chapter 28

Labor Supply over the Life Cycle

The neoclassical labor-supply model, presented in Chapter 4, offers valuable insights into the effects of changes in the hourly wage, the level of unearned income, and an assortment of welfare policies on the number of hours that people work. The framework also sheds light on two theoretical concepts of fundamental importance: the *income* and *substitution* effects. Nevertheless, despite its strengths, the approach suffers from several weaknesses. From the perspective of this chapter, an especially pressing one is that it is *static* in nature: individuals maximize their utilities over a single time period.[1] Of course, in reality, not only do people make a series of decisions over their life cycles but their behavior also responds to the evolution of their economic circumstances. For example, most people tend to work during their most productive middle-aged years and retire when they are old.

In response to these shortcomings, economists have increasingly eschewed simple static models in favor of dynamic ones in which individuals—just like you and I—make a sequence of decisions over time. As we will show in this chapter, the approach has revolutionized the way economists think about labor supply (and consumption) choices over the life cycle.

In Sections 28.1 and 28.2, we begin by outlining the basic theory. The principal result that emerges from the analysis is the *intertemporal substitution hypothesis* (ISH), which is captured by the aphorism: *make hay when the sun shines*. In other words, workers are predicted to adapt their labor-supply decisions to best take advantage of changes in their environments. In Section 28.3, we show how the dynamic labor-supply model has been applied to shed light on two of the most fundamental social and economic developments in contemporary U.S. history: the *decrease* in the labor-force participation rates of men, and the stunning *increase* in the participation rates of women—especially married

LEARNING OBJECTIVES

By reading this chapter, you should be able to:

- Recognize the strengths and weaknesses of the static labor-supply model presented in Chapter 4.
- Understand the core principles that underlie decision making over the life cycle and be able to explain why permanent and temporary shocks can elicit such dramatically different behavioral responses.
- Explain the meaning of the intertemporal elasticity of substitution (IES).
- Understand the root causes of the profound shifts in male and female participation rates that have occurred over the last 70 years or so.
- Understand why the supply of labor moves in tandem with the pulse of economic activity that occurs over the course of the business cycle.

women. Finally, Section 28.4 examines how the framework can be used to explain how the aggregate supply of labor responds to the pulse of economic activity that occurs over the course of the business cycle.

28.1 Theoretical Issues

The simplest way of modeling the dynamic choices of men and women would be to repeatedly apply the simple static labor-supply model to each period of their lives. One would analyze their behavior when they are 16, and then 17, and so on up until they reach 99 (or some other plausible endpoint). This approach, however, is reasonable only if people are—in a decision-making sense—*myopic*; they must ignore the consequences of their current actions (how much to consume, how hard to work and so on) on their future well-being.

Nevertheless, a virtual bounty of evidence suggests that individuals make decisions in the present with an eye on their future consequences. Examples include investments in education (which increase future earnings), pensions (which provide a future source of income), and insurance (which offers relief in the event of some unfortunate future mishap). It is therefore disquieting, then to say the least, to assume away forward-looking behavior when it seems to matter the most. After all, for many workers, this week's choices of how much to consume and how many hours to work have an immediate, ineluctable, and arithmetic impact on what is available to them next week.

The Foundations

Because a fully fledged formal analysis of intertemporal decision making is quite technical, we shall focus on just the essentials of the subject.[2] The following three points appear to be essential for thinking about how people behave over time:

- **H1. Current actions** Workers recognize **current** actions have **future** consequences.
- **H2. Future events** More subtly, workers take into account the **future** when making their **current** decisions. For example, there is little doubt that your own current behavior would be affected were you to suddenly discover that you will receive $20 million in 5 years' time because of the generosity of a long-lost aunt or uncle.
- **H3. Uncertainty** Related to **H2**, the future is inherently uncertain. Absent a handy crystal ball, we simply do not know with any certainty what next month, next year, 10 years from now will bring. As a result, the decisions we make today are shrouded in a cloud of doubt, risk, and uncertainty concerning what the future holds for us. To mitigate these effects, workers will try to use the information they do have at their disposal in the most effective manner.

It is a great testimonial to the power of economics to learn that, over the past 40 years or so, great progress has been made in cleverly weaving all of these elements into a coherent view of intertemporal behavior. The models that have emerged not only have offered important theoretical insights into human behavior but also have provided an excellent forum for conducting rigorous empirical work.[3]

Foremost among the theoretical insights is the *intertemporal substitution hypothesis* (ISH) that says individuals substitute leisure and consumption across different periods of their lives to make the most expedient use of any changes in their economic circumstances.

Perfect Foresight

To understand the essential features of intertemporal decision making, it is helpful to begin by exclusively focusing on hypotheses **H1** and **H2**. (The effects of uncertainty, **H3**, are discussed in Section 28.2.) With this in mind, consider a 20-year-old worker who, as luck would have it, is endowed with a handy crystal ball that gives him **perfect foresight** about how his future will unfold.

Suppose his crystal ball tells him that, after he has turned 21, he has $S = 70$ remaining years to live and that his potential hourly wage profile, w_t, is given by Figure 28.1.[4] His wage profile exhibits a classic cap-shaped ($\cap$) pattern, which is explained as follows. When he is young, his wage is initially low because he has very little labor-market experience and few skills that employers value. Over time, however, he acquires experience and his skill set grows, which explains why his wage increases along the segment *AB*. Eventually (so the story goes), his age begins to tell, and his wage begins to decline (see point *C* in the figure). Finally, the individual shuffles off his mortal coil at the grand old age of $S + 20 = 90$.

Given this wage profile and the assumption of perfect foresight, what is the individual's optimal lifetime plan? To provide a detailed answer to this question it is first necessary to describe his lifetime objective and the lifetime constraints that limit his actions. This is done later, but the reader who is pressed for time can skip ahead to the subsection titled *life cycle behavior,* which discusses the results.

FIGURE 28.1 The Age-Earnings Profile

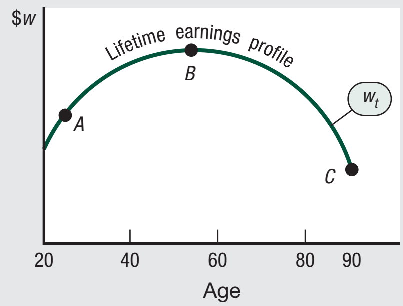

The Objective. Suppose that within each of the $t = 0, 1, 2, \ldots, S$ periods (years) of his life he derives utility $u_t = u(c_t, \ell_t)$ from consumption, c_t, and from leisure, ℓ_t. Assume further that $u(\cdot)$ is characterized by standard indifference curves of the sort depicted in Figure 4.3.

The individual derives utility during each period of his life. Hence the arguments presented in Appendix C suggest that his lifetime utility, denoted U_0, equals the discounted sum of the utilities available in each period:

$$U_0 = \underbrace{u_0}_{present} + \underbrace{\delta u_1 + \delta^2 u_2 + \ldots + \delta^S u_S}_{future} \tag{28.1}$$

where $0 \leq \delta \leq 1$ is his *discount factor*, which provides a numerical representation of his impatience. If $\delta < 1$, then the utility that accrues sooner is more valuable to him than the utility that accrues later. Notice that his lifetime utility, U_0, consists of two separate parts: his current utility, u_0, plus the present discounted value of his *future utility stream*, which is denoted $\delta u_1 + \delta^2 u_2 + \ldots + \delta^S u_S \equiv \sum_{t=1}^{t=S} \delta^t u_t$. It is immediate from 28.1 that to maximize his lifetime utility, U_0, not only must the individual evaluate the effects of his actions on his current utility but he must also anticipate their effects on his future well-being.

The formulation also makes quite clear the *intratemporal* (within period) and *intertemporal* (across period) utility-substitution possibilities that are available to him. Notice that, within a given period t, his lifetime utility remains unchanged if he substitutes consumption for leisure along any given indifference curve. Likewise, his lifetime utility, U_0, remains unchanged if his current utility, u_0, decreases by x utils, provided his discounted future utility, $\sum_{t=1}^{t=S} \delta^t u_t$, increases by x utils.

