Volume Title: National Saving and Economic Performance

Volume Author/Editor: B. Douglas Bernheim and John B. Shoven, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-04404-1

Volume URL: http://www.nber.org/books/bern91-2

Conference Date: January 6-7, 1989

Publication Date: January 1991

Chapter Title: Consumption Growth Parallels Income Growth: Some New Evidence

Chapter Author: Christopher D. Carroll, Lawrence H. Summers

Chapter URL: http://www.nber.org/chapters/c5995

Chapter pages in book: (p. 305 - 348)

10 Consumption Growth Parallels Income Growth: Some New Evidence

Christopher D. Carroll and Lawrence H. Summers

10.1 Introduction

The idea that consumers allocate their consumption over time so as to maximize a stable individualistic utility function provides the basis for almost all modern work on the determinants of consumption and saving decisions. The celebrated life-cycle and permanent income hypotheses represent not so much alternative theories of consumption as alternative empirical strategies for fleshing out the same basic idea. While tests of particular implementations of these theories sometimes lead to statistical rejections, life-cycle/permanent income theories succeed in unifying a wide range of diverse phenomena. It is probably fair to accept Franco Modigliani's (1980) characterization that "the Life Cycle Hypothesis has proved a very fruitful hypothesis, capable of integrating a large variety of facts concerning individual and aggregate saving behaviour."

This paper argues, however, that both permanent income and, to an only slightly lesser extent, life-cycle theories as they have come to be implemented in recent years are inconsistent with the grossest features of cross-country and cross-section data on consumption and income and income growth. There is clear evidence that consumption growth and income growth are much more closely linked than these theories predict. It appears that consumption smoothing takes place over periods of several years not several decades.

Christopher D. Carroll is an economist at the Federal Reserve Board in Washington, D.C. Lawrence H. Summers is professor of economics at Harvard University and a research associate of the National Bureau of Economic Research.

The authors gratefully acknowledge funding from the National Science Foundation. Also thanked are David Cutler, James Poterba, Greg Mankiw, Olivier Blanchard, Daniel Feenberg, and the members of the M.I.T. Public Finance Lunch, Money Lunch, and Money Workshop for constructive comments. Thanks especially to Knut Morck for generous help in obtaining data, and to Steve Zeldes for discussions on uncertainty. Remaining errors are of course the authors'. Views expressed here do not necessarily reflect those of the Federal Reserve Board or its staff.

These results confirm Milton Friedman's (1957) initial view that "the permanent income component is not to be regarded as expected lifetime earnings. . . . It is to be interpreted as the mean income at any age regarded as permanent by the consumer unit in question, which in turn depends on its horizon and foresightedness." They call into question the usefulness of standard representative consumer approaches to the analysis of saving behavior. And they call for increased emphasis on liquidity constraints and short-run precautionary saving as determinants of consumption behavior.

This paper is divided into five sections. Section 10.2 presents the rational expectations version of the permanent income hypothesis, which has been increasingly popular in empirical macroeconomics recently, and draws out the low frequency implications of this hypothesis. The principal implications on which we concentrate are, first, that (absent capital market imperfections) the anticipated rate of growth of income should be unrelated to the rate of growth of consumption and, second, that the rate of interest should be a powerful determinant of the rate of growth of consumption. We present evidence that challenges both of these propositions. We demonstrate that over periods of several years there is nearly perfect equality between rates of income growth and consumption growth. These facts hold both across countries and, within countries, across different eras when productivity increased at different rates. The prediction of the permanent income hypothesis that consumption growth and anticipated income growth are unrelated is clearly refuted. We next argue that these facts cannot be explained by imperfections in the international capital market, since there is no evidence that countries with more rapid consumption growth have higher rates of return on bonds or other assets.

Section 10.3 asks whether recognizing that consumers have finite lifetimes helps in understanding these stylized facts. This is plausible a priori. Because the gap in lifetime income between old and new generations ones is greater in rapidly growing than in slowly growing countries, the life-cycle hypothesis would predict that consumption growth should equal income growth looking across countries with permanently different productivity growth rates. We find, however, that the life-cycle story is not consistent with the data. Contrary to the predictions of the theory, individual consumers in rapidly growing countries like Japan have had more rapid consumption growth rates than consumers in the United States, where income growth is slower. Indeed, where life-cycle theory predicts that longitudinal age-consumption profiles should be similar in countries with different growth rates, the fact is much more nearly that point-in-time cross-sectional age consumption profiles are similar across countries.

The close international linkages between consumption growth and income growth could arise either because some common factor causes some countries both to defer consumption and to grow rapidly or because individual consumers display more sensitivity to current income than theory suggests they should. In Section 10.4 we seek to distinguish these alternative views by looking at the relationship between income growth and consumption growth for consumers in different occupations and educational categories. Using data from several American Consumer Expenditure Surveys, we discover that there is considerable variation in the lifetime profile of income across categories, and that the lifetime profiles of consumption track the profiles of income very closely.

Section 10.5 uses information on saving rates to confirm the inference drawn in the previous sections that consumers are not responsive to changes in their long-run future income. First, we show that there is no tendency for countries that experience reductions in their expected growth rate to experience short-run increases in saving as theory would predict. Second, we test the pure life-cycle theory's prediction, that when a country experiences a sharp productivity slowdown as the United States has in recent years, there should be a tendency for the relative saving rate of the young to increase greatly. This prediction is not borne out. Third, we document that, contrary to the theory's prediction, there is no tendency for young people in occupations where income rises rapidly to have lower saving rates than those in occupations where income rises slowly.

Section 10.6 discusses the implications of these results for consumption theory. We suggest that both our data and the available time-series evidence are consistent with Milton Friedman's view that people save to smooth consumption over several years in the face of uncertain income but, because of liquidity constraints, caution, or shortsightedness do not seek to smooth consumption over longer horizons. We follow the recent work of Deaton (1989) in arguing for a "buffer stock" view of saving as appropriate for most consumers. This view is supported by tabulations from a longitudinal data set on tax returns suggesting that about 40% of the population never earned more than $100 in dividend and interest income over a six-year period, 30% of the population earned more than $100 in every year, and 30% earned more than $100 in some but not all years. The buffer stock view of saving is attractive in another respect. If the size of the stock is proportional to income, then one would expect to observe the close relation that is actually observed between saving rates and income growth. We also present evidence, however, that suggests that even if the typical consumer may be accurately described by the buffer stock model, the typical saver may not be. This discrepancy is possible if the distribution of saving is more unequal than the distribution of consumption, so that the great majority of dollars saved are not saved by the typical consumer but rather by a small number of very wealthy consumers who have very high saving rates. We argue that the apparent importance of the distinction between the typical consumer and the typical saver is large enough to justify more attention and perhaps to justify different models for the two groups.

Section 10.7 concludes the paper. We begin by discussing the destructive implications of the results for representative consumer approaches to the study of asset pricing, economic growth, and economic fluctuations. We then suggest some constructive implications of the results for understanding international differences in saving rates, takeoffs of economic growth, and the effects of tax policies. Finally, we suggest some directions for future research.

10.2 International Evidence on Consumption and Growth

The representative agent infinite-horizon consumer model is the simplest and probably most commonly used model in studies of intertemporal issues. The Ramsey model (as we will refer to it throughout) provides the basis for the large body of work on consumption that has emanated from the seminal analysis of Hall (1978). The increasing popularity of this framework for analyzing intertemporal income and consumption behavior is suggested by the large literature surveyed in Campbell and Mankiw (1989). The focus of the research described there has been on the relationship between short-run fluctuations in consumption and income and on the nature of substitution between present and future income. Here we focus instead on long-term predictions of the theory.

In the commonly used constant relative risk-aversion formulation, solution of the model gives rise to the first-order condition for a consumer operating under certainty:

$$\dot{c}/c = \sigma(r-\delta), \tag{1}$$

where σ is the elasticity of substitution of consumption, δ is the consumer's subjective discount rate, and r is the interest rate. Under uncertainty, it will continue to be the case that the interest rate is a sufficient statistic for predicting consumption growth. In a world with a well-functioning capital market that equates returns on the safe asset in different countries, the simple model of (1) predicts that consumption growth rates averaged over long time periods should be equalized around the world if tastes for present as opposed to future consumption do not vary across countries.[1] It certainly would not imply that consumption growth rates should bear any particular relation to income growth rates. We shall now argue that this prediction is obviously and dramatically falsified by the recent experience of industrialized economies.

We have gathered data on income and consumption for 15 OECD countries for the period 1960–85.[2] Our sample includes all the major Western European economies, Japan, the United States, and Canada as well as all of the smaller economies for which relatively complete data was available for the entire period. We study the effects of low-frequency variations by looking at differences both across countries and across different time periods in individual countries. For these comparisons, the issues of measurement and time aggregation that have been discussed in the literature on the time-series properties of consumption are not very important. In order to highlight the strength of the patterns in the data we present them graphically.

Figure 10.1, panels a–d, document a stylized fact that any theory of consumption should account for: *at low frequencies there is near perfect equality between consumption growth rates and income growth rates.* When consumption growth rates are plotted against income growth rates the result is almost precisely a 45° line. While figure 10.1*a–c* documents this fact looking across the entire 1960–85 period and two different subperiods, *d* compares the change in income growth with the change in consumption growth between the 1960–73 and 1980–85 periods. We choose these periods so as to avoid the difficulty of assessing when during the 1970s expectations became entrenched

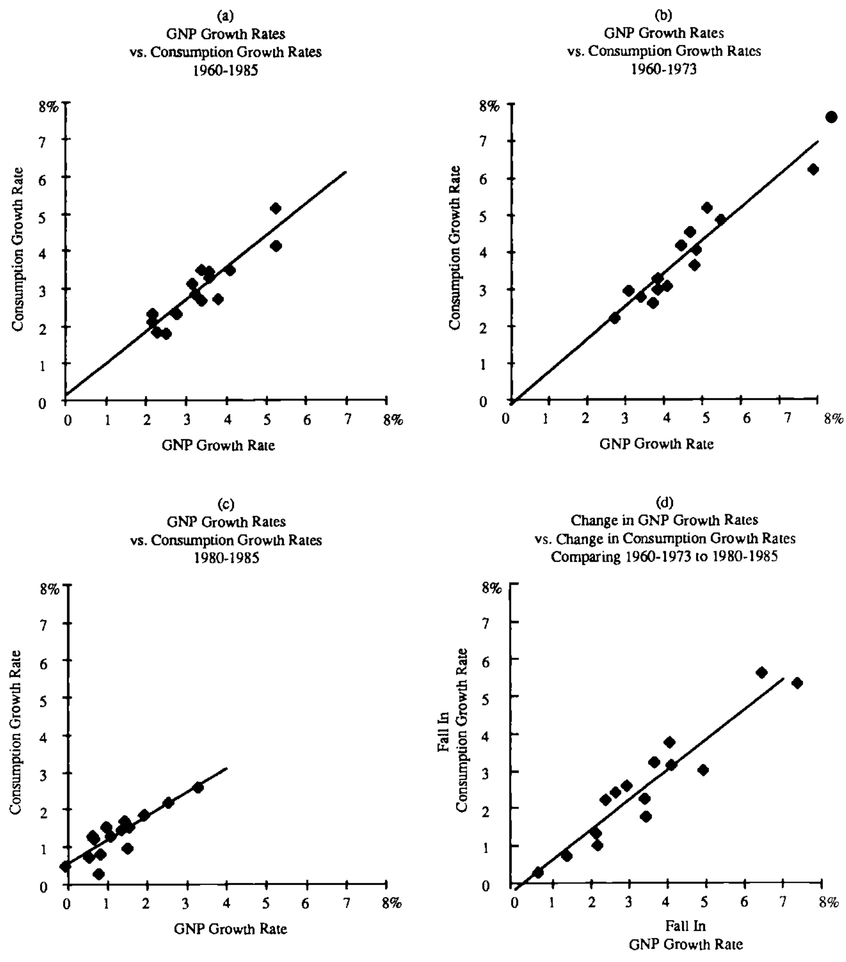

Fig. 10.1 Comparing income growth and consumption growth in the OECD countries, 1960–85

Source: OECD National Income Accounts data.

that the productivity slowdown would last. Again the result is close to a 45° line.

While we have used GDP growth in these comparisons rather than the disposable income measures that would be more appropriate on some views, this and other measurement issues cannot be important. It is easy to see that the consumption growth–income growth regularity has to hold up using almost any measure. Suppose that over a 25-year period a country's saving rate changed by 15 percentage points. This would only alter its consumption growth rate by .6 percentage points, a rather small difference compared to the spread of growth experiences illustrated in figure 10.1*a–d*. In fact, the striking thing about saving rates, whether measured on a private or a national basis, is their stability through time. Comparing the saving rates of the countries in our sample before and after 1973, no country experienced a change of more than 5% in either its private or its national saving rate.[3] This compares with a range of saving rates across countries of over 10%.

Returning now to the Ramsey model, figure 10.1*a–d* appears anomalous in light of the model's implication that the expected rate of growth of consumption should be the same across countries and should be unrelated to the rate of growth of income. We therefore consider in turn whether income surprises, imperfect capital markets, or international differences in tastes can explain the consumption/income parallel within roughly a Ramsey framework.

