PART I: Neoclassical Price Theory

I. Introduction: A Short History of Price Theory

The logic of supply and demand is ancient. It would, for example, have been "intuitive" to sellers of goods and services for thousands of years before economic models were worked out in the nineteenth century. The prices at which they could sell their goods would obviously be higher when buyers had relatively intense desires for their goods and lower when they had to compete with other craftsmen and merchants for the purchases of potential buyers. Business necessities and common sense would have linked their minimum selling prices to their costs of production. To be profitable, the selling price had to be greater than a good's cost, otherwise there would be no benefit associated with bringing it to market. Ideally—from the perspective of sellers—the price would be much higher than their costs, but prices could clearly be too high without reducing sales to a pittance or to zero. Thus, pricing was in practice constrained to a relatively narrow band of "reasonable" or "fair" prices by the limitations of consumer demand, the cost of producing goods and services, and competition from rival craftsmen and merchants.

Once general pricing or markup rules of thumb were developed, they would be passed on to successive generations of artisans, merchants, and middlemen. These rules were arguable the first economic price theory.

More general principles for understanding market prices were subsequently developed by entrepreneurs with a scholarly bent and by students of markets. Among the latter, early and relatively well-known examples can be found in Aristotle's books on politics and ethics. The former provides a theory of money goods and inflation. The latter provides a theory of relative prices that is worked out to illustrate his concept of relative justice. His brief discussion is largely compatible with current theories of general equilibria in perfectly competitive markets. In his theory, the band of "fair" relative prices was a narrow one—essentially that associated with zero profits for speculators and middlemen. Athens and several other Greek city states were centers of commerce during Aristotle's lifetime, with trading posts distributed around the shores of the Mediterranean Sea. Athens was also a center of intellectuals, schools, and scholarship, and thus it is natural that at least a few Greek scholars were interested in the properties of market networks.

Contemporary efforts to develop principles or economic laws that could describe prices in general began with Adam Smith's most famous book published in 1776, *the Wealth of Nations*. In that book, Adam Smith develops a theory of long-run price determination based on the number of hours of labor embodied in the product—the more hours of work required to produce something, the more it would cost relative to products that required fewer hours to build. He also noted that the number of hours required varied with technology and specialization. Much of the book is devoted to empirical demonstrations of that theory and also of the effect of money goods on long term prices and price fluctuations. It should be kept in mind that Smith was not attempting to provide a model that would precisely describe prices—possibly because he did not believe that one was possible—but rather to characterize in a rough way equilibrium prices. He did, however, suggest that an invisible hand guided the economy towards prosperity if not unduly restrained by governmental regulations.

It should also be kept in mind that during Smith's time, most skills could be learned through a relatively short period of apprenticeship and thus differences among types of labor could be neglected without much loss or assumed to be produced by labor during the period of apprenticeship. Commerce was growing in Western Europe, but most people were still farmers, farm workers, or servants in this period, although commerce was about to expand rapidly. Economics, to the extent it was studied, was undertaken by philosophers such as Adam Smith, clergymen such as Thomas Malthas, and practical businessmen such as David Ricardo and John Stuart Mill who had sufficient time and interest in markets to speculate a bit about how they operated. It was a field of study, but not one favored by many academics.

Three quarters of a century after Adam Smith's treatise, three scholars independently worked out economic theories based on the logic marginal utility at about the same time: William Jevons (1862), Carl Menger (1871) and Lèon Walras (1874). The concept of utility was at the heart of utilitarianism, an important strand of normative theory worked out Jeremy Bentham and many others in the previous half century or so. It argued that all humans attempt to maximize utility—a single index of all the human aims and purposes—and that the best society was the one that produced the greatest sum of utility for its residents (or the greatest average happiness). Several of its leading utilitarian theorists and practitioners wrote economic textbooks including

Bentham (xxxx) and Mill (xxxx). They had understood that utility was the basis of demand, but not that marginal utility was the relevant factor.

In Jevon's case, he had worked out the mathematics behind utility maximization which led directly to the marginal utility concept—as the first derivative of a utility function, which is directly relevant for an individual's decision about the utility maximizing quantity of a good an utility maximizing individual will be inclined to purchase. Menger had emphasized the subjective nature of both utility and marginal utility. Walras (1877) subsequently extended the idea of a general equilibrium for an entire marketplace, with prices simultaneously clearing all markets. Prices, he argued would adjust through trial and error (e. g. a *tâtonnement* process).

It can be argued that after these three contributions by marginalists and the earlier ones by classical economists, neoclassical economics had been given its general form and focus. However, it took another 75 years for the neoclassical synthesis to be worked out. Indeed, it is still being worked out.

II. Organization and Focus of Part I

Part I of this text reviews the core ideas that emerged from that synthesis. Most of these ideas were worked out in the post-war period, between 1945 and 1975. It was also in this period that contemporary microeconomics textbooks came to have their present focus on topics and that use of calculus to demonstrate important economic relationships replaced geometric intuitive representations in mainstream economic journals. This caused economic theory to look like a "hard science" like physics (This may be partly because Samuelson, one of the pioneers of the neoclassical synthesis, was trained as a physicist.)

Part II explores extension of the core to take account of time, uncertainty, information asymmetries, and innovation. Parts III and IV explore extensions to fields outside of economics per se. Many relationships in social science can be understood as outcomes of more or less rational decisions by the individuals involved in settings of scarcity.

Part I consists of 3 chapters. Chapter 2 develops the neoclassical theory of demand. Chapter 3 develops the neoclassical theory of supply, and Chapter 4, develops their implications for equilibrium prices. Although general equilibrium models exist, the models reviewed in these chapters are of the partial equilibrium variety. The Appendix to Chapter 4 provides a short overview of general equilibrium analysis. Each chapter in Part I begins with geometric illustrations, then develops mathematical results using concrete functional forms, and ends with more abstract forms and more general results.

Chapter 2: Routine Consumer Choices and Market Demand

I. Introduction:

The neoclassical synthesis emerged after World War II. It used mathematics to systematically linked routine production and consumption decisions to produce a general theory of price determination. For the most part, it implicitly focused on routine choices—that is to say, choices with respect to well-understood goods and services that are produced by methods that are also well-understood. Such choices are in a sense perfectly informed because those making the choices have completely accurate expectations about the nature of the goods and services at issue and their manner of production. Not all choices are routine, but the models worked out clearly illustrated that fundamental relationships existed between consumer preferences, production methods (technology) and the prices of such goods.

It should be acknowledged that the neoclassical synthesis predicts a narrower band of prices than actually observed in practice. This is, course, because models of economic tendencies are being used to analyze what are believed to be fundamental properties of market systems. As models, they necessarily abstract from many details of economic life with the hope that the most important relationships might be worked out. As a collection of model, the neoclassical synthesis does not attempt to explain every possible event in an economy, only the ones that are commonplace or that tend to be commonplace. It should also be acknowledged that it's lack of precision—in the sense of departures from realism—is generated by the assumptions and mathematical tools used to characterize choices and outcomes. This may be why the acceptance of mathematical models took so long to emerge among economists—and remains somewhat controversial among a subset of economists today.

Nonetheless, by rigorously illustrating how equilibrium prices emerge and affect (coordinate) the decisions of consumers and producers, neoclassical models demonstrate that an internally consistent theory of prices is possible. Verbal and geometric illustrations had been undertaken before, as in Marshall's textbooks and other that followed his lead. The ew mathematical models demonstrated that earlier intuitive and geometric claims about equilibrium were not ideological or logically incoherent. A theory of equilibrium prices could be grounded plausible mathematical characterizations of choices that produced logically rigorous and clear implications about prices in well-functioning market networks.

