Draft – please do not cite without permission
June 2009
Development Economics: Theory, Empirical Research and Policy Analysis
Julie Schaffner

Chapter 6
Consumption, Time Allocation and Production Basics

People participate actively in development and respond to incentives. Understanding how people make choices and how they respond to new challenges and opportunities is fundamental for understanding development and analyzing policy. We thus begin building our analytical toolkit by examining three fundamental sets of choices: choices regarding consumption, time allocation and production. We review the basic theories that economists bring to the study of these choices, and demonstrate their usefulness in applications relating to nutrition policy, child labor and the Green Revolution.

In Chapter 7 we will argue the need to combine, modify and extend these basic theories to achieve more complete understanding of developing country households and their choices. The basic theories nonetheless illuminate many insights of enduring value in development analysis.

Basic consumer theory: allocating budgets

A motivation. Nearly 40 percent of the developing world's population consumes insufficient quantities of calories or micronutrients such as iron or Vitamin A. Nearly 30 percent of children are seriously malnourished.[1] Malnourished children are much more likely to die, and malnourished mothers are much more likely to give birth to children with serious defects. Even among those who make it to adulthood without serious disabilities, inadequate nutrition renders them more susceptible to illness and reduces their productivity and earnings (Barrett, 2002). Thus reducing malnutrition rightly demands the attention of policymakers.

Several obvious approaches to reducing malnutrition suggest themselves: (a) giving poorly nourished people cash with which to buy more or better food, (b) giving them food, and (c) reducing the economy-wide price of staple foods, rendering it easier for people to buy the foods they need. The impacts such policies have on nutrition, however, are mediated by peoples' choices regarding what foods to purchase and consume. Thus analysis of nutrition policy must be informed by knowledge of how people make these

[1] Twenty-seven percent of children in developing countries are "stunted," meaning that their heights are at least two standard deviations below the mean height in a well-nourished population of the same age, and 23 percent are "underweight," meaning that their weights are at least two standard deviations below the mean weight in a well-nourished population of the same age. These and other statistics reported in this paragraph are taken from World Bank (2006).

choices. In what follows we review the basic economic theory of consumer choice, and then demonstrate its application in the analysis of nutrition policy.[2]

Economic analysis of choices. A basic premise in economic analysis is that people make choices with the aim of making themselves as well off as possible given the constraints they face. This suggests that we can gain insight into the forces that shape their choices by studying the nature of the constraints they face and the way they evaluate their options. In general, as we saw in Chapter 2, people make many inter-related choices, perceive their well-being to be a function of a wide array of life circumstances, and face constraints that are numerous and complicated. It is often possible, however, to gain useful insight by imagining a stylized decision-maker who faces just a small subset of choices, identifying what constrains those choices, and making simple but reasonable assumptions about how the decision-maker evaluates how "well off" any choice would make her.

The three basic "theories" discussed in this chapter should be understood as three of the most frequently employed applications of the more general economic approach to the analysis of choice. Each is defined by (1) what it identifies to be the relevant set of choices, (2) the main constraints on those choices, and (3) the objective that guides the decision-maker's choice. In each of the theories a decision-maker is construed as having to choose the values of several variables, and a "choice" is a set of values for those variables. For example, a consumer may choose the quantities of rice, vegetables and other goods to consume. The decision-maker is assumed to maximize her **objective,** which is a function of the choice variables that indicates how "well off" any choice would leave her. The **constraints** generally point to the ways in which the decision-maker's feasible choices are limited by her **endowments,** or the assets and other resources over which she has control, as well as the socio-economic conditions she faces in markets and non-market institutions.

The three theories exposited in this chapter should be seen not as analytical straightjackets imposing economists' beliefs onto the world, but as starting points for the development of analytical frameworks tailor-made for the analysis of a wide range of choices in a wide range of contexts. A good economist observes context closely in order to make wise choices in the construction of an adequate analytical framework.

Assumptions of basic consumer theory. Basic consumer choice theory focuses on a decision-maker called a "consumer," who is endowed with a fixed amount of income, which she must allocate across the purchase of diverse goods. Her objective is to maximize her **utility,** which is assumed to be a function only of the quantities of goods

[2] This chapter's discussion of these policies will be incomplete. The analysis is taken further in Chapter 7, and more comprehensive analysis of agricultural pricing policies is taken up in Chapter 14, while more comprehensive analysis of cash and food distribution efforts are taken up in Chapters 21 and 22.

she consumes.[3,4] She maximizes utility subject to her **budget constraint**, which states that the total value of her purchases must not exceed her income.

Consider the simple case of a consumer who consumes only two goods, "food" and "non-food." She chooses a quantity of food, *F*, and a quantity of non-food, *N*, measured in kilograms, to maximize her utility, *U(F,N)*. Utility is assumed to rise as either *F* or *N* rises (while holding the other constant), indicating that "having more is better." The extent to which utility rises with *F* (or *N*) is assumed to decline as *F* (*N*) increases. This assumption of "diminishing marginal utility" says that each additional unit of a good contributes less to utility than the one before. (Put another way, this says that a consumer values a kilogram of food the most when she has little of it to start with.) If her income measured in pesos is *Y*, and she faces prices for food and non-food of P_f and P_n pesos per kilogram, then her budget constraint may be expressed as

$$P_f F + P_n N \leq Y.$$

Figure 6.1a offers a useful summary of this simple theory and its implications. The two axes measure *F* and *N*. All "consumption pairs" (*F,N*) that satisfy the budget constraint lie either along the diagonal "**budget line"** in the diagram, or in the shaded region that lies "under" that line. Since our consumer always prefers to consume more of either good, she always chooses to consume at a point on her budget line (rather than at a point inside the shaded region). The budget line touches the horizontal axis at the point Y/P_f, indicating that if she spent all of her income on food, she could purchase Y/P_f kilograms of food.[5] Similarly, the budget line touches the vertical axis at Y/P_n. The slope of the budget line is $-P_f/P_n$, indicating that for each unit of food she gives up (moving one unit to the left in the diagram), she frees up P_f dollars, which she can use to buy P_f/P_n units of non-food (moving up by that amount in the diagram).

To depict the utility our consumer associates with various consumption pairs (*F,N*) (without adding a third dimension to our graph), we draw in **indifference curves**, such as the curves I_1, I_2 and I_3. Each indifference curve identifies all the combinations of *F* and *N* that would yield the consumer a particular level of utility. (Choice between two such pairs on I_1 would be a matter of "indifference" to the consumer.) Every point in the graph lies on some indifference curve. I_1, I_2 and I_3 are just a few of the infinite number of indifference curves that fill the graph. Indifference curves that lie further from the origin pertain to higher levels of utility. Individual indifference curves are downward-sloping,

---

[3] While the notion of a consumer's utility is related to the broader and deeper notion of well-being discussed in Chapter 2, we give it a distinct name, highlighting the basic theory's simplified assumptions regarding what people care about.

[4] While this framework is useful for analyzing many choices, it is inadequate, and must be extended, for analyzing some choices. For example, in this simple framework the utility a consumer achieves today is a function only of the quantities of goods and services she consumes today. When analyzing choices regarding goods with addictive properties, like cigarettes, we would need to complicate the framework, introducing the possibility that a consumer who has consumed more cigarettes in the past may place higher weight on cigarette consumption when evaluating today's utility.

[5] Dividing Y pesos of income by P_f pesos per kilogram of food yields the kilograms of food that can be purchased with Y pesos.

indicating that if *F* increases, then utility can remain constant only if *N* decreases. The curves are "convex" (i.e. bowed toward the origin, with slopes that get shallower as we move to the right in the diagram), reflecting the assumption of diminishing marginal utility. When the consumer has little food and much non-food (and is thus consuming nearer to the vertical axis), she would be willing to give up quite a lot of non-food in exchange for an extra unit of food, while when she has much food and little non-food (and is thus consuming nearer the horizontal axis), she would be willing to give up little non-food for an additional unit of food.

Figure 6.1a

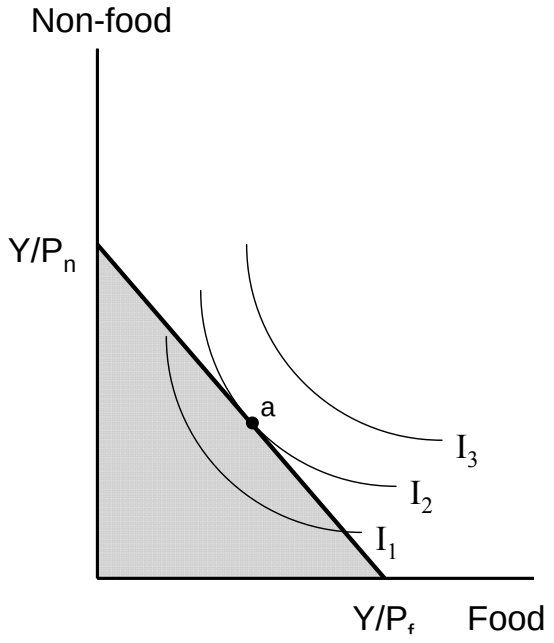

The choice of food and non-food consumption that maximizes the consumer's utility is found where an indifference curve is tangent to her budget constraint.

Figure 6.1b

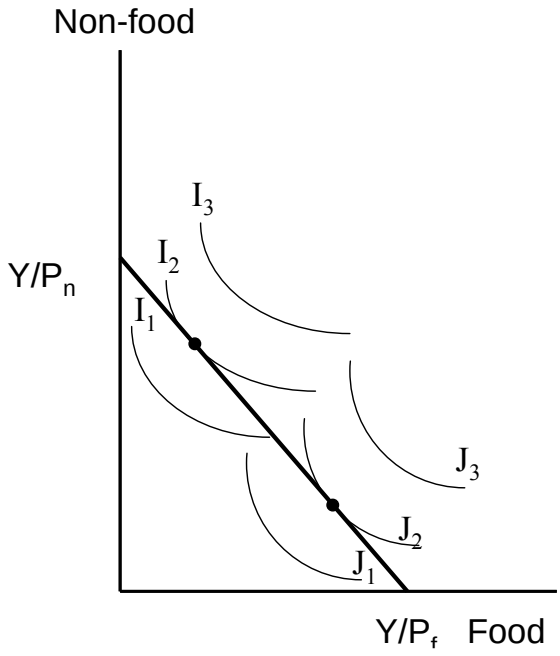

Consumers with different preferences facing the same budget constraint choose different combinations of food and non-food.

The quantities *F* and *N* that maximize our consumer's utility while satisfying the budget constraint are found by identifying the point on the budget line that lies on the highest indifference curve (i.e. the indifference curve that lies furthest from the origin). In Figure 6.1a, the highest indifference curve touching the budget line is I_2, which is just tangent to the budget line at point *a*. The consumer's utility is maximized by consuming the quantities of food and non-food associated with point *a*.[6]

---

[66] Mathematically, the consumer chooses *F* and *N* to maximize $U(F,N)$ subject to $P_fF + P_nN = Y$. Rearranging the budget constraint and substituting into the maximand, the problem reduces to one of choosing *F* to maximize $U(F,(Y-P_fF)/P_n)$. Taking the derivative of this function with respect to *F*, setting it equal to zero and re-arranging, we find that the solution must be characterized by the requirement that

$$U_f(F,N)/U_n(F,N) = P_f/P_n,$$

where $U_f(F,N)$ and $U_n(F,N)$ are the first partial derivatives of $U(F,N)$ with respect to *F* and *N*. The left hand side of this expression is equal to slope of the indifference curve associated with the utility-maximizing choice. The right hand side is the slope of the budget constraint. The two are equal where an indifference curve is tangent to the budget constraint.

The theory summarized in Figure 6.1a points to four reasons why food consumption choices may differ across consumers: differences in (1) Y, (2) the P_f, (3) P_n and (4) the nature of the consumer's preferences. Differences in income and prices cause consumers to face different budget constraints, as we discuss below. Differences in preferences cause them to make differing choices even when facing the same budget constraint. For example, in Figure 6.1b, indifference curves I_1, I_2 and I_3 describe the preferences of consumer i, while curves J_1, J_2 and J_3 describe the preferences of consumer j, who evidently places greater value on food consumption relative to non-food consumption. Facing the same budget constraint, consumer j chooses to consume more food (and less non-food) than consumer i. Such differences in preferences might arise out of differences in the size and demographic structure of the consumers' households (which shape household nutritional requirements), as well as differences in culture, beliefs, knowledge and habits, and for more idiosyncratic reasons.

Effects of an increase in income. Consider the impact of increasing our consumer's income, while holding both prices constant.[7] Recall that the endpoints of the budget constraint are given by Y/P_f and Y/P_n, and that the slope of the budget constraint is P_f/P_n. If Y increases while prices hold constant, both endpoints of the budget constraint shift further out from the origin, while the slope remains the same. That is, an increase in income causes the budget constraint to shift out from the origin in a parallel fashion. For example, the parallel budget lines in Figure 6.2a pertain to two levels of income and identical prices. As income increases (from Y_1 to Y_2), the level of well-being our consumer attains must increase, because the range of feasible consumption pairs expands. What happens to the specific quantities of food and non-food that she chooses is less clear, however. If her preferences are like those illustrated in Figure 6.2a, then as income increases, her consumption choice changes from point a_1 to a_2. The quantities she consumes of both food and non-food increase, with the F rising more slowly than N. If her preferences were different, however, increases in income might cause F to rise more rapidly than N, and or could even cause F to fall. (See Problem 2 at the end of the chapter.) The theory tells us that an additional dollar of income is likely (but not guaranteed) to lead to an increase in spending on both goods, and that empirical research is required to determine the most likely way in which the extra dollar will be spent by a member of a particular population.

Figure 6.2a **Figure 6.2b**

[7] The next several sub-sections offer graphical descriptions of the model's implications for how the consumer's utility-maximizing choices of F and N change in response to changes in Y, P_f, P_n. For mathematical treatment of these results, see problem 1 at the end of the chapter.

