Intermediate Microeconomics: Choice

Simon Naitram, PhD

October 9, 2020

Chapter 5: Choice

READ SECTIONS 5.1 TO 5.5. This is by far the most important chapter in this text. Don't focus too heavily on 5.4, but have a look at it. Understand the special cases in Section 5.3, but don't memorise them. Focus on understanding the optimality conditions, particularly in Section 5.1. Read Sections 5.5 very carefully and make sure you understand it entirely.

"Individuals choose the most preferred alternative available to them" is the behavioural principle of rationality. In this economic model of consumer choice, that translates to "consumers choose the most preferred bundle from their budget sets".

Combining What We Know We've examined the set of choices available to the consumer. This set of choices is the individual's budget set. It is defined by market prices and the individual's income. This gives us the budget set. We've examined the consumer's preferences. Consumer preferences are represented by a set of preference relations between all pairs of goods. We can describe these preferences numerically using utility functions. We now ask the question: given the set of choices available to you, which one do you *optimally* choose? This chapter goes through the method of choosing optimally, combining the budget set and the indifference curve.

Figure 1: All the indifference curves.

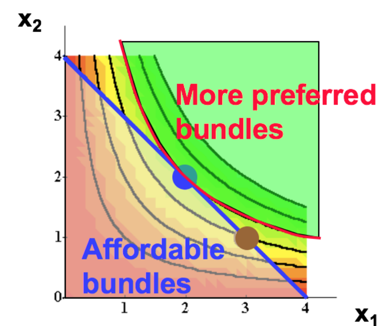

Graphical Analysis Let's first take this graphically. We restrict our attention to well-behaved preferences[1], so we focus on nicely curved indifference curves. Both out budget line and our indifference curve have been graphed on a chart with x_1 on the x-axis and x_2 on the y-axis.

Figure 2: The most preferred affordable bundle.

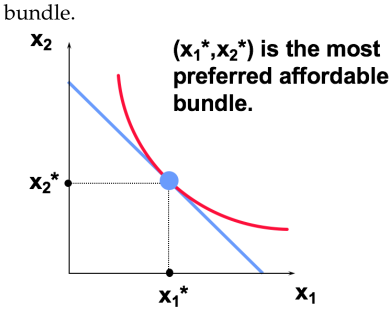

First we draw the budget constraint, which gives us the set of all combinations of (x_1, x_2) that the consumer can afford on their income. Now how do we know which combination (bundle) the consumer prefers the most? Well second, we draw the consumer's indifference curves. Now remember that the further to the right the indifference curve is, the more the consumer prefers it. In utility terms, the further to the right the indifference curve is, the higher the consumer's utility from any budget along that line.

[1] You remember the properties of well-behaved preferences right?

Optimality Condition If you draw all the indifference curves of the consumer, each parallel to one another, where do you find the bundle the consumer prefers the most? Well, the highest indifference curve the consumer can reach *on their budget* is the one that is furthest to the right while still being within the budget set. *As you move to the right, the highest indifference curve that is still affordable is the indifference curve that just touches the budget line.*

The consumer's optimal bundle is the point where the indifference curve of the consumer is tangent to the budget line. Let's think about how maximums work. If you move a little to the right, the bundle is no longer affordable. If you move a little to the left, then you're on a lower indifference curve. If you move a little up or down the budget line then you move to a lower indifference curve as well.[2] There is no direction you can move which both makes the consumer better off, and is affordable.

[2] And remember you can't simply shift the indifference curve up or down—indifference curves can't cross! Refresh your memory of why this is.

Focus on Simple Cases There are loads of weird special cases in the text. Read them but don't memorize them. Understand what they're saying, but we won't focus on them. We focus on the case of *strictly* convex preferences. This restricted focus gives us a useful outcome—there is only one unique optimal choice for every budget line given a consumer's preferences. Let's call this bundle $(x_1^*(p_1, p_2, m), x_2^*(p_1, p_2, m))$. This makes it clear that the consumer's optimal choice directly depends on the price of the commodities and their income.

Slopes Now let's begin to think about the slopes of our lines. Remember that when a straight line is tangent to a curve at a point, it means that in that very small area, the slopes of those lines are approximately equal.

What is the slope of the budget line? It is the rate at which *the market* is willing to substitute one unit of Good 1 for Good 2.[3] What is the slope of the indifference curve? It is the rate at which *the consumer* is willing to substitute one unit of Good 1 for Good 2.[4] *At the optimal, the rate at which the market is willing to substitute good 1 for good 2 is exactly equal to the rate at which the consumer is willing to substitute good 1 for good 2.*

[3] If I give up one unit of good 1, I can buy p_1/p_2 units of good 2.]

