cient should be zero.[9] The estimated coefficient is in fact 0.89, with a standard error of 0.46. Thus Shea also finds a quantitatively large (though only marginally statistically significant) departure from the random-walk prediction.

Recall that in our analysis in Sections 7.1 and 7.2, we assumed that households can borrow without limit as long as they eventually repay their debts. One reason that consumption might not follow a random walk is that this assumption might fail—that is, that households might face *liquidity constraints.* If households are unable to borrow and their current income is less than their permanent income, their consumption is determined by their current income. In this case, predictable changes in income produce predictable changes in consumption.

Shea tests for liquidity constraints in two ways. First, following Zeldes (1989) and others, he divides the households according to whether they have liquid assets. Households with liquid assets can smooth their consumption by running down these assets rather than by borrowing. Thus if liquidity constraints are the reason that predictable wage changes affect consumption growth, the prediction of the permanent-income hypothesis will fail only among the households with no assets. Shea finds, however, that the estimated effect of expected wage growth on consumption is essentially the same in the two groups.

Second, following Altonji and Siow (1987), Shea splits the low-wealth sample according to whether the expected change in the real wage is positive or negative. Individuals facing expected declines in income need to save rather than borrow to smooth their consumption. Thus if liquidity constraints are important, predictable wage increases produce predictable consumption increases, but predictable wage decreases do not produce predictable consumption decreases.

Shea's findings are the opposite of this. For the households with positive expected income growth, the estimated impact of the expected change in the real wage on consumption growth is 0.06 (with a standard error of 0.79); for the households with negative expected growth, the estimated effect is 2.24 (with a standard error of 0.95). Thus there is no evidence that liquidity constraints are the source of Shea's results.

7.4 The Interest Rate and Saving

An important issue concerning consumption involves its response to rates of return. For example, many economists have argued that more favorable

[9]An alternative would be to follow Campbell and Mankiw's approach and regress consumption growth on actual income growth by instrumental variables, using the constructed wage growth measure as an instrument. Given the almost one-for-one relationship between actual and constructed earnings growth, this approach would be likely to produce similar results.

tax treatment of interest income would increase saving, and thus increase growth. But if consumption is relatively unresponsive to the rate of return, such policies would have little effect. Understanding the impact of rates of return on consumption is thus important.

The Interest Rate and Consumption Growth

We begin by extending the analysis of consumption under certainty in Section 7.1 to allow for a nonzero interest rate. This largely repeats material in Section 2.2; for convenience, however, we quickly repeat that analysis here.

Once we allow for a nonzero interest rate, the individual's budget constraint is that the present value of lifetime consumption cannot exceed initial wealth plus the present value of lifetime labor income. For the case of a constant interest rate and a lifetime of T periods, this constraint is

$$\sum_{t=1}^{T} \frac{1}{(1+r)^t} C_t \leq A_0 + \sum_{t=1}^{T} \frac{1}{(1+r)^t} Y_t, \tag{7.24}$$

where r is the interest rate and where all variables are discounted to period 0.

When we allow for a nonzero interest rate, it is also useful to allow for a nonzero discount rate. In addition, it simplifies the analysis to assume that the instantaneous utility function takes the constant-relative-risk-aversion form used in Chapter 2: $u(C_t) = C_t^{1-\theta}/(1-\theta)$, where θ is the coefficient of relative risk aversion (the inverse of the elasticity of substitution between consumption at different dates). Thus the utility function, (7.1), becomes

$$U = \sum_{t=1}^{T} \frac{1}{(1+\rho)^t} \frac{C_t^{1-\theta}}{1-\theta}, \tag{7.25}$$

where ρ is the discount rate.

Now consider our usual experiment of a decrease in consumption in some period, period t, accompanied by an increase in consumption in the next period by $1+r$ times the amount of the decrease. Optimization requires that a marginal change of this type has no effect on lifetime utility. Since the marginal utilities of consumption in periods t and $t+1$ are $C_t^{-\theta}/(1+\rho)^t$ and $C_{t+1}^{-\theta}/(1+\rho)^{t+1}$, this condition is

$$\frac{1}{(1+\rho)^t} C_t^{-\theta} = (1+r)\frac{1}{(1+\rho)^{t+1}} C_{t+1}^{-\theta}. \tag{7.26}$$

We can rearrange this condition to obtain

$$\frac{C_{t+1}}{C_t} = \left(\frac{1+r}{1+\rho}\right)^{1/\theta}. \tag{7.27}$$

This analysis implies that once we allow for the possibility that the real interest rate and the discount rate are not equal, consumption need not be a random walk: consumption is rising over time if r exceeds ρ and falling if r is less than ρ. In addition, if there are variations in the real interest rate, there are variations in the predictable component of consumption growth. Mankiw (1981), Hansen and Singleton (1983), Hall (1988b), Campbell and Mankiw (1989a), and others therefore examine how much consumption growth responds to variations in the real interest rate. For the most part they find that it responds relatively little, which suggests that the intertemporal elasticity of substitution is low (that is, that θ is high).

The Interest Rate and Saving in the Two-Period Case

Although an increase in the interest rate causes the path of consumption to be more steeply sloped, it does not necessarily follow that the increase reduces initial consumption and thereby raises saving. The complication is that the change in the interest rate has not only a substitution effect, but also an income effect. Specifically, if the individual is a net saver, the increase in the interest rate allows him or her to attain a higher path of consumption than before.

The qualitative issues can be seen in the case where the individual lives for only two periods. For this case, we can use the standard indifference-curve diagram shown in Figure 7.2. Assume, for simplicity, that the individual has no initial wealth. Thus in (C_1, C_2) space, the individual's budget constraint goes through the point (Y_1, Y_2): the individual can choose to consume his or her income each period. And the slope of the budget constraint is $-(1 + r)$: giving up one unit of first-period consumption allows the individual to increase second-period consumption by $1 + r$. When r rises, the budget constraint continues to go through (Y_1, Y_2) but becomes steeper; thus it pivots clockwise around (Y_1, Y_2).

In Panel (a), the individual is initially at the point (Y_1, Y_2); that is, saving is initially zero. In this case the increase in r has no income effect—the individual's initial consumption bundle continues to be on the budget constraint. Thus first-period consumption necessarily falls, and so saving necessarily rises.

In Panel (b), C_1 is initially less than Y_1, and thus saving is positive. In this case the increase in r has a positive income effect—the individual can now afford strictly more than his or her initial bundle. The income effect acts to decrease saving, whereas the substitution effect acts to increase it. The overall effect is ambiguous; in the case shown in the figure, saving does not change.

Finally, in Panel (c) the individual is initially borrowing. In this case both the substitution and income effects reduce first-period consumption, and so saving necessarily rises.