10.2.1 Income Growth Surprises

One possible objection to direct tests of the independence proposition arises from the possibility that differences in income growth over time were largely unexpected. If the consumer receives information about present or future income she will adjust her *level* of consumption discontinuously to be consistent with her new intertemporal budget constraint. From this new level the proposition will again apply, but if we calculate consumption growth between the period before the information arrived and the period after it arrived we will not observe a growth rate of $\sigma(r-\delta)$. Moving from the abstract to the concrete, this point would be important if, for instance, Japan's continued growth over the postwar period constituted a succession of pleasant surprises that successively caused Japanese consumers to adjust consumption upward in accordance with their new, surprisingly higher, lifetime income.

A first bit of evidence on the plausibility of this scenario is given by figure 10.2, which plots Data Resource, Inc.'s (DRI's) projected income growth for our sample of 15 countries from 1988 to 2000 against their actual growth rates over the period 1976–88. The figure illustrates that there are major differences in expected rates of growth of income across countries. Furthermore, expected future income growth is clearly correlated with past income growth. This suggests that the simplest version of a "surprise" theory, in which any deviation from the average growth rate is unanticipated, is very hard to sustain.

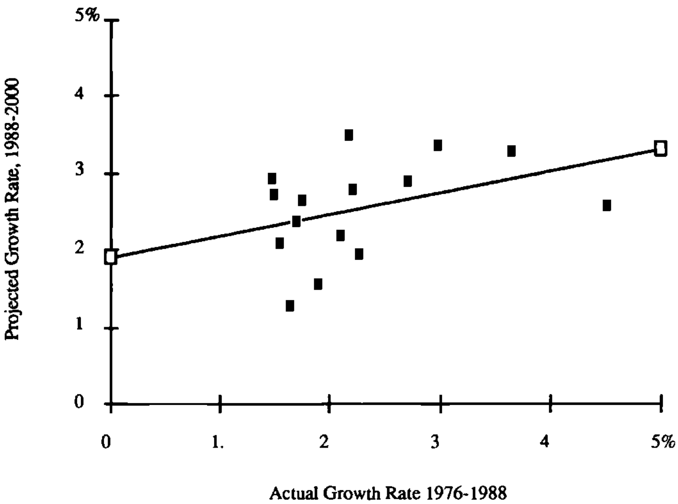

Fig. 10.2 Projected per capita income growth, 1988–2000, versus actual income growth, 1976–88

Source: DRI International Economic Model Database (workspace @INTL/MODELBANK).

Table 10.1 Regressions of Consumption Growth on Income Growth (standard errors in parentheses)

Income Growth Measure	Coefficient on Income Growth[a]	Coefficient on Lagged Income Growth[b]
Current income (OLS)	.601	.253
	(.037)	(.048)
Past 3 years	.725	1.101
	(.220)	(.388)
Past 5 years	.964	.97
	(.194)	(.237)
Past 10 years	1.000	1.14
	(.524)	(.595)

Note: These equations were run using the 15 countries described in the text. Data for 1960–85 were used, and dummies for each year (not reported) were included in all regressions. Eq. (1) runs current consumption growth on current income growth. Eq. (2) forms an expectation of current income growth using the average income growth over the past 3 years. Eq. (3) and (4) form expectations using previous 5-year and previous 10-year growth rate.

[a]This column gives the coefficient when the right-hand-side variable is as just described.

[b]This column gives the coefficient using a one-year lag of the variable just described.

Table 10.1 presents some more formal tests of the idea that the close international correlation between income growth and consumption growth reflects the effects of income surprises. We estimate an international cross section relating consumption growth to measures of expected income growth formed on the basis of past income growth. Each equation includes year dummies, so the identifying variation comes from variations across countries in consumption growth and lagged income growth. The results using measures of income growth over long past periods suggest a nearly one-to-one relationship between expected income growth and consumption growth.[4]

The results using only a single lag of income growth are less strong. However, this is accounted for by the fact that lagged income growth over a long period is a better predictor of contemporaneous income growth than is lagged income growth over a short time period. When past income growth is used as an instrument for expected income growth all specifications suggest a very strong relationship between consumption growth and income growth.

10.2.2 Imperfect Capital Markets and Different Interest Rates

Consider a set of independent closed economies with different rates of exogenous productivity growth. Then theory predicts that each would converge to a steady state with consumption growth equal to income growth. The first-order condition (1) would be satisfied in each country because of differences across countries in the steady-state real rate of interest. More rapidly growing countries would have higher real interest rates. It is possible therefore that the close correlation between consumption growth and income growth is a consequence of imperfections in the international capital market. In this case, one would expect to observe a close relationship between consumption growth rates and rates of return.

Figure 10.3*a–d* illustrates, however, that there is essentially no evidence, looking across countries, that differences in consumption growth rates across countries are explained by differences in real interest rates or other proxies for ex ante returns. This point may be seen most easily by comparing the United States and Japan. It is almost inconceivable that a plausible measure could be found on which ex ante returns were higher in Japan than in the United States in recent years. This evidence is reinforced by figure 10.3*e* which asks whether changes in consumption growth rates in different countries between the pre-1973 period and the post-1980 period are predicted by changes in real interest rates. Perhaps surprisingly, the countries with the greatest declines in consumption growth rates had the smallest increases in real interest rates.

The point that differences in average returns across countries cannot account for differences in consumption growth can be made another way. The range of consumption growth rates in our sample of countries is 3.4%. Most estimates of the intertemporal elasticity of substitution put it at below .25. Even taking the high rate of .25, and assuming that differences in consumption growth rates were perfectly explained by differences in rates of return, the range of rates of return would have to be 13.6%. Persistent differences in safe rates of return of this magnitude over a 25-year period are implausible on even strong views about world capital immobility.

In an influential paper, Mehra and Prescott (1985) have raised questions about the ability of the representative consumer model to account for the risk

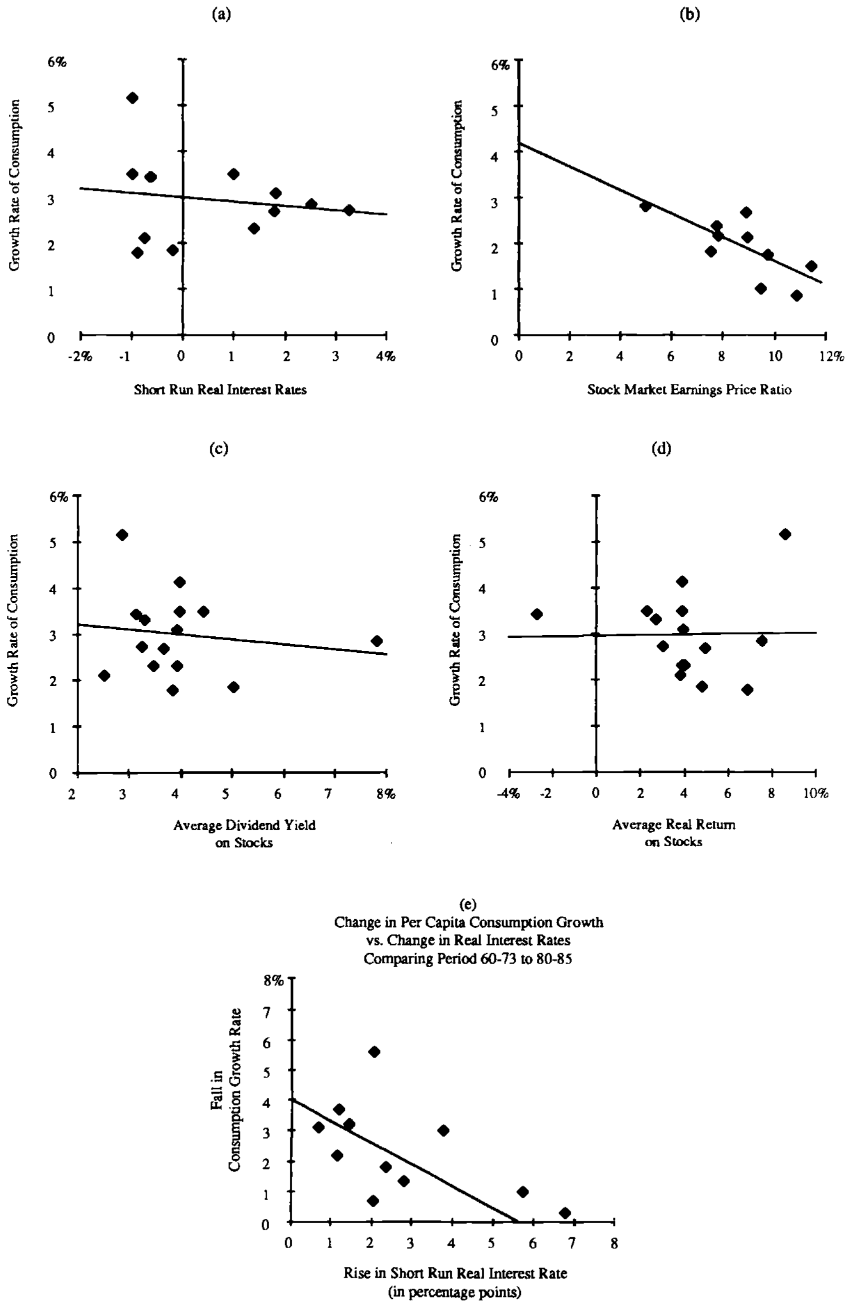

Fig. 10.3 Per capita consumption growth and real rates of return
Source: See data appendix.

premium between debt and equity. This problem is deepened by the apparent absence of correlation between safe interest rates and consumption growth rates across countries. It appears that any successful attempt to rationalize differences in consumption growth rates across countries with fairly similar interest rates would involve postulating a high intertemporal elasticity of substitution. This deepens the difficulty of accounting for the equity risk premium.

10.2.3 Variation in Tastes

One potential channel for reconciling the Ramsey formulation with these facts is to assert that discount rates, δ, differ across countries. If the production technology is of the *Ak* variety discussed by Barro (in this volume), differences in δ would also be associated with differences in steady state growth rates. The same would be true in endogenous growth models relying on increasing returns of the type developed by Paul Romer (1986) and others. Even if there were diminishing returns, one would expect that low δ countries would grow more rapidly while in transition to their steady states (assuming countries started with equal, below-steady-state capital intensity).

We are skeptical that differences in growth across countries and across time primarily reflect taste differences. It seems very implausible to suppose that the primary reasons for the worldwide slowdowns in economic growth rates between the 1960–73 and 1980–87 periods was a taste shock reflecting increased impatience. Yet, since the growth rate of consumption in (1) depends only on tastes and the interest rate, a simultaneous worldwide increase in impatience would be necessary to account for the simultaneous slowing of consumption and income growth.

Even returning to the cross-country consumption growth/income growth relation, the "tastes" theory has a problem. If differences in tastes were a dominant explanation for differences in growth rates there should be a strong tendency for low δ (fast-growing) countries to lend to high δ (slow-growing) countries. As table 10.2 makes clear, this tendency is not apparent in the data. No matter how the data are disaggregated by time there is apparently little or no correlation between trade balances and growth rates.

Note finally that unless an extremely high value of σ is selected, enormous differences across countries in subjective rates of discount are needed to account for the wide range of observed consumption growth rates.

10.2.4 Conclusion

We conclude that there do not appear to be plausible ways of squaring the independence proposition with our facts. While some story involving both variations in r and in δ could be used to account for differences in consumption growth across countries, the problem of explaining why they are so nearly equal to differences in income growth would remain.

Table 10.2 The Relationship between Trade Balances and Growth Rates

Sample Period	Cross-Country Correlation between Average Trade Balance and Average Growth over Sample Period
1961–85	.051
1961–73	.213
1974–85	.045
1961–65	.113
1966–70	.265
1971–75	−.116
1976–80	−.327
1981–85	.222

Sources: DRI @IMF database for trade balance; DRI @OECDNIA database for real GDP Growth. Country sample same as for table 10.1. See appendix for details.

10.3 The Life Cycle and the Consumption/Income Parallel

As a matter of logic, the life-cycle hypothesis is consistent with both the stylized fact that consumption and income growth rates are equated across a sample of countries and the fact that saving and growth rates are positively correlated. To see this, think of a very simple life-cycle model where individuals seek level consumption over their lifetimes. Even though individuals would have level consumption over their lifetimes regardless of their income growth rates, it will nonetheless be true that in steady state, total consumption will grow at the same rate as total income. This is because the gap in lifetime income between old and young generations is greater in rapidly than in slowly growing countries.

Consider the modern life-cycle hypothesis's explanation of the equality between consumption and income growth rates across countries with different growth rates. The essence of the theory (assuming common tastes worldwide and the irrelevance of rate of return differences) is that the rate of growth of consumption for all individuals is the same in all countries. (Implicitly we are assuming rational expectations rather than the myopic expectations assumed by Modigliani in some early statements of the life-cycle hypothesis.) Countries differ in their consumption growth rates only because of the differential effect of the continuous replacement of old, lifetime poor individuals by young, lifetime rich ones.

This argument has two essentially equivalent testable implications. First, tracking the consumption of a given cohort, say those who were 25 in 1950, one should find no difference across countries in the rate of growth of consumption. Second, at a point in time the age-consumption cross-section profile should be less positively sloped in a rapidly growing country than in a slowly growing country. This is because in more rapidly growing countries the old are much lifetime poorer than the young; thus consumption of the old will be much lower relative to consumption of the young. This point is illustrated graphically in figure 10.4*a*. This figure supposes that each individual desires a rate of growth of consumption over his lifetime of 2% annually and demonstrates what the age/consumption cross-section profile should look like in steady state across countries with different growth rates, normalizing the consumption of all individuals by the consumption of individuals at age 20.

