Microeconomics with Ethics

by Steven Suranovic (2022) – George Washington University

Chapter 7
Theory of Consumer Demand

In this chapter we begin a long journey of model development with the final outcome being the mechanics of supply and demand curves in Chapter 14. This chapter focuses on the development of a model of market demand for a given product.

Demand, or more specifically consumer demand, refers to the desire to purchase goods or services by households, or consumers. A model of consumer demand seeks to identify the variables that can affect consumers' decisions about how much to buy of a particular product like bread or meat or beer.

A real world market rarely operates exactly like the previous pure exchange model in which Smith and Jones engage in barter exchange, trading one product for another (oranges for apples) in a market. Instead markets involve two types of individuals, those who come to a market to sell their goods and others who come to the market to buy those goods. The exchange involves a consumer (a demander) trading money for the product that is up for sale by the merchant (supplier). The standard economic model of supply and demand separates the decision process of the buyer and the seller and investigates what motivates the behavior on both sides of the transaction. At this stage, all thought will be on what motivates a consumer to go to the market to buy something. All worries about whether the product will be there or what motivates a merchant to put something up for sale is saved for later.

Building on the model of Smith and Jones, the theory of demand will first focus on the market for only one of the two products, say oranges. Jones is the individual who comes to the market wishing to buy oranges so he is a consumer demanding oranges. Smith comes to the market wishing to sell oranges so he is the merchant supplying oranges. In the next section, we'll consider more details about what affects Jones decision about how many oranges to demand. In the subsequent section, we will add up the individual demands of many consumers like Jones to construct a market demand function comprising numerous households who may wish to purchase a good.

Reality Check

Although in the real world, barter is fairly uncommon today, still every purchase in the market involves a kind of barter, or quid pro quo, in which one of the traded "goods" is money, or currency. Money is often said to "grease the wheels of exchange" because it makes it possible to enter any market as a consumer and purchase a desired good. In a barter economy, if Jones wants to buy Smith's oranges, he has to hope that Smith has a simultaneous desire to buy his apples. In other words, barter requires a "double coincidence of wants."

Money eliminates this problem as long as it satisfies an important criteria, namely money must be widely accepted in exchange for goods and services. If Smith knows that he can use the money given to him by Jones to later buy something else for himself, then he is more likely to accept that money in exchange for his oranges. Potentially a modern counterexample is the case of cryptocurrencies. Although many people have invested in cryptocurrencies today, they still fail to satisfy this criteria since cryptocurrencies cannot be used to but many things. For example, currently you cannot use it to buy your groceries. That may change in the future, but for now cryptocurrencies can at best be classified as a near-money.

There are several other important features of money, including its use as a store of value and a unit of account, but these are issues covered more carefully in a macroeconomics course, so I will leave that for you to investigate elsewhere.

7.1 Theory of Individual Demand

Learning Objectives

1. Learn how to construct an individual budget constraint in a two good economy.
2. Learn how a utility maximizing individual responds to changes in prices and income

In a typical market there are many individuals who consider purchasing a particular good that is available for sale. The total demand for a good in the market is thus made up of the summation of the demands by all of these individuals. We will assume that all of these individuals are affected by similar concerns that affect their decision to purchase the good, although each one of these concerns may influence different people to different degrees. This similarity allows us to initially focus attention on the behavior of just one individual and what might affect his or her desire to purchase, or demand, a particular good.

Suppose there is an individual, named Smith, who has a standard set of preferences over apples and oranges. By standard preferences, I mean those that include the "more is better" and the "diminishing marginal utility" assumptions, and which can be displayed as convex indifference curves in a diagram with the quantities of oranges and apples measured along the axes.

The Budget Constraint

Suppose Smith has a monthly income, I, measured in dollars, that he can use to purchase apples and oranges in the market. To keep things simple, let's also assume that there are no other products besides apples and oranges that Smith can purchase.

Next, let's assume that apples and oranges are for sale in the market at a fixed price. Let P_A and P_O be the prices of apples and oranges, both measured in $/lb. This means that Smith does not bargain, or negotiate, with the sellers of apples and oranges. Instead he must take the price as given. (Note: this assumption in consistent with the competitive market model that will be developed later.)

Finally, let's assume that Smith wishes to spend his income so as to maximize his utility from the consumption of apples and oranges. With all of these assumptions in mind, we can now investigate Smith's behavior in the market.

Define Smith's budget constraint as the set of all quantities of apples and oranges that Smith can purchase while spending all of his fixed income I. These purchase possibilities are described by the following equation, which is called his budget constraint.

$$P_O Q_O + P_A Q_A = I$$

where

P_O = price of oranges in \$/lb

Q_O = quantity of oranges purchased (demanded) in pounds per month.

P_A = price of apples in \$/lb

Q_A = quantity of apples purchased (demanded) in pounds per month.

I = income in \$/month

Note that the product of price and quantity, P_OQ_O, represents the total amount of money that Smith must pay for oranges, if he were to purchase Qo of oranges at the price P_O. For example, if Smith bought 5 pounds of oranges per month at the price \$3/lb, then at checkout his bill would be \$15 per month.

Since P_OQ_O is the amount Smith pays for oranges, then P_AQ_A is the amount he must pay for apples and $P_OQ_O + P_AQ_A$ is the total he pays for both goods. If Smith spends all of his income, then $P_OQ_O + P_AQ_A$ must be equal to his income I.

QUICK SIMPLIFICATION: In most textbooks, repeating the unit of time in the market quickly becomes cumbersome. Thus, for simplicity it is usually dropped. Thus, it is acceptable to talk about the amount Smith pays for oranges as \$15 rather than the more precise \$15 per month. Similarly, Smith's income might be reported as \$1000 instead of \$1000/month. From this point forward, I'll follow the usual simplifying convention, but do remember that implicitly there is always some unit of time in the model.

Figure 7.1 shows a graph of Smith's budget constraint and a plausible set of two indifference curves describing his preferences for apples and oranges. (Notice the simplification that quantities on the axes are measured in lbs now, rather than lbs/month.)

