Life Cycle Time Allocation and Saving in an Imperfect Capital Market

Patricia Apps and Ray Rees University of Sydney and Australian National University University of Munich and University of York

December 2003

Abstract

This paper combines income and expenditure with time use data to provide a unique picture of the time paths of labour supplies, saving and full consumption for two-adult households over the life cycle. These data are used to test the life cycle model presented in the paper, at the core of which is the hypothesis that households face a borrowing interest rate that rises sharply with the amount of non collateral based borrowing. The household members jointly choose time paths of time use, consumption and saving over their life cycle in the face of this capital market imperfection. This model explains the data much better than does the alternative hypothesis of a perfect capital market. Finally, households are shown to differ significantly in their saving behaviour in a way that depends on secondary earner labour supply, with a strong positive association between saving and the secondary earner's income.

JEL classification: D13, D91, H31, J2

Keywords: saving; labour supply; imperfect capital market; lifecycle.

Acknowledgements: We are grateful to participants at the NBER Summer Institute 2003, in the session on Aggregate Implications of Microeconomic Consumption Behavior, as well as to Jonathan Skinner and Achim Wambach, for helpful comments on earlier versions of this paper. The usual disclaimer applies. We would also like to thank Margi Wood for programming assistance and data management. The research was supported by an Australian Research Council DP grant.

1 Introduction

The defining characteristic of the standard model of consumption choice over the life cycle¹ is, as a result of the assumed separability of leisure and consumption, that the income generation process is *effectively* exogenous to the household. As Heckman (1974) showed, if this separability assumption is relaxed, the great central "puzzle" of the literature based on this model - why current consumption tracks current income so closely in the data - is rather easily resolved. Nonetheless, a controversy continues over how to resolve this puzzle within a model that takes the only household decision variables to be its dated consumptions, with its income stream treated effectively as exogenous.

The leading contenders for resolution of this (model-contingent) puzzle seem to be precautionary or buffer-stock saving,² liquidity constraints in the extreme form of the complete absence of borrowing possibilities,³ and demographic effects, especially the presence of children.⁴

The first of these argues that, to avoid the consequences of adverse random shocks to income in the future, households in the earlier phase of the life cycle build up buffer stocks of assets, and then, in their mid-forties to fifties, begin accumulating savings for retirement and possibly for bequests. In its purest form, this approach seems to be capable of fully explaining the data on household expenditure and saving, in particular the tendency for consumption to track income in the early phase, while leaving virtually no explanatory role for liquidity constraints on the one hand, and demographic factors on the other. Indeed, it implies that liquidity constraints, even if they exist, are non-binding. The household does not want to borrow, completely deterred by the risk that its future income will fall to zero indefinitely. However, if there is a positive lower bound on income (social security, support from other family members), the theory allows that consumers might indeed want to borrow, though never more than the present value of this

¹For a comprehensive survey of the theory and evidence on this model see Deaton (1992). Browning and Lusardi (1996) provide a more concise survey of saving behaviour. Browning and Crossley (2001) and Carroll (2001) give shorter surveys of recent work.

²See for example Carroll (1992), (1994), (1997), Gourinchas and Parker (2002), Hubbard, Skinner and Zeldes (1994), and Zeldes (1989).

³See for example Deaton (1991).

⁴See for example Attanasio and Browning (1995), Attanasio, Banks, Meghir and Weber, (1999), Blundell, Browning and Meghir (1994), and Browning and Ejrnaes (2002).

lower bound on the income stream.⁵ In this case there may be room for other explanations of consumption behaviour.

Under absolute no-borrowing constraints, an impatient household's current consumption will be constrained by its income, and so will track it over time. As opposed to the buffer stock model, households do not borrow because they cannot.

Finally, the demographics approach suggests that if consumption is deflated for family size, it shows the relatively flat time profile consistent with the permanent income hypothesis, under which a household uses the perfect capital market⁶ to decouple its consumption and income streams in such a way as to maintain constancy of its discounted marginal utility of income over time. Browning and Ejrnaes (2002) argue that in this way the data can be fully explained without introducing a precautionary motive.

In his recent survey, Carroll remarks that the development of the precautionary savings approach brings the life cycle model back to its roots in the work of Milton Friedman in the 1950's.⁷ Friedman's analysis was called into question by the results of the models of the 1970's and 80's, based on explicit intertemporal optimization under uncertainty. Carroll argues convincingly that in fact Friedman's intuitions were more closely consistent with the data, and that the recent precautionary savings models provide a superior theoretical underpinning for these intuitions.

However, we should take notice of the fact that one of the single most important socio-economic developments in the forty-five years or so since Friedman's work, has been the large expansion in female labour force participation, with its far-reaching implications for the household's labour supply and income generation process.⁸ The point which motivates the present pa-

⁵See Carroll (1997).

⁶By which we mean one in which the interest rate is the same for borrowing and lending, is invariant to the amount borrowed or lent and to the identity of the economic agent, and has no quantity restrictions of any kind.

⁷Though it may be worth noting that Keynes also identified the precautionary motive as an important reason for saving. It is the first on his list of the motives for saving, see p.107 of Keynes (1936).

⁸Interestingly enough though, Friedman (1957), in defining his Permanent Income Hypothesis, is careful to refer to the "earners" in his "consumer unit" in the plural. This is perhaps because he does not derive his hypothesis from an explicit model of the utility maximising household. Ando and Modigliani (1963), on the other hand, in the formulation of their Life Cycle Hypothesis, do so, and so treat the household as a single individual, which tradition has been followed in the literature ever since.

per is that it no longer makes any sense, if it ever did, to take the household's labour income as effectively exogenous.

As long as models of consumption choices are estimated solely on the type of income and expenditure data available in family expenditure surveys, it does not seem possible to reject the claims made by any of the parties in contention just discussed.⁹ However, when we expand the data set to include the household's time allocation and labour supply decisions, as this paper does, we see that, precisely because of the importance of female labour supply in the modern household, the assumption of an exogenous process of household income determination is no longer sustainable. In other words, possible exogenous uncertainty in the income of the primary household earner may be small beer compared to the variations in household income generated by *endogenous* choices of secondary earner labour supply.¹⁰ This leads to a model which integrates life cycle choices of time allocation, labour supply and consumption.¹¹

In following up this approach, we do incorporate elements of both demographics and capital market imperfections. Decisions on female labour supply are closely related to the presence of children and the choice of sources of supply of child care. Moreover, it seems possible to explain the data only by assuming some kind of capital market imperfection, though our data set, which gives detailed information for each household on purpose, source, amount and cost of borrowing, does not support the extreme assumption of no borrowing. Also, we certainly would not rule out the possibility that some saving could be precautionary in nature, but do not believe, from our inspection of the data, that this can be anything like a *complete* explanation of household consumption behaviour over the life cycle.

An important feature of our modelling approach is the characterisation of the life cycle not in terms of calendar years, but rather in terms of the

⁹Thus Browning and Ejrnaes conclude "the data are not informative enough to allow us to convincingly distinguish between different explanations for the tracking of income by consumption seen in the earlier stages of the life-cycle".

¹⁰Interestingly, Heckman remarks at the conclusion of his paper "It is also relatively straightforward to generalize our results to multiple worker households, although few new analytical insights emerge." This is true only as long as the second worker is just a formal replication of the first, which is not actually the case when childbirth is a possibility.

¹¹The paper by Attanasio, Low and Sanchez-Marcos (2003) adopts an approach similar in spirit to this paper, but quite different in detail. It uses a life cycle model with endogenous female labour force participation, consumption and saving decisions to explain the changes in female participation across three age cohorts in the US.

phases through which a typical family goes over its lifetime. Essentially we are saying that the important differences between households at different stages of the life cycle are not captured sufficiently sharply by differences in calendar age of the head of the household, but rather depend more on whether or not they have children, and on what stage the children are at. By organising the data in this way we are trying to bring out more clearly than in the existing literature the effects of children on the time allocation and labour supply decisions of the household, and, through that, on its income stream and saving decisions. Thus, we argue that the time paths of saving and consumption of market goods reflect the movements in household income that are determined by changes in female labour supply over time, which in turn are determined by the process of substitution between market and household work associated with bringing up children.¹² We then go on to argue that the data strongly suggest that some form of imperfect capital market assumption is indispensable to explaining what happens to household consumption, saving, labour supply and leisure in the early stages of the life cycle. There may appear to be some evidence of "precautionary" saving, in the form of a high level of household saving before the advent of children, but, at least in the context of the present model, this would be better characterized as "anticipatory" saving.¹³ In anticipation of the major impact that the arrival of children will have on family resources, and faced with a capital market that does not offer unsecured loans at a reasonable interest rate, young households save at a higher rate than at any other time in their lives.

