CHAPTER 1

Labor Supply

Contents

In this chapter we will see:

- How people make choices between consumption, leisure, and household production
- What the reservation wage is
- How the shape of the labor supply curve results from the combination of substitution effects and income
- When and why people decide to retire
- The principles guiding the econometrics of labor supply and the main empirical results
- Examples of natural experiments

INTRODUCTION

To hold a paid job, you must first have decided to do so. This is the starting point of the so-called ''neoclassical'' theory of the labor supply. It posits that each individual

disposes of a limited amount of time, which he or she chooses to allocate between paid work and leisure. Evidently the wage an individual can demand constitutes an important factor in the choice of the quantity of labor supplied. But it is not the only factor taken into account. Personal wealth, income derived from sources outside the labor market, and even the familial environment also play a decisive role.

In reality the allocation of one's time depends on trade-offs more complex than a simple choice between work and leisure. In the first place, the counterpart of paid work is not simply leisure in the usual sense, for much of it consists of time devoted to ''household production'' (the preparation of meals, housekeeping, minor repairs and upkeep, the raising of children, etc.), the result of which substitutes for products available in the consumer goods market. This implies that the supply of wage labor takes into account the costs and benefits of this household production, and that most often it is the result of planning, and even actual negotiation, within the family. The family situation, the number of children, the income a person enjoys apart from any wage labor (personal wealth, illegal work, spousal income, etc.), all weigh heavily in this choice. Decisions concerning labor supply also depend on trade-offs over the course of time that make the analysis of the decisions of agents richer and more complex.

Empirical studies on labor supply have also multiplied in the course of the last twenty years. The development of these studies—exhaustively reviewed in Blundell and MaCurdy (1999)—has profited from advances made in the application of econometric methods to individual data, and from a desire to evaluate public policies that attempt to influence labor supply directly. A number of countries have set up programs explicitly aimed at increasing labor supply among the most disadvantaged, rather than park them on the welfare rolls. These ''welfare to work'' programs, sometimes abbreviated as *workfare*, so as to contrast them with more traditional programs called simply *welfare*, have given a powerful incentive to empirical research on labor supply in the United States and Great Britain, as well as in certain continental countries like Sweden and France.

The first section of this chapter lays out the principal elements of the neoclassical theory of labor supply. This approach is based on the traditional microeconomic model of consumer choice. The basic model explains the choice between the consumption of products available in the marketplace and leisure. This simple model is then extended in such a way as to take into account household production and intrafamilial decisions. The basic model is also enhanced into a ''life-cycle'' model integrating the decisions taken by agents over the course of time. This enhancement is particularly important from the point of view of economic policy, for most employment policy measures aim to modify the behavior of agents permanently. It also furnishes an adequate framework for analyzing decisions taken from the onset of a career to retirement. The second section of this chapter is devoted to empirical matters. It begins by laying out the main lines of the econometrics of labor supply, elucidates the principles that guide empirical studies in this area, and concludes with a review of the principal quantitative results arrived at by studies of labor supply.

1 THE NEOCLASSICAL THEORY OF LABOR SUPPLY

The theory of labor supply is based on the model of a consumer making a choice between consuming more goods and consuming more leisure. With it, we can elucidate the properties of labor supply and begin to understand the conditions of participation in the labor market. The model has been variously enhanced to make the theory of labor supply more precise, and sometimes to modify it profoundly, principally by taking into account household production, the collective dimension of decisions about labor supply (most often within the family), and the life-cycle aspect of these decisions.

1.1 The Choice Between Consumption and Leisure

The basic model of a trade-off between consumption and leisure gives us the principal properties of the supply of labor. In particular, it shows that labor supply is not necessarily a monotonic function of wages. It suggests that labor supply grows when the wage is low, and subsequently diminishes with the wage when the latter is sufficiently high. Further, the study of the trade-off between consumption and leisure makes it possible to grasp the factors that determine participation in the labor market.

1.1.1 The Basic Model

We indicated, in the general introduction to this chapter, that the traditional approach to labor supply arises, fundamentally, out of the idea that each of us has the possibility to make trade-offs between the consumption of goods and the consumption of leisure, this last being defined as time not spent at work. The analysis of this choice makes it possible to pinpoint the factors that determine labor supply, first at the individual, then at the aggregate, levels.

Preferences

The trade-off between consumption and leisure is shown with the help of a utility function proper to each individual, that is, $U(C, L)$, where C and L designate respectively the consumption of goods and the consumption of leisure. Given that an individual disposes of a total amount of time, L_0, the length of time worked, expressed, for example, in hours, is then given by $h = L_0 - L$. It is generally supposed that an individual desires to consume the greatest possible quantity of goods and leisure; his or her utility function therefore increases with each argument. Moreover, the same individual is capable of attaining the same level of satisfaction with much leisure and few goods, or little leisure and many goods. The set of pairs (C, L) by which the consumer obtains the same level of utility $\overline{U}$, i.e., such that $U(C, L) = \overline{U}$, is called an *indifference curve*. A curve of this type is shown in figure 1.1. Its properties follow directly from those of the utility function (for more detail, consult Varian, 1992, and Mas-Colell et al., 1995). In particular, the properties listed below will be useful for what follows:

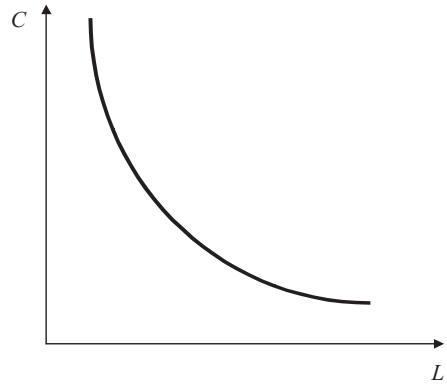

Figure 1.1
An indifference curve.

(i) Each indifference curve corresponds to a higher level of utility, the farther out the curve is from the origin. Hence the consumer will prefer indifference curves situated farther out from the origin.

(ii) Indifference curves do not intersect. If they did, the point of intersection would correspond to a combination of leisure and consumption through which the individual would have two different levels of satisfaction. Incoherence in preferences of this kind is excluded.

(iii) The increase in the utility function in relation to each of its components implies that the indifference curves are negatively sloped (see appendix 1 at the end of this chapter). The slope of an indifference curve at a given point defines *the marginal rate of substitution between consumption and leisure*. It represents the quantity of goods which a consumer must renounce in exchange for an hour of supplementary leisure, for his or her level of satisfaction to remain unchanged.

(iv) It is assumed that the individual is ready to sacrifice less and less consumption for an extra hour of leisure when the amount of time dedicated to leisure rises. This property signifies that the marginal rate of substitution between consumption and leisure diminishes with leisure time, or again that the indifference curves are convex, which is equivalent to the hypothesis of the quasi-concavity of the utility function (the relation between the shape of the indifference curves and the utility function is studied in appendix 1 at the end of this chapter).

Choices

An individual's income derives from his or her activity as wage-earner and from his or her activity (or inactivity) outside the labor market. If we designate the real hourly wage by w, the income from wages totals wh. Investment income, transfer income, even gains deriving from undeclared or illegal activities are examples of what an individual may acquire outside the labor market. We will designate the set of these resources expressed in real terms by the single scalar R.

Note that for a married or cohabiting person, a part of the income of his or her partner is capable of being integrated into this set. Thus the budget constraint of the agent takes the form:

$$C \leq wh + R$$

This constraint is also expressed in the following manner:

$$C + wL \leq R_0 \equiv wL_0 + R \tag{1}$$

In this way we arrive at the standard concepts of the theory of the consumer. The fiction is that the agent disposes of a potential income R_0 obtained by dedicating his entire endowment of time to working, and that he or she buys leisure and consumer goods using this income. From this point of view, the wage appears to correspond equally to the *price* and the *opportunity cost* of leisure. The solution of the consumer's problem then follows the path of utility optimization subject to the budget constraint. We thus derive the functions of demand for consumer goods and leisure (for more details, see the microeconomics textbooks by, for example, Varian, 1992; Mas-Colell et al., 1995). The decision of the consumer is expressed:

$$\underset{\{C,L\}}{\text{Max}}\ U(C,L) \quad \text{subject to the budget constraint} \quad C + wL \leq R_0$$

We begin by studying the so-called "interior" solutions, such as $0 < L < L_0$ and $C > 0$.

The Interior Solutions

For an interior solution, the consumer puts forth a strictly positive supply of labor. Using $\mu \geq 0$ to denote the Lagrange (or Kuhn and Tucker) multiplier associated with the budget constraint, the Lagrangian of this program is[1]:

$$\mathscr{L}(C,L,\mu) = U(C,L) + \mu(R_0 - C - wL)$$

Designating the partial derivatives of the function U by U_L and U_C, the first-order conditions are expressed as:

$$U_C(C,L) - \mu = 0 \qquad \text{and} \qquad U_L(C,L) - \mu w = 0$$

On the other hand, the complementary-slackness condition is expressed as:

$$\mu(R_0 - C - wL) = 0 \qquad \text{with} \qquad \mu \geq 0$$

This relation, and the hypothesis that the utility function increases with each of its components, imply that the budget constraint is binding, since the first first-order condition is equivalent to $\mu = U_C(C,L) > 0$. Thus, the solution is situated on the budget line of equation $C + wL = R_0$. We obtain the optimal solution (C^*, L^*) by using this last equality and eliminating the Kuhn and Tucker multiplier μ of the first-order conditions, so that:

$$\frac{U_L(C^*,L^*)}{U_C(C^*,L^*)} = w \qquad \text{and} \qquad C^* + wL^* = R_0 \tag{2}$$

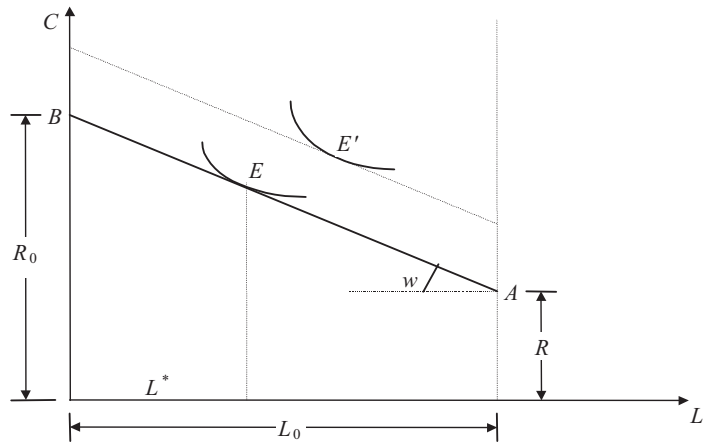

FIGURE 1.2
The trade-off between consumption and leisure.

Figure 1.2 proposes a graphic representation of this solution. It shows that the optimal solution is situated at a tangency point between the budget line *AB*, whose slope is *w*, and the indifference curve corresponding to the level of utility obtained by the consumer. For the comparative statics of the model, it is worth noting that any increase in *w* results in a clockwise rotation of the line *AB* around point *A*, of abscissa L_0, and of ordinate *R*, and that a rise in non-wage income corresponds to an upward shift of this budget line.

The Reservation Wage

For relation (2) actually to describe the optimal solution of the consumer's problem, point *E* has to lie to the left of point *A*, otherwise labor supply is null ($L = L_0$). Now, the convexity of indifference curves implies that the marginal rate of substitution between consumption and leisure, U_L/U_C, decreases as one moves to the southeast along an indifference curve (see appendix 1 at the end of this chapter).

Since this marginal rate of substitution also represents the slope of the tangent to an indifference curve, an agent offers a strictly positive quantity of hours of work if and only if the following condition is met:

$$\left(\frac{U_L}{U_C}\right)_A < w$$

The marginal rate of substitution at point *A* is called the *reservation wage*. It is thus defined by:

$$w_A = \frac{U_L(R, L_0)}{U_C(R, L_0)} \tag{3}$$

According to this model, assuming that the allocation of time L_0 designates a fixed quantity, the reservation wage depends only on the form of the function *U* at

point A and on the value R of non-wage income. It determines the conditions of *participation* in the labor market. If the current wage falls below it, the agent does not supply any hours of work; we then say that he or she is not participating in the labor market. The decision to participate in the labor market thus depends on the reservation wage. Hence its determinants deserve special attention. In this model, setting aside any change in the consumer's tastes, the only parameter capable of modifying the reservation wage is non-wage income R. If, with respect to this last variable, we derive the relation (3) that defines the reservation wage, we can easily verify that the latter rises with R if, and only if, leisure is a *normal*[2] good (one, that is, the consumption of which increases with a rise in income). Under these conditions, an increase in non-wage income increases the reservation wage, and thus has a *disincentive* effect on entry into the labor market.

1.1.2 The Properties of Labor Supply

The properties of the supply of individual labor result from the combination of a substitution effect and income effect. The combination of these effects seemingly leads to a nonmonotonic relation between wages and the individual supply of labor. We shall see as well that, by starting with individual decisions and taking into account the heterogeneity of individuals, we will be able to grasp the factors that determine the aggregate supply of labor.

Substitution Effect and Income Effect

For an interior solution, the demand for leisure L^* is implicitly defined by relations (2). It is a function of the parameters of the model, which can conveniently be written in the form $L^* = \Lambda(w, R_0)$. The corresponding labor supply, i.e., $h^* = L_0 - L^*$, is often called the "Marshallian" or "uncompensated" labor supply. The impact of an increase in non-wage income R on time given over to leisure is indicated by the partial derivative of the function $\Lambda(w, R_0)$ with respect to its second argument, i.e., $\Lambda_2(w, R_0)$. It may be positive or negative. By definition, leisure is a *normal good* if its demand rises with R_0 (see appendix 2 to this chapter). In the opposite case, in which the time dedicated to leisure decreases with non-wage income, leisure is an *inferior good*. The consequences of an increase in non-wage income are represented in figure 1.2 by the shift from point E to point E'.

The impact of a variation in wages is obtained by differentiating function $\Lambda(w, R_0)$ with respect to w. Taking account of the fact that $R_0 = wL_0 + R$, we arrive at:

$$\frac{dL^*}{dw} = \Lambda_1 + \Lambda_2 \frac{\partial R_0}{\partial w} \qquad \text{with} \qquad \frac{\partial R_0}{\partial w} = L_0 > 0 \tag{4}$$

Figure 1.3 traces the movement of the consumer's equilibrium when wages go from a value of w to a value of $w_1 > w$. The partial derivative of the function Λ with respect to w, denoted Λ_1, corresponds to the usual compound of substitution and income effects in the theory of the consumer (the calculations are presented in appendix 2). To learn the sign of this derivative, it is best to reason in two stages. In the first stage, we suppose that the potential income R_0 does not change: the

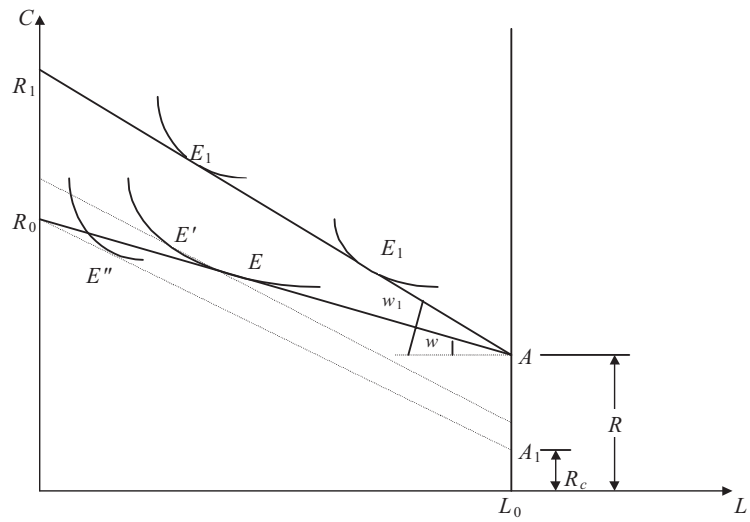

Figure 1.3
The effects of a wage increase.

consumer then faces a new budget line A_1R_0. For him or her, it is as though his or her non-wage income had decreased from R to $R_c = R - (w_1 - w)L_0$. Income R_c is described as *compensated* income and the line A_1R_0 is called the *compensated* budget constraint. In the second stage, we assume that the potential income grows from R_0 to $R_1 = R + w_1L_0$.

Reckoning first with R_0 as a given, we discover the usual compound of substitution and income effects of the theory of the consumer. When the initial equilibrium lies at point E, the substitution effect moves it to point E' offering the same degree of utility as at E, but with the wage now worth w_1 (at point E' the tangent to the indifference curve is parallel to the budget line A_1R_0). The shift from point E to point E' corresponds to a ''Hicksian'' or ''compensated'' modification of the labor supply, obtained by minimizing the outlay of the consumer under the constraint of reaching a given level of utility. The substitution effect thus implies a reduction of leisure. Starting from point E', and assuming that the wage keeps the value w_1, the income effect shifts the equilibrium of the consumer to point E''. If leisure is a normal good, the shift from E' to E'' being the consequence of a fall in income, the demand for leisure must diminish. Thus, the substitution effect and the (indirect) income effect work to produce the same result: an increase in wage leads to a diminution of the time allotted to leisure, or in other words, to an increase in labor supply. Consequently, in relation (4), we will have $\Lambda_1 < 0$ if leisure is a normal good. Finally, the increase in potential income from R_0 to R_1 causes the equilibrium to shift from point E'' to point E_1. What we have is a *direct* income effect identified by the partial derivative Λ_2 of the demand for leisure with respect to R_0 in relation (4). If leisure is a normal good, then by defi-

nition Λ_2 is positive and any rise in wage leads to a rise in the consumption of leisure, and thus to a fall in labor supply. This direct income effect runs counter to the usual substitution and ''indirect'' income effects of the theory of the consumer. In sum, a wage increase has an ambivalent effect on labor supply. In figure 1.3 the abscissa of point E_1 can as easily lie to the left as to the right of that of E.