A lifetime plan for the individual is just a list (sequence) of consumption–leisure pairs for each of the remaining years of his life: $(c_0, \ell_0), (c_1, \ell_1), \ldots, (c_s, \ell_s)$. The individual's goal is to choose the particular plan that maximizes his lifetime utility U_0.

The Constraints. Naturally, not every potential consumption–leisure plan is feasible. For instance, few of us can afford indolent lives in which we spend \$2 million a year sunning ourselves in the proverbial tropical paradise. Hence we must describe the constraints that limit the individual's choices. With this in mind, let $r > 0$ denote the rate of interest. For the moment, suppose that his lifetime wealth, Ω_0, is exogenously given. The individual's **lifetime budget constraint** is:

$$\Omega_0 = c_0 + \sum_{t=1}^{t=S} \frac{c_t}{(1+r)^t} \tag{28.2}$$

This equation has a surprisingly simple interpretation: it says that he cannot consume more than he has. The left-hand side represents the funds that are available to him, and the right-hand side describes the use of these funds. Notice that all future consumption values are properly discounted to determine their present values.

Equation 28.2 embodies the core intertemporal consumption trade-offs that are available to the individual. For example, it is easy to see that, given Ω_0, a \$1 increase in his current consumption, $\$c_0$, must call for an equal and opposite

\$1 reduction in the present value of his future consumption (i.e., over the periods $t = 1, 2, \ldots, S$). Hence current actions have future consequences. Notice too that because of the interest that accrues on savings, current consumption can be very expensive when it is measured in terms of forgone future consumption opportunities.

To better see this, suppose that the individual increases his current consumption by \$1, and that he plans to compensate for the change by reducing his consumption in (say) 25 years. By how much must he actually reduce c_{25}—denoted Δc_{25}—to make up for the shortfall? To answer this, note that the present value of his proposed change is $\Delta c_{25}/(1 + r)^{25}$ and that to balance his lifetime budget this change must equal \$1. Hence $\Delta c_{25} = (1 + r)^{25}$. In the United States, the interest rate tends to fluctuate around 5%. If $r = 5\%$, then $\Delta c_{25} = (1.05)^{25} \approx \$3.39$. Thus \$1 of consumption today costs the individual a remarkable \$3.39 of consumption 25 years from now.

In reality, the individual's wealth, Ω_0 is, of course, not exogenously given but instead depends on the number of hours he works each period, h_t:

$$\Omega_0 \equiv A_0 + w_0 \cdot h_0 + \sum_{t=1}^{t=S} \frac{(w_t \cdot h_t)}{(1+r)^t} \tag{28.3}$$

This says his wealth, Ω_0, equals his initial wealth A_0 (which also includes any unearned income); plus his first-period earnings, $\$w_0 \cdot h_0$; and, finally, plus the discounted sum of his future earnings: $\sum_{t=1}^{t=S} (w_t \cdot h_t)/(1 + r)^t$. Notice that dollars received in the future are properly discounted to determine their present values.

Equation 28.3 points to myriad work–leisure substitution possibilities that are available to the worker. For example, suppose that he decides to enjoy an extra leisure hour when he is 20 years old and plans to make up for the loss in earnings by working an extra Δh_{45} hours in 45 years time (when he is 65 years old)—rather than by altering his consumption plan $c_0, c_1, \ldots, c_S$. If (say) $\$w_0 = \$w_{45} = \$15$ per hour, how many additional hours must he work? Since $\$w_0 = \$15$, the additional youthful hour of leisure blows a \$15 hole in his lifetime budget. If he works an additional Δh_{45} hours when he is 65, then his actual earnings increase by $\$15 \times \Delta h_{45}$, but the present value of his earnings increase by only $(\$15 \cdot \Delta h_{45})/(1 + r)^{45}$. Consequently, to make up for the \$15 shortfall, he must pick Δh_{45} to ensure that $\$15 \equiv (\$15 \cdot \Delta h_{45})/(1 + r)^{45}$, implying $\Delta h_{45} = (1 + r)^{45}$. If we again take an interest rate of 5%, then $(1 + r)^{45} \approx 9$. This implies that he must forgo a remarkable 9 hours of leisure, when he is 65 to compensate for the loss income resulting from an extra leisure hour at the tender age of 20. From this perspective, an indolent youth is very costly indeed.

In fact, given that enjoying consumption and leisure early in life are so expensive, the reader may be forgiven for wondering why people don't behave as workaholic misers during their youths. Surely, this strategy would allow them each to enjoy a superb old age. Indeed, the wonders of compound interest do tell us that hard graft today may well translate into fabulous opportunities in the future. However, the individual's objective is to maximize his lifetime utility, which depends upon the discounted utility available in each and every period of his life. Consequently, he might place a relatively low weight on the utility available in the distant future because of his impatience. This creates a fantastic decision-making tension that is so admirably captured by the model. On the one hand, his impatience encourages him to seek immediate gratification. On the other, the accrual of interest encourages him to defer his enjoyment of utility to the future. His optimal behavior is guided by judiciously balancing these opposing forces.

Life Cycle Behavior

Given the individual's initial wealth, A_0, and the wage rates $w_0, w_1, \ldots, w_s$, suppose that he formulates an optimal lifetime consumption and leisure plan, and suppose that his maximized lifetime utility is U_0^*. Let us now examine the implications of this optimal plan for his life cycle behavior.

In Figure 28.2, we depict the individual's lifetime wage profile by the solid black line and his predicted optimal labor-supply behavior by the solid green one. (For the moment, disregard the dashed lines.) To understand why he behaves in this manner, consider the segment *AB* of his life cycle, along which he is enjoying substantial wage growth. At first glance, it might seem that understanding his behavior in this region is as easy as pie—calling for nothing more than the elementary application of the methods set out earlier in Chapter 4. In fact, it is tempting to suspect that the increase in the wage, over the region *AB*, is associated with the standard *income* and *substitution* effects that were described at length in that chapter. Consequently, whether the individual supplies more or less labor simply depends on the relative strengths of these two opposing forces.[5] Nevertheless, this conclusion is **incorrect**: the (anticipated) increase in the wage along *AB*—which is called an *evolutionary wage change*—is associated with an income effect of precisely zero! This result often flabbergasts the unsuspecting.[6]

FIGURE 28.2 The Life Cycle

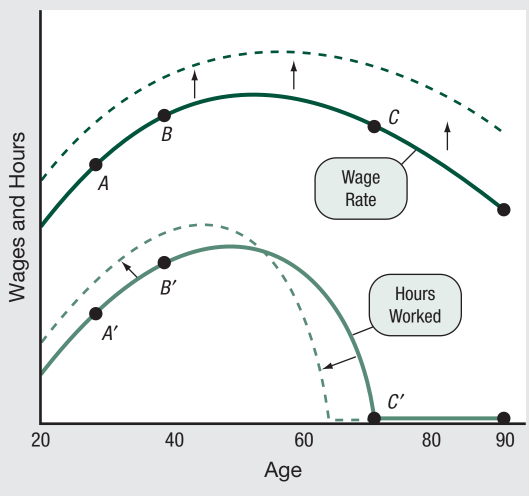

The explanation hinges on the following observation: the income effect is unleashed by a change in an individual's economic environment only if it affects his *maximized* utility. In the case of the static theory of labor supply, an increase in the wage has an income effect precisely because it enhances the opportunities available to the worker and increases his utility.[7] Contrast this situation with the life cycle model of labor supply. The individual's maximized lifetime utility, U_0^*, depends on the interest rate, his initial wealth, and the entire sequence of wages $w_0, w_1, w_2, \ldots, w_s$. Consequently, as the wage increases along the segment AB this *cannot* increase U_0^*, for the simple reason that the projected wage increase is already included as a determinant of U_0^*. Hence there is no income effect because the individual has already accounted for the movement along AB as part of his initial economic opportunities.

