Intermediate Microeconomics: Demand and Market Demand

Simon Naitram, PhD

October 23, 2020

Chapter 6: Demand

Read Sections 6.0 to 6.5 and 6.8 carefully. Get a quick understanding of what's in Sections 6.6 and 6.7, but don't spend too much time on it. Don't waste too much time on perfect substitutes, perfect complements, and discrete goods. Make sure you work through all the math I've done here. **Don't take it for granted that you know what I did!**

Now that we've figured out what the equilibrium *endogenous* variables are, let's see how they change when we alter the *exogenous* variables. Remember the exogenous variables in this model of consumer theory are prices (p_1, p_2) and income m.

Comparative Statics Comparative statics is another useful tool the economist has. It's a thought experiment: "what would the new world look like if I change one thing and leave everything else constant?" Comparative statics focuses on how people's optimal decisions change when an exogenous variable is changed. Most importantly, it ignores *how* we get from point A to point B. It just compares an old point A to a new point B. Is the new equilibrium higher? Is it lower?

Income Change

Mathematically, how would we investigate the effect of a change in income on the choice of x_1? Well, yes, we have a new function $x_1^*(p_1, p_2, m)$, so we can obviously take the partial derivative of this function with respect to m. That's the correct calculus way to do it! Now you're beginning to think like an analytical economist.

But let's also think through the process of what happens. An increase in an individual's income, holding all prices constant, moves the budget line outwards in a parallel fashion. It means that the consumer's set of choices has expanded, and they're not able to achieve a higher level of utility. This implies that they'll move onto a new utility curve, and relocate at a new optimal point. Relative to the original optimal point, the new optimal point is *higher*. This means that an increase in income increases utility.

Under well-behaved preferences, a higher utility curve means that the consumer is now consuming more of both good 1 and 2. This means that:

$$\frac{\partial x_1^*(p_1, p_2, m)}{\partial m} > 0 \tag{1}$$

$$\frac{\partial x_2^*(p_1, p_2, m)}{\partial m} > 0. \tag{2}$$

That is, a small increase in income increases the optimal choices of both x_1 and x_2. This holds under normal circumstances, and so we would call good 1 and good 2 **normal goods**.

Inferior Goods What would happen if the optimal choice of either of these goods *decreased* when we increased income? It would mean we were dealing with an inferior good. An inferior good is one that you want to consume *less of* as you become richer. What is defined as an inferior good might be different at different levels of income.[1]

[1] There may be some really rich people who think that caviar is beneath them!

Examples Can you think of any examples of normal and inferior commodities? Tell me why. There's one important and large class of services that can be considered to be inferior. Guess what it is!

Income Offer Curve What if we kept increasing a person's income by a dollar, and seeing what their new optimal choice is each time? We get a bunch of points that are basically walking outwards (away from the origin) if both goods are normal. We can join these points with a nice line. This line is what we call the *income offer curve* or the *income expansion path*. The income expansion path is then positively sloped for two normal goods.

Engel Curve We can plot what happens to the demand for one of the goods. Change income by small increments and see how $x_1^*(p_1, p_2, m)$ changes as we change m. We can plot this on a new graph. Since we're dealing with changes in m and how it affects x_1, we plot the consumer's optimal choice of x_1^* for every value of m.[2]

[2] Note that economists usually plot this chart backwards, with the independent variable (m) on the y-axis and the dependent variable x_1 on the y-axis. I explain a bit more later.

What does the Engel curve look like? As always, it depends on preferences. If it's a normal good, it'll be upward sloping. Now let's remember that the demand for good 1 using a Cobb-Douglas utility function gives us a demand for x_1 as:

$$x_1^* = \left(\frac{\alpha}{\alpha + \beta}\right) \frac{m}{p_1} \tag{3}$$

Let's make this a little clearer. Can we express x_1 as a function of m? Well yes:[3]

[3] We've used $\alpha + \beta = 1$ to make this a little clearer.

$$x_1^* = \left(\frac{\alpha}{p_1}\right) \cdot m, \tag{4}$$

it already is! And once again, holding prices constant, x_1^* is a linear function of m. That means that the slope of the line is simple α/p_1. The Engel curve for good 1 with Cobb-Douglas preferences is simply a straight line with a slope based on the share of a consumer's income they spend on good 1 divided by the price of good 1. As the price of good 1 goes up, the line gets a bit flatter.

