5 Consumer Theory

Consumer theory is to demand as producer theory is to supply. The major difference is that producer theory assumes that sellers are motivated by profit, and profit is something that one can usually directly measure. Moreover, the costs that enter into profit arise from physical properties of the production process – how many coffee cups come from the coffee cup manufacturing plant? In contrast, consumer theory is based on what people like, so it begins with something that we can't directly measure, but must infer. That is, consumer theory is based on the premise that we can infer what people like from the choices they make.

Now, inferring what people like from choices they make does not rule out mistakes. But our starting point is to consider the implications of a theory in which consumers don't make mistakes, but make choices that give them the most satisfaction.

Economists think of this approach as analogous to studying gravitation in a vacuum before thinking about the effects of air friction. There is a practical consideration that dictates ignoring mistakes. There are many kinds of mistakes, e.g. "I meant to buy toothpaste but forgot and bought a toothbrush," a memory problem, "I thought this toothpaste was better but it is actually worse," a learning issue, and "I meant to buy toothpaste but I bought crack instead," a self-control issue. All of these kinds of mistakes lead to distinct theories. Moreover, we understand these alternative theories by understanding the basic theory first, and then seeing what changes these theories lead to.

5.1 Utility Maximization

Economists use the term *utility* in a peculiar and idiosyncratic way. Utility refers not to usefulness but to the flow of pleasure or happiness that a person enjoys – some measure of the satisfaction a person experiences. Usefulness might contribute to utility, but so does style, fashion, or even whimsy.

The term *utility* is unfortunate not just because it suggests usefulness, but because it makes the economic approach to behavior appear more limited than it actually is. We will make very few assumptions about the form of utility that a consumer might have. That is, we will attempt to avoid making value judgments about the preferences a consumer holds – whether they like smoking cigarettes or eating only carrots, watching Arnold Schwarzenegger movies or spending time with a hula hoop. Consumers like whatever it is that they like; the economic assumption is that they attempt to obtain the goods that they like. It is the consequences of the pursuit of happiness that comprise the core of consumer theory.

In this chapter, we will focus on two goods. In many cases, the generalization to an arbitrary number of goods is straightforward. Moreover, in most applications it won't matter because we can view one of the goods as a "composite good" reflecting consumption of a bunch of other goods.[46]

---

[46] Thus, for example, savings for future consumption, or to provide for descendents, or to give to your alma mater, are all examples of consumption. Our consumer will, in the end, always spend all of her

As a starting point, suppose the two goods are X and Y. To distinguish the quantity of the good from the good itself, we'll use capital letters to indicate the good and a lower case letter to indicate the quantity consumed. If X is rutabagas, a consumer who ate three of them would have $x=3$. How can we represent preferences for this consumer? To fix ideas, suppose the consumer is both hungry and thirsty and the goods are beer and pizza. The consumer would like more of both, reflected in greater pleasure for greater consumption. Items one might consume are generally known as "bundles," as in bundles of goods and services, and less frequently as "tuples," a short-form for the "n-tuple," meaning a list of n quantities. Since we will focus on two goods, both of these terms are strained in the application; a bundle because a bundle of two things isn't much of a bundle, and a tuple because what we have here is a "two-tuple," also known as a pair. But part of the job of studying economics is to learn the language of economics, and bundles it is.

One might naturally consider measuring utility on some kind of physical basis – production of dopamine in the brain, for example – but it turns out that the actual quantities of utility don't matter for the theory we develop. What matters is whether a bundle produces more than another, or less, or the same. Let $u(x, y)$ represent the utility a consumer gets from consuming x units of beer and y units of pizza. The function u guides the consumer's choice, in the sense that, if the consumer can choose either (x_1, y_1) or (x_2, y_2), we expect him to choose (x_1, y_1) if $u(x_1, y_1) > u(x_2, y_2)$.

But notice that a doubling of u would lead to the same choices, because

$u(x_1, y_1) > u(x_2, y_2)$ if and only if $2u(x_1, y_1) > 2u(x_2, y_2)$.

Thus, doubling the utility doesn't change the preferences of the consumer. But the situation is more extreme than this. Even exponentiating the utility doesn't change the consumer's preferences, because

$u(x_1, y_1) > u(x_2, y_2)$ if and only if $e^{u(x1, y1)} > e^{u(x2, y2)}$.

Another way to put this is that there are no natural units for utility, at least until such time as we are able to measure pleasure in the brain.

It is possible to develop the theory of consumer choice without supposing that a utility function exists at all. However, it is expedient to begin with utility, to simplify the analysis for introductory purposes.

5.1.1 Budget or Feasible Set

Suppose a consumer has a fixed amount of money to spend, M. There are two goods X and Y, with associated prices p_X and p_Y. The feasible choices the consumer can make satisfy $p_X x + p_Y y \leq M$. In addition, we will focus on consumption and rule out

---

income, although this happens because we adopt a very broad notion of spending. In particular, savings are "future spending."

negative consumption, so $x \geq 0$ and $y \geq 0$. This gives a *budget set* or *feasible set* illustrated in Figure 5-1.

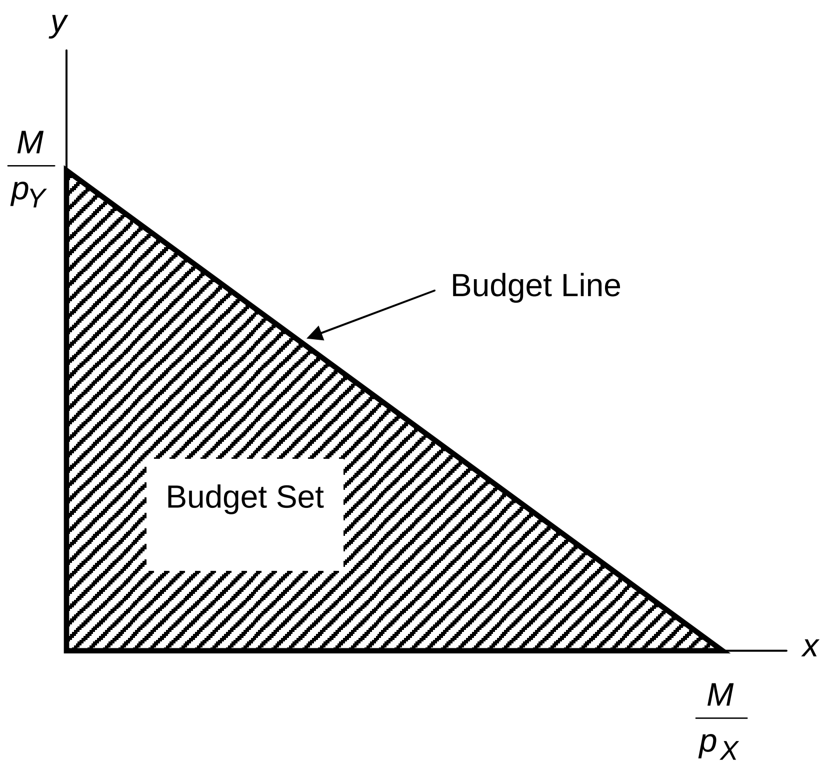

Figure 5-1: Budget Set

In this diagram, the feasible set of purchases that satisfy the budget constraint are illustrated with shading. If the consumer spends all her money on *X*, she can consume the quantity $x = \frac{M}{p_X}$. Similarly, if she spends all of her money on *Y*, she consumes $\frac{M}{p_Y}$ units of *Y*. The straight line between them, known as the budget line, represents the most of the goods she can consume. The slope of the budget line is $-\frac{p_X}{p_Y}$.

An increase in the price of one good pivots or rotates the budget line. Thus, if the price of *X* increases, the endpoint $\frac{M}{p_Y}$ remains the same, but $\frac{M}{p_X}$ falls. This is illustrated in Figure 5-2.

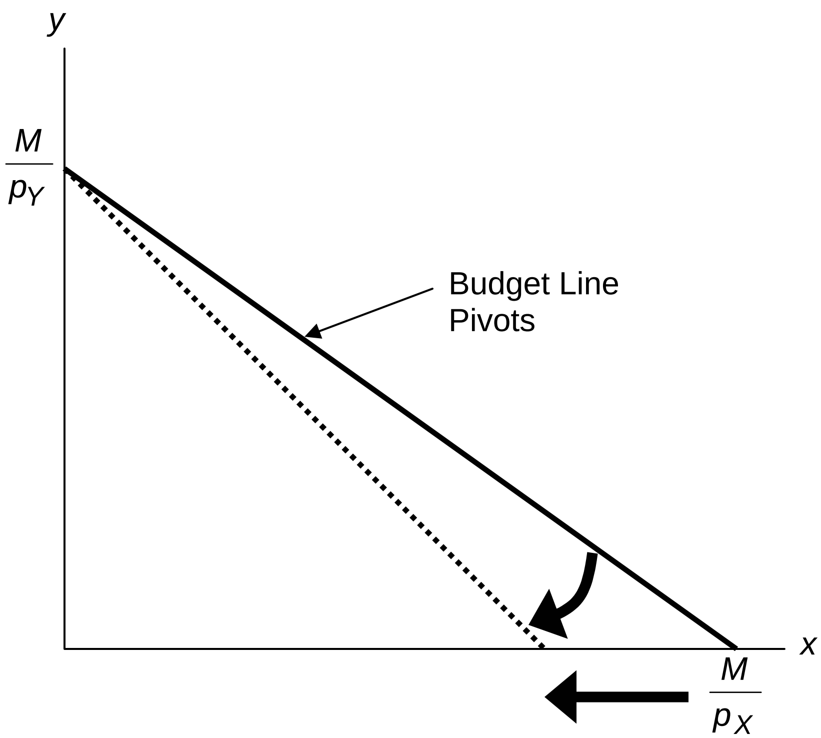

Figure 5-2: Effect of an Increase in Price on the Budget

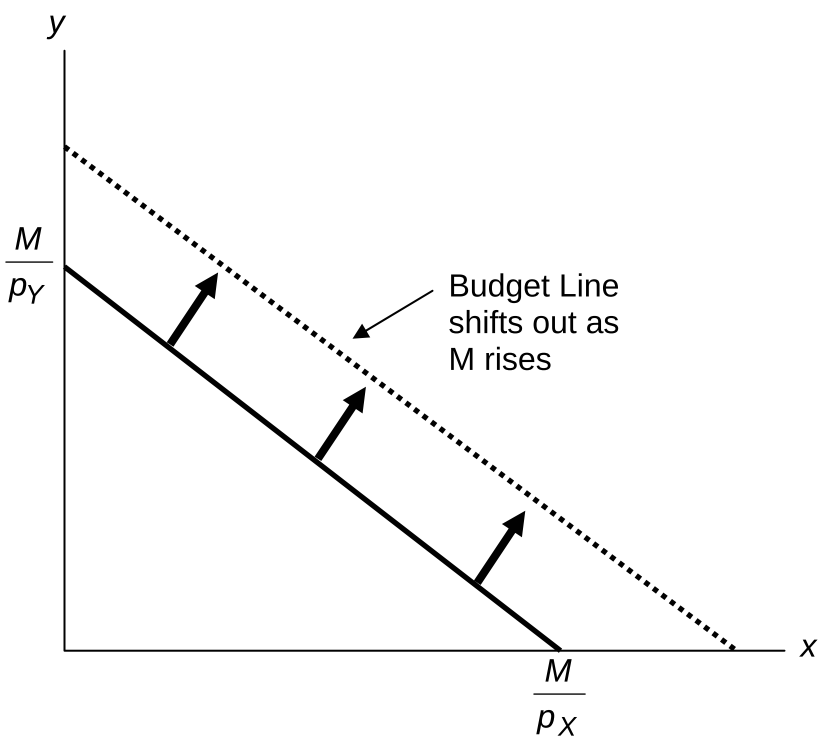

Figure 5-3: An Increase in Income

The effect of increasing the available money M is to increase both $\frac{M}{p_X}$ and $\frac{M}{p_Y}$ proportionately. This means an increase in M shifts the budget line out (away from the origin) in a parallel fashion, as in Figure 5-3.

An increase in both prices by the same proportional factor has an effect identical to a decrease in income. Thus, one of the three financial values – the two prices and income – is redundant. That is, we can trace out all the possible budget lines with any two of the three parameters. This can prove useful; we can arbitrarily set p_X to be the number one without affecting the generality of the analysis. When setting a price to one, that related good is called the *numeraire*, and essentially all prices are denominated with respect to that one good. A real world example of a numeraire occurred when the currency used was based on gold so that the prices of other goods are denominated in terms of the value of gold.

Money is not necessarily the only constraint on the consumption of goods that a consumer faces. Time can be equally important. One can own all the compact discs in the world, but they are useless if one doesn't actually have time to listen to them. Indeed, when we consider the supply of labor, time will be a major issue – supplying labor (working) uses up time that could be used to consume goods. In this case there will be two kinds of budget constraints – a financial one and a temporal one. At a fixed wage, time and money translate directly into one another and the existence of the time constraint won't present significant challenges to the theory. The conventional way to handle the time constraint is to use as a baseline working "full out," and then view leisure as a good which is purchased at a price equal to the wage. Thus, if you earn $20/hour, we would set your budget at $480/day, reflecting 24 hours of work, but then permit you to buy leisure time, during which eating, sleeping, brushing teeth and every other non-work activity is accomplished at a price equal to $20 per hour.

5.1.1.1 (Exercise) Graph the budget line for apples and oranges, with prices of $2 and $3 respectively and $60 to spend. Now increase the price of apples from $2 to $4 and draw the budget line.

5.1.1.2 (Exercise) Suppose that apples cost $1 each. Water can be purchased for 0.5 cents per gallon up to 20,000 gallons, and 0.1 cent per gallon for each gallon beyond 20,000 gallons. Draw the budget constraint for a consumer who spends $200 per month on apples and water.

5.1.1.3 (Exercise) Graph the budget line for apples and oranges, with prices of $2 and $3 respectively and $60 to spend. Now increase expenditure to $90 and draw the budget line.

5.1.2 Isoquants

With two goods, we can graphically represent utility by considering the contour map of utility. Utility contours are known as *isoquants,* meaning "equal quantity," and are also known as *indifference curves,* since the consumer is indifferent between points on the line. We have met this idea already in the description of production functions, where the curves represented input mixes that produced a given output. The *only* difference here is that the output being produced is consumer "utility" instead of a single good or service.

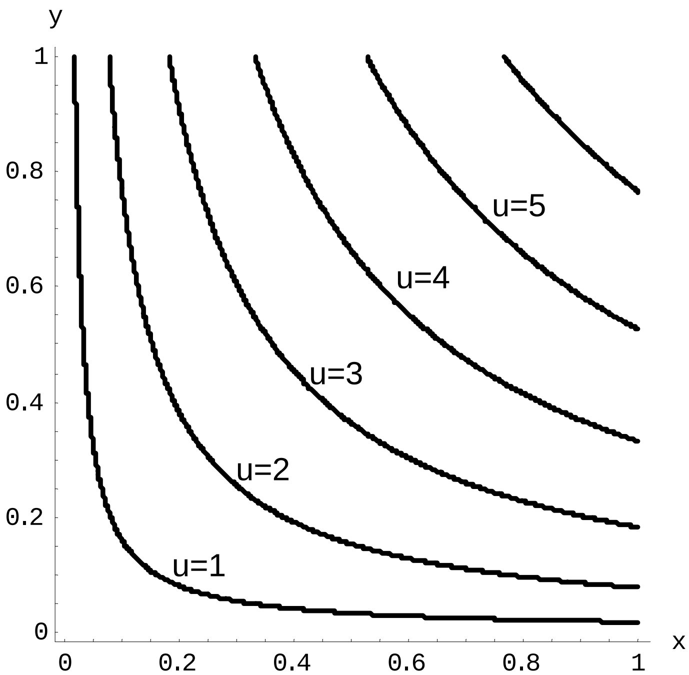

Figure 5-4: Utility Isoquants

Figure 5-4 provides an illustration of isoquants or indifference curves. Each curve represents one level of utility. Higher utilities occur to the northeast, further away from the origin. As with production isoquants, the slope of the indifference curves has the interpretation of the tradeoff between the two goods. The amount of *Y* that the consumer is willing to give up to obtain an extra bit of *X* is the slope of the indifference curve. Formally, the equation

$$u(x, y) = u_0$$

defines an indifference curve for the reference utility u_0. Differentiating in such a way as to preserve the equality, we obtain the slope of the indifference curve:

$$\frac{\partial u}{\partial x}dx + \frac{\partial u}{\partial y}dy = 0 \text{ or } \left.\frac{dy}{dx}\right|_{u=u_0} = -\frac{\partial u / \partial x}{\partial u / \partial y}.$$

This slope is known as the marginal rate of substitution and reflects the tradeoff, from the consumer's perspective, between the goods. That is to say, the marginal rate of substitution (of *Y* for *X*) is the amount of *Y* the consumer is willing to lose to obtain an extra unit of *X*.

An important assumption concerning isoquants is reflected in the diagram: "midpoints are preferred to extreme points." Suppose the consumer is indifferent between (x_1, y_1)

and (x_2, y_2), that is, $u(x_1, y_1) = u(x_2, y_2)$. Then we say *preferences are convex* if any point on the line segment connecting (x_1, y_1) and (x_2, y_2) is at least as good as the extremes. Formally, a point on the line segment connecting (x_1, y_1) and (x_2, y_2) comes in the form

$$(\alpha x_1 + (1 - \alpha)\, x_2, \alpha y_1 + (1 - \alpha)\, y_2),$$

for α between zero and one. This is also known as a "convex combination" between the two points. When α is zero, the segment starts at (x_2, y_2) and proceeds in a linear fashion to (x_1,y_1) at α equal to one. Preferences are convex if, for any α between 0 and 1,

$$u(x_1, y_1) = u(x_2, y_2) \text{ implies } u(\alpha x_1 + (1 - \alpha)\, x_2, \alpha y_1 + (1 - \alpha)\, y_2) \geq u(x_1, y_1).$$

This property is illustrated in Figure 5-5. The line segment that connects two points on the indifference curve lies to the northeast of the indifference curve, which means the line segment involves strictly more consumption of both goods than some points on the indifference curve, which means that it is preferred to the indifference curve. Convex preferences mean that a consumer prefers a mix to any two equally valuable extremes. Thus, if the consumer likes black coffee and also likes drinking milk, the consumer prefers some of each (not necessarily mixed) to only drinking coffee or only drinking milk. This sounds more reasonable if you think of the consumer's choices on a monthly basis; if you like drinking 60 cups of coffee, and no milk, per month the same as 30 glasses of milk and no coffee, convex preferences entails preferring 30 cups of coffee and 15 glasses of milk to either extreme.

