

Coleman, Andrew

Working Paper

Household Savings: A Survey of Recent Microeconomic Theory and Evidence

New Zealand Treasury Working Paper, No. 98/08

Provided in Cooperation with:

The Treasury, New Zealand Government

Suggested Citation: Coleman, Andrew (1998) : Household Savings: A Survey of Recent Microeconomic Theory and Evidence, New Zealand Treasury Working Paper, No. 98/08, New Zealand Government, The Treasury, Wellington

This Version is available at:

<https://hdl.handle.net/10419/205402>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>

TREASURY WORKING PAPER

98/8

Household Savings: A Survey of Recent Microeconomic Theory and Evidence

Andrew Coleman

Abstract

This paper summarises recent theoretical and empirical developments in the vast literature that has examined the microeconomic determinants of household saving. It is designed as a primer to provide a basic understanding of some of the developments in the literature in the last decade. The paper uses the standard intertemporal optimising model as the basic organisational framework, and includes a discussion of precautionary savings, liquidity constraints and bequests. The empirical data is examined in light of this framework. Three key issues related to superannuation provision arise. First, evidence suggests that in most countries households continue to save after retirement. Second, in several countries there is evidence that most people hold very few financial assets at any stage of their lives, and a large number of people hold very few assets, either financial or real. Third, intergenerational transfers appear to provide a motive for the lack of consumption for many of the elderly, particularly the wealthy.

Disclaimer: The views expressed are those of the author(s) and do not necessarily reflect the views of the New Zealand Treasury. The Treasury takes no responsibility for any errors or omissions in, or for the correctness of, the information contained in these working papers.

Introduction

There are many varied reasons why people save. Keynes (1936) produced the following celebrated list.¹

Precaution: To build up a reserve against unforeseen contingencies;

Foresight: To provide for an anticipated future relation between income and the needs of the individual or family different from that which exists in the present;

Calculation: To enjoy interest and appreciation;

Improvement: To enjoy gradually increasing expenditure;

Independence: To enjoy a sense of independence and a power to do things;

Enterprise: To secure a masse de manoeuvre to carry out speculative or business projects;

Pride: To bequeath a fortune;

Avarice: To satisfy pure miserliness, i.e. unreasonable but insistent inhibitions against acts of expenditure as such.

Corresponding to these reasons for saving, Keynes provided a list of motives for consumption: Enjoyment, Shortsightedness, Generosity, Miscalculation, Ostentation, and Extravagance.

This number of items in this list is the source of much empirical, theoretical, and policy-making bother. There are many reasons why an individual might save; and many reasons he or she might not. Worse, an individual might save (or not save) for different reasons at different stages of his or her life. Consequently, it is unreasonable to expect a single theory to be able to explain a large proportion of saving behaviour facts - indeed, with so many reasons, it has been difficult to even establish facts about saving behaviour from the small microeconomic datasets currently available.

Despite this variety of savings reasons, there is a standard theoretical model which explains much household saving behaviour: it is the intertemporal utility optimising agent model. In this model a household chooses current consumption and saving, and an asset portfolio, so as to smooth utility over time. Most theoretical advances of the last decade have taken place within the confines of this framework; for instance, the emphasis on precautionary saving, and liquidity constraints. Somewhat paradoxically, these advances have lessened intuitive understanding about the importance of lifecycle saving within the framework, for it has become apparent that there is a wide range of behaviour that is consistent with individuals following an optimal saving strategy. This point deserves emphasising: *theoretical work in the last decade has upset much traditional intuition about household savings behaviour.*

¹ J M Keynes (1936) The General Theory of Employment, Interest, and Money. London: McMillan and Co Chapter 9.

The most notable area where empirical progress has been made concerns apparent anomalies from the traditional view of the life-cycle hypothesis. It is now clear that a large fraction of elderly people continue to save out of their annuity and investment income in retirement; a mixture of habit formation, bequest behaviour, and precautionary savings (for large medical bills) seems to explain this behaviour. It is also clear that a large number of people at all ages have very few savings, data consistent with both liquidity constrained behaviour and precautionary savings behaviour when individuals have high discount rates. Finally, the importance of transfers and bequests among the wealthy is better documented. These transfers account for a surprising large fraction of the stock of capital, even though they affect a small fraction of the population.

This survey has been written with policy issues in mind. Policy makers tend to be concerned that people do not save enough; that aggregate household savings are insufficient; and that Government policies exacerbate both of these failings. These issues are not directly addressed, but the survey is written in a manner that is intended to be accessible to those with these concerns in mind.

The survey concentrates on the economic reasons for savings. Those items on Keynes's lists that can be broadly considered non-economic — the independence and avarice motives, for instance — are not discussed, nor is the growing literature examining the psychological determinants of ways people choose to save. (For a very brief introduction to this literature, see the Thaler 1994.) Nor is there a discussion on the ways saving incentives affect people; for a review see the discussion in the Fall 1996 Journal of Economic Perspectives, or Bernheim 1996.

The survey is structured as follows. Section 2 presents the standard theoretical model in which uncertainty and capital market imperfections are unimportant. Section 3 is a discussion about the importance of bequest motives. Section 4 examines the theoretical basis of two of the major advances of the last decade, the effects of precautionary savings and liquidity constraints. These are treated together because many of the implications of the model are similar. Section 5 presents some of the empirical data relevant to the discussion, largely from the US but also from other countries. Section 6 acts as a summary, and includes a selected personal view on the issues that are most important to consider when designing a superannuation scheme.

Section 2 Traditional Theory: The Intertemporal Optimising Agent.

The standard theoretical model is that of a household solving an intertemporal utility maximisation problem. The household solves the following problem:

$$\max_{C_t, A_t} E_t \left[U(C_t, C_{t+1}, \dots, C_T; B) \mid Z_t, Z_{t+1}, \dots, Z_T \right] \quad \text{such that} \quad C_{t+s} + A_{t+s+1} = Y_{t+s} + (1 + r_{t+s})A_{t+s}$$

where

- C_{t+s} = consumption at time $t+s$
- Y_{t+s} = income at time $t+s$
- A_{t+s} = assets held at the beginning of period $t+s$
- Z_{t+s} = variables affecting utility of consumption such as household size
- B = a bequest left to younger generations, $B \geq 0$
- r_{t+s} = interest rate at time $t+s$
- U = a utility function
- E_t = statistical expectation operator taken at time t

This framework is very general, and in fact all of the analysis in this paper can be expressed in terms of this equation. In this section we impose the following simplifications.

Assumption 1

Utility is intertemporally additive.

$$U(C_t, C_{t+1}, \dots, C_T \mid Z_t, \dots, Z_T) = \sum_{s=0}^T \beta^s v(C_{t+s})$$

This assumption means that consumption is enjoyed when it takes place, independent of consumption at other times.

The discount rate β indicates the patience of the consumer. Typically β is assumed to be less than one, so that people prefer a marginal dollar spent now to a marginal dollar spent later. Suppose $\beta=0.98$. This means that at the same level of consumption at age 65 and age 25, someone values the marginal dollar of consumption at age 65 44 percent ($= 0.98^{40}$) as much as they do at age 25. The assumption of impatience gives people a bias against savings.

Assumption 2

Capital markets are perfect, and agents can lend or borrow at interest rate r in order to smooth consumption.

Assumption 3

Either the utility function $v(C_t, Z_t)$ is quadratic in C_t , or there is no uncertainty about future income streams. When quadratic utility is assumed, households

are risk neutral. This model is called the certainty equivalence model. It dominated theory from 1978 to 1990.