II. The Geometry of Net Benefit Maximization and Consumer Demand

The geometric approach to economics began in the late nineteenth century and continues through to today, although in published work diagrams are not the main engine of analysis as they often were in the century between the marginal revolution and the 1970s. Instead, diagrams are used to illustrate what is being derived using other mathematical tools. Such illustrations are useful because most people's geometric intuition is better than their calculus or real analysis intuition. It is for that reason that the chapters of Part I begin with geometric derivations and diagrams. It turns out that the same diagrams can be used to illustrate the results derived using calculus in the sections that follow those illustrations.

In general, geometric derivations assume that decision makers are either "net benefit" maximizers or "utility" maximizers. A person that maximizes net benefits maximizes the difference between the benefits measured in some currency, as with dollars, euros, yen, or yuan,, etc. and its costs measured in the same currency. Although the geometric approach can deal with discrete units as well as continuous or infinitely divisible units of goods and services, assuming that goods and services are infinitely divisible allows us to more easily connect up the geometry of this section with the mathematics of the sections that follow.

The term "marginal" is an adjective and means the change in something—such as benefit or cost or utility—associated with a one unit change in quantity (in the case of goods sold only in discrete units) or as the rate at which that variable (benefit, cost, utility, etc.) changes as quantity increases (in the case of goods sold in continuous units). In practice these days, goods tend to be sold in prepackaged discrete units or sizes, although it is still possible to purchase fruit and in some cases nails by the pound or kilogram, which would be a case where any among that can be measured can be purchased or sold. It is this later case that most mathematical models assume, partly because it makes it possible to use calculus to build models as seen in the next section.

Each quantity that might be purchased by a consumer is assumed to have a particular benefit level associated with it and a particular cost associated with it. Generally, this allows one to represent the benefits associated with the quantity of a particular good to be represented as a function that maps quantities into benefits. That function is generally assumed to be strictly concave. **A strictly concave function has the property that** a line (cord) connecting any two points on a graph of the function will lie below the function. When a benefit function is strictly, concave, continuous, and differentiable, the marginal benefit function is downward sloping. Strict concavity is an implication of both diminishing marginal utility and diminishing marginal benefit.

The intuition behind diminishing marginal utility is that an individual will use the first unit of a good for purpose that generates the highest benefit, the second unit for the second highest benefit and so on. Differentiability is not necessary for this relationship to hold, although there are other assumptions even in this case.[1] Similarly, the marginal cost of a good that sold without quantity discounts or surcharges is simply its price. It is the amount that one has to pay for an additional unit of the good, or the rate at which one's cost increases as quantity increases in the continuous case—which amounts to the same thing, e.g. $\Delta C = P\Delta Q$.

Figure 2.1: Maximizing Net Benefits

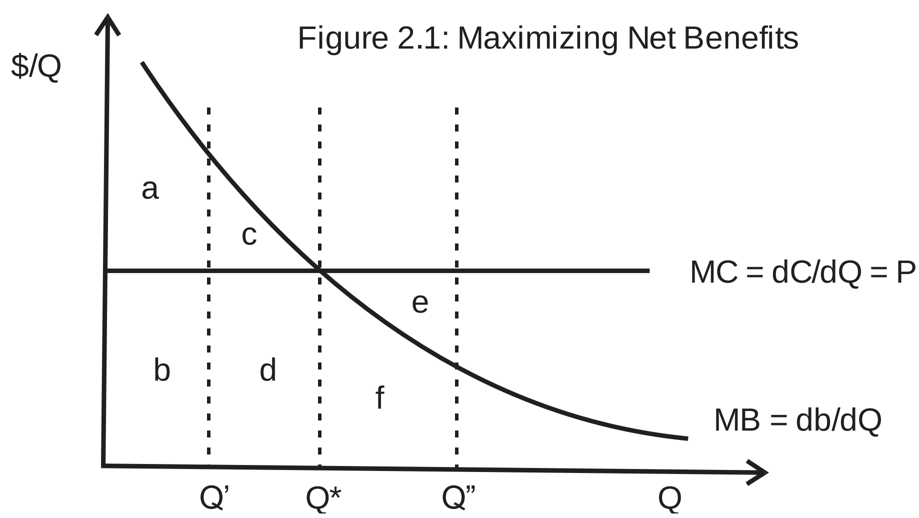

[1] For example, it assumes that one unit of a good can be applied to any use. However, if a particular number of the good completes a set and therefore allows a previously impossible use to be realized, it is possible that the units that complete a set have a higher marginal utility and marginal benefit than the previous unit(s).

Figure 2.1 illustrates a typical marginal cost curve and price line for a representative individual, who will be called Al (short for Allan or Alice). It turns out that the area under the marginal benefit curve from 0 to any quantity Q is the total benefit of Q units of the good or service of interest [b(Q)] and that the area under the marginal cost curve is the total cost of Q units of the good [c(Q)]. Thus, subtracting this cost from the benefit produces Al's net benefit for purchasing particular quantities of the good or service depicted in the diagram. The relevant areas are labeled with letters so the net benefit levels can be calculated for quantities Q', Q*, and Q".

The total benefits associated with quantity Q' is a+ b—the area under the MB curve from 0 to Q'. The total cost of Q' units is b, so the net benefits of purchasing Q' units at price P is area a—the area under the MC curve from 0 to Q'. Similarly, the total benefits associated with quantity Q* is a+ b + c + d. The total cost of Q* units is b+d, so the net benefits of purchasing Q' units at price P is area a+ c. The total benefits associated with quantity Q" is a+ b +c +d +f. The total cost of Q' units is b+ d + e + f , so the net benefits of purchasing Q' units at price P is area a + c - e. Notice that the net benefit realized at Q* is larger than the other two, which implies that net benefits (here, consumer surplus) is maximized at Q*, because any amount less than Q* will entail the loss of net benefits c and any amount more than Q* will entail losses of area e. A net-benefit maximizing consumer that purchase positive amounts of a good will select a quantity where his or her marginal benefits equal his or her marginal costs.

An exception to his rule occurs when the MC curve is above the MB curve for all quantities, in which case an individual will purchase zero units of the good. At zero, there are no benefits and no costs, so net benefits are also zero. This is larger than any positive quantity, because all positive quantities have costs that are greater than their benefits and thus negative net benefits associated with them.[2] Note also, that such cases are not "weird" or "unusual" cases,

[2] Another exception occurs when goods are available only in discrete units. In that case the Benefit and cost functions are a sequence of points rather than curves, and individuals will purchase the last amount where the marginal benefit is larger than marginal cost. In the rare case where there is nonetheless a unit where marginal benefit equals marginal cost, the individual will be indifferent between that unit and one fewer units of the good—and may simply "flip a coin" to determine which quantity to purchase.

but the most common most case experienced in well developed markets. For example, the average consumer at a grocery store purchases zero of far more goods than he or she purchases positive amounts of.

The above diagram assumes that the marginal benefit curve exhibits diminish returns and thus is monotonically decreasing in the quantity purchased or used. Other cases may be of interest, but these are the choice settings that attract most of the attention of neoclassical economics and so is the main focus of this chapter.

The result characterized by figure 2.1 can be used to help us determine Al's demand for the good of interest (burgers, bread, bananas, beers, beans, beds, bras, bugles, boomerangs, bungee cords, banjo lessons, bus tickets, etc.) An individual's demand function is a mapping from prices into quantities purchased. Given an individual's marginal benefit curve, one can determine this mapping by choosing a price, finding the quantity at which the implied marginal cost curve equals his or her marginal benefits, and plotting that price and quantity on another diagram. This process is illustrated for three different prices in Figure 2.2. To generate a complete mapping would require this to be done for all possible prices.

