

Chapter Two

Labour Supply: Individual Attachment to the Labour Market

Main Questions

- *How do we measure labour market attachment? How did labour market attachment evolve over the twentieth century?*
- *What is the labour force participation decision, and how does it fit into the general theory of labour supply?*
- *How can we incorporate the possibly different factors that determine men's and women's labour supply decisions?*
- *Is there any theoretical reason to believe that labour supply increases with the market wage?*
- *What is the evidence regarding the responsiveness of individual labour supply to changes in the wage?*
- *Why are workers paid an overtime premium, rather than just a higher wage, in order to induce them to work overtime?*

At its core, the subject matter of labour supply is how individuals earn a living by selling labour services through the labour market. While there are important quality dimensions to labour supply, such as the levels of skill that someone brings to the market, the focus of the next three chapters is on the quantity of homogeneous labour offered. The guiding theoretical question is whether labour supply is an upward-sloping function of the wage rate.

Why might market labour supply curves increase with the wage rate? In introductory courses, supply curves are commonly drawn as upward-sloping functions of the market price. Why not borrow from this, and merely assert that labour supply, as a special case of "supply," is also upward-sloping in its price? Of course, such reasoning is circular and

provides no insight as to why more labour would be offered at a higher wage. The classical economists provide one possible explanation. They argued that labour supply was upward-sloping at least in the long run, because higher wages would spur population growth, increasing the number of labourers. Modern labour economists do not pursue that line of reasoning. Instead, we take the population as fixed at a point in time, and focus on individual decisions of how much to work. These decisions can be broken down into participation (whether to work) and hours (how much to work). Why might both of these dimensions of individual labour supply be increasing in the wage rate? Intuition suggests that people will want to work more if the return for doing so increases. As we shall see, however, intuition alone can be misleading. Nevertheless, incentives to work lie at the heart of the study of labour supply.

In this chapter we develop a model of individual labour supply. The basic theoretical framework is the income- or labour-leisure choice model, and it can be applied to both the participation and the hours dimension. Within this framework, we analyze the effect of changes in wage rates and other economic variables on preferred hours of work, and thus derive the individual's supply curve of labour. An important strength of this model is that it can easily be adapted to explore the impact of more general types of incentives, such as those provided by government tax and transfer programs. Accordingly, we exploit the labour-leisure model in the following chapters. Before turning to the theoretical model, however, we define some important concepts and investigate several labour supply features of the Canadian labour market.

A GUIDE TO THIS CHAPTER

This chapter is a “spinach” one. Much of the material in labour economics, even topics not obviously related to labour supply, such as individual schooling decisions, require an understanding of this theory. This chapter develops the theory of a worker as a rational decision-maker, making choices between various opportunities (good or bad) presented by the labour market. The chapter is divided into two sections:

Section One: The Theory of Labour Supply

The objective of Section One is to begin with a model of individual behaviour, and trace through the implications for the market supply of labour. Ultimately, we derive a labour supply function relating labour supply to the market wage. We begin with an explanation of how economists quantify labour market attachment, dividing this attachment into participation (whether to work) and hours (how much to work) components. After exploring participation and hours patterns in Canada, we develop the basic income-leisure model that underlies both dimensions of labour supply. Using this model, we investigate how participation and hours respond to changes in the economic environment (like the wage). Finally, we summarize the empirical evidence on the effect of wages on labour supply.

Section Two: Extensions and Applications

The objective of Section Two is to use the labour supply model to explore a number of features of the labour market, exclusive of government tax and transfer programs, which we treat separately in the next chapter. In particular, we look at the impact of labour market constraints (like unemployment) on individual labour supply, and the means by which the model can account for multiple jobholding (moonlighting) and the structure of overtime premiums.

Finally, the appendix to this chapter provides a review of the general consumer theory that underlies the income-leisure at the heart of labour supply.

Section One: The Theory of Labour Supply

QUANTIFYING LABOUR MARKET ATTACHMENT

Many students are initially surprised to see the relative attention paid to labour supply by economists. Perhaps based on their own experience, their caricature of a labour supply decision is that virtually everyone wishes a full-time job or career. However, is it really the case that most people participate in the labour market, working the standard work week of around 40 hours? As we shall see, there is actually considerable variation in the degree of attachment of individuals to the labour market.

Labour Force Participation

The **labour force participation decision** is basically a decision to participate in paid labour market activities as opposed to other activities such as unpaid work in the home, volunteer work, education, or retirement. As such, it influences the size and composition of our labour force and it has an impact on household activities, education, and retirement programs.

The policy implications of these changes can be dramatic. Changes in the size and composition of our labour force affect our growth and unemployment rates, as well as the occupational and sex composition of the labour force. The latter, in turn, affect such factors as relative wages, demands for unionization, daycare, and equal pay and equal employment opportunity legislation. Changes in household activities can involve family formation and mobility. Retirement programs can be affected insofar as new labour force participants will add contributions to pension funds, while those who retire (i.e., do not participate in the measured labour force) will be a drain on the funds.

As illustrated in Figure 2.1, the **labour force** consists of those persons in the eligible population who participate in labour market activities, as either employed or unemployed. The eligible population is that portion of the population that is surveyed as potential labour force participants (i.e., civilian noninstitutional population, 15 years and over, excluding the Yukon, Northwest Territories, and those living on Indian reservations). Persons from that potential population of labour force participants (POP) are categorized as either in the labour force (LF) or not in the labour force (NLF). Those in the labour force are either **employed** (E), or **unemployed** (U), with the latter being not employed but seeking work. People are categorized as employed if they are normally employed but they happen not to be at work at the time of the survey because, for example, they are ill or on strike. Individuals on temporary layoff, and those not employed but who have a job to start in the next month, are classified as unemployed. Those not in the labour force are usually students, retired people, persons in the household, those unable to work, or some “discouraged” workers who have simply given up looking for work. According to labour force definitions, the latter are not categorized as unemployed because they are not seeking work.

The **labour force participation rate** (LFPR) is the fraction of the eligible population that participates in the labour force ($LFPR = LF/POP$). The unemployment rate (UR) is the proportion of the labour force that is unemployed ($UR = U/LF$).

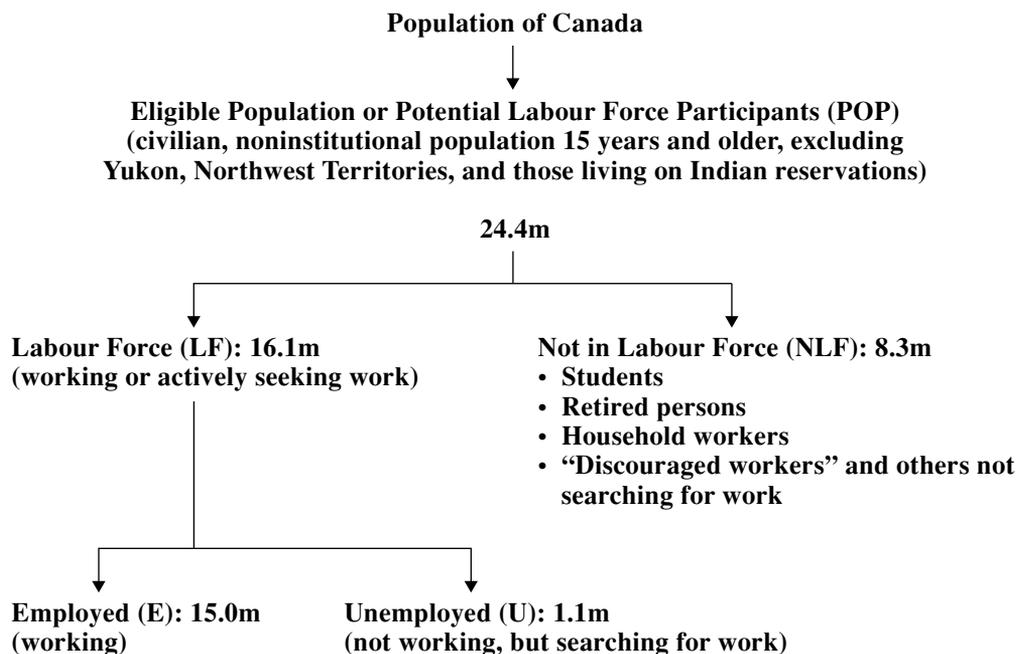
The Canadian Labour Force Survey (LFS), conducted by Statistics Canada, is currently based on a monthly sample of approximately 62,000 households. The results are published monthly in *The Labour Force* (Statistics Canada, Catalogue No. 71-001), with the survey described in more detail in that publication. The survey began in 1945, first on a quarterly, and by 1952, on a monthly basis. In 1976, and again in 1997, the survey was expanded substantially to collect more detailed information on labour market attachment.

While the LFS provides the most frequently used estimates for our labour force,



Figure 2.1**Labour Force Concepts (numbers refer to November 2000)**

The labour force is the sum of those individuals either working (employed) or not working, but searching for work (unemployed). The unemployment rate is the percentage of the labour force that is unemployed, while the labour force participation rate is the percentage of the eligible population in the labour force.

*Notes:*

"m": million

Labour force participation rate: $LFPR = LF/POP = 16.1/24.4 = 66.0\%$.

Unemployment rate: $UR = U/LF = 1.1/16.1 = 6.8\%$.

Source: Adapted from the Statistics Canada publication "Labour Force", Catalogue 71-001, November 2000.

employment, and unemployment figures, other sources are available. In particular, the Canadian Census is now conducted every five years, the most recent being in 1996, referring to activity in 1995. The Census is more comprehensive (not being based on a sample from a larger population), and consequently includes richer details on such factors as employment and unemployment by industry and occupation. However, its use is limited because it is conducted only every five years, there is a considerable lag before the results are published, and its reliability on labour force questions may be questioned because, unlike the Labour Force Survey, it does not focus only on labour force activity.

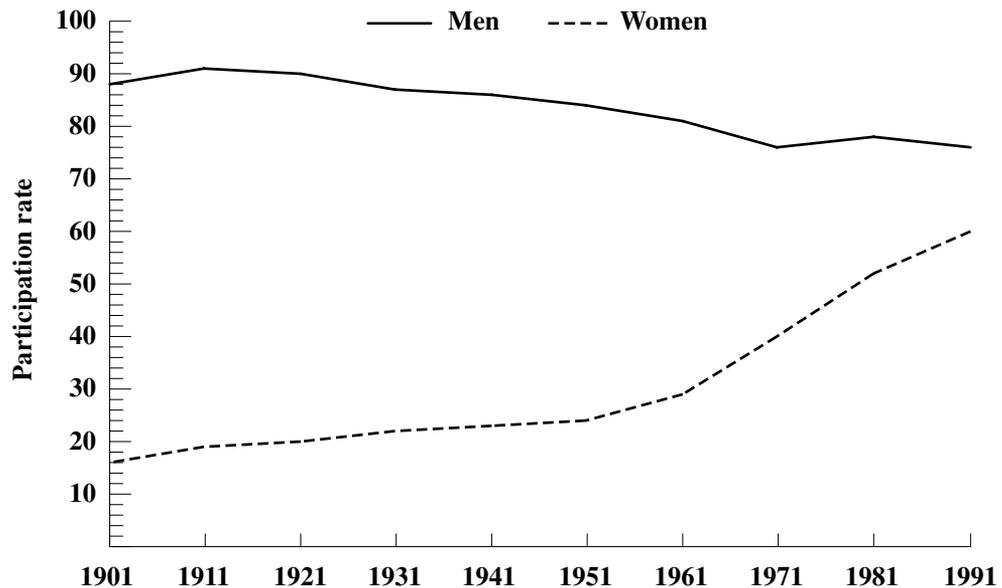
Labour force participation rates have changed significantly over the twentieth century, especially for women. Figure 2.2 shows trends in participation based on census data. As we can see, male participation has declined from just over 90 percent of the eligible population to 76 percent by 1991. Women's rates were below 30 percent as recently as 1961, but they have increased by 10 percentage points per decade to 60 percent in 1991. Results from the 1996 census show that the male participation rate has continued to decline, from 76.4 percent in 1991 to 72.7 percent in 1996. The female rate, on the other hand, has plateaued at 58.6 percent, as against 59.9 percent in 1991. What factors might account for these changes? Was it primarily improvements in contraception or changing attitudes towards women's work that led to increased female participation? Did women's wages rise since 1960, inducing more women to enter the labour force? Did their husbands' wages fall, requiring more women to work to support their families? A well-developed economic model would aid in analyzing these explanations, and also provide alternative testable hypotheses.



www.statcan.ca/english/census96/list.htm

Figure 2.2**Labour Force Participation Rates by Sex, Canada, Census Years 1901–1991**

Labour force participation rates for men and women are plotted by census years, beginning in 1901. The general trends suggest a gradual reduction of labour force attachment for men, in contrast to a sharp increase for women, beginning in the 1950s.



Sources: M. Gunderson, "Work patterns," in *Opportunity for Choice: A Goal for Women in Canada*, ed. G. Cook (Ottawa: Statistics Canada, 1976), p. 97 for 1901–1971, reproduced by permission of the Minister of Supply and Services Canada. Figures for 1981 were computed from the 1981 Census, Catalogue 92-915, Vol. 1, Table 1, p. 1-1. Figures for 1991 were computed from the 1991 Census, Catalogue 93-324, *The Nation: Labour Force Activity*, Table 1.

Figure 2.2 hides as much variation in participation as it shows, as participation also varies significantly by age and marital status. We focus on these demographic and life-cycle elements of labour supply in Chapter 4. However, Table 2.1 highlights another important aspect of cross-sectional variation in participation: variation across countries. The Canadian rates are at the higher end of the range of countries reported, just below that of the United States for both men and women. Generally speaking, the European countries outside Scandinavia have lower participation rates for both sexes (especially men) than North America and Australia, while the less developed countries have higher participation rates for men. Do these participation rates lie along an upward-sloping labour supply curve, with participation rates higher in those countries with higher wages? In fact, this table suggests that the effect of wages may not be uniform: the poorest countries have the highest male participation rates. This table also raises another question: Why do women's participation rates vary so much? Bangladesh and Pakistan have similar levels of development and economic conditions, but quite different patterns of female employment. Is this due to cultural differences, or are other economic variables more important? Clearly, there are several factors at play, and no single theory is likely to explain the intricacies of these numbers. What this table shows, however, is that the relationship between "wages" and participation may not be a simple one.

Hours

The **hours-of-work aspect** of labour supply has a variety of dimensions including hours per day, days per week, and weeks per year. Changes in any or all of these dimensions can alter the hours-of-work aspect of the labour supply decision. Phenomena such as the eight-hour day, the shorter work week, and increased vacation time are institutional embodiments of a reduction in hours of work. Similarly, moonlighting, overtime, flexible

Table 2.1 Labour Force Participation Rates by Sex, Various Countries, 1998

Country	Male	Female	Both Sexes
Denmark	83.1	73.2	78.2
Zimbabwe ^a	79.4	67.5	73.1
Bangladesh ^b	88.8	55.9	72.6
Barbados ^a	73.6	62.1	67.5
USA ^c	74.9	59.8	67.1
Canada ^c	72.4	58.1	65.1
Australia ^c	72.9	53.9	63.3
Hong Kong ^c	75.5	48.5	62.0
Korea ^c	75.2	47.0	60.7
Russia ^d	68.1	50.4	58.5
Argentina ^e	76.2	41.3	58.2
Germany	67.6	47.8	57.3
France	62.2	47.8	54.7
Pakistan ^f	83.1	15.0	50.1
Italy ^{a, c}	61.3	34.8	47.5

Notes: Participation rates for the population 15 years and older. Ranked in descending order, from highest to lowest labour force participation rate for both sexes. Some figures are preliminary estimates.

a. 1997.

b. 1995–96; excludes armed forces.

c. Excludes armed forces.

d. 1996.

e. 1995.

f. 1996–97; excludes certain regions.

Source: Based on *ILO Yearbook of Labour Statistics*, 1998, 1999.

working time, and compressed work weeks are institutional arrangements that alter the typical pattern of hours of work.

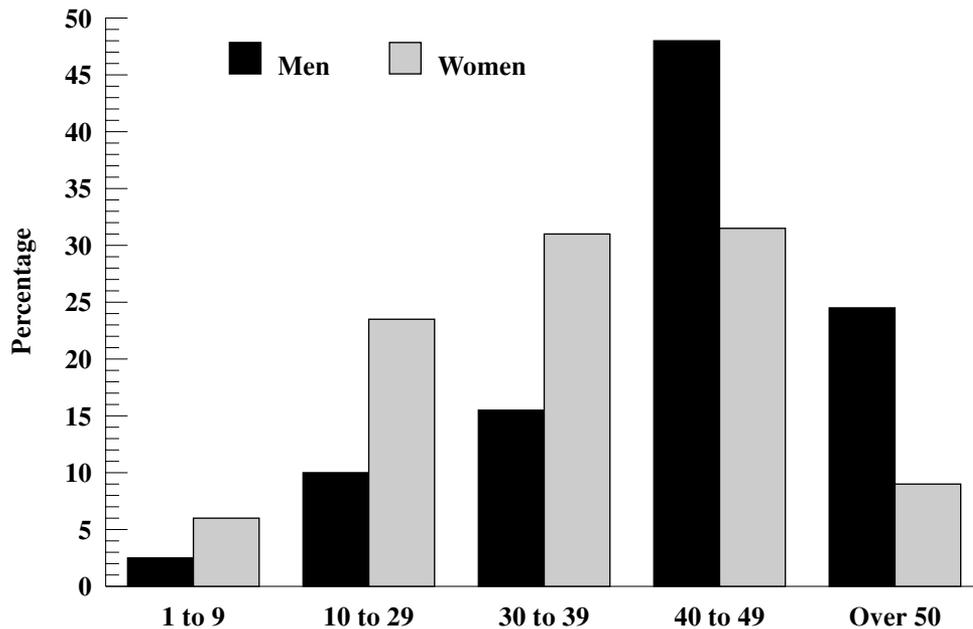
The policy importance of the hours-of-work decision is illustrated in a variety of ways. Changes in hours of work can affect not only the quantity but also the quality of our overall labour supply (and hence national output), as well as absenteeism, turnover, employment opportunities, and the demand for related activities, notably those involving leisure time and flexible working hours. Changes in hours of work, in turn, can be affected by changes in the age and sex structure of the labour force, the prominence of two-earner families, as well as government policies and laws and, of course, the wage rate.

In the short run, hours of work appear to be relatively fixed with little scope for individual variation. The eight-hour day, five-day work week, and fixed vacation period are fairly common for many wage and salary earners. However, the increased importance of flexible working hours is altering these arrangements. In addition, occupational choice provides an element of flexibility as people choose jobs partly on the basis of the hours of work required. Individuals may also be able to combine jobs, perhaps two part-time jobs or a full- and a part-time job, in order to work the desired number of hours for a given wage rate.

Figure 2.3 illustrates the apparent degree of flexibility that Canadian workers have in the number of hours they work per week. The work patterns of women are quite varied. Men are more likely to work the typical 40-hour week. However, even for working men, less than half work 40 to 49 hours per week. Almost a quarter work more than 50 hours, and an equal proportion work fewer than 40 hours. Women, on the other hand, are more likely to work part-time than 40 hours. Might these patterns be explained by labour supply theory? Perhaps relatively higher wages for men lead to higher average hours worked? Alternatively, women may “prefer” to work part-time. As with participation, a good theory of labour supply should be able to accommodate these features of labour market attachment.