Furthermore, the data indicate that households exhibit very considerable heterogeneity in their consumption, labour supply and saving decisions, within and across phases of the life cycle. In particular, saving behaviour

¹²This bears a superficial resemblance to the model of Baxter and Jermann (1999). They explain the tendency for consumption of market goods to track income by arguing that as the wage rate rises over the life cycle, goods produced in the household (of which the most important is surely child care) become more expensive, and therefore substitution toward market goods takes place. In a sense they are spelling out a source of the nonseparability between consumption (of market goods) and household non-labour time that was the basis for Heckman's (1974) contribution. However, the problem with this theory is that the domestic production is carried out predominantly by the female, whose wage does not rise - if anything it tends to fall on average due to depreciation of human capital with nonparticipation - over the life cycle. It is important to model two-person households, as we do here.

¹³Of course, since in our model there is no uncertainty, precautionary saving in the sense of Carroll and Kimball just does not arise.

depends very closely on female labour supply. For example, households with no significant female labour supply do virtually no saving once they have children, other than that involved in house purchase and superannuation schemes. Controlling for primary earner income, there is a high propensity to save out of secondary earner income.

The paper is set out as follows. In the next section we present two models of the household's decisions on consumption, saving and time allocations over the life cycle, which assume respectively perfect and imperfect capital markets. Section 3 then presents empirical life cycle profiles of consumption, saving, labour supply and domestic work, obtained by combining information on income, household expenditure and time use. The results suggest a pattern of full consumption¹⁴ over the life cycle that is very different from that obtained by studies of expenditure on market goods alone. Section 4 presents the empirical specification of the models. Section 5 gives the parameter estimates for the within-period demand system and compares simulated consumption profiles for the perfect and imperfect capital market models in terms of how well they predict the data. As we would expect, the perfect capital market model predicts smooth profiles of consumption that in no way matches the data. We then show that the life cycle profiles of full consumption and leisure can be closely approximated by the imperfect capital market model. Section 6 concludes.

2 The Models

The household has a lifetime of T + 1 periods, with t = 0, 1, ..., T denoting the period. As discussed in the Introduction, and spelled out more fully in the next section, we assume that this lifetime is partitioned into phases corresponding to whether children are present in the household and, if so, what types of demands they are placing on household resources.

It is useful in developing the models to view the household as solving its lifetime allocation problem in three steps. First, within each period, it chooses an optimal time allocation, given the wage rates it faces and its technology of household production. This essentially determines an implicit price of the domestic good. Next, still within each period, it chooses a Pareto efficient allocation of consumption of market and household goods and of

¹⁴Defined as the value of consumption of market and domestically produced goods, but excluding the value of leisure.

leisure among all current members, including children. This then yields an indirect utility function for the household, the key component of which is the total income available for full consumption and leisure in each period. The household then uses the capital market to choose an optimal intertemporal allocation of this income, which determines its borrowing/saving behaviour. Only in this third stage is it necessary to distinguish between the cases of perfect and imperfect capital markets.

In each period t = 0, ..., T, for any given output of the domestic good y_{ht} , household h chooses its allocations of time inputs a_{iht} to solve:

$$minC_{ht} = \sum_{i=1}^{2} w_{it}a_{iht} \tag{1}$$

s.t.
$$y_{ht} = f(a_{1ht}, a_{2ht}; k_{ht})$$
 (2)

where w_{it} are net market wage rates. In this problem the subscript i = 1, 2refers to the adults in the household - only these are assumed to work. The household production function may well vary both across time and across households, due to variations in human and physical capital, and this is expressed by including the parameter k_{ht} . We assume the production function is linear homogeneous and strictly quasiconcave. This problem therefore yields input demand functions $a_{iht}(w_{1t}, w_{2t}; k_{ht})y_{ht}$, and total cost functions $p_{ht}(w_{1t}, w_{2t}; k_h)y_{ht}$, with p_{ht} the implicit price of the household good in period t.

Defining c_{ht} as household total income¹⁵ in each period, and taking this as fixed for the moment, the household solves its within period allocation problem

$$max \sum_{i=1}^{K} \varphi_{iht} u_{it}(x_{iht}, y_{iht}, z_{iht}) = u_{ht}$$
(3)

s.t.
$$\sum_{i=1}^{K} x_{iht} + p_{ht}y_{ht} + \sum_{i=1}^{2} w_{it}z_{iht} = c_{ht}$$
 (4)

where i = 1, 2 again denotes the adults and i = 3, ..., K the children, when present in the household. Here x denotes a market consumption good, and

¹⁵The total value of the household's time endowments at net of tax market wage rates, plus any net transfers from government, and plus (minus) any borowing (lending). It is thus the total available to be spent in each period on consumption of market and household goods and leisure.

z leisure.¹⁶ The φ_{iht} are the household's welfare weights, which will determine the particular Pareto efficient allocation chosen. Given the standard assumptions¹⁷ on the individual utility functions, $u_{it}(.)$, it is well known that an alternative interpretation of this household allocation process is that the household first shares its full income among its members, each of whom then maximises his or her individual utility.¹⁸ This is in fact the approach adopted in the empirical analysis below, in section 4.1. The main result in the present context is that this problem yields an indirect utility function $u_{ht}(p_{ht}, w_{1t}, w_{2t}, c_{ht}; \varphi_{iht})$, with, given the fixed welfare weights, standard properties. It is convenient to suppress the exogenously given prices and wages in the indirect utility function, as well as the distributional weights, and to write it simply as $v_{ht}(c_{ht})$. We now consider the two models of the household's intertemporal choices.

2.1 Model 1: A Perfect Capital Market

There is a single market interest rate at which all households borrow and lend, and which is invariant to the amounts borrowed or lent. The budget constraints in each period are then

$$c_{ht} = A \sum_{i=1}^{2} w_{it} + (1+r)s_{h,t-1} - s_{ht} + P_t \quad t = 0, ..., T$$
(5)

where A is the total time endowment of an adult in each period, s_{ht} is saving (> 0) or dissaving (< 0) at t = 0, 1, ..., T, $P_t \ge 0$ is a lump sum government transfer in each period, which in the retirement phase is the pension payment, and r is the one-period market interest rate, assumed constant over time. To be consistent with the assumption that there is no bequest motive, which implies saving at zero in the last period of life, we also assume there is no inherited wealth, so that assets are also zero at the beginning of period 0. These constraints can be collapsed in the usual way into the wealth constraint

$$\sum_{t=0}^{T} \delta^{t} \left[\sum_{i=1}^{K} x_{iht} - A \sum_{i=1}^{2} w_{it} - P_{t} \right] = 0$$
(6)

¹⁶The quantities of leisure for children are set at their total time endowments.

¹⁷Strictly increasing in all arguments, strictly quasi concave.

¹⁸See Apps and Rees (2002) for the details of this formulation in a multi-person household with children.

where $\delta = (1 + r)^{-1}$ is the market discount factor. Introducing a "felicity discount factor" ρ , the household chooses its time stream of full consumption to solve

$$maxu_h = \sum_{t=0}^{T} \rho^t v_{ht}(c_{ht})$$

subject to its wealth constraint. This is a perfectly standard problem, at least in its intertemporal aspect. The non-standard aspects, involving multiperson households and domestic production, are subsumed in the form of v_{ht} , reflecting as it does the outcome of the within-period allocation.

