DEATON

CHAPTER 2

Consumption, Growth, and Interest

That there should be a positive relationship between productivity growth and the ratio of savings to income in the aggregate economy is an early and justly celebrated prediction of the life-cycle hypothesis. Although various theories of economic growth also predict that saving (or investment) and growth should be positively related, the life-cycle model is unique among models of consumer behavior in predicting a causal relationship that runs from faster growth to greater household saving. The result is not implied by the permanent income theory of consumption nor by the earlier and less sophisticated models that postulated a simple relationship between income and outlay. The growth-to-saving prediction offers a possible explanation for the fact that virtually all developed economies have experienced simultaneous falls in the rate of productivity growth and in national saving rates during the 1970s and 1980s. In Section 2.1, I discuss the basis in life-cycle theory for the growth–saving relationship, and look at some of the recent international evidence. While there seems little doubt that growth and saving are indeed linked, it cannot be established that life-cycle saving is the cause.

The effect of interest rates on saving, discussed in Section 2.2, has always been a central issue in political economy. Aggregate saving and capital accumulation are society's provision for the future, so that, for many people, failures or distortions in saving behavior are seen as compromising the welfare of future generations. There is also a commonly held view that identifies saving with growth, regarding both as measures of a country's economic performance. If interest rates have an effect on saving, then there is a direct link between policy, particularly monetary and fiscal policy, and economic performance. There are many versions of the story: taxation of capital income lowers real interest rates and stifles the incentive to save; taxation of capital income distorts saving and generates deadweight losses; financial 'repression' in developing countries lowers returns, depresses saving, and retards growth. All such arguments depend on the existence of a positive response of saving to higher interest rates. The theory of Chapter 1, while relevant to these questions, hardly suggests a definitive answer, so that much hinges on the empirical evidence. However, the interpretation of the data is not straightforward, particularly once we recognize the importance of the aggregation from microeconomic theory to macroeconomic data.

2.1 Saving and growth

That saving will be generated as a result of productivity growth is an insight that comes from even the simplest model of life-cycle saving and consumption, so that such a model is a good place from which to begin. Armed with the basic ideas, I discuss various elaborations of the model, and their likely effects on its predictions. I then turn to the empirical evidence, and to its implications, both for the relationship between saving and growth, and for the life-cycle model itself.

The stripped-down life-cycle model

Begin by considering the simplest version of the life-cycle model without uncertainty, where the only change in income is when the consumer retires from work, and where consumption is constant over life. This can be formally justified using the theory of Chapter 1, or we can simply assert, along with Modigliani (1986), 'the self-evident proposition that the representative consumer will choose to consume at a reasonably stable rate, close to his anticipated average life consumption,' a proposition that is a good deal more general than intertemporal additivity of preferences. Figure 2.1 illustrates this case, which Modigliani refers to as the 'stripped-down' version of the life-cycle model. Labor income is constant throughout the L (= 40) years of working life, at one unit per period, and then falls to zero through the R (= 10) years of retirement. The real interest rate is zero so that consumption is constant at $L/(L+R)$ per period, or 80% of income through the working life. Assets accumulate at $R/(R+L)$ (= 20%) of income per period, reaching a maximum of $RL/(R+L)$ (= 8) times income immediately before retirement. Throughout life, the average ratio of assets to labor income is $0.5R(L/(R+L))$ (= 4), the first of many such numbers that conform to reality (at least roughly) under even the simplest assumptions.

The stripped-down model predicts that both demographic and productivity growth will generate saving, and that without either there will be no net saving in the economy as a whole. Saving is done by young people, and dissaving by the old. If the population is stationary, and if the incomes of the young are the same as were the incomes of the old, saving and dissaving are equal and opposite. With productivity growth, the younger are richer than were their parents at the same age, their saving is on a larger scale than was that of their parents, and net saving is positive. The faster

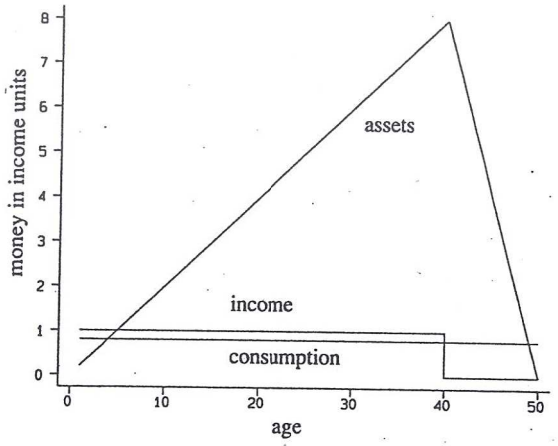

Figure 2.1: Consumption, income, and assets in the stripped-down life-cycle model

the growth, the higher the saving rate. Population growth works exactly the same way; if there are more young people than old people, their total saving outweighs the total dissaving by members of their parents' generation, so that once again there is positive saving in society as a whole.

