Simon Fraser University Prof. Karaivanov
Department of Economics Econ 201

THE CONSUMER PROBLEM AND HOW TO SOLVE IT

In these notes I provide a detailed description of the main problem in consumer theory - the "consumer problem" as well as the algorithm to solve it. Remember what consumer theory was all about: "The consumer chooses her most preferred bundle from her budget set". The question is how is this bundle chosen? How do we know how much of good 1 (x_1) and how much of good 2 (x_2) to consume? Thus is exactly what the consumer problem is.

The first part of the consumer theory statement says that a consumer chooses her "most preferred" bundle. This is the same as saying that she would choose the bundle (x_1, x_2) that provides her with the highest level of utility. However, not any bundle can be chosen. The consumer can only choose among the bundles from her budget set as all others are unaffordable for her, i.e. even though they may give her higher utility she cannot afford to buy them. We call the most preferred bundle in the budget set **the optimal bundle.**

Thus the question is how to choose the bundle from the budget set that yields maximum utility. When preferences are monotonic (which is what we assume) meaning that "more is better than less" it is clear that this maximum utility bundle must in fact lie on the budget line, i.e. it must be one of the bundles that completely exhaust the person's budget. If this were not true and since there are only two goods to spend money on, it would mean that the consumer is not spending all of her budget i.e. she can buy more of each good and hence get to a higher utility level. In other words, if we suppose that the optimal bundle (x_1^*, x_2^*) is not on the budget line, we would have: $p_1x_1^* + p_2x_2^* < m$ but then we can buy a little bit more of each good with the remainder of our money which would make us happier, i.e. (x_1^*, x_2^*) cannot be optimal.

What the above tells us is that the consumer problem is: "From all bundles on the budget line, find the one that gives you maximum utility", i.e. "find the quantities of good 1 and good 2 (x_1, x_2) which maximize the utility function $u(x_1, x_2)$ and which lie on the budget line, i.e. for which it is true that $p_1x_1 + p_2x_2 = m$. Let us write this down using mathematical notation. The problem looks as follows:

$$\max_{x_1, x_2} \; u(x_1, x_2)$$
$$s.t. \; p_1x_1 + p_2x_2 = m$$

Let us explain what the above symbols mean. The "max" word on the first line means that we want to maximize what follows to the right of it, i.e. the utility function $u(x_1, x_2)$. The x_1, x_2 below the "max" say that the maximization is performed by searching over (choosing) values for the quantities consumed of the two goods x_1 and x_2. Indeed, this is the only thing that the consumer can actually choose - the prices and her income are given - they cannot be influenced by what the consumer does, i.e. by how much of the goods she buys. The second line starts with the abbreviation "s.t." which means "subject to" - what it says is that the quantities that we can choose (x_1 and x_2) cannot take any arbitrary values but must be such that they satisfy a constraint - they must be on the budget line (i.e. satisfy the **budget constraint**). The consumer's choice of quantities to consume must be such that her total expenditure equals her income. She can't spend more than what she has and since she likes more better than less, it will be stupid to spend less than what she has (remember these two goods are everything you could ever buy in her world).

So how do we solve the above problem? How do we find the values of x_1 and x_2 that maximize utility and that can be just afforded? To explain the procedure we will take a specific example.

Suppose the utility function is $u(x_1, x_2) = x_1^\alpha x_2^{1-\alpha}$, where $\alpha \in (0, 1)$, i.e. α is some fraction, like 1/2, 1/3, 3/8, etc. The consumer's problem is then:

$$\begin{aligned}&\max_{x_1,x_2}\ x_1^\alpha x_2^{1-\alpha}\\ &s.t.\ p_1x_1 + p_2x_2 = m\end{aligned}$$

That's a tough problem to solve. We have to maximize a function of two variables (x_1 and x_2) and on top of that they have to satisfy a constraint. What do we do? Well, the best thing to do when you have a hard problem to solve is simplify it, i.e. replace it with an equivalent problem that is much easier to solve. The following steps show how to do that:

Step 0 (OPTIONAL): Check if the utility function can be simplified (made easier to work with) by applying a monotonic transformation to it.