2.2 Model 2: An Imperfect Capital Market

There is clearly a range of possibilities in modelling an imperfect capital market. As a minimum, we would set the interest rate on saving below that on borrowing. An extreme version of an imperfect capital market would have an upper bound on borrowing, possibly at zero, as for example in Deaton (1992). However, the following formulation would seem both more realistic and consistent with the data we have. All households face the same saving interest rate r_s , and a borrowing rate r_{ht} which is an increasing function of the amount borrowed, $b_{ht} \geq 0$, such that

$$r_{ht} = r(b_{ht}) \ r'(.) > 0, \quad r''(.) \ge 0$$
(7)

and

$$b_{ht} > 0 \Rightarrow r_{ht} > r_s \tag{8}$$

for all h, t. Thus households can borrow, but at an increasing interest rate that is always higher than the lending rate. There is no capital rationing in the sense of an absolute upper bound on borrowing, but of course the function may increase very sharply and b'(.) could approach infinity asymptotically. Realistically, this borrowing function could vary across time and could also contain as arguments the household's income and/or assets, reflecting its default risk and ability to put up collateral for loans. However, on grounds of tractability we stay with this simple formulation. Its implication is that in equilibrium households may face different borrowing rates at the margin, and these rates may vary across periods, depending on the household's borrowing in each period. Fortunately the data set we have allows us to handle this in the estimation procedure. The utility function, time and household production constraints remain as in the previous model. We just have to reformulate the household's budget constraints. We now let s_{ht} denote saving alone. We then have

$$c_{ht} + s_{ht} - b_{ht} = A \sum_{i=1}^{2} w_{it} + (1+r_s)s_{h,t-1} - (1+r_{h,t-1})b_{h,t-1} + P_t$$
(9)

$$s_{hT} = 0 = s_{h,-1} \tag{10}$$

$$b_{hT} = 0 = b_{h,-1} \tag{11}$$

$$b_{ht} \geq 0, \quad s_{ht} \geq 0 \quad all \ h, t \tag{12}$$

The intertemporal problem is then

$$maxu_h = \sum_{t=0}^{T} \rho^t v_{ht}(c_{ht}) \tag{13}$$

subject to the constraints (9) to (12).

Associating Lagrange multipliers λ_{ht} with the constraints (9), the first order (Kuhn Tucker) conditions (assuming full consumption is always positive) are

$$\rho^t \frac{\partial v_{ht}}{\partial c_{ht}} - \lambda_{ht}^* = 0 \tag{14}$$

$$(1+r_s)\lambda_{h,t+1}^* - \lambda_{ht}^* \leq 0 \quad s_{ht}^* \geq 0 \quad [(1+r_s)\lambda_{h,t+1}^* - \lambda_{ht}^*]s_{ht}^* = 0 \quad (15)$$

$$\lambda_{ht}^* - m_{ht}^* \lambda_{h,t+1}^* \leq 0 \quad b_{ht}^* \geq 0 \quad [\lambda_{ht}^* - m_{ht}^* \lambda_{h,t+1}^*] b_{ht}^* = 0 \tag{16}$$

together with the constraints. Here $m_{ht}^* \equiv 1 + r(b_{ht}^*) + r'(b_{ht}^*)b_{ht}^*$ is the marginal cost of borrowing to household h at time t, and $r_{ht}^* = r(b_{ht}^*)$ can be called the household's marginal borrowing rate. Asterisks denote values at the optimum. We can immediately establish the intuitively reasonable¹⁹

Lemma: The household never both saves and borrows in the same time period.

¹⁹Though of course the data have households both saving, through compulsory superannuation payments and financing house purchase, and borrowing short term. The former can best be thought of as exogenous amounts that are subtracted from income in the preretirement phases, and added back in to P_t in the retirement phase, before the household solves its intertemporal optimisation problem. Overall this situation of "lending long" and "borrowing short" can be thought of as a further expression of the imperfection of the real capital market.

Proof: Suppose not, so that $s_{ht}^* > 0$, $b_{ht}^* > 0$, for some t. Then the first order conditions imply

$$(1+r_s)\lambda_{h,t+1}^* = \lambda_{ht}^* = m_{ht}^*\lambda_{h,t+1}^*$$
(17)

But this contradicts the assumption that $r_{ht}^* > r_s$ and r'(.) > 0.

Figure 1 illustrates the nature of the solution. Point γ is the initial endowment point. The household may lend from there at a constant interest rate r_s to reach an equilibrium at a point such as α , characterised by the first order condition

$$\frac{\partial v_{ht}/\partial c_{ht}}{\rho \partial v_{h,t+1}/\partial c_{h,t+1}} = 1 + r_s \tag{18}$$

Alternatively, according to its preferences, the household may borrow along the curve rightward from γ to reach equilibrium at a point such as β , characterised by the condition

$$\frac{\partial v_{ht}/\partial c_{ht}}{\rho \partial v_{h,t+1}/\partial c_{h,t+1}} = m_{ht}^* \tag{19}$$

where m_{ht}^* is the slope of the curve at the optimal point. Clearly household borrowing will be less than if it were possible to borrow at a constant rate equal to r_s (as indicated by the broken line). Our contention is that this reduced borrowing accounts for the large reduction in leisure and full consumption in the second phase of the household's life cycle indicated by the data.

Figure 1 about here.

The key difference between the perfect and imperfect capital market models lies in the impact of changes in household per capita income in a given period on consumption in that period. In a perfect capital market model that impact is diffused over the entire lifetime, leading as it does to a shift in the overall wealth constraint. The effect on consumption in the period in which the income change takes place will therefore be relatively small. In an imperfect capital market on the other hand, a change in the initial endowment point such as γ in Figure 1 can have a large effect on optimal consumption in that period, the more so, the greater the difference between lending and borrowing rates.²⁰ As we show in the next section, the arrival of

 $^{^{20}}$ The extreme case is of course that where there is a no-borrowing constraint, since then the change in income translates exactly into a change in consumption when the constraint binds.

children creates just such an income change. For the average household, the reallocation of time from market labour to household production resulting from the advent of children causes a significant fall in household income, and a large increase in total hours worked, which can only be explained in terms of capital market imperfections that do not allow the impact of the income change to be smoothed over the entire lifetime. We now go on to provide a descriptive picture of life cycle profiles of consumption, saving and time allocation, before proceeding to the empirical analysis.

3 Evidence on Life Cycle Profiles

The first step in our approach is to define eight phases which seem to us to represent the key transitions in the life cycle of a typical household. Given the decision to have children, which we take here as exogenous, the life cycle evolves in a way which seems to be determined by them. This view of the life cycle leads to a representation of the data on labour supplies, consumption and saving for the average household which is as familiar to everyday experience as it is foreign to the economics literature on lifetime consumption decisions.²¹ Before they have children, both household members have high labour supplies, high saving and plenty of leisure. The presence of preschool children dramatically changes the pattern of labour supply, leading to large falls in market labour supply of the secondary earner (usually female), saving, and leisure.²² As the children grow up these changes are gradually

²¹This is not to say that the importance of "demographics" has been ignored, as we hope our discussion in the Introduction has made clear. Our contention is that the effects of having children on female labour supply choices, in the presence of an imperfect capital market and costly market child care, are much more significant than seems to be recognised in the literature, and that this significance is made clearer by the way in which we organise the data.

²²As shown also in Apps and Rees (2001) using 1993-4 data, the true profiles of net income and consumption are lost in studies that define the life cycle on the age of the male (or female) partner. The key problem is that this definition leads to the aggregation of two-income phase 1 households and single-income phase 2 and 3 households with very young children. Averaging across these households produces single humped net income and consumptions profiles (see, for example, Figure 2 in Gourinchas and Parker, 2002). Blundell et al (1994) note specifically that what is interesting in their results is that although female participation falls in the early years, household income does not. However, this finding is we would argue an artifact of aggregating phase 1 couples and phase 2 and 3 young families.

reversed, with the state, through the public education system, taking over a large part of the burden of child minding and education, allowing increases in secondary earner labour supply. Consumption of market goods steadily increases and borrowing falls, with high saving levels again being achieved in the phase immediately after the children have left home. A notable feature of the household's capital market behaviour is the substantial long term saving in the form of house purchase, usually mortgage financed, and saving for retirement in a (possibly compulsory or strongly tax-advantaged) contractual scheme, combined with short term borrowing, often at high interest rates, which is at its peak when the children are young. Our contention is that if the capital market were perfect, the effects of children on labour supplies and on the market/domestic consumption mix would be much less dramatic, with higher borrowing in the early years allowing substantial smoothing of these paths. We now go on to fill in the details of this picture.

