CHAPTER 5

CHOICE

In this chapter we will put together the budget set and the theory of preferences in order to examine the optimal choice of consumers. We said earlier that the economic model of consumer choice is that people choose the best bundle they can afford. We can now rephrase this in terms that sound more professional by saying that "consumers choose the most preferred bundle from their budget sets."

5.1 Optimal Choice

A typical case is illustrated in Figure 5.1. Here we have drawn the budget set and several of the consumer's indifference curves on the same diagram. We want to find the bundle in the budget set that is on the highest indifference curve. Since preferences are well-behaved, so that more is preferred to less, we can restrict our attention to bundles of goods that lie *on* the budget line and not worry about those *beneath* the budget line.

Now simply start at the right-hand corner of the budget line and move to the left. As we move along the budget line we note that we are moving to higher and higher indifference curves. We stop when we get to the highest indifference curve that just touches the budget line. In the diagram, the bundle of goods that is associated with the highest indifference curve that just touches the budget line is labeled (x_1^*, x_2^*).

The choice (x_1^*, x_2^*) is an **optimal choice** for the consumer. The set of bundles that she prefers to (x_1^*, x_2^*)—the set of bundles *above* her indifference curve—doesn't intersect the bundles she can afford—the bundles *beneath* her budget line. Thus the bundle (x_1^*, x_2^*) is the best bundle that the consumer can afford.

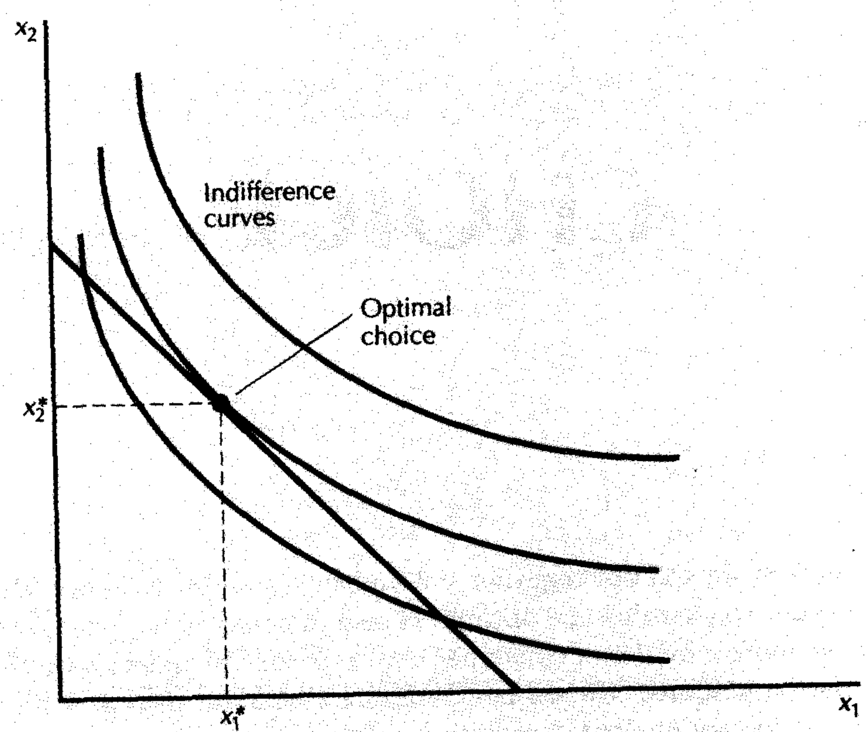

Optimal choice. **The optimal consumption position is where the indifference curve is tangent to the budget line.**

Note an important feature of this optimal bundle: at this choice, the indifference curve is *tangent* to the budget line. If you think about it a moment you'll see that this has to be the case: if the indifference curve weren't tangent, it would cross the budget line, and if it crossed the budget line, there would be some nearby point on the budget line that lies above the indifference curve—which means that we couldn't have started at an optimal bundle.

Does this tangency condition really *have* to hold at an optimal choice? Well, it doesn't hold in *all* cases, but it does hold for most interesting cases. What is always true is that at the optimal point the indifference curve can't cross the budget line. So when does "not crossing" imply tangent? Let's look at the exceptions first.

First, the indifference curve might not have a tangent line, as in Figure 5.2. Here the indifference curve has a kink at the optimal choice, and a tangent just isn't defined, since the mathematical definition of a tangent requires that there be a unique tangent line at each point. This case doesn't have much economic significance—it is more of a nuisance than anything else.

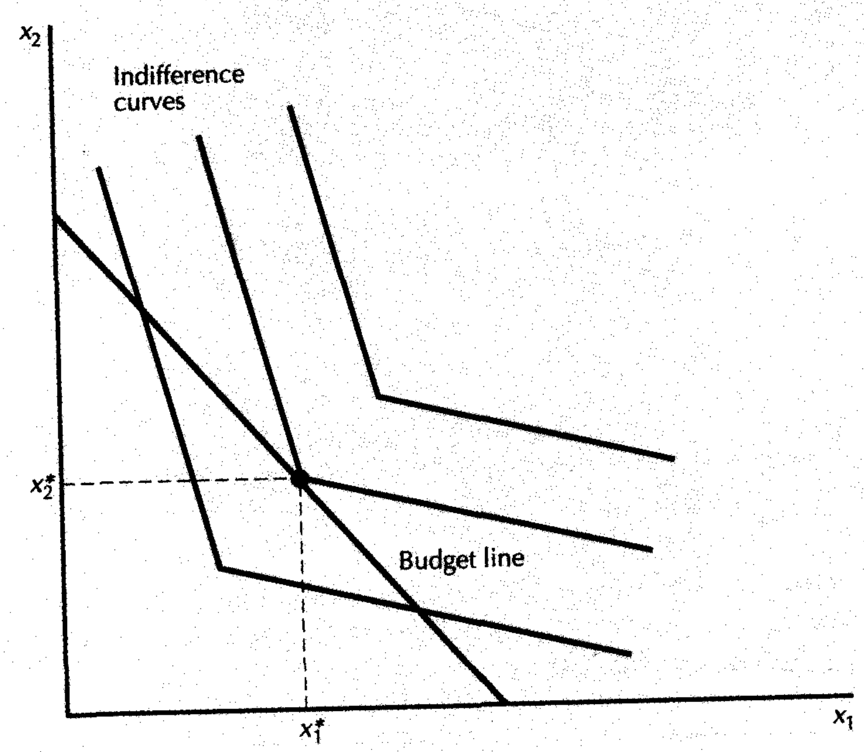

Kinky tastes. Here is an optimal consumption bundle where the indifference curve doesn't have a tangent.