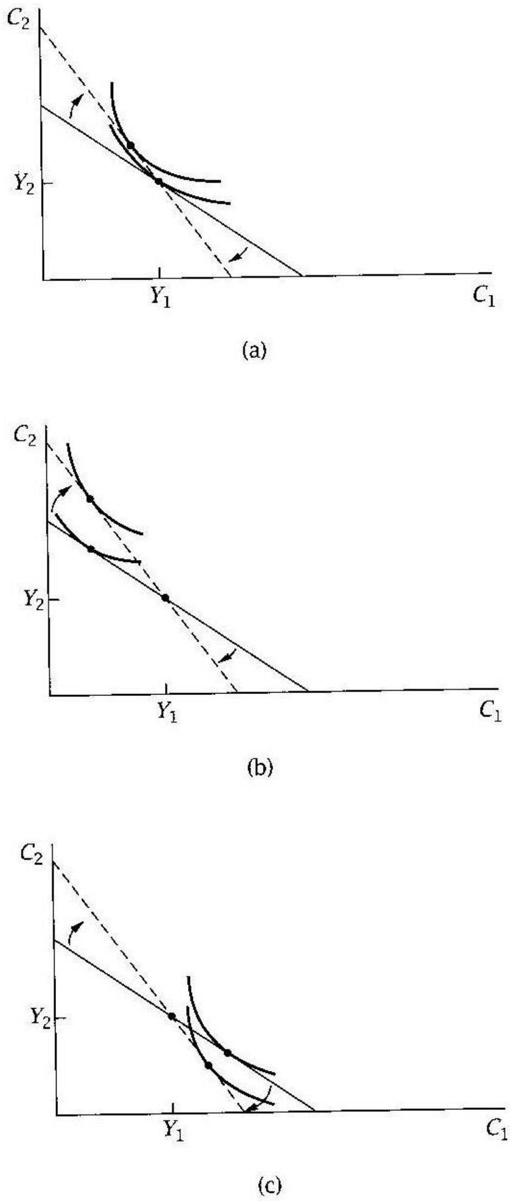

FIGURE 7.2 The interest rate and consumption choices in the two-period case

Since the stock of wealth in the economy is positive, individuals are on average savers rather than borrowers. Thus the overall income effect of a rise in the interest rate is positive. An increase in the interest rate thus has two competing effects on overall saving, a positive one through the substitution effect and a negative one through the income effect.

Complications

This discussion appears to imply that, unless the elasticity of substitution between consumption in different periods is large, increases in the interest rate are unlikely to bring about substantial increases in saving. There are two reasons, however, that the importance of this conclusion is limited.

First, many of the changes we are interested in do not involve just changes in the interest rate. For tax policy, the relevant experiment is usually a change in composition between taxes on interest income and other taxes that leaves government revenue unchanged. As Problem 7.5 shows, such a change has only a substitution effect, and thus necessarily shifts consumption toward the future.

Second, and more subtly, if individuals have long horizons, small changes in saving can accumulate over time into large changes in wealth (Summers, 1981a). To see this, first consider an individual with an infinite horizon and constant labor income. Suppose that the interest rate equals the individual's discount rate. From (7.27), this means that the individual's consumption is constant. The budget constraint then implies that the individual consumes the sum of interest and labor incomes: any higher steady level of consumption implies violating the budget constraint, and any lower level implies failing to satisfy the constraint with equality. That is, the individual maintains his or her initial wealth level regardless of its value: the individual is willing to hold any amount of wealth if $r = \rho$. A similar analysis shows that if $r > \rho$, the individual's wealth grows without bound, and that if $r < \rho$, his or her wealth falls without bound. Thus the long-run supply of capital is perfectly elastic at $r = \rho$.

Summers shows that similar, though less extreme, results hold in the case of long but finite lifetimes. Suppose, for example, that r is slightly larger than ρ, that the intertemporal elasticity of substitution is small, and that labor income is constant. The facts that r exceeds ρ and that the elasticity of substitution is small imply that consumption rises slowly over the individual's lifetime. But with a long lifetime, this means that consumption is much larger at the end of life than at the beginning. But since labor income is constant, this in turn implies that the individual gradually builds up considerable savings over the first part of his or her life and gradually decumulates them over the remainder. As a result, when horizons are finite but long, wealth holdings may be highly responsive to the interest rate in the long run even if the intertemporal elasticity of substitution is small.[10]

7.5 Consumption and Risky Assets

In practice, individuals can invest in many assets, almost all of which have uncertain returns. Extending our analysis to account for multiple assets and risk raises some new issues concerning both household behavior and asset markets.

The Conditions for Individual Optimization

Consider our usual experiment of an individual reducing consumption in period t by an infinitesimal amount and using the resulting saving to raise consumption in period $t+1$. If the individual is optimizing, this change leaves expected utility unchanged regardless of which asset the increased saving is invested in. Thus optimization requires

$$u'(C_t) = \frac{1}{1+\rho} E_t[(1+r^i_{t+1})u'(C_{t+1})] \qquad \text{for all } i, \tag{7.28}$$

where r^i is the return on asset i. Since the expectation of the product of two variables equals the product of their expectations plus their covariance, we can rewrite this expression as

$$u'(C_t) = \frac{1}{1+\rho}\{E_t[1+r^i_{t+1}]E_t[u'(C_{t+1})] + \mathrm{Cov}_t(1+r^i_{t+1}, u'(C_{t+1}))\} \qquad \text{for all } i, \tag{7.29}$$

where $\mathrm{Cov}_t(\bullet)$ is covariance conditional on information available at time t.

If we assume that utility is quadratic, $u(C) = C - aC^2/2$, then the marginal utility of consumption is $1 - aC$. Using this to substitute for the covariance term in (7.29), we obtain

$$u'(C_t) = \frac{1}{1+\rho}\left\{E_t[1+r^i_{t+1}]E_t[u'(C_{t+1})] - a\mathrm{Cov}_t(1+r^i_{t+1}, C_{t+1})\right\}. \tag{7.30}$$

Equation (7.30) implies that in deciding whether to hold more of an asset, the individual is not concerned with how risky the asset is: the variance of the asset's return does not appear in (7.30). Intuitively, a marginal increase in holdings of an asset that is risky, but whose risk is not correlated with

[10] Carroll (1992) shows, however, that the presence of uncertainty weakens this conclusion somewhat.

the overall risk the individual faces, does not increase the variance of the individual's consumption. Thus in evaluating that marginal decision, the individual considers only the asset's expected return.