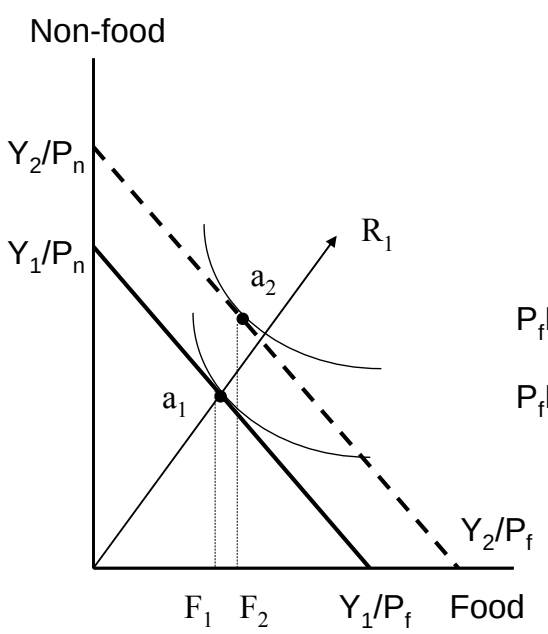

An increase in income induces a parallel shift out of the budget constraint. Whether the share of income spent on food rises or falls is an empirical question, as either impact is theoretically possible.

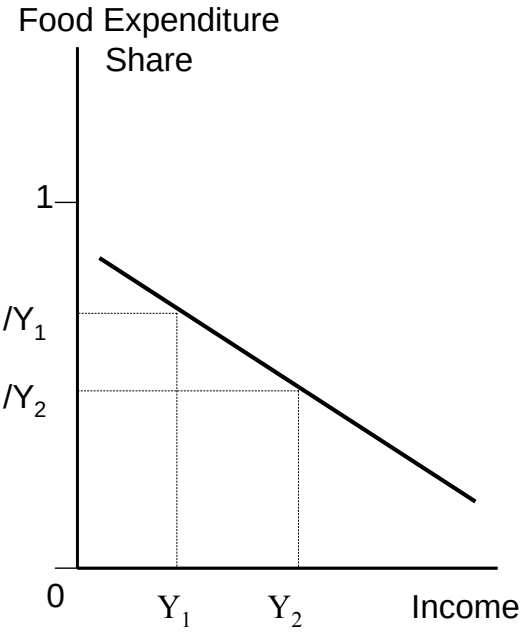

A consumer's Engel Curve for food indicates how the share of income spent on food changes as the level of income rises.

Normal and inferior goods. A good whose consumption increases when income increases is called a **normal good**, while a good whose consumption decreases is called an **inferior good**. When Y rises, the consumption of at least one good must rise, and we would expect that most goods are normal. Inferior goods are exceptional goods that consumers would rather replace by higher quality items when their incomes rise. For example, brown sugar may be an inferior good for some consumers, who shift from the use of brown sugar to more refined (and more expensive) white sugar as their income levels rise.

Income shares and Engel curves. In a diagram like Figure 6.2a, any (*F,N*) pair falling on a single ray from the origin (such as the ray R_1) involves the same ratio of *F* to *N*. As long as prices remain constant, the share of Y spent on food (or non-food) is also identical along the same ray. Rays from the origin with higher slopes involve higher ratios of *N* to *F*, and thus smaller shares of total consumption expenditure devoted to food consumption.

We can trace out a consumer's **Engel curve** for food in a graph like Figure 6.2b, in which we measure the consumer's income along the horizontal axis and the share of her expenditure devoted to food (a fraction between 0 and 1) measured along the vertical axis. While tracing out the curve, we hold P_f and P_n constant. If the consumer described in Figure 6.2a received income Y_1, she would choose to consume F_1 units of food. We thus plot the point (Y_1, $P_f F_1/Y_1$) in Figure 6.2b. Similarly, with income Y_2 she would choose to spend a share of income P_fF_2/Y_2 on food. As income rises for the consumer depicted here, her total peso expenditure on food rises, but the share of total consumption

expenditure that she devotes to food consumption falls (as witnessed by the fact that when the budget line shifts out, the new consumption point lies on a ray from the origin with higher slope). Thus this consumer's Engel curve for food slopes downward in Figure 6.2b. Basic consumer theory allows for a good's Engel curve to rise or fall as income increases. Empirically, Engel curves for staple foods tend to slope downward (at least after sufficiently high income levels are reached. (Why might this be?)

Effects of a reduction in the price of food. A reduction in P_f (from P_{f1} to P_{f2}, holding Y and P_n constant) pushes the endpoint of the budget constraint on the horizontal axis out from the origin while leaving the other endpoint fixed in place, as illustrated in figure 6.3a. (To confirm this, recall that the endpoints of the budget constraint are given by Y/P_f and Y/P_n.) After a reduction in price, the new budget constraint lies entirely outside the original budget constraint (except for a single shared endpoint), thus the consumer's well-being must increase. For the consumer whose preferences are illustrated in Figure 6.3a, consumption of both F and N rises when the price of food falls, though F rises much more than N. Again, a wide range of responses to the price change is theoretically possible. The consumption of an item may even fall when its price falls. In most cases we expect that demand rises when the price falls. Whether the demand response is large or small is a matter for empirical study, however.

Figure 6.3a

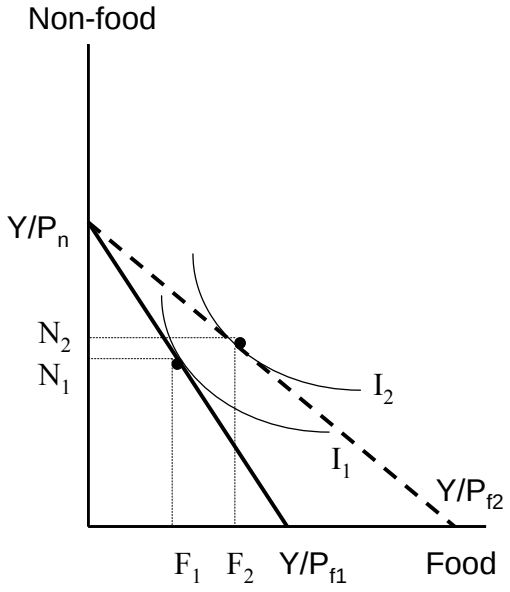

A reduction in the price of food causes the endpoint of the budget line on the food axis to move to the right, while leaving the endpoint on the non-food axis unchanged.

Figure 6.3b

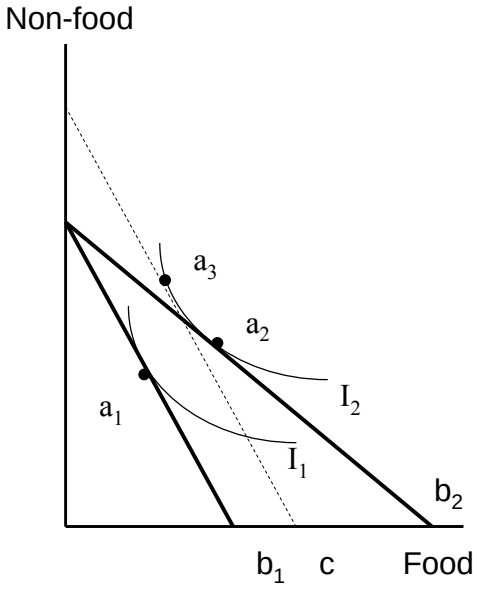

The overall effect on consumption choices of a reduction in the food price is indicated by the change from point a_1 to a_2. This may be decomposed into the income effect (the change from a_1 to a_3) and the substitution effect (the change from a_3 to a_2).

Purchasing power transfer implicit in a price reduction. The solid budget lines in Figure 6.3b are identical to the budget lines in Figure 6.3a, but for convenience have endpoints on the horizontal axis re-labeled as b_1 and b_2. Notice that our consumer's well-being

could have been raised from the level associated with the initial consumption choice (a_1) to the level associated with indifference curve I_2 through an increase in income rather than a reduction in the price of food. Such an increase in Y would have shifted the budget constraint out in a parallel fashion, from the solid line ending at b_1 to the dashed line ending at c. Thus we can think of a price reduction as containing within it an increase in purchasing power much like an income increase.

We can assess the approximate magnitude of the purchasing power transfer to a household implicit in a price reduction by performing the following mental exercise. Suppose the consumer initially consumes F_1 units of food and N_1 units of non-food. If the consumer were to continue to consume these quantities, then a reduction in the price from p_1 to p_2 would "free up" $(p_1 - p_2)*F_1$ pesos of purchasing power, which the consumer may now use to purchase more of whatever she wishes.[8] (This is only an approximate indication of the ultimate transfer size, because the consumer is likely to respond to the lower price of food by buying more of it at the subsidized price, increasing the size of the implicit transfer.) The more this consumer purchases of the relevant food, whether because the consumer has higher income or has stronger preferences for this food, the greater the absolute size of the implicit transfer. The greater the percentage share of income typically devoted to purchases of this food, the greater the implicit transfer as a percentage of income.[9]

Income and substitution effects. We can gain insight into the comparative impacts on consumers of income increases and food price reductions by decomposing the overall effect of a price reduction into two components: the "income effect" and the "substitution effect." The **income effect** is the change in consumption that would have occurred had the consumer's utility been raised to the new level through an explicit transfer of income (rather than through the price reduction). The **substitution effect** is the *additional* change in consumption that arose because of the price change. In Figure 6.3b, the overall effect of a reduction in the price of food is to shift consumption from a_1 to a_2. To decompose this overall change into income and substitution effects, we draw in the hypothetical budget constraint that the consumer *would have faced* had her utility been increased to the higher level through an increase in income rather than through a price reduction. Such a budget constraint would be parallel to the original budget constraint but tangent to the new higher level of utility associated with the price reduction, like the one ending at c in the diagram. The income effect of the price

[8] If the price of food had risen, then the expenditure required to continue purchasing F_1 units of food and N_1 units of non-food would have risen, while the consumer's income remained the same. The gap would represent a tax, or reduction in purchasing power, to which the consumer must respond by reducing consumption of one or both goods.

[9] The approximate implicit transfer expressed as a percentage of income is $((p_1 - p_2)*F_1/Y)*100$. If we multiply and divide this expression by p_1, we can re-express it as $[(p_1-p_2)/p_1)*100]*\ [p_1F_1/Y]$, where the expression in the first set of square brackets is the percentage reduction in the price and the expression in the second square bracket is the fractional share of income this consumer devotes to food purchases. If the price of a good falls by 20 percent, and the consumer typically spends 50 percent of her income on this good, then the implicit income transfer she receives is equivalent to an approximately 10 percent income increase.

reduction is thus illustrated by the hypothetical shift in consumption from a_1 to a_3, while the substitution effect is illustrated by the shift from a_3 to a_2.

While the income effect of a food price reduction on the consumption of food may be positive or negative, depending on whether the food is a normal or inferior good, the substitution effect *must* be positive. To see this notice that the budget constraint ending at b_2 has a shallower slope than the budget constraint ending at c. It reflects a smaller ratio P_f/P_n, because food has become cheaper relative to non-food. The substitution effect is thus associated with a shift along an indifference curve from a point at which the slope is steeper to a point at which the slope is flatter. Given the convex shape of any indifference curve, this shift must be associated with a movement to the right in the diagram, and thus with an increase in food consumption and reduction in non-food consumption. This captures our intuition that when prices change people try to substitute more of the goods that have become "better deals" for goods that have become worse deals. As a result, price reductions will tend to increase consumption of items that have become cheaper by *more* than a comparable increase in income would have increased it.

Nominal and real price changes. If Y, P_f and P_n (measured in pesos) all rose by the same percentage, the budget line in Figure 6.1a would remain unchanged (explain how you know this!), and we would see no change in F, N or our consumer's well-being. Such a change would thus have no **real** effect. Real effects on consumption and well-being happen only when prices and incomes change *relative* to one another. In periods of inflation, it is possible for all prices and incomes to rise by great amounts, without changing prices and income relative to one another, and thus without causing real effects. It is useful, then, when studying the effects of price and income movements on consumption choices and well-being, to distinguish between changes in the *level* of all income and prices, and changes in their relative values**.** Standard practice is to measure what is happening to the *level* of prices and income by constructing a **price index**, which is a weighted average of prices throughout the economy, where the weights reflect the relative importance of the associated goods in peoples' consumption expenditures. The actual price of rice in pesos, called the **nominal** price of rice, may rise either because the price level is increasing or because the price of rice is rising relative to other prices and income (or both are true). We define the **real** price of rice to be the ratio of the nominal price of rice to the overall price index, and use measures of the real price of rice to describe how the price of rice is changing relative to other prices in the economy. For the real price of rice to rise, the nominal price of rice must rise by more than the weighted average of other prices. Similarly, for real income to fall, the nominal value of income must rise by less (or fall by more) than the price index. For most of the remainder of this text, the prices and incomes of relevance to our analysis will be real rather than nominal prices and incomes.

Choices among multiple foods. New questions about the effect of food price changes on food consumption choices are raised when we recognize that real people allocate consumption expenditure not just across "food" and "non-food", but across many foods and many non-food items. As incomes rise, households may increase their consumption of some foods more than others, and will even reduce their consumption of inferior foods.

We can think of their food consumption choices as being shaped by their more fundamental demands for calories and other nutrient intakes, taste, variety and convenience. At low income levels they may be able to afford little in the way of taste, variety, convenience or important nutrients such as those found in meat, dairy and fresh vegetables, needing to spend large fractions of their food budget on unprocessed grains that offer them the best chance of meeting their caloric requirements on a low budget. It is reasonable to assume, however, that variety, flavor and convenience are normal goods, in which case as incomes rise, consumers shift toward the purchase of more expensive meat and dairy products, fresh vegetables, processed foods and prepared meals, which provide them with more flavor, status and convenience, but also provide them with fewer calories per dollar spent.