Figure 2.3**Distribution of Hours Worked per Week by Sex, May 1996**

This histogram shows the distribution of usual hours worked per week for men and women. Less than 50 percent of men or women work 40 to 49 hours per week. Women are more likely to work part-time (under 30 hours per week), while men are more likely to work overtime (more than 50 hours per week).



Source: Reproduced by authority of the Minister of Industry, 1996, Statistics Canada, *The Labour Force*, May 1996, Catalogue No. 71-001, Table 20.

Table 2.2 documents further evidence that “typical” hours worked per week are not immutable. This table traces the standard work week in Canadian manufacturing and shows a pronounced and continuous decline over time in hours of work. Between 1901

Table 2.2 Standard Weekly Hours in Manufacturing, Canada, 1901–1981

Year	Standard Weekly Hours ^a	Hours Net of Vacations & Holidays ^b
1901	58.6	N/A
1911	56.5	N/A
1921	50.3	N/A
1931	49.6	N/A
1941	49.0	N/A
1951	42.6	40.7
1961	40.4	38.1
1971	39.8	36.7
1981	39.2	34.8

Notes:

- a. Standard hours are usually determined by collective agreements or company policies, and they are the hours beyond which overtime rates are paid. The data apply to nonoffice workers.
 b. Standard hours minus the average hours per week spent on holidays and vacations.

Sources: Figures for 1901–1971 for standard weekly hours are from S. Ostry and M. Zaidi, *Labour Economics in Canada*, 3rd ed. (Toronto: Macmillan, 1979), pp. 80-81. Figures for 1951–1971 for hours net of vacations and holidays are from Labour Canada (1974, p. 6). Both sources used as their primary data the Survey of Working Conditions conducted annually by the Canada Department of Labour and published as *Wage Rates, Salaries, and Hours of Labour*. This survey is also the source for the 1981 figures. Unfortunately, the Survey was discontinued in the early 1980s. Reproduced with the permission of Public Works and Government Services Canada, 2001.

and 1981 the standard work week declined from almost 60 hours to less than 40 hours. The decline slowed down in the depression years of the 1930s, and the war years of the 1940s, and it appears to be slower in the postwar period. However, as the last column illustrates, when vacations and holidays are considered the decline in average working hours is more noticeable. In essence, in recent years the work force has reduced its working hours more in the form of increased vacations and holidays rather than a reduction in hours worked per week. The decline in net weekly hours in the postwar period and in standard hours prior to World War II give a long-run trend reduction of about two hours per decade.

Since real wages have risen over the century, this long-run decline in hours worked appears inconsistent with an upward-sloping labour supply function. Instead, it suggests an independent effect of increased wages: as societies become wealthier, they need not toil as hard, and can afford to take more time off. Hopefully, the theoretical model will clarify the ways in which wages can affect labour supply.

BASIC INCOME-LEISURE MODEL

The objective of the **labour supply model** is to represent an individual's choice of hours worked given her market opportunities and the value she places on her nonmarket time. We wish to model this person as doing the best she can, subject to the constraints of the labour market and the limited availability of time. The model is an extension of standard microeconomic consumer theory. As such, we divide her decision-making problem into two parts: what she would like to do (preferences); and the choices available to her (constraints). The assumption of rationality brings the two parts together, yielding a unique characterization of her choice as a function of her preferences and market constraints.

Preferences

We assume that individual preferences can be distilled into **preferences** over two “goods”: consumption of goods and services, and “leisure.” **Leisure** embodies all non-labour market activities, including household work and education, as well as pure leisure activities. The phrase “leisure” is somewhat of a misnomer since it includes activities that are not leisure activities. The term is used here, however, because it is the phraseology used by Robbins (1930) in his original work on the subject and it has been retained in the literature as a summary term for “non-labour market activities.”

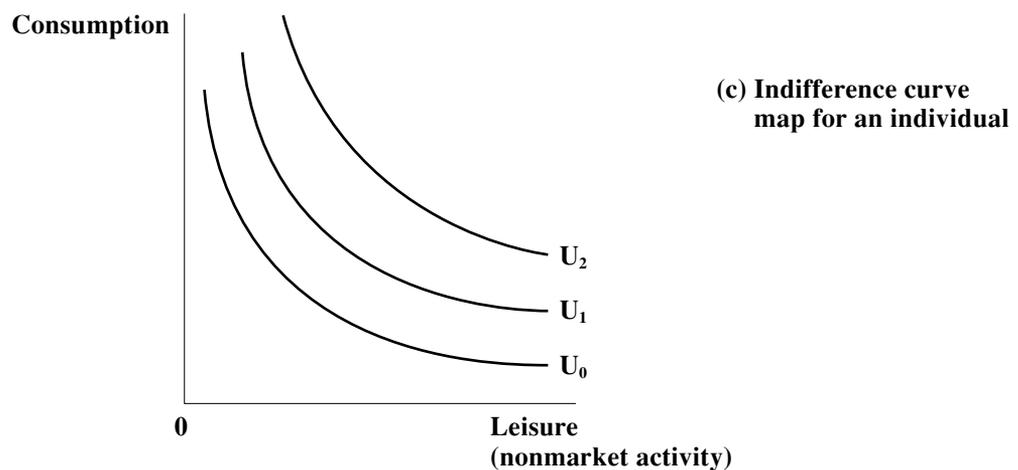
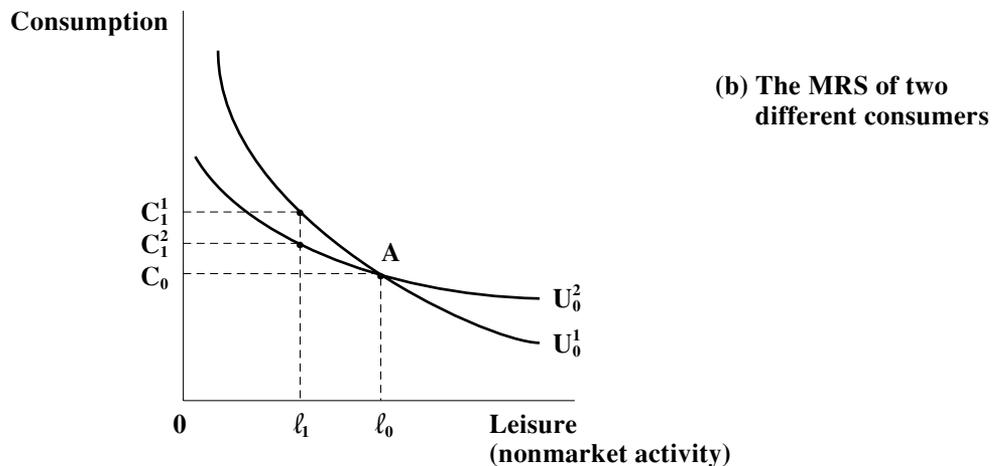
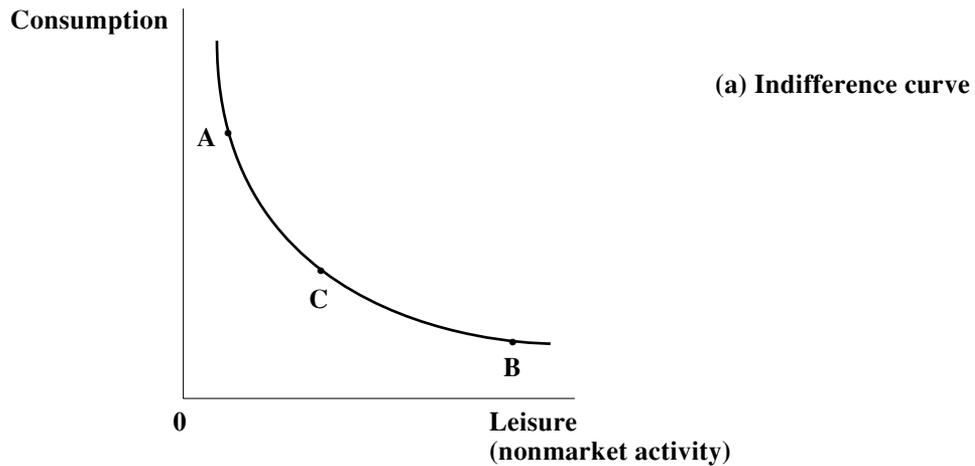
We can graphically represent consumer preferences with **indifference curves**. Indifference curves and budget constraints are used extensively in developing the theory of labour supply. Those unfamiliar with indifference curve analysis should review the appendix to this chapter before proceeding further.

The individual is indifferent (has the same utility, or welfare) between various combinations of consumption and leisure as given by the indifference curve U_0 in Figure 2.4(a). The slope of the indifference curve exhibits a diminishing **marginal rate of substitution** (MRS) between consumption and leisure. For example, at point A, the individual has an abundance of consumption and hence is willing to give up considerable consumption to obtain more leisure; hence the steeply sloped indifference curve at point A. At point B, the individual has an abundance of leisure and hence is willing to give up considerable leisure (i.e., work more) to obtain more labour income and thus consumption; hence the relatively flat indifference curve at point B. At intermediate points such as C, consumption and leisure are more substitutable because the individual is not saturated with either.

While we cannot compare **utility** or welfare levels across individuals, we can compare marginal rates of substitution across consumers. This allows us to represent different preferences, at least partially, by indifference curves with different MRS at various combinations of consumption and leisure. For example, in Figure 2.4(b) we depict two different consumers. For consumer 1, bundle A lies on indifference curve U_0^1 , while for person 2,

Figure 2.4 Consumer Preferences: Consumption–Leisure Indifference Curves

Indifference curves plot combinations of consumption and leisure that yield the consumer equal levels of utility. The absolute value of the slope of an indifference curve gives the marginal rate of substitution (MRS), that is, the amount of leisure the consumer is willing to accept in exchange for giving up some consumption. The indifference curve in panel (a) exhibits diminishing MRS, as the consumer requires more leisure to offset the decline in consumption as his or her level of consumption falls (from A to C, versus C to B). In panel (b), we compare the MRS's of two consumers. (Note: Indifference curves cannot cross for a single individual.) Beginning at A, person 2 has a lower valuation of leisure, requiring a smaller increase in consumption (C_1^2 versus C_1^1) to offset the decline in leisure from l_0 to l_1 . Finally, panel (c) shows representative indifference curves for an individual. Utility is increasing as the person has more leisure and consumption (i.e., as combinations move further from the origin).



the same bundle lies on indifference curve U_0^2 . Person 1 has a higher MRS at A than person 2. This implies that person 1 places a higher value on leisure (at A) than person 2. To induce person 1 to give up leisure from ℓ_0 to ℓ_1 requires compensating him with an increase in consumption from C_0 to C_1^1 to keep him equally happy, whereas it would only require an increase to C_1^2 to compensate person 2. In this way we can incorporate those factors that affect the value of **nonmarket time**, such as the value of work at home, into individual preferences and manifested in the MRS.

We assume that consumers have well-defined preferences over all the conceivable combinations of consumption and leisure. This implies that all combinations lie on some indifference curve. Consumer preferences can then be represented by an indifference curve map, as illustrated in Figure 2.4(c). Higher indifference curves such as U_1 and U_2 represent higher levels of utility since they involve more of both consumption and leisure.

Constraints

The individual will try to reach the highest indifference curve possible, constrained by economic opportunities provided by the labour market. Let the price of consumption be denoted by P , so that the value of consumption is PC . In this simple model we ignore saving, so that we can set PC equal to income. This allows a simple transformation of our model from **consumption-leisure** to **income-leisure**. We now want to summarize the income (or consumption) and leisure combinations from which the consumer can choose. In Figure 2.5 we show a few examples of **potential income constraints**. These are “potential” income constraints because they indicate varying amounts of income that can be obtained by giving up leisure and working. The actual amount of income earned will depend on the chosen amount of work in addition to the amount of income received from other sources, referred to as “nonlabour” income and shown as the amount Y_N . Alternative phrases for the potential income constraint include **budget**, **income**, and **full-income constraint**.

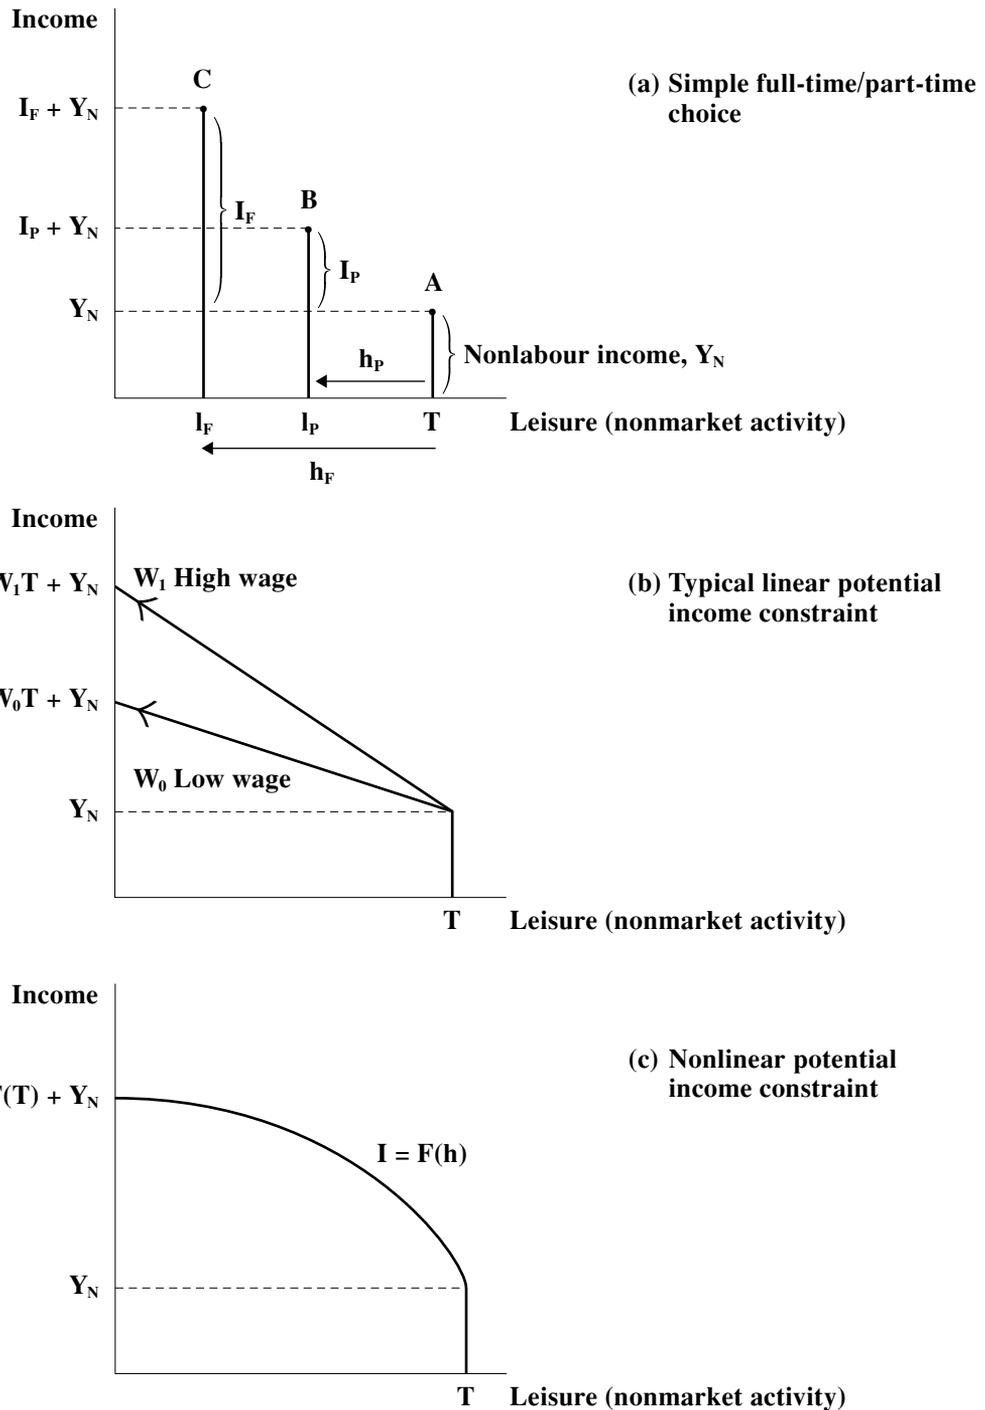
Consider first a very simple situation where the individual has only three discrete choices, as depicted in Figure 2.5(a). She can do no paid work, spending all of her available time, T , engaged in non-marketed activities, including leisure. In this case, the most that she can consume is given by her nonlabour income, Y_N (which might be 0). This combination is denoted by A. She could work part-time, increasing her income by I_p to $I_p + Y_N$, reducing her leisure by h_p to $\ell_p = T - h_p$ (point B). Alternatively, she could work full-time, earning $I_F + Y_N$, and consuming $\ell_F = T - h_F$ units of leisure (point C). In this case she need only compare the utility of three bundles, choosing the one on the highest indifference curve.

The more conventional potential income constraint is depicted in Figure 2.5(b). This linear budget constraint allows the individual to choose from a continuum of income-leisure bundles, ranging from 0 hours of work to T hours of work. The slope of the constraint depends on the person’s wage rate, W . For each hour worked the individual gives up one hour of leisure, but earns W of additional income. The constant slope reflects an assumed constant wage rate for each hour worked. Persons with a higher market wage, such as W_1 where $W_1 > W_0$, will be able to earn more income by giving up leisure and working more; hence, the slope of W_1 is steeper than the slope of W_0 .

Perhaps the simplest way to understand how wage changes (as well as other factors such as taxes and transfers, which will be examined later) affect the potential income constraint is to first mark the endpoint on the leisure axis. This endpoint is the maximum amount of leisure available (T); depending on the units in which leisure is measured, it could be 24 hours per day, 7 days per week, 52 weeks per year, or some combination such as 8736 hours per year. The potential income constraint is then derived by plotting the income people receive as they give up leisure and work additional hours; that is, they move leftward from the maximum leisure endpoint. The endpoint on the income axis would be the maximum income the individual could attain by working all of the available time—that

Figure 2.5 Potential Income Constraints: Summarizing Individual Market Opportunities

Budget sets, or potential income constraints, show the combinations of income (or consumption) and leisure available to an individual. In panel (a), the individual has only three choices: not working (point A), and obtaining T hours of leisure, no labour earnings, and non-labour income of Y_N ; working part-time for h_p , obtaining l_p hours of leisure, and earnings I_p (plus non-labour income Y_N); or working full-time, obtaining l_f leisure and I_f earnings (plus Y_N). Panel (b) shows the more conventional opportunity set, where the individual is free to choose any number of hours to work at the wage rate, W . Higher wages yield a steeper budget constraint, and higher potential income. Panel (c) illustrates the possibility that the “wage rate” is not constant per hour worked, as would be the case for a self-employed person.



is, by having zero leisure. If individuals choose to work T , the most they can earn is $WT + Y_N$, commonly called **full income**.