Equation (7.30) implies that the aspect of riskiness that matters to the decision of whether to hold more of an asset is the relation between the asset's payoff and consumption. Suppose, for example, that the individual is given an opportunity to buy a new asset whose expected return equals the rate of return on a risk-free asset that the individual is already able to buy. If the payoff to the new asset is typically high when the marginal utility of consumption is high (that is, when consumption is low), buying one unit of the asset raises expected utility by more than buying one unit of the risk-free asset does. Thus (since the individual was previously indifferent about buying more of the risk-free asset), the individual can raise his or her expected utility by buying the new asset. As the individual invests more in the asset, his or her consumption comes to depend more on the asset's payoff, and so the covariance between consumption and the asset's return becomes less negative. In the example we are considering, since the asset's expected return equals the risk-free rate, the individual invests in the asset until the covariance of its return with consumption reaches zero.

This discussion implies that hedging risks is crucial to optimal portfolio choices. A steel worker whose future labor income depends on the health of the American steel industry should avoid—or better yet, sell short—assets whose returns are positively correlated with the fortunes of the steel industry, such as shares in American steel companies. Instead the worker should invest in assets whose returns move inversely with the health of the U.S. steel industry, such as foreign steel companies or American aluminum companies.

The Consumption CAPM

This discussion takes assets' expected returns as given. But individuals' demands for assets determine these expected returns. If, for example, an asset's payoff is highly correlated with consumption, its price must be driven down to the point where its expected return is high for individuals to hold it.

To see the implications of this observation, suppose that all individuals are the same, and return to the first-order condition in (7.30). Solving this expression for the expected return on the asset yields

$$E_t[1 + r^i_{t+1}] = \frac{1}{E_t[u'(C_{t+1})]}[(1 + \rho)u'(C_t) + a\mathrm{Cov}_t(1 + r^i_{t+1}, C_{t+1})]. \quad (7.31)$$

Equation (7.31) states that the higher the covariance of an asset's payoff with consumption, the higher its expected return must be.

We can simplify (7.31) by considering the return on a risk-free asset. If the payoff to an asset is certain, then the covariance of its payoff with consumption is zero. Thus the risk-free rate, $\overline{r}_{t+1}$, satisfies

$$1 + \overline{r}_{t+1} = \frac{(1+\rho)u'(C_t)}{E_t[u'(C_{t+1})]}. \tag{7.32}$$

Subtracting (7.32) from (7.31) then gives us

$$E_t[r^i_{t+1}] - \overline{r}_{t+1} = \frac{a\mathrm{Cov}_t(1 + r^i_{t+1}, C_{t+1})}{E_t[u'(C_{t+1})]}. \tag{7.33}$$

Equation (7.33) states that the expected-return premium an asset must offer relative to the risk-free rate is proportional to the covariance of its return with consumption.

This model of the determination of expected asset returns is known as the *consumption capital-asset pricing model*, or *consumption CAPM*. The covariance between an asset's return and consumption is known as its *consumption beta*. Thus the central prediction of the consumption CAPM is that the premiums that assets offer are proportional to their consumption betas (Breeden, 1979; see also Merton, 1973, and Rubinstein, 1976).[11]

Empirical Application: The Equity-Premium Puzzle

One of the most important implications of this analysis of assets' expected returns concerns the case where the risky asset is a broad portfolio of stocks. To see the issues involved, it is easiest to return to the Euler equation, (7.28), and to assume that individuals have constant-relative-risk-aversion utility rather than quadratic utility. With this assumption, the Euler equation becomes

$$C_t^{-\theta} = \frac{1}{1+\rho} E_t[(1 + r^i_{t+1})C_{t+1}^{-\theta}], \tag{7.34}$$

where θ is the coefficient of relative risk aversion. If we divide both sides by $C_t^{-\theta}$ and multiply both sides by $1+\rho$, this expression becomes

$$1 + \rho = E_t\left[(1 + r^i_{t+1})\frac{C_{t+1}^{-\theta}}{C_t^{-\theta}}\right]. \tag{7.35}$$

Finally, it is convenient to let g^c_{t+1} denote the growth rate of consumption from t to $t+1$, $(C_{t+1}/C_t) - 1$, and to omit the time subscripts. Thus we have

[11] The original CAPM assumes that investors are concerned with the mean and variance of the return on their portfolio rather than the mean and variance of consumption. That version of the model therefore focuses on *market betas*—that is, the covariances of assets' returns with the returns on the market portfolio—and predicts that expected-return premia are proportional to market betas (Lintner, 1965; Sharpe, 1964).

$$E[(1 + r^i)(1 + g^c)^{-\theta}] = 1 + \rho. \tag{7.36}$$

To see the implications of (7.36), we take a second-order Taylor approximation of the left-hand side around $r = g = 0$. Computing the relevant derivatives yields

$$(1 + r)(1 + g)^{-\theta} \simeq 1 + r - \theta g - \theta g r + \frac{1}{2}\theta(\theta + 1)g^2. \tag{7.37}$$

Thus we can rewrite (7.36) as

$$\begin{aligned} &E[r^i] - \theta E[g^c] - \theta\{E[r^i]E[g^c] + \mathrm{Cov}(r^i, g^c)\} \\ &\quad + \frac{1}{2}\theta(\theta + 1)\{(E[g^c])^2 + \mathrm{Var}(g^c)\} \simeq \rho. \end{aligned} \tag{7.38}$$

When the time period involved is short, the $E[r^i]E[g^c]$ and $(E[g^c])^2$ terms are small relative to the others.[12] Omitting these terms and solving the resulting expression for $E[r^i]$ yields

$$E[r^i] \simeq \rho + \theta E[g^c] + \theta \mathrm{Cov}(r^i, g^c) - \frac{1}{2}\theta(\theta + 1)\mathrm{Var}(g^c). \tag{7.39}$$

Again, it is helpful to consider a risk-free asset. For such an asset, (7.39) simplifies to

$$\bar{r} \simeq \rho + \theta E[g^c] - \frac{1}{2}\theta(\theta + 1)\mathrm{Var}(g^c). \tag{7.40}$$

Finally, subtracting (7.40) from (7.39) yields

$$E[r^i] - \bar{r} \simeq \theta \mathrm{Cov}(r^i, g^c). \tag{7.41}$$

In a famous paper, Mehra and Prescott (1985) show that it is difficult to reconcile observed asset returns with equation (7.41). Mankiw and Zeldes (1991) report a simple calculation that shows the essence of the problem. For the United States during the period 1890–1979 (which is the sample that Mehra and Prescott consider), the difference between the average return on the stock market and the return on short-term government debt—the *equity premium*—is about six percentage points. Thus if we take the average return on short-term government debt as an approximation to the average risk-free rate, the quantity $E[r^i] - \bar{r}$ is about 0.06. Over the same period, the standard deviation of the growth of consumption (as measured by real purchases of nondurables and services) is 3.6 percentage points, and the standard deviation of the return on the market is 16.7 percentage points; the

[12] Indeed, for the continuous-time case, one can derive equation (7.39) without any approximations.

correlation between these two quantities is 0.40. These figures imply that the covariance of consumption growth and the return on the market is 0.40(0.036)(0.167), or 0.0024.