Absent the income effect, the change in the wage along the segment of the curve AB is associated with a *pure substitution effect*.[8] In fact, without the work-retarding income effect, the work-enhancing substitution effect is given free rein to influence his labor supply. This is the essence of the intertemporal substitution hypothesis: individuals choose the timing of their leisure so as to take full advantage of the evolution in their economic circumstances. In particular, the individual will enjoy lots of leisure when it is relatively cheap (i.e., when the wage is low, such as at point A) and by working hard when leisure is expensive (i.e., when the wage is high, such as at point B). Together, these observations explain the steeply increasing segment of the solid green hours-of-work schedule $A'B'$.

Parallel arguments explain the general pattern of the hours he works over his entire life cycle, including the declining region depicted in the figure. Thus the ∩-shaped wage profile is predicted to induce a corresponding ∩-shaped labor-supply pattern as represented by the green curve. He increases his labor supply as the wage rises, until it reaches its peak. Eventually, his wage begins to decline, and he responds by increasing his enjoyment of leisure and working fewer hours. As shown, he retires (dropping out of the labor force altogether), once his wage falls to the critically low level depicted at point C.

The reader may be forgiven for wondering if the income effect has been lost forever. It has not. Remember, the essence of the income effect: it emerges whenever there is a change in the environment that alters the individual's maximized utility. An increase in the individual's initial wealth, A_0, is precisely such a change and, in fact, it is associated with a *pure wealth effect*. Because leisure is a normal good, this change would induce the individual to reduce his labor supply at each point over his life cycle.

Likewise, an upward shift in the age-earnings profile (depicted by the dashed black line in the figure) unleashes both income and substitution effects. The income effect encourages the individual to work fewer hours at each point over his life cycle because leisure is a normal good. The substitution effect captures the increased opportunity cost of leisure and so has precisely the opposite effect. The outcome of the tug-of war between these forces determines the overall impact on the number of hours the individual works. It is conceivable that the relative strength of the two effects may vary over the course of the worker's life cycle. Hence, as illustrated by the dashed-green curve, the upward shift in the wage profile may result in an increase in the hours he works over some periods of his life and a reduction over others.

Applications

Once one accepts the worldview that anticipated future events affect current actions, then certain other implications are simply inescapable. For instance, Ziliak and Kniesner (1999) study how the income tax structure affects life cycle labor-supply decisions. These tax effects arise because most people begin their careers earning a low wage (and are in a low tax bracket) and end up earning a higher wage (and move to a higher tax bracket). According to the dynamic theory of labor supply, individuals are predicted to anticipate this pattern and to adjust their behavior accordingly. The authors find that these tax effects have a substantial impact on the life cycle supply of labor. In fact, they estimate that the tax reforms that occurred during the 1980s, which lowered marginal income tax rates—increased male labor supply by almost 3%.

Trostel (1993) examines the effect of income taxes on human-capital investments, in the context of a life cycle model. He finds that a 1% increase in the income tax rate reduces the stock of human capital by almost 1%. The primary reason is that a tax increase reduces the supply of labor over the course of the life cycle, which then reduces the present value of the returns to investments in education. Finally, one of the most important facets of the theory is that it suggests that both pension reforms and the anticipated crisis in Social Security funding will telescope back to the present, affecting the behavior of workers even if they are many years away from retirement.

28.2 Uncertainty

So far, we have assumed that the individual is endowed with a handy crystal ball that he uses to perfectly foretell his future (in particular, the sequence of wages over the course of his known life span). Yet, common sense tells us that life is not so neat (see **H3** on page 2). For most, it is capricious, involving its fair share of random ups and downs. In this section, let's examine how the individual behaves in such an environment.

Random Shocks

Suppose that we return the individual's crystal ball to him but muddy it so that he can just see the wage he can expect to earn at each point over his life cycle. His actual wage may lie above or below its expected value for a variety of different reasons.[9] Negative *shocks*, which reduce his actual wage below its expected value, can result from

- An illness, which reduces his productivity.
- A technological change, which reduces (or renders obsolete) the value of his skill set (e.g., the invention of the Internet and the associated growth of

outsourcing implies that now is probably not the best time to work in a U.S. call center).

- The removal of a trade barrier, which adversely impacts his occupation or industry.
- A recession, which reduces the value of his labor to employers.

Of course, by going into reverse gear, each of these shocks can also raise an individual's hourly wage above its expected value. Moreover, the shocks may either be **permanent** (chronic) or **temporary** (acute) in nature. As such, they may have a permanent or temporary effect on his behavior.

The extension of the life cycle model to include random shocks of the sort just described, adds great richness to the approach. It also raises several interesting questions: (a) What is the individual's objective? (b) How do individuals respond to temporary shocks? (c) How do they respond to permanent ones? In regard to the first question, economists envision the individual as choosing a lifetime plan to maximize his average—or expected—discounted lifetime utility. Most important, the individual's optimal plan is now made contingent upon the events that unfold over the course of his life. For example, if he subsequently discovers that a sudden upswing in economic activity has increased his hourly wage above its expected value, then his plan tells him precisely what to do in this situation. Likewise, it also tells him what to do if, instead, he experiences a negative shock that adversely affects his wage.

Permanent and Temporary Shocks

To fix ideas about the differential effects of temporary and permanent shocks on optimal behavior, suppose that an individual is employed by a firm that experiences a sudden temporary increase in the demand for its product. For instance, one of its primary customers might have placed a large order that has a tight 3-month completion deadline. Assume that, during this period, the firm offers a generous wage \$$w_1$, to ensure the project is completed in a timely manner.

The (unanticipated) wage increase unleashes both income and substitution effects. However, the income effect is liable to be very modest because it applies to only 3 months out of the worker's entire career. This leaves the substitution effect, which encourages the worker to temporarily increase the number of hours he works (as his labor is most valuable during this time) and to reduce his labor supply in the future (where it is less so). Hence the worker intertemporally substitutes his labor across the different periods of his life and, in essence, *makes hay when the sun shines*.

The idea is depicted in Figure 28.3. The wage is initially constant and equal to w_0 in every period of the worker's life. At this wage, the worker optimally works for h_0 hours. The *temporary* increase in the wage to w_1, however, has a tremendous impact on the relative value of working during the 3-month period. As shown,

FIGURE 28.3 Intertemporal Substitution

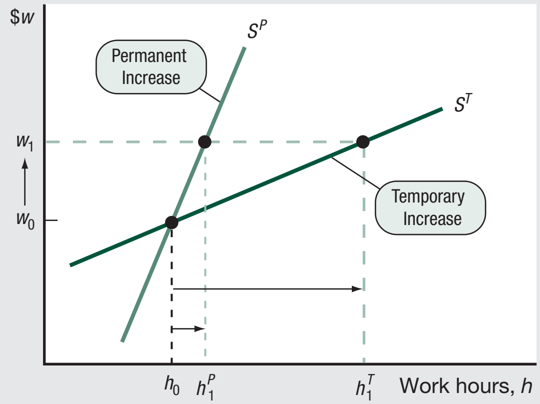

the worker responds accordingly via a substantial *temporary* increase in his labor supply from h_0 to h_1^T.

It is instructive to contrast this situation with one in which the wage increase is permanent. Continuing with the example, suppose instead that there is a sudden permanent increase in the demand for the entire industry's product, which causes the wage to permanently increase from $\$w_0$ to $\$w_1$. The size of the substitution effect is the same as for the temporary increase: it depends on the relative wage increase $\$(w_1 - w_0)$. This time, however, the offsetting, work-retarding, income effect is now much stronger since the increase persists over his remaining career. As shown, because of the conflict between these two forces, the permanent increase in the wage has only a modest effect on the worker's labor supply, which increases from h_0 to h_1^P. This important finding is summarized in Major Result 28.1.

MAJOR RESULT 28.1

Temporary vs. Permanent Shocks

The individual's supply of labor is much more responsive to temporary wage changes than to permanent ones.

The Intertemporal Elasticity of Substitution (IES)

An individual's willingness to intertemporally substitute leisure over time is captured by the **intertemporal elasticity of substitution** (IES). The IES measures the strength of the substitution effect. It equals the percentage change in leisure induced by a temporary 1% increase in the wage. As we shall see, in Section 28.4, a sizable value of the IES is essential for the empirical relevancy of an important branch of macroeconomics called real-business cycle theory: if the IES is too small, then this approach cannot satisfactorily account for the changes in the supply of labor that are witnessed over the course of the U.S. business cycle.