The second exception is more interesting. Suppose that the optimal point occurs where the consumption of some good is zero as in Figure 5.3. Then the slope of the indifference curve and the slope of the budget line are different, but the indifference curve still doesn't *cross* the budget line.

We say that Figure 5.3 represents a **boundary optimum**, while a case like Figure 5.1 represents an **interior optimum**.

If we are willing to rule out "kinky tastes" we can forget about the example given in Figure 5.2.[1] And if we are willing to restrict ourselves only to *interior* optima, we can rule out the other example. If we have an interior optimum with smooth indifference curves, the slope of the indifference curve and the slope of the budget line must be the same ... because if they were different the indifference curve would cross the budget line, and we couldn't be at the optimal point.

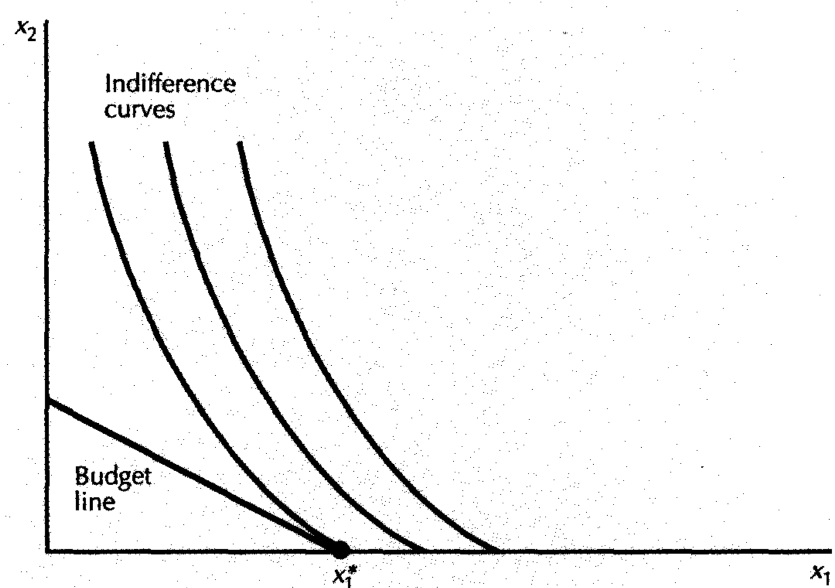

Boundary optimum. The optimal consumption involves consuming zero units of good 2. The indifference curve is not tangent to the budget line.

We've found a necessary condition that the optimal choice must satisfy. If the optimal choice involves consuming some of both goods—so that it is an interior optimum—then necessarily the indifference curve will be tangent to the budget line. But is the tangency condition a *sufficient* condition for a bundle to be optimal? If we find a bundle where the indifference curve is tangent to the budget line, can we be sure we have an optimal choice?

Look at Figure 5.4. Here we have three bundles where the tangency condition is satisfied, all of them interior, but only two of them are optimal.

[1] Otherwise, this book might get an R rating.

So in general, the tangency condition is only a necessary condition for optimality, not a sufficient condition.

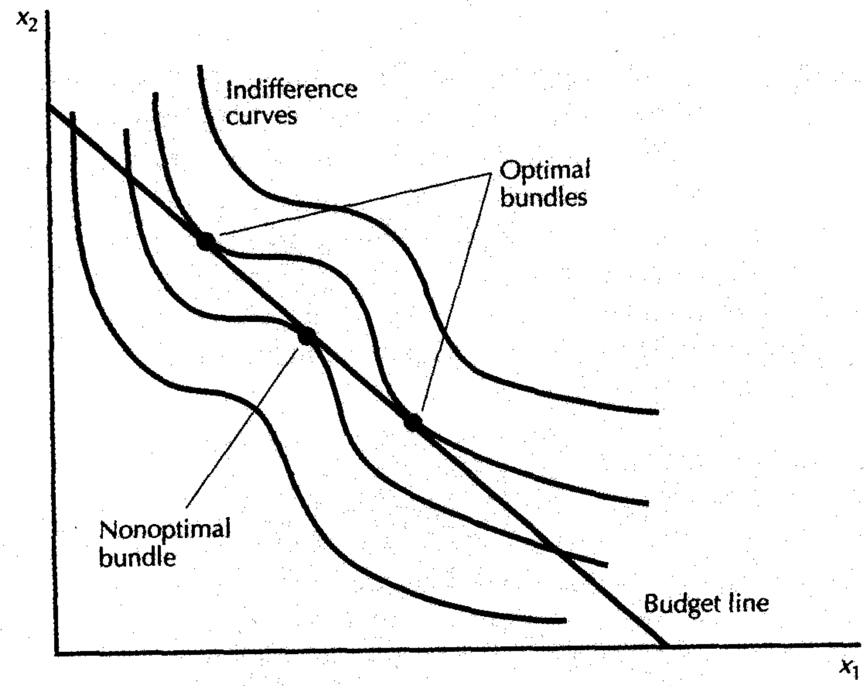

More than one tangency. Here there are three tangencies, but only two optimal points, so the tangency condition is necessary but not sufficient.

However, there is one important case where it is sufficient: the case of convex preferences. In the case of convex preferences, any point that satisfies the tangency condition must be an optimal point. This is clear geometrically: since convex indifference curves must curve away from the budget line, they can't bend back to touch it again.

Figure 5.4 also shows us that in general there may be more than one optimal bundle that satisfies the tangency condition. However, again convexity implies a restriction. If the indifference curves are *strictly* convex—they don't have any flat spots—then there will be only one optimal choice on each budget line. Although this can be shown mathematically, it is also quite plausible from looking at the figure.