[4] If the consumer has to give up one unit of good 1, how much of good 2 do you need to give them in order to keep them exactly the same amount of satisfied?

Put mathematically, it says that the Marginal Rate of Substitution (MRS)[5] must be equal to the slope of the budget line at the optimal. That is:

$$\text{MRS} = -\frac{p_1}{p_2}. \tag{1}$$

[5] The slope of the indifference curve.

What happens if these aren't equal? A person may be willing to give up a unit of good 1 for 2 units of good 2. The market may value these goods at 1 unit of good 1 for 3 units of good 2. If a consumer gives up one unit of good 1, the market says that they can get 3 units of good 2, more than how much they were actually bargaining for. This would mean that they'd be moved to a higher indifference curve by $(-1, +3)$[6] since the swap of $(-1, +2)$ would've kept them on the same indifference curve. It means that they can afford a better bundle! It means that they weren't at an optimum if MRS $\neq -p_1/p_2$.

[6] Using (x_1, x_2).

Obviously in reality people don't reach that optimal point immediately. Maybe one week they buy two trays of eggs and one gallon of milk at the given prices. The next week maybe they decide that at those prices they'd rather have a bit more milk and a few less eggs. Maybe after a couple months they figure out a nice routine of exactly how many eggs and how much milk is optimal for them given their prices and their budget.

What if we all did this? We swap 1 unit of good 1 for p_1/p_2 units of good 2 to see if it moves us onto a higher or lower indifference curve. We swap and chop and change until we all find exactly the right balance—the optimal bundle given our preferences, prices, and income. Now both your marginal rate of substitution and my marginal rate of substitution are equal to $-p_1/p_2$. That means that at the optimal, both your and my marginal rates of substitution are exactly equal! This means that in the market, everyone adjusts their marginal valuation of milk and eggs up until the point where it's equal to everyone else's and equal to prices.

Equality of MRS The cool thing now is that regardless of everyone's different preferences and tastes and incomes, they all end up with the exact same marginal rate of substitution. Even better, since we know this rate must be equal to the ratio of prices, and we know what prices are, then we know what everyone's marginal rate of substitution is at the optimum! Everyone agrees on how much of good 2 they'd be willing to give up in order to get an extra unit of good 1.

This idea that we can measure everyone's marginal valuation of a good by observing prices is extremely important. It turns out that in microeconomics, prices are the key to unlocking the balance of the universe[7].

[7] Maybe not the *entire* universe, but almost.

Chapter 5 Appendix: Choice

READ THE APPENDIX to Chapter 5. Work through it carefully.

Calculus First we discussed a graphical analysis of the consumer's optimal choice given their budget set. Now we want to analyse the same question using the tools of calculus. Why is this important? Well mostly because it is extraordinarily difficult to actually give an answer to any real problem using graphical methods. With graphs we can only sketch the idea. Further, we can only sketch graphs of two goods. With calculus we can describe a person's entire consumption bundle. Further, we can do so in interesting ways, incorporating their income, leisure, and even saving! And with all of that, calculus gives us *exact* answers. This means we can actually use examples with numbers!

Setup We're going to set up our problem in exactly the same way we discussed during the tutorial. We first need to figure out what our objective function is. The objective function is function that we're trying to maximise. For the consumer, this objective function is their utility function. Remember that the consumer is trying to get as much utility as possible! Then, we need to figure out what the consumer's constraint is. That's easy—it's their budget constraint. So we have an objective function and a constraint. Remember that the variables the consumer is choosing are the amounts of each commodity they want to consume. Remember also that the prices are fixed, and their income is also fixed.

The Problem The consumer's problem is therefore to maximise utility $u(x_1, x_2)$ by choosing the amount of goods 1 and 2 (i.e. x_1 and x_2) subject to their budget constraint $p_1x_1 + p_2x_2 = m$. We are going to use the Lagrangian method of solving the consumer's problem. We can write this as:

$$\begin{aligned} \max_{x_1, x_2} \quad & u(x_1, x_2) \\ \text{s.t.} \quad & p_1x_1 + p_2x_2 = m \end{aligned} \tag{2}$$

Step 1 Rewrite the constraint by bringing everything onto one side. This gives $p_1x_1 + p_2x_2 - m = 0$.