3.1 Data

Ideally, panel data are required to construct life cycle consumption and saving profiles. Since they are not available, we use micro-level cross section data. We construct life cycle profiles using information from two complementary surveys, the Australian Bureau of Statistics (ABS) 1998 Household Expenditure Survey (HES) and the ABS 1997 Time Use Survey (TUS).²³ The HES contains data collected by interview on household consumption expenditure and individual incomes, earnings and hours of work. The TUS provides detailed information collected by diary on time allocations to ten activities,²⁴ as well information collected by interview on individual incomes and "usual hours of work". We aggregate the ten time use activities into three general categories: market work, domestic work and leisure. In the

 $^{^{23}}$ The analysis is, in effect, based on a single cross section (all results are presented in 1998 prices) and therefore does not take account of cohort effects. While we recognise that cohort effects can be important, it does not seem to us that they would alter the direction of our key results.

²⁴The activity episode classification distinguishes between labor market activities and nine major categories of non-market activities. Market hours are calculated as the sum of time allocations to all subcategories of labor market activities excluding travel to work and job search. Domestic work is computed as the sum of time allocations to the categories "domestic activities", "purchasing goods and services" and "child care/minding". For each episode, information is recorded for a "primary" and, if relevant, a "secondary" activity. Where primary and secondary activities are reported, the weighting used is 0.6:0.4.

domestic work category there are two sub-categories: child care and time spent on the other domestic work activities. Both surveys provide data on a common set of demographic, education and occupation variables.

We select matching samples of two-adult households from these datasets. All two-adults households are included except for those who do not have children and the female partner is aged from 40 to 44 years. Our reason for excluding these households is that they are likely to represent couples who have decided not to have children and, ideally, we would like to exclude all such households. The sample drawn from the HES contains 4016 records and from the TUS, 1938 records.²⁵

In addition to income and expenditure data, the HES provides detailed information on household debt, house price, mortgage and loan repayments and contributions to mandatory retirement saving and to life insurance. The information on loans is highly disaggregated, for example, by purpose, type of lender, term of loan, etc. The HES also includes estimates of indirect government taxes and benefits as well as the usual detailed data on direct taxes and benefits.

The two samples are split into the eight phases on criteria that capture the presence and age of children and the later transition of their parents from work to retirement. The criteria are also chosen to give phases of nearly equal cell size, for the purpose of comparisons across the life cycle. Phase 1 is limited to couples with no dependent children and a female partner aged under 40 years. Phase 2 represents families with children under 5. Records in this phase are selected on the criteria that at least one child under 5 years is present, there are no older children unless there is a child under 2 years but no child over 9 years. Phase 3 families have at least one child aged 5 to 9 years and may have a younger child or children in the 10 to 11 year age group. In phase 4 the children are predominantly in the 12 to 14 year age group. In phase 5 families have older dependent children still living at home. There are no children present in phase 6 to 8. Phase 6 is defined to include couples in which a partner is aged under 55 years or the male partner is under 60 and has a significant workforce attachment. Phase 7 is pre-retirement, and represents couples in which the male partner is aged under 65, or at least one partner is not fully retired. In phase 8 both partners are retired.

²⁵There are 102 records excluded from the HES full sample of two-adult households and 52 from the TUS sample, on the criteria that no children are present and the female partner is aged from 40 to 44 years.

3.2 The Average Household

3.2.1 Income, Consumption and Saving

Table 1 reports life cycle profiles of median net household income,²⁶ expenditure on market goods and saving, in columns 1 to 3 respectively, using the HES sample. Column 4 gives the average number of dependent children in each phase and column 5 lists the cell size of each phase.

The excess sensitivity puzzle is confirmed by the profiles in columns 1 and 2, which show the strong tendency of household consumption to track net household income, with the highest median consumption expenditure coinciding with the highest net income in phase 5. This is brought out clearly in Figure 2. The figure shows that there is first a sharp fall in median income as the household moves from phase 1 to phase 2. This is then followed by hump shaped profiles of net income and consumption from phases 2 to 8. Saving is at its highest in the pre-children phase, drops sharply in phase 2 with the arrival of children, and fails to rise to near its phase 1 level until phase 6 when the children have left home.

Table 1 and Figure 2 about here

While it is clear that net income and consumption are strongly associated with the presence of children, data on these variables alone can give an entirely misleading picture of the true paths of income and consumption, and of the impact of demographic variation, because they exclude the household's implicit income from, and expenditure on, domestic production. The time use data we now present give an indication of the importance of household production.

3.2.2 Time Allocation

Table 2 reports life cycle profiles of time allocations to market and domestic work. The table lists TUS weighted data means for male and female market hours, domestic hours and total hours of work, in columns 1 to 6, respectively. Comparing these profiles with those for net income and consumption in Table 1, it is immediately apparent that much of the variation in net income across phases 1 to 6 reflects changes in female labour supply or, more specifically, the reallocation of time from market to domestic work by the secondary

²⁶Net household income includes all government direct (cash) benefits but not indirect benefits through, for example, the eduation and health systems.

earner after the arrival of children. Across these phases there is relatively little variation in male market hours but large changes in the hours of the female partner, which are negatively related to domestic hours of work. This is shown graphically in Figure 3.

Table 2 and Figure 3 about here

The strong negative relationship between female market and domestic hours suggests that the two types of work become close substitutes after the arrival of children. The most dramatic substitution occurs in phase 2, reflecting the fact that young children generate a high demand for care. This can in general be provided at home or bought on the market, but the time use data show that there is a very large domestic supply of child care.

In phases 2 to 5 the female partner allocates, on average, 2198, 1980, 1069 and 471 hours to child care and the male partner, 935, 860, 520 and 180 hours, respectively. Even though family size is larger is phase 3, more time is spent on child care in phase 2 because of the predominance of children aged 0-4 in that phase.²⁷ The data show that *total* hours of work rise, and therefore that leisure falls dramatically, with the arrival of children, and this is then steadily reversed over successive phases of the life cycle.

3.2.3 Full Consumption, Taxes and Benefits

When consumption expenditure includes the time cost of domestic production, its profile tends to track total hours of work, rather than net income, as shown in Table 3. Column 1 of the table presents a profile of domestic consumption expenditure, computed as the product of time allocated to domestic work (including child care) and the net wage.²⁸ Column 2 reports the average cost of parental time allocated to child care in these phases.

Studies of life cycle consumption usually omit indirect government benefits. These are important because they are large (averaging over \$11,500 per household in the sample), they vary quite dramatically across phases, and they tend to vary inversely with the household's cost of domestic child care.

 $^{^{27}}$ The average number of children aged 0-4 in phase 2 is 1.43 and in phase 3, 0.75. There are no children in this age group present in subsequent phases.

 $^{^{28}}$ On the assumption of constant returns to scale of time inputs, the expenditure on domestic consumption at the implicit price of domestic output is given by the value of time (measured here as the net wage) spent in household production. To obtain the net wage we intrument for the gross wage and compute a marginal tax rate from the data on direct taxes and cash benefits. Data on earnings and hours are used to compute hourly earnings as the measure of the gross wage. For further details, see Apps and Rees (2002).

Column 3 of Table 3 reports the profile of average indirect government benefits.²⁹ The very high levels in phases 3 to 5 are due to public spending on education and the child care it provides. Government indirect child care and education benefits are shown separately in column 4. Families with children at school or in tertiary education receive by far the larger share of support. Relatively little is spent on children in the pre-school phase.³⁰ In addition, families can face an effective tax penalty in this phase if the mother goes out to work.³¹ Column 5 of the table reports data means for taxes net of government benefits (direct and indirect in both cases) in each phase. The profile captures the net effect of larger indirect benefits for families with school aged and older children and the lower average tax liabilities of families in the early child rearing phases due to the withdrawal of female labour supply.