Figure 2.2: A Consumer's Demand Curve

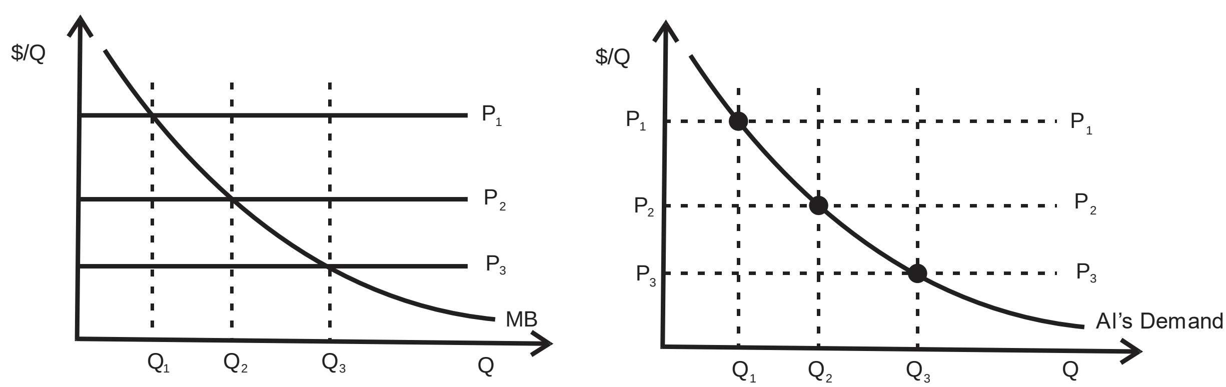

Recall that when consumer is a "price taker" he or she has no bargaining power and adjusts his or her purchases to the market price and the market price is his or her marginal cost for the good of interest. An implication of figure 2.1 in combination with the definition of a demand curve is that when prices fall, as from P_1 to P_2, the consumer will purchase additional units, $Q_2 > Q_1$. This implies that demand curves and their associated functions are "downward sloping" whenever an individual's marginal benefit curve is downward sloping or equivalently monotonically decreasing in the quantity purchased.

Notice also, that in such cases, an individual's demand curve goes through exactly the same points. (P_1, Q_1), (P_2,Q_2), and (P_3,Q_3) are all points on both the marginal benefit curve and the individual's associated demand curve. This is true of all the other points on the demand and marginal benefit curves as well. However, the marginal benefit function and the demand functions are different. Marginal benefits are a mapping from quantities into benefits per unit. Demand is mapping from prices into quantities. They are inverse functions for one another.

This basic model of consumer behavior can easily be extended to other walks of life. For example, a person that is running for an elective office in a democracy may be modeled as a "net vote maximizer." Candidates may take policy positions that maximize their net votes—realizing that as they shift among positions, they will lose some votes (their marginal cost in terms of votes) but gain others (their marginal benefits in terms of votes). In setting where risk or time are important, the persons may be modeled as maximizing "expected" net benefits or the "present value" of net-benefit flows. Cost-Benefit analysis is grounded in such extended models of net-benefit maximizing choice.

III. Concrete Functions, Net Benefit Maximization, and Consumer Demand

Geometric derivations of demand and other relationships are not as dependent on mathematically convenient assumptions as calculus-based derivations, but calculus derivations are often shorter, and often provide clearer, if less general, results than geometry allows. They also allow more considerations to be simultaneously taken into account in one's models, as with exercises in comparative statics and proofs of the prices exist that will generate a general equilibrium in any countable number of markets. Calculus-based derivations also demonstrates that the mathematical methods used in the "hard" sciences and engineering can also be applied in the social sciences—although it should be admitted, often with less realistic models and implications. Such derivations and those based on other tools from real analysis and topology also appear more rigorous, although it should be acknowledged that geometric proofs are no less rigorous than those associated with more advanced mathematics.

In the hard sciences, greater rigor is often associated with more predictive power and precision, although that is not always true in the social sciences. There are more factors that tend to influence individual decisions and social outcomes than which determine the orbit of the moon around the Earth or of the Earth around the Sun. Newtonian mechanics and gravity do not provide as much insight into the paths of creatures that are self-propelled and make their own independent decisions—although such physical laws are not entirely irrelevant. No human can leap over a tall building with a single bound. On the other hand, no rock, however large or small, makes a decision about when to fall or where to land.

To illustrate how a net benefit maximizing choice can be modelled using functions and calculus, suppose that Al's benefit function for the good of interest is $B= 5Q^{.8}$. If Al can buy as much of Q as he or she wishes at price P, than Al's net benefit function can be written as

$$N = B - C = 5Q^{.8} - PQ \qquad (2.1)$$

Its first derivative with respect to Q is:

$$dN/dQ = 4Q^{-.2} - P \qquad (2.2)$$

The first term ($4Q^{-.2}$) is Al's marginal benefit, and the second term (P) is Al's marginal cost.

The second derivative of the net benefit function is:

$$dN^2/dQ^2 = -.8Q^{-1.2} \qquad (2.3)$$

Notice that this is the slope of Al's marginal benefit curve, which is downward sloping for this particular function for any $Q>0$. The marginal cost curve, in contrast, is not in this case affected by quantity, and so is a horizontal line of slope 0.

Calculus implies that **a function is strictly concave if** its second derivative is always less than zero in the domain of interest. Such functions will satisfy the geometric definition—a cord connecting any to points on the function will lie below it. Such functions have at most one maximum value. (They will not have a maximum value if they are monotone increasing, as the benefit function is in this case.) Calculus also implies that a function is at a maximum or minimum when ever the first derivative of the function of interest has the value 0. It is at a maximum if the second derivative is negative at that point.

Thus, if there is a quantity in the positive domain at which the derivative of the net benefit function is zero, then there is a net-benefit maximizing quantity of this good—and as a net-benefit maximizer, it will be the quantity that Al purchases. To see if such a point exists, set the derivative of N equal to zero and determine if there is a Q with this property.

$$0 = 4Q^{-.2} - P \qquad (2.4)$$

A bit of algebra shows:

$$4Q^{-.2} = P \qquad (2.5)$$

$$Q = (P/4)^{-5} \qquad (2.6)$$

(Recall that raising a variable (such as X) with an exponent (such as a) to some power (such as b), has the property that $(X^a)^b = X^{ab}$).

Equation 2.6 shows that there will be a quantity that maximizes net benefits for any price P that might exist. It describes what that quantity will be for every P. Since these are that quantities that Al would purchase at those prices, equation 2.6 is Al's demand function for the good of interest—potato chips, burgers, bicycles, apples, plane tickets, etc.. Equation 2.5 is also of interest; it shows that at the net benefit maximizing quantity—if one exists—then it will be at a point (quantity) where marginal benefits ($4Q^{-.2}$) equal marginal cost (P). The quantity that solves the "first order condition," (sets the first derivative of an objective function equal to zero) is often denoted with an asterisk (*), to indicate an ideal value, as with $Q^* = (P/4)^{-5}$. Note also that this demand function is monotone decreasing in price. Note that $dQ^*/dP = -5(P/4)^{-6}$ which is less than 0 for all positive prices).

Thus, if we know the specific function that describes an individual's benefit function, we can derive his or her demand function—which he or she should follow if a net benefit maximizer. However, we may not know the specific benefit function, but rather believe that it belongs to a particular family of concrete functions, not necessarily $B=5Q^{.8}$, but belong to the family of strictly concave exponential functions, $B = aQ^b$ with $a> 0$ and $0<b<1$. (If "a" were not greater than 0, it would not be a good, and if b were not between zero and one, the benefit function would not be strictly concave.)

To drive a more abstract version of Al's demand curve if his or her benefit function belongs to this family of functions, we can repeat the steps above.

$$N = B - C = aQ^b - PQ \tag{2.7}$$

Its first derivative with respect to Q is:

$$dN/dQ = abQ^{b-1} - P \tag{2.8}$$

The first term (abQ^{b-1}) is Al's marginal benefit, and the second term (P) is Al's marginal cost.

The second derivative of the net benefit function is:

$$dN^2/dQ^2 = ab(b-1)Q^{b-2} \tag{2.9}$$

The second derivative is negative ($b-1<0$, $b>0$, $a>0$, and $Q>0$) and so this entire family of net benefit functions is strictly concave.