The magnitude of the IES depends on the precise relationship between utility and leisure within each time period of the individual's life. To understand why this is so, suppose that during each period the economy is in one of either two states: good (G) or bad (B). Furthermore, assume that, in each period, there is a 50–50 chance the state is either good (and the wage is high) or bad (and the wage is low). This is clearly a world in which individuals are constantly buffeted by a sequence of good and bad shocks over their lifetimes and must respond accordingly. How great will their responses be?

FIGURE 28.4 The Intertemporal Substitution Elasticity

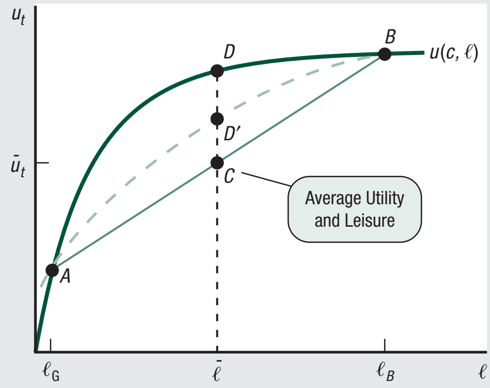

In Figure 28.4, the curved line depicts the relationship between Dougal's utility and his leisure time, holding constant the level of his consumption, c. (For the moment disregard the dashed curve, which is Betsy's utility schedule.) The relationship is increasing, because Dougal's utility increases with his leisure. It is also concave, which reflects leisure's diminishing marginal utility. Suppose that Dougal contemplates a plan in which he proposes to take full advantage of the variation in the wage. Thus, during good times (when the wage is high), he chooses the low level of leisure ℓ_G, and during bad times (when the wage is low) he chooses to enjoy the high level of leisure ℓ_B. Is this a good plan? The answer is that it is and it isn't. It is good because he takes full advantage of the change in his economic opportunities. There is, however, a catch: the huge variation in his leisure is itself very costly.

To see why this is so, notice that his plan provides him with an average utility of $\bar{u}_t$ and an average level of leisure of $\bar{\ell}$. This outcome is located at point C in the figure (given the 50% chance of either state it is located at the midpoint of the chord connecting points A and B). Compare this plan with the one in which Dougal picks $\bar{\ell}$ in both the good and the bad states, which eliminates the variation in his leisure altogether. The resulting outcome is located at point D, which lies *on* the curve $u(c, \ell)$. Notice that point D lies above point C: Dougal apparently gains DC in additional utility by keeping his leisure constant, rather than by varying it according to whether the state is good or bad. Given this apparent gain, why does he not just keep his leisure fixed at $\bar{\ell}$? The reason is simple enough. The posited thought experiment is predicated on holding fixed the level of his consumption. However, if he fails to make hay when the sun shines and work harder when times are good, then his consumption will suffer.

The basic economic tension is now clear enough. On the one hand, he strives to keep the variation in his hours worked—and hence leisure—to a minimum, for the reasons just described. On the other hand, he strives to make the variation in the hours he works as large as possible to take full advantage of the change in his economic circumstances. His optimal behavior is derived by judiciously balancing these opposing forces.

Finally, (and this is the key point) the value of ironing out the fluctuations in leisure depends directly on the **curvature** of the relationship between u and ℓ. To see why, let us return to Figure 28.4. The dashed curve represents Betsy's utility schedule. Notice that its curvature is far less pronounced than Dougal's. Because of this property, Betsy loses only CD' if she varies her leisure to take full advantage of the fluctuations in the wage, whereas Dougal loses the much larger amount CD. Hence, following a temporary change in the wage, Betsy's labor-supply response is predicted to be much greater than Dougal's: they both enjoy the same benefits, but it is much cheaper in utility terms for Betsy (CD') to vary her leisure than it is for Dougal to vary his (CD). In short, Betsy's IES is greater than Dougal's. This finding is general. Consider Major Result 28.2.

MAJOR RESULT 28.2

The Intertemporal Elasticity of Substitution

Ceteris paribus, the **smaller** is the curvature of $u(c, \ell)$ with respect to ℓ the **greater** is the IES.

Later, in Section 28.4, we examine some of the empirical work that has sought to estimate the magnitude of the IES.

28.3 Labor Supply over the Life Cycle

In this section, we present some of the empirical evidence that pertains to the labor-supply choices that people make over the course of their respective life cycles.[10] As we have seen, one of the primary insights of the model, embodied in the intertemporal substitution hypothesis, is that individuals supply the greatest amount of labor when it is most propitious for them to do so. The $\cap$-shaped age-earnings profile implies that earnings are greatest in the middle part of the life cycle. Consequently, individuals are predicted to be most likely to participate in the labor force and to work longer hours in the mid part of their careers. In Figure 28.5, we depict the relationship between the participation rate and age for both men (panel (a)) and women (panel (b))—across several different U.S. cross-sections.[11] The $\cap$-shaped age-participation pattern is clearly discernible. In the remainder of this section, we take a closer look at this evidence.

Evolutionary Wage Changes

In this subsection, we discuss changes in participation rates *within* each of the cross-sections depicted in Figure 28.5. (For the moment, ignore the various shifts that have occurred over time—the reasons for their occurrence are discussed later.)

Men. For men, a $\cap$-shaped participation pattern is clearly discernible (for both the 1960 and 2000 cohorts). In each cohort, a striking age-related increase in participation rates for younger men is clearly visible. Indeed, over the 10-year period, between the ages of 16 and 26, the participation rate more than doubles. Notice

FIGURE 28.5 Labor-Participation Rates by Age and Gender

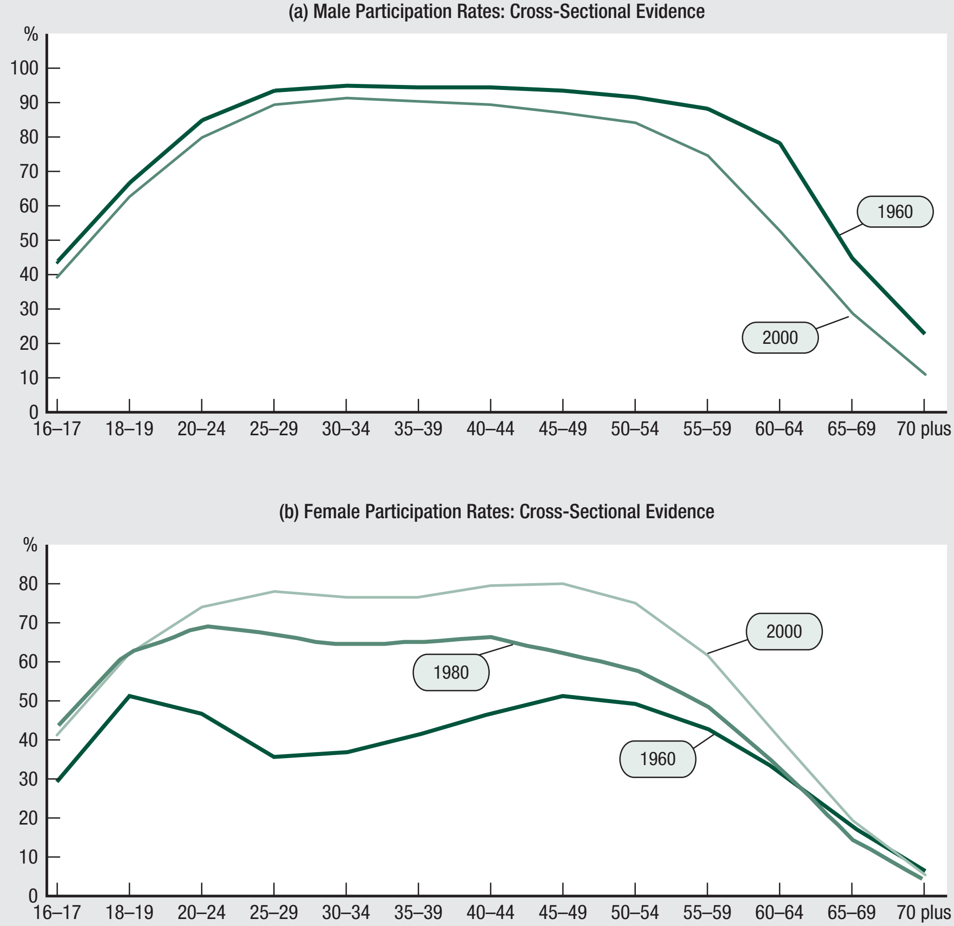

Source: Szafran (2002), original BLS data 1960–2002.