Table 3 and Figure 4 about here

The life cycle profile of household full consumption, computed to include market and domestic consumption³² and indirect benefits, is presented in column 6 of the table. A profile of the parents' full consumption, obtained by subtracting child costs from household full consumption, is reported in column 7. Both profiles are depicted in Figure 4. Child costs are calculated as the sum of the parents' time costs of child care and government spending on child care and education, plus a share of market and domestic consumption (excluding child care) computed for an "equivalence scale" that sets the cost of a child to 0.4 that of an adult.³³

The U-shaped profile of the adults' full consumption matches their leisure

²⁹The HES estimate of indirect government benefits covers non-cash benefits and services for education, health, housing and social security and welfare. For details of the calculation of these benefits, see ABS (2001).

³⁰Note that the data mean of \$9818 for indirect government benefits in phase 2 includes medical costs for the birth of a child and post natal care. Note also that part of the \$2368 of government spending on child care and education in this phase goes to children of school age who are also present in some families with a child under 2.

 $^{^{31}{\}rm This}$ is due to the withdrawal of tax benefits (specifically, Family Tax Benefit Part B) on the income of the spouse alone.

³²To include domestic production expenditure, we combine information on time use from the TUS with the consumption data for each record in the HES, instrumenting for male and female leisures. For further details, see Apps and Rees (2002).

³³This scale is used elsewhere in the life cycle literature (see for example Blundell et al., 1994) to deflate household consumption expenditure. We obtain an estimate of the average full consumption costs of children per family of around half that of the adults. Child costs of this order are consistent with results for a "sharing rule" in a multi-person model estimated on time use data in Apps and Rees (2002).

profiles, suggesting that parents cut back on both consumption and leisure, instead of borrowing more, in order to support their children in the early child rearing phases. The explanation for this that we suggest is that parents face higher interest rates in the early child rearing phases, particularly in phase 2, together with a lack of access to good quality, affordable market child care.³⁴ Because there is very little government support for child care and high effective tax rates can apply to the incomes of two-earner families,³⁵ the household's optimal choice is, first, to reallocate the mother's time from market to household work, since she generally faces a lower wage rate, and secondly, for both parents, but especially the mother, to work longer hours in total, and so reduce leisure, in phase 2.

In later years, the cost of children to parents is substantially reduced by public funding of education. In other words, when the child reaches school age the public education system takes over many of the child-minding activities that the household itself has to undertake for pre-school children. This allows the female partner to expand her market labour supply, while simultaneously reducing total hours of work once the preschool phase is over. This effect is evident in phase 4 and is accentuated in phase 5. Household income, labour supply and market consumption expenditure all peak in phase 5, with teenaged and older children living at home, while saving is at its peak in the following phase, when the children have left home but market labour supply is still high. Thus, the profile of total hours of work, together with that of adult full consumption, is, we argue, to a significant extent an outcome of an imperfect capital market and variations in the public funding of the costs of children. Once the children have reached school age, access to public education allows parents to maintain family consumption without cutting back excessively on leisure.

³⁴To appreciate the inefficiencies and consequent high cost of market child care, one need only consider the impact that government financial support, central planning and regulation has had on primary school care and education, and what would have happened to female labour supply and school attendance if that sector had been treated in the same way as child care.

³⁵Australia now has a combined income tax and family tax benefit system that is effectively a system of joint taxation, as, for example, in the US and Germany. However, there is a difference. The system applies only to families with children and imposes additional penalties on two-earner families with dependent children under 5. Under this relatively new regime, married mothers who work can lose around half their earnings in taxes and reduced family payments, and many cannot meet the cost of formal child care out of their net incomes.

3.2.4 Long-term Saving and Short-term Borrowing

This argument is supported by the data on saving and borrowing and on housing available in the HES. These show, on the one hand, how much families must save under a mandatory system of superannuation. They also have an overwhelming incentive to invest in owner occupied housing if, ultimately, they are to buy housing over their life time at an affordable (and in fact very low) user cost. Table 4 lists contributions to superannuation and life insurance (column 1) and mortgage repayments of capital (column 2) by phase. When the sum of these is subtracted from saving, many households are found to be in the position of having to borrow short term to finance these forms of long term contractual saving. The median of the amounts they must borrow short term in each phase, calculated as the difference between saving (column 3 of Table 1) and the sum of mortgage repayments of capital and superannuation contributions, is shown in column 3.

Table 4 about here

The imperative to save for house purchase is indicated by the dramatic decline in debt to house price ratio from phases 1 to 8, shown in column 4, which is matched by a rise in the percentage of households who own their homes, from 59 per cent in phase 1 to 95 per cent in phase 8. It is straightforward to show that the user cost of owner occupied housing, obtained by discounting repayments of capital and the initial equity at the time of purchase, becomes negative over time, due primarily to capital gain in an imperfect capital market and also to low transactions costs relative to private rental.³⁶

3.3 Within-Phase Heterogeneity

The data show that there is a very high degree of heterogeneity in respect of female labour supply and savings behaviour across households with the same wage rates and demographics, which is concealed in the overall average figures considered above. The underlying idea in the models in section 2 is that households choose lifetime paths of male and female labour supplies,

³⁶The data suggest that, under these conditions, home ownership is analogous to an annuity with a very high rate of return, especially if households minimise transactions costs by rarely moving over the life cycle. The preferential tax treatment of owner occupied housing is also a contributing factor but cannot alone explain the differential between owning and renting over time if one assumes, implausibly, a perfect capital market.

saving and consumption of household and market goods, given wage rates (net of taxes), interest rates and productivities in household production. Differences in domestic productivities across households lead to differences in choices of these endogenous variables, for households facing the same net wage and interest rates and capital market conditions.

To give an indication of the empirical importance of this heterogeneity we construct life cycle profiles for two groups defined according to female labour supply as an indicator of domestic productivities. We are limited to this strategy for defining household types because of missing data on domestic output.

Ideally, we would like to distinguish between those households in which female labour supply is zero or "marginal"³⁷ throughout the life cycle, and those in which it is significant and relatively large over the entire life cycle. This categorisation requires panel data. Since we have access to cross section data only, we present profiles for a sample of "in work" households, with those in phases 2 to 7 partitioned into two groups of equal size according to the female partner's "usual hours of work". We label those in which the female partner is a non-participant or supplies relatively little market labour as "Type I: Traditional" and those in which she is employed full-time or works relatively long part-time hours as "Type II: Non-traditional".³⁸ Heterogeneity in female labour supply is strongly evident only after the arrival of children, and so we do not split phase 1.

The sample of "in work" households is selected on the criterion that the male partner reports positive hours of work. This yields a sample of 2992 records from the HES dataset, and gives cell sizes for phases 1 to 7 of 428, 468, 459, 456 458 435 and 288, respectively. In the results to follow, the data means and medians for phase 8 are also included, to give complete life cycle profiles. Given our data, we are limited to making life cycle comparisons between the two types on the basis of the assumption that means and medians in phases 1 and 8 are representative of both.

3.3.1 Time Allocations

Table 5 presents data means for time allocations across the eight phases in the same format as Table 2, but with separate results for type I and type II

 $^{^{37}}$ In the Heckman (1993) sense.

 $^{^{38}}$ Part-time employment status is defined as 1-35 hours of work per week and full-time as 35 hours of work or more per week.

households in phases 2 to 7. The means for female hours of work reveal a high degree of polarisation across these phases, a result we would expend since in the majority of type I households the female partner is a non-participant, and in an almost equally large proportion of type II households the female partner is employed full time.

Table 5 about here

3.3.2 Income and Saving

As for the average household, saving tends to track net income which in turn tracks female labour supply. As a consequence, the observed heterogeneity of female labour supply is strongly associated with differences in household net incomes and saving, as indicated in Table $6.^{39}$ Column 1 of the table lists the median net incomes of the two household types in each phase, and column 2 reports median saving net of mortgage repayments and superannuation contributions. The saving profiles are depicted in Figure 5. Holding wage rates constant, non-traditional households are found to have higher net incomes due to longer hours of work, and also to have much higher levels of saving.⁴⁰

Table 6 and figure 5 about here

The gap between the net incomes of the two household types is much narrower than between their earnings, due to the tax-benefit system. This becomes evident when we compare net incomes in column 1, household private incomes⁴¹ in column 3, and male and female earnings in columns 4 and 5. Data means for direct taxes net of benefits (direct and indirect) are listed in column 6. In phase 3, for example, the traditional household receives, on average, a net benefit of \$10,012 whereas the non-traditional household receives only \$600, a difference of almost \$9,500. In phases 2 and 4 type I households gain by over \$9,000 relative to type II.