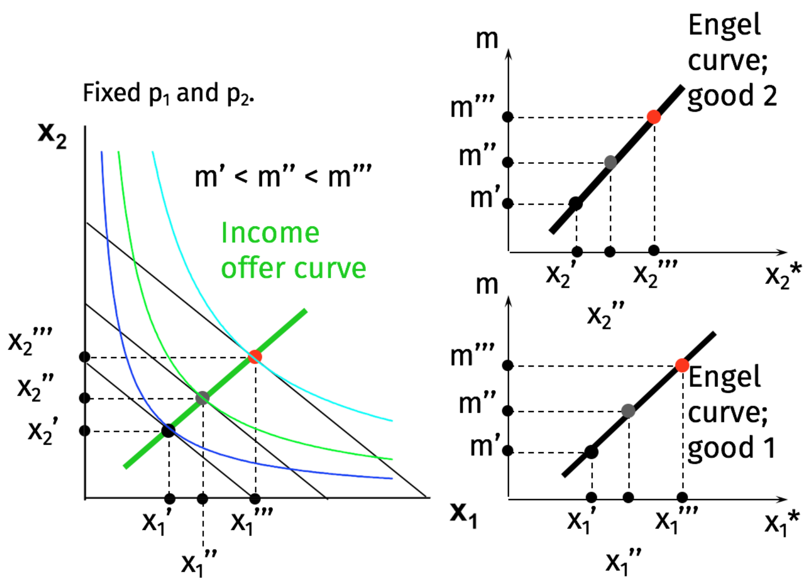

Figure 1: Mapping Income Changes to Engel Curve

Homothetic Preferences The implications of this straight line is that regardless of an individual's income, they *always* spend the same **proportion** of their income on good 1. Is that natural? Unlikely! The share of your income you spend on a single commodity typically increases or decreases as you become richer. For example, you don't buy more toilet paper as you get richer. Why would you? You may, however, begin to purchase fancy cheeses like Grano Padano or Emmental. In a case like that, the proportion of your income you spend on Cheddar cheese likely decreases as your income increases. So no, it's not normal that we might expect a person's consumption pattern to remain exactly the same as their income increases. We call this **homothetic preferences** where a person's consumption *pattern* remains the same regardless of income. What it means mathematically is that we can double a person's income, and they'd automatically exactly double their consumption of both x_1 and x_2.

When does it hold? Two goods it *might* hold for are housing and vehicles. People tend to spend a fixed portion of their income on these, buying more expensive cars or houses as their income increases. This is one of the main reasons you often hear people far richer than yourself complaining that things are tight.[4] This might also hold for small increases in your income. If I give you an extra dollar, it's unlikely to make you increase your demand for villas in Port Ferdinand.

[4] Rich people really are ridiculous sometimes: CNBC on why you need $350,000 a year to be middle class.

Cobb-Douglas Even though the Cobb-Douglas utility function is a bit weird in the *qualitative predictions* it makes about a person's consumption patterns, we use it really often. Importantly, it's a special case of one of the more useful utility functions: the Constant Elasticity of Substitution utility function. We won't use it in this course, but at least know that we can generalise the Cobb-Douglas function to a more realistic utility function very easily.

Example 2 Let's use the quasi-linear utility function. Let's set the quasi-linear utility function to $u(x_1, x_2) = \ln(x_1) + x_2$. Note that the reason we say it's quasi-linear is because one term enters the function linearly (x_2) and one term enters non-linearly ($\ln(x_1)$). The problem can be set up as

$$\begin{aligned} \max \quad & u(x_1, x_2) = \ln(x_1) + x_2 \\ \text{s.t.} \quad & p_1 x_1 + p_2 x_2 = m \end{aligned} \tag{5}$$

The constraint is first re-written as $p_1 x_1 + p_2 x_2 - m = 0$. We then form the Lagrangian:

$$\max_{x_1, x_2} \quad \mathcal{L} = \ln(x_1) + x_2 - \lambda(p_1 x_1 + p_2 x_2 - m) \tag{6}$$

The first order conditions of this problem are:

$$\frac{\partial \mathcal{L}}{\partial x_1} = \frac{1}{x_1} - \lambda p_1 = 0 \tag{7}$$

$$\frac{\partial \mathcal{L}}{\partial x_1} = 1 - \lambda p_2 = 0 \tag{8}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = p_1 x_1 + p_2 x_2 - m = 0 \tag{9}$$