Equation (7.41) therefore implies that the coefficient of relative risk aversion needed to account for the equity premium is the solution to $0.06 = \theta(0.0024)$, or $\theta = 25$. This is an extraordinary level of risk aversion; it implies, for example, that individuals prefer a 17% reduction in consumption with certainty to a 1-in-2 chance of a 20% reduction. As Mehra and Prescott describe, other evidence suggest that risk aversion is much lower than this. Among other things, such a high degree of aversion to variations in consumption makes it puzzling that the average risk-free rate is close to zero despite the fact that consumption is growing over time.

In addition, the problem becomes even more severe if we focus on the postwar period. Mankiw and Zeldes report that for the 1948-1988 period, the average equity premium is 8 percentage points, the standard deviation of consumption growth is 1.4 percentage points, the standard deviation of the market return is 14.0 percentage points, and the correlation of consumption growth and the market return is 0.45. These numbers imply a value of θ of $0.08/[0.45(0.014)(0.140)] \simeq 91$.

The large equity premium, particularly when coupled with the low risk-free rate, is thus difficult to reconcile with household optimization. This *equity-premium puzzle* has stimulated a large amount of research, and many explanations for it have been proposed. No clear resolution of the puzzle has been provided, however.[13]

7.6 Alternative Views of Consumption

The permanent-income hypothesis provides appealing explanations of many important features of consumption. For example, it explains why temporary tax cuts appear to have much smaller effects than permanent ones, and it accounts for many features of the relationship between current income and consumption, such as those described in Section 7.1.

Yet there are also important features of consumption that appear inconsistent with the permanent-income hypothesis. For example, as described in Section 7.3, both macroeconomic and microeconomic evidence suggest that consumption responds to predictable changes in income. And as we just saw, simple models of consumer optimization cannot account for the equity premium.

---

[13] Cochrane and Hansen (1992) provide an overview of work on the puzzle and a framework for thinking about proposed explanations. For some proposed explanations, see Mankiw (1986b); Mankiw and Zeldes (1991); Constantinides (1990); Campbell and Cochrane (1995); Weil (1989b); Epstein and Zin (1991); and Problem 7.10.

Because of these and other difficulties, there has been considerable work on extensions or alternatives to the permanent-income hypothesis. This section touches on some of the issues raised by these theories.[14]

Precautionary Saving and the Growth of Consumption

Recall that our derivation of the random-walk result in Section 7.2 was based on the assumption that utility is quadratic. Quadratic utility requires, however, that marginal utility reaches zero at some finite level of consumption and then becomes negative. It also implies that the utility cost of a given variance of consumption is independent of the level of consumption. Since the marginal utility of consumption is declining, individuals have increasing absolute risk aversion: the amount of consumption they are willing to give up to avoid a given amount of uncertainty about the level of consumption rises as they become wealthier. These difficulties with quadratic utility suggest that marginal utility falls more slowly as consumption rises; that is, the third derivative of utility is probably positive rather than zero.

To see the effects of a positive third derivative, assume that both the real interest rate and the discount rate are zero, and consider again the Euler equation relating consumption in consecutive periods, equation (7.20): $u'(C_t) = E_t[u'(C_{t+1})]$. As described in Section 7.2, if utility is quadratic, marginal utility is linear, and so $E_t[u'(C_{t+1})]$ equals $u'(E_t[C_{t+1}])$; thus in this case, the Euler equation reduces to $C_t = E_t[C_{t+1}]$. But if $u'''(\bullet)$ is positive, then $u'(C)$ is a convex function of C. Thus in this case $E_t[u'(C_{t+1})]$ exceeds $u'(E_t[C_{t+1}])$. But this means that if C_t and $E_t[C_{t+1}]$ are equal, $E_t[u'(C_{t+1})]$ is greater than $u'(C_t)$, and so a marginal reduction in C_t increases expected utility. Thus the combination of a positive third derivative of the utility function and uncertainty about future income reduces current consumption, and thus raises saving. This saving is known as *precautionary saving* (Leland, 1968).

Panel (a) of Figure 7.3 shows the impact of uncertainty and a positive third derivative of the utility function on the expected marginal utility of consumption. Since $u''(C)$ is negative, $u'(C)$ is decreasing in C. And since $u'''(C)$ is positive, $u'(C)$ declines less rapidly as C rises—that is, $u'(C)$ is convex. If consumption takes on only two possible values, C_A and C_B, each with probability $\frac{1}{2}$, the expected marginal utility of consumption is the average of marginal utility at these two values. In terms of the diagram, this is shown by the midpoint of the line connecting $u'(C_A)$ and $u'(C_B)$. As the

[14]Three extensions of the permanent-income hypothesis that we will not discuss are durability of consumption goods, habit formation, and nonexpected utility. For durability, see Mankiw (1982); Caballero (1990a, 1993); Eberly (1994); and Problem 7.6. For habit formation, see Deaton (1992, pp. 29-34, 99-100) and Campbell and Cochrane (1995). For nonexpected utility, see Weil (1989b, 1990) and Epstein and Zin (1989, 1991).

