Consumption and Children.

Martin Browning
Institute of Economics
University of Copenhagen
Martin.Browning@econ.ku.dk

Mette Ejrnæs
Institute of Economics
University of Copenhagen
Mette.Ejrnes@econ.ku.dk

January 2000

Abstract

Several recent papers have concluded that we need to allow for a precautionary saving motive to reconcile data on lifetime patterns of consumption and income with a standard optimising model. In this paper we contest this and show that if we take proper account of the presence of children then we have consumption 'tracking' income exactly as in the data.

1 Introduction.

For over fifty years there has been a recognition that the life-cycle (or low frequency) relation between household consumption and income is of prime importance in the determination of aggregate saving and economic growth. There are three elements to this: allocation between a period of human capital formation and later; allocation within the working life and allocation between the pre-retirement and retirement period. In this paper we shall be concerned with the second of these. Considering consumption, there is widespread agreement that household consumption over the 'working life-cycle' displays an inverted U-shape (see, for example, Thurow (1969), Browning, Deaton and Irish (1985) and Carroll and Summers (1991)). In figure 1 we present smoothed consumption against age for our UK FES data from 1968 to 1995 which displays the familiar pattern[1]. The same sources also show that household income displays a very similar pattern to consumption over the working life. Any contribution to understanding why consumption

[1] We shall present a thorough description of our data below.

and income have such a high life-cycle correlation is of critical importance in many policy debates[2].

There have been five broad responses to the observed 'tracking' of consumption and income over the life-cycle. The first response is that this is evidence that households use some '*rule of thumb*' for consumption that sets it close to current income (the simplest rule being that households simply spend all they earn in the planning period). In the current context this is widely rejected as an explanation since to be valid it has to be true of almost everyone to give the observed mean correlations (see Carroll and Summers (1991), section 10.6 for an elaboration of this argument). Moreover, his explanation is inconsistent with any standard optimising model of intertemporal allocation that allows for forward looking agents (although it must be said that this makes it more attractive for many researchers). Although the observed patterns are inconsistent with the simplest standard model with quadratic preferences and perfect capital markets most investigators are reluctant to abandon the standard framework altogether and the other four responses all involve variants of the standard model.

Thurow (1969) suggested that households are impatient and *liquidity constrained*. Nagatani (1972) showed that even without liquidity constraints, the presence of income uncertainty and the coincidence of a high discount rate, a utility function with a positive third derivative and income growth over the early part of the life-cycle gives a *precautionary motive* which induces a high correlation between consumption and income over simple simulated life-cycles. Heckman (1974) goes one step further and shows that even without liquidity constraints or uncertainty, *non-separabilities between consumption and labor supply* can lead to the observed patterns. If consumption and labor supply are Frisch complements (because of the costs of going to work and the possibility of substituting market goods for home production) then consumption and income will move together over the life-cycle. Moreover both will display an inverted U-shape if the pattern of 'anticipated' (discounted) wages over the life-cycle is inverted U-shaped[3]. The final response to the inverted U-shape for consumption is that the path of *demographics* over the life-cycle display similar patterns to that of consumption so that if house-

[2] Many issues of cyclical policy turn on the high frequency correlation between consumption and income; we shall not address this in the paper.

[3] This intertemporal substitution explanation argument is closely related to the explanation of cyclical variations in employment being due to cyclical variations in wages; see Lucas and Rapping (1970) for the original analysis and Blundell and MaCurdy (1999) and Browning, Hansen and Heckman (1999) for recent surveys of the evidence. In both cases it seems that implausibly high elasticities of intertemporal substitution for labor are required to reconcile the theory with the evidence.

holds increase consumption when children are present then the associated pattern is consistent with a life-cycle model with no liquidity constraints. Tobin (1967) was the first to incorporate realistic patterns of demographics into simulated life-cycle allocation models (see, in particular, Tobin's figure 5). The first formal incorporation of demographics in micro-estimation of intertemporal consumption relationships is due to Browning, Deaton and Irish (1985). Attanasio and Weber (1995), Blundell, Browning and Meghir (1994) and Attanasio and Browning (1995) all argue that using estimation based estimates of the impacts of children on consumption we remove most of the inverted-U shape for consumption.

The relative importance of these different factors is still a matter of considerable dispute. For example, Carroll and Summers (1991) consider and discount the non-separability argument and argue for a form of precautionary motive (the 'buffer stock' model of Deaton (1991)) whilst remaining agnostic on the importance of demographics. Carroll (1994) and Hubbard, Skinner and Zeldes (1994) present evidence based on U.S. data that income processes estimated from micro-data and a precautionary motive can lead to the observed inverted U-shape for consumption without any need to account for demographics. Attanasio, Banks, Meghir and Weber (1999) use a simulation model with parameter estimates from U.S. quasi-panel consumption data and find that allowing for family size gives a peak in consumption at the same age as observed in the data. They also find, however, that income uncertainty and a precautionary motive is needed to match the observed ratio of peak consumption to consumption at age 25. In a closely related paper, Gourinchas and Parker (1999) argue that whilst accounting for family size can go some way to removing the 'excessive' correlation between consumption and income we also need to introduce some precautionary motive.

One common conclusion in all these papers is that some form of precautionary motive is needed to explain the coincidence of income and consumption, particularly in the early part of the life-cycle. In the analysis developed below we contest this and argue that if we take proper account of the effects of the numbers and ages of children then there is no need to introduce uncertainty. We stress that we do not thus claim that uncertainty is unimportant but merely that it is not necessary to rationalise the observed low frequency mean data.