The potential income constraint can be generalized even further, allowing the returns to work to vary with the amount worked. One example would be the case of a self-

employed individual (like a doctor or lawyer), who sells output that she produces to the market. Assume that she produces income as a function of hours worked, $I = F(h)$. If she has diminishing marginal productivity for each hour worked beyond a certain point, then her potential income can be depicted by Figure 2.5(c). Here, her income is also a function of the number of hours she works, but the wage rate is the value the market places on an hour's worth of the goods she produces. The treatment of consumer choice in this case would be no different than that in panel (b), though it is analytically more complicated.

The Consumer's Optimum

By putting the individual's potential income constraint and indifference curves together, we can obtain the **consumer's optimum** amount of income and leisure (and hence we can obtain the optimal amount of work, or labour supply). The **utility-maximizing** individual will reach the highest indifference curve possible, constrained by the labour market opportunities as given by the potential income constraint. Figure 2.6 illustrates two types of outcomes to this choice problem.

Panel (a) shows an individual who will not participate in the labour market, given the individual's preferences, nonlabour income Y_N , and market wage rate W_0 . The highest possible utility is attained at the point on the budget constraint corresponding to maximum leisure (T), or zero hours of work. This outcome is referred to as a **corner solution** because the individual equilibrium occurs at one of the two extreme points on the potential income constraint.

Panel (b) shows an individual who will participate, given the individual's preferences and constraints. In this case the optimum occurs in between the two extreme points on the potential income constraint, and is referred to as an **interior solution**. With an interior solution, the equilibrium is characterized by a tangency between the budget constraint and the highest attainable indifference curve. At a corner solution, the tangency condition usually does not hold and the slopes of the highest attainable indifference curve and the budget constraint differ. For example, in the situation shown in Figure 2.6(a) the indifference curve is more steeply sloped than the budget constraint. In Figure 2.6(b), the tangency E_0 involves optimal leisure of l_0 , and labour supply or work $h_0 = T - l_0$, yielding income $W_0 h_0 + Y_N$. One can easily verify that E_0 is the utility-maximizing equilibrium by seeing what would happen if the individual were at any point other than E_0 .

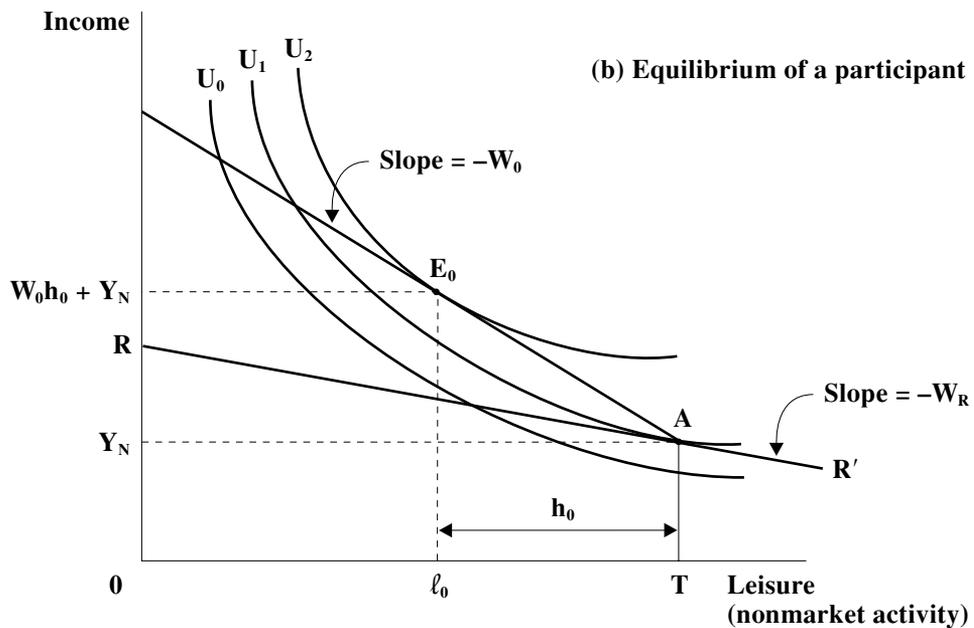
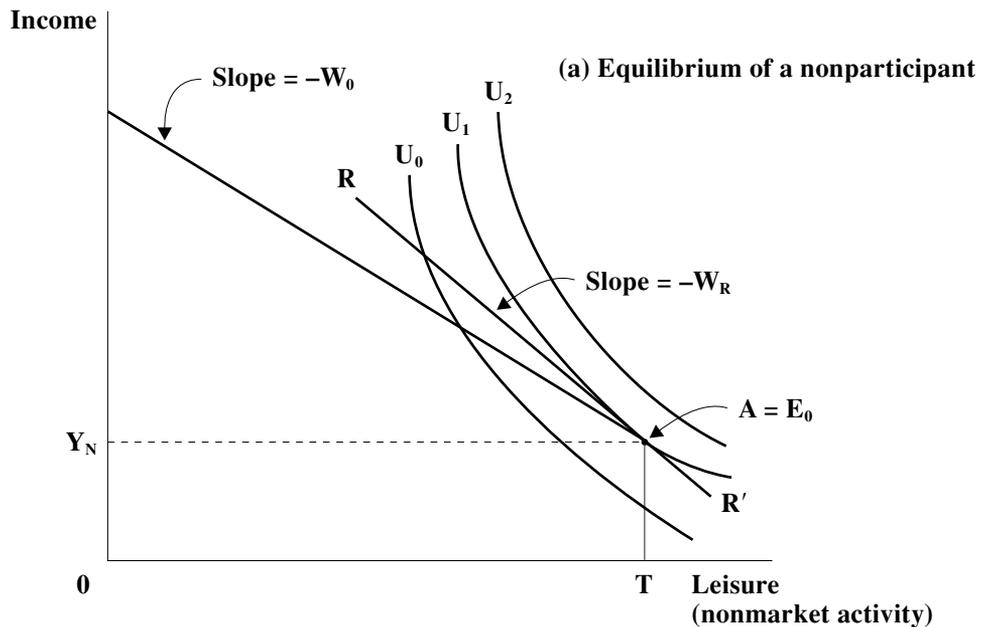
A useful way to characterize the consumer's optimum is to compare the individual's marginal rate of substitution (MRS) with the market wage rate. The former measures the individual's preferences or *willingness* to exchange nonmarket time for income, while the market wage measures the individual's *ability* to exchange leisure for income. Consider first the participation decision, and focus on the point of maximum leisure, point A in panels (a) and (b) of Figure 2.6. When the MRS at zero hours of work (slope of indifference curve at A) exceeds the wage rate (slope of budget constraint), as at point A in Figure 2.6(a), the individual's implicit value of nonmarket time is high relative to the explicit market value of that time and the individual will therefore not participate. When the MRS at zero hours of work is less than the wage rate, as at the point A in Figure 2.6(b), the individual's implicit value of leisure time is less than the explicit market value of that time and the individual will participate in the labour market. In this case, the individual increases hours of work until the MRS between income and leisure equals the wage rate, thereby exhausting all "gains from trade" associated with exchanging nonmarket for market time.

In analyzing labour force participation behaviour, economists often use the concept of a **reservation wage**, which is defined as the wage rate at which an individual would be indifferent between participating and not participating in the labour force; that is, the wage at which an individual would be indifferent between work in the labour market as opposed to engaging in nonlabour market activities such as household work, retirement, or leisure activities, all of which require time. If the market wage rate were equal to the

Figure 2.6

The Consumer's Labour Supply Decision

The utility-maximizing choice of hours worked is illustrated. In panel (a), the consumer's optimum occurs at A, with leisure equal to T and consumption Y_N yielding the highest utility. This person does not work. The reservation wage is given by the slope of RR' , and is the MRS at this point. If the wage were higher than RR' , the individual would participate. In panel (b), E_0 is an interior solution, where the indifference curve U_2 is tangent to the budget constraint. This consumer chooses l_0 units of leisure, and consumption (income) of $Y_N + W_0 h_0$.



reservation wage W_R , the individual's potential income constraint would be given by the line RR' in Figures 2.6(a) and 2.6(b). That is, the reservation wage is the slope of the individual's indifference curve at zero hours of work. If the market wage is less than the reservation wage, as in Figure 2.6(a), the individual will not participate in the labour force; if the market wage exceeds the reservation wage, as in Figure 2.6(b), the individual will participate in the labour force since the return from engaging in labour market activity exceeds that individual's valuation of time in non-labour market activities.

COMPARATIVE STATICS

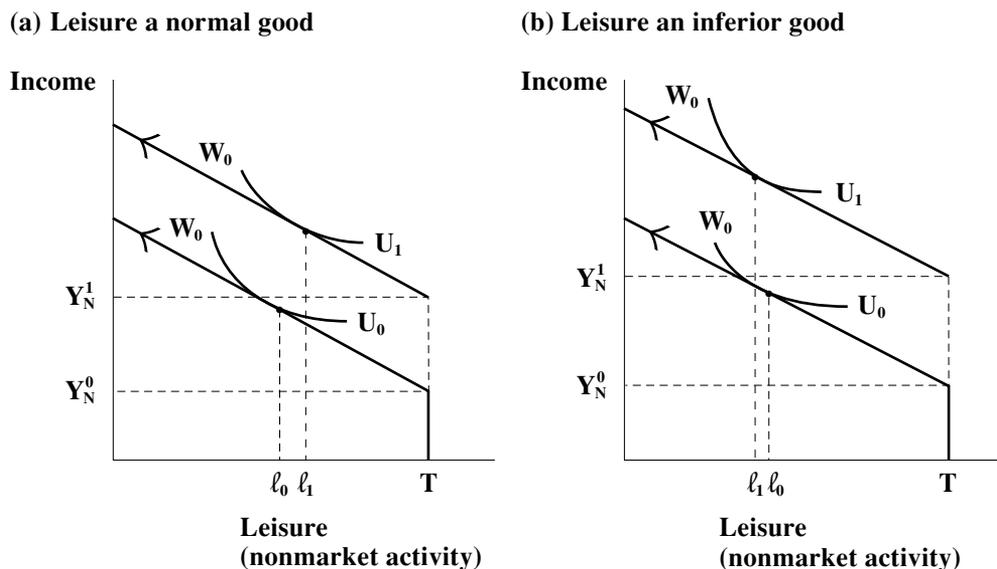
We now wish to analyze the effects of changes in key features of the economic environment on an individual's labour supply decision. We focus on the impact on labour supply of changing the wage, holding nonlabour income constant, and similarly, the effect of changing nonlabour income, holding the wage rate constant. The theoretical methodology is quite simple: given our assumption of optimizing behaviour, we need only compare the consumer equilibrium under alternative conditions, comparing the resulting labour supply choice. As we saw above, the decisions on participation and hours are integrated into the same theoretical model. However, given the continuous nature of the hours decision, versus the discrete nature of the participation decision, it is worth treating the two sides of the labour supply problem in sequence.

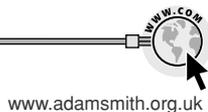
Hours of Work for Participants

The easiest change to examine in the budget constraint is the effect of changing nonlabour income on desired hours of work. An increase in nonlabour income results in a parallel shift outward of the budget constraint, as depicted in Figure 2.7. This shift outward means that the consumer can purchase more of both goods and leisure. Whether or not she decides to spend more on both depends on her preferences. Goods of which she consumes more in response to an increase in nonlabour income are called normal goods. If leisure is a **normal good**, she will consume more leisure (and work less), such as is depicted in Figure 2.7(a). In this figure, she increases her consumption of leisure (and thus decreases labour supply) from l_0 to l_1 as a result of the increase in nonlabour income. On the other hand, if leisure is an **inferior good**, she will actually consume less leisure, concentrating the additional income (and more) on the purchase of consumption. This possibility is shown in Figure 2.7(b). Economic theory makes no predictions as to whether leisure is normal or inferior. It is an entirely empirical question, depending on individual preferences. However, most available empirical evidence suggests that inferior goods are rare, and in particular, that leisure is not inferior. Through most of our discussion we will thus treat leisure as a normal good—though it must be emphasized this is a matter of pure, albeit reasonable, assumption.

Figure 2.7 The Effect of an Increase in Nonlabour Income on Labour Supply

Panels (a) and (b) compare the impact of an increase in non-labour income (from Y_0 to Y_1) on hours worked. In (a), the consumer "buys" more leisure (works less), as leisure is a normal good. In (b), the consumer "buys" more consumption, but less leisure (and works more), as leisure is an inferior good.





What will happen to the equilibrium amount of effort if wages are increased? Among the classical economists, Adam Smith discussed the “short-run” response of labour to an increase in the wage rate:

The wages of labour are the encouragement of industry, which, like every other human quality, improves in proportion to the encouragement it receives. ... Some workmen, indeed, when they can earn in four days what will maintain them through the week, will be idle the other three. This, however, is by no means the case with the greater part. Workmen, on the contrary, when they are liberally paid by the piece, are very apt to over-work themselves, and to ruin their health and constitution in a few years.¹

As this quote suggests, a change in the wage rate has two effects. On the one hand, the higher wage rate means that for each quantity of work, the worker now has more income from which to buy more of all goods, including leisure. This is termed the **income effect**, and in the case of a wage increase it leads to reduced work, assuming that leisure is a normal good. On the other hand, the individual may work more because the returns for work are greater; that is, the opportunity cost of leisure or the income forgone by not working is higher and hence the person may substitute away from leisure. This is termed the **substitution effect**, and in the case of a wage increase, it leads to increased work. In the case of a wage change, therefore, the income and substitution effects work in opposite directions; hence it is ultimately an empirical proposition as to whether a wage increase would increase or decrease the supply of work effort.

The income and substitution effects of a wage change are illustrated in Figure 2.8. After the wage increases from W_0 to W_1 , the new equilibrium is E_1 . This is the net or bottom-line

Exhibit 2.1

Importance of Paying Attention to Both Income and Substitution Effects

The capital of Brazil, Brasilia, was constructed as a planned new city. It was built in the undeveloped heartland of Brazil in the Amazon jungle area. This movement of the capital from the popular coastal area was done, in part, to encourage development in the otherwise isolated area.

Because of that isolation, it was difficult to recruit civil servants to work in the new capital, and wages were raised in order to encourage recruitment. This worked; in essence, it induced a substitution effect whereby the higher wages elicited a voluntary increase in “labour supply” to this region.

Unfortunately, the higher wages also induced an income effect that worked at cross purposes to the objective of encouraging the civil service to permanently move to the new capital and become part of a new, integrated community. That is, with their new, higher income, many of the civil servants could afford to maintain another residence in Rio de Janeiro on the coast. They could also afford to fly regularly out of Brasilia, leaving it to spend their greater wealth.

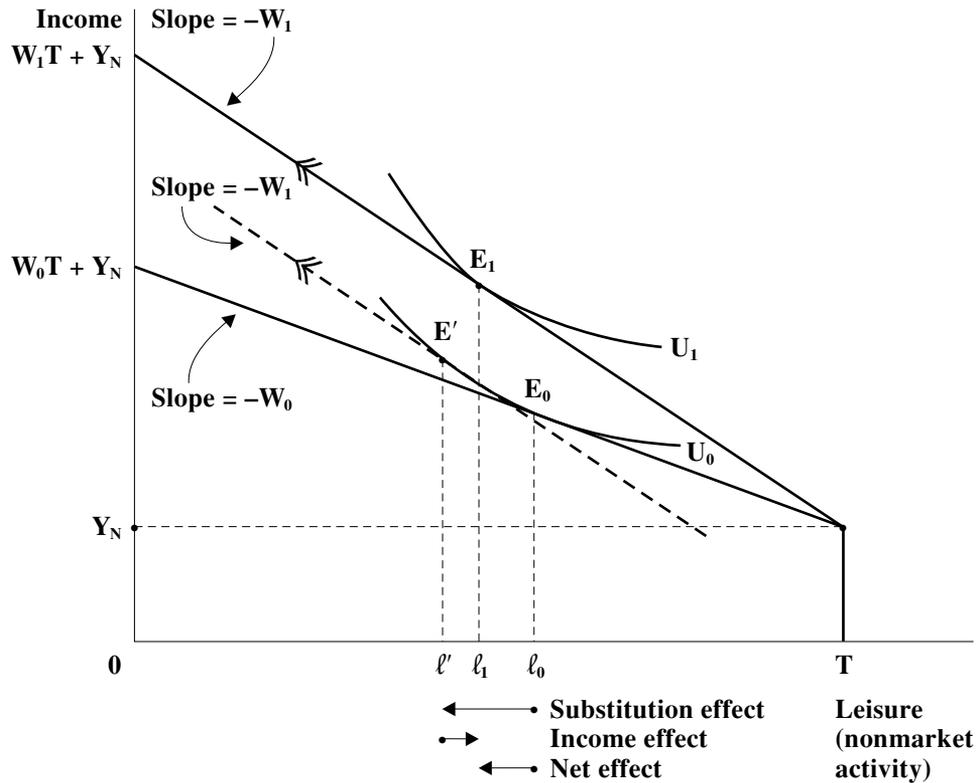
Over time, a more permanent, committed community developed. Nevertheless, it highlights the importance of paying attention to both the income and substitution effects of wage changes. The emphasis on incentive effects in many similar policy discussions often ignores the possibility that an income effect can undo a perfectly reasonable substitution effect.

¹Adam Smith, *The Wealth of Nations* (ed. E. Cannan), Book I (London: Methuen and Company, Ltd., 1930), p. 83.

Figure 2.8

Income and Substitution Effects of a Wage Increase

A wage increase from W_0 to W_1 has an ambiguous impact on hours worked. On the one hand, leisure is more expensive, and the substitution effect leads the consumer to switch away from leisure toward consumption (E_0 to E'). On the other, the increase in full income makes the consumer "richer." With leisure a normal good, the income effect leads the consumer to purchase more leisure (E' to E_1). The net result depends on the relative magnitudes of the income and substitution effects. In this example, the substitution effect dominates, and the consumer works more.



effect of the wage change. What we wish to know is: Does economic theory allow us to state whether this new equilibrium entails more or less labour supply? Like all price changes, this increase in the price of leisure can be broken down into an income effect and a substitution effect. This is critical in understanding how wage changes affect labour supply.

The substitution effect represents the effect of a pure increase in the relative price of leisure. That is, the substitution effect is that part of the consumer's adjustment that would occur if she were forced (or allowed) to remain on the original indifference curve U_0 , but maintain a tangency with a budget line with slope given by the new wage, W_1 . In consumer theory, the substitution effect is sometimes called the compensated price effect, since it refers to the effect of a price change whereby the individual is kept as well-off (i.e., compensated) as she was before the price change. Given our assumption of diminishing MRS, economic theory unambiguously predicts that leisure demand should fall with an increase in its price, holding utility constant. This is depicted in Figure 2.8 with the movement from E_0 to E' , and the resulting change in leisure demand, l_0 to l' . In fact, this is the only empirically testable prediction of labour supply theory: the substitution effect for leisure demand is negative, or equivalently, the substitution effect for labour supply is positive. The theory predicts that compensated labour supply curves must slope upward.

The income effect is given by the movement from E' to E_1 . The wage rate is held constant at the new wage rate, W_1 , but potential income increases to reflect the higher wage rate. The implicit shift in the budget constraint is given by the parallel shift of the constraint from the dashed line to the actual final budget constraint.

