Introduction: What Determines Savings?

Economists have long been interested in savings for several important reasons. Savings constitutes the supply of capital, which with labor represent the two primary inputs to production. Savings or nonhuman wealth, together with human wealth, determine which individuals or families, which regions, and which nations are rich or poor. And savings, which equals the sum of past saving (nonhuman wealth accumulation), influences current saving of both human and nonhuman wealth; indeed, savings is the central economic link connecting the past, the present, and the future.

The study of savings brings together analysis of consumption choice, labor supply, demographics, economic growth, and government policy. This joint analysis raises, in turn, other issues, such as the nature of preferences, the rationality of expectations, the degree of economic risks, the completeness of insurance markets, and the role of credit institutions. While it would be convenient if each of these influences on savings could be understood one at a time and in partial equilibrium, such is typically not the case. On the contrary, since savings depends, in part, on a large array of interrelated economic choices and since these choices are intertemporal in nature, understanding savings requires a dynamic general equilibrium framework.

Recent research on savings has built on the seminal contributions of Solow (1956) and Modigliani and Brumberg (1954) and the growth theory literature of the 1960s. Some of this research involves new models of consumption choice [e.g., Barro (1974)]; some involves new tests of intertemporal optimization [e.g., Hall (1981)]; some involves new empirical approaches [e.g., Feldstein (1974c)]; some uses simulation analysis [e.g., Auerbach and Kotlikoff (1987)]; and some is based on economic experiments [e.g., Johnson, Kotlikoff, and Samuelson (1987)]. The new research has made real progress: The simulation analyses have helped to distinguish potentially important from potentially unimportant determinants of savings; the empirical research has demonstrated a much greater role of intergenerational transfers in US savings than was previously believed to be the case; the theoretical research has demonstrated the importance for savings of economic ties within the family; and the experimental research has raised questions about the ability of individuals to make rational intertemporal choices.

Fortunately or unfortunately, this research progress has, on balance, increased the number of factors that have been identified as playing a major role in determining the amount of savings. Take, as an example, the recent analysis of nonaltruistic bequests that arise when annuity insurance cannot be purchased on the market and is not provided by the government or by employers. Kotlikoff and Spivak (1981), Davies (1981), Eckstein, Eichenbaum, and Peled (1983), Hubbard (1984a,b), Abel (1985), and Kotlikoff, Spivak, and Shoven (this volume) demonstrate the considerable additional savings that can result because of nonaltruistic bequests. For example, according to Kotlikoff, Spivak, and Shoven's results, the absence of annuity insurance or a close substitute can mean as much as a doubling of savings in life cycle economies.

Discovering more factors that can significantly affect savings exacerbates the problem of determining empirically which factor or combination of factors is most important in understanding any particular change in saving behavior. Consider, for instance, the question of why the US saving rate in the postwar period is roughly two-thirds the rate observed between 1870 and 1930. The list of possible explanations includes (1) increased provision of health insurance, disability insurance, unemployment insurance, and annuity insurance that reduced the need for precautionary savings, (2) an increase in government consumption and distortionary taxation, (3) a decline in altruistically motivated bequests to children, (4) an increase in subjective time preference rates, (5) deficit finance by the federal government, both explicit and implicit (such as unfunded social security), (6) a reduction in the labor supply of younger workers because of increased college attendance, (7) welfare programs that reduce the risk of abject poverty and precautionary savings to avoid poverty, and (8) the asset tests of Medicaid and other welfare programs that reduce incentives to save.

Each of these factors may have played a major role in reducing saving. Indeed, as a group they may more than explain the decline in saving. Hence one also needs to consider events that may have increased saving and offset this set of potential saving deterrents. Postwar factors that should have increased US saving include (1) the prolonged baby bust that started in the early 1960s, (2) the increase in life expectancy, (3) the increase in early retirement, (4) the increased earnings of young and middle-age women, (5) the decline in family financial support, (6) the increased risk and cost of nursing home institutionalization, (7) government subsidization of pension and similar forms of saving, and (8) historically high real after-tax returns during the 1980s.

Extended lists of this kind are a start in considering questions such as the postwar decline in US saving, but it is a far cry from such lists to the empirical knowledge needed to assess to relative contributions of different factors to changes in savings. Such empirical knowledge is a long way off not only because more and better data are needed but also because some issues of great importance to savings are surprisingly difficult to test. The value of a review of knowledge concerning savings is not, therefore, to say precisely what determines actual savings. The value is to identify key determinants, to indicate what is known about their savings effects, and to indicate areas for future research. In this survey I first identify eight principal determinants of savings within a general framework. I then review the state of knowledge concerning each of these savings determinants. In the final section I indicate directions for future research.

A Framework for Analysis

To organize the discussion, I find it useful to consider the accumulation of wealth in a two-period overlapping generations (OLG) model with a single good. The one-good two-period OLG model is the simplest way to consider intertemporal choices. Its use places no restrictions on behavior; for example, in making consumption choices agents can have life cycle intertemporal preferences à la Modigliani and Brumberg (1954), they can have altruistic preferences à la Barro (1974), or they can have Keynesian preferences.

Consider first the model in the absence of government, and assume that members of each generation are identical. Agents are adults for two periods, corresponding to youth and old age. New adults appear at the beginning of each period at which time they, together with old adults (their parents), go to work with the capital supplied by the old adults. At the end of each period output is produced, factors are paid, consumption occurs, and intergenerational and intragenerational private transfers are made. At the beginning of any period the young adults have no wealth, since they do not receive any income or transfers until the end of the period. All wealth at the beginning of any period is held by the current old generation. The assets of each old person at time $t + 1$, A_{t+1}, equal the assets they accumulated when young, i.e., at the end of time t. As indicted in equation (1), this accumulation equals the labor earnings received by each young adult at time t, E_{yt}, plus the net intergenerational transfers received by each young adult at time t, I_t, less the consumption by each young adult at time t, C_{yt}.

$$A_{t+1} = E_{yt} + I_t - C_{yt}. \tag{1}$$

The consumption when young, C_{yt}, of each member of generation t (those who are young adults at time t) is constrained by

$$C_{yt} + A_{t+1} \leqslant E_{yt} + I_t. \tag{2}$$

If young adults at time t are unable to borrow against their future labor earnings or their possibly positive net future transfers from their children ($I_{t+1} < 0$), constraint (3) will also apply:

$$A_{t+1} \geqslant 0. \tag{3}$$

Th second-period budget constraining consumption of members of generation t when old, C_{ot+1}, and their transfers I_{t+1} to their children is

$$C_{ot+1} + I_{t+1} \leqslant A_{t+1}(1 + r_{t+1}) + E_{ot+1}, \tag{4}$$

where E_{ot+1} is the labor earnings of the elderly at $t + 1$ and r_{t+1} is the second period return on savings accumulated in the first period. If the elderly in period $t + 1$ are retired, E_{ot+1} equals 0. Second-period labor earnings and the rate of return may be uncertain from the perspective of the first period.

While this framework is simple, it is sufficient to point out a number of determinants of savings. To focus attention on the determinants of the economy's savings at time $t + 1$, A_{t+1}, let us assume that (2) holds with equality and rewrite the equation as

$$A_{t+1} = w_{yt}L_{yt} + I_t - C_{yt}, \tag{2'}$$

where we replace E_{yt} by the wage at time t paid to young workers, w_{yt}, multiplied by the supply of labor of young workers at time t, L_{yt}.

The Nature of the Consumption Function

There are three principal alternative theories of consumption: the life cycle model of Modigliani and Brumberg (1954) and Ando and Modigliani (1963), the Barro (1974) model, and the Keynesian model. Friedman's (1957) permanent-income model of consumption might be viewed by some as a distinct theory of consumption, but it seems consistent with either the life cycle or Barro model. In brief, the life cycle model, or at least a stylized version of it, views young adults as caring only about their first- and second-period levels of consumption and leisure. Since there is no bequest motive, in the absence or presence of life span uncertainty but with perfect annuity markets, I_t in (2′) will be 0. If life cycle agents are not liquidity constrained, i.e., (3) is not binding, they will make their decision about C_{yt} and L_{yt} taking full account of the second-period budget constraint (4). In the absence of uncertainty, taking account of equation (4) means properly discounting future labor earnings in determining one's first-period consumption and labor supply decisions.

The Barro model views individuals as caring not only about their own welfare but also about their children's. Since children care about grandchildren and grandchildren care about great grandchildren, etc., caring about one's children translates into caring about all future descendants. This intergenerational altruism toward one's descendants provides a motive for positive values of intergenerational transfers I_t. Of course, altruism could run in the opposite direction, with children caring about their parents, in which case one would expect negative values of I_t. The degree of foresight assumed in the Barro model is far greater than that in the life cycle model; in the Barro model young individuals at time t have to take into account their own current and future labor earnings and inheritance in deciding how much to consume, but they also have to take into account the future labor earnings and consumption of all their descendants.