To properly account for children we need to construct the counter-factual: how much would households that have children have spent if they never had children? We also examine the closely related question of the relationship between children and income paths. Finally, we examine the consumption paths of a sample of 'married' households who do not currently have children. In any period this sample is comprised of three groups: those who never have children, those who will have children in the future and those who have had children who have now left home. If we could follow the first group through time then we could investigate directly whether consumption 'tracks' income independently of children. Unfortunately long panels with good consumption information are not available. The best we can do is to construct a sample of households that do not currently have children present in the household and then to adjust for the fact that some of these will have or have had children. We present techniques for doing this and the consequent adjusted paths of consumption.

2 Constructing counter-factual consumption.

We present here a basic model of lifetime consumption, income and children. It is important to stress the nature of the 'thought experiment' that we shall be undertaking. We shall derive an optimal consumption program in a *very* simplified context with perfect foresight, perfect capital markets and exogenously given children. We then ask, how well does this simple model fit the known facts concerning mean consumption, the timing and spacing of children and income over the life-cycle? We find that it does a very good job of 'explaining' the data. Thus our 'null' hypothesis is that children account for *all* of the observed inverted U-shape of consumption.

We assume that households have perfect foresight and they face perfect capital markets so that the path of income over time is irrelevant for the consumption allocation decision. The household lives for T years and has either zero or one child (in the empirical section below we extend this to allow for different numbers and ages of the members of the household). Let π be the proportion of households that have a child at some time and p_t be the proportion of households that currently have a child present (so that $p_t \leq \pi$). If household h has a child at some time then the child is in the household for d_h periods (where $d_h < T$). We also assume that the discount rate and interest rate are equal so that households that never have children have a constant consumption level, denoted c^0. To capture the effects of children on consumption we need to specify the within period utility function. Let z be a binary variable denoting whether there is a child currently present in the household. We assume that:

$$u(c, z) = \nu\left(c \exp(-f(z))\right)\exp(f(z)) + \Psi(z) \tag{1}$$

with the normalisation $f(0) = \Psi(0) = 0$[4]. The first order conditions (Euler equation) for intertemporal allocation implicitly defines the consumption

[4] The exponenetial form is convenient since we shall be working with log consumption.

when a child is present $\hat{c}$ with respect to the consumption when there is no child present, here denoted c^1. Formally:

$$\begin{aligned} u_c(\hat{c}(c^1, z), z) &= u_c(c^1, 0) \Rightarrow \\ \nu'(\hat{c}(c^1, z)\exp(-f(1))) &= \nu'(c^1) \end{aligned} \tag{2}$$

If $\nu(.)$ is strictly concave then this implies:

$$\ln \hat{c}(c^1, z) = f(1) + \ln c^1 \tag{3}$$

so that log consumption when children are present is an additive function of the children function and log consumption when the children are not present. All empirical models of the effects of children on consumption take this form since only the additive form is tractable in an empirical Euler equation analysis. The utility function in (1) is chosen to give this form (and is the only functional form consistent with the additive form). If a household h has a child and consumes c^1 before and after the child then they consume $\exp(f(1))c^1$ when the child is present, so that $\exp(f(z))$ is a (multiplicative) adult equivalence scale. It is usual to require that $f(1) > 0$ but there is no logical necessity for this. It might be, for instance, that parents choose to spend more when children are *not* present (for example, on going out or having expensive holidays). As we shall see below this issue becomes important when we allow for the dependence on the children's age.

Finally we allow that the lifetime incomes of the two types of household may differ because children have a direct effect on income. We denote by Δ the ratio of lifetime income if the household does not have a child to the lifetime income if they do. Putting all this together, we have the following relationship between the consumption levels of households that never have children (c^0) and the consumption of child households that do not have children present (c^1):

$$c^0 = \Delta \left(1 + (\exp(f(1)) - 1)\frac{d_h}{T}\right) c^1 \tag{4}$$

This relationship is the basis of the derivation of counter-factuals below.

The important features of all this are that households that have children at some time spend more when children are present than when they are

For simplicity we assume here that this scale is independent of the child's age but in the empirical analysis below we allow for the dependence of the scale on the ages of any children in the household. Note, as well, that since z is binary we could simply use $f(z) = kz$ but we shall use a more general formulation below when we allow for many children of different ages.

not and that households that never have children may consume a different amount than otherwise similar 'children' households that do not currently have children present. The latter may result from the differential spending paths or from the two sorts of households having different lifetime wealth (in mean). To model these features at a household level is a formidable undertaking, particularly when we allow for the possibility of many children. First we need a model of the timing and spacing of births of children and a model of children leaving home to construct the analogue of the proportion of the lifetime that children are present, d_h/T. Second, we need a model of the effect of children on lifetime income. Although there are some estimates of the impact of childbearing on women's earnings in the literature (see, for example, Calhoun and Espenshade (1988) for the U.S. and Joshi (1990) for the U.K.) we are still a long way from having reliable estimates of the impact on household income. This includes not only the impact on women's earnings but also the possibility, due to specialisation, that men's earnings may increase if they and their partner have more children. Finally we need an adult equivalence (ae) scale $f(z)$ (that allows for different numbers and ages of children in the data). We shall show below that at an aggregate (or cohort) level much weaker informational requirements are required.

To illustrate our objective we present some simulations based on this simple model. We take $T = 40$ and $d_h = 18$ for all households that have children. The adult equivalence scale is set to $f(1) = \ln(1.25)$. We simulate a population of $4,000$ households all with the same birth date ('from the same cohort'), of whom 80% have one child. For those who have a child, the age at the birth of the child is taken to be a Beta distribution on $[1, 20]$ rounded to the nearest integer. The Beta parameters used are $a = 2$ and $b = 3$ so that the resulting distribution is uni-modal but slightly skewed. Given our assumptions, no households have children in the final two periods of the lifetime. We set the interest rate to zero and assume that households that do not have children have lifetime income that is 10% above the lifetime income of households that do have children; that is $\Delta = 1.1$. Thus we have:

$$c^0 = \Delta \left(1 + (\exp(f(1)) - 1)\frac{d_h}{T}\right) c^1 = 1.155c^1 \tag{5}$$

Thus households that never have a child have expenditures that are 15.5% higher than 'child households' who do not currently have a child present.