The condition that the MRS must equal the slope of the budget line at an interior optimum is obvious graphically, but what does it mean economically? Recall that one of our interpretations of the MRS is that it is that rate of exchange at which the consumer is just willing to stay put. Well, the market is offering a rate of exchange to the consumer of $-p_1/p_2$—if you give up one unit of good 1, you can buy p_1/p_2 units of good 2. If the consumer is at a consumption bundle where he or she is willing to stay put, it must be one where the MRS is equal to this rate of exchange:

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

Another way to think about this is to imagine what would happen if the MRS were different from the price ratio. Suppose, for example, that the MRS is $\Delta x_2/\Delta x_1 = -1/2$ and the price ratio is $1/1$. Then this means the consumer is just willing to give up 2 units of good 1 in order to get 1 unit of good 2—but the market is willing to exchange them on a one-to-one basis. Thus the consumer would certainly be willing to give up some of good 1 in order to purchase a little more of good 2. Whenever the MRS is different from the price ratio, the consumer cannot be at his or her optimal choice.

5.2 Consumer Demand

The optimal choice of goods 1 and 2 at some set of prices and income is called the consumer's **demanded bundle**. In general when prices and income change, the consumer's optimal choice will change. The **demand function** is the function that relates the optimal choice—the quantities demanded—to the different values of prices and incomes.

We will write the demand functions as depending on both prices and income: $x_1(p_1, p_2, m)$ and $x_2(p_1, p_2, m)$. For each different set of prices and income, there will be a different combination of goods that is the optimal choice of the consumer. Different preferences will lead to different demand functions; we'll see some examples shortly. Our major goal in the next few chapters is to study the behavior of these demand functions—how the optimal choices change as prices and income change.

5.3 Some Examples

Let us apply the model of consumer choice we have developed to the examples of preferences described in Chapter 3. The basic procedure will be the same for each example: plot the indifference curves and budget line and find the point where the highest indifference curve touches the budget line.

Perfect Substitutes

The case of perfect substitutes is illustrated in Figure 5.5. We have three possible cases. If $p_2 > p_1$, then the slope of the budget line is flatter than the slope of the indifference curves. In this case, the optimal bundle is where the consumer spends all of his or her money on good 1. If $p_1 > p_2$, then the consumer purchases only good 2. Finally, if $p_1 = p_2$, there is a whole range of optimal choices—any amount of goods 1 and 2 that satisfies the budget constraint is optimal in this case. Thus the demand function for good 1 will be

$$x_1 = \begin{cases} m/p_1 & \text{when } p_1 < p_2; \\ \text{any number between 0 and } m/p_1 & \text{when } p_1 = p_2; \\ 0 & \text{when } p_1 > p_2. \end{cases}$$

Are these results consistent with common sense? All they say is that if two goods are perfect substitutes, then a consumer will purchase the cheaper one. If both goods have the same price, then the consumer doesn't care which one he or she purchases.

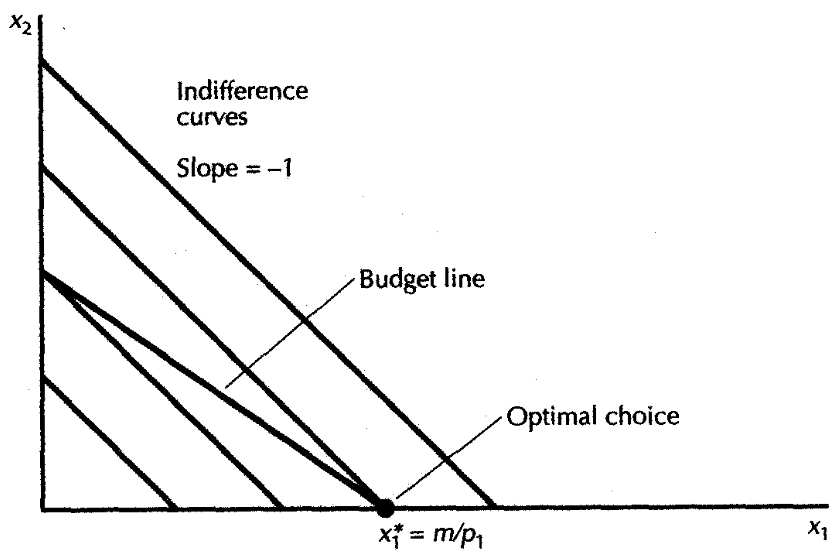

Optimal choice with perfect substitutes. If the goods are perfect substitutes, the optimal choice will usually be on the boundary.

Perfect Complements

The case of perfect complements is illustrated in Figure 5.6. Note that the optimal choice must always lie on the diagonal, where the consumer is purchasing equal amounts of both goods, no matter what the prices are.

In terms of our example, this says that people with two feet buy shoes in pairs.[2]

Let us solve for the optimal choice algebraically. We know that this consumer is purchasing the same amount of good 1 and good 2, no matter what the prices. Let this amount be denoted by x. Then we have to satisfy the budget constraint

$$p_1 x + p_2 x = m.$$

Solving for x gives us the optimal choices of goods 1 and 2:

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

The demand function for the optimal choice here is quite intuitive. Since the two goods are always consumed together, it is just as if the consumer were spending all of her money on a single good that had a price of $p_1 + p_2$.

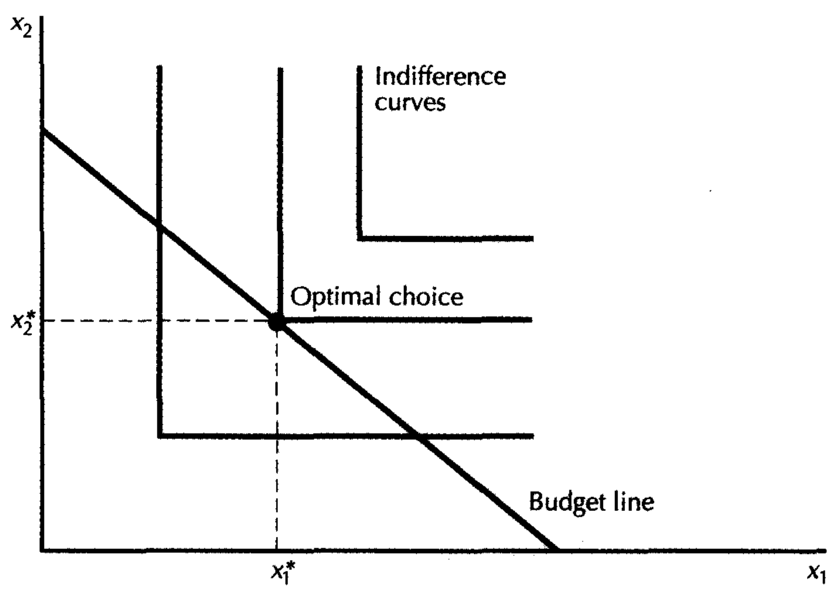

Optimal choice with perfect complements. If the goods are perfect complements, the quantities demanded will always lie on the diagonal since the optimal choice occurs where x_1 equals x_2.