For convenience, we can aggregate the two income effects by retaining only the shift from E' to E, in which case we refer to the global income effect. This allows us to analyze a rise in the hourly wage with the help of only two effects. In the first place, there is an incentive to increase labor supply, since this factor is better remunerated (the substitution effect). But equally there is an opportunity to consume the same quantity of goods while working less, which motivates a diminution of labor supply (the global income effect) if leisure is a normal good.

Compensated and Noncompensated Elasticity of Labor Supply

Along with the Marshallian supply of labor h^* considered to this point, we can also make use of the Hicksian supply of labor; it is arrived at by minimizing the consumer's expenditure, given an exogenous minimal level of utility $\bar{U}$. The Hicksian supply of labor, denoted $\hat{h}$, is then the solution of the problem:

$$\underset{(L,C)}{\text{Min}}\ C + wL \quad \text{subject to constraint} \quad U(C,L) \geq \bar{U}$$

The Marshallian supply depends on the wage and on non-wage income, whereas the Hicksian supply of labor depends on the wage and on the level of utility $\bar{U}$. The Hicksian elasticity of the labor supply, defined by $\eta_w^{\hat{h}} = (w/\hat{h})(d\hat{h}/dw)$, represents the percentage of variation of the Hicksian supply of labor that follows from a 1% rise in wage. It corresponds to the variation in labor supply for a shift from point E to point E' in figure 1.3. Hicksian elasticity is called ''compensated'' elasticity because it posits that the income of the consumer varies in order for him to stay on the same indifference curve. The Marshallian elasticity of labor supply, defined by $\eta_w^{h^*} = (w/h^*)(dh^*/dw)$, represents the percentage of variation of the Marshallian supply of labor that follows from a 1% rise in wage. It corresponds to the variation in the labor supply for a shift from point E to point E_1 in figure 1.3. Marshallian elasticity is also called noncompensated elasticity because it takes into account the real variation in income resulting from the variation in wages.

Marshallian and Hicksian elasticities are linked by the Slutsky equation, which is written thus:

$$\eta_w^{h^*} = \eta_w^{\hat{h}} + \frac{wh^*}{R_0}\eta_{R_0}^{h^*}$$

A demonstration of this equality is presented in appendix 3 at the end of this chapter. The Slutsky equation shows that Marshallian elasticity is to be interpreted as the sum of two effects. The substitution effect, represented by the Hicksian elasticity $\eta_w^{\hat{h}}$, is necessarily negative. The (global) income effect, represented by the term $(wh^*/R_0)\eta_{R_0}^{h^*}$, is positive if leisure is a normal good.

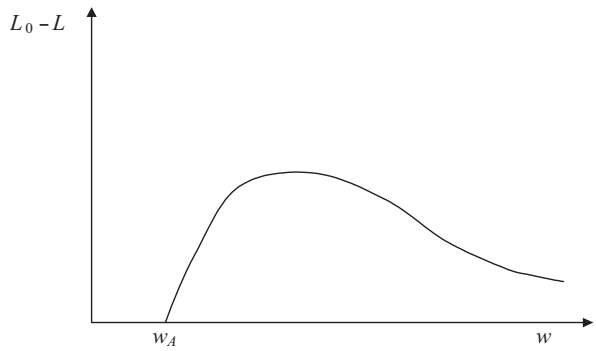

Figure 1.4
The individual labor supply.

The Shape of the Labor Supply Curve
We can now offer a plausible graph of labor supply. It is shown in figure 1.4. When the hourly wage rises just above the reservation wage, the substitution effect prevails over income effects, and labor supply grows. But the global income effect swells with the wage, and it is reasonable to believe that when the latter reaches a certain level, it will dominate the substitution effect. The supply of labor then begins to shrink. This is the reason why it is generally thought to turn down, as shown in figure 1.4.

Supplementary Constraints
The preceding analysis leaves out many elements that may play a part in the trade-off between work and leisure. For example, the budget constraint is actually piecewise linear, since on the one hand, overtime hours are not remunerated at the same rate as normal ones, and on the other hand income tax is progressive. This constraint may even present nonconvexities related to the ceilings on various social security contributions. Neither does the model hitherto presented take into account the fact that most often the decision to take a job entails a fixed cost independent of the number of hours worked, such as, for example, the purchase of a second vehicle, or the cost of child care. All these elements pose serious problems for empirical assessment (see below, section 2.1.3).

Another element that may alter the foregoing analysis comes from the relative absence of freedom of choice in the number of hours worked. The majority of wage-earners hold full-time employment, other workers hold part-time jobs, but the reality is always a far cry from a hypothetical complete flexibility in hours worked. To illustrate the effects of a rigidity constraint on hours worked, we present a situation in figure 1.5 in which the agent has a choice between working during a set period, represented by the abscissa point L_f, or not working at all.

Let us designate by E the nonconstrained optimum of the problem of the agent. If this point is situated to the left of E_f, the agent agrees to furnish $(L_0 - L_f)$ hours of

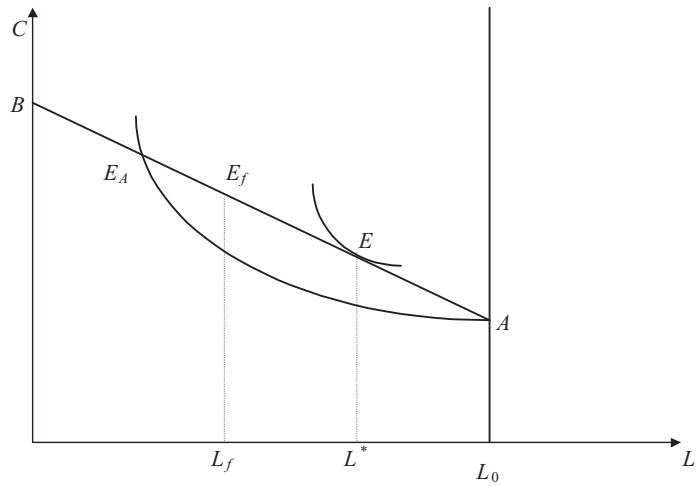

Figure 1.5
Constraint on hours of work.

work; in this situation, he or she would simply have liked to work more. Vice versa, when the point E lies to the right of E_f, he or she agrees to work the quantity of fixed hours offered if, and only if, the point E_A—corresponding to the intersection of the indifference curve passing through A with the budget line—lies to the left of E_f. In this case, he or she obtains a level of utility superior to what he or she would have attained by not participating at all in the labor market. The agent then works more than he or she would have wished to (since $L^* > L_f$). On the other hand, if the point E_A were to lie to the right of E_f, he or she would choose not to participate, since he or she would have preferred to supply $(L_0 - L^*) > 0$ hours of work. This individual is in a situation that we can call ''involuntary nonparticipation,'' since he or she does wish to supply a certain quantity of work at the current wage and faces constraints that keep him or her from supplying them. The abscissa and the ordinate of point E_f being equal respectively to L_f and $w(L_0 - L_f) + R$, the reservation wage—which we will still denote by w_A—is defined by the equality:

$$U[R + w_A(L_0 - L_f), L_f] = U(R, L_0)$$

Aggregate Labor Supply and the Labor Force Participaton Rate

We arrive at the aggregate labor supply, for a wage level of w, by adding up the total number of hours supplied by each individual. The existence of indivisibilities in the amounts of working hours offered to agents implies that the elasticity of the aggregate supply differs from that of the individual supply. To show this result, let us take the case envisaged previously, in which each agent has the choice between working for a fixed length of time $\bar{h} = L_0 - L_f$ and not working at all. In a population of large size, the reservation wages differ from one individual to another, for preferences and

non-wage incomes are not identical. Let us imagine that this diversity of reservation wages $w_A \in [0, +\infty)$ may be represented by a cumulative distribution function $\Phi(\cdot)$. By definition, the quantity $\Phi(w)$ represents the *participation rate*, that is, the proportion of individuals in the population of working age whose reservation wage is below the current wage w. Since the function Φ is increasing, the participation rate climbs as the wage increases. If the size of the total population is N, the quantity $N\Phi(w)$ represents the labor force and the aggregate labor supply is equal to $\Phi(w)$. Supposing that the size of the population N does not vary, the wage elasticity of the aggregate supply of labor is identical to that of the participation rate. This elasticity is positive, since a rise in wages draws workers into the labor market.

This result extends far beyond the example given; it is confirmed whenever the hours offered to workers are not entirely flexible. From an empirical point of view this result has a certain importance, since it implies that the aggregate supply of labor or the global supply of labor of a subpopulation may be sensitive to changes in the wage, even if the labor supplied by most of the individual agents is not. We shall discover below that the elasticity of the individual's supply of labor is indeed slight, but that decisions to participate in the labor market turn out to be extremely sensitive to the various incentives, particularly fiscal ones, that suppliers of labor are faced with. In this case the total aggregate supply, or the supply of a given subpopulation, ought to follow the fluctuations in the participation rate (a point emphasized particularly by Heckman, 1993).

1.2 Labor Supply with Household Production and Within the Family

The basic model of a trade-off between consumption and leisure neglects numerous elements that may influence labor supply. In this subsection we extend the model in two important directions. By allocating time not dedicated to wage labor to leisure, the basic model fails to take account of production within households—production that represents a substitute for wage income from work. Furthermore, decisions about labor supply frequently result from bargaining involving several members of the household.

1.2.1 Household Production

The dichotomy between leisure and wage labor masks an important part of the complexity of individual decisions concerning the allocation of time. In reality, leisure is not the sole alternative to wage labor. Time devoted to household tasks is (generally) distinguished from leisure. Now, these tasks are not always unavoidable. The bulk of the goods and services produced domestically can be purchased. It is possible, for example, to eat a meal that one has prepared oneself, or go to a restaurant, or telephone a caterer, or hire a cook. Clearly each alternative entails a different expense, and an individual's choice depends on his or her preferences, effectiveness at performing household chores versus doing paid work, income, and prices. We can ana-

lyze the consequences of time devoted to household production by modifying our basic model of labor supply at the margin.

The Consumer's Program

Individual preferences are always represented by the utility function $U(C, L)$. Goods consumed may be purchased, in quantity C_M, or produced domestically, in quantity C_D, with $C = C_D + C_M$. The total endowment of time available L_0 breaks down into paid working time h_M, household working time h_D, and leisure L, hence $L_0 = h_M + h_D + L$. The efficiency of household tasks is represented by a "production function," $C_D = f(h_D)$, linking the amount of the good produced to the time spent on household work. This production function is increasing and concave; thus we will have $f' > 0$ and $f'' < 0$. Income is made up of wage earnings, wh_M, and non-wage ones, R. The consumer must choose the quantities C_M, C_D, h_D, h_M, and L that maximize his or her utility under the budget constraint $C_M \leq wh_M + R$. Let us further designate potential income as $R_0 = wL_0 + R$; since $h_M = L_0 - h_D - L$, the budget constraint is again written $C_M + wL \leq wh_D + R_0$. Taking into account the identity $C_M = C - f(h_D)$, the consumer's program then takes the following form:

$$\underset{\{C, L, h_D\}}{\text{Max}} \; U(C, L) \quad \text{subject to the budget constraint} \quad C + wL \leq [f(h_D) - wh_D] + R_0$$

In this program the choice variables of the consumer are *total* consumption C, leisure L, and the time h_D given over to household production. Additionally, the budget constraint shows that the total income of the consumer is equal to the sum of the potential income R_0 and the "profit" derived from household activities. Since household production only comes into the consumer's program through the expression of this profit, its optimal value h_D^* is that which maximizes the value of this profit; hence it is defined by $f'(h_D^*) = w$. Given time h_D^* dedicated to household activities, the consumer's program becomes formally equivalent to that of the basic model, as long as we replace potential income R_0 by $\tilde{R}_0 \equiv R_0 + f(h_D^*) - wh_D^*$. The optimal solutions $C^* = C_M^* + f(h_D^*)$ and L^* are then defined by the equalities:

$$\frac{U_L(C^*, L^*)}{U_C(C^*, L^*)} = w = f'(h_D^*) \qquad \text{and} \qquad C^* + wL^* = \tilde{R}_0 \tag{5}$$

This result is close to the one described by equation (2) in the basic model. At the optimum, the marginal rate of substitution between consumption and leisure is equal to the wage. As previously, this condition describes the division between the consumption of goods and that of leisure. The equality $f'(h_D^*) = w$ shows that the allocation of working time between household and waged activities is determined by the relative productivities of the two types of activity. Consequently the wage reflects the individual productivity of wage labor. The agent thus has an interest in devoting his or her working time to household activities to the extent that the marginal productivity $f'(h_D)$ of an hour of this type of work is superior to an hour's wage. Therefore he or she augments the length of time given to household work to the point where $f'(h_D^*) = w$.

Elasticity of the Labor Supply

The possibility of making trade-offs between household and waged activities alters the elasticity of the labor supply curve. The system of equations (5) allows us to write the optimal demand for leisure in the form $L^* = \Lambda(w, \tilde{R}_0)$. Differentiating this equality with respect to w, we get:

$$\frac{dL^*}{dw} = \Lambda_1 + \Lambda_2 \frac{d\tilde{R}_0}{dw} \qquad \text{with} \qquad \frac{d\tilde{R}_0}{dw} = L_0 - h_D^*$$

Because $f'(h_D^*) = w$ implies that $dh_D^*/dw = 1/f''(h_D^*)$, the identity $h_M^* \equiv L_0 - h_D^* - L^*$ entails:

$$\frac{dh_M^*}{dw} = -(\Lambda_1 + \Lambda_2 L_0) + \left[\Lambda_2 h_D^* - \frac{1}{f''(h_D^*)}\right] \tag{6}$$

The term $-(\Lambda_1 + \Lambda_2 L_0)$ represents the impact of a variation in the wage on the supply of wage labor for a given amount of household activity. It corresponds to the set of effects discussed in the basic model—see equation (2) above and the accompanying remarks. We have seen, in particular, that a change in the wage has an ambivalent impact on labor supply. The second term of the right-hand side of equation (6) is positive if leisure is a normal good (that is, if $\Lambda_2 > 0$). Consequently the possibility of making trade-offs against household activity ought to increase the wage elasticity of the labor supply. This result might explain why empirical studies show that the wage elasticity of the supply of female labor is generally higher than that of the supply of male labor (see section 1.4.1 below). For men, the trade-off between household and waged activity is often marginal. An instructive limit case is that of an optimal "corner solution," with a null supply of domestic labor $h_D^* = 0$. This might be the case if the productivity of household work were far below the current wage. A high proportion of men would then trade leisure off against wage labor only, whereas many women, whose household productivity is high in relation to the wage that they could get, would trade off among leisure, household activity, and wage labor.

Taking household activity into consideration allows us to make the predictions of the basic model richer. It should be emphasized, however, that the model presented here remains very rudimentary. For one thing, it rests on the hypothesis of an identical disutility of work for waged and household activities. In reality, the inconvenience arising from these activities is different. A more general approach, proposed by Becker (1965) consists of taking into account the disutility (or the utility) associated with each activity by distinguishing the diverse kinds of work done in the home. Such an approach has the merit of analyzing the choices underlying the allocation of time among different activities with great precision (on this subject, see the syntheses of Gronau, 1986 and 1997).

1.2.2 Intrafamilial Decisions

The family has considerable influence on the behavior of its members. The supply of labor is not exempt from this rule, and the basic model has to be adapted so as to take

into account the influence of family structures. The question bears an important empirical aspect, for numerous data (in particular those on consumption) only describe the behavior of the household, so we require a theory that goes beyond the basic individual frame of reference and gets us to a point where our estimates make some sense. The analysis of family choices has developed along two different lines. The first, known as the ''unitary'' model, starts from the principle that the family can be likened to a sole agent having its own utility function. The second, known generically as the ''collective'' approach, postulates that making choices is fundamentally something individuals do, and that the family is no more than a particular framework that enlarges (or constrains) the range of choices of each individual member of it.

The Unitary Model

This approach extends, as simply as possible, the basic model proposed hitherto. Let us imagine a family composed of two persons: we then postulate that the preferences of this entity are representable by a utility function $U(C, L_1, L_2)$, where C represents the *total* consumption of goods by the household and L_i $(i = 1, 2)$ designates the leisure of individual i.[3] This formalization assumes that the satisfaction attained through the consumption of a good depends solely on its total amount, and not on the manner in which it is shared among the individual members. For agent i, let us denote his or her wage and non-wage income respectively as w_i and R_i; the optimal choices are then determined by maximizing utility under a single budget constraint. The program of the household is written as:

$$\underset{\{C, L_1, L_2\}}{\text{Max}} \; U(C, L_1, L_2) \quad \text{s.c.} \quad C + w_1 L_1 + w_2 L_2 \leq R_1 + R_2 + (w_1 + w_2) L_0$$

Scrutiny of this program reveals that the unitary representation of the household implies that the *distribution* of non-wage incomes has no importance; the only thing that counts is their *sum* $R_1 + R_2$. This hypothesis, known in the literature as ''income pooling,'' signifies, for example, that it is not necessary to know which member of the couple is the beneficiary of transfer income. Now, the fact is that empirical studies refute this hypothesis for large segments of the population. For example, Fortin and Lacroix (1997) find that the unitary model only fits couples with pre-school-age children (see Blundell and MaCurdy, 1999, for a general overview). This invalidation is one of the reasons why the unitary model of the household is not completely satisfactory and is giving way to the collective model for the purpose of describing decisions made within a family.