The Certainty Equivalence Model. (Hall, 1978; Flavin 1981)

For simplicity, let $v(C_t, Z_t) = v(C_t/\alpha(Z_t))$, where $\alpha(Z_t)$ is an equivalence scale for the demographic characteristics of the household.² The solution to the above maximisation problem equates discounted marginal utility across time:

$$\frac{v'(C_t/\alpha(Z_t))}{v'(C_t)} = E_t \left[b(1+r) \frac{v'(C_{t+1}/\alpha(Z_{t+1}))}{v'(C_{t+1})} \right]$$

When utility is quadratic, marginal utility is linear in consumption, and the equation becomes

$$\frac{1}{b(1+r)} \frac{v'(C_{t+1})}{v'(C_t)} = \frac{1}{b(1+r)} \frac{v'(C_t)}{v'(C_t)} + \varepsilon_{t+1}$$

where ε is a random shock, normally an unanticipated change in income.

The power of this equation is its simplicity. In the simplest possible case, where the household demographics are stable so that $\alpha(Z)$ is constant, optimal consumption depends solely on a person's level of impatience and lifetime resources. In each period, the agent plans to equate discounted consumption over time; in making this decision, what is important to the agent is lifetime income, not current income. Optimal consumption follows a random walk if the discount rate is the reciprocal of the interest rate, $\beta = 1/(1+r)$; if consumers are impatient, $\beta < 1/(1+r)$, consumption gradually declines; and if they are patient, $\beta > 1/(1+r)$, it rises.

As the following three graphs indicate, optimal lifetime consumption proves to be very sensitive to the parameter β . The graphs show optimal lifetime consumption, saving, and asset accumulation for three different agents with $\beta = 1/(1.01)$, $\beta = 1/1.02$, and $\beta = 1/1.03$, but all with identical lifetime income and the same interest rate 0.02. All three earn approximately \$500 per year while working, and are entitled to a \$300 pension from age 65. The first, patient, agent starts by saving \$123 in the first year, and eventually amasses \$8240 before spending it in her final years; the second borrows \$84 in the first year, but saves in the middle years accumulating a peak of \$3023; while the third borrows \$348 in the first year - 75 percent of initial income - before starting to pay back. This person pays off the final dollar aged 85, so never has positive savings.

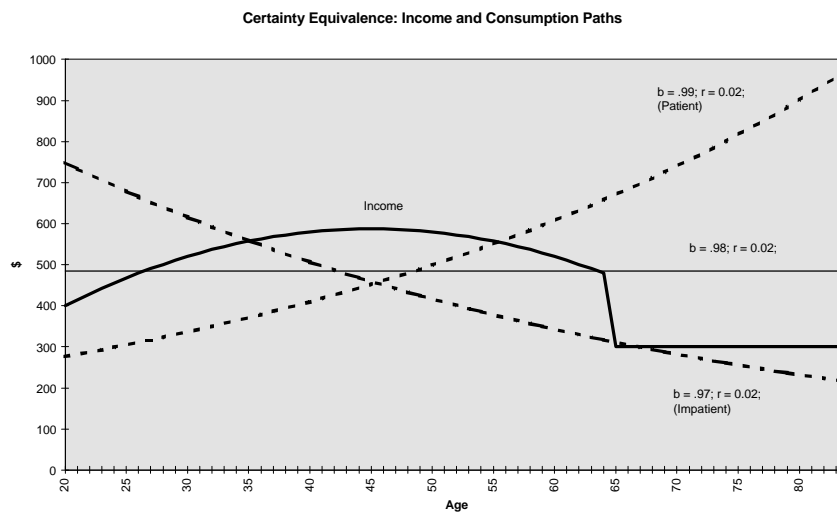
² For instance, for a household with 1 adult, $\alpha(Z) = 1$; with two adults and an infant $\alpha(Z) = 2.2$ (say), etc.

This example is important because it shows how a considerable array of behaviour is consistent with the certainty equivalence model. All three agents optimise their consumption paths, but differ in terms of the way they value present and future income. All three agents have the same zero level of assets at the end of their lifetimes, but accumulate substantially different levels of assets during their lifetimes.


When $\alpha(Z)$ varies over the lifecycle, consumption varies accordingly, rising when there are children and falling as children become independent. This variation is also important, for the consumption paths of agents with and without children are very different. Consider the middle case above, where $\beta = 1/(1+r)$, and contrast it with the stylised example where the individual supports a child from years 30 - 50 of her life, but has no change in income. For concreteness, $\alpha(Z) = 1$ for the adult only, and $\alpha(Z) = 1.5$ when there is a child as well. In the base (adult only) case, consumption is equal every year, at \$484. In the second case, consumption is the same from age 21 - 30 and 51 - 85, but increases by 50 percent in the middle years. The consumption level before the adult has a child is \$409.

The optimal behaviour for an adult planning on a child is to save up for the child, and then run down assets until the child leaves home when the adult is aged 50. From ages 51 to 65 assets are accumulated again, until they are slowly run down in retirement. Peak assets in this case are \$1780, versus \$3020 in the case where there is no child.


Certainty Equivalence Behaviour: No Children



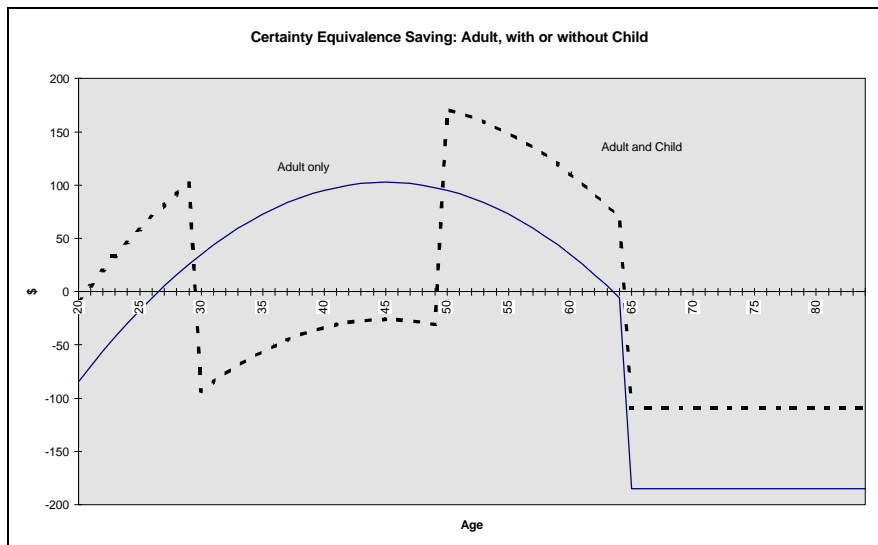
Certainty Equivalence: Saving Behaviour




Certainty Equivalence: Asset Accumulation



Certainty Equivalence Behaviour: With or without Children





Several implications of the certainty equivalence model should be stressed, as they form the basis of the life-cycle model.