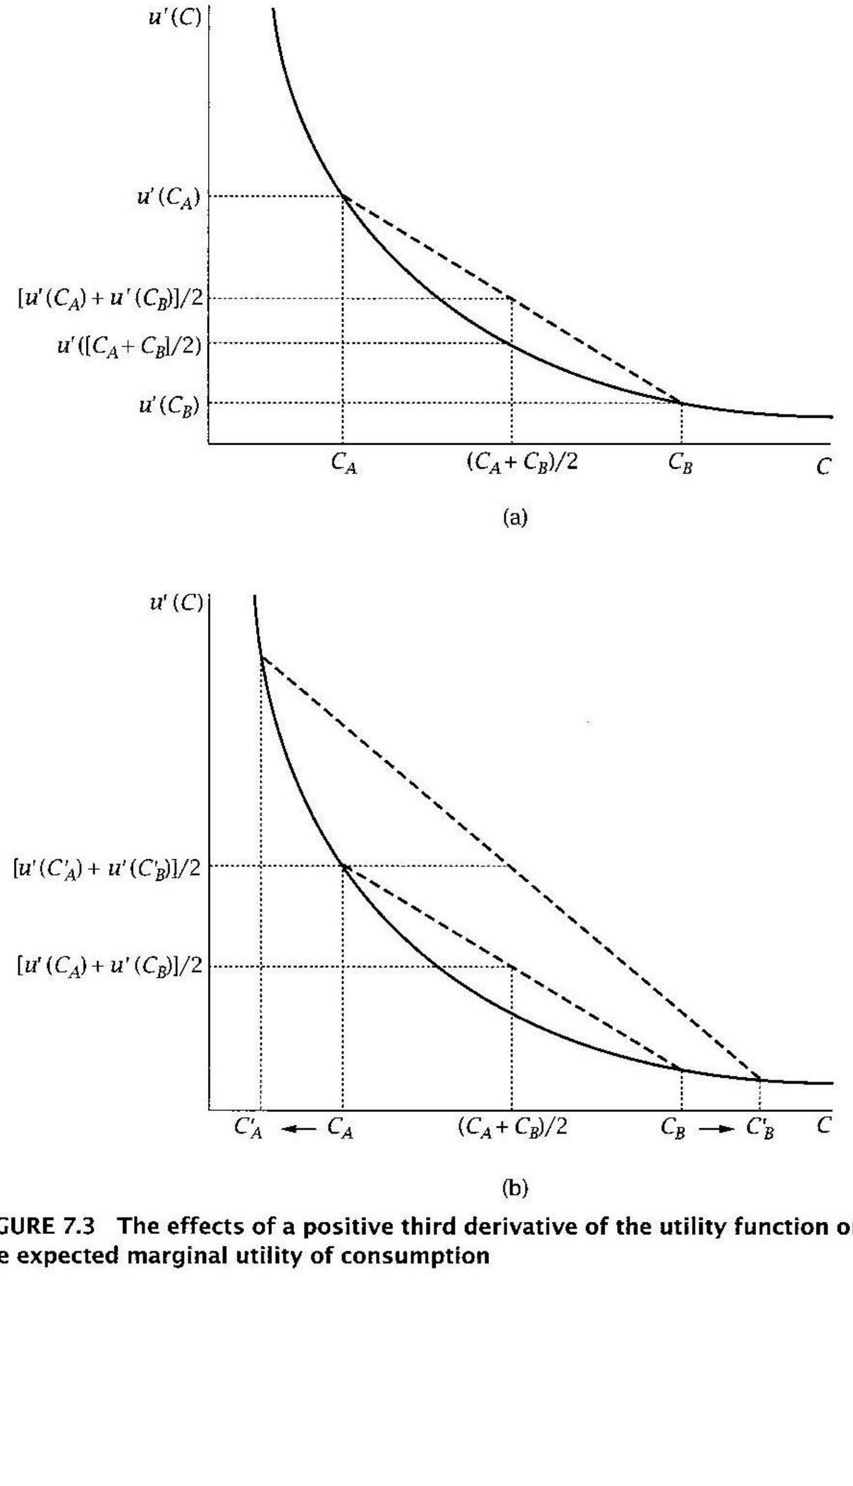

FIGURE 7.3 The effects of a positive third derivative of the utility function on the expected marginal utility of consumption

diagram shows, the fact that $u'(C)$ is convex implies that this quantity is larger than marginal utility at the average value of consumption, $(C_A+C_B)/2$.

Panel (b) shows the effects of an increase in uncertainty. When the high value of consumption rises, the fact that $u'''(C)$ is positive means that marginal utility falls relatively little; but when the low value falls, the positive third derivative magnifies the rise in marginal utility. As a result, the increase in uncertainty raises expected marginal utility for a given value of expected consumption. Thus the increase in uncertainty raises the incentive to save.

An important question, of course, is whether precautionary saving is quantitatively important. To address this issue, recall that in our analysis of the equity premium we found that the Euler equation for the risk-free asset is $\overline{r} \simeq \rho + \theta E[g^c] - \theta(\theta+1)\text{Var}(g^c)/2$ (see [7.40]). For the case of $\overline{r} = \rho$, this becomes

$$E[g^c] \simeq \frac{1}{2}(\theta+1)\text{Var}(g^c). \tag{7.42}$$

Thus the impact of precautionary saving on expected consumption growth depends on the variance of consumption growth and the coefficient of relative risk aversion.[15] If both are substantial, precautionary saving can have a large effect on expected consumption growth. If the coefficient of relative risk aversion is 4 (which is toward the high end of values that are viewed as plausible), and the standard deviation of households' uncertainty about their consumption a year ahead is 0.1 (which is consistent with the evidence in Dynan, 1993, and Carroll, 1992), (7.42) implies that precautionary saving raises expected consumption growth by $(1/2)(4+1)(0.1)^2$, or 2.5 percentage points.[16]

Finally, the presence of precautionary saving implies that not just expectations of future income but also uncertainty about that income affects consumption. C. Romer (1990), for example, argues that the tremendous uncertainty generated by the stock-market crash of 1929 and by the subsequent gyrations of the stock market was a major force behind the sharp fall in consumption spending in 1930, and thus behind the onset of the Great Depression. To give another example, Barsky, Mankiw, and Zeldes (1986) show that the combination of a current tax cut and an offsetting increase in future tax rates reduces households' uncertainty about their lifetime after-tax resources. Thus when there is precautionary saving, this change raises current consumption. More generally, Caballero (1990b) observes that, for a given level of expected lifetime resources, uncertainty is likely to be larger when

[15]For a general utility function, the $\theta+1$ term is replaced by $-Cu'''(C)/u''(C)$. In analogy to the coefficient of relative risk aversion, $-Cu''(C)/u'(C)$, Kimball (1990) refers to $-Cu'''(C)/u''(C)$ as the coefficient of relative prudence.

[16]For more on the impact of precautionary saving on the level of aggregate consumption, see Skinner (1988); Caballero (1991); and Aiyagari (1994).

more of those resources are expected to come in the future. As a result, precautionary saving can help to account for the fact that when income is expected to rise, consumption is also expected to rise. Finally, Dynan (1993) and Carroll (1994) investigate the empirical relationship between households' uncertainty about their future income and consumption growth; they reach conflicting conclusions, however.

Liquidity Constraints

The permanent-income hypothesis assumes that individuals can borrow at the same interest rate at which they can save as long as they eventually repay their loans. Yet the interest rates that households pay on credit-card debt, automobile loans, and other borrowing are often much higher than the rates they obtain on their saving. In addition, some individuals are unable to borrow more at any interest rate.

A large literature investigates the causes, extent, and effects of such liquidity constraints. They are potentially important for many aspects of consumption. As described in Section 7.3, they can produce excess sensitivity of consumption to predictable changes in income. If individuals face high interest rates for borrowing, they may choose not to borrow to smooth their consumption when their current resources are low. And if they cannot borrow at all, they have no choice but to have low consumption when their current resources are low. Thus liquidity constraints can cause current income to be more important to consumption than is predicted by the permanent-income hypothesis.