We can solve these by setting $\lambda = 1/p_2$, and substituting into equation (12). This gives:

$$\begin{aligned} \frac{1}{x_1} - \frac{p_1}{p_2} &= 0 \\ \frac{1}{x_1} &= \frac{p_1}{p_2} \\ \frac{p_2}{p_1} &= x_1^* \end{aligned} \tag{10}$$

Substituting into the third FONC, we get:

$$\begin{aligned} p_1 \frac{p_2}{p_1} + p_2 x_2 - m &= 0 \\ p_2 + p_2 x_2 &= m \\ p_2 x_2 &= m - p_2 \\ x_2 &= \frac{m - p_2}{p_2} \\ x_2^* &= \frac{m}{p_2} - 1 \end{aligned} \quad (11)$$

This gives an interesting pair of demand functions. We see that x_1^* is purely a function of prices. It doesn't make a difference what the consumer's income is, they'll always demand the same amount of good 1. The demand for good 1 is constant. It means that whatever extra income they get, it is spent on good 2.

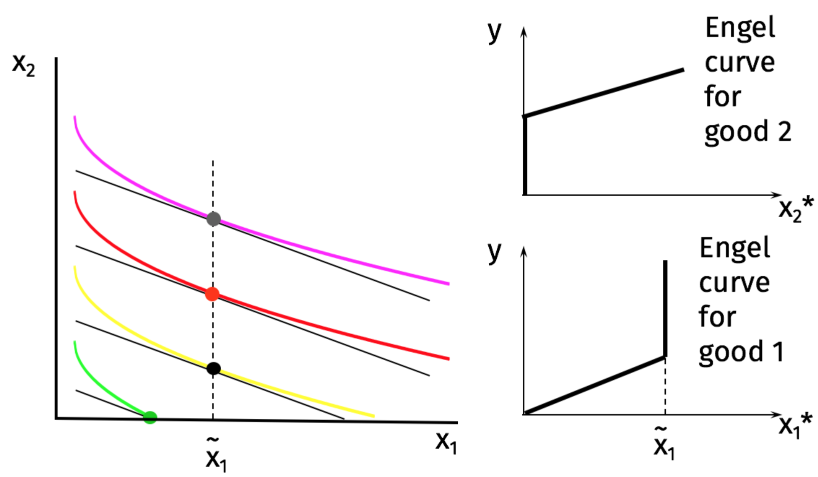

Figure 2: Engel Curve with Quasilinear Preferences

However, if a person's income is below the money needed to buy the amount of good 1 they want, then an increase in income means they'll spend all the additional money on good 1. So the consumer is single-minded in their goal of achieving a specific quantity of good 1. At first their Engel curve is sloping upwards since they spend all their money on good 1. Once they have enough money to buy how much of good 1 is optimal for them, they begin spending their money on good 2.

Now if good 2 is simply 'money' or 'all other goods', then it makes sense. You can think of good 1 as toilet paper. If you don't have enough money to buy enough toilet paper, then if I give you an extra $10, you'll probably spend all of it on toilet paper. If I give you another $100 after that, you probably won't buy any more toilet paper! You'll use that money on 'all other things'.

Qualitative v Quantitative Predictions Qualitative predictions are the more general predictions we make about the effect of one thing on another. It's less about giving a precise number but more about describing in simple terms what you would expect to happen. Being able to make qualitative predictions is important for explaining economics to non-economists, and helpful for thinking through problems. For example, a qualitative prediction would be the answer to: *does a higher price mean an increase or decrease in demand for a good?* A quantitative prediction would be the answer to: *by how much does a price increase of $1 increase or decrease the demand for good 1?* You'll likely get questions asking for both types of predictions on your exams.

Price Changes

Now let's do out comparative statics using price changes. So let's first consider what happens if we change *one* price. So we consider a small decrease in the price of good 1, holding the price of good 2 constant. A change in one price implies a change in the slope of the budget line. A reduction in the price of good 1 will cause the budget line to *pivot* to the right.[5]

Different MRS This means that the market is now willing to substitute good 1 for good 2 at a different rate. Specifically, in order to get an additional unit of good 1, you have to give up less of good 2. It sounds like it means good 1 is less valuable to the market now than it was before. In fact, that's exactly what a price decrease implies.