The Collective Model

The most highly elaborated form of the collective model is due to Chiappori (1988, 1992). This model starts from the principle that household choices must arise out of individual preferences. In making the household the sole locus of decisions, the unitary model arbitrarily aggregates the preferences of its members, and hence does not respect the basic principle of ''methodological individualism.'' Conversely, if one does adhere to this principle, it appears natural to assume that decisions made within

a household are efficient in the Pareto sense, meaning that the possibility of mutually advantageous allocation does not occur. If we use $U_i(C_i, L_i)$, $i = 1, 2$, to designate the individual preferences of the persons composing the household, the efficient allocations will be the solutions of the following program:

$$\underset{\{C_1, C_2, L_1, L_2\}}{\text{Max}} \; U_1(C_1, L_1)$$

Subject to constraints:

$$U_2(C_2, L_2) \geq \bar{U}_2$$

$$C_1 + C_2 + w_1 L_1 + w_2 L_2 \leq R_1 + R_2 + (w_1 + w_2) L_0$$

In this program the parameter $\bar{U}_2$ represents a given level of utility, and we may suppose that it depends on the parameters w_i and R_i. Chiappori (1992, proposition 1) then shows that the efficient allocations are also the solutions of *individual* programs in which each person would be endowed with a specific non-wage income and which would depend on the overall income of the household. More precisely, the program of agent i takes the following form:

$$\underset{\{C_i, L_i\}}{\text{Max}} \; U_i(C_i, L_i) \quad \text{subject to constraint} \quad C_i + w_i L_i \leq w_i L_0 + \Phi_i$$

where Φ_i is a ''sharing rule,'' depending on the parameters w_i and R_i, and such that $\Phi_1 + \Phi_2 = R_1 + R_2$. In other words, it is as if each member of the household received a fraction of the total non-wage income of the household. In a way this approach reinforces the basic model of choice between the consumption of goods and leisure by specifying, for the budget constraint of the individual, the composition of his or her non-wage income. It is possible to expand the collective model by taking into account the ''public'' goods pertaining to the household and the household production of its members.

From the empirical point of view, the collective model has the advantage of not adopting the hypothesis of ''income pooling'' a priori; the latter is no more than a particular case of this model. Moreover, Chiappori (1992) shows that this formulation of the decision-making process within a household allows us to deduct individual consumption—which is not, for the most part, observable—using the individual supplies of labor and the total consumption of the household, which are observable entities. Hence, the simple observation of the supplies of labor and individual incomes allows us to determine the sharing rules within households. Knowing these rules, it becomes possible to assess the consequences of public policies for each member of the household using available data. In this context, Browning et al. (1994) have shown, on the basis of Canadian data, that differences of age and income among the members of households, as well as the wealth of households, appear to be the sole elements that affect the sharing rules Φ_i.

The Additional Worker Effect

Models of intrafamilial choice throw a revealing light on decisions to participate in the labor market. Taking into account the familial dimension does indeed allow us to

explain why certain members of the household specialize in household production, while others offer their services on the market for wage labor. From whatever angle the household is viewed, the choices of different members are interdependent, and an individual's fluctuations in income will have an impact on his or her own supply of labor, but also on that of the spouse or other members of the household, for example working-age children. This interdependence of choices may lead an individual to increase his or her supply of labor when the household income declines. It might even motivate him or her to participate in the labor market if he or she was not already doing so before the income fell. In principle, a fall in wages may thus entail an increase in the labor force, by spurring additional workers to enter the market for the precise purpose of making up for the loss of income in their household. From the empirical point of view, this *additional worker effect* seems to have little weight (see, for example, Lundberg, 1985). It is interesting to note that the additional worker effect implies a negative relationship between the participation rate and the average wage. When we constructed the aggregate supply of labor out of individuals making decisions in isolation, we obtained a positive relationship between the average wage and the participation rate (see above, section 1.2.2). In practice, this second relationship turns out to be dominant, and we do indeed observe a positive correlation between wages and the participation rate.

1.3 Life Cycle and Retirement

The static models utilized to this point obviously do not allow us to understand how agents substitute for their consumption of leisure over time when their flow of income undergoes transitory or permanent shocks. Taking into explicit account a succession of periods does not markedly alter the conclusions of the static model, but it does provide an adequate framework within which to scrutinize certain theories about the business cycle. The decision to go into retirement—in other words, the definitive end of participation in the labor market—can also be analyzed suitably using a dynamic model of labor supply within which we have redefined the flow of income and legal constraints.

1.3.1 Intertemporal Labor Supply

The dynamic theory of labor supply gives a central role to the possibility of substituting for the consumption of physical goods and leisure over time. We highlight this possibility using a dynamic model in discrete time. This model likewise allows us to grasp the contrasting effects caused by a transitory change in wages or a permanent modification of the wage profile, and thus to examine critically certain aspects of the theory of ''real business cycles.''

A Dynamic Model of Labor Supply

In a dynamic perspective, a consumer must make his or her choices over a ''life cycle'' represented by a succession of periods that start with an initial date, conventionally taken as equal to 0, and end with an independent terminal date, annotated T. Assuming

that the period t unfolds between the dates $(t-1)$ and t, the succession of periods is then given by the index $t = 1, 2, \ldots, T$. The date t is also used as an indicator of the age, professional experience, or seniority of an individual, according to the subjects under study. In a very general way, the preferences of the consumer must be represented by a utility function of the form $\mathscr{U}(C_1, \ldots, C_t, \ldots, C_T; L_1, \ldots, L_t, \ldots, L_T)$, where C_t and L_t designate respectively the consumption of physical goods and the consumption of leisure for the period t. But this very general form does not permit us to obtain analytically simple and easily interpretable results. That is why it is often assumed that the utility function of the consumer is temporally separable, in which case it is written $\sum_{t=1}^{t=T} U(C_t, L_t, t)$. Under this hypothesis, the term $U(C_t, L_t, t)$ represents simply the utility obtained by the consumer in the course of period t. It is sometimes called the ''instantaneous'' utility of the period t. We must bear in mind, however, that this representation of preferences is very restrictive: in particular, it does not allow us to take into account the inertia of habits of consumption, or ''habit persistence,'' that empirical studies reveal (see Hotz et al., 1988). To bring out this phenomenon, the influence of past consumption on the utility of the current period would have to be incorporated. Another important limitation of the model presented here has to do with the absence of decisions about training. Training increases the human capital of an individual and raises his or her wage-earning prospects, so trading off must take place between leisure, working time, and time dedicated to training (we examine this question in detail in chapter 2, section 1).

In this dynamic model, we will assume that individuals have the opportunity to save, and we will use r_t to denote the real rate of interest between the dates $t-1$ and t. For each period, the endowment of time is an independent constant to which we shall give the value 1 in order to simplify the notation. On this basis, the hours worked during a period t are equal to $(1 - L_t)$. If we use A_t to designate the consumer's assets on date t, and B_t to designate his or her income apart from wages and the yield on savings on the same date, for a given initial value A_0 for the assets, the evolution of the wealth of the consumer is described by:

$$A_t = (1 + r_t)A_{t-1} + B_t + w_t(1 - L_t) - C_t, \qquad \forall t \geq 1 \tag{7}$$

This equation can easily be understood as follows: on date t, the increase in wealth $A_t - A_{t-1}$ is due to income $w_t(1 - L_t)$ from wage labor, to income $r_t A_{t-1}$ from savings, and to other income B_t. Consumption C_t for the period has to be deducted from these gains. The *non-wage* income R_t for the period t is thus equal to $B_t + r_t A_{t-1}$.

Optimal Solutions and Demands in Frisch's Sense

The consumer attempts to maximize his or her intertemporal utility subject to the budget constraint described, on each date, by equation (7). If we use ν_t to denote the multiplier associated with this equation, the Lagrangian of the consumer's problem takes the form:

$$\mathscr{L} = \sum_{t=1}^{t=T} U(C_t, L_t, t) - \sum_{t=1}^{t=T} \nu_t [A_t - (1 + r_t)A_{t-1} - B_t - w_t(1 - L_t) + C_t]$$

The first-order conditions are obtained by equating the derivatives of this Langrangian to zero with respect to variables C_t, L_t, and A_t. After a few simple calculation, we arrive at:

$$U_C(C_t, L_t, t) = v_t \qquad \text{and} \qquad U_L(C_t, L_t, t) = v_t w_t \tag{8}$$

$$v_t = (1 + r_{t+1})v_{t+1} \tag{9}$$

Relations (8) imply $U_L/U_C = w_t$. The equality between the marginal rate of substitution and the current wage is thus maintained at every date, but this result is not general, it is a direct consequence of the hypothesis of the separability of the utility function. Limiting ourselves to interior solutions, the optimal consumptions of physical goods and leisure are implicitly written in the following manner:

$$C_t = C(w_t, v_t, t) \qquad \text{and} \qquad L_t = L(w_t, v_t, t) \tag{10}$$

For a given level of marginal utility of wealth, in other words, for a given v_t,[4] these equations define the ''Frischian'' demands for period t. The elasticity of labor supply in Frisch's sense is then equal to the current wage elasticity of function $h(w_t, v_t, t) = 1 - L(w_t, v_t, t)$, assuming that v_t remains constant. This elasticity is often called ''intertemporal substitution elasticity.'' If we take into account the fact that v_t is really an endogeneous variable depending on, among other things, the current wage, by analogy with the static case we may define the ''Marshallian'' elasticity of labor supply as being the current wage elasticity of function $h(w_t, v_t, t)$, taking into account the dependence between v_t and w_t. In order to define this elasticity, it is necessary to specify this dependence.

Equation (9), which is known as the Euler equation, shows that the multipliers v_t depend solely on the interest rate. More precisely, successive iterations of the logarithms of equation (9) entail:

$$\ln v_t = -\sum_{\tau=1}^{\tau=t} \ln(1 + r_\tau) + \ln v_0 \tag{11}$$

This way of writing the law of motion of v_t proves extremely interesting from the empirical point of view, since it shows that v_t can be broken down into a fixed individual effect v_0 and an age effect $-\sum_{\tau=1}^{\tau=t} \ln(1 + r_\tau)$ common to all agents (see subsection 2.1 below on the econometrics of the labor supply). Introducing uncertainty into this model, for example concerning wages, does not change the essential results notably. We can verify that the first-order conditions (8) remain true, whereas the marginal utility of wealth v_t becomes a random variable, following a stochastic process described by equation (11), with an error term with zero average appearing on the right-hand side of this equation (see Blundell and MaCurdy, 1999).

A priori, the value of v_0 depends on all the wages received by an individual during his or her lifetime. If we want to estimate the effects of a modification of the wage profile, and not just those due to a change in the current wage, then we have to take account of the dependence of v_0 on all wages. On the other hand, variation in a single wage, for example w_t, ought to have little influence on v_0 and elasticity in

Frisch's sense will certainly measure the effect of a change in a single wage w_t on labor supply $h(w_t, v_t, t)$. This difference, fundamental on the level of economic policy, between a modification of the wage profile and a change in a particular wage, emerges clearly with the help of the following example, taken from Blanchard and Fischer (1989, chapter 7, section 7.2).

Transitory Shock Versus Permanent Shock

Let us suppose that the real interest rate is constant ($r_t = r, \forall t \geq 0$), that the consumer is receiving no exogeneous income ($B_t = 0, \forall t \geq 0$), and that his or her instantaneous utility takes the explicit form:

$$U(C_t, L_t, t) = (1+\rho)^{-t}\left(\ln C_t + \frac{\sigma}{\sigma - 1} L_t^{(\sigma-1)/\sigma}\right), \qquad \sigma > 1,\ \rho \geq 0$$

The constant factor ρ represents the psychological discount rate. The Frischian demand functions are then written:

$$C_t = \frac{1}{(1+\rho)^t v_t} \qquad \text{and} \qquad L_t = \left[\frac{1}{(1+\rho)^t v_t w_t}\right]^{\sigma}$$

We may note that the intertemporal elasticity of substitution of leisure—in other words, elasticity in Frisch's sense—is equal, in absolute value, to the constant coefficient σ. With a constant interest rate, the Euler equation (9) then gives $v_t = v_0/(1+r)^t$, and the demand functions are expressed, as a function of v_0, in the form:

$$C_t = \frac{1}{v_0}\left(\frac{1+r}{1+\rho}\right)^t \qquad \text{and} \qquad L_t = \left[\frac{1}{v_0 w_t}\left(\frac{1+r}{1+\rho}\right)^t\right]^{\sigma} \tag{12}$$

In order to obtain an implicit equation giving the value of v_0, we have to write the intertemporal budget constraint of the consumer. This constraint is arrived at by eliminating assets A_t through successive iterations of the accumulation equation (7). With $r_t = r$ and $B_t = 0$ for all $t \geq 0$, we arrive at:

$$\sum_{t=1}^{T}(1+r)^{-t}(C_t + w_t L_t) = \sum_{t=1}^{T}(1+r)^{-t} w_t \tag{13}$$

This expression generalizes the budget constraint (1) of the static model: it states that the discounted present value of expenditure for the purchase of consumer goods and leisure cannot exceed the discounted present value of global income.

The value of v_0 is obtained by bringing the expressions of C_t and L_t given by (12) into the intertemporal budget constraint (13). It is implicitly defined by the following equation:

$$\sum_{t=1}^{T}(1+\rho)^{-t}\left\{1 + \left[\left(\frac{1+r}{1+\rho}\right)^{-t} v_0 w_t\right]^{1-\sigma} - \left(\frac{1+r}{1+\rho}\right)^{-t} v_0 w_t\right\} = 0 \tag{14}$$

It emerges clearly that the multiplier v_0 depends on all wages over the life cycle of the individual. For sufficiently large T, this multiplier is affected very little by

changes in a particular wage: what we have in that case is a transitory shock. On the other hand, it is affected by a change affecting all wages: what we have then is a modification of the wage profile, or a permanent shock. To grasp clearly the difference between these two types of shock, let us imagine that a permanent shock corresponds to a multiplication of all wages by a single positive quantity; relation (14) shows that v_0 will be divided by this quantity. But relation (12) then indicates that the optimal level of leisure—and therefore that of hours worked—remains unchanged. In this model, a permanent shock has no influence on labor supply, since the income effect and the substitution effect cancel each other out exactly. Let us now consider a transitory shock that causes only the wage w_t to change. This shock has only slight influence on the value of v_0, and relation (12) shows that leisure at date t diminishes, while leisure at all other dates remains unchanged. This particular model thus succeeds in conveying the notion that the permanent component of the evolution of real wages has no effect on labor supply, whereas the transitory component affects the level of supply immediately through the optimal response of agents who adjust their supply of labor in response to *temporary* changes in the wage.

Labor Supply and Real Business Cycles

Since the first publications of Lucas and Rapping (1969), a number of authors have studied changes in the labor supply as a function of movements in the real wage. The goal of these studies is to explain a striking fact of major importance, which is that aggregate employment fluctuates a great deal in the course of a cycle, whereas the transitory component of changes in the real wage proves limited in scope. At the outset, the theory referred to as that of ''real business cycles'' saw the mechanism of intertemporal substitution of leisure as the principal cause of fluctuations in the level of employment. Following this line of thought, the economy is always the object of multiple shocks (on technology, or on preferences) that have repercussions on the remuneration of labor and capital; to these agents respond in an optimal manner by instantaneously adjusting their supply of labor. More precisely, a favorable shock, one perceived as transitory, would motivate agents to increase their supply of labor today and to reduce it tomorrow when the shock has passed (for a comprehensive evaluation of the implications of the theory of real business cycles for the labor market, see Hall, 1999). This theory is simple, even seductive, but it runs up against a sizable obstacle. If it is to agree with empirical findings, it must explain how *small* movements in the real wage could entail *large* variations in the level of employment.

Hence in its original version, the theory of real business cycles requires employment to be very sensitive to small changes in the wage. Relation (12) shows that this will be the case if the absolute value of the intertemporal elasticity of substitution of leisure σ is large. Now, the majority of empirical studies arrive instead at small values (Hall, 1980, estimates that a value of 0.4 might apply at the macroeconomic level; Pencavel, 1986, suggests values even lower than that for men, while Blundell et al., 1993, find levels ranging from 0.5 to 1 for married women in the United Kingdom). In these circumstances, variations in the labor supply in response to transitory

changes in the wage cannot serve as a sufficient basis for a theory of the business cycle. Relation (12) does indicate, however, that transitory shocks might influence the level of employment via interest rates. Since these variables are noticeably more volatile than wages, there would thus be another way to reproduce the stylized facts in question. This trail, however, also comes to a dead end. To demonstrate this, let us suppose that the intertemporal utility function of the consumer is temporally separable; the first-order conditions (8) then imply:

$$\frac{u_L(C_t, L_t, t)}{u_C(C_t, L_t, t)} = w_t \qquad \forall t = 1, \ldots, T$$

If the wage does not change, it can easily be verified that this expression defines an increasing relation between consumption and leisure if these are normal goods. In this case, movements in labor supply supposedly due to the variability of interest rates alone would be accompanied by an *inverse* movement of consumption. Here too we run up against contradictory empirical observations, which show a positive correlation between levels of employment and consumption. Faced with this fresh setback, one might try out other modifications of the formulation of the problem of the trade-off over time between consumption and leisure, such as, for example, giving up the hypothesis of separability, or introducing fixed costs into the decision to participate. To this day, no way has really been found to escape the substantially negative verdict that hangs over explanations of variability in employment based on the sole mechanism of intertemporal substitution of leisure (see the discussion and proposals of Hall, 1999).