What do we mean "simplify the utility function"? We will maximize the function so it must be in a form that is easy to differentiate. Notice that this step is **optional**! We do not have to do this in all cases. Actually, it is probably only in the Cobb-Douglas case that there will be a need to do a monotonic transformation!

In our particular example since sums are easier to differentiate than products, it is a good idea to take log of the Cobb-Douglas utility function $x_1^\alpha x_2^{1-\alpha}$. Since log is a monotonic transformation, the resulting function, $\log(u(x_1, x_2)$ corresponds to the exactly the same preferences as the original one and hence we will get exactly the same answer in terms of the optimal x_1 and x_2.

Doing the log monotonic transformation we transform the consumer problem into the following equivalent problem:

$$\begin{aligned}&\max_{x_1,x_2}\ \alpha \log x_1 + (1-\alpha)\log x_2\\ &s.t.\ p_1x_1 + p_2x_2 = m\end{aligned}$$

Remember that: $\log(xy) = \log(x) + \log(y)$ and $\log(x^a) = a\log(x)$ which is what we used above!

Ok, we made things simpler but we still have a function of two variables and a constraint. It would be good if we can somehow turn our problem into a maximization of a function of just one variable with no constraint! Seems like this is too much to ask for but we can use a trick and achieve that in the following step.

Step 1: Simplify the problem (further) by expressing x_2 from the budget constraint and plug it into the utility function

This is the first necessary step to solve the consumer problem. What we do is notice that once we are thinking of choosing a value for x_1 this automatically implies the value that x_1 must have if the budget has to be exhausted. Graphically, if you choose an x_1 and draw a vertical line upwards from it, the point where this line crosses the budget line determines the value of x_2 that you can just afford. Algebraically, since the optimal bundle must lie on the budget line we must have:

$$p_2x_2 = m - p_1x_1$$

or,

$$x_2 = \frac{m}{p_2} - \frac{p_1}{p_2}x_1 \tag{1}$$

What is the above? It is exactly the equation of our budget line! Plug in for example $x_1 = 0$ and we get $x_2 = \frac{m}{p_2}$ - the vertical intercept. Alternatively, if we set $x_2 = 0$ we get $x_1 = \frac{m}{p_1}$ - the horizontal intercept of the BL. Since the optimal bundle *must lie* on the budget line, the above relationship between x_2 and x_1 must hold.

How does this help us? Well, we can use (1) and substitute the expression on the right hand side into the utility function we want to maximize. With this we kill two birds with the same stone - first, we take care of the budget constraint and second, we get to maximize a function of a single variable only, x_1. The consumer's problem then becomes:

$$\max_{x_1} \ \alpha \log x_1 + (1-\alpha)\log(\frac{m}{p_2} - \frac{p_1}{p_2}x_1)$$

subject to nothing! Why? Because we already used the constraint to express x_2 in terms of x_1! Notice how much simpler this problem is - we have a function of a single variable and there are no constraints to mess up with our maximization. This is also a problem that we know how to solve! How do we maximize a function of a single variable - take the first derivative and set it to zero[1].

Step 2: Maximize the function of one variable obtained by plugging x_2 expressed in terms of x_1 into the utility function by taking its first derivative and setting it to zero to find the optimal amount of good 1 consumed.

[1] In general we also need to verify that the second derivative is positive in order to be sure that we have found a maximum. In economics however we will most often deal with strictly concave functions that automatically has positive second derivatives at any point so this is not a problem.

Let us do that. The equation that we obtain by setting the first derivative to zero is called "**the first order condition**". In our example it looks as follows:

$$\alpha\frac{1}{x_1} + (1-\alpha)\frac{1}{\frac{m}{p_2} - \frac{p_1}{p_2}x_1}(-\frac{p_1}{p_2}) = 0$$

Why is that? Remember than the derivative of $\log x$ is $\frac{1}{x}$. Also to differentiate the second log we use the chain rule - the derivative of $g(f(x))$ where f and g are two functions is the derivative of g at $f(x)$, i.e. $g'(f(x))$ times the derivative of $f(x)$. In our case the derivative of the log (the g) is $\frac{1}{\frac{m}{p_2}-\frac{p_1}{p_2}x_1}$ and the derivative of $\frac{1}{\frac{m}{p_2}-\frac{p_1}{p_2}x_1}$ (the f) is $-\frac{p_1}{p_2}$.