too that participation rates gradually begin to decline between the ages of 50 and 55. Upon reaching the age of 55 in the 2000 cohort (62 in the 1960 cohort), however, the decline becomes precipitous, and men begin to retire in droves. In point of fact, the consensual view among economists is that the observed decline in the participation rate during this period is too rapid to be explained by the corresponding change in the wage over this part of the life cycle. As discussed in Chapter 30, the Social Security system and private pension plans offer a possible explanation for this pattern of behavior. On balance, the evidence provides rather strong prima-facie support for the view that men do respond in the manner predicted by the theory to evolutionary changes in their life cycle earnings.

Women. Broadly speaking, the same patterns are observed for women, except that the three cohorts exhibit a pronounced *M*-shaped participation pattern. More specifically, in accordance with the basic theory, women are least likely to participate at either ends of their life cycles (i.e., when they are young and old) because here their earnings are at their lowest.[12] Likewise, with an important caveat, women are more likely to participate in the middle segments of their life cycles because here their earnings are at their greatest. The caveat is easy enough to see. It is the slight dip in participation that occurs among 18–25 year olds in the 1960 cohort, 21–30 year olds in the 1980 cohort, and 25–35 year olds in the 2000 cohort. This reflects the effects of childbearing and child rearing during these years.[13] In conclusion, with the exception of issues related to retirement and the family concerns just noted, people respond to evolutionary wage changes in broad accordance with the basic precepts of the theory. They supply the greatest amount of labor when it is most advantageous for them to do so.

Changes in Male and Female Participation Rates

Arguably, two of the most fundamental social and economic developments in contemporary U.S. history are the *decrease* in the labor force participation rates of men, and the stunning *increase* in the participation rates of women—especially among married women. The life cycle model offers valuable insights into the proximate causes of these changes. It is fair to say, however, that the model provides a better account of the behavior of men than it does the behavior of women. In fact, from a theoretical perspective, providing a satisfactory explanation for the changes in the observed participation rates of men calls for little more than the light field artillery of the life cycle model presented in this chapter (together with some modest emendations that account for the effect of pensions and Social Security payments). Accounting for the changes in the participation rates of women is a different matter altogether. Although the life cycle model does provide some illumination, a fully satisfactory account calls for the deployment of much heavier guns in one's theoretical arsenal. Both the guns and, indeed, the shells are provided in Chapter 29, where we take a closer look at the economics of the family.

The Decline in the Male Participation Rate. It is now time to examine the reasons for the shift in male participation rates depicted in Figure 28.5. As is evident from a comparison of the 1960 and 2000 cohorts, the participation rate declined at all points over the male life cycle. Two distinct phases are, however, clearly discernible that correspond to the 16–50 and 50-plus age groups. The former group exhibited a fairly uniform but small (absolute) decrease in their participation rates. The latter group is characterized by a much larger shift. In fact, as illustrated, the participation rate for men in their 60s fell from over 80% in 1960 to barely 55% in 2000, which is a huge change. Two proximate factors are thought to be responsible for the decline. First, between 1960 and 2005 the average *real* level of total hourly compensation rose by almost 60%, which unleashed both income and substitution effects.[14] The evidence suggests that among men, the income effect must have been the dominant force because of the observed decline in their participation rate. Second, the liberalization of Social Security eligibility requirements and the growth in private pension plans encouraged greater numbers of men to retire and do so earlier in their careers.[15]

The Increase in the Female Participation Rate. The increase in the female participation rate is one of the great marvels of our time. As depicted in Figure 28.5, between 1960 and 2000 it increased substantially in every age group but one: those in the 65 plus retirement category. Moreover, a comparison of these two cohorts with the 1980 cohort suggests that much of the transformation in participation rates of younger women (aged 16 to 24) had already occurred by 1980. Consequently, the greater part of the most recent increase in female participation rates—over the last 25 years or so—resulted from the increased labor force attachment of women in the 25 to 60 age group.[16]

At first glance, the differential experiences of men and women might appear somewhat enigmatic. Women too saw their wages increase over the same period, and they were also affected by the change in both Social Security eligibility and by the growth in private-sector pension plans. This raises a fundamental question: why didn't their labor force participation rate also decline over time? This is a difficult question to fully address using just the intertemporal labor-supply model. Nevertheless, the approach does provide two important insights. First, from the material presented in Chapter 4, recall that among nonparticipants an increase in the wage has a zero income effect, which leaves only the work-promoting substitution effect. Since, historically, a large proportion of women were nonparticipants, the secular increase in the real wage is predicted to increase their participation rate.

Second, a possible interpretation of the behavior of women in the 60 plus age group is that it reflects the stalemate between two conflicting forces. On the one hand, the increase in the real wage tended to increase the participation rate (just as for the younger group of women). On the other hand, the enhanced access to Social Security benefits and the increase in the real values of these benefits tended to reduce it (because of the income effect). If these two forces happened to just cancel each other out, then there would be little observed change in the overall participation rate of this group of women.

28.4 Labor Supply over the Business Cycle

One of the most important and fertile applications of the ISH is helping to explain the observed changes in labor supply that occur over the course of the business cycle. The reason is that the aggregate economy is constantly buffeted by a

FIGURE 28.6 U.S. Real GDP 1960–2008

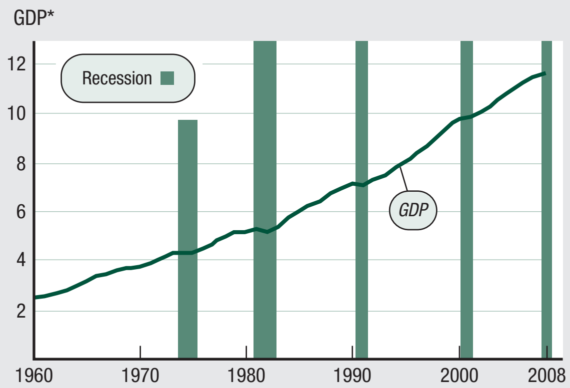

* In trillions of constant 2000 dollars.
Source: Bureau of Economic Analysis.

variety of shocks that might be expected to elicit the sorts of behavior discussed in Section 28.2. These shocks arise from positive or negative technological innovations that affect production opportunities and from policy changes that either foster or retard economic activity.[17]

Figure 28.6 depicts U.S. GDP over the period 1960–2008 (measured in constant 2000 dollars). The positive trend—reflecting economic growth—is clearly visible. Notice, too, the vicissitudes of economic activity. Some epochs lie above trend; others lie below it. (Recessions are depicted between the vertical colored bands.) Much of modern macroeconomics seeks to understand the patterns depicted in Figure 28.6 in general and its relation to labor-market activity in particular (see Figure 2.4*d* for the corresponding fluctuations in the unemployment rate).

Real Business Cycle (RBC) Theory

Lucas and Rapping (1969) is the seminal work that examined the connection between macroeconomic shocks and labor supply. The most recent work, pioneered by the 2004 Nobel laureates Fin Kydland and Edward Prescott, flies under of the general banner of **real business cycle (RBC) theory**.[18] The hallmark of this revolutionary research program is a careful adherence to the principles of economic theory. Thus the analysis adopts a *general-equilibrium perspective*, which means that economists worry about how all of the elements of the economy coherently fit together, and it presumes that both firms and workers seek to maximize their respective discounted profits and utilities. In this setting, fluctuations in output, employment, and unemployment reflect agents' *optimal-equilibrium* responses to the myriad economic shocks that occur over the course of the business cycle.