³⁹Note that income generated by household production is inherently non-saveable, though to the extent that it substitutes for market consumption it may permit higher saving.

⁴⁰It may be argued that the gap between the saving of the two types is overstated by these results, because households may switch "type". For example, married women who work and save more in the early child rearing phases may become nonworkers in the later phases. Studies of the persistence of female labour supply indicate strongly that this is not the case (see, for example, Shaw, 1994).

⁴¹The ABS (2001) defines private income as income from all sources before tax and excluding government transfers.

These differences reflect the fact that, *ceteris paribus*, non-traditional households pay much more in income taxes, while receiving relatively little in family payments or compensation for the cost of child care.⁴² The differences become even greater when indirect taxes are included.⁴³ In addition, in the retirement phase, with higher levels of saving, non-traditional households are less likely to be eligible for the income tested age pension. In effect, non-traditional households save for their own retirement and contribute to financing transfers and pensions for traditional households, by working longer hours and paying higher taxes.

Under this type of regime, small differences in domestic productivities are likely to be sufficient to give rise to the considerable heterogeneity in female labour supply decisions that we observe.

3.3.3 Full Consumption and Leisure

Table 7 compares the consumption and leisure profiles of the two household types. Columns 1 to 3 list expenditures on market goods, domestic consumption and leisure⁴⁴, respectively. Traditional households have, on average, much higher expenditures on domestic consumption and leisure but lower levels of spending on market goods. The total consumption spending of the household, obtained as the sum of columns 1 to 3 and indirect government benefits, is shown in column 4. Although the household types have close to the same gross wage rates and non-labour incomes, the total consumption spending of type I is higher in all phases due to the tax-benefit system.

To make living standard comparisons, these expenditures need to be deflated by prices and demographics⁴⁵. Here we adjust for the latter by sub-

⁴²This highly unequal distribution of the tax burden between non-traditional and traditional households is a relatively recent phenomenon in Australia, and has been largely a consequence of reducing the overall progressively of the tax-transfer system, as in other OECD countries, notably the US. In effect, lower rates at the top of the distribution of income have been funded by raising taxes on working married women. It is important to see the issue in this context, and not in terms of a conflict between non-traditional and traditional households.

⁴³In all the child rearing phases non-traditional households effectively pay more than \$10,000 in taxes than traditional households when indirect taxes are included.

⁴⁴The leisure expenditures are computed for a time constraint of 14 hours per day. For further details, see section 5.1.

⁴⁵The average number of children of traditional household is slightly higher than that of

tracting child costs. The result, adult total consumption, is reported in column 5. Again, the traditional household is ahead. However, if leisure is omitted, to give adult full consumption shown in column 6, the profiles of the two types tend to converge.

Table 7 about here

3.3.4 Conclusions on Across-Household Heterogeneity

The data we have presented here show that heterogeneity is important. Variation in female labour supply, which we hypothesise is due to differences in household human and physical capital, is associated with significant differences in saving and the division of consumption between household and market goods, though not with wage rates or numbers of children. Both types of households of course are affected by the inability to use the capital market to smooth the time profile of leisure and full consumption. The tax/benefit system however, rather than correcting for this by supporting all households in the early child-rearing phases, simply brings about very large transfers from non-traditional to traditional households. The implied high marginal tax rates on working wives are clearly very questionable on efficiency grounds, while it is not a *priori* clear whether there are gains in equity of the income distribution. For this we would need to know exactly how the household productivity variations, created by variations in human and physical capital, are correlated with female labour supply across households, something about which virtually nothing is known empirically. Our own judgement is that policy changes to reduce the tax burden on working married women and increase support for families with pre-school children would significantly improve both efficiency and equity, as well as increase fertility.⁴⁶

4 Empirical Specification of the Models

To estimate the life cycle consumption choice models presented in section 2 on the data described in the preceding section, a number of simplifying assumptions are required. In both the HES and the TUS, as in all household

non-traditional households. Traditional households in the "in work" sample have an average of 1.71, 2.42, 2.13 and 1.79 children in phases 2 to 5, respectively. The corresponding figures for non-traditional households are 1.47, 2.07, 1.97 and 1.73.

⁴⁶See Apps and Rees (1999), (2004), for further analysis and discussion of these points.

survey datasets to the best of our knowledge, information on individual consumptions of market and domestic goods is missing. While the TUS provides data on adult leisures and child care, these alone do not allow the identification of individual preference parameters. Nor do they permit the estimation of the parameters of an intra-household sharing rule.⁴⁷ We therefore assign shares of full consumption between adults and children prior to estimation. We set the children' share to the child costs calculated as outlined above, and treat these costs as a lump sum transfer from parents. This leaves the within-phase leisure and market and domestic goods demands of parents for estimation as a two-adult household model.

Given that we are estimating a two-adult demand system, aggregation restrictions are required that are valid only if family members face the same prices. This condition is not satisfied if adult members face different prices (wage rates) for leisure. To deal with this problem, we specify the leisures as inputs to the production of a private leisure good, z, that can be consumed by either family member.

4.1 Within-period Demand System

In our view, changes in household preferences over the life cycle are likely to reflect the changing needs of children rather than changes in the preferences of the adults. Ideally, therefore, we would like to specify a system in which adult preference parameters are constrained to be identical across phases, with variation in within-period demands explained by exogenous prices and the within-period total consumption expenditure variable. However, with missing data on output, prices can only be set on the basis of some essentially arbitrary assumption on productivity. The approach we take here is to allow prices to depend on household specific production parameters as well as wage rates and, on the demand side, to introduce preference heterogeneity so that we can predict the data.

We specify Cobb-Douglas (CD) production functions for leisure, z, and the domestic good, y, as

$$s_z = \sigma_z + \xi_z \tag{20}$$

$$\sigma_y = \sigma_y + \xi_y \tag{21}$$

where $s_z = w_2 z_2/(w_1 z_1 + w_2 z_2)$ and $s_y = w_2 a_2/(w_1 a_1 + w_2 a_2)$. Within-period prices, q and p, are then computed for each record as functions of wage

 $^{^{47}}$ For a proof, see Apps and Rees (1997).

rates and the production parameters specific to each record, together with a scaling factor, consistent with the CD form. Using record specific parameters implies that unobserved domestic productivity is systematically related to the error term of the relevant production share equation. Thus, for example, households in which the female partner specialises in domestic work will, ceteris paribus, be found to have a larger production share, and therefore a lower domestic price, due to a higher domestic productivity.⁴⁸

We select the Almost Ideal (AI) demand system for estimation of withinperiod preference parameters. Suppressing the household type and phase subscripts, the indirect utility function for adult i, i = f, m, takes the form

$$u_i(q, p, c_i) = (\ln c_i - \ln a_i(q, p))/b_i(q, p)$$
(22)

where c_i is adult *i*'s total consumption expenditure. The price indexes $a_i(q, p)$ and $b_i(q, p)$ are given by

$$\ln a(q,p) = \alpha_0 + \alpha_z \ln q + \alpha_y \ln p + 0.5\gamma_{zz} \ln^2 q + \gamma_{zy} \ln q \ln p + 0.5\gamma_{yy} \ln^2 p \quad (23)$$

$$\ln b(q, p) = \beta_z \beta_y \ln q \ln p \tag{24}$$

where α_0 , α_j , γ_{jl} and β_j , j, l = x, y, z, are parameters and the α_j contain a dummy variable for the presence of dependent children and an error term capturing preference heterogeneity. The restrictions for adding up are $\sum \alpha_j = 1$, $\sum \beta_j = 0$ and $\sum \gamma_{jl} = 0$, for symmetry, $\gamma_{jl} = \gamma_{lj}$, and for homogeneity, $\sum \gamma_{jl} = 0$. Household demands in share form are

$$S_z = \alpha_z + \gamma_{zz} \ln q + \gamma_{zy} \ln p + \beta_z \ln(c/a(q, p)) + \varepsilon_z$$
(25)

$$S_y = \alpha_y + \gamma_{yy} \ln p + \gamma_{yz} \ln q + \beta_y \ln(c/a(q, p)) + \varepsilon_y$$
(26)

$$S_x = \alpha_x + \gamma_{xz} \ln q + \gamma_{xy} \ln p + \beta_x \ln(c/a(q, p)) + \varepsilon_x$$
(27)

where $S_x = x/c$, $S_z = qz/c$ and $S_y = py/c$, and $c = \sum c_i$, i = f, m. Given adding up, we need only estimate the share equations for leisure and the domestic good.