If the rates of population and productivity growth are constant, and if the real interest rate is zero, the stripped-down model generates simple formulas for the relationship between saving rates and growth. Suppose that population is growing at rate n, and productivity at rate g, so that national income is growing at rate $n+g$. There are $n_0 e^{nt}$ consumers born at date t, and each has a *constant* income of $y_0 e^{gt}$ throughout life, so that productivity growth takes place between cohorts, but not within the lifetime of each individual. Just as in Figure 2.1, each person consumes a fraction $L/(R+L)$ of income per year. Hence at date t, total income is the income of all those born since $t-L$, while consumption is the consumption of those born since $t-R-L$. The aggregates are therefore given by

$$C = \int_{t-L-R}^{t} \frac{L}{L+R} y_0 e^{g\tau} n_0 e^{n\tau} d\tau, \quad Y = \int_{t-L}^{t} y_0 e^{g\tau} n_0 e^{n\tau} d\tau. \qquad (2.1)$$

If we compute the integrals, the consumption to income ratio is given by

$$\frac{C}{Y} = \frac{L}{L+R} \frac{1 - \exp\left(-(g+n)(L+R)\right)}{1 - \exp\left(-(g+n)L\right)}. \qquad (2.2)$$

Given the assumptions, equation (2) can be used to investigate the behavior of the saving ratio as the growth rates change. Note first that the two growth rates, of population and productivity growth, do not appear separately in (2), but only as the sum $n+g$, so that it is only *aggregate* growth that matters for the saving ratio, not whether it is population growth at constant per capita income, or income growth in a stationary population. For given work and retirement spans, (2) implies that saving is zero if growth is zero and that saving is a concave increasing function of the growth rate of aggregate income. The formula generates realistic saving rates at realistic growth rates; for $L = 40$ and $R = 10$, as in Figure 2.1, the saving rate rises from zero at zero growth, to 4.5% at 1% growth, 8.2% at 2%, 11.1% at 3%, and 15.1% at 5%. The slope of the saving-growth relationship is $R/2$ at the origin, and for the (40,10) combination for (R, L) is an easily remembered figure of 2, so that if growth increases from, say, 3.5 to 4.5%, the saving ratio will increase by 2 percentage points.

Complications to the basic model

These back-of-the-envelope calculations suggest that even the stripped-down model generates the sorts of results and predictions that are well worth checking out against the data. However, we must first examine how many of the qualitative results are basic features of the life-cycle model, and how many are artefacts of the very special and clearly unrealistic assumptions of the stripped-down model, particularly the assumptions that income is constant until retirement, and that interest rates are zero.

The introduction of a positive interest rate does a great deal to complicate the algebra, because we now have to keep track of capital income as well as labor income, but the main features are not seriously affected. Positive real interest rates will tip consumption paths downwards in the early years and upwards in old age, as agents adjust their intertemporal paths to the intertemporal incentives. But this only implies that the young will save relatively more, and it is the saving of the young that is the fulcrum upon which the growth effects operate.

More serious consequences follow from the recognition that labor incomes are not constant throughout the working life. Even in occupations where there is very little training, and productivity depends on brawn rather than brain, incomes typically start out at low levels, increasing with age before eventually declining. The longer the period of training, and the

greater the return to human capital, the later is the peak, but there is nearly always some period of income growth at the beginning of the life cycle. Hence, if consumption is constant over life, it is possible that young consumers may want to *borrow*, not save, in the early years of their careers, especially if they are in occupations where the educational and training period is long. If so, then at rapid enough growth rates, additional growth will *decrease* saving, as higher growth rates magnify early borrowing relative to later repayment.