Figure 7.1 Utility Maximization with a Budget Constraint

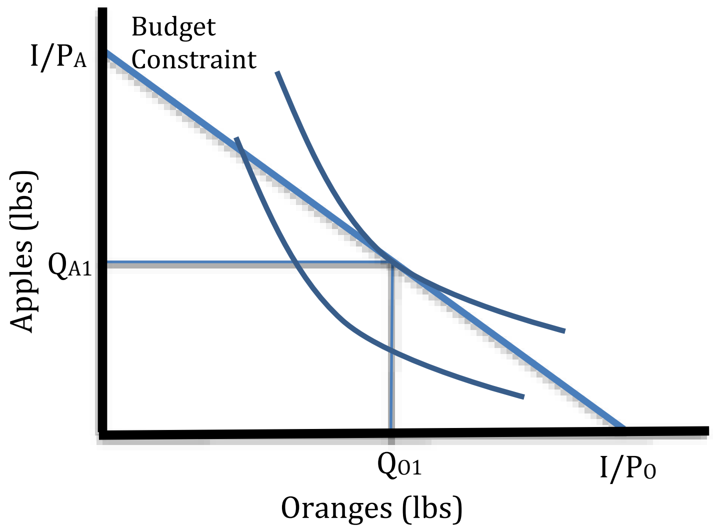

Smith's budget constraint is graphed as the line connecting the points, I/P_A on the vertical axis, and I/P_O on the horizontal axis. These endpoints can be found quickly by setting one quantity equal to zero in the budget equation and solving for the other quantity. This process is like answering the following question: If Smith did not purchase any oranges, → $Q_O = 0$, then how many apples could he buy with his income I? To answer, set $Q_O = 0$ → $P_AQ_A = I$ → $Q_A = I/P_A$. This is the vertical intercept value. Similarly, if Smith didn't buy any apples, then set $Q_A = 0$ → $P_OQ_O = I$ → $Q_O = I/P_O$. This is the horizontal intercept value. Finally, draw a straight line between the two points to graph the budget constraint. The budget constraint displays all of the combinations of orange quantities and apple quantities that Smith can purchase with his income, I, when the prices are P_O and P_A.

Notice also the slope of the budget constraint. It can be found using the rise over run slope formula using the endpoints on the two axes. In this case the rise is $- (I/P_A)$ and the run is (I/P_O), therefore the slope $= - (I/P_A)/(I/P_O) = - (I/P_A) \times (P_O/I) = - (P_O/P_A)$. This was defined as the terms of trade between oranges and apples in Chapter 3 and has units of pounds of apples per pound of oranges.

Side Note: The budget constraint is defined as the points Smith can purchase when he spends all of his income. But what if he didn't spend all his income? In that case, his purchases would lie somewhere inside (i.e., to the left or down) of the budget constraint. The set of all points that are feasible to purchase is called the *budget set*. The budget set consists of all points both on and inside the budget constraint.

Also, you may notice that this budget constraint resembles Smith's production possibility frontier described in Chapter 4. They are both straight lines with a negative slope in a diagram with quantities of oranges and apples on the axes. However, they are not the same thing. The budget line describes purchase possibilities for Smith who comes to the market with a fixed income. The PPF, in contrast, describes Smith's production possibilities in a model in which he can chose how much to produce of the two goods using his own labor input. Although Smith's two graphs are different, the maximization condition will turn out to be the same.

Utility Maximization

Smith is assumed to maximize his utility in this market situation. To do so he will choose a bundle of apples and oranges that achieves the highest utility among the bundles that he can afford to buy. The affordable bundles are all those contained within his budget set. The bundle with the highest utility can be found by superimposing his indifference curve map on top of the budget set. In Figure 7.1 the utility maximizing bundle is on the budget constraint with Q_{O1} pounds of oranges and Q_{A1} pounds of apples. Notice that this bundle also satisfies our earlier utility maximizing condition, namely,

$$ToT = \frac{P_O}{P_A} = \frac{MU_O}{MU_A} = MRS$$

Recall that the slope of an indifference curve at a particular point represents the marginal rate of substitution (MRS) which is the ratio of Smith's marginal utilities of oranges to apples respectively, at that point. Smith's MRS will equal the terms of trade only at the point (Q_{O1} , Q_{A1}) and thus this is his utility maximum.

Visually, you can see on the diagram that every other point in Smith's budget set would have to be on a lower indifference curve and therefore would have to give him lower utility. You can also see that, given these assumptions which in turn determines the shapes of the lines in the diagram, Smith must choose a point on the budget constraint itself if he is to maximize utility. Any point in which he does not spend all of his income would have another feasible point that gives him higher utility.

Evaluating the Effect of a Change in the Price of Oranges

We can use this simple model to evaluate how changes in some parameters would affect Smith's decision about how many pounds of oranges and apples he would purchase, or demand. The first change we'll consider is a change in the market price of oranges themselves. Suppose we begin with the equilibrium bundle (Q_{O1} , Q_{A1}) and assume that the price of oranges rises, say from P_{O1} to P_{O2}. To keep the analysis of the effect simple we also assume that no other exogenous variable in the model changes at the same time the price changes. Economists use the term *ceteris paribus* to signify that other variables remain at their original values. Thus we would say, assume the price of oranges rises, ceteris paribus.

When the price of oranges rises the budget constraint rotates leftward around the apple price intercept, as shown in Figure 7.2. This is because the orange price intercept decreases from I/P_{O1} to I/P_{O2} . Note that $I/P_{O2} > I/P_{O1}$ because $P_{O2} > P_{O1}$ and the price is in the denominator of the expression. After the orange price increases to P_{O2}, the consumer maxes utility by choosing (Q_{O2}, Q_{A2}). The effect of the orange price increase is to reduce the demand of oranges.

From this exercise we derive the negative relationship between the price of a good and demand for a good. If the price of oranges were to decrease, then the quantity of oranges demanded would increase and vice versa.