If there is a value of Q that sets equation 2.8 equal to zero, then every benefit function from this strictly concave exponential family of functions will have an ideal or optimal quantity of the good of interest that maximizes net benefits for every possible price greater than zero. To determine this set equation 2.8 equal to zero and attempt to solve for Q.

$$abQ^{b-1} - P = 0 \tag{2.10}$$

A couple of algebraic steps are sufficient to solve for Q.

$$abQ^{b-1} = P \tag{2.11}$$

$$Q^* = (P/ab)^{1/b-1} \tag{2.12}$$

Such a Q does exist for this family of functions. Note that it is always a positive quantity and that according to equation 2.11 it always occurs where marginal benefit equals marginal cost. Equation 2.12 characterizes Al's demand function for all the possible combinations of a and b that characterize the benefit functions in the assumed family of functions.

The above results for exponential benefit functions are useful for studies that attempt to estimate demand functions, because the demand function is "log linear." If you take the log of both sides the equation 2.12, the result resembles the equation of a straight line, $\log(Q^*) = (1/b-1) \log(P/ab) = (-1/b-1) \log(ab) + (1/b-1) \log(P)$. The first term is a constant, $[(-1/b-1) \log(ab)]$ and the second is a constant times the log of price $[(1/b-1) \log(P)]$. Thus, Al's particular demand curve for the good of interest can be estimated in logs using the conventional linear methods from econometrics given a data set for prices and purchases by Al or a group of consumers with similar benefit functions.

IV. Net Benefit Models with More General Families of Functions

The problem with results obtained from models grounded on explicit functional forms is that they are often special cases. The function families focused on tend to have properties that can be easily worked out using algebra and calculus—which is to say they are mathematically very tractable functions. Of course, it is the tractability that is their main attraction because this allows one to focus most of one's attention on key features of the choice setting of interest rather than on the mathematics of the models developed. They same functions often have implications about the algebraic structure of demand and other functions of interest that are fairly easy to estimate with well-understood econometric methods.

Models grounded in more general families of functions may also be mathematically tractable, although the results tend to be abstract and provide less guidance about how to estimate the model developed or a subset of the key relationships worked out. However, they are also less likely to have odd properties that seem unlikely to characterize all choices in the circumstances being analyzed. For example, it would be surprising if all demand curves were linear in logs as implied by the model developed in the previous section—although this could be a useful first approximation for many demand curves. Moreover, many economists are more interested in qualitative relationships—like demand curves slope downward—than in estimates of particular demand function, because they may help explain a broad range of phenomena. If a general analysis implies that most demand curves slope downward, then one can use that result for an informal analysis of most markets, without having to estimate demand curves for each the markets of interest. Often qualitative—intuitive—analysis is sufficient to understand, for example, roughly how a particular type of public policy is likely to affect a particular market or network of linked markets.

A general model of net benefit choices by consumer in markets in which the are price takers is easy to develop. Let b(Q) be the individual of interest's total benefit function and C=PQ be his or her total cost function. Assume that function b is strictly concave (e.g. has positive first derivatives and negative second derivatives) in the domain of interest ($Q \geq 0$). Net benefits can be written as:

$$N = b(Q) - PQ \tag{2.13}$$

Differentiating with respect to Q and Net benefits are maximized when:

$$dB/dQ - P = 0 \tag{2.14}$$

(Note that $dN^2/dQ^2 = dB^2/dQ^2 < 0$, since B is assumed to be strictly concave. So, if a quanty Q exists that satisfies equation 2.14, it is the one that maximizes net benefits.) The first term, as in the previous cases, is the marginal benefit of the good or service and the second is its marginal cost.

The **implicit function theorem** states that given a differentiable function h(a, b, c, d) = 0, then a function h exists such that a = f(b, c, d). In other words, given a first order condition (as in equation 2.14), it is possible to develop a function that describes how any one of the variables in the first order condition is affected by the others in the first order condition. There are just two variables in the first order condition, Q and P, so a function f exists that can describe how Q is affected by P.

$$Q = f(P) \tag{2.15}$$

That function is Al's demand function for the product or service of interest—the one for which we have assumed a strictly concave benefit function.

The implicit function differentiation rule can be used when h(a, b, c, d) = 0, and the implicit function theorem is used to characterize "a" as a function of the other variables in function h, a = f(b, c, d). In this case, the derivative of function f with respect to variable b is df/db = dh/db / (- dh/da). (This rule can be derived by finding the total derive of function h, assuming that only variables a and b change, combined with a bit of algebra.) In the case of the function f described by equation 2.15, equation h is that characterized by equation 2.14.

So the slope of the demand curve is just:

$$df/dP = dh/dP \,/\, (-dh/dQ) = -1/-(db^2/dQ^2) < 0 \qquad (2.15)$$

So, it turns out that the geometry of figure 2.2 is quite general. As long as the marginal benefit curve is downward sloping (e.g. the benefit function is strictly concave), then demand curves derived from the net-benefit maximizing model of consumer choice are always downward sloping.

V. Utility Functions and Choosing a Bundle of Goods

The net-benefit maximizing model of rational consumers has clear implications and, arguably, is the basis for much of an economist's intuition about how consumer demand operates. However, there are choice settings in which several goods are chosen simultaneously and there are also occasions when one's budget constrains what one can afford to purchase. Neither are obviously the case in the net-benefit maximizing models worked out above.

Another—and more widely used—model of rational consumer behavior assumes that individuals have unified objective normally referred to as "utility," a concept introduced by utilitarian philosophers such as Jeremy Bentham in the late eighteenth century. Initially, utility was interpreted as a synonym for the (net) happiness or satisfaction associated with a particular pattern of behavior or outcomes. Utilitarians were and continue to be more interested in appraising the merits of alternative public policies and institutions than in describing consumer behavior. From the utilitarian perspective, one policy is better than another if it increases the sum of utility in a community or the average level of utility in a community more than another. Insofar as these included economic policies, various ideas about market networks were developed by utilitarians and several utilitarians wrote economic textbooks including James Mill (1821), John Stuart Mill (1848), and William Stanly Jevons (1871), who developed the idea of a utility function. Many of the most remembered "economists" from the late nineteenth century were utilitarians. The idea of an indifference curve was worked out by Edgeworth (1881) although not drawn until Pareto (1906).

When applied to consumer choice, an individual's utility is normally characterized as a monotonically increasing function of the quantities of goods and services that are at an individual's disposal, or the quantity consumed—although not always literally consumed. One does not eat an automobile but may gradually use it up through wear and tear, although most folks "trade their cars in" well before they wear out. A typical general utility function for individual i is written as $U_i = u_i(A_i, B_i, C_i, \ldots)$, where $A_i, B_i, C_i, \ldots$ are quantities of goods and services consumed and the first partial derivatives of u_1 are all greater than zero and the second partial derivatives are negative, so that marginal utility curves are all downward sloping.

However, rather than skip to the general, we'll again start with geometry. The geometry that can be drawn clearly on a sheet of paper is limited to two-dimensions. To do so for a choice among various combinations of two goods, normally one draws a series of indifference curves. An indifference curve plots the various combinations of the two goods that generate exactly the same utility. In effect, they characterize a topological map of a utility function. (Hikers may be familiar with topological maps if they've climbed mountains and wanted to know how far above sea level they were along the way to the top.) The "utility mountain" is the missing third dimension and readers can imagine that "mountain" as rising out of the page toward their nose, where the height of the mountain is in "utils" rather than feet or meters above sea level. Although conceptually indifference curves can take any shape that a mountain can, the usual assumption about utility functions limits them to a series of more or less c-shaped curves that represent higher utility levels as one moves to the north west (upper right) of the diagram.