[2] Don't worry, we'll get some more exciting results later on.

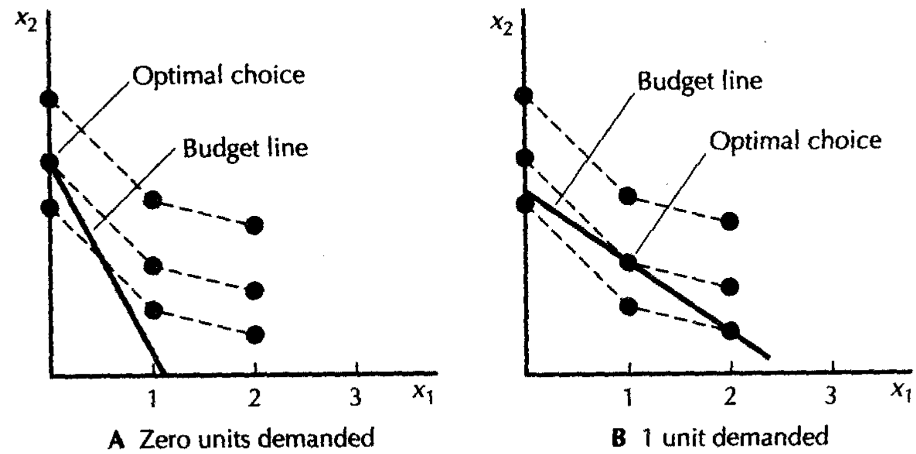

Discrete goods. In panel A the demand for good 1 is zero, while in panel B one unit will be demanded. Figure 5.7

Neutrals and Bads

In the case of a neutral good the consumer spends all of her money on the good she likes and doesn't purchase any of the neutral good. The same thing happens if one commodity is a bad. Thus, if commodity 1 is a good and commodity 2 is a bad, then the demand functions will be

$$x_1 = \frac{m}{p_1}$$
$$x_2 = 0.$$

Discrete Goods

Suppose that good 1 is a discrete good that is available only in integer units, while good 2 is money to be spent on everything else. If the consumer chooses $1, 2, 3, \cdots$ units of good 1, she will implicitly choose the consumption bundles $(1, m - p_1)$, $(2, m - 2p_1)$, $(3, m - 3p_1)$, and so on. We can simply compare the utility of each of these bundles to see which has the highest utility.

Alternatively, we can use the indifference-curve analysis in Figure 5.7. As usual, the optimal bundle is the one on the highest indifference "curve." If the price of good 1 is very high, then the consumer will choose zero units of consumption; as the price decreases the consumer will find it optimal to consume 1 unit of the good. Typically, as the price decreases further the consumer will choose to consume more units of good 1.

Concave Preferences

Consider the situation illustrated in Figure 5.8. Is X the optimal choice? No! The optimal choice for these preferences is always going to be a boundary choice, like bundle Z. Think of what nonconvex preferences mean. If you have money to purchase ice cream and olives, and you don't like to consume them together, you'll spend all of your money on one or the other.

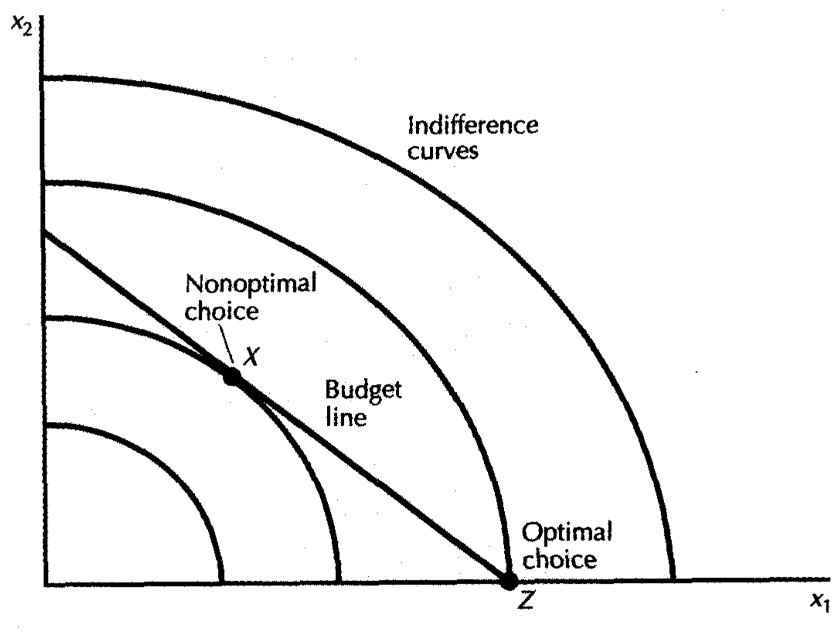

Optimal choice with concave preferences. The optimal choice is the boundary point, Z, not the interior tangency point, X, because Z lies on a higher indifference curve.

Cobb-Douglas Preferences

Suppose that the utility function is of the Cobb-Douglas form, $u(x_1, x_2) = x_1^c x_2^d$. In the Appendix to this chapter we use calculus to derive the optimal choices for this utility function. They turn out to be

$$x_1 = \frac{c}{c+d}\frac{m}{p_1}$$
$$x_2 = \frac{d}{c+d}\frac{m}{p_2}.$$

These demand functions are often useful in algebraic examples, so you should probably memorize them.