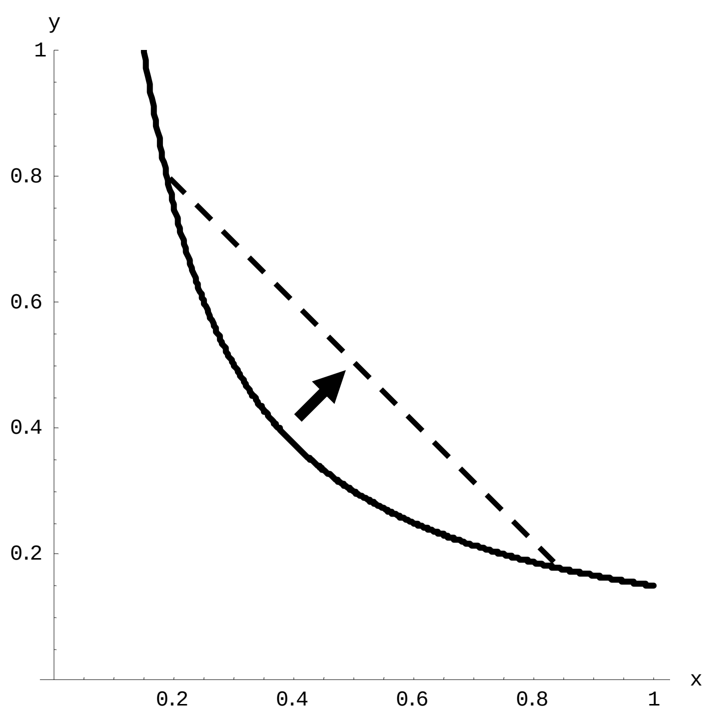

Figure 5-5: Convex Preferences

How does a consumer choose which bundle to select? The consumer is faced with the problem of maximizing $u(x, y)$ subject to $p_X x + p_Y y \leq M$.

We can derive the solution to the consumer's problem as follows. First, "solve" the budget constraint $p_X x + p_Y y \le M$ for y, to obtain $y \le \frac{M - p_X x}{p_Y}$. If Y is a good, this constraint will be satisfied with equality and all the money will be spent. Thus, we can write the consumer's utility as

$$u\left(x, \frac{M - p_X x}{p_Y}\right).$$

The first order condition for this problem, maximizing it over x, has

$$0 = \frac{d}{dx} u\left(x, \frac{M - p_X x}{p_Y}\right) = \frac{\partial u}{\partial x} - \frac{p_X}{p_Y}\frac{\partial u}{\partial y}.$$

This can be re-arranged to obtain the marginal rate of substitution (MRS).

$$\frac{p_X}{p_Y} = \frac{\partial u / \partial x}{\partial u / \partial y} = -\left.\frac{dy}{dx}\right|_{u=u_0} = MRS\ .$$

The first order condition requires that the slope of the indifference curve equals the slope of the budget line, that is, there is a tangency between the indifference curve and the budget line. This is illustrated in Figure 5-6. Three indifference curves are drawn, two of which intersect the budget line, but are not tangent. At these intersections, it is possible to increase utility by moving "toward the center," until the highest of the three indifference curves is reached. At this point, further increases in utility are not feasible, because there is no intersection between the set of bundles that produce a strictly higher utility and the budget set. Thus, the large black dot is the bundle that produces the highest utility for the consumer.

It will later prove useful to also state the second order condition, although we won't use this condition now:

$$0 \ge \frac{d^2}{(dx)^2} u\left(x, \frac{M - p_X x}{p_Y}\right) = \frac{\partial^2 u}{(\partial x)^2} - \frac{p_X}{p_Y}\frac{\partial^2 u}{\partial x \partial y} + \left(\frac{p_X}{p_Y}\right)^2 \frac{\partial^2 u}{(\partial y)^2}\ .$$

Note that the vector $(u_1, u_2) = \left(\frac{\partial u}{\partial x}, \frac{\partial u}{\partial y}\right)$ is the gradient of u, and the gradient points in the direction of steepest ascent of the function u. Second, the equation which characterizes the optimum,

$$0 = p_X \frac{\partial u}{\partial y} - p_Y \frac{\partial u}{\partial x} = \left(\frac{\partial u}{\partial x}, \frac{\partial u}{\partial y} \right) \bullet (-p_Y, p_X),$$

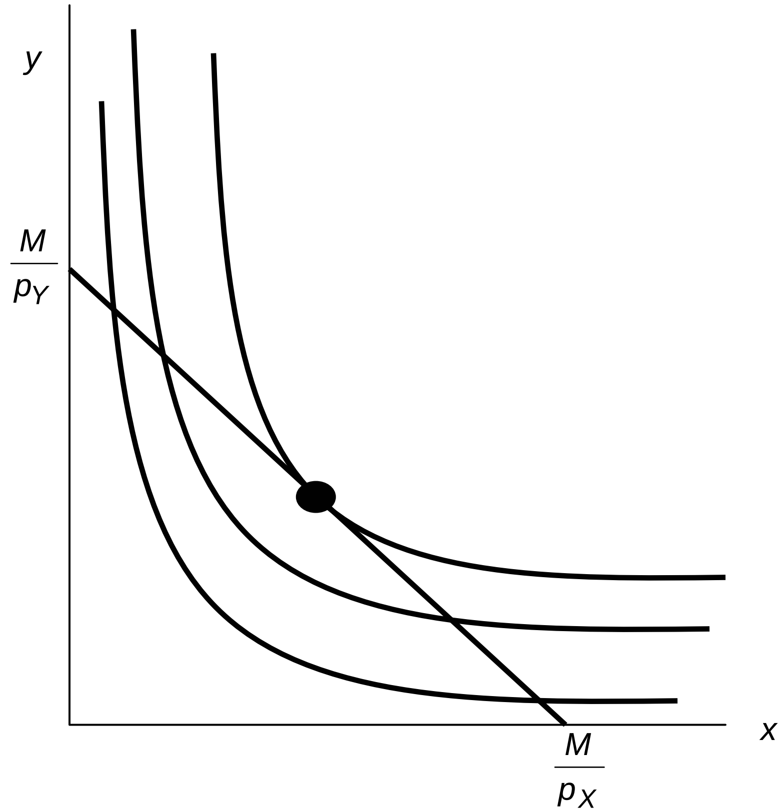

Figure 5-6: Graphical Utility Maximization

where • is the "dot product" which multiplies the components of vectors and then adds them, says that the vectors (u_1, u_2) and ($-p_Y$, p_X) are perpendicular, and hence that *the rate of steepest ascent of the utility function is perpendicular to the budget line.*

When does this tangency approach fail to solve the consumer's problem? There are three ways it can fail. First, the utility might not be differentiable. We will set aside this kind of failure with the remark that fixing points of non-differentiability is mathematically challenging but doesn't lead to significant alterations in the theory. The second failure is that a tangency didn't maximize utility. Figure 5-7 illustrates this case. Here, there is a tangency, but it doesn't maximize utility. In Figure 5-7, the dotted indifference curve maximizes utility given the budget constraint (straight line). This is exactly the kind of failure that is ruled out by convex preferences. In Figure 5-7, preferences are not convex, because if we connect two points on the indifference curves and look at a convex combination, we get something less preferred, with lower utility, not more preferred as convex preferences would require.

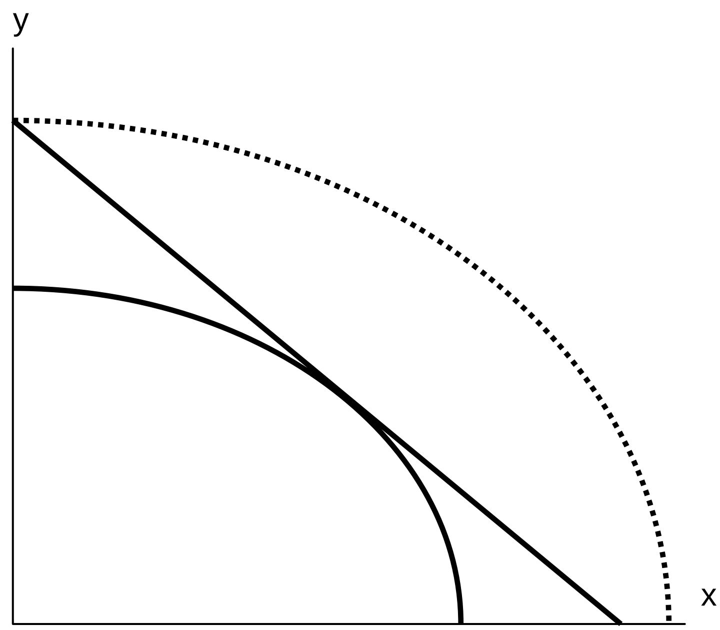

Figure 5-7: "Concave" Preferences, Prefer Boundaries

The third failure is more fundamental: the derivative might fail to be zero because we've hit the boundary of $x=0$ or $y=0$. This is a fundamental problem because in fact there are many goods that we do buy zero of, so zeros for some goods are not uncommon solutions to the problem of maximizing utility. We will take this problem up in a separate section, but we already have a major tool to deal with it: convex preferences. As we shall see, convex preferences insure that the consumer's maximization problem is "well-behaved."

5.1.3 Examples

The Cobb-Douglas utility function comes in the form $u(x,y)=x^{\alpha}y^{1-\alpha}$. Since utility is zero if either of the goods is zero, we see that a consumer with Cobb-Douglas preferences will always buy some of each good. The marginal rate of substitution for Cobb-Douglas utility is

$$-\left.\frac{dy}{dx}\right|_{u=u_0}=\frac{\partial u/\partial x}{\partial u/\partial y}=\frac{\alpha y}{(1-\alpha)x}.$$

Thus, the consumer's utility maximization problem yields

$$\frac{p_X}{p_Y}=-\left.\frac{dy}{dx}\right|_{u=u_0}=\frac{\partial u/\partial x}{\partial u/\partial y}=\frac{\alpha y}{(1-\alpha)x}.$$

Thus, using the budget constraint, $(1-\alpha)xp_X=\alpha y p_Y=\alpha(M-xp_X)$.

This yields $x = \frac{\alpha M}{p_X}$, $y = \frac{(1-\alpha)M}{p_Y}$.

Cobb-Douglas utility results in *constant expenditure shares*. No matter what the price of X or Y, the expenditure xp_X on X is αM. Similarly, the expenditure on Y is $(1-\alpha)M$. This makes the Cobb-Douglas utility very useful for computing examples and homework exercises.

5.1.3.1 (Exercise) Consider a consumer with utility $u(x,y) = \sqrt{xy}$. If the consumer has $100 to spend, and the price of X is $5 and the price of Y is $2, graph the budget line, and then find the point that maximizes the consumer's utility given the budget. Draw the utility isoquant through this point. What are the expenditure shares?

5.1.3.2 (Exercise) Consider a consumer with utility $u(x,y) = \sqrt{xy}$. Calculate the slope of the isoquant directly, by solving $u(x,y) = u_0$ for y as a function of x and the utility level u_0. What is the slope $-\left.\frac{dy}{dx}\right|_{u=u_0}$? Verify that it satisfies the formula given above.

5.1.3.3 (Exercise) Consider a consumer with utility $u(x,y) = (xy)^2$. Calculate the slope of the isoquant directly, by solving $u(x,y) = u_0$ for y as a function of x and the utility level u_0. What is the slope $-\left.\frac{dy}{dx}\right|_{u=u_0}$? Verify that the result is the same as in the previous exercise. Why is it the same?

When two goods are perfect complements, they are consumed proportionately. The utility that gives rise to perfect complements is in the form $u(x, y) = \min\{x, \beta y\}$ for some constant β (the Greek letter beta). First observe that with perfect complements, consumers will buy in such a way that $x = \beta y$. The reason is that, if $x > \beta y$, some expenditure on x is a waste since it brings in no additional utility, and the consumer gets higher utility by decreasing x and increasing y. This lets us define a "composite good" which involves buying some amount y of Y and also buying βy of X. The price of this composite commodity is $\beta p_X + p_Y$, and it produces utility $u = \frac{M}{\beta p_X + p_Y}$. In this way, perfect complements boil down to a single good problem.

5.1.3.4 (Exercise) The case of perfect substitutes arises when all that matters to the consumer is the sum of the products – e.g. red shirts and green shirts for a color-blind consumer. In this case, $u(x, y) = x + y$. Graph the isoquants for perfect substitutes. Show that the consumer maximizes utility by spending their entire income on whichever product is cheaper.

If the only two goods available in the world were pizza and beer, it is likely that *satiation* would set in at some point. How many pizzas can you eat per month? How much beer can you drink? [Don't answer that.]

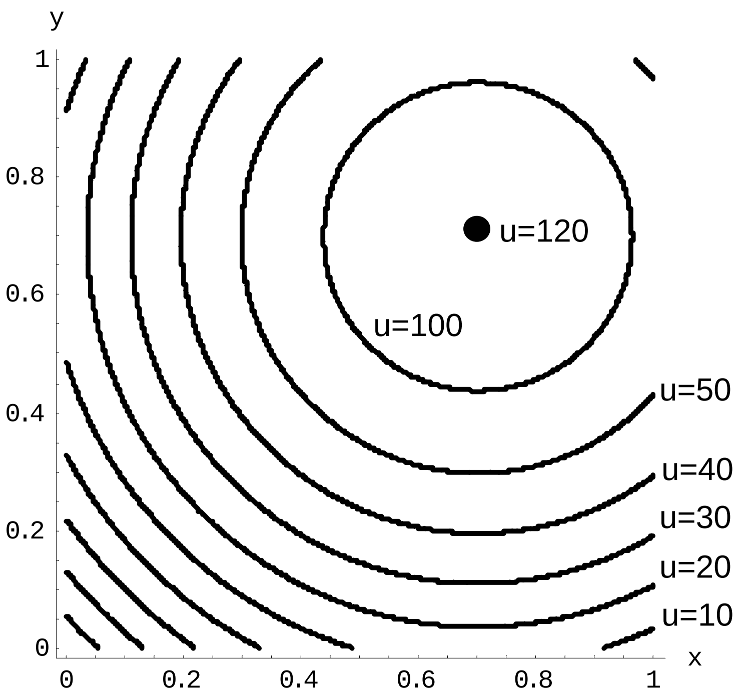

Figure 5-8: Isoquants for a Bliss Point

What does satiation mean for isoquants? It means there is a point that maximizes utility, which economists call a *bliss point*. An example is illustrated in Figure 5-8. Near the origin, the isoquants behave as before. However, as one gets full of pizza and beer, a point of maximum value is reached, illustrated by a large black dot. What does satiation mean for the theory? First, if the bliss point isn't within reach, the theory behaves as before. With a bliss point within reach, consumption will stop at the bliss point. A feasible bliss point entails having a zero value of money. There may be people with a zero value of money, but even very wealthy people, who reach satiation in goods that they personally consume, often like to do other things with the wealth and appear not to have reached satiation overall.

5.1.3.5 (Exercise) Suppose $u(x,y) = x^{\alpha} + y^{\alpha}$ for $\alpha<1$. Show

$$x = \frac{M}{p_X\left(1+\left(\frac{p_Y}{p_X}\right)^{\alpha}\right)}, \text{ and } y = \frac{M}{p_Y\left(1+\left(\frac{p_X}{p_Y}\right)^{\alpha}\right)}.$$

5.1.3.6 (Exercise) Suppose one consumer has the utility function u (which is always a positive number), and a second consumer has utility w. Suppose, in addition, that for any x, y, $w(x, y) = (u(x, y))^2$, that is, the second person's utility is the square of the first. Show that these consumers make the same choices – that is, $u(x_a, y_a) \geq u(x_b, y_b)$ if and only $w(x_a, y_a) \geq w(x_b, y_b)$.

5.1.4 Substitution Effects

It would be a simpler world if an increase in the price of a good always entailed buying less of it. Alas, it isn't so, as the following diagram illustrates. In this diagram, an increase in the price of *Y* causes the budget line to pivot around the intersection on the *X* axis, since the amount of *X* that can be purchased hasn't changed. In this case, the quantity *y* of *Y* demanded rises.

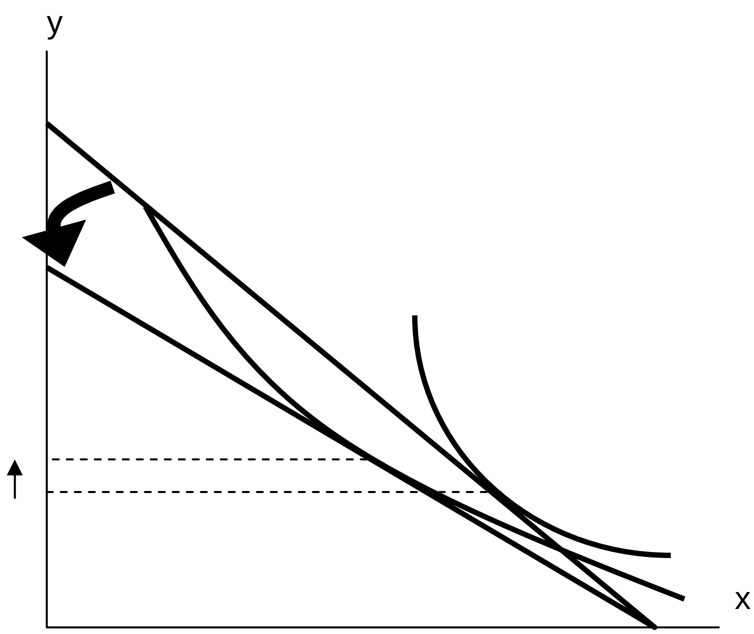

Figure 5-9: Substitution with an Increase in Price

At first glance, this increase in the consumption of a good in response to a price increase sounds implausible, but there are examples where it makes sense. The primary example is leisure. As wages rise, the cost of leisure (forgone wages) rises. But as people feel wealthier, they choose to work fewer hours. The other examples given, which are hotly debated in the "tempest in a teapot" kind of way, involve people subsisting on a good like potatoes but occasionally buying meat. When the price of potatoes rises, they can no longer afford meat and buy even more potatoes than before.

Thus, the logical starting point on substitution – what happens to the demand for a good when the price of that good increases? – does not lead to a useful theory. As a result, economists have devised an alternative approach, based on the following logic. An increase in the price of a good is really a composition of two effects: an increase in the relative price of the good, and a decrease in the purchasing power of money. As a result, it is useful to examine these two effects separately. The substitution effect considers the change in the relative price, with a sufficient change in income to keep the consumer on the same utility isoquant.[47] The income effect changes only income.