The Keynesian model, in contrast to the other two, is not based on an optimization problem; rather it simply asserts that agents consume some fraction of their current incomes and some other fraction of their assets. In time series regressions the estimated marginal propensity to consume out of income is typically close to unity, while the estimated marginal propensity to consume out of wealth is a fairly small fraction. Since Keynesian agents apparently do not consume all their lifetime resources, the Keynesian model also suggests a positive value of I_t in (2′).

While there is no formal justification for the Keynesian model, the informal argument rests on some combination of liquidity constraints [equation (3)] and myopia. Clearly, in an economy with positive wealth, not all consumers can be liquidity constrained, but, depending on the distribution of wealth and the timing of labor earnings over the life cycle, a significant fraction of consumers may be liquidity constrained. The notion that consumers are to some extent myopic seems to reflect the view that many consumers will fully or partially ignore the future because the future is highly uncertain and because planning for an uncertain future is a complex and difficult problem.

The Shape of the Age-Consumption Profile

Regardless of which consumption model governs choices, the ratio of C_{yt} to C_{ot+1}, holding constant the present value of lifetime consumption expenditures $[C_{yt} + C_{ot+1}/(1 + r_{t+1})]$, the level of transfers I_t, the interest rate r_{t+1}, and the values of earnings E_{yt} and E_{ot+1}, is an important determinant of savings. According to (2′), a decision by agents to consume more when young, which they finance by consuming less when old, means a lower stock of savings. In simple formulations of the life cycle model in which labor supply is exogenous, one can change just the shape of the age-consumption profile by changing the rate of time preference or the intertemporal elasticity of substitution. In the Barro model there are also ways of changing the shape of the age-consumption profile without changing E_{yt}, E_{ot+1}, r_{t+1}, or I_t. For example, in the Boskin and Kotlikoff (this volume) version of the Barro model, the shape of the age-consumption profile is controlled by age-specific utility weights.

The Keynesian model does not share this feature (often referred to as Ricardian equivalence) of non-liquidity-constrained neoclassical models that the timing of consumption over the life cycle is potentially independent of the timing of exogenous income. In the Keynesian model a present value neutral change in exogenous income that increases income when young and decreases income when old produces an increase in consumption when young. Stated differently, in the Keynesian model the shape of the age-consumption profile depends critically on the shape of the age-income profile.

In addition to changes in agents' general preferences for current as opposed to future consumption, there are other kinds of preference changes that may change the ratio of consumption when young to consumption when old and influence savings. Suppose, for example, that individuals are interested in consuming particular goods in their second period, such as the services from an owner-occupied house. As the demand for second-period housing services rises, agents will save more when young by reducing C_{yt}. As a consequence, total savings in the economy will rise.

The Shape of the Age-Earnings Profile

Just as the time path of consumption can change in neoclassical models without altering the time path of labor earnings, the time path of labor earnings can change with no change in the time path of consumption or the amount of intergenerational transfers. Present-value neutral changes in the timing of labor earnings that involve increased earnings when young, E_{yt}, but leave C_{yt} and I_t unchanged will, according to (2′), raise savings.

Of course, many determinants of the shape of the age-earnings profile also affect the present value of earnings. One example is investment in human capital. The decision to spend more time in school or in training when young may be viewed in this model as lowering E_{yt} in order to raise E_{ot+1}. The decline in E_{yt} lowers savings, and, if the human capital investment has the effect of raising the present value of lifetime earnings, the investment will also likely increase C_{yt} in the life cycle and Barro models. In the Keynesian model the consumption response to increased human capital investment could well be a decline in C_{yt}. In all three models, however, increased human capital investment is likely to imply decreased investment in nonhuman capital, i.e., decreased savings.

The age of retirement is a second important determinant of the age-earnings profile. A decline in the age of retirement can be viewed within the two-period model as a reduction in E_{ot+1}. Presumably, agents who retire earlier will consume less (save more) when young to make up for the loss in earnings in old age (at least in neoclassical models). The reduction in C_{yt} means an increase in savings, holding E_{yt} and I_t fixed. Viewing bequests and inheritances as fixed in response to changes in lifetime earnings may not, however, be appropriate. If bequests are desired because of intergenerational altruism, agents are likely to reduce their bequests in old age and their first- and second-period consumption in response to an earlier retirement. The effect of lower bequests and inheritances could outweigh the reduction in C_{yt} and spell an ultimate net increase in savings (Skinner 1985b). On the other hand, if inheritances arise because annuity insurance is incomplete and agents die before they consume all their assets, the initial higher savings of the first generation that begins to retire early will lead to more, not fewer, bequests to the next generation, as some of the additional assets are passed on to the next generation. This process could lead to more, not less savings in the long run.

A third determinant of the shape of the age-earnings profile is the rate of secular productivity growth. Clearly productivity growth raises savings by raising real wages. In terms of (2′) the value of W_{yt} will be larger because of productivity growth. But higher rates of productivity growth also mean steeper age-earnings profiles; wages in the second period of each agent's life will be higher than in the first because of the secular productivity growth. As a consequence, C_{yt} will likely rise by a greater percentage than W_{yt}. In addition, the labor supply response to higher wages in the second period is likely to be a decline in first-period labor supply and an increase in second-period labor supply. Hence, the increase in savings resulting from productivity growth is likely to be less than proportional to the increase in the first-period wage W_{yt}.

Changes in Real Interest Rates

As stressed by Summers (1981a), changes in real interest rates, holding constant the present value of resources, represent uncompensated changes in the relative price of current and future consumption. But real interest rate changes also change the present value of human wealth through the discount factors [$1 + r_{t+1}$ in equation (4)]. In conventional life cycle and Barro models an increase in interest rates typically leads to a decline in consumption when young. In terms of equation (2′), this means an increase in savings. If labor supply is also subject to choice, an increase in interest rates is likely to increase current labor supply relative to future labor supply. Thus E_{yt} is likely to rise when interest rates, ceteris paribus, increase.

The effect of higher interest rates on savings in the Keynesian model is less clear. For young agents who initially have no assets, a higher interest rate will not mean higher first-period income and consequently will not affect C_{yt} through that channel. On the other hand, the amount of inheritance I_t received from one's parents might increase.

In thinking about interest rate changes one must keep in mind that, except for small open economies in which the interest rate is set from abroad, a change in real interest rates is typically associated with a simultaneous change in real wages. Since real wage changes can independently influence savings, analyses of the effects of interest rate changes on savings need to control for changes in real wages.

Intergenerational Transfers

If the intergenerational transfer to the young at time t increases because of increased altruism by their parents, then I_t in (2′) will rise. If the other terms on the right-hand side of (2′) remain unchanged, savings will increase as well. The response of young agents to a larger inheritance is likely to be an increase in their own consumption when young and old and an increase in the bequests to their children. The young may also reduce somewhat their labor supply when young and old. Since the young will spend only a fraction of the increased inheritance on C_{yt}, the net impact on savings of an increase in I_t is likely to be positive. Of course, the young at time t may also experience an increased desire to bequeath. In this case the increase in savings associated with a higher value of I_t may be larger.

Intergenerational transfers may also reflect the absence of an actuarially fair market in annuity insurance. In the two-period model presented so far, there is no uncertainty about the data of death. But the model can easily be altered to include the case when a fraction of each generation dies at the beginning of the second period. Obviously, those who die will have zero end-of-second-period consumption and earnings. Hence the budget constraints (2) and (3) need to be reconsidered. Suppose that we view the variables in these equations as describing the average amounts of consumption, earnings, and assets per individual in generation t. Then the constraints are valid, and savings is still determined by equation (2′).

To see this point, suppose that the strict (no intentional bequest) life cycle model holds and that there is an actuarially fair annuities market. In this case no one will make bequests, and I_t and I_{t+1} will be 0. If one takes equations (2) and (4) as holding with equality and combines them by substituting out for A_{t+1}, the resulting equation states that the average (expected) value of consumption when young of members of generation t plus the average (expected) value of consumption when old of members of generation t equals the average (expected) present value of generation t's earnings when young and old. This is precisely the result one expects of a fair annuities market; rather than leaving any resources to the next generation, a fair annuities market transfers the assets of those in generation t that die early to those that survive to the end of their second period. Hence the present value of total generation t realized consumption divided by the size of generation t equals the present value of total generation t realized earnings divided by the size of generation t.