The Cobb-Douglas preferences have a convenient property. Consider the fraction of his income that a Cobb-Douglas consumer spends on good 1. If he consumes x_1 units of good 1, this costs him p_1x_1, so this represents a fraction p_1x_1/m of total income. Substituting the demand function for x_1 we have

$$\frac{p_1x_1}{m} = \frac{p_1}{m}\frac{c}{c+d}\frac{m}{p_1} = \frac{c}{c+d}.$$

Similarly the fraction of his income that the consumer spends on good 2 is $d/(c+d)$.

Thus the Cobb-Douglas consumer always spends a fixed fraction of his income on each good. The size of the fraction is determined by the exponent in the Cobb-Douglas function.

This is why it is often convenient to choose a representation of the Cobb-Douglas utility function in which the exponents sum to 1. If $u(x_1, x_2) = x_1^a x_2^{1-a}$, then we can immediately interpret a as the fraction of income spent on good 1. For this reason we will usually write Cobb-Douglas preferences in this form.

5.4 Estimating Utility Functions

We've now seen several different forms for preferences and utility functions and have examined the kinds of demand behavior generated by these preferences. But in real life we usually have to work the other way around: we observe demand behavior, but our problem is to determine what kind of preferences generated the observed behavior.

For example, suppose that we observe a consumer's choices at several different prices and income levels. An example is depicted in Table 5.1. This is a table of the demand for two goods at the different levels of prices and incomes that prevailed in different years. We have also computed the share of income spent on each good in each year using the formulas $s_1 = p_1x_1/m$ and $s_2 = p_2x_2/m$.

For these data, the expenditure shares are relatively constant. There are small variations from observation to observation, but they probably aren't large enough to worry about. The average expenditure share for good 1 is about 1/4, and the average income share for good 2 is about 3/4. It appears

Some data describing consumption behavior.

Year	p_1	p_2	m	x_1	x_2	s_1	s_2	Utility
1	1	1	100	25	75	.25	.75	57.0
2	1	2	100	24	38	.24	.76	33.9
3	2	1	100	13	74	.26	.74	47.9
4	1	2	200	48	76	.24	.76	67.8
5	2	1	200	25	150	.25	.75	95.8
6	1	4	400	100	75	.25	.75	80.6
7	4	1	400	24	304	.24	.76	161.1

that a utility function of the form $u(x_1, x_2) = x_1^{\frac{1}{4}} x_2^{\frac{3}{4}}$ seems to fit these data pretty well. That is, a utility function of this form would generate choice behavior that is pretty close to the observed choice behavior. For convenience we have calculated the utility associated with each observation using this estimated Cobb-Douglas utility function.

As far as we can tell from the observed behavior it appears as though the consumer is maximizing the function $u(x_1, x_2) = x_1^{\frac{1}{4}} x_2^{\frac{3}{4}}$. It may well be that further observations on the consumer's behavior would lead us to reject this hypothesis. But based on the data we have, the fit to the optimizing model is pretty good.

This has very important implications, since we can now use this "fitted" utility function to evaluate the impact of proposed policy changes. Suppose, for example, that the government was contemplating imposing a system of taxes that would result in this consumer facing prices $(2, 3)$ and having an income of 200. According to our estimates, the demanded bundle at these prices would be

$$x_1 = \frac{1}{4}\frac{200}{2} = 25$$
$$x_2 = \frac{3}{4}\frac{200}{3} = 50.$$

The estimated utility of this bundle is

$$u(x_1, x_2) = 25^{\frac{1}{4}} 50^{\frac{3}{4}} \approx 42.$$

This means that the new tax policy would make the consumer better off than he was in year 2, but worse off than he was in year 3. Thus we can use the observed choice behavior to value the implications of proposed policy changes on this consumer.

Since this is such an important idea in economics, let us review the logic one more time. Given some observations on choice behavior, we try to determine what, if anything, is being maximized. Once we have an estimate of what it is that is being maximized, we can use this both to predict choice behavior in new situations and to evaluate proposed changes in the economic environment.

Of course we have described a very simple situation. In reality, we normally don't have detailed data on individual consumption choices. But we often have data on groups of individuals—teenagers, middle-class households, elderly people, and so on. These groups may have different preferences for different goods that are reflected in their patterns of consumption expenditure. We can estimate a utility function that describes their consumption patterns and then use this estimated utility function to forecast demand and evaluate policy proposals.

In the simple example described above, it was apparent that income shares were relatively constant so that the Cobb-Douglas utility function would give us a pretty good fit. In other cases, a more complicated form for the utility function would be appropriate. The calculations may then become messier, and we may need to use a computer for the estimation, but the essential idea of the procedure is the same.

5.5 Implications of the MRS Condition

In the last section we examined the important idea that observation of demand behavior tells us important things about the underlying preferences of the consumers that generated that behavior. Given sufficient observations on consumer choices it will often be possible to estimate the utility function that generated those choices.

But even observing *one* consumer choice at one set of prices will allow us to make some kinds of useful inferences about how consumer utility will change when consumption changes. Let us see how this works.

In well-organized markets, it is typical that everyone faces roughly the same prices for goods. Take, for example, two goods like butter and milk. If everyone faces the same prices for butter and milk, and everyone is optimizing, and everyone is at an interior solution ... then everyone must have the same marginal rate of substitution for butter and milk.

This follows directly from the analysis given above. The market is offering everyone the same rate of exchange for butter and milk, and everyone is adjusting their consumption of the goods until their own "internal" marginal valuation of the two goods equals the market's "external" valuation of the two goods.

Now the interesting thing about this statement is that it is independent of income and tastes. People may value their *total* consumption of the two goods very differently. Some people may be consuming a lot of butter and a little milk, and some may be doing the reverse. Some wealthy people may be consuming a lot of milk and a lot of butter while other people may be consuming just a little of each good. But everyone who is consuming the two goods must have the same marginal rate of substitution. Everyone who is consuming the goods must agree on how much one is worth in terms of the other: how much of one they would be willing to sacrifice to get some more of the other.

