A review of the effects of taxes on consumption and saving*

Preliminary and Incomplete.

Not for citation.

Orazio P. Attanasio†

and

Matthew Wakefield†

* This is a preliminary draft of a paper being prepared as a 'review chapter' for: "Reforming the tax system for the 21st Century: The Mirrlees Review" (http://www.ifs.org.uk/mirrleesreview/).

† University College, London, and The Institute for Fiscal Studies.

1. Introduction

Set context by discussing possible reasons for reducing tax on return to saving:

- Theoretical presumption for expenditure tax / to avoid double tax (cf. Meade).
- To encourage individuals who are not saving enough, to save more.
- Because want individuals to hold certain otherwise unattractive assets (eg. pensions)

UK setup probably motivated by all three. The first is focus of Banks, Diamond & Mirrlees chapter, will not repeat arguments here (although some of the evidence we survey may be relevant). What we say will have more or less direct implications about whether policies can effectively achieve last two aims.

This paper is intended to review what we know about how responsive consumption and savings choices are to tax incentives (interest rate changes). We will describe the economic framework that is useful for analysing these issues, highlighting particularly what this framework indicates about which factors might be particularly important in shaping how choices respond to the rate of return. We will then attempt to summarise what current empirical knowledge tells us about the actual relative importance of these different factors.

The structure of the paper is as follows. Section 2 describes some features of the UK institutional setup; this description is useful both to provide context for our discussion, and because knowledge of some of these institutions is useful for the discussion in later sections of the paper. Section 3 …

2. Brief description of UK institutions

The past 25 years have seen multiple reforms to the tax treatment of different assets in the U.K. While these reforms have often been motivated by different goals, certain trends can now be discerned. For example, in section [X.X] of their paper on *Taxation in the U.K.*, Adam, Browne and Heady describe how these reforms have contributed to a 'significant reduction in the extent to which the tax system distorts the return on different savings vehicles' [p. reference here]. In this section we highlight some of the patterns that we think are important. The selective discussion of reforms and institutions is intended both to provide some context for the discussion of the chapter and to familiarise readers with certain elements of the U.K. tax system that will be relevant later in the chapter. The section focuses exclusively on financial assets, though Adam et al

(ibid) do provide some details about the tax treatment of housing and other physical assets.

One broad direction of reform to the tax treatment of financial assets over the past 25 years has been a trend to exempt the return on savings from taxation. Approved Personal Pensions were, for example, introduced in 1988. These are individual accounts in which working age adults can save for retirement. The accounts are 'tax deferred' in the sense that tax is only paid once an income stream is drawn from the fund in retirement (and even then one-quarter of the fund may be taken as a tax free lump sum): money paid into the accounts is paid out of income before any tax or National Insurance Contributions (NICs[1]), and the returns on the assets held in a pension fund are exempt from tax. The advent of personal pensions meant that the opportunity to save for retirement without facing tax on the return to those savings was extended to workers who were not covered by the pre-existing employer-provided occupational pensions.

There has also been a move away from the taxation of the return on more accessible forms of saving, first through the introduction of Personal Equity Plans (PEPs, in 1987 [check date!]) and Tax Exempt Special Saving Accounts (TESSAs, in 1991), and subsequently through Individual Savings Accounts (ISAs) which superseded PEPs and TESSAs in April 1999. A PEP was a savings vehicle within which funds had to be held in equities either directly or in trust. Contributions into PEPs were paid from net income, but the accounts were tax privileged since any interest income or capital gains accruing to the fund was tax exempt.[2] TESSAs were a form of bank or building society cash deposit account. Contributions to TESSAs were also paid from net income and these accounts provided tax relief for interest income accruing to funds held in the designated account as long as the capital remained untouched for five years.[3] The ISA is a tax privileged savings vehicle for cash deposits, or for holdings of stocks and shares either directly or in trust, or for both cash and equities. Like both TESSAs and PEPs, contributions to ISAs are paid from net income and returns on funds are not taxed. The absence of a statutorily fixed holding period is the main difference between a cash ISA and a TESSA. The option of holding cash or safe interest bearing accounts and so avoiding stock market risk, is what differentiates the product from a PEP.

[1] NICs is the name for social insurance contributions in the UK, although the contributions look increasingly like a payroll tax rather than an insurance premium.
[2] PEPs also attracted a "dividend tax credit", as subsequently (until 2004 – [check date!]) did ISAs. For an example explaining the relative generosity of the credit in the two policies, see Banks and Tanner, 1999, pp. 91-2.
[3] Early withdrawals would pay back the tax advantage but attracted no further penalty.

Although ISAs, like TESSAs and PEPs before them, are restricted in terms of the total amount that could be invested in an account in any one year, the limits imposed are high relative to the holdings of accessible financial assets of most households in the U.K. [cite figure from Banks, Oldfield & Wakefield, 2002]. Thus ISAs potentially allow many – probably the majority – of households to invest all of their accessible savings in a form that is free from a tax on the return to savings.[4] Personal and Occupational Pensions imply a similar situation is true for retirement saving. It is beyond the scope of this paper to assess the extent to which this implies that households in the U.K. are faced by an expenditure tax system (though see [cross refer Adam et al and Banks, Diamond and Mirrlees chapters if relevant]), but it is worth noting that pensions and ISAs still face differential tax treatment, with pension saving being tax preferred relative to ISA saving. The most obvious reason why this is true is the tax-free lump sum mentioned above. Additionally, employer contributions to pensions are particularly tax favoured because they are not subject to either employee or employer National Insurance, and pensions may also offer opportunities to 'tax-rate smooth' whereby contributions to the pension while working are exempted from a higher marginal tax rate than is eventually paid on income received in retirement.

These tax advantages for pensions may, at least in part, be thought of as a means of ensuring that individuals will save in pensions even though they represent a relatively inflexible savings vehicle in which funds must be locked away until a certain age and can then only be used in a restricted way.[5] Differences in tax treatment between assets, and particularly opportunities to tax rate smooth, may also mean that when making portfolio decisions individuals face quite complicated choices about which will be the best part of the lifetime in which to build up any pension funds (cf. Blundell, Emmerson and Wakefield, 2006). These incentives may have become more important with the recent (April 2006) increase in the annual limits on pension saving which have increased the flexibility individuals have about the time structure of any contributions to (personal) pensions.

As well as the general trend away from the taxation of the return to financial assets, certain other more specific trends have emerged during the last ten or so years. One has

[4] While this statement is certainly true if we consider personal taxes, the 2004 [check year!] abolition of the dividend tax credit for stocks and shares held in ISAs means it may not be considered quite accurate once corporate taxes are taken into account [check interpretation here!]. A similar tax credit for pension funds was abolished in 1997 [check year!].

[5] Essentially, the funds must be turned into a retirement income stream. [Check – rules have changed q recently].

been an increasing focus on whether individuals' are accumulating 'adequate' (retirement) savings. So, for example, publications from the DWP ([2002 Green Paper]) and from a recent government appointed 'Pensions Commission' ([1st report of pensions commission]) have presented evidence on how many individuals might be at risk of having inadequate resources in retirement.[6] Interestingly in both cases the definition of adequacy focused on whether individuals look set to achieve a certain income replacement rate once they retire and start drawing a pension. One way of viewing such a benchmark is that it concerns whether people are going to be adequately able to smooth their consumption (or standard of living) into retirement, and so the definition seems closely related to what the lifecycle model (reviewed in the next section) tells us is the aim of saving (for retirement).

The focus on whether individuals are doing 'enough' saving was evident even before 2002, in reforms that focused on specific groups of the population where there were perceived gaps in participation in the savings market. This was perhaps most clear with the introduction of Stakeholder Pensions in April 2001. Stakeholder Pensions were a new type of no frills personal pension, and they were, at least rhetorically, targeted at 'middle earners' earning between around £10,000 and £20,000 a year. This was a group in which take-up of pensions was perceived to be low [add govt quote here].[7] The replacement of PEPs and TESSAs with ISAs could also be seen as a policy that aimed to expand saving further down the income distribution: by being more flexible than their predecessors, it was hoped that ISAs would be more appealing to lower income households. [Add something on Saving Gateway here?].

It is not only through creating certain savings vehicles that Government policy can affect incentives to save. For lower income individuals the incentive to save for retirement may also be affected by government transfer payments if means-tested retirement benefits mean that small retirement incomes are effectively taxed away through benefit withdrawal. This is another area in which reforms since 1997 – notably the introduction of a taper on the main means tested benefit for pensioners in the UK which is now called the Pension Credit – have aimed to alter incentives to save for individuals lower down

[6] [DWP thought 2-8 million; Pensions Commission thought 12million.]

[7] Although even when the policy was announced some argued that take-up in the target group was not especially low ([Disney, Emmerson and Tanner, 1999]). Recent evidence suggests that if it affected anyone, the stakeholder pension reform package may have affected the pension saving of a lower income group, and individuals in this group may have been responding to a tax reform that happened at the same time that Stakeholder Pensions were introduced ([Disney, Emmerson and Wakefield, XXXX]).

the income distribution. The change in incentives to save was though ambiguous at different points in the income distribution ([Disney and Emmerson, FS, and Sefton, van de Ven and Weale, forthcoming]). This is a clear example of the tradeoff between providing clear incentives and achieving redistributive goals through the pension system, a topic to which we will return.