Figure 7.2 Effects on Demand of an Increase in the Price of Oranges

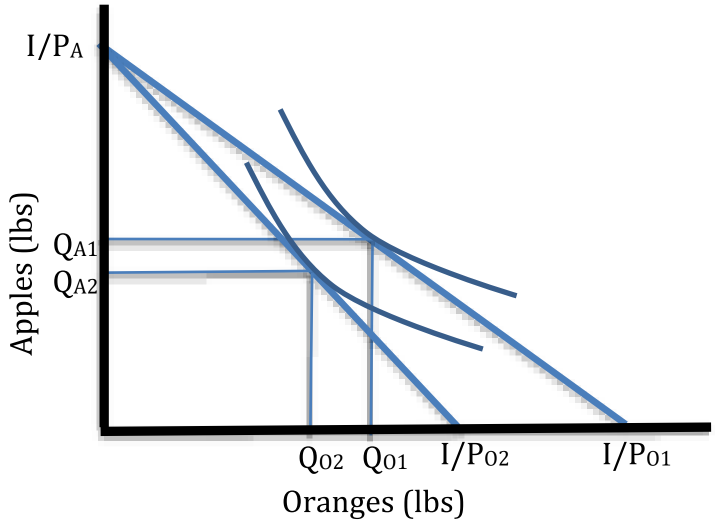

Evaluating the Effect of a Change in Income

We can also use this diagram to assess the effects of changes in other exogenous variables in the model. For example, suppose Smith's income rises, ceteris paribus. Again, ceteris paribus means that other variables, such as the price of oranges, do not change in value when income rises.

Let the original income level be I_1. In Figure 7.3, the consumer maximizes utility by choosing (Q_{O1}, Q_{A1}). Next assume there is an increase in income (from I_1 to I_2 such that $I_2 > I_1$).   The income increase will shift the budget constraint up and to the right – and parallel to the original budget constraint because the terms of trade, which is given by the slope of the budget line, does not change.

As shown, the new utility maximum is now at (Q_{O2}, Q_{A2}), Notice that the quantity Smith purchases, or demands, of both apples and oranges, increases after the increase in income. This is a very typical outcome that we would expect to see in most cases, namely that when income rises an individual will buy more of a good.   In this case, when demand for a good increases after an increase in income, we call that good a **normal good**.

However, such an outcome need not always arise as shown graphically in the Appendix below. It is possible for an increase in income to cause a decrease in demand for a good. In this case, when demand falls after income rises, we call that good an **inferior good**.

Figure 7.3 Effects on Demand of an Increase in Income

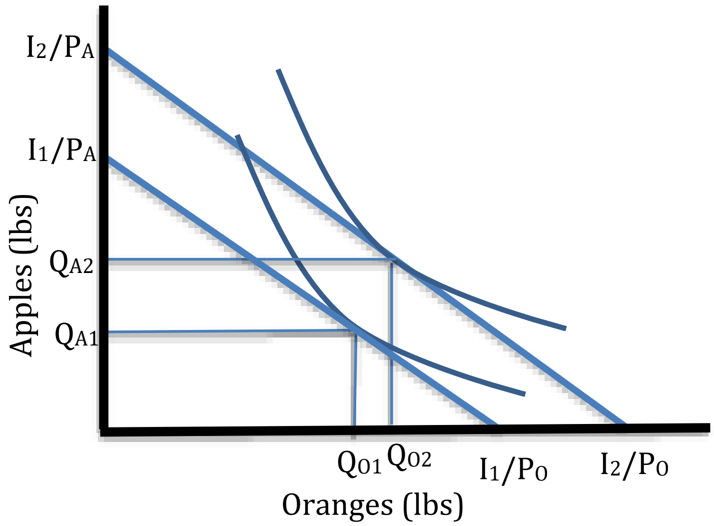

Key Takeaways

1. An individual budget constraint shows the combinations of quantities of two goods that can be purchased with a fixed income.
2. A utility maximizing consumer chooses the quantities of two goods to consume by finding the indifference curve that is tangent to their budget constraint.
3. A utility maximizing consumer would reduce the quantity of a good demanded when the price of that good increases under standard assumptions about preferences.
4. A utility maximizing consumer would increase (decrease) the quantity of a good demanded when their income increases and when preferences are such that the good is normal (inferior).
5. Ceteris paribus means that the values of other variables are held fixed when a change in one of the variables is analyzed.

7.2 Market Demand Determinants

Learning Objective

1. Learn the variables that can affect the market demand for a product and the direction of the cause and effect relationships.

The previous section demonstrated how a utility maximizing consumer would adjust their demand for a good when the price of the good changes, or when the individual's income changes, ceteris paribus. The exercises are meant to illustrate the decision process at the level of an individual. Going forward we will create a model of demand for a particular product, like oranges, at the level of the market, imagining that the market may consist of thousands or even

millions of individual consumers. Each one of these individuals can be thought to behave similarly to the individual in the exercises above.
Differences between individuals are allowable in the magnitudes of the effects of various changes. Thus, we will allow different people to respond to different degrees to a price change. For example, while Smith might reduce his demand for oranges by 1 pound when the price rises by $1, Jones might reduce demand by 1.5 pounds given the same price change. Also, while for Smith, oranges may be a normal good and a $1000 increase in income induces him to increase his demand for oranges, the same $1000 increase in Jones' income may induce Jones to reduce his demand. The market effects will be determined as the summation of all the individual responses to a change in one of the exogenous variables.

Although we could evaluate more variable changes using the utility maximizing model from section 7.1, instead we will proceed with a simplification by imagining a market demand function that is influenced by many other market variables beyond the two that were evaluated above.

Begin by defining a variable Q^D to represent the market demand for a particular good, let's say coffee, measured in pounds per month. (As before, we will drop the unit of time going forward for convenience, but it is important to remember that time is implicit in the model.)

We will assume market demand is affected by the following variables and can be written as a function of these variables, with units of each variable specified.

P - Own Price (i.e., the price of coffee ($/lb))
I - The consumer's income ($/month)
PS – the price of substitute goods (e.g. tea ($/lb))
PC – the price of complement goods (e.g. cream and sugar ($/lb))
T – tastes or basic desire for coffee (shape of indifference curves) (not directly measurable)
NB – number of coffee buyers in the market (# of coffee consumers)

Also, sometimes included are expectations of the future

FP – future price of coffee ($/lb)
FI – future income ($ per month) (when in the future usually unspecified)

A demand function in general form can now be written as,

$$Q^D = f(P-, I+ \text{ or } -, PS+, PC-, T+, NB+ \mid FP+, FI+)$$

Here, f(.) is meant to represent some unspecified functional form, called the demand function. The variable on the left, Q^D, is called the dependent variable because its value depends on the values of the other variables on the right. The variables on the right are called independent variables. To determine market demand one would first need to know the values of all the independent variables, then plug those into the demand function (if known), to determine the value for market demand.

However, even without knowing the correct functional form, or the specific values for the independent variables, we can still proceed with the model as long as we specify the direction of the cause and effect relationships. Those are given as the + or – signs after each independent variable on the right of the expression above. Next, we'll discuss each cause and effect relationship in turn.