1.3.2 Economic Analysis of the Decision to Retire

Economic analysis of the process by which a person ends his or her labor market participation fits well into the life-cycle model offered above, provided that legal constraints and the flow of income specific to retirement are brought into clear focus. In an uncertain environment, the process of making this decision can be analyzed with the help of the ''option value'' associated with the choice not to go into retirement today. Empirical studies show that workers generally react in a significant fashion to the financial incentives that accompany either early retirement or continued wage-earning.

Social Security and Private Pensions

Most countries in the OECD zone have put in place pension systems, public and private, enabling workers to receive income when they retire from the labor market. For example, in the United States there exists a public system (Social Security) funded by mandatory contributions coming from employers, which gives around 41% of his or her last wage to the median worker retiring at age 62. This ratio increases by 6.67% each year between 62 and 65. Every individual has the opportunity to supplement this public retirement payout with private pensions, contributions to which are negotiated between employer and employee at the moment the labor contract is signed. Taken

as a whole, these contributions represent considerable financial accumulations—the celebrated pension funds—managed by specialized insurance companies that pay out retirement pensions to their members that vary according to the return their investments have made. In other countries like the Netherlands and France, the private system is practically nonexistent, and the replacement rate offered by the public pensions is, in these two countries, on the order of 91% for a person who ends his or her wage-earning activity at age 60 (for a comparative international perspective, see Gruber and Wise, 1999 and 2001, from which these isolated figures are taken).

The system of public and private pensions, to which we must add the tax system, creates incentives for workers to take their retirement earlier or later. Most retirement systems specify a legal age, sometimes called the ''normal'' age, past which a worker is obliged to end his or her wage-earning activity (for example, normal retirement falls at 65 in the United States and Japan, and 70 in the United Kingdom). But every individual obviously has the right to retire before this legal age. As a general rule, he or she receives a smaller pension the farther the age at which he or she ceases to work is from the legal age. Hence the decision to retire brings into play a number of elements that emerge very clearly with the help of the life-cycle model, significantly modified.

Option Value in the Life-Cycle Model

Let us consider a person employed on date τ—this date represents, if you like, the age of this person—and let us suppose that this person decides to retire on date $s \geq \tau$. The evolution of his or her wealth starting from date τ is always given by equation (7), provided that we redefine certain variables of this equation. So, to simplify, we will suppose that the agent does not work at all after date s; we will then have $L_t = 1$ for $t \geq s$. In practice, the process of ceasing to work can be gradual, and for that matter the legislation sometimes permits work to continue while the agent is receiving a retirement pension. We will use $B_t(s)$ to denote the income expected in the period $t \geq s$, composed of pension payments over the period t and other income which the agent may happen to have. Most often, this income is an increasing function of age s from career onset to retirement. To avoid any confusion, we will use $B_t(0)$ to designate the non-wage income of the agent while he or she is still working, hence for $t < s$, and we will use C_{et} and C_{rt} respectively to designate his or her consumption of physical goods before and after retirement. For given s, the agent solves the following problem:

$$\underset{C_{et}, C_{rt}, L_t}{\text{Max}} \left[\sum_{t=\tau}^{s-1} U(C_{et}, L_t, t) + \sum_{t=s}^{T} U(C_{rt}, 1, t) \right]$$

Subject to constraints:

$$A_t = \begin{cases} (1 + r_t)A_{t-1} + B_t(0) + w_t(1 - L_t) - C_{et} & \text{if } \tau \leq t \leq s - 1 \\ (1 + r_t)A_{t-1} + B_t(s) - C_{rt} & \text{if } s \leq t \leq T \end{cases}$$

Let us designate the value of the welfare of the consumer at the optimum of this problem by $V_\tau(s)$, and finally let us denote the legal age of retirement by T_m, after

which it is not possible to work any more. An agent age τ chooses the date s on which to end his or her working life by solving the following problem[5]:

$$\underset{s}{\text{Max}}\ V_\tau(s) \quad \text{subject to constraint} \quad T_m \geq s \geq \tau \tag{15}$$

These problems never lend themselves to an explicit resolution and are generally solved numerically. In practice, we have to specify the utility function and the manner in which the replacement income is assembled to arrive at a model capable of being simulated or estimated empirically (one of the first attempts is found in Gustman and Steinmeier, 1986). Moreover, the decision to retire is made in an environment marked by numerous uncertainties (changes in one's professional and married life starting from date t, the chances of illness, changes in taste, retirement systems, etc.) that steadily subside as the legal age approaches. In order to simplify the explanation, we have written the agent's program without taking these uncertainties into account, but it is easy formally to introduce random factors into the utility function and into the equation for the evolution of wealth so as to obtain a stochastic model that fits reality more closely. In this case, $V_\tau(s)$ represents the intertemporal utility *expected* by an agent of age τ. Supplementary information may be acquired that will cause the decision taken at age $(\tau + 1)$ to be different from the decision taken at age τ. Let us denote by s^* the optimal solution of problem (15); for every period, the program (15) allows the agent to choose between two possibilities: retire today—the optimal solution of the problem of the agent is a corner solution such as $s^* = \tau$—or continue to work until age $(\tau + 1)$ and reconsider his or her decision then, in which case the optimal solution is of the kind $s^* > \tau$.

This way of envisaging the process of ending one's working life leads us to examine the *option value* attached to the decision not to take retirement right now (Stock and Wise, 1990). Supposing that the decision to retire is irreversible, we have just shown that if $s^* = \tau$, the agent stops working immediately, and on the other hand if $s^* > \tau$, the agent continues to work and reconsiders his or her decision at age $(\tau + 1)$ in light of the new situation that he or she will be in when that date comes. The option value of not retiring today is thus equal to $V_\tau(s^*) - V_\tau(\tau)$. If it is positive, the agent continues to work. If it is not, he or she goes into retirement. At the empirical level, this approach suggests that we estimate the probability of retirement at a given age by taking the option value as our principal explanatory variable. In order to obtain an indicator of this variable, we have to choose an explicit utility function, then estimate the option value tied to this utility function on the basis of a set of relevant variables, among which are income from public and private pensions and the wage outlook (readers may consult the survey of Lumsdaine and Mitchell, 1999, for more details). In general, the indicator of the option value strongly influences decisions about retiring.

Some Facts About the Impact of Eligibility Rules

Empirical studies carried out in the United States have shown that changes made to the eligibility rules regarding Social Security pensions (the elimination of means test-

ing, extension of the normal age for stopping work) have had little effect. The reason perhaps lies in the fact that private pension plans encourage workers to take their retirement starting at age 55, whereas Social Security only pays retirement pension starting at age 62. If one looks only at private pensions, Gustman et al. (1994) have shown that individuals with the highest pensions are those who retire soonest. But this income effect is relatively feeble, since at age 60, a 10% increase in expected income over the entire (expected) duration of retirement reduces the length of working life by less than two months. Conversely, workers under financial pressure to postpone their retirement do in fact extend their working lives. Here, too, the quantitative effects are faint: a 10% rise in expected income over the entire (expected) duration of retirement prolongs working life by less than six months.

These results reveal the effects of retirement plans entered into at the time the worker was hired. But it is possible that, for reasons of productive efficiency, firms may offer pension plans that make it advantageous to take retirement sooner. Such firms will therefore attract workers who have a stronger inclination to retire sooner. In this case, the observed correlation between the financial incentives and the age at which retirement is taken do not reveal a causality; they simply show a property of an optimal contract between particular types of firms and particular workers. In order to eliminate this endogenous bias, numerous studies analyze the behavior of workers in the face of *unanticipated* changes in their retirement conditions. For example, Lumsdaine et al. (1990) studied a large American firm that, in 1982, offered a ''window'' to its employees over 55 and enrolled in the pension plan, through which they could retire early; the financial bonus offered exceeded a year's worth of wages for certain categories of worker. By definition, this window of opportunity was of limited duration and had not been anticipated by the employees. Clearly, it therefore counts as an exogenous shock. Lumsdaine et al. (1990) found that, in the case of the workers most advantaged by the new arrangement, the rate of leaving more than tripled. For the overall workforce, this study estimates that, for a worker aged 50 employed in the firm, the likelihood of his or her retiring at age 60 was 0.77 under the new arrangement, whereas it was only 0.37 before it was put in place. These results are confirmed by Brown (1999), who systematically examined the effect of ''windows'' utilizing data on the entire American population provided by the Health and Retirement Study (HRS).

The effects of this type of financial incentive can also be studied through international comparisons. The studies of Gruber and Wise (1999, 2001) on a number of OECD countries show that financial incentives have, as a general rule, important impacts on the decision to retire.

2 EMPIRICAL ASPECTS OF LABOR SUPPLY

The supply of labor is probably the area of labor economics in which the greatest number of empirical studies have been carried out over the last twenty years.

Advances in econometric methods have accompanied and made possible this increase. The reason for this trend is that, for those whose job it is to plan employment policies or reforms of the fiscal system, the response of labor supply is a primary consideration. The econometrics of labor supply today rests on a solid foundation, of which we shall give the essential aspects. A retrospective of the principal results will complete this empirical tableau.

2.1 Introduction to the Econometrics of Labor Supply

The econometrics of labor supply is today a domain of study in its own right, and we shall merely sketch the problems that arise within it and the principles that govern their resolution. For a comprehensive account, the reader will profit from consulting the survey of Blundell and MaCurdy (1999).

2.1.1 The Principal Ingredients of a Labor Supply Equation

The principal goal of empirical models of the individual labor supply is to furnish an estimate of the wage elasticity of this supply. But the preceding theoretical analyses have taught us that there are several possible definitions of this elasticity, according to whether or not we integrate a temporal dimension into the choices of consumers. On the empirical level, it is primarily the way an indicator of income from sources other than the current wage is constructed that permits us to discriminate between the different definitions of elasticity. Based on the preceding theoretical analyses, in what follows wages will be treated as exogenous or independent variables. This hypothesis is not entirely satisfactory. From the dynamic point of view in particular, an individual's wage must depend on, among other things, the training he or she has decided to acquire and his or her seniority. Because these considerations belong more properly to the theory of human capital than to that of labor supply, we shall return to them later in chapter 2.

The Basic Equation and the Specification of Control Variables

As a general rule, estimates of labor supply are made on the basis of cross-sectional data (perhaps with temporal elements as well) produced by investigating a population of large size, out of which a number of individuals or households are sampled. The empirical models which the econometrician tries to estimate always rest on a basic equation relating hours h worked by a given individual at hourly wage w at each date. The following log-linear relation is a typical form of this basic equation:

$$\ln h = \alpha_w \ln w + \alpha_R \ln \mathscr{R} + x\theta + \varepsilon \tag{16}$$

In this expression, $\mathscr{R}$ is a measure of income other than the current wage, x is a vector of dimension $(1, n)$—one row and n columns—describing the n individual charateristics or control variables used, and θ is a vector of dimension $(n, 1)$ comprising n parameters to be estimated. The coefficients α_w and α_R are also parameters to be estimated, and finally, ε designates a random term reflecting individual heterogeneity that is not observed. Certain studies take h as a dependent variable rather than $\ln h$ and/or income w and $\mathscr{R}$ rather than $\ln w$ or $\ln \mathscr{R}$. These different specifications corre-

spond to different restrictions on preferences (see Blundell and MaCurdy, 1999) that do not alter the principles guiding the estimation of equation (16). In order to fit theoretical models, such as, for example, the one in section 1.1.1, it is also possible to introduce a polynomial form of wage into the right-hand side of equation (16) so as to avoid postulating a priori that hours worked are a monotonic function of the hourly wage.

Parameter α_w measures the wage elasticity of labor supply. This elasticity can be interpreted in several different ways according to the hypotheses made and the model utilized: a diversity of interpretation present here in the manner in which $\mathscr{R}$, indicating income apart from the current wage, is specified. The theoretical models taught us that individual labor supply at a given period was a function of the hourly wage for that period and other elements forming the expected wealth of an agent, such as, for example, his or her anticipated income from savings or work. If we limit ourselves to an equation of type (16), these elements have to be incorporated into variable $\mathscr{R}$. The important thing is to know how to carry out this incorporation.

One solution might be to consider only non-wage income for the period under investigation. During our study of the life-cycle model in section 1.3.1, we made it clear that this income, denoted by R_t, is composed of income from savings, which, for date t, are denoted by $r_t A_{t-1}$ (denoting by r_t the rate of interest between periods $t-1$ and t, and by A_{t-1} the assets of the agent in period $t-1$), and exogenous income B_t. To set $\mathscr{R}_t = R_t = r_t A_{t-1} + B_t$ amounts to supposing that agents make their choices in a myopic fashion, with no opportunity to save today for consumption tomorrow. But this hypothesis of total myopia is not in the least realistic, for agents largely make choices with an eye to the future, so that to estimate coefficient α_w while taking $\mathscr{R}$ to be only non-wage income at the date of the investigation does not give pertinent information about the real reactions of labor supply. It is possible to make up for this drawback by defining indicator $\mathscr{R}$ differently. To that end, it will help to return to what we learned from the life-cycle model laid out in section 1.3.1.

A Reexamination of the Life-Cycle Model

If, in the life-cycle model in section 1.3.1, the utility function is temporally separable, we have seen that the first-order condition (8) always implies equality between the marginal rate of substitution between consumption and leisure and the current wage *at each date*. This property suggests a two-stage resolution of this model, known in the literature as ''two-stage budgeting.'' In the first stage, analogous to the basic static model, we define a potential income Ω_t for each period t in such a way that the consumer's program consists of maximizing his or her instantaneous utility for the period t under a budget constraint, of which the non-wage income would be exactly Ω_t. In the second stage, the consumer optimizes the series of Ω_t, given the resources, present or anticipated, at his or her disposal. To arrive at such a program, we must first point out that the intertemporal budget constraint (7) of the life-cycle model can be rewritten in the following way:

$$C_t + w_t h_t = (1 + r_t)A_{t-1} + B_t - A_t$$

Let us set $\Omega_t = (1 + r_t)A_{t-1} + B_t - A_t$; the two-stage procedure by which the consumer resolves the program then emerges quite naturally. In the first stage, the consumer makes his or her choices for period t while maximizing instantaneous utility $U(C_t, 1 - h_t, t)$ subject to the "static" budget constraint $C_t + w_t h_t = \Omega_t$. At the conclusion of this first stage, the consumer thus attains a level of indirect utility $V(\Omega_t, t)$. In the second stage, he or she selects the optimal path for his or her assets A_t by solving the program:

$$\underset{\{A_t\}}{\text{Max}} \sum_{t=0}^{t=T} V(\Omega_t, t) \quad \text{s.c.} \quad \Omega_t = (1 + r_t)A_{t-1} + B_t - A_t, \ \forall t$$

This two-stage procedure evidently gives the same solutions as the solution (in one stage) employed in section 1.3.1.

Changes in a Wage

On the empirical level, we should first note that the econometrician can know the values of Ω_t when he or she can observe the value of the consumption of physical goods C_t and the hours worked h_t, since $\Omega_t = C_t + w_t h_t$. If that is not the case, or if they cannot be known precisely enough, it is possible to estimate Ω_t by taking as explanatory variables the value A_{t-1} of assets at the outset of period t, the interest rate r_t, exogenous income B_t, all or part of the control variables of vector x, and the expectation of all these independent variables (inasmuch as the value A_t of the assets at the end of the period t is not necessarily known, and depends on expectations of future resources). Hence, if we wish to make a relevant assessment (that is, one that avoids the supposition that individuals are completely myopic) of the reactions of labor supply to changes in a given wage, it is best to take $\mathscr{R}$ as an estimator of potential income Ω. In other words, if t designates the date of the survey, the income indicator $\mathscr{R}_t$ to be taken into account in the basic equation (16) must then be estimated by a relation of the type:

$$\mathscr{R}_t = \mathscr{R}(A_{t-1}, r_t, B_t, x_t, Z_t)$$

Here, Z_t represents the vector of the anticipated values of r, w, B and x. Note that, according to the procedure of "two-stage budgeting," potential income is an *endogenous* variable, since its value depends on choices made by the consumer during the allocation through time of his or her wealth. Hence it is best to apply methods based on instrumental variables in order to estimate equation (16). The "two-stage budgeting" procedure allows us to estimate, in a pertinent manner, the elasticity of labor supply with respect to *one* particular wage (or *one* expected wage), but does not allow us to know the effects of a change in the *overall* wage profile, since under this hypothesis, potential income Ω_t changes as well. Now, it is indispensable to study the overall wage profile if one wants to know, for example, the impact of a reform of the tax system, or more generally any measure of economic policy likely to become permanent. Before answering this question, we will show how to measure elasticity in Frisch's sense.

Estimating Elasticity in Frisch's Sense

The dynamic model of section 1.3.1 has much to teach us. In particular, relations (10) and (11), defining its solutions, reveal that labor supply h_t depends on the current wage w_t and the marginal utility of wealth v_t, so that $h_t = h(w_t, v_t, t)$. According to relation (11) of this model, the logarithm of v_t breaks down into an individual fixed effect equal to $\ln v_0$ and an age effect $\sum_{\tau=1}^{\tau=t} \ln(1 + r_\tau)$, common to all agents and which may be written in the form ρt, supposing that r_τ is constant. To obtain the elasticity of the labor supply in Frisch's sense, also called the intertemporal elasticity of substitution, we view the marginal utility of wealth v_t as exogenously given, independent of the current wage. Following relation (11), we see that that amounts to supposing that $\ln v_0$ is also independent of the current wage, but evidently does depend on individual characteristics. This property suggests substituting $\ln v_0 + \rho t$ for $\ln \mathscr{R}$ in equation (16) to estimate Frischian elasticity. If we have longitudinal data available, we can eliminate individual fixed effects by taking the basic relation (16) in first-differences, which is written:

$$\Delta \ln h_t = \rho + \Delta x_t \theta + \alpha_w \Delta \ln w_t + \Delta \varepsilon_t$$

This equation allows us to estimate the elasticity of labor supply in Frisch's sense in a coherent manner, that is, the impact of a transitory change in the wage. It does not, however, allow us to evaluate the impact of a change in the overall wage profile, for a change of this type causes the marginal utility of wealth to vary a priori.