In our empirical work below we adopt two sampling schemes. First, we take means across the whole sample in each period; we refer to this as the 'unconditional mean' sampling scheme. To approximate this sampling scheme in our simulated data we draw, without replacement, a sample of size 100 in each period t (so that each of our $4,000$ households is sampled exactly once) and record log consumption ($\ln c_{ht}$) and a dummy for the presence of a child (z_{ht}) for each household. The other sampling scheme we use is to take means for the sub-sample of households that do not currently have any children present; we refer to this as 'no child present' sampling scheme. In figure 2 we display the unconditional mean paths of average log consumption[5] (averaged over the 100 households in the sample at each age) and the proportion of households having children from one run of our simulation. Two features emerge from this. First, there is a pronounced hump shape to consumption that mirrors the shape of the path of children even though the underlying allocation scheme has consumption being flat in the distinct demographic regimes. The second feature is that there is a high frequency correlation between consumption and children (that is, there are coincident 'spikes' in the children and log consumption series). This is due to our ('repeated cross-section') sampling scheme: if, in a particular period, we happen to sample a high proportion of households with children (relative to the population at that age) then average consumption will also be higher than the population mean. In our econometric analysis below we shall take account of this. Our goal for the 'unconditional mean' sampling scheme is to find the correct way to adjust consumption so that the adjusted path is flat.

Turning to the 'no child present' sampling scheme, we note that the composition of this sub-sample changes over time since at early and late ages the sample includes almost everyone whereas around 'middle' age (about 20 in our simulations) the sample consists mostly of households that never have children. In figure 3 we graph the lifetime path of sample means from one simulation. The hump shape in this figure is more pronounced than for the previous figure. This is because the samples at the younger and older ages contain proportionately more 'children' households who not only have lower consumption when children are not present but also have lower lifetime wealth. One feature of this figure is somewhat misleading in that the variance of the mean over age does not vary much. In fact, there are far fewer households in the sub-sample around age 20 so the variance should be higher. In our simulations, however, all 'no child' households have the same consumption so this reduces the variance. In our empirical analysis below, in which different 'no child' households have different levels of lifetime wealth, the series analogous to figure 3 is much noisier in the middle ages. Once again, our goal is to find some way of adjusting consumption so that the 'no children mean' path is constant. We shall show that not only is this possible but it requires quite different information than for the adjustment for the unconditional mean path. Thus for the two sampling schemes we have

[5] The scale of log consumption has been changed for convenient graphing.

similar objectives (constructing a 'flat' adjusted path) but the informational requirements are quite different.

We are now in a position to derive our desired counter-factuals. The first sampling scheme is to take means over the whole sample in each period. This sample consists of three distinct groups. First, a proportion $(1-\pi)$ never have children and have log consumption $\ln c^0$. The second group who have children present have log consumption equal to $(f(1)+\ln c^1)$; the proportion of such households is p_t. Finally, a proportion $(\pi - p_t)$ of households who will have children but for which the children are not currently present, have consumption $\ln c^1$. Thus mean log consumption in period t is given by:

$$\begin{aligned} E_t(\ln c_{ht}) &= (1-\pi)\ln c^0 + p_t(f(1)+\ln c^1) + (\pi - p_t)\ln c^1 \\ &= \left\{(1-\pi)\ln c^0 + \pi \ln c^1\right\} + f(1)p_t \end{aligned} \tag{6}$$

Define the time invariant variable 'adjusted mean log consumption' by:

$$\ln \hat{c} = E_t(\ln c_{ht}) - f(1)p_t = \left\{(1-\pi)\ln c^0 + \pi \ln c^1\right\} \tag{7}$$

Thus the counter-factual level of log consumption, $\ln \hat{c}$, is the population weighted mean of the levels of consumption when children are not present. To construct this counter-factual we need p_t which is given by the data and the ae scale $f(1)$ which can be consistently estimated from the data at hand; details are given in the next section.

Our second sampling scheme takes means over households with no children currently present. This is made up of those who never have children with proportion $(1-\pi)/(1-p_t)$ and 'children' households who do not have children present who are in the proportion $(\pi - p_t)/(1-p_t)$. Thus the period t conditional mean log consumption is:

$$\begin{aligned} E_t(\ln c_{ht} \quad | \quad z_{ht} = 0) &= \frac{(1-\pi)\ln c^0 + (\pi - p_t)\ln c^1}{(1-p_t)} \\ &= \ln c^0 + m_t(\ln c^1 - \ln c^0) \end{aligned} \tag{8}$$

where $m_t = (\pi - p_t)/(1-p_t)$. In the econometric section below we show how to estimate the proportion m_t from available data and how to use this to estimate the time invariant quantity $(\ln c^1 - \ln c^0)$. Given such estimates we can construct the time invariant 'adjusted conditional mean of log consumption' as:

$$\ln \tilde{c}_t = E_t(\ln c_{ht} \mid z_{ht} = 0) - m_t(\ln c^1 - \ln c^0) = \ln c^0 \tag{9}$$

In this case the counter-factual level of consumption is the (mean log) consumption of households that never have children. Note that to make this adjustment we do not need the ae scale (as in equation (7)) but we do need information on completed fertility and the differences in consumption levels.