Of course, positive interest rates, the precautionary motive, restrictions on borrowing, or the effects of habits may act so as to restrain early consumption, so that young people with hump-shaped income profiles may not wish to borrow, or be able to do so. However, it must also be recognized that productivity growth may generate income growth *within* individual life cycles, and not just across them. If so, and if consumers anticipate real growth, as surely they must, there is again an incentive to borrow against that growth, and the borrowing will be larger the larger is the growth rate. The general point is that growth will increase aggregate saving if life-cycle saving occurs at earlier ages than life-cycle dissaving. The arguments above suggest that such a result is far from automatic, but it is nevertheless plausible, and will be the case even if there is only modest saving in late middle age, followed by dissaving after retirement.

The dependence of the saving ratio on total growth, and not on how it is divided between population and productivity growth, is another result that does not survive more realistic modelling. The stripped-down model recognizes old age, but not childhood. Workers spring from the womb, tools in hand, and immediately begin accumulating wealth for their retirement. If instead, they are born as dependent children into the households of those who are in the early years of their own working life, there is a further reason to expect consumption to be high and saving low in the first years of the working life. The presence of children, by placing an additional burden on young workers, may precipitate borrowing in the early years of the life cycle, and again reverse the postulated effect of productivity growth on saving. Faster population growth, if long enough maintained, increases the ratio of workers to the retired, but it also increases the ratio of children to workers, so that the net effect on saving is not necessarily positive, nor is there any longer a simple link between the effects of population and productivity growth. The stripped-down model is suggestive, but many of its predictions depend on its special structure.

The empirical evidence

The qualifications to the stripped-down model should only warn us that, as a matter of theory, there is no simple relation between national saving rates and growth; they should certainly not discourage us from looking at the evidence, nor from trying to interpret it in terms of the life-cycle story. Indeed, empirical studies have repeatedly shown that there is indeed a positive relationship across countries between saving rates and growth rates of national income. The relationship is clearest among the more developed countries, but the positive correlation also exists in a somewhat weaker form among less developed economies. Figure 2.2 shows a typical scatter diagram using data from (version 5 of) the Penn World Tables, see Summers and Heston (1991). There are 120 countries in the scatter, which plots the average value over the years 1981 through 1985 of the national 'saving ratio,' defined as 100 minus the percentage share of consumption in gross domestic product versus the average rate of growth of real GDP from 1965 to 1980. The slope of the regression line through these points is 1.34 with a standard error of 0.33, close enough to the prediction of the stripped-down model. Similar scatters can be generated using other data sets such as the annual data in the *World Development Report*, World Bank (annual)—see for example Deaton (1990). Of course, there are other possible explanations for these results; the share of investment in GDP is positively correlated across countries with the share of saving in GDP—see Feldstein and Horioka (1980)—and virtually all growth models predict that growth should respond to the share of investment. But this does not detract from the fact that the prediction of the life-cycle model is supported both qualitatively and quantitatively.

The clearest and most up-to-date survey of the cross-country evidence is by Modigliani (1990), who looks at OECD data from twenty-one developed countries from 1961 through 1987 (Canada, the US, Japan, Australia, Austria, Belgium, Denmark, Finland, France, Germany, Greece, Iceland, Ireland, Italy, The Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and the UK), and, in a separate exercise, at data from eighty-five developing countries from 1982 to 1988, provided from (a revised version) of Aghevli *et al.* (1990). The OECD data show a marked reduction in both saving and growth rates from the 1960s through to the 1980s; Table 2.1 shows the averages over the twenty-one countries by the three 'decades' 1961–70, 1971–80, 1981–7. As Modigliani emphasizes, the

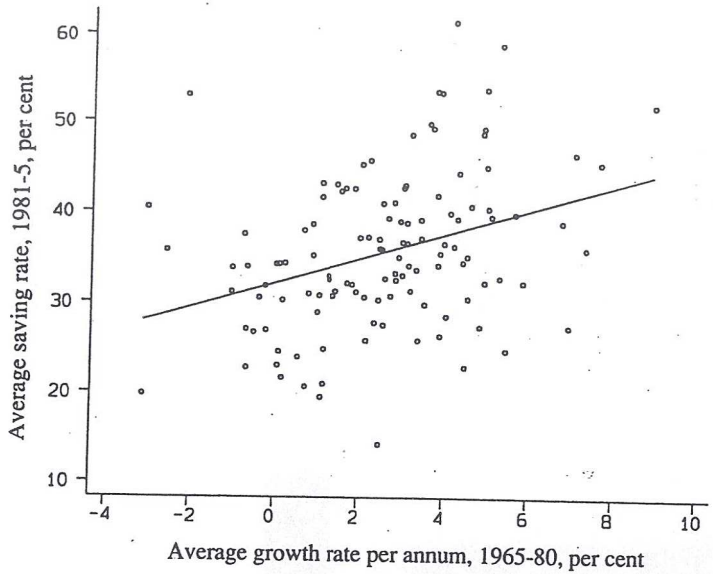

Figure 2.2: Cross-country saving and income growth: 120 countries from the Penn World Tables, Summers and Heston (1991)