Compound interest produces dramatic results here. The ratio of the consumption of the 65-year-olds to the consumption of 25-year-olds should be more than twice as great in countries growing at a 4% rate as in countries growing at a 2% rate. Given the large differences in growth rates illustrated in section 10.2 above, if the life-cycle hypothesis is even approximately accurate, some tendency for consumption of the elderly to be relatively low in rapidly growing countries ought to show up in the international comparisons.

In order to test this proposition, we have obtained cross-sectional point-in-time consumer expenditure profiles by age for Canada, Denmark, Japan, Norway, the United Kingdom, and the United States.[5] Our estimates of the age-consumption profiles are provided in figure 10.4*b*. We have carried the profiles only up to age 65 because of concern that measures of the consumption of the aged are distorted in some countries by the tendency of the poorer elderly to move in with their children.

The results are at odds with the life-cycle hypothesis, since the profiles look quite similar across countries. The similarity of these profiles means that there is no evidence that old people in the slow-growing countries have relatively higher consumption than those in the fast-growing countries. To take a specific example, the profile is more positively sloped in Japan than in the United States, exactly the opposite of what the theory would predict given Japan's much more rapid growth rate. Norway, which has also grown relatively rapidly, also has relatively higher consumption among the aged than the United States. Deaton (1989), using a sample of LDC age-consumption profiles, reaches conclusions similar to those reached here.

This comparison is very crude. But it is instructive to observe how large the differences in age-consumption profiles predicted by the theory would be. Over the 25-year period 1960–85, per capita GNP in Japan grew at 5.2% as compared with 2.1% in the United States. Suppose that we take the Japanese steady-state growth rate to be 4.0% and the U.S. steady-state growth rate to be 2.5%. Then the lifetime income of 30-year-olds in Japan should be 3.94 times the lifetime income of 65-year-olds, compared with a ratio of 2.37 in the United States. This is a difference equal to more than 150% of the income of the average 65-year-old. It is large enough that one would expect it to show up even in our crude measures of age-consumption profiles.[6]

What about the experience of individual cohorts? The longitudinal evidence that we would like to have to answer this question is not available. However,

(a)

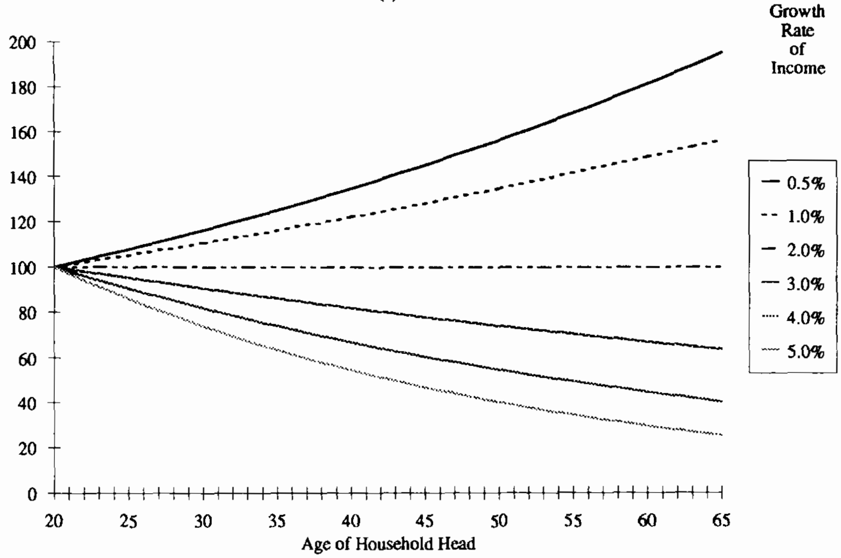

(b)

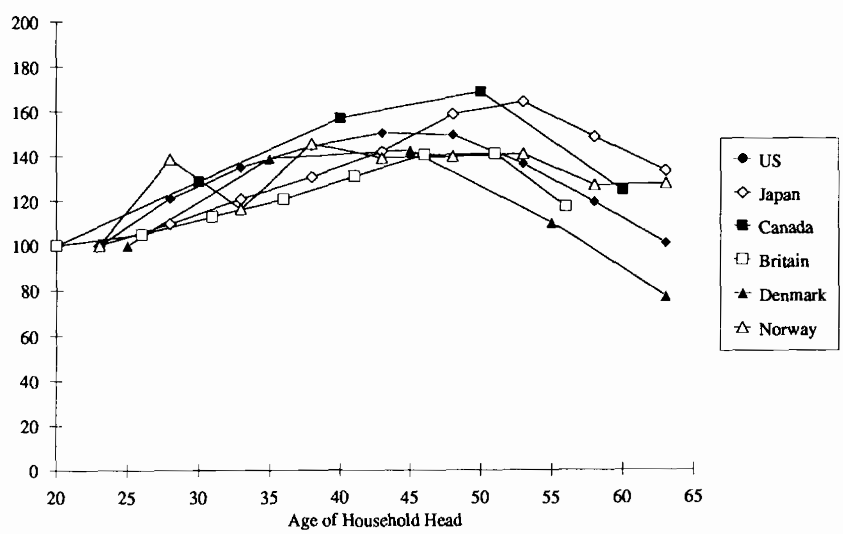

Fig. 10.4 Theoretical and empirical age-consumption cross-section profiles in countries with differing rates of income growth

Source: (*a*) Theoretical calculations described in the text, (*b*) empirical calculations described in the data appendix.

evidence discussed by Kotlikoff and Summers (1981) for the United States and by Ando and Kennickell (1986) for Japan suggests that the shape of age-expenditure profiles is quite stable through time. Figures 10.5*a* and 10.5*b* for these two countries confirm that, between the dates for which we have specific data, the profiles have been fairly stable. If we make the stability assumption for all the countries in our sample, it is possible to trace the consumption of individual cohorts by using data on aggregate consumption and the age structure of the population. If c_i indicates the relative consumption of people in age group i, p_{it} indicates the number of people in this age group in year t, and y_t is total real personal consumption in year t, then we calculate a scaling factor s_t from the equation:

$$y_t = s_t \sum_i c_i p_{it}. \tag{2}$$

Using the scaling factor s_t we calculate real consumption of people of age group i in year t, cr_{it}, from $cr_{it} = c_i s_t$. The results are shown in figure 10.6. Not surprisingly given our results so far, this technique indicates that individuals in fast-growing countries like Japan have enjoyed much more rapid growth in consumption than individuals in slower-growing countries like the United States. How much more rapid? Given that the cross-section profiles are very similar across the whole range of countries in figure 10.4*b*, it follows that *none of the difference in aggregate consumption growth rates across countries can be explained by life-cycle replacement effects.*

10.3.1 Conclusion

While there are obviously many measurement problems here, the data suggest that demographic replacement of the low-consuming aged by the high-consuming young cannot account for the correlation between income growth and consumption growth across countries. If this were the explanation for the correlation there would be large differences across countries in the ratio of the consumption of the old to the consumption of the young. These are not observed.

These results call into question the life-cycle hypothesis' interpretation of the positive correlation between saving and growth. The life-cycle explanation as described, for example, by Modigliani (1967) relies on differences in the ratio of lifetime income among the old and the young to account for the positive relation between saving and income growth. It is not consistent with the observation that individuals in rapidly growing countries enjoy more rapid consumption growth over their lifetimes than individuals in slowly growing countries.

10.4 Tests Using Individual Data

Section 10.2 demonstrated that consumption growth has been very closely related to income growth across both countries and time and argued that this

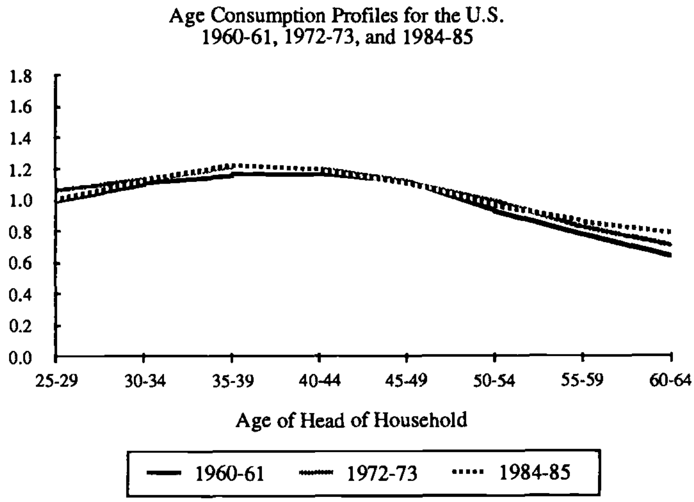

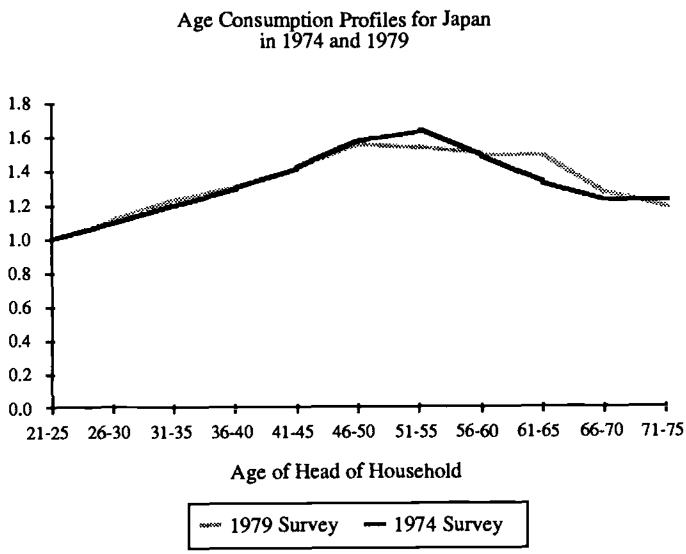

Fig. 10.5 Cross-section age-consumption profiles for the United States and Japan over time

Source: For the United States, direct calculations with the 1960–61, 1972–73, and 1985–86 CES tapes; for Japan, Ando and Kennickell (1986).

was not consistent with the standard Ramsey model. Section 10.3 argued that the consumption/income parallel could not be explained by life-cycle considerations. This leaves two classes of explanations for the apparent international association of consumption growth and income growth. A first possibility is that because consumers are myopic or liquidity constrained or operate on the

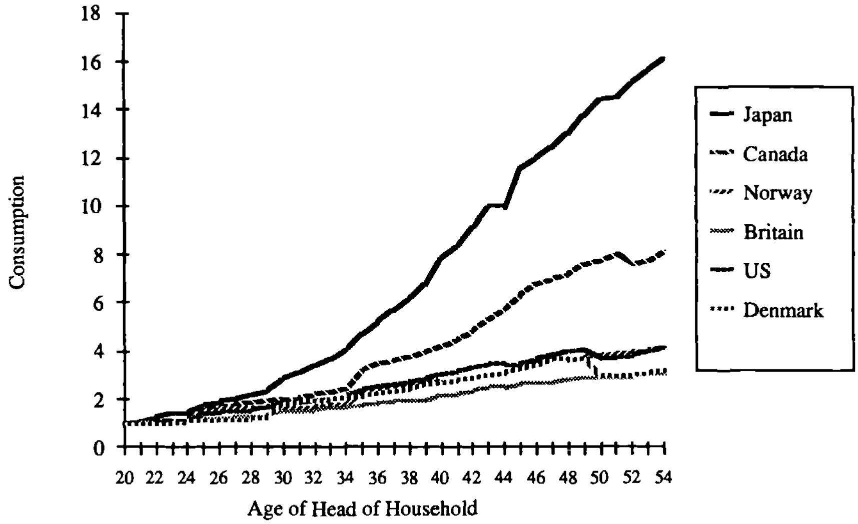

Fig. 10.6 Estimated consumption over the lifetime for individual consumers in different countries

Source: Calculations by the authors using data described in the data appendix.

basis of rules of thumb, consumption and income are strongly associated. A second possibility is that some common cause of both rapid income growth and rapid consumption growth operates across countries.

In an effort to distinguish these possibilities, this section uses information on income growth and consumption growth for individuals in different occupations and with different educational backgrounds. Liquidity constraints, myopia, or the like would be expected to create an association between age-consumption and age-income profiles across different occupations. On the other hand, theories of growth that might apply at the international level would not imply that individual age-income and age-consumption profiles should move together.

Anecdotal evidence about sports stars and medical students suggests that consumption is closely tied to current income, but for a more formal test we turned to the Bureau of Labor Statistics Consumer Expenditure Surveys (CES) of 1960–61 and 1972–73. These studies, originally done for the purpose of calculating consumer price indices, contain detailed expenditure and income accounts for a large representative sample of households (13,000 in 1960–61, 20,000 in 1972–73) and so are an ideal source for comparing income and consumption of households at different ages. For our income measure we took the total after-tax income of the household. Results were similar using several definitions of consumption and expenditure, ranging from total expenditures of the household (including payments for social security and prearranged pension plans) to just consumption of nondurable goods. The consumption measure chosen does not include payments for social security, private pensions, or home mortgages, but does include gifts and contributions to private charities and to other households, as well as insurance premia.