This exercise shows that, at the cohort level, there are two ways of deriving a counter-factual consumption that should be constant in a life-cycle model. If we use a sample of all households we need to know the ae scale and the proportions of households who currently have children. If, instead, we consider a sample of those without children currently present then we need to know $(\ln c^1 - \ln c^0)$ and the complete fertility of households which at present may not have all their children present. These corrections cannot be performed on the household level unless we have information on current household fertility, completed fertility, when children leave home, how children impact on lifetime income and an ae scale for each household. By taking (cohort) means, many of these factors become time invariant values that are subsumed into the desired counter-factual. Thus the informational requirements at the cohort level are much less demanding than at the individual household level.

3 Econometric issues

3.1 Estimating ae scales.

This sub-section discusses how to estimate the ae scale from observed repeated cross-section data on consumption. We begin by showing that the ae scale *cannot* be estimated on the basis of individual cross-section observations. Continuing with our one-child model, consider a sample of households, all observed at the same age t. The difference in log consumption between households with a child and those (currently) without a child is (using equation (8):

$$\begin{aligned} E_t(c_{ht}|z_{ht} &= 1) - E_t(c_{ht}|z_{ht} = 0) \\ &= \left(\ln c^1 + f(1)\right) - \left(\ln c^0 + m_t(\ln c^1 - \ln c^0)\right) \\ &= f(1) + (1 - m_t)(\ln c^1 - \ln c^0) \end{aligned} \tag{10}$$

This expression shows clearly that the cross-section estimate of the ae scale (the difference in the conditional means) is downwards biased unless $m_t = 1$ or $\ln c^1 = \ln c^0$. Neither restriction is plausible: the first implies that everyone has a child at some time ($\pi = 1$, see the definition of m_t after equation (8)). The second restriction implies that childless households have lower lifetime income than child households ($\Delta < 1$, see equation (4)). Moreover, the bias varies with age, being largest when the proportion of households who never will have children is large. Since the bias changes with age in a nonlinear way, using first differences will not eliminate the bias. We shall now show how we can obtain consistent estimates of the ae scale using quasi-panel data.

Consistent estimation of the ae scale is based on equation (6). This indicates that a simple regression of mean log consumption on the proportion of households with children will identify the ae scale as the coefficient on the latter. In practice, we regress first differences of mean log consumption on first differences in changes in demographics to remove any cross-section correlation between the numbers of children and the level of lifetime income. That is, we allow that the selection into the 'child' or 'no child' group is correlated with tastes over the difference $(\ln c^1 - \ln c^0)$. We also have to take account of the fact we have a time series of sample means so that we have a spurious correlation between consumption and children due to the sampling; as discussed in the last section a higher proportion of children households will lead to a higher level of observed mean consumption. Two solutions suggest themselves. First we could use outside information (for example, from the census) on the proportion of households in the cohort that have children in any period as an instrument for the proportion with children in our data. An alternative, which we adopt, uses only the data to hand. This uses the levels of the proportions of children in the relevant cohort in periods outside the current one as instruments for the current cohort proportion. These should be good instruments in the sense that the sample means predicted from periods around the current one are likely to be quite accurate but their measurement error (induced by the sampling procedure) will be uncorrelated with the measurement error on the current proportion.

3.2 Estimating completed fertility

Turning to our second sampling scheme, we see from equation (9) that we need estimates of the proportion of currently childless households that never have children, m_t. Since p_t can be observed, this requires an estimate of π, the proportion of households that have children at some point. Again we have two choices, namely using outside (census) information or using the data to hand and once again we choose the latter option. When using the repeated cross-section data, we have to impose some assumptions to be able to identify π. The details of the estimation procedure will be discussed further in the empirical section (where we allow for different parities) but here we simply note that the basic identification assumption is that at some age we observe the completed fertility in the household. Suppose we have a consistent estimate of π and hence of m_t. Referring to equation (8) we see that a regression using currently childless households of (the first difference) of mean log consumption on the (first difference of) m_t has a coefficient which can be interpreted as the mean (over the population) of $(\ln c^1 - \ln c^0)$. Thus a by-product of our estimation procedure is an estimate of the relative expenditures of 'childless' and 'child' households when no children are present. As mentioned above, in our empirical work we allow for different levels of completed fertility so this allows us to generate estimates of these ratios for different parities.

4 The data.

4.1 Sample selection.

The data used in this study come from the U.K. Family Expenditure Survey (FES). The FES contains information on expenditures on different consumption items, income and household characteristics. In particular, the detailed information about the household composition such as the number and the age of the children make these data attractive from our point of view; the U.S. CEX does not give details on children's' age and this, as we shall show, leads to serious problems for the analysis. In addition, the FES has been conducted since 1968, allowing us to follow some birth cohorts for more than 25 years. The sample used in this paper covers the period 1968 to 1995. As usual in studies of household consumption we limit ourselves to households consisting of married couples. The potential bias introduced by the selection may be quite important in the current context. If the members of high income households marry later then we would see an increasing path in mean consumption and mean income at early ages, even if every household held consumption constant over time. The sign of the correlation between being married and life time income is ambiguous. In studies of wages, a positive marriage premium is often found, which suggests a positive correlation. In Van Der Klaauw (1996), it is shown that the marital status of women strongly depends on the her own and (potential) husband's earnings. However, the effect is ambiguous; an increment in woman's earnings will increase the age of marriage and increase the risk of divorce while an increment in the husband earnings increase the predicted number of years of marriage. Generally, there is a distinct lack of theoretical and empirical research on consumption before, during and after marriage and we shall not have anything to add here. We do note, however, that the sample selection here works 'against us' in the sense that we are trying to find corrections for consumption that do not take account of the marriage selection that nevertheless give a flat path. We also restrict attention to households in which the wife is aged 20 or more and stratify on education groups to mitigate the selection effects. However, we acknowledge that conditioning on educational attainment might not totally eliminate the problem of sample selection, especially for the more educated group, because the educational attainment varies a lot within this group. To properly deal with the selection bias we need a model of the relationship between being married and lifetime income but we leave this to future work.