When the price of a specific food like corn falls, consumption of corn is likely to rise, but consumption of other foods may fall, rendering the overall impact on food expenditure and calorie consumption unclear. Consumption of foods that people like to consume in combination with corn, which are called **complements** to corn, will increase. But consumption of foods that satisfy needs very similar to those satisfied by corn (such as wheat), called **substitutes** for corn in consumption, will fall. (Consumption of yet other goods may change little. Such items and corn are said to be **independent** in consumption.)

Changing social norms. The way a consumer evaluates the utility of diverse consumption choices may be influenced by her perception of **social norms**, or the rules of behavior that people believe they must follow to achieve social approval. These perceptions are shaped by the behaviors and beliefs of people around her, especially co-members of the peer, community, ethnic and religious communities with which she identifies. As social norms evolve, her indifference curves might shift and take new shapes, modifying the way she allocates income across food and non-food, and across specific foods. For example, a community might initially associate consumption of processed white rice with higher status than consumption of unprocessed brown rice ("poor peoples' food"), but information campaigns promoting the belief that brown rice consumption is healthier might render a community more comfortable with brown rice consumption, leading the consumer to increase brown rice consumption and reduce white rice consumption even if facing an unchanged budget constraint.

Elasticities. While basic consumer theory points to many logical effects of income, prices and other socio-economic circumstances on a consumer's purchases of various foods, it does not (and cannot!) tell us how large or small these effects will be. Only empirical research can shed light on those magnitudes. A concept commonly employed in discussing the magnitudes of these relationships is **elasticity**. Let R be the quantity of rice consumed. The elasticity of R with respect to Y measures the percentage change in R brought about by a 1 percent increase in Y (while holding prices constant). This **income elasticity** should be positive for normal goods. Similarly, the elasticity of R with respect to P_R measures the percentage change in R brought about by a 1 percent increase in P_R (while holding Y and other prices constant). Except in unusual cases, we expect such an **own price elasticity** to be negative.

Elasticities are useful measures, because we can roughly assess their size without reference to the units in which we measure foods or the typical level of prices, simply by comparing their values to benchmarks of 1, 0 or -1. If the income elasticity of rice consumption is greater (less) than 1, food consumption rises more (less) than in proportion to income. If the income elasticity were exactly equal to one, consumption would rise exactly in proportion to income, so that the ratio of rice consumption expenditure to income remains constant. If the own price elasticity of rice consumption is less than -1 (or "greater than 1 in absolute value") then when P_R rises by 10 percent, R falls by more than 10 percent, so that the expenditure on rice ($P_R R$) falls. In such a case rice demand is said to be "price elastic". If the own price elasticity were less than -1 in absolute value, then R falls by less than 10 percent, expenditure on rice increases, and rice demand is considered "price inelastic".

Application of Basic Consumer Theory to the Study of Nutrition Policy

When viewed from a "macro" perspective, malnutrition and famine may appear to result from a shortage of available food supplies relative to a population's food requirements. The "micro" perspective described above reveals, however, that malnutrition and famine are possible even when ample food supplies are available. People with low incomes facing high prices may simply be unable to purchase what they need, even if ample food supplies are available. For an eloquent elaboration of this theme, see Sen (1981).

Distributing cash, increasing incomes. If people are malnourished because they are poor, it is natural to consider improving their nutrition by increasing their income, whether directly through a cash transfer program or indirectly through development policies aimed at broad-based income improvements. Basic consumer theory reminds us that households use additional income in the ways they consider best, and do not necessarily spend it all on improving nutrition.[10] Thus the theory raises a first important empirical question: How effective is additional income in improving the nutrition of the poor?

Fortunately, basic consumer theory offers structure to guide empirical study of such questions. It reminds researchers that when seeking to measure the causal effect of income on consumption using data from multiple households or time periods, they must "control for" or "hold constant" all other socio-economic circumstances (other than income) that might differ across observations in the dataset and might help to explain differences in consumption. If households with higher incomes tend to live in communities with lower real food prices, then simple correlations between income and food consumption (which do not control for variation in food prices) would tend to pick up not only the effect of income on food consumption, but also some of the effects of lower food prices. Such estimates might incorrectly indicate that income is a potent force for increasing food consumption, when in fact income alone has little impact.

[10] We later learn additional reasons why the receipt of additional income may tend to increase food consumption expenditures by only a fraction of the cash received from a transfer program. Other households who had been providing help to a recipient household may decide to withdraw their support, and the recipients themselves may decide to reduce their efforts to earn income from other sources.

Efforts to estimate income effects while controlling properly for other socio-economic circumstances often involve multiple regression analysis. Basic consumer theory tells us that the list of "all other socio-economic circumstances," for which we need to include controls on the right hand side of our regressions, should include prices of food and other items, and demographic and cultural variables that affect households' preferences. The theory also points to the importance of measuring income and prices in real rather than nominal terms. For more on econometric estimation of the impacts of income and prices on consumer choices, see Pollak and Wales (1992).

Empirical studies from around the world, guided by the basic theory, have demonstrated that, even in very poor populations, additional cash leads to significantly less than a one-for-one increase in expenditures on food. Often even hungry people will spend 10 to 20 percent of any additional cash on non-food items, because they have urgent non-food needs, perhaps for home repairs, basic cooking utensils, children's clothing or feed for goats or chickens (Alderman, 1986; Barrett, 2002). Studies demonstrate further that calorie consumption rises much more slowly than total food consumption expenditures, because as incomes rise people dedicate more of their food budget to foods that provide fewer calories per peso of expenditure but offer more in the way of flavor, status or convenience. For example, Bouis and Haddad (1992) estimate that the elasticity of calorie intake with respect to income among low-income households in the Philippines is very low: in the 0.08-0.14 range. Other researchers argue that the elasticity is somewhat higher than this, in the .3 to .5 range (e.g. Subramanian and Deaton, 1996), but all agree that the elasticity is quite significantly below 1. As a result, while general processes of broad-based economic growth, and more specific efforts to give cash to poor households, can and do improve nutrition, they tend to do so only slowly (World Bank, 2006).

Distributing food. If giving cash to households has only modest impacts on nutrition, should policymakers consider giving households food rather than cash? Basic consumer theory offers some surprising answers. Consider Figure 6.4, in which we examine the impact of giving free rice to a consumer, under two alternative assumptions about the nature of the consumer's preferences. In both panels the horizontal axis measures consumption of rice (R), while the vertical axis measures consumption of all other items (O), including other foods as well as non-food items. The dark diagonal line segments extending from Y/P_o to Y/P_r indicate the budget constraints prior to the distribution of rice. The post-distribution budget constraint is given by the dashed line segments in Figure 6.4a. To understand why the new budget constraint has this "kinked" form, consider the following. If, after the rice distribution, the household chose to spend all of its cash income (Y) on items other than rice, it would be able to consume Y/P_o units of other items, but, as the result of the rice distribution, would also be able to consume R units of rice. Thus instead of starting at Y/P_o on the vertical axis, the relevant budget constraint starts at the point *a.* Starting from *a,* if the household wishes to consume more rice, it faces the same trade-off as before the distribution: it must give up P_r/P_o units of other items to be able to purchase one more peso's worth of rice. Thus to the right of *a*

the post-distribution budget constraint has the same slope as the pre-distribution constraint (lying *R* units to the right of it).[11]

Figure 6.4a

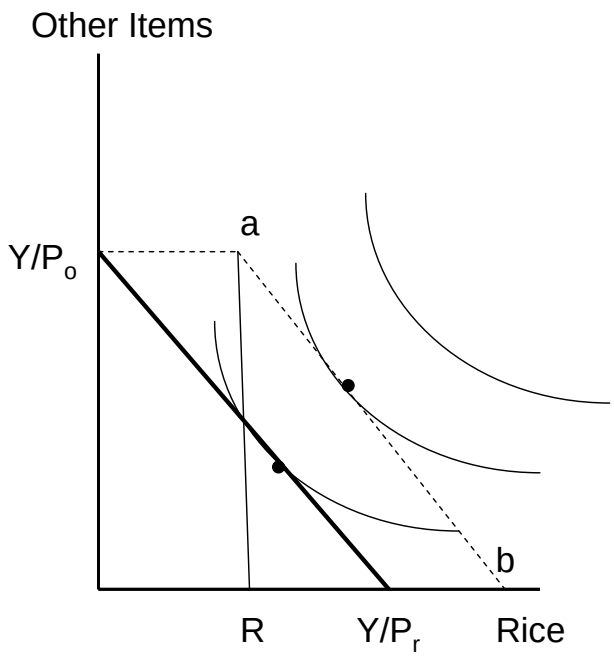

The receipt of R kilograms of free rice leaves the consumer facing a "kinked" budget constraint. For this consumer the impact of the food receipt is exactly the same as the impact of receiving the cash value of the free food.

Figure 6.4b

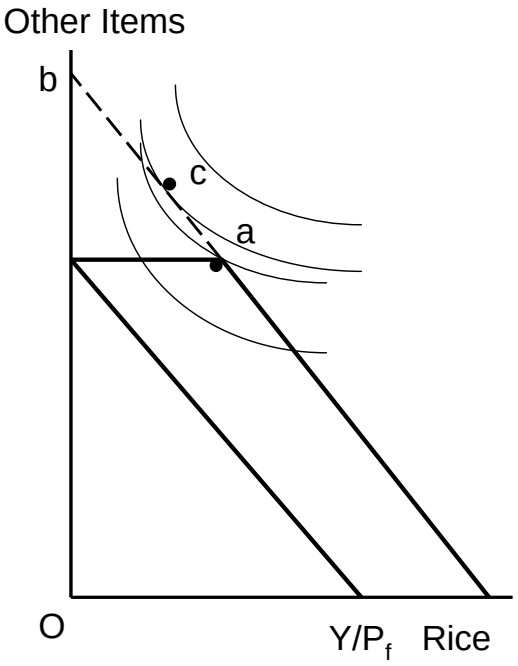

For this consumer the receipt of free rice causes rice consumption to rise by more than the receipt of the cash value of the rice, and causes the consumer's utility to rise by less.

In both panels of Figure 6.4 we identify post-distribution consumption choices by identifying the point on the kinked budget constraint at which the household achieves the highest level of utility. In panel a, the highest utility is found at a tangency between an indifference curve and the sloped segment of the budget constraint. At such a point the consumer consumes all the rice distributed (*R*) *plus* some additional rice that it purchases out of its cash income (*Y)*. In panel b, depicting a consumer that places greater value on consumption of Other relative to Rice, the highest utility is found at the kink point. The consumer consumes the entire rice distribution, but purchases no additional rice, instead spending all its non-program income on Other items.

Regardless of which diagram we use to analyze the consumption impact of the rice transfer, the basic theory forces us to recognize an important fact: The quantity of rice the recipient consumes is likely to rise by *less* than the quantity of rice it receives in the rice distribution. Why? We can think of the distribution as allowing them to replace

[11] We are holding the market prices of rice and other items constant in our analysis here. In later chapters we will identify circumstances under which these prices might change as a result of the food distribution, requiring adjustments to the analysis.

costly purchased rice by free rice, freeing up income that the household may now choose to spend on whatever it wishes, and the increased income will tend to increase purchases of all normal goods (and not just rice).

In fact, the basic theory leads to the even more extreme conclusion: If rice recipients consume more rice than the quantity distributed, then (under the assumptions of the basic theory) the impact of the rice transfer on consumption choices is *exactly the same* as that of a cash transfer of equal value. To see this, consider Figure 6.4a. Use your hand to cover the portion of the diagram to the left of the kink point in the budget constraint. What does the un-covered portion of the diagram look like? It looks exactly like a diagram analyzing the impact of a cash transfer, with a new budget constraint parallel to the original budget constraint. How much cash would be required to push the budget constraint to this position? Knowing that the budget constraint passes through point a, we know that the household's total income (after an infusion of cash) would have to be enough to purchase R units of rice as well as Y/P_0 units of other items. Thus the cash transfer that would cause the budget constraint to lie along the *ab* segment is $R*P_r$, which is the market value of the distributed rice. Every unit of free rice distributed replaces a unit of rice the household would have purchased, thus freeing up income equal to the full value of the food distribution. The impact on consumption of this implicit income increase is the same as if the consumer had been given cash, because "on the margin" (i.e. when consumers are thinking about how to spend their last pesos of cash income) they are still interested in buying both rice and other items, and they must pay market prices for both.

Empirical studies do indeed bear out the theoretical prediction that food distributions increase food consumption by less than the quantities of food distributed. But they also often find that food distributions tend to increase consumption of the distributed foods by more than equivalent transfers of cash would (Barrett, 1996), contradicting the extreme conclusion of the basic theory. Such results contribute to the motivation for more careful thought about household decision-making, such as we present in Chapter 7. It remains the case, however, that food transfer impacts are much more similar to cash transfer impacts than we might have guessed.

Consider now panel 6.4b. For consumers with preferences like this, the rice distribution does indeed increase their rice consumption by more than if they had received a cash transfer of comparable value, as we will argue below. Even so, the rice distribution still increases consumption of other items as well as consumption of rice.

Had a rice recipient with preferences as described in Figure 6.4b been given sufficient cash to consume at point a (rather than being given rice), the horizontal portion of their post-distribution budget constraint would have been replaced by the segment of long dashes between points *a* and *c*. A cash transfer just large enough to allow purchase of the distributed quantity of rice (*R)* at the market price would lead them to a budget constraint that runs through the kink point *a*, but is a straight line with slope given by the market prices of rice and other items. Facing such a budget constraint, the consumer would consume at point *c*, where an indifference curve is tangent to the dashed portion of the

budget constraint. Point *c must* correspond to a higher level of well-being than point *a*, but a lower level of rice consumption, because the slope of the indifference curve touching point a must be less than the slope of the cash transfer budget constraint.[12] This means that the cash transfer budget constraint cuts through the indifference curve at point *a*, heading to the north-west, and cutting from lower levels of utility than that associated with point a to higher levels. The consumer finds a consumption combination involving less rice consumption and more consumption of other items (relative to point *a)* that is feasible under the cash transfer alternative *and* leaves the household better off than it would be with the rice distribution.