Figures 10.7*a,b* and 10.8*a,b* present mean income and consumption profiles for the nine occupational groups and the five educational levels that could usefully be distinguished with the CES.[7] The data's suggestion that saving for almost all groups increased between the first and second survey is almost certainly a consequence of changes in measurement procedure. What is more interesting is the figures' apparent refutation of the simple life-cycle/permanent income view that the shape of the path of income should not have an effect on the shape of the consumption path. In life-cycle terms, these graphs indicate that people in occupational or educational groups with income peaks late in life do not borrow significantly against those future earnings in order to finance higher consumption when they are young. Conversely, people with income peaks relatively early in life do not appear to save much in anticipation of lower future income. These observations appear inconsistent with life-cycle theory.

It is possible to imagine some combinations of circumstances that can explain some of the apparent correlation above while remaining roughly within a life-cycle framework. For instance, suppose that each cohort in a category consumes its permanent income and that the differences in income across categories and age groups are the result of idiosyncratic shocks to cohorts. Then we would observe the pattern that the income and consumption of households of any given age within a category would be closely related, as we see in the figures. This explanation works, however, essentially by denying any element of predictability in income profiles. But at least across educational categories there is a very strong resemblance of the age-income profiles in the 1972–73 CES to those in the 1960–61 CES—surely a strong refutation of the "no predictability" hypothesis. And, informally, we surely believe that people with college and postgraduate educations can expect higher wage growth over their lifetimes than those with only grade school educations, so that there is some degree of predictability. Although the degree of similarity of 1960–61 and 1972–73 income profiles is smaller across occupations than across educational categories, it is still the case that several occupations, particularly professionals, managers, operatives, and unskilled workers, have quite similar, and thus presumably predictable, profiles in the different years.

The calculations here do not take account of changes in family composition. By calculating consumption on an equivalence scale basis it is possible to create consumption profiles that do not follow estimated income profiles. But it is not clear what this proves, since total consumption spending does follow income. More relevant is the observation that there do not appear to be large differences in average family sizes at different ages among different edu-

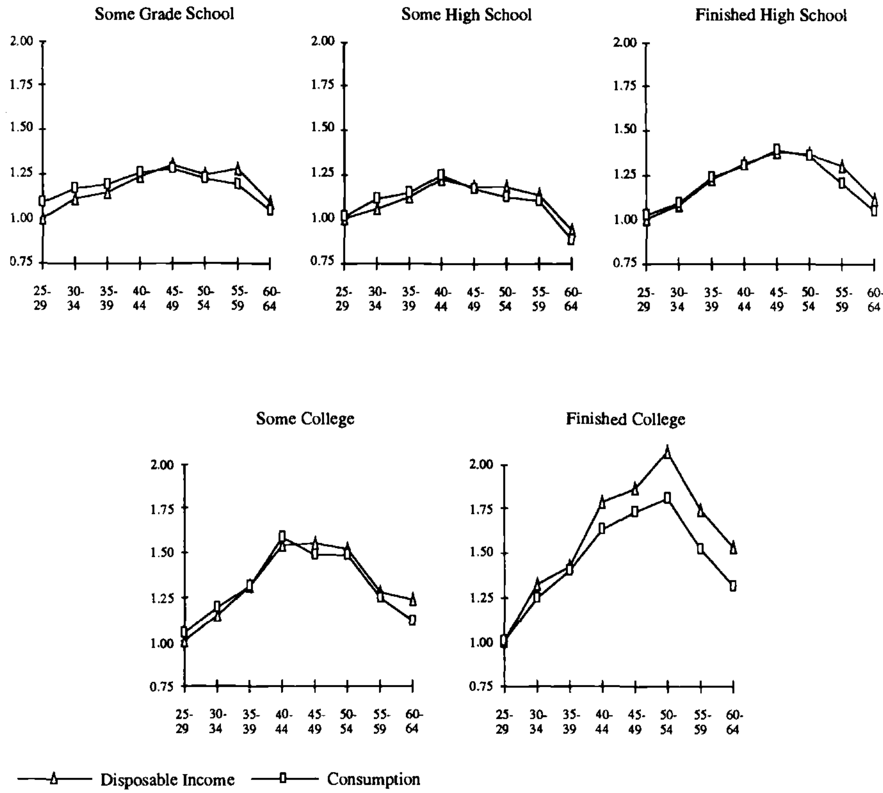

Fig. 10.7*a* Income and consumption profiles by educational group, 1960–61 CES

Source: Calculations by authors using CES tapes.

cational and occupational groups. While the issue deserves further research, our tentative conclusion is that parallel movements in income and consumption cannot be explained by family size considerations.

Another explanation of the consumption/income parallel was provided by Ghez (1975). Using the 1960 CES, Ghez prepared a figure for all consumers similar to our figures 10.7 and 10.8 for subcategories of consumers and sought to explain the observed close correlation between income and consumption using a "family production function" model of the type advocated by Becker (1965). Suppose, for example, that utility is a function both of consumption c and hours of leisure h. Suppose further that, because of the accumulation of experience or other human capital, hourly wages grow over the life cycle. Then individuals will have an incentive to work the longest hours when they are most productive, late in life. But this extra work takes away leisure time, giving the consumer an incentive to consume more time-substituting goods.

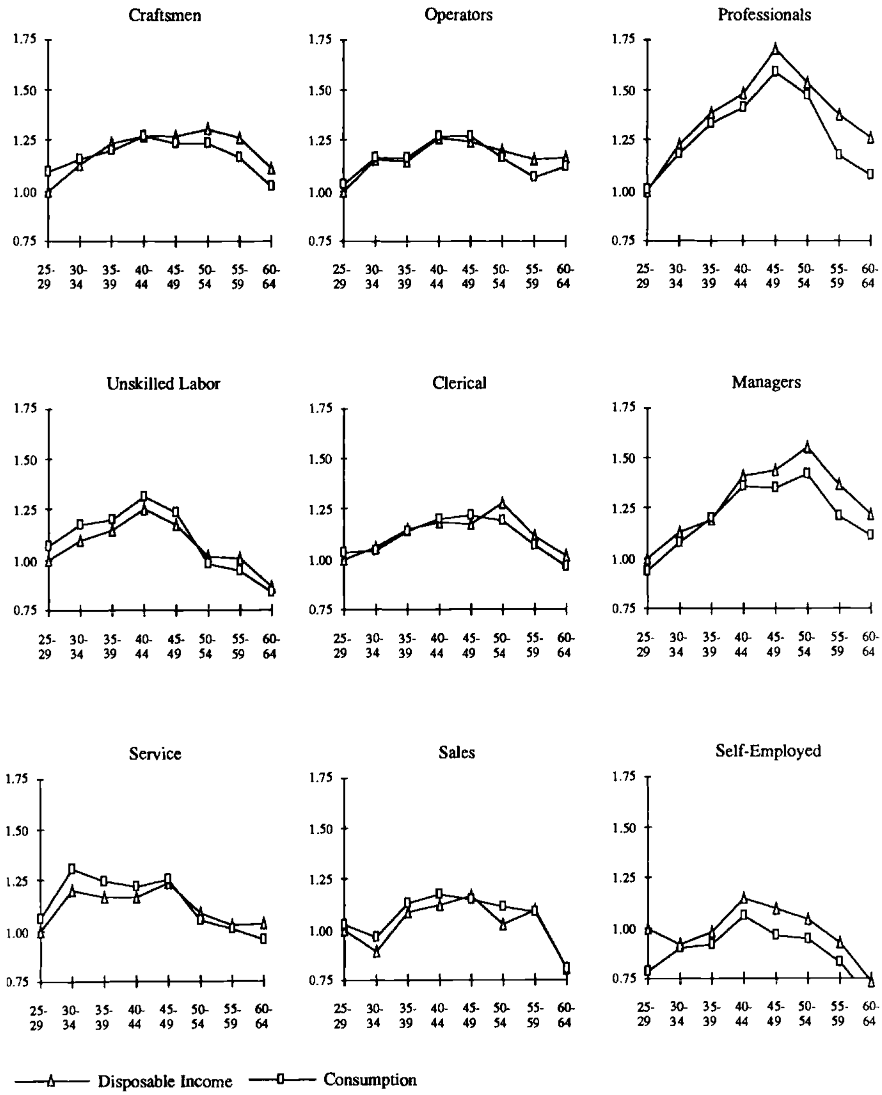

Fig. 10.7*b* Income and consumption profiles by occupational group, 1960–61 CES

Source: Calculations by authors using CES tapes.

The consumer will therefore be observed consuming more during those periods of life when he works most and earns the most income. To be more specific, this model would suggest that busy executives late in life would be more likely to have a maid to do housekeeping chores and more likely to send out their laundry than young people with (presumably) more time on their hands.

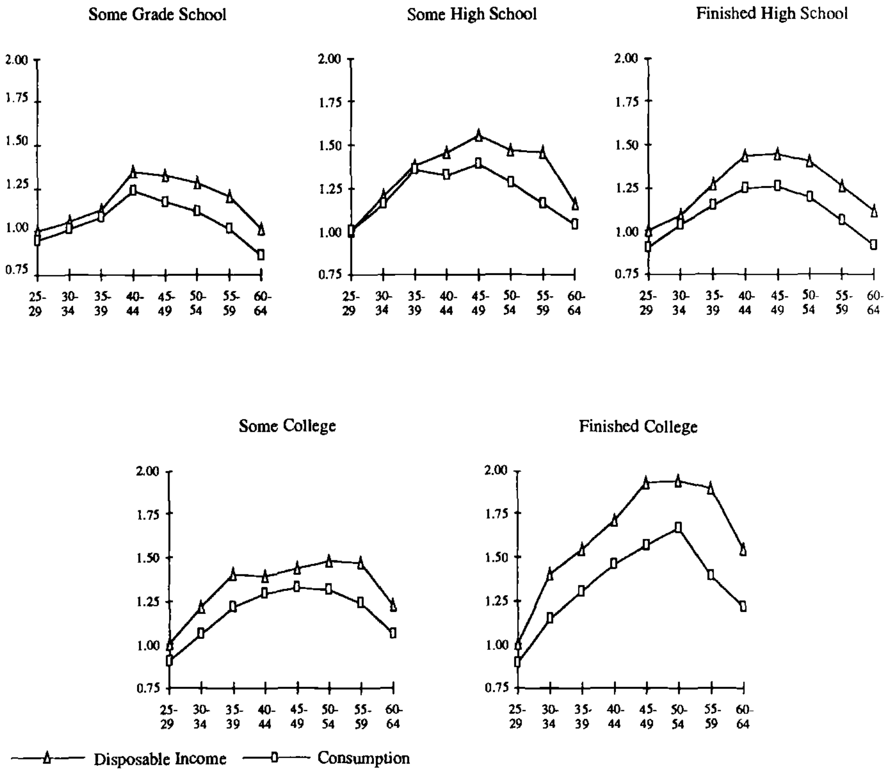

Fig. 10.8*a* Income and consumption profiles by educational group, 1972–73 CES

Source: Calculations by authors using CES tapes.

The Ghez model seems unlikely to be a satisfactory explanation for the close consumption/income parallel observed in figures 10.7 and 10.8 for several reasons. First, it is not even obvious that consumption and hours are substitutes rather than complements. With more leisure time one can engage in expensive activities, such as foreign travel, that may not be possible at all in busier periods of life. Ghez himself makes the point that if time is very valuable one may eat more fast food (presumably inexpensive) and fewer elaborate meals out (presumably expensive). Further, even if we accept that consumption and hours are substitutes, the Ghez model only makes predictions about the sign of the relationship between income and consumption, not about its size. There is no reason in his model to expect that the relationship between income and consumption will be one-for-one as we observe. Finally, the Ghez explanation relies heavily on the assumption that hours and income move exactly in parallel. Figure 10.9, which is reproduced from a book by Ghez and Becker (1975), plots hours worked and hourly earnings at each age across the

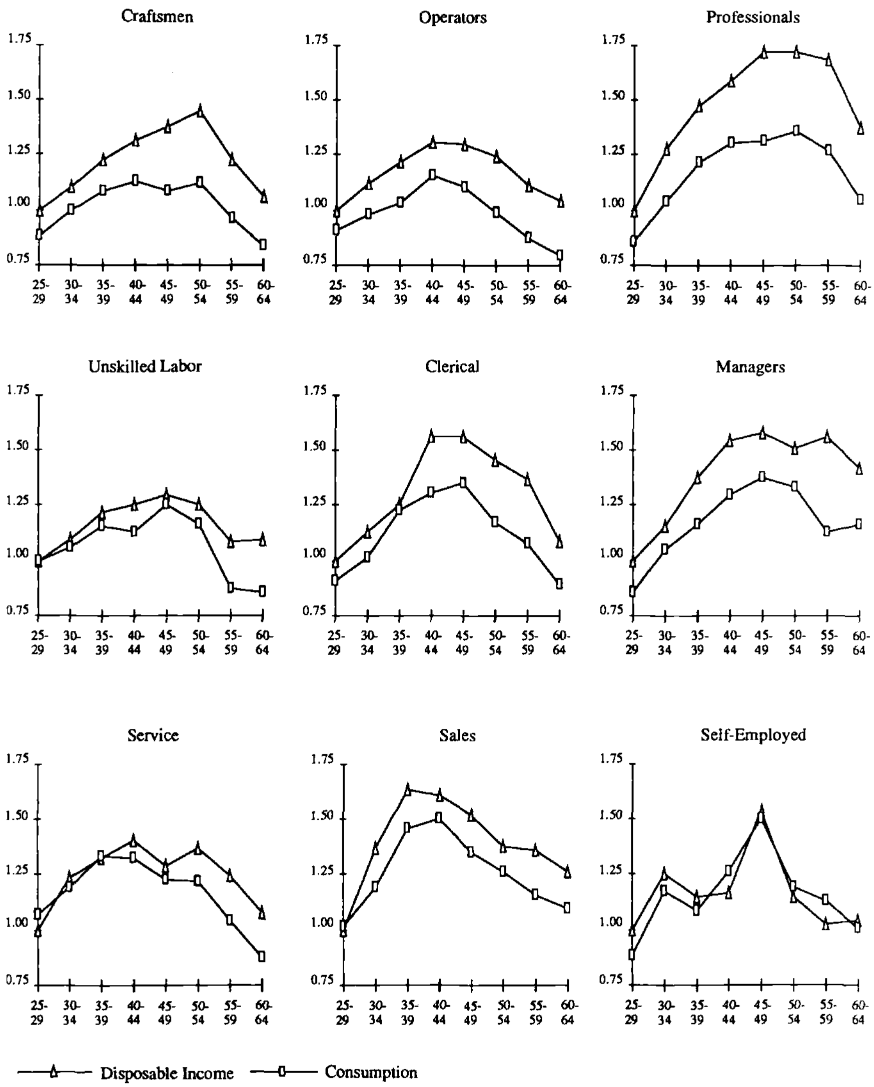

Fig. 10.8*b* Income and consumption profiles by occupational group, 1972–73 CES

Source: Calculations by authors using CES tapes.