The measure of consumption used in this paper is total consumption, which is defined as nominal expenditure on all goods purchased during a two-week period. In the following the nominal expenditure is deflated by the consumer price index for total expenditure. Furthermore, to remove year effects the consumption is adjusted by using the residuals from a regression of household consumption on year dummies.

4.2 Constructing cohorts

In the previous section, we showed that constructing counter factual consumption is really only feasible using cohort mean data. The cohorts are based on year of birth and education but contrary to most other studies the cohorts are defined on the basis of the wife's birth year and the husband's education[6]. The main reason for using the birth year of the wife instead of the husband is that we want homogenous groups in terms of consumption and household composition, and the number and the age of the children are more correlated with the age of the wife than with the age of the husband. As we argued in the theoretical section, the number of children is an important determinant in consumption and therefore making homogenous cohorts in terms of household composition might also contribute to homogenous cohorts in terms of consumption[7]. For constructing birth cohorts we use five-year bands where the wife in the eldest cohort is born between 1910-1914 and the youngest cohort is born in 1965-1969.

Another important aspect in constructing homogenous groups is educational attainment. As is well known (see, for example, Attanasio *et al.* 1999 for the U.S.) there are strong differences in the number and timing of children across education groups. Educational attainment also seems to be important

[6] A more satisfactory solution would be to condition on the birth dates and education of both husband and wife. For example, one cohort would be "the wife was born in 1946 and had twelve years of education and the husband was born in 1948 and had ten years of education". Since this will result in rather small cell sizes we could use some non-parametric smoother across cells. A version of this is implemented in a compananion paper, Browning and Ejrnæs (1999).

[7] A variance analysis confirms that birth cohorts based on the wife crossed with year dummies capture more of the variation in household consumption than birth cohorts, based on the husband's birth year, crossed with year dummies.

for the selection into marriage. We control for educational differences by dividing the households according to the educational attainment of the husband since only the age at which the husband stopped education is reported in the data. In order to keep a sufficiently large number of households in each cohort the sample is split into two education groups: those who have the official minimum level of education and those with more than the official minimum level of education. Unfortunately, information on educational attainment is only available after 1977 which limits the number of observations in the synthetic panel of cohort means. In order to be able to use the data from the period 1968-1977 we introduce imputed education.

4.3 Imputed education

The underlying idea in dealing with the non-observation of education in the years 1968 to 1977 is to use household characteristics which are available in the entire sample period to predict the probability that the husband belongs to the more educated group. Based on the period 1978-1995 we estimate a logit model for the husband's education, where the explanatory variables contain information on occupation, household tenure, region of residence and employment status of husband and wife and the age difference between husband and wife. The logit model is estimated on a sample of households where the husband is aged between 20 and 64. The logit estimation will not be discussed in detail but the sign and the significance of the estimates seem reasonable. A detailed description of the logit estimation is provided in the appendix. Based on the logit estimation it is possible to obtain the predicted probabilities of belonging to one of the two education groups.

We tried two alternative ways of constructing cohorts using the imputed information on education . In the first approach we construct a binary variable on the basis of the predicted probability and use this new imputed education to split the sample. The new imputed education for household h, e_h, is defined in the following way. Let $\hat{p}_h$ be the estimated probability for household i belong to the more educated group. The binary variable e_h is defined such that

$$e_h = \begin{cases} 0 \text{ if } \hat{p}_h < K_c \\ 1 \text{ if } \hat{p}_h > K_c \end{cases}$$

where K_c is a constant depending on the birth cohort for household h. We choose K_c so that for each cohort, the mean of the imputed variable equals the mean of the original education variable. If we denote the sample of the cohort in period t by χ_t then the cohort means for the two education groups can be constructed in the following way:

$$\bar{x}^0_{h,t} = \frac{\sum_{h\in\chi_t}(1-e_h)x_{h,t}}{\sum_{h\in\chi_t}1-e_h}$$

$$\bar{x}^1_{c,t} = \frac{\sum_{h\in\chi}e_h x_{h,t}}{\sum_{h\in\chi}e_h}. \tag{11}$$

where the superscripts denote the education group. The alternative way of constructing cohort means is simply by replacing the imputed education, e_h, in (11) with the estimated probability, $\hat{p}_h$. The cohort mean is then an weighted average, where the weights reflect the likelihood that the household belongs to a certain education group.

Before constructing the imputed cohort sample means, we need to validate the proposed method. First of all, it is important to have a good prediction of the education level. A goodness of fit test (shown in the appendix) confirms a very high degree of correspondence between the imputed and actual education level and a formal test for no dependency is heavily rejected. By construction the actual and imputed education are equal for each cohort, but there might be some deviations over the years or ages. However, we do not find any systematical deviations over the year or age of the wife. In figure 4, the consumption profiles are constructed on the basis of actual education, imputed education and estimated probabilities of education. To avoid a too messy picture we have only shown the consumption profiles for a single birth cohort (born 1940-1944). The consumption profiles are shown as a function of years and before 1978 only imputed consumption paths are shown. A comparison shows that both the consumption paths based on imputed and weighted education track the consumption path based on the actual education quite closely. In general, the difference between the two education groups are larger when the imputed education is used instead of the weighting procedure. Using imputed education overpredicts the difference between the education groups while using the weighting procedure underpredicts. On the basis of this picture and similar pictures for the remaining birth cohorts, we conclude that the two alternative ways of imputing cohort sample means work equally well in this context. For convenience, we have chosen to construct the cohort sample means on the basis of the imputed education. The imputed education is used in the entire sample to avoid any systematic differences before and after 1977.