This chapter will not provide a thorough treatment of liquidity constraints.[17] Instead, as with our discussion of precautionary saving, we will focus on the potential effects of liquidity constraints on the level of consumption.

Liquidity constraints can raise saving in two ways. First, and most obviously, whenever a liquidity constraint is binding, it causes the individual to consume less than he or she otherwise would. Second, as Zeldes (1989) emphasizes, even if the constraints are not currently binding, the fact that they may bind in the future reduces current consumption. Suppose, for example, that there is some chance of low income in the next period. If there are no liquidity constraints, and if income in fact turns out to be low, the individual can borrow to avoid a sharp fall in consumption. If there are liquidity constraints, however, the fall in income causes a large fall in consumption unless the individual has savings. Thus the presence of liquidity constraints

[17] See Deaton (1992, pp. 194–213) for a general introduction to liquidity constraints. In addition, Section 8.7 presents a model of capital-market imperfections in the context of loans to firms rather than to households.

causes individuals to save as insurance against the effects of future falls in income.

These points can be seen in a three-period model. To distinguish the effects of liquidity constraints from precautionary saving, assume that the instantaneous utility function is quadratic. In addition, continue to assume that the real interest rate and the discount rate are zero.

Begin by considering the individual's behavior in period 2. Let A_t denote assets at the end of period t. Since the individual lives for only three periods, C_3 equals $A_2 + Y_3$, which in turn equals $A_1 + Y_2 + Y_3 - C_2$. The individual's expected utility over the last two periods of life as a function of his or her choice of C_2 is therefore

$$U = (C_2 - \frac{1}{2}aC_2^2) + E_2[(A_1 + Y_2 + Y_3 - C_2) - \frac{1}{2}a(A_1 + Y_2 + Y_3 - C_2)^2]. \quad (7.43)$$

The derivative of this expression with respect to C_2 is

$$\begin{aligned} \frac{\partial U}{\partial C_2} &= 1 - aC_2 - (1 - aE_2[A_1 + Y_2 + Y_3 - C_2]) \\ &= a(A_1 + Y_2 + E_2[Y_3] - 2C_2). \end{aligned} \quad (7.44)$$

This expression is positive for $C_2 < (A_1 + Y_2 + E_2[Y_3])/2$, and negative thereafter. Thus, as we know from our earlier analysis, if the liquidity constraint does not bind, the individual chooses $C_2 = (A_1 + Y_2 + E_2[Y_3])/2$. But if it does bind, he or she sets consumption to the maximum attainable level, which is $A_1 + Y_2$. Thus,

$$C_2 = \min\left[\frac{A_1 + Y_2 + E_2[Y_3]}{2}, A_1 + Y_2\right]. \quad (7.45)$$

Thus the liquidity constraint reduces current consumption if it is binding.

Now consider the first period. If the liquidity constraint is not binding that period, the individual has the option of marginally raising C_1 and paying for this by reducing C_2. Thus if the individual's assets are not literally zero, the usual Euler equation holds. With the specific assumptions we are making here, this means that C_1 equals the expectation of C_2.

But the fact that the Euler equation holds does not mean that the liquidity constraints do not affect consumption. Equation (7.45) implies that if the probability that the liquidity constraint will bind in the second period is strictly positive, the expectation of C_2 as of period 1 is strictly less than the expectation of $(A_1 + Y_2 + E_2[Y_3])/2$. A_1 is given by $A_0 + Y_1 - C_1$, and the law of iterated projections implies that $E_1[E_2[Y_3]]$ equals $E_1[Y_3]$. Thus,

$$C_1 < \frac{A_0 + Y_1 + E_1[Y_2] + E_1[Y_3] - C_1}{2}. \quad (7.46)$$

Adding $C_1/2$ to both sides of this expression and then dividing by $3/2$ yields

$$C_1 < \frac{A_0 + Y_1 + E_1[Y_2] + E_1[Y_3]}{3}. \tag{7.47}$$

Thus even when the liquidity constraint does not bind currently, the possibility that it will bind in the future reduces consumption.

Finally, if the value of C_1 that satisfies $C_1 = E_1[C_2]$ (given that C_2 is determined by [7.45]) is greater than the individual's period-1 resources, $A_0 + Y_1$, the first-period liquidity constraint is binding; in this case the individual consumes $A_0 + Y_1$.[18]

Empirical Application: Liquidity Constraints and Aggregate Saving

As we have just seen, liquidity constraints can raise saving. Jappelli and Pagano (1994) investigate empirically whether cross-country differences in liquidity constraints are important to cross-country differences in aggregate saving.

Jappelli and Pagano begin by arguing that there are important differences in the extent of liquidity constraints across countries. In Spain and Japan, for example, home purchases generally require down payments of 40% of the purchase price, whereas in the United States and France they require 20% or less. Similarly, Korea strongly restricts the availability of consumer credit, but the Scandinavian countries do not. Bankruptcy and foreclosure laws also vary greatly. In Belgium and Spain, for example, it takes two years or more to foreclose on a mortgage, whereas in Denmark and the Netherlands it takes less than six months. Greater legal barriers to foreclosure are likely to discourage lending.

Jappelli and Pagano then ask whether these differences in credit availability are associated with differences in saving rates. They first examine the relationship between the loan-to-value ratio for home purchases (that is, one minus the required down payment) and the saving rate. As Figure 7.4 shows, there is a clear negative association. They then add the loan-to-value ratio to a regression of saving rates on measures of government saving, the demographic composition of the population, and income growth. The loan-to-value ratio enters negatively and significantly. In a typical specification, the point estimates imply that an increase in the required down payment of 10 percent of the purchase price is associated with a rise in the saving rate of 2 percent of NNP. They also find that using a measure of the availability of consumer credit in place of the loan-to-value ratio yields similar results.

[18]Because both present and future liquidity constraints potentially affect behavior, complete solutions of models with liquidity constraints usually require the use of numerical methods (see, for example, Deaton, 1992, pp. 180–189).

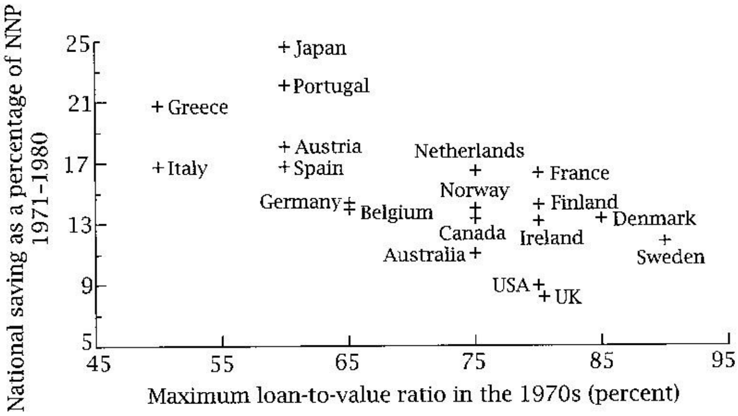

FIGURE 7.4 The loan-to-value ratio for home purchases and the saving rate (from Jappelli and Pagano, 1994; used with permission)

In sum, their evidence suggests that liquidity constraints are important to aggregate saving.[19]

Empirical Application: Buffer-Stock Saving

A central prediction of the permanent-income hypothesis is that there should be no relation between the expected growth of an individual's income over his or her lifetime and the expected growth of his or her consumption: consumption growth is determined by the real interest rate and the discount rate, not by the time pattern of income.