If the slope of the budget line changes[6], then at the point of optimality the slope of the indifference curve must change as well[7]. Consider well-behaved preferences that give a non-linear utility function and a non-linear indifference curve. Not only do we now move to a higher utility function, but we also move to a different marginal rate of substitution.

Graphically As the budget line pivots to the right, we will get new optimal points of tangency between the indifference curve and the budget line. We draw a new indifference curve parallel to the first, and find the new point where it is tangent to the new budget line. Let's compare it to the old ordinary demand point. Under well-behaved preferences, we will see the consumer demanding more of good 1. Therefore, the key qualitative prediction is that with well-behaved preferences, *a decrease in the price of good 1 leads to an increase in the demand for good 1.* That is, when the budget line pivots to the right, the consumer can now afford better bundles than they were

[5] I'm sure you know what this pivot looks like. Mathematically, we can describe it as the vertical intercept remaining the same and the slope of the line changing (using $= mx + c$). This means the horizontal intercept changes. For a change in the price of good 2, it'd be the opposite.

[6] The rate at which the *market* is willing to substitute good 1 for good 2.

[7] The rate at which the *consumer* is willing to substitute good 1 for good 2.

able to afford before. You can think of this as an increase in their purchasing power.

Ordinary v Giffen If a reduction in the price of good 1 leads us to consume *more* of good 1, then we consider good 1 an **ordinary good**. If a reduction in the price of good 1 leads us to consume *less* of good 1, then we consider good 1 a **Giffen good**. Giffen goods are named after a Scottish economist called Sir Robert Giffen. The original example of a Giffen good was bread. Bread is considered a necessity, and an increase in the price of bread makes it difficult for a family to buy other foodstuffs, like meat. So rather than buying less bread, they double down on their purchases of bread in order to satisfy their nutritional needs when the price of bread goes up. That's weird to think about. I'm not even sure if it entirely makes sense. It's difficult to consider any rational reason why a person would consume less of a good as it becomes cheaper.[8]

[8] Unless, of course, you're simply buying something for the highest price possible to impress others. Then it is a Veblen good: https://en.wikipedia.org/wiki/Veblen_good.

Demand Curve We can, once again, draw a line through these optimal points. This new line is our **price offer curve**. This isn't very interesting, however. What we can do, instead, is to plot a new chart with all of these new points on it. We want to know what the demand for good 1 is at different prices of good 1. On the y-axis we'll put the price of good 1 p_1. On the x-axis we'll put the quantity demanded of good 1. **Note that this chart is backwards!** To figure out every point, we can simply use our demand function derived from the optimal choice problem.

For the Cobb-Douglas utility function, we have a demand function for x_1 of the form:

$$x_1^* = \left(\frac{\alpha}{\alpha+\beta}\right)\frac{m}{p_1}$$
$$x_1^* = \frac{\alpha m}{p_1} \tag{12}$$

We can simply plug in different values of p_1 and see what the optimal demand of x_1 is. This will give us a chart of the demand for x_1 at different prices p_1. Remember that we plot them backwards, so its (p_1, x_1) rather than (x_1, p_1). This backwards line is now our **demand curve**. The function above tells us that as we increase p_1, the value of x_1 is going to fall. That is, there is a negative relationship between x_1 and p_1. *The demand curve is almost always negatively sloped.*

Inverse Demand Function Because we've plotted our graph backwards, with the independent variable on the y-axis instead of the x-axis, we can invert our demand function too. By inverting our demand function, it means that we move from a simple function of the