Step 2 Set up the Lagrangian in the form of L = objective - λ(constraint). This means that we can set up the problem as:

$$\max_{x_1, x_2} \mathcal{L} = u(x_1, x_2) - \lambda(p_1x_1 + p_2x_2 - m) \tag{3}$$

Step 3 Take partial derivatives of the Lagrangian with respect to x_1, x_2, and λ. This gives us:

$$\frac{\partial \mathcal{L}}{\partial x_1} = \frac{\partial u(x_1^*, x_2^*)}{\partial x_1} - \lambda p_1 = 0 \tag{4}$$
$$\frac{\partial \mathcal{L}}{\partial x_2} = \frac{\partial u(x_1^*, x_2^*)}{\partial x_2} - \lambda p_2 = 0 \tag{5}$$
$$\frac{\partial \mathcal{L}}{\partial \lambda} = p_1 x_1^* + p_2 x_2^* - m = 0 \tag{6}$$

Note that I'm using x_1^* and x_2^* now. This is because the above conditions identify the values of x_1 and x_2 that define a maximum of the Lagrangian. Therefore, we're now working with our optimal values.

Step 4 We can now use these three equations to solve for our three unknowns. Usually we don't care about the value of λ. We typically ignore it and simply solve for x_1 and x_2. However, I'll let you in on a little secret that helps you understand it a little better. We consider the variable λ to be the 'shadow price' of the fixed variable in the constraint. So in this case, λ would be the shadow price of income. That is, how much would a tiny increase in the consumer's income would increase their optimal utility. Or alternatively, how much the consumer would be willing to pay to relax the constraint (increase income) slightly. Basically think of it as how much they would value a raise!

MRS In this case, we can set:

$$\frac{\partial u(x_1^*, x_2^*)}{\partial x_1} \times \frac{1}{p_1} = \lambda \tag{7}$$
$$\frac{\partial u(x_1^*, x_2^*)}{\partial x_2} \times \frac{1}{p_2} = \lambda \tag{8}$$
$$\tag{9}$$

Since they're both equal to λ, then they're both equal to one another. Setting these two equations equal to one another, we get:

$$\frac{\partial u(x_1^*, x_2^*)}{\partial x_1} \times \frac{1}{p_1} = \frac{\partial u(x_1^*, x_2^*)}{\partial x_2} \times \frac{1}{p_2}$$
$$\frac{\partial u(x_1^*, x_2^*)/\partial x_1}{\partial u(x_1^*, x_2^*)/\partial x_2} = \frac{p_1}{p_2} \tag{10}$$

which is exactly our graphical condition for optimality. Remember we said that the point of tangency between the budget line and the utility curve is exactly where they just touch, and the slopes are exactly equal only at that point? Well, we've just derived that exact same condition from our calculus problem.[8]

[8] Remember that the MRS is $-\mathrm{MU}_1/\mathrm{MU}_2$, and the slope of the budget line is also $-p_1/p_2$. They're both negative. But it also means that we can ignore the negative signs since both sides are negative. That is, multiply both sides by -1.

Demand We get our first taste of demand. Note that our optimal values of x_1 and x_2 depend on the market prices and the consumer's income. This means that we can rewrite the optimal choices as function of these exogenous variables. It's easy to see that since the budget line pivots in response to changes in prices, then a change in prices will change the optimal allocation of x_1 and x_2. It's also easy to see that since an increase in income shifts the budget line inward or outward, a change in income will also affect this optimal allocation of (x_1, x_2). This means we can rewrite the optimal values as functions:

$$x_1^*(p_1, p_2, m) \tag{11}$$

$$x_2^*(p_1, p_2, m) \tag{12}$$

These are now our **demand functions**.

Example 1 Let's work with the common Cobb-Douglas utility function. That is given by $u(x_1, x_2) = x_1^\alpha x_2^\beta$. Importantly, we can take logs of the utility function and this is a monotonic transformation. It means that logs do not change the ordering of the utility function, so that it represents the exact same set of preferences over x_1 and x_2. We can rewrite the utility function as: $u(x_1, x_2) = \alpha \ln(x_1) + \beta \ln(x_2)$. The constraint is $p_1 x_1 + p_2 x_2 = m$. We rewrite this as $p_1 x_1 + p_2 x_2 - m = 0$. We then form the Lagrangian as L = objective $-\lambda$(constraint). This gives:

$$\max_{x_1, x_2, \lambda} \quad \mathcal{L} = \alpha \ln(x_1) + \beta \ln(x_2) - \lambda(p_1 x_1 + p_2 x_2 - m) \tag{13}$$

Let's take the partial derivatives of this function with respect to x_1, x_2, and λ. This gives us three first order conditions:

$$\frac{\partial \mathcal{L}}{\partial x_1} = \frac{\alpha}{x_1} - \lambda p_1 = 0 \tag{14}$$

$$\frac{\partial \mathcal{L}}{\partial x_2} = \frac{\beta}{x_2} - \lambda p_2 = 0 \tag{15}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = p_1 x_1 + p_2 x_2 - m = 0 \tag{16}$$