If annuity markets are imperfect or, in the limit, nonexistent, bequests are likely to arise despite the lack of a bequest motive. In this case I_t and I_{t+1} will be positive and should be viewed as the average amount of inheritance received and the average amount of bequests made by members of generation t. Thus, in contrast to the case of a perfect annuity market, A_{t+1} in (2′) is likely to be larger because I_t is positive. True, the average value of C_{yt} in (2′) may also be larger, reflecting the positive average value of inheritance I_t and the possible desire to hedge the uncertainty of the date of death by consuming more when young. But, on balance, the effect of imperfect annuity insurance is likely to raise savings, potentially by a large amount.

Uncertainty

While the positive effect on savings of lifetime uncertainty in the absence of annuity insurance works primarily through an increase in I_t, a reduced value of C_{yt} is often the source of precautionary saving in response to other types of uninsured uncertainties. Earnings uncertainty is one example; consider a mean preserving spread of the value of E_{ot+1}. Although Sandmo (1970) has shown that the response of C_{yt} to uncertainty in E_{ot+1} is theoretically ambiguous, the response that seems most plausible empirically is a decline in C_{yt}, and, according to (2′), an increase in savings. Other types of uncertainty, such as uncertainty concerning future government taxes and transfers and uncertainty concerning future health expenditures, is also likely to lead to significant precautionary savings (e.g., see chapter 6).

Clearly, the provision of various forms of insurance, such as unemployment and disability insurance, is likely to reduce the amount of precautionary savings. But even when formal insurance markets do not exist, insurance may nonetheless be available. The reason is that families can self-insure; that is, they can form an implicit, incomplete insurance market (see chapter 4 for an example). For instance, with earnings uncertainty the ability of a spouse to increase his or her earnings in case the other spouse suffers a loss in earnings constitutes implicit earnings insurance. Family insurance of this kind can reduce the precautionary demand for savings.

Demographics

Obviously, since equation (2′) describes the assets per elderly in the two-period economy, the product of the population size of generation t and A_{t+1} determines the total amount of savings. Since the real wage and the real interest rate depend on capital per worker, assuming a constant returns to scale technology, the amount of savings per worker a_{t+1} expressed in (5) is the focus of most discussions of savings and population growth:

$$a_{t+1} = (E_{yt} + I_t - C_{yt})/(1 + n). \tag{5}$$

In (5) the term n stands for the population growth rate; that is, for each old person there are $1 + n$ young people at any point in time. Holding E_{yt}, I_t, and C_{yt} constant, an increase in n leads to a decline in assets per worker. In general equilibrium this effect would feed back to lower real wages and raise the real rate of return. The net impact of these factor price changes is likely to be a further reduction in savings per worker.

A second important effect of population growth on saving per worker involves the added consumption expenditures and loss of female earnings from raising more children. While the model so far has ignored children and has treated agents as effectively being born into adulthood, this unrealistic treatment can be easily modified. Suppose that we model young adults as having young children whose consumption is included in C_{yt}. Then the increase in children per young adult will mean an increase in C_{yt} and, possibly, a decline in E_{yt} because of mothers' time out of the labor force. The explanation for the higher value of C_{yt} is the following: While the increased consumption expenditures on children will necessitate a reduction in the adult parent's consumption both when young and when old, on balance, C_{yt}, which now includes expenditures on children, will increase. Stated differently, the increase in children per young adult leads to a shift in the age-consumption profile toward more consumption in the first period. The higher value of C_{yt} and the possibly lower value of E_{yt} provide an additional reason why population growth reduces saving per worker.

While these effects of demographics on savings are easy to see, demographic change can have more subtle effects on saving. Consider the effect on savings of an increase in divorce rates. The uncertainty about the longevity of one's marriage and therefore the extent of support by one's spouse may lead each spouse to accumulate more human capital when young to hedge the possibility of being divorced and self-sufficient in the future. As mentioned, increased human capital accumulation is likely to occur at the expense of nonhuman capital accumulation. On the other hand, uncertainty about divorce could raise, rather than lower savings; young agents may try to hedge themselves against future divorce by lowering C_{yt}. In general, the actual or potential size of one's family helps determine the extent of possible implicit family insurance. And the smaller the extent of family insurance, the greater the motive for precautionary savings.

Fiscal Policy

Extending the Analytical Framework

So far I have not introduced fiscal policy into the savings framework. While there are several notational choices for discussing fiscal policy, the use of certain notation and accounting conventions can blur the fundamental similarities of a variety of fiscal policies (see chapters 7 through 10). To avoid any possibility of fiscal illusion, assume that fiscal policy works in the following manner: In each period t the government simply orders the young generation to purchase an amount G_t per young person of consumption for the government. The government in period t also orders the old generation to transfer an amount Z_t per young person to the young generation. With these assumptions and the simplifying assumption of zero population growth, the amount of savings per old person at time $t + 1$ can be written as

$$A_{t+1} = w_{yt}L_{yt} + I_t + Z_t - C_{yt} - G_t. \tag{6}$$

And, assuming either that A_{t+1} is positive or that there are no liquidity constraints, the present value budget constraint is now

$$C_{yt} + \frac{C_{ot+1}}{1 + r_{t+1}} + \frac{I_{ot+1}}{1 + r_{t+1}} + G_t + \frac{Z_{t+1}}{1 + r_{t+1}}$$

$$= w_{yt}L_{yt} + \frac{w_{ot+1}L_{ot+1}}{1 + r_{t+1}} + I_t + Z_t. \tag{7}$$

Equation (6) suggests two important possible effects of fiscal policy on savings. One involves the effects of redistributing across generations, and the other involves the effects of increases in government consumption.

Intergenerational Redistribution

Consider first the issue of intergenerational redistribution in a pure life cycle model in which I_t and I_{t+1} are 0. Holding G_t fixed, suppose that at the beginning of time t the government orders the elderly to increase their end of time t required transfers Z_t to the young. Also suppose that there is no change in Z_{t+1}. This policy will then lead to an increase in A_{t+1}; to see this, note that Z_t in (6) rises. Although C_{yt} will also increase, if second-period consumption is a normal good, the increase in C_{yt} will be less than the increase in Z_t, and A_{t+1} will increase. Next, examine the effect of an ongoing transfer from the old to the young in which Z_t, Z_{t+1}, Z_{t+2}, etc. are all increased by a fixed amount. Equations (6) and (7) indicate that this policy will also increase savings but in future periods as well as period $t + 1$. According to (7) an equal increase in Z_t and Z_{t+1} implies an expansion of the present-value budget opportunities and will lead to an increase in C_{yt}; but the increase in C_{yt} will be smaller than when Z_{t+1} is fixed. Again, from (6), A_{t+1} rises because Z_t increases more than C_{yt} (and more than $w_{yt}L_{yt}$ falls, if L_{yt} is variable). The same reasoning indicates that A_{t+2}, A_{t+3}, etc. are all larger because of the government's intergenerational redistribution. The change in savings will have general equilibrium implications for the time path of wage rates and interest rates, but these changes in factor prices are likely to reinforce the partial equilibrium effects on the time path of savings.

In the Barro model temporary or permanent increases in the time path of Z_t will have no effect on savings because the time path of I_t will adjust to exactly offset changes in Z_t. For example, suppose that the government orders an increase in the time path of the Z_t's in which Z_t, Z_{t+1}, Z_{t+2}, etc.

all increase by a fixed amount. In the Barro model, I_t, I_{t+1}, I_{t+2}, etc. will all be reduced by the same fixed amount, leaving the time path of assets (the A_t's) unchanged. The intuition here is that in the Barro model families are otherwise transferring resources across generations. Since the government's policy does not alter the total resources of the intergenerationally altruistic extended family, the family will simply offset the government's policy by returning in terms of private transfers the resources it was forced to transfer by the government.

If one models Keynesian consumption behavior as involving agents who consume a fixed fraction, say θ, of their income plus government transfers (but not including private transfers), then an increase in Z_t in the Keynesian model will likewise have no effect on A_{t+1}. To see this, note that C_{yt} will rise by $\theta \Delta Z_t$. But I_t in (6) will fall by $(1 - \theta)\Delta Z_t$. And since Z_t in (6) rises by ΔZ_t, the net impact on A_{t+1} is 0. Unlike the life cycle model in which the marginal propensity to consume of the old is unity and exceeds the marginal propensity to consume of the young, in the Keynesian model the marginal propensities to consume of the young and old are equal; hence redistribution between the young and the old has no effect on savings.[1]

Increases in Government Consumption

The second policy to consider is changes in G_t. In practice, increases in government consumption will likely involve concomitant changes in the time path of the Z_t's; that is, the government will typically spread the burden of paying for an increase in its consumption over all generations, including the initial old generation. But since changes in the time path of the Z_t's have just been discussed, it may help to clarify effects by discussing changes in G_t holding the time path of the Z_t's fixed.