4.4 Consumption profiles

The cohort sample means are constructed by averaging over all households belonging to the same birth cohort and education group observed in a given year. On the basis of imputed education we define our final sample. We are using 13 birth cohorts crossed by two education groups. The synthetic panel contains 438 cohort-year observations. We have further limited the sample in two ways. First, the average age of the wife in each cell (cohort-year combination) should be above 20 and below 60. Secondly the minimum number of household in each cell should be above 75. The average cell size is 231. The final sample consists of 101,385 households of whom 63% belong to 'the minimum education group'.

The cohort sample means are plotted against the average age of the wife in figure 5. The consumption of the more educated group is higher than the group with minimum education. Furthermore, the graph shows that the consumption profile of the educated group is increasing faster in the early years and is peaking later. The more educated group reaches it maximum consumption at about the age of 50 while the less educated group peaks at the age of 45. In the following we examine if these differences can be explained by differences in household compositions.

In figure 6, we have present the consumption profiles for households where there are no children (or other adults) currently living in the household. Since we are only averaging over households with no children, the cell size is much lower and the picture becomes more noisy. This is particularly true for the age-group 30 to 45 where only a small fraction of the households remains childless. As shown in section 2, the consumption profile of households which at the present are childless should reflect the selection process. In the earlier age the 'no children group' contains households who are going to have children later and for the older group some of the households might have had children who have already left home. For 'the minimum education group' the consumption profile is still hump-shaped although the picture is not that clear. A closer look at the graph reveals that the consumption for 'no children' peaks around the age of 40, which is five years earlier than the similar graph for 'all households'. The peak around 40 corresponds to the fact that at that age the 'no children group' is most likely to only consist of households who are never going to children. When examining the 'above minimum education group' the graph is very noisy due to the very low number of households in each cell. From the picture it is very difficult to see a hump-shaped consumption profile, however a test for no age effects (fitting a polynomial in age and testing for the age coefficients) is heavily rejected.

5 Modeling the demographic effects

5.1 Estimating adult equivalent scales

In this section we estimate the effects of the demographics. This subsection concerns estimating the ae scale, following the strategy outlined in the econometric section. Compared to the simple one-child model we have extended the model by allowing that the ae scale depends on the age of the child and for economies of scale. In the second sub-section we estimate the completed fertility from the currently observed fertility. Based on the estimates of the completed fertility we derive an expression of the total cost of children.

While the National research council (1995) stress the importance of taking account of the economies of scale, little attention has been paid to the relation between the ae scale and the age of the child. In this study, we deal with these issues by allowing that the extra expenditures on children depend on the number and age of the children in a flexible form. First, we assume that consumption depends on the age of the children in a continuous way, which here is approximated by a third order polynomial. However, the data do not contain information on the age of children aged more than 18 and thereby we are not able to distinguish between children aged more than 18 and other adults living in the households. In this study we assume that the cost for extra adults (including children aged more than 18) equals the cost of an 18 year old. One important issue that we cannot deal with here is the effects of increasing autonomy for grown up children who are still living in the parental home[8]. If such children have their own incomes and make their own consumption decisions (subject to the pooling of some expenditures with their parents) then the 'unitary' life-cycle model we have used as the basis for our analysis may not be appropriate. That is, we can no longer define a household marginal utility of money which is held constant (in expectation) from period to period. This raises the important question of how we model the decision making process of the individuals in the household. There is very little in the literature on this topic: Browning (1999) analyses a theoretical non-unitary model of intertemporal allocation in a two person household and Schultz (1999) discuses the importance of this in the context of explaining the saving behaviour of households in low income countries. This is obviously an important area for future research but we can do little here beyond noting that these effects may differ between the education groups and across cohorts and may be responsible for some of the differences we identify below.

Denote the number of adults and children living in household h in period

[8] Or the impact of older children who are not resident in the household by the FES definition but still receive some 'help' (in consumption terms) from their parents.

t by n^a_{ht} and n^c_{ht} respectively. We define the number of *equivalent adults*, n_{ht}, as:

$$n_{ht} = n^a_{ht} + \sum_{j=1}^{n^c_{ht}} \left(\alpha_0 + \alpha_1 \left(\frac{\text{age}_j}{18} \right) + \alpha_2 \left(\frac{\text{age}_j}{18} \right)^2 + \alpha_3 \left(\frac{\text{age}_j}{18} \right)^3 \right) \tag{12}$$

To ensure that an 18 year old has the same effect as an adult, we impose

$$\alpha_3 = 1 - \alpha_0 - \alpha_1 - \alpha_2$$

Some authors use restricted versions of this formulation. For example, Attanasio et al (1998) do not allow for age effects for children but do allow that adults and children can have different effects; this is equivalent to assuming:

$$\alpha_1 = \alpha_2 = \alpha_3 = 0 \Rightarrow n_{ht} = n^a_{ht} + \alpha_0 n^c_{ht} \tag{13}$$

The authors of National Research Council (1995) make the same suggestion with a value of α_0 equal to 0.7. Gourinchas and Parker (1999) further restrict household composition effects and use only family size which is equivalent to:

$$\alpha_0 = 1 \text{ and } \alpha_1 = \alpha_2 = \alpha_3 = 0 \Rightarrow n_{ht} = n^a_{ht} + n^c_{ht} \tag{14}$$

In our empirical work below we shall test for the validity of these constraints.