- (1) The response to both short term fluctuations in income — say temporary periods of unemployment — and long term fluctuations in income — retirement — is governed by the same principle. People always plan to keep marginal utility constant.
- (2) The *shape* of the lifetime consumption profile is independent of income. It depends solely on an individual's intertemporal substitution choices — whether β is less than, equal to or greater than $1/(1+r)$ — and the demographic cycle.
- (3) The *level* of the consumption profile depends on expected lifetime income and resources. The marginal propensity to consume from current income is the same as the marginal propensity to consume from expected future income; for instance, young people expecting higher income later in life will borrow against this income.
- (4) The marginal propensity to consume from an unanticipated change in income depends on whether or not the income change is seen to be transient or permanent. A short spell of unemployment should lead to little change in consumption; but an increase in salary that is permanent should be consumed immediately.
- (5) Consumption growth is independent of past, current, or predictable changes in income.
- (6) Consumption patterns are independent of the riskiness of income.
- (7) Retired people should dissave, as labour income declines.

The certainty equivalence model has a compelling logic, and one which has appeared to be the basis for most discussion about retirement income policy in many countries. It also provides a large number of testable propositions. These implications have led to a large amount of empirical work testing various implications of the model. Unfortunately, most of the tests have proved to be of low power, and while the evidence is broadly against the model, it has not been conclusively rejected. One of the difficulties, for instance, is obtaining a reliable method to incorporate demographic data into the estimations to account for variation in Z .

The major empirical shortcomings of the model appear to be as follows.

- (1) Consumption seems to be *excessively sensitive* to changes in income: people change consumption more out of transitory income changes than predicted by the model.
- (2) Consumption growth seems to be correlated with predictable income changes, and with previous income changes.
- (3) There is considerable evidence that liquidity constraints matter: people, especially the young, do not borrow against future income. Moreover, changes in laws affecting the availability of credit have real effects, which would not be the case if liquidity constraints did not matter.
- (4) The elderly do not seem to dissave nearly as much as the model predicts.

That said, in the last three years several careful studies have caused some reappraisal of the evidence. In particular, it seems that while most studies report a strong link between income and consumption (which is prima-facie evidence against the certainty equivalence model) it seems that demographic variables that affect consumption are correlated with age and income in a manner that makes consumption appear correlated with income. While this supports the certainty equivalence and life cycle hypotheses, it does so in a manner which suggests that low saving at ages up to 40 - 45 is rational. More generally, it points to the greatest weakness in the current state of theory: the endogeneity of the overall relationship between family structure, family income, family consumption and age.

Most of the advances in the theoretical literature in the last ten years have come from extending the above model to account for liquidity constraints, and the interaction of uncertainty and precautionary savings. These two models are discussed in Section 4. It is worth noting here, however, that many of the insights from the certainty equivalence model do not carry over; the results of the certainty equivalence case are in many ways special.

Section 3: Bequests.

The certainty equivalence model is consistent with agents leaving bequests because they are concerned with the welfare of others (usually their children). There are other reasons why bequests may be left, however, related to imperfections in annuity markets. As discussed below, bequests and transfers may account for around fifty percent of accumulated wealth. As such, it is important to understand why bequests are made, because if a large fraction of wealth is accumulated for reasons of altruism, saving for retirement cannot be the prime motivating factor behind wealth accumulation.

(1) Reasons for Bequests.

There are several reasons why a bequest could be made. These reasons can be divided into two broad classes: unintentional; and intentional.

Unintentional Bequests

In the certainty equivalence model with perfect markets, a selfish individual would only hold annuities in retirement³. Clearly, for an unintentional bequest to occur, there has to be a demand for ordinary (non-annuity type) assets by the elderly; the primary reason is the interaction of uncertainty and imperfect annuity markets. Three examples are usually suggested to explain why a selfish consumer would want to hold ordinary assets in retirement.

- (a) An individual may not wish to purchase a nominal annuity when inflation is uncertain, and the price of real annuities may make holding an ordinary asset more attractive.
- (b) An individual wishing to live in his or her existing home may prefer to own it if the price of home-equity based annuities is very high.
- (c) An individual may demand ordinary assets because of uncertainty about future medical and nursing home expenses, given the inability to cash up an annuity at actuarially fair prices in these circumstances.

In each of these cases, an individual will retain some ordinary assets, leaving open the possibility of an accidental bequest.

Intentional Bequests

Five main reasons are given for intentional bequests.

- (a) Altruism.
- (b) Consumption satiation among the very wealthy.

³ Capital markets do not even need to be perfect. Annuities will be preferred even if their price is not actuarially fair, so long as the annual return from an annuity exceeds the annual return from holding an ordinary asset.

- (c) Payment to children for services rendered (help with shopping, medical problems etc)
- (d) Payment for an implicit annuity or insurance type arrangement, whereby children support the parents in cases of living longer than expected.
- (e) Bribes to ensure continued contact with children.

It is not clear which of these reasons is predominant.

The debate about the extent to which bequests are intentional or unintentional is both important and contentious. Three types of evidence have been used.

First, Hurd (1987) examined the dissaving behaviour of elderly people who had children, and elderly people who did not have children. He found that on average the wealth decumulation and dissaving behaviour was similar, and argued that this fact suggests most bequests are unintentional (because the elderly without children have little motive for bequests.)

Two forms of contrary evidence have been provided. Bernheim (1991), using the same data set as Hurd showed that there was a negative relationship between the level of social security that a household has and the amount of term life insurance held.⁴ He also showed that the demand for term insurance is higher for the elderly who have children than those who do not. Since around 16 percent of the elderly population without children appear to hold term life insurance to leave a bequest, compared to about 25 percent of those with children, Hurd's identification assumption that those without children have no bequest motive appears incorrect.

The third piece of evidence comes from analysis by Gale and Scholz (1994) of inter-vivos transfers — those between live people — and intentional bequests. The data in this case is from the US Survey of Consumer Finances, taken in the mid 1980's, which asks questions about income, wealth, and transfers given and received. The data suggests that the annual flow of bequests is approximately \$105 billion. It is not known whether these bequests are accidental or not. In addition, \$35 billion is transferred by parents as college payments, and a further \$63 billion is made as intended transfers. The latter comprise inter-vivos transfers from the elderly to the young, transfers from trusts, and transfers from term life insurance. Most of these transfers are made from the wealthy to the wealthy, and the average size is large. The large amount of these intended transfers quite clearly shows that there is some bequest motive, so it seems unreasonable to claim that most bequests are unintentional.

⁴ Term life insurance is a way of converting annuity income into bequestable wealth.

(2) The Importance of Bequests to Wealth Accumulation.

The motives behind bequests and transfers are important because the evidence suggests a very large fraction of the accumulated capital stock, perhaps 50 percent, is a result of bequests. This number is subject to much controversy, with estimates ranging from 20 percent to 80 percent. (See Modigliani (1988), Kotlikoff (1988) and Gale and Scholz (1994).) There are two main methods used to calculate this number. First, data on consumption and income by cohort are used to give a saving rate figure by cohort. These saving rates are then aggregated across cohorts. Various assumptions are made along the way; but in general, because consumption closely follows income for most people aged less than 40 - 45, capital accumulated in this manner seems to account for less than 50 percent of wealth. The second method is to calculate the stock of wealth necessary to generate an annual flow of transfers and bequests. Gale and Scholz, suggest that 20 percent of wealth can be attributed to identifiably intentional transfers, a further 12 percent to transfers to pay for college, and 30 percent for bequests.

To quote Deaton (1992: p217)

“All of this spells trouble for the long held belief that life-cycle wealth can account for most of the wealth in the US, and although Kotlikoff and Summers’ (1981) original estimates have not gone unchallenged, it now seems that bequest motives are a good deal more important than we used to think. It is hard to think of issues that are more important than understanding why people are prepared to hold the aggregate wealth stock of the economy.”