In the life cycle model a one-time increase in G_t holding the Z_t's fixed will, in general, lead to a temporary decline in savings. In (6), A_{t+1} declines because G_t rises by more than the associated decline in C_{yt}; if C_{ot+1} is a normal good, members of generation t will pay for the higher G_t by reducing C_{ot+1} and C_{yt}. Hence C_{yt} will fall by less than G_t. In addition, L_{yt} is likely to increase if first-period leisure is also a normal good. In contrast to a one-time increase in government consumption, a permanent increase in the time path of the G_t's will permanently lower savings in the life cycle model. There will be temporary changes in factor prices for the temporary increase in government consumption and permanent changes in factor prices for permanent increases in government consumption. These general equilibrium factor price response will, in general, reinforce the partial equilibrium effects of changes in the time path of the G_t's on savings.

The impact of a temporary increase in G_t in the Barro model is also a short-term reduction in savings. The Barro dynastic family will adjust their intergenerational transfers, the time path of the I_t's, to spread the burden of paying of the temporary increase in G_t over all current and future family members. As a consequence, in the formula for A_{t+1} the reduction in G_t will lead to a much smaller offsetting reduction in C_{yt} than in the life cycle model. In addition, I_t will increase somewhat, but the likely net impact will be a decline in A_{t+1}. On the other hand, if the increase in government consumption is permanent, there may be no effect on savings. Suppose, for example, that labor is inelastically supplied and that the Barro family responds to a permanent increase of ΔG in the time path of the G_t's by reducing its total consumption each period ($C_{yt} + C_{ot}$) by the same amount ΔG. In this case total private plus government consumption is unchanged, as is national income. Since national saving equals national income less total consumption, national saving is unaffected by the policy. In terms of equation (6), the increase in G_t is exactly offset by a decline in C_{yt} and an increase in I_t, leaving A_{t+1} unaffected. If labor supply is variable, the Barro family may increase its labor supply rather than fully absorb the increase in government consumption through a decline in its consumption. In this case savings may again be unaffected, depending on the precise responses.

The Keynesian model predicts temporary and permanent declines in savings from temporary and permanent increases in government consumption, respectively, still assuming no change in the time path of the Z_t's. The decrease in disposable income leads young Keynesian agents to reduce C_{yt} by a fraction θ of the increase in G_t, but the net effect, according to equation (6), is a reduction in A_{t+1} by $(1 - \theta)\Delta G_t$.

Intragenerational Redistribution

In addition to the income effects of intergeneration redistribution (the Z_t's) policy and government consumption (the G_t's) policies, lump sum intragenerational redistribution has income effects that may alter savings. Consider at time t lump sum transfers from type A young individuals to type B young individuals. In the life cycle model if the B individuals have larger marginal propensities to consume goods (a flatter age-consumption profile) and leisure (a steeper age-earnings profile), this policy will raise the average value of C_{yt} and lower aggregate savings.

In the Barro model the story is quite similar, but one must consider how the preferences of the Barro family to which individual A belongs differ from those of the Barro family to which individual B belongs. If the Barro family containing individual B has larger marginal consumption and leisure propensities than the Barro family containing individual A, the redistribution will also reduce A_{t+1}. For example, if the Barro family receiving the transfer discounts the utility of future generations at a higher rate than the Barro family providing the transfer, the redistribution will reduce subsequent aggregate bequests and lower current saving.

Finally, in the Keynesian model, the time t redistribution will lower A_{t+1} if person B's marginal propensity to consume exceeds person A's. Even if the policy is discontinued after time t, it is likely to permanently reduce the time path of savings through its impact on the time path of the average value of the I_t's.

Distortionary Taxation

To this point the issue of distortionary taxation has been skirted by assuming that the government can simply command each young generation to purchase G_t and each old generation to transfer Z_t to the contemporaneous young generation. To add distortionary taxation to the analysis, consider distortionary taxes that are compensated in a lump sum manner to the individuals paying the distortionary taxes in the period the taxes are paid. In contrast to the Z_t's and G_t's, policies that have only income effects, these compensated distortionary taxes have only incentive effects. In combination with a proper choice of the Z_t's, G_t's, and intragenerational redistribution policies, these compensated distortionary taxes can be used fully to describe any actual fiscal policy. To see this point most clearly for neoclassical consumption theories, recall that any fiscal policy will ultimately affect individual behavior through its effects on the individual's budget constraint. Changes in an individual's budget constraint will produce movements from one indifference curve to another. But such movements can always be decomposed into a substitution effect plus an income effect.

The compensated distortionary taxes that are of most interest to a discussion of savings are compensated capital income taxes, compensated labor taxes, and compensated consumption taxes. Consider first the effects of compensated taxes in the life cycle model, and ignore changes through time in distortionary taxes. In a two-period model in which labor supply is exogenous, a compensated capital income tax leads to an increase in C_{yt} and a decrease in C_{ot+1} and, according to (6), decreases savings. If labor supply is endogenous, compensated capital income taxes are likely to increase C_{yt} and to decrease L_{yt}, also implying a decline in savings. Time-invariant compensated labor income taxes and compensated consumption taxes distort consumption-leisure choices and are likely to lead to a reduction in L_{yt} and C_{yt}. The likely net effect on savings from time-invariant compensated labor or consumption taxes is negative.

If the compensated labor or consumption taxes differ across the two periods, these taxes will distort not only the choice between consumption and leisure at a point in time but also the choice between consumption when young and consumption when old and the choice between leisure when young and leisure when old. For example, if the compensated tax rate on second-period consumption exceeds that on first-period consumption, this difference in tax rates acts, in part, like a compensated capital income tax in that it raises the price of future consumption relative to current consumption.

The effects of compensated taxes in the Barro model on first- and second-period consumption and leisure are similar. But in the Barro model one must also consider how voluntary intergenerational transfers will be affected. Since such transfers facilitate the consumption of future members of the altruistic family, a time-invariant capital income tax raises the price of Barro family consumption beyond period $t + 1$ relative to consumption at time t. Hence the impact of a time-invariant compensated tax on capital income is likely to be a reduction in I_t in equation (6) that reinforces the likely increase in C_{yt} and decline in L_{yt} in reducing savings. Time-invariant compensated labor or consumption taxes are also likely to reduce the time path of the I_t's. The reason is that the Barro family will substitute leisure, both today and in the future, for consumption, both today and in the future.

In the Keynesian model, or at least a simplified version of the Keynesian model, relative prices of current and future consumption and leisure do not influence consumption decisions. Hence, since they also do not affect disposable income, compensated taxes appear to have no savings effects in the Keynesian model.

Translating Actual Fiscal Policies into Paths of Z_t's and G_ts, Intragenerational Redistribution, and Compensated Distortionary Taxes

While space does not permit a lengthy categorization of fiscal policies in terms of the fundamental components just discussed, it may be useful to provide a few translations. To make the examples concrete, consider the recent Reagan fiscal policy. It combined income tax cuts, major reductions in future Social Security benefits, reductions in tax progressivity, a short-lived (six-year) shift from income toward consumption taxation, and a reduction in federal government consumption as a share of net national product (Boskin and Robinson 1987).

The last of these five policies can be disposed of quickly since it corresponds to an at least temporary decline in the time path of the G_t's. The income tax cut that produced the large official federal deficits can be viewed as a combination of an at least temporary reduction in compensated capital income and labor income taxes combined with an ongoing reduction in the Z_t's; that is, the deficit can be viewed as transferring to current generations from future generations by reducing the amount each generation is forced by the government to hand back to the succeeding generation.

The reduction in future Social Security benefits, in contrast, can be viewed as increasing the Z_t's, starting at time $t + 1$ (where a period stands for a generation). Since Social Security is unfunded, the 1983 cut in the future Social Security benefits of baby boomers means a smaller redistribution to them from their children and smaller subsequent redistributions from later to earlier generations; it means larger values of the Z_t's starting with Z_{t+1}. Using this more basic fiscal policy language makes clear that the 1983 Social Security policy is an intergenerational redistribution policy that offsets the intergenerational redistribution associated with the Reagan tax cuts.

The changes in tax progressivity can be described as compensated reductions in capital and labor income tax rates plus intragenerational redistribution from high- to low-income households. And the 1981 shift toward consumption taxation in the Economic Recovery Tax Act, which was reversed in the 1986 Tax Reform Act, can be understood as combining a short-lived reduction in compensated capital income taxes with a short-lived increase in Z_t that reflects the policy's short-lived likely effects on the asset values of existing capital [see Summers (1981b) and chapter 9].

As discussed in Kotlikoff (1988), the combination of these policies appears, at least from a life cycle perspective, to have provided a small stimulus to savings. This view differs, of course, from most popular accounts of the Reagan fiscal policy that have focused almost exclusively on the federal deficits of this period.