This leads to a stark and somewhat surprising conclusion of basic consumer theory: While (under the theory's assumptions) food transfers have the potential to increase consumption of specific foods by more than cash transfers of comparable value, they do so only under conditions in which the food transfers leave recipients *worse off* than comparable cash transfers. The conditions required for rice distributions to increase rice consumption by more than cash transfers are those in which beneficiary households do not choose to spend any of their own income on rice after receiving free rice (as in Figure 6.4b). This is most likely when rice distributions are large relative to household incomes (e.g. it is a true emergency situation, in which households have little income with which to purchase food), or when the program distributes foods that are not commonly used and not much preferred locally (e.g. when U.S. wheat received as food aid is distributed in regions of Africa where wheat is not commonly consumed), so that households would rather use increased resources to increase purchases of other food or non-food items. According to the basic analysis, under these conditions, the food distribution does indeed increase their consumption of the distributed food by more than a comparable cash transfer would, but the households would have preferred to receive the cash equivalent, because that would have given them more freedom to use additional resources to consume other items they value more highly than the program foods.[13,14]

Anecdotal evidence suggests that for analysis of some food transfers, especially when foreign grains are distributed in poor regions, Figure 6.4b is indeed relevant. Recipients of food transfers sometimes report wishing that they had received the transfer in the form

---

[12] Somewhat more precisely, there is a very small chance that the indifference curve touching point *a* just happens to be tangent there, indicating that the consumer would use a cash transfer to purchase exactly the amount of rice distributed by the rice distribution. But there is a much greater chance that the slope at point a will be less than that.

[13] If the utility that households could achieve at point *c* is much higher than the utility they achieve at point *a*, then they may be willing to engage in costly activities to turn some of their free food receipts into cash, by re-selling them. Such re-sale transactions are indeed costly. For many households the costs outweigh the benefits, and they content themselves with consuming at point *a*. For some the benefits may outweigh the costs, but the costs are still significant. Thus while they may manage to consume somewhere to the northwest of point *a*, they exhaust some of their income in carrying out re-sale transactions, and thus end up worse off than if they had been given cash in the first place.

[14] As indicated above, empirical studies often indicate that food transfers have larger impacts on food consumption than cash transfers of the same value. While the discussion of Figure 6.4b offers a possible explanation for this result, it often falls short of providing a convincing explanation, because often food transfer sizes are small relative to what households are already consuming of the distributed items, suggesting that they should be found in circumstances better described by Figure 6.4a than Figure 6.4b.

of cash rather than food, and even engage in costly efforts to sell the foreign grain in order to garner cash for purchasing more familiar local grains and other items. Even very poor recipients of cash transfers also sometimes use the cash to purchase feed for livestock, purchase fertilizer or send children to school, which would not have been feasible without the distribution (MCDSS/GTZ 2007, Sadoulet et al. 2001), suggesting a longer-term potential benefit to distributing cash rather than food.[15]

Reducing food prices. For many years explicit targeting of cash or food transfers to needy households was considered beyond the bureaucratic capacity of many developing countries. (See Chapter 21 on some of the difficulties involved.) Policymakers thus pursued an alternative approach to improving nutrition: reducing the economy-wide prices of key foods, thereby making the acquisition of adequate nutrition easier without the need for bureaucratic transfer distributions. Indeed, developing country governments have a long history of orchestrating "general consumer food subsidies"(see Chapter 15).

Basic consumer theory points out several weaknesses of using food price reductions rather than targeted income increases to improve nutrition. First, while a wheat price reduction increases a consumer's wheat consumption by more than a cash transfer of equivalent value, it increase the consumer's well-being by less. (See problem 3 at the end of the chapter, which employs logic very similar to the logic employed above for comparing the impacts of cash and food distributions.)

A second weakness of general food subsidies arises because reducing the price of a food channels costly benefits to *everyone* who purchases that food. Wealthy and well-nourished households often purchase larger absolute quantities of wheat than do the poor and malnourished. As a result, the non-poor often enjoy large fractions of the implicit income transfers embodied in general food subsidies, and the cost of such policies per dollar of benefit transferred to *poor* households at risk of malnutrition tends to be very high.[16]

Finally, while a general consumer wheat subsidy increases consumption of wheat, it may also reduce consumption of foods that are substitutes for wheat, raising empirical questions about the ultimate impact on total food consumption and nutrition. Results can sometimes be perverse. For example, a wheat subsidy in Brazil is widely thought to have *reduced* the nutritional status of the poor, because it caused them to shift consumption from rice toward bread, which offered fewer calories per dollar spent (Calegar and Schuh, 1988).

---

[15] We will also learn in later chapters that in some circumstances cash transfers will be preferable to food transfers because cash transfers will tend to stimulate local agricultural production, while food distributions may tend to depress local production. On the other hand, where food markets are not well established, food distributions may have the advantage of bringing more rapid improvements in consumption, because they do not require recipients to wait for the private sector to develop to turn what they have received into actual consumption.

[16] It may be possible to channel a larger fraction of the benefits of a general food subsidy to the poor by choosing to subsidize particular foods that are consumed very little by the rich. Unfortunately, as Alderman and Lindert (1998) point out, the foods that are consumed little by the rich also tend to constitute only a small fraction of the budget of the poor as well, limiting the poverty reduction potential of a subsidy.

Deeper questions. If consumers *choose* to spend only a fraction of a cash transfer on improved nutrition, the theory forces us to ask: On what grounds might policymakers claim that it would have been "better" for households to use the additional resources in a more nutrition-enhancing way? That is, on what grounds might policymakers claim that they are in a better position to judge what is best for a household than the household itself?

The most plausible ground on which to make these claims arises because policymakers may have superior knowledge regarding the implications of nutrient intakes for consumers' health and happiness.[17] We might wish to define "true" well-being as the utility households *would* derive from consumption *if* they had full understanding of the connections between food consumption, nutrient intakes, health, life expectancy and well-being. Lacking such information, the utility they maximize differs from true well-being, in a way that undervalues good nutrition. Studies indeed verify that households in which mothers have more schooling or greater nutrition knowledge acquired through other sources use identical incomes to purchase a more nutritious basket of consumption items relative to households with less knowledgeable mothers (Block, 2004). Thus well-informed policymakers may indeed have better understanding of nutrition implications than do some poorly educated consumers.

Nutrition education. But if the crux of the problem (of consumers spending "too small" a fraction of cash transfers on nutritious food) is that consumers lack adequate nutrition knowledge, then policymakers ought to consider attacking the problem directly, by educating consumers about the nature and effects of good nutrition. If the cost and difficulty of imparting nutrition knowledge are not too great, then programs offering cash transfers together with nutrition education may be superior to either food-based policies or simple cash transfers. A variety of cash transfer programs around the world now require beneficiaries to receive instruction in good nutrition. In the United States the Women, Infants and Children (WIC) program combines the distribution of vouchers to use in purchasing food with nutrition education, and is widely considered a great success (Barrett, 2002).

Basic labor supply theory: allocating time

A motivation. The International Labour Organization (ILO) estimates that 211 million children aged 5 to 14, and another 141 million children aged 15 to 17 are involved in some form work (ILO, 2004). Many observers see such child labor as depriving children of their right to a proper childhood, exposing them to harsh working conditions, and preventing them from improving their future prospects by attending school. Such concerns have fueled efforts to prohibit the employment of child labor, condition international trade agreements on the enforcement of such prohibitions, and convince

[17] Experts point out that it is often very difficult for parents to discern whether infants and young children are consuming enough calories until problems have emerged and it is too late. Many other forms of malnutrition and their implications for health and well-being are even more subtle.

developed country consumers to buy only imported goods that have been certified as "child labor free" (Edmonds, 2008).

In weighing the likely benefits and costs of child labor prohibitions, we must assess the likely consequences for children and their families of eliminating child labor opportunities. Before demanding child labor prohibitions, we should also consider alternative policy approaches to reducing child labor or reducing the ills associated with child labor (by, e.g., helping more working children attend school or improving their working conditions). In this section we develop analytical tools to aid us in these activities: basic economic theories of labor supply choices, first by individuals and then by multi-member families.

Basic assumptions. Except in cases of slavery or other forms of coercion, it seems reasonable to assume that people work as much as they do because they believe that the benefits of working this much outweigh the costs. Micro-economists point out that the primary benefit of working is the ability to consume goods and services purchased with the income work generates. The primary cost of devoting time to work is the foregone opportunity to use that time for other valued activities, such as rest, recreation, cooking, childcare or schooling. Basic labor supply theory highlights the implications of this by focusing on a "labor supplier" maximizes his utility subject to his **total time constraint**, which recognizes that the total amount of time allocated to all activities cannot exceed the total time he has available (e.g. the number of waking hours per week). His objective is to maximize his utility, which is assumed to be a function of the goods and services he can acquire with income, as well as his enjoyment of non-work activities.

Consider a labor supplier who must allocate a total of T hours per week to just two activities: home time and wage labor. "**Home time**" refers to time spent in recreation, education, child care and other home activities, which enters directly into his utility calculation as a "good" he values consuming. He values time spent in wage labor only because it allows him to earn a wage of w pesos per hour, with which he can purchase goods and services. He may also receive M pesos of **non-labor income,** which is income he receives regardless of how many hours he works, including transfer payments from the government or an NGO, gifts from other households, and receipts of rental payments on properties he owns and rents out. If he supplies S hours to wage labor, he is able to enjoy a total consumption expenditure of $C = M + wS$ pesos, but is left with only $H=T-S$ hours of home time. Thus his decision regarding how many hours to work determines what combination of consumption expenditure and home time he will "consume." He chooses C and H to maximize utility, $U(C,H)$, subject to the constraint that $C=M + w(T-H)$.

In Figure 6.5a the horizontal axis measures home time in hours per week, while the vertical axis measures total consumption expenditure in pesos (C).[18] Let's first trace out

[18] For simplicity of expositing this diagram, we initially hold constant all prices of consumption goods, so that increases in the nominal wage and nominal consumption expenditure always represent increases in real purchasing power and real consumption. This will allow us to discuss the diagram using simple language involving wages and consumption expenditure in "pesos." Later we will recognize that what really matters

the implications of the time constraint for our labor supplier's ability to consume combinations of *C* and *H*. If he chose to spend all available time in home time, then he would enjoy *T* hours of home time, but would earn no labor income and would be able to purchase only *M* pesos of goods and services, as indicated by point a in the diagram. If, at the other extreme, he devoted all time to wage labor, he would consume no home time but would consume *M+wT* in consumption expenditure, as indicated by point *b* in the diagram. In between these extremes, each time he reduces his home time by one hour (moving to the left by one unit in the diagram) he increases his possible consumption expenditure by *w* pesos (rising by that amount in the diagram), thus the slope of the time constraint connecting *a* and *b* is given by *–w*. We can think of this line as representing a budget constraint limiting the combinations of *H* and *C* that the labor supplier can "purchase" given the time, non-labor income and wage available to him. It is as if the labor supplier is endowed with an income of *M+wT*, which he may use to "buy" either consumption expenditure of goods and services (at a price of one dollar of income per dollar of consumption expenditure) or home time (at a price of w dollars per hour).

Figure 6.5a

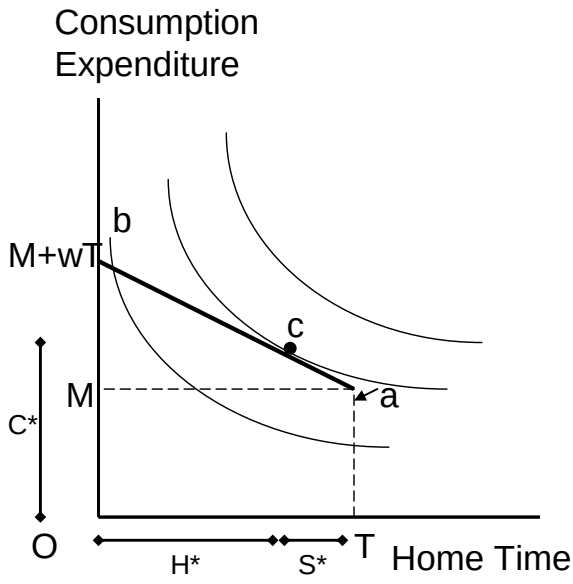

This labor supplier chooses to allocate *S** hours of time to labor and *H** hours to home time.

Figure 6.5b

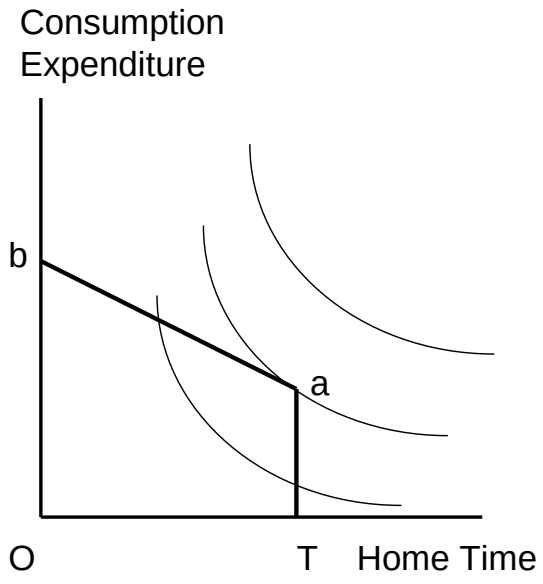

This labor supplier chooses to supply no labor, and instead to allocate all available time to home time.