(3) Annuity Demand

If the life-cycle hypothesis were true, and the only uncertainty was the time of death, on retirement households would have an incentive to convert wealth into annuities.⁵ In fact, there is very little demand for private annuities, even when they are actuarially fair. (See Hurd 1990, Benheim 1991). Moreover, most private annuities are a result of employment based savings schemes, not a market in which an individual approaches an insurance company, stumps up a sum of money, and asks to purchase an annuity. These facts raise several questions, both about the nature of the annuity market, the demand for ordinary assets, and the demand for bequeathable wealth.


There are two key aspects of annuities that fundamentally shape their demand. First, the market for annuities is not fairly priced. The main problem is adverse selection: people who expect to die young will not wish to hold an annuity, whereas people who expect to live a long time will demand them. The second problem is also caused by information problems: it is very difficult to borrow

⁵ Actually, even this qualified statement is too strong. It would be true for households comprised of singles. It would not necessarily be true for households comprised of couples facing different times of death.

against an annuity income stream, as the income stream is conditional on the life of its holder. This means people with uncertain consumption paths — particularly the possibility of high medical expenses or nursing home care — will want to hold some ordinary assets (or well priced medical insurance) to meet these contingencies.

The following analysis framework is based on Bernheim (1991). An individual has wealth W , which can be split between purchasing annuities and purchasing ordinary assets. Ordinary assets can be left as a bequest, and the demand for both annuities and ordinary assets is assumed to be normal. The return to ordinary assets is β ; and the actuarially fair rate of return on an annuity is $\alpha > \beta$. Annuities can be purchased at rate $\alpha_b < \alpha$, or sold at a rate $\alpha_s > \alpha$. (An annuity is sold when the person purchases term life insurance, which is left as a bequest.) The demand curve and budget constraint is pictured below.


Initial Situation: No Government Annuities



Note the kinked budget curve due to $\alpha_b < \alpha < \alpha_s$

Now suppose the Government withholds an Amount G from the household, and gives them the actuarially fair annuity αG . Their budget constraint is given by the bold line:

Final Situation: Government Annuity Provided



The budget constraint has changed because the household has an actuarially fair annuity, rather than one bought at α_b or sold at α_s . Note that the household can purchase more annuities, or sell some of the Government annuity. However, for any level of provision αG , there will be many households that only hold that level of provision because of the kink in the budget curve. For these households - in fact for all households that want more annuities than level A^* - the provision of the actuarially fair annuity is welfare improving.⁶

This framework helps explain why so many people only have Government annuities. According to the US Social Security Administration, in 1986, 74 percent of households had no private annuities, whereas 92 percent of households had some Government annuity income. For only 13 percent of households did private annuity income comprise at least 20 percent of household income. (Hurd, 1990)

Bernheim used this framework to estimate the demand for private annuities (and term life insurance) as a function of the level of Government annuity income. He finds strong evidence that the higher the level of Government annuity the higher the demand for term insurance and the lower the demand for private annuities. As previously said, he uses this evidence to argue that bequest motives are important.

The fact that annuity demand is so low raises questions about the level of imperfections in the annuity and health insurance markets. It may be the case that adverse selection problems are so bad that individuals are only offered annuities and insurance at such high prices that they prefer to hold alternative assets. In this case, accidental bequests can be considered a preferred option to poorly priced annuities. Whether or not this is the case, the low level of

⁶ That is, there is a welfare improvement from having a compulsory actuarially fair annuity; the way in which G is collected may be welfare reducing.

demand for annuities and the high level of ordinary assets is consistent with precautionary demand theories of savings, as the elderly demand assets in cases of high medical expenses.

Section 4 : Precautionary Savings and Liquidity Constraints.

The certainty equivalence model implies both that consumers are risk neutral and that they can borrow and lend at a fixed interest rate. Empirical evidence suggests that neither of these assumptions are particularly appropriate. For instance:

- (a) in panel data and macroeconomic data, it is apparent that a large proportion of consumers consumed their income each period, rather than consumption smoothing. Hall and Mishkin (1982) estimated this fraction was over 20 percent;
- (b) in a panel data study, Zeldes (1989) showed that the poorest third of households behaved as if they were liquidity constrained, and could not borrow;
- (c) Elderly households hold large number of assets, seemingly in case of large medical expenses.

The precautionary savings literature argues that risk averse agents suffer a greater utility decline from a decline in consumption than they obtain a utility increase from a similarly sized increase in consumption. When this is the case, agents have a preference to hold assets (or borrow less) and have consumption increase over time (as uncertainty is resolved) rather than have a consumption path that is level over time. Mathematically, this type of behaviour is the solution to the Euler equation

$$\frac{\partial v(C_t/a(Z_t))}{\partial C_t} = E_t \left[b(1+r) \frac{\partial v(C_{t+1}/a(Z_{t+1}))}{\partial C_{t+1}} \right]$$

when the marginal utility function $\partial v/\partial C$ is strictly convex. (Skinner 1988, Zeldes 1989, Caballero 1990).

The solution to this problem is rather complicated, and can typically found only by numeric means. The solutions found by these early authors suggested that precautionary savings motives could have a large influence at all stages of life. Students, for instance, would be loathe to get into debt because of concern that they may not earn what they expect later in life; and elderly would hold high levels of assets as insurance against unexpected expenses. In both cases — in fact for all people, no matter their discount rate — the precautionary saving literature implies that people will borrow less or save more early in life than implied by the certainty equivalence model. Moreover, the elderly will dissave less quickly than the certainty equivalence model suggests.

A similar type of solution exists when an individual is liquidity constrained. In essence, the maximisation is solved with the additional constraint that assets have to be non-negative. This means:

$$\frac{f v(C_t/a(Z_t))}{f C_t} \geq E_t \left[b(1+r) \frac{f v(C_{t+1}/a(Z_{t+1}))}{f C_{t+1}} \right]$$

When an individual has positive assets, he or she will equate marginal utility between periods. It may be the case that he or she has no assets, and would like to borrow but is unable to. In this case, consumption this period is less than expected consumption next period, and marginal utility is higher.

The two concepts are closely related. In each case, the agent consumes less now than they would like to either because they cannot borrow (liquidity constraints) or they are scared of borrowing because of concern that they will be unable to pay back in the future without having a substantial decline in consumption. The concepts may also be complementary. Agents are much more likely to face liquidity constraints early in life when they have no assets, or are likely to have no assets, than they are later in life. However later in life, when most agents are not liquidity constrained, they may have a precautionary savings motive because of potential large medical expenses.

It is worth stressing that an agent's behaviour will be affected by *potential* liquidity constraints even if they are not currently liquidity constrained. The fact that they may not be able to borrow for an emergency in the future affects consumption and saving in the present.

The literature made considerable progress when the above problems were first solved in simple dynamic settings in 1991 (Deaton 1991; Carroll 1991). The focus of both problems were impatient agents with a discount rate $\beta < 1/(1+r)$. In the certainty equivalence model, such agents would choose declining consumption over their lifetimes. At the first stages of adulthood, these people would typically borrow, gradually paying off their debts as they grew older before amassing a small pool of savings to eke out their elderly years.⁷ When they are liquidity constrained, however, this consumption path is not available to them.