Figure 2.3 depicts a typical consumer's choice (Al's) between goods X and Y for the case in which the price of good X is P_x and the price of good Y is P_y and the consumer has amount W to spend on the two goods in the period of interest. Given those assumptions about prices and the total amount to be spent on the two goods implies that Al has a budget constrain that can be represented algebraically as $W = P_xX + P_yY$, which appears in figure 2.3 as a diagonal line running from W/P_y on the vertical axis to W/P_x on the horizontal axis. That line is sometime called a budget line, but more often referred to as Al's budget constraint. The end points characterize the maximum quantity of good X or Y that could be purchased if Al spends all of his or her money on just one of the goods. The slope of the budget line is determined by the relative prices of the two goods, which in this case is $-P_x/P_y$.

Figure 2.3: The Geometry of a Utility Maximizing Choice

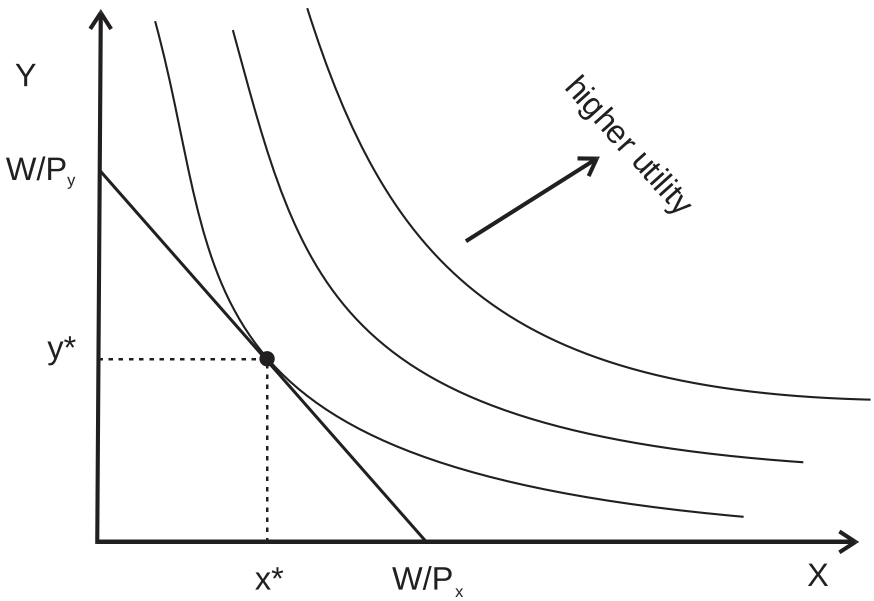

The highest indifference curve that can be reached is one that is tangent to the budget line. At that combination of goods X and Y, Al has achieved the highest utility that is feasible given the amount that he or she has to spend, W.

A demand curve remains, as above, a mapping of prices into quantities purchased, but the derivation is a bit different than in the net benefit maximizing case. To derive a demand curve, one of the prices is varied and the other is held constant as is the consumers budget (W). To illustrate the process, the demand for good X is derived below. To do so, various prices for X are tried and the price and quantity of good X purchased at the price are plotted in the diagram to the right to trace out Al's demand curve for good X. In each case, Al purchased the quantities of X and Y that maximize his or her utility, given the "new' price for good X.

Figure 2.4 illustrates how this process operates geometrically. Three prices good X are selected with P_1 the highest and P_3 the lowest. This is evident in that the highest quantity that Al could purchase of good X rises as prices fall from P_1 to P_2 to P_3. Every time price changes, Al changes his "bundle" of X and Y, because his budget line (opportunity set) changes. Note that both X and Y change as the price of X changes, not simply X as might be suggested by a straight-forward net benefit maximizing model. The result illustrated is "qualitative" because no specific values for W, P_Y and P_X have been specified, nor a specific function for the utility function. If those had been specified, particular values for X and Y would have emerged from a careful diagram of the possible indifference curves that might be reached, and the budget lines associated with different prices.

Figure 2.4: Indifference Curves and Demand

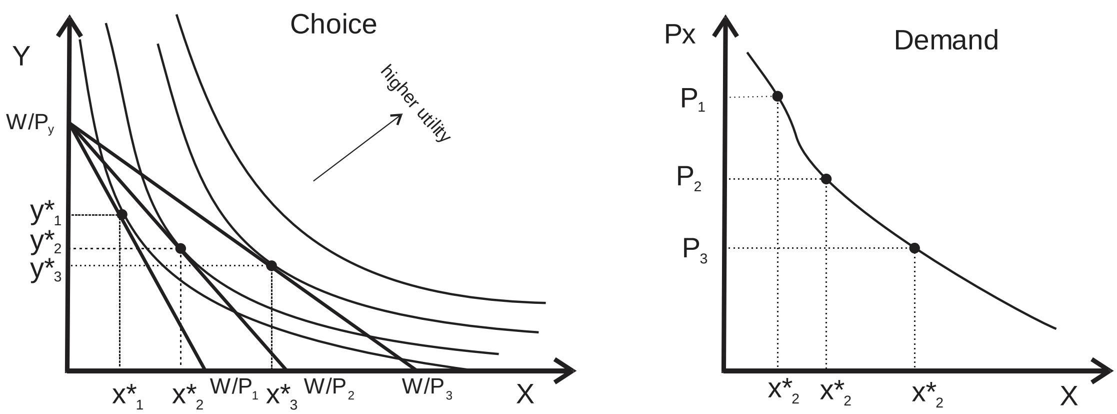

In the case depicted, demand is downward sloping (monotone decreasing) although this is not always the case with indifference-curve based analysis. It is, however, by far the most common type of demand curve or function found in empirical research. However, it bears noting that a demand curve is not necessarily linear or smoothly downward sloping as often assumed in empirical work.

VI. Cobb-Douglas Utility Functions and Consumer Demand

Mathematical models of utility maximizing choices can be more general than the two dimensional one illustrated in the previous section, but for two-good models are sufficient to illustrate the most common methods for deriving a demand curve using calculus rather than geometry. The work horse of this section of chapter 2 will be generalizations of the Cobb-Douglas (1928) family of functions, which can be termed a multiplicative-exponential functions. The basic algebraic structure for the three-good version of a utility function defined over goods X, Y, and Z is $U = eX^aY^bZ^c$ with a, b, c > 0. The Cobb-Douglas function is the special case of that family of functions, where a+b+c = 1. Cobb-Douglas functions exhibit constant returns. If one multiplies the quantity of each good by factor Δ, then utility also increases Δ-fold. The other possible relationships among exponents are also of interest. For example, if $0< a+b+c \leq 1$, the utility function may exhibit constant or diminishing returns. With $0< a+b+c < 1$, the utility necessarily exhibits diminishing marginal returns and the utility function will be strictly concave, whereas if $0< a+b+c \leq 1$, the utility function is concave rather than strictly concave.

A Short Digression on Mathematical Methods

However, to derive demand curves using even relatively straightforward generalizations of the exponential model of net-benefit based demand developed above, requires somewhat more sophistical tools from calculus to be employed, namely partial derivatives and the Lagrangian method. A **partial derivative** characterized the effect of a change in one variable in a multi-variate function on the function value, holding all the other variables constant. For example, the partial derivative of $U = eX^aY^bZ^c$ with respect to X is just $dU/DY = aeX^{a-1}Y^bZ^c$. In effect, Y^bZ^c is treated as just another constant. Similarly, $dU/dY = beX^aY^{b-1}Z^c$, and $dU/dZ = ceX^aY^bZ^{c-1}$. A function with three variable factors (here, the quantities of goods X, Y, and Z) has three partial derivatives—one for each of the variable factors that influence the value of the function.

The **Lagrange method** was developed in 1764 by Joseph Louis Lagrange, one of many mathematicians that can be regarded as a genius. The method is relatively straight forward. Given a strictly concave differentiable objective functions, such as U=u(X, Y, Z), and a constraint in a form equal to zero, such as 0=W-PxX-PyY-PzZ, the combination of X, Y, and Z that maximizes U can be characterized as follows. First, form the "Lagrangian function" which consists of the objective function and the constraint function multiplied by parameter λ, as with $\mathcal{L} = u(X, Y, Z) + \lambda[W-PxX-PyY-PzZ]$. Second, take the partial derivatives of the Lagrangian function with respect to the variables that determine the values of the objective function (the control variables) and also the partial derivative with respect to the Lagrangian multiplier, λ. Third, set all of the partial derivatives equal to zero. The resulting system of equations describes the optimal values of the "control variables" (here, X, Y, and Z) given the constraint that must be satisfied (here the budget constraint). In some cases, one will be able to solve for explicit functions that characterize the ideal levels of each of the control variables—given the constraint. It turns out that the family of multiplicative exponential functions is one such case.