The fact that price ratios measure marginal rates of substitution is very important, for it means that we have a way to value possible changes in consumption bundles. Suppose, for example, that the price of milk is \$1 a quart and the price of butter is \$2 a pound. Then the marginal rate of substitution for all people who consume milk and butter must be 2: they have to have 2 quarts of milk to compensate them for giving up 1 pound of butter. Or conversely, they have to have 1 pound of butter to make it worth their while to give up 2 quarts of milk. Hence everyone who is consuming both goods will value a marginal change in consumption in the same way.

Now suppose that an inventor discovers a new way of turning milk into butter: for every 3 quarts of milk poured into this machine, you get out 1 pound of butter, and no other useful byproducts. Question: is there a market for this device? Answer: the venture capitalists won't beat a path to his door, that's for sure. For everyone is already operating at a point where they are just willing to trade 2 quarts of milk for 1 pound of butter; why would they be willing to substitute 3 quarts of milk for 1 pound of butter? The answer is they wouldn't; this invention isn't worth anything.

But what would happen if he got it to run in reverse so he could dump in a pound of butter get out 3 quarts of milk? Is there a market for this device? Answer: yes! The market prices of milk and butter tell us that people are just barely willing to trade one pound of butter for 2 quarts of milk. So getting 3 quarts of milk for a pound of butter is a better deal than is currently being offered in the marketplace. Sign me up for a thousand shares! (And several pounds of butter.)

The market prices show that the first machine is unprofitable: it produces \$2 of butter by using \$3 of milk. The fact that it is unprofitable is just another way of saying that people value the inputs more than the outputs. The second machine produces \$3 worth of milk by using only \$2 worth of butter. This machine is profitable because people value the outputs more than the inputs.

The point is that, since prices measure the rate at which people are just willing to substitute one good for another, they can be used to value policy proposals that involve making changes in consumption. The fact that prices are not arbitrary numbers but reflect how people value things on the margin is one of the most fundamental and important ideas in economics.

If we observe one choice at one set of prices we get the MRS at one consumption point. If the prices change and we observe another choice we get another MRS. As we observe more and more choices we learn more and more about the shape of the underlying preferences that may have generated the observed choice behavior.

5.6 Choosing Taxes

Even the small bit of consumer theory we have discussed so far can be used to derive interesting and important conclusions. Here is a nice example describing a choice between two types of taxes. We saw that a **quantity tax** is a tax on the amount consumed of a good, like a gasoline tax of 15 cents per gallon. An **income tax** is just a tax on income. If the government wants to raise a certain amount of revenue, is it better to raise it via a quantity tax or an income tax? Let's apply what we've learned to answer this question.

First we analyze the imposition of a quantity tax. Suppose that the original budget constraint is

$$p_1x_1 + p_2x_2 = m.$$

What is the budget constraint if we tax the consumption of good 1 at a rate of t? The answer is simple. From the viewpoint of the consumer it is just as if the price of good 1 has increased by an amount t. Thus the new budget constraint is

$$(p_1 + t)x_1 + p_2x_2 = m. \tag{5.1}$$

Therefore a quantity tax on a good increases the price perceived by the consumer. Figure 5.9 gives an example of how that price change might affect demand. At this stage, we don't know for certain whether this tax will increase or decrease the consumption of good 1, although the presumption is that it will decrease it. Whichever is the case, we do know that the optimal choice, (x_1^*, x_2^*), must satisfy the budget constraint

$$(p_1 + t)x_1^* + p_2x_2^* = m. \tag{5.2}$$

The revenue raised by this tax is $R^* = tx_1^*$.

Let's now consider an income tax that raises the same amount of revenue. The form of this budget constraint would be

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

or, substituting for R^*,

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

Where does this budget line go in Figure 5.9?

It is easy to see that it has the same slope as the original budget line, $-p_1/p_2$, but the problem is to determine its location. As it turns out, the budget line with the income tax must pass through the point (x_1^*, x_2^*). The way to check this is to plug (x_1^*, x_2^*) into the income-tax budget constraint and see if it is satisfied.

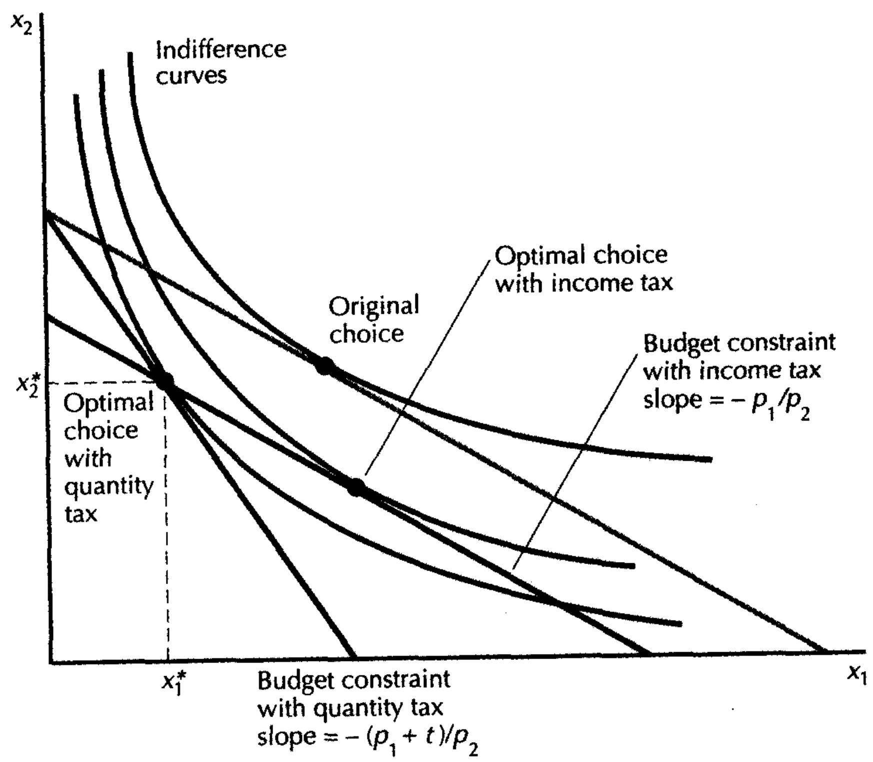

Income tax versus a quantity tax. Here we consider a quantity tax that raises revenue R^* and an income tax that raises the same revenue. The consumer will be better off under the income tax, since he can choose a point on a higher indifference curve.