change is widespread over these countries. Only in Portugal, where it was constant, was there no decline in the saving rate between the 1960s and 1970s and only in Switzerland and Norway was there an increase from the 1970s to the 1980s. Moreover, the 6.3 point drop in the saving ratio for a 2.5 point drop in growth rates is close to the 2 for 1 effect that comes out of the simple stripped-down model. Modigliani treats the decades as sample points and estimates by ordinary least squares a pooled regression over countries and periods:

$$\frac{s}{y} = \underset{(4.2)}{0.06} + \underset{(6.1)}{1.81} g_{-1} \qquad \bar{R}^2 = 0.37,\ s.e. = 0.041, \tag{2.3}$$

where the figures in brackets are *t*-values. The growth coefficient is well determined and close to the value predicted from the simple theory. (Modigliani presents regressions containing a number of other variables, but the coefficient on growth is robust and remains within the range suggested by the model.)

Table 2.1: Saving and growth in twenty-one developed countries, 1961–1987. (averages in percentages)

	Growth	National saving rate
1961–70	4.9	16.6
1971–80	3.4	15.3
1981–7	2.4	10.3

Source: Modigliani (1990)

For the eighty-five less developed countries, a comparable regression is

$$\frac{s}{y} = \underset{(5.5)}{0.068} + \underset{(1.3)}{1.30}\, g + 0.17\, active + \ldots \qquad \bar{R}^2 = 0.59,\ s.e. = 0.0549 \tag{2.4}$$

where *active* is the proportion of the population aged 15–64, and there are other regressors (results not shown) for the reciprocal of the level of national income, the terms of trade (both significant and both negative), and a dummy for whether the country has a debt-servicing problem (significant and positive.) The demographic variable is not well determined, perhaps not surprisingly given the theoretical arguments, and this sort of finding is typical of the literature. Although there are some studies and some equations that find an influence of population growth or demographic effects, the results are typically not robust and there is no consensus on the direction of the effect on saving—see Gersovitz (1988) for a survey. However, the growth variable enters as predicted, and is again both well determined and in the range suggested by the theory. For the poorer countries, increases in the *level* of national income also appear to enhance saving, perhaps because saving for old age is unnecessary in the early stages of economic development.

These results would seem to provide a striking endorsement of the theoretical predictions of the simplest, stripped-down life-cycle model. Whatever the selection of countries, growth and saving are positively linked, and the much discussed recent reduction of saving rates in the developed countries can be attributed, along with much else, to the slow-down in productivity growth that began in the early 1970s.

Contradictory evidence?

In some ways, it is surprising that the evidence is quite so favorable. The predictions of the stripped-down life-cycle model owe as much to the simplifying assumptions as to the more basic supposition that consumers make sensible life-cycle plans. Hence the very success of the predictions suggests that other factors might be at work. Even in Modigliani's results, there is some hint that all is not well. Regressions of differences in saving ratios from one decade to the next show a weaker and less significant effect of changes in growth rates than did the earlier level on level regressions. Undoubtedly some of the change is attributable to the smaller sample size in these differenced regressions, although one might also suspect the influence of country-specific fixed effects that are removed by the differencing. If fixed effects are important, the cross-sectional results will attribute to growth what are in reality long-established differences between countries that are correlated with growth in the cross-section. As a result the cross-sectional results will not be consistent with the response of saving to the slow-down in productivity. In fact, this turns out not to be a problem. The inclusion of country dummies into (3) reduces the growth coefficient to approximately 1.5, with a *t*-value of 4.1, which, although showing some effect in the direction predicted, hardly suggests that country fixed effects are a major source of error.