Price - The first independent variable in the demand curve is the own price, P. If this is the demand for coffee, then the own price is the price of coffee. P has a negative sign next to it to indicate that the relationship between the price of coffee and the demand for coffee is assumed to be negative. This means that the two variables move in opposite directions. Thus, if the price of coffee rises, ceteris paribus, then the demand for coffee will fall. Also, if the price of coffee falls, then the demand for coffee rises.

This assumption is derived from the previous model described in section 7.1. That model shows that if a consumer maximizes utility, has standard preferences, and comes to the market to spend a fixed income, then there will be a negative relationship between price changes and demand. If all consumers in the market are similar to this representative consumer, then the market demand will display the same relationship. In section 7.1, the price-demand relationship is an implication of that model. However, if we begin the model with the demand equation above, then the negative relationship is an assumption.

Income - The second independent variable is income, I. Since the equation derives market demand and since the market has numerous consumers, the appropriate income is the summation of the incomes of all potential consumers in the market. A good way to describe this is as community income.

In section 7.1 we demonstrated that an individual who experiences an increase in their income may, under the circumstances shown, increase their demand for a good. This describes a normal good with a positive relationship between community income and the quantity demanded, if all consumers behave this way. However, we also noted that it is possible to draw the indifference curves differently to produce a negative relationship between income and the quantity demanded. This outcome is shown in the appendix and describes an inferior good with a negative relationship between income and demand. In the market with many consumers, whether a good is normal or inferior will depend on the aggregated effects across all consumers. If the consumers who behave as if the good is normal dominate, then the good would be measured as normal, even though some consumers may act as if it were inferior.

Side Note: Oftentimes economists constructing models will simplify the model to avoid discussing these complications. For example, if we assume all consumers have identical preferences, then all consumers will treat the good as normal, or, everyone will treat it as inferior, in which case the complication is swept under the rug, so to speak.

There is an intuitive explanation for why a good may be inferior for a consumer that is more important to know then the technical reason involving oddly shaped indifference curves. The term inferior is applied here because there are some products that consumers use only when they are unable to afford something better. For example, when I was young my family usually purchased margarine instead of butter because margarine was considerable less expensive. In this circumstance, a notable increase in family income may induce a switch from margarine demand to the better tasting product butter. In this case, margarine would be an inferior good since my family's demand for it would fall with increases in income. Note too that this effect might not be valid for all families. Some families might choose margarine over butter because of health reasons and thus may not switch between the two when income rises.

In summary, if the good is a normal good, then when community income rises, market demand for the good will also increase, and vice versa. However if the good is inferior, then when community income rises, market demand will fall and vice versa. This means that to answer any question that involves this relationship requires you to know whether the good is normal or inferior. If you don't know, or if it is unspecified in the question, then you must add an assumption to continue.

Substitute Prices - The next variable affecting demand is the price of substitute goods, PS. A substitute is any good that an individual might purchase instead of the good in the demand function. One might imagine what a person with a desire for something, such as coffee, would purchase if coffee were not available. Those other goods, and there may be many such goods, are substitutes. For coffee, tea is a likely substitute for many people. For those who dislike tea, hot chocolate may be a substitute. For most products, there are many potential substitute goods.

The independent variable is the price of these goods. We expect that those prices can affect demand for the original product in the following way. Suppose the price of tea were to increase. Because of the extra cost, individuals who normally drink tea might then look for substitute goods to purchase. If coffee were that substitute, then the individual would increase her demand for coffee when the price of tea increases because she reduces purchases of tea and buys more coffee instead. Thus we expect there to be a positive relationship between the price of a substitute good (tea) and the quantity demanded of the other good (coffee). In summary, when PS rises, ceteris paribus, coffee demand rises, and vice versa. Remember that in the aggregate with many consumers, there may be many substitute goods prices affecting the demand for coffee.

Side Note: Although we did not present the case of substitutes in the utility maximization model in section 7.1, the substitution effect is described in more detail in the Appendix. In brief, a change in the price of a substitute good causes a change in the terms of trade between the two products and causes consumption to shift along an individual's indifference curve. Thus as the price of the substitute good rises, the consumer naturally reduces demand for that good and increases demand for the other in order to maximize their utility. We will not use the utility maximization model to describe the remaining effects on demand, but you may see these presented in an intermediate microeconomics course.

Complement Prices - The next variable affecting demand is the prices of complement, or complementary, goods, PC. A complement is any good that an individual might purchase together with, or in conjunction with, the good in the demand function. In the case of coffee, many people use sugar and cream in their coffee. So imagine how coffee demand might change if there were a sizable increase in the price of sugar. Consumers might respond by switching to sugarless substitute drinks and thus their demand for coffee would decrease. This effect would be noticeable in the aggregate even if all coffee drinkers do not use sugar. The same effect would occur if the price of cream increased, so there may be several complementary goods for every product. Notice that the expected relationship is a negative one. If the price of a complement good rises, ceteris paribus, then the demand for the good falls, and vice versa.

Tastes - The next variable affecting demand is the tastes or preferences for the good, T. This variable is difficult to describe because it is not obvious how to measure it. It is included because it is clear that different people have different likes or dislikes for a particular good. Some people love coffee and some dislike it an never consume it. For those who love coffee we could say that their taste for coffee is high. We can also use the taste variable to discuss how advertising might affect consumer preferences for a good. Since advertising makes people more aware of a product and attempts to induce a stronger preference to buy it, we can say that advertising increases the consumer taste variable. Thus, as tastes, or preferences, for a good increases, then the demand for the good will increase as well, and vice versa.

Number of Buyers - The next variable affecting demand is the total number of consumers in the market, NB. This variable is much easier to measure than tastes or preferences and the cause and effect is obvious. If there is an increase in the number of coffee consumers, ceteris paribus, then there will be an increase in the total demand for coffee, and vice versa. This is a positive relationship and not much more needs to be said.

Future Price and Future Income – These final two variables are sometimes included in a demand function specification, and sometimes not; namely, the future price of the good, FP and the future community income, FI. By including them, one recognizes that time may be an important component of present demand because of the way that expectations can affect behavior. Suffice it to say that introducing time greatly complicates the model, so much so, that most traditional economic analysis is called "static," which means that time does not play a role. Time and expectations does have a large literature in economics, but because of its complexity it is usually only introduced in this simple way in microeconomics principles courses.