In the latest set of proposed pension reforms [pensions white paper, 2006, building on main pensions commission report] the aim of targeting certain groups where there might be a particularly high risk of 'inadequate' saving has been supplemented by a reform which, it is hoped, might target specifically those who are not currently saving in pensions. The idea is to switch from a system in which individuals have to actively choose to 'opt in' to pension saving, to one in which the default is to be opted in (with a given level of employer and employee contributions and tax relief) to the pension. The hope is that this might get more people saving in pensions simply due to 'inertia' in behaviour whereby individuals do not tend to deviate from the current (or default) option. While the idea of saving 'adequacy' that this policy aims to address may well be closely related to the consumption smoothing behaviour which results in the lifecycle model, the idea of 'inertial' behaviour may well lie outside at least standard versions of this model. The next sections of this paper look at what is and is not known about the lifecycle model and how well it can describe behaviour.

[Add something on tax treatment of annuities?? Is it a bit anomalous??]

3. The life cycle model and the response of saving to taxation.

In this section we provide a simple theoretical framework that allows us to organize the discussion of the effect of taxation on saving and consumption. The behavioural model we use, the so-called life cycle model of Modigliani and Brumberg (1954), is not meant to be a literal description of reality but can capture some important features of individual behaviour. It also constitutes a useful benchmark against which we can contrast the implications of more complex and possibly more realistic models.

We start with a discussion of a simple version of the life cycle model, we then use it to simulate the behaviour of hypothetical individuals faced with different incentives to save. This exercise will give us the theoretical effect that one would observe as individual responses to changes in the taxation of saving if individuals behaved according to such a simple model. We conclude this section with a mention of alternative theoretical

frameworks that have recently been proposed to analyze individual behaviour. One point we will make is that the life cycle model is a potentially very rich structure that can encompass and explain different types of behaviour. Different effects will be obtained with different specifications of the life cycle model. The issue of which version of the model is relevant is ultimately therefore an empirical question. In the section 4 we discuss the available empirical evidence.

3.1. The life cycle model.

The life cycle model is a workhorse of modern macroeconomics and public finance. Proposed by Modigliani and Brumberg (1954) in a seminal contribution, it shares with the Permanent Income Hypothesis of Friedman (1957) the idea that individuals allocate resources over time (and therefore determine current consumption and saving – or, equivalently, present and future consumption) taking into account the resources available over a long period of time and the possibility of moving them over time. The latter might mean saving to move resources from the present to the future or borrowing to move resources from the future to the present. The attractiveness of the life cycle model and of the permanent income hypothesis is that they treat the allocation of consumption over time (saving) in a manner analogous to that of the allocation of a certain amount of expenditure among several commodities. This powerful intuition is behind some of the optimal taxation results that Banks and Diamond discuss in their chapter. Unlike the permanent income model, which is usually thought of as a model with an infinite horizon, the life cycle model focuses explicitly on a finite horizon and is therefore particularly useful to study retirement saving and related issues.

Although the model has been around for some time, it is worth being precise about what we mean by it. We interpret the life cycle model as one in which an individual maximizes life time utility derived from consumption using her life time resources and taking into account the available inter-temporal trades. At this level of generality, the model is not very useful. For it to have some bearing on reality we need to make a number of assumptions about the details of the various components of the model. In the basic version that we will be simulating we assume that individuals maximize expected utility and that this is intertemporally separable in the sense that consumption in a given year does not affect utility in future period. Future utility is discounted at a geometric rate. The income process is modelled as being made of a deterministic component, assumed to be known perfectly by the individual, and a stochastic component. Our individuals are

assumed to be aware of the process that generates it. The fact that income is exogenous implicitly implies that we are not considering the process that generates labour supply. This can be justified if one is willing to assume that consumption and leisure are additively separable. After a certain age, individuals are assumed to retire and to be entitled to benefits that are a fixed proportion of the last earnings of their working life.

We assume that life has a fixed duration which is known to individuals and that there is no bequest motive. Moreover, individuals are not allowed to die in debt. In terms of the inter-temporal trades available to the individual, we assume that individuals have a single asset that pays a fixed interest rate to move resources in the future. We can also restrict the amount individuals can borrow.[8]

In the simplest version of the model, we assume that the household whose consumption we study is made of a single individual. We then introduce more realistic settings where household size, and therefore the needs individual households have, vary deterministically with time.

Some of the assumptions we make are very stark and unrealistic. However, many of them can be relaxed in a relatively simple fashion and/or have minimum effect on the main results we obtain. A good example of such an assumption is that of certain life length. One can introduce the possibility that individual dies with a certain probability (possibly increasing in age) without much complication. Other assumptions, however, are crucial for the results we will be discussing and cannot be relaxed without changing the nature of the model and the implied results substantially. One of these examples is that of inter-temporal separability. This assumption implies that consumption at a certain date does not affect the marginal utility derived from consumption at different dates. We therefore rule out both durability and habit formation. We will briefly discuss models with habits below.

Some of the basic insights of the life cycle model are well known and do not need to be discussed here in detail. Consumption at any point in time depends not only on current income, but on the total amount of resources available to an individual over her life cycle. Consumption should react to the interest rate, although, at least theoretically, the effect of changes in interest rates on consumption is ambiguous. As is well known, 'substitution effects' call for a decrease in current consumption because the price at

[8] Under certain assumptions on preferences, individuals will never want to borrow. This occurs when the marginal utility of consumption goes to infinity at very low level of consumption. Individuals will not want to borrow if this induces a positive probability of having extremely low levels of consumption.

which consumption can be transferred into the future has been decreased, but this might be counteracted by income effects since a given target level of future consumption can be achieved with less saving when the interest rate is higher. As noted by Summers (1981), wealth effects concerning the amount that expected future incomes are discounted, tend to reinforce substitution effects and also call for a decrease in consumption or increase in saving. Ultimately, which of these forces dominates depends on preference parameters and is, therefore, an empirical issue.

A parameter that is crucial to determine the size and even the sign of the effect is the elasticity of intertemporal substitution, which may be thought of as measuring the value that individuals place on achieving a smooth standard of living[9] across the lifecycle.[10] If this elasticity is small, income effects could be relatively important. Obviously such a parameter is crucial if one wants to determine the size of the effect of a change in the taxation of saving on current consumption and saving. A remarkable feature of the sizeable recent literature on the effect of preferential taxation on retirement saving on personal and national saving is that it never refers to the literature that has studied the life cycle model and estimated preference parameters, including the EIS.

One possible reason for this neglect might be the feeling that the life cycle model constitutes a poor approximation of reality. An alternative reason for the disjuncture of the two literatures, however, could be that modern versions of the life cycle model that explicitly incorporate uncertainty, for the most part do not deliver a closed form solution for consumption. That is, one can write down optimality conditions that can be used in the empirical analysis, but cannot use them to solve for an expression for consumption as a function of, say, present and future income, current wealth and the level of current and future (expected) interest rates. This lack of functions that could be used for policy analysis has led public finance economists to neglect the consumption literature. One of the things we do in this chapter is to start filling this gap by providing some simple comparative statics exercises.

The computations we provide have to be obtained by numerical methods. These, however, have become increasingly standard and affordable, both because of the increased power of computers and because of a small literature that, starting with Deaton (1991) contribution, has looked at the solution of life cycle models of increasing degree

[9] Or, more properly, marginal utility of consumption.
[10] In our set up the same parameter is also the reciprocal of the coefficient of relative risk aversion.

of complexity and realism. Much of the existing literature has been devoted to the plausibility of the standard LC-PI and of slightly different models, such as the so called buffer stock model. Here what we do is to simulate some versions of the standard model to show two things. First, we want to see how the responsiveness of savings to the interest rate varies with preference parameters and other features of the model, including the institutional set up in which the behaviour of our consumers is imbedded. Second, we want to compute the elasticities implied by parameters estimated from micro data that do not reject some realistic versions of the model.

3.1.1. Models details

In Table 3.1, we report the features of the baseline model and of the alternatives we simulate. Key parameters in the simulations include the EIS which, as we already mentioned, is set to 1 for our baseline results. The discount rate which is set at 2.5%, which is equal to the larger of the two real interest rates for which we run simulations.

The process generating income is also important to our results. There is a deterministic component of income which is hump-shaped during working life and assumed to be known perfectly by the individual. The hump shape is such that income is expected to grow at approximately 4% per year at the beginning of the working life, and the income growth then gradually declines until the expected to income peak of approximately twice initial income is reached at age 55. There is also a stochastic component of income and individuals are assumed to be aware of the process that generates this. It is made up of a transitory (iid) shock, and a permanent (random walk) component, with variances of shocks to these components as given in the table. Finally, income in retirement is certain, and in our baseline run is equal to half the final income draw of the working life.