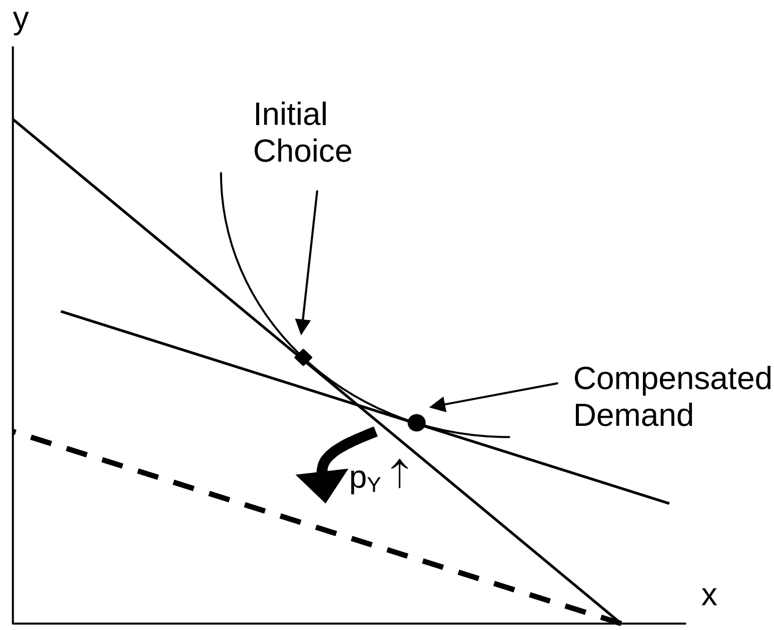

Figure 5-10: Substitution Effect

To graphically illustrate the substitution effect, consider Figure 5-10. The starting point is the tangency between the isoquant and the budget line, denoted with a diamond shape and labeled "Initial Choice." The price of Y rises, pivoting the budget line inward. The new budget line is illustrated with a heavy, dashed line. To find the substitution effect, increase income from the dashed line until the original isoquant is reached. Increases in income shift the budget line out in a fashion parallel to the original. We reach the original isoquant at a point labeled with a small circle, a point sometimes called the *compensated* demand, because we have compensated the consumer for the price increase by increasing income just enough to leave her unharmed, on the same isoquant. The substitution effect is just the difference between these points – the substitution in response to the price change, holding constant the utility of the consumer.

We can readily see that the substitution effect of a price increase in Y is to decrease the consumption of Y and increase the consumption of X.[48] The income effect is the change in consumption resulting from the change in income. The effect of any change in price can be decomposed into the substitution effect, which holds utility constant and changes the relative prices, and the income effect, which adjusts for the loss of purchasing power arising from the price increase.

---

[47] Some authors instead change the income enough to make the old bundle affordable. This approach has the virtue of being readily computed, but the disadvantage that the substitution effect winds up increasing the utility of the consumer. Overall the present approach is more economical for most purposes.
[48] To construct a formal proof, first show that if p_Y rises and y rises, holding utility constant, the initial choice prior to the price increase is feasible after the price increase. Use this to conclude that after the price increase it is possible to have strictly more of both goods, contradicting the hypothesis that utility was held constant.

Example (Cobb-Douglas): Recall that the Cobb-Douglas utility comes in the form $u(x,y)=x^{\alpha}y^{1-\alpha}$. Solving for x, y we obtain

$$x=\frac{\alpha M}{p_X},\quad y=\frac{(1-\alpha)M}{p_Y}\text{, and } u(x,y)=\alpha^{\alpha}(1-\alpha)^{1-\alpha}\frac{M}{p_X^{\alpha}p_Y^{1-\alpha}}.$$

Thus, consider a multiplicative increase Δ in p_Y, that is, multiplying p_Y by $\Delta>1$. In order to leave utility constant, M must rise by $\Delta^{1-\alpha}$. Thus, x rises by the factor $\Delta^{1-\alpha}$ and y falls, by the factor $\Delta^{-\alpha}<1$. This is the substitution effect.

What is the substitution effect of a small change in the price p_Y for any given utility function, not necessarily Cobb-Douglas? To address this question, it is helpful to introduce some notation. We will subscript the utility to indicate partial derivative, that is,

$$u_1=\frac{\partial u}{\partial x},\quad u_2=\frac{\partial u}{\partial y}.$$

Note that, by the definition of the substitution effect, we are holding utility constant, so $u(x, y)$ is being held constant. This means, locally, that

$$0=du=u_1\,dx+u_2\,dy.$$ [49]

In addition, we have $M=p_X x+p_Y y$, so

$$dM=p_X dx+p_Y dy+y\,dp_Y$$

Finally, we have the optimality condition

$$\frac{p_X}{p_Y}=\frac{\partial u/\partial x}{\partial u/\partial y}$$

which is convenient to write as $p_X u_2=p_Y u_1$. Differentiating this equation, and letting $u_{11}=\frac{\partial^2 u}{(\partial x)^2},\quad u_{12}=\frac{\partial^2 u}{\partial x\partial y}\ and\ u_{22}=\frac{\partial^2 u}{(\partial y)^2}$, we have

$$p_X(u_{12}dx+u_{22}dy)=u_1\,dp_Y+p_Y(u_{11}dx+u_{12}dy).$$

[49] Writing dx for an unknown infinitesimal change in x can be put on a formal basis. The easiest way to do so is to think of dx as representing the derivative of x with respect to a parameter, which will be p_Y.

For a given dp_Y, we now have three equations in three unknowns dx, dy, and dM. However, dM only appears in one of the three. Thus, the effect of a price change on x and y can be solved by solving two equations:

$0 = u_1\, dx + u_2\, dy$ and

$$p_X (u_{12} dx + u_{22} dy) = u_1\, dp_Y + p_Y (u_{11} dx + u_{12} dy)$$

for the two unknowns dx and dy. This is straightforward and yields:

$$\frac{dx}{dp_Y} = -\frac{p_Y u_1}{p_X^2 u_{11} + 2 p_X p_Y u_{12} + p_Y^2 u_{22}} \text{ and}$$

$$\frac{dy}{dp_Y} = \frac{p_Y u_2}{p_X^2 u_{11} + 2 p_X p_Y u_{12} + p_Y^2 u_{22}}.$$

These equations imply that x rises and y falls.[50] We immediately see

$$\frac{\frac{dy}{dp_Y}}{\frac{dx}{dp_Y}} = -\frac{u_1}{u_2} = -\frac{p_X}{p_Y}.$$

Thus, the change in (x,y) follows the budget line locally. (This is purely a consequence of holding utility constant.)

To complete the thought while we are embroiled in these derivatives, note that $p_X u_2 = p_Y u_1$ implies that $p_X dx + p_Y dy = 0$.

Thus, the amount of money necessary to compensate the consumer for the price increase, keeping utility constant, can be calculated from our third equation:

$$dM = p_X dx + p_Y dy + y dp_Y = y dp_Y .$$

The amount of income necessary to insure the consumer makes no losses from a price increase in Y is the amount that lets them buy the bundle they originally purchased, that is, the increase in the amount of money is precisely the amount needed to cover the increased price of y. This shows that locally there is no difference from a substitution effect that keeps utility constant (which is what we explored) and one that provides sufficient income to permit purchasing the previously purchased consumption bundle, at least when small changes in prices are contemplated.

[50] This is a consequence of the fact that $p_X^2 u_{11} + 2 p_X p_Y u_{12} + p_Y^2 u_{22} < 0$, which follows from the already stated second order condition for a maximum of utility.

5.1.5 Income Effects

Wealthy people buy more caviar than poor people. Wealthier people buy more land, medical services, cars, telephones, and computers than poorer people, because they have more money to spend on goods and services, and overall, buy more of them. But wealthier people also buy fewer of some goods, too. Rich people buy fewer cigarettes and processed cheese food. You don't see billionaires waiting in line at McDonald's, and that probably isn't because they have an assistant to wait for them. For most goods, at a sufficiently high income, the purchase tends to trail off as income rises.

When an increase in income causes a consumer to buy more of a good that good is called a *normal good* for that consumer. When the consumer buys less, the good is called an *inferior* good, which is an example of sensible jargon that is rare in any discipline. That is, an inferior good is any good whose quantity demanded falls as incomes rise. At a sufficiently low income, almost all goods are normal goods, while at a sufficiently high income, most goods become inferior. Even a Ferrari is an inferior good against some alternatives, such as Lear jets.

The curve that shows the path of consumption as incomes rise is known as an Engel curve.[51] An Engel curve graphs $(x(M), y(M))$ as M varies, where $x(M)$ is the amount of X chosen with income M, and similarly $y(M)$ is the amount of is the amount of Y. An example of an Engel curve is illustrated in Figure 5-11.

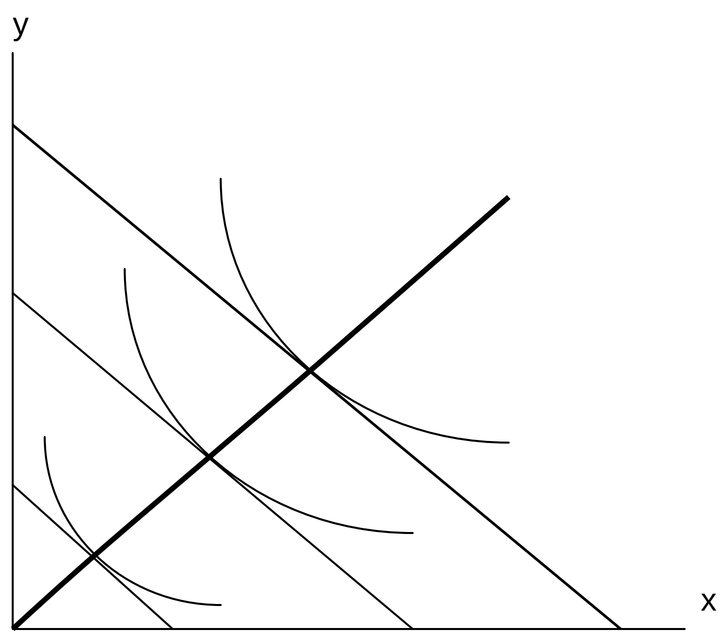

Figure 5-11: Engel Curve

[51] The Engel curve is named for Ernst Engel (1821-1896), a statistician, not for Friedrich Engels, who wrote with Karl Marx.

Example (Cobb-Douglas): Since the equations $x = \frac{\alpha M}{p_X}$, $y = \frac{(1-\alpha)M}{p_Y}$ define the optimal consumption, the Engel curve is a straight line through the origin with slope $\frac{(1-\alpha)p_X}{\alpha\, p_Y}$.

5.1.5.1 (Exercise) Show that, in the case of perfect complements, the Engel curve does not depend on prices.

An inferior good has the quantity fall as incomes rise. Note that, with two goods, at least one is normal good – they can't both be inferior goods, for otherwise when income rose, less of both would be purchased. An example of an inferior good is illustrated in Figure 5-12. Here, as incomes rise, the consumption of *x* rises, reaches a maximum, then begins to decline. In the declining portion, *X* is an inferior good.

The definition of the substitution effect now permits us to decompose the effect of a price change into a substitution effect and an income effect. This is illustrated in Figure 5-13.

What is the mathematical form of the income effect? This is actually more straightforward to compute than the substitution effect computed above. As with the substitution effect, we differentiate the conditions $M = p_x x + p_y y$ and $p_x u_2 = p_y u_1$, holding p_X and p_Y constant, to obtain:

$$dM = p_X dx + p_Y dy \text{ and } p_X(u_{12}dx + u_{22}dy) = p_Y(u_{11}dx + u_{12}dy).$$

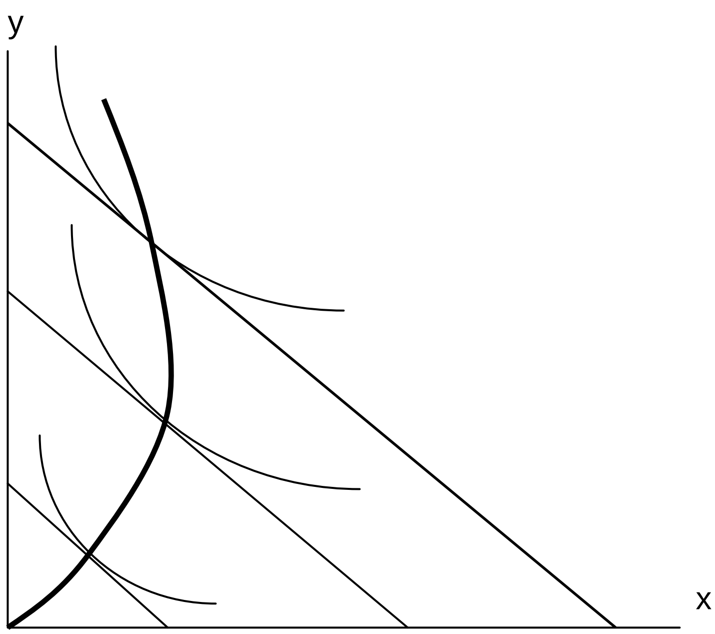

Figure 5-12: Backward Bending – Inferior Good

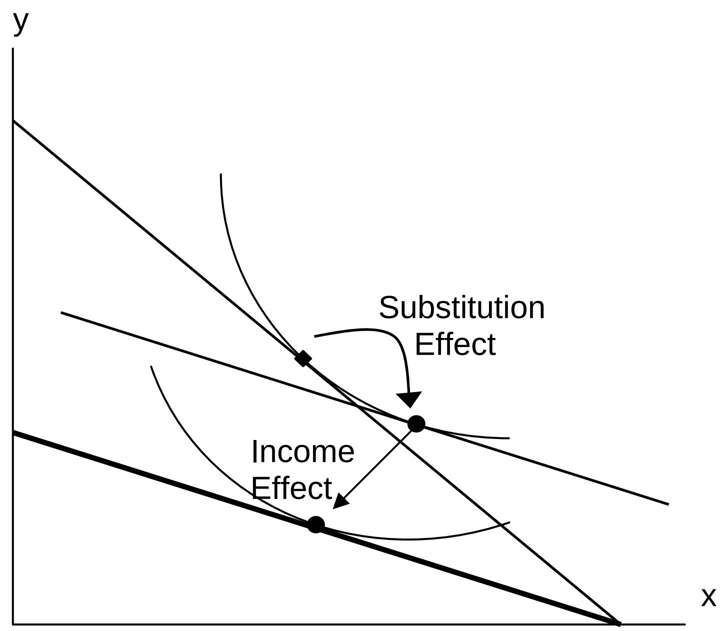

Figure 5-13: Income and Substitution Effects

The second condition can also be written as

$$\frac{dy}{dx} = \frac{p_Y u_{11} - p_X u_{12}}{p_X u_{22} - p_Y u_{12}}.$$

This equation alone defines the slope of the Engel curve, without determining how large a change arises from a given change in *M*. The two conditions together can be solved for the effects of *M* on *X* and *Y*. The Engel curve is given by

$$\frac{dx}{dM} = \frac{p_Y^2 u_{11} - 2 p_X u_{12} + p_X^2 u_{22}}{p_X u_{22} - p_Y u_{12}} \text{ and } \frac{dy}{dM} = \frac{p_Y^2 u_{11} - 2 p_X u_{12} + p_X^2 u_{22}}{p_Y u_{11} - p_X u_{12}}.$$

Note (from the second order condition) that good *Y* is inferior if $p_Y u_{11} - p_X u_{12} > 0$, or if $\frac{u_{11}}{u_1} - \frac{u_{12}}{u_2} > 0$, or $\frac{u_1}{u_2}$ is increasing in *x*. Since $\frac{u_1}{u_2}$ is locally constant when *M* increases, equaling the price ratio, and an increase in *y* increases $\frac{u_1}{u_2}$ (thanks to the second order condition), the only way to keep $\frac{u_1}{u_2}$ equal to the price ratio is for *x* to fall. This property characterizes an inferior good – an increase in the quantity of the good increases the marginal rate of substitution of that good for another good.

5.1.5.2 (Exercise) Compute the substitution effect and income effect associated with a multiplicative price increase Δ in p_Y, that is, multiplying p_Y by Δ>1, for the case of Cobb-Douglas utility $u(x,y)=x^{\alpha}y^{1-\alpha}$.

5.2 Additional Considerations

Let us revisit the maximization problem considered in this chapter. The consumer can spend M on either or both of two goods. This yields a payoff of $h(x)=u\left(x,\frac{M-p_X x}{p_Y}\right)$. When is this problem well behaved? First, if h is a *concave* function of x, which implies $h''(x)\le 0$,[52] then any solution to the first order condition is in fact a maximum. To see this, note that $h''(x)\le 0$ entails $h'(x)$ decreasing. Moreover, if the point x^* satisfies $h'(x^*)=0$, then for $x\le x^*$, $h'(x)\ge 0$, and for $x\ge x^*$, $h'(x)\le 0$, because $h'(x)$ gets smaller as x gets larger, and $h'(x^*)=0$. Now consider $x\le x^*$. Since $h'(x)\ge 0$, h is increasing as x gets larger. Similarly, for $x\ge x^*$, $h'(x)\le 0$, which means h gets smaller as x gets larger. Thus, h concave and $h'(x^*)=0$ means that h is maximized at x^*.

Thus, a *sufficient* condition for the first order condition to characterize the maximum of utility is that $h''(x)\le 0$, for all x, p_X, p_Y, and M. Letting $z=\frac{p_X}{p_Y}$, this is equivalent to

$$u_{11}-2zu_{12}+z^2u_{22}\le 0 \text{ for all } z>0.$$

In turn, we can see that this requires (i) $u_{11}\le 0$ ($z=0$) and (ii) $u_{22}\le 0$ ($z\to\infty$), and (iii) $\sqrt{u_{11}u_{22}}-u_{12}\ge 0$ $\left(z=\sqrt{u_{11}/u_{22}}\right)$. In addition, since

$$-\left(u_{11}+2zu_{12}+z^2u_{22}\right)=\left(\sqrt{-u_{11}}-z\sqrt{-u_{22}}\right)^2+2z\left(\sqrt{u_{11}u_{22}}-u_{12}\right),$$

(i), (ii) and (iii) are sufficient for $u_{11}+2zu_{12}+z^2u_{22}\le 0$.

Therefore, if (i) $u_{11}\le 0$ and (ii) $u_{22}\le 0$, and (iii) $\sqrt{u_{11}u_{22}}-u_{12}\ge 0$, a solution to the first order conditions characterizes utility maximization for the consumer. We will assume that these conditions are met for the remainder of this chapter.

5.2.1 Corner Solutions

When will a consumer specialize and consume zero of a good? A necessary condition for the choice of x to be zero is that the consumer doesn't benefit from consuming a very small x, that is, $h'(0)\le 0$. This means

[52] The definition of concavity is that h is concave if $0<a<1$ and for all x, y, $h(ax+(1-a)y)\ge ah(x)+(1-a)h(y)$. It is reasonably straightforward to show this implies the second derivative of h is negative, and if h is twice differentiable, the converse is true as well.

$$h'(0) = u_1\left(0, M/p_Y\right) - u_2\left(0, M/p_Y\right) p_X/p_Y \le 0 ,$$

or

$$\frac{u_1\left(0, M/p_Y\right)}{u_2\left(0, M/p_Y\right)} \le p_X/p_Y .$$

Moreover, if the concavity of *h* is met, as assumed above, then this condition is sufficient to guarantee that the solution is zero. To see that, note that concavity of *h* implies h' is decreasing. Combined with $h'(0) \le 0$, that entails *h* maximized at 0. An important class of examples of this behavior are *quasilinear* utility. Quasilinear utility comes in the form $u(x, y) = y + v(x)$, where *v* is a concave function ($v''(x) \le 0$ for all *x*).