⁴⁸With this specification there is the potential for parameter bias due to the endogeneity of time allocations. However, with missing data on domestic output there is inevitably a trade-off between this problem and parameter bias due to omitted domestic price variables.

4.2 Intertemporal Demand System

As outlined above, we distinguish eight phases of the household's life cycle, which are defined, and can be broadly described, as:

- $\phi_1 = \{0, ..., \tau_1\}$: the two-person household has no children;
- $\phi_2 = \{\tau_1 + 1, .., \tau_2\}$: the children are of pre-school age;
- $\phi_3 = \{\tau_2 + 1, .., \tau_3\}$: the children are of primary school age;
- $\phi_4 = \{\tau_3 + 1, .., \tau_4\}$: the children are predominantly in the 10-14 year age group;
- $\phi_5 = \{\tau_4 + 1, .., \tau_5\}$: the household has older dependent children living at home
- $\phi_6 = \{\tau_5 + 1, .., \tau_6\}$: the children have left home and the male partner is under 55, or under 60 and not retired;
- $\phi_7 = \{\tau_6 + 1, .., \tau_7\}$: the adults are under 65, or over 65 and not retired; and
- $\phi_8 = \{\tau_7 + 1, .., T\}$: the adults are over 65, and are both retired.

We assume that within each given phase, the parameters of the utility functions, as well as the household welfare weights, remain constant, though they may change between phases. The subscript j = 1, ..., 8 will refer to the phase. Introducing the phase subscripts into the above indirect utility function we can write it as

$$u_t = \hat{a}_j(q_t, p_t) + \frac{\ln c_t}{b_j(q_t, p_t)} \quad t \in \phi_j, \ j = 1, .., 8$$
(28)

with

$$\hat{a}_j(q_t, p_t) \equiv \frac{-\ln a(q_t, p_t)}{b_j(q_t, p_t)}$$
(29)

The solution to the household's problem yields the life cycle profile of total income, and the estimated demand and labour supply functions within periods can then be used to derive profiles of market and domestic consumption, saving and secondary earner labour supplies, for the perfect and imperfect capital market models respectively.

4.2.1 Perfect Capital Market

Given the assumed functional form for indirect utility, the first order conditions for this problem in the perfect capital market case are

$$\frac{\rho^t}{\delta^t b_j(q_t, p_t)c_t} = \lambda \quad t \in \phi_j, \ j = 1, .., 8$$
(30)

$$\sum_{t=0}^{T} \delta^{t} c_{t} = W \equiv \sum_{t=0}^{T} \delta^{t} (A \sum_{i=1}^{2} w_{it} + P_{t})$$
(31)

where W is "full wealth". The important thing to note is that the marginal utility of total consumption expenditure in each period depends on the prices of the domestic good, p_t , and leisure, q_t , and therefore on the wage rates and the domestic productivity. Thus the entire time profile of total consumption, as well as its allocation within each period as between market and domestic consumption, depends on this productivity. The solution of the system is given very simply by

$$c_t = \alpha_t c_T \tag{32}$$

$$c_T = \frac{W}{\sum_{t=0}^{T-1} \delta^t \alpha_t + \delta^T}$$
(33)

$$\alpha_t \equiv \left(\frac{\rho}{\delta}\right)^{t-T} \frac{b_8(q_T, p_T)}{b_j(q_t, p_t)} \quad t \in \phi_j, \ j = 1, .., 7$$
(34)

4.2.2 Imperfect Capital Market

In principle, this problem could be fairly complicated to solve. However, from the data, we can establish that, at the margin, the average household is in equilibrium at the saving interest rate in phases 1 and 6 to 8, and at borrowing interest rates in phases 2 to 5. It can also be established from the data that the latter interest rates are higher than the former. Denoting the discount factors by $\delta(t, j)$, j = 1, ..., 8, t = 0, ..., T, we can use these to collapse the single period budget constraints into a lifetime wealth constraint, which we write as

$$\sum_{t,j} c_t \delta(t,j) = W \tag{35}$$

where wealth W is computed from the full income data and the discount rates. The household again maximises utility subject to this wealth constraint, yielding the first order conditions

$$\frac{\rho^t}{\delta(t,j)b_j(q_t,p_t)c_t} = \lambda \quad t \in \phi_j, \ j = 1,..,8$$
(36)

together with the wealth constraint. We then have to solve the equations

$$c_t = \hat{\alpha}_t c_T \tag{37}$$

$$c_T = \frac{W}{\sum_{t,j} \delta(t,j)\hat{\alpha}_t + \delta(T,8)}$$
(38)

$$\hat{\alpha}_t \equiv \frac{\rho^{t-T}\delta(T,8)b_8(q_T,p_T)}{\delta(t,j)b_j(q_t,p_t)}$$
(39)

for the optimal time path of total income.

5 Results

Equations (37) to (39) show how the optimal path of life cycle total income, or total consumption expenditure, c_t , $t \in \phi_j$, j = 1, ..., 8, depends on the marginal utility of consumption in each phase, as a function of the discount factors $\delta(t, j)$, t = 0, ..., T, and the price index $b_j(q_t, p_t)$. We have data that support our hypothesis that the majority of households in the child rearing phases, particularly those in the earlier phases, borrow short term at an interest rate above the lending rate, and we select discount rates consistent with this hypothesis. We compute the prices, p_t and q_t , in the price index as outlined in Section 4.1

5.1 Demand System Parameters

We estimate the within-period demand system on data for a more restricted sample of "in-work" households than that used in the empirical analysis of heterogeneity in section 3.3. The sample is selected on the additional criteria that the male partner's usual hours of work are equal to or greater than 25 per week, partners in work report earnings from wages/salaries as the primary source of income, and neither partner has negative earned or unearned incomes. The sample contains 2151 records.

To avoid parameter bias arising from the endogeneity of earnings, the system is estimated on wage rates, net of taxes, predicted from regression models corrected, in the case of the female partner, for selectivity. Full income is defined on the basis of a time constraint of 14 hours per day, which means that each adult is given a fixed allocation of 10 hours of "own time" (pure leisure and/or sleep), with the residual of leisure time beyond own time being treated as an input to the general leisure good, z. Thus, total consumption within each phase is the sum of the household's expenditure on market consumption and on the domestic and leisure goods with own time omitted in the latter.

We estimate the system on all records, ignoring corner solutions, on the assumption that domestic work is analogous to a particular type of employment. Under this assumption, corner solutions are potentially a general problem, arising in respect of both market and domestic work choices. Dealing with the issue here is outside the scope of the present study. Table 8 reports the parameter estimates of the system. All are significant at well above the 5 per cent level. The intercept term, α_0 , is set at log(20,000). The cost function is concave at data means.

Table 8 about here.

5.2 Intertemporal Profiles of Consumption

The approach we apply to evaluate the alternative capital market models is to see how well the life cycle profiles of total and full consumption generated by the data can be predicted, using the estimated preference parameters and selected discount rates. The first step is the construction of a reference adult total consumption profile. For this we compute c_t , $t \in \phi_j$, j = 1, ..., 8, using data means for time allocations, wage rates and the tax-transfer system in each phase, and median saving in phases 1 to 7. Consumption in phase 8 is obtained by compounding up previous saving/borrowing at the relevant interest rates. We then compute reference consumption and leisure profiles using the parameters of the within-phase production and demand system.