A Review of Knowledge Concerning Savings Determinants

Time Series Evidence Concerning the Consumption Function

Time series consumption regressions represent one set of studies bearing on the true nature of the consumption function. These regressions typically relate aggregate consumption to disposable income, aggregate wealth, the real interest rate, and measure of government policy, such as government debt. As a rule these studies report significant coefficients for disposable income and wealth, but the coefficients on the interest rate and government policy variables are sensitive to the precise specification of the regression and the dates of the sample.

It is quite difficult to interpret the results of the standard time series regressions with respect to alternative models of savings. The reason is that these regressions seem to mix various models rather than choose one as a null hypothesis that can be tested. An example of this problem is the time series literature that purports to test the effects of Social Security on savings by including Social Security wealth among the set of regressors. This literature includes the studies of Feldstein (1974c), Barro (1978a,b), Darby (1979), and Leimer and Lesnoy (1980). The results of this body of research can be summarized with one word, ambiguous. Even if the results all agreed, it would be difficult to know precisely what had been learned; as pointed out in chapter 19 and Williamson and Jones (1983), if the life cycle model is taken as the null hypothesis in these studies, the models are misspecified because of the inability to aggregate the behavior of different age groups. Chapter 19 shows that these time series regression procedures would reject the life cycle model even using data taken from a pure life cycle economy. An alternative view of these regressions is that the Barro model is the null hypothesis. But in this case the regressions are also misspecified because they use disposable income rather than the present value of human wealth and because they ignore the government's intertemporal budget constraint.

A different time series literature that is relevant for distinguishing neoclassical models from the Keynesian model is the Euler equation study of Hall (1978) and closely related studies of Sargent (1978), Flavin (1981), Mankiw, Rotemberg, and Summers (1982, 1983), Davidson and Hendry (1981), Daly and Hadjimatheou (1981), Cuddington (1982), Bilson (1980), Muellbauer (1982), and others. These papers test intertemporal expected utility maximization, specifically its implication that the Euler error is uncorrelated with previous information. A rejection of this null hypothesis would rule out neoclassical consumption models. But, as stressed by King (1983), tests of the Euler equation require specifying the explicit form of preferences, and rejection of the Euler equation may simply reflect an incorrect formulation of preferences. In addition, analysis of the Euler relationship using time series data requires the questionable assumption that one can aggregate the Euler equations of millions of households into a single relationship that is based on per capita consumption.

The time series tests of the Euler equation have provided mixed results. Hall's (1978) results are generally supportive of the Euler equation's implication that past information does not predict the evolution of the marginal utility of consumption. On the other hand, Flavin (1981) and Mankiw, Rotemberg, and Summers (1982) reject the equation's implication. While the studies subsequent to Hall's use the same consumption time series and as such are not independent tests of the Euler proposition, it is clear that the time series Euler relationship is not robust.

Another time series approach to testing neoclassical consumption functions that is closely related to the Euler equation tests involves estimating the volatility of income changes and determining whether consumption responses to that income volatility are too small or too large. Flavin's (1981) study provides an early comparison of income and consumption changes. She finds that consumption responses are excessively volatile. Other studies, such as Kotlikoff and Pakes (1988), West (1987), and Campbell and Deaton (1987) show that consumption is not sufficiently sensitive to changes in income. The jury is clearly still out on this question. The research by Mankiw and Shapiro (1985), Deaton (1986), Campbell and Deaton (1987), and West (1987) points out that conclusions about excessive sensitivity of consumption to income changes hinge critically on how one models the income process.

A fourth approach is to use time series data to estimate a structural model of consumption. This is possible for certain simple versions of the Barro model because this model aggregates much more nicely than the life cycle model. Chapter 18, a joint study will Michael Boskin, illustrates this approach. This study determines the optimal consumption plan for a Barro dynasty, with a known future course of demographic change, that faces uncertainty with respect to labor earnings and rates of return. The joint distribution of earnings and rates of return is estimated empirically and used to calculate the optimal annual consumption path. This predicted time path of consumption is then compared with actual consumption. While the fit is fairly good, the study goes on to consider whether the deviation between the actual time path of consumption and the predicted time path depends on proxies for the age-distribution of resources. As demonstrated in chapter 3, in the Barro model or other models in which individuals are altruistically linked, the distribution of resources within altruistically linked clans of individuals does not influence the distribution or level of consumption of such individuals. Chapter 18 reports a rejection of this proposition; the level of aggregate consumption appears, in part, to be determined by the age-distribution of resources. While the chapter 18 results cast doubt on the Barro model, they are certainly not definitive; the analysis may be subject to aggregation bias if different Barro clans have different preferences. In addition, the specification of preferences and uncertainty may be inappropriate.

A fifth type of time series research on consumption includes articles by Poterba (1987a,b) and Wilcox (1987a,b). These event studies compare monthly changes in consumption expenditure to specific one-time changes in disposable income. The goal is to determine whether the consumption changes are larger than would be suggested by neoclassical consumption functions. Poterba (1987a) and Wilcox (1987b) report evidence of sizable consumption responses that are consistent with a view that roughly one-quarter of households are Keynesian consumers. Wilcox (1987a), on the other hand, reports small consumption responses to disposable income changes—the kind of responses suggested by neoclassical models. Even if these three studies agreed, one might question whether the result was picking up something other than Keynesian behavior. The aggregation problems in these kinds of analyses seem as troublesome as the problems in other time series consumption studies.

In sum, problems of aggregation make reliance on macroeconomic time series tests of microeconomic behavior highly suspect. And the time series evidence that seems least suspect permits no strong conclusion about the nature of consumption preferences.

Cross Section and Panel Studies of Consumption

Much of the recent consumption research based on microeconomic data involves tests of Hall's (1978) Euler error proposition. This research has paid special attention to the possibility that some consumers are liquidity constrained. The conclusion that emerges from a number of studies including Hall and Mishkin (1982), Hayashi (1982, 1984, 1985), Shapiro (1986), Zeldes (1985), Runkle (1983), and Lawrence (1983) is that roughly 20 percent of US households are liquidity constrained. These households are among the poorest in the economy and presumably account for much less than 20 percent of US consumption.

Not all studies support the conclusion that even 20 percent of households are liquidity constrained or otherwise act like Keynesian consumers. In a recent and careful panel analysis of consumption, Altonji and Siow (1986) develop a regression equation that nests both the life cycle and Keynesian models. They conclude that "the vast majority of households obey the life cycle model" (p. 319). They also point out that failure to account for measurement error in income may lead one inappropriately to accept the Keynesian model. Finally, their tests of the pure life cycle model against a life cycle model with liquidity constraints "do not show much evidence against the perfect capital markets assumption" (p. 322).

Most of the Euler tests based on microeconomic data can be viewed as testing both the Barro and the life cycle models against the Keynesian liquidity-constrained alternative. What is missing in this literature are tests of the Barro model against the life cycle model. The reason for this is the lack of data on either the consumption or income of extended families. As suggested in chapter 3 and by Bernheim and Bagwell (1986), who independently reached the same conclusion, the extended family may be very, very large because of altruistic linkages arising from marriage across otherwise distinct extended families. One approach to testing the Barro model against alternatives that does not require extended family data is the study by Abel and Kotlikoff (1988). This study shows that, under the assumptions of homothetic and time separable utility, the Barro clan will respond to shocks to the incomes of its members by having the consumption of all members increase or decrease by the same percentage. With some weak additional assumptions, this proposition implies that the average percentage change in household consumption within an age cohort should be the same for all age cohorts.

Testing the Barro model by comparing average percentage changes in consumption across age cohorts is particularly advantageous because it is nonparametric; in determining whether the average consumption of different age cohorts move together, Abel and Kotlikoff place no restrictions on preferences beyond the assumptions of homotheticity and time separability. In particular, each Barro clan can have quite different preference parameters.

The null hypothesis of their test is that cohort differences in the average percentage change in consumption are due simply to sampling and measurement error. Alternative hypotheses, suggested by the life cycle model, are that (1) the percentage changes in the average consumption of any two cohorts are more highly correlated the closer in age are the two cohorts and (2) the variance in the percentage change in consumption is a monotone function of the age of the cohort. Their US Consumer Expenditure Survey data for the 1980s fail to reject the null hypothesis of equal average percentage changes in consumption. Indeed, their results provide fairly strong support for the intergenerational altruism model as opposed to the life cycle model.