Changes in the Wage Profile

In order to evaluate the consequences of a change in the overall wage profile, we have to take into consideration variations in the marginal utility of wealth. The initial value of the marginal utility of wealth v_0 depends on individual preferences and all anticipated income; it may be approximated by the equation:

$$\ln v_0 = y\alpha_y + \sum_{i=0}^{T} \gamma_i E_0 \ln w_i + \phi A_0$$

In this expression, y, α_y and T ($T \geq \tau$) designate respectively a vector of individual characteristics relating to the onset of working life, a vector of parameters to be estimated, and the duration of working life (putatively known). The term A_0 designates the initial value of the stock of wealth, ϕ is a parameter, and E represents the expectation operator. Replacing $\ln \mathscr{R}$ by $\ln v_0 + \rho t$ in the basic equation (16), this equation becomes:

$$\ln h_t = \alpha_w \ln w_t + x\theta + y\alpha_y + \sum_{\substack{i=0 \\ i \neq t}}^{T} \gamma_i E_0 \ln w_i + \phi A_0 + \rho t + \varepsilon_t \tag{17}$$

Expected wages, which are evidently not observed, can themselves be approximated by an equation of the form:

$$E_0 \ln w_t = a_0 + a_1 t + a_2 t^2 + u_t \tag{18}$$

In this equality we have set $a_j = z\alpha_j,\ j = 0, 1, 2$, where z is a vector of observable characteristics unchanging over time, α_j is a vector of parameters, and u_t is a random element. The term t^2 is introduced to account for a possible nonlinearity in the relation between wages and experience, which is generally confirmed by empirical work on this subject (see below, chapter 6, section 4.3). The simultaneous estimation of equations (17) and (18) allows us to obtain the parameters needed to assess the impact of an overall change in wages on labor supply. Parameter α_w measures the impact of a change in the current wage w_t, while parameters γ_i measure the consequences of changes in the overall wage profile (see Blundell and MaCurdy, 1999, pp. 1600–1603, for more details).

To sum up, it is necessary to define precisely the set of variables that explain labor supply—in particular, the indicators of non-wage income—in order to see what type of elasticity the model utilized allows us to estimate. Having thus set out the ingredients that go to make up an empirical labor supply equation of type (16), we can now present the principles that guide this estimation.

2.1.2 A Short Guide to Estimating Labor Supply

Estimating the basic equation by ordinary least squares leads to biased results, since it neglects to take into account participation decisions. If we want to obtain unbiased estimators of the elasticity of labor supply, we have to estimate jointly decisions to participate and decisions about the number of hours worked. These estimates oblige us to attribute a fictitious wage to those who do not participate in the labor market.

What We Must Not Do

The first idea that comes to mind is to apply the method of ordinary least squares to equation (16) alone. Until the 1970s most studies proceeded in this way. But it is not a correct method, for it fails to distinguish decisions about *participation* in the labor market from those about the number of hours an agent is prepared to offer. The question that faces the econometrician is, given a sample of individuals, how to take into account persons who do not work (or episodes during which an agent has not worked if the data are equally temporal)? Certain studies subsequent to the 1970s simply set $h_i = 0$ for these persons. In other words, these studies took the view that certain workers choose exactly $h_i = 0$, just like any other value of h_i, which entails that equation (16) holds for any wage value of h_i and w_i. It is precisely this last hypothesis that is false. Equation (16) is only valid for wages *above* the reservation wage, and for *all other* wages, labor supply is null. Making do with equation (16) and setting $h_i = 0$ for episodes of nonwork thus leads to specification errors. An alternative solution was simply to exclude the unemployed, and nonparticipants in the labor market, from the sample. But in this case the econometrician commits a selection bias, forgetting that not to supply any hours of work is a decision in the same way that supplying them is. The fact that this type of decision is not described by equation (16) does not authorize us to set it aside purely and simply. The solution is to employ an empirical

model that, like the basic model of 1.1.1, describes participation and hours decisions *jointly*.

What We Must Do

The approach most often utilized today is ''structural.'' It combines an explicit functional form for the direct utility function of consumers, dependent in parametric fashion on the different observable characteristics of an individual, and a random term representing the nonobserved heterogeneity among individuals. We then write the budget constraint, from which we deduce, by the usual maximization procedure, labor supply and the reservation wage. The participation condition is then arrived at using the probability distribution of the random term, by positing that the wage offered must be superior to the reservation wage. We estimate the model at which we arrive using cross-sectional data that specify, for each individual, the values of every variable we are interested in, and his or her decisions to participate or not in the labor market. Let us illustrate this approach using an example, for purely pedagogic purposes, based on the static model of section 1.1.1, with a utility function of the Cobb-Douglas type.

The utility of a consumer will then take the form $C^{1-\beta}L^{\beta}$, $1 > \beta > 0$, and the budget constraint continues to be written $C + wL = wL_0 + R$. We assume that the explanatory variables and the random term intervene via the coefficient β according to the linear form $\beta = x\theta + \varepsilon$. Following the static model of section 1.1.1, we know that the reservation wage w_A is equal to the marginal rate of substitution U_L/U_C taken at point (R, L_0) and that the maximization of utility subject to the budget constraint gives the optimal value of leisure. After several simple calculations, we find that:

$$w_A = \frac{\beta}{1-\beta}\frac{R}{L_0} \qquad \text{and} \qquad L = \begin{cases} \beta\left(L_0 + \dfrac{R}{w}\right) & \text{if } w \geq w_A \\ L_0 & \text{if } w \leq w_A \end{cases}$$

Since the coefficient β is a function of the random term ε, the inequality $w \geq w_A$ is equivalent to an inequality on the values of ε, which is written:

$$w \geq w_A \Leftrightarrow \varepsilon \leq \frac{wL_0}{R + wL_0} - x\theta$$

In conclusion, the decisions concerning labor supply $h = L_0 - L$ and participation may be summed up in this fashion:

$$h = \begin{cases} L_0 - (x\theta + \varepsilon)\left(L_0 + \dfrac{R}{w}\right) & \text{if } \varepsilon \leq \dfrac{wL_0}{R + wL_0} - x\theta \\ 0 & \text{if } \varepsilon \geq \dfrac{wL_0}{R + wL_0} - x\theta \end{cases} \tag{19}$$

This expression of labor supply is related, as regards the interior solution, to the basic equation (16). But we see that taking account of participation decisions constrains the variations of the random term, making them depend on explanatory

variables. In these circumstances, the use of ordinary least squares is seen to be inadequate.

Joint Estimations of Hours Worked and Participation Decisions

Let us suppose that the econometrician has at his or her disposal a sample of individuals, N in size, specifying that individuals $i = 1, \ldots, J$ have worked h_i hours and that individuals $i = J + 1, \ldots, N$ have not worked. Let us denote by $F(.)$ and $f(.)$ respectively the cumulative distribution function and the probability density of the random term ε (the random term is most often assumed to follow a normal distribution). It is then possible to write the likelihood of the sample. Following rule (19) giving the optimal decisions of an agent, when an individual i has worked h_i hours, that means that the random term has taken the value $\varepsilon_i = w_i(L_0 - h_i)/(R_i + w_i L_0) - x_i\theta$. In this case its contribution to the likelihood of the sample is equal to $f(\varepsilon_i)$. If agent i has not worked, that means that the random term is bounded above by the value $\tilde{\varepsilon}_i = [w_i L_0/(R_i + w_i L_0)] - x_i\theta$. In this case, its contribution to the likelihood of the sample is given by $\Pr\{h_i = 0\} = 1 - F(\tilde{\varepsilon}_i)$. Setting $\bar{F} = 1 - F$, the likelihood function of the sample is written in logarithmic form:

$$\mathscr{L} = \sum_{i=1}^{i=J} \ln f\left[\frac{w_i(L_0 - h_i)}{R_i + w_i L_0} - x_i\theta\right] + \sum_{i=J+1}^{i=N} \ln \bar{F}\left[\frac{w_i L_0}{R_i + w_i L_0} - x_i\theta\right] \tag{20}$$

The maximization of the likelihood function by appropriate techniques (in this case of the *probit* type, since there is a mixture of continuous and discrete variables) furnishes estimates of the parameters in which we are interested. The expression of the likelihood function also permits us to understand clearly the mistakes made in failing to formalize participation decisions completely. If we set $h_i = 0$ for persons who do not work, that amounts to believing that their contribution to the likelihood is equal to $f[(w_i L_0/(R_i + w_i L_0)) - x_i\theta]$, which comes down to substituting function f for function $\bar{F}$ in the second term of the right-hand side of relation (20). If we exclude persons who do not work from the sample, then we are neglecting to take account of the second term on the right-hand side of relation (20). These two solutions result in biased estimators.

A Nonparticipant's Wage

The expression (20) of the likelihood function also highlights a delicate problem. By definition, the econometrician does not observe the wages of individuals $i = J + 1, \ldots, N$ who do not work. However, relation (20) shows that it is necessary to attribute a fictitious *wage* to these individuals if we want to maximize the likelihood function. We thus have to be able to assign a quantity to the (unobserved) wage notionally offered to an individual, which he or she has refused. The most common solution at present consists of deducing the wage of a nonparticipant using the wage received by participants with similar characteristics in terms of educational qualification, experience, age, and so on. In practice we can explain the wages of individuals participating in the labor market by a regression of the type $w_i = y_i\theta_p + u_i$ in which

the vector y_i represents the characteristics of an individual i participating in the labor market, and θ_p designates the vector of the parameters to be estimated. Let us use $\hat{\theta}_p$ to denote the vector of the estimates of θ; we can then use this vector $\hat{\theta}_p$ to calculate the wage w_k of a nonparticipant k, using the vector y_k of his or her characteristics and setting $w_k = y_k\hat{\theta}_p$. This simple technique unfortunately presents a *selection bias*, since it assumes that the regression equation $w_i = y_i\theta_p + u_i$ also applies to the notional wages of nonparticipants. This hypothesis is highly likely to be mistaken, inasmuch as participants in the labor market must on average have nonobserved characteristics that allow them to demand wages higher than those that nonparticipants can demand. Formally this means that the distribution of the random disturbance u_i should not be the same for participants and nonparticipants. The distribution that applies to participants ought to weight the high values of the random factor more strongly than the one that applies to nonparticipants, and consequently the estimation procedure described previously will *overestimate* the notional wage attributable to a nonparticipant. One way to correct this bias consists of making simultaneous estimations of equations explaining wages and decisions to supply labor (see Heckman, 1974, for an application).

2.1.3 Nonlinear Budget Constraint

In practice, the budget constraint of an agent does not come down to a simple segment of a line, as in the basic model of section 1.1.1. Mandatory contributions and transfers make this constraint (at best) piecewise linear. The estimation of labor supply then runs into a new problem, that of the endogeneity of the choice of the ''piece'' on which an agent will settle. The method of virtual incomes and the construction of a differentiable approximation of the budget constraint make such an estimation possible, however.

The Method of Virtual Incomes

In all countries, the systems of tax and subsidy that agents come under present important differences according to income, so that, from the point of view of empirical estimations, it is not possible to assume that the budget constraint of an agent is represented by a single segment of a line, as in the basic model of section 1.1.1. In practice, the different schedules of marginal rates according to income brackets, and the different deductions to which certain contributors are entitled, imply that the budget constraint of an agent is piecewise linear. By way of illustration, let us consider the example of a tax system in which an agent whose income does not exceed an exogenous threshold $R_{\max}$ is not taxed, whereas if his or her income crosses this threshold, his or her wage will be taxed at rate τ. Let us use w and R to denote respectively the wage and the non-wage income of this agent. Our example of a fiscal system starts to tax the consumer from the point at which his or her working time surpasses the value $h_{\max}$ defined by $wh_{\max} + R = R_{\max}$. To this maximal value of working time there corresponds a value for leisure of $L_{\min} = L_0 - h_{\max}$. Figure 1.6 represents the budget constraint associated with this rudimentary fiscal system in the plane (L, C). In reality, the

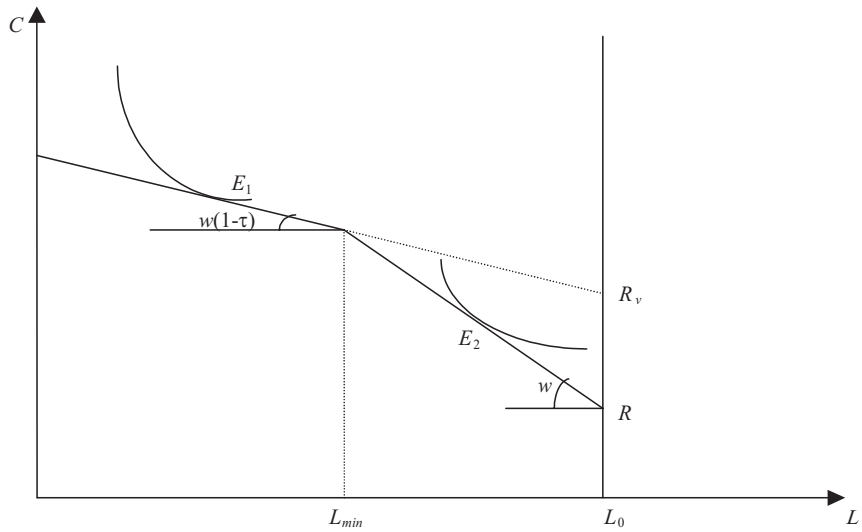

Figure 1.6
Piecewise linear budget constraint.

budget constraint is made up of more than two segments, and the set situated under the budget constraint can even present nonconvexities, due, for example, to the rate applied to overtime hours. The coherence of the tax system dictates, however, that the budget constraint should be continuous. Under this hypothesis, this constraint is characterized in the following manner:

$$C = \begin{cases} wh + R & \text{if } h \leq h_{\max} \\ wh(1-\tau) + R + w\tau h_{\max} & \text{if } h \geq h_{\max} \end{cases}$$

When the consumer chooses what he or she will consume in such a way as to maximize his or her utility $U(C, L)$ subject to his or her budget constraint, figure 1.6 shows that an interior solution may be situated at points E_1 or E_2, according to whether or not labor supply is such that the consumer is taxed. This figure also indicates that point E_1 corresponds to the optimum of the consumer whose hourly wage would be equal to $w(1-\tau)$ and who would receive a *virtual* non-wage income R_v equal to $R + w\tau h_{\max} = R + \tau(R_{\max} - R)$. It should be noted that this virtual income is perfectly well known, so it forms part of the ''observations'' available to the econometrician. Let us denote by $\varphi(w, R)$ the expression of labor supply if there were no taxes, that is to say, its value at point E_2. Let us again denote by w_A the reservation wage, which is once more equal to the marginal rate of substitution between leisure and consumption evaluated at the point of nonemployment. Since $h_{\max} = (R_{\max} - R)/w$, labor supply in the presence of our rudimentary fiscal system is then written:

$$h = \begin{cases} 0 & \text{if } w \leq w_A \\ \varphi[w(1-\tau), R_v] & \text{if } \varphi(w, R) \geq (R_{\max} - R)/w \\ \varphi(w, R) & \text{if } \varphi(w, R) \leq (R_{\max} - R)/w \end{cases}$$

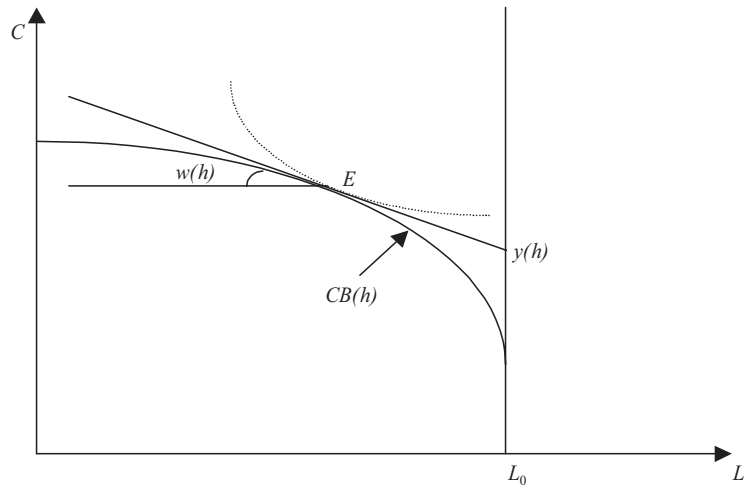

Figure 1.7
Differentiable budgetary constraint.

If we add other explanatory variables and a random term, which we have not done here so as not to burden the presentation, we arrive at an empirical model formally rather close to that described by equation (19). Here again, labor supply takes different values according to the values of the random term, and so can be estimated by the same methods as those envisaged above.