Table 3.1 Model Properties

Parameter	Baseline	Variants
Utility Parameters		
EIS	1	½, ¼
Discount rate	2.5%	
Income process		
Initial income growth rate (expected)	4.2%	0
Age at peak expected income	55	n/a
Ratio of peak income to initial income	2.04	1
Variance of permanent shock (to log income)	0.030	
Variance of transitory shock (to log income)	0.015	
Replacement rate in retirement	0.5	0.25, 0.75
Retirement age	65	
Others		
Real interest rate	2%	2.5%
Lifespan	60 yrs (21 – 80)	

3.1.2. Model's results

Probably the best way to introduce our results is to present some pictures of life cycle profiles for consumption and wealth under different scenarios. In Figure 3.1, we plot the life cycle profile for consumption in our baseline run. This picture, as the others reported below, is obtained simulating the life cycle of 10,000 individuals who receive a particular realization of the assumed income shocks[11] and averaging the relevant variable of interest. The figures report plots of these averages. Initially, the average profile for consumption is upward sloping because the individuals are relatively impatient and, given their income profiles, they would prefer to move resources from the future to the present. However, they cannot borrow so that in the early years of the life cycle consumption tracks income. At some point in the late 40s income becomes sufficiently high relative to current resources that individuals can achieve the desired consumption levels without borrowing, and instead start saving for retirement, as well as for precautionary reasons. Approaching retirement, individuals are not constrained by the inability to borrow. Nonetheless, there is a change in the consumption growth rate at the retirement date (barely visible in the chart) because, in this simple set up all uncertainty ceases.

[11] The shocks are assumed to be independent across individuals.

Figure 3.1 Profile of average consumption: baseline case

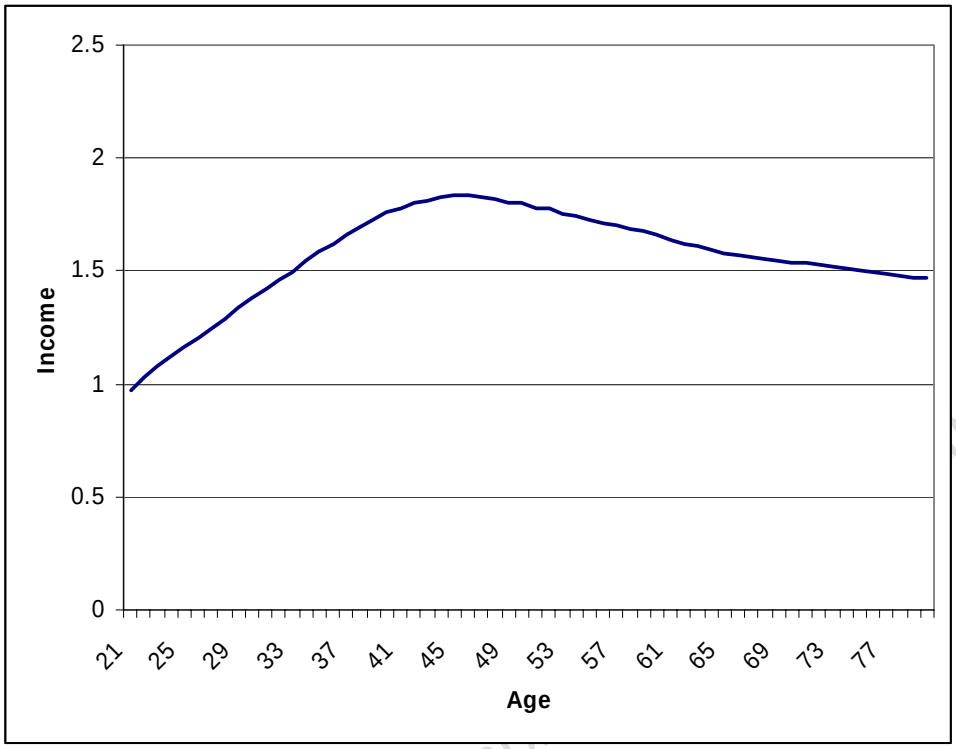

Note: Income is normalised by dividing through by expected income at age 21.

Figure 3.2 Profile of average asset holdings: baseline case

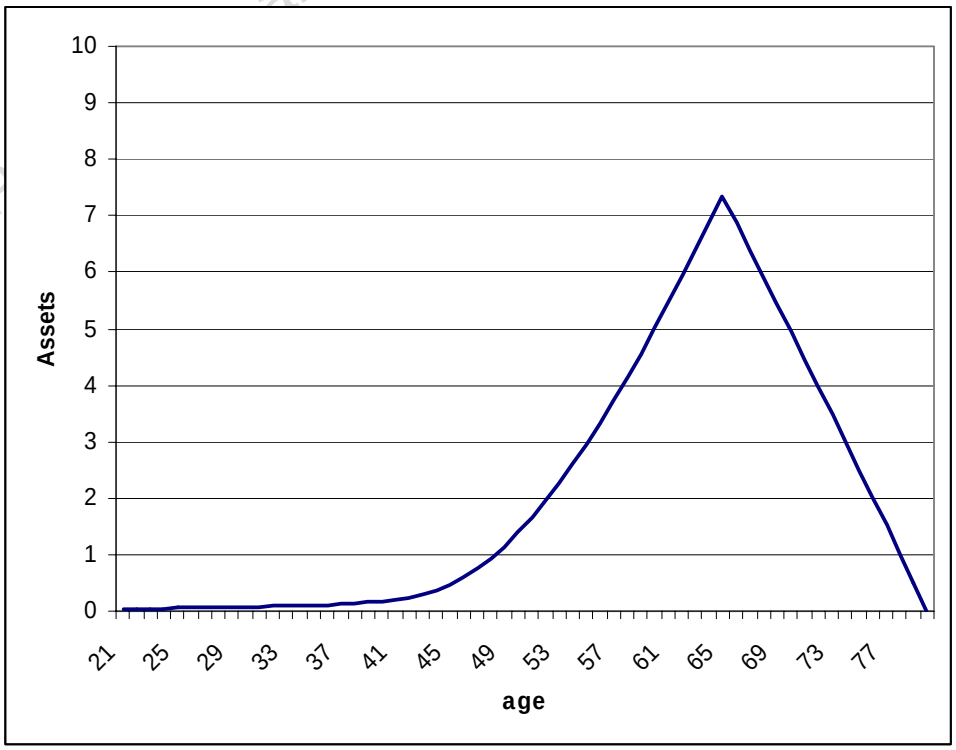

Note: Assets are normalised by dividing through by expected income at age 21.

These factors also show up in the profile of expected asset holdings across the lifecycle, which is plotted in Figure 3.2. In this picture, the level of asset is scaled by expcted income at age 21. Notice the very low level of average assets in the early part of the life cycle (corresponding to the period in which individuals are liquidity constrainted), followed by a sharp increase corresponding to the accumulation of assets for retirement and the subsequent running down of assets.

Our main interest is to determine how the profile asset holding changes when the real interest rate is increased from 2% real to 2.5%. This would roughly correspond to a case in which the pre-tax rate is 0.033 and the tax rate on saving is reduced from 40% to 25%. In Figure 3.3, we plot two average profiles corresponding to these two different values of the interest rate. In Figure 4, shows how the level of accumulated assets changes, on average, in each period of life, which may be thought of as the effect of the change in the interest rate on asset holdings. Obviously, until the early forties, there is no effect, because our agents are not saving in the first part of their life. Removing the liquidity constraint would reduce the amount these individuals borrow. Starting in the mid forties, however, the effects of this change in real interest rates are not negligible, at least for the level of the EIS considered in this baseline simulation. (EIS=1.) .