In addition to the microeconomic-level Euler equation literature, there is a short literature that tests the life cycle model's implication that Social Security reduces private savings. This literature includes Feldstein and Pellechio (1979), chapter 15, Kurz (1981), Blinder, Gordon, and Wise (1981), and Diamond and Hausman (1984). The results here are mixed; some studies and aspects of some studies support the life cycle model's predictions. Others are at odds with its predictions. For example, in considering the savings response to Social Security, Blinder, Gordon, and Wise also test whether the marginal propensity to consume rises with age as predicted by the life cycle model (at least under certainty). They find no evidence that such is the case. In contrast to other microeconomic empirical research, much of the data used in this literature are quite weak. Since the key social security variables are constructed based on estimates of future labor earnings and since many of these data sets report only current earnings, the critical Social Security variables are often estimated based on current earnings. Hence these variables are measured with error—error that is correlated with other resource variables in the regressions.

To summarize the findings based on microeconomic data, the strongest empirical evidence supports neoclassical models over the Keynesian liquidity-constrained alternative for most US households. For a minority, about 20 percent, the Keynesian liquidity-constrained model is probably most appropriate. Determining which neoclassical formulation best describes the consumption behavior of most US households remains an important task but a task that may require extended family data as well as the construction of statistical frameworks that nest both the life cycle and Barro models.

Testing the Rationality of Life Cycle Consumption Choice

The assumption that consumers make rational choices is a maintained hypothesis that is rarely tested in empirical research, although it is viewed by many as implausible [e.g., Thaler and Sheffrin (1981)]. One way to explore the issue of rationality is through experiments. Johnson, Kotlikoff, and Samuelson (1987) report an experiment in which subjects are paid to answer a computerized consumption questionnaire. The experiment is nonsalient in that subjects are not rewarded for their answers. Despite this feature, subjects appeared to take the experiment seriously and took their time to answer the questionnaire. The questionnaire asks subjects to place themselves in a simple life cycle setting, one in which they are single, are of a given age, face no uncertainty, face a given life span, have a given amount of initial assets and a given stream of labor earnings, and can borrow and lend at a fixed real interest rate. As the computer changes the parameters of the life cycle setting, it solicits consumption choices.

The questions in the experiment are designed to test rational choice and to elicit information about preferences. The questionnaire intersperses, in a subtle manner, seventeen pairs of cases in which the present value of resources is the same but the mix between human and nonhuman wealth differs. Many of the subjects clearly undervalued future earnings streams, a result that is suggested by the Keynesian model. Indeed, the subject's responses suggest a widespread inability to make coherent and consistent consumption decisions. Errors in consumption decision making appear to be substantial and, in many cases, systematic. In addition, the experiment's data strongly reject the standard time-separable homothetic life cycle model of consumption choice.

While this type of research is in its infancy, these findings raise serious questions about the ability of the typical person to make consumption choices that, at least on average, mimic the results of highly complex dynamic optimization problems.

A second way to consider consumer rationality is to examine whether households' saving and insurance decisions are grossly at odds with what one would expect sensible people to do. Diamond (1977) and Kotlikoff, Spivak, and Summers (this volume) examine the adequacy of savings, and Holden, Burkhauser, and Myers (1986), Myers, Burkhauser, and Holden (1986), Auerbach and Kotlikoff (1987), and Hurd and Wise (1987) consider whether the amount of life insurance households purchase is adequate. The method Kotlikoff, Spivak, and Shoven used to consider the adequacy of savings is to compare, albeit inferentially, the level of consumption in old age with the level of consumption when young. Evidence of dramatically lower levels of consumption when old relative to consumption when young would suggest either remarkably poor planning or extreme poor fortune. While Diamond reports that a significant fraction of US elderly have small or zero amounts of net worth, Kotlikoff, Spivak, and Summers show that, once one takes account of Social Security, pensions, and labor earnings, there is no evidence of widespread irrational undersaving. Quite the contrary, they find surprisingly few elderly households who appear to have saved far too little.

The research on life insurance purchase leads, however, to the opposite conclusion. Here it seems quite clear that most US couples, in which one spouse is the principal earner, are underinsured and a significant fraction are dramatically underinsured. In addition, households fail to adjust their purchase of life insurance in response to survivor insurance provided by Social

Security, suggesting that they improperly value Social Security's provision of insurance.

In sum, the initial experimental evidence and the analyses of life insurance holdings raise real questions as to the ability of consumers to make complex intertemporal consumption choices consistently and coherently. When juxtaposed with the microeconomic and time series evidence, the resulting picture concerning the true nature of the consumption function is quite mixed. Given the reservations about the time series studies, one should probably given strongest weight to the cross section and panel studies, which, on balance, strongly support the neoclassical view; but research on consumer rationality and casual introspective give one considerable pause.

How Changes in the Shapes of Age-Earnings and Age-Consumption Profiles Affect Savings

An Illustration from the Auerbach-Kotlikoff Simulation Model

The importance of the shape of consumption profiles is demonstrated in Tobin's (1969) and subsequent simulation studies of life cycle savings. There has been much less analysis of the savings effects of changes in the shapes of age-earnings profiles on savings. The Auerbach-Kotlikoff dynamic life cycle simulation model (henceforth, the AK model) is convenient for illustrating the effects of changes in preferences that influence both age-earnings and age-consumption profiles. The AK model, which is described in detail in Auerbach and Kotlikoff (1987), is a pure life cycle model, which solves for the economy's perfect foresight general equilibrium transition path from one steady state to another. Agents life for 55 periods, face no uncertainty, and have a time-separable CES utility function in consumption and leisure. There are two factors of production, capital and labor, and the production function is also CES.

Equations (8) and (9) indicate, respectively, how the growth of consumption and leisure as one ages in the AK model are related to the model's preference parameters and factor prices:

$$C_t = [(1 + r_t)/(1 + \delta)]^{\gamma}(v_t/v_{t-1})C_{t-1}, \tag{8}$$

$$l_t = [(1 + r_t)/(1 + \delta)]^{\gamma}(v_t/v_{t-1})^{-\rho}(w_t^*/w_{t-1}^*)^{-\rho}l_{t-1}, \tag{9}$$

where

$$v_t = [1 + \alpha\rho w_t^{*(1-\rho)}]^{(\rho-\gamma)/(1-\rho)} \tag{10}$$

and

$$w_t^* = (w_t e_t + \mu_t). \tag{11}$$

In these equations w_t and r_t are the wage rate per effective unit of labor and the interest rate, respectively, in period t. In the presence of taxes these factor prices need to be modified appropriately. The terms γ, ρ, δ, and α are, respectively, the intertemporal elasticity of substitution, the static elasticity of substitution, the time preference rate, and a leisure share parameter that enters the CES utility function. The term w_t^* is the worker's effective wage at time t. It equals the sum of μ_t, the shadow wage on the constraint on nonnegative labor supply plus the product of the economy-wide wage per unit of effective labor w_t and e_t, which determines the age profile of effective labor supply; that is, the e_t profile permits workers to become more productive over the life cycle.

According to equations (8) and (9), the larger δ is, the time preference rate, the flatter the age-consumption profile and the steeper the age-earnings profile (since the age-leisure profile is flatter)—both of which reduce savings. The values of the other taste parameters influence the effects of changes in δ on the slopes of the age-consumption and age-earnings profiles. Next consider simultaneous increases in γ and ρ under the assumptions that γ and ρ are equal and that w_t^* does not change over the life cycle. In this case the growth rates of consumption and leisure over the life cycle are identical, and, assuming that r_t exceeds δ, increases in γ steepen the age-consumption profile and flatten the age-earnings profile over the life cycle—both of which stimulate savings. When γ and ρ are not equal, the effect of increases in γ will also depend on the shape of the age-effective wage profile. The shape of the e_t profile, on which the shape of w_t^* in part depends, is based on estimates for the United States by Welch (1979). The e_t profile peaks at age 30 (corresponding to a real world age of about 50), and wages at that age are 45 percent greater than at age 1 (a real world age of 21).

Table 1 reports the impact in the AK model of different values of γ and δ on steady-state aggregate savings, national income, aggregate labor supply, the wage per effective unit of labor input, and the interest rate. The table assumes a 15 percent proportional income tax, a value of ρ of 0.8, a value of α of 1.5, and a Cobb-Douglas production function in which capital's income share is 0.25. The changes in preference parameters in the table are not small in terms of their effects on behavior. For example, when γ equals 0.25, a value of ρ of 5 percent rather than 1.5 percent means a growth rate of consumption that is almost 60 percent larger.

Table 1
Effects on steady-state savings and other variables of changes in the shapes of age-earnings and age-consumption profiles

Parameters γ	δ	Savings	Income	Labor supply	Wage rate	Interest rate (%)
0.25	0.015	95	25	19.1	1.00	7
0.10	0.015	45	21	19.3	0.82	12
0.50	0.015	148	29	19.2	1.11	5
0.25	0.050	55	22	19.2	0.87	10
0.25	−0.030	202	31	19.3	1.20	4

As demonstrated by table 1, savings is highly sensitive to variations in preferences that influence the shapes of age-earnings and age-consumption profiles. When the intertemporal elasticity of substitution γ increases from 0.25 to 0.50, savings increases by more than 50 percent. When the rate of time preference δ increases from 1.5 percent to 5 percent, savings declines by almost half.