Other evidence comes from following through the life-cycle explanation of saving and growth, and checking, not just the final result, that national saving rates should respond to national growth rates, but also the intermediate implications for the cross-sectional behavior of consumption and saving. If saving occurs at earlier stages in the life cycle than does dissaving, then increased growth, by magnifying the scale of activities of relatively younger consumers, will generate additional saving. So the life-cycle explanation for the cross-country relationship between saving and growth can only be correct if there is evidence for at least some 'hump' saving in the cross-section. Of course we do not have evidence for all the countries that appear in the cross-section regressions. But for the countries where there are data, and we shall see several examples below, it is typically the case that consumption and income are much more closely associated than is the case in stylized diagrams like Figure 2.1. The typical hump-shaped profile of income is closely matched by a corresponding hump-shaped profile of consumption, so that the smoothing of consumption over the life

cycle, sometimes referred to as long-term or 'low-frequency' smoothing, if it takes place at all, takes place only on a limited scale. Life-cycle saving, when it occurs, takes place in middle or late middle age, not long prior to retirement. Whether there is saving among the young, or dissaving among the old, is something that varies from data set to data set, and has been the subject of a good deal of controversy.

These findings have implications, not only for the relationship between saving and growth, but also for the question of how much of total wealth can be accounted for by life-cycle saving. In particular, Kotlikoff and Summers (1981) have argued that the cross-sectional evidence for the US which follows the general pattern described above, does not generate enough life-cycle saving to justify the common belief since Tobin (1967) that this form of saving can at least approximately account for total wealth holdings in the US, a belief that again owes much to the stripped-down model of Figure 2.1. These topics have been the source of much discussion—see in particular Modigliani (1988) and Kotlikoff (1988)—but it would probably be fair to conclude that the observation of apparently widespread 'tracking' of income by consumption has led to a downward revaluation of the likely fraction of wealth that can be attributed to life-cycle saving, both in the US and elsewhere. As for saving and growth, if life-cycle saving means saving in late middle age followed by limited dissaving in retirement, then increased growth can be expected to generate increased saving. But the evidence for dissaving in retirement is at best mixed, and if, in addition, there is borrowing by young consumers, the life-cycle effects of growth may be to magnify borrowing, not saving.

Growth and the age-profile of consumption

What is to me the most persuasive evidence against the life-cycle interpretation of the cross-country relationship between growth and saving comes from cross-country comparisons of cross-sections of consumption and income. The idea comes from Carroll and Summers (1991). Consider two otherwise identical economies, one of which has had no income growth for a very long time, while the other has been growing steadily for an equally long time. In each we collect household survey data and compute consumption age-profiles, graphs of consumption against the age of the household head. In the no-growth economy, young consumers have on average the same lifetime resources as their parents and grandparents, while in the

expanding economy, children are richer over their lifetimes than were their parents, and much richer than were their grandparents. In the absence of growth, the life-cycle consumption age-profile can take any shape at all, depending on needs and tastes over the life cycle; the argument does not require that there is a preference for constant consumption over life. However, in the growing economy, the ratio of the lifetime resources of the young to that of the old must be greater than in the stagnant economy, so that, since consumption is determined by lifetime resources, the consumption age-profile must be relatively tipped towards the young in the more rapidly growing economy. According to life-cycle theory, consumption depends on lifetime resources, not on current resources, and in rapidly growing economies the life-time resources of the young are larger relative to those of their parents and grandparents than is the case in a more slowly growing economy.

Carroll and Summers calculate age consumption profiles for the US in 1960, 1973, and 1985, for Japan in 1974 and 1979, and for various years for Canada, the UK, Denmark, and Norway. In spite of differences in growth experiences, the profiles are quite similar from one country to another. For the US and Japan, growth rates of real per capita GDP from 1960 to 1985 have been 2.1% per annum and 5.2% per annum respectively, so that if these rates were maintained indefinitely, a 25-year-old Japanese would be 12.5 times richer than his 75-year-old grandfather, whereas a 25-year-old American would be only 2.8 times richer than his grandfather. In spite of this, the Japanese consumption age-profile, although quite similar to the American, actually peaks slightly later, in direct contradiction to the theoretical prediction.