Approximation of the Budget Constraint by a Derivable Function

Another, more recent, method, relies on an approximation of the budget constraint by a derivable function (see, for example, MaCurdy et al., 1990, to see how such an approximation is constructed). The curve denoted $CB(h)$ in figure 1.7 represents a function of this type. Point E of this curve, where hours worked are equal to h, can be linked to a *virtual* wage, denoted $\omega(h)$, equal to the slope of the curve at this point, and a *virtual* non-wage income, denoted $y(h)$, corresponding to the intersection of the tangent of this curve with the vertical line with abscissa L_0. Note that this virtual wage and virtual income are ''observable'' by the econometrician from the moment he or she has been able to construct the curve $CB(h)$. All the optima of the consumer's program are then obtained by maximizing his or her utility under a (virtual) budget constraint written $C = \omega(h)h + y(h)$. For the interior solutions, the hours worked are then given by the implicit equation:

$$h = \varphi[\omega(h), y(h)]$$

This equation suggests a procedure for estimating labor supply: after having approximated the budget constraint by a derivable function, one ''observes'' the virtual wages and incomes and regresses the actual hours of work onto these virtual wages and incomes. Because these explanatory variables are manifestly not independent of hours worked, one has to resort to procedures utilizing instrumental variables.

This strategy, though simple in principle, poses problems owing to measurement errors that are almost always present in data relating to hours worked and wages. Thus, the dependent variable represents a priori the number of hours worked during the year—a piece of information that is rarely available. If, for example, we know the number of hours worked every week, then we multiply this figure by the number of weeks worked during the year. But this procedure is very arbitrary: in particular, it takes no account of voluntary or involuntary absences. As regards wages, the available data most often yield no more than a gross annual or monthly wage, when the explanatory variable that really counts ought to be the net hourly wage. Here again, the passage from the available data to the explanatory variable is a potential source of measurement errors (it should be noted that these problems of measurement errors extend to all the procedures by which labor supply is estimated, and not solely the one under study here). The upshot is that virtual wages and incomes are themselves the object of measurement errors. In these conditions, one solution lies in estimating a system of equations that takes the following form (see, for example, Bourguignon and Magnac, 1990):

$$h = \varphi(\omega, R, x_h, \varepsilon_h), \qquad \omega = \zeta(h, x_\omega, \varepsilon_\omega) \qquad \text{and} \qquad y = \xi(h, x_y, \varepsilon_y)$$

In this system, x_ω and x_y are two vectors of control variables that do not necessarily coincide with vector x_h of the control variables that appear in the equation defining labor supply. The random terms $(\varepsilon_h, \varepsilon_\omega, \varepsilon_y)$ capture the measurement errors and the nonobserved heterogeneity among individuals.

Having presented the problems encountered in the estimation of labor supply and the methods by which they can be solved, it is now time to examine the main results to which these estimates lead.

2.2 MAIN RESULTS

The econometric methods laid out above have made it possible to know better the value of the elasticity of labor supply. At the present time, the results obtained have converged toward a relative consensus. ''Natural experiments'' are another source of knowledge of the properties of the labor supply. The evolution of the amount of time worked and the participation rates fill out this factual panorama.

2.2.1 Form and Elasticity of Labor Supply

A consensus is emerging around the idea that movements in labor supply are principally owing to variations in the participation rate, and that the elasticity of the supply of female labor, especially that of married women, is greater than that of men.

The Hump-Shaped Curve

Does an individual's supply of labor take the form of a hump-shaped curve, as depicted in figure 1.4? The study of Blundell et al. (1992) suggests that it does. Using data from research on the expenditures of British families, these authors focus on a sample of single mothers, whose weekly supply of labor they estimate, distinguishing

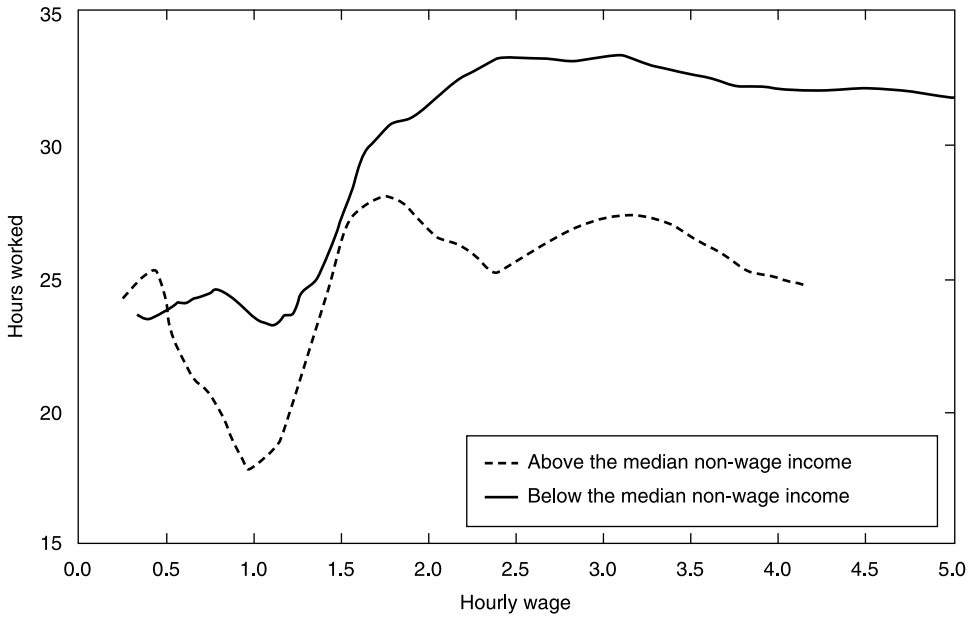

Figure 1.8
The labor supply of single mothers.

Source: Blundell et al. (1992).

between those who have non-wage income R greater than the median of the sample and those whose non-wage income is less than the median. The results of this study are represented in figure 1.8.

Scrutiny of this graph confirms, in the first place, that the hypothesis that leisure is a normal good is well-founded. We see that for practically all values of hourly wage, individuals in the sample who dispose of a non-wage income exceeding the median work less than the others. This graph also shows that the labor supply curve can indeed present a maximum (and even local maxima). Excluding wage values that are too low, we see that the labor supply curve for individuals whose non-wage income is less than the median strongly resembles the theoretical form of figure 1.4. For other individuals in the sample the resemblance is less marked, but the essential point remains: for low hourly wages (on the order of £1 to £1.5), there is little supply, and the substitution effect prevails, whereas for higher wages (from around £3 on up), the global income effect overrides the substitution effect.

The Elasticity of Labor Supply

The distinctive features and adaptations of the different fiscal systems found in different countries are often used to estimate the elasticity of labor supply of certain groups belonging to the population of working age (see Heckman, 1993, and Blundell and MaCurdy, 1999). These estimates run up against numerous difficulties. We have

Table 1.1
The elasticity of the labor supply of married women.

Authors	Sample	Uncompensated wage elasticity	Income elasticity
Hausman (1981)	U.S.	0.995	−0.121
Arrufat and Zabalza (1986)	U.K.	2.03	−0.2
Blundell et al. (1988)	U.K.	0.09	−0.26
Arellano and Meghir (1992)	U.K. (young children)	0.29	−0.40
Triest (1990)	U.S.	0.97	−0.33
Bourguignon and Magnac (1990)	France	[0.05; 1]	[−0.2; −0.3]

Source: Blundell and MaCurdy (1999, table 2, pp. 1649–1651).

already noted, for example, the need to distinguish clearly between decisions to participate, and decisions by people who already have a job about how many hours to work, and between hours freely supplied and ones that workers are forced to supply; and further, the complexity of budget constraints arising from different fiscal systems, the presence of fixed costs, the need to attribute a fictitious wage to nonparticipants, and so on.

Although the range of estimated elasticities is very broad, there is a relative consensus stressing the preponderance of variations in the participation rate over variations in hours. More precisely it is the variations in the rate of participation of a given group that explain the core of the elasticity of this group's labor supply. Another consensus emerges regarding the elasticity of labor supply by married women, which is demonstrably positive and greater than that of their spouses.

Table 1.1 furnishes some estimates obtained from empirical models utilizing methods set forth in sections 2.1.2 and 2.1.3. In this table, uncompensated elasticity designates the global effect of a wage change highlighted in equation (4) in section 1.2.2, that is, $(w/h^*)/(dh^*/dw)$. Income elasticity measures the impact of a change in income on labor supply, that is, with the notations in 1.2.2, $(R_0/h^*)(\partial h^*/\partial R_0) = (R_0/h^*)\Lambda_2$. Table 1.1 shows that the income elasticity of labor supply is negative, which means that leisure is a normal good (its consumption rises with income). Vice versa, wage elasticity is positive, so substitution effects prevail over income effects. Attention must be drawn to the large range of the estimates, however.

Table 1.2 shows that the wage elasticity of the labor supply is much weaker for married men, while income effects are, in general, more significant. If we turn to theoretical models, these results indicate that within the household, fiscal reforms affect principally the participation decisions of women, since on average they have access to lower wages than those of men and in all likelihood possess a comparative advantage when it comes to household production.

Table 1.2
The elasticity of the labor supply of married men.

Authors	Sample	Uncompensated wage elasticity	Income elasticity
Hausman (1981)	U.S.	[0; 0.03]	[−0.95; −1.03]
Blomquist (1983)	Sweden	0.08	[−0.03; −0.04]
Blundell and Walker (1986)	U.K.	0.024	−0.287
Triest (1990)	U.S.	0.05	0
Van Soest et al. (1990)	Netherlands	0.12	−0.01

Source: Blundell and MaCurdy (1999, table 1, pp. 1646–1648).

2.2.2 Natural Experiments

When a change is made to some aspect of economic policy, the econometrician has a chance to perform a ''natural experiment.'' The basic idea is to compare the reactions of a group affected by the change with those of another group having similar characteristics but that is untouched by the change. The second group is the ''control group.'' Changes in the fiscal system often provide a chance to apply this methodology to the study of labor supply behavior in a well-defined subpopulation (chapter 11, section 3, probes the question of the evaluation of economic policies in detail and problems arising from the utilization of the results of natural experiments; see also the surveys of Heckman et al., 1999, and Rosenzweig and Wolpin, 2000). Within the framework of a natural experiment, the effect of a change in economic policy is most often assessed with the help of an estimator called a ''difference-in-differences estimator.'' Blundell and MaCurdy (1999, section 5) have shown that this estimator corresponds, under certain conditions, to the estimator of ordinary least squares of a standard model with fixed individual effect. What follows derives from their work.

The Methodology of Natural Experiments

Let us take a population of individuals of size N, out of which a group of size N_M has been affected by a change in economic policy, while the control group of size N_C has not been so affected. Suppose that we want to find out the effects of this change on a variable y (for example, hours worked or participation in the labor market). Let us denote by y_{it} the observed value of this variable on an individual i at date t, and let us use δ_{it} to designate the dummy variable, which equals 1 if the policy change applies to individual i at date t, and 0 if it does not. We can then try to evaluate the impact of the policy by estimating the following equation:

$$y_{it} = \alpha\delta_{it} + x_{it}\theta + \gamma_i + \zeta_t + \varepsilon_{it} \tag{21}$$

Parameter α is an indicator of the impact of the change, γ_i is a fixed effect proper to individual i, ζ_t is a temporal effect proper to all agents, x_{it} is the vector of the

observable characteristics of individual i at date t, θ is a vector of parameters to be estimated, and ε_{it} designates an error term distributed independently among the individuals and also independent of γ_i and ζ_t.

Let us denote by Δ the difference operator; by definition $\Delta\kappa_t = \kappa_t - \kappa_{t-1}$ for any variable κ. When confronted with an equation like (21), the usual method consists of applying this operator to both sides of the equation in order to eliminate fixed individual effects γ_i. We thus obtain:

$$\Delta y_{it} = \alpha\Delta\delta_{it} + (\Delta x_{it})\theta + \Delta\zeta_t + \Delta\varepsilon_{it} \tag{22}$$

The general principles of econometrics with longitudinal data could be applied to equation (22), but the dummy variables δ_{it} have an interesting peculiarity that in certain cases lets us uncover simple expressions of the estimators. Let us therefore suppose that the observations concern only two periods. In period $(t-1)$ the same economic policy applies to all individuals, while in period t, economic policy is altered for individuals $i \in M$. For individuals $i \in C$ of the control group, there is no alteration. Since the model has only two periods, we can leave the time indexes out of equation (22). Let us suppose for simplicity's sake that individual characteristics do not vary ($\Delta x_i = 0$), and let us posit $\beta = \Delta\zeta_t$, $u_i = \Delta\varepsilon_{it}$. Equation (22) is now written:

$$\Delta y_i = \beta + \alpha\Delta\delta_i + u_i$$

By definition, the estimator of ordinary least squares of coefficient α is then given by:

$$\hat{\alpha} = \frac{cov(\Delta\delta, \Delta y)}{var(\Delta\delta)} = \frac{\sum_{i\in M}(\Delta\delta_i - \overline{\Delta\delta})(\Delta y_i - \overline{\Delta y}) + \sum_{i\in C}(\Delta\delta_i - \overline{\Delta\delta})(\Delta y_i - \overline{\Delta y})}{\sum_{i\in M}(\Delta\delta_i - \overline{\Delta\delta})^2 + \sum_{i\in C}(\Delta\delta_i - \overline{\Delta\delta})^2}$$

where $\overline{\Delta\delta}$ and $\overline{\Delta y}$ designate respectively the average values of $\Delta\delta$ and Δy. Since $\Delta\delta_i = 1$ for $i \in M$ and $\Delta\delta_i = 0$ for $i \in C$, after several simple calculations we get:

$$\hat{\alpha} = \frac{\sum_{i\in M}\Delta y_i}{N_M} - \frac{\sum_{i\in C}\Delta y_i}{N_C} \tag{23}$$

Estimator $\hat{\alpha}$ is called a ''difference-in-differences'' estimator. To construct it, we first calculate the average within each group of the differences between the dates $(t-1)$ and t of the dependent variable y, then we calculate the difference between these two averages. Its interpretation is very intuitive: if $\hat{\alpha}$ is equal to 0, that is because on average, the dependent variable y has undergone the same variations in the treated group (M) and in the control group (C). We may then conclude that the change of economic policy has had no effect. It is necessary, however, to look at the order of magnitude of $\hat{\alpha}$ carefully, for a change of economic policy often affects certain components of vector x_{it} of observed explanatory variables (for example, wages). It is also possible that the nonobserved heterogeneity included in the disturbance ε_{it} depends on variations in economic policy (for example, the entry into the labor market of less motivated persons may be favored by an increase in unemployment benefit). It is best,

Table 1.3
Participation rates of single women.

	Pre-TRA86	Post-TRA86	Difference	$\hat{\alpha}$
Treated group	0.729 (0.004)	0.753 (0.004)	0.024 (0.006)	
Control group	0.952 (0.001)	0.952 (0.001)	0.000 (0.002)	0.024 (0.006)

Standard errors in parentheses.
Source: Eissa and Liebman (1996, table 2).

therefore, to specify carefully the content of exogeneous variables in the estimation of equations grounded on natural experiments (the survey of Blundell and MaCurdy, 1999, clarifies in detail many points concerning the application of this methodology to labor supply; see also chapter 11, section 3, of the present work, which is dedicated to the problem of evaluating labor market policies and discusses the conditions under which the difference-in-differences estimator is valid).

Examples of Natural Experiments
Eissa and Liebman (1996) have studied the effects of the fiscal reform carried out in the United States in 1986 on labor force participation rates and hours worked.

The Tax Reform Act (TRA) of 1986 profoundly altered the system of earned income tax credits (EITC) by giving greater financial encouragement to take a low-wage job, but only to those with children in their care. To avoid difficulties arising from intrafamilial decisions (see section 1.2.2), Eissa and Liebman studied only single women. The control group therefore consisted of single childless women, while the treated group comprised single women with at least one child to care for. Eissa and Liebman (1996) then estimated the changes in the participation rate of each of these two groups. The data utilized were those of the March Current Population Survey for the years 1985–1991 (excluding 1987, which was considered the year of the change-over). The treated and control groups comprised respectively 20,810 and 46,287 individuals. The stages by which the difference-in-differences estimator $\hat{\alpha}$ was calculated are summarized in table 1.3.

The first two columns of table 1.3 represent the average of the participation rates for the periods 1984–1986 and 1988–1990, respectively. The third column shows, for each group, the difference between these averages after and before the reform. In this column, the figures 0.024 and 0.000 thus respectively represent the terms $(\sum_{i \in M} \Delta y_i)/N_M$ and $(\sum_{i \in C} \Delta y_i)/N_C$ of relation (23). The difference-in-differences estimator is then deduced and reported in column 4. In order to guarantee the robustness of their results, Eissa and Liebman then estimated an equation of the probit type analogous to (21). In their study, y_{it} is a dummy variable equal to 1 if person i has worked (for at least one hour) during period t, and equal to 0 if he or she has not. The dummy

variable δ_{it} is equal to 1 if person i is eligible for EITC during period t, and equal to 0 in all other cases; the term ζ_t is captured by the indicator variables relative to the years covered in the study; and vector x_{it} of observable characteristics contains indications of the number of children in school and not, the size of the family, level of education, age, and race. The estimation of this equation leads to the conclusion that single women caring for at least one child saw their probability of participating in the labor market grow, on average, by 1.9 percentage points (which is of the same order of magnitude as the 2.4 percentage points appearing in the third column of table 1.3). The further studies of Meyer and Rosenbaum (2000, 2001) on the same subject confirm the results of Eissa and Liebman (1996) and underline even more the importance of financial incentives in decisions to return to employment.

For France, an example of this approach grounded in a comparison between a treated group and a control group is the study of Piketty (1998) of the consequences of the extension of the parental education allowance (Allocation Parentale d'Education, or APE) starting in 1994. The APE is a monthly allowance of 3000 French francs (about 40% of the median wage) paid to a spouse who accepts leaving the workforce. Beginning in 1994, this measure was applied to families with two children (one of them under 3), whereas before that date a family had to have at least three children in order to be eligible. This ''natural experiment'' permits a precise analysis of the labor supply behavior of the subpopulation of mothers of two children (one of them under 3), taking as a control group the subpopulation of mothers with at least three children. Piketty (1998) shows that the fall in the participation rate, which was around 16% between 1994 and 1997 for the treated group, is entirely explained by the extension of the APE. He estimates that at least 35% of the mothers of young children would not have stopped working without this measure. The wage elasticity of the participation rate thus turns out to be particularly high for this category of women.