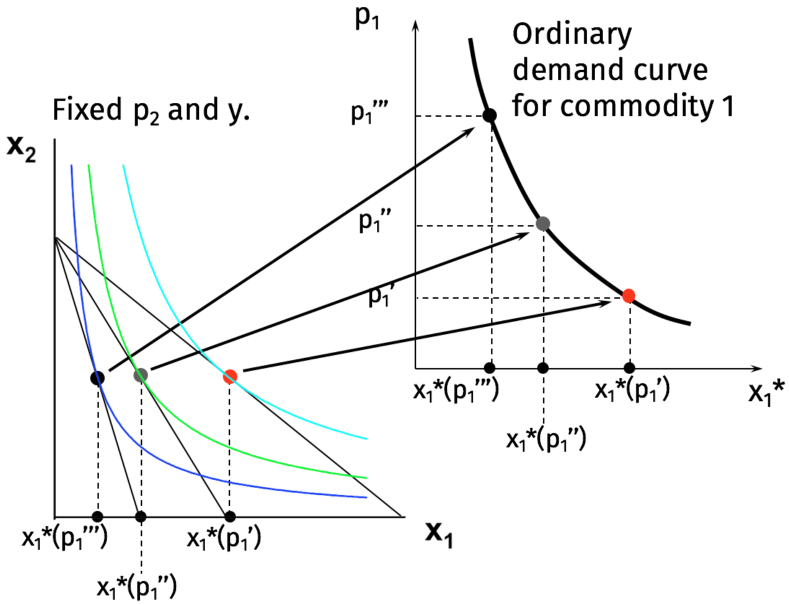

Figure 3: Mapping Price Changes to Demand Curve

form $y = f(x)$ to an inverse function $x = f^{-1}(y)$. Weird right? But since we've flipped our graph, we could as well flip our demand function too. To do this then, we "simply" solve our demand function for p_1 as a function of x_1. For the Cobb-Douglas utility function, we'd get:

$$x_1^* = \frac{\alpha m}{p_1} \tag{13}$$

$$p_1 = \frac{\alpha m}{x_1^*} \tag{14}$$

Once again, we see an inverse relationship between p_1 and x_1. But remember, the consumer can't change p_1, so this isn't a behavioural function, just a rearrangement.

What it allows us to do, is to go back to our original conditions for optimality. We had that MRS $= p_1/p_2$. If we solve this for p_1 to match above, we get $p_1 = p_2$MRS. Now let's suppose that good 2 is the numeraire, and it represents 'money' or 'all other goods' or 'money to spend on all other goods'. This means that $p_2 = 1$ and we get that $p_1 =$ MRS. That says that p_1 is exactly equal to the marginal rate of substitution between good 1 and money.

The marginal rate of substitution has a nice easy interpretation now: *how much money are you willing to give up to get a little more of good 1*? Or conversely, *how much money do I need to give you in order for you to give up a little bit of good 1*? This means we can now think of the MRS as the marginal willingness to pay!

Thought Experiment We've already said that the consumer *cannot set the price*. However, if we imagined a world for a moment where they *could* influence the price. That is, a world where the inverse demand function is actually a behavioural equation. Therefore, for each chosen value of x_1, the consumer faces a new price which we'll call $\hat{p}_1$[9]. We established above that this price the consumer faces is the consumer's marginal willingness to pay (WTP) for a bit more of good 1. This price is equal to the consumer's inverse demand function. We can then use this inverse demand function to figure out what the consumer's marginal WTP is for different values of x_1! As you'd expect, the marginal willingness to pay is *higher* when you have less of a good. The marginal willingness to pay decreases as you get more of the good.

[9] We call it this to differentiate it from the actual market price. This is just a counterfactual price in the consumer's mind.

Mathematically, this is:

$$\hat{p}_1 = \frac{\alpha m}{x_1^*}$$
$$\hat{p}_1 = d\text{WTP} = \frac{\alpha m}{x_1^*} \tag{15}$$

It means that if we can figure out a person's inverse demand function based on market prices, we can figure out how much they'd be willing to pay for various amounts of the good. It effectively gives us a way to put a dollar value on how consumers value commodities. Pretty tricky, but pretty cool.

Who Reversed My Axes?! You've got accustomed to the mathematically correct way of drawing charts. If you have a function $y = f(x)$, then you draw the chart with y on the y-axis, and x on the x-axis. It's not particularly rocket science. So why do we draw our demand curves with p, the independent variable, on the x-axis? Alfred Marshall wrote the most important economics textbook of his time, "Principles of Economics", in 1890. He was the possibly the person who first drew the demand curve with the price on the y-axis instead of the x-axis, but certainly the person who made it popular. There are competing explanations as to why he did so. Some think that he viewed price as the dependent variable. This view would mean that *price adjusts to demand*, rather than the other way around. Of course we know that individual demand must react to price and not the other way around. But does market demand react to price? Or does the price react to market demand? What do you think? Here are some interesting links on the topic:

- https://historyofeconomics.wordpress.com/tag/reversed-axes/
- http://gregmankiw.blogspot.com/2006/09/who-invented-supply-and-demand.html

Undoubtedly, it's simply a matter of how you want to think about it!