Deriving Demand Curves from Multiplicative Exponential Utility Functions

With these two mathematical methods in mind, we are now in position to derive a demand curve from a utility function in algebraic form. A two good utility function is sufficient to demonstrate the technique and some of the main results, although utility functions can include any number of goods and services. Let Al's utility function be $U=aX^bY^c$ and his or her budget constraint be $0 = W\text{-}PxX\text{-}PyY$. Given these assumptions, the associated Lagrangian function is:

$$\mathcal{L} = aX^bY^c + \lambda[W\text{-}PxX\text{-}PyY] \qquad (2.16)$$

Differentiating with respect to X, Y, and λ, then setting the results equal to zero, generates the following family of equations that characterize Al's ideal levels of X and Y

$$d\mathcal{L}/dX = abX^{b\text{-}1}Y^c - \lambda Px = 0 \qquad (2.17a)$$

$$d\mathcal{L}/dY = acXbUc\text{-}1 - \lambda Py = 0 \qquad (2.17b)$$

$$d\mathcal{L}/d\lambda = W\text{-}PxX\text{-}PyY = 0 \qquad (2.17c)$$

Note that the first term in equations 2.17a and 2.17b are the marginal utility functions for goods X and Y respectively. he marginal utility of X is $abX^{b\text{-}1}Y^c$. It is partly determined by its own consumption level (X) and partly by the consumption level of Y. The more of good Y that is purchased and consumed, the greater is the marginal utility of X. The same is true for good Y. Goods are, in this sense, gross complements in utility functions based on multiplicative exponential functions, including the Cobb-Douglas function. Note also that the last term in both equations is lambda times the marginal cost of those goods, which in this case is their price.

A bit of algebra allows both demand curves to be derived and also some interesting results to be worked out. The following steps are often used when working with this family of functions and constraints. First, add λPx to both sides of eq 2.17a and add λPy to both sides of equation 2.17b. Second, divide the lefthand side of first resulting equation by the lefthand term, and the righthand term of the first equation by the righthand term of the other to obtain:

$$\frac{abX^{b-1}Y^{c}}{acX^{b}Y^{c-1}} = \frac{\lambda \mathrm{Px}}{\lambda \mathrm{Py}} \tag{2.18a}$$

Note that the numerator and denominator on the left include many of the same terms and can be divided out. A bit of algebra allows equation 2.18a to be simplified to:

$$\frac{by}{cX} = \frac{\mathrm{Px}}{\mathrm{Py}} \tag{2.18b}$$

This is the tangency condition between the indifference curves in the XY plane and the budget constraint.

Third, in order to find the demand function for X, we need to isolate X by solving equation 2.18b for X.

$$X = \frac{bPy}{cPx}Y \tag{2.18c}$$

Fourth, use the constraint (equation 2.17c) to characterize Y in terms of X, which in this case is: Y=[W-PxX]/Py. Substitute that relationship into equation 2.18c for Y and simplify:

$$X = \frac{bPy[W-PxX]}{cPxPy} = \frac{bW}{cPx} - \frac{bPxX}{cPx} \tag{2.18d}$$

Fifth, add the positive of term with X in it on the right (the last term) to both sides.

$$X + \frac{bPxX}{cPx} = \frac{bPyW}{cPx} \tag{2.18e}$$

Isolate X and simplify:

$$X\left[1 + \frac{bPx}{cPx}\right] = \frac{bPyW}{cPx} \xrightarrow{yields} X\left[\frac{cPx+bPx}{cPx}\right] = X[\frac{c+b}{c}] = \frac{bW}{cPx} \tag{2.18f}$$

Cross-multiply to isolate X on the lefthand side, then simplify.

$$X = \frac{bW}{cPx}\left[\frac{c}{c+b}\right] \xrightarrow{yields} X^* = \frac{bW}{(c+b)Px} \tag{2.19}$$

A similar series of steps yields:

$$Y^* = \frac{cW}{(c+b)Py} \tag{2.20}$$

Both demand curves are remarkably simple for this family of functions. Al spends fraction b/(b+c) on good X and fraction c/(b+c) on good Y. To find the quantity purchased, just divide those amounts by the relevant prices. In the Cobb-Douglas case, b+c = 1, the demand curves are even simpler X* = bW/Px and Y* = (1-b)W/Py.[3]

Note that these demand curves are all downward sloping. Note also that each spends amounts on the two goods that are proportional to the exponents of the utility function. The demands for each good are thus independent of one another's prices. The latter also implies that when income (W) increases, expenditures increase proportionately. Thus, the income expansion paths for persons with utility functions from the multiplicative exponential family of functions are a straight line, with a slope that varies with the relative size of the exponents.

All these are rather "special" properties, but they allow one to use fairly straightforward techniques from econometrics to estimate the relevant demand and Engel's functions. One can estimate demands for single goods without paying much attention to the prices of other goods. It turns out that for many goods, the implications of this family of utility functions provide reasonably good approximations of both the demand for those goods and the relationship between an average consumer's income and his or her allocation of income of those goods.

However, this is not true for all goods. There are cases in which the prices of other goods—substitutes and complements—clearly affect both demand and expenditures. And there are goods for which the relationship between income and expenditures is non-linear—inferior

---

[3] Students should now repeat the derivation on their own—without using the text, except to help overcome roadblocks. Once, the simplicity of the solution is known, most student with a bit of algebraic skill and intuition will be able to derive this result through similar steps. The text simply provides one of the many series of steps that will generate the surprisingly simple algebraic expression for the demand functions associated with this family of multiplicative exponential utility functions.

and superior goods. More general families of utility functions imply demand functions that reflect such possibilities as well as the simpler ones implied by multiplicative exponential functions such as the Cobb-Douglas utility function.

VII. General Utility Functions and Demand Functions

A Digression on the Limited Domain of Utility Functions Imposed by Calculus

As in the case of net benefit maximizing models of the demand for goods and services (and most other things), utility-based analysis can be based on functions from a very broad family of functions with particular shapes. There is, however, a mathematical constraint on the types of functions that can potentially characterize utility if one wants to derive a demand function from them. A demand function is a one-to-one relationship between prices and amounts purchased. Every price induces a unique quantity to be purchased by the consumer of interest. Thus, for the purpose of characterizing that relationship, it is important that there be a unique utility-maximizing choice of bundles of goods associated with every vector of prices. Strictly concave utility functions have this property when the opportunity set is convex, which is true of the budget line faced by consumers who are price takers (e.g., purchasers adapt to prices over which they have no control).[4] Thus, it is normally assumed that abstract utility functions are strictly concave functions. In order to use calculus to derive demand functions it is also necessary that the utility functions be differentiable—ideally at least twice differentiable.[5]

Deriving a Demand Function from an Abstract Strictly Concave Utility Functions

A general abstract utility function has the form U = u(A, B, C, … X, Y) where the arguments of the function are all quantities of a good or service, partial derivatives with respect

[4] A function is a mapping from a set into a point. Thus, mapping from one set into another set (e.g. several points) would not be a function. It would be a "correspondence," rather than a function, although a function is a special case of a correspondence. Price theory can be worked out for correspondences as well as functions, as was shown in Debreu (xxxx) among many other works in general equilibrium theory. A convex set is a set in which any line segment connecting two points within the set is entirely contained in the set of interest.