life cycle for two educational groups using 1960 census data. It is apparent that there is very little variability in hours worked over the lifetime in either group. Furthermore, hours seem to decline after roughly age 35, while income and consumption both peak in the CES data roughly at age 50. Finally, there is no clear difference across the two educational groups in the age profile

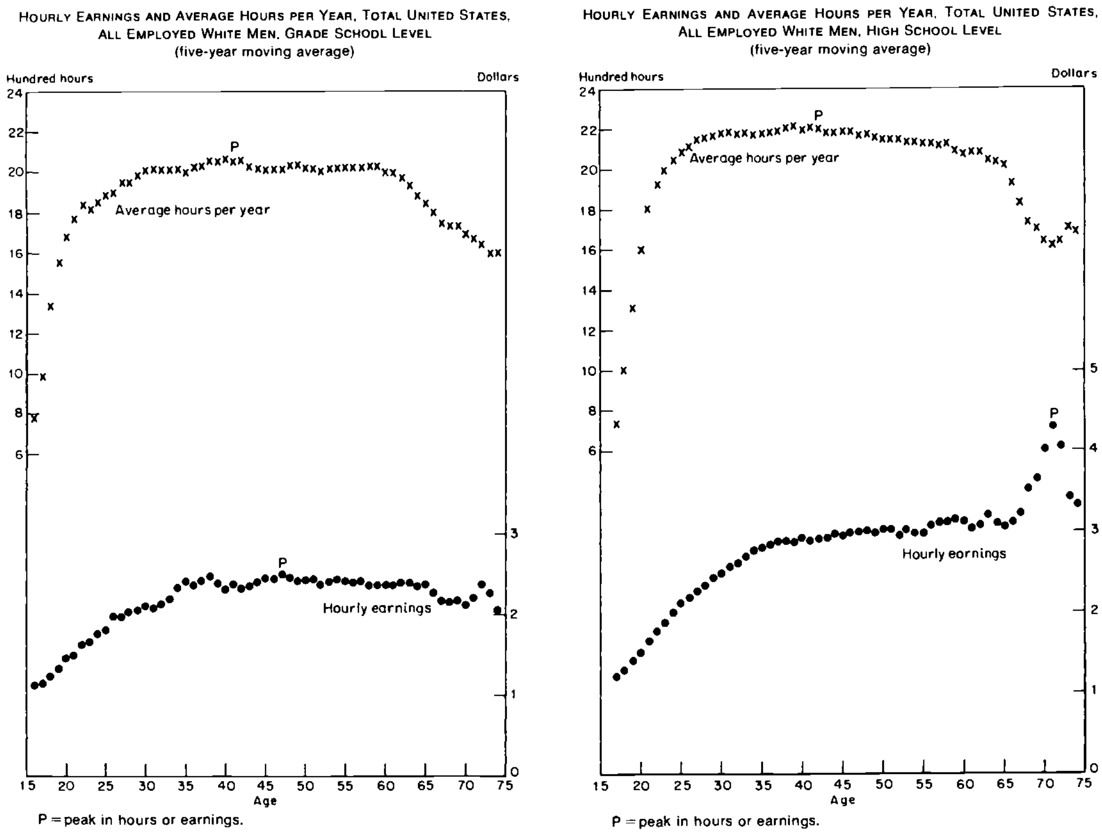

Fig. 10.9 Average earnings and hours by age
Source: Ghez and Becker (1975, 86–87).

of hours worked in spite of a noticeable difference in the profile of wages. We conclude that consumption/hours substitution is not a viable explanation for the consumption/income parallel.

10.4.1 Conclusion

This evidence on individuals suggests that explaining why consumers should allow their consumption to be heavily influenced by current income is a more plausible way to explain the international correlations with which we began the paper than is seeking an endogenous growth theory that could explain a high correlation between consumption growth and income growth. The behavior of these profiles suggests that the excess association of income and consumption is stronger at the low frequencies considered here than it is in the higher frequency contexts that have been more extensively studied.

10.5 Saving and Expected Income Growth

The analysis so far has suggested that both internationally and across individuals consumption and income growth are much more closely associated than standard theories would predict. A different way of stating the same point is to observe that saving decisions appear to be less responsive to expected long-term growth rates of income than simple theories would predict. In this section we examine the response of saving to differences in expected income growth using several different types of data.

The worldwide productivity slowdown after 1973 provides one natural test of the proposition that a decline in growth should lead to reduced human wealth and increased saving. As figure 10.10 demonstrates, the life-cycle hypothesis predicts that a two percentage point decline in expected income growth should have dramatic effects on saving, particularly for young consumers. Young consumers targeting even a 3% annual consumption growth rate are predicted to raise their saving ratio out of income by 20%. For the population as a whole the saving rate should increase by about 10% since the human wealth effect is less important for older consumers.

As figure 10.11*a–d* demonstrates, these predictions are not borne out. Saving rates around the world did not rise following the productivity slowdown. If anything they have fallen. Moreover, there is no tendency for the countries that have suffered the greatest declines in growth to have had greater increases in saving.

This failure of the theory might be due to other shocks that have changed saving behavior. A further test using information derived from the productivity slowdown focuses on its effects on consumers in different age groups. A decline in growth reduces expected future income by much more for young consumers than for older ones and not at all for those who have retired. Whatever happened to overall saving, one would expect to observe a tendency for the *relative* saving rate of the young to rise following the productivity slow-

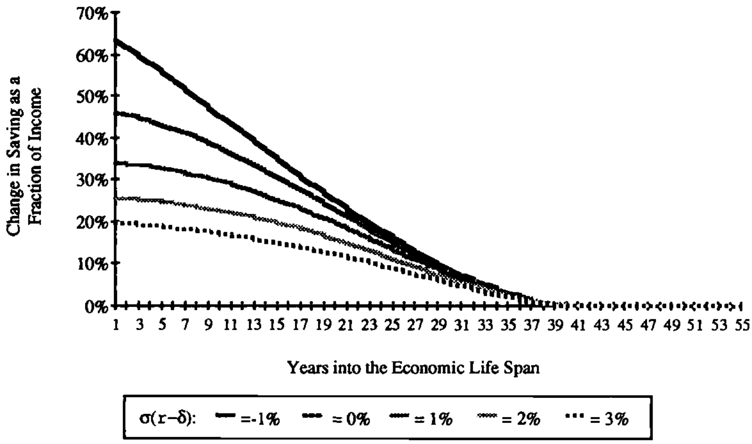

Fig. 10.10 Change in saving as a fraction of income if the expected growth rate of income changes from 3% to 1% per capita per year

Source: Calculations by the authors described in the text.

Note: The assumed growth rate of consumption $\sigma(r-\delta)$, ranges from -1% to 3% annually. Consumers are assumed to retire 40 years into their 55-year economic life span.

down if consumers were farsighted. This tendency should have been reinforced by declining fertility. It is borne out only to a very slight extent in figure 10.12. (Again, because of changes in measurement procedures, nothing can be inferred from the position of these profiles, only their shape.) This finding is perhaps not so surprising given that the shape of the age-saving rate profiles in figure 10.12 is not really consistent with the predictions of the life-cycle theory in the first place.

Information on the shape of occupational income profiles can also be used to test the life-cycle theory. It predicts a tendency for those in occupations where income can be expected to rise rapidly to save less than those in occupations where income can be expected to rise slowly. The profiles from figures 10.7 and 10.8 can be used to calculate a ratio of future income to current income for young people in different occupational groups, and the results can then be compared with observed saving rates.

Figure 10.13 plots, for each occupation in 1960, the ratio (future income/current income) against the saving rate of young people in that occupation, where "future income" is defined as the sum of income for people age 30–65, "current income" is the sum of income for people age 25–29, and "young" refers to people in the age group 25–29. The slope of these lines should be strongly negative because high-future-income occupations should be low-saving occupations. Instead, the slope seems to be positive. This evidence is also consistent with the view that consumption is excessively sensitive to current income, though this cannot explain the positive association in the data.

Overall information on saving supports the conclusion reached in earlier

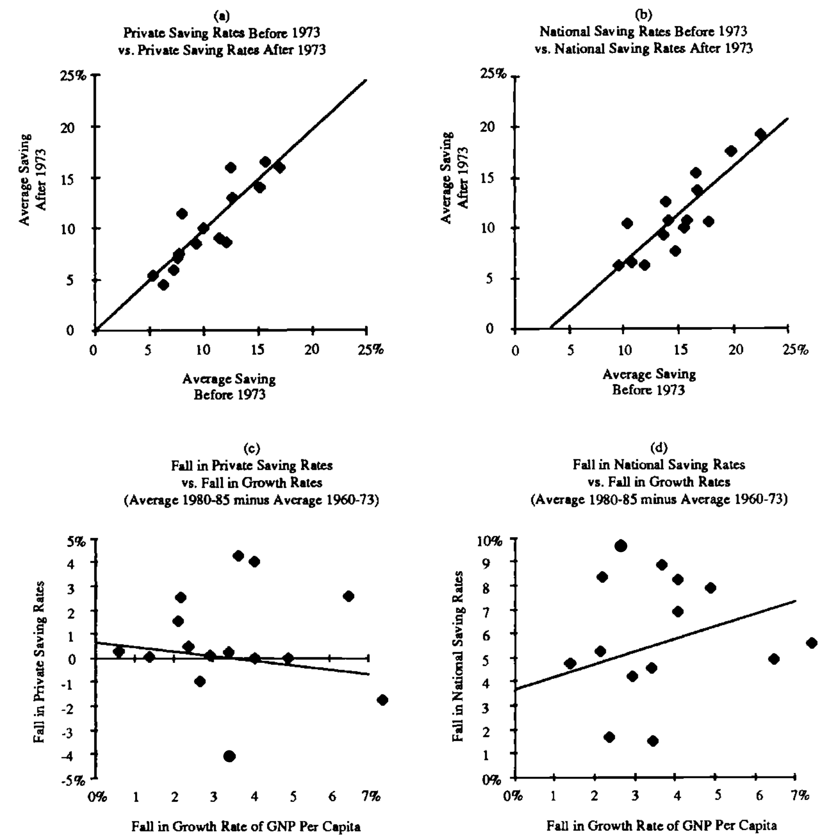

Fig. 10.11 Changes in saving rates since 1973
Source: OECD National Income Accounts.

sections that consumption is much more closely tied to current income than strong forms of the life-cycle or permanent income hypotheses would predict. While reassuring, this evidence is of course not independent of the earlier evidence on the behavior of measured consumption.

10.6 Liquidity Constraints, Myopia, and Uncertainty

One obvious interpretation of the close link between consumption growth and income growth is that consumers are liquidity constrained or myopic. This would "explain" why consumption and income growth are so closely associated. The principal difficulty with this line of thought is that in order to account for the observed equality of consumption and income growth rates one would have to assume that essentially all consumers were liquidity constrained or myopic.

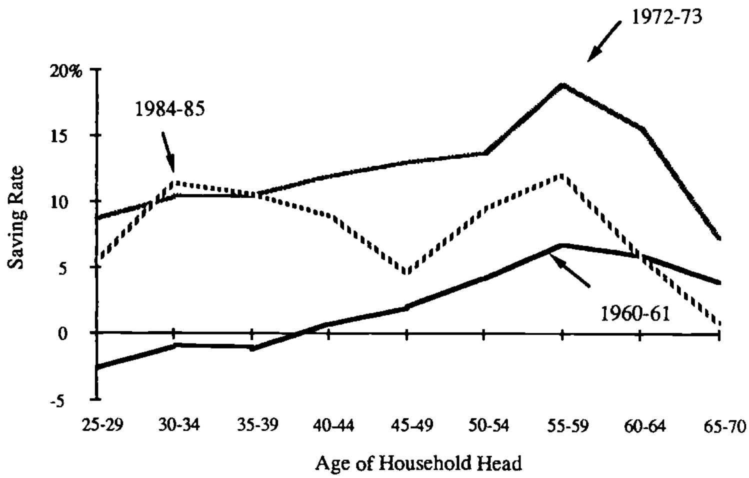

Fig. 10.12 Saving rates by age in 1960–61, 1972–73, and 1984–85
Source: Calculations by the authors using the 1960–61, 1972–73, and 1984–85 CES tapes.