Chapter 15: Market Demand

READ SECTIONS 15.1, 15.5, 15.6, AND 15.11 carefully. Less carefully, ensure that you read 15.2, 15.4, and 15.8.

Market The market can refer to a group of consumers. In this case, it means that we want to take what we know about the individual consumer and aggregate it. Aggregating means that we add everyone's behaviour together to see how the entire group behaves. Specifically, we want to aggregate everyone's demand functions to see how the entire group of consumers behave.

Individual Demand Function We start from the standard demand function $x^1(p_1, p_2, m)$. Now we'll make two adjustments. First, instead of putting the 1 for good 1 as a subscript, we put it as a superscript. Then we say that there are n individuals in the market. We consider each individual i, for $i \in \{1, n\}$. An individual i's demand function is given by:

$$x_i^1(p_1, p_2, m_i). \tag{16}$$

So notice we now have a superscript for the good and a subscript for the individual. Note that each individual can have a different demand function since they have different utility functions! In addition, income also has an individual subscript on it. Why? Because each individual can have a different income too.

Aggregate Demand Aggregate demand for good 1 is the sum of the demand functions for each individual, from 1 to n.

$$X^1 = \sum_{i=1}^{n} x_i^1(p_1, p_2, m_i). \tag{17}$$

Note that X^1 is effectively a function of prices and of *everyone's* income. So the aggregate demand function is $X^1(p_1, p_2, m_1, m_2, ..., m_n)$. We say that aggregate demand is a function of prices and the *distribution* of incomes.

Market Demand Curve Let's combine everyone's demand function for varying price p_1. It's a simple matter of adding demand functions. Consider price $p_1 = \bar{p}_1$. Paula's demand function is:

$$x_P = \frac{0.5m_P}{\bar{p}_1} \tag{18}$$

while Rachel's demand function is

$$x_R = \frac{0.3m_R}{\bar{p}_1} \tag{19}$$

The sum of their demand functions at price $\bar{p}_1$ is $(0.5m_P)/\bar{p}_1 + (0.3m_R)/\bar{p}_1$. Market demand is therefore:

$$X^1 = \frac{0.5m_P + 0.3m_R}{\bar{p}_1}. \tag{20}$$

If we know everyone's demand function and their incomes, we can construct a market demand curve in this way. This is the thorough and direct way to construct an aggregate demand function from microfoundations. The demand function

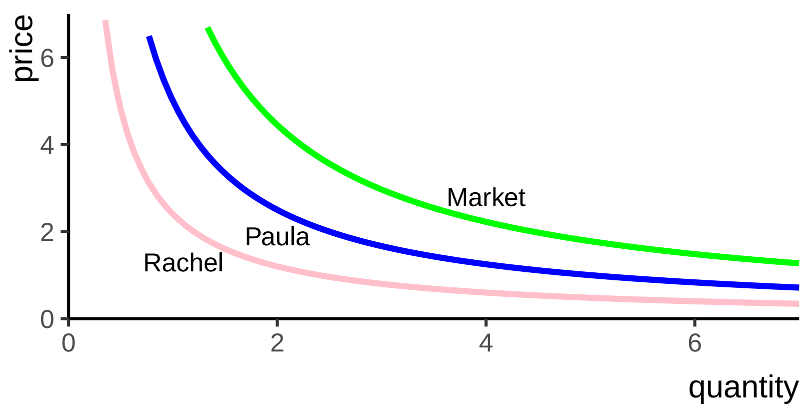

Figure 4: Note that I've assumed some incomes for Rachel and Paula to plot this graph.

Representative Consumer What if, instead of a distribution of incomes, we made demand a function of aggregate income? To make the aggregate demand function simpler, we convert individual incomes to an aggregate income: $M = \sum_{i=1}^{n} m_i$. We would get an aggregate demand function:

$$X^1(p_1, p_2, M). \tag{21}$$

Now the aggregate economy has a demand function that looks a lot like the individual consumer's. This is cool and useful since it means we can analyse the aggregate economy in much the same way we analyse the consumer. This is the *representative consumer* assumption. Now when we talk about an increase in income, we mean an increase in aggregate income.[10]

[10] Just for clarity, can you think about what the representative consumer assumption implies when you've got serious inequality in the country?