The solution technique is as follows. Suppose income y_t is randomly distributed with a known distribution. Let A_t equal cash on hand,

$$A_t = A_{t-1}(1+r) + y_t$$

⁷ After all, it is well known that you cannot enjoy the beach at Rio when you are 60 in the same way you can when you are 22.

and call the marginal utility function $U_c(C)$. Optimal consumption will be a function of cash on hand that solves the following equation:

$$U_c(C_t) = \max[U_c(A_t), \mathbf{b}(1+r)E_t(U_c(C_{t+1}))]$$

The first term on the right hand side is the marginal utility of consuming all cash on hand now and taking no assets forward to the next period. We do this if consumption is *lower* than we expect it to be next period, so that marginal utility is higher. The second term on the right hand side is the marginal utility that we get if we take assets into the future. In this case we equalise current marginal utility with discounted marginal utility next period.

This equation is not trivial to solve because consumption next period will be a function of the cash on hand next period, which depends on the decisions made this period (ie the assets taken forward to next period). The simplest way to solve it is to assume that income each period is an independently and identically distributed. In this case,

$$\begin{aligned} U_c(C(A_t)) &= \max[U_c(A_t), \mathbf{b}(1+r)E_t[U_c[C((1+r)(A_t - C(A_t)) + y_{t+1})]]] \\ &= \max[U_c(A_t), \mathbf{b}(1+r) \int [U_c[C((1+r)(A_t - C(A_t)) + y_{t+1})]f(y_{t+1})dy_{t+1}] \end{aligned}$$

This formulation has the same interpretation as above: agents either consume everything now, because expected consumption is higher next period, or they equate expected marginal utility between periods. The wriggle comes as expected marginal utility depends on the number of assets we take into the next period.

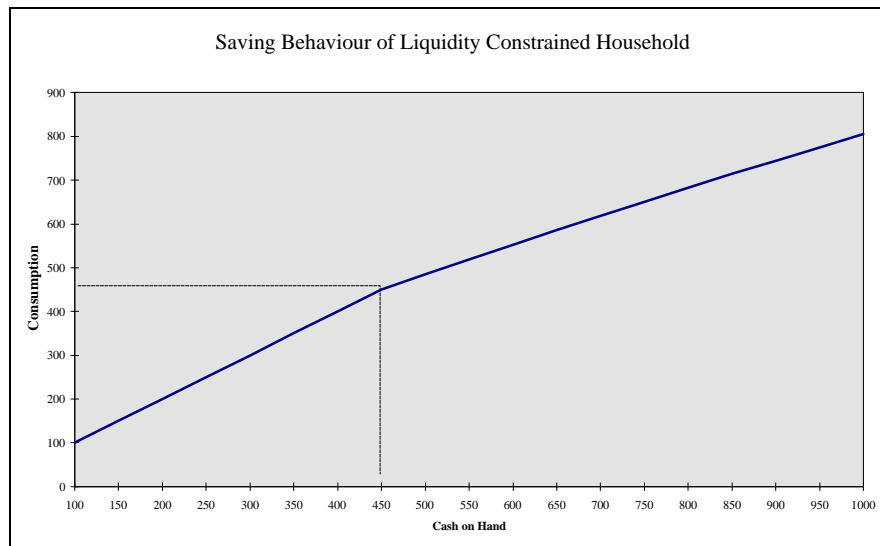
It turns out that these problems can be solved numerically.⁸ The solution is a consumption function $C(A)$ that maximises expected utility: it is a rational expectation solution. Since the first problem was solved, solution techniques have evolved to allow for more complicated income uncertainty; an important series of problems was solved by Hubbard, Skinner and Zeldes (1994, 1995) that had income uncertainty, time of death uncertainty, and health uncertainty.

The solution in this simple case is of considerable interest. There is a threshold of cash-on-hand below which it is all consumed. (In the diagram below, the threshold is \$450.)

⁸ Deaton (1991,1992) describes the process quite well. I have similar programmes available for the truly interested.

Above this level, some of the cash on hand is saved; while the amount is not a linear function of cash-on-hand, it is very close to a linear function. This is important because the optimal saving path in this case is very close to the simple rule of thumb:

- if cash on hand < \$450, consume everything;
- if cash on hand > \$450, save 30 percent of the balance above \$450.




In this model, impatient consumers save a small buffer stock. If they were not liquidity constrained they would have no savings: rather they would borrow to have high consumption. Liquidity constraints mean that they amass a small amount of cash to ease the worst income fluctuations. This type of behaviour seems to closely model the actual behaviour of a large number of people, particularly at young ages.

In the more complicated models solved by Hubbard, Skinner, and Zeldes (1994, 1995), which allow for a government provided social security benefit in old age, a consumer does not even have to be impatient to never accumulate any assets. If a consumer faces either income tested benefits, or asset tested benefits, the incentive to save is destroyed if the benefit level is relatively high compared to income. There are two effects here: first, the provision of a guaranteed income floor reduces the precautionary incentive to save. This provision is typically welfare enhancing, as it means people can consume along a profile closer to what they would if there was no uncertainty. Secondly, the high marginal tax rates implicit in asset testing reduce the incentive to save compared to the level when there is no uncertainty. This second effect would appear to have the welfare costs usually associated with moral hazard problems.


The following simulated diagrams show the optimal consumption, saving, and asset accumulation behaviour of a low income person. Income has the same mean as used in the certainty equivalent case, but in this case it is random until

age 65; thereafter, the person receives a pension of about 50 - 60 percent peak earnings. (In the graph, this is modelled to be random to represent the pension left after uncertain medical expenses.) For this person, the optimal utility path involves very little asset accumulation, particularly early in life. Rather, it is rational to spend almost all of income as it arrives, saving just a small buffer stock.⁹

Simulated Income and Consumption Paths, with Pension



Simulated Saving and Assets Profile, with pension



⁹ These graphs are approximations. I don't have the programmes to generate them myself. However, they are reasonable stylised versions of the real things.

The effects of precautionary savings, uncertainty, and government provided income insurance depend on wealth and income. An essential point to note is the difference in the behaviour of the low lifetime income and high lifetime income agents. High life-time income agents are unlikely to find the consumption floor offered by the Government to be very attractive, and thus will save to accumulate levels of wealth close to those they would have accumulated in the absence of a Government plan. Low income agents will have a disincentive to save, however, particularly if there is an income or asset threshold beyond which they lose entitlement to the benefit.

Summary

Precautionary savings models and liquidity constrained models solve the standard optimising agent model under more realistic assumptions. For impatient consumers, the results are strikingly different than the usual certainty equivalence solution. Optimising consumers will save little for the early years of their life; they will basically try to keep a small supply of assets available to smooth the worse fluctuations in income, and consume everything else. For these people, saving is welfare reducing. If a person has low income, and if a Government pension is available, little saving will take place at any stage in life.

Even for patient people, the precautionary savings literature has different implications than the certainty equivalence model because of uncertainty about the time of death and large medical expenses. In both models, households save throughout life to enjoy interest income, but in the precautionary model people save more. Precautionary households are also much more reluctant to dissave, because of the ever present worry that will need their money later on.

The literature has yet to come to a consensus about the importance of precautionary saving. Carroll and Sandwick (1995) suggest that precautionary saving motives may be responsible for 40 percent of wealth in the US; similar estimates exist for other countries, but other authors using different methodology place the figure as low as 2 percent. Given the reluctance of high lifetime income students to borrow, and the demand for non-annuity assets, the high figures seem more plausible than the low figures; but further research is clearly necessary to get better estimates.