Figure 3.3

Profile of average asset holdings: baseline case, different interest rates

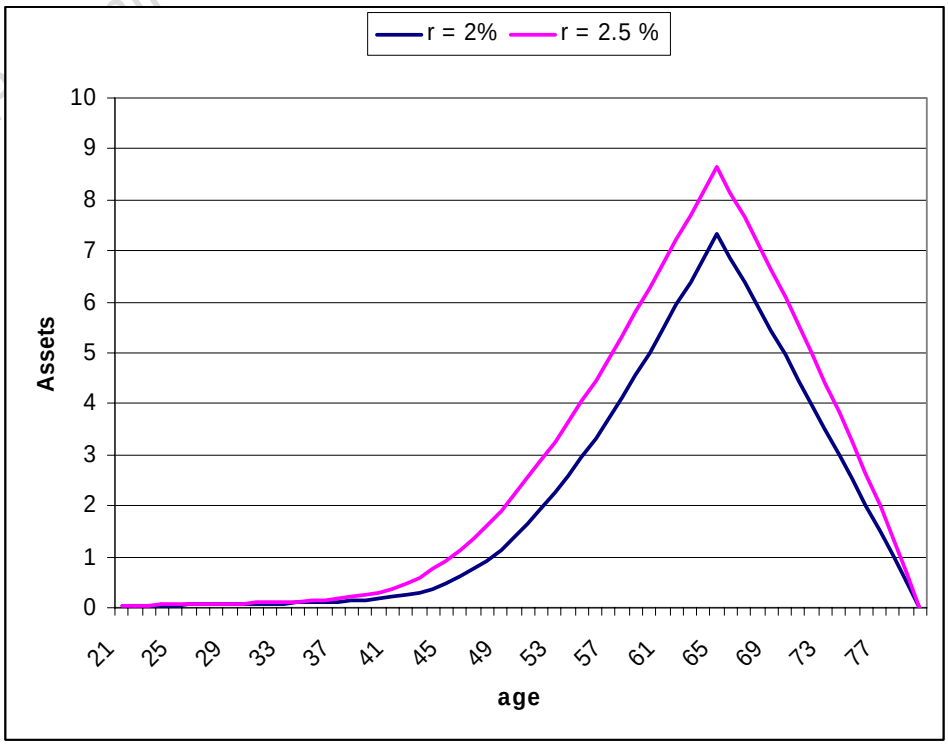

Figure 3.4 Average change in asset holdings when r changes: baseline case

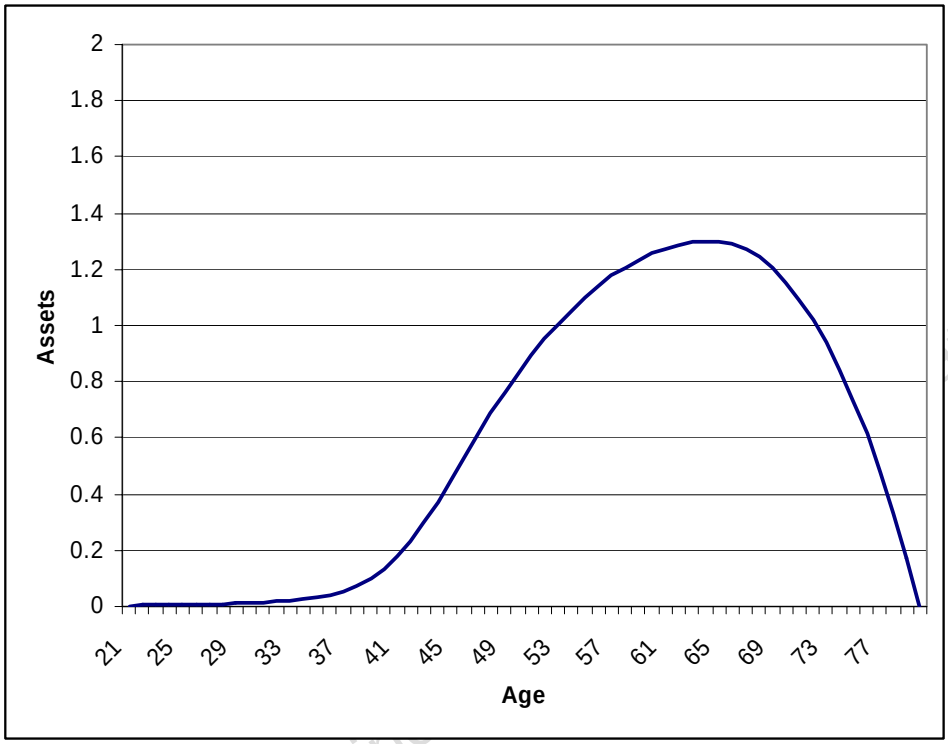

Note: Assets are normalised by dividing through by expected income at age 21.

Our next question is how much these effects change with different levels of the EIS. For this reason, we repeat the exercise whose results were plotted in Figures 3.3 and 3.4 for three different levels of the EIS. We report the graphs corresponding to this exercise in Figures 3.5 and 3.6. In particular, in Figure 3.5 we plot the average level of assets corresponding to three different levels of the EIS (1, ½ and ¼), while in Figure 3.6 we report the effect of the change in interest rate corresponding to these same three levels of the EIS. Not surprisingly, assets are in general higher for lower levels of the EIS. The effect of a change in the interest rate, however, is considerably smaller for lower levels of the EIS. For instance, halving the EIS from 1 to ½ reduces the peak effect from about 1.3 to about 0.7.

Figure 3.5 Profile of average asset holdings: baseline case, different preferences

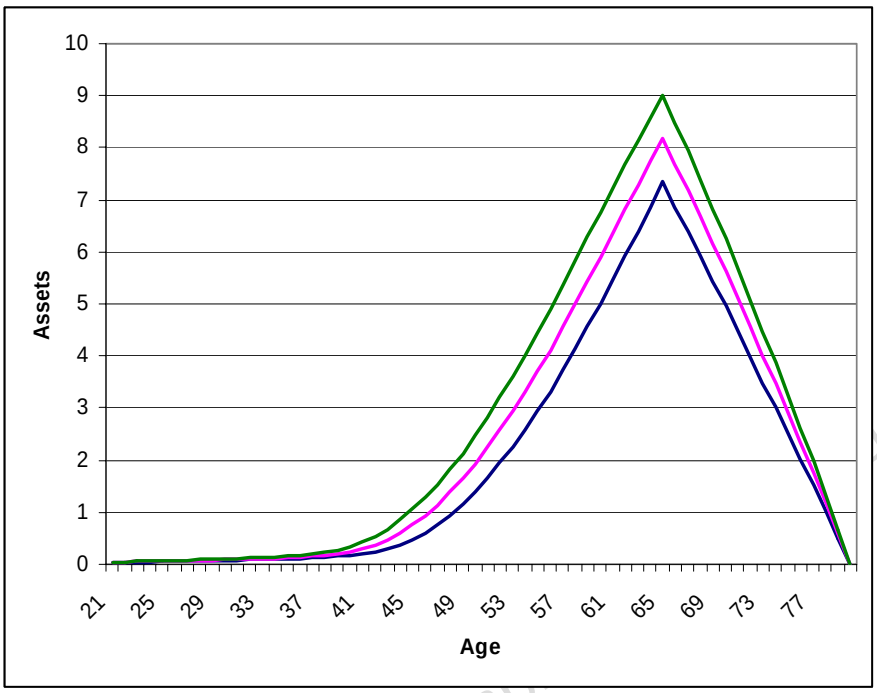

Figure 3.6. Average change in asset holdings when r changes: baseline case, different preferences

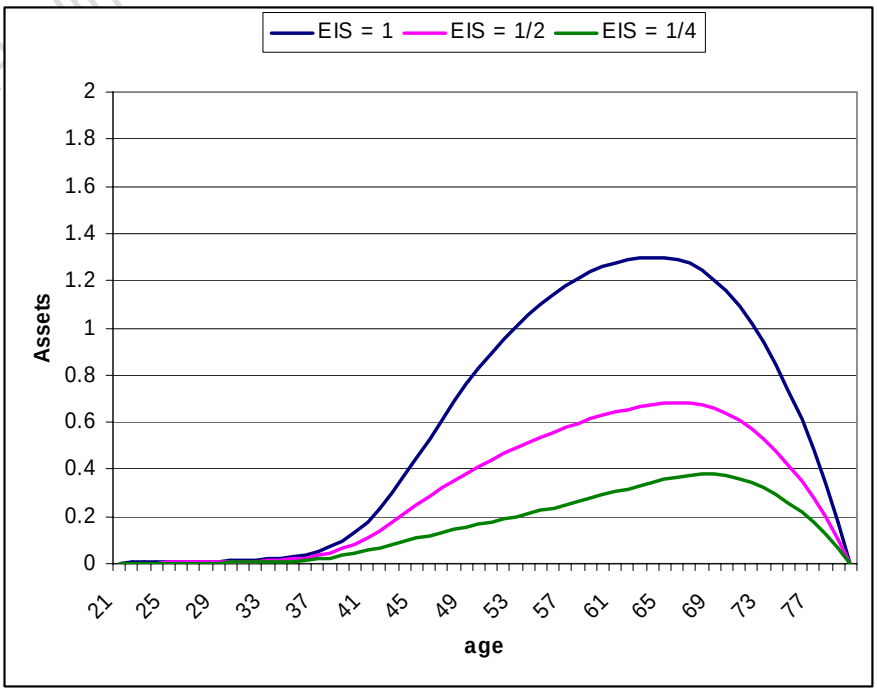

The model we have considered so far is very unrealistic in that it does not consider changes in consumption needs over the life cycle. Our first extension to the basic model is to introduce such needs. Utility is now defined in terms of consumption per adult equivalent, where we consider [add details of equivalence scale here. At the minute it rises from x=1 (two adults) to x=1.8 (two adults + two kids) smoothly between ages 20 and 42, and then declines back to 1 from age 65. Utility is a function of c/x. Might want to check x against average profiles in, say FRS].

In Figure 3.7, we plot, together with the baseline consumption profile plotted in Figure 3.1, the consumption profile implied by a model with changing consumption needs. As noted in Attanasio and Browning (1995) and Attanasio, Banks, Meghir and Weber (1999), the explicit consideration of changing family size and composition generates a hump shaped life cycle consumption profile. Relative to the case in which needs are constant and consumption tracks income because of liquidity constraints in the first part of the life cycle, the hump is much more pronounced in this context, because of the hump shaped nature of the adult equivalents profile over the life cycle.