We may note that experiments can be carried out on purpose by the authorities, in which case we refer to ''social'' or ''controlled'' experiments. The Self-Sufficiency Project launched in the Canadian provinces of New Brunswick and British Columbia falls into this class. First, 6000 single parents who had been receiving only minimal social assistance for at least a year were selected at random. Then, from among these 6000, 3000 were picked at random and offered a bonus (amounting to around C$500 per month) which doubled the difference in disposable income between inactivity and employment if they found a full-time job. A year later more than 25% of the treated group were in full-time employment, as opposed to less than 11% of the control group (all the other results of the Self-Sufficiency Project can be found in Card and Robbins, 1996; see also Blundell et al., 1995, for studies of natural experiments in the United States and Great Britain).

To complete this rapid survey, we must note that natural experiments are not confined to the evaluation of public policies; they may also be applied to spontaneous events such as climate change. In this case economists sometimes speak of ''natural'' natural experiments (see Rosenzweig and Wolpin, 2000). In the domain of labor supply, studies have evaluated the consequences of meteorological change on the behav-

ior of farm families, while others have focused on the impact of children on the working lives of women: the treated group consisted of women who had had twins at their first childbirth, and the control group consisted of women who had had a single child at first childbirth. These studies generally bring out a negative effect of parenthood on labor supply by women.

Value and Limits of the Methodology of Natural Experiments

At first sight the methodology of natural experiments constrasts, by its simplicity, with the structural or econometric approach presented above, which consists of specifying a model and deriving from it equations that are estimated by an appropriate statistical method. The methodological simplicity of natural experiments is an undeniable advantage. Furthermore, this approach makes it possible rigorously to identify the consequences of a particular event, if it is properly conducted. But it has its limitations. For one thing, situations capable of giving rise to natural or controlled experiments are few. For another, each natural experiment constitutes, by definition, a very particular event, the consequences of which cannot be generalized into other contexts in the absence of theory. From this perspective, the structural approach and the methodology of natural experiments are complementary. The structural approach, starting from an explicit model and relying by definition on particular hypotheses, leads to the estimation of elasticities that allow us to evaluate the effects of numerous changes in the economic environment, the fiscal system for example, on behaviors and welfare. The structural approach is thus a valuable aid to decision-making in matters of public policy, since it has the power to *predict*, given well-defined hypotheses, the consequences of different public initiatives. The methodology of natural experiments assists in testing, a posteriori and in a particular context, the success of the predictions of the theoretical models and to some extent the impact of public policies.

2.2.3 Amount of Time Worked and Labor Force Participation Rate

The neoclassical model of labor supply discussed thus far throws light on significant shifts in participation rates, the amount of time worked, and the part-time work of women.

The Evolution of Participation Rates

Figure 1.9 traces the evolution of male and female participation rates in the United States labor market since 1947. The participation rate is equal to the ratio between the labor force (composed of employed workers and the unemployed) and the total population for the category concerned. This figure brings out an important characteristic of the industrialized countries as a group, which is the continuing rise in the participation rate of women for the last several decades. This rise is surely explained by the profound changes in our way of life, but it also corresponds to a steep rise in the wages available to women, accompanied by a fall in the relative price of goods that can replace household work (washing machines, child care, etc.). In these conditions we have seen that, in the model with household production, the substitution effect

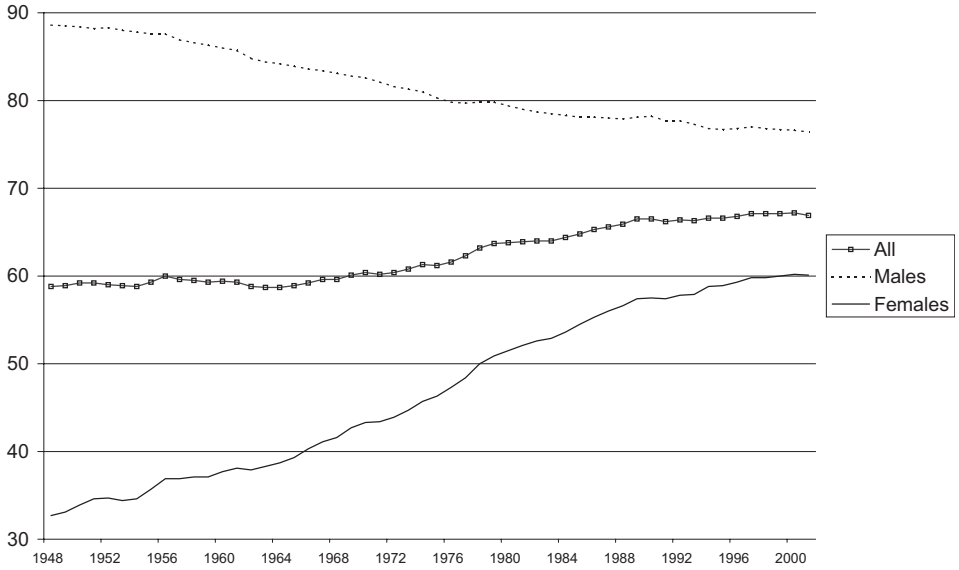

Figure 1.9
The evolution in participation rates in the United States for persons 16 years of age and older, 1948–2001.

Source: Bureau of Labor Statistics.

near the borderline of nonparticipation is very important, and induces a rise in participation in the labor market.

Figure 1.10 presents the evolution of participation rates for the whole of the population aged 15 to 64 in the United States, continental Europe (Germany, France, Italy), and Japan since the beginning of the 1960s. It is apparent that the participation rate of men has clearly diminished since the beginning of the 1960s in continental Europe and the United States. For example, it fell 17 points between 1960 and 2000 in the European countries and around seven points in the United States. On the other hand, the participation rate for women did not stop growing over the same period, having gained around seven points in the whole of the European Union and growing by more than 29% in North America. It should be noted that Japan forms an exception to the rule, inasmuch as its participation rates, both male and female, do not show a regular trend over this period. The male participation rate rose by 1.5 points, while for women it rose by five points. We also observe that, for the European countries, the contrary movements of the male and female participation rates approximately cancel each other out, and the total participation rate fell only slightly, by about two points. This observation does not apply to North America, where the very strong rise in the female participation rate has regularly caused the overall rate of participation to advance.

The data on labor force participation also confirm certain predictions of the model of the trade-off between consumption and leisure. Under the hypothesis that leisure is a normal good, we have seen that this model forecasts an increase in the

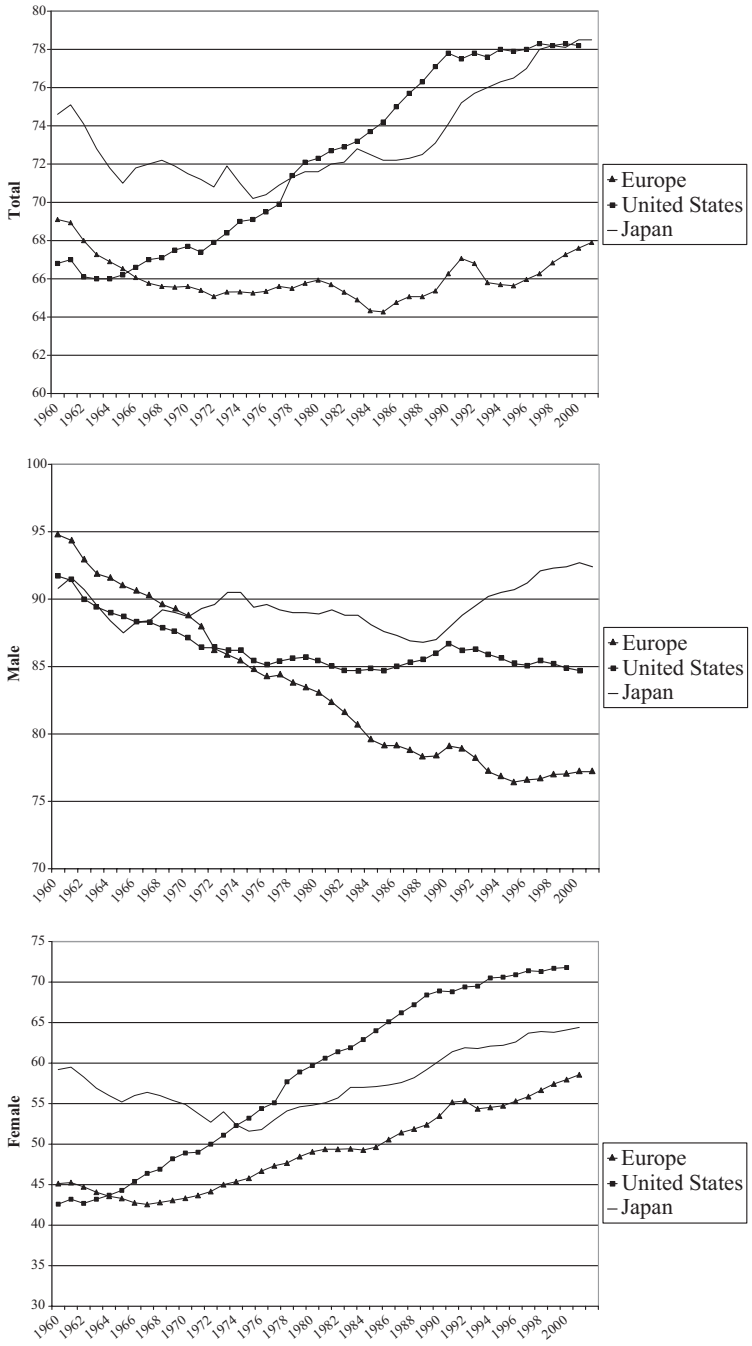

Figure 1.10
Participation rates in the United States, Europe (Germany, France, Italy), and Japan.

Source: OECD data.

Table 1.4

Participation rates of women classified by their marital status in the United States.

	Single	Married
1900	45.9	5.6
1950	53.6	21.6
1988	67.7	56.7
2001	78.7	69.6

Source: Ehrenberg and Smith (1994, table 6.1, p. 165) for 1900, 1950, and 1988, and Bureau of Labor Statistics for 2001.

reservation wage when the non-wage income of an individual climbs. Considering that within a couple, the non-wage income of one partner is often linked to the income of the other, the participation rate of married women ought to fall below that of single women. Table 1.4 shows that married North American women do in fact have a weaker rate of participation in the labor market than single women, even if the difference between these rates has a tendency to diminish over the long term. Additionally, empirical studies generally find that if a husband's income rises, his wife's labor supply falls off.

The Trend in the Amount of Time Worked

The long-term trend in the amount of time worked illustrates certain important characteristics of labor supply. Table 1.5 shows that labor productivity, which over the long term shapes the trend of real wages, has not stopped growing since the 1870s, though at a pace that varies at different times and in different countries. Production per hour worked was around 15 times greater in 1997 than in 1870 in Germany, France, and Sweden. It has multiplied by (only) six in the United States, and seven in the United Kingdom over the same period, since these two countries had much higher levels of productivity than the others at the end of the nineteenth century. In fact, before the agricultural and industrial revolutions, productivity had varied very little for several centuries. Likewise, until the industrial revolution, the amount of time worked probably remained stable, coinciding more or less with the hours of daylight. Subsequently, the onset of the industrial revolution saw longer hours: in the factories, we sometimes find that people were present at work for up to 17 hours per day. To work for 14 hours was normal, and a working day of 13 hours was considered short (Marchand and Thélot, 1997).

The historical movement in the amount of time worked can be grasped by using the same elements that allowed us to specify the form of the labor supply curve presented in figure 1.4. The substitution effect was probably prevalent for a few years during the economic take-off, as rural workers abandoned the countryside and went into the factories. But the number of hours worked rose so quickly, along with some growth in labor productivity, that the global income effect came to prevail. Hence the

Table 1.5

Hours worked annually per person and real hourly wages in the manufacturing sector.

	Amount of time worked				
	1870	1913	1938	1997	2000
Germany	2941	2584	2316	1507	1467
United States	2964	2605	2062	1850	1821
France	2945	2588	1848	1603	1532
United Kingdom	2984	2624	2267	1731	1711
Sweden	2945	2588	2204	1629	1603
	Wages				
Germany	100	185	285	1505	1569
United States	100	189	325	586	605
France	100	205	335	1579	1785
United Kingdom	100	157	256	708	819
Sweden	100	270	521	1601	1839

Source: Maddison (1995) for 1870, 1913, and 1938, and OECD data for 1997 and 2000.

diminution in hours of work after the industrial revolution can be interpreted as the consequence of an income effect due to a strong increase in the real wage.

Nevertheless, hours worked have undergone shifts less marked, and differing from one country to another, since the 1970s. In some countries the amount of time worked fluctuates, while it continues to shrink overall in others. Figure 1.11 shows that the annual amount of time worked has slightly increased in the United States and Sweden over this period, while it has diminshed in Germany, France, and the United Kingdom. These aggregate figures, which portray the global trend in the amount of time worked, are, however, difficult to interpret without further ado using the labor supply model, inasmuch as they result from different composition effects owing to important changes in the structure of the labor force by age and sex that vary from country to country.

Part-Time Work by Women

For the same amount of work, women's wages are generally noticeably lower than men's (see chapter 5). We have observed that when an individual decides to participate in the labor market, the number of hours that he or she wants to provide decreases with his or her non-wage income. Supposing that for a married woman, non-wage income often corresponds to her husband's income, the model immediately implies that women ought more frequently to be found in jobs with reduced hours than men. Table 1.6 indicates that this is indeed the case, for in the majority of the industrialized countries, women's share of part-time work often exceeds 70%. Of

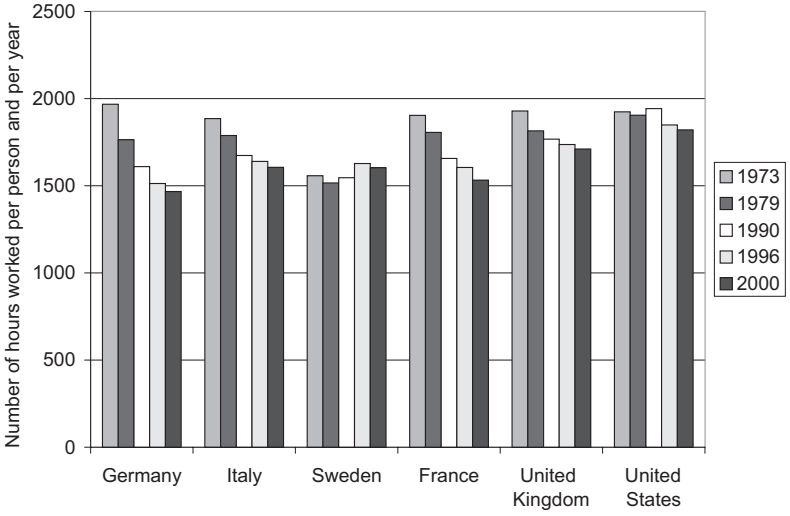

FIGURE 1.11
Amount of time worked annually in six OECD countries over the period 1973–2000 (total number of hours worked during the year divided by the average number of persons holding a job).

Source: OECD data.

Table 1.6
Women's share of part-time labor (in percentage terms).

	1979	1990	2000
Belgium	88.9	89.6	81.1
Canada	72.1	71.0	69.1
France	82.2	83.1	80.4
Germany	91.6	90.5	84.5*
Japan	70.1	73.0	67.5
Sweden	87.5	83.7	79.2
United Kingdom	92.8	87.0	79.9*
United States	68.0	67.6	67.5

Source: OECD data.
*1999.

course, other factors come into play to explain this state of affairs—in our day, household chores and the raising of children are still most frequently the tasks of women—but the value of women's relative wage must not be left out of account.

3 SUMMARY AND CONCLUSION

- According to the neoclassical theory of labor supply, every individual trades off between consuming a good and consuming leisure. The supply of individual labor is positive if the current wage exceeds the *reservation wage*, which depends on preferences and non-wage income. If labor supply is positive, the marginal rate of substitution between consumption and leisure is equal to the hourly wage.
- The relation between the individual supply of labor and the hourly wage is the result of combined substitution and income effects. The substitution effect implies an increasing relation between the wage and labor supply, while the income effect works in the opposite direction if leisure is a normal good. The supply of labor generally rises with the wage at low wage levels (the substitution effect prevails) and falls off when the wage reaches higher levels (the income effect prevails).
- In the neoclassical theory of labor supply, the labor force participation rate corresponds to the proportion of individuals whose reservation wage is less than the current wage. The fact that hours of work are offered to agents in indivisible blocks implies that the elasticity of the aggregate supply of labor may be very different from that of the individual supply of the majority of workers.
- When an individual has the opportunity to devote a part of his or her endowment of time to household production, at the optimum, the hourly wage is equal to the marginal productivity of household work. Household production increases the elasticity of the individual supply of wage work.
- As a general rule, the mechanism of substitution of leisure over time implies that the permanent component of the evolution of real wages has a feeble effect on labor supply, whereas the transitory component affects this variable more strongly.
- The elasticity of labor supply by women is, in general, greater than that of men, which is small. Moreover, variations in the total number of hours worked in an economy flow principally from variations in participation rather than from variations in hours worked by individuals.
- The methodology of natural experiments confirms the results of more traditional econometric studies, showing that financial incentives significantly influence labor supply by women.

- Finally, the neoclassical theory of labor supply permits the explanation of certain characteristics of long-term tendencies in amount of time worked and male and female participation rates.