Section 5: Empirical Evidence

There are a variety of surveys on income, consumption, and wealth in most countries that can be used to examine saving behaviour. Three generic warnings must be stressed about all of these surveys.

- (1) Most surveys do not measure saving directly. Saving rates are calculated as the difference of two amounts, each measured with error, so they are prone to error. There is a some measure of internal consistency in the surveys, so the problems are not so acute that the data are worthless; but the results can not be considered conclusive either.
- (2) The surveys tend to under sample the wealthy, the self-employed, and the very poor. They are of most use describing the behaviour of the middle 80 percent. Because they under sample the wealthy, it should be noted that the behaviour of the survey sample is not characteristic of the behaviour of the aggregate economy: typically the wealthiest ten percent own well over half the non-housing wealth of the economy.
- (3) The cross-sectional data collected by the surveys — people of different ages at the same date — cannot be interpreted as the way a particular cohort will behave as it ages. The different lifetime income experiences of people born at different times means that there are considerable differences between cross-sectional data and cohort data.

Most of the survey data has been used to test aspects of the lifecycle hypothesis. In general, I shall not concentrate on the conclusions of these studies, which were briefly mentioned in section 2. Rather, seven aspects of the data will be addressed:

- (1) the median saving rates by age across different countries;
- (2) the difference of saving behaviour across different income groups;
- (3) the saving behaviour of the elderly;
- (4) the importance of housing as the asset most people accumulate;
- (5) the link between consumption and income over the lifecycle;
- (6) the number of people with few assets;
- (7) the number of people who seem to save for retirement.

(1) Cross Section Evidence on Saving Rates.

The table below gives the mean age specific saving rates in various countries during the 1980s.

Age Specific Personal Saving Rates						
Age Group	Canada	Germany	Italy	Japan	UK	US
<30	0.0	9.8	10.0	17.9	5.0	-2.2
30-34	3.0	9.8	20.0	27.4	8.0	7.1
35-39	3.0	10.6	26.0	31.8	12.0	9.4
40-44	5.0	10.2	22.0	31.8	12.0	9.8
45-49	5.0	10.2	23.0	28.5	11.0	11.2
50-54	8.0	10.4	31.0	31.5	10.0	13.9
55-59	11.0	11.0	32.0	34.5	13.0	16.6
60-64	9.0	12.2	34.0	31.7	6.0	8.6
65-69	6.0	9.2	36.0	32.0	2.0	7.1
70-74	6.0	9.7	31.0	33.8	9.0	1.1
75+	8.0	10.2	na	31.1	na	na

James Poterba (ed) *International Comparisons of Household Saving*

Points to note are:

- (1) elderly people save out of income in all countries;
- (2) saving was considerably higher at all age groups in Italy and Japan than in the other countries;
- (3) In all countries, the highest saving rates were recorded for people aged 55-64 (but this is a cross-section fact; these people may have always saved at high rates).

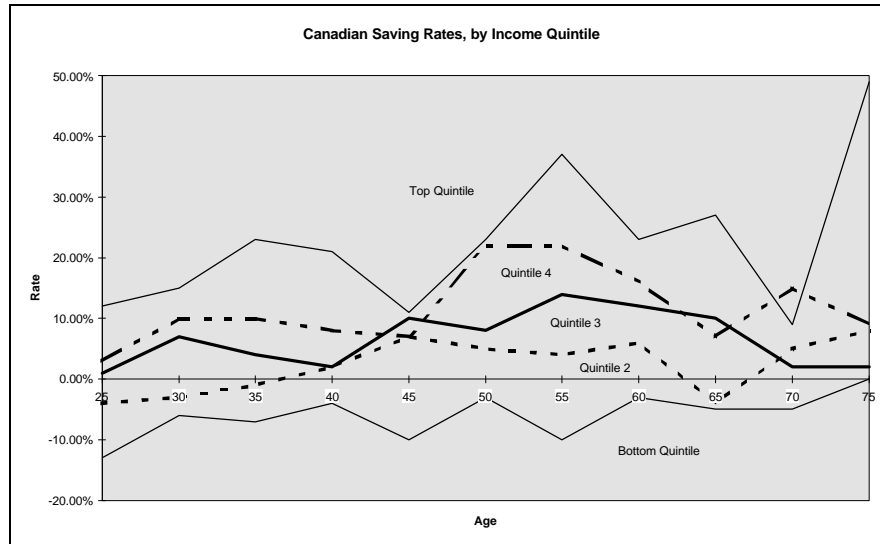
(2) The Difference of Saving Behaviour across Different Income Groups.

In several countries, it is apparent that saving behaviour is determined more than income than by age. Saving rates for high income people are typically similar no matter

age; and saving rates for low income people are typically negative at all ages.¹⁰

¹⁰ One needs to be very careful interpreting this fact, as the permanent income hypothesis makes it quite clear that in a cross section one would expect low income people to dissave, and high income people to save, if low or high income was expected to be transient (Friedman, 1957).

The graph below shows median saving rates by income quintile and age for Canada. It is probably the least distinctive of similar graphs for different countries, but even so shows that saving rises with income (in the cross section).



Much of the US data is classified by the education level as a proxy for lifetime income. Average permanent income in 1984 dollars for those with the lowest educational attainment (“high school drop-outs”) is \$17200; for high school graduates, \$22200; and for college graduates, \$32100. (Hubbard, Skinner, and Zeldes, 1995). It is possible to calculate median saving rates by age group for these three groups. This data, in the table below, clearly shows that saving rises with lifetime income, although again it should be stressed that these are cross-section not cohort figures. The interpretation is that in 1990, at any age people with high lifetime income typically saved a greater percentage of their income than people with low lifetime income.

The lifecycle hypothesis suggests that people with the same discount rate should have similar saving rates no matter the income. It is not clear that this data leads to this interpretation, even if we allow for consistent differences in discount rates between the groups. The data is somewhat consistent with the Hubbard, Skinner and Zeldes (1995) argument that low income groups do not save much because means-tested government social security is relatively attractive, whereas high income groups do save because the low level of social security is not attractive.

US Saving Rates by Education Level			
Age	High School Dropouts	High School Graduates	College Graduates
21	0.7	-18.8	2.2
26	1.6	6.9	14.8
31	-7.2	6.0	11.5
36	-2.0	7.7	16.1
41	-0.2	7.0	19.0
46	-13.4	13.6	14.0
51	8.8	14.4	17.7
56	16.6	21.5	10.6
61	10.1	-2.0	15.2
66	4.2	6.0	7.7
71	1.7	5.3	-10.0

Median rates. Attanasio (1994 p77)

The fact that saving rates do seem to vary so strongly with income raises another possibility: that saving and asset acquisition is a normal good, held for reasons other than those of the life-cycle. This argument is consistent with people having strong bequest motives. Takayama and Kitamaru (1994: p141) in their survey of Japanese saving behaviour say

“ As income grows and wealth accumulation increases, the richer households depart from the pattern of the lifecycle hypothesis and follow what the intergeneration transfer hypothesis would suggest. As society in general becomes wealthier, average households behave more the way richer households do.”

(3) The Saving Behaviour of the Elderly.

With one important qualification, the data in the above tables is consistent with a variety of data that indicates that a large proportion of the elderly save out of their income. The important qualification is that by “income” is meant annuity payments plus other income; annuity payments are of course in part the dissaving of accumulated capital, so it is probable that the elderly are in fact dissaving. What is meant by the above statement is that a large proportion of the elderly continue to increase their non-annuity assets after retirement. This is most notable in Japan and Germany, but it also occurs in Canada, the UK and the US. The following figures give the saving rates in the US by income quartile as an example.