5.2.1.1 (Exercise) Demonstrate that the quasilinear consumer will consume zero *X* if and only if $v'(0) \le \frac{p_x}{p_y}$, and that the consumer instead consumes zero *Y* if $v'\left(M/p_X\right) \ge \frac{p_x}{p_y}$. The quasilinear utility isoquants, for $v(x) = (x + 0.03)^{0.3}$, are illustrated in Figure 5-14. Note that even though the isoquants curve, they are nonetheless parallel to each other

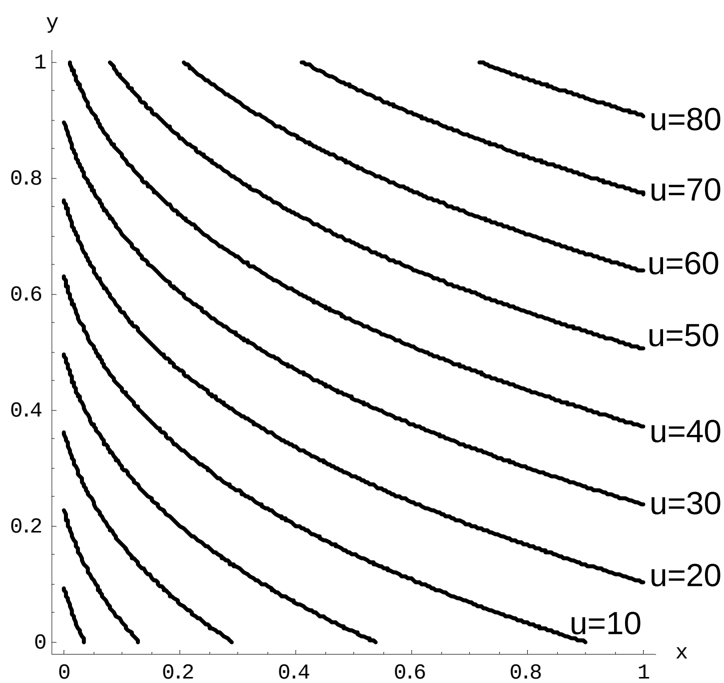

Figure 5-14: Quasilinear Isoquants

The procedure for dealing with corners is generally this. First, check concavity of the *h* function. If *h* is concave, we have a procedure to solve the problem; when *h* is not concave, an alternative strategy must be devised. There are known strategies for some cases that are beyond the scope of this text. Given *h* concave, the next step is to check the endpoints, and verify that $h'(0) > 0$ (for otherwise *x*=0 maximizes the consumer's utility) and that $h'\left(M\big/ p_X \right) < 0$ (for otherwise *y*=0 maximizes the consumer's utility). Finally, at this point we seek the interior solution $h'(x) = 0$. With this procedure we can insure we find the actual maximum for the consumer, rather than a solution to the first order conditions that doesn't maximize the consumer's utility.

5.2.2 Labor Supply

Consider a taxi driver who owns a car or convenience store owner, or anyone else who can set his own hours. Working has two effects on this consumer – more goods consumption, but less leisure consumption. To model this, we let *x* be the goods consumption, *L* the amount of non-work time or *leisure*, and working time *T* – *L*, where *T* is the amount of time available for activities of all kinds. The variable *L* includes a lot of activities that aren't necessarily fun, like trips to the dentist and haircuts and sleeping, but for which the consumer isn't paid, and which represent choices. One could argue that sleeping isn't really a choice, in the sense that one can't choose zero sleep, but this can be handled by adjusting *T* to represent "time available for chosen behavior" so that *T* – *L* is work time and *L* the chosen non-work activities. We set *L* to be leisure rather than labor supply because it is leisure that is the good thing, whereas most of us view working as something we are willing to do provided we're paid for it.

Labor supply is different from other consumption because the wage enters the budget constraint twice – first as the price of leisure and second as income from working. One way of expressing this is to write the consumer's budget constraint as

$$p\mathrm{x} + wL = M + wT.$$

Here, *M* represents non-work income, such as gifts, government transfers, and interest income. We drop the subscript on the price of *X*, and use *w* as the wage. Finally, we use a capital *L* for leisure because a small el looks like the number one. The somewhat Dickensian idea is that the consumer's maximal budget entails working the total available hours *T*, and any non-worked hours are purchased at the wage rate *w*. Alternatively, one could express the budget constraint so as to reflect that expenditures on goods *px* equals the total money, which is the sum of non-work income *M* and work income *w*(*T* – *L*), or

$$p\mathrm{x} = M + w(T - L).$$

These two formulations of the budget constraint are mathematically equivalent.

The strategy for solving the problem is also equivalent to the standard formulation, although there is some expositional clarity used by employing the budget constraint to eliminate *x*. That is, we write the utility *u*(*x*,*L*)

$$h(L) = u\left(\frac{M + w(T - L)}{p}, L\right).$$

As before, we obtain the first order condition

$$0 = h'(L^*) = -u_1\left({}^{w}\!/_{p}\right) + u_2$$

where the partial derivatives u_1 and u_2 are evaluated at $\left(\frac{M + w(T - L^*)}{p}, L^*\right)$. Note that the first order condition is the same as the standard two-good theory developed already. This is because the effect so far is merely to require two components to income: M and wT, both of which are constant. It is only when we evaluate the effect of a wage increase that we see a difference.

To evaluate the effect of a wage increase, differentiate the first order condition to obtain

$$0 = \left(u_{11}\left({}^{w}\!/_{p}\right)^2 - 2u_{12}\left({}^{w}\!/_{p}\right) + u_{22}\right)\frac{dL}{dw} - \frac{u_1}{p} - \left(\frac{w}{p}\right)u_{11}\frac{T - L}{p} + u_{12}\frac{T - L}{p}$$

Since $u_{11}\left({}^{w}\!/_{p}\right)^2 - 2u_{12}\left({}^{w}\!/_{p}\right) + u_{22} < 0$ by the standard second order condition,

$\frac{dL}{dw} > 0$ if, and only if, $\frac{u_1}{p} + \left(\frac{w}{p}\right)u_{11}\frac{T - L}{p} - u_{12}\frac{T - L}{p} < 0$, that is, these expressions are equivalent to one another. Simplifying the latter, we obtain

$$\frac{-\left(\frac{w}{p}\right)u_{11}\frac{T - L}{p} + u_{12}\frac{T - L}{p}}{\frac{u_1}{p}} > 1\text{, or,}$$

$$(T - L)\frac{-\left(\frac{w}{p}\right)u_{11} + u_{12}}{u_1} > 1\text{, or,}$$

$$\frac{\partial}{\partial L}Log(u_1) > \frac{1}{T - L} = -\frac{\partial}{\partial L}Log(T - L)\text{, or,}$$

$$\frac{\partial}{\partial L}Log(u_1) + \frac{\partial}{\partial L}Log(T - L) > 0\text{, or,}$$

$$\frac{\partial}{\partial L} Log\left(u_1(T-L)\right) > 0 .$$

Since the logarithm is increasing, this is equivalent to $u_1(T-L)$ being an increasing function of L. That is, L rises with an increase in wages, and hours worked falls, if the marginal utility of goods times the hours worked is an increasing function of L, holding constant everything else, but evaluated at the optimal values. The value u_1 is the marginal value of an additional good, while the value T-L is the hours worked. Thus, in particular, if goods and leisure are substitutes, so that an increase in L decreases the marginal value of goods, then an increase in the wage must decrease leisure, and labor supply increases in the wage. The case where the goods are complements holds a hope for a decreasing labor supply, so we consider first the extreme case of complements.

Example (perfect complements): $u(x, L)$= Min $\{x, L\}$

In this case, the consumer will make consumption and leisure equal to maximize the utility, so

$$\frac{M + w(T - L^*)}{p} = L^*$$

or

$$L^* = \frac{\frac{M + wT}{p}}{1 + \frac{w}{p}} = \frac{M + wT}{p + w}.$$

Thus, L is increasing in the wage if $pT>M$, that is, if M is sufficiently small that one can't buy all one's needs and not work at all. (This is the only reasonable case for this utility function.) With strong complements between goods and leisure, an increase in the wage induces fewer hours worked.

Example (Cobb-Douglas): $h(L) = \left(\frac{M + w(T-L)}{p}\right)^{\alpha} L^{1-\alpha}$.

The first order condition gives

$$0 = h'(L) = -\alpha\left(\frac{M + w(T-L)}{p}\right)^{\alpha-1} L^{1-\alpha}\frac{w}{p} + (1-\alpha)\left(\frac{M + w(T-L)}{p}\right)^{\alpha} L^{-\alpha}$$

or

$$\alpha L \frac{w}{p} = (1-\alpha)\frac{M + w(T-L)}{p}$$

$$\frac{w}{p} L = (1-\alpha)\frac{M + wT}{p}$$

or $L = (1-\alpha)\left(\frac{M}{w} + T\right)$

If M is high enough, the consumer doesn't work but takes $L=T$; otherwise, the equation gives the leisure, and labor supply is given by

$$T - L = Max\left\{0, \alpha T - (1-\alpha)\left(M/w\right)\right\}$$

Labor supply increases with the wage, no matter how high the wage goes.

5.2.2.1 (Exercise) Show that an increase in the wage increases the consumption of goods, that is, x increases when the wage increases.

The wage affects not just the price of leisure, but also the income level; this makes it possible that the income effect of a wage increase dominates the substitution effect. Moreover, we saw that this is more likely when the consumption of goods takes time, that is, the goods and leisure are complements.

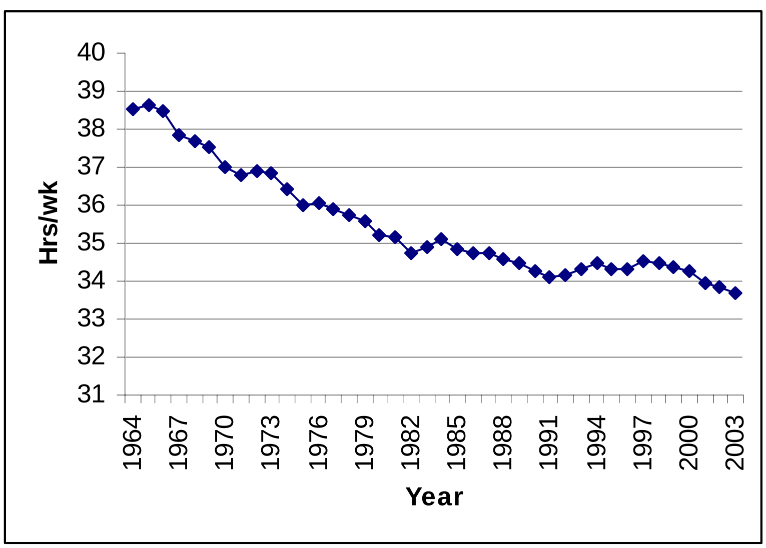

Figure 5-15: Hours per Week

As a practical matter, for most developed nations, increases in wages are associated with fewer hours worked. The average workweek prior to 1950 was 55 hours, which fell to 40 hours by the mid-1950s. The workweek has gradually declined since then, as Figure 5-15 illustrates.

Thought Question: Does a bequest motive – the desire to give money to others – change the likelihood that goods and leisure are complements?

5.2.3 Compensating Differentials

A number of physicists have changed careers, to become researchers in finance or financial economics. Research in finance pays substantially better than research in physics, and yet requires many of the same mathematical skills like stochastic calculus. Physicists who see their former colleagues driving Porsches and buying summer houses are understandably annoyed that finance research – which is intellectually no more difficult or challenging than physics – pays so much better. Indeed, some physicists say that other fields – finance, economics, and law – "shouldn't" pay more than physics.

The difference in income between physics researchers and finance researchers is an example of a *compensating differential.* A compensating differential is income or costs that equalize different choices. There are individuals who could become either physicists or finance researchers. At equal income, too many choose physics and too few choose finance, in the sense that there is a surplus of physicists, and a shortage of finance researchers. Finance salaries must exceed physics salaries in order to induce some of the researchers capable of doing either one to switch to finance, which compensates those individuals for doing the less desirable task.

Jobs that are dangerous or unpleasant must pay more than jobs requiring similar skills but without the bad attributes. Thus, oil field workers in Alaska's North Slope, well above the Arctic Circle, earn a premium over workers in similar jobs in Houston, Texas. The premium – or differential pay – must be such that the marginal worker is indifferent between the two choices – the extra pay compensates the worker for the adverse working conditions. This is why it is known in economics jargon by the phrase of a *compensating* differential.

The high salaries earned by professional basketball players are not compensating differentials. These salaries are not created by a need to induce tall people to choose basketball over alternative jobs like painting ceilings, but instead are payments that reflect the rarity of the skills and abilities involved. Compensating differentials are determined by alternatives, not by direct scarcity. Professional basketball players are well-paid for the same reason that Picasso's paintings are expensive: there aren't very many of them relative to demand.

A compensating differential is a feature of other choices as well as career choices. For example, many people would like to live in California, for its weather and scenic beauty. Given the desirability of California over, say, Lincoln, Nebraska or Rochester, New York, there must be a compensating differential for living in Rochester, and two significant ones are air quality and housing prices. Air quality worsens as populations rise, thus tending to create a compensating differential. In addition, the increase in housing prices also tends to compensate – housing is inexpensive in Rochester, at least compared to California.[53]

[53] There are other compensations besides housing to living in Rochester – cross-country skiing, proximity to mountains and lakes. Generally employment is only a temporary factor that might compensate,

Housing prices also compensate for location within a city. For most people, it is more convenient – both in commuting time and for services – to be located near the central business district than in the outlying suburbs. The main compensating differentials are school quality, crime rates, and housing prices. We can illustrate the ideas with a simple model of a city.

5.2.4 Urban Real Estate Prices

An important point to understand is that the good in limited supply in cities is not physical structures like houses, but the land on which the houses sit. The cost of building a house in Los Angeles is quite similar to the cost of building a house in Rochester, New York. The big difference is the price of land. A $1 million house in Los Angeles might be a $400,000 house sitting on a $600,000 parcel of land. The same house in Rochester might be $500,000 – a $400,000 house on a $100,000 parcel of land.

Usually, land is what fluctuates in value, rather than the price of the house that sits on the land. When the newspaper reports that house prices rose, in fact what rose was land prices, for the price of housing has changed only at a slow pace, reflecting increased wages of house builders and changes in the price of lumber and other inputs. These do change, but historically the changes have been small compared to the price of land.

We can construct a simple model of a city to illustrate the determination of land prices. Suppose the city is constructed in a flat plane. People work at the origin (0,0). This simplifying assumption is intended to capture the fact that a relatively small, central portion of most cities involves business, with a large area given over to housing. The assumption is extreme, but not unreasonable as a description of some cities.

Suppose commuting times are proportional to distance from the origin. Let $c(t)$ be the cost to the person of a commute of time t, and let the time taken be $t = \lambda r$, where r is the distance. The function c should reflect both the transportation costs and the value of time lost. The parameter λ accounts for the inverse of the speed in commuting, with a higher λ indicating slower commuting. In addition, we assume that people occupy a constant amount of land. This assumption is clearly wrong empirically, and we will consider making house size a choice variable later.

A person choosing a house priced at $p(r)$ at distance r thus pays $c(\lambda r) + p(r)$ for the combination of housing and transportation. People will choose the lowest cost alternative. If people have identical preferences about housing and commuting, then house prices p will depend on distance, and will be determined by $c(\lambda r) + p(r)$ equal to a constant, so that people are indifferent to the distance from the city center – decreased commute time is exactly compensated by increased house prices.

---

because employment tends to be mobile, too, and move to the location the workers prefer, when that is possible. It is not possible on Alaska's North Slope.

The remaining piece of the model is to figure out the constant. To do this, we need to figure out the area of the city. If the total population is N, and people occupy an area of one per person, the city size r_{max} satisfies $N = \pi r_{\max}^2$, and thus

$$r_{\max} = \sqrt{\frac{N}{\pi}}$$

At the edge of the city, the value of land is given by some other use, like agriculture. From the perspective of the determinant of the city's prices, this value is approximately constant. As the city takes more land, the change in agricultural land is a very small portion of the total land used for agriculture. Let the value of agricultural land be v per housing unit size. Then the price of housing $p(r_{max}) = v$, because that is the value of land at the edge of the city. This lets us compute the price of all housing in the city:

$$\mathrm{c}(\lambda r) + \mathrm{p(r)} = \mathrm{c}(\lambda r_{\max}) + \mathrm{p}(r_{\max}) = \mathrm{c}(\lambda r_{\max}) + v = \mathrm{c}\left(\lambda\sqrt{\frac{N}{\pi}}\right) + v$$

or

$$\mathrm{p(r)} = \mathrm{c}\left(\lambda\sqrt{\frac{N}{\pi}}\right) + v - \mathrm{c}(\lambda r) \quad .$$

This equation produces housing prices like those illustrated in Figure 5-16, where the peak is the city center. The height of the figure indicates the price of housing.

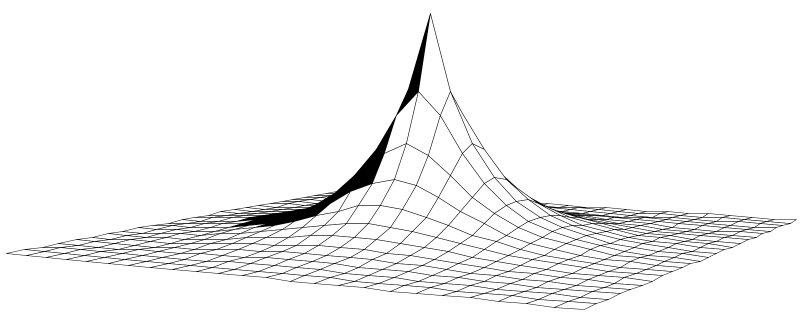

Figure 5-16: House Price Gradient

It is straightforward to verify that house prices increase in the population N and the commuting time parameter λ, as one would expect. To quantify the predictions, we consider a city with a population of 1,000,000, a population density of 10,000 per square mile, and an agricultural use value of $6 million per square mile. To translate these assumptions into the model's structure, first note that a population density of 10,000 per square mile creates a fictitious "unit of measure" of about 52.8 feet, which we'll call a purlong, so that there is one person per square purlong (2788 square feet). Then the agricultural value of a property is v = $600 per square purlong. Note that this density requires a city of radius r_{max} equal to 564 purlongs, which is 5.64 miles.