Denoting the consumption of a household with n equivalent adults by $c(n)$ we set (see equation (3)):

$$c(n) = e^{k(n)} c(1) \tag{15}$$

To allow for economies of scale effects we use a Box-Cox transformation of the number of equivalent adults:

$$\begin{aligned} k_{ht} = k(n_{ht}) &= \frac{(n_{ht})^\rho - 1}{\rho} \text{ if } \rho \neq 0 \\ &= \ln(n_{ht}) \text{ otherwise} \end{aligned} \tag{16}$$

The degree of economies of scale is captured by the parameter ρ. If the parameter ρ is equal to zero there are no economies of scale ($c(n) = nc(1)$) which should represent a bound on the ae scale. Despite this many studies implicitly use a value of ρ of unity which actually implies that there are diseconomies of scale. If ρ is negative then there are economies of scale with the extreme case being attained as ρ tends to $-\infty$ which represents the case that 'n can live as cheaply as one'. The authors of the National Research

Council (1995) proposal suggest taking a value equivalent to approximately -0.6. Gourinchas and Parker (1999) use dummies for household size which is more general than our formulation and allows for any economies of scale.

The estimation of the parameters of the cost of children is based on the analysis presented in the econometrics section. That is, we run a regression of first differences in cohort mean log consumption on first differences of the mean scale. :

$$\Delta E_{\chi t}(\ln c_{ht}) = \delta_0 + \delta_1 \Delta E_{\chi t}(k_{ht}) + \varepsilon_{\chi t}. \tag{17}$$

where $E_{\chi t}$ denotes the cohort mean and Δ denotes the first difference operator. This is a conventional Euler equation formulation but since the parameters of the ae scale enter in a nonlinear way we have to perform the estimation in two steps. Given the parameters of the ae scale function, the ae scale for each household can be calculated and then aggregated into cohort sample means. On the basis of cohort sample means δ_0 and δ_1 can be estimated and the cohort residuals can be determined. The parameters of the ae scale are then determined such that the sum of square of the cohort residuals is minimized. Let $\theta = \{\delta_0, \delta_1, \rho, \alpha_{0,}\alpha_1, \alpha_2\}$ and the criterion function is then given by

$$K(\theta) = \sum_{\chi,t} (\varepsilon_{\chi t})^2$$

The covariance matrix of the parameters is estimated by:

$$\begin{aligned} V(\theta) &= (D'D)^{-1} D'(\varepsilon\varepsilon') D (D'D)^{-1} \\ D &= \frac{\partial \varepsilon}{\partial \theta} \end{aligned}$$

To estimate the parameters we grid search over the ρ parameter, estimating the remaining parameters using conventional optimisation procedures. When estimating the parameters of the ae scale we have limited the sample further. For the less educated group we use a sample where the average cohort age of the wife is between 20 and 55. The reason for not using the observations where the wife is aged 55 to 60 is that we do not want to consider retirement. For the more educated group a serious problem is the selection in to marriage. For that reason we have chosen to only look at households where the average cohort age of the wife is above 25. The grid search of ρ was performed by varying ρ from -2 to 1 in steps of 0.1. The minimum is obtained for ρ equals -1.6 for the less educated group and 0.3 for the more educated group. In table 1, the estimates are reported. On the basis of the results we conclude

Table 1: Estimates of the ae scale [a] : A cohort analysis[b]

	Min. education	Above min education
ρ	-1.6	0.3
α_0	-0.183	-0.052
	(0.003)	(0.006)
α_1	2.262	0.553
	(0.006)	(0.006)
α_2	-5.848	-1.455
	(0.016)	(0.021)
δ_1	4.075	0.844
	(0.078)	(0.043)
Number of cohort obs.	186	184

a: Δ log consumption deflated with the consumer price index and adjusted for year effects

b: The cohorts are based on the year of birth of the wife (5 year bands)

* indicating significantly different from zero at a 5 per cent level

The numbers in the brackets are standars errors

that the economies of scale is very important for the less educated group while there is a modest degree of economies of scale for the more educated group. For the less educated group, the consumption will increase by 23 per cent if the family increases from two adults to three adults while for the more educated group the similar number is 69 per cent. To illustrate the impact of age on the consumption the age profiles for both education groups are pictured in figure 7. We find that having small children actually lowers consumption by about 20 per cent for the less educated group and 5 per cent for the more educated group. For the less educated group the cost of children is almost constant at ten per cent from the age of 5 to 12, while the cost increases from the age of 12. For the more educated group the cost is increasing for the entire period but at a faster rate after the age of 12. These estimates indicate that the age of the child is important when considering ae scales. In the next section we illustrate the importance of allowing the ae scale to vary with the age.

5.2 Completed fertility

In the econometric section we outlined that estimating completed fertility in an one-child model was fairly simple. Unfortunately, when extending the model to deal with more than two types of completed fertility the estimation procedure becomes more complicated and we will need more assumptions. In our framework, we consider four types of households, which are based on completed fertility: no children, one child, two children and more than two children. For the further analyses we need the distribution (for each cohort) of these four types given that no children are currently present in the household. To construct the conditional distribution either external census information or internal information from the data can be used. We have chosen the latter. Ideally, we need a model for the timing and spacing of births and a model explaining when children leave home. We have used a somewhat simpler approach where only information on the current household composition is needed. To estimate this distribution two identifying assumptions are needed. Firstly, we assume that at one point (and this is the same for all households in the cohort) during the life cycle the completed fertility is revealed. We have chosen this point to be when the wife is aged 37 or 38. This means, for instance, that if we observe a household with the wife aged 37 and with no children living in the household we assume that this household never has had and never will have children. This is off course a very restrictive assumption because women can have children later or some might have had children very early who already have left home. However, the data show that the fraction of households with no children is lowest when the wife is 37 years old and the fraction of households with more than two children is highest at the same age, which provides some support for our assumption. The second assumption is that number of children living in a household can only change by one child per year; this means that only one child can be born per year and only one child can leave home per year. Given these two assumptions, it is possible to estimate the conditional distribution of the four types from the observed number of children living in the household. We allow the distribution to differ across age, education groups and birth year of the wife. To reduce the influence of sampling variation on the estimates we have used a smoothing procedure; the details of the estimation are given in the appendix .