Expectations of the future price of coffee can affect today's decision about how much to buy in the following way. Suppose you are a regular coffee drinker and you learn that because of a drought in coffee growing regions, the price of coffee is likely to rise in the near future. In order to reduce your annual cost of coffee consumption, you might rush out and purchase more coffee than usual while the price remains low. In other words, stock up on coffee while the price is low and then purchase less when the price rises later. If many coffee consumers hear the news and react in the same way, then when the expected future price rises, current demand for coffee also rises, and vice versa.

Expectations of future income can also affect a person's decision about how much to buy of things today. To illustrate allow me to change the good being demanded to automobiles. Next suppose that as college students prepare for graduation each year, they also receive job offers for work that will begin several months later after graduation. These students, anticipating a regular paycheck will begin soon, may rush out now and purchase a new car, using their job offer as proof of their ability to pay back a new car loan. Thus, in general, if consumers in a market expect their income will rise in the near future, they may use their credit cards and demand more now, and vice versa. Hence a positive relationship between future income expectations and current demand for goods of all sorts.

Key Takeaways

1. Eight variables were identified whose changes could cause a change in the market demand for a good. These are (with the direction of the relationship indicated), Own Price (-), Community Income (+ or -), the Prices of Substitutes (+), the Pries of Complements (-), Tastes (+), the Number of Buyers (+), Future price (+) and Future Income (+).
2. The demand effects from changes in community income depend on whether the good is normal or inferior.

7.3 Graphing a Market Demand Curve

Learning Objectives

1. Learn how to graph a demand function using a simplified linear form and a generalized form.
2. Learn what variable changes cause a shift in the demand curve and in what direction.

Now that we have specified the variables that can affect demand for a product in the market, it will be useful to display this relationship in a graph. However immediately there is a problem. The demand function specified above is a function with one dependent variable and eight independent variables. To graph this relationship completely we would need a nine-dimensional graph. A graph on paper though, only has two dimensions, and at best we can only imagine and draw a three-dimensional relationship. So what to do?

The answer is to draw a graph in two dimensions by choosing the most important cause and effect relationship we care to display. But to do this, we need to make some additional assumptions.

First, of all the most important relationship is between the own price of the good and the quantity demanded of that good. Although it is typical in mathematics to display the independent variable on the horizontal axis and the dependent variable on the vertical axis, in economics the tradition has been to reverse this. Hence the dependent variable in a demand function, the quantity demanded Q^D, is plotted on the horizontal axis and the dependent variable, the price P is plotted on the vertical axis.

From the model description and the explanation given in section 7.2, we know that these two variables are negatively related to each other. That means that a graph of the demand function must be a curve that has a negative slope.

Figure 7.4 Possible Demand Curves with a Negative Slope

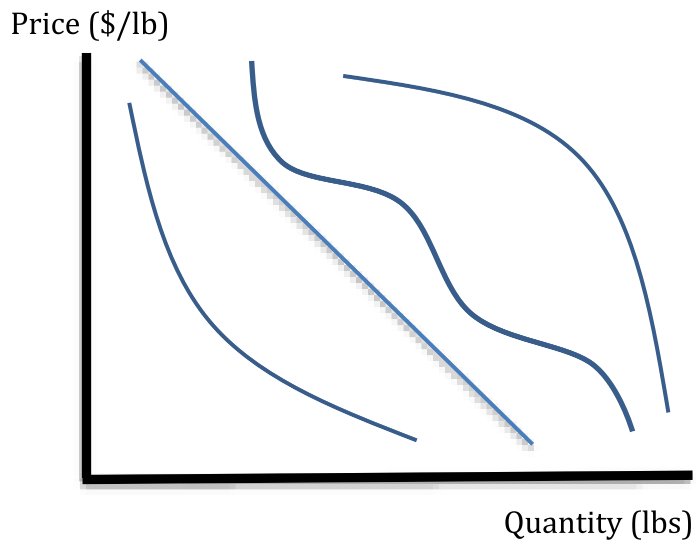

Each of the lines in Figure 7.4 satisfies the negative relationship between the price of coffee and the demand for coffee, namely as the price falls, the quantity demanded increases along each drawn line. But this raises a question. Which line is the best representation of the demand for the product? And, how would we determine the appropriate market demand function for the product we are interested in? The answer is somewhere between very difficult to almost impossible to discern. And so we will proceed by making some additional simplifying assumptions.

Reality Check

For students new to economics one might imagine that because economists are known to quantify and measure many things in the economy, and because economics is perhaps the most scientific of the social sciences, that economics would be able to identify with some precision the demand curve for many different goods. One might imagine that if you went into a large firm producing and selling coffee, or butter, or computers, that they would be able to show you the equation with the demand function for their product. The reality is quite a bit different though.

One reason is that the demand function is little more than a hypothetical demand that has never been actually realized. For example, if Figure 7.4 were displaying the demand curve for ground coffee, then the prices shown on the vertical axis might range from $0 to $150. It would show the demand for coffee at prices that have never been seen in the world. There would be demand at a price of $100 per pound, but having never witnessed that price, we could only guess what the demand would be. In fact, the most that economists can do is to estimate the sensitivity of demand changes to changes in prices and only in the historical range of prices. (More on sensitivities in Chapter 8). Even if that measurement is accurate, there is no guarantee that relationship holds over the entire range of potential prices, or that the market demand today is the same as the demand that existed when the data was measured.

The point to understand here is that there are numerous measurement problems that must be dealt with in economic models. Some problems can be "solved," but even these solutions are little more the educated guesses, or estimates, that often require us to make even more simplifying assumptions. In physics, there are inviolable laws because atoms never feel and act differently today than they did 10 years ago. The social sciences study people though and people do behave differently at different times and that means that even if we could use data to construct an accurate demand curve, that curve would be true of the past, but not necessarily true of the present and future.

Many people believe that economics intends to predict economic outcomes today and in the future. Some economists may believe that too. However, the problems inherent in something as simple as specifying the demand curve for coffee are multiplied many times over in more complex economic relationships and should serve as a warning that effective prediction is unlikely to occur anytime soon.

All hope need not be lost though. Even if we can't quantify and measure the precise demand relationship in a market, we can still gain a lot of understanding of how markets function by making some simplifying assumptions and pushing forward. And so that's what we'll do.