Figure 3.7. Profile of average consumption, with & without family

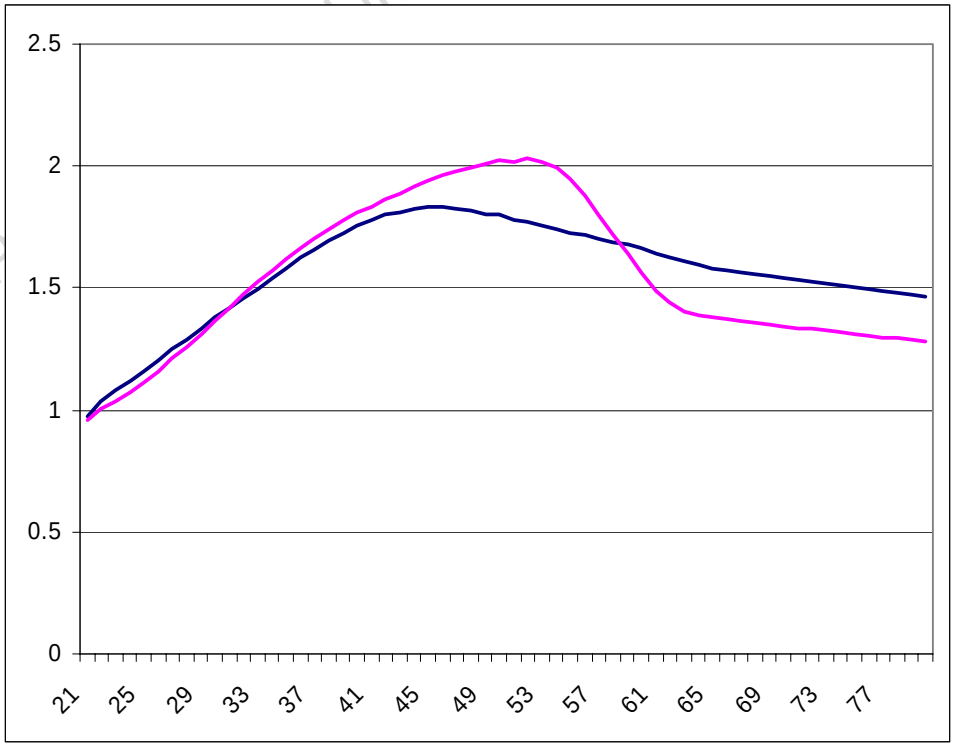

In Figure 8, for the same two cases, we plot the implied level of assets. In the changing needs model, the level of assets is much reduced, a direct consequence of the hump shaped nature of consumption, which is much closer to the profile for income.

Figure 3.8. Profile of average asset holding, with & without family

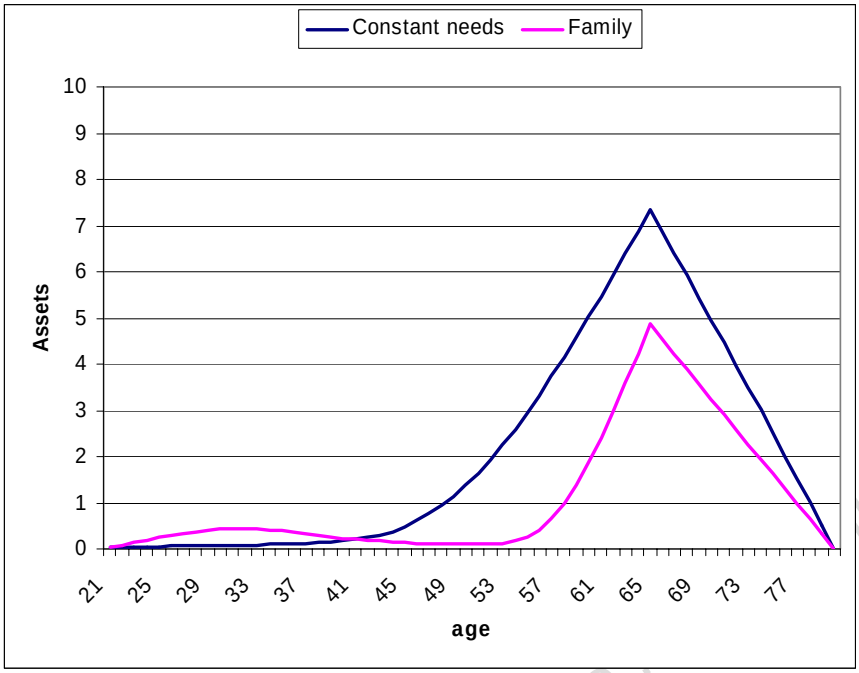

Finally, for the case of changing needs, we also perform the comparative statics exercise of increasing the interest rate. The effect of the change in the interest rate is now much reduced.

Figure 3.9 Profile of average change in assets, with & without family

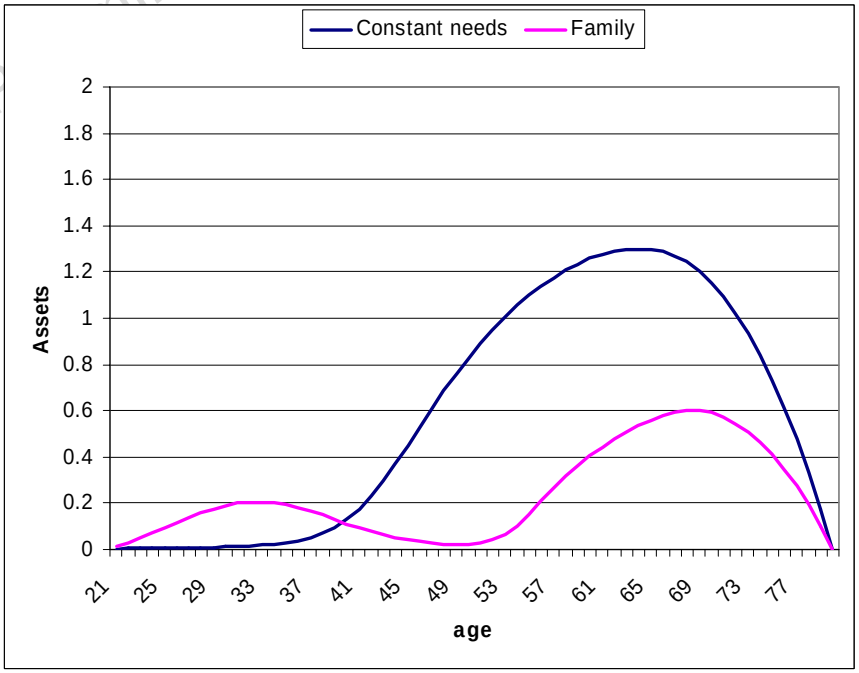

Note that when we introduce family needs, the family is not compensated and the income profile considered remains the same, so that welfare is lower. The thought experiment in simply how is the profile changed if needs that vary over the lifecycle are taken into account.

Finally, we compare the baseline case to two cases with different levels of income replacement rate in retirement. The replacement rate is given by some 'pension income' relative to the final income received during the working life and excludes completely interest income. In the base case retirement income is half the final working age income, and in the comparison cases is 25% or 75%. It should be stressed that an increase (decrease) in replacement income represents a direct increase (decrease) in the amount of life time resources that can be enjoyed by our individuals. Not surprisingly, therefore, consumption is higher (lower) and saving is lower (higher) for the higher (lower) replacement case. The more interesting exercise is to check how the changes induced by changes in interest rates are affected by different replacement rates. In Figure 3.10, we plot again the changes in asset holding induced by the higher interest rate for the baseline case, together with the same changes in the two new comparison cases. As can be seen, changes in replacement rates do not affect dramatically the effect that interest rate changes have on asset holdings.

Figure 3.10: Average change in asset holdings due to change in the interest rate from 2% to 2.5%: different replacement rates

[This for EIS = ½ : need to do runs for EIS = 1]

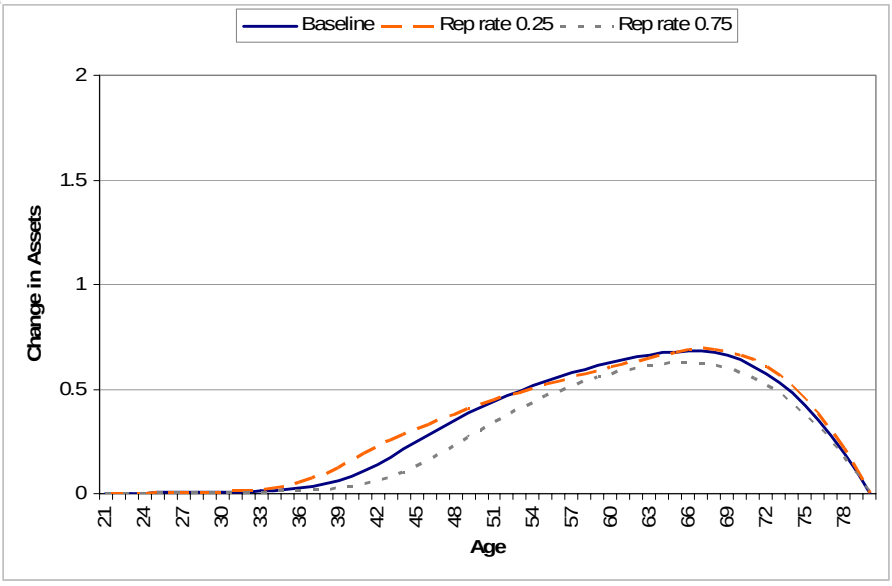

Note: Assets are scaled by dividing through by expected income at age 21

One question we can ask is how the increase in interest rates affects the well-being of individuals. To quantify this, we convert the change in expected utility at the beginning of life that results from the interest rate change, into a "compensating variation" (CV) consumption value. This is the proportionate increase in consumption each period that would be required to make the individual indifferent between accepting the lower interest rate with this boost to consumption, or accepting the higher interest rate. Table A1 displays such CV values for several of the different examples that are described in the above pictures.