The Effects of Liquidity Constraints on the Age-Consumption Profile and Savings

A recent simulation study of savings and liquidity constraints by Hubbard and Judd (1986a,b) provides a sense of the possible importance of liquidity constraints. This research, which follows in the footsteps of earlier work by Tobin and Dolde (1971a,b), indicates that liquidity constraints can play a major role in the determination of savings. Simulations in the Hubbard and Judd 55-period life cycle model in which young individuals are liquidity constrained for nine years (from age 20 to 29) indicate that the liquidity constraints could raise aggregate savings by as much as one-third. These results, like all simulation analyses, are sensitive to assumptions. In this case the assumption that all young individuals are liquidity constrained for extended periods seems rather strong. It ignores the possibility that the young can borrow from their parents and other relatives. By omitting intergenerational transfers, their model also ignores the possibility that some of the young inherit sufficient assets or receive large enough gifts (such as college tuition) that liquidity constraints are not binding.

The data, at least for the United States, indicate that young households on average have positive net worth. Hence it may be more appropriate to model liquidity constraints along the lines of Lawrence (1983), in which only a segment of individuals, those in occupations with steep age-earnings profiles, are liquidity constrained. Lawrence's simulation analysis suggests a smaller effect on aggregate savings of liquidity constraints.

Zeldes's (1986) analysis, on the other hand, suggests that certainty models are likely to understate the savings effects of liquidity constraints. Zeldes demonstrates that, when future labor earnings are uncertain, consumers may increase their current saving, not because they are currently liquidity constrained but because they may be liquidity constrained in the future. With the use of simulation techniques, Zeldes demonstrates that possible future liquidity constraints can have a nontrivial impact on current consumption and aggregate savings.

The Importance of Intergenerational Transfers to Savings

Chapter 2 provides an extensive review of recent research on the importance of intergenerational transfers to savings. As the chapter demonstrates, a wide variety of evidence suggests that intergenerational transfers play a major role, if not a predominant role in US wealth accumulation. Other studies, including Hayashi (1986), indicate that the same is true for Japan. The literature indicates (1) that the shapes of age-consumption and age-earnings profiles are far different from those needed to generate considerable life cycle savings [e.g., Darby (1979) and chapter 1], (2) that the elderly dissave far less than is predicted by the life cycle model [e.g., Mirer (1979), Bernheim (1986), and Danziger et al. (1983)], (3) that the distribution of wealth can only be explained by significant intergenerational transfers [e.g., Atkinson (1971), Oulton (1976)], (4) that the elderly fail to purchase annuity insurance even on actuarially fair terms (Bernheim et al. 1985), (5) that savings has been unresponsive to changes in the length of retirement, and (6) that life cycle simulation studies with reasonable parameter values predict far less wealth than is actually observed (e.g., chapter 13).

Knowing that intergenerational transfers are important and knowing why they arise are two different things. As mentioned, the recent simulation literature on nonaltruistic bequests suggests that much of non-life-cycle intergenerational transfer wealth could potentially be traced to imperfections in annuity insurance arrangements. But true intergenerational altruism —the desire to improve the circumstances of one's descendants—is surely a major, if not the major, factor in explaining savings resulting from intergenerational transfers. As chapter 2 observes, sorting out altruistic from nonaltruistic motives in intergenerational transfers appears to require data on the extended family, not simply data on individual households.

In general, the limited data on bequests and inter vivos transfers also makes studying intergenerational transfers difficult. But one source of data on bequests appears to have been underutilized. This is data taken from living respondents on their current net worth plus face value of life insurance. Together these pieces of information determine the bequest a living survey respondent would leave (ignoring burial and related costs) if the respondent were to die at the time of the survey. As chapter 16 stresses, since life span is uncertain, individuals choose a level of bequests at each point in their lives; that is, there is not a single planned terminal bequest but rather age-specific bequests contingent on dying at that age. Chapter 16 examines such contingent bequests and reaches a surprising conclusion, namely, that the elasticity of bequests with respect to lifetime resources is less than unity. This finding implies that redistribution from the lifetime rich to the lifetime poor will increase saving. It also suggests that, as countries grow, their wealth to income ratios will decline.

Precautionary Savings

The nascent literature on precautionary savings represents one of the most promising frontiers of research on savings. Life cycle simulation studies have examined precautionary savings resulting from uncertain life span [e.g., Davies (1981) and chapters 4 and 5], labor earnings uncertainty [e.g., Barsky et al. (1987), Zeldes (1986), and Skinner (1986)], and uncertain health expenditures (chapter 6). The conclusion flowing from these studies is that each of these forms of uncertainty can explain a substantial amount of precautionary savings. For example, Zeldes's (1986) and Skinner's (1986) results suggests earnings uncertainty could easily raise total savings by 25 percent. The simulations reported in chapter 6 suggest that precautionary savings resulting from health expenditure uncertainty could raise savings by one-third.

The literature also indicates that the degree of insurance against income, expenditures, life span, and other risks can greatly influence the amount of precautionary savings. Take, for example, the case of uncertain health expenditures. Table 2, extracted from chapter 6, indicates how insurance and Medicaid arrangements can alter aggregate savings. The table reports results on a 55-period life cycle simulation model in which individuals between ages 20 and 55 have a constant annual probability of becoming ill. If they become ill once and are cured, they will not become ill again. The cost of the cure of the illness is parameterized as a multiple of annual earnings. The table considers four different situations. The first

Table 2
Aggregate savings under different health expenditure regimes

Regime	Base case[a]	Cure costs = 2.5 × annual earnings	Annual illness probability = 0.01
Self-payment	1,008,670	869,710	1,136,140
Fair insurance	891,521	828,212	822,061
Live with it	527,017	527,017	682,301
Medicaid	222,062	325,871	626,383

a. Cure costs equal 5 times annual earnings, and annual illness probability equals 0.05.

case is "self-payment." In this economy there is no health insurance and individuals self-insure. In the second case actuarially fair health insurance is available. In the third case individuals choose (because of a preference parameter) to live with the sickness if they become ill rather than take the cure. And in the fourth case there is no insurance but there is a Medicaid program that has a 100 percent asset tax on its recipients.

Table 2 considers the base case in which the probability of the illness is 5 percent per year starting after age 20, and the cost of the cure is five times annual earnings. The second case involves a cost of the cure equal to $2\frac{1}{2}$ year's earnings. And the third case lowers the annual probability of the illness from 5 percent to 1 percent.

A comparison of the "live with it" regime and the "self-payment" regimes indicates the potential importance of precautionary savings. In the "live with it" regime there is no uncertainty about future health expenditures because, when people become ill, they simply live with the ailment rather than have the cure. The "fair insurance" regime has much less savings than the "self-payment" regime reflecting the reduction in precautionary saving when insurance is available. Savings in the "fair insurance" regime exceeds savings in the "live with it" regime because, even though agents are insured, there is a steepening of the age-consumption profile in the "fair insurance" regime relative to the "live with it" regime, reflecting the consumption health expenditures on the cure in old age. And the increased savings of young agents explains the increase in savings in "fair insurance" relative to "live with it." Savings in the Medicaid regime is remarkably small, even when the illness probability is only 1 percent per year. The prospect of having all one's assets confiscated by a Medicaid system is clearly a major saving disincentive.

Given the ability of families to self-insure and the lack of data exploring implicit family insurance arrangements, it is difficult to know how much to make of the precautionary savings motive. The extent of precautionary savings becomes even more difficult to assess in light of another form of insurance, namely government insurance. Government insurance includes social security, which often provides real annuities, life insurance, and disability insurance, Medicaid and Medicare, and welfare programs, which insures against poverty. Some fiscal policies can provide significant if subtle forms of insurance. For example, Merton (1983) points out that an unfunded social security system in which benefits are adjusted in response to labor earnings can pool earnings risks across generations. And Eaton and Rosen (1980), Varian (1980), Chan (1983), Barsky et al. (1987), and Kimball and Mankiw (1987) show that a progressive income tax can provide earnings insurance.

More theoretical and simulation research is needed to explore how well families and government policies hedge uncertain earnings and health expenditures and other uncertainties, such as uncertain rates of return and uncertain dates of retirement because of disability. But pinning down empirically the extent of precautionary savings will require new surveys that examine two issues: first, the nature of implicit family insurance arrangements and, second, the extent of subjective uncertainties. This latter issue has been thoroughly finessed in the precautionary savings literature by simply assuming the nature of subjective probability distributions.