Carroll and Summers (1991) present extensive evidence that this prediction of the permanent-income hypothesis is incorrect. For example, individuals in countries where income growth is high typically have high rates of consumption growth over their lifetimes, and individuals in slowly growing countries typically have low rates of consumption growth. Similarly, typical lifetime consumption patterns of individuals in different occupations tend to match typical lifetime income patterns in those occupations. Managers and professionals, for example, generally have earnings profiles that rise steeply until middle age and then level off; their consumption profiles follow a similar pattern.

More generally, most households have little wealth (see, for example, Deaton, 1991, and Hubbard, Skinner, and Zeldes, 1994a). Their consumption

[19] Jappelli and Pagano go on to investigate the relationship between liquidity constraints and aggregate growth. They find that even when they control for investment, liquidity constraints are positively related to growth. Given that the way that liquidity constraints most plausibly affect growth is through their effect on saving (and hence investment), this finding is difficult to interpret.

approximately tracks their income, but they have a small amount of saving that they use in the event of sharp falls in income or emergency spending needs. In the terminology of Deaton (1991), most households exhibit *buffer-stock* saving behavior. As a result, a small fraction of households hold the vast majority of wealth.

At least three explanations of buffer-stock saving have been proposed. First, Shefrin and Thaler (1988) argue that consumption behavior is not well described by complete intertemporal optimization (see also Laibson, 1993). Instead, individuals have a set of rules of thumb that they use to guide their consumption behavior. Examples of these rules of thumb are that it is usually reasonable to spend one's current income, but that assets should be dipped into only in exceptional circumstances. Such rules of thumb may lead consumers to use saving and borrowing to smooth short-run income fluctuations, and thus cause consumption to follow the predictions of the permanent-income hypothesis reasonably well at short horizons. But they may also cause consumption to track income fairly closely over long horizons.

Second, Deaton (1991) and Carroll (1992) argue that buffer-stock saving arises from a combination of a high discount rate, a precautionary-saving motive, and some reason that households do not go heavily into debt. In Deaton's analysis, the reason for the absence of debt is the presence of liquidity constraints. In Carroll's, it is that the marginal utility of consumption approaches infinity as consumption becomes sufficiently low; as a result, households are unwilling to risk the very low consumption that would occur if they were in debt and their future income turned out to be low. The combination of the high discount rate and the inability or unwillingness to go into debt causes households' wealth to be approximately zero, and thus causes consumption to approximately track income. But even with a relatively high discount rate, the positive third derivative of the utility function causes households to view the risks of sharp falls in consumption and sharp rises as asymmetric; as a result, they typically keep a small amount of savings to use in the event of large falls in income.

Third, Hubbard, Skinner, and Zeldes (1994a, 1994b) suggest an explanation of buffer-stock saving that is close in spirit to the permanent-income hypothesis. The key elements of their explanation, aside from intertemporal optimization, are a precautionary-saving motive and the fact that welfare programs provide insurance against very low levels of consumption. For households that face a nonnegligible probability of going on welfare, the presence of welfare discourages saving in two ways: it directly provides insurance against unfavorable realizations of income, and it imposes extremely high implicit tax rates on asset holdings. Nonetheless, the precautionary-saving motive causes these households to typically hold some assets when their consumption is above the guaranteed floor. For households whose income prospects are favorable enough that the possibility of going on welfare is negligible, on the other hand, consumption is determined by conventional intertemporal optimization; thus they exhibit conventional life-cycle saving. Hubbard, Skinner, and Zeldes therefore argue that the different patterns of wealth accumulation of the poor and the rich can be explained without appealing to differences in their preferences.

Problems

7.1. The average income of farmers is less than the average income of non-farmers, but fluctuates more from year to year. Given this, how does the permanent-income hypothesis predict that estimated consumption functions for farmers and nonfarmers differ?

7.2. **The time-averaging problem.** (Working, 1960.) Actual data do not give consumption at a point in time, but average consumption over an extended period, such as a quarter. This problem asks you to examine the effects of this fact.

Suppose that consumption follows a random walk: $C_t = C_{t-1} + e_t$, where e is white noise. Suppose, however, that the data provide average consumption over two-period intervals; that is, one observes $(C_t + C_{t-1})/2$, $(C_{t+2} + C_{t+3})/2$, and so on.

(*a*) Find an expression for the change in measured consumption from one two-period interval to the next in terms of the e's.

(*b*) Is the change in measured consumption uncorrelated with the previous value of the change in measured consumption? In light of this, is measured consumption a random walk?

(*c*) Given your result in part (*a*), is the change in consumption from one two-period interval to the next necessarily uncorrelated with anything known as of the first of these two-period intervals? Is it necessarily uncorrelated with anything known as of the two-period interval immediately preceding the first of the two-period intervals?

(*d*) Suppose that measured consumption for a two-period interval is not the average over the interval, but consumption in the second of the two periods. That is, one observes C_{t+1}, C_{t+3}, and so on. In this case, is measured consumption a random walk?

7.3. (This follows Hansen and Singleton, 1983.) Suppose instantaneous utility is of the constant-relative-risk-aversion form, $u(C_t) = C_t^{1-\theta}/(1-\theta)$, $\theta > 0$. Assume that the real interest rate, r, is constant but not necessarily equal to the discount rate, ρ.

(*a*) Find the Euler equation relating C_t to expectations concerning C_{t+1}.

(*b*) Suppose that the log of income is distributed normally, and that as a result the log of C_{t+1} is distributed normally; let σ^2 denote its variance conditional on information available at time t. Rewrite the expression in part (*a*) in terms of $\ln C_t$, $E_t[\ln C_{t+1}]$, σ^2, and the parameters r, ρ, and θ. (Hint: if a variable x is distributed normally with mean μ and variance V, $E[e^x] = e^{\mu}e^{V/2}$.)

(c) Show that if r and σ^2 are constant over time, the result in part (b) implies that the log of consumption follows a random walk with drift: $\ln C_{t+1} = a + \ln C_t + u_{t+1}$, where u is white noise.