We select a relatively low real lending rate of a quarter of one per cent, which we consider plausible for the average lender whose capital income is taxed at a relatively high marginal rate. Compulsory superannuation payments and the capital component of mortgage repayments are treated as exogenous amounts subtracted from income in the preretirement phases (i.e., as taxes), compounded up and added back into P_t in phase 8. We compound up contributions to superannuation at a real rate below the lending rate, of 0.1 per cent, to take account of the risk associated with choice of fund (which is now becoming increasingly evident) in addition to transactions costs. In contrast, the capital component of mortgage repayments is compounded up at 1.0 per cent, to make an adjustment for the high rate of return to owner occupied housing. The results for the imperfect capital market model are derived for real borrowing rates of 1.4, 1.6, 1.4 and 1.1 per cent in phases 2 to 5, respectively. Though these rates may appear low in absolute terms, they are proportionately much higher than the lending interest rate. It is also important to keep in mind that they represent across-household averages, and that there is considerable heterogeneity across households in saving/borrowing behaviour within each phase, particularly in the early phases. Most households will in fact face either much higher or lower rates.

We compare reference profiles constructed in this way with the predictions of the models based on the same data means for time allocations, wage rates, taxes and benefits, and the saving profiles generated by the models. The preference parameters, β_j , j = z, y, x, are applied to obtain c_t for each model (as set in (32) to (34) and in (37) to (39), respectively) and the full set of parameters are then used to predict within-period consumption and leisure expenditures. The costs of children computed for the reference case are held constant across models.

Table 9 presents the reference profiles for 2-adult total consumption and 2-adult full consumption, in columns 1 and 2, respectively. Column 3 lists household full consumption in each phase, computed as the sum of that of the adults and the transfer they make to the children. Table 10 reports corresponding profiles for the perfect capital market model, and Table 11, for the imperfect capital market model.

Tables 9, 10 and 11 about here

The profiles of 2-adult total consumption, c_t , are compared graphically in Figure 6. The reference profile is strongly U-shaped across the phases in which dependent children are present, as we would expect from the evidence in Section 3.2. The imperfect capital market model predictions match very closely those of the reference profile. In contrast, the perfect capital market model predicts that the household will smooth total consumption expenditure on an adult per capita basis.

Figure 6 about here

Figures 7 and 8 show graphically the full consumption profiles predicted by the two models. Figure 7 illustrates some of the implications of the perfect capital market hypothesis. Because the model generates a relatively smooth profile of 2-adult full consumption, adding in the costs of the children's full consumption gives a more strongly humped profile of household full consumption across phases 2 to 5 than indicated by the data. In other words, evidence of a more humped profile of household full consumption across the phases in which children are present is required in order to support the perfect capital market hypothesis. The imperfect capital market model clearly predicts the data far better.

Figures 7, 8 about here

6 Conclusions

Our descriptive picture of a household's life cycle time allocation, income and consumption, defined in terms not of calendar years, but of key phases in the evolution of the family, helps resolve some of the "puzzles" that have been noted in the existing literature, but suggests a new one: Why, in the phase in which the household has pre-school children, are there such dramatic changes in time allocations, consumption and saving? The data on borrowing and interest rates suggest that the standard assumption of a perfect capital market is untenable, but so is the hypothesis that households do not borrow short-term. By modelling household life cycle choices under respectively perfect and imperfect capital markets, we show that in the former case we cannot reasonably explain the data, in the latter case we can.

More generally, we are proposing an approach to life cycle saving and consumption behaviour, which sees the endogenisation of the income process via female labour supply choices, as essential. Although in this paper we have found it useful to ignore uncertainty, we certainly would not want to claim that it is unimportant in reality, nor that a precautionary motive may not be operative in those phases in which households save. Indeed we would see that as an interesting direction in which to develop the model we have presented here.

Our results have interesting implications for public policy, at a time when declining fertility is seen as the major cause of population ageing and consequential problems in sustaining social security programmes, such as Pay-As-You-Go pension systems. Greater support for households during the critical early childhood phase could help overcome the problems presented by an imperfect capital market and reduce the costs of having children. This should be a fruitful area for future research.

References

- [1] Australian Bureau of Statistics (2001), Government Benefits, Taxes and Household Income, Cat No 6537.0, ABS, Canberra.
- [2] Ando, A. and F. Modigliani, (1963), The "Life Cycle" Hypothesis of Saving: Aggregate Implications and Tests, *American Economic Review*.
- [3] Apps, P. F., and R. Rees, (1997), Collective Labor Supply and Household Production", *Journal of Political Economy*, 105, 178-190.
- [4] Apps, P. F., and R. Rees, (1999), On the Taxation of Trade Within and Between Households, *Journal of Public Economics* 73, 241-263.
- [5] Apps, P. F., and R. Rees, (2001), Household Saving and Full Consumption over the Life Cycle, Discussion Paper 280, IZA, Bonn.
- [6] Apps, P. F., and R. Rees, (2002), Household Production, Full Consumption and the Costs of Children, *Labour Economics*, 8, 621-648.
- [7] Apps, P. F., and R. Rees, (2004), Fertility, Female Labour Supply and Public Policy, *Scandinavian Journal of Economics*, forthcoming.
- [8] Attanasio, O.P. and J. Banks (1998), Trends in Household Saving: A Tale of Two Countries, Working Paper Series No W98/15, *Institute of Fiscal Studies*, London.
- [9] Attanasio, O.P., J. Banks, C. Meghir, and G. Weber (1999), Humps and Bumps in Lifetime Consumption," *Journal of Business and Economic Statistics*, 17, 22-35.
- [10] Attanasio, O. P. and M. Browning, (19995), Consumption over the Life Cycle and Over the Business Cycle, *American Economic Review*, 85, 1118-1137.
- [11] Attanasio, O. P., H. Low and V Sanchez-Marcos (2003), Explaining Changes in Female Labour Supply in a Life-Cycle Model, mimeo.
- [12] Banks, J.,R. and R. Blundell (1994), Household Saving Behavior in the United Kingdom, in J. Poterba (ed.), *International Comparisons* of Household Saving, Chicago: Chicago University Press.

- [13] Baxter, M., and U. J. Jermann, (1999), Household Production and the Excess Sensitivity of Consumption to Current Income, *American Economic Review*, 89, 4, 902-920.
- [14] Becker, G.S., (1976), The Economic Approach to Human Behavior, University of Chicago Press, Chicago/London
- [15] Blundell, R., M. Browning and C. Meghir (1994), Consumer Demand and the Life-Cycle Allocation of Household Expenditures, *Review of Economic Studies*, 61, 57-80.
- [16] Browning, M and TF Crossley, (2001), The Life-cycle Model of Consumption and Saving", Journal of Economic Perspectives, 15, 3-22.
- [17] Browning, M and A Lusardi, (1996), Household Saving: Micro theories and Micro Facts, *Journal of Economic Literature*, XXXIV, 1797-1855.
- [18] Browning, M. and M Ejrnaes, (2002), Consumption and Children. CAM Discussion Paper 2002-06, University of Copenhagen
- [19] Carroll, C.D., (1992), The Buffer Stock Theory of Saving: Some Macroeconomic Evidence, Brookings Papers on Economic Activity, 61-156.
- [20] Carroll, C.D., (1994), How Does Future Income Affect Current Consumption?, The Quarterly Journal of Economics, CIX, 111-148.
- [21] Carroll, C. D., (1997), Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis, The Quarterly Journal of Economics, CXII, 1-55.
- [22] Carroll, C.D., (2001), A Theory of the Consumption Function, With and Without Liquidity Constraints, *Journal of Economic Perspectives*, 15, 3, 23-46.
- [23] Deaton, A., (1991), Saving and Liquidity Constraints, *Econometrica* 59, pp1221-48.
- [24] Deaton, A., (1992), Understanding Consumption, Clarendon Press, Oxford.
- [25] Friedman, M., (1957), A Theory of the Consumption Function, Princeton University Press for NBER, Princeton NJ.

- [26] Gourinchas, P-O., and J.A. Parker (2002), Consumption over the Life Cycle, *Econometrica*, 70, 47-89.
- [27] Heckman, J., (1974), Life Cycle Consumption and Labor Supply, American Economic Review, 64, 188-194
- [28] Keynes, J. M., (1936), A General Theory of Employment, Interest and Money, MacMillan, London.
- [29] Shaw, K., (1994), The Persistence of Female Labor Supply: Empirical Evidence and Implications, *Journal of Human Resources*, 29, 348-378.
- [30] Zeldes, S. P., (1989), Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence", *Quarterly Journal of Economics*, CIV, 275-98.