Within the RBC framework, intertemporal substitution possibilities take center stage. Both the employment and unemployment rates vary over the course of the business cycle because of the optimal responses of workers to changes in their economic circumstances. Hence, as implied by the ISH, during good times—when the wage is high—workers increase their supply of labor hours to take advantage of the temporary favorable shift. In contrast, during bad times—such as a recession—they respond to temporary reductions in the wage by increasing their enjoyment of leisure and reducing the number of hours they work. Hence unemployment rises during a recession because it represents the most prudent time for workers to consume their leisure. Likewise, employment rises during an economic upswing because this is the best for them to forego their leisure and to work harder. The main insights of the RBC approach are summarized in Major Result 28.3.

MAJOR RESULT 28.3

Macroeconomic Implications

According to the intertemporal substitution hypothesis, the observed business cycle fluctuations in employment and unemployment stem from the rational adjustments by workers in the use of their time. Thus the increase in unemployment during a recession occurs because workers substitute toward leisure and away from working for pay in the market. Similarly, employment rises during booms, as this is the most advantageous time for workers to supply their labor.

At the time Lucas and Rapping (1969) first advanced their hypothesis, which essentially asserts that unemployment arises because workers are enjoying leisure time, it stirred a hornet's nest of controversy. The reason was that the prevailing view held that unemployment is involuntary and that, by definition, people are unemployed because they cannot find work at the going wage. Over the intervening years, however, the controversy has diminished somewhat, as the RBC approach has continued to attract adherents. In fact, from one perspective, it is now those theories that fail to adhere to the tenets of the RBC framework that are viewed as controversial, since they often depart from the canons of received economic wisdom with little or no apparent justification.

One of the main reasons the RBC approach has attracted such favorable attention is that it does a good job explaining some of the core facets of economic activity observed over the course of the business cycle. Another, as explained in Economic Application 28.1, is that it has shed light on the nature of **hidden unemployment** and hence the proper way to measure the unemployment rate.

Nevertheless, while its qualitative findings are compelling enough, the RBC approach has perennially been dogged by the problem of establishing its *quantitative* significance. Put simply, over the course of the business cycle, the U.S. economy experiences relatively modest changes in the wage but sometimes huge fluctuations in that commonly used measure of national economic virility: the unemployment rate $u = U/L$ (where, to recap, U is the *level* of unemployment, and L is the *size of the labor force*). Hence, for a given value of L, the empirical success of the theory requires that small changes in the wage must induce a huge substitution between work and leisure among workers; only then is it possible to generate the necessary fluctuations in the level of unemployment, U, required to explain the correspondingly large-scale fluctuations observed in u. The validity of this story requires (as explained on page 10) that the intertemporal elasticity of substitution (IES) must be extremely large: workers must be really willing to make hay when the sun shines! Yet, as we shall shortly see, the evidence indicates that the IES is not up to the task because it is apparently too small.

The Added Worker and Discouraged Worker Effects

The unemployment rate is defined as $u = U/L$. It follows immediately that the observed fluctuations in u could result from changes in L rather than from changes in U. Specifically, if L is *procyclical*—rising in booms and falling in recessions—then this might explain the observed *countercyclical* behavior of the unemployment rate u. In fact, in 2005, a full third of the U.S. working population—amounting to some 75 million persons—did not belong to the labor force (N).

Furthermore, there is a growing consensus among economists that the boundary between the N and L states is extremely porous. Given the huge number of people occupying the N state, it is conceivable that, in accordance with the RBC approach, business cycle fluctuations could potentially induce large swings in the size of the labor force as workers switch between the two states. In turn, this could account for the sizable countercyclical fluctuations that are observed in the unemployment rate.

But are there grounds for suspecting that the size of the labor force is procyclical in practice? According to the **discouraged worker effect,** there are. More specifically, during a recession, workers are predicted to drop out of the labor force (and to enter the N state) to enjoy more leisure time following the deterioration in their employment opportunities. Exactly the opposite occurs during an upswing in activity because some workers switch from the N to the L state to take advantage of the improved conditions. The upshot of this theory is that L is procyclical: it rises in booms and falls in recessions, which then induces a corresponding countercyclical pattern in the unemployment rate $u = U/L$.

ECONOMIC APPLICATION 28.1

Hidden Unemployment

It is sometimes argued that the measured rate of unemployment understates its true level, because it fails to account for the hidden unemployed, who include among their members, marginally attached workers in general and discouraged workers in particular. For example, according to BLS statistics, in December 2003 the unemployment rate was 5.7%. The total number of discouraged workers at that time was 433,000 persons, while the number of marginally attached workers was about 1.5 million persons. Adding discouraged workers would raise the unemployment rate to 6.0%; adding all marginally attached workers would increase it to 6.7%—which is a 17% increase in the measured rate of unemployment.

There is a heated debate concerning the prudence of using these more encompassing measures of unemployment. One reason for the reticence stems from the intertemporal substitution hypothesis considered in this chapter. Recall that, from this theoretical perspective, nonworkers (whether unemployed, marginally attached, or discouraged) are responding optimally to an adverse change in their economic circumstances by voluntarily substituting work for leisure. Thus, from a policy perspective, it makes little sense to worry about counting people as unemployed, if they are optimally choosing not to work.

To get a firmer handle on just how marginal the marginally attached really are, Flaim (1984) studied their behavior and characteristics more closely. His evidence (though admittedly 25 years old) is striking. For instance, he finds that a full two thirds of discouraged workers reported (in roughly equal proportions) that for the entire 1979–1983 period (i) they had either never worked at all or that more than 5 years had elapsed since their last job, and (ii) they had held their last job 1 to 5 years before being interviewed. Indeed, only one third of them had held a job in the year before being identified as discouraged workers. He concludes from this:

> Several sets of data-covering periods of relatively low as well as very high unemployment show that a large proportion of persons classified as discouraged workers . . . have rare contacts with the job market. . . . Many of the discouraged, although expressing their desire for a job and their intention to look for one, find it very difficult to translate their sentiments into concrete and productive job-seeking efforts.[19] ■

Benati (2001) finds that discouraged workers do act in precisely this manner. His evidence indicates that there is a clear countercyclical pattern in which the number of discouraged workers *declines* during booms (as they enter the labor force in search of work) and *rises* during recessions (as they abandon their job search efforts).

Economists have, however, identified another possible business cycle effect, called the **added worker effect**. This refers to a situation in which a family member enters the labor market in response to a change in the economic circumstances faced by another family member. The canonical—but hopefully soon-to-be-dated—example is one in which a married woman temporarily enters the labor force to make up for the shortfall in the family's income after her husband loses his job. The incentive for her to do so presumably depends on the net loss in the family income. As a consequence, the generosity (or otherwise) of unemployment benefits—a core component of the realized net loss in familial income following a job loss—could exert a strong influence on the size of this effect. (As we shall see shortly, this prediction offers a simple test of the theory.) Finally, it is worth stressing that involuntary job loss is neither necessary nor sufficient for the existence of the added worker effect. For example, according to the ISH, a temporary reduction in the primary bread winner's earnings may induce the temporary substitution of labor within the household and result in the labor force participation of the spouse.

A number of empirical studies sought to test for the presence of the added worker effect, and several of them have confirmed its presence. Moreover, Cullen and Gruber (2000) find that its magnitude is, as predicted by theory, substantially diminished by generous unemployment benefit payments. Nevertheless, the available evidence indicates that the added worker effect is much smaller than the discouraged worker effect. Hence, on balance, the size of the labor force is procyclical.

The Estimated Size of the IES

As noted earlier, the IES measures an individual's willingness to intertemporally substitute leisure over time. An empirically sizable IES is essential for the empirical relevance of the RBC approach. If it is too small, then these models cannot explain the sizable fluctuations in hours worked that occur over the course of the business cycle. For instance, if it were discovered that the IES equals zero, then business cycle fluctuations would have no effect at all on the number of hours worked. Yet, given the large employment fluctuations that are actually observed in the data, this would then cast serious doubt on the practical usefulness of the approach.