US Saving Rates by Income Quartile				
Age	Quartile 1	Quartile 2	Quartile 3	Quartile 4
66-70	-28.4	4.2	16.9	24.4
71-75	-8.1	-15.8	14.5	21

Median rates. Attanasio (1994 p77)

This data shows that well over 50 percent of US households save out of their income between ages 66 and 75, and that for the upper income households, savings rate are substantial. These rates are not high by international standards: in Japan they exceed 30 percent.

The evidence is possibly less conclusive than it first appears, as Hurd (1986) shows that there was a decline in the quantity of assets held by the elderly over time from 1969 to 1979. This decline may be linked to oscillating capital values, however, as saving data from the 1970s is consistent with a large proportion of elderly households continuing to save. It also should be noted that the difference between cohort and cross-sectional data is acute, as the wealthy tend to live longer, so that in cross-section data average assets invariably increase with age.

The fact that a large proportion of the elderly accumulate non-annuity assets in retirement is consistent with both precautionary saving and bequest motives. The relative importance of these two motives is unclear; but there is other evidence, discussed in section 3, which indicates bequests are important. The German experience is relevant here: saving among retired people is high in Germany, despite the government provision of very good medical insurance.

(4) The Importance of Housing Capital.

In each of the six OECD countries covered in Poterba (1994), the dominant form of asset accumulation was housing. Home ownership rates are high — by age 60, they reach 60 percent in the UK, 70 percent in Canada, 80 percent in the US, and 90 percent in Japan — and for most people of all ages and all income classes form the overwhelming majority of their assets. Consider the following table for Canada, which reports financial and total wealth (primarily housing wealth) by income class and age.

Median Canadian Asset Holdings by Income Quintile						
Age	Quintile 2		Quintile 3		Quintile 5	
	Financial	Total	Financial	Total	Financial	Total
25-29	22	2781	-59	10890	3207	44948
30-34	1035	16211	1132	32319	10088	90433
35-39	647	33419	1849	55829	11624	119720
40-44	1373	40472	3765	60125	15138	142909
45-49	2889	64787	5348	107990	27146	197880
50-54	5220	66394	10424	98395	39264	194045
55-59	5314	66204	15720	112536	37001	200990
60-64	13104	81775	19399	104943	66126	263967
65-69	8154	53246	17947	85124	66504	178567
70-74	6848	28996	11407	78628	74818	171917
75+	5186	35177	12092	47612	65161	158001

1984 data; Burbidge and Davies 1994

The table shows that the second and third income quintiles hold relatively few financial assets at any age, and very few at ages less than 50.¹¹ Nonetheless, wealth is reasonably high, as young Canadians typically save by purchasing a house and paying off the mortgage. In the US the pattern is the same: in fact most households have financial assets equal to only 10 percent of income. In Japan the story is a little different: housing wealth is still very important, but households also have much higher holdings of financial assets.

(5) The link between Consumption and Income.

In most countries, the data strongly suggests that consumption expenditure tracks income. This data is not prima facie evidence against the life-cycle hypothesis, even in its certainty equivalence version, for it is not known whether the changes in income are permanent or temporary. Nonetheless, most econometric studies suggest that consumption is excessively sensitive to income, in that changes in income lead to bigger changes in consumption than would be expected if people were certainty equivalence consumers. It was this sort of evidence that made academics take liquidity constraints and precautionary savings seriously.

Two aspects of this evidence should be mentioned. First, consumption typically falls in retirement as income (including annuity income) falls. People smooth consumption — but they do not completely smooth consumption. Secondly, consumption strongly rises with income for young people.¹² This data is consistent with either liquidity constrained or precautionary young people avoiding borrowing when they would really like to, or with demographic-related

¹¹ The first and fourth income quintiles are similar to the second and third in this case. This data has been omitted for tabular convenience.

¹² The author reserves the right to use “young” to mean people under 40.

expenses rising at the same time as income, increasing the marginal utility of expenditure. This latter topic is the focus of much current research, as it is poorly understood. It is important, for if consumption and income are correlated for people aged less than 40 not because people do not save enough for retirement but because of family formation requirements, mandated saving programmes have the potential to be welfare reducing.

(6) Many people have no assets.

A large fraction of households in all countries have very few assets. This is particularly the case for households under the age of 40, and also the case for low life-time income families. The situation is best documented in the US. By age 40, 40 percent of low education people have net worth less than \$17000 (one year's income); 20 percent have near zero. 40 percent of high school education people have less than \$20 000 net worth, as do 20 percent of college education households. Median financial assets in the US are less than \$3000 at all ages up to 60, and forty percent of retirees have no income from assets.¹³

This data is consistent with the buffer stock saving model as applied to impatient consumers, and in fact is one of the motivations for these models.

(7) Many People Save for Retirement.

To some extent, people tend to save for retirement. For instance, in all countries saving takes place at high rates in the fifteen years before retirement. A large fraction of people in all countries seem to have sufficient assets to be comfortable in retirement. This is particularly true for higher income, better educated people, who seem able to consume at levels similar to those they were used to. To contrast the above figures, in the US the ratio of non-annuity household wealth to annual lifetime income exceeded 6 for 20 percent of low education households and 40 percent of high school and college graduates at the age of 60.

It is not clear whether better educated people save more in line with the life-cycle hypothesis than others because they are better educated and find it easier to calculate the appropriate saving levels, or because they have higher incomes and find it easier to save.¹⁴ There is growing evidence that people save more for retirement when they have specific financial education in the workplace (Bernheim and Garrett (1995); Bayer, Bernheim and Scholz (1996)), or when they are presented with incentives to save specifically for retirement. The literature examining this issue is vast, and is not reviewed here.

¹³ Deaton suggests that up to 25 percent of US households have no bank accounts.

¹⁴ Other possibilities suggest themselves. For instance, people with a natural inclination to save may choose to become better educated.

Section 6: Highlights and Conclusions

(1) How Useful is the Life-cycle Hypothesis?

At the theoretical level, the lifecycle model is a very useful way of organising thoughts about why people save. It seems that many people, particularly the better educated, smooth their consumption over their lifetime, at least in part. To some extent, people seem to save for retirement, particularly once they are aged over 50. The lifecycle model, at least in its precautionary savings mode, is also consistent with the behaviour of many elderly people in the face of uncertain medical expenses, and the behaviour of impatient people adopting buffer stock saving behaviour.

However, it is not clear that the lifecycle model of saving as traditionally presented explains a particularly large fraction of saving, and for this reason it should not be the only model used to think about policy. There are several concerns:

- (1) lifecycle savings probably account for only 50 - 60 percent of a country's wealth;
- (2) precautionary motives may have a greater influence on the pattern of saving than strict certainty equivalence motives;
- (3) bequest motives seem very important;
- (4) to some extent, saving seems to be a normal good (an increasing function of income), held for its own sake, possibly for bequests, possibly for independence motives; and,
- (5) for most people, asset accumulation takes place through home ownership, behaviour which is consistent with many other motives. For instance, a desire for economic independence may be the underlying motivating force behind saving, and this may be best expressed through land ownership.