Similarly dramatic results can be obtained from examination of the consumption age-profiles in less developed countries. I illustrate using household survey data from Thailand, a rapidly growing economy, and Ivory Coast, which has experienced very little growth in the last quarter century. According to the Summers-Heston data, 1960 real national income per capita was 8% greater in Ivory Coast than in Thailand; by 1985, the Thai per capita income was twice that in Ivory Coast. The average annual growth rates over the period were 4.1% for Thailand and 0.9% for Ivory Coast, so that the corresponding ratios for lifetime resources of 25 to 50-year-olds are 7.11 times for Thailand and 1.64 times for Ivory Coast. Figures 2.3 and 2.4 show consumption age-profiles together with income-age profiles for the two countries. The Thai profiles are taken from the

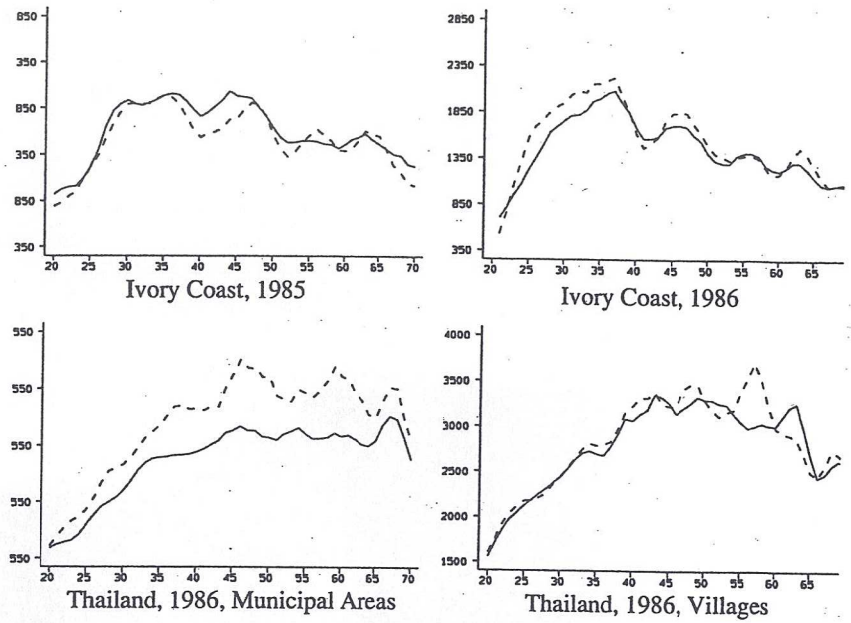

Figure 2.3: Age-profiles of income (broken lines) and consumption (solid lines), Ivory Coast 1985 and 1986, and Thailand 1986, urban and rural.

1986 Socioeconomic Survey, and are shown separately for municipal areas (urban) and villages (rural), and are averaged over 3,589 urban and 5,012 rural households. For Ivory Coast, the data are for the two years 1985 and 1986, and come from the Living Standards Surveys carried out jointly by the World Bank and the Government of Ivory Coast. The sample sizes are smaller than for Thailand, 1,600 households per year, so I have not attempted to show the rural and urban results separately. The graphs have been smoothed over age-groups, so that each point shows the average of consumption and income for the age and the two (Thailand) or three (Ivory Coast) ages on either side, with decreasing triangular weights. Hence, for example, average consumption at 30 years of age in Thailand is one-third of the mean for 30-year-olds, plus two-ninths of each of the means for 29- and 31-year-olds, plus one-ninth of the means for 28- and 32-year-olds. In principle, the Thai samples are large enough for one year averages to be sufficient, but in these data, as in many others from LDCs, people tend to

round their ages to numbers ending in 5s and 0s, and the people who do so tend to have lower consumption and income than those who report non-rounded ages. Smoothing over five-year bands removes the irregularities that would otherwise be generated by such anomalies.

The Thai consumption age-profile peaks much later in the life cycle than does that for the Ivory Coast; Ivorian households reach their peak consumption levels when the head is aged around 35 years, while in Thailand, consumption continues to grow until at least age 45, and perhaps later. Once again, as in Carroll and Summers' comparison between the US and Japan, this is exactly the wrong way round. Thailand is growing much more rapidly than is Ivory Coast, and should have a consumption profile that is relatively tipped toward the young, whereas, in reality, it is the Ivorian profiles that peak at the younger ages. Of course, these comparisons of total household expenditure make no allowance for the different demographic compositions of households across the two countries. However, although household size is larger in Ivory Coast, maximum household size is attained at around the same age in the two countries, so that differences in the relationship between age and household size cannot by themselves explain the differences in the consumption age-profiles. Of course, total household size is only a crude measure of needs, but it seems unlikely that more sophisticated measures can account for the differences in the profiles.