The only remaining structure to identify in the model is the commuting cost *c*. To simplify the calculations, let *c* be linear. Suppose that the daily cost of commuting is $2 per mile (roundtrip), so that the present value of daily commuting costs in perpetuity is about $10,000 per mile.[54] This translates into a cost of commuting of $100.00 per purlong. Thus, we obtain

$$p(\mathrm{r}) = \mathrm{c}\left(\lambda\sqrt{\frac{N}{\pi}}\right) + v - \mathrm{c}(\lambda r) = \$100\left(\sqrt{\frac{N}{\pi}} - r\right) + \$600 = \$57{,}000 - \$100r .$$

Thus, the same 2788 square foot property at the city edge sells for $600, versus $57,000 less than six miles away at the city center. With reasonable parameters, this model readily creates dramatic differences in land prices, based purely on commuting time.

As constructed, a quadrupling of population approximately doubles the price of land in the central city. This probably understates the change, since a doubling of the population would likely increase road congestion, increasing λ and further increasing the price of central city real estate.

As presented, the model contains three major unrealistic assumptions. First, everyone lives in an identically-sized piece of land. In fact, however, the amount of land used tends to fall as prices rise. At $53 per square foot, most of us buy a lot less land than at twenty cents per square foot. As a practical matter, the reduction of land per capita is accomplished both through smaller housing units and through taller buildings, which produce more housing floor space per acre of land. Second, people have distinct preferences, and the disutility of commuting, as well as the value of increased space, vary with the individual. Third, congestion levels are generally endogenous – the more people that live between two points, the greater the traffic density and consequently the lower the level of λ. The first two problems arise because of the simplistic nature of consumer preferences embedded in the model, while the third is an equilibrium issue requiring consideration of transportation choices.

This model can readily be extended to incorporate different types of people, different housing sizes, and endogenous congestion. To illustrate such generalizations, consider making the housing size endogenous. Suppose preferences are represented by the utility function:

$$u = H^{\alpha} - \lambda r - p(r)H ,$$

where *H* is the house size that the person chooses, and *r* is the distance they choose. This adaptation of the model reflects two issues. First, the transport cost has been set to be linear in distance, for simplicity. Second, the marginal value of housing decreases in the house size, but the value of housing doesn't depend on distance from the center. For

[54] Figure 250 working days per year, for an annual cost of about $500 per mile, yielding a present value at 5% interest of $10,000. See Section 4.3.1. With a time value of $25 per hour, and an average speed of 40 mph (1.5 minutes per mile), the time cost is 62.5 cents per minute. Automobile costs (gas, car depreciation, insurance) are about 35-40 cents per mile. Thus the total is around $1 per mile, which doubles with roundtrips.

these preferences to make sense, $\alpha<1$ (otherwise either zero or an infinite house size emerges). A person with these preferences optimally would choose a house size of

$$H=\left(\frac{\alpha}{p(r)}\right)^{\frac{1}{1-\alpha}}$$

resulting in utility

$$u^*=\left(\alpha^{\frac{\alpha}{1-\alpha}}-\alpha^{\frac{1}{1-\alpha}}\right)p(r)^{\frac{-\alpha}{1-\alpha}}-\lambda r$$

Utility at every location is constant, so $\left(\frac{\alpha^{\frac{\alpha}{1-\alpha}}-\alpha^{\frac{1}{1-\alpha}}}{u^*+\lambda r}\right)^{\frac{1-\alpha}{\alpha}}=p(r)$.

A valuable attribute of the form of the equation for p is that the general form depends on the equilibrium values only through the single number u^*. This functional form produces the same qualitative shapes as in Figure 5-16. Using the form, we can solve for the housing size H.

$$H(r)=\left(\frac{\alpha}{p(r)}\right)^{\frac{1}{1-\alpha}}=\alpha^{\frac{1}{1-\alpha}}\left(\frac{u^*+\lambda r}{\alpha^{\frac{\alpha}{1-\alpha}}-\alpha^{\frac{1}{1-\alpha}}}\right)^{\frac{1}{\alpha}}=\left(\frac{u^*+\lambda r}{\alpha^{-1}-1}\right)^{\frac{1}{\alpha}}=\left(\frac{\alpha}{1-\alpha}(u^*+\lambda r)\right)^{\frac{1}{\alpha}}.$$

The space in the interval $[r, r+\Delta]$ is $\pi(2r\Delta+\Delta^2)$. In this interval, there are approximately $\frac{\pi(2r\Delta+\Delta^2)}{H(r)}=\pi(2r\Delta+\Delta^2)\left(\frac{1-\alpha}{\alpha(u^*+\lambda r)}\right)^{\frac{1}{\alpha}}$ people. Thus, the number of people within r_{max} of the city center is

$$\int_0^{r_{max}} 2\pi r\left(\frac{1-\alpha}{\alpha(u^*+\lambda r)}\right)^{\frac{1}{\alpha}}dr=N \quad .$$

This equation, when combined with the value of land on the periphery:

$$v = p(r_{\max}) = \left(\frac{\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}}{u^* + \lambda r_{\max}} \right)^{\frac{1-\alpha}{\alpha}}$$

jointly determine $r_{\max}$ and u^*.

5.2.4.1 (Exercise) For the case of $\alpha = ½$, solve for the equilibrium values of u^* and $r_{\max}$.

When different people have different preferences, the people with the highest disutility of commuting will tend to live closer to the city center. These tend to be people with the highest wages, since one of the costs of commuting is time that could have been spent working.

5.2.5 Dynamic Choice

The consumption of goods doesn't take place in a single instance, but over time. How does time enter into choice? We're going to simplify the problem a bit, and focus only on consumption and set aside working for the time being. Let x_1 be consumption in the first period, x_2 in the second period. Suppose the value of consumption is the same in each period, so that

$$u(x_1, x_2) = v(x_1) + \delta v(x_2),$$

where δ is called the rate of "pure" time preference. The consumer is expected to have income M_1 in the first period and M_2 in the second. There is a market for loaning and borrowing, which we assume has a common interest rate r.

The consumer's budget constraint, then, can be written

$$(1+r)(M_1 - x_1) = x_2 - M_2.$$

This equation says that the net savings in period 1, plus the interest on the net savings in period 1 equals the net expenditure in period 2. This is because whatever is saved in period 1 earns interest and can then be spent in period 2; alternatively, whatever is borrowed in period 1 must be paid back with interest in period 2. Rewriting the constraint:

$$(1+r)x_1 + x_2 = (1+r)M_1 + M_2.$$

This equation is known as the *intertemporal budget constraint*. It has two immediate consequences. First, $1+r$ is the price of period 2 consumption in terms of period 1 consumption. Thus, the interest rate gives the relative prices. Second, the relevant income is "permanent income" rather than "current income." That is, a change in incomes that leaves the present value of income the same should have no effect on the choice of consumption.

Once again, as with the labor supply, a change in the interest rate affects not just the price of consumption, but also the budget for consumption. Put another way, an increase in the interest rate represents an increase in budget for net savers, but a decrease in budget for net borrowers.

As always, we rewrite the optimization problem to eliminate one of the variables, to obtain

$$u = v(x_1) + \delta v\big((1+r)(M_1 - x_1) + M_2\big)$$

Thus the first order conditions yield

$$0 = v'(x_1) - (1+r)\delta v'(x_2)$$

This condition says that the marginal value of consumption in period 1, $v'(x_1)$, equals the marginal value of consumption in period 2, $\delta v'(x_2)$, times the interest factor. That is, the marginal present values are equated. Note that the consumer's private time preference, δ, need not be related to the interest rate. If the consumer values period 1 consumption more than does the market, so $\delta(1+r) < 1$, then $v'(x_1) < v'(x_2)$, that is, the consumer consumes more in period 1 than in period 2.[55] Similarly, if the consumer's discount of future consumption is exactly equal to the market discount, $\delta(1+r) = 1$, the consumer will consume the same amount in both periods. Finally, if the consumer values period 1 consumption less than the market, $\delta(1+r) > 1$, the consumer will consume more in period 2. In this case, the consumer is more patient than the market.

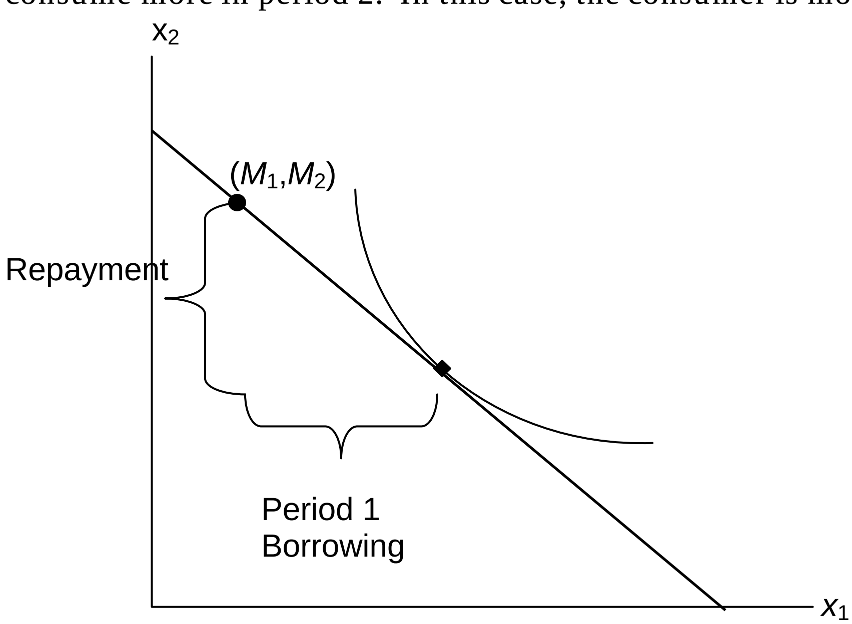

Figure 5-17: Borrowing and Lending

[55] As usual, we are assuming that utility is concave, which in this instance means the second derivative of *v* is negative, which means the derivative of *v* is decreasing. In addition, to insure an interior solution, it is useful to require the *Inada* conditions: $v'(0) = \infty$, $v'(\infty) = 0$.

Whether the consumer is a net lender or borrower depends not just on the preference for earlier versus later consumption, but also on incomes. This is illustrated in Figure 5-17. In this figure, the consumer's income mostly comes in the second period. As a consequence, the consumer borrows in the first period, and repays in the second period.

The effect of an interest rate increase is to pivot the budget constraint around the point (M_1, M_2). Note that this point is *always* feasible – that is, it is feasible to consume one's own endowment. The effect of an increase in the interest rate is going to depend on whether the consumer is a borrower or a lender. As Figure 5-18 illustrates, the net borrower borrows less in the first period – the price of first period consumption has risen and the borrower's wealth has fallen. It is not clear whether the borrower consumes less in the second period because the price of second period consumption has fallen even though wealth has fallen, too, two conflicting effects.

An increase in interest rates is a benefit to a net lender. The lender has more income, and the price of period 2 consumption has fallen. Thus the lender must consume more in the second period, but only consumes more in the first period (lends less) if the income effect outweighs the substitution effect. This is illustrated in Figure 5-19.

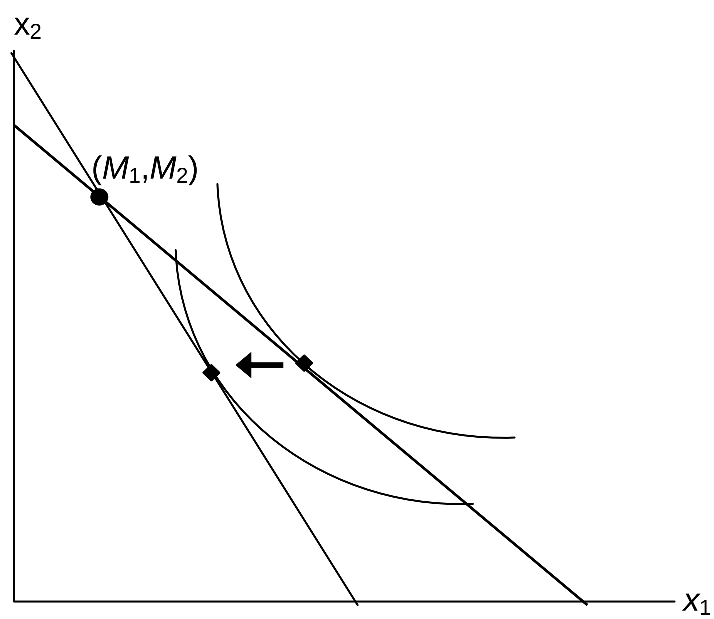

Figure 5-18: Interest Rate Change

The government from time to time will rebate a portion of taxes to "stimulate" the economy. An important aspect of the effects of such a tax rebate is the effect to which consumers will spend the rebate, versus savings the rebate, because the stimulative effects of spending are thought to be larger than the stimulative effects of savings.[56] The

[56] This belief shouldn't be accepted as necessarily true; it was based on a model that has since been widely rejected by the majority of economists. The general idea is that spending creates demand for goods, thus

theory suggests how people will react to a "one-time" or transitory tax rebate, compared to a permanent lowering of taxes. In particular, the budget constraint for the consumer spreads lifetime income over the lifetime. Thus, for an average consumer that might spend a present value of $750,000 over a lifetime, a $1,000 rebate is small potatoes. On the other hand, a $1,000/year reduction is worth $20,000 or so over the lifetime, which should have twenty times the effect of the transitory change on the current expenditure.

Tax rebates are not the only way we receive one-time payments. Money can be found, or lost, and we can have unexpected costs or windfall gifts. From an intertemporal budget constraint perspective, these transitory effects have little significance, and thus the theory suggests people shouldn't spend much of a windfall gain in the current year, nor cut back significantly when they have a moderately-sized unexpected cost.

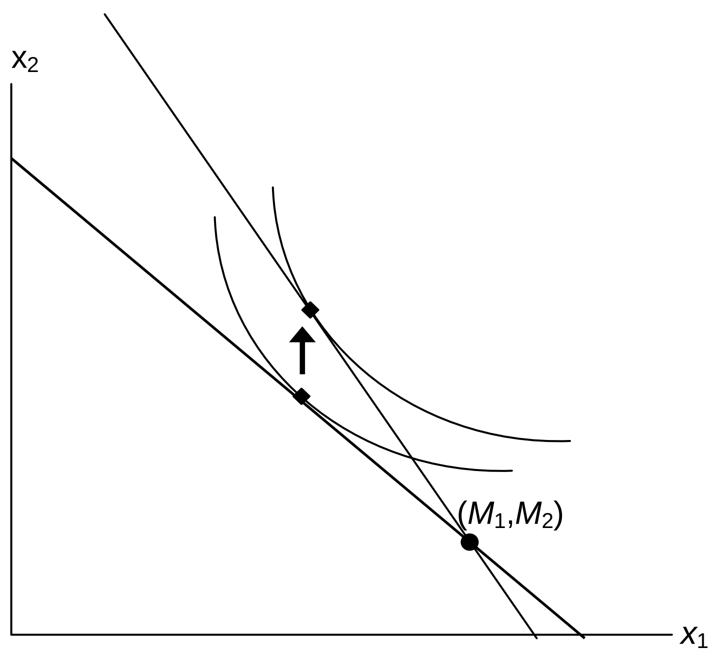

Figure 5-19: Interest Rate Increase on Lenders

As a practical matter, most individuals can't borrow at the same rate at which they lend. Many students borrow on credit cards at very high interest rates, and obtain a fraction of that in interest on savings. That is to say, borrowers and lenders face different interest rates. This situation is readily explored with a diagram like Figure 5-20. The cost of a first period loan is a relatively high loss of x_2, and similarly the value of first period savings is a much more modest increase in second period consumption. Such effects tend to favor "neither a borrower nor a lender be," as Shakespeare recommends, although it is still possible for the consumer to optimally borrow in the first period (e.g. if M_1=0) or in the second period (if M_2 is small relative to M_1).

encouraging business investment in production. However, savings encourage investment by producing loanable funds, so it isn't at all obvious whether spending or savings have a larger effect.

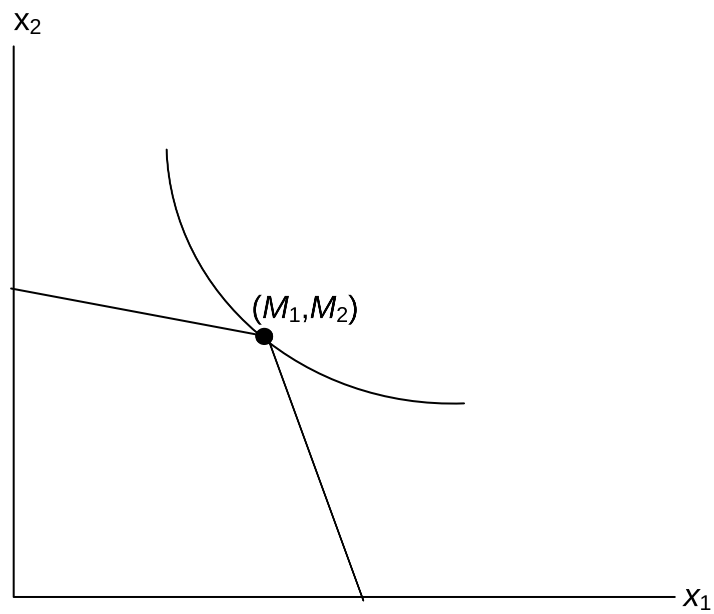

Figure 5-20: Different Rates for Borrowing and Lending

Differences in interest rates causes transitory changes in income to have much larger effects than the intertemporal budget constraint would suggest, and may go a long way to explaining why people don't save much of a windfall gain, and suffer a lot temporarily, rather than a little for a long time, when they have unexpected expenses. This is illustrated in Figure 5-21.

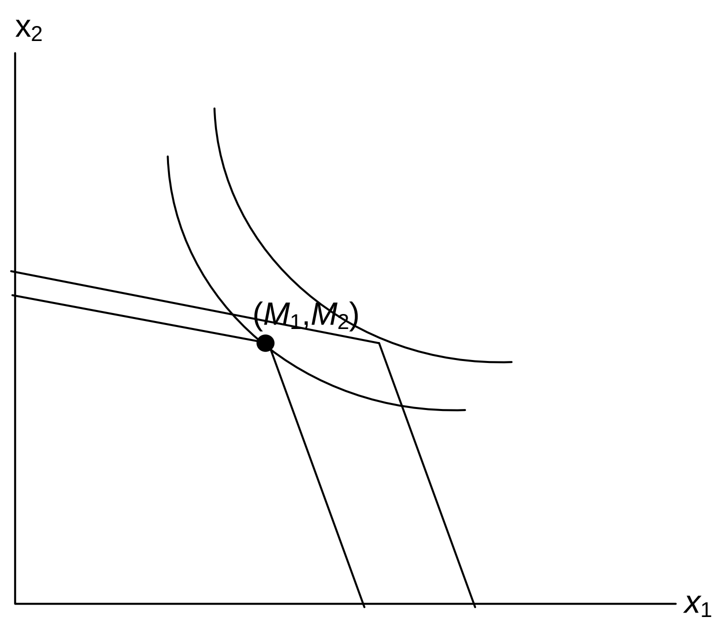

Figure 5-21: The Effect of a Transitory Income Increase

5.2.6 Risk

There are many risks in life, even if one doesn't add to these risks by intentionally buying lottery tickets. Gasoline prices go up and down, the demand for people trained in your major fluctuates, house prices change. How do people value gambles? The starting point for the investigation is the *von Neumann-Morgenstern* utility function. The idea of a von Neumann-Morgenstern utility function for a given person is that for each possible outcome x, there is a value $v(x)$ assigned by the person, and the average value of v is the value the person assigns to the risky outcome. This is a "state of the world" approach, in the sense that each of the outcomes is associated with a state of the world, and the person maximizes the expected value of the various possible states of the world. Value here doesn't mean a money value, but a psychic value or utility.