These factors should be taken seriously. It is not clear that the life-cycle hypothesis should be the dominant paradigm by which we think about saving, although it is clearly a useful paradigm that is relevant to various stages of life of many people. Unfortunately the data are not good enough to provide firm grounds for distinguishing between competing theories. Personally, I believe that many households, particularly those with higher incomes, have a dynasty view of the lifecycle, and are aiming to start or build upon an asset portfolio that may be left to enhance the lives of future generations; while this is consistent with a very long term life-cycle model, it has quite different implications to those usually considered, and the behaviour of such people can be expected to be different than those who mainly try to smooth consumption over their own lifetimes.

(2) Annuities, Bequests, and the Saving of the Elderly.

It is quite apparent that the elderly demand a mixed portfolio of assets: houses, annuities, bequeathable wealth (term insurance), financial assets. For most people the government provision of actuarially fair annuity income is a good deal, because it changes their budget constraint in a manner that is welfare enhancing simply because annuities can be purchased at actuarially fair rates. Even if the level of Government provision is not optimal, and the household uses some of the annuity income to purchase other assets (ie saving out of annuity income, or purchasing life insurance), welfare may still be enhanced compared to the situation whereby only privately priced annuities are available.

These two facts — that households demand a mixture of assets, and that adverse selection problems affect the private provision of annuities — should influence the design of any Government retirement income scheme. It is clearly possible for the Government to target a level of annuity income that is too high for most preferences; consequently the scheme should be designed to enable those who wish to choose a different portfolio to undo the Government provided portfolio at reasonable cost. The key design issue is in the phrase “at reasonable cost”. For instance, I think it is reasonable for the Government to insist on purchasing a moderate level of annuities at an actuarially fair price, and letting households undo it at private market prices; but I am not so convinced that it is reasonable to insist the same level be purchased at private actuarial prices and then resold at a different set of private actuarial prices. Alternately, the Government could specify a graduated scale whereby people choose an age at which to become eligible for a pension, and either get a different (actuarially determined) level of payment, or have to provide a different level of cash for the same level of payment. Such variations may be worth considering.

(3) Liquidity Constraints and Precautionary Savings.

Precautionary savings are important. It is reasonable to believe that many elderly people hold non-annuity assets for precautionary reasons; and that many young people avoid borrowing either because of precautionary reasons or because of capital market imperfections. It is unfortunate that there are not good estimates of the relative importance of precautionary savings and lifecycle savings in national wealth in different countries, because this could be one of the determinants of differences in national savings.

Precautionary savings can be considered a substitute for insurance, and so it can be expected that the demand for precautionary savings is affected by the availability and cost of insurance alternatives. To state the obvious, government provision of consumption insurance reduces the demand for precautionary savings. The welfare implications of such interventions have not typically been fully explored; the studies of the interaction of means and asset tested benefits on the saving of low income people in the US clearly show the difficulties of

designing consumption insurance that does not provide substantial moral hazard problems.

(4) Home Ownership

Most people save by purchasing a house and repaying the loan. It is not clear why this is so. Houses appear to have some tax advantages (particularly if there is no property tax) as imputed rent is not taxed; but they may be desirable because of economic independence motives, or simply because people like to shape their surroundings in their own image, and this is not possible for tenants. Whatever the cause, any retirement policy should be shaped in the knowledge that most people have a strong preference to purchase houses, and policies that forcibly change households portfolio allocation away from houses are likely to be welfare reducing.

Bibliography

Abel, Andrew B. (1990) Asset Prices under Habit Formation and Catching up with the Jones. American Economic Review 80(2) 38 - 42

Attanasio Orazio P (1994) "Personal Saving in the United States" pp 57-123 in James Poterba (ed) (1994) International Comparisons of Household Saving (Chicago and London: The University of Chicago Press).

Banks, James, and Richard Blundell (1994) "Household Saving Behaviour in the United Kingdom" pp 169 - 205 in James Poterba (ed) (1994) International Comparisons of Household Saving (Chicago and London: The University of Chicago Press).

Bayer, Patrick, B Douglas Bernheim and John Karl Scholz (1996) The Effects of Financial Education in the Workplace: Evidence from a survey of Employers. NBER Working Paper 5655

Bernheim, B Douglas (1991) How Strong are Bequest Motives? Evidence based on Estimates of the Demand for Life Insurance and Annuities. Journal of Political Economy 99(5) 899 - 927

Bernheim, Douglas B and Daniel Garrett (1996) The Determinants and Consequences of Financial Education in the Workplace: Evidence from a Survey of Households. NBER Working Paper 5667

Browning, Martin, and Annamaria Lusardi (1996) Household Saving: Micro Theories and Micro Facts. Journal of Economic Literature 34 1797-1855

Burbidge, John, and James B Davies (1994) "Household Data on Saving Behaviour in Canada", pp 11-56 in James Poterba (ed) (1994) International Comparisons of Household Saving (Chicago and London: The University of Chicago Press).

Caballero Ricardo J (1990) Consumption Puzzles and Precautionary Savings Journal of Monetary Economics 25(1) 113 - 136

Carroll, CA (1991) Buffer Stock Saving and the Permanent Income Hypothesis mimeo, Board of Governors of the Federal Reserve System

Carroll, CA and Andrew Sandwick (1995) How Important is Precautionary Saving? NBER working paper 5194.

Deaton, Angus. (1991) Saving and Liquidity Constraints Econometrica 59 1121 - 1142.

Deaton, Angus. (1992) Understanding Consumption (Oxford: Oxford University Press.)

- Flavin, Marjorie (1981) The Adjustment of Consumption to Changing Expectations about Future Income. Journal of Political Economy 89(5) 974-1009
- Friedman, Milton (1957) A Theory of the Consumption Function (Princeton: Princeton University Press)
- Gale, WG and JK Scholz (1994) Intergenerational Transfers and the Accumulation of Wealth Journal of Economic Perspectives 8(4) 145 - 160
- Hall, Robert E (1978) Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. Journal of Political Economy 86(6) 971-987.
- Hall, RE and Frederic Mishkin (1982) The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households. Econometrica 50(2) 461-481.
- Hubbard, R Glenn, Jonathon Skinner, and Stephen Zeldes (1994) Expanding the Life Cycle Model: Precautionary Saving and Public Policy American Economic Review 84(2) 174-179
- Hubbard, R Glenn, Jonathon Skinner, and Stephen Zeldes. (1975) Precautionary Savings and Social Insurance. Journal of Political Economy 103(2) 360 - 399
- Hurd, Michael (1987) Savings of the Elderly and Desired Bequests. American Economic Review 77(3) 298 - 312
- Hurd, Michael (1990) Research on the Elderly: Economic Status, Retirement, and Consumption and Saving. Journal of Economic Literature 28 565-637
- Kotlikoff, Laurence J (1988) Intergenerational Transfers and Savings Journal of Economic Perspectives 2(2) 41-58
- J M Keynes (1936) The General Theory of Employment, Interest, and Money. London: McMillan and Co
- Poterba, James ed (1994) International Comparisons of Household Saving (Chicago and London: The University of Chicago Press).
- Skinner, Jonathan (1988) Risky Income, Life Cycle Consumption, and Precautionary Savings. Journal of Monetary Economics 22 237 - 255
- Takayama, Noriyuki and Kitamura Yukinobu (1994) "Household Saving Behaviour in Japan" pp125 - 167 in James Poterba (ed) (1994) International Comparisons of Household Saving (Chicago and London: The University of Chicago Press).

Thaler, Richard. (1994) Psychology and Saving Policies American Economic Review 84(2) 186-192

Zeldes, Stephen P (1989) Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence. Quarterly Journal of Economics 275 - 298