[5] It is interesting to note that none of these assumptions are necessary for geometric models—although such models cannot be depicted for more than two goods. Thus, for some purposes, geometric models are easier to construct and analyze than calculus-based models.

to the quantity of each good is greater than zero (otherwise the arguments would not be goods), and their second derivatives are negative (which characterizes diminishing marginal utility and also contributes to the strict concavity of the utility functions). The associated price vector is (Pa, Pb, Pc, … Px, Py) and its budget constraint is W = PaA + PbB+ PcC … + PxX, +PyY). One can characterize a consumers choice using the Lagrange method of the previous section. However, no neat solutions will be possible, and determining the effects of prices and income on the quantities demanded require the use of matrix methods. Both the methods and results of such analyses tend to be "messy" and cumbersome, and the associated mathematics shed little new light on the core properties of demand relationships beyond those associated with two good versions of general models. They do, however, show that it is possible to model such very general choice settings with calculus-based mathematical models. (For an overview of the required matrix methods, consult a mathematical economics textbook on bordered Hessian matrices and Taylor's matrix rule.)

A general abstract utility function for two-good choice settings has the form U=u(X, Y) with $\delta U/\delta X > 0$ and $\delta U^2/\delta Y^2 < 0$ and $\delta U/\delta Y > 0$ and $dU^2/dY^2<0$. The associated budget line is the same as in the previous section, W = PxX + PyY. The best combination of X and Y to purchase can again be characterized with the Lagrange method, which again will require the use of matrix methods to undertake comparative statics (analysis of how changes in prices or income affect demand).

However, in the two good case, there is another method that is mathematically tractable and less cumbersome than the usual matrix methods that I call the substation method. If one is interested in the demand for X, first solve the budget line for Y as a function of X, Y = (W-PxX)/Py along the upper bound of the utility function. Since X and Y are both goods, it is along this line that the optimal combinations will be found. Substitute this equation for Y into the utility function to create a composite utility function that evaluates utility levels along the budget constraint.

$$U = u(X, W\text{-}PxX)/Py) \qquad (2.21)$$

To characterize the ideal level of X for given prices and income (or wealth), differentiate 2.21 with respect to X and set the result equal to zero. Because Y is now a characterized as a function of X, as X changes so will Y. The composite differentiation rule, thus, has to be applied to capture that effect.

$$dU/dX = \delta U/\delta X + (\delta U/\delta Y)\ (-Px/Py) = 0 \equiv H \text{ at } X^* \qquad (2.22)$$

Notice that the first term ($\delta U/\delta X$) is the subjective marginal cost of good X and the second term, ($\delta U/\delta Y$) (Px), is the subjective marginal cost of X from the associated reduction in the consumption of good Y associated with any increase in X—holding prices and income constant. These marginal costs are in terms of "utils" rather than dollars, but the relationship implied is otherwise similar to the relationship characterized by Figure 2.1—although because of diminishing marginal utility, the marginal cost curve in this case will be upward sloping rather than horizontal. This is the marginal opportunity cost of consuming more X.

Although this first order condition is abstract, it is possible to characterize the slope of the demand curve and the effect of changes in the price of Y or income on demand. Recall that the **implicit function theorem** states that given a function h(A, B, C) = 0, there exists a function the characterizes the relationship of each variable in terms of the others, as with A = f(B, C). Note that the first order condition (equation 2.22) is an instance of function h. Thus, there exists a function f that characterizes purchases of X* as a function of the other parameters of Al's purchasing decision.

$$X^* = f(Px, Py, W) \qquad (2.23)$$

That function is, of course, Al's demand function for X.

Recall that the utility function was assumed to be twice differentiable, its first order condition is also differentiable, and so the implicit function differentiation rule can be applied. The implicit function differentiation rule states that the new function (here, f) has derivatives that can be characterized with derivatives of the original function (h) as long the original function was differentiable. The utility function was assumed to be twice differentiable (e.g. to have second derivatives), so, the first order condition characterized by equation 2.22 is also differentiable. In the illustrative implicit function of the previous paragraph, the derivative of A with respect to B is dA/dB = δh/dA/-δh/δA.

The slope of Al's demand function for X is somewhat complex. To characterize that slope, it is important to remember that each of the marginal utility functions in equation 2.22 includes the same arguments as the original (parent) function. Each has two arguments and X is in both. This implies that the derivatives tend to be more complex than one might guess based on the notation used in equation 2.23.

It is:

$$\frac{dX^*}{dPx} = \frac{\frac{dH}{dPx}}{-\frac{dH}{dX}} = \frac{\frac{dU^2}{dXdY}\left(\frac{-X}{Py}\right) + \frac{dU^2}{dY^2}\left(\frac{X}{Py}\right)\frac{Px}{Py} - \frac{dU}{dY}\frac{1}{Py}}{-[\frac{dU^2}{dX^2} + \frac{dU^2}{dXdY}\left(-\frac{Px}{Py}\right) + \frac{dU^2}{dYdX}\left(-\frac{Px}{Py}\right) + \frac{dU^2}{dY^2}(\frac{Px}{Py})^2]}$$

(2.24)

This equation is, of course rather abstract and complex, but each of the component terms has a sign (is greater or less than zero) and in the case where the cross partials are greater than zero—e.g., $(\frac{dU^2}{dXdY} > 0$ –this expression has clear sign. It is less than zero, thus Al's demand curve for X is downward sloping, given the assumption about cross-partials.

The Importance of Cross Partials

To see this, first, notice that the term inside the brackets in the denominator is simply the second derivative of the utility function with respect to X and has to be negative is the utility function is strictly concave. Both second derivatives are negative (there is diminishing marginal utility, by assumption), and the two cross partials both have negative post-multipliers and so are overall negative terms. So, together with the assumptions about diminishing marginal utility, the assumption of positive cross partials is sufficient to assure concavity. The negative sign in front of the bracket, thus, assures that the denominator has a positive sign. The sign of the slope of A's demand curve is thus determined by the numerator. Notice that each term in the numerator has a value less than zero. The first term is a positive (the cross partial) times a negative. The second term is a negative (the second derivative of U with respect to Y) times a positive and so is also negative. The third term is a positive (marginal utility is greater than zero for all goods) times a positive (a ratio of two prices) but has a negative sign in front of it and so is also negative overall. Thus, strict concavity of the utility function together with the assumption of positive cross partials (gross complementarity, as in the multiplicative exponential utility functions in the previous section) are **sufficient to assure that demand curves are downward sloping**.

However, if cross partial's are negative, (if the two goods are gross substitutes) it is possible that the opposite result may occur, even if the utility function is strictly concave, although for most purposes such possibilities can be ignored because empirically essentially all demand curves have been downward sloping in prices—with the exception of cases in which price is used by a consumer as an index of quality. In Part I of the text, we are assuming uniform and known quality for both goods X and Y. This is the usual "full" or "perfect" information assumption of perfectly competitive markets.

On the Importance of Demand Functions

The theory of demand is one of more important departures from classical economics, which stressed the linkages between prices and production costs, especially in the long run. Those connections were not discarded by neoclassical economics but came to be regarded as only part of the process through which prices are determined in both the long and the short run, as demonstrated in chapter 3.

The theory of consumer demand for products produced and sold in markets, however, is only one application of this very general model of human behavior. Very similar models of human decision making can be developed for household production, for investment and savings decisions, for political choices, and for decisions about war and peace, and even with respect to the development and use of ethical and normative theories on one's own life. It has proven to be the most generalizable of the essential models and results from neoclassical economics. Economics and the "rational choice" strands of other social sciences use "methodological individualism" as the conceptual foundation for their theories and models. This means that social phenomena such as markets and political systems are regarded to be joint consequences of individual decisions and actions. It is the model of demand and extensions of it that serve as the work horse of such approaches to large scale social phenomena.

VIII. Appendix to Chapter 2: Related Technical Mathematical Terminology

A. **A function** is a mapping from one set (often Q or quantity in economics) into another set (such as net benefits, costs, benefits, utility, profits, revenue, etc.) Many of the mathematical properties of a given function can be deduced from its "shape."