A Linear Demand Function

The best way to simplify the demand curve so as to preserve its key characteristic (the negative relationship between price and quantity) but make it easy to work with, is to assume that demand is linear, i.e., described by a straight line. The second thing we must do is explain what happened to all the other demand determinants.

Suppose the demand for coffee can be written as the following linear equation.

$$Q^D = -P + (1/10)I + (1/2)PS - (1/4)PC + 5$$

On the left side is the dependent variable Q^D, measured in pounds and on the right side are several of the independent variables described in the general demand function in section 7.2. Some things to note:

1) Each variable appears with a parameter value in front of it. The P has a -1, PS has +(1/10), etc. This form means the equation is linear with respect to every independent variable.
2) The sign of the parameter indicates the relationship between the variable and Q^D. Thus, because P has a -1 parameter value, it means that as P rises, Q^D falls. Also, since I has a positive parameter value it means we are also assuming the good is normal.
3) The parameter values, and the values of the independent variables, may or may not correspond to the values these might take in the real world.
    a. For example, later I will use and income value of 300, but that value does not match any real world measurement of income.
    b. We only care here about the direction of changes, not so much if the values used are realistic. However, in some examples and for some models in economics courses, we do try to make the values at least partially realistic.
    c. Unrealistic values are used primarily to make computations easier.
4) Since there is only one PS entry and one PC entry, it means we are assuming this product has only one substitute good and one complement good.
5) The +5 parameter value at the end of the equation is meant to incorporate the realized values of the remaining unspecified independent variables. In other words, the values for tastes, the number of consumers, the future price and the future income all combine to generate the value +5.

Next let's graph the equation. However, there remains a problem. We cannot graph the demand curve on a two-dimensional diagram unless we specify values for the variables that do not appear on the axes of the graph. In other words, to proceed we must assume some value for I, PS and PC. So, let's assume I = 300, PS = 20, and PC = 20. Also, I am not specifying the units of each variable mainly because the scale of these variables are not realistic. This leads to an important realization about economics: units are very important in economic models, except when they're not!

Plugging these values into the equation above yields,

$$Q^D = -P + (1/10)*300 + (1/2)*20 - (1/4)*20 + 5$$

$$= -P + 30 + 10 - 5 + 5$$
$$= -P + 40$$

$$Q^D = 40 - P$$

Notice that this equation is now written as a function with only two variables. However, the values of all the other variables combine implicitly to generate the value of 40 in the equation. This equation tells us how Q^D will respond to changes in P, but it does so while assuming ceteris paribus, meaning that the values of all the other independent variables are remaining fixed, or constant, at their original values.

Side Note: Sometimes in other economics texts ceteris paribus is defined as meaning "all else equal" which often generates some confusion. "All else" clearly refers to all the other independent variables, but what are they all equal to? They can't be assumed equal to each other, because they are all measured differently. Instead, what it means is that they are all equal to their original values. To me, it is clearer to say "all else fixed," rather than "all else equal."

The demand curve can be displayed in several ways. One way is to create a demand schedule by plugging in various prices and indicated what demand would be, as shown in the following Table

Price ($/lb)	Quantity Demanded (lbs)
$40	0
$30	10
$20	20
$10	30
$0	40

We can also plot the function or demand schedule as shown in Figure 7.5. Note that a reduction in the price of coffee raises demand by moving along the curve (no shift occurs because the coffee price is plotted on the axis)

Figure 7.5 Graph of a Linear Demand Function

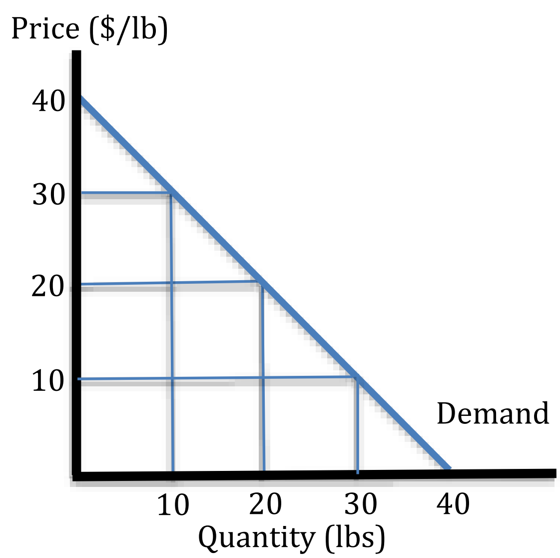

Shifts in the Demand Function

The demand function as graphed above will shift its position whenever one of the other independent variables, i.e., those not drawn on the axes, changes. To illustrate, suppose the variable income changes in value from 300 units to 200. To adjust the curve, we first must go back to the more general linear demand function, plug in the new value for I, and calculate the new demand function equation.

$$Q^D = -P + (1/10)I + (1/2)PS - (1/4)PC + 5$$

$$Q^D = -P + (1/10)*200 + (1/2)*20 - (1/4)*20 + 5$$

$$= -P + 20 + 10 - 5 + 5$$

$$= -P + 30$$

$$Q^D = 30 - P$$

The red values indicate the values that have changed from the original exercise. The new demand schedule can now be written alongside the original as,

Price ($/lb)	Quantity Demanded (I=300) (lbs)	Quantity Demanded (I=200) (lbs)
$40	0	-
$30	10	0
$20	20	10
$10	30	20
$0	40	30

The new demand curve graph, labeled D2, is shown below in Figure 7.6, also alongside the original demand curve, labeled D1. Notice that the effect of the decrease in income is to shift the demand curve to the left. Also recall that this shift has already incorporated the assumption that the good is normal.