To illustrate the assumption, consider equal probabilities of winning \$100 and winning \$200. The expected outcome of this gamble is \$150 – the average of \$100 and \$200. However, the expected value of the outcome could be anything between the value of \$100 and the value of \$200. The von Neumann-Morgenstern utility is ½ v(\$100) + ½ v(\$200).

The von Neumann-Morgenstern formulation has certain advantages, including the logic that what matters is the average value of the outcome. On the other hand, in many tests, people behave in ways not consistent with the theory.[57] Nevertheless, the von Neumann approach is the prevailing model of behavior under risk.

To introduce the theory, we will consider only money outcomes, and mostly the case of two money outcomes. The person has a Neumann-Morgenstern utility function v of these outcomes. If the possible outcomes are $x_1, x_2, \ldots, x_n$ and these occur with probability $\pi_1, \pi_2, \ldots, \pi_n$ respectively, the consumer's utility is

$$u = \pi_1 v(x_1) + \pi_2 v(x_2) + \ldots + \pi_n v(x_n) = \sum_{i=1}^{n} \pi_i v(x_i)$$

This is the meaning of "having a von Neumann-Morgenstern utility function" – that utility can be written in this weighted sum form.

The first insight that flows from this definition is that a individual dislikes risk if v is concave. To see this, note that the definition of concavity posits that v is concave if, for all π in [0,1], and all values x_1 and x_2,

$$v(\pi x_1 + (1-\pi)x_2) \geq \pi v(x_1) + (1-\pi) v(x_2)$$

For smoothly differentiable functions, concavity is equivalent to a second derivative that is not positive. Using induction, the definition of concavity can be generalized to show:

[57] For example, people tend to react more strongly to very unlikely events than is consistent with the theory.

$$v(\pi_1 x_1 + \pi_2 x_2 + \ldots + \pi_n x_n) \geq \pi_1 v(x_1) + \pi_2 v(x_2) + \ldots + \pi_n v(x_n)$$

That is, a consumer with concave value function prefers the average outcome to the random outcome. This is illustrated in Figure 5-22. There are two possible outcomes, x_1 and x_2. The value x_1 occurs with probability π and x_2 with probability $1-\pi$. This means the average or expected outcome is $\pi x_1+(1-\pi)x_2$. The value $v(\pi x_1+(1-\pi)x_2)$ is the value at the expected outcome $\pi x_1+(1-\pi)x_2$, while $\pi v(x_1)+(1-\pi)v(x_2)$ is the average of the value of the outcome. As is plainly visible in the picture, concavity makes the average outcome preferable to the random outcome. People with concave von Neumann-Morganstern utility functions are known as *risk averse* people.

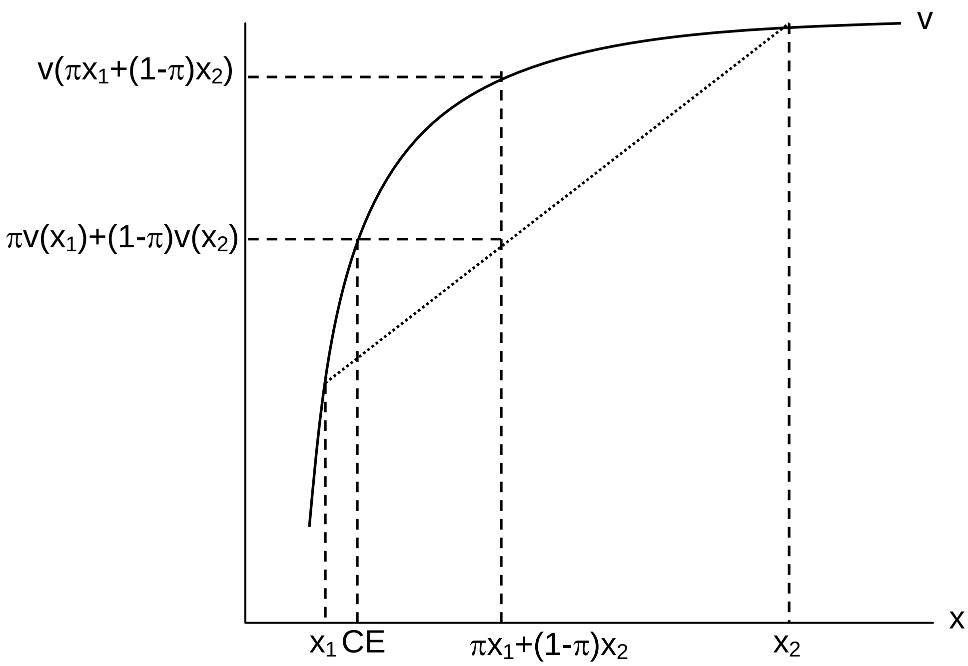

Figure 5-22: Expected Utility and Certainty Equivalents

A useful concept is the *certainty equivalent* of a gamble. The certainty equivalent is an amount of money that provides equal utility to the random payoff of the gamble. The certainty equivalent is labeled *CE* in the diagram. Note that CE is less than the expected outcome, if the person is risk averse. This is because risk averse individuals prefer the expected outcome to the risky outcome.

The *risk premium* is defined to be the difference between the expected payoff (in the graph, this is expressed as $\pi x_1 + (1 - \pi)x_2$) and the certainty equivalent. This is the cost of risk – it is the amount of money an individual would be willing to pay to avoid risk. This means as well that the risk premium is the value of insurance. How does the risk premium of a given gamble change when the base wealth is increased? It can be shown that the risk premium falls as wealth increases for any gamble if, and only if,[58]

[58] R. Preston McAfee and Daniel Vincent, The Price Decline Anomaly, *Journal of Economic Theory 60*, June, 1993, 191-212.

$$-\frac{v''(x)}{v'(x)} \text{ is decreasing.}$$

The measure $\rho(x) = -\frac{v''(x)}{v'(x)}$ is known as the Arrow-Pratt[59] measure of risk aversion, and also as the measure of absolute risk aversion. To get an idea why this measure matters, consider a quadratic approximation to v. Let μ be the expected value and σ^2 be the expected value of $(x - \mu)^2$. Then we can approximate v(CE) two different ways.

$$v(\mu) + v'(\mu)(CE - \mu) \approx v(CE) = E\{v(x)\} \approx E\{v(\mu) + v'(\mu)(x - \mu) + ½\, v''(\mu)(x - \mu)^2\},$$

thus

$$v(\mu) + v'(\mu)(CE - \mu) \approx E\{v(\mu) + v'(\mu)(x - \mu) + ½\, v''(\mu)(x - \mu)^2\}.$$

Canceling $v(\mu)$ from both sides and noting that the average value of x is μ, so $E(x-\mu)=0$, we have

$$v'(\mu)(CE - \mu) \approx ½\, v''(\mu)\sigma^2.$$

Then, dividing by $v'(x)$,

$$\mu - CE \approx -½ \frac{v''(\mu)}{v'(\mu)}\sigma^2 = ½\rho(\mu)\sigma^2.$$

That is, the risk premium, the difference between the average outcome and the certainty equivalent, is approximately equal to the Arrow-Pratt measure, times half the variance, at least when the variance is small.

5.2.6.1 (Exercise) Use a quadratic approximation on both sides to sharpen the estimate of the risk premium. First, note

$$\begin{aligned} v(\mu) + v'(\mu)(CE - \mu) + ½\, v''(\mu)(CE - \mu)^2 &\approx v(CE) \\ = E\{v(x)\} &\approx E\{v(\mu) + v'(\mu)(x - \mu) + ½\, v''(\mu)(x - \mu)^2\}. \end{aligned}$$

Conclude that $\mu - CE \approx \frac{1}{\rho}\left(\sqrt{1 + \rho^2\sigma^2} - 1\right)$. This approximation is exact to the second order.

The translation of risk into dollars, by way of a risk premium, can be assessed even for large gambles if we are willing to make some technical assumptions. Suppose the utility

[59] The measure was named after its discoverers Nobel laureate Kenneth Arrow and John Pratt.

has constant absolute risk aversion or CARA, that is $\rho = -\frac{v''(x)}{v'(x)}$ is a constant. This turns out to imply, after setting the utility of zero to zero, that

$$v(x) = \frac{1}{\rho}(1 - e^{-\rho x}).$$

(This formulation is derived by setting $v(0)=0$ handling the case of $\rho=0$ with appropriate limits.) Now also assume that the gamble x is normally distributed with mean μ and variance σ^2. Then the expected value of v(x) is

$$Ev(x) = \frac{1}{\rho}\left(1 - e^{-\rho\left(\mu - \frac{\rho}{2}\sigma^2\right)}\right).$$

It is an immediate result from this formula that the certainty equivalent, with CARA preferences and normal risks, is $\mu - \frac{\rho}{2}\sigma^2$. Hence the risk premium of a normal distribution for a CARA individual is $\frac{\rho}{2}\sigma^2$. This formulation will appear when we consider agency theory and the challenges of motivating a risk averse employee when outcomes have a substantial random component.

An important aspect of CARA with normally distributed risks is that the preferences of the consumer are linear in the mean of the gamble and the variance. In fact, given a choice of gambles, the consumer selects the one with the highest value of $\mu - \frac{\rho}{2}\sigma^2$. Such preferences are often called "mean variance preferences," and they comprise the foundation of modern finance theory.

5.2.6.2 (Exercise) Suppose $u(x) = x^{0.95}$ for a consumer with a wealth level of \$50,000. Consider a gamble with equal probability of winning \$100 and losing \$100 and compute the risk premium associated with the gamble.

5.2.6.3 (Exercise) Suppose $u(x) = x^{0.99}$ for a consumer with a wealth level of \$100,000. A lottery ticket costs \$1 and pays \$5,000,000 with the probability $\frac{1}{10{,}000{,}000}$. Compute the certainty equivalent of the lottery ticket.

5.2.6.4 (Exercise) The return on U.S. government treasury investments is approximately 3%. Thus, a \$1 investment returns \$1.03 after one year. Treat this return as risk-free. The stock market (S&P 500) returns 7% on average and has a variance that is around 16% (the variance of return on a \$1 investment is \$0.16). Compute the value of ρ for a CARA individual. What is the risk premium associated equal probabilities of a \$100 gain or loss given the value of ρ?

5.2.7 Search

In most communities, every Wednesday grocery stores advertise sale prices in a newspaper insert, and these prices vary from week to week. Prices can vary a lot from week to week and from store to store. The price of gasoline varies as much as fifteen cents per gallon in a one mile radius. Decide you want a specific Sony television, and you may see distinct prices at Best Buy, Circuit City, and other electronics retailers. For many goods and services, there is substantial variation in prices, which implies that there are gains for buyers to search for the best price.

The theory of consumer search behavior is just a little bit arcane, but the basic insight will be intuitive enough. The general idea is that, from the perspective of a buyer, the price that is offered is random, and has a probability density function $f(p)$. If a consumer faces a cost of search (e.g. if you have to visit a store, in person, telephonically or virtually, the cost includes your time and any other costs necessary to obtain a price quote), the consumer will set a *reservation price*, which is a maximum price they will pay without visiting another store. That is, if a store offers a price below p^*, the consumer will buy, and otherwise they will visit another store, hoping for a better price.

Call the reservation price p^* and suppose that the cost of search is c. Let $J(p^*)$ represent the expected total cost of purchase (including search costs). Then J must equal

$$J(p^*) = c + \int_0^{p^*} pf(p)dp + \int_{p^*}^{\infty} J(p^*)f(p)dp .$$

This equation arises because the current draw (which costs c) could either result in a price less than p^*, in which case observed price, with density f, will determine the price paid p, or the price will be too high, in which case the consumer is going to take another draw, at cost c, and on average get the average price $J(p^*)$. It is useful to introduce the cumulative distribution function F, with $F(x) = \int_0^x f(p)dp$. Note that something has to happen, so $F(\infty)=1$.

We can solve the equality for $J(p^*)$,

$$J(p^*) = \frac{\int_0^{p^*} pf(p)dp + c}{F(p^*)}.$$

This expression has a simple interpretation. The expected price $J(p^*)$ is composed of two terms. The first is the expected price, which is $\int_0^{p^*} p\frac{f(p)}{F(p^*)}dp$. This has the interpretation of the average price conditional on that price being less than p^*. This is because $\frac{f(p)}{F(p^*)}$ is in fact the density of the random variable which is the price given that the price is less than p^*. The second term is $\frac{c}{F(p^*)}$. This is the expected search costs, and it arises because $\frac{1}{F(p^*)}$ is the expected number of searches. This arises because the odds of getting a price low enough to be acceptable is $F(p^*)$. There is a general statistical property underlying the number of searches. Consider a basketball player who successfully shoots a free throw with probability y. How many throws on average must he throw to sink one basket? The answer is $1/y$. To see this, note that the probability that exactly n throws are required is $(1-y)^{n-1}y$. This is because n are required means n-1 must fail (probability $(1-y)^{n-1}$) and then the remaining one go in, with probability y. Thus, the expected number of throws is

$$y + 2(1-y)y + 3(1-y)^2 y + 4(1-y)^3 y + \ldots$$

$$= y(1 + 2(1-y) + 3(1-y)^2 + 4(1-y)^3 + \ldots)$$

$$= y\big((1 + (1-y) + (1-y)^2 + (1-y)^3 + \ldots) + (1-y)(1 + (1-y) + (1-y)^2 + (1-y)^3 + \ldots)$$

$$+ (1-y)^2(1 + (1-y) + (1-y)^2 + (1-y)^3 + \ldots) + (1-y)^3(1 + (1-y) + (1-y)^2 + \ldots) + \ldots$$

$$= y\left(\frac{1}{y} + (1-y)\frac{1}{y} + (1-y)^2\frac{1}{y} + (1-y)^3\frac{1}{y} + \ldots\right) = \frac{1}{y}.$$

Our problem has the same logic, where a successful basketball throw corresponds to finding a price less than p^*.

The expected total cost of purchase, given a reservation price p^* is given by

$$J(p^*) = \frac{\int_0^{p^*} pf(p)dp + c}{F(p^*)}.$$

But what value of p^* minimizes cost? Let's start by differentiating:

$$J'(p^*) = p^* \frac{f(p^*)}{F(p^*)} - \frac{f(p^*)\int_0^{p^*} pf(p)dp + c}{F(p^*)^2}$$

$$= \frac{f(p^*)}{F(p^*)}\left(p^* - \frac{\int_0^{p^*} pf(p)dp + c}{F(p^*)} \right) = \frac{f(p^*)}{F(p^*)}\left(p^* - J(p^*)\right).$$

Thus, if $p^* < J(p^*)$, J is decreasing, and it lowers cost to increase p^*. Similarly, if $p^* > J(p^*)$, J is increasing in p^*, and it reduces cost to decrease p^*. Thus, minimization occurs at a point where $p^* = J(p^*)$.

Moreover, there is only one such solution to the equation $p^* = J(p^*)$ in the range where f is positive. To see this, note that at any solution to the equation $p^* = J(p^*)$, $J'(p^*) = 0$ and

$$J''(p^*) = \frac{d}{dp^*}\left(\frac{f(p^*)}{F(p^*)}\left(p^* - J(p^*)\right) \right)$$

$$= \left(\frac{d}{dp^*}\frac{f(p^*)}{F(p^*)} \right)(p^* - J(p^*)) + \frac{f(p^*)}{F(p^*)}(1 - J'(p^*)) = \frac{f(p^*)}{F(p^*)} > 0.$$

This means that J takes a minimum at this value, since its first derivative is zero and its second derivative is positive, and that is true about any solution to $p^* = J(p^*)$. Were there to be two such solutions, J' would have to be both positive and negative on the interval between them, since J is increasing to the right of the first (lower) one, and decreasing to the left of the second (higher) one. Consequently, the equation $p^* = J(p^*)$ has a unique solution that minimizes the cost of purchase.

Consumer search to minimize cost dictates setting a reservation price equal to the expected total cost of purchasing the good, and purchasing whenever the price offered is lower than that level. That is, it is not sensible to "hold out" for a price lower than what you expect to pay on average, although this might be well useful in a bargaining context rather than in a store searching context.

Example (Uniform): Suppose prices are uniformly distributed on the interval $[a,b]$. For p^* in this interval,

$$J(p^*) = \frac{\int_0^{p^*} pf(p)dp + c}{F(p^*)} = \frac{\int_a^{p^*} p\frac{dp}{b-a} + c}{\frac{p^* - a}{b-a}}$$

$$= \frac{½(p^{*2} - a^2) + c(b-a)}{p^* - a} = ½(p^* + a) + \frac{c(b-a)}{p^* - a}.$$

Thus, the first order condition for minimizing cost is

$$0 = J'(p^*) = ½ - \frac{c(b-a)}{(p^*-a)^2}\text{, implying } p^* = a + \sqrt{2c(b-a)}.$$

There are a couple of interesting observations about this solution. First, not surprisingly, as $c \to 0$, $p^* \to a$, that is, as the search costs go to zero, one holds out for the lowest possible price. This is sensible in the context of the model, but in the real search situations delay may also have a cost that isn't modeled here. Second, $p^* < b$, the maximum price, if $2c<(b - a)$. Put another way, if the *most* you can save by a search is twice the search cost, don't search, because the expected gains from search will be half the maximum gains (thanks to the uniform distribution) and the search unprofitable.

The third observation, which is much more general than the specific uniform example, is that the expected price is a concave function of the cost of search (second derivative negative). That is in fact true for any distribution. To see this, define a function

$$H(c) = \min_{p^*} J(p^*) = \min_{p^*} \frac{\int_0^{p^*} pf(p)dp + c}{F(p^*)}.$$

Since $J'(p^*) = 0$,

$$H'(c) = \frac{\partial}{\partial c} J(p^*) = \frac{1}{F(p^*)}.$$

It then needs only a modest effort to show p^* is increasing in c, from which it follows that H is concave. This means that the effects of an increase in c are passed on at a decreasing rate. Moreover, it means that a consumer should rationally be risk averse about the cost of search.