To see this, consider the formulation of Hall (1978) in which the population is divided into two classes. A fraction α of the population consumes all its income and no more each year because of liquidity constraints and/or myopia. The remaining fraction $(1 - \alpha)$ behaves according to the first-order condition in (1). Assuming that the optimizing non-liquidity-constrained latter group enjoys consumption growth at the same rate in every country at the rate $(\dot{c}^*/c^*)$, the growth rate of consumption will be given by:

$$\dot{c}/c = \alpha(\dot{y}/y) + (1-\alpha)(\dot{c}^*/c^*) \tag{3}$$

In order to account for the unit slope observed in figure 10.1, it is necessary to postulate that $\alpha = 1$ so that the entire population is liquidity constrained. This assumption robs the permanent income theory of any content. In addition, it leaves unanswered the unquestion of where savings come from. Of course it is also contradicted by all of the evidence supporting the permanent income hypothesis. The challenge is finding a theory that can account for the apparent absence of pervasive liquidity constraints or myopia in high frequency tests but can still account for our low frequency facts.

However, the possibility that most consumers act as if they were liquidity constrained or expected to be in the future should not be ruled out. Studies such as Campbell and Mankiw (1989), which seek to estimate the fraction of rule-of-thumb or liquidity-constrained consumers by applying time-series techniques, are likely to understate it for three reasons. First, the specification adopted assumes a restrictive form of liquidity-constrained behavior. It would be more difficult to demonstrate conclusively the existence of an economically significant population of myopic consumers if the myopes were assumed

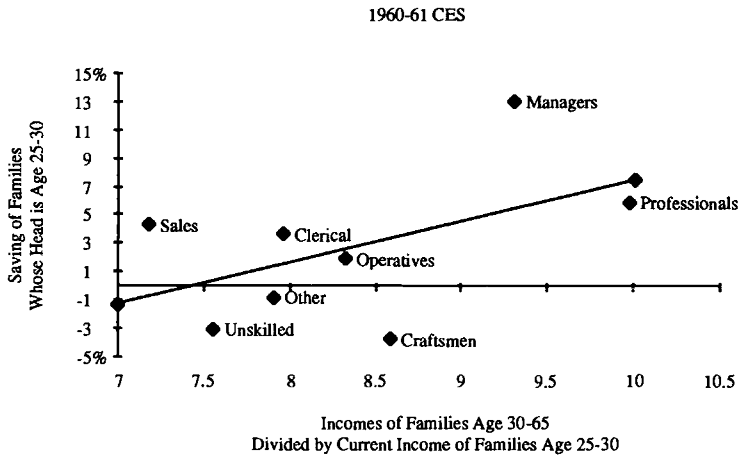

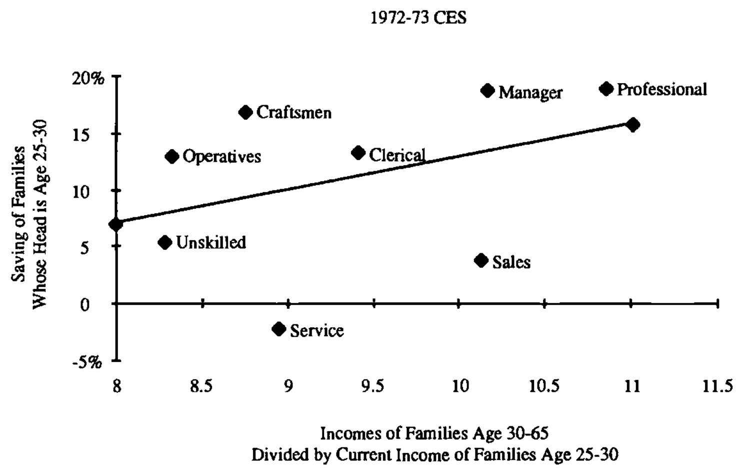

Fig. 10.13 Young families' saving as a fraction of income versus future income streams in their occupations

Source: Calculations described in the text using the income profiles of figures 10.7 and 10.8 calculated from the 1960–61 and 1972–73 CES.

to follow a rule in which consumption responded to income and its lags. Second, the assumption that liquidity constrained consumers spend a fixed fraction of their income on nondurable consumption rules out the possibility that these consumers cut durable spending disproportionately when income declines. If this is in fact the case, standard methods will understate the liquidity-constrained fraction of the population. Third, most recent research efforts have focused on the post-war period where income is close to a random walk. DeLong and Summers (1986) present evidence that in the prewar period, when income fluctuations were more transitory, the fraction of liquidity-constrained consumers was greater.

In spite of the considerable evidence that liquidity constraints are important, the assertion that people spend their incomes is not a rich enough theory of saving. We are attracted by Angus Deaton's (1989) view of savings as a "buffer stock" for contingencies. As he shows, if liquidity-constrained consumers facing risky income are both risk averse and impatient, they will maintain a small "buffer stock" of assets to insulate consumption against transitory income but will not engage in long-horizon borrowing or lending. The buffer stock view has the appeal of predicting (or at least labeling) the consumption smoothing that goes on at high frequency while not implying that consumption smoothing should go on over long horizons. It also has the potential to explain the observed correlation between saving and growth. If consumers desire (as financial planners recommend) a buffer stock equal to a certain number of months' income, saving will be greater for consumers with rapidly growing incomes than for those with slowly growing incomes. Essentially, the accelerator mechanism will create a positive growth-saving relationship.

Table 10.3 presents some empirical evidence, drawn from panel data on tax returns for the period 1979–84, that supports the buffer stock idea. For persons under and over 65, it presents estimates of the fraction of people, fraction

Table 10.3 The Incidence of Liquidity Constraints

Number of Years with < $100 in Interest and Dividend Income	Fraction of the Population Falling in This Category	Fraction of Total Labor Income That Goes to People in This Category	Fraction of Total Income That Goes to People Who Fall in This Category	Fraction of Total Capital Income That Goes to People in This Category
A. Total population:				
0	27.6	37.3	41.4	90.4
1	6.5	8.0	7.7	4.5
2	5.5	6.0	5.7	2.0
3	5.9	7.0	6.5	1.6
4	6.9	7.0	6.6	0.9
5	8.8	8.3	7.6	0.5
6	38.7	26.7	24.5	0.2
B. Population under age 65:				
0	35.2	39.7	47.4	92.8
1	6.2	7.8	7.1	2.8
2	5.2	6.0	5.3	1.7
3	5.3	6.3	5.8	1.0
4	6.1	6.7	5.9	.7
5	7.8	7.9	6.7	.4
6	34.2	25.6	21.8	.3

Source: Calculations by Daniel Feenberg of the National Bureau of Economic Research. See appendix for more detailed discussion of calculations.

of labor income, fraction of total income, and fraction of interest and dividend income going to persons with less than $100 in interest and dividend income in various numbers of years. The results suggest that liquidity constraints are potentially very important. More than half of total income went to persons who usually (three years or more out of six) had less than $100 of interest and dividend income. Furthermore, the fraction of total interest and dividend income received by those who do not always have such income is quite small. This suggests that even in years when such people have over $100 of interest and dividend income they do not have very large amounts of such income. Interestingly, whatever weights are used it appears that about a third of households have minimal interest and dividend income in some but not all years. This is what one expects on the buffer stock view. It suggests that "snapshot" evidence estimating the fraction of the population without assets is likely to underestimate the potential significance of borrowing constraints.

The view that borrowing constraints are important for a large fraction of consumers is also supported by the observation that a large majority of American households report that they have substantial amounts of consumer debt. The interest rate on this debt is typically considerably greater than the rate on safe assets like Treasury bills. Simultaneously borrowing at high rates and holding safe assets is difficult to square with the Ramsey model view of consumption decisions. As Julio Rotemberg and others have argued, it is rational for a consumer who believes he may be liquidity constrained in the future. Such a consumer would also tend to allow his consumption to closely follow his income.

It is also important to recall that typical consumers and typical savers may behave very differently. This point is illustrated by table 10.4. The conceptual unit in this table is the typical dollar of income rather than the typical taxpayer. If the distribution of property income is very unequal we should expect the median or mean dollar of property income to accrue to a person with a very large amount of such income. This is exactly what the table shows. Although the median dollar amount of interest and dividend income was $185, the median dollar of such income went to someone with property income of $16,100. Furthermore, although the mean amount of interest and dividend income was $2,755, the mean dollar went to a taxpayer earning $46,533 of property income. (See the appendix for details.)

The numbers become even more striking when we use assumed rates of return to convert statements about capital income into statements about liquid assets (see the appendix for details). When we do this we discover that the median dollar of (estimated) assets is held by a person holding $274,893 and that the mean dollar is held by a person with three quarters of a million dollars of assets. The general picture of extreme inequality in the distribution of wealth painted by these numbers is borne out by an analysis of some evidence from the Federal Reserve's Survey of Consumer Finances (SCF) in a recent paper by Avery and Kennickell (1988). The SCF allows a direct calculation

Table 10.4 Sources of Dividend and Interest Income

	Whole Population		Population Excluding Elderly	
	Mean	Median	Mean	Median
A. Interest and dividend income weighted by:				
Adjusted gross income	9,344	544	7,878	364
Taxpayers	2,755	185	1,600	113
Interest & dividend income	46,533	16,100	62,515	12,657
Estimated assets	43,840	14,930	58,401	11,457
B. Adjusted gross income weighted by:				
Adjusted gross income	62,910	38,537	63,279	38,773
Taxpayers	30,069	24,693	30,481	25,468
Interest & dividend income	101,983	45,728	150,050	56,695
Estimated assets	99,797	43,883	148,073	53,676
C. Wage income weighted by:				
Adjusted gross income	42,940	32,923	45,327	35,248
Taxpayers	25,212	20,995	27,616	23,439
Interest & dividend income	28,198	6,051	45,110	25,960
Estimated assets	27,701	6,361	44,750	26,920
D. Estimated assets weighted by:				
Adjusted gross income	162,342	9,966	137,393	6,735
Taxpayers	48,914	3,588	28,282	2,183
Interest & dividend income	778,317	287,375	1,032,177	224,299
Estimated assets	753,831	274,893	995,144	212,415

Source: Calculations by Daniel Feenberg of the NBER from IRS tax panel data.
Note: See appendix for detailed description of calculations. All figures are in 1988 dollars.

of net saving via a comparison of families' net worth in 1983 and 1986. In their table 12, the authors estimate the fraction of aggregate positive saving between 1983 and 1986 that was done by the members of each 1983 wealth decile. They estimate that nearly 70% of all the positive saving between 1983 and 1986 was done by families in the top 1983 wealth decile. Using crude smoothing techniques (see appendix), we calculated that the median dollar of saving was done by a family roughly at the 94th percentile in the wealth distribution. Smoothing again, we estimated that a person at the 94th percentile in the 1983 wealth distribution had \$661,000 (1988 dollars) of net wealth. This compares with an estimated median 1983 net wealth of \$46,800 (1988 dollars).[8] Again it would appear that wealth and saving are extremely unequally distributed.

Taken together, this evidence along with tables 10.3 and 10.4 suggest that there are two kinds of consumers. The great majority of consumers are liquidity constrained and have only small amounts of liquid assets, which they keep as a buffer against uncertainty. A small minority of consumers, however, have very substantial assets and are not liquidity constrained. These wealthy consumers are the source of most of the net dollars saved in the economy.

10.6.1 Conclusion

The broad picture painted above suggests that focusing separately on two different models, one for the liquidity-constrained majority of consumers who save little outside of housing equity and one for the small but wealthy minority who seem to do most of the saving, will yield more empirical success than continuing to work with a single model postulating identical unconstrained consumers. These are not new ideas: in arguing for a typically short horizon, Milton Friedman (1957) observed, "The appropriate definition of the permanent component [of income] is a period of three years or slightly longer. This is the same as the conclusion reached earlier from [different] data on urban families. It is also consistent with the time series data. It is encouraging to find such a close agreement in the precise definition of permanent components suggested by three independent bodies of data." And the idea that accumulation is chiefly an activity of the already wealthy goes back at least to Pareto.

10.7 Conclusions

Recent studies of consumption behavior have tested increasingly subtle implications of the life-cycle/permanent income hypothesis using increasingly sophisticated time-series techniques with increasingly ambiguous results. Many existing estimates suggest that at least a large fraction and possibly all of consumption is done by optimizing nonmyopic non-liquidity-constrained consumers maximizing individualistic utility functions with long or infinite horizons. We believe this conclusion is not correct. It seems to us that the wide variety of evidence presented here is much more robust to the possibility of measurement or specification error than the numerous complex econometric tests that have been performed. We regard our evidence as decisively refuting the low frequency predictions of standard intertemporal theories.

As we emphasized in the introduction, the evidence here is generally consistent with the life-cycle and permanent income hypotheses as they were originally advanced. Indeed, Milton Friedman explicitly rejected the idea that consumers had horizons as long as a lifetime in discussing the permanent income hypothesis. And Modigliani relied on myopic expectations in some early development of his theory. What is decisively rejected here is the modern representative consumer versions of these theories, not the core idea that people seek to smooth consumption.