One common reaction to these figures is to argue that it is unreasonable to suppose that tastes are the same in the two countries. It is certainly possible that the life-cycle model is true for each country separately, each with its different preferences. But it is not possible to make such arguments while maintaining that the life-cycle model is the explanation for the cross-country correlation between saving and growth. Without the assumption that the structure of intertemporal tastes is the same in different countries, the life-cycle story delivers no predictions about the effects of saving on growth in an international cross-section.

Does consumption track income too closely?

Where then are we left? There are two separate issues. The first is the validity or otherwise of the life-cycle explanation for the link between saving and growth, and the second is the validity of the life-cycle hypothesis itself. The cross-country evidence on consumption age-profiles makes

it difficult to believe that the correlation between growth and saving is a consequence of hump-saving by youngish consumers saving for their old age. There is too little hump-saving and too little dissaving to make the story plausible. More importantly, consumption profiles for economies growing at different rates are not consistent with the basic hypothesis that consumption at all ages is determined by *lifetime* resources. But this does not by itself mean that the life-cycle hypothesis must be rejected. Inter-country taste-variation can be admitted, and other explanations sought for the correlation between saving and growth. There is no lack of contenders. There is a strong correlation between national saving rates and national investment rates, so that we may simply be observing the relationship between investment and growth, either lagged as in (3), or current, as in (4). While the standard neoclassical Solow growth model does not predict a relationship between saving rates and growth in the long run, the dynamics of adjustment may be sufficiently slow for the cross-country correlations to be dominated by the transitional behavior, although see the counter-arguments in King and Rebelo (1989). Models such as that of Rebelo (1991) postulate that there are constant returns to a (broadly defined) concept of capital, and are therefore consistent with a positive relationship between the investment share and growth even in the long run. Alternatively, accounts of growth under increasing returns, for example Romer (1990), emphasize the role of preferences that embody a willingness to wait, and that allow countries to accumulate the human capital that is the decisive catalyst for growth. Countries with low rates of time-preference will have high saving and high growth, and vice versa, although the causality is neither directly from saving to growth, nor from growth to savings.

As far as the life-cycle hypothesis itself is concerned, there is certainly a sharp contrast between the life-cycle consumption and income profiles of the stripped-down model in Figure 2.1 and the actual paths for the US, or for Thailand and Ivory Coast as illustrated in Figures 2.3 and 2.4. I have already cited the findings of Kotlikoff and Summers (1981), who found little hump-saving in the US. Figures 2.3 and 2.4 show little evidence of saving at any age in either Ivory Coast or in rural Thailand, and while urban Thais do save over the life cycle, the heaviest savers are households headed by older heads. Deaton and Paxson (1992) look at households in Taiwan, one of the fastest growing and highest saving economies in the world. Although Taiwanese households save a great deal, there is once again very little that could be described as hump-saving. Households save

at all ages, and if anything, the saving rate tends to *increase* with age. Ando and Kennickell (1987), after examining six sets of survey data from the US, conclude that 'most families save a relatively small portion of their income throughout the period of their active participation in the labor force, and after they retire, they dissave very little, keeping their assets more or less at the same level.' Other authors, Danziger *et al.* (1983) and Diamond and Hausman (1984) have not found the relationship between wealth and age that is predicted by hump-saving.

Perhaps most notably, Carroll and Summers (1991) use data from the US Consumer Expenditure Surveys of 1960–1 and 1972–3 to draw age-profiles of consumption and income for five educational groups and nine occupational groups. The patterns differ markedly from one group to another, although they are much more stable across time, and in each case the consumption profile is close to the income profile. Those who are in educational or occupational groups where income peaks late, in their late forties or early fifties, have consumption paths that also peak late, while those with little education and unskilled jobs have income and consumption profiles that peak early, that are flat, or that fall with age. Browning, Deaton, and Irish (1985), using British data, also find that consumption and income profiles are synchronized over the life cycle, both for manual and nonmanual workers. In the Michigan Panel Study of Income Dynamics, Lawrance (1991) finds that food consumption of poorer households rises less rapidly with age than does food consumption of richer households, a result that she interprets as showing that the rate of time-preference is inversely related to income, see equation (1.10). But poorer, less well-educated households also have less rapidly rising incomes, so that once again the evidence is consistent with consumption tracking income over the life cycle. These findings suggest a much closer association between consumption and income than might be thought to be compatible with the life-cycle hypothesis. While it is clear that a simple relationship between consumption and income is not a good alternative, if only because consumption is much smoother than income, these results suggest that it is worth considering hypotheses that link consumption and income over shorter periods than complete life cycles; I shall turn to this task in Chapter 6.