Is it true that

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

Yes it is, since this is just a rearrangement of equation (5.2), which we know to be true.

This establishes that (x_1^*, x_2^*) lies on the income tax budget line: it is an *affordable* choice for the consumer. But is it an optimal choice? It is easy to see that the answer is no. At (x_1^*, x_2^*) the MRS is $-(p_1 + t)/p_2$. But the income tax allows us to trade at a rate of exchange of $-p_1/p_2$. Thus the budget line cuts the indifference curve at (x_1^*, x_2^*), which implies that there will be some point on the budget line that will be preferred to (x_1^*, x_2^*).

Therefore the income tax is definitely superior to the quantity tax in the sense that you can raise the same amount of revenue from a consumer and still leave him or her better off under the income tax than under the quantity tax.

This is a nice result, and worth remembering, but it is also worthwhile understanding its limitations. First, it only applies to one consumer. The argument shows that for any given consumer there is an income tax that will raise as much money from that consumer as a quantity tax and leave him or her better off. But the amount of that income tax will typically differ from person to person. So a *uniform* income tax for all consumers is not necessarily better than a *uniform* quantity tax for all consumers. (Think about a case where some consumer doesn't consume any of good 1 – this person would certainly prefer the quantity tax to a uniform income tax.)

Second, we have assumed that when we impose the tax on income the consumer's income doesn't change. We have assumed that the income tax is basically a lump sum tax—one that just changes the amount of money a consumer has to spend but doesn't affect any choices he has to make. This is an unlikely assumption. If income is earned by the consumer, we might expect that taxing it will discourage earning income, so that after-tax income might fall by even more than the amount taken by the tax.

Third, we have totally left out the supply response to the tax. We've shown how demand responds to the tax change, but supply will respond too, and a complete analysis would take those changes into account as well.

Summary

1. The optimal choice of the consumer is that bundle in the consumer's budget set that lies on the highest indifference curve.

2. Typically the optimal bundle will be characterized by the condition that the slope of the indifference curve (the MRS) will equal the slope of the budget line.

3. If we observe several consumption choices it may be possible to estimate a utility function that would generate that sort of choice behavior. Such a utility function can be used to predict future choices and to estimate the utility to consumers of new economic policies.

4. If everyone faces the same prices for the two goods, then everyone will have the same marginal rate of substitution, and will thus be willing to trade off the two goods in the same way.

REVIEW QUESTIONS

1. If two goods are perfect substitutes, what is the demand function for good 2?

2. Suppose that indifference curves are described by straight lines with a slope of $-b$. Given arbitrary prices and money income p_1, p_2, and m, what will the consumer's optimal choices look like?

3. Suppose that a consumer always consumes 2 spoons of sugar with each cup of coffee. If the price of sugar is p_1 per spoonful and the price of coffee is p_2 per cup and the consumer has m dollars to spend on coffee and sugar, how much will he or she want to purchase?

4. Suppose that you have highly nonconvex preferences for ice cream and olives, like those given in the text, and that you face prices p_1, p_2 and have m dollars to spend. List the choices for the optimal consumption bundles.

5. If a consumer has a utility function $u(x_1, x_2) = x_1 x_2^4$, what fraction of her income will she spend on good 2?

6. For what kind of preferences will the consumer be just as well-off facing a quantity tax as an income tax?

APPENDIX

It is very useful to be able to solve the preference-maximization problem and get algebraic examples of actual demand functions. We did this in the body of the text for easy cases like perfect substitutes and perfect complements, and in this Appendix we'll see how to do it in more general cases.

First, we will generally want to represent the consumer's preferences by a utility function, $u(x_1, x_2)$. We've seen in Chapter 4 that this is not a very restrictive assumption; most well-behaved preferences can be described by a utility function.

The first thing to observe is that we already *know* how to solve the optimal-choice problem. We just have to put together the facts that we learned in the last three chapters. We know from this chapter that an optimal choice (x_1, x_2) must satisfy the condition

$$\mathrm{MRS}(x_1, x_2) = -\frac{p_1}{p_2}, \tag{5.3}$$

and we saw in the Appendix to Chapter 4 that the MRS can be expressed as the negative of the ratio of derivatives of the utility function. Making this substitution and cancelling the minus signs, we have

$$\frac{\partial u(x_1, x_2)/\partial x_1}{\partial u(x_1, x_2)/\partial x_2} = \frac{p_1}{p_2}. \tag{5.4}$$

From Chapter 2 we know that the optimal choice must also satisfy the budget constraint

$$p_1 x_1 + p_2 x_2 = m. \tag{5.5}$$

This gives us two equations—the MRS condition and the budget constraint—and two unknowns, x_1 and x_2. All we have to do is to solve these two equations to find the optimal choices of x_1 and x_2 as a function of the prices and income. There are a number of ways to solve two equations in two unknowns. One way that always works, although it might not always be the simplest, is to solve the budget constraint for one of the choices, and then substitute that into the MRS condition.

Rewriting the budget constraint, we have

$$x_2 = \frac{m}{p_2} - \frac{p_1}{p_2}x_1 \tag{5.6}$$

and substituting this into equation (5.4) we get

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

This rather formidable looking expression has only one unknown variable, x_1, and it can typically be solved for x_1 in terms of (p_1, p_2, m). Then the budget constraint yields the solution for x_2 as a function of prices and income.

We can also derive the solution to the utility maximization problem in a more systematic way, using calculus conditions for maximization. To do this, we first pose the utility maximization problem as a constrained maximization problem:

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

$$\text{such that } p_1x_1 + p_2x_2 = m.$$

This problem asks that we choose values of x_1 and x_2 that do two things: first, they have to satisfy the constraint, and second, they give a larger value for $u(x_1, x_2)$ than any other values of x_1 and x_2 that satisfy the constraint.

There are two useful ways to solve this kind of problem. The first way is simply to solve the constraint for one of the variables in terms of the other and then substitute it into the objective function.