Figure 7.6   A Shift in the Demand Curve

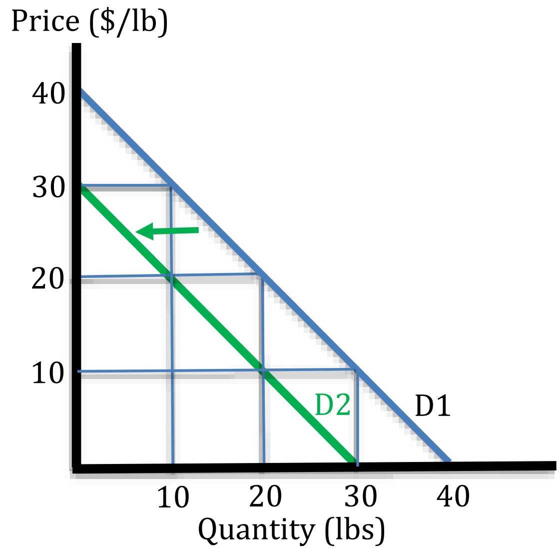

Here are some things to note:

1) The plot of the original demand function D1 assumes that income has the value 300, while the new demand function assumes an income level of 200.
2) Any change in the price of the good, P, changes the quantity demanded along the demand curve and assumes ceteris paribus, meaning that income is assumed fixed.
3) A change in income shifts the entire demand curve in the direction of the relationship. Since the good is assumed to be normal, the relationship between income and demand is positive. Thus, when income falls as in this exercise, the entire demand function shifts down (same direction) along the quantity axis, meaning leftward.
4) The exercise generalizes for all the other independent variables
    a. If PS rises, ceteris paribus, then because PS has a positive relationship with the quantity demanded, the demand curve will shift upward along the quantity axis, meaning to the right.
    b. If PC rises, ceteris paribus, then because PC has a negative relationship with the quantity demanded, the demand curve will shift downward along the quantity axis, meaning to the left.
    c. Ditto, for the remaining variables.

It should now be clear how a demand function with multiple independent variables can be represented in a diagram with only two dimensions. The two most important dimensions are plotted explicitly (P and Q), while the remaining variables implicitly define the position of the curve on the diagram. Whenever there is a change in one of the implicit variable values, it must cause a shift of the plotted demand relationship.

Demand Shift Summary

We summarize all of the possible shifts that may occur below. However, before doing so, it is worth commenting on the distinction between specific demand functions and general ones. In the previous exercise we wrote down a specific linear demand function with specific parameter values and then worked with it to illustrate how a demand curve is drawn and how it may shift. Such an exercise is useful because one can see the mechanics of the model more clearly.

However most of the time we will not need that level of detail, because it will quickly become cumbersome. Instead what we will need to illustrate most of the principles in economics is a graph that displays the directions of the cause and effect relationships rather than the magnitudes. In other words, we need only draw a negatively-sloped line, or curve, to keep track of the fact that demand increases with decreases in the price. What the actual values being measured are, will be immaterial. We can call this a conceptual model because it focuses on the general principles rather than specific values.

For example, we need not use an explicit function to illustrate the effects of a decrease in income on demand for a good. Instead, we just need to know that the positive relationship (assuming the good is normal) implies that a decrease in income will cause the demand curve to shift left. Later when we combine demand with supply, we will only seek to learn how changes in demand affects the direction of the market price change (does it rise or fall) and will not ask by how much does the price rise or fall.

As we did above, occasionally we'll illustrate the mechanics of the models by using specific functions and we'll ask students to work with these functions to learn the underlying relationships. But, most of the key results can be illustrated merely by displaying the general relationships and that is how we'll conduct much of the analysis that follows.

Figure 7.7 displays two general linear demand functions. They are general because there are no explicit values shown for the price and quantity. They are linear because they have no curvature. The negative slope incorporates the negative relationship between changes in the price of a good and changes in the quantity demanded in the market. To the right of the graph is shown all of the other variables that are implicitly defined in the model and the changes in those variables that would cause the demand curve to shift leftward from its original position at D1 to its new position at D2.

Figure 7.7 Demand Left-Shifters

A change in the following variables in the direction indicted would cause the demand curve to shift to the left.

↓ I income decrease (if a normal good)
↑ I income increase (if an inferior good)
↓PS substitute good price decrease
↑ PC complement good price increase
↓ T tastes for the good decreases
↓ NB number of buyers decreases

↓ FP expected future price decreases
↓ FI expected future income decreases

Figure 7.8 displays two general linear demand functions. To the right of the graph is shown all of the other variables that are implicitly defined in the model and the changes in those variables that would cause the demand curve to shift rightward from its original position at D1 to its new position at D2.

Figure 7.8 Demand Right-Shifters

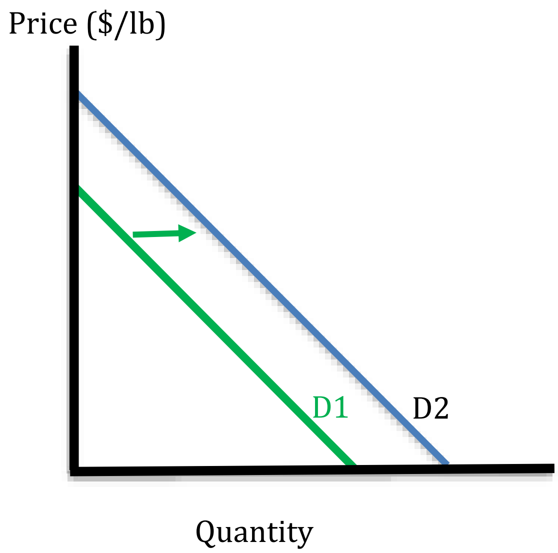

A change in the following variables in the direction indicted would cause the demand curve to shift to the right.

↑ I   income increase (if a normal good)
↓ I   income decrease (if an inferior good)
↑PS   substitute good price increase
↓ PC complement good price decrease
↑ T   tastes for the good increases
↑ NB   number of buyers increases

↑ FP   expected future price increases
↑ FI   expected future income increases

Side Note: Some textbooks use the word "demand" to mean the demand curve and to distinguish it from the quantity demanded. Under this naming convention, an increase in the price of a substitute causes an increase in demand for coffee, whereas a decrease in the price of coffee does not affect demand. However, others may use demand to mean the quantity demanded, in which case a decrease in the price of coffee will increase demand. There is no "right" way to define these terms, only different conventions. In this text I will always be explicit by using the phrase demand curve when I am referring to the curve.   If I ask what happens to the demand for coffee when the price of coffee falls, the answer is that the demand increases (as illustrated by moving along the curve). Thus, I'll take demand to be synonymous with the quantity demanded.

Key Takeaways

1. The market demand curve is a negatively sloped line which indicates that market demand rises as the price of the product falls and vice versa, while holding all the other independent variables fixed in value.
2. A linear demand function can be used to work numerically with the cause and effect relationships.
3. A general demand function can be used to work conceptually with the cause and effect demand relationships.
4. All variables that can cause shifts of the demand curve and the directions of the shifts are summarized in Figures 7.7 and 7.8.