Assuming that both *C* and *H* are "goods," we describe the labor supplier's preferences using indifference curves with similar properties to those employed in basic consumer theory, and assume he consumes at the point along the budget constraint that lies on the highest indifference curve. With preferences like those described by indifference curves in Figure 6.5a, he consumes at point *c,* where an indifference curve is just tangent to the

---

here is the real wage and real consumption expenditure. The real wage may rise either because the nominal wage rises or because the prices of some goods and services fall.

budget constraint. He chooses to consume H^* hours of home time, and devote $S^*=T-H^*$ hours to work, using the resulting wage income for consumption expenditure of C^*.

Corner solutions. As long as our labor supplier receives some non-labor income, his budget constraint has a "corner" to it at point *a*, rather than simply ending in an intercept along the horizontal axis. In economics jargon, this budget constraint is "kinked." When budget constraints are kinked like this, it becomes possible that the choice that maximizes utility is found at the kink or corner. An example of such a "**corner solution**" (to a labor supplier's utility maximization problem) is depicted in Figure 6.5b. The budget constraint there is identical to the budget constraint in Figure 6.5a. The figure differs only in that it describes a labor supplier with different preferences, who apparently places greater value on home time (perhaps because he has more young children requiring care). Of all the points on this person's budget constraint, the one associated with the highest level of utility is point *a*. Thus this labor supplier chooses to supply no labor, allocating all available time to home activities.

Effects of an increase in non-labor income. If we increase non-labor income (M) while holding the wage constant, then the kink point *a* and the vertical axis intercept $wT + M$ both rise by the same amount in either panel of Figure 6.5, and the budget constraints shift up (in a parallel fashion) by the amount of the non-labor income increase. (Try drawing this.) Whether our labor supplier is initially supplying labor (as in panel a) or not (as in panel b), this increase certainly increases his utility, because it expands the range of feasible consumption pairs. If the decision-maker initially supplies some labor as in Figure 6.5a, and if H and C are normal goods, then both H and C increase. As home time increases, the supply of time to wage labor falls. Because wage labor income falls, consumption expenditure (C) rises by less than the increase in M. In the case of the corner solution (Figure 6.5b), however, the increase in M has no impact on labor supply, which remains at zero hours, and every additional dollar of M translates into an additional dollar of C.

Effects of an increase in the wage. When we hold M constant and increase the wage w, the budget line's kink point a remains in the same place, while its intercept along the vertical axis ($M+wT$) rises. (Try drawing this.) Again, the range of feasible consumption pairs expands and utility tends to increase (though it may stay the same if the labor supplier starts at the corner solution). The wage change induces an "income effect" (associated with the increase in the cash value of the labor supplier's time endowment) and a "substitution" effect (associated with the increased cost, in terms of consumption expenditure foregone, of consuming home time) on time allocation, which are likely to work in opposite directions. (See Problem 4 at the end of the chapter.) If H is a normal good, then the income effect causes H to rise and S to fall. But the wage increase raises the price of home time relative to work time, inducing a substitution effect away from home time and into work time. Thus it is theoretically possible for labor supply (S) to increase, decrease, or be little changed. In most populations we expect labor supply to rise when the wage rises, though whether the effect will be large or small in any given context is an empirical question.

Multiple uses of time. In reality, people must allocate their time across more than two uses. Much of what we learn from the model with only labor and home time remains relevant and extrapolates in a straightforward way to the more complicated case. Increases in non-labor income increase the allocation of time to any uses that can be thought of as "normal goods." An increase in the wage brings an income effect that tends to increase all those same "normal" activities while reducing labor supply, and also substitution effects that tend to increase labor supply relative to other time uses.

Acknowledging multiple time uses does, however, modify the way we think about interactions among time uses. When time is allocated only between working and attending school, then when an increase in the wage increases labor supply, school attendance certainly falls. Once we recognize that time may be allocated to household chores, rest and recreation, as well as to work and school, however, it is no longer inevitable that increases in work time are associated with reductions in school attendance. The additional time that goes into wage labor may come out of time allocated to household chores or recreation rather than out of school time.

Effects of programs that free up more time. Households in developing countries often spend long hours carrying water from distant sources or searching for fire wood. The time spent in such activities is unavailable either for income-generating work or for enjoying home time. Within the simple labor supply framework developed here, we can model the effect of infrastructure projects that provide households with clean water or energy sources close to their homes as expanding the total amount of time (T) that is available for wage work and home time. Expanding total available time from T to T^* would shift the vertical portion of the budget constraint to the right, and would increase the budget constraint's intercept on the vertical axis by $w(T^*-T)$. (Try drawing this.) This represents an expansion of the feasible set of consumption pairs, and leaves the labor supplier better off. How this extra time will be allocated across increases in home time and work (and other time uses) will depend on the nature of the labor supplier's preferences.

More generally, the basic theory highlights the important observation that time is valuable. Remembering this simple point will be of great importance later in the text, when we contemplate households' choices about whether or not to take advantage of new opportunities created by policies and programs. When they weigh the benefits and costs of participating, households will consider the time required or freed up by participation, in addition to monetary benefits and costs.

Norms and institutions. Social norms and other non-market institutions may help shape labor supply choices. Cultural or legal prohibitions may prevent households from allocating any time of its female members to wage labor outside the home, essentially limiting the total time T that may be allocated to wage labor activities. Social norms may also shape the value parents place on sending their children to school rather than keeping them home to work in the fields.

Multi-member households. Most potential workers live in multi-member households; and members of the same household are likely to make time allocation decisions in an interactive way. For example, they may choose to have some members specialize in wage labor outside the home, while others work in a home-based enterprise or undertake the production of child care, cooking and cleaning services for the household. To understand the labor supply behavior of any one member, then, it is important to consider the work opportunities available to all members, and the value the household places on the enjoyment of non-work time by various members.

If we are willing to assume that a single decision-maker within the household (or the entire household acting unanimously) makes a joint choice about the allocation of every member's time, then we can analyze such choices by treating home time for each member as a separate good that the household values consuming. Members who face lower wages or higher productivity in home-based work or schooling, or whose rest and recreation the household values more highly, will be less likely to work than other members. To the extent that the household treats the non-work time of a particular member as a normal good, increases in income will tend to reduce that member's labor supply. (We will point out some inadequacies of this view of multi-member households in the next chapter.)

Application of Basic Labor Supply Theory to the Study of Child Labor

The Progessive agenda in the late nineteenth century encouraged a "sacralization" of childhood in the United States, in which childhoods of schooling and play took on special moral and sentimental significance, and came to be considered a human right (Zelizer, 1994). Increased recognition in recent decades of education's vital role in development further fueled concern that child labor might prevent the next generation from attaining greater well-being than the last. Thus is it not surprising that many activists seek to prohibit child labor around the world. Basic labor supply theory, however, urges careful empirical study before concluding that the prohibition of child labor would improve life for the world's children or increase investments in education. If children work because the benefits of increased family consumption (made possible by their earnings) outweigh the costs, then eliminating the opportunity for child labor may leave children worse off, even when it succeeds in preventing them from working. Furthermore, given children's potential to allocate time to multiple activities, working need not preclude school attendance, and eliminating child labor opportunities might not increase school enrollment. The comprehensive review of empirical research on child labor presented in Edmonds (2008) sheds light (albeit imperfectly) on these issues. Here we point out a few key empirical observations.

Table 6.1 presents a small sample of the statistics Edmonds (2008) computes from household survey data collected in many countries around the world. For the purposes of this table, a child is recorded as engaged in a work activity if he or she devoted at least one hour to it in the last week. Children are coded as attending school if they have attended school in the last year. "Market work inside the home" involves participation in the production of goods and services on a family farm or in a family business. "Market

work outside the home" involves working for pay for other employers. "Domestic work" includes cooking, cleaning, carrying water, caring for siblings and similar activities. While domestic work is not included in the forms of child labor that activists usually wish to prohibit, it can sometimes absorb many hours of children's time, and helps households generate income by freeing up adults for self- or wage employment. Children are recorded as involved in "any work" if they are involved in any of these three work activities.

Table 6.1
Percentages of Children Aged 5 to 14 Undertaking Various Activities

Country	Any Work	Market Work Outside Home	Market Work Inside Home	Domestic Work	School	Any Work and No School	No Work and No School
Albania	62.7	3.5	29.5	56.1	54.7	32.6	12.7
Cameroun	85.2	30.8	42.7	81.1	94.5	5.0	0.4
Kenya	66.8	2.2	1.0	66.3	95.9	2.8	1.0
Mongolia	91.7	1.4	20.6	91.2	95.2	4.6	0.2
Venezuela	64.6	4.5	3.9	62.4	92.0	3.8	4.2
Vietnam	57.8	1.9	23.4	51.7	95.1	4.3	0.6

Source: Table 4 in Edmonds (2008).

Most child labor is quite unlike the image of children working long hours in sweatshops that we might pick up from the media. In the countries listed in Table 6.1, as in many other countries, the great majority of work performed by children is carried out within their own homes. In many cases they are working on family farms or in family businesses, where they are likely to be working shoulder to shoulder with their parents and where they may be learning how to run the family enterprise. Other data sources document that a very large fraction of child market work is in agriculture, forestry and fishing, while only a very small fraction is in manufacturing. Moreover, many working children are engaged in fewer than 10 hours of market work per week (see Figure 1 in Edmonds, 2008).

Poverty appears to play an important role in decisions regarding child labor, and children's contributions to family income are sometimes substantial. Cross country studies demonstrate that child labor is more prevalent in poorer countries, and studies that follow individual families over time often find that child participation in work rises significantly when the household is hit by an economic downturn, such as a crop failure or the loss of the household head's job (Beegle, et al., 2006; Duryea, et al., 2007). Studies estimate that labor by 13-year-olds contributed 13 percent of household income in Bolivia (Psacharolpolos, 1997) and that children in Nepal contribute 11 percent of the value of agricultural production (Menon, et al., 2005), suggesting that the loss of their income could be quite significant for poor households. A concern that child labor prohibitions might push some children deeper into poverty appears warranted.

Evidence furthermore casts doubt on the necessity and effectiveness of child labor prohibitions for increasing school enrollment. The majority of working children attend

school (see Table 6.1), and conditional cash transfer programs that pay benefits to families whose children attend school tend to raise schooling rates by significantly more than they reduce rates of child labor (de Janvry, et al., 2006). In a similar vein, where families have greater farm or business assets, which increase both income and the value of child labor in family businesses, children tend to work more, but also go to school more (Shafiq, 2006). Thus child labor need not preclude school attendance. Furthermore, in the countries in which larger fractions of working children do not attend school, large fractions of non-working children also remain out of school (see Table 6.1), suggesting that obstacles other than child labor may be more important in preventing school attendance. In some cases work by older siblings appears to help cover the cost of sending younger siblings to school (Edmonds, 2008). Thus, a child labor prohibition might even cause siblings of working children to drop out of school.

Given this evidence, we may wish to consider alternative approaches to reducing child labor or increasing schooling. Again, the basic time allocation model offers useful guidance. First, it points out that if families consider child non-work time a normal good, then raising family income might reduce child labor. Empirical study of the effect of household income growth on child labor is complicated by the fact that the incomes of the poor often grow precisely because wages for unskilled labor are rising. Thus wages paid for child labor may be rising at the same time as household income rises. While income growth might tend to reduce child labor, rising wages might tend to increase it. In Brazil, a coffee boom seems to have given rise to wage effects that outweighed the income effects for poorer households, causing child labor to increase. As Kruger (2007) points out, wage effects on child labor may have been especially strong because the coffee boom was expected to be temporary, and families expected children to be able to return to school after the boom ended. In Vietnam, however, trade reform-driven increases in rice prices, in the presence of an unusually egalitarian distribution of rice farming land, led to income increases that were broadly spread, including among many poor households. This gave rise to income effects that outweighed the wage effects and caused child labor to diminish (Edmonds and Pavcnik, 2006). A 30 percent increase in the price of rice was associated with a 9 percentage point drop in the share of children who were working, and the impact was bigger among households with more land, for whom the income effect would tend to be larger.

The basic time allocation model also suggests that schooling may be increased, with or without a reduction in child labor, by reducing the cost or increasing the perceived benefits of schooling. Schooling might thus be encouraged by improving school quality, improving parents' understanding of the value of education, reducing school fees, or increasing the immediate benefits of schooling by offering free school meals or conditional cash transfers (while holding constant the wage paid for child labor). We review evidence on such efforts in Chapters 18 and 21). The model furthermore suggests that projects building water and power infrastructure closer to children's homes can improve enrollment by freeing children from spending hours each day carrying water or gathering firewood. Anecdotes suggest this might be particularly important for raising enrollments of girls, whose responsibility it often is to collect water and fuel. Finally, practical contemplation of time allocation choices suggests that school enrollments might

be increased by efforts to render schooling and work more compatible. For example, many NGOs working to draw unschooled children into school argue the importance collaborating with parents to create school schedules that conflict as little as possible with the children's work responsibilities at home.

Despite casting some doubts on the advisability of generalized child labor prohibitions, the empirical evidence reported thus far falls far short of assuring us that decisions to put children to work are always made with children's best interest in mind, and that intervention is never required. The ILO estimated that 8.4 million children were involved in "unconditional worst forms of child labor", which include slavery and practices similar to slavery, as well as pornography, sex trafficking, and production or processing of drugs (ILO 2002). Cases of coercion and involvement in illegal activities certainly demand special attention.

Furthermore, even when children are working within their own homes, we cannot know that their interests are being fully weighed by household decision-makers. The basic time allocation model assumes that a single decision-maker within the household maximizes some notion of family well-being, but remains silent about exactly how he or she evaluates well-being. If decision-makers take children's interests fully into account, value children's well-being at least as highly as the well-being of other household members, and are fully informed about the impacts of work on their children's health and development, then they may make reasonable choices regarding when to put children to work. If, however, household decision-makers exercise coercive power over children, and seek to maximize only their own selfish interests, then allowing them to put their children to work may increase the suffering of children. Unfortunately, the basic time allocation model offers little guidance for empirical attempts to uncover such power relationships within households. In Chapter 7 we will examine more advanced "non-unitary" models of household decision-making, which researchers have developed in their efforts to understand better such intra-household relationships.