Price Elasticity

We've been able to describe the effect of changes in prices and income on demand. We discussed these intuitively, saying that an increase in prices leads to a reduction in demand and an increase in income likely leads to an increase in demand. However, the more mathematically correct way to do it is obviously to take the derivative of the demand function with respect to either price or income. That would give us dx_1/dp_1.

At first glance, it looks like this derivative tells us the responsiveness of demand to a change in price. The problem is that the derivative is entirely in the units of the good. For example, what if we measures quantity as *number* of saltbreads rather than *packs* of saltbreads. Suddently, your derivative would increase by six.[11] This means that the derivative is kinda useless, since it depends entirely how we measure quantity.

[11] Assuming, of course, that you did get six saltbreads in your pack.

How do we get away from this? Yes, by percentage changes! Percentage changes completely ignore the units that we're using. For example, increasing from 4 packs of saltbreads to 5 is exactly the same as increasing from 24 saltbreads to 30 when you measure it in percentage change.

How do we convert a derivative to a percentage change? It turns out to be really easy. Remember that dp_1 is a change in p_1. The percentage change in p_1 is therefore simply dp_1/p_1. Similarly, the percentage change in x is dx_1/x_1. We can then simply divide these percentage changes by one another to get a ratio of the two percentage changes.

$$\frac{\%\text{ change in } x_1}{\%\text{ change in } p_1} = \frac{dx_1/x_1}{dp_1/p_1} \tag{22}$$

This is an elasticity! Since we know how to work out dx_1/dp_1 directly from the demand function, let's separate this equation so it's a) the derivative, plus b) something else. Here's what I mean:

$$\epsilon = \frac{dx_1}{dp_1} \times \frac{p_1}{x_1}. \tag{23}$$

This is your magic formula for the elasticity, which we'll call ϵ. It is the derivative of the demand function with respect to price times the price divided by the quantity. We call this the **price elasticity of demand**, or the **elasticity of demand with respect to price**. It tells us the percentage change in demand when price increases by 1 percent.

Characteristics The price elasticity of demand is usually negative. We already know this because we know an increase in the price usually leads to a decline in demand. Most of the time, when an economist says the elasticity is smaller, they mean *closer to zero* rather than *more negative*. So an elasticity that is more negative[12] would lead to a larger decline in demand when price increases by 1%. We typically say that this more negative elasticity is *more elastic* than one that is closer to zero.

[12] -6 is more negative than -3.

Straight Line We learnt that the slope of a straight line is the same everywhere on the line. So if the demand curve is a straight line, is the elasticity the same everywhere? Well the answer is no. The slope is the same everywhere, so it means the derivative is the same everywhere. However, the second part of the equation changes depending on what part of the line you're on. Remember there's a negative relationship between price and demand. Let's consider the derivative or slope to be equal to a constant $c = dx_1/dp_1$. There are five cases to consider for the second part of the elasticity:[13]

[13] Check Page 276 of the Varian for a graphical representation.

1. When price is zero, then demand is very high so we have $p_0/x_{\text{VH}} = 0$.

2. When price is low, then demand is high so we have a small elasticity $p_L/x_{\text{H}} < 1$.

3. When ratio of the price to quantity demanded is exactly equal to the inverse of the slope (or derivative) c, then the elasticity is equal to -1.

4. When price is high, then demand is low so we have a large elasticity $p_H/x_{\text{L}} > 1$.

5. When the price is exorbitantly high, then demand is zero so we have $p_{\text{VH}}/x_0 = \infty$.

Remember that the derivative will usually be negative, so the price elasticity of demand will usually be negative.

Elasticity Chat So how do we talk about elasticities? Here's the terminology you'll frequently hear me use.

Elastic: When the elasticity is $\epsilon < -1$, or equivalently that the absolute value is greater than 1 $|\epsilon| > 1)$, we say that demand is elastic. This means that a small change in the price will generate a large demand response. Can you think of any goods like this?

Inelastic: When the elasticity is between -1 and 0, or the absolute value is between 0 and 1, we say that demand is inelastic. This means a small change in the price doesn't generate much of a change in quantity demanded. What goods have inelastic demand?

Unit Elastic: If the elasticity is exactly -1 or the absolute value is exactly 1, we say demand is unit elastic. This means that an increase in price will generate a decrease in demand of the exact same percentage value. Are there any goods that might behave like this?

