than water because of their value at the margin which is caused by their scarcity.

Tax policy

Quantity taxes impact the budget constraint by increasing the price of one good, and the optimal bundle will satisfy:

$$(p_1 + t)x_1^* + p_2 x_2^* = m$$

This tax raises total revenue $R^* = t x_1^*$

Income taxes impact the budget constraint by reducing income (for this example, we'll use a revenue-equivalent **lump sum** tax).

$$p_1 x_1 + p_2 x_2 = m - R^*$$

Which can be rewritten as:

$$p_1 x_1 + p_2 x_2 = m - t x_1^*$$

Tax policy

Under the income tax the budget line will have the same slope as it did initially.

And a revenue-equivalent, lump-sum income tax will ensure that the optimal bundle under the quantity tax will still be available.

This is true because when we rearrange the budget constraint under the quantity tax, we get:

$$p_1x_1 + p_2x_2 = m - tx_1^*$$

Thus, it will be affordable under the income tax, but will it be optimal?

Tax policy

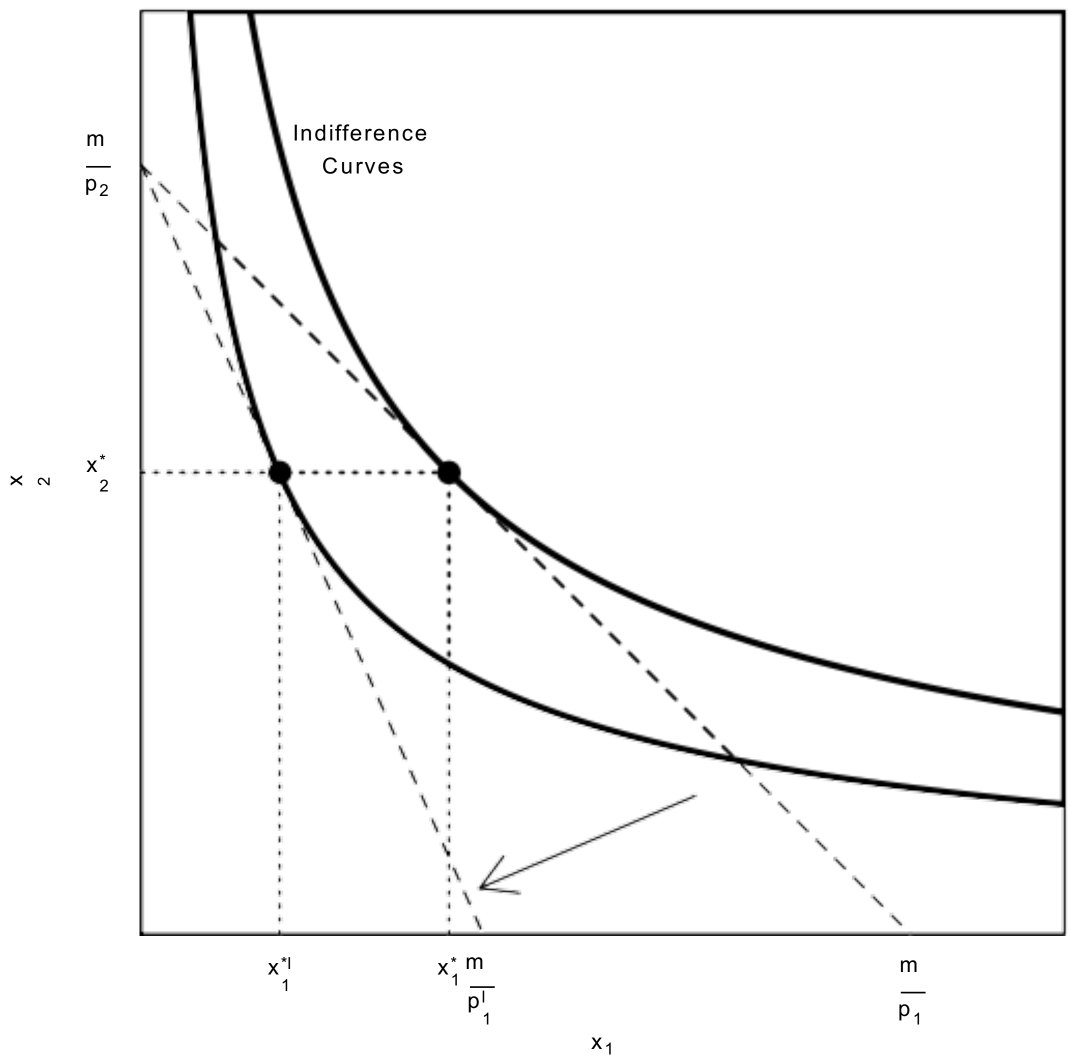

This shows the effect of a tax on good 1.