If leisure is a normal good, the income effect offsets the substitution effect, since it leads to an increase in the demand for leisure. In Figure 2.8 the substitution effect dominates the

income effect; hence in this illustration the wage increase resulted in a net increase in the amount of work. However, the income effect can dominate the substitution effect, in which case a wage increase will result in a decrease in the amount of work supplied.

The overall income effect can itself be broken down into two parts. First, the increase in the price of leisure *decreases* purchasing power, as it would for any other good. With a price increase, all else equal, our command over resources falls. This conventional income effect is offset, however, by the fact that full income also increases with the wage increase. Maximum potential earnings rise by $(W_1 - W_0)T$, the consumer's available time, as valued by the increase in the wage rate. This increase in full income leads to a net increase in income available for the purchase of goods and leisure.

Participation

The effects of changes in the wage rate and nonlabour income on labour force participation are illustrated in Figure 2.9. For a participant an increase in the wage rate has both income and substitution effects; because these are opposite in sign, hours of work may either increase or decrease. However, even if the income effect dominates the substitution effect so that hours of work decline, an increase in the wage rate can never cause a participant to withdraw from the labour force. This prediction of income-leisure choice theory is illustrated in Figure 2.9(a). As the wage increases, the household can achieve higher levels of utility and would never choose to not participate in the labour force which yields the utility level U_0 .

For a nonparticipant, an increase in the wage rate may result in the individual entering the labour force or it may leave labour force participation unchanged, depending on whether the now higher market wage exceeds or is less than the individual's reservation wage. As illustrated in Figure 2.9(b), the increase in the wage from W_0 to W_1 leaves the individual's preferred outcome at E_0 —that is, out of the labour force. However, the further wage increase to W_2 results in the individual entering the labour force (point E_2). This prediction can be explained in terms of income and substitution effects. An increase in the wage rate raises the opportunity cost of leisure time and tends to cause the individual to substitute market work for nonmarket activities. However, because the individual works zero hours initially, the income effect of an increase in the wage rate is zero. Thus for nonparticipants a wage increase has only a substitution effect which tends to increase hours of work.

These predictions can also be expressed in terms of the reservation wage. An increase in the market wage will not alter the labour-force status of participants because the market wage already exceeds their reservation wage. However, an increase in the market wage may cause some nonparticipants to enter the labour force if the increase in the wage is sufficiently large that the now-higher market wage exceeds the reservation wage.

Together these two predictions of the income-leisure choice model imply that an increase in the wage rate can never reduce and may increase labour force participation.

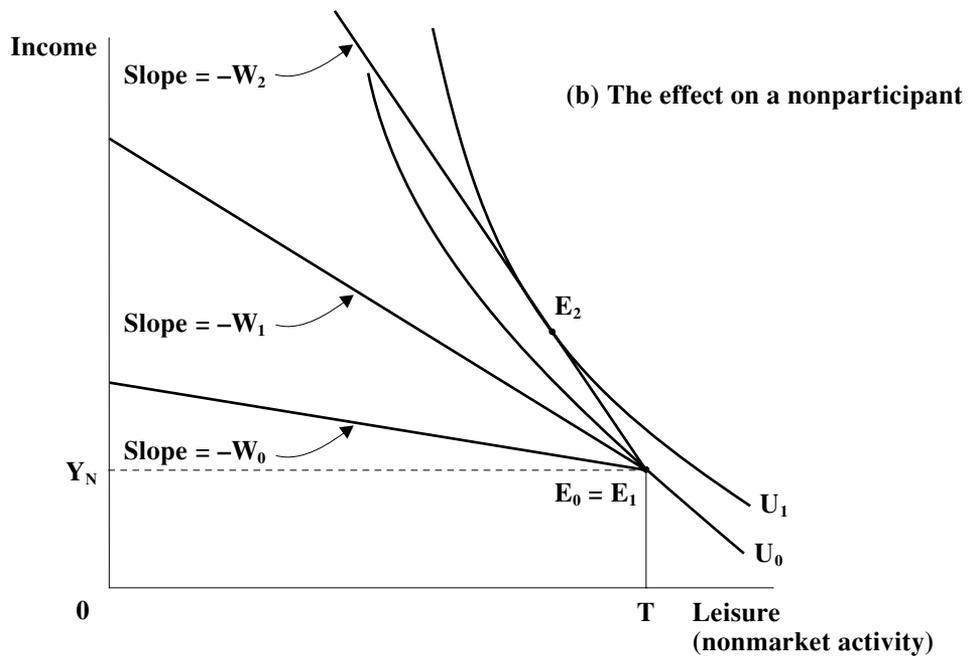
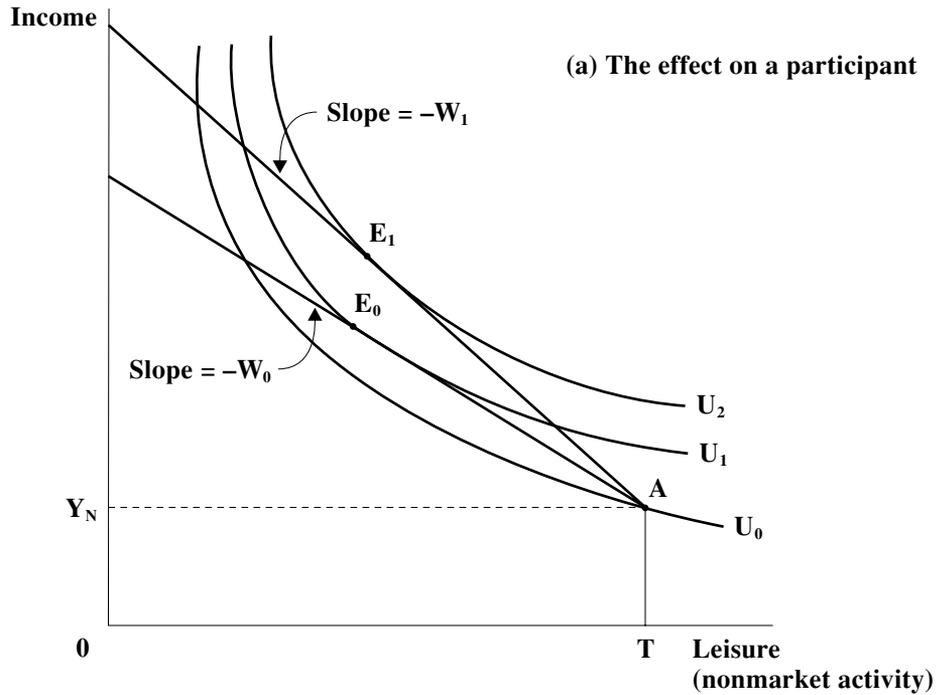
The effects of an increase in nonlabour income on labour force participation are opposite to those of an increase in the wage rate. If leisure is a normal good, an increase in nonlabour income will never cause a nonparticipant to enter the labour force and may cause some participants to withdraw from the labour force. These predictions follow from the fact that an increase in nonlabour income has a pure income effect. Thus if leisure is a normal good, the amount of time devoted to nonmarket activities must either rise (if the individual was a labour force participant in the initial equilibrium) or not decline (if the individual was already at the maximum leisure point initially).

These predictions of the income-leisure choice model can alternatively be stated using the concept of a reservation wage. If leisure is a normal good, an increase in nonlabour income will increase the individual's reservation wage. (Proof of this statement is left as an

Figure 2.9

The Effect of a Wage Increase on Participation

As a wage increase may lead to a reduction in hours worked, it will never induce a participant (in panel (a)) to stop working: the wage is still higher than the reservation wage. A non-participant, however, may be induced to work if the wage rises above the reservation wage. In panel (b), the consumer chooses to work at W_2 , but not the lower wages W_0 and W_1 .



exercise; see Review Question 5 at the end of this chapter.) Thus some participants will withdraw from the labour force, their reservation wage having risen above the market wage, while nonparticipants will remain out of the labour force, their reservation wage having been above the market wage even prior to the increase in nonlabour income.

The determinants of an individual's labour force participation decision can be conveniently categorized as to whether the variable affects the individual's reservation wage, market wage or both. Other things being equal, a variable that raises the individual's reservation wage would decrease the probability of participation in labour market activities. Such variables could be observable in the sense that data are often available on characteristics such as the presence of children, the availability of nonlabour market income, and the added importance of household work as on a farm. Or the variables could be unobservable in the sense that data are not directly available on the characteristic, such as a preference for household work.

While some such variables may affect primarily an individual's reservation wage and hence have an unambiguous effect on their labour force participation decision, others may affect both their market wage and their reservation wage and hence have an indeterminate effect on their labour force participation decision. An increase in a person's age, for example, may be associated with a higher market wage that makes participation in labour market activities more likely. However, it may also raise their reservation wage if the disutility associated with work increases relative to leisure time, and this may induce them to retire from the labour force.

DERIVING THE INDIVIDUAL SUPPLY CURVE OF LABOUR

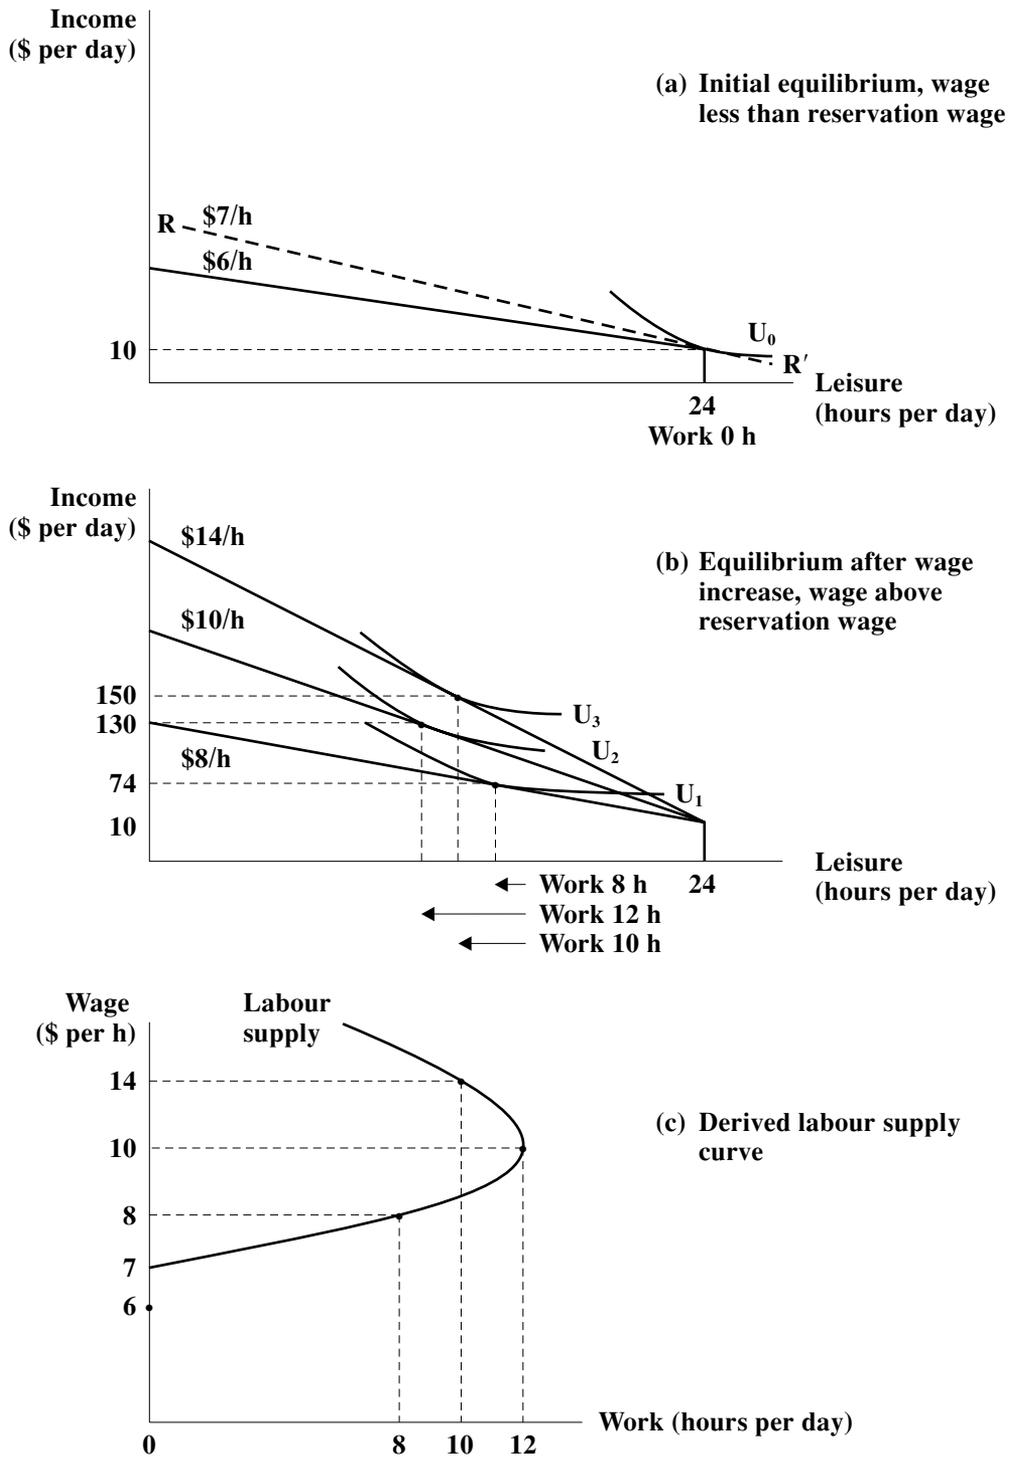
We have now met the principal theoretical objectives of this chapter: establishing the avenues through which wage changes affect labour supply, and investigating the slope of the **labour supply schedule**. As the early writings of Robbins (1930) indicate, the income-leisure choice framework can be used to derive the individual's labour supply schedule, which indicates the amount of labour that will be offered at various wage rates. By varying the wage rate, we can obtain a schedule of corresponding desired amounts of work, and hence trace out the individual's labour supply curve.

Derivation of an individual's labour supply schedule is further illustrated by the specific example of Figure 2.10 where leisure and work are measured in units of hours per day and wages in units of wages per hour. Income is the hourly wage rate multiplied by the hours of work. Figure 2.10(a) illustrates the desired hours of work for a wage rate of \$6 per hour. In this case, the market wage is below the individual's reservation wage of \$7, so the individual does not participate and desired hours of work is zero. The corresponding income is just the value of nonlabour income, \$10 per day. This gives us one point on the individual's labour supply schedule of Figure 2.10(c): at \$6/h the labour supply would be zero hours. Indeed, this would be the labour supply until the market wage reaches \$7/h. In Figure 2.10(b) we see how hours evolve as wages are raised above the reservation wage. At \$8/h, the equilibrium hours of work is 8 hours per day, and the corresponding daily income is \$64 of earnings, plus nonlabour income, for a total of \$74. This gives us the second point on the labour supply schedule: at \$8/h labour supply would be 8 hours. As the wage rises to \$10/h, hours increase to 12 hours per day. In this particular case, the substitution effect of the wage increase outweighs the income effect so that hours of work increase. This yields a third point on the supply schedule. Finally, in Figure 2.10(b), wages are raised to \$14/h and the individual is observed to decrease work effort to 10 hours per day. This yields a fourth point on the supply curve in Figure 2.10(c). At the higher wage rate, the income effect of the wage increase begins to dominate the substitution effect and the labour supply schedule becomes backward-bending. In fact, this may be a reasonable approximation for many individuals: at low wage rates they have an abundance of unmet needs, so that higher wages will induce them to work more in order to fulfill these needs; at higher wages many of their needs are fulfilled, so that additional wage increases will be used to purchase leisure.

Figure 2.10

Deriving the Individual Supply Curve of Labour

The individual labour supply curve, relating desired hours of work to the wage rate, can be derived by tracing out the labour supply choices (tangencies) in response to different wages, reflected in the slope of the budget constraint. Labour supply is zero until the wage equals the reservation wage. For higher wages, the slope of the labour supply function depends on the relative magnitudes of the income and substitution effects.



EMPIRICAL EVIDENCE

One of the most active areas of empirical research in labour economics has been the estimation and testing of models of labour supply. The research has focused on two sets of questions: First, does labour supply behaviour conform with the predictions of economic theory, and second, how responsive is labour supply to changes in the wage?

Participation, Married Women

The basic theory of labour force participation is often tested with data on married women since they may have more flexibility to respond to the determinants of labour force participation. In addition, their participation has increased dramatically over the past few decades, and this has policy implications for such things as sex discrimination, family formation, and the demand for child-care facilities.

Empirical evidence in both Canada and the United States tends to confirm the predictions based on the economic theory of participation.² For Canada, some illustrative figures are given in Table 2.3. We expect those variables that increase the market wage will be positively associated with participation, while those factors that increase the reservation wage will be negatively related to participation. In the table we show three columns. The first gives the average participation rate by the various categories. The second shows the difference in participation between a given category and the base category. The final column presents estimates of the same differentials as in the second column, but it shows regression-adjusted estimates that control for factors such as age, education, number of children, and husband's income.

Participation for married women appears to peak in the 35–44 age group, with the steepest decline for those over 55 years of age. Whether this can be interpreted as a pure aging or retirement effect will be addressed in more detail in Chapter 4. The effect of children on women's participation is puzzling at first blush: women with no children work less than those with more children. However, this raw correlation confounds the fertility effect with other factors, notably age (married 15-to-24-year-olds may still be in school). In the third column (adjusted differences), we see in fact that labour force participation is strictly declining with the number of children at home.

Exhibit 2.2

The Economics of Sleep

Sleep occupies more of our time than any other activity, with approximately one-third of our adult life devoted to sleeping. Casual observation suggests that many of our students undertake that same time allocation during class!

Presumably, we derive satisfaction from sleeping and it is an important component of leisure, and an input into our production of other activities. Hence, it should be amenable to economic analysis.

Biddle and Hamermesh (1990), in fact, provide empirical evidence from a variety of sources indicating that the demand for sleep is negatively related to the price or cost of sleeping. That is, other things equal, people of higher potential earnings power sleep less because it costs them more (in terms of forgone income) to sleep longer.

²Nakamura and Nakamura (1985) and Killingsworth and Heckman (1986) provide a survey of these and related studies.

Table 2.3 Labour Force Participation Rates of Married Women, Canada, 1996

	Participation Rate	Difference from Base Group ^a	Adjusted Difference from Base Group ^b
All Women, Total	65.0		
<i>Age</i>			
15–24	72.2	N/A	N/A
25–34	77.0	4.7	5.5
35–44	79.1	6.9	10.2
45–54	73.9	1.7	2.0
55–64	38.1	–34.1	–31.3
65–74	7.5	–64.7	–60.1
<i>Children Under 17 Years Old at Home</i>			
No children	56.5	N/A	N/A
One child	76.8	20.2	–4.6
Two children	75.7	19.2	–8.4
Three children	69.0	12.5	–14.9
Four or more children	57.9	1.4	–24.4
<i>Education</i>			
Less than Grade 9	29.4	N/A	N/A
Grade 9–13 without certificate or diploma	51.5	22.1	8.9
Grade 9–13 with certificate or diploma	66.2	36.8	17.6
Trades (with or without certificate)	74.5	45.1	23.9
Some university (without certificate or diploma)	73.5	44.1	22.2
University with certificate or diploma	77.2	47.8	27.3
University degree	83.4	54.0	30.5
<i>Husband's Income</i>			
Under \$10,000	57.8	N/A	N/A
\$10,000–\$19,999	52.3	–5.4	3.5
\$20,000–\$29,999	62.9	5.1	5.3
\$30,000–\$39,999	69.6	11.8	6.0
\$40,000–\$49,999	71.5	13.8	5.1
\$50,000–\$59,999	72.0	14.2	3.4
\$60,000–\$99,999	71.5	13.7	2.0
Over \$100,000	66.6	8.9	–1.5

Notes: Sample based on all married women between the ages of 15 and 75 with husband present.

a. Difference in participation rate from the base group (the lowest category).

b. Difference in participation rates from the base group (the lowest category), controlling for the other factors (age, education, fertility, or husband's income), estimated by ordinary least squares.