Table 2: Compensation variation in consumption required to make an agent indifferent between a 2% real interest rate with this compensation, and a 2.5 % real interest rate

[These have been computed assuming (despite liquidity constraint) that everything is homothetic, which is probably slightly wrong and can be correctly later.]

Case	CV value (% increase in consumption)
Baseline	1.92
EIS 1	1.04
EIS 4	4.76
Flat income profile	To follow
Family needs	0.48
Replacement rate 0.025	To follow
Replacement rate 0.075	To follow

4. The life cycle model: where do we stand?

The model that we have used to perform the simple exercises in Section 3 is obviously, even in the most complicated version we have used, a fairly simple and to some extent unrealistic framework. Nonetheless we would argue that it constitutes a useful framework to address the issues we have been discussing. Such a stand must be justified on two grounds: first, the model is, at some level, an appropriate conceptualization of individual behaviour; second, the many simplifications we have introduced do not constitute a fundamental drawback of the structure we have used and the latter can provide useful insights. In this section, we will start by discussing the evidence on the plausibility of the life cycle model. We will first consider structural estimates and tests of the model. We will then move on to consider indirect evidence on the relevance of some implications of the model that can be obtained from the reaction of individuals or groups of individuals to policy reforms. Finally, we will discuss some important aspects that

were not included in the model used in Section 3 as well as some important alternatives to the life cycle model.

4.1. The evidence on the life cycle model.[12]

In our opinion, since its initial formulation of the 1950s, there have been two important development in the literature on consumption and the life cycle. First, since the 1970s, economists have learned to introduce uncertainty in a rigorous fashion in the theoretical framework. The assumption of rational expectations has had some important applications in the consumption literature, as we will discuss. While this has made the analytical properties of the model much harder, it has also revealed some important issues. While the precautionary saving motive was discussed already in Dreze and Modigliani (1972), its implications have only been extensively explored much later (see, for instance, Carroll, (1992)). Second, household level data have been brought to bear on the empirical relevance of the model and a number of contributions have made a serious attempt at complex specifications of preferences that could be brought to the data.

Since Hall (1978) and Sargent (1978) and Flavin (1981), the empirical analysis of the life cycle model has taken two main strategies. A group of contributions, has based the analysis on the so-called 'Euler equation' for consumption, that is the first order condition that links the marginal utilities of consumption at two different dates to the interest rate and the discount factor. This equation reflects what is one of the main implications of the model: the consumer will choose consumption and saving states that at the optimum the ratio of expected marginal utilities is equal to the ratio of relative prices, that is the interest rate. This equation has powerful implications. Together with an assumption about the nature of expectations it implies a relationship among observable variables that can be used to estimate the preference parameters and to test the model. The first focus of this literature, starting with Hall (1978) was on testing the model. In that first paper, the implication tested was that, conditional on current consumption, other current variables, including income, do not help in predicting future consumption. Subsequent papers (see, for instance, Campbell and Mankiw (1989)) tried to interpret rejection of this strong hypothesis in terms of some type of market failures (liquidity constraints that force some individuals who would like to borrow against future income to consume no more than current income), or rationality failure ('rule of thumb'

[12] Some of these issues are discussed in Attanasio (1999) and, more recently, in Attanasio and Weber (in progress).

consumers who consume a fraction of their current income). Most of these papers (an exception being Hall and Miskin (1982)) used aggregate level data and neglected completely aggregation issues. In the late 1980s and in the 1990s, some contributions started to fit the Euler equation to individual level data.

Some of the relevant papers were Attanasio and Weber (1989, 1993, 1995), Blundell, Browning and Meghir (1994), Banks, Blundell and Preston (1994), Attanasio and Browning (1995). A short and definitely subjective and not unbiased summary of these results is that a rich version of the life cycle model can fit the available data, especially if one focuses on households headed by prime aged individuals. In these papers, that typically used time series of cross sections (often the U.K. Family Expenditure Survey – FES) to generate pseudo panels, the emphasis was on rich and theory-consistent specification of preferences that used the insights of two-stage budgeting and focus on the allocation of resources over time. The specifications used allowed for a flexible role of demographics and family composition in preferences. Several of these papers also allowed for non separabilities between leisure (and in particular female participation) and consumption. The results showed that on one side the data failed to reject the model and on the other the approach provided useful estimates of preference parameters. Attanasio and Weber (1993) also showed that aggregation biases could explain some of the results in the macro literature, such as the celebrated Campbell and Mankiw (1989) paper: when 'wrongly' aggregated the micro data delivered estimates that were very similar to those obtained by Campbell and Mankiw (1989). Moreover, they showed that the rejection of the standard model could be attributed to the dynamics of the cross sectional second moments of consumption.

Aside from the role of demographics that could be interpreted as complex adult equivalence schemes, the specification often used was of the iso-elastic type. The implied EIS obtained in these papers was considerably higher than those implied in the macro level data. Hall (1988) for instance, reported elasticities of intertemporal substitution around 0.1 and not statistically different from zero. The papers cited above reported estimates between 0.65 and 1. The implications of these much higher estimates for the taxation of saving are obvious, as showed by the simple computations in Section 3. And yet these findings have not been consistently used in the public finance literature.

The big attraction of the Euler equation approach is that it can deliver an empirically treatable specification without necessarily making very strong assumptions. The approach

is robust to the presence of several imperfections in different markets in which the individual acts, it is possible to control for non separable leisure and consumption choice, it can control in a reasonably flexible way for unobserved heterogeneity and, above all, it is not necessary to specify the whole stochastic environment in which the individual operates. The price one pays for this is that the approach does not deliver a 'consumption function'. It is therefore not possible to establish how consumption or saving will change in reaction to changes in the various variables faced by the individual. This is obviously an important limitation for policy analysis and probably explains the dichotomy mentioned above between the empirical consumption literature and the public finance literature. There are two possibilities to overcome this difficulty. One is to use numerical methods to obtain consumption functions; the other is to use approximations to the consumption function. We will come back to this issue below.

The other approach used in the empirical consumption literature was that based on simple versions of the model that would yield a consumption function. One if the few such cases is the one in which utility is quadratic and the only uncertainty comes from income. In such a case, Flavin (1981) and subsequently Campbell (1987) and Campbell and Deaton (1989) derived the cross equations restrictions implied by the model on the time series representation of income and consumption. Most of the papers that used this type of approach (including West (1988) and Hansen, Roberds and Sargent (1991)) used time series data. Virtually no papers have looked at this type of restrictions on micro data (one exception is Nalewaik (2004)). Even though subsequent papers obtained consumption function with slightly more general preference specifications (see for instance Caballero, (1991)), the type of assumptions needed to get a closed form solution are restrictive if the aim is to make such an approach viable as a realistic one relevant for policy analysis, and especially if the focus is on variation in the interest rate.

These considerations lead us back to the use of numerical methods (or approximations) to address this problem. And indeed, starting with Deaton (1991), Hubbard, Skinner and Zeldes (1996), Attanasio, Banks, Meghir and Weber (1999) and others, have developed methods to solve these models. The simulations presented in Section 3 are, in effect, an application of this approach. The big difficulty of this approach, if it wants to be realistic and of policy relevance, is that one needs to specify each single detail of the stochastic environment in which the consumer lives. It should be clear from even the simple exercises reported in Section 3, some of these details are quantitatively and qualitatively important for the results one obtains. Moreover, even simple modifications of the basic

model may introduce considerable complications at the numerical level. This was evident in some of the subsequent papers that have taken this approach, such as Palumbo (1999), who looked at health and consumption, Low (2000) who introduced endogenous labour supply choices, Pijoan-Mas (2003) who introduced habits and, more recently, Attanasio, Bottazzi, Low, Neisham and Wakefield (2007) and Bottazzi, Low and Wakefield (2007) who introduce endogenous housing choices and housing and labour supply choices respectively. We discuss these issues below. It is clear, however, that, perhaps not surprisingly, when start using numerical methods one looses the big advantages of the Euler equation approach: robustness and simplicity.

The difficulty in getting numerical solutions of even modestly realistic models implies that this approach can be used only with great difficulty for the estimation of parameters. A possibly more productive approach, which is effectively the one used in Section 3, is to obtain estimates for some of the parameters from formal estimation, possibly using robust methods such as Euler equations, and others possibly from matching specific data moments.

Before concluding this subsection, it is worth mentioning the use of approximate solutions to the life cycle problem. These approximations have recently been used extensively in some of the papers that have looked at the evolution of inequality. Blundell, Low and Preston (2006) show that some of these approximations work well in some situations.