A vast number of studies have sought to estimate the IES.[20] (Economic Application 28.2 offers some evidence on the size of the IES from two interesting sources: Stadium vendors and New York cabdrivers.) However, the basic conclusion of most of these studies is that the estimated elasticity—which, in most cases lies between 0 and 0.5—*is too small* for RBC models to fully account for the fluctuations in either hours worked or in the level of unemployment that are observed to occur over the business cycle. In a nutshell, *the theory fails to fully explain the data*.

Now, these negative findings did not cause (many) proponents of the RBC framework to run home sobbing. Rather, it is inherent in the nature of science that negative findings can be as rewarding—if not more so—than positive ones that apparently confirm a particular theory. In this light, there are at least three main accomplishments of the theory. First, it is no small undertaking to construct a theoretical apparatus that is precise enough to be tested. Second, the approach has provided valuable insights into the behavior of the supply of labor over the business cycle and the nature of unemployment.

Finally, the negative empirical findings led to perhaps one of the most significant recent developments in modern macroeconomics: the integration of *household production theory* with real business cycle models.[21] The thrust of the new approach is based on the fundamental recognition that we have been rather cavalier in our interpretation of leisure time, which has simply been defined as time that is *not spent working for pay*. Yet, much of what we have then called leisure is actually a form of work that is carried out at home. Examples include doing the laundry, cleaning the house, cooking a meal at home, and so on and so forth. It turns out—and we discuss further in Chapter 29—once these richer uses of time are admitted, the (correctly) estimated IES rises substantially, and, as consequence, RBC models do a much better job at explaining the evidence.

ECONOMIC APPLICATION 28.2

Stadium Vendors and Cabdrivers

As noted in the text, a major theme that emerges from the intertemporal labor-supply model is that the response to an *unanticipated* temporary change in the wage may have a sizable impact on the supply of labor within a given period. The reason is temporary changes are subject to the full work-enhancing force of the substitution effect, with little negative attenuation from the income effect.

Oettinger (1999) conducted an interesting empirical study that examines the daily labor supply of stadium vendors. This population is (statistically) attractive for two main reasons. First, vendors freely choose whether to work on a given day. Second, there are substantial and predictable shifts in product demand (according to the number of tickets sold at the game) that generate equally large and predictable changes in the vendor's earnings. As suggested by the intertemporal substitution hypothesis, wage elasticities are extremely large. Controlling for the fact that the vendor's wage depends on both supply and demand conditions, he estimates participation elasticities of about 0.55.

Camerer, Babcock, Loewenstein, and Thaler (1997) empirically investigate the daily labor supply of New York cabdrivers. This is an interesting group to study. They can set their own hours. Moreover, there are numerous transitory shocks (the weather being one prime case in point) that influence the number of people traveling by cab on any given day, and hence their earnings. Nevertheless, the authors' results indicate a transitory increase in the wage has a negative effect on the labor supplied by cabdrivers. This is at direct odds with the intertemporal substitution hypothesis because it suggests the income effect (which should be close to zero) is large and, in fact, dominant. What is more, they find that drivers act as if they are not forward-looking because they take things one day at a time and quit after they have attained a target level of earnings. This behavior may even contradict the hypothesis that they behave rationally.

However, in a recent study, Farber (2005) reexamines the evidence, and finds, in contrast to Camerer, Babcock, Loewenstein, and Thaler (1997), that the labor supply of taxi drivers is consistent with the existence of intertemporal substitution. As predicted from the basic theory, income effects are small. He attributes the very different findings of the two studies to differences in empirical methods and to possible problems with Camerer and co-authors' wage measurements. One suspects, however, this is not the end of the debate. ■

SUMMARY

- According to the dynamic theory of labor supply, individuals pick their consumption and leisure plans so as to maximize the discounted values of their respective lifetime utilities.
- In formulating their plans, workers recognize that their current actions have future consequences. Anticipated future events also affect current actions.
- According to the *intertemporal substitution hypothesis* (ISH), workers carefully choose the timing of their consumption and leisure activities in order to make best use of the economic opportunities that are available to them.
- Given the typical ∩-shaped age-earnings profile, the ISH predicts that the number of hours worked over the life cycle is also ∩-shaped.
- An important insight of the ISH is that workers are predicted to be more responsive to *temporary* wage changes than to *permanent* ones. This corresponds to the aphorism *Make hay when the sun shines*.
- Two of the most profound changes in contemporary U.S. social history are the reduction in male participation rates and the sizable increase in female participation rates. The theory of dynamic labor supply provides important insights into the proximate causes of these events, based on observed changes in earnings, pensions, and the Social Security system.
- The intertemporal elasticity of substitution (IES) measures the strength of the substitution effect. It equals the percentage change in leisure induced by a temporary 1% impulse in the wage.
- Real business cycle theory (RBC) is a revolutionary new approach to macroeconomics. The hallmark of the theory is that it attributes the fluctuations in output, employment, and unemployment observed over the business cycle to agents' rational responses to real shocks of one sort or another.
- According to the RBC, unemployment tends to rise during a recession, as this is the most propitious time for workers to enjoy their leisure. For similar but opposite reasons, employment rises and unemployment declines during an economic expansion.

PROBLEMS

P1. Consider an individual who decides to take a day (8 hours) off work when he is 20 years old, and to make up for the earnings shortfall by working an extra Δh hours in 30 years' time. Suppose that the interest rate is $r = 5\%$, and that the constant hourly wage is $W = \$20$. What is Δh?

P2. Briefly describe the distinction between **evolutionary** and **parametric** wage changes. Suppose that Dougal anticipates that his (real) wage will gradually grow and have doubled when he turns 50. In contrast, Betsy receives a surprise promotion when she is 50 that doubles her earnings. Describe their respective life cycle behaviors.

P3. Even today, many people hold much of their wealth in the form of their home equity. Given this, what is the likely effect of the sudden recent housing crunch on their labor-supply decisions?

P4. Suppose that the government suddenly borrows a trillion or so dollars to finance an unanticipated war and that everyone expects that, as a result, income tax rates will increase markedly in the future. How might this affect labor supply?

P5. Explain the **added** and **discouraged** worker effects. What does the evidence suggest about their relative importance in the United States?

P6. What is meant by **hidden unemployment**? What are the pros and cons of including measures of hidden unemployment in official unemployment statistics?

P7. Is it reasonable to suppose that unemployment tends to increase during recessions because workers are enjoying more leisure?

NOTES

1. Another problem is that the theory emphasizes unilateral decision making. Most labor-supply decisions, however, are made within the broader context of the family to which the individual belongs. This issue forms much of the subject matter of Chapter 29.
2. An excellent account of the dynamics of labor supply is provided in Killingsworth (1983), ch. 5.
3. The current state of the field is admirably surveyed by Blundell and MaCurdy (1999).
4. The term *potential* is used because the individual may choose not to work. Note, w_0 is his potential wage in the first year (when he is 20), w_1 is his wage in the second year (when he is 21), and so forth for all the years $t = 2, 3, \ldots, S$.
5. According to the simple static theory of labor supply, the individual works longer in response to a wage increase if the substitution effect dominates the income effect.
6. An anticipated wage change is called evolutionary, and (as we shall see in the next section), an unanticipated change is called parametric. Pistaferri (2003) uses data collected on workers' subjective wage expectations to empirically separate the two effects.
7. As was discussed in Chapter 4, the income effect itself is calculated by examining how the individual responds to the increase in utility—holding constant the wage at its initial value.
8. Recall that the substitution effect captures how an individual responds to a change in the opportunity cost of some or another action for a given level of utility.
9. Previously, deterministic movements along the age-earnings profile were termed evolutionary changes. Shocks that cause the wage to move away from its expected value are called parametric changes.
10. An enormous number of empirical studies have examined the intertemporal life cycle model. Pioneering work includes Heckman (1974, 1976) and MaCurdy (1981). Card (1994) offers an excellent purview of the earlier work on the intertemporal substitution hypothesis. Blundell and MaCurdy (1999) provide an excellent account of more recent research.
11. Recall that a cross section refers to a snapshot at a given point in time. Thus, in Figure 28.5*b*, the curve labeled 1980 refers to the age participation relationship observed for women in 1980.
12. Just as was the case for men, the decline in labor force participation among elderly women is too rapid to be explained by a decline in the wage alone.
13. The evidence hints at a sizable change in fertility patterns over the past 40 years or so. As indicated, it appears that there has been a secular increase in the age at which women have children. This important topic is discussed in Chapter 29.
14. Nevertheless, the increase was not uniform across all worker groups. Some witnessed both protracted and substantial declines in their compensation levels over the period. Juhn (1992)

and Welch (1997) examine this issue carefully and find that the decline in the participation rate of Black men was especially pronounced.