Source: Statistics Canada, Individual Use Microdata File: 1996 Census of Population.

Education has a pronounced effect on participation. In principle, higher education should be associated with higher market wages, and an increased likelihood of working. Alternatively, education may be correlated with preferences for working, or reduce desired fertility, so that the reservation wages of more educated women may also be lower. The

education coefficients alone are not enough to disentangle these competing explanations, though both are consistent with the model of participation. Finally, the results for husband's income illustrate some of the difficulties in identifying "pure" income effects. We would expect that the reservation wages of women with high-income husbands would be higher, at least if the wife's non-market time is a normal good. However, the results in column 3 show that this is the case only for women with very high incomes (over \$100,000), at least using women with very-low-income husbands as the omitted category. Using women with husbands earning \$30,000–\$39,000 as the base category, we would see participation declining with husband's income, at least for incomes \$40,000 and higher. Part of the problem in isolating a possibly larger income effect is a consequence of the fact that more-educated, higher-income working women tend to marry more-educated, higher-income husbands. This "assortative mating" confounds the income effect.

Evidence on the Elasticity of Labour Supply

A large number of econometric studies have estimated the "shape" of the labour supply schedule—that is, the responsiveness of labour supply to changes in the wage rate. As discussed previously, there are a number of components to that responsiveness. The total or gross or **uncompensated elasticity** of labour supply is the percentage change in labour supply that results from a 1 percent increase in the wage rate. Its sign is theoretically indeterminate because it reflects both the expected negative income effect of the wage change and the expected positive substitution effect. The **income elasticity** of labour supply is the percentage change in wages that results from a 1 percent increase in non-labour income. It is expected to be negative, reflecting the leisure-inducing effect of the wage increase. The **compensated elasticity** of labour supply is the percentage increase in labour supply that results from a 1 percent increase in the wage rate, after compensating the individual for the increase in income associated with the wage increase (i.e., after subtracting the income effect of the wage increase). The compensated wage elasticity of labour supply is expected to be positive since it reflects the pure substitution effect (movement along an indifference curve) as the wage rate increases the price of "leisure," inducing a substitution away from more expensive leisure and into labour market activities. For this reason, the income-compensated wage elasticity is also often called the pure substitution elasticity.

Knowledge of the separate components of the labour supply elasticity may be important for policy purposes. For example, the uncompensated wage elasticity would be used to predict the labour supply response to a wage subsidy (since it has both income and substitution effects), but the compensated wage elasticity would be used if it was assumed that the higher income would be taxed back. The income elasticity would be relevant to predicting the labour supply response to a lump-sum government transfer.

As indicated, a large number of econometric studies have estimated these various elasticities that constitute the labour supply response to a wage change. The results differ substantially depending upon such factors as the econometric technique and data. Especially for women, the results can differ depending upon whether the labour supply response refers to the participation decision, the hours-of-work decision conditional upon participation, or a combination of both. For women, this can be important because there tends to be more flexibility in their participation decision and hence it is important to take account of potential econometric problems. In particular, the subsample of labour force participants may be a select sample in terms of unobservable characteristics (like attitudes toward career or work) that can influence wages. These characteristics may, in turn, confound the estimated wage elasticity if conventional regression analysis is used. As well, their labour supply decision may be more affected by discontinuities associated, for example, with fixed costs of entering the labour market, such as daycare costs. The differences

in female labour supply responses often reflect differences in the extent to which these factors are taken into account in the estimation procedure.

In spite of the substantial variation that exists in the results of the different studies, a number of generalizations can be made. Table 2.4 provides an illustrative representation of those results based on a number of reviews that have been done in the literature (as cited in the source to the table). For example, Hansson and Stuart (1985) review approximately 50 labour supply studies for men and women. They calculate total elasticities for both sexes by weighing the male and female elasticities by their respective share of earnings. They calculate that for both sexes the overall elasticities were: uncompensated 0.10; compensated 0.25; and income elasticity -0.15 . However, on the basis of the 13 newer studies, which used more sophisticated econometric techniques to account for many of the previously discussed issues, the elasticities were: uncompensated 0.44; compensated 0.52; and income elasticity -0.08 . (These numbers illustrate how the uncompensated or gross elasticity is simply the sum of the compensated and income elasticities.) The numbers used in the first row of Table 2.4 for both sexes are simply based on a “rounded approximation” of those elasticities from the 50 studies and the 13 newer ones. For that reason they are meant to be illustrative and representative rather than strict averages of the results of the different studies. The separate figures for males and females are also meant to be only illustrative.

Exhibit 2.3

Take Me Out to the Ball Game ...

What lessons can economists draw by studying the labour supply of baseball stadium vendors? There is a long history in labour economics of trying to estimate “the” elasticity of labour supply, since it is such an important parameter for public policy. However, the labour supply elasticity almost certainly varies across groups in the population, and also over time. Increasingly, researchers focus their efforts on estimating the impact of specific changes in labour market opportunities on employment decisions, instead of estimating the more general “labour supply elasticity.” For example, they may focus on identifying the impact of a tax change or income support scheme on the hours or participation decision. The economic theory of income and substitution effects is still a useful tool in interpreting the results from such exercises.*

But there still remains an important role for trying to estimate “the” elasticity of labour supply, especially in evaluating the one truly testable implication of labour supply theory: that the substitution effect is positive. This is not an easy task for a number of reasons, but especially because of the difficulty of isolating the substitution effect from the income effect. Consider the apparently simple exercise of regressing the number of hours worked on an individual’s wage and non-labour income. Ideally, we would like to use the estimated wage coefficient in order to construct the substitution elasticity, under the assumption that we can estimate the income effect from the non-labour income elasticity. A number of statistical problems make this difficult.

First, it is difficult to measure non-labour income. Almost all forms of income are somehow related to labour earnings: government transfer income depends on having low labour earnings (because of either low wages or low hours), and investment income is generally higher for those with higher wages and hours. This could lead to either an over- or an understatement of the true income effect. Second, individual preferences may be correlated with wages. For example, ambitious people may have both higher wages and higher hours. We would thus exaggerate the

The following generalizations are illustrated in Table 2.4:

1. The overall labour supply schedule for both sexes is likely to be slightly upward-sloping; that is, a wage increase does lead to a slight increase in the amount of labour supplied to the labour market. The representative elasticity of 0.25 indicates that a 1 percent increase in real wages would lead to a one-quarter of 1 percent increase in labour supply. This small uncompensated total elasticity is a result of the positive pure compensated (substitution) elasticity slightly outweighing the negative income elasticity.
2. For males, however, the labour supply schedule is likely to be slightly backward-bending; that is, real wage increases are associated with a reduction in the amount of labour supplied to the labour market. This overall effect is very small and could well be zero (i.e., vertical or perfectly inelastic labour supply) or even slightly forward-sloping. The small overall negative elasticity is a result of a weak positive substitution elasticity being outweighed by a weak but slightly larger negative income elasticity.
3. For women, the labour supply schedule is more strongly forward-sloping; that is, an increase in real wages is associated with a more substantial increase in the amount of labour supplied to the labour market. This is the result of a strong positive substitution elasticity outweighing the weak negative income elasticity.

Exhibit 2.3

Take Me Out to the Ball Game ... (*continued*)

impact of wages on their labour supply decision. Third, an increase in someone's wage may have a long-run impact on his lifetime earnings potential, leading to a potentially large income effect. Thus, without a clean estimate of either the wage or the income elasticities, it is difficult to test whether the wage response reflects a positive substitution effect. An ideal "experimental design" would involve comparing an individual's labour supply decisions on a day-by-day basis, when the individual has complete knowledge of his or her "long-run" income, and the patterns of wage rates across different days. In this way, a researcher could attribute higher effort on high-wage days to a pure (positive) substitution effect.

That's where the baseball vendors come in. In a clever study that follows this ideal procedure, Gerald Oettinger (1999) examines the labour supply of vendors in a professional baseball stadium (in Arlington, Texas). He finds that vendors are more likely to work on days where the expected returns (sales, and thus wages) are higher. Apparently, vendors respond to the size of the anticipated crowd, on the basis of the opposing teams, the Texas Rangers' performance, the weather, and promotional events (such as "Free Hat Day"). Even taking into account the fact that the labour supply response also leads more vendors to show up (and thus soak up some of the additional sales), Oettinger estimates that a 1.00 percent increase in the expected daily wage leads to a 0.60 percent increase in the number of vendors working. Because he can identify each vendor's daily choices, he can hold constant individual preferences, non-labour income, and the value of the time-endowment over the entire baseball season, and this response can reliably be interpreted as a pure substitution effect. While the estimated labour supply elasticities may not generalize to other types of workers, Oettinger's paper shows how the careful study of even a very specific group of workers can, more generally, shed light on the possible validity of labour supply theory.

*See Blundell and MaCurdy (1999).

Table 2.4 Compensated and Income Elasticity of Labour Supply

Sex	Uncompensated (Gross, Total) Wage Elasticity of Supply	Compensated (Substitution) Wage Elasticity of Supply	Income Elasticity of Supply
Both sexes	0.25	0.40	-0.15
Men	-0.10	0.10	-0.20
Women	0.80	0.90	-0.10

Sources: As discussed in the text, these are “representative illustrative” numbers based on different reviews of over 50 econometric studies of labour supply. These reviews include Hansson and Stuart (1985), Killingsworth (1983), Killingsworth and Heckman (1986), and Pencavel (1986).

4. The strong positive total elasticity for females is sufficiently strong to outweigh the weak negative total elasticity for males, so that the aggregate supply schedule for both sexes is likely to be forward-sloping, as discussed.
5. The substantial variation in the magnitudes of these effects across studies suggests that these representative numbers be used with caution. As well, there is some evidence³ that the female labour supply response is closer to the male response than portrayed in many studies—a result that may not be surprising if female labour market behaviour is becoming more like male behaviour over time.

Section Two: Extensions and Applications

One of the most important assumptions of the labour supply model is that individuals can freely choose any package of consumption and leisure along the linear potential income constraint. Unfortunately, this assumption may not be valid for those individuals facing constraints in the labour market. There may be no work available at the going wage rate, resulting in unemployment for the individual. Alternatively, the individual may have only limited choice over the number of hours worked. As we suggested previously, however, the model is quite flexible, and some of these issues can be incorporated into the existing framework.

ADDED AND DISCOURAGED WORKER EFFECTS

Recall that the definition of the labour force includes both the employed (those who wish to work, and have jobs) and the unemployed (those who wish to work, but do not have jobs). One interesting question pertains to how labour force participation responds to changes in the unemployment rate. Specifically, in periods of high unemployment people may become discouraged from looking for work and drop out of the labour force, returning to household activities or to school or perhaps even entering early retirement. This is termed the **discouraged worker** effect. On the other hand, in periods of high unemployment, some may enter the labour force to supplement family income that may have deteriorated with the unemployment of other family members. This is termed the **added worker** effect.

Changes in unemployment are a proxy for transitory changes in expected wages and other income, and thus the discouraged and added worker effects can be interpreted as short-run substitution and income effects respectively. That is, if unemployment is high,

³Nakamura, Nakamura, and Cullen (1979) and Robinson and Tomes (1985) for Canadian women. Nakamura and Nakamura (1981) for both Canadian and U.S. women.

then opportunities in the labour market are lowered temporarily: the price of leisure (opportunity cost of forgone income from not working) is reduced temporarily. The wage might fall below an individual's reservation wage, leading him to drop out of the labour force. On the other hand, the high unemployment means that it is more likely that family income is lowered temporarily, lowering other household members' reservation wages and inducing them to participate in labour market activities so as to restore that income.

Since the added and discouraged worker effects operate in opposite directions, one lowering the reservation wage and the other lowering the returns to working, it is necessary to appeal to the empirical evidence to see which dominates. Again, most of the empirical tests have been based on data of married women since they are more likely to respond to changes in unemployment. In the United States, the empirical evidence clearly indicates the dominance of the discouraged worker effect for most married women. That is, in periods of high unemployment, women become discouraged from entering the labour force to look for work, and this dominates any tendency to add themselves to the labour force to maintain their family income. In Canada, however, the empirical evidence is mixed.

HIDDEN UNEMPLOYMENT

The discouraged worker effect also gives rise to the problem of **hidden unemployment**—a topic that we will return to in a later chapter on unemployment. During a recession when the unemployment rate is high, there may also be a large number of discouraged workers who do not look for work because they believe that no work is available. They are considered to be outside the labour force because they are not working or actively looking for work. However, because they would look for work were it not for the high unemployment rate, they are often considered as the hidden unemployed—people whom the unemployment figures tend to miss because they are considered outside the labour force rather than unemployed persons in it.

In this sense, in times of high unemployment our unemployment figures may understate the true unemployment by missing those discouraged workers who constitute the hidden unemployed. The Survey of Job Opportunities (SJO) is an annual supplement to the labour force survey designed to measure the importance of this group. It provides information on the desire for work among nonsearchers, and the reasons for nonsearching for those who desire, but are not looking for work.

The notion of the discouraged worker is especially important as groups like married women and teenagers become a larger portion of our potential labour force. Such persons often have a loose attachment to the labour force, moving in and out in response to changing economic conditions. They often, but not always, have other family income upon which to rely. For these reasons they are often labelled “secondary” or “marginal workers”—unfortunate misnomers if they are used to belittle their employment problems or the contribution to the family income of such workers. This is especially the case if they are contrasted with “primary workers” or “breadwinners,” terms often used to describe the employment position of males.

However, because of their flexibility with respect to labour market activities, married women and teenagers do constitute a large portion of discouraged workers who would enter the labour force and look for work if employment prospects increased. Hence, there is the recent emphasis in labour force surveys (and supplements like the SJO) to find out exactly what these people are doing, and why they are not looking for work.

Clearly the decision to include such persons either in the category of unemployed or as outside of the labour force is a difficult one. Recent work by Jones and Riddell (1999) suggests that while the marginally attached are closer in behaviour to the unemployed than the remainder of the nonemployed, it is not strictly appropriate to pool them with either group.



www.statcan.ca/english/Dli/ftp.htm

Exhibit 2.4**Are Women Working to Pay the Mortgage?**

In the labour supply model, individuals are assumed free to choose whether to work. As we saw with the added worker effect, someone's labour supply may increase in response to unexpected declines in family income caused by the unemployment of a breadwinner. In such circumstances, increased female participation need not be associated with improvements in family well-being.

Mortgage pre-commitments and borrowing constraints provide another reason why a woman's labour market participation may reflect household constraints, as much as enhanced opportunities for women. Nicole Fortin (1995) examines the labour force attachment of a sample of Canadian women according to the mortgage-holding status of the household. She finds that as a family's mortgage commitments approach the fraction of family income permitted by the bank, the wife's participation rate increases sharply. One (extremely simplified) way to model this is to treat the mortgage payment as a pre-determined financial commitment, outside the household's current choice set. In that case, larger mortgage payments have a similar effect on participation as a reduction in family income. More generally, higher mortgage commitments, especially those close to the maximum that banks will permit, reduce the wife's reservation wage. Of course, this ignores the fact that house purchases are a choice that households make, distinguishing the interpretation from the added worker effect. The joint aspect of the housing and labour supply decision is treated more thoroughly in Fortin's paper. Furthermore, interest rate or housing price swings may lead to genuine changes in the budget constraint that look like typical income effects.

Fortin estimates that the effect of mortgage commitments can be quite large. For example, the effect on a woman's participation of moving into a category where the ratio of mortgage payments to husband's income is close to the maximum of what banks will tolerate has as large an effect as the presence of two pre-school children (though in the opposite direction).

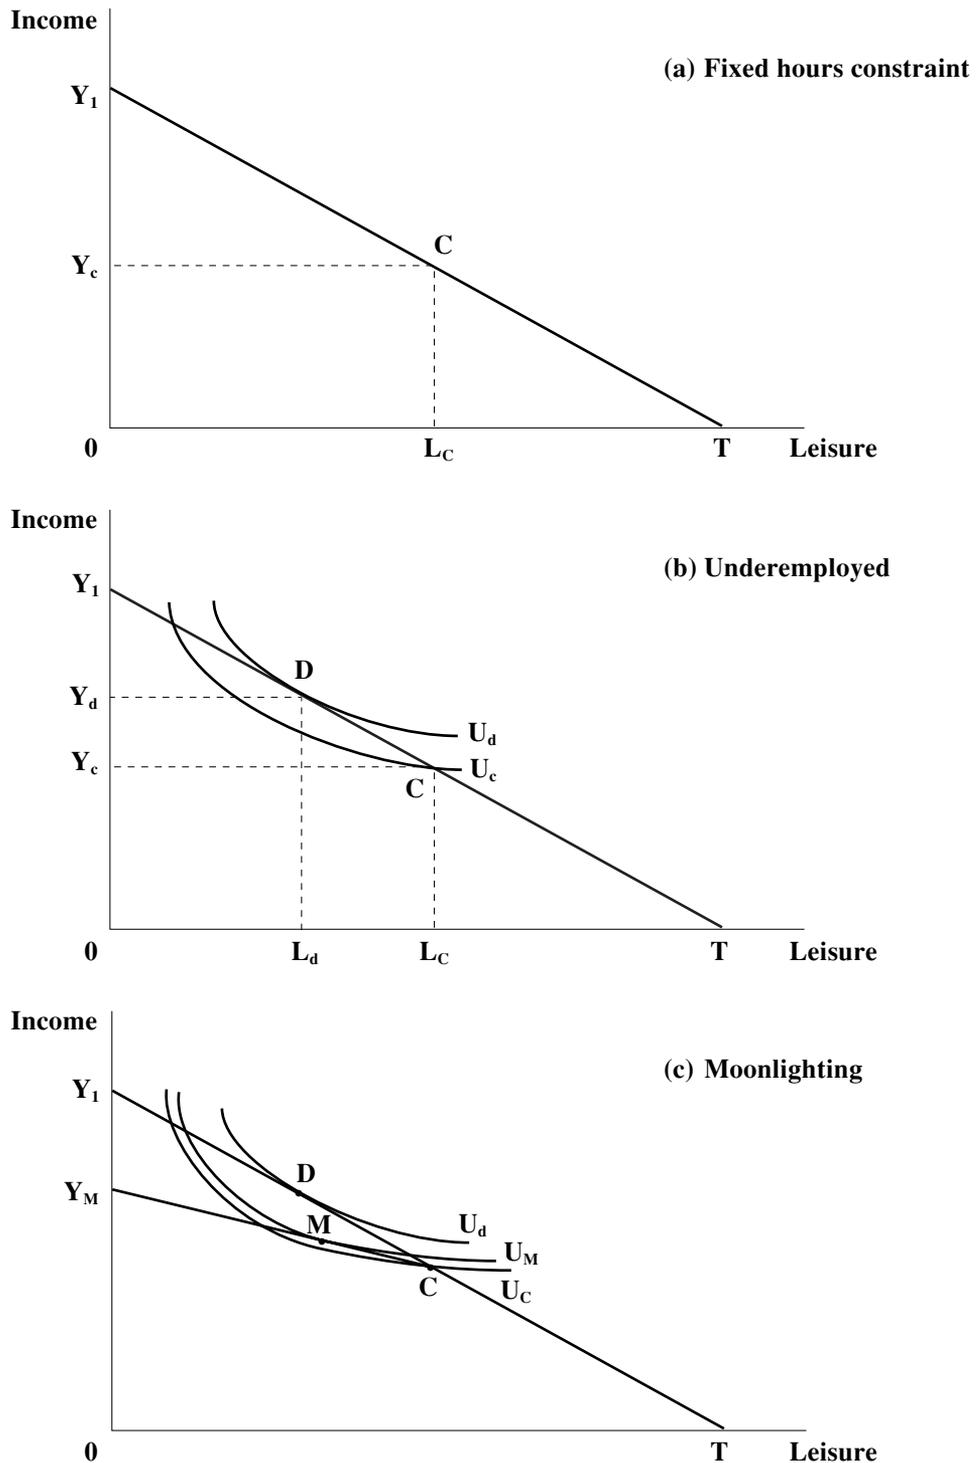
MOONLIGHTING, OVERTIME, AND FLEXIBLE WORKING HOURS

Any analysis of the hours-of-work decision must confront the following basic question: Why is it that some people moonlight at a second job at a wage less than their market wage on their first job, while others require an overtime premium to work more? This apparent anomaly occurs because people who moonlight are underemployed at the going wage on their main job while people who require an overtime premium are already overemployed at the going wage on their main job. Underemployment and overemployment, in turn, occur because different workers have different preferences and they tend to be confronted with an institutionally fixed work schedule. The fixed hours of work, in turn, can arise because of such factors as legislation, union pressure, or company personnel policy. The need for interaction among employees (or between employees and the firm's physical capital) may cause the employer to set standardized hours of work, rather than allowing each employee to work his or her preferred number of hours. Again, economic theory will enable us to analyze the labour market impact of these important institutional constraints.