We need to solve the above equation for the quantity of good 1 consumed at the optimum, x_1. So let's do a little algebra. The equation can be written as:

$$\frac{\alpha}{x_1} - \frac{(1-\alpha)p_1}{(\frac{m}{p_2} - \frac{p_1}{p_2}x_1)p_2} = 0$$

or,

$$\frac{\alpha}{x_1} = \frac{(1-\alpha)p_1}{m - p_1 x_1}$$

which is:

$$\alpha m - \alpha p_1 x_1 = p_1 x_1 - \alpha p_1 x_1$$

from where we get:

$$x_1^* = \frac{\alpha m}{p_1}$$

where x_1^* is the optimal amount of good 1 that is chosen by the consumer. How do we find the amount of good 2 consumed?

Step 3: Find the optimal quantity of good 2 consumed by plugging the optimal quantity of good 1 found in Step 2 back into the budget constraint.

Fortunately we don't have to go through this whole process again but need only remember that we know how to express x_2 in terms of x_1 from the budget constraint. Thus we obtain:

$$x_2^* = \frac{m}{p_2} - \frac{p_1}{p_2}x_1^* = \frac{m}{p_2} - \frac{p_1}{p_2}\frac{\alpha m}{p_1} = \frac{(1-\alpha)m}{p_2}$$

The optimal quantities x_1^* and x_2^* are called the **demand functions** for good 1 and good 2. They tell us how much of the two goods the consumer would choose (demand) given her income, m and the prices p_1 and p_2. Notice that the demand functions make a lot of sense: if income m is increased people would demand more of each of the goods, while if a price of a good increases people would demand less.

There is another property of the demand functions resulting from the Cobb-Douglas preferences that needs mentioning. Notice that the total expenditure (price times quantity demanded) of good 1 is $p_1 x_1^* = \alpha m$ which is a fraction α of total income m. Similarly, the total expenditure on good 2 at the optimal bundle is: $p_2 x_2^* = (1-\alpha)m$, i.e. a share $1-\alpha$ of total income is spent on good 2. But α was the exponent on x_1 in the utility function and $1-\alpha$! So this provides us with an easy way of remembering what the Cobb-Douglas demands equal to. What it also means is that a person with C-D preferences spends **constant shares** of her income on each of the goods, i.e. if her income increases the share of her spending increases proportionally.

Another example: The utility function doesn't have to be Cobb-Douglas to follow the above solution algorithm - it applies to whatever utility function we have. For example let us take: $u(x_1,x_2)=\sqrt{x_1}+\sqrt{x_2}$ and go through the above steps.

Step 0: This function is already in nice form to differentiate (it is a sum) so we skip this optional step.

Step 1: Everything is exactly as before, $x_2=\frac{m}{p_2}-\frac{p_1}{p_2}x_1$ is substituted into the utility function to arrive at the following simple problem:

$$\max_{x_1}\ \sqrt{x_1}+\sqrt{\frac{m}{p_2}-\frac{p_1}{p_2}x_1}$$

Step 2: Now we need to take the first derivative of the above and set it to zero. Remember that the derivative of the function $\sqrt{x}$ is $\frac{1}{2\sqrt{x}}$. In general the derivative of x^a is ax^{a-1} and for the square root function we just have $a=1/2$. Using this knowledge we obtain the following first order condition:

$$\frac{1}{2\sqrt{x_1}}+\frac{1}{2\sqrt{\frac{m}{p_2}-\frac{p_1}{p_2}x_1}}(-\frac{p_1}{p_2})=0$$