Appendix 7.1 Example of an Inferior Good

As discussed in section 7.1, it is possible for an increase in income to cause a decrease in demand for a good. We call such a good an **inferior good**. Inferior goods arise out of a particular type of preferences as illustrated in Figure 7.9.

Figure 7.9 Income Increase for an Inferior Good

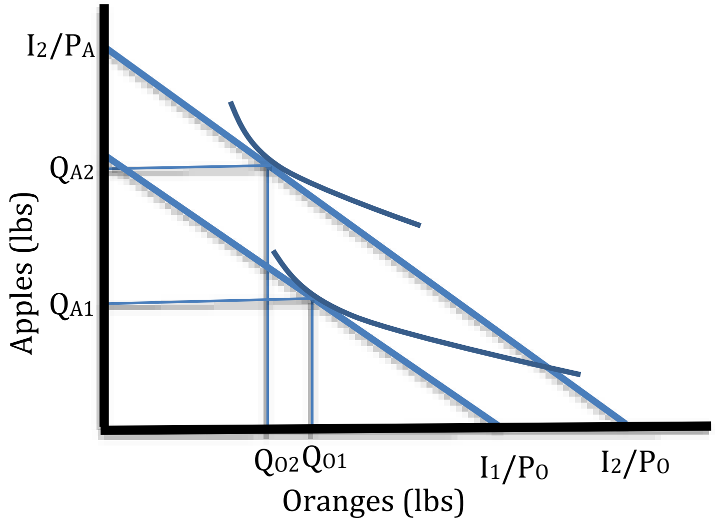

Let the original income level be I_1. In Figure 7.9, the consumer maximizes utility by choosing (Q_{O1}, Q_{A1}). Next assume there is an increase in income (from I_1 to I_2 such that $I_2 > I_1$). The income increase will shift the budget constraint up and to the right – and parallel to the original budget constraint because the terms of trade, which is given by the slope of the budget line, does not change.

As shown, the new utility maximum is now at (Q_{O2}, Q_{A2}), Notice that the quantity Smith purchases, or demands, of oranges, decreases after the increase in income. Thus, oranges would be called an inferior good because of the negative relationship between income and the quantity demanded.

Notice that to produce such an effect the indifference curves need to be drawn in a particular way. They still conform to the standard assumptions that more is better and exhibit diminishing marginal utility, but they are skewed such that an increase in orange consumption does not raise utility very much compared to the increase in utility caused by a comparable increase in apple consumption. Graphically, the indifference curves are much closer together along the orange axis which means that this person does not get nearly as much happiness from an extra orange compared to the happiness from an extra apple. Thus, as income increases, this person would prefer to shift consumption away from oranges and purchase the more desirable apples instead.

Note also that apples are a normal good in this example because as income increases the consumers demand for apples also increases.

Appendix 7.2 Income and Substitution Effects

It is possible to decompose the effects of a price change into two separate effects; an income effect and a substitution effect.   The total effect of a price change is the summation of these two effects.

Consider Figure 7.10 depicting an increase in the price of oranges from P_{O1} to P_{O2}. The total effect of the price change is to shift the demand for oranges from Q_{O1} to Q_{O3} and the demand for apples from Q_{A1} to Q_{A3}.

Figure 7.10 Decomposing Income and Substitution Effects

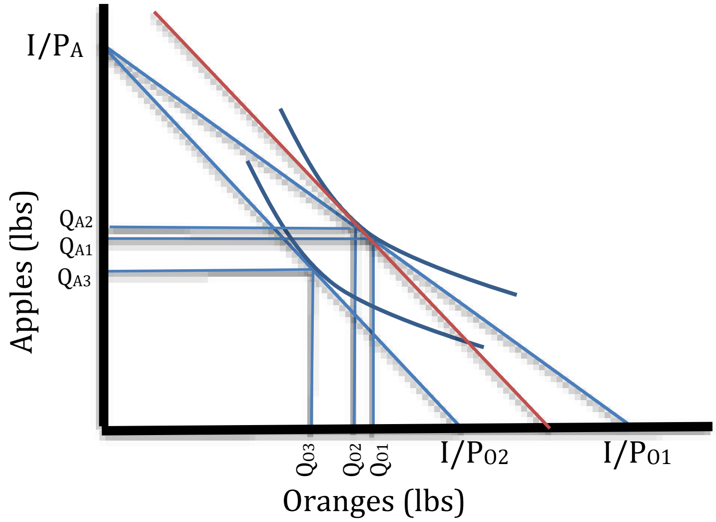

The price increase changes the terms of trade (the slope of the budget constraint) from P_{O1}/P_A to the steeper P_{O2}/P_A. This price increase also lowers real income of the individual. Real income is defined as the purchasing power of the nominal (or money) income I.

The real income with respect to apples is the quantity of apples that can be purchased with one's total income and is given by the apple intercept value I/P_A.   In the example, this does not change when the price of oranges rises.

The real income with respect to oranges is the quantity of oranges that can be purchased with one's total income and is given by the orange intercept value.   This value falls from I/P_{O1} to I/P_{O2} when the price of oranges rises.

Thus, the real income of the individual falls with a price increase because purchasing power of one good falls while purchasing power of the other good stays the same.

The substitution effect of a price change is found by answering the following question: How does demand for oranges and apples change because of the change in the terms of trade IF real income were to stay the same.

Real income can be hypothetically kept constant by finding the new level of income that would allow the individual to reach the same level of utility as was realized before the price of oranges changed. We find that hypothetical budget line by pushing the new budget line with endpoints

I/P_A and I/P_{O2} upward (and to the right) until it is tangent to the original indifference curve. The new hypothetical budget line is depicted by the red line on the graph.

This shows that if the individual faced a change in the price of oranges, but no change in real income, then demand for oranges would fall from Q_{O1} to Q_{O2} and the demand for apples would rise from Q_{A1} to Q_{A2}. In other words, the individual "substitutes" apples for oranges when the relative price of oranges rises. Hence it is called the substitution effect.

The income effect component is represented by the rest of the movement to the final equilibrium. That is, from Q_{O2} to Q_{O3} and Q_{A2} to Q_{A3}. The income effect corresponds to the changes in quantities demanded caused solely by the decrease in real income after the substitution effect is already accounted for. The income effect is such that a decrease in income cause a reduction in demand for both apples and oranges. In other words, both are normal goods in this example.