Overall, the theory of labor supply sheds much light, often in agreement with empirical observations, on the manner in which agents decide how long to be active as wage-earners. It does not, however, allow us to understand why there should be unemployed people looking for work, since this category of the population has no reason to exist in a universe where information is perfect. The theory of the job search abandons the hypothesis of such a universe and succeeds in explaining the simultaneous presence of unemployed people and nonparticipants. It marks an important advance in the analysis of the functioning of the labor market, and forms the subject of the chapter 3.

4 RELATED TOPICS IN THE BOOK

- Chapter 2, section 2: Human capital and wage-earnings prospects
- Chapter 3, section 1: The choice between nonparticipation, job-search and employment
- Chapter 3, section 2.2: Optimal unemployment insurance
- Chapter 6, section 4: The relation between experience and wage
- Chapter 10, section 2.4: Migrations
- Chapter 11, section 3: The evaluation of active labor market policies

5 FURTHER READINGS

Blundell, R., and MaCurdy, T. (1999), ''Labor supply: A review of alternative approaches,'' in Ashenfelter, O., and Card, D. (eds.), *Handbook of Labor Economics*, vol. 3A, chap. 27, Amsterdam: Elsevier Science/North-Holland.

Heckman, J. (1974), ''Shadow prices, market wages and labor supply,'' *Econometrica*, 42(4), pp. 679–694.

Heckman, J. (1993), ''What has been learned about labor supply in the past twenty years?'' *American Economic Review, Papers and Proceedings*, 83(2), pp. 116–121.

Lumsdaine, R., and Mitchell, O. (1999), ''New developments in the economic analysis of retirement,'' in Ashenfelter, O., and Card, D. (eds.), *Handbook of Labor Economics*, vol. 3, chap. 49, pp. 3261–3307, Amsterdam: Elsevier Science/North-Holland.

6 APPENDICES

6.1 Appendix 1: Properties of Indifference Curves

If we suppose that the satisfaction of an agent increases with leisure and consumption, so that $U_C(C,L) > 0$, and $U_L(C,L) > 0$, the indifference curves are then negatively sloped. Consequently, the indifference curve associated with level of utility $\bar{U}$ is composed of the set of couples (C,L) satisfying $U(C,L) = \bar{U}$. This equality implicitly defines a function $C(L)$, which satisfies $U[C(L),L] = \bar{U}$. Differentiating this last expression with respect to L, we get:

$$C'(L) = -\frac{U_L(C,L)}{U_C(C,L)} < 0 \tag{24}$$

The indifference curves are indeed negatively sloped. We observe that the absolute value of the slope $C'(L)$ of an indifference curve is equal to the marginal rate of substitution U_L/U_C between consumption and leisure.

The hypothesis of the convexity of indifference curves is equivalent to the property of quasi-convexity of the utility function. Indifference curves are convex if and only if $C''(L)$ is positive. This second derivative is calculated using the equality $U(C,L) = \bar{U}$ and equation (24). We thus get:

$$C''(L) = \frac{U_L\left[2U_{CL} - U_{LL}\left(\frac{U_C}{U_L}\right) - U_{CC}\left(\frac{U_L}{U_C}\right)\right]}{(U_C)^2} \tag{25}$$

Since $C''(L)$ is of the sign of the term between square brackets of the numerator of the right-hand side of equation (25), the quasi-concavity of the utility function corresponds to the condition:

$$U(C,L) \text{ quasi-concave} \Leftrightarrow 2U_{CL} - U_{LL}\left(\frac{U_C}{U_L}\right) - U_{CC}\left(\frac{U_L}{U_C}\right) > 0 \tag{26}$$

6.2 Appendix 2: Properties of the Labor Supply Function

For an interior solution, relations (2) allow us to obtain the demand for leisure L^*. We thus have:

$$wU_C(R_0 - wL^*, L^*) - U_L(R_0 - wL^*, L^*) = 0 \tag{27}$$

This equation implicitly defines L^* as a function of $R_0 = wL_0 + R$ and of w. We denote this function $\Lambda(w, R_0) = L^*$. Its partial derivatives are obtained by differentiating equation (27), which implies:

$$dL^*(-w^2U_{CC} + 2wU_{CL} - U_{LL}) + dw[U_C - L(wU_{CC} - U_{CL})] + dR_0(wU_{CC} - U_{CL}) = 0 \tag{28}$$

By replacing the value w defined by (27), so that $w = U_L/U_C$ in (28), we get the expressions of the partial derivatives of function Λ:

$$\Lambda_1 = \frac{\partial L^*}{\partial w} = \frac{-L\left(\frac{U_{CL}U_C - U_{CC}U_L}{U_L}\right) - U_C\left(\frac{U_C}{U_L}\right)}{\left[2U_{CL} - U_{LL}\left(\frac{U_C}{U_L}\right) - U_{CC}\left(\frac{U_L}{U_C}\right)\right]} \tag{29}$$

$$\Lambda_2 = \frac{\partial L^*}{\partial R_0} = \frac{\frac{U_{CL}U_C - U_{CC}U_L}{U_L}}{\left(2U_{CL} - U_{LL}\left(\frac{U_C}{U_L}\right) - U_{CC}\left(\frac{U_L}{U_C}\right)\right)} \tag{30}$$

According to relation (26), the quasi-concavity of the utility function implies that the denominator of the right-hand side of equations (29) and (30) is positive. Λ_2 then has the sign of $U_{CL}U_C - U_{CC}U_L$. It is positive if and only if leisure is a normal good (L^* then grows with R_0). If Λ_2 is negative, leisure is an inferior good. Scrutiny of equation (29) shows that an increase in the wage entails an income effect (which we have described as indirect) and a substitution effect corresponding to the first and second terms in square brackets of the numerator of the right-hand side. If leisure is a normal good, $U_{CL}U_C - U_{CC}U_L > 0$, the two effects work in the same way and Λ_1 is negative. If leisure is an inferior good, Λ_1 has an ambiguous sign.

6.3 APPENDIX 3: COMPENSATED AND NONCOMPENSATED ELASTICITY

The Hicksian demand functions of leisure and of consumption good are obtained by minimizing the expenditures of the consumer under the constraint of a minimal exogenous level of utility, denoted $\bar{U}$. They are thus solutions of the problem:

$$\underset{(L,C)}{\text{Min}}\ C + wL \quad \text{subject to constraint} \quad U(C,L) \geq \bar{U} \tag{31}$$

Let us use $\hat{L}(w,\bar{U})$ and $\hat{C}(w,\bar{U})$ to designate the solutions of this problem; the expenditure function, denoted $e(w,\bar{U})$, is defined by the identity $e(w,\bar{U}) = \hat{C}(w,\bar{U}) + w\hat{L}(w,\bar{U})$. By construction, the Hicksian and Marshallian demand functions, respectively $\hat{L}(w,\bar{U})$ and $L^* = \Lambda(w,R_0)$, given by the equation (2), satisfy the identity $\Lambda[w,e(w,\bar{U})] = \hat{L}(w,\bar{U})$. If we derive this identity with respect to w, we get:

$$\Lambda_1[w,e(w,\bar{U})] + e_1(w,\bar{U})\Lambda_2[w,e(w,\bar{U})] = \hat{L}_1(w,\bar{U}) \tag{32}$$

We may point out that function $d(w) \equiv \hat{C}(x,\bar{U}) + w\hat{L}(x,\bar{U}) - e(w,\bar{U})$ reaches a minimum for $w = x$, which implies $d'(w) = 0$ for $w = x$, and thus $e_1(w,\bar{U}) = \hat{L}(w,\bar{U})$. In order to simplify these notations, let us simply use L and $h = L_0 - L$ to designate the solutions of problem (31). Multiplying both sides of relation (32) by w/h, we get:

$$\frac{w}{h}\Lambda_1 + \frac{wL}{h}\Lambda_2 = \frac{w}{h}\hat{L}_1 \tag{33}$$

Moreover, since $L^* = \Lambda(w,R + wL_0)$ and $\hat{L} = \hat{L}(w,\bar{U})$, the Marshallian and Hicksian elasticities of labor supply are respectively defined by:

$$\eta_w^{h^*} = -\frac{w}{h}\frac{\partial L^*}{\partial w} = -\frac{w}{h}(\Lambda_1 + L_0\Lambda_2) \qquad \text{and} \qquad \eta_w^{\hat{h}} = -\frac{w\hat{L}_1}{h} \tag{34}$$

Comparing (33) and (34), we finally arrive at the equality:

$$\eta_w^{h^*} = \eta_w^{\hat{h}} + \frac{wh}{R_0}\eta_{R_0}^{h^*} \tag{35}$$

In this expression $\eta_w^{h^*} = -R_0\Lambda_2/h$ represents the Marshallian elasticity of the labor supply with respect to potential income. Identity (35) is the Slutsky equation. It links the Hicksian elasticity $\eta_w^{h^*}$ (also called compensated elasticity) to the Marshallian elasticity $\eta_w^{\hat{h}}$ (also called noncompensated elasticity).

REFERENCES

Arellano, M., and Meghir, C. (1992), ''Female labour supply and on-the-job search: An empirical model using complementary data sets,'' *Review of Economic Studies*, 59, pp. 537–559.

Arrufat, J., and Zabalza, A. (1986), ''Female labour supply with taxation, random preferences, and optimization errors,'' *Econometrica*, 54, pp. 47–63.

Becker, G. (1965), ''A theory of the allocation of time,'' *Economic Journal*, 75, pp. 493–517.

Blanchard, O., and Fischer, S. (1989), *Lectures on Macroeconomics*, Cambridge, Mass.: MIT Press.

Blomquist, N. (1983), ''The effect of income taxation on the labour supply of married men in Sweden,'' *Journal of Public Economics*, 22, pp. 169–197.

Blundell, R., and MaCurdy, T. (1999), ''Labor supply: A review of alternative approaches,'' in Ashenfelter, O., and Card, D. (eds.), *Handbook of Labor Economics*, vol. 3A, chap. 27, Amsterdam: Elsevier Science.

Blundell, R., and Walker, I. (1986), ''A life cycle consistent empirical model of labour supply using cross section data,'' *Review of Economic Studies*, 53, pp. 539–558.

Blundell, R., Duncan, A., and Meghir, C. (1992), ''Taxation and empirical labour supply models: Lone parents in the UK,'' *Economic Journal*, 102, pp. 265–278.

Blundell, R., Duncan, A., and Meghir, C. (1995), ''Estimation of labour supply responses using tax policy reforms,'' IFS Working Paper 95/17.

Blundell, R., Meghir, C., and Neves, P. (1993), ''Labor supply: An intertemporal substitution,'' *Journal of Econometrics*, 59, pp. 137–160.

Blundell, R., Meghir, C., Symons, E., and Walker, I. (1988), ''Labour supply specification and the evaluation of tax reforms,'' *Journal of Public Economics*, 36, pp. 23–52.

Bourguignon, F., and Magnac, T. (1990), ''Labour supply and taxation in France,'' *Journal of Human Resources*, 25, pp. 358–389.

Brown, C. (1999), ''Early retirement windows,'' in Mitchell, O., Hammond, B., and Rappaport, A. (eds.), *Forecasting Retirement Needs and Retirement Wealth*, Philadelphia: University of Pennsylvania Press.

Browning, M., Bourguignon, F., Chiappori, P.-A., and Lechène, V. (1994), ''Income and outcomes: A structural model of intrahousehold allocation,'' *Journal of Political Economy*, 102, pp. 1067–1096.

Card, D., and Robbins, P. (1996), ''Do financial incentives encourage welfare participants to work? Initial 18 months findings from the Self-Sufficiency Project,'' Vancouver, B.C., Canada: Social Research Corporation.

Chiappori, P.-A. (1988), ''Rational household labor supply,'' *Econometrica*, 56(1), pp. 63–89.

Chiappori, P.-A. (1992), ''Collective labor supply and welfare,'' *Journal of Political Economy*, 100, pp. 437–467.

Ehrenberg, R., and Smith, R. (1994), *Modern Labor Economics*, 5th ed., New York: HarperCollins.

Eissa, N., and Liebman, J. (1996), ''Labour supply responses to the earned income tax credit,'' *Quarterly Journal of Economics*, 112(2), pp. 605–607.

Fortin, B., and Lacroix, G. (1997), ''A test of neoclassical and collective models of household labor supply,'' *Economic Journal*, 107, pp. 933–955.

Gronau, R. (1986), ''Home production,'' in Ashenfelter, O., and Layard, R. (eds.), *Handbook of Labor Economics*, vol. 1, chap. 4, Amsterdam: Elsevier Science/North-Holland.

Gronau, R. (1997), ''The theory of home production: The past ten years,'' *Journal of Labor Economics*, 15(2), pp. 197–205.

Gruber, J., and Wise, D. (eds.) (1999), *Social Security and Retirement Around the World*, Chicago: University of Chicago Press.

Gruber, J., and Wise, D. (2001), ''An international perspective on policies for an aging society,'' NBER Working Paper 8103.

Gustman, A., and Steinmeier, T. (1986), ''A structural retirement model,'' *Econometrica*, 54(3), pp. 555–584.

Gustman, A., Mitchell, O., and Steinmeir, T. (1994), ''The role of pensions in the labor market,'' *Industrial and Labor Relations Review*, 47(3), pp. 417–438.

Hall, R. (1980), ''Labor supply and aggregate fluctuations,'' in Brunner, K., and Meltzer, A. (eds.), *On the State of Macroeconomics*, Carnegie-Rochester Conference Series on Public Policy, Amsterdam/North-Holland.

Hall, R. (1999), ''Labor market frictions and employment fluctuations,'' *Handbook of Macroeconomics*, vol. IB, chap. 17, Amsterdam: Elsevier Science/North-Holland.

Hausman, J. (1981), ''Labour supply,'' in Aaron, H., and Pechman, J. (eds.), *How Taxes Affect Economic Behaviour*, Washington, D.C.: Brookings Institution.

Heckman, J. (1974), ''Shadow prices, market wages and labor supply,'' *Econometrica*, 42(4), pp. 679–694.

Heckman, J. (1993), ''What has been learned about labor supply in the past twenty years?'' *American Economic Review, Papers and Proceedings*, 83(2), pp. 116–121.

Heckman, J., Lalonde, R., and Smith, J. (1999), ''The economics and econometrics of active labor market programs,'' in Ashenfelter, A., and Card, D. (eds.), *Handbook of Labor Economics*, vol. 3, chap. 31, pp. 1865–2097, Amsterdam: Elsevier Science/North-Holland.

Hotz, V., Kydland, F., and Sedlacek, G. (1988), ''Intertemporal substitution and labor supply,'' *Econometrica*, 56, pp. 335–360.

Lucas, R., and Rapping, L. (1969), ''Real wages, employment and inflation,'' *Journal of Political Economy*, 77, pp. 721–754.

Lumsdaine, R., and Mitchell, O. (1999), ''New developments in the economic analysis of retirement,'' in Ashenfelter, O., and Card, D. (eds.), *Handbook of Labor Economics*, vol. 3, chap. 49, pp. 3261–3307, Amsterdam: Elsevier Science/North-Holland.

Lumsdaine, R., Stock, J., and Wise, D. (1990), ''Efficient windows and labor force reduction,'' *Journal of Public Economics*, 43, pp. 131–159.

Lundberg, S. (1985), ''The added worker effect,'' *Journal of Labor Economics*, 3, pp. 11–37.

MaCurdy, T., Green, D., and Paarsch, H. (1990), ''Assessing empirical approaches for analyzing taxes and labour supply,'' *Journal of Human Resources*, 25, pp. 415–490.

Maddison, A. (1995), *The World Economy, 1820–1992*, Paris: OECD.

Marchand, O., and Thélot, C. (1997), *Le travail en France (1800–2000)*, Paris: Nathan.

Mas-Colell, A., Whinston, M., and Green, J. (1995), *Microeconomic Theory*, Oxford, U.K.: Oxford University Press.

Meyer, B., and Rosenbaum, D. (2000), ''Making single mothers work: Recent tax and welfare policy and its effects,'' *National Tax Journal*, 53, pp. 1027–1162.

Meyer, B., and Rosenbaum, D. (2001), ''Welfare, the earned income tax credit, and the labor supply of single mothers,'' *Quarterly Journal of Economics*, 116, pp. 1063–1114.

OECD (1991), *Employment Outlook*, Paris.

OECD (1995), *Employment Outlook*, Paris.

OECD (1999), *Employment Outlook*, Paris.

Pencavel, J. (1986), ''Labor supply of men: A survey,'' in Ashenfelter, O., and Layard, R. (eds.), *Handbook of Labor Economics*, vol. 1, pp. 3–102, Amsterdam: Elsevier Science/North-Holland.

Piketty, T. (1998), ''L'impact des incitations financières au travail sur les comportements individuels: Une estimation pour le cas français,'' *Economie et Prévision*, 132–133, pp. 1–36.

Rosenzweig, M., and Wolpin, K. (2000), ''Natural 'natural experiments' in economics,'' *Journal of Economic Literature*, 38, pp. 827–874.

Stock, J., and Wise, D. (1990), ''Pension, the option value of work and retirement,'' *Econometrica*, 58(5), pp. 1151–1180.

Triest, R. (1990), ''The effect of income taxation on labor supply in the US,'' *Journal of Human Resources*, 25, pp. 491–516.

Van Soest, A., Woittiez, I., and Kapteyn, A. (1990), ''Labor supply, income taxes and hours restrictions in the Netherlands,'' *Journal of Human Resources*, 25, pp. 517–558.

Varian, H. (1992), *Microeconomic Analysis*, 3rd ed., New York: W.W. Norton.