5.2.7.1 (Exercise) Suppose that there are two possible prices, 1 and 2, and that the probability of the lower price 1 is x. Compute the consumer's reservation price, which is the expected cost of searching, as a function of x and the cost of search c. For what values of x and c should the consumer accept 2 on the first search, or continue searching until the lower price 1 is found?

5.2.8 Edgeworth Box

The Edgeworth[60] box considers a two person, two good "exchange economy." That is, two people have utility functions of two goods and endowments (initial allocations) of the two goods. The Edgeworth box is a graphical representation of the exchange problem facing these people, and also permits a straightforward solution to their exchange problem.

[60] Francis Edgeworth, 1845-1926, introduced a variety of mathematical tools including calculus for considering economics and political issues, and was certainly among the first to use advanced mathematics for studying ethical problems.

The Edgeworth box is represented in Figure 5-23. Person 1 is "located" in the lower left (southwest) corner, and person 2 in the upper right (northeast). The *X* good is given on the horizontal axis, the *Y* good on the vertical. The distance between them is the total amount of the good they have between them. A point in the box gives the allocation of the good – the distance to the lower left to person 1, remainder to person 2. Thus, for the point illustrated, person 1 obtains (x_1, y_1), and person 2 obtains (x_2, y_2). The total amount of each good available to the two people will be fixed.

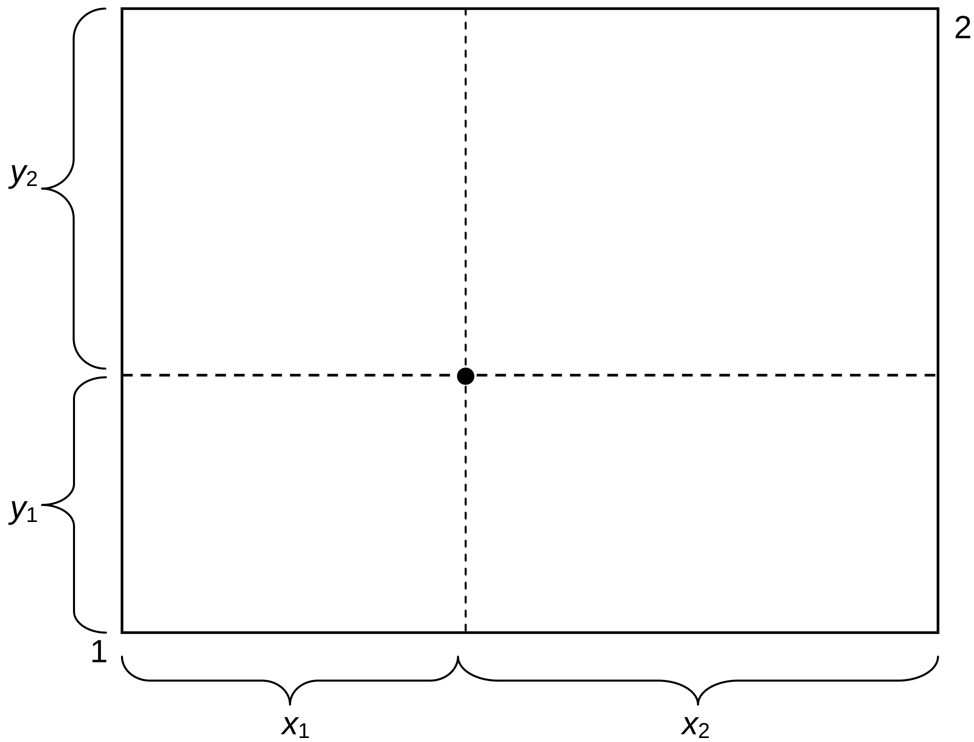

Figure 5-23: The Edgeworth Box

What points are efficient? The economic notion of efficiency is that an allocation is efficient if it is impossible to make one person better off without harming the other, that is, the only way to improve 1's utility is to harm 2, and vice versa. Otherwise, if the consumption is inefficient, there is a re-arrangement that makes both parties better off, and the parties should prefer such a point. Now, there is no sense of fairness embedded in the notion, and there is an efficient point in which one person gets everything and the other nothing. That might be very unfair, but it could still be the case that improving 2 must necessarily harm 1. The allocation is efficient if there is no waste or slack in the system, even if it is wildly unfair. To distinguish this economic notion, it is sometimes called *Pareto Efficiency*.[61]

We can find the Pareto-efficient points by fixing person 1's utility and then asking what point, on the indifference isoquant of person 1, maximizes person 2's utility? At that

[61] Vilfredo Pareto, 1848-1923, was a pioneer in replacing concepts of utility with abstract preferences, which was later adopted by the economics profession and remains the modern approach.

point, any increase in person 2's utility must come at the expense of person 1, and vice versa, that is, the point is Pareto-efficient. An example is illustrated in Figure 5-24.

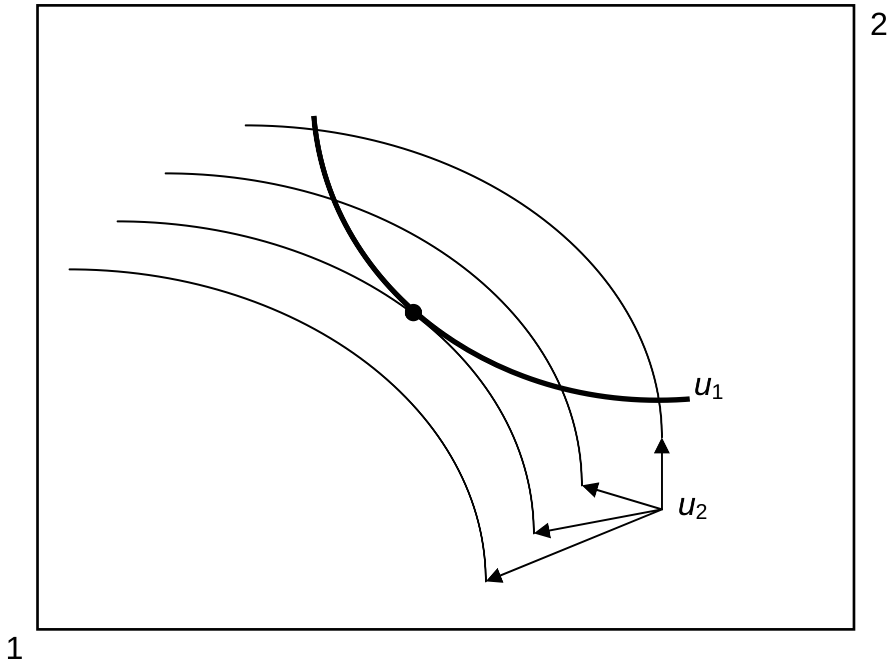

Figure 5-24: An Efficient Point

In Figure 5-24, the isoquant of person 1 is drawn with a dark thick line. This utility level is fixed. It acts like the "budget constraint" for person 2. Note that person 2's isoquants face the opposite way because a movement southwest is good for 2, since it gives him more of both goods. Four isoquants are graphed for person 2, and the highest feasible isoquant, which leaves person 1 getting the fixed utility, has the Pareto-efficient point illustrated with a large dot. Such points occur at tangencies of the isoquants.

This process, of identifying the points that are Pareto-efficient, can be carried out for every possible utility level for person 1. What results is the set of Pareto-efficient points, and this set is also known as the *contract curve*. This is illustrated with the thick line in Figure 5-25. Every point on this curve maximizes one person's utility given the other, and they are characterized by the tangencies in the isoquants.

The contract curve need not have a simple shape, as Figure 5-25 illustrates. The main properties are that it is increasing and goes from person 1 consuming zero of both goods to person 2 consuming zero of both goods.

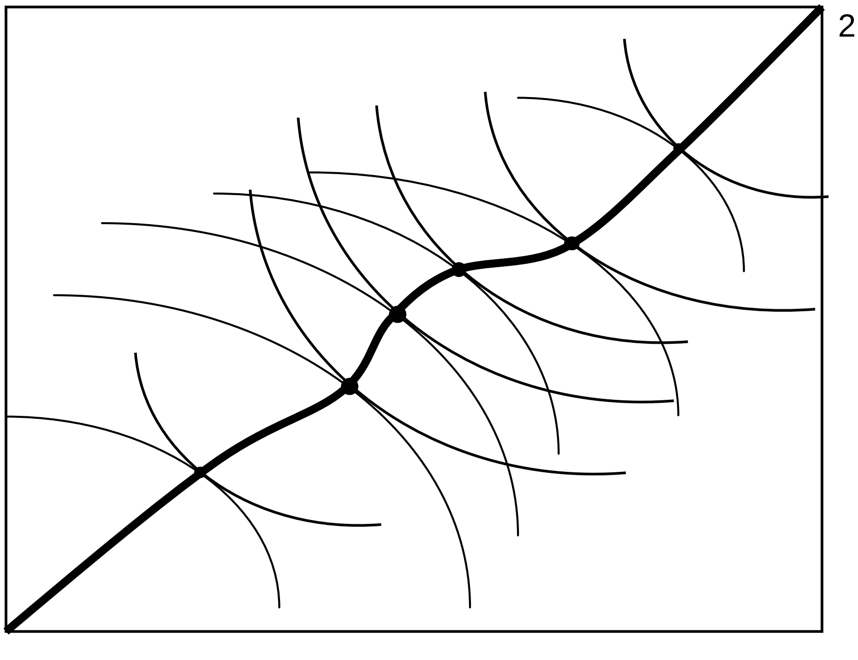

Figure 5-25: The Contract Curve

Example: Suppose both people have Cobb-Douglas utility. Let the total endowment of each good be 1, so that $x_2 = 1 - x_1$. Then person 1's utility can be written as $u_1 = x^{\alpha} y^{1-\alpha}$, and 2's utility is $u_2 = (1-x)^{\beta} (1-y)^{1-\beta}$. Then a point is Pareto-efficient if

$$\frac{\alpha y}{(1-\alpha)x} = \frac{\partial u_1 / \partial x}{\partial u_1 / \partial y} = \frac{\partial u_2 / \partial x}{\partial u_2 / \partial y} = \frac{\beta(1-y)}{(1-\beta)(1-x)}.$$

Thus, solving for y, a point is on the contract curve if

$$y = \frac{(1-\alpha)\beta x}{(1-\beta)\alpha + (\beta-\alpha)x} = \frac{x}{\frac{(1-\beta)\alpha}{(1-\alpha)\beta} + \frac{\beta-\alpha}{(1-\alpha)\beta} x} = \frac{x}{x + \left(\frac{(1-\beta)\alpha}{(1-\alpha)\beta}\right)(1-x)}.$$

Thus, the contract curve for the Cobb-Douglas case depends on a single parameter $\frac{(1-\beta)\alpha}{(1-\alpha)\beta}$. It is graphed for a variety of examples (α and β) in Figure 5-26.

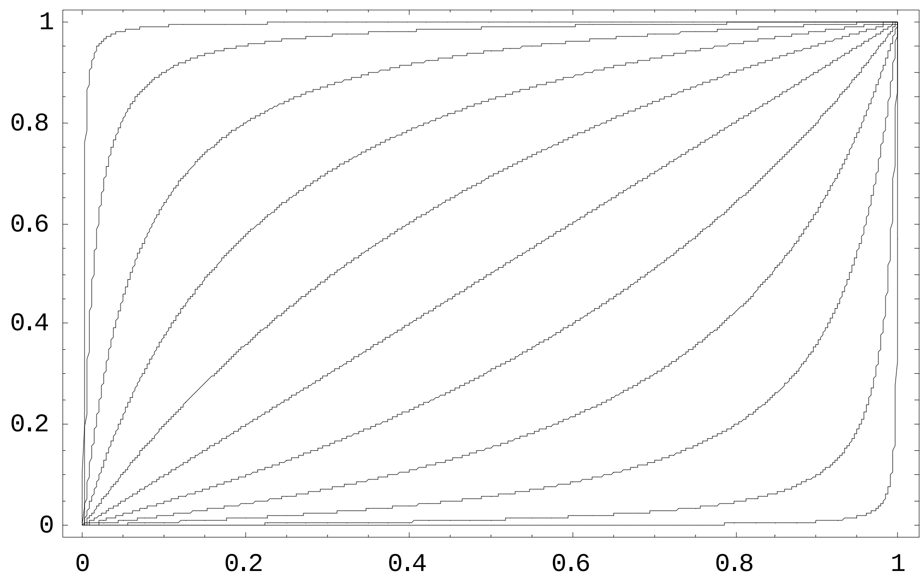

Figure 5-26: Contract Curves with Cobb-Douglas Utility

5.2.8.1 (Exercise) If two individuals have the same utility function concerning goods, is the contract curve the diagonal? Why or why not?

5.2.8.2 (Exercise) For two individuals with Cobb-Douglas preferences, when is the contract curve the diagonal?

The contract curve provides the set of efficient points. What point will actually be chosen? Let's start with an endowment of the goods. An endowment is just a point in the Edgeworth box, which gives the initial ownership of both goods for both people. The endowment is marked with a triangle in Figure 5-27. Note that this point gives the endowment of both person 1 and 2, because it shows the shares of each.

Figure 5-27 also shows isoquants for persons 1 and 2 going through the endowment. Note that the isoquant for 1 is concave toward the point labeled 1, and the isoquant for 2 is concave toward the point labeled 2. These utility isoquants define a reservation utility level for each person – the utility they could get alone, without exchange. This "no exchange" state is known as *autarky*. There are a variety of efficient points that give these people at least as much as they get under autarky, and those points are along the contract curve but have a darker line.

In Figure 5-27, starting at the endowment, the utility of both players is increased by moving in the general direction of the southeast, that is, down and to the right, until the contract curve is reached. This involves person 1 getting more X (movement to the right) in exchange for giving up some Y (movement down). Thus, we can view the increase in utility as a trade – 1 trades some of his Y for some of 2's X.

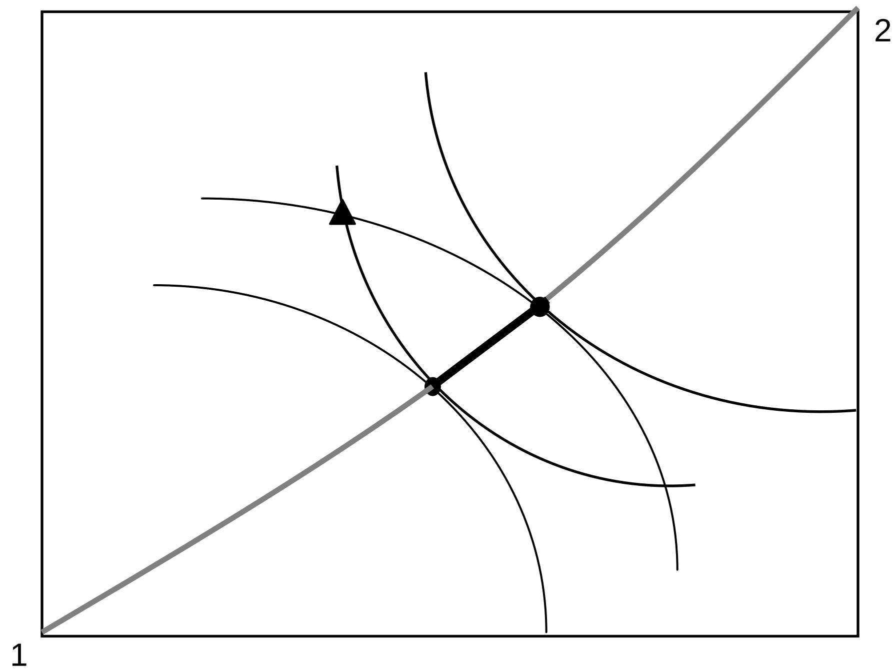

Figure 5-27: Individually Rational Efficient Points

In principle, any of the darker points on the contract curve, which give both people at least as much as they achieve under autarky, might result from trade. The two people get together and agree on exchange that puts them at any point along this segment of the curve, depending on the bargaining skills of the players. But there is a particular point, or possibly a set of points, that result from exchange using a *price system*. A price system involves a specific price for trading *Y* for *X*, and vice versa, that is available to both parties. In this diagram, prices define a straight line, whose slope is the negative of the *Y* for *X* price (the *X* for *Y* price is the reciprocal).

Figure 5-28 illustrates trade with a price system. The O in the center is the point on the contract curve connected to the endowment (triangle) by a straight line (the price line), in such a way that the straight line is tangent to both 1 and 2's isoquants at the contract curve. This construction means that, if each person took the price line as a budget constraint, they would maximize their utility function by choosing the point labeled O.

That a price line that (i) goes through the endowment and (ii) goes through the contract curve at a point tangent to both people's utility exists is relatively easy to show. Consider lines that satisfy property (ii) and let's see if we can find one that goes through the endowment. Start on the contract curve at the point that maximizes 1's utility given 2's reservation utility, and you can easily see that the price line through that point passes above and to the right of the endowment. The similar price line maximizing 2's utility given 1's reservation utility passes below and to the left of the endowment. These price lines are illustrated with dotted lines. Thus, by continuity, somewhere in between is a price line that passes through the endowment.