The fixed-hours-of-work phenomenon is illustrated in Figure 2.11(a). In this case, the worker is faced with two constraints. The first is the usual budget constraint TY_1 as

Figure 2.11 Fixed Hours Constraint, Underemployment, and Moonlighting

Consumers may be constrained in their choice of hours worked, and confronted with the point C, a fixed hours and income “package.” As shown in panel (b), given the wage, the consumer would rather choose D, working $T - L_D$ hours. As this exceeds the available hours of work, the consumer is underemployed. As shown in panel (c), if a second job is available, even one paying a lower wage, the consumer may choose M, working $T - L_C$ on the main job, and moonlighting on the second.



determined by the person's hourly wage. This restricts the worker's maximum choice set to the triangular area TY_1O , where O denotes the origin. The second constraint of the fixed workday of TL_c hours (recall work is measured from right to left) restricts the worker's maximum choice set to the area L_cCY_cO : the worker can take no more leisure than OL_c (work no more than TL_c), and earn no more income than OY_c even if he worked more than TL_c . In effect, this reduces the worker's realistic choice set to the point C since C will always be preferred to other points within L_cCY_cO .

Moonlighting and Underemployment

Some individuals, however, may prefer to work more hours at the going wage rate. Figure 2.11(b) illustrates the case for an individual whose preferences (indifference curve U_d) are such that he would prefer to be at D , working TL_d hours at the going wage and taking home an income of Y_d . However, because of the hours of work constraint, the worker must be at C , obviously being less well off since $U_c < U_d$. In fact, the difference between U_d and U_c is a measure of how much the worker would be willing to give up in order to have the fixed-hours constraint relaxed.

A variety of implications follow from this analysis. The worker is **underemployed** because he would like to work more at the going wage rate. Because of the additional constraint of the fixed working hours, the worker is also less well off ($U_c < U_d$) and may be seeking a different job that would enable him to achieve his desired equilibrium at D . In addition, the worker may be willing to engage in **moonlighting** and do additional work at a wage rate that is lower than the wage rate of the first job.

This moonlighting equilibrium is illustrated in Figure 2.11(c) by the budget constraint CY_m . (To simplify the diagram the details of Figure 2.11(b) have been omitted.) This new budget constraint rotates downward from CY_1 because the moonlighting wage, which is less than the regular wage as given by the slope of CY_1 , applies only to hours of work beyond TL_c . In spite of the lower moonlighting wage, the worker is willing to work more hours (move from C to M) because of the greater utility associated with the move ($U_m > U_c$). That is, workers who are underemployed at the going wage rate would be willing to moonlight at a lower wage rate on their secondary job.

Overtime and Overemployment

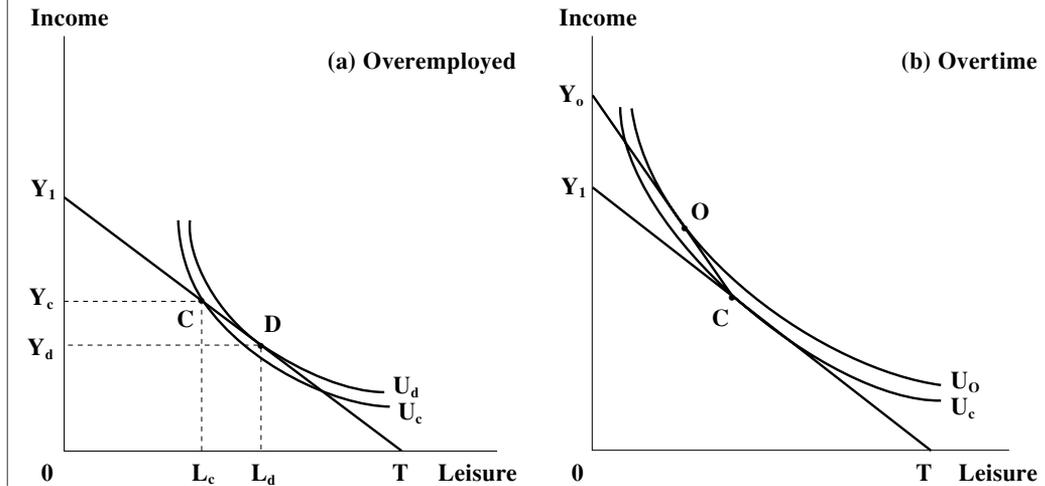
Other individuals, however, may prefer to work fewer hours at the going wage rate. Figure 2.12(a) illustrates the situation where the worker would prefer (maximize utility) to be at D , working TL_d hours for an income of Y_d . However, because of the institutionally fixed work week, she is compelled to be at C , being less well off ($U_c < U_d$) even though she takes home more income ($Y_c > Y_d$).

Such a worker is **overemployed** at the going wage rate and consequently would support policies to reduce the institutionally fixed work week. In addition, she also may be seeking a different job that would enable her to work fewer hours at the going wage rate, and she may even exhibit absenteeism and tardiness as ways of moving toward her desired work week. Because such a worker is already overemployed at the going wage rate, she would not willingly work more hours at the wage rate; however, she may be induced to do so by an **overtime premium**.

This overtime premium is illustrated in Figure 2.12(b) by the budget constraint CY_o that rotates upward from CY_t because the overtime premium, which is greater than the normal wage that determines the slope of TY_1 , applies only to overtime hours of work beyond TL_c . If, for example, the overtime premium is time-and-a-half, then the overtime budget constraint has a 50 percent greater slope than the regular straight-time budget constraint. As long as the worker is willing to give up some leisure for additional income, then there is an overtime premium that will induce her to work more, for example, to move to point O .

Figure 2.12**Overemployment and Overtime**

A consumer may be constrained to work more hours than desired. In this case, he or she would prefer to work at D, at fewer hours than C, though C is still preferred to not working at all. However, this consumer could be induced to work more hours, corresponding to point O, if offered a higher wage rate for hours worked after C.



on Figure 2.12(b) (on the indifference curve U_o). Because the worker is overemployed at the going wage rate, the overtime premium is necessary to get her to work more hours.

The person works longer hours even though he is overemployed at the going wage rate because the overtime premium results in a large substitution effect, by making the price (opportunity cost, income forgone) of leisure higher only for the overtime hours. The overtime premium has a small income effect because the budget constraint rotates upward *only* for the overtime hours; consequently, it does not have an income effect for the normal straight-time hours. Recall that the substitution effect was illustrated by a changed *slope* in the budget constraint, while the income effect was illustrated by a *parallel shift* of the constraint. Since the overtime premium changes the slope for the overtime hours, it is primarily a work-inducing substitution effect, with little leisure-inducing income effect.

Overtime Premium Versus Straight-Time Equivalent

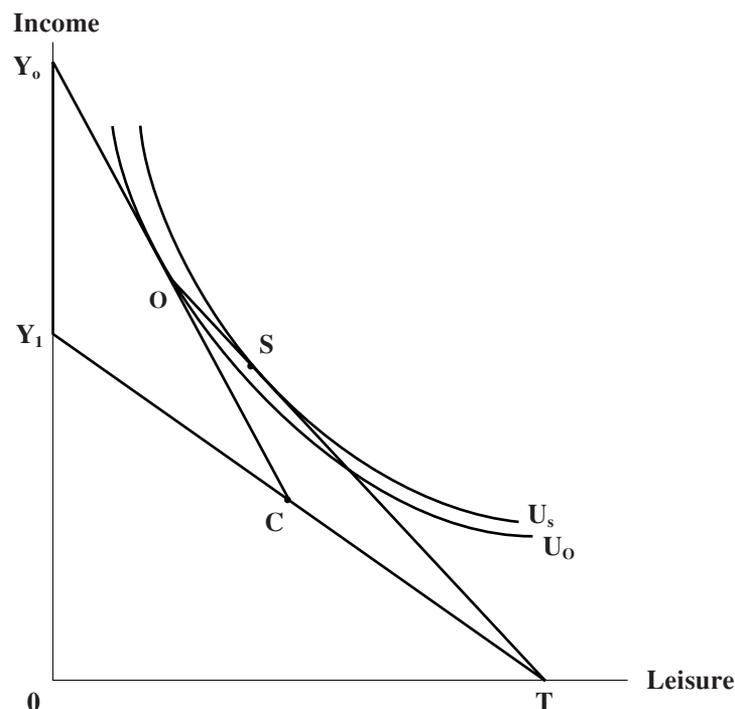
The importance of the relative absence of the income effect in the overtime premium can be illustrated by a comparison of the overtime premium with the straight-time equivalent. One might logically ask the question: If workers are constantly working overtime, why not institutionalize that into a longer workday and pay them the straight-time equivalent of their normal wage plus their overtime wage?

This alternative is illustrated in Figure 2.13. The overtime situation is illustrated by the budget constraint TCY_o with TC being the normal wage paid during the regular work period and CY_o being the overtime premium paid for overtime hours. The normal wage is assumed not to change as a result of the overtime premium. The regular work period would be TL_c hours and overtime hours would be L_cL_o . (To simplify the diagram, these points are not shown; however, as in Figure 2.12(a), they are simply the points on the horizontal leisure axis vertically below their corresponding equilibrium points.) The straight-time hourly equivalent for TL_o hours of work is given by the budget constraint TO , the slope of which is a weighted average of the slopes of the regular wage TC and the overtime premium CO . The straight-time hourly equivalent is derived by simply taking the earnings associated with the overtime plus regular time hours of work, TL_o , and determining the straight-time wage, TO , that would yield the same earnings.

A worker who is paid the straight-time equivalent, however, would not voluntarily remain at O, but rather would move to the point S, which involves less work. This is so

Figure 2.13**Overtime Premium Versus Straight-Time Equivalent**

An overtime premium can be used to induce a worker to choose hours corresponding to the point O, after paying the straight-time wage to point C. Why not simply offer the consumer a straight-time equivalent, equal to the average wage associated with the earnings and hours at O? The income effect associated with the higher wage received from the first hour worked will lead the consumer to choose fewer hours worked (at S) than when the higher (overtime premium) is only paid on hours in excess of those at C.



because the wage line TSO has a larger leisure-inducing income effect, whereas the overtime premium COY_0 is dominated by the work-inducing substitution effect (rotation of the budget constraint). In essence, since the overtime premium is paid only for hours *beyond* the regular workday, the person has to work more to get the additional income.

Overtime premiums, therefore, may be a rational way for employers to get their existing work force to voluntarily work more hours, even if they are overemployed at their regular workday. Employers, in turn, may want to work their existing work force longer hours, rather than hiring additional workers, so as to spread their fixed hiring costs over longer hours. The importance of these fixed hiring costs will be analyzed in more detail when labour demand is discussed.

Workers need not be overemployed at their going wage rate for the overtime premium to work. One can easily portray the situation for a worker in equilibrium at the regular wage rate who would unambiguously work more when offered an overtime wage premium, but who may work less if offered a straight-time hourly wage increase. Firms that want to work their existing work force longer hours may prefer a wage package that involves a low regular wage and a high overtime premium to a package that costs them the same but that involves a straight-time wage increase.

Again, what at first glance appear to be costly and irrational actions on the part of firms—in this case the coexistence of overtime and moonlighting rates and a preference for overtime premiums over straight-time equivalent earnings—may well be rational actions when viewed in the larger picture where the parties are optimizing with respect to legal institutional constraints and when the varying preferences of individual workers are considered. Rather than rendering economic theory irrelevant in the face of such constraints, they highlight the usefulness of economics in analyzing the impact of the

constraints and in explaining why, in fact, they may arise as an endogenous institutional response to the peculiarities of the labour market.

Allowing Choice in Working Hours

As illustrated previously, the composition of the work force has been changing dramatically in recent years. This is evidenced by such phenomena as the increased labour force participation of women, the dominance of the two-earner family, and the aging of the work force. Given such changes it is not surprising that these different groups would have different tastes and preferences for alternative worktime arrangements in the labour market. As well, they will face different household constraints.

For example, two-earner families may prefer part-time employment for one party, reduced hours of work for the other party, and flexible worktime arrangements, as well as the right to refuse overtime work for both parties. This would enable them to better combine labour market work with their household activities. In contrast, the one-earner family may prefer a long work week (e.g., with regular overtime) for the single earner so as to earn a family income similar to that of the dual-earner family. In essence, the growing diversity of the work force has given rise to a growing diversity of preferences for alternative work-time arrangements. Preferences for such arrangements are no longer dictated by the former stereotypical male “breadwinner” in a single-earner family.

This diversity of preferences is illustrated in the results of the Statistics Canada Survey of Work Reductions, conducted as a supplement to the June 1985 Labour Force Survey (Benimadhu, 1987). According to that survey only about one-third of the work force was content with its worktime arrangements. The two-thirds who were discontented were about evenly divided between wanting more work for more pay and wanting less work for a corresponding reduction in pay. The strongest preference for worktime reductions was from women in the childrearing years (25 to 34) and women with children under the age of five.

The basic income-leisure choice framework can be used to illustrate that there are gains to be had by employers providing alternative worktime arrangements to meet the divergent tastes and preferences of an increasingly heterogeneous work force. These benefits are illustrated in Figure 2.14. Point C illustrates where workers are constrained to operate given the all-or-nothing choice of working TL_C hours (points on the leisure axis are not marked to simplify the diagram) at the going wage (slope of TY_C). However, many workers have different preferences. Some, for example, may prefer to be at point D. Because they are overemployed at the going wage rate, their discontent ($U_C < U_D$) may be exhibited in the form of costly absenteeism, high turnover, and perhaps reduced morale and productivity.

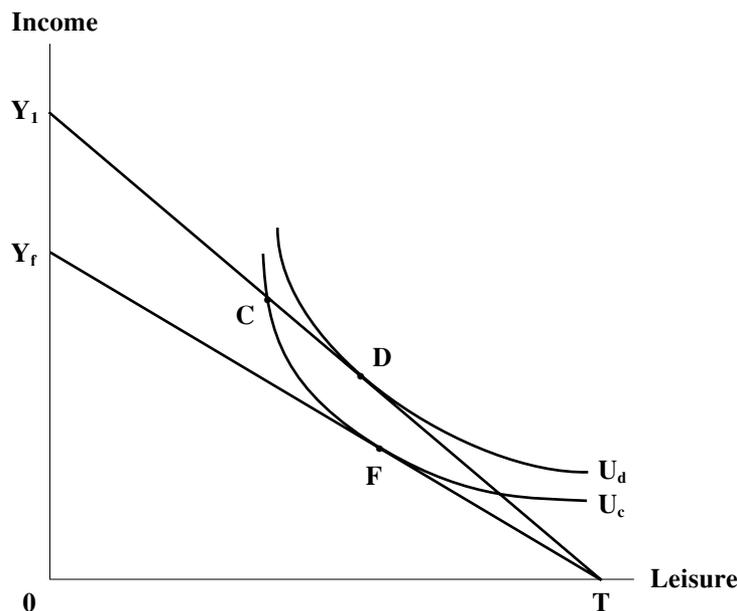
Obviously firms that allowed such workers to work their desired hours of work could save on these costs. Alternatively, such firms could lower their wage rates and still retain their work force. This is illustrated by the wage line TY_F , which could be lowered until the point of tangency, F, with the original utility curve U_C . Workers are equally well off at C and F (same level of utility U_C) even though F implies a lower wage rate, simply because they are at an equilibrium with respect to their hours of work. In essence, they are willing to give up wages in return for a work schedule that meets their preferences.

Competition for such jobs would ensure that firms could offer lower wages in return for more flexible work schedules. In this sense, the gains from flexible work schedules could be recouped by the firm to cover the other costs that may be associated with such schedules. Firms that offer more flexible hours need not lower wages, but may take the benefits in the form of reduced absenteeism, lower turnover, and improved worker morale. Various combinations of reduced wages (downward-rotated wage line) and improved worker morale (higher indifference curve), of course, are possible.

While there are benefits from allowing workers to choose their preferred hours of

Figure 2.14**Gains from Alternative Work Schedules**

Consumers may be willing to pay for flexibility in choosing their hours worked. A consumer constrained to working at C would be indifferent to any lower wage and hours package lying on the indifference curve U_c . The wage could fall as low as that represented by TY_f , where the consumer would choose hours associated with point F.



work, there are costs, especially if such flexibility is given to different workers in the same establishment. Individual differences in hours worked can give rise to problems of monitoring, supervision, communication, scheduling, and coordination. One compromise is flextime, whereby the workers are required to work a fixed number of hours (e.g., eight per day), a certain “core” hours (e.g., 10:00 a.m. to 3:00 p.m.), but the beginning and ending times are flexible. This helps meet the divergent preferences of different workers (e.g., late risers versus early risers); it enables some with childraising responsibilities to be home when children return from school; and it facilitates avoiding rush-hour commuting problems.

Compressed work weeks are another alternative worktime arrangement. Common compressed schedules include four 10-hour days or even three 12-hour days. These are often attractive to employees because of the longer “weekends” that are involved. They also enable the amortizing or spreading of fixed daily commute costs over a longer workday. For example, a two-hour commute over three 12-hour days is six hours of commute time per week, as against ten hours based on a five-day week. As well, the 12-hour day may avoid rush-hour commutes. Such compressed work weeks have been used in some areas (e.g., nursing) as a recruiting device to help attract and retain personnel.