(d) How do changes in each of r and σ^2 affect expected consumption growth, $E_t[\ln C_{t+1} - \ln C_t]$? Interpret the effect of σ^2 on expected consumption growth in light of the discussion of precautionary saving in Section 7.6.

7.4. **A framework for investigating excess smoothness.** Suppose that C_t equals $[r/(1+r)][A_t + \sum_{s=0}^{\infty} E_t[Y_{t+s}]/(1+r)^s]$, and that $A_{t+1} = (1+r)(A_t + Y_t - C_t)$.

(a) Show that these assumptions imply that $E_t[C_{t+1}] = C_t$ (and thus that consumption follows a random walk) and that $\sum_{s=0}^{\infty} E_t[C_{t+s}]/(1+r)^s = A_t + \sum_{s=0}^{\infty} E_t[Y_{t+s}]/(1+r)^s$.

(b) Suppose that $\Delta Y_t = \phi \Delta Y_{t-1} + u_t$, where u is white noise. Suppose that Y_t exceeds $E_{t-1}[Y_t]$ by one unit (that is, suppose $u_t = 1$). By how much does consumption increase?

(c) For the case of $\phi > 0$, which has a larger variance, the innovation in income, u_t, or the innovation in consumption, $C_t - E_{t-1}[C_t]$? Do consumers use saving and borrowing to smooth the path of consumption relative to income in this model? Explain.

7.5. Consider the two-period setup analyzed in Section 7.4. Suppose that the government initially raises revenue only by taxing interest income. Thus the individual's budget constraint is $C_1 + C_2/[1+(1-\tau)r] \leq Y_1 + Y_2/[1+(1-\tau)r]$, where τ is the tax rate. The government's revenue is zero in period 1 and $\tau r(Y_1 - C_1^0)$ in period 2, where C_1^0 is the individual's choice of C_1 given this tax rate. Now suppose the government eliminates the taxation of interest income and instead institutes lump-sum taxes of amounts T_1 and T_2 in the two periods; thus the individual's budget constraint is now $C_1 + C_2/(1+r) \leq (Y_1 - T_1) + (Y_2 - T_2)/(1+r)$. Assume that Y_1, Y_2, and r are exogenous.

(a) What condition must the new taxes satisfy so that the change does not affect the present value of government revenues?

(b) If the new taxes satisfy the condition in part (a), is the old consumption bundle, (C_1^0, C_2^0), not affordable, just affordable, or affordable with room to spare?

(c) If the new taxes satisfy the condition in part (a), does first-period consumption rise, fall, or stay the same?

7.6. **Consumption of durable goods.** (Mankiw, 1982.) Suppose that, as in Section 7.2, the instantaneous utility function is quadratic and the interest rate and the discount rate are zero. Suppose, however, that goods are durable; specifically, $C_t = (1-\delta)C_{t-1} + E_t$, where E_t is purchases in period t and $0 \leq \delta < 1$.

(a) Consider a marginal reduction in purchases in period t of dE_t. Find values of dE_{t+1} and dE_{t+2} such that the combined changes in E_t, E_{t+1}, and E_{t+2} leave the present value of spending unchanged (so $dE_t + dE_{t+1} + dE_{t+2} = 0$) and leave C_{t+2} unchanged (so $(1-\delta)^2 dE_t + (1-\delta)dE_{t+1} + dE_{t+2} = 0$).

(b) What is the effect of the change in part (a) on C_t and C_{t+1}? What is the effect on expected utility?

(*c*) What condition must C_t and $E_t[C_{t+1}]$ satisfy for the change in part (*a*) not to affect expected utility? Does C follow a random walk?

(*d*) Does E follow a random walk? (Hint: write $E_t - E_{t-1}$ in terms of $C_t - C_{t-1}$ and $C_{t-1} - C_{t-2}$.) Explain intuitively. If $\delta = 0$, what is the behavior of E?

7.7. Consider a stock that pays dividends of D_t in period t and whose price in period t is P_t. Assume that consumers are risk-neutral and have a discount rate of r; thus they maximize $E[\sum_{t=0}^{\infty} C_t/(1+r)^t]$.

(*a*) Show that equilibrium requires $P_t = E_t[(D_{t+1} + P_{t+1})/(1+r)]$ (assume that if the stock is sold, this happens after that period's dividends have been paid).

(*b*) Assume that $\lim_{s\to\infty} E_t[P_{t+s}/(1+r)^s] = 0$ (this is a *no-bubbles* condition; see the next problem). Iterate the expression in part (*a*) forward to derive an expression for P_t in terms of expectations of future dividends.

7.8. Bubbles. Consider the setup of the previous problem without the assumption that $\lim_{s\to\infty} E_t[P_{t+s}/(1+r)^s] = 0$.

(*a*) **Deterministic bubbles.** Suppose that P_t equals the expression derived in part (*b*) of Problem 7.7 plus $(1+r)^t b$, $b > 0$.

(*i*) Is consumers' first-order condition derived in part (*a*) of Problem 7.7 still satisfied?

(*ii*) Can b be negative? (Hint: consider the strategy of never selling the stock.)

(*b*) **Bursting bubbles.** (Blanchard, 1979.) Suppose that P_t equals the expression derived in part (*b*) of Problem 7.7 plus q_t, where q_t equals $(1+r)q_{t-1}/\alpha$ with probability α and equals zero with probability $1-\alpha$.

(*i*) Is consumers' first-order condition derived in part (*a*) of Problem 7.7 still satisfied?

(*ii*) If there is a bubble at time t (that is, if $q_t > 0$), what is the probability that the bubble has burst by time $t+s$ (that is, that $q_{t+s} = 0$)? What is the limit of this probability as s approaches infinity?

(*c*) **Intrinsic bubbles.** (Froot and Obtsfeld, 1991.) Suppose that dividends follow a random walk: $D_t = D_{t-1} + e_t$, where e is white noise.

(*i*) In the absence of bubbles, what is the price of the stock in period t?

(*ii*) Suppose that P_t equals the expression derived in (*i*) plus b_t, where $b_t = (1+r)b_{t-1} + ce_t$, $c > 0$. Is consumers' first-order condition derived in part (*a*) of Problem 7.7 still satisfied? In what sense do stock prices overreact to changes in dividends?

7.9. The Lucas asset-pricing model. (Lucas, 1978.) Suppose the only assets in the economy are infinitely-lived trees. Output equals the fruit of the trees, which is exogenous and cannot be stored; thus $C_t = Y_t$, where Y_t is the exogenously determined output per person and C_t is consumption per person. Assume that initially each consumer owns the same number of trees. Since all consumers are assumed to be the same, this means that, in equilibrium, the behavior of the price of trees must be such that, each period, the representative