B. One of the most widely used characterizations of a function's shape in economics is *concavity*. There are **three notions** of concavity used in economics, although in this course, only the first and second are used outside of this chapter.

i. DEF: **Strictly Concave**: function f is *strictly* concave iff

$$\alpha f(X_1) + (1-\alpha)f(X_2) < f(\alpha X_1 + (1-\alpha)X_2) \quad \text{where } 0 < \alpha < 1.$$

- Geometrically this means a function is strictly concave "if and only if" all the points on a line segment connecting any two points on a function always lies "beneath" the function of interest.
- **Strict Concavity is the assumption about functions that is most often used in this course.**
- A strictly concave function has at most one maximum, which allows us to characterize choices that are very specific—whether the choices are unconstrained or constrained by some convex set.

ii. DEF: Concavity: function f is concave iff

$$\alpha f(X_1) + (1-\alpha)f(X_2) \leq f(\alpha X_1 + (1-\alpha)X_2) \quad \text{where } 0 < \alpha < 1.$$

- Geometrically this means that points on a cord (line segment) connecting any two points of a function always lies on or beneath the function of interest.
- A straight line is concave, but not strictly concave.
- Other strictly concave function are also concave.

iii. DEF: Quasi-Concave: Concavity: function f is *quasi* concave iff

$f(X_1) < f(\alpha X_1 + (1-\alpha)X_2))$ where: and $f(X_1) < f(X_2)$ and $0 < \alpha < 1$.

- The values of a quasi-concave function always lies above the lower of the two end points of a cord connecting any two points on the function.
- Any monotone increasing function is quasi-concave, but it is not necessarily concave or strictly concave, because it may increase at an increasing rate.

Figure 2.5: Concavity

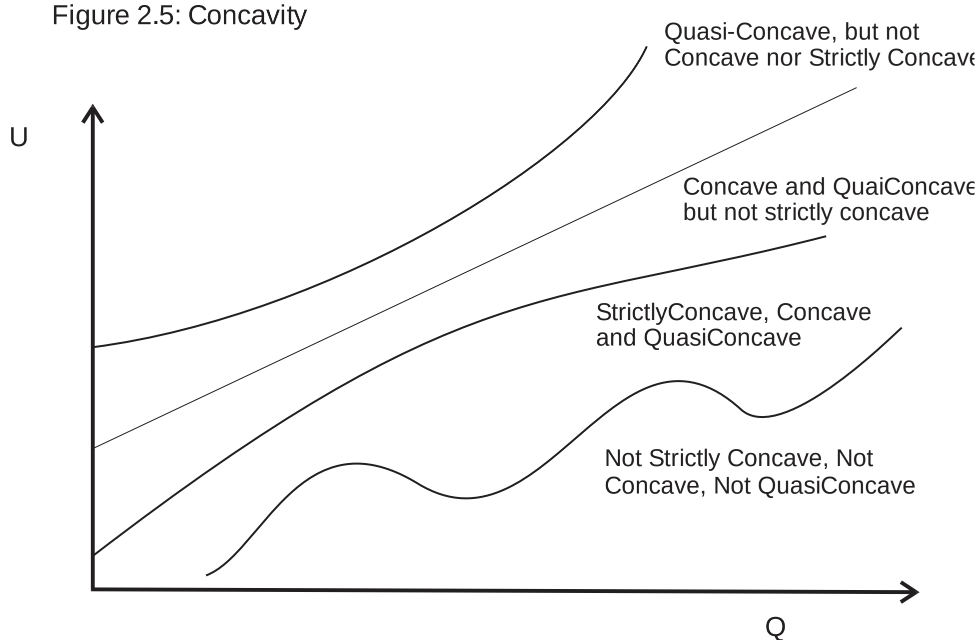

Maxima and Minima of Functions

A. Strictly concave functions have a number of useful properties in the context of "optimizing" behavior.

i. A *strictly* concave function has at most one maximum. (Draw some pictures to see why.)

ii. However, a **concave function** may have an infinite number of global maxima, but if there is more than one maximum, they make up a continuous linear interval. (A horizontal line is concave, but not strictly concave.)

B. DEF: The **global maximum** of a function, f(x), is a value, f(x*), that exceeds all others over the entire range of the function (e. g. for every neighborhood of x*).

C. DEF: A **local maximum** of a function, f(x), has a value which exceeds those of other points within a finite neighborhood of x*. That is, f(x*) is a local maximum if f(x*+e) < f(x*) and f(x*-e) < f(x*) for 0<e<E, for some E>0.

- Note that if a function has a global maximum, then that global maximum is also a local maximum.
- However, because a function may have many local maxima, only one of those can be a global maximum.
- Derivatives of functions can be used to characterize *sufficient* conditions for concavity, strict concavity, and therefore also for global maxima and minima.

i. A function is *strictly concave* if its first derivative(s) is positive, and its second derivative(s) is negative over its entire domain.

ii. A function is concave if its first derivative is positive, and its second derivative is less than zero over its entire domain.

D. Functions may have local and global maxima, although most of the functions used in economic model-building are assumed to be strictly concave and so have at most one maximum (e.g. only one local maximum, which is also its global maximum).

i. A function is at local maximum at point Q* if and only if (iff) its first derivative at Q* has the value zero and its second derivative is negative within a finite neighborhood around Q*.

ii. A point, Q*, is the global maximum of function f(Q) if its first derivative has the value zero at Q* and its second derivative is negative throughout the domain of the function. (Notice that in this case function f(Q) is strictly concave.)

iii. Maxima are, as it turns out, important for constrained optimization.

- With a particular domain, as with 0<Q<2, any function, f(Q) will have a highest value.
    - This would be the constrained optima or maximum for function f(Q) within the domain from 0 to 2. It would be the "constrained" optimum.
    - Note that there may be more than one such optima, as when f(Q) is a horizontal straight line or a simple sine curve.

- However, at least one maximum will always exist. This is simply a property of real numbers, within any set there will always be a largest value (number).

iv. When the function is strictly concave and the constraint is a convex set (as with the interval example above, $0<Q<2$) there will be a unique maximum (optimum).

v. Thus, when a consumer has a constraint set that is convex (e.g. can choose any "bundle" within a particular convex set such as budget set), and attempts to maximize a strictly concave objective function such as a utility function, net benefit function, or profit function, the unique maximum will be the rational individual's choice.

- Virtually all of the models of consumer choice are grounded in this very general property.
- It is used, for example, to derive demand and supply curves from consumer and firm choices.
- It is also used in Game theoretic settings where strategy sets are continuous rather than discrete.

E. DEF. A function is said to be **homogeneous of degree k**, if and only if whenever

$$Y = f(X) \text{ , then } \quad f(\beta X) = \beta^k Y$$

i. A production function that is homogenous of degree 1 exhibits constant returns to scale. Doubling all inputs, exactly doubles outputs.

ii. Cobb-Douglas functions and linear functions through the origin ($Y = ax$) are homogeneous of degree 1.

iii. Occasionally, utility and production functions are assumed to be *homothetic*, a somewhat more general family of functions than homogeneous functions.

F. **DEF**. A **homothetic function** is a composite function of the form $H = h(Q(a,b))$ where Q is a homogeneous function and $dH/dQ > 0$ over the entire domain of h or $dH/DQ < 0$ of the entire domain of h. (E.g. a homothetic function is a monotone increasing or decreasing transformation of a homogeneous function.)

i. Not all homothetic functions are homogeneous.

ii. Homothetic utility functions have linear income expansion paths.

iii. Similarly, homothetic production functions have linear output expansion paths. (*The slopes of the isoquants are the same along any straight line through the origin*.)

iv. Assumptions of homogeneity and homotheticity make models and their implication less general than they would have been with assumptions of monotonicity, concavity, and strict concavity, but the clarity of the results is often felt to warrant such assumptions.