While the evidence here does not undercut the usefulness of the life-cycle and permanent income theories in explaining some broad features of consumption behavior, it does cast serious doubt on modern uses of these theories, which take the idea of a representative forward-looking consumer very seriously. The absence of any relation between rates of return on a variety of assets and consumption growth rates across countries makes us skeptical of the use of consumption information in explaining risk premia on different capital assets. The absence of any clear tendency for consumption to respond to expectations of future income growth leads us to doubt that models that assume consumers optimize over long or infinite horizons will give very good predictions about the effects of various tax changes. And we suspect that those concerned with modeling the determinants of income growth should build in a different consumption function than the one suggested by the Ramsey model. Finally, we note that a major claim of real business-cycle theorists is that their models are calibrated on the basis of noncyclical phenomena. It does not appear that the representative consumer approach used in most real business-cycle models is consistent with low frequency evidence.

We argued in section 10.6 that Deaton's notion of the saving of the typical consumer as a buffer stock to smooth consumption over short horizons and to prepare for temporary sharp declines in income was consistent with both the evidence usually cited in favor of life-cycle permanent income theories and our low frequency evidence. We argued further, however, that although the buffer stock model may describe the typical consumer well, it may not accurately describe the typical saver because saving and wealth are extremely unequally distributed. Further research is needed to determine how the behavior of the typical consumer differs from the behavior of the typical saver.

Even though it may not apply to all consumers, we are attracted to the buffer stock model for several reasons. It provides a natural explanation for the correlation between saving and income growth both across countries and across occupational groups. If consumers desire to hold a cash reserve equal to a certain number of months of income, they will have higher saving rates the more rapid their income growth.

This notion raises a number of interesting possibilities for the growth process. If, as recent studies have argued, steady growth rates are increasing functions of saving rates, and if, as we have just suggested, saving rates are positive functions of growth rates, there is a clear possibility of multiple equilibria. This idea might be relevant to the experience of nations like Taiwan and Korea where actual and expected growth rates have increased sharply and at the same time that saving rates have soared.

The buffer stock model, if correct, also has implications for certain tax policy issues. In the United States there has been considerable controversy about the efficacy of IRAs and other saving vehicles. Critics allege that individuals transfer money from one account to another to realize tax benefits without doing any incremental saving. To the extent that, because of its illiquidity, IRA saving is not a substitute for buffer stock saving, it may be incremental even for households that have liquid assets.

Our future research in this area will proceed in two directions. First we need to refine our knowledge about the behavioral differences between the typical consumer and the typical saver. Second, we will try to develop models that can explain the differences between typical consumers and typical savers and models that are consistent both with the high frequency evidence that some consumption smoothing exists and the low frequency evidence that consumption growth tracks income growth. Although a single unified model may be desirable as an eventual goal, it may turn out to be more fruitful in the meantime to pursue separate models to explain the consumption/income parallel and the consumption/saving divergence. We hope that this multifaceted approach will eventually succeed both in explaining international differences in saving rates and in making predictions about the response of saving to policy changes.

Appendix
Data Sources and Methods

This appendix describes the sources and methods used to prepare the data charts and tables of the paper. We proceed roughly in the order in which the data appear.

OECD Data on Income, Consumption, and Interest Rates

OECD data come from the DRI @OECDNIA, @IMF, and @OECDMEI data bases. Data for most countries for most series begin in 1960. Gross Domestic Product (GDP) is given by the series VAGDPA, personal consumption is given by AGPC, real personal consumption by AGPCR. We derived the CPI deflator and hence inflation rates by dividing AGPC/AGPCR (for some reason the direct data on deflators is less complete than this indirect source). Population figures come from the @IMF database, series 199z. Trade balance data were taken from the @IMF database series 177 ac&d or the nearest existing equivalent. The 15 countries that appear in most of the figures are the United States, the United Kingdom, Austria, Belgium, France, West Germany, Italy, Norway, Switzerland, Canada, Japan, Finland, Greece, Australia, and Sweden.

For short-run interest rates we generally used the rate of return on three-month Treasury-bills, except in Italy where the only series was for six-month Treasury-bills (with a few missing observations that we filled from other interest rate series), and France and Germany where we used call money rates because there was no three-month Treasury-bill data before the early 1980s. The other rate of return data are courtesy David Cutler, who calculated them from the Morgan Stanley *Capital International Perspective*.

International Cross-Section Data on Income and Consumption

Gathering the data for figure 10.4*b* sent us far and wide. For Japan we used the profiles given in Ando and Kennickell, (1986, 194), specifically the data on mean CONSM in the working families. For Canada we used data taken from the Statistics Canada publication *Family Expenditure in Canada* (1989), kindly provided to us by Harry Champion of Statistics Canada prior to publication. For Norway we used unpublished data from government consumer surveys, graciously provided by Knut Morck. For Denmark we used data from the *Statistisk årbog* (Statistical yearbook) *1988* (Danmarks Statistik 1988, 171). Data for Great Britain were taken from Browning, Deaton, and Irish (1985, 503).

To generate figure 10.6 we used the above-described cross-section age-consumption data from all our countries, cohort population data from *Global Estimates and Projections of Population by Sex and Age* (United Nations 1987), and real personal consumption data from the DRI OECD databases mentioned above. We imputed family consumption by age of head of household by assuming that the relative magnitudes of consumption of typical families at different ages did not change over time (see eq. 2 and the description of the calculations in the text).

U.S. Cross-Section Data on Income and Consumption

All the microdata for the U.S. presented in figures 10.5*a*, 10.7, 10.8, and 10.12 were calculated from the Consumer Expenditure Survey (CES) Bureau of Labor Statistics (BLS) for the 1960–61, the 1972–73, and the 1985 and 1986 surveys. These surveys attempt to construct a complete balance sheet for the households surveyed over a one-year period, including information on changes in assets and liabilities that should balance the difference between income and consumption. Fortunately the definitions of variables have not changed much between the surveys so we are able to calculate income and consumption measures that should correspond over time. The 1960 survey, however, differed from the later surveys in at least two respects. First, each household was interviewed only once, at the end of the survey year, and asked to recall income and expenditures for the preceding year. In the later surveys each household was interviewed quarterly for five quarters in a row and asked about consumption over the preceding three months. Second, in the 1960 survey the interviewers made a greater effort to ensure that the family balance sheets actually balanced, so that if income exceeded consumption by $1,000 the interviewer tried to make sure that net assets rose by $1,000. There was less emphasis on such balance in the later surveys.

The figures result from straightforward calculations from the 1960–61, 1972–73, and 1985 CES tapes. In all years our income measure was disposable income after tax, calculated in the earlier surveys by subtracting all taxes from the total income variable; disposable income exists directly in the 1980s tapes so was not calculated. As our measure of consumption we took the variable called "current consumption expenditures" in the 1960 and 1972 surveys and added insurance premia and cash contributions and gifts. To construct the same variable from the 1980s surveys we took the "total expenditures" variable and subtracted contributions to pensions, retirement funds, and social security. The 1972–73 survey presented a particular problem because income numbers below $2,000 or above $35,000 were not reported. By comparing means of our tape sample with means in the BLS's printed summaries of the 1972–73 CES, however, we were able to calculate the average income of the bottom-coded individuals as $973.18 and the average income of the top-coded consumers as $54,942. The disposable income figures were $897.14 and $44,057, respectively. For consumers whose income was top- or bottom-coded, we assumed an income equal to the average income of their group. A final adjustment to the 1972 and 1985 samples was necessary because a small fraction of the people did not provide complete information about income; these were excluded from the sample altogether.

The basic patterns presented here were robust to the few reasonable variations in calculation technique we could think of, which consisted of excluding people from the sample for various plausible reasons and of considering different definitions of consumption and income (e.g., nondurables consumption, pretax income, wage income, etc). Detailed charts for 1985 analogous to those from 1960–61 and 1972–73 were not presented for two reasons. First, the 1985 data seemed to have much higher variability. This is partly due to a smaller sample size (about half as large) and partly (we think) due to a new processing methodology devised by the BLS. Second, CES occupational group classifications in the 1980s series are much less detailed, and occupations within each group seem less similar, than is the case with the 1960–61 and 1972–73 surveys.

Liquidity Constraints Tax Panel Data

The liquidity constraints tax panel is a random sample (based on primary taxpayer's Social Security number) of tax returns. It includes single and joint returns, but women drop from the sample when they marry and return when they divorce or widow. The sample was maintained for 1979 to 1984. Of the total set of tax returns in the data set, there were 5,997 taxpayers with positive adjusted gross income in all six years. This is the sample we used in preparing tables 10.3 and 10.4. The calculations for the tables were performed by Daniel Feenberg of the NBER.

The procedure for estimating liquid assets from capital income was simple. To estimate the market value of the stock portfolio we took dividend income and divided by the dividend/price ratio on the stock market as a whole for the appropriate year. To estimate the dollar value of interest-bearing assets we divided by the average interest rate on interest-bearing assets and cash. The latter was estimated by taking total personal interest earnings from the NIPA and dividing by the sum of cash and interest-bearing assets taken from the *Balance Sheets for the U.S. Economy* (Board of Governors of the Federal Reserve 1989). The latter figure yields interest rates in the 8%–10% range, probably much higher than the actual interest rate on the typical dollar of interest-bearing assets and cash. Overestimating the interest rate should cause us to underestimate associated wealth, however, so whatever error exists here biases our results against finding the extreme inequality in wealth that we do in fact find. A better interest rate measure should only intensify our findings about inequality.

The rates used in these calculations are given in table 10A.1. The dividend/price ratios were taken from *The Dow Jones-Irwin Business and Investment Almanac, 1986* Levine (1986).

A brief word about the interpretation of the numbers in table 10.4 is in order. Consider, for example, the part of the table concerning adjusted gross income (AGI) for everyone excluding the elderly. We claim that the median AGI weighted by AGI is $38,537. What this means is that if we were to sort all taxpayers by AGI and then to find the taxpayer such that the sum of the AGIs of the taxpayers with less AGI than his equals the sum of the AGIs of the taxpayers with more AGI than his, that taxpayer has an AGI of $38,537. This is what we mean when we say that the median dollar of AGI goes to a taxpayer with AGI $38,537. The meaning of the mean dollar of AGI weighted by AGI is less intuitive, but can be understood by analogy with calculation of mean tax rates. Suppose we knew income and total taxes paid by a set of individuals, and we wanted to calculate the average tax rate on all the dollars of income in the group. Simply taking the average of the tax rates across individuals would be inappropriate because the tax rate on individuals with high incomes clearly has more influence on the tax rate on the average dollar of income than the rate on low-income individuals. The appropriate procedure is to take a weighted mean of all the tax rates, where the weights are given by the incomes of the individuals. By analogy, the appropriate procedure to find the "typical" dollar of income in the mean sense is to take a weighted mean of income where the weights are also given by income.

Wealth Calculations from Avery and Kennickell

Avery and Kennickell (1988) present tables drawn from the 1983 and 1986 Federal Reserve *Survey of Consumer Finances,* which is virtually the only reinterview wealth survey containing a large number of high-income families. This survey allows a direct calculation of net saving via a comparison of each family's net worth in 1983 and 1986. In their table 12 the authors estimate the fraction of aggregate positive saving between 1983 and 1986 that was done by the members of each 1983 wealth decile. We used this table to generate a crude approximation to the distribution function for saving by wealth decile.

Table 10A.1 Rates of Return on Equities and on Interest-bearing Assets, 1979–84

Year	Dividend Price Ratio	Average Interest Rate
1979	5.47	7.8
1980	5.26	8.4
1981	5.20	9.4
1982	5.81	9.3
1983	4.40	8.8
1984	4.64	8.9

The technique was as follows. The graph of saving by wealth decile appeared to be close to exponential, so we assumed that the function log(saving) = f(1983 wealth decile) was exactly linear. Using two points, the saving of the first decile and the saving of the last decile, we calculated the slope and the intercept for the line passing through those two points. This technique should substantially underestimate the inequality of the wealth distribution because research (as well as the simple graph of log saving against wealth decile) suggests that wealth is even more unequally distributed in the upper income brackets than the log assumption suggests. Since the results indicate a high degree of inequality in spite of this bias we are confident that our figures do not overstate the degree of inequality.

Given a continuous function for the distribution of saving as a function of wealth, it is a simple matter of numerical integration to find the point at which saving below that point equals saving above the point. This is the point that defines the amount of saving done by what we call in the text the "median" saver. The procedure described above was repeated using Avery and Kennickell's (1988) table 10 to produce a distribution of wealth by wealth decile, and the resulting function was used to calculate the estimated wealth of someone at the 94th percentile in the wealth distribution, the point that the previous function identified as being associated with the median saver.

Notes

1. We comment below on the possibility that differences in tastes can explain our observations.
2. Because of data limitations we do not carefully distinguish durable and nondurable consumption as theory would suggest. Given that durables are a relatively stable share of consumption in the United States at least, we doubt that this has much impact on our results.
3. We use both private and national saving measures in order to avoid taking a stand in the Ricardian equivalence debate.
4. Note that this test differs from the popular Hall-style tests by focusing on low-frequency measures of income growth rates like the geometric average over the previous five years rather than very high frequency variables like previous quarter's income growth. If we believe there is long-term dependence in growth rates, then this is an appropriate variable to use as a proxy for expected current and future growth. We recognize that the previous discussion does not fully address the implications of uncertainty, because the model that produces (1) is a perfect certainty model. We address the implications of a model that incorporates important uncertainty below.
5. See the data appendix for details on data sources and methods.
6. Given the large differences in lifetime income between cohorts it is also surprising under the life-cycle theory that the consumption of 30-year-olds is not much greater than the consumption of 65-year-olds in both countries.
7. The unused occupational groups were retired people, nonresponses, and others. The unused educational group was "none, nonresponse, or other." The figures grouped by occupation are in order of increasing standard deviation of the mean level of in-

come, so more credibility should probably be ascribed to inferences drawn from figures near the top of the page than those near the bottom. The difference in variance across educational groups was substantially less (the groups are closer in size) so the figures grouped by education are ordered by increasing educational level.