4.2. Evidence from policy reforms

In the previous section we cited empirical evidence that has attempted to directly test certain implications of the lifecycle model, or to estimate certain parameters of the model. In this section we refer to evidence that is largely based on policy reforms and which does not so directly test or estimate the model described in section three, but which may nonetheless be interpreted as shedding light on the plausibility of some of the model's assumptions. The evidence we draw on does not include that on responses of consumption and savings to reforms to the taxation of the return to savings: as we pointed out above this evidence has been somewhat disjoint from the consumption literature and we postpone discussion until section 5.X, below. [This structuring may need revisiting].

As the discussion of section 2 indicated, during the last 25 years the U.K. has experienced a plethora policy reforms affecting pensions and other savings products.

Evidence on individuals' responses to at least some of these reforms suggests that individuals have responded to the incentives that have been created, often in ways at least superficially consistent with optimising behaviour.

One striking example of responses to the incentives came when personal pensions were introduced in 1988. When this happened, individuals were provided with a financial incentive to opt out of the State Earnings Related Pension Scheme (SERPS) in favour of these private savings vehicles. These incentives were better than actuarially fair, and structured to be stronger for younger individuals (see Disney and Whitehouse, 1992). The response to these incentives was clear: within five years take-up of the new accounts exceeded Department of Social Security projections by a factor of more than 4 [check], and more than half of those who opted into the new accounts were aged under 30. Other (albeit less striking) examples of responses of savings behaviour to incentives include data showing that large balances were quickly transferred into PEPs and TESSAs and subsequently ISAs when these products were introduced (a fuller discussion is included in section [5.X]), and evidence that when stakeholder pensions were introduced in 2001 this may have affected the saving behaviour of lower income individuals who were affected by a change in the limits on tax-relieved saving in pensions (see Disney, Emmerson and Wakefield, 2007).

While these episodes indicate that individuals do respond to financial incentives when making retirement savings choices, the data that have been available to analyse them have precluded a clean test against the predictions of the lifecycle model. One area where this has been more possible has been in assessing whether changes in state pension wealth have led to offsetting changes in private savings. If state and private pensions are good substitutes, then the lifecycle model suggests that changes in the former should be offset through changes in the latter, and if wealth effects from changes in state pensions are adequately controlled for this offset should approach pound-for-pound as the degree of substitutability between the two types of wealth increases. This insight has been exploited in a number of international studies that have attempted to assess the extent of offset between these two forms of wealth [cite US papers, also Attanasio and Brugiavini, 2003].

Recent[13] UK evidence on this issue is presented by Attanasio and Rohwedder (2003), who look at the extent to which the consumption and saving behaviour of past cohorts during their working lives was affected by the announcement of reforms to SERPS and

[13] The discussion in the remainder of this subsection is adapted from Blundell, Emmerson and Wakefield, (2006).

the basic state pension during the 1970s and 1980s. They exploited the fact that these reforms hit different cohorts at different ages to estimate the degree to which public pensions crowd out private saving at different points in the life-cycle. Their evidence, depicted in figure 4.1, suggests that the introduction of SERPS did lead to an offsetting reduction in private saving at least for middle aged and older working age households, and point estimates of this effect were 65% and 75%, respectively, for the two older age-groups depicted in the figure. However, while Attanasio and Rohwedder find a significant impact of the introduction of SERPS on saving, they find little evidence of any impact of the 1981 decision to reduce the generosity of the basic state pension by indexing it to price growth rather than to the greater of the growth in prices or earnings.

To understand this evidence that individuals did offset an increase in state pension wealth by reducing their saving did not do the reverse following a decrease in public pensions, we probably need to consider extensions to the simplest complete markets, full knowledge version of the lifecycle framework. One possible explanation put forward by Attanasio and Rohwedder, is that liquidity constraints may have been important. Credit constraints could explain why young individuals did not seem to respond to the SERPS reform, and if the young individuals who bore the biggest financial consequences from the basic state pension reform were credit constrained even before the change, then this could also explain the limited response of their consumption behaviour to that reform. An alternative explanation is that the SERPS reform – which was introduced with cross-party support and with some publicity – was more widely understood and anticipated to be long lasting, than was the change in indexation for the basic state pension.

Figure 4.1: Evidence on the substitutability of state pensions and private saving from Attanasio and Rohwedder (2003)

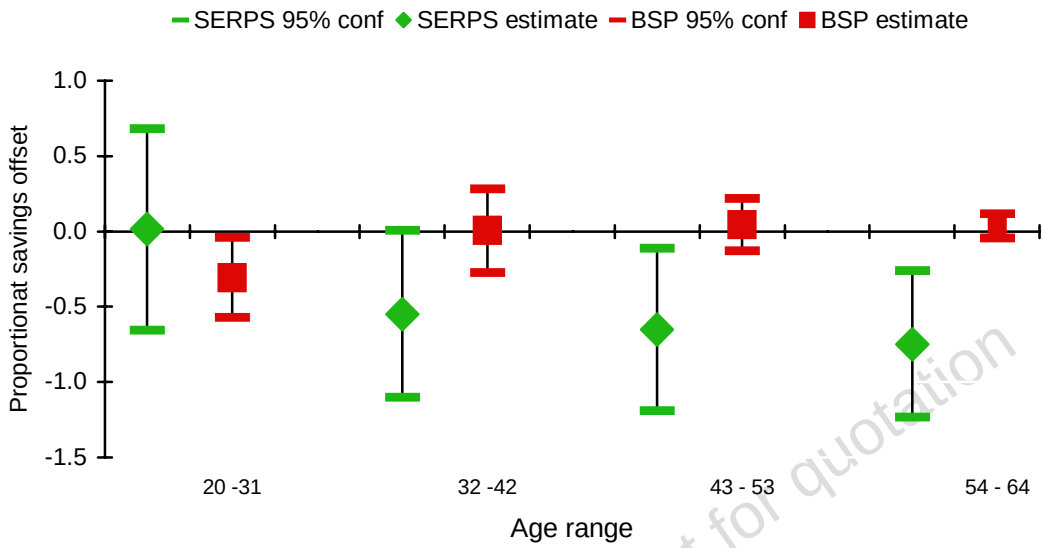

Source: Blundell, Emmerson and Wakefield, 2006. Authors' calculations based on Attanasio and Rohwedder, (2003).

As suggested in the previous paragraph, if individuals' beliefs and expectations do not wholly reflect the current rules of the pension or tax and benefit systems, then we cannot expect to observe responses to all elements of these system that would accord with the predictions of an economic theory that is predicated on a full understanding and belief of current rules. Good data on expectations could therefore be valuable in assessing *why* individuals respond to the pension system in the ways that they do. For example, by using data on a recent period of state pension reforms in Italy, Bottazzi, Jappelli and Padula (2006) are able to assess how these reforms affected expectations and then to infer how fully individuals' changes in behaviour reflected their new beliefs. To do this, these authors exploit data that include information on expected retirement age and replacement rates, as well as on household wealth and demographics. They conclude that individuals did not immediately fully internalise the implications of a series of pension reforms in their expectations of retirement outcomes, and that even their expectations about changes in social security wealth were not fully accommodated through changes in private wealth accumulation.

Evidence on how expectations adapt to pension systems and reforms probably cannot be directly generalised, but rather is specific to the particular institutions under analysis. To our knowledge relatively little research on pension expectations and savings has been conducted in the U.K., probably due to a lack of good evidence on expectations of

retirement outcomes. New data, for example from the English Longitudinal Study of Ageing (ELSA), should help to fill this gap for current and future reforms (see chapter 5 of Banks et al, 2005).

If, whether for reasons of expectations or due to other factors, individuals do not fully adjust their personal wealth to accommodate social security wealth, then this may show up not only at the time of reforms but also in consumption patterns later in life. Banks, Blundell and Tanner (1998) exploit exactly this insight and examine how consumption changes around the time of retirement. A fall in consumption at retirement does not necessarily imply that individuals had not saved enough – some part of the drop in consumption may be planned. Banks et al find that by modelling individuals' life-time consumption plans, around two-thirds of the drop in consumption growth at retirement that occurred for those U.K. cohorts retiring in between the 1970s and the early 1990s, can be explained within an optimal consumption plan. The residual third remains a puzzle, with one possible explanation being that at least some individuals had not saved enough. On the other hand the drops in consumption may be precipitated by unforeseen shocks affecting some individuals, rather than by a lack of forward planning. For example, evidence from panel data (Smith, 2006) suggests that it is those who left the labour market as a result of an employment or health shock who experienced a decline in their food spending and potential indicators of their well-being around the time they left paid employment.

[Should also cite Scholz et al JPE 2006. Here or in 4.1?]

[Conclude the section here]

4.3. When the simple model does not work.

As we mentioned several times, the model presented in Table 3 makes some very stark assumptions. In this subsection we analyze the following situations, in which the predictions of the standard life cycle model considered above would not apply. [14] What they have in common is the failure of some of the assumptions made above.

(i) Leisure and consumption are not separable so that income is not an exogenously given process.

[14] Effectively in section 3 we did consider an important deviation from the standard LC model: we did not allow consumers to borrow. Whether this makes a difference to the results we discussed depends on how binding the exogenously given borrowing restriction is, given the income profile and the needs of the individual family.

(ii) Individuals receive utility from housing which also constitute a lumpy asset subject to transaction costs.