Tax policy

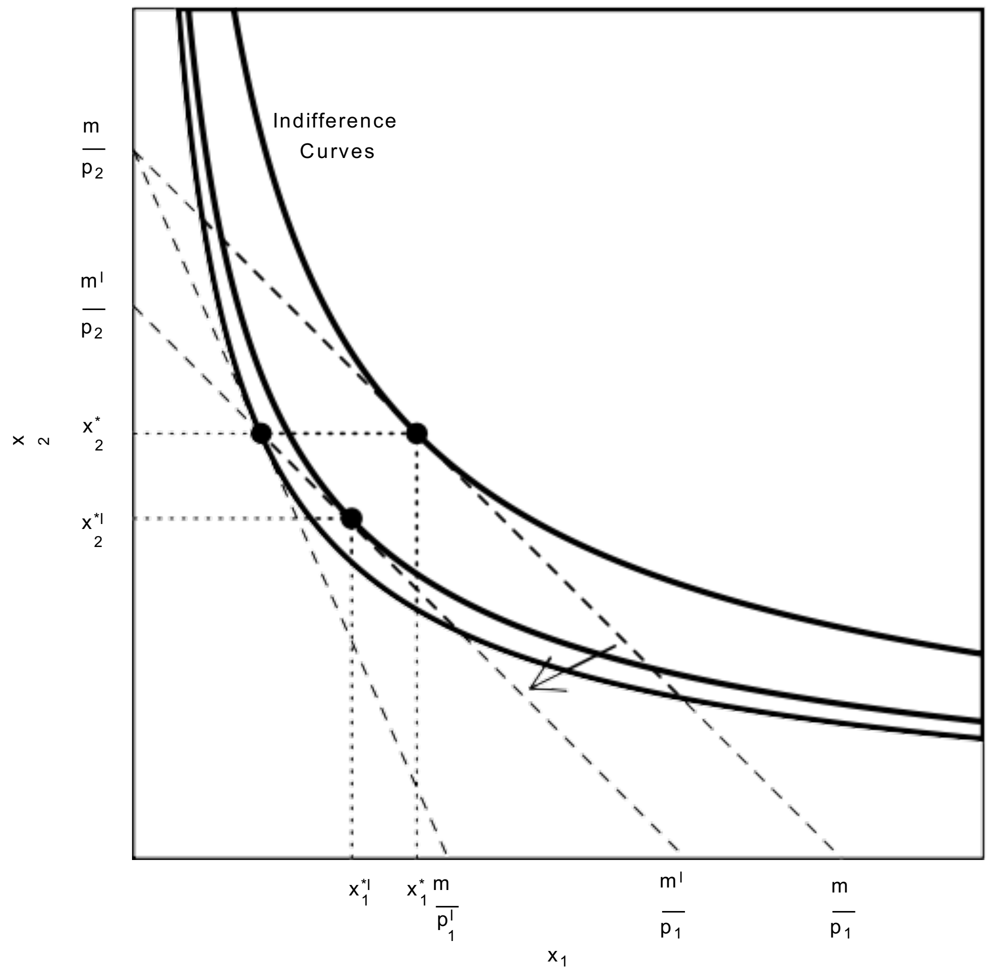

Answer: No, the consumer is better off with an income tax (and the government gets the same revenue)!

Optimal choice

- The optimal choice is the bundle in the budget set that puts the consumer on the highest indifference curve. This is what we mean by the *best bundle a consumer can afford*.
- For well-behaved preferences the optimal bundle will usually require $MRS = -p_1/p_2$ = slope of the budget line.
- We can use choice data to estimate utility functions (see the textbook for an example), and these functions can be used to evaluate policy proposals.
- When everyone faces the same prices for two goods, then everyone will have identical MRSs and will be willing to trade off the two goods in the same way!