Even so, it should be admitted that with enough ingenuity, a good deal of this evidence can be made compatible with the life-cycle model, or indeed with the theory of Chapter 1. One route is to look for factors that condition preferences and that are correlated with labor incomes over the life cycle. Children are one possibility, and family size and the expenditures that go with it tend to peak in middle or late middle age, at around the same time that family income peaks. Somewhat more heretically, preferences may be directly affected by income, if income and lifestyles come with particular types of jobs. Yet again, hours worked tend to follow labor income over the cycle, so that if consumption is a substitute for leisure, it would make sense to use consumption to compensate for the limited leisure opportunities of middle age. Such an explanation would hold whether or not the consumer chooses how many hours to work (although in either case the theory of Chapter 1 would have to be extended.) However, hours worked and wage rates also have their own life cycle and business cycle patterns, and it turns out that neither substitutability nor complementarity of leisure for goods can account for the co-movement of hours, wages, and consumption—see Browning, Deaton, and Irish (1985) for Britain, and Ando and Kennickell (1987) who report the same finding for the US. However, the precautionary and habit-formation models discussed in Chapter 1 would go some way to reconciling the theory with the evidence, since both tend to depress consumption early in life, when income is also low. Interest rates may play a similar role, and if real returns are higher in more rapidly growing economies, there may be a partial explanation for why consumption in rapidly growing economies does not more heavily favor the young. Some of these hypotheses have problems of their own, and none has the simple appeal of the direct link between consumption and income that is so characteristic of the data. However, a great deal more research remains to be done before we have the decisive evidence that would sharply discriminate these hypotheses one from the other, or indeed clearly invalidate the life-cycle hypothesis.

2.2 Saving and interest rates

The last decade has seen a renewal of interest in the empirical analysis of the relationship between consumption and asset returns. There has also been a change of focus, away from the traditional question of the effects of interest rates on saving, and towards the direct investigation of the intertemporal relationships discussed in Chapter 1. However, these studies are typically less concerned with the effects of returns on the shape of the intertemporal profile of consumption than with the implications for asset prices of their role in enabling agents to allocate consumption over time in

an uncertain and risky environment. Before turning to this more modern literature, it is worth reviewing briefly the conclusions of the earlier work that tried to resolve the theoretical ambiguity by direct empirical analysis of the effects of interest rates on saving.

Traditional analyses of saving and interest

In the literature on the time-series consumption function, there are many studies where an interest rate is included in the model, along with income and other variables. There are also studies from cross-sections of countries that typically also incorporate the rate of growth effects discussed in the previous section. My reading of this literature is that the empirical results are as ambiguous as is the theory, or more positively, that the empirical results confirm the lack of invariance to time, place, and other variables that the theory predicts. Many of the studies have serious reproducibility problems, in that the results appear to depend on particular data sets, particular sample choices, particular specifications, or particular econometric techniques. For example, the widely cited study by Boskin (1978), which estimates a large response of saving to interest rates, uses data definitions that are far from standard, and the results do not recur in studies using more familiar data. Certainly, no one has generated the sort of robust finding that has commanded widespread assent among other researchers. Much of the difficulty reflects the fact that the amount of information in aggregate time-series data is rather low, so that it is generally difficult to measure precisely the effects of variables other than income, and to a lesser extent wealth.

The study by Blinder and Deaton (1985) can be used to document many of the problems. Reasonably precise estimates can be obtained for the effects on purchases of non-durables and services of current and lagged income, as well as, at least in this study, of the relative price of durable and non-durable goods, and of the change in wealth, a variable that is dominated by changes in the value of the stock market, and is essentially orthogonal to its own past and to other actual and potential regressors. Beyond that, there is a long list of other variables that might be included on theoretical grounds. Blinder and Deaton examine the effects of stocks of durable goods, of interest rates, of nominal inflation (see Deaton 1977), of government budget surpluses or deficits (Barro 1974), of temporary taxes, and of the timing of tax payments, and even this list hardly exhausts