Clearly, there may be costs associated with such alternative worktime arrangements and these have to be weighed against the potential benefits. The point made here, and illustrated in the income-leisure framework, is that there are potential benefits to meeting the divergent tastes and preferences of workers. As well, these benefits are likely to be growing as a result of the increasing diversity of the work force.

APPENDIX: CONSUMER CHOICE THEORY

The economic theory of consumer behaviour and labour supply has two main building blocks: preferences and constraints. Preferences summarize what the individual or household wishes to achieve. Constraints summarize what is feasible, that is, what the individual is able to achieve. The key assumption of the theory is that individuals choose from among

the feasible outcomes that outcome which yields the highest level of satisfaction or well-being. In this appendix, we review the theory of consumer behaviour, and show how it can be adapted for labour supply.

Constraints

Assume that there are two goods, X and Y, that the consumer can buy in the market. The budget set, or budget constraint, summarizes the set of consumption bundles of X and Y that she can consume. Her expenditure must be less than or equal to her budget M, so that purchases must satisfy

$$P_X X + P_Y Y \leq M$$

where P_X and P_Y are the prices of X and Y purchased in the market. This set of feasible consumption packages is illustrated by the shaded area of Figure 2.A1(a). Because we assume that “more is better” for the consumer, she will exhaust her budget and we can focus on the outer frontier of the budget set, replacing the inequality with equality:

$$P_X X + P_Y Y = M$$

This is simply the equation of a straight line, and we call it the budget constraint.

In labour economics it is important to understand how to manipulate budget constraints, so it is worth reviewing the key features of this very simple budget constraint. It is a simple exercise of high school algebra to see that the slope of the budget constraint is $-P_X/P_Y$ (i.e., rise/run). This slope represents the relative price of X to Y, and it indicates that if the consumer wishes to increase her consumption of X by 1 unit, she would have to give up P_X/P_Y units of Y. The endpoints of the budget line occur where the budget constraint intersects with the X and Y axes. If the consumer spends all her budget on X, she will be able to consume M/P_X units of X, and zero Y. Alternatively, she can buy nothing but Y, consuming a maximum of M/P_Y units of Y.

A simple extension to this framework is illustrated in Figure 2.A1(b). In this case, we assume that the consumer has an initial endowment of \bar{X} units of X. We assume that she can sell all or part of her endowment on the market at the market price for X, P_X . The easiest way to see how this affects her budget set is to assume that she first sells her entire endowment, and then adds the proceeds $P_X \bar{X}$ to her previous budget, M. She can then decide how to spend her budget on X and Y. It is possible to simplify the consumer's problem in this manner because there are no transaction costs associated with buying or selling X. Her new budget constraint is then given by

$$P_X X + P_Y Y = M + P_X \bar{X} = \bar{M}$$

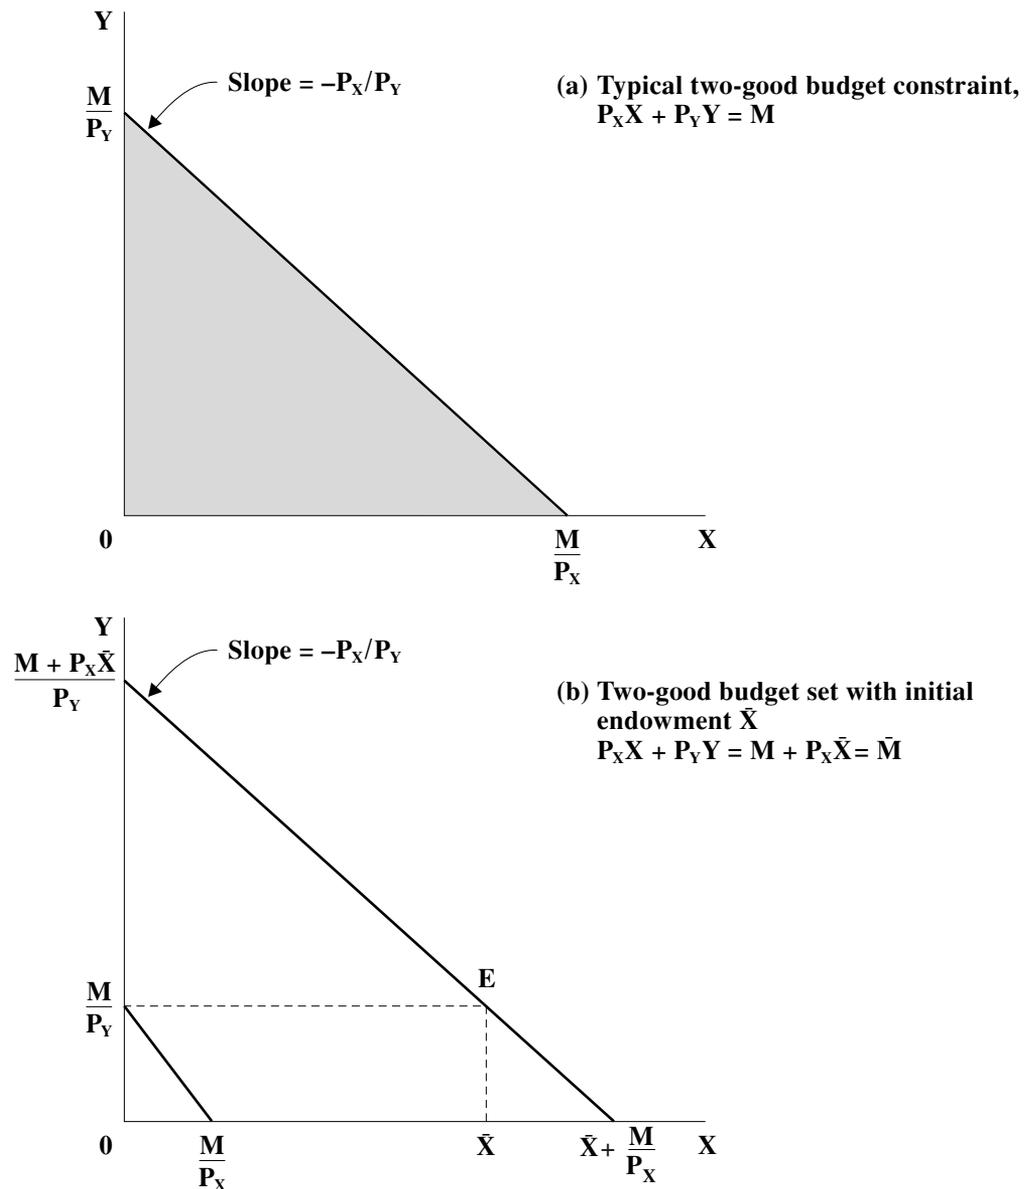
The relative price of X and Y are unaffected, so the slope of this line remains the same. All that has happened to her budget constraint is that it is shifted out by the value of her endowment of X. We can simply relabel her new budget as \bar{M} , and treat this budget constraint as if she had no endowment.

Preferences

The individual's preferences can be expressed formally in terms of a **utility function** that shows the level of satisfaction or utility associated with any specific basket of X and Y. The utility function can be written as $u = U(X, Y)$, where X and Y are the quantities of X and Y consumed. In order for the theory to make some testable predictions about behaviour, it is necessary to make some assumptions about the general nature of the preferences of individuals. The following four fundamental assumptions form the basis of the theory of consumer choice:

Figure 2.A1**Consumer Budget Constraints**

The feasible combinations of X and Y available to the consumer are represented by the shaded area in panel (a). Panel (b) shows how this set is changed when consumers have an initial endowment of \bar{X} , which they are free to consume or sell at P_X . The opportunity set is expanded by the market value of the endowment, $P_X\bar{X}$, while the tradeoff between consuming X and Y remains the relative price of X and Y (P_X/P_Y). An endowment of X thus results in a parallel shift of the budget constraint.



- A1. *Complete ordering.* The individual can rank all possible bundles (combinations of X and Y). This assumption means that for any two combinations of X and Y, say bundle A and bundle B, she can state that A is preferred to B, B is preferred to A, or she is indifferent between A and B. In other words, the individual is assumed to know what she derives satisfaction from and, therefore, to be able to say which bundle would yield higher utility or whether both would yield the same utility.
- A2. *More of X or Y is preferred.* Both X and Y are assumed to be “goods,” not “bads,” in the sense that the more of X consumed (holding constant the amount of Y consumed), the higher the level of satisfaction. Similarly, the more Y consumed (holding constant the consumption of X), the higher her level of utility.

- A3. *Transitivity.* If bundle A is preferred to bundle B and bundle B is preferred to C, then bundle A must be preferred to bundle C. This assumption implies that preferences have a logical consistency similar to that which holds to objects of different weights: if object A is heavier than B, and B weighs more than C, then A is heavier than C.
- A4. *Diminishing marginal rate of substitution.* The individual becomes less willing to substitute X for Y the lower the quantity of X is consumed. More precisely, holding utility constant, the smaller the amount of X consumed, the greater is the amount of Y that would be needed to compensate the individual for one less unit of X.

These four assumptions imply that the individual's utility function has certain properties. Graphically, these assumptions imply that individual preferences can be illustrated in the form of an indifference map as shown in Figure 2.A2(a). Each of the curves labelled U_0 , U_1 , U_2 , and U_3 is called an indifference curve. An indifference curve shows combinations of X and Y that yield the same utility or satisfaction to the consumer. These curves are downward-sloping and have a convex shape. They do not cross or intersect. Indifference curves above and to the right of other indifference curves show combinations of X and Y that are preferred to those on curves lower and to the left (that is, all bundles on U_2 are preferred to any bundle on U_1).

The convex shape of the indifference curves results from the assumption of diminishing marginal rate of substitution. This is illustrated in Figure 2.A2(b). At point A, the individual consumes a large amount of Y (24 units) but has limited X (15 units). Her utility is unchanged if she gives up six units of Y in exchange for one unit of X, thereby moving to point B on the indifference curve. Thus, beginning at the bundle A, she is willing to exchange six units of Y for one additional unit of X. However, beginning at bundle C (which consists of 13 units of Y and 18 units of X), she is only willing to sacrifice one unit of Y for an additional unit of X (i.e., moving from bundle C to bundle D). At point E, which consists of 22 units of X and only eight units of Y, she is only willing to give up a half-unit of Y for one extra unit of X. Thus, the fewer units of Y consumed, the less willing she becomes to give up additional units of Y for extra units of X. As can also be seen in this example, the marginal rate of substitution is equal to the absolute value of the slope of the indifference curve. Thus, the slope of the indifference curve represents the psychic tradeoff that the consumer is willing to make between X and Y and remain equally well off.

The Consumer Optimum and Comparative Statics

The consumer's optimal, utility-maximizing choice of X and Y is given by that combination of X and Y on the budget constraint that lies on the highest indifference curve. This is illustrated in Figure 2.A3. Here, at A, the consumer chooses X^* , Y^* , achieving utility level represented by U^* . Notice that the indifference curve U^* is tangent to the budget constraint at X^* , Y^* . This implies that the slopes of these lines are equal at the optimum, that is:

$$MRS = P_X/P_Y$$

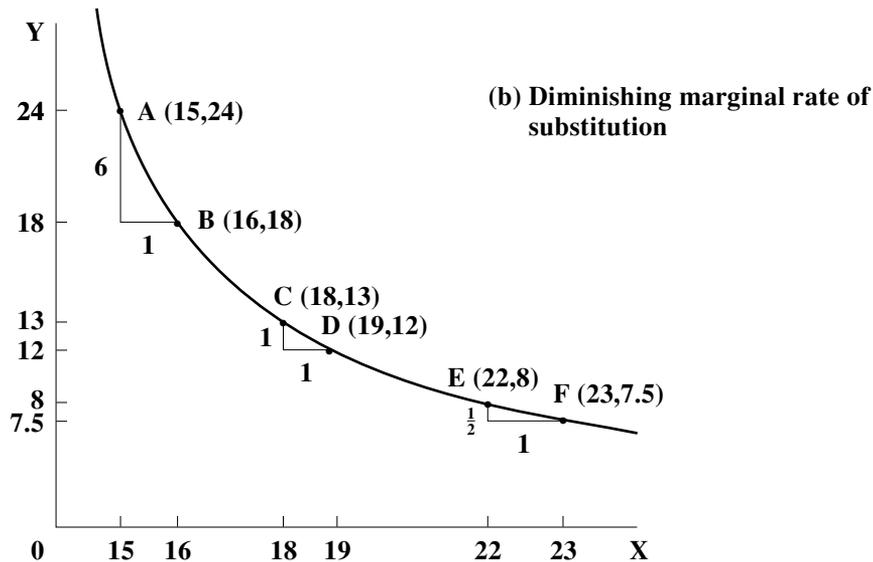
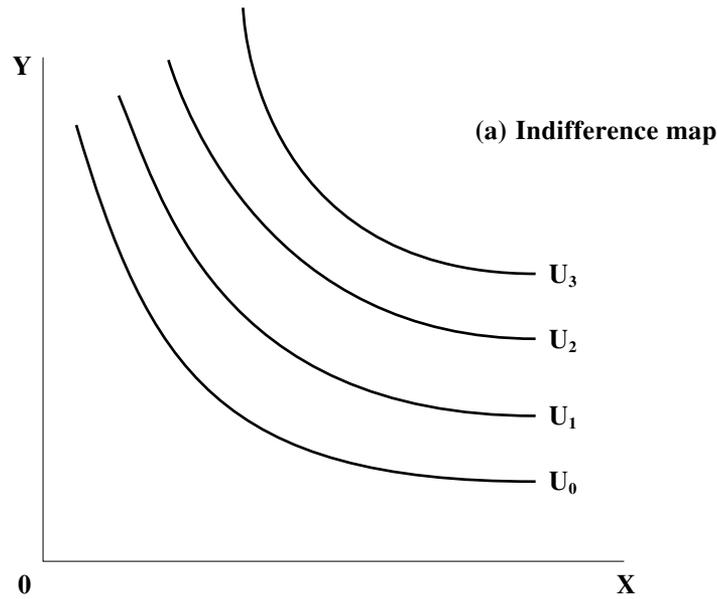
This tangency condition has economic content: it indicates that at the optimum, the individual's psychic (or internal) tradeoff between X and Y equals the external, or market, rate of exchange between X and Y. The reader is advised to consider why this would not be an optimum if this condition did not hold.

The objective of the comparative statics exercise is to now examine what happens to the optimal demand for X if the price of X changes, holding M and P_Y constant. Figure 2.A4(a) shows the case of an increase in P_X to P'_X , and the decomposition of the change in optimal X into income and substitution effects. The increase in P_X leads to an inward rotation of the budget constraint. If the consumer wished to buy only Y, her consumption possibility would be unaffected. On the other hand, if she wanted to buy only X, the maximum she could buy is now reduced. The budget constraint is also now steeper, with slope

Figure 2.A2

Consumer Preferences

Indifference curves provide a graphical representation of consumer preferences, which rank various bundles (combinations) of X and Y. These indifference curves exhibit diminishing marginal rate of substitution. At higher levels of X, the consumer requires less Y to be compensated for a unit reduction in X. At B, where consumption of X is low, the consumer requires 6 units of Y to compensate for a unit reduction in X, whereas at F, the consumer needs only one-half of a unit of Y to compensate for the same reduction of X.

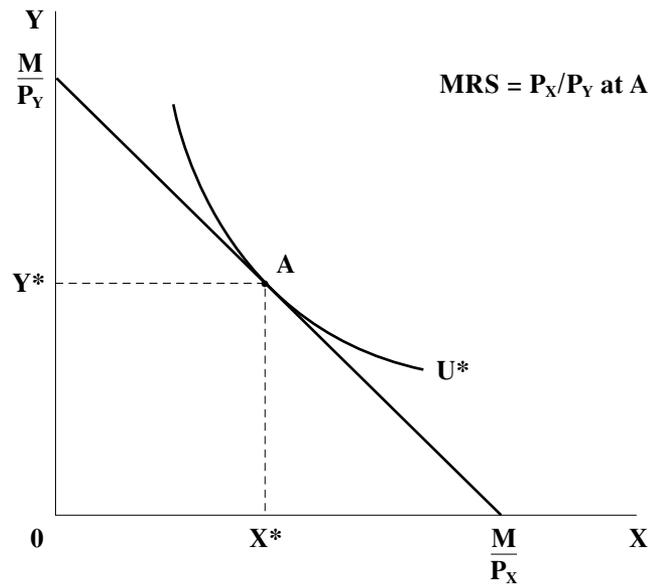


$-P'_X/P_Y$. The optimal consumption bundle moves from A to B. The question we wish to ask is whether this new optimum involves a reduced quantity of X, that is, whether the demand for X is decreasing in P_X .

We can conceptually divide the adjustment into two steps. First, we consider what bundle the consumer would buy if she faced the new relative prices, but had a budget large enough to allow her to attain her original level of utility, U_0 . In this case, her consumption of X would decline from X_0 to X' . The shape of the indifference curve insures that a relative price increase results in substitution away from the now more expensive good. This change, called the substitution effect, predicts an unambiguous decline in the demand for

Figure 2.A3**The Consumer's Optimum**

The consumer's favourite feasible combination of X and Y is that point on the budget constraint that lies on the highest indifference curve. This point A (corresponding to the choice X^* , Y^*) is characterized by the tangency between U^* and the budget constraint. At A, the marginal rate of substitution between X and Y equals the ratio of prices of X and Y.



X, at least holding utility constant. Unfortunately for the consumer, in reality she is not compensated for the price increase. In fact, her budget set has been shrunk by the increase in the cost of X. The budget constraint shifts inward, parallel to the original budget constraint, because only income or wealth, and not relative prices have changed. This decline in “purchasing power” leads to another reallocation of her expenditure, denoted the income effect. If X is a normal good, then a decline in purchasing power (or “income”) will lead to a further reduction in X consumed, leading her to X_1 . Thus, in the case of a normal good, an increase in the price of X leads to an unambiguous decline in the demand of X: the income effect and substitution effect reinforce each other, both leading to a movement away from X.