This can be simplified to:

$$\frac{1}{\sqrt{x_1}}=\frac{p_1}{p_2\sqrt{\frac{m}{p_2}-\frac{p_1}{p_2}x_1}}$$

or, squaring both sides:

$$\frac{1}{x_1}=\frac{p_1^2}{p_2^2(\frac{m}{p_2}-\frac{p_1}{p_2}x_1)}$$

This in turn is equivalent to:

$$p_2m-p_2p_1x_1=p_1^2x_1$$

or,

$$x_1^*=\frac{p_2m}{p_1(p_1+p_2)}$$

which is the demand for good 1. Notice again that it is increasing in income, m and decreasing in the price of the good, p_1 but in this case it also depends on the price of good 2!

Step 3: We use the budget constraint again to find the demand for good 2:

$$x_2^*=\frac{m}{p_2}-\frac{p_1}{p_2}x_1^*=\frac{m}{p_2}-\frac{p_1}{p_2}\frac{p_2m}{p_1(p_1+p_2)}=\frac{p_1m}{p_2(p_1+p_2)}$$

Notice that this time the expenditure shares on good 1 and good 2 are not constant but move with the prices.

An alternative (sometimes short-cut) way of solving the consumer problem

We can use our insight that if preferences are well-behaved (strictly monotonic and convex in addition to reflexive, transitive and complete) the indifference curve through the optimal bundle and the budget line must be tangent i.e. we must have:

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

at the optimum. But we saw that we can express the MRS (the slope of the IC) in terms of the **marginal utilities** for the two goods:

$$MRS = -\frac{MU_1}{MU_2}$$

Thus we could just start from the following condition:

$$\frac{MU_1}{MU_2} = \frac{p_1}{p_2}$$

and use it together with the budget constraint to find the optimal bundle. How to find the marginal utilities however? Remember they were equal to the **partial derivatives** of the utility function with respect to x_1 and x_2.

Let us then look at our example above. For the C-D utility $u(x_1, x_2) = x_1^\alpha x_2^{1-\alpha}$ we have that:

$$MU_1 = \alpha x_1^{\alpha-1} x_2^{1-\alpha}$$

Remember, we take partial derivatives by treating all variables we are not differentiating with respect to as constants. In this case x_2 is such variable so we just copy $x_2^{1-\alpha}$ in the expression for the partial derivative with respect to x_1. Similarly

$$MU_2 = (1-\alpha) x_1^\alpha x_2^{-\alpha}$$

Taking the ratio of the marginal utilities and setting it to the price ratio we get:

$$\frac{MU_1}{MU_2} = \frac{\alpha x_1^{\alpha-1} x_2^{1-\alpha}}{(1-\alpha) x_1^\alpha x_2^{-\alpha}} = \frac{p_1}{p_2}$$

That can be simplified to:

$$\frac{\alpha x_2}{(1-\alpha) x_1} = \frac{p_1}{p_2}$$

or,

$$\alpha p_2 x_2 = (1-\alpha) p_1 x_1$$

As before the optimal bundle must be on the budget line so we still must have that: $x_2 = \frac{m}{p_2} - \frac{p_1}{p_2} x_1$. We can plug that into the condition for an optimum above to get:

$$\alpha p_2 \left(\frac{m}{p_2} - \frac{p_1}{p_2} x_1\right) = (1-\alpha) p_1 x_1$$

or, opening brackets and cancelling p_2 :

$$\alpha m - \alpha p_1 x_1 = p_1 x_1 - \alpha p_1 x_1$$

so

$$x_1^* = \frac{\alpha m}{p_1}$$

just as before. We can then follow Step 3 from the original algorithm to solve for x_2^*.

Thus using this way we obtain the same answer. One can use any correct way they wish to solve for the optimal bundle (the demand functions, x_1^* and x_2^*). If your algebra is correct you should get exactly the same answer. The first way is more general and more intuitive (we see what is being maximized subject to what) and does not require knowledge of partial derivatives. The second way has sometimes the advantage of less messy algebra but this is by no means the general case and this method will fail when the IC is not tangent to the budget line and/or when there is a corner solution.