Perfectly Inelastic: When the elasticity is zero, we say that demand is perfectly inelastic. This means that a change in the price generates absolutely no change in demand.

Perfectly Elastic: When the elasticity is $\epsilon = -\infty$, we say that demand is perfectly elastic. This means that any change in price at all will lead to an infinite change in demand. This isn't quite realistic, but some goods might behave approximately like this.

Usefulness Elasticities are an extremely useful way to think about effects because it's unit-free. Elasticities are excellent for translating data into math or math into data since it's such an intuitive measure and is still based on calculus. Elasticities are a useful way to think about explaining economic effects to non-economists.

Constant Elasticity Since we begin to work in the world of elasticities, how do we simplify the amount of work we have to do with elasticities? One really useful type of demand function is one that gives exactly the same elasticity everywhere on the line. The general

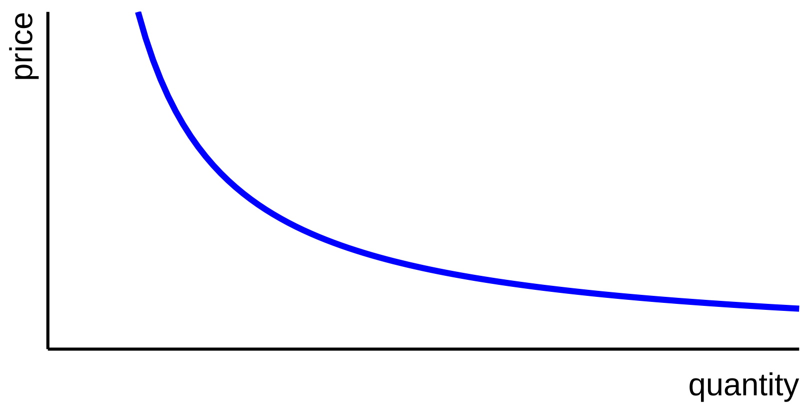

formula for such a demand function looks like

$$x = Ap^{\epsilon}$$
$$\ln(x) = \ln(A) + \epsilon \ln(p) \tag{24}$$

Obviously since a straight line gives varying elasticities, we need a nonlinear demand function to give us a constant elasticity. This function is very useful in applying the idea of elasticities to reality since it allows us to measure a *single* elasticity for a good rather than trying to estimate it for each different price-quantity combination.

You might notice that the Cobb-Douglas demand function can fit into this mould. For example, a demand function of $x = (\alpha m)/p$ can be re-written as: $x = \alpha m p^{-1}$. This implies that $A = \alpha m$ and the price elasticity of demand is -1. In fact, a Cobb-Douglas utility function is a special case where the price elasticity of demand is $\epsilon = -1$. Can you explain why?

Income Elasticity

We can do exactly the same gig for income. In the previous chapter we considered a change in income. We found that sometimes and increase in income increases demand (for normal goods) and sometimes it reduces income (for inferior goods). We measure the income elasticity of demand as:

$$\frac{\text{percentage change in demand}}{\text{percentage change in income}} = \frac{dx}{dm} \times \frac{m}{x} \tag{25}$$

Income Elasticities Income elasticities are usually positive rather than negative.

- For normal goods, the elasticity is between 0 and 1.
- For inferior goods, the elasticity is negative.
- For luxury goods, the elasticity is greater than 1.

On average, income elasticities will need to be just about 1. If an increase in income leads to a reduction in demand for an inferior good, it must mean an increase in demand for a luxury good. You must spend your money somewhere. That's simply by definition of the income elasticity, and has no real deeper economic meaning.

Notes

Flipped Remember that economists are silly and continue to flip everything. Your demand charts have price on the wrong axis, but you need to be able to read them both ways. Similarly, the text starts to pretend that price elasticities of demand are positive. Don't let it trick you. They're talking about absolute values. Why? Anyone's guess. This becomes especially complex when you begin to consider income elasticities which can easily be either negative or positive.

Marginal Propensity to Consume The derivative of the demand function for good x with respect to income $\partial x/\partial m$ can be called the *marginal propensity* to consume good x. It tells us how we expect the consumption of a good to change when the consumer's income changes. Again, this is very useful policy parameter because it can give us clues about economic policy changes.