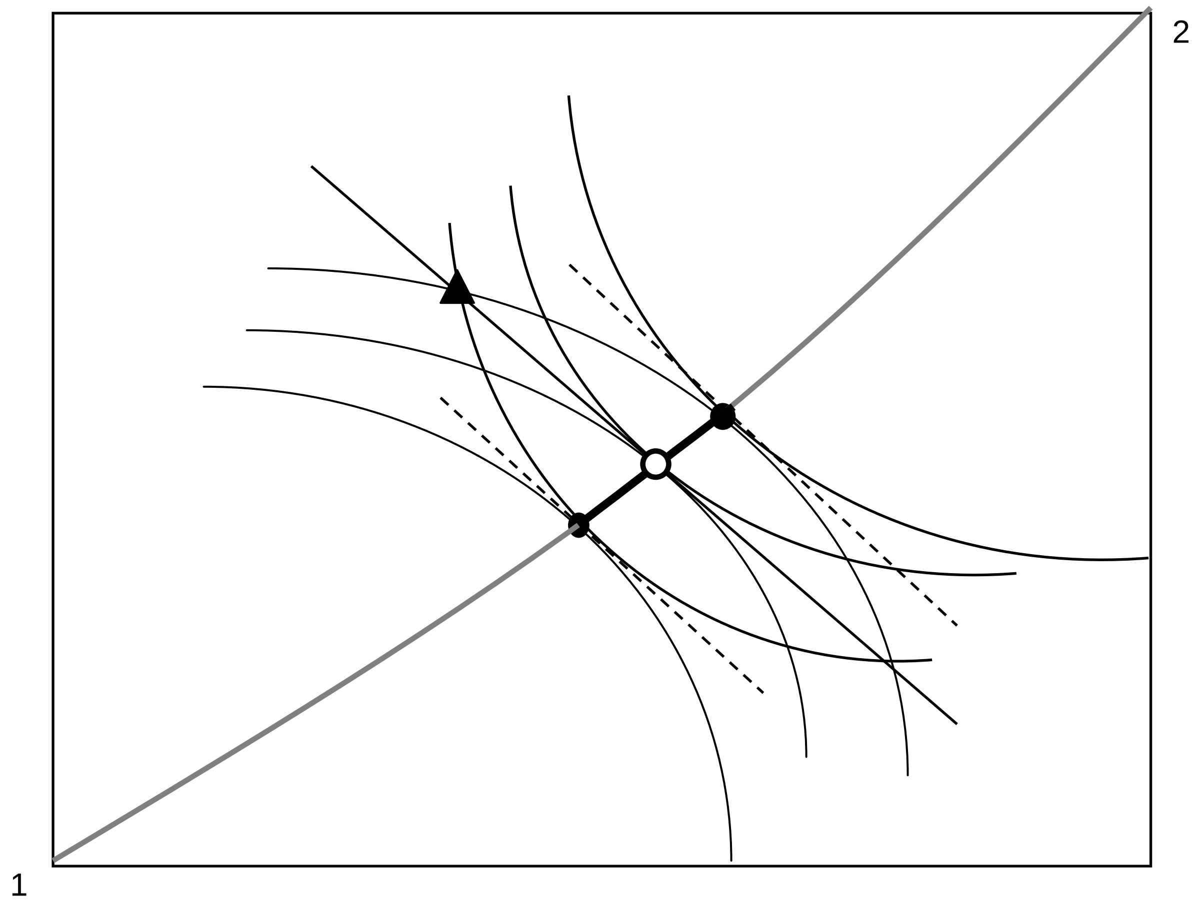

Figure 5-28: Equilibrium with a Price System

The point marked with the O represents an *equilibrium* of the price system, in the sense that supply and demand are equated for both goods. Note that, given the endowment and given the price through the endowment, both parties maximize utility by going to the O. To see this, it may help to consider a version of the picture that only shows person 2's isoquants and the price line.

Figure 5-29 removes player 1's isoquants, leaving only player 2's isoquants and the price line through the endowment. The price line through the endowment is the budget facing each player at that price. Note that, given this budget line, player 2, who gets more the less player 1 gets, maximizes utility at the middle isoquant, given the budget. That is, taking the price as given, player 2 would choose the O given player 2's endowment. The logic for player 1 is analogous. This shows that, if both players believe that they can buy or sell as much as they like at the tradeoff of the price through the O, both will trade to reach the O. This means that, if the players accept the price, a balance of supply and demand emerges. In this sense, we have found an equilibrium price.

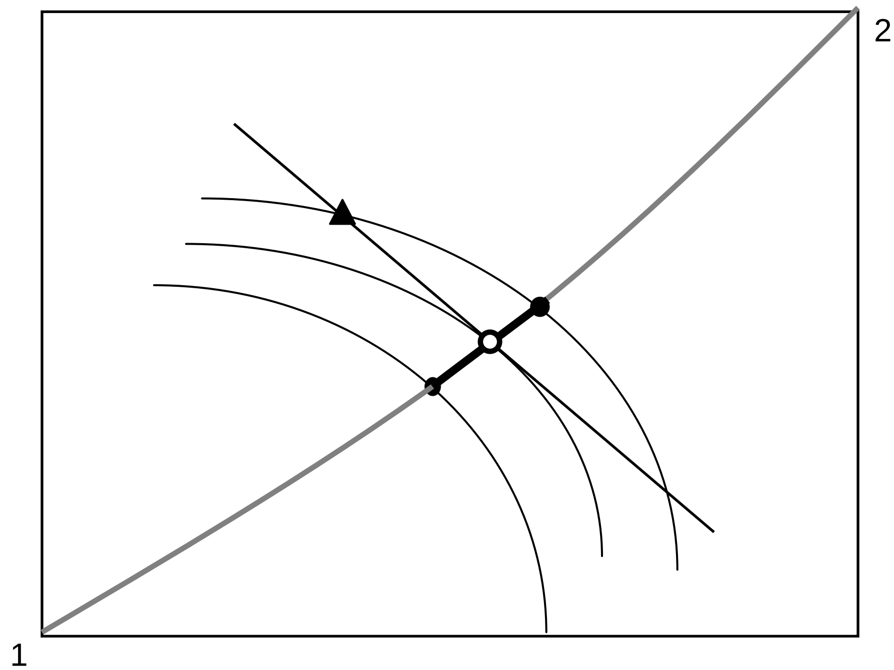

Figure 5-29: Illustration of Price System Equilibrium

In the Edgeworth box, we see that, given an endowment, it is possible to reach some Pareto-efficient point using a price system. Moreover, any point on the contract curve arises from as an equilibrium of the price system for some endowment. The proof of this proposition is startlingly easy. To show that a particular point on the contract curve is an equilibrium for some endowment, just start with an endowment equal to the point on the contract curve. No trade can occur because the starting point is Pareto-efficient – any gain by one party entails a loss by the other.

Furthermore, if a point in the Edgeworth box represents an equilibrium using a price system (that is, if the quantity supplied equals the quantity demanded for both goods), it must be Pareto-efficient. At an equilibrium to the price system, each player's isoquant is tangent to the price line, and hence tangent to each other. This implies the equilibrium is Pareto-efficient.

Two of these three propositions – any equilibrium of the price system is Pareto-efficient, any Pareto-efficient point is an equilibrium of the price system for some endowment, are known as the first and second *welfare theorems* of *general equilibrium*. They have been demonstrated by Nobel laureates Kenneth Arrow and Gerard Debreu, for an arbitrary number of people and goods. They also demonstrated the third proposition, that for any endowment, there exists an equilibrium of the price system, with the same high level of generality.

5.2.9 General Equilibrium

We will illustrate general equilibrium, for the case when all consumers have Cobb-Douglas utility in an exchange economy. An exchange economy is an economy where the supply of each good is just the total endowment of that good and there is no production. Suppose there are N people, indexed by $n = 1, 2, \ldots, N$. There are G goods, indexed by $g = 1, 2, \ldots, G$. Person n has Cobb-Douglas utility, which we can represent using exponents $\alpha(n,g)$, so that the utility of person n can be represented as $\prod_{g=1}^{G} x(n,g)^{\alpha(n,g)}$, where $x(n,g)$ is person n's consumption of good g. Assume that $\alpha(n,g) \geq 0$ for all n and g, which amounts to assuming that the products are in fact goods. Without any loss of generality, we can require

$$\sum_{g=1}^{G} \alpha(n,g) = 1,$$

for each n. (To see this, note that maximizing the function U is equivalent to maximizing the function U^{β} for any positive β.)

Let $y(n,g)$ be person n's endowment of good g. The goal of general equilibrium is to find prices $p_1, p_2, \ldots, p_G$ for the goods, in such a way that demand for each good exactly equals supply of the good. The supply of good g is just the sum of the endowments of that good. The prices yield a wealth for person n equal to

$$W_n = \sum_{g=1}^{G} p_g y(n,g).$$

We will assume that $\sum_{n=1}^{N} \alpha(n,g) y(n,i) > 0$ for every pair of goods g and i. This assumption states that for any pair of goods, there is at least one agent that values good g and has an endowment of good i. The assumption insures that there is always someone willing and able to trade if the price is sufficiently attractive. The assumption is much stronger than necessary but useful for exposition. The assumption also insures the endowment of each good is positive.

Cobb-Douglas utility simplifies the analysis because of a feature that we already met in the case of two goods, but which holds in general: the share of wealth for a consumer n on good g equals the exponent $\alpha(n,g)$. Thus, the total demand for good g is

$$X_g = \sum_{n=1}^{N} \frac{\alpha(n,g) W_n}{p_g}.$$

The equilibrium conditions, then, can be expressed as saying supply (sum of the endowments) equals demand, or, for each good g,

$$\sum_{n=1}^{N} y(n,g) = X_g = \sum_{n=1}^{N} \frac{\alpha(n,g) W_n}{p_g}.$$

We can rewrite this expression, provided $p_g>0$ (and it must be for otherwise demand is infinite), to be

$$p_g - \sum_{i=1}^{G} p_i \frac{\sum_{n=1}^{N} y(n,i)\alpha(n,g)}{\sum_{n=1}^{N} y(n,g)} = 0 .$$

Let **B** be the $G \times G$ matrix whose (g, i) term is

$$b_{gi} = \frac{\sum_{n=1}^{N} y(n,i)\alpha(n,g)}{\sum_{n=1}^{N} y(n,g)} .$$

Let **p** be the vector of prices. Then we can write the equilibrium conditions as

$$(\mathbf{I}-\mathbf{B})\mathbf{p}=\mathbf{0},$$

where $\mathbf{0}$ is the zero vector. Thus, for an equilibrium (other than $\mathbf{p}=\mathbf{0}$) to exist, **B** must have an eigenvalue equal to 1, and a corresponding eigenvector **p** that is positive in each component. Moreover, if such an eigenvector, eigenvalue pair exists, it is an equilibrium, because demand is equal to supply for each good.

The actual price vector is not completely identified, because if **p** is an equilibrium price vector, so is any positive scalar times **p**. Scaling prices doesn't change the equilibrium because both prices and wealth (which is based on endowments) rise by the scalar factor. Usually economists assign one good to be a numeraire, which means all other goods are indexed in terms of that good, and the numeraire's price is artificially set to be 1. We will treat any scaling of a price vector as the same vector.

The relevant theorem is the Perron-Frobenius theorem.[62] It states that if **B** is a positive matrix (each component positive), then there is an eigenvalue $\lambda>0$ and an associated positive eigenvector **p**, and moreover λ is the largest (in absolute value) eigenvector of **B**.[63] This conclusion does most of the work of demonstrating the existence of an

---

[62] Oskar Perron, 1880 - 1975 and Georg Frobenius, 1849 – 1917.

[63] The Perron-Frobenius theorem, as usually stated, only assumes that **B** is non-negative and that **B** is irreducible. It turns out that a strictly positive matrix is irreducible, so this condition is sufficient to invoke the theorem. In addition, we can still apply the theorem even when **B** has some zeros in it, provided that it is irreducible. Irreducibility means that the economy can't be divided into two economies, where the people in one economy can't buy from the people in the second because they aren't endowed with anything the people in the first value. If **B** is not irreducible, then some people may wind up consuming zero of things they value.

equilibrium. The only remaining condition to check is that the eigenvalue is in fact 1, so that $(\mathbf{I}-\mathbf{B})\mathbf{p}=\mathbf{0}$.

Suppose the eigenvalue is λ. Then $\lambda\mathbf{p} = \mathbf{Bp}$. Thus for each g,

$$\lambda p_g = \sum_{i=1}^{G} \frac{\sum_{n=1}^{N} \alpha(n,g) y(n,i)}{\sum_{m=1}^{N} y(m,g)} p_i ,$$

or

$$\lambda p_g \sum_{m=1}^{N} y(m,g) = \sum_{i=1}^{G} \sum_{n=1}^{N} \alpha(n,g) y(n,i) p_i .$$

Summing both sides over g,

$$\lambda \sum_{g=1}^{G} p_g \sum_{m=1}^{N} y(m,g) = \sum_{g=1}^{G} \sum_{i=1}^{G} \sum_{n=1}^{N} \alpha(n,g) y(n,i) p_i$$

$$= \sum_{i=1}^{G} \sum_{n=1}^{N} \sum_{g=1}^{G} \alpha(n,g) y(n,i) p_i = \sum_{i=1}^{G} \sum_{n=1}^{N} y(n,i) p_i .$$

Thus $\lambda=1$ as desired.

The Perron-Frobenius theorem actually provides two more useful conclusions. First, the equilibrium is unique. This is a feature of the Cobb-Douglas utility and does not necessarily occur for other utility functions. Moreover, the equilibrium is readily approximated. Denote by $\mathbf{B}^t$ the product of $\mathbf{B}$ with itself t times. Then for any positive vector $\mathbf{v}$, $\lim_{t\to\infty} \mathbf{B}^t \mathbf{v} = \mathbf{p}$. While approximations are very useful for large systems (large numbers of goods), the system can readily be computed exactly with small numbers of goods, even with a large number of individuals. Moreover, the approximation can be interpreted in a potentially useful manner. Let $\mathbf{v}$ be a candidate for an equilibrium price vector. Use $\mathbf{v}$ to permit people to calculate their wealth, which for person n is $W_n = \sum_{i=1}^{G} v_i y(n,i)$. Given the wealth levels, what prices clear the market? Demand for good g is

$$x_g(v) = \sum_{n=1}^{N} \alpha(n,g) W_n = \sum_{i=1}^{G} v_i \sum_{n=1}^{N} \alpha(n,g) y(n,i)$$

and the market clears, given the wealth levels, if $p_g = \dfrac{\sum_{i=1}^{G} v_i \sum_{n=1}^{N} \alpha(n,g) y(n,i)}{\sum_{n=1}^{N} y(n,g)}$, which is equivalent to $\mathbf{p} = \mathbf{Bv}$. This defines an iterative process. Start with an arbitrary price vector, compute wealth levels, then compute the price vector that clears the market for the given wealth levels. Use this price to recalculate the wealth levels, and then compute a new market-clearing price vector for the new wealth levels. This process can be iterated, and in fact converges to the equilibrium price vector from any starting point.

We finish this section by considering three special cases. If there are two goods, we can let $a_n = \alpha(n, 1)$, and then conclude $\alpha(n, 2) = 1 - a_n$. Then let

$$Y_g = \sum_{n=1}^{N} y(n,g)$$

be the endowment of good g. Then the matrix $\mathbf{B}$ is

$$\mathbf{B} = \begin{pmatrix} \frac{1}{Y_1}\sum_{n=1}^{N} y(n,1)a_n & \frac{1}{Y_1}\sum_{n=1}^{N} y(n,2)a_n \\ \frac{1}{Y_2}\sum_{n=1}^{N} y(n,1)(1-a_n) & \frac{1}{Y_2}\sum_{n=1}^{N} y(n,2)(1-a_n) \end{pmatrix}$$

$$= \begin{pmatrix} \frac{1}{Y_1}\sum_{n=1}^{N} y(n,1)a_n & \frac{1}{Y_1}\sum_{n=1}^{N} y(n,2)a_n \\ \frac{1}{Y_2}\left(Y_1 - \sum_{n=1}^{N} y(n,1)a_n\right) & 1 - \frac{1}{Y_2}\sum_{n=1}^{N} y(n,2)a_n \end{pmatrix}$$

The relevant eigenvector of $\mathbf{B}$ is

$$\mathbf{p} = \begin{pmatrix} \sum_{n=1}^{N} y(n,2)a_n \\ \sum_{n=1}^{N} y(n,1)(1-a_n) \end{pmatrix}.$$

The overall level of prices is not pinned down – any scalar multiple of $\mathbf{p}$ is also an equilibrium price – so the relevant term is the price ratio, which is the price of good 1 in terms of good 2, or

$$\frac{p_1}{p_2} = \frac{\sum_{n=1}^{N} y(n,2) a_n}{\sum_{n=1}^{N} y(n,1)(1-a_n)}.$$

We can readily see that an increase in the supply of good 1, or a decrease in the supply of good 2, decreases the price ratio. An increase in the preference for good 1 increases the price of good 1. When people who value good 1 relatively highly are endowed with a lot of good 2, the correlation between preference for good 1 a_n and endowment of good 2 is higher. The higher the correlation, the higher is the price ratio. Intuitively, if the people who have a lot of good 2 want a lot of good 1, the price of good 1 is going to be higher. Similarly, if the people who have a lot of good 1 want a lot of good 2, the price of good 1 is going to be lower. Thus, the correlation between endowments and preferences also matters to the price ratio.

In our second special case, we consider people with the same preferences, but who start with different endowments. Hypothesizing identical preferences sets aside the correlation between endowments and preferences found in the two good case. Since people are the same, $\alpha(n, g) = A_g$ for all n. In this case,

$$b_{gi} = \frac{\sum_{n=1}^{N} y(n,i)\alpha(n,g)}{\sum_{n=1}^{N} y(n,g)} = A_g \frac{Y_i}{Y_g},$$

where as before $Y_g = \sum_{n=1}^{N} y(n,g)$ is the total endowment of good g. The matrix **B** has a special structure, and in this case, $p_g = \frac{A_g}{Y_g}$ is the equilibrium price vector. Prices are proportional to the preference for the good divided by the total endowment for that good.

Now consider a third special case, where no common structure is imposed on preferences, but endowments are proportional to each other, that is, the endowment of person n is a fraction w_n of the total endowment. This implies that we can write $y(n, g) = w_n Y_g$, an equation assumed to hold for all people n and goods g. Note that by construction, $\sum_{n=1}^{N} w_n = 1$, since the value w_n represents n's share of the total endowment. In this case, we have

$$b_{gi} = \frac{\sum_{n=1}^{N} y(n,i)\alpha(n,g)}{\sum_{n=1}^{N} y(n,g)} = \frac{Y_i}{Y_g} \sum_{n=1}^{N} w_n \alpha(n,g) .$$

These matrices also have a special structure, and it is readily verified that the equilibrium price vector satisfies

$$p_g = \frac{1}{Y_g} \sum_{n=1}^{N} w_n \alpha(n,g) .$$

This formula receives a similar interpretation – the price of good g is the strength of preference for good g, where strength of preference is a wealth-weighted average of the individual preference, divided by the endowment of the good. Such an interpretation is guaranteed by the assumption of Cobb-Douglas preferences, since these imply that individuals spend a constant proportion of their wealth on each good. It also generalizes the conclusion found in the two good case to more goods, but with the restriction that the correlation is now between wealth and preferences. The special case has the virtue that individual wealth, which is endogenous because it depends on prices, can be readily determined.

5.2.9.1 (Exercise) Consider a consumer with Cobb-Douglas utility,

$$U = \prod_{i=1}^{G} x_i^{a_i} ,$$

where $\sum_{i=1}^{G} a_i = 1$, and facing the budget constraint $W = \sum_{i=1}^{G} p_i x_i$. Show that the consumer maximizes utility by choosing $x_i = \frac{a_i W}{p_i}$ for each good i. Hint: Express the budget constraint as $x_G = \frac{1}{p_G}\left(W - \sum_{i=1}^{G-1} p_i x_i \right)$, and thus utility as $U = \left(\prod_{i=1}^{G-1} x_i^{a_i} \right) \left(\frac{1}{p_G}\left(W - \sum_{i=1}^{G-1} p_i x_i \right) \right)^{a_G}$. This function can now be maximized in an unconstrained fashion. Verify that the result of the maximization can be expressed as $p_i x_i = \frac{a_i}{a_G} p_G x_G$, and thus $W = \sum_{i=1}^{G} p_i x_i = \sum_{i=1}^{G} \frac{a_i}{a_G} p_G x_G = \frac{p_G x_G}{a_G}$, which yields $p_i x_i = a_i W$.