(iii) Intertemporal trades available to individuals are more complex than the stark assumption made in Section 3. One can consider situations in which individuals can borrow against some assetts (typically housing) or at higher interest rates and can even interact these restrictions with earnings (for instance one could borrow an amount that is a function of earnings. Finally, one can consider situations in which the market structure is endogenously determined by the specific imperfections one consider.

(iv) Preferences are not intertemporally separable, either because of durability of consumption or because of habits.

(v) Consumers do not discount future utility geometrically but hyperbolically. Another set of preferences that could be explored is those proposed by Gul and Pesendorfer (2002).

All these points will be briefly discussed in terms of the likely implications they might have for the type of exercises we discussed in Section 3 and more generally for saving taxation.

4.3.1. Labour supply and saving

An important assumption in the simulations we have reported so far is that the income process considered is exogenous. As we mentioned above, this assumption can be justified with the assumption of separability of leisure and

4.3.2. Housing

4.3.3. Intertemporal trades

4.3.4. Habits and durability

4.3.5. Hyperbolic discounting and temptations

5. Discussion of reforms

Discussion of evidence from reforms that altered tax on return to saving. We divide our discussion of the available empirical evidence into two parts. First we discuss evidence that can be inferred from large changes in legislation that affects aspects of saving behaviour directly. The most obvious example is pension reform. While the evidence is not often presented within the framework of the life cycle model (or any specific model for that matter), it is obvious that some of the results obtained in this analysis have direct bearings both for the plausibility of different theoretical models and for the size of the relevant elasticities. In the second part of this section, instead, we discuss the evidence on the plausibility

5.1 Evidence from reforms past

- Begin with UK, since review is UK focussed.
    - For UK reforms, draw on Attanasio, Banks and Wakefield (2005).
    - Weakness of UK data leading to tentative conclusions (can't properly frame tests exploiting economic insights).

- Better data, and fuller literature in US.
    - Start with evidence on 401ks (Jim Poterba has rightly pointed out that there is better cross section variation to exploit here than there was for IRAs). Engen, Gale & Scholz; Poterba, Venti & Wise.
    - IRAs, Engen, Gale & Scholz; (Poterba), Venti & Wise; Attanasio & Deleire.
    - Also evidence from Duflo, Gale, Liebman, Orszag &, Saez.

- [A summary can be drawn from Blundell, Emmerson and Wakefield:

  The UK also provides examples of the introduction of non-pension accounts that have provided tax relief on the returns accruing to savings, first with Personal Equity Plans (PEPs) and Tax Exempt Special Savings Accounts (TESSAs), and more recently with the more flexible Individual Savings Accounts (ISAs). Data limitations make it hard to conduct thorough econometric analysis of the UK experience of these accounts, but descriptive evidence is available in Attanasio, Banks and Wakefield (2005). Aggregate evidence on the balances held in TESSAs indicates that these balances tended to jump immediately when new contributions to accounts could be made in a new year, and also that average contributions were often close to the maximum amounts that could be

deposited in accounts. This, it is argued, is at least consistent with a pattern of individuals largely reshuffling existing wealth into TESSAs, rather than making new savings.

The authors also present microeconomic data on the experience of ISAs which indicate that while the take-up of ISAs was quite high, there is no strong evidence that this had much affect on overall ownership of non-pension financial assets or on levels of saving among those with such assets. However there is some evidence of an increase in ownership of financial assets among low education groups and the young. This could suggest that ISAs were successful in being more attractive to low income savers than TESSAs or PEPs had been (Chapter 5 of HM Treasury, 2000). There has, however, been at least one study using micro-data to examine the effects of ISAs which argued that the products were little better than TESSAs and PEPs at reaching some low-income groups (Paxton, 2003).

Attanasio, Banks and Wakefield argue that the UK evidence that they present is consistent with US evidence concerning tax deferred "Individual Retirement Accounts" (IRAs) that is drawn directly from Attanasio and DeLeire (2002). Attanasio and DeLeire use short horizon panel data on consumption and argue that while around 40% of contributions to IRAs in the 1980s may have been additional personal saving, once account is taken of the cost of tax relief then less than one dollar in ten of contributions could be considered to be new national saving. Thus, for the cases that they consider, Attanasio, Banks and Wakefield conclude that "only relatively small fractions of the funds going into tax-advantaged savings vehicles can be considered to be 'new' saving. As such, the best interpretation of the evidence is that such policies are expensive ways of encouraging savings."

[*Last paragraph out of context if put before US evidence – instead go from second para, into discussion of weakness of UK data leading to tentative conclusions (can't properly frame tests exploiting economic insights). Link into US literature to survey.*]

[5.1 (b) Should we insert a subsection here on responses through retirement behaviour?

Again, BEW has a discussion of retirement behaviour & incentives (§4.2.2). Perhaps just refer to the issue and existing surveys?]

5.2 Evidence and simulations concerning reforms present and future

- Increased scope for transferring money in near retirement now have lifetime limits:

  - Include some evidence on income profiles. Particularly on scope for individuals to transfer money into a pension rapidly near to the time of retirement. e.g. Look in ELSA at how pension wealth of individuals aged 50-64 compares to current income.

- Sefton, vandeVen and Weale (EJ, forthcoming) is a good study of how pension credit does and might affect saving (and retirement).
- There is also a simulation study about tax deferred accounts in US by Gaobo Pang (student of Rust).

6. Conclusions

Some generalisations based on what we have written. Perhaps – people definitely respond to incentives, at least in portfolio. Level of saving may be relatively insensitive to policies that increase return.

Some straightforward policy conclusions:

- Any 'silly' distortions we can recommend scrapping.
- Say that although we are about recommending reforms, a key point implicit in the lifecycle framework is that stability would really help people planning their saving behaviour!

References

Attanasio, O. P., J. Banks and M. Wakefield, (2005), 'Effectiveness of tax incentives to boost (retirement) saving', OECD Economic Studies, Special Issue: Tax-favoured retirement saving, No.39, pp.145-172.

Attanasio and Brugiavini, (2003), *QJE*.

Attanasio, O. P., and T. Deleire (2002), "The effect of Individual Retirement Accounts on Household Consumption and National Saving", *The Economic Journal*, 112, pp.504-538

Attanasio, O. P., and H. Low (2004/5) …

Attanasio, O. P., and S. Rohwedder, (2003), "Pension Wealth and Household Saving: Evidence from Pension Reforms in the U.K.", American Economic Review, 93(5), December, 1499–1521.

Attanasio O. P., and G. Weber, (1995) …

Attanasio O. P., and G. Weber, (2002) …

Banks, Emmerson, Oldfield and Tetlow (2005), Prepared for Retirement? The adequacy and distribution of retirement resources in England, London: Institute for Fiscal Studies.

Blundell, R., M. Browning and C. Meghir, (1994) …

Blundell, Emmerson and Wakefield (2006), http://www.ifs.org.uk/wps/wp0609.pdf

Bottazzi, R., T. Jappelli and M. Padula, (2006), 'Retirement Expectations, Pension Reforms and their Effect on Private Wealth Accumulation', Journal of Public Economics, 90(12), pp.2187-2212

Disney, Emmerson and Wakefield, (2007), 'Pension Provision and Retirement Saving: Lessons from the United Kingdom', SEDAP working paper, http://socserv.mcmaster.ca/sedap/p/sedap176.pdf.

Disney and Whitehouse, (1992), *The Personal Pensions Stempede*

Duflo, E., Gale, W., Liebman, J., Orszag, P., Saez, E. (2005), 'Saving Incentives for Low- and Middle-Income Families: Evidence from a Field Experiment with H&R Block', NBER Working Paper No. 11680 (http://www.nber.org/papers/w11680).

Elmendorf, D. W., (1996), *The effect of interest rate changes on household saving and consumption*, mimeo, Washington DC: The Federal Reserve Board.

Engen, E., W. G. Gale and J. K. Scholz, (1996), 'The Illusory Effects of Saving Incentives on Saving', Journal of Economic Perspectives, 10(4), Fall, pp.113-138. [And references contained therein]

Hubbard, R. G. and J. S. Skinner, (1996), 'Assessing the Effectiveness of Saving Incentives', Journal of Economic Perspectives, 10(4), Fall, pp.73-90. [And references contained therein]

Low, H (Labour supply and savings paper. Also Josep?)

Poterba, J. M., S. F. Venti and D. A. Wise (1996), 'How Retirement Saving Programs Increase Saving', Journal of Economic Perspectives, 10(4), Fall, pp.91-112. [And references contained therein]

Martin Salm (2006) http://www.iza.org/index_html?lang=en&mainframe=http%3A//www.iza.org/iza/en/webcontent/publications/papers&topSelect=publications&subSelect=papers

Sefton, J., van de Ven, J., Weale, M. (2005), The Effects of Means-Testing Pensions on Savings and Retirement, Discussion Paper Number 265, London: National Institute for Economic and Social Research (http://www.niesr.ac.uk/pubs/dps/dp265.pdf). [Latest is EJ, forthcoming].

Smith, S. (2006), 'The retirement consumption puzzle and involuntary early retirement: Evidence from the British Household Panel Survey', *Economic Journal*, Conference papers, c130 – 189.

Cross check survey with BEW article!