One approach to determining the extent of subjective uncertainty that does not require eliciting survey responses about probability distributions is proposed by Eden and Pakes (1981) and Kotlikoff and Pakes (1988). These papers demonstrate how one can infer the degree of earnings uncertainty from information on consumption. Their technique, however, does not appear to generalize to other types of uncertainty. And in the presence of other uncertainties besides labor earnings, their technique may give an inaccurate picture of earnings uncertainty.

Demographic Change and Savings

Simulation analyses have been the main vehicle for demonstrating the importance of demographics to savings. Table 3 provides an illustration; it summarizes some findings of chapter 14, which considers the response of savings to changes in life span and the age of retirement in a 55-period life cycle model. The table reports capital-labor ratios for different ages of death and retirement. The parametrization of this model is similar to that of the AK model and is detailed in chapter 14. The table indicates that savings can be quite sensitive to the lengths of work and life spans.

Table 3
Capital-labor ratios for various life spans and work spans

Age of retirement	Age of death 50	70	80	100
40	6.89	13.29	15.99	20.46
50		10.23	11.27	17.02
60		8.02	10.13	14.22
70			8.80	11.85
80				10.15

Table 4
Effects of a baby bust on saving rates

	Fraction of population at specified ages				
Year	1–20	21–40	41–60	61–75	Saving rate (%)
0	0.50	0.28	0.15	0.07	7.6
20	0.37	0.36	0.20	0.09	7.9
50	0.28	0.28	0.28	0.16	3.0
70	0.26	0.27	0.27	0.21	0.0
110	0.27	0.27	0.27	0.21	−1.5
150	0.27	0.27	0.27	0.20	0.0

Holding the age of retirement at 60, an increase in the model's age of death from 70 to 80 raises the capital-labor ratio and quantity of savings (since labor supply is inelastic in this model) by more than 25 percent. Reductions in the age of retirement can also have important savings effects. Holding the age of death at 70, a reduction in the age of retirement from 60 to 50 leads to a 27 percent increase in the capital-labor ratio.

A second illustration of the impact of demographics on savings is given in table 4, which draws on findings from chapter 13. This chapter uses the AK model to study how baby booms and baby busts affect the economy's general equilibrium transition path. The parameters of this model are described in chapter 13, but they are quite similar to those underlying table 1. Table 4 depicts the effects on the age structure of the economy over a 150-year period of a year zero reduction in fertility rates from one that produces a 3 percent annual population growth rate to one that produces no population growth. The table also indicates how the demographic transition affects national saving rates. Once the transition begins, saving rates in the first few years (not included in the table) fall slightly. Then they rise through year 20 to a value above that in the initial steady state. There follows a decline in saving rates, which reach negative values in the year 110. Between years 110 and 150 the saving rate rises to its ultimate steady-state value of 0.

The initial drop in the saving rate is unrelated to concurrent demographic changes, which in the first few years are unimportant, but is related to anticipated general equilibrium increases in future wages. These projected increases in budget opportunities resulting from capital deepening stimulate higher current consumption and lower current saving. The subsequent increase in the saving rate between years 1 and 20 reflects the drop in fertility. The drop reduces the number of children and the importance of their dissaving, that is, their consumption. By year 20 the fraction of the population between 20 and 60 has increased from 45 percent to 56 percent, and this group is doing more saving because there are fewer children to feed. By year 70, however, the decline in birth rates has affected the size of the young and middle-age adult saving population so that there are relatively more elderly dissavers. This leads temporarily to a slightly negative saving rate. The ultimate steady-state saving rate is 0 because the population growth rate is 0, and savings per capita is constant in the steady state. The final steady-state wage rate is 11 percent larger than the initial steady-state wage, reflecting the permanent capital deepening resulting from the permanent drop in fertility.

Despite the importance of demographic change to savings, as indicated by these two illustrations, there has been limited research on this subject. Research by Barro and Becker (1987, 1988) and several other researchers has just begun to link fertility decisions and long-run savings behavior. Research by Noguchi (1986) has begun to explore the interactions of nonaltruistic bequests and demographic change. And studies such as that in chapter 13 have begun to examine the connections among demographic change, fiscal policy (including social security), and savings. But little or no work has been done on the more subtle influences of demographics on savings. For example, no one has yet considered how changes in family size influence family insurance and thereby the extent of precautionary savings. No one has explored how increases in divorce rates affect saving behavior. And no one has explored how a strong trend toward the elderly living alone, rather than with their children, affects the demand for old age housing and aggregate savings. These and related questions present some fruitful areas for additional research.

Fiscal Policy's Effects on Savings

Chapter 7 presents an extensive review of fiscal policy effects on savings. The main message of this chapter is that fiscal policies can greatly influence savings. This message may be contrasted with an earlier view that fiscal policies, such as reductions in capital income tax rates, have ambiguous saving effects. Chapter 7 points out that many policies, such as a shift in the tax structure, can be decomposed into compensated tax changes plus income effects. The impact of compensated tax changes on savings are unambiguous, and the income effects of shifts in the tax structure often reinforce the effects of the compensated tax changes. The discussion distinguishes four types of policies: changes in the tax structure holding constant the time path of government consumption, tax-financed changes in government consumption, intergenerational redistribution, and intragenerational redistribution. Each of these policies, with the exception of intragenerational redistribution, appears to be capable of generating major changes in national savings.

The fiscal policy literature, since chapter 7 was written, has begun to explore models with uncertainty [e.g., Barsky et al. (1984)] and models in which individuals are liquidity constrained [e.g., Hubbard and Judd (1986)]. Both of these elements can significantly influence fiscal policy effects on savings. Take, as an example, an increase in income tax progressivity. In a certainty model increased tax progressivity would likely reduce savings because of the increased disincentives to work and save. But there is a second reason savings might fall in a world of earnings uncertainty, namely, the increased tax progressivity may improve risk sharing and thereby reduce precautionary savings.

More research along these lines can be expected. In addition, it seems likely that uncertain fiscal policy will, itself, be identified as a reason for additional precautionary savings [e.g., Eaton (1981)]. Understanding the source of changes in fiscal policy may lead to the incorporation of explicit models of government decisions. Chapter 11 presents one such model—a life cycle model in which government decisions reflect the outcome of bargaining between young and old generations. Fully modeling government decisions can lead to some surprising results. In chapter 11 the specification of government decision making implies that debt is neutral even in the life cycle model.

Conclusion and Suggestions for Additional Research

One way to describe the current state of knowledge about savings is that a great deal is known at a theoretical level about savings determinants taken one at a time. Much less is known about the interactions of these determinants, and too little is known at an empirical level about the true causes of savings. As is often the case in economics, the theory seems to have advanced well beyond the empirical research. Unfortunately the theoretical progress has dimmed the prospects for quick empirical resolution of the major outstanding questions concerning savings. The theory has pointed out a number of subtle but powerful determinants of savings on which little data are available. These include the extent of family and government insurance, the degree of intra- and intergenerational altruism, and the nature of subjective probability distributions. New data will be needed to assess these issues, but some of this data will be difficult to collect. For example, it will be a challenge to elicit meaningful statements from workers about their subjective probabilities of getting laid off from their job, of finding a new job if they are laid off, of the salary on their new job, etc. It will also be a challenge to elicit information on who is and who is not altruistically linked; the extent to which family members are currently making transfers to one another is insufficient for assessing altruistic linkages since family members may plan to make such transfers in the future or be willing to make such transfers in emergencies.

Even with the best data, answering some savings questions seems quite difficult. Consider trying to distinguish altruistic family behavior from selfish family risk pooling. In both settings one will observe intrafamily transfers. In both settings one will observe that each family member's consumption depends on collective resources. And in both settings family members will surely state that they care for one another. The selfish risk-sharing family will differ from the altruistic family in that the distribution of consumption across different members will in the selfish model depend on the distribution of resources; for example, members with more wealth will consume more. But in the altruistic model different members may consume more than others because their utility is weighted differently. If these utility weights are correlated with resource position, it may be quite hard to tell whether, for example, the elderly members of the extended family are consuming more because they are old or because they are the ones with most of the wealth.

In sum, the question of what determines savings is like a good jigsaw puzzle. It has a large number of pieces. Some of the pieces have been found.

Those that are available do not immediately fit together, but not all the combinations have yet been tried. The pieces that are missing are not necessarily in the bottom of the box and indeed may be mixed up with identical looking pieces in some other jigsaw puzzle. While it is easiest to keep playing with the pieces at hand, the puzzle may never be solved without the tedious task of looking in the other boxes. While the outlines of the puzzle are getting clearer, the precise picture it displays remains well worth the search.

Note

1. If the Keynesian model specifies that the young consume a fraction θ of their disposable income, including their private transfers, while the old consume a fraction θ of their income, before private transfers, there will be an increase in A_{t+1}, when Z_t increases, by an amount $\theta^2 \Delta Z_t$.