For example, for any given value of x_1, the amount of x_2 that we need to satisfy the budget constraint is given by the linear function

$$x_2(x_1) = \frac{m}{p_2} - \frac{p_1}{p_2}x_1. \tag{5.7}$$

Now substitute $x_2(x_1)$ for x_2 in the utility function to get the *unconstrained* maximization problem

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

This is an unconstrained maximization problem in x_1 alone, since we have used the function $x_2(x_1)$ to ensure that the value of x_2 will always satisfy the budget constraint, whatever the value of x_1 is.

We can solve this kind of problem just by differentiating with respect to x_1 and setting the result equal to zero in the usual way. This procedure will give us a first-order condition of the form

$$\frac{\partial u(x_1, x_2(x_1))}{\partial x_1} + \frac{\partial u(x_1, x_2(x_1))}{\partial x_2}\frac{dx_2}{dx_1} = 0. \tag{5.8}$$

Here the first term is the direct effect of how increasing x_1 increases utility. The second term consists of two parts: the rate of increase of utility as x_2 increases, $\partial u/\partial x_2$, times dx_2/dx_1, the rate of increase of x_2 as x_1 increases in order to continue to satisfy the budget equation. We can differentiate (5.7) to calculate this latter derivative

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

Substituting this into (5.8) gives us

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

which just says that the marginal rate of substitution between x_1 and x_2 must equal the price ratio at the optimal choice (x_1^*, x_2^*). This is exactly the condition we derived above: the slope of the indifference curve must equal the slope of the budget line. Of course the optimal choice must also satisfy the budget constraint $p_1x_1^* + p_2x_2^* = m$, which again gives us two equations in two unknowns.

The second way that these problems can be solved is through the use of **Lagrange multipliers**. This method starts by defining an auxiliary function known as the *Lagrangian:*

$$L = u(x_1, x_2) - \lambda(p_1x_1 + p_2x_2 - m).$$

The new variable λ is called a **Lagrange multiplier** since it is multiplied by the constraint.[3] Then Lagrange's theorem says that an optimal choice (x_1^*, x_2^*) must satisfy the three first-order conditions

$$\begin{aligned}
\frac{\partial L}{\partial x_1} &= \frac{\partial u(x_1^*, x_2^*)}{\partial x_1} - \lambda p_1 = 0 \\
\frac{\partial L}{\partial x_2} &= \frac{\partial u(x_1^*, x_2^*)}{\partial x_2} - \lambda p_2 = 0 \\
\frac{\partial L}{\partial \lambda} &= p_1x_1^* + p_2x_2^* - m = 0.
\end{aligned}$$

There are several interesting things about these three equations. First, note that they are simply the derivatives of the Lagrangian with respect to x_1, x_2, and λ, each set equal to zero. The last derivative, with respect to λ, is just the budget constraint. Second, we now have three equations for the three unknowns, x_1, x_2, and λ. We have a hope of solving for x_1 and x_2 in terms of p_1, p_2, and m.

Lagrange's theorem is proved in any advanced calculus book. It is used quite extensively in advanced economics courses, but for our purposes we only need to know the statement of the theorem and how to use it.

In our particular case, it is worthwhile noting that if we divide the first condition by the second one, we get

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

which simply says the MRS must equal the price ratio, just as before. The budget constraint gives us the other equation, so we are back to two equations in two unknowns.

[3] The Greek letter λ is pronounced "lamb-da."

EXAMPLE: Cobb-Douglas Demand Functions

In Chapter 4 we introduced the **Cobb-Douglas utility function**

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

Since utility functions are only defined up to a monotonic transformation, it is convenient to take logs of this expression and work with

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

Let's find the demand functions for x_1 and x_2 for the Cobb-Douglas utility function. The problem we want to solve is

$$\max_{x_1, x_2} c \ln x_1 + d \ln x_2$$

$$\text{such that } p_1 x_1 + p_2 x_2 = m.$$

There are at least three ways to solve this problem. One way is just to write down the MRS condition and the budget constraint. Using the expression for the MRS derived in Chapter 4, we have

$$\frac{cx_2}{dx_1} = \frac{p_1}{p_2}$$
$$p_1 x_1 + p_2 x_2 = m.$$

These are two equations in two unknowns that can be solved for the optimal choice of x_1 and x_2. One way to solve them is to substitute the second into the first to get

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

Cross multiplying gives

$$c(m - x_1 p_1) = dp_1 x_1.$$

Rearranging this equation gives

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

or

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

This is the demand function for x_1. To find the demand function for x_2, substitute into the budget constraint to get

$$\begin{aligned} x_2 &= \frac{m}{p_2} - \frac{p_1}{p_2} \frac{c}{c+d} \frac{m}{p_1} \\ &= \frac{d}{c+d} \frac{m}{p_2}. \end{aligned}$$

The second way is to substitute the budget constraint into the maximization problem at the beginning. If we do this, our problem becomes

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

The first-order condition for this problem is

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

A little algebra—which you should do!—gives us the solution

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

Substitute this back into the budget constraint $x_2 = m/p_2 - x_1 p_1/p_2$ to get

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

These are the demand functions for the two goods, which, happily, are the same as those derived earlier by the other method.

Now for Lagrange's method. Set up the Lagrangian

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

and differentiate to get the three first-order conditions

$$\begin{aligned} \frac{\partial L}{\partial x_1} &= \frac{c}{x_1} - \lambda p_1 = 0 \\ \frac{\partial L}{\partial x_2} &= \frac{d}{x_2} - \lambda p_2 = 0 \\ \frac{\partial L}{\partial \lambda} &= p_1 x_1 + p_2 x_2 - m = 0. \end{aligned}$$

Now the trick is to solve them! The best way to proceed is to first solve for λ and then for x_1 and x_2. So we rearrange and cross multiply the first two equations to get

$$\begin{aligned} c &= \lambda p_1 x_1 \\ d &= \lambda p_2 x_2. \end{aligned}$$

These equations are just asking to be added together:

$$c + d = \lambda(p_1 x_1 + p_2 x_2) = \lambda m,$$

which gives us

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

Substitute this back into the first two equations and solve for x_1 and x_2 to get

$$\begin{aligned} x_1 &= \frac{c}{c+d} \frac{m}{p_1} \\ x_2 &= \frac{d}{c+d} \frac{m}{p_2}, \end{aligned}$$

just as before.