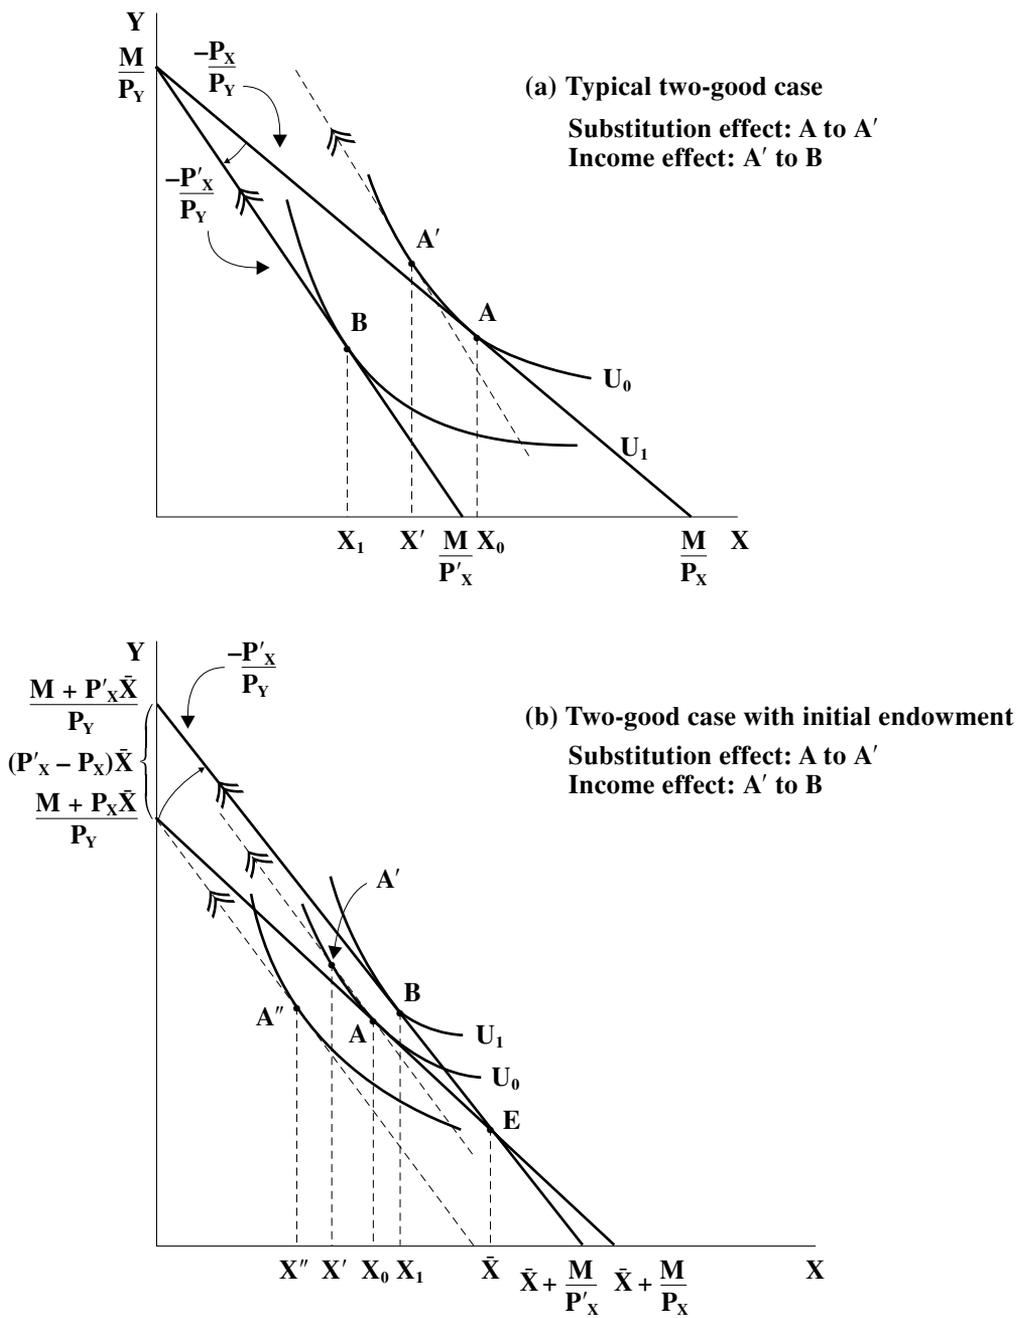
The exercise is complicated somewhat when we consider the case where the individual has an endowment of X. This case is illustrated in Figure 2.A4(b). Again we can decompose the change into income and substitution effects, but the income effect is more complicated by the fact that the consumer has an endowment of X. The increase in P_X results in a rotation of the budget constraint through the point containing \bar{X} (point E). The slope of the budget constraint has changed exactly as it did in the previous case where the consumer had no endowment. The shift of the budget constraint, and the more complicated income effect is most easily seen by examining the extreme points of the budget constraint. If the consumer wished to consume nothing but Y, the increase in the price of X is a pure windfall: her endowment increases in value, and she devotes all of the proceeds to consumption of Y which has not changed in price. This means that her feasible set has actually increased to the left of point E. On the other hand, if she wants to consume only X, she suffers a net decline in purchasing power because X is more expensive. She can still consume her endowment of \bar{X} , but she can purchase fewer units of X from the market.

The substitution effect from the price increase is no different from the previous example, and is given by the movement from A to A'. It unambiguously suggests a decline in the demand for X. We then divide the income effect into two parts. The first is the conventional decline in purchasing power that results from the increase in P_X . We hold constant

Figure 2.A4

The Effect on an Increase in P_X

An increase in P_X has income and substitution effects. As X is more expensive, the consumer will substitute away from X toward Y. This is given by the movement from A to A', where the consumer is allowed to remain equally happy as before the price increase. The income effect results from the reduction of purchasing power, which for a normal good leads to a further reduction in X demanded, from A' to B. In panel (b), there is an additional income effect for consumers who have an additional endowment of X. As X is now worth more, the consumer is richer, and the additional income effect may offset both of the "original" income and substitution effects, leading to a net increase in X consumed.



the person's budget, ignoring the fact that he is "richer" because his endowment is worth more. This involves the shift inward of the budget constraint, and the further decline in X as represented by the movement from A' to A'' (assuming that X is a normal good). To this point, we have the conventional result that income and substitution effects reinforce the decline in demand for X. With the increase in the value of the endowment, however, we

now consider the consequences for demand for X . This endowment “revaluation” results in the shift of the budget constraint by $P_X^1 X - P_X X$, and the movement from $A \leq$ to B . If X is a normal good, this will involve an increase in the demand for X . This effect counteracts the previous income and substitution effects. The net effect is ambiguous, though in the case illustrated, the final income effect was large enough to result in a net increase in the demand for X .

Mapping to Labour Supply

The labour supply model can be seen to be a simple relabelling of the standard consumer model with endowments. In the specific notation of the model in this section, Y can be relabelled as consumption of a composite bundle of goods and services. X can be relabelled leisure, where the individual begins with an endowment of T units of leisure that must be allocated between market or nonmarket time. The price of consumption in this case can be normalized to one, while the price of leisure is the wage rate. M becomes non-labour income, while \bar{M} is full income. The one additional wrinkle in the labour supply problem is that there is an additional constraint: total time use cannot exceed the time endowment of T units of time, so that the budget set is truncated at T units of leisure.

Summary

- Individual attachment to the labour market is measured in two ways: first, whether an individual is working, or searching for work (the participation dimension); and second, how many hours he or she works (the hours dimension).
- The microeconomic model of consumer choice can accommodate both dimensions of the labour supply decision. Consumers choose their preferred combination of income (consumption) and leisure (non-market time), as represented by their opportunity set, or budget constraint. If this optimum occurs at zero hours of work, the individual does not participate. If optimal hours are positive, the individual participates, and the marginal rate of substitution between leisure and consumption equals the wage rate.
- The reservation wage is the critical wage to the participation decision: for wage rates above the reservation wage, the consumer will choose to work; for wages below, the consumer will not participate. The reservation wage is given by the marginal rate of substitution between leisure and consumption, at zero hours of work—that is, the person’s value of non-market time (in terms of consumption) at zero hours worked.
- The consumer choice model can be used to build an individual’s labour supply curve. By varying the wage rate, we can trace out the consumer’s optimal choice of hours worked, holding all other factors constant. For wages below the reservation wage rate, labour supply is zero. The consumer moves to positive hours supplied as the wage exceeds the reservation wage. The impact of increased wages on hours worked will then depend on the relative magnitudes of income and substitution effects. If the substitution effect is largest, wage increases lead to increases in labour supply; if the income effect dominates, wage increases lead to decreases in labour supply.
- The labour-leisure choice model can be used to investigate a variety of labour market phenomena, such as the “added worker effect,” the structure of overtime wage premiums, and the willingness of individuals to accept flexibility in work hours in exchange for lower wages.

REVIEW QUESTIONS

1. What would happen to an individual's labour supply schedule if leisure were an inferior good? Use an income-leisure diagram to illustrate your argument.
2. A labour force survey yields the following estimates:

Population 15 and older	30m
Employed	22m
Not working, but actively seeking work	1m
Full-time students	2m
Retired	3m
Not working, discouraged because of lack of jobs	0.5m
Not working (household workers)	1.5m

Calculate the labour force participation rate and the unemployment rate.
3. Use the basic income-leisure choice framework to analyze the possible labour supply response of various groups to changes in their wage rate. The different groups could include the following:
 - a. The poor who are at a minimum subsistence, and who aspire to middle-class consumption patterns
 - b. The wealthy who have acquired an abundance of material goods and who now aspire to be members of the idle rich
 - c. Workers who have a fairly strong attachment to the labour force and who are reluctant to change their hours of work
 - d. Workers who have a weak attachment to the labour force and who have viable alternatives to market work
 - e. "Workaholics" who have strong preferences for labour market work
4. Illustrate the case where an individual responds differently to a wage increase and a wage decrease of the same magnitude. Specifically, have the person become "locked in" to a certain consumption pattern associated with the higher wage.
5. Use the income-leisure choice model to show that an increase in non-labour income will increase the individual's reservation wage if leisure is a normal good.
6. On the basis of the diagrams of Figure 2.11, illustrate how an underemployed worker would respond to:
 - a. An offer to work as many more hours as the worker would like at the going wage
 - b. Payment of an overtime premium for hours of work beyond C
 - c. An offer to work an additional fixed number of hours, as determined by the employee at the going wage
7. On the basis of the diagrams of Figure 2.12, illustrate how an overemployed worker would respond to:
 - a. An offer to work as many hours as the worker would like at the going wage
 - b. Payment of the moonlighting rate for hours of work beyond C
8. On the basis of Figure 2.12(b), precisely illustrate the following overtime rates for hours worked beyond TL_C :
 - a. Time-and-a-half
 - b. Double-time
 - c. Time-and-a-half for the first two hours of overtime, and double-time thereafter

PROBLEMS

1. Assume that the following regression equation has been estimated, where P_W is the labour force participation rate of married women (measured as a percentage with

average $\bar{P} = 35.0$), Y_H is husband's wage income (measured in thousands of dollars, with average $\bar{Y}_H = 10$), Y_W is the wife's expected "wage" (expected income from working a fixed number of hours, measured in thousands of dollars, with average $\bar{Y}_W = 6$), and u_H is the male unemployment rate (measured as a percentage, with average $\bar{u}_H = 6.0$):

$$P_W = -7Y_H + 18Y_W - 0.5u_H$$

- a. What is the expected effect of an increase of \$1000 in the income of the husbands on the participation rate of their wives?
 - b. What is the expected effect of an increase of \$1000 in the wage of the wives themselves?
 - c. Break the latter down into its separate income and substitution effects.
 - d. Given the magnitude of the latter two effects, what would be the impact on female participation of an equal pay policy that increased the expected wages of females by \$1000 while at the same time decreasing the expected earnings of their husbands by \$1000?
 - e. Calculate the pure income and the gross or uncompensated wage elasticities of participation, evaluated at the means.
 - f. Does this equation shed any light on why the labour force participation of married women has increased over time, even though their non-labour wage has also increased?
2. Consider two individuals with endowments of $T = 60$ hours (per week) of leisure, non-labour income of Y , and facing a wage of \$7.50 per hour. At this wage, assume that workers are constrained by their employers to work 40 hours per week, or not at all.
 - a. On a carefully labelled diagram, show the equilibrium for a worker for whom 40 hours is the optimum labour supply; and a worker who would like to work 20 hours, but still prefers the 40-hour week to not working. Compare the marginal rates of substitution for these individuals at 40 hours per week.
 - b. The average part-time wage is \$7 per hour, in contrast to \$7.50 wage for full-time workers. Using the above model, provide an explanation for this difference in wage rates.
 3. Assume that women's marginal rate of substitution of leisure for consumption is given by the following function:

$$MRS = A(x) \frac{C}{\ell}$$

where C is consumption (with price = 1) and ℓ is leisure. $A(x)$ is a taste-shifting function of the following form:

$$A(x) = \exp(\beta_1 x_1 + \beta_2 x_2 + \dots + \beta_K x_K + \varepsilon)$$

where the x_j are K different observable factors that affect her preferences, and ε represents unobservable factors that affect the MRS. ε has the property that it can be described as a probability density function, such as the normal or uniform densities.

- a. Show that this functional form for the MRS represents preferences that exhibit a diminishing marginal rate of substitution. Give specific examples of x_j that may affect the MRS, and explain how they will affect the MRS. Give a graphical example.
- b. Assume that an individual woman has non-labour income y (including her husband's income) and a time endowment of T . Derive an expression for this woman's reservation wage, w^* . Show that if the market wage is w , this expression implies that a woman will participate if

$$\ln w > \ln y - \ln T + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_K x_K + \varepsilon$$

Rewrite this as an expression of the form: Participate if $\varepsilon < Z$, where Z depends on the market wage, non-labour income, and preferences. *Hint:* Remember $\ln(ab) = \ln a + \ln b$; $\ln(a/b) = \ln a - \ln b$, etc.

- c. Assume that Z has a standard normal distribution. Using your results from part (b), graphically show and explain how non-labour income, the market wage, and the various taste shifters affect the probability of a woman participating in the labour market.
4. Using carefully labelled diagrams, indicate the expected impact on the labour force participation of married women of changes in each of the following factors, other things held constant:
 - a. An increase in the education of women
 - b. A more equal sharing of household responsibilities between husband and wife
 - c. A reduction in the average number of children
 - d. An increased tendency to have children spaced more closely together
 - e. An increase in the earnings of husbands
 - f. Daycare paid out of general tax revenues
 - g. Allowing daycare expenses to be tax-deductible
 - h. Paying housewives a fixed sum out of general tax revenues for household work
5. Using the income-leisure framework, formally analyze the “added worker” effect, that is, the impact on the wife’s labour supply of an adverse shock to her husband’s job. For concreteness, assume that the wife has a wage rate of \$10 per hour, and that her husband’s employer has forced him to take a pay cut from \$20 per hour to \$15 per hour, but allowed him to continue working at 40 hours per week (if he wishes). Consider two approaches:
 - a. In the first, take the husband’s hours as given, and analyze the household choice over consumption and wife’s leisure.
 - b. In the second, focus on the choice of husband versus wife’s labour supply, allowing both to adjust their hours, and putting the consumption decision in the background.
 - c. Compare the pros and cons of the two approaches. In particular, use the two approaches to compare the impact on the wife’s labour supply of unemployment compensation for the husband equal to \$200 per week.
6. Susan claims labour supply theory is nonsense. She determines how much income she needs to support her “addiction” to maintaining and insuring her 1967 Mustang convertible. She then works as many hours as necessary. “No crazy income and substitution effects for me,” she asserts. Is she right? Depict her labour supply choice in an income-leisure diagram, and break a wage increase down into its constituent income and substitution effects.
7. Curious George must decide how much to work. He has 60 hours per week available that he can spend either working or engaged in leisure (which for him is creating various kinds of mischief). He can work at a wage rate of \$5 per hour. The Man with the Yellow Hat (who looks after George) also gives him an allowance of \$100 per week, no matter how much George works. George’s only source of income that he can use for consumption (mostly bananas) is this allowance plus his wage earnings.
 - a. In a carefully labelled diagram, draw George’s consumption-leisure budget constraint. Show an equilibrium where George chooses to work 40 hours per week.
 - b. In an effort to have George pay for other household expenses, the Man with the Yellow Hat decides to tax George 50 percent of his *wage income*. Using the same diagram you drew in part (a), where George works 40 hours, show what happens to his labour supply. To do this, show one possible outcome, and break the change down into income and substitution effects.

- c. Instead of the wage tax, the Man with the Yellow Hat could impose a “poll tax,” a lump-sum tax independent of George’s wage earnings. This tax must raise the same revenue as the wage tax in part (b), and could be accomplished by reducing George’s allowance. Draw the budget constraint with the new tax *and* the wage tax, and compare the work incentive effects of the poll tax to the wage tax in part (b). As in part (b), assume that George was working 40 hours before any taxes were imposed.

KEYWORDS

labour force participation decision	33	corner solution	42
labour force	33	interior solution	42
employed	33	reservation wage	42
unemployed	33	normal good	44
labour force participation rate	33	inferior good	44
hours-of-work aspect	35	income effect	45
labour supply model	38	substitution effect	45
preferences	38	labour supply schedule	49
leisure	38	uncompensated elasticity	53
indifference curves	38	income elasticity	53
marginal rate of substitution	38	compensated elasticity	53
utility	38	discouraged worker	56
nonmarket time	40	added worker	56
consumption-leisure	40	hidden unemployment	57
income-leisure	40	underemployed	60
potential income constraint	40	moonlighting	60
budget constraint	40	overemployed	60
full income	41	overtime premium	60
consumer’s optimum	42	utility function	65
utility-maximizing	42		

REFERENCES

- Benimadhu, P. 1987. *Hours of Work: Trends and Attitudes in Canada*. Ottawa: Conference Board.
- Biddle, J., and D. Hamermesh. 1990. Sleep and the allocation of time. *JPE* 98 (October):922-43.
- Blundell, R., and T. MaCurdy. 1999. Labor supply: A review of alternative approaches. In *Handbook of Labor Economics*, eds. O. Ashenfelter and D. Card. New York and Oxford: Elsevier Science, North Holland.
- Fortin, N. M. 1995. Allocation inflexibilities, female labor supply, and housing assets accumulation: Are women working to pay the mortgage? *JOLE* 13 (July):524-57.
- Hansson, I., and C. Stuart. 1985. Tax revenue, and the marginal cost of public funds in Sweden. *JPubEc* 27 (August):333-53.
- Jones, S.R.G., and W.C. Riddell. 1999. The measurement of unemployment: An empirical approach. *Ecta* 67 (January):147-61.
- Killingsworth, M. 1983. *Labor Supply*. Cambridge: Cambridge University Press.
- Killingsworth, M., and J. Heckman. 1986. Female labor supply: a survey. In *Handbook of Labor Economics*. eds. O. Ashenfelter and R. Layard. New York: Elsevier.
- Labour Canada. (1974). *Trends in working time*. Ottawa: Wages Research Division, Economics and Research Branch.
- Nakamura, A., and M. Nakamura. 1981. A comparison of the labour force behavior of married women in the United States and Canada, with special attention to the impact of income taxes. *Ecta* 49 (March):451-89.
- Nakamura, A., and M. Nakamura. 1985. A survey of research on the work behavior of Canadian women. In *Work and Pay: The Canadian Labour Market*, ed. W.C. Riddell. Toronto: University of Toronto Press.
- Nakamura, A., M. Nakamura, and D. Cullen. 1979. Job opportunities, the offered wage, and the labour supply of married women. *AER* 69 (December):785-805.
- Oettinger, G. S. 1999. An empirical analysis of the daily labor supply of stadium vendors. *JPE* 107 (April):360-92.
- Pencavel, J. 1986. Labor supply of men: a survey. In *Handbook of Labor Economics*. Vol. 1, eds. O. Ashenfelter and R. Layard. New York: Elsevier.
- Robbins, L. 1930. On the elasticity of demand for income in terms of effort. *Economica* (June):123-29.
- Robinson, C., and N. Tomes. 1985. More on the labour supply of Canadian women. *CJE* 18 (February):156-63.
- Smith, A. 1776. *The Wealth of Nations*. London: Methuen and Company.