15. There is evidence that changes in access to disability insurance also exerted a strong influence on the participation rates of younger men in the 35–65 age group. See, for example, Parsons (1980).
16. The figure depicts the changes in participation rates among all women. Great care, however, must be exercised when analyzing these sorts of aggregate statistics because they can mask important differences among subgroups. For example, Coleman and Pencavel (1993) show that the participation rate actually decreased for less educated women.
17. The electric restroom hand dryer notwithstanding, the reader may wonder what on earth is meant by a negative technology shock. This simply refers to a situation in which a sub-par level of innovative activity occurred over a given period of time.
18. See, for example, Kydland and Prescott (1982). Plosser (1989) offers a very nice introduction to the literature. The labor-market implications of the theory are examined by Rogerson (1988). Rogerson (1997) offers an excellent discussion off the main issues.
19. Flaim (1984), p. 8.
20. Early, pioneering studies include Altonji (1982) and MaCurdy (1983). Research in this area is still very active and opinions differ on the actual value of the elasticity. Recent work includes Ham and Reilly (2002), French (2004), and Imai and Keane (2004). French (2004) estimates the elasticity is close to zero. Imai and Keane (2004) find that, in general, previous studies have underestimated the elasticity because they fail to account for the fact that the wage itself is endogenous since it depends on the accumulation of human capital over the life cycle. After these factors are accounted for, they estimate an elasticity of 3.82, which is large enough for real business cycle models to provide a satisfactory account of the data.
21. See, for example, Benhabib, Rogerson, and Wright (1991); Greenwood, Rogerson, and Wright (1993); and Nosal, Rogerson, and Wright (1992). Rupert, Rogerson, and Wright (2000) show that failing to account for home production leads to a substantial downward bias in estimates of the IES.

REFERENCES

Altonji, Joseph G. (1982). "The Intertemporal Substitution Model of Labor Market Fluctuations: An Empirical Analysis." *Review of Economic Studies* 49(5): 783–824.

Benati, Luca. (2001). "Some Empirical Evidence on the 'Discouraged Worker' Effect." *Economic Letters* 70(3): 387–395.

Benhabib, Jess, Richard Rogerson, and Randall Wright. (1991). "Homework in Macroeconomics: Household Production and Aggregate Fluctuations." *Journal of Political Economy* 99(6): 1166–1187.

Blundell, Richard W., and Thomas E. MaCurdy. (1999). "Labor Supply: A Review of Alternative Approaches." In *Handbook of Labor Economics*. Vol. 3A. Ed. Orley Ashenfelter and David Card. Amsterdam, North Holland: Elsevier, pp. 1559–1695.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. (1997). "Labor Supply of New York City Cabdrivers: One Day at a Time." *Quarterly Journal of Economics* 112(2): 407–441.

Card, David. (1994). "Intertemporal Labor Supply: An Assessment." In *Advances in Econometrics, Sixth World Congress*. Vol. II. Ed. Christopher Sims. New York: Cambridge University Press, pp. 49–80.

Coleman, Mary T., and John H. Pencavel. (1993). "Trends in Market Work Behavior of Women since

1940." *Industrial and Labor Relations Review* 46(4): 653–676.

Cullen, Julie Berry, and Jonathan Gruber. (2000). "Does Unemployment Insurance Crowd Out Spousal Labor Supply?" *Journal of Labor Economics* 18(3): 546–572.

Farber, Henry S. (2005). "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers." *Journal of Political Economy* 113(1): 46–82.

Flaim, Paul O. (1984). "Discouraged Workers: How Strong Are Their Links to the Job Market?" *Monthly Labor Review* 107(8): 8–11.

French, Eric. (2004). "The Labor Supply Response to (Mismeasured but) Predictable Wage Changes." *Review of Economics and Statistics* 86(2): 602–613.

Greenwood, Jeremy, Richard Rogerson, and Randall Wright. (1993). "Putting Home Economics into Macroeconomics." *Federal Reserve Bank of Minneapolis Quarterly Review* 17(3): 2–11.

Ham, John C., and T. Reilly Kevin. (2002). "Testing Intertemporal Substitution, Implicit Contracts, and Hours Restriction Models of the Labor Market Using Micro Data." *American Economic Review* 92(4): 905–927.

Heckman, James J. (1974). "Life Cycle Consumption and Labor Supply: An Explanation of the Relationship between Income and Consumption over the Life Cycle." *American Economic Review* 64(1): 188–194.

———. (1976). "A Life-Cycle Model of Earnings, Learning, and Consumption." *Journal of Political Economy* 84(4): S11–S44.

Imai, Susumu, and Michael P. Keane. (2004). "Intertemporal Labor Supply and Human Capital Accumulation." *International Economic Review* 45(2): 601–641.

Juhn, Chinhui. (1992). "Decline of Male Labor Force Participation: The Role of Declining Market Opportunities." *Quarterly Journal of Economics* 107(1): 79–121.

Killingsworth, Mark R. (1983). *Labor Supply*. Cambridge, UK: Cambridge University Press.

Kydland, Finn E., and Edward C. Prescott. (1982). "Time to Build and Aggregate Fluctuations." *Econometrica* 50(6): 1345–1370.

Lucas, Robert E. Jr., and Leonard A. Rapping. (1969). "Real Wages, Employment, and Inflation." *Journal of Political Economy* 77(5): 721–754.

MaCurdy, Thomas E. (1981). "An Empirical Model of Labor Supply in a Life-Cycle Setting." *Journal of Political Economy* 89(6): 1059–1085.

———. (1983). "A Simple Scheme for Estimating an Intertemporal Model of Labor Supply and Consumption in the Presence of Taxes and Uncertainty." *International Economic Review* 24(2): 265–289.

Nosal, Ed, Richard Rogerson, and Randall Wright. (1992). "The Role of Household Production in Models of Involuntary Unemployment and Underemployment." *Canadian Journal of Economics* 25(3): 507–520.

Oettinger, Gerald S. (1999). "An Empirical Analysis of the Daily Labor Supply of Stadium Vendors." *Journal of Political Economy* 107(2): 360–392.

Parson, Donald O. (1980). "The Decline in Male Labor Force Participation." *Journal of Political Economy* 88(1): 117–134.

Pistaferri, Luigi. (2003). "Anticipated and Unanticipated Wage Changes, Wage Risk, and Intertemporal Labor Supply." *Journal of Labor Economics* 21(3): 729–754.

Plosser, Charles I. (1989). "Understanding Real Business Cycles." *Journal of Economic Perspectives* 3(3): 51–77.

Rogerson, Richard. (1988). "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics* 21(1): 3–16.

———. (1997). "Theory Ahead of Language in the Economics of Unemployment." *Journal of Economics Perspectives* 11(1): 73–92.

Rupert, Peter, Richard Rogerson, and Randall Wright. (2000). "Homework in Labor Economics: Household Production and Intertemporal Substitution." *Journal of Monetary Economics* 46(3): 557–579.

Trostel, Philip A. (1993). "The Effect of Taxation on Human Capital." *Journal of Political Economy* 101(2): 327–350.

Welch, Finis. (1997). "Wages and Participation." *Journal of Labor Economics* 15(1): S77–S103.

Ziliak, James P., and Thomas J. Kniesner. (1999). "Estimating Life Cycle Labor Supply Tax Effects." *Journal of Political Economy* 107(2): 326–359.