Basic producer theory: buying inputs to produce and sell outputs

A motivation. Over the last half of the twentieth century, global incomes rose while the world population doubled, nearly tripling the world demand for food. World food supply kept pace as a result of technical change in agriculture, which allowed large increases in per-acre yields of basic food crops (Ruttan, 2002). Much of the increased food production took place in the developing world, as a result of the "Green Revolution", the adoption of high yielding varieties of rice, corn and wheat, which were developed by international crop research institutes funded by the Ford and Rockefeller Foundations.

The Green Revolution contributed to GDP growth and prevented real food prices from rising. Yet for many years it remained controversial because of suspicions that it channeled disproportionate benefits to better-off farmers and even left some of the rural poor worse off. One suspicion was that it encouraged labor-displacing mechanization, which reduced (or at least prevented increases in) the demand for agricultural labor, preventing low-skilled and landless rural workers from sharing in the benefits of local

growth. Thus if we are to grasp fully the impacts of the Green Revolution, we must look closely at its impact on the demand for labor. The basic producer theory exposited in this section offers useful guidance for study of this and many other issues.

Basic assumptions. Basic producer theory focuses on a "producer" who hires labor and other inputs, uses them to produce goods or services, and sells the output. She chooses input and output quantities that maximize her **profits** subject to the constraints imposed by the technology available to her and the competition she faces. Her profits are just the difference between the **revenue** she derives from selling her outputs and her **expenditures** on inputs. The technology, or array of relevant technical options, she faces is summarized by a **production function**. (You may wish to review the introduction to production functions found in Chapter 3.) If she is a small participant in large, competitive markets for inputs and outputs, then she is a "**price taker**," and market conditions simply dictate the prices she must charge for outputs and the prices she must pay for inputs. We can think of her as endowed with a technology, managerial skills and perhaps the ownership of some factors of production such as land or tools. She makes decisions about how to put her endowments to work, and what additional inputs to combine them with.

Basic producer theory highlights the importance of information on producers' technological options, as well as the prices they face for outputs and inputs, in efforts to understand changes in their use of inputs and supply of outputs. The theory yields diverse insights and predictions, depending on the nature of the specific technologies involved, whether we are studying the "short run" or the "long run" (as defined below), and the nature of the markets in which the producer might participate. In this chapter we draw out some key insights by examining just a few relatively simple cases.

The case of one output, one variable input, perfectly competitive markets. Consider first a producer who manufactures a single output, cloth, according to the production function

(2) $Q = F(L;K, A)$

where Q is yards of cloth, L is hours of low-skill labor, K is the dollar value of capital equipment, A indexes the current state of technology, and the producer operates in perfectly competitive markets for cloth and labor. We focus on the **short run** in which the producer may choose the hours of labor L, but takes the quantity of capital (K) and level of technology (A) as given and unchanging. For reasons discussed in Chapter 3, we expect the function $F(.)$ to exhibit **diminishing marginal returns** to either input. That is, we expect $F(,)$ increases, at a decreasing rate, as either K or L increases (while holding the other constant). If p is the price per yard of cloth and w is the hourly wage for low-skill labor, then the producers' profit (when using L hours of labor and selling Q yards of cloth) is pQ-wL. Because the producer would never want to waste any inputs, she will always choose a combination of Q and L that is "on" the production function (2), which tells us the *maximum* quantity of cloth that can be obtained from L hours of labor, given the technological options and capital currently available.

This producer maximizes profits by deciding how much labor to employ, and thus how much output to produce, subject to the constraint represented by the production function. The two panels of Figure 6.6 offer two useful ways to summarize basic producer theory for this simple case. In Figure 6.6a the horizontal axis measures hours of labor (*L)*. The vertical axis is measured in pesos, the unit of measure for revenues, expenditures or costs, and profits. The downward-sloping line describes the contribution to total revenue of each successive unit of labor employed. It is equal to the value (at market price *p*) of the additional yards of cloth produced when each unit of labor is added to production (while holding *K* and *A* constant). That is, it is equal to the value of the marginal product of labor (VMP_L), which is just the (physical) marginal product of labor (MP_L) multiplied by the price of a unit of output (*p*). Our assumption of diminishing marginal returns to the use of either input implies that the function slopes downward. The horizontal line in the diagram is drawn at the height of the wage (*w*). It represents the marginal cost of a unit of labor, or the amount by which total costs increase when an additional unit of labor is added to the payroll.

Figure 6.6a

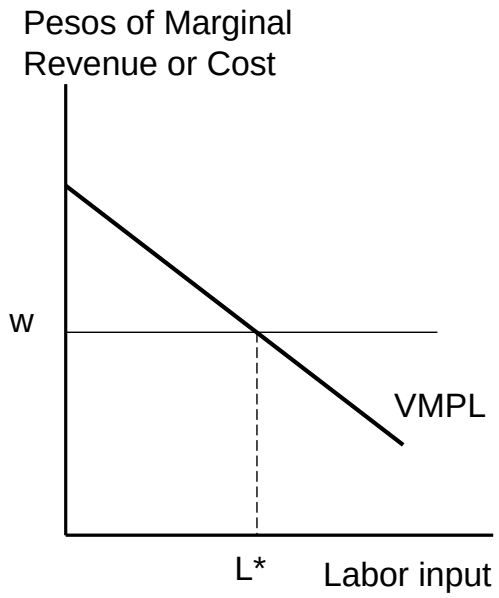

Short-run profit maximization, when labor is the only variable input, and the producer is a price taker in output and labor markets, requires that the producer employ the quantity of labor at which the VMPL is equal to the wage.

Figure 6.6b

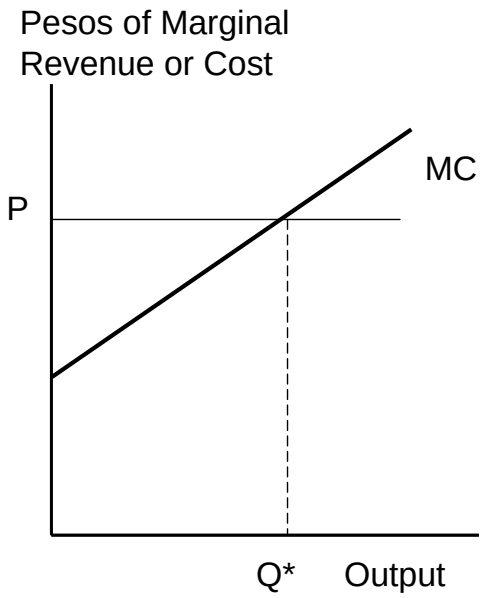

Short-run profit maximization, when labor is the only variable input, and the producer is a price taker in output and labor markets, requires that the producer produce the quantity of output at which the marginal cost equals the price.

The profit-maximizing level of labor use (*L**) is found at the intersection of the *VMPL* and *w* lines. If the producer chose to employ a quantity of labor less than *L**, and thus found herself using a quantity of labor to the left of *L** in the diagram, the *VMPL* would be greater than the wage. This would imply that total profits could be increased by increasing *L*, because adding another unit of labor would increase revenue by the *VMPL*, which is more than the associated increase in total costs (*w*). Similarly, if the producer started with more than *L** units of labor (thus operating at a point to the right of *L** in the

diagram), she would be able to increase profits by reducing L, because a reduction in L would reduce revenue by less than it would reduce cost. Only at L^* is there no potential to increase profits by increasing or decreasing L.

Figure 6.6b describes the very same profit-maximizing choice in another way. Here the horizontal axis measures yards of cloth Q rather than L. Again the vertical axis is measured in pesos, allowing us to examine the increases in revenues and costs associated with each additional unit of output the producer might contemplate. The line drawn at the height p indicates the addition to revenue associated with each additional yard of cloth produced. The marginal cost (MC) line indicates the addition to cost associated with each additional yard of cloth produced. The marginal cost equals the product of the wage (w) and the quantity of labor that would be required to produce the additional yard of cloth, and thus equals $w(1/MP_L)$. Knowing that the marginal product of labor declines as L increases, we know that the marginal cost of producing an additional yard of cloth must increase as Q increases, thus the marginal cost schedule rises as we move to the right in the diagram. As long as the marginal cost lies below the marginal revenue (p), producing more cloth would increase profits. Profits are maximized at the quantity Q^*, where the marginal cost just equals p. Additional increases in Q would add more to cost than to revenue and would thus reduce profits. This is just another way of looking at the same decision described in Figure 6.6a. As a result, we know that Q^* must equal $F(L^*;K,A)$.

Effects of changes in output price, wage and technology. Using either of the diagrams in Figure 6.6, you should be able to demonstrate the following.

- An increase in p, which shifts up the *VMPL* curve in Figure 6.6a and the price line in Figure 6.6 b, while leaving the wage line and marginal cost schedule unchanged, leads to an increase in both the quantity of cloth supplied and the quantity of low-skill labor demanded.
- An increase in w, which shifts up the wage line in Figure 6.6a and the marginal cost schedule in Figure 6.6b leads to reductions in both cloth supply and labor demand.
- Increases in K or A that increase the marginal productivity of labor, causing the *VMPL* schedule to shift up in Figure 6.6a and the marginal cost schedule to shift down in Figure 6.6b, would lead to increases in labor demand and output supply.

Multiple variable inputs, single output. When low-skill labor is the only variable input (as in the model just described), the only way to expand production is to employ more labor. Thus any changes in prices, wages, capital or technology that prompt the producer to expand production also increase her demand for low-skilled labor. When producers instead use skilled as well as unskilled labor, and may also vary the use of other inputs, the situation becomes more complicated. Producers may be able to increase production without increasing low-skilled labor. Producers for whom multiple inputs are variable must make choices about both the **scale of production** (i.e. how many units of cloth to produce) and the **factor proportions** to use in producing each unit (e.g. whether to use lots of capital and little low-skill labor, or little capital and lots of low-skill labor).

Changes in prices and fixed factors revisited. For producers with multiple variable inputs, any changes in price or economic conditions that increase the value of a yard of cloth relative to the marginal cost of producing it – such as increases in the price of cloth,

reductions in the price of inputs, capital accumulation or technical change – will increase the desired scale of production (just as we saw in Figure 6.6b for the simpler case). Many of these changes, however, also cause some variable inputs to become cheaper or more productive relative to others, motivating the producer to shift toward production methods that make more use of the inputs that have become more attractive. For example, when the wage for low-skill labor falls, the reduction in the cost of low-skill labor relative to the cost of other inputs may lead producers to employ more low-skill labor (and less of some other inputs) per unit of output produced. We can think of decomposing the overall effect of such changes on the demand for low-skilled labor (or any other input) into a "scale effect" and an "input substitution effect." The **scale effect** refers to the increase in the demand for low-skill labor that would occur *if* the producer chose to increase the scale of production while continuing to use the same factor proportions. The **input substitution effect** refers to the additional increase in the demand for labor that would arise out of the producer's efforts to increase the relative importance of low-skilled labor in producing each unit of output (while holding total quantity of output fixed).

The scale and input substitution effects on the demand for low-skill labor might work in opposite directions. For example, a reduction in the wage paid for skilled labor would tend to increase the desired scale of production, because it implies an overall reduction in marginal cost, but will also increase the employers' interest in using production methods that make greater use of skilled labor. If skilled and low-skill labor are **complements in production**, then the reduction in the skilled labor wage encourages greater use of both skilled and low-skill labor in producing a unit of output (while reducing the use of other factors). In such a case both scale and substitution effects lead the demand for low-skill labor to increase. If, however, skilled and low-skill labor are **substitutes in production**, then the reduction in the skilled labor wage would encourage a shift away from low-skill labor toward more skilled labor use. In such a case the scale effect works to increase low-skill labor demand, while the substitution effect works to reduce such demand, and the net effect of the skilled wage change on low-skill labor demand is theoretically ambiguous. Whether skilled and low-skill labor are complements or substitutes in production is a technological question, the answer to which is likely to vary across production sectors with different technologies.

Labor-using and labor-saving technical change. By definition, a technological improvement makes it possible to produce a unit of the same output at lower cost or produce a higher value output at the same cost. Holding output prices, input prices and fixed factors constant, technical change thus increases marginal revenue relative to marginal cost, encouraging an expansion of production, bringing a scale effect that tends to increase the demand for all inputs. Often technical change also induces input substitution effects, however, modifying the factor proportions used to produce a unit of output. The ultimate impact on the demand for low-skill labor (or any other input) thus depends on the exact nature of the technological improvement. For example, sometimes the adoption of a new technology in manufacturing requires the use of more sophisticated machinery that automates some tasks previously performed by low-skill labor. Once the new technology is installed, the marginal product of skilled labor capable of running the

machines rises (because the new machines would be of little use without competent operators), while the marginal product of low-skill labor falls (because the tasks such labor is now restricted to performing tasks that are less important to the overall productivity of the operation). The adoption of the new technology would bring both scale and input substitution effects tending to increase the demand for skilled labor. For low-skill labor, however, the scale effect would tend to increase demand, while the input substitution effect would tend to reduce it. If the input substitution effect is strong enough, the technical change reduces the demand for low-skill labor. We will call a technological change **low-skill labor using** if it would tend to increase the demand for low-skill labor (while holding prices and fixed factors constant), and will call it **low-skill labor saving** if it would tend to reduce the demand for low-skill labor.