8. Note that these wealth estimates include housing equity, which accounts for the discrepancy between the estimated median wealth here and in table 10.4.

References

Ando, A., and Arthur B. Kennickell. 1986. How Much (or Little) Life Cycle Is There in Micro Data? The Cases of the United States and Japan. In *Macroeconomics and Finance: Essays in Honor of Franco Modigliani*. Cambridge, Mass.: MIT Press.

Avery, Robert B., and Arthur B. Kennickell. 1988. Savings and Wealth: Evidence from the 1986 Survey of Consumer Finances. Paper presented at the 50th Anniversary Conference on Research in Income and Wealth of the National Bureau of Economic Research. Washington, D.C., May.

Becker, Gary. 1965. A Theory of the Allocation of Time. *Economic Journal* 73, no. 5 (September): 493–517.

Board of Governors of the Federal Reserve. 1989. *Balance Sheets for the U.S. Economy, 1949–88*. Washington, D.C., April.

Browning, Martin, Angus Deaton, and Margaret Irish. 1985. A Profitable Approach to Labor Supply and Commodity Demands Over the Life Cycle. *Econometrica* 53, no. 3 (May).

Campbell, John Y. and N. G. Mankiw. 1989. Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence" NBER Working Paper no. 2924. Cambridge, Mass., March.

Danmark Statistik. 1988. *Statistisk årbog* (Statistical yearbook). Copenhagen: Danmark Statistik.

Deaton, Angus. 1989. Saving in Developing Countries: Theory and Review. Paper prepared for the First Annual World Bank Conference on Economic Growth. Washington, D.C., April.

DeLong, J. Bradford, and Lawrence H. Summers. 1986. Is Increased Price Flexibility Stabilizing? *American Economic Review* 76, no. 5 (December): 1031–44.

Friedman, Milton. 1957. *A Theory of the Consumption Function*. Princeton, N.J.: Princeton University Press.

Ghez, Gilbert. 1975. Education, the Price of Time, and Life-Cycle Consumption. In *Education, Income, and Human Behavior*, ed. F. Thomas Juster, 295–312. New York: Maple Press for the Carnegie Commission on Higher Education and the National Bureau of Economic Research.

Ghez, Gilbert, and Gary S. Becker. 1975. *The Allocation of Time and Goods Over the Life Cycle*. Studies in Human Behavior and Social Institutions, no. 6. New York: Columbia University.

Hall, Robert E. 1978. Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence. *Journal of Political Economy* 86 (December): 971–87.

Kotlikoff, Laurence J., and Lawrence H. Summers. 1981. The Role of Intergenerational Transfers in Aggregate Capital Accumulation. *Journal of Political Economy* 89: 706–32.

Levine, Sumner N. 1986. *The Dow Jones-Irwin Business and Investment Almanac, 1986*. Homewood, Ill.: Dow Jones-Irwin.

Mehra, Rajnish, and Edward Prescott. 1985. The Equity Premium: A Puzzle. *Journal of Monetary Economics* 15, no. 1 (January): 145–61.
Modigliani, Franco. 1970. The Life Cycle Hypothesis of Saving and Intercountry Differences in the Saving Ratio. In *Induction, Growth, and Trade: Essays in Honour of Sir Roy Harrod.* Oxford: Clarendon Press.
———. 1980. The Life Cycle Hypothesis of Saving Twenty Years Later. In *The Collected Papers of Franco Modigliani,* vol. 2: *The Life Cycle Hypothesis of Saving,* ed. Andrew Abel. Cambridge, Mass.: MIT Press.
———. 1986. Life Cycle, Individual Thrift, and the Wealth of Nations. *American Economic Review* 76, no. 3 (June): 297–313.
Morgan Stanley. 1968 to present. *Capital International Perspective.* New York: Sladkus.
Ramsey, Frank. 1928. A Mathematical Theory of Saving. *Economic Journal* 38 (December): 543–59.
Romer, Paul. 1986. Increasing Returns and Long-Run Growth. *Journal of Political Economy* 94 (October): 1002–37.
Statistics Canada. 1989. *Family Expenditure in Canada.* Statistics Canada catalog no. 62–555. Ottawa.
United Nations, Department of International Economic and Social Affairs. 1987. *Global Estimates and Projections of Population by Sex and Age: The 1984 Assessment.* New York: United Nations.

Comment N. Gregory Mankiw

Christopher Carroll and Lawrence Summers present us with a collage of facts about consumption and income. They give us scatter plots from aggregate cross-country data, tabulations from individual tax return data, and profiles of consumption and income from consumer survey data. Although they do not give us a model to explain all these data, the myriad pieces of evidence they present form a compelling image of how consumers behave.

Most important, the image of the consumer that arises from the paper contrasts sharply with the modern renditions of the permanent income hypothesis that pervade much of macroeconomics. According to these modern theories, consumers are rational, forward looking, and able to borrow and lend to smooth consumption over time. In the Carroll and Summers collage, we see consumers who, because of myopia or liquidity constraints, do not set their consumption on the basis of the present value of expected future income. Instead, current income exerts a larger influence on consumption than many modern theories imply.

I find myself sympathetic to many of the conclusions of this paper. In our joint work on the time-series properties of consumption and income, John Campbell and I also found that current income is a more important determinant of consumer spending than the permanent income hypothesis suggests. I

N. Gregory Mankiw is professor of economics at Harvard University and a research associate of the National Bureau of Economic Research.

think that Carroll and Summers are right that these findings call for more work on liquidity constraints and precautionary saving.

Because I agree with Carroll and Summers on the implications of these findings for the theory of the consumer, I would like to discuss the implications for another topic: the theory of economic growth. Their cross-country evidence provides as serious a challenge for those economists trying to model economic growth as it does for those trying to model consumer spending.

The Central Fact

From the standpoint of growth theory, the central fact in the paper is found in figure 10.1: countries with high growth in income have high growth in consumption. This fact is compelling because it is simple and because it is robust. Much of the paper is aimed at arguing that this fact is a problem for standard theories of the consumer, because most of the natural explanations of it do not hold up under close scrutiny.

Although I am convinced that this fact is correct, I am left somewhat skeptical of the authors' interpretation. They would like us to believe that, because they are looking at averages over long periods, the differences in mean income growth reflect differing steady-state growth rates rather than differing shocks across countries. This is probably not completely true, however. For example, much of the Japanese growth miracle was unanticipated—otherwise it would not be called a miracle.

The question Carroll and Summers do not fully answer is how much of the cross-country variation was unanticipated. If the differences in income growth across countries were mostly due to luck, the Carroll and Summers facts would be far less interesting. For example, the correlation between income and consumption growth would be explained trivially by the revision in permanent income. The distinction between anticipated and unanticipated growth is thus important for how we interpret these facts.

For now, however, I will assume that this fact will stand up to closer empirical scrutiny. Like the authors, I will assume that the observed differences in sample means reflect differences in steady-state growth rates. I want to discuss what their cross-country evidence implies for our theories of economic growth.

A Diamond Model

In order to think about this fact systematically, let us consider a couple of simple growth models. Let me begin with a standard Diamond model with some plausible and convenient functional forms. Suppose the production function is Cobb-Douglas

$$Y = AK^{\alpha}L^{1-\alpha}, \tag{1}$$

where Y is output, K is capital, which lasts for one period, and L is labor. Suppose that people live for two periods, supply labor inelastically when young, and consume in both periods according to the utility function

(2) $$U = \gamma \log C^y + (1-\gamma) \log C^o.$$

Suppose further that the population is constant, and that labor supply is normalized to be

(3) $$L = 1.$$

Not surprisingly, these assumptions make the model easy to solve. Log utility tells us that consumption when young is a constant fraction γ of labor income, and the Cobb-Douglas production function tells us that labor income is a constant fraction $(1-\alpha)$ of total income. In addition, consumption when old equals capital income, which is also a constant fraction α of total income. Thus, most quantity variables in this economy are simply proportional to total income.

Can this sort of model mimic the close connection between consumption growth and income growth that Carroll and Summers document? As stated, the model is not even a growth model: it reaches a steady state with fixed level of capital and income. The standard way to get steady-state growth into such a model is to assume that the technological parameter A grows exogenously over time:

(4) $$A = a(1+g)^t.$$

Countries will have different steady-state growth rates if the parameter g varies across countries.

This model can now explain the observed relation between consumption growth and income growth. Summing the consumption of the young and old shows that the steady-state level of aggregate consumption is

(5) $$C = [\gamma(1-\alpha) + \alpha]\, Y.$$

Because aggregate consumption is proportional to aggregate income, high growth in aggregate income leads to high growth in aggregate consumption. In addition, if we look at growth in consumption over an individual's life, we obtain

(6) $$C^o/C^y = [\alpha/((1-\alpha)\gamma)]\, (Y_{+1}/Y).$$

Individual consumption also grows more quickly if aggregate income is growing quickly. Hence, the growth in aggregate consumption and the growth in individual consumption in this model appear to be in line with the Carroll and Summers findings.

This model begins to have problems when we turn to examining rates of return. The steady-state real interest rate is

(7) $$1 + r = [\alpha/((1-\alpha)(1-\gamma))]\, (Y_{+1}/Y).$$

High growth should lead to a high real interest rate. Thus, the absence of any correlation between real interest rates and growth, which Carroll and Summers document, appears to be evidence against this traditional growth model.

Finally, I should note that this growth model does not run into problems because it adopts the life-cycle theory of the consumer. If instead I had supposed that young consumers in this economy obeyed an arbitrary Keynesian consumption function according to which consumption was proportional to income, the model would be little changed—it would merely turn into a textbook Solow growth model. Again, the model would predict, counterfactually, a correlation between growth rates and real interest rates.

A Romer Model

Let us now consider a second growth model—identical to the first except for the determination of technological change. In particular, I want to replace the assumption of exogenous technological change with an assumption of endogenous technological change along the lines pioneered by Paul Romer. Suppose that total factor productivity is given by

$$A = aK^{\beta}. \tag{8}$$

The state of technology evolves not as a function of time, but rather as a function of the level of capital.

I will not go into detail about why technology evolves in this way. The key underlying assumption is that there are externalities to capital accumulation. One possible story is that when a firm builds a factory, it thinks up good ideas that become part of the general pool of knowledge. Alternatively, there may be network externalities or external benefits to specialization and product differentiation that, because of scale economies, are only possible as the economy grows larger.

The crucial implication of these externalities is that the private and social production functions now diverge. The economy faces the aggregate production function

$$Y = aK^{\alpha+\beta}L^{1-\alpha}. \tag{9}$$

Individual firms, however, ignore the external effects and view themselves as facing the Cobb-Douglas production function (1). It is therefore the Cobb-Douglas production function that governs the distribution of income between capital and labor.

To turn this model into one of endogenous steady-state growth, let us take Romer's suggestion and assume that $\alpha + \beta = 1$, so that the aggregate production function exhibits constant returns to scale in capital. Under these assumptions, it is straightforward to show that the steady-state growth rate is

$$Y_{+1}/Y = (1-\gamma)(1-\alpha)a. \tag{10}$$

In contrast to traditional growth models, the steady-state growth rate in this model depends on preferences. If we view all countries as obeying this model and differing by their rate of time preference γ, we obtain different equilibrium growth rates. Impatient countries such as the United States have high γ

and thus low growth rates; patient countries such as Japan have low γ and thus high growth rates.

The appealing feature of this model is that it mimics some of the facts documented by Carroll and Summers. To see this, note first that the relation between consumption and income remains the same:

$$C = [\gamma(1-\alpha) + \alpha]\,Y. \tag{5}$$

This implies that countries with high income growth also have high consumption growth (and also a high rate of saving). Inferring the growth in individual consumption from equations (6) and (10), we obtain

$$C^o/C^y = [\alpha a(1-\gamma)/\gamma]. \tag{11}$$

Because high growth countries have low γ, they also have high growth in individual consumption.

In addition to producing the positive correlation between income growth and consumption growth, the model also mimics the observed patterns in real interest rates. The real interest rate in this economy, which is determined by the private marginal product of capital, is easily shown to be:

$$1 + r = \alpha a. \tag{12}$$

Note that the real interest rate is independent of γ and thus will not vary systematically with the growth rate. In this model, the externalities associated with capital accumulation imply that higher saving causes higher growth, and they also prevent higher saving from lowering capital's rate of return.

The Real Interest Rate Puzzle

From the standpoint of the theory of economic growth, the puzzling fact in the Carroll and Summers paper is not the high correlation between consumption growth and income growth. This correlation will arise in almost any growth model. The puzzle is the absence of any correlation between growth and real interest rates. I have shown that one can explain this fact by appealing to an endogenous growth model that assumes constant returns to capital. Yet many economists (including myself) will find this assumption unappealing.

The findings in this paper therefore call for two directions of future research. First, as Carroll and Summers emphasize, economists need better models of consumer spending. Second, as I have emphasized, economists need better models of growth—in particular, models to explain why real interest rates fail to vary across countries with growth rates.