So we have three equations which define the conditions for utility to be at a constrained maximum. We now need to solve for the main choice variables x_1 and x_2. How do we do so? Well let's try to solve the first two equations for λ and then set them equal to one another.

$$\frac{\alpha}{x_1} - \lambda p_1 = 0$$

$$\frac{\alpha}{x_1} = \lambda p_1$$

$$\frac{\alpha}{p_1 x_1} = \lambda \tag{17}$$

We do the same for the second first order condition, giving us

$$\frac{\beta}{p_2 x_2} = \lambda \tag{18}$$

We can then set them equal to one another. This gives us

$$\begin{aligned} \frac{\alpha}{p_1 x_1} &= \lambda = \frac{\beta}{p_2 x_2} \\ \frac{\alpha}{p_1 x_1} &= \frac{\beta}{p_2 x_2} \end{aligned} \tag{19}$$

Now we have eliminated λ, we can solve for either x_1 as a function of x_2, or vice versa. Let's try:

$$\begin{aligned} \frac{\alpha}{x_1 p_1} &= \frac{\beta}{p_2 x_2} \\ \frac{\alpha}{p_1} &= x_1 \frac{\beta}{p_2 x_2} \\ \frac{\alpha}{\beta p_1} &= x_1 \frac{1}{p_2 x_2} \\ \frac{\alpha p_2 x_2}{\beta p_1} &= x_1 \end{aligned} \tag{20}$$

Then we can take this and substitute it into the third first order condition. That is:

$$p_1 \frac{\alpha p_2 x_2}{\beta p_1} + p_2 x_2 - m = 0 \tag{21}$$

Now that we have eliminated x_1, we can simply re-arrange this equation to solve for optimal x_2 as a function of (p_1, p_2, m). Let's try!

$$\begin{aligned} p_1 \frac{\alpha p_2 x_2}{\beta p_1} + p_2 x_2 - m &= 0 \\ \frac{\alpha p_2 x_2}{\beta} + p_2 x_2 - m &= 0 \\ p_2 \left(\frac{\alpha x_2}{\beta} + x_2 \right) - m &= 0 \\ p_2 \left(\frac{\alpha x_2}{\beta} + x_2 \right) &= m \\ \frac{\alpha x_2}{\beta} + x_2 &= \frac{m}{p_2} \\ x_2 \left(\frac{\alpha}{\beta} + 1 \right) &= \frac{m}{p_2} \\ x_2 \left(\frac{\alpha}{\beta} + \frac{\beta}{\beta} \right) &= \frac{m}{p_2} \\ x_2 \left(\frac{\alpha + \beta}{\beta} \right) &= \frac{m}{p_2} \\ x_2^* &= \left(\frac{\beta}{\alpha + \beta} \right) \frac{m}{p_2} \end{aligned} \tag{22}$$

If we substitute this value for x_2 into the equation for x_1, we get:

$$
\begin{aligned}
x_1 &= \frac{\alpha p_2 \left(\frac{\beta}{\alpha+\beta}\right) \frac{m}{p_2}}{\beta p_1} \\
x_1 &= \frac{\alpha p_2}{\beta p_1} \left(\frac{\beta}{\alpha+\beta}\right) \frac{m}{p_2} \\
x_1 &= \frac{\alpha p_2}{p_1} \left(\frac{1}{\alpha+\beta}\right) \frac{m}{p_2} \\
x_1 &= \frac{\alpha}{p_1} \left(\frac{1}{\alpha+\beta}\right) \frac{m}{1} \\
x_1^* &= \left(\frac{\alpha}{\alpha+\beta}\right) \frac{m}{p_1} \qquad (23)
\end{aligned}
$$

Which is in exactly the same format as for x_2. These two functions are the consumer's demand functions for x_1 and x_2. They're both clearly expressed as a function of p_1 or p_2, and m.

Now how do we interpret this? We know that $\alpha + \beta = 1$.[9] This means that $\alpha/(\alpha+\beta)$ is the share of an individual's consumption that they dedicate to good 1. Similarly, $\beta/(\alpha+\beta)$ is the share the consumer spends on good 2. This means that regardless of the consumer's income or the price of the goods, the share the consumer spends on each good remains constant!

[9] We can show that we can always do a monotonic transformation of a Cobb-Douglas utility function that makes the exponents α and β sum to 1. This means that $\beta = 1 - \alpha$. Check page 65 of Varian.