Investment decisions and the long run. In the discussion thus far we have treated increases in fixed factors and changes in technology as if they were gifts received by producers who have no power to make them happen. In the real world, of course, producers help make these changes happen through investment in the purchase or construction of new machines, buildings and other assets, and in the adoption of new technologies. We will take a detailed look at the complex of circumstances that shape such investment decisions in Chapter 10. Here we simply point out that in the **long run**, in which producers have the ability to modify their technology and physical assets, changes in prices and wages may induce additional effects on labor demands over and above the kinds of changes we have already examined. Increases in output prices, for example, might initially motivate producers to put more labor to work using the same fixed factors and technology they are already using, but in the longer run also to invest in more fixed factors and technological improvements, which will bring a second round of impacts on labor demand. Similarly, increases in unskilled wages might reduce the scale of production and employment in the short run, but in the longer run also stimulate investment in capital equipment that can take the place of relatively more expensive labor and the adoption of labor-saving new technologies.

Agricultural production decisions. To understand agricultural production decisions, which impact the lives of large fractions of developing country populations, it is useful to add two more dimensions to our discussion of production decisions. First, we must consider not just total quantities of low-skill labor, skilled labor, fertilizer, pesticides, machinery, and irrigation water employed as inputs in any year, but the quantities employed in various seasons of the year. That is, we must take the **seasonality** of input demands and output production into account. A calendar year may be broken down into one, two or three agricultural crop cycles, and each crop cycle may be broken down into a series of activities, which are paced by the seasons. Within each crop cycle land preparation activities (tilling, burning), are followed by planting (seeding fields, nurturing seedlings and transplanting to fields), crop care (fertilizing, weeding, thinning, applying pesticides), and harvesting activities (cutting, picking, threshing).

Labor requirements often differ greatly by season, and a change in any one price, fixed factor or dimension of technology might have diverse implications for labor use across the seasons. For example, a reduced price for chemical herbicides may reduce the

demand for manual labor during the weeding phase of the agricultural cycle, but, by increasing the resulting yields of output per acre, increase the demand for harvest season labor. Construction of new irrigation equipment may increase yields and increase the demand for labor not only in the seasons of the year in which agricultural activities have traditionally taken place, but, by making it possible to complete a second or third crop cycle each year, greatly increase the demand for labor in months in which farm land was previously left idle.

Second, farmers often produce (or at least have the potential to produce) more than one crop. One choice farmers must make is how to allocate land across diverse activities (e.g. corn cultivation or bean cultivation). For any given plot of land on their farm, we can think of them as comparing the profits from devoting that land to diverse crops, and will choose to allocate that land to cultivation of the crop that yields the highest profit. If the plot cannot be cultivated profitably under any crop (given current prices and technology) it will be left **fallow** or idle. Changes in prices, fixed factors and technology may now influence production of multiple crops, and their potential impacts on labor demand become more complicated. An increase in the price of corn may lead farmers to increase production of corn, not only by working land already devoted to corn cultivation more intensively (i.e. devoting more labor, fertilizer and other inputs per acre), but also by expanding corn cultivation onto land previously left fallow or by inducing **crop substitution,** in which land previously devoted to bean cultivation is now allocated to corn. The increase in the corn price may thus cause a reduction in the supply of beans. While the scale effect in corn cultivation expands the demand for labor in corn cultivation, it is less clear what happens to the farmers' total demand for labor in all crops. If bean cultivation tends to require more labor per acre of land than corn cultivation (when facing the same wages and other input prices), then when land is shifted from bean cultivation to corn cultivation in response to the corn price increase, the demand for labor may fall!

Derived demands and linkages. Basic producer theory highlights that labor demand is a **derived demand:** producers' demand for labor is derived from consumers' demand for the goods and services the producers sell. Increases in the demand for goods and services, by increasing the prices at which producers may sell their output, tend to increase their demands for labor, as well as other inputs. This is the source of an important **linkage** among markets. When policies or circumstances increase the demand for agricultural goods, driving up prices in agricultural output markets, they also induce increases in the demand for labor, which may increase wages in labor markets. (Whether these effects are large or small will depend on many factors in which we will take great interest in subsequent chapters.)

Basic producer theory also raises the possibility of another less obvious but important linkage. The same policies that directly increase the demand for agricultural goods and labor may also indirectly stimulate demand for goods and services by **rural non-farm** businesses. Farmers facing higher output prices may demand more agricultural inputs produced by the rural non-farm sector, and farmers and workers with rising incomes are likely to demand more non-agricultural consumer goods and services. Indeed,

agricultural expansion tends to bring with it expansion of non-agricultural production as well (Reardon and Timmer, 2007). The nonagricultural expansion, in turn, may increase further the local demand for labor.

Application of Basic Producer Theory to Study of the Green Revolution's Impact on Rural Labor

One early criticism of the Green Revolution (GR) pointed to its impact on income distribution. Larger farmers were quicker to adopt GR methods than smaller farmers, and adoption of GR seed varieties was often closely followed by the adoption of mechanized production methods, which was seen as displacing agricultural labor and worsening the plight of landless rural workers (Lipton and Longhurst 1989, Hazell and Ramasamy 1991). Chapter 20 offers a comprehensive view of Green Revolution impacts. Here we simply discuss guidance offered by basic producer theory for study of GR impacts on labor demand.

Basic producer theory tells us that if we want to understand what is happening to the demand for labor in agriculture, we must first understand the array of technological options available to farmers and then seek to discern how and why their choices are changing over time. The main GR technological advance was the development through selective plant breeding of new varieties of rice, corn and wheat that were more efficient at turning fertilizer and water into edible grain, and that had shorter stems strong enough to prevent larger heads of grain from falling over and lodging in the soil. They made higher yields per acre possible, but only when their use was accompanied by greater fertilizer use and more carefully controlled watering. Initial GR varieties were also more sensitive to pests and disease, though later varieties were selected for greater resistance. In some regions the shorter growing cycles and reduced photosensitivity of the new varieties also made it possible to add a second or third cropping season to the agricultural year.

Most analysts agree that if prices and wages had held constant, adoption of the GR package of new seed varieties and greater fertilizer use would increase yields, and *increase* the demand for labor, especially in certain seasons of the year. While the seeds required little change in land preparation and planting activities, they required increased labor per acre during the period of crop care (for applying fertilizer and dealing with weeds and insects), and increased labor per acre during the harvest, because of the increased quantities of grain being produced per acre. In some places the adoption of GR technology also increased demand for labor in traditionally slack seasons by allowing the introduction of a second or third crop cycle.

In some places the increased profitability of rice, corn or wheat cultivation led farmers to start cultivating these crops on land that had previously been devoted to other crops. In a few cases the displaced crops used more labor per acre than the GR crops, implying that the expansion of GR seed usage might have reduced the demand for labor per acre. This may have been the case, for example, in parts of Turkey (Bayri and Furtan, 1989). But in many cases GR seed usage simply displaced cultivation of the same crops with traditional

varieties and would be expected to increase the annual demand for labor by perhaps 20 to 30 percent (Lipton and Longhurst, 1989). Where the construction of irrigation systems and adoption of GR technologies allowed expansion of second and third season crops, it also tended to reduce fluctuations in the demand for labor over the course of the year.

Despite this expectation, the change in agricultural labor demanded per acre over periods of GR expansion in many areas was much smaller than 20 percent or even negative (Otsuka, et al., 1994). This was technically possible only because (at least eventually) farmers adopting GR technology also made other changes in their production choices that were labor saving. For example, among rice producers in the Philippines, the period of GR expansion also brought adoption of labor-saving technologies in land preparation, planting and harvesting. Tractors replaced oxen in land preparation. Labor intensive practices of establishing seedlings in specialized seed beds and then transplanting them into fields by hand were replaced by "direct planting," in which rice was seeded directly into the field, and hand threshing was replaced by mechanical threshing.

There was nothing about the GR seed varieties that required these new labor saving practices, so what motivated their adoption? In some cases, independent technological developments made these innovations newly attractive. For example, transplanting had long been preferred over direct seeding because it allowed rice seedlings to be planted in fields where flooding prevented the development of weeds. The development of new chemical herbicides made the planting in flooded fields less important. In the Philippines the development of a smaller, more portable thresher made mechanical threshing more feasible on small farms. But in many cases the adoption of labor-saving technologies was fueled by rising wages (Otsuka et al. 1994, Lipton and Longhurst 1989).

The very fact that wages were rising must give us pause. If the *only* force for change during the period were the introduction of GR technologies, and if the GR technologies led to a net reduction in the demand for labor, then wages should have been falling rather than rising. Thus a simple story, in which the GR necessarily forced labor displacement and impoverishment of rural labor, cannot be complete.

Wage increases were brought about in part as a direct result of GR adoption, which increased the demand for labor per acre. In the Philippines, urban development and improved infrastructure connections between rural and urban areas also encouraged commuting and migration into urban employment by rural workers, placing upward pressure on rural wages (Otsuka, et al., 1994). In many places expansion of rural non-farm sectors increased rural competition for low-skilled labor, also tending to raise rural wages and encourage adoption of labor-saving technologies (Hazell and Ramaswamy, 1991). Regression analysis indicates that adoption of labor-saving practices and reductions in the demand for labor are better explained by rising wages and independent introduction of labor-saving technologies than by adoption of GR seed varieties (Otsuka, et al., 1994).

Interestingly, researchers believe that the GR was responsible for stimulating much of the growth in rural non-farm sectors that helped drive rural wages and employment up.

Adoption of GR technologies and consequent rising incomes increased the demand for both agricultural inputs and consumer goods and services produced by the rural non-farm sector, raising their prices and stimulating expansion of the sector. The derived demand for labor in these sectors thus increased. This suggests that the total impact of the GR on demand for rural labor, including agricultural and non-agricultural demand, was even greater than the increases implied by technological studies of impacts on per-acre labor requirements in agriculture. Many researchers now agree that the creation and adoption of GR seed varieties brought net benefits for landless rural labor in many places (see, e.g. Hazell and Ramaswamy, 1991).

Moving forward

Understanding how people make choices is absolutely fundamental to economic analysis. But it is only the starting point. If we are to understand the forces the determine peoples' *choices,* we must understand how the prices, wages, and social norms and opportunities that shape their choices are determined. For this we must study what happens when active decision-makers interact with each other in markets and non-market institutions. That is the task of Chapters 8 through 13.

Before moving from the "micro" level of how decisions are made to the "meso" level of markets and non-market institutions, however, we must recognize that real people seldom function as lone individuals making only narrow choices regarding consumption, time allocation *or* production. Rather, most people live in households, composed of diverse members, within which they must achieve joint decisions, and which often make simultaneous choices regarding consumption, labor supply *and* production. Acknowledging diversity across households in the sets of markets in which they participate as buyers and sellers, and potential conflicts of interest among household members, is of great importance for studying how the impacts on *wellbeing* of shared socio-economic changes are likely to be differentiated across households of different types and across members of different standing within households. Chapter 7 develops tools for such analysis, extending and modifying the basic theories of this chapter for use in studying developing country households.

References

Alderman, H. (1986), The effect of food price and income changes on the acquisition of food by low-income households, Washington, D.C.: International Food Policy Research Institute.

Alderman, H. and K. Lindert (1998), The potential and limitations of self-targeted food subsidies, *The World Bank Research Observer 13(2)*, The World Bank, pp. 213-29.

Barrett, C.B. (2002), Food security and food assistance programs, in B. Gardner and G. Rausser, ed., *Handbook of Agricultural Economics 2*, Amsterdam, New York: Elsevier Science B.V.

Bayri, T.Y. and W.H. Furtan (1989), The impact of new wheat technology on income distribution: a Green Revolution case study, Turkey, *Economic Development and Cultural Change 38(1)*, pp. 113-127.

Beegle, K., R. Dehejia and R. Gatti (2006), Child labor and agricultural shocks, *Journal of Development Economics 81,* pp. 80-96.

Block, S. (2004), Maternal nutrition knowledge and the demand for micronutrient-rich foods: Evidence from Indonesia, *The Journal of Development Studies 40(6),* pp. 82-105.

Bouis, H.E. and Haddad, L.J. (1992), Are estimates of calorie-income elasticities too high? A recalibration of the plausible range, *Journal of Development Economics 39,* pp. 333-364.

Calegar, G. and E. Schuh (1988), The Brazilian wheat policy: Its costs, benefits, and effects on food consumption, *International Food Policy Research Institute Research Report No .66*, Washington, D.C.: International Food Policy Research Institute.

De Janvry, A., F. Finan, E. Sadoulet and R. Vakis (2006), Can conditional cash transfers serve as safety nets in keeping children at school and from working when exposed to shocks? *Journal of Development Economics 79,* pp. 349-373.

Duryea, S., D. Lam and D. Levison (2007), Effects of economic shocks on children's employment and schooling in Brazil, *Journal of Development Economics 84,* pp. 188-214.

Edmonds, E. (2008), Child labor, in ed. Chenery, H. and T.N. Srinivasan, *Handbook of Development Economics, vol. 4,* Elsevier, Amsterdam.

Edmonds, E. and N. Pavcnik (2005), The effect of trade liberalization on child labor, *Journal of International Economics 65,* pp. 401-419.

Edmonds, E. and N. Pavcnik (2006), Trade liberalization and the allocation of labor between households and markets in a poor country," *Journal of International Economics 69*, pp. 272-295.

Hazell, P.B.R. and C. Ramasamy (1991), *The Green Revolution Reconsidered: The Impact of High-Yielding Rice Varieties in South India,* Baltimore: Johns Hopkins University Press.

International Labour Office (2004), *Child Labour: A Textbook for University Students,* International Labour Organization, Geneva.

International Labour Organization (1999), C 182 Worst Forms of Child Labour Convention, ILO, Geneva.

International Labour Organization (2002), *Every Child Counts: New Global Estimates on Child Labour,* ILO, Geneva.

Kruger, D. (2007), Coffee production effects on child labor and schooling and rural Brazil, *Journal of Development Economics 82,* pp. 448-463.

Lipton, M. with Longhurst, R. (1989), *New Seeds and Poor People,* Baltimore, Johns Hopkins University Press.

MCDSS/GTZ (2007), *The Pilot Social Cash Transfer Scheme – Zambia: Summary report, 5th edition,* Lusaka.

Menon, M., F. Pareli and F. Rosati (2005), Estimation of the contribution of child labour to the formation of rural incomes: An application to Nepal, Working paper No. 10, Centre for for Household Income, Labour, and Demographic Economics, Rome, Italy.

Otsuka, K. F. Gascon and S. Asano (1994), Green Revolution and labour demand in rice farming: the case of central Luzon, 1966-90, *The Journal of Development Studies 31(1),* pp. 82-109.

Pollak, R.A. and T.J. Wales (1992), *Demand System Specification and Estimation,* Oxford: Oxford University Press.

Psacharopoulos, G. (1997), Child labour versus educational attainment: Some evidence from Latin America, *Journal of Population Economics 10*, pp. 377-386.

Reardon, T. and Timmer, C. P. (2007), Transformation of markets for agricultural output in developing countries since 1950: How has thinking changed?, in R. Evenson and P. Pingali, eds., *Handbook of Agricultural Economics 3*, Amsterdam, New York: Elsevier Science B.V.

Ruttan, V.W. (2002), Productivity growth in world agriculture: sources and constraints, *The Journal of Economic Perspectives 16(4),* pp. 161-184.

Sadoulet, E., A. de Janvry and B. Davis (2001), Cash transfer programs with income multipliers: PROCAMPO in Mexico, *World Development 29(3)*, pp. 467-80.

Sen, A.K. (1981), *Poverty and Famines: an Essay on Entitlement and Deprivation,* New York: Oxford University Press.

Shafiq, M. (2006), "Household schooling and child labor decisions in rural Bangladesh," unpublished paper, Washington and Lee University.

Subramanian, S. and A. Deaton (1996), The demand for food and calories, *The Journal of Political Economy 104(1),* pp. 113-162.

World Bank (2006), *Repositioning Nutrition as Central to Development: A Strategy for Large-Scale Action,* Washington, D.C.: The World Bank.

Questions for Review

1. What is the basic premise of economic theories regarding the way people make choices?
2. Define the terms **objective**, **constraint** and **endowment** as they are used in economic theories of how people make choices.
3. What assumptions define basic consumer choice theory? What is the consumer's **utility,** and how does it relate to consumption of food and non-food in the simple application of the theory discussed in the text? What is the consumer's **budget constraint**?
4. Describe how to determine the position and slope of the **budget line** in a diagram like Figure 6.1.
5. How do indifference curves (in a graph like Figure 6.1a) relate to utility? What properties do indifference curves have? What is the significance of these properties?
6. Using a graph like 6.1a, demonstrate the effect on a consumer's choice of: (a) an increase Y, (b) a reduction in P_f, (c) a reduction in P_n, and (d) a change in culture or beliefs that causes the consumer to place greater value on providing adequate nutrition for her children.
7. What is a **normal good**? What is an **inferior good**? What is an **Engel curve** and how could you use a diagram like 6.1a to derive a consumer's Engel curve?
8. What is meant by the statement: "A reduction in the price of food implicitly transfers income to food consumers, with larger implicit transfers going to those who consume more of the relevant food"?
9. Define the **income effect** and **substitution** effect associated with a change in the price of a consumption good. What does theory tell us about the signs of the income and substitution effects?
10. For a consumer like that described in Figure 6.1a, what does it mean for a change in circumstances to have a real effect?
11. What do we hope to measure by constructing a **price index**? What is the distinction between **nominal** price or income and **real** price or income?
12. What does it mean for two goods to be **complements** in consumption? **Substitutes**? **Independent**?
13. What are **social norms**, and how might their evolution affect consumption choices?
14. Define the **income elasticity** of rice consumption and the **own price elasticity** of rice consumption? What do we know from theory about these elasticities? What must we learn about them from empirical research?

15. What guidance does basic consumer theory offer to research regarding the efficacy of income increases for improving nutrition among the poor? What are some key empirical findings in this area?
16. Explain why a rice distribution leads to the kinked budget constraint shown in Figure 6.4a.
17. Compare the impacts on a consumer's consumption choices and well-being of the receipt of a bag of rice versus the cash value (at local market price) of a bag of rice, being careful to distinguish whether the consumer would purchase any additional rice in the case of the rice distribution. Use diagrams like 6.4a and b to explain your answer.
18. Discuss the weaknesses revealed by basic consumer theory of using economy-wide food price reductions to improve nutrition.
19. Under what conditions might nutrition education improve consumers' true well-being?
20. What assumptions define basic labor supply theory? What is the labor supplier's **total time constraint**? What is **home time**? What is **non-labor income**? What is meant by a **corner solution** to the labor supply decision?
21. Using graphs like Figures 6.5a and b, demonstrate the effects on the choices of potential labor suppliers who do and do not initially supply any labor of (a) an increase in M, (b) an increase w, (c) an increase in T, and (d) a change in cultural norms that increases the relative value the consumer places on consumption expenditure relative to home time.
22. Contrast the implications of basic labor supply theory regarding the impact of an increase in the value of schooling on time allocation under two sets of circumstances: (a) the only two uses to which time is allocated are school and work, and (b) time may be allocated to school, work and at least one more activity.
23. Discuss some of the major lessons to be drawn from Table 6.1.
24. What questions does basic labor supply theory raise about the advisability of child labor prohibitions? What does empirical evidence have to say about these questions?
25. What assumptions define basic producer theory? Define the terms **profit**, **revenue**, **expenditure**, **production function**, **price taker**, and **diminishing marginal returns.** Define the distinction between the **short run** and the **long run.**
26. Using first a graph like Figure 6.6a and then a graph like Figure 6.6 b, discuss the impacts on the production decision of (a) an increase in the p, (b) an increase in the w, (c) an increase in the K or an increase in A that increases the marginal productivity of labor.
27. Define **scale of production, factor proportions, scale effects**, **input substitution effects**, and inputs that are **complements** and **substitutes** in production. Define **low-skill labor-using** and **low-skill labor saving** technical change.
28. Discuss the **seasonality** of agricultural production and its implications for production decision analysis.
29. Define the terms **fallow** and **crop substitution**. How does the possibility of crop substitution modify our predictions regarding the effect of an increase in the price of corn on the demand for low-skill labor?
30. What is the significance of the statement that that labor demand is a **derived demand**? What is a **linkage** between markets? What are some sources of linkage between agricultural markets and **rural nonfarm** markets?
31. Discuss the impacts of the Green Revolution on rural markets for low-skill labor.

Questions for Discussion

1. Offer a logical argument as to why two of the indifference curves describing a single consumer's preferences cannot cross.
2. Define a consumer's "true utility" to be the utility she would associate with consumption choices if she had complete understanding of the consequences of those choices for her health and nutrition. Define her "perceived utility" to be the utility she associates with consumption choices in her current state of imperfect knowledge regarding these consequences. Assume that because she lacks complete knowledge of health and nutrition consequences, she tends to under-value the consumption of vitamin-rich vegetables relative to the consumption of other items. Using a diagram similar to 6.1a, demonstrate that when she maximizes her perceived utility she fails to maximize her true utility.
3. Suppose a labor supplier allocates time to only two activities: work and home time, and has a total of 80 hours available per week to be allocated across these two activities. Show that under the

assumptions of basic labor supply theory the following two circumstances would lead a labor supplier to make the same time allocation choice (because they lead to identical budget constraints): (a) the labor supplier receives 20 pesos in non-labor income and is offered a wage of 5 pesos per hour of wage labor; and (b) the consumer is offered 420 pesos for a contracted work week of 80 hours, but may buy time off (to allocate to home time) at a price of 5 pesos per hour.

4. Using diagrams like Figure 6.5a and b, describe the kinds of changes to a woman's labor supply choices that might be induced by her receipt of additional education. How might her budget constraint change? Why? How might her preferences change? Why?
5. Consider a farmer who uses all of his available land to cultivate corn and beans, and assume that bean cultivation requires more labor per acre than corn production. For each of the following changes, discuss the likely sign of the effect on corn production, bean production, and labor demand: (a) a reduction in the price of beans, (b) an increase in the price of corn, (c) an increase in the wage, and (d) a reduction in the price of chemical herbicides. Explain your answers.
6. When studying any issue, good economists instinctively identify the relevant *decisions* and *decision-makers*, seek to understand the *constraints* and *preferences* that shape those decisions, work out the *logically possible implications*, and then undertake *empirical research* to determine which of the theoretical possibilities are important in practice. Suppose you are asked to study the likely effects on teacher and student performance of an effort to link teacher pay to how well their students perform on standardized tests. Which of the three basic models described in this chapter would serve as the best starting point for developing an appropriate analytical framework? Who is the most relevant decision-maker? What is the most salient set of choices? What constrains those choices? What does the framework suggest regarding the possible impacts of the teacher performance pay proposal?

Problems

1. Consider a consumer who chooses F and N to maximize $U(F,N)=F^{.75}N^{.25}$ subject to the budget constraint that $P_fF + P_nN = Y$. Making an appropriate substitution based on the budget constraint, this consumer may be construed as choosing F to maximize $F^{.75}[(Y-P_fF)/P_n]^{.25}$. (a) Write down the first order condition that must characterize the utility-maximizing choice F^*. (b) Re-arrange this condition to derive an equation for the utility-maximizing F^* as a function of Y, P_f, and P_n. (c) According to this equation, how does the consumer's choice of F* change as Y increases? As P_f increases? As P_n increases? (d) Re-arrange the condition one more time to derive an equation describing how the utility-maximizing food expenditure share, $(P_fF^*)/Y$, relates to the model's parameters. What do you learn from this expression?
2. Consider a consumer who chooses F and N to maximize utility U(F,N)=FN + F subject to the budget constraint that $P_fF + P_nN = Y$. (a) We can find all combinations of F and N that yield this consumer the same level of utility, U, by finding all combinations of F and N that satisfy FN + F = U. We can graph the associated indifference curve, in a quadrant with N on the vertical axis and F on the horizontal axis, by re-expressing this equality as N = (U-F)/F = (U/F) – 1. Graph indifference curves for U=1, U=2 and U=3. (b) Making an appropriate substitution based on the budget constraint, we find that this consumer chooses F to maximize $F[(Y-P_fF)/P_n]+F$. Write down the first order condition for this maximization problem and re-arrange to find an expression for the utility-maximizing food consumption share, $(P_fF^*)/Y$, as a function of Y, P_f, and P_n, and graph this consumer's Engel curve for food, assuming that Y is always greater than or equal to 1 and P_n is always less than or equal to 1.
3. Draw three identical graphs with axes labeled as in Figure 6.1a. Draw into the three graphs identical budget constraints, with a point *a* located in the same place on the budget constraint (like point *a* in Figure 6.1). In all three graphs draw in the new budget constraint that would become relevant after an increase in income (of the same amount in all diagrams). Now draw (different) indifference curves into the three diagrams to indicate the preferences of three different consumers, in which their preferences lead to the following outcomes: All three consumers choose to consume at point *a* when facing the initial budget constraint. The first consumer spends the entire increase in income on food. The second consumer increases food and non-food consumption in the same proportion. The third consumer reduces wheat consumption after the income increase. Draw at least two indifference curves in each diagram, making sure that they satisfy all the properties required by the basic theory.
4. Draw a diagram like Figure 6.3a. Suppose this diagram illustrates the impact on a consumer of a general food subsidy's reduction in the price of food (to the level associated with the dashed budget

line). (a) If the consumer still faced the initial prices, how much cash would she have to be given to be able to consume F_2 units of food and N_2 units of non-food? Please answer using notation defined in the graph. (b) Show that your answer to part (a) is equal to $F_2(P_{f2} - P_{f1})$. (c) Now suppose the consumer had been given $F_2(P_{f2} - P_{f1})$ in cash while continuing to face the initial prices. Draw the budget constraint she would face into your diagram, and explain why you draw it as you do. (d) Making reference to your graph and to any relevant assumptions of the model, demonstrate that if given the cash transfer described in part c, the consumer would consume less than F_2 units of food but would achieve a higher level of utility than that associated with consumption of F_2 and N_2.

5. Draw a diagram like Figure 6.5a, including the solid budget constraint and the indifference curves. Indicate what would happen to the budget constraint and the labor supplier's choice if the wage increased. Think of the labor supplier as a consumer of C and H with income of M+wT and the increase in the wage as a reduction in the price of home time. Explain how you would divide the overall change in choices of C and H (as a result of the increase in the wage) into "income" and "substitution" effects. What must be the sign of the substitution effect? What does theory tell us about the sign of the income effect?
6. Draw and label completely a diagram describing a consumer's choice regarding consumption of rice and other items when she has Y rupees of income and faces prices P_r and P_o in rupees per unit of rice and other items. Construct the new budget constraint she would face if she were offered the opportunity to purchase a limited quantity R of rice at a subsidized price P_s, which is lower than P_r.
7. Consider a manufacturer who produces cloth using low-skill labor, high-skill labor and several other inputs. According to basic producer theory, what do we know (or not know) about the signs of the scale effect and input substitution effect on the demands for low-skill and high-skill labor associated with each of the following changes: (a) a reduction in the wage for low-skilled labor, (b) a reduction in the cost of an input that is a complement to low-skill labor, (c) a reduction in the cost of an input that is a substitute for low-skill labor, (d) low-skill labor using technical change, and (e) low-skill labor saving technical change.
8. Suppose every farmer in a region cultivates one acre of land. The quantity of output they produce, Y, as a function of the labor they devote to cultivation, L, is given by the production function $Y = e^{\alpha Q}L^{\beta}$. The price of output is p, the wage is w and farm profits are given by pY-wL. Take Q as given and treat L as a variable input. Derive an expression for the profit-maximizing value of labor input, L*, as a function of Q. According to this expression, how does the value of labor input per acre vary across farmers with different land quality? Explain.