In figure 8, the predicted ratio of households who, at some point, will have two children but at the given age do not have children in the household is presented. The similar ratio of household which never will have children is shown in the same diagram. The probabilities are calculated for a household belonging to the educated group and for which the wife was born in 1940. The way to interpret the figure is, for example, at the age of 30 only 19 per cent of the childless households will end up having two children while 43 per cent will remain childless. By construction the ratio of households which never will have children should be one at the age 37 and 38. Since these ratios are used to construct the counter factual consumption, it is important that they are reliable. To check these figures, another source of information has been used, namely the age of the oldest child currently living in the household. The age of the oldest child can be used to impute the conditional distribution for household aged below 37. A comparison of the two estimated distributions reveals a high degree of similarity. Unfortunately the assumption of completed fertility being observed at age 37 is needed in both estimation procedures, and therefore the sensitivity of this assumption cannot be examined.

Table 2: Total cost of children [a] : A cohort analysis[b]

	Min. education	Above min education	
Cost of one child	0.102	0.321	
	(0.313)	(0.373)	
Cost of two children	0.260	-0.345	
	(0.337)	(0.363)	
Cost of more than two children	0.478	*1.356	0.823
	(0.522)	(0.669)	(0.445)
Number of cohort obs.	215	198	198

a: Δ log consumption deflated with the consumer price index and adjusted for year effects

b: The cohorts are based on the year of birth of the wife (5 year bands)

* indicating significantly different from zero at a 5 per cent level

The numbers in the brackets are standars errors

Based on the estimated conditional distribution of completed fertility we are in principal able to estimate the relative expenditures of 'childless' and 'children' households from the consumption of the 'no children' group. The estimation is based on equation (8), which in the framework with four types of households can be generalised to

$$\begin{aligned} E_t(\ln c_{ht}|z_{ht} &= 0) = \ln c^0 + m_t^1(\ln c^1 - \ln c^0) \\ &\quad + m_t^2(\ln c^2 - \ln c^0) + m_t^3(\ln c^3 - \ln c^0). \end{aligned}$$

m_t^j is the conditional probability of having j children at some point but at time t non of the children are living in the household. Given that we can obtain estimates of the conditional probabilities: m_t^1 m_t^2 and m_t^3 the relative expenditures can be estimated from the equation above. However, we may expect imprecise estimates due to the fact that we have very few observations in each cell. In table 2, the estimates are reported. The estimates indicate for the less educated group that the cost of children demands that a household which will have one child should lower the consumption with about ten per cent when the child is not present in order to save for the higher consumption and lower income when the child is present. In the estimation based on the more educated group the results seem a bit surprising, since the results indicate that there is a gain from having two children compared to not having children[9]. Whether this is a fact or simply due to statistical uncertainty is difficult to tell. In the last column we have restricted cost of one child and two children to be zero. Given these restrictions we find that the cost of having more than two children means that these households have to lower their consumption to about 80 per cent when no children are present. In the following we will use the last column for correction of the consumption for the more educated group.

6 Adjusted consumption.

6.1 Full sample results.

In this section we construct the adjusted consumption paths. As shown in the theory section we can construct the counter-factual cohort consumption up to a constant, which may depend on the cohort. This means that the adjusted cohort consumption should be constant over the life cycle. The counter-factual consumption can be constructed either from the total sample or from the sample of the no children' group. In the following adjusted consumption is established on the basis of the correction outlined in the previous section.

When adjusting the total sample, the estimates of the adult equivalent scales and economies of scale are used. The adjusted consumption paths are shown in figure 9. The consumption paths of both education groups seem be almost constant for each cohort, although the consumption of the less educated group is declining from the age of 50. For the less educated group, we find only weak evidence of cohort effects, however the graphs of the more educated group reveal strong cohort effects. The youngest cohort in the more educated group (born 1960-1964) has a consumption which is about 35 per cent higher than the oldest cohort (born 1910-1914). The differences

[9] We conjecture that some of this is due to the differences in correlated heterogeneity between work and children between the two education groups.

between the cohort effects for the two education groups are consistent with recent findings on the divergence between the wages of different education groups.

For comparison we have pictured the adjusted consumption where only the family size has been used for correction. This model is nested in to the model given by (12) with the restriction given in (14). In this exercise, ρ has been fixed to zero, which corresponds to using the log of household size[10]. By neglecting the age of the children we impose that consumption peaks earlier, specifically around the age of 40 instead of in the mid-forties, see figure 10. This means that it is not possible to obtain a constant consumption paths when only household size is used for the correction. This exercise emphasises the importance of taking account of the age of the children living in the household.

6.2 The 'no children' sample

The no children sample is adjusted by the estimated total cost of children, see figure 11. For the more educated group we use the estimate based on the last column in table 2. The adjusted consumption of the less educated group is almost constant and the hump shape has been removed (compared with figure 6). For the more educated group it is difficult to say if the consumption is more flat after the adjustment because the graphs are very noisy.

7 Conclusion

Many studies have concluded that a precautionary saving motive is needed to explain the lifetime path of household consumption even if we allow for the effects of demographics. In this study, we demonstrate that when taking proper account of the presence of children, family composition can explain the hump-shape in consumption. The empirical findings stress the importance of allowing the adult equivalent scale to depend on the age of the child. For example, we find that having small children in the household actually lowers consumption.

Furthermore, we show that adjusting consumption for household composition is very difficult to do for individual households unless long panel data observations on household consumption are available. We also show that estimating the impact of children from a cross section will lead to downwards biassed estimates. However, both estimating ae scales and adjust-

[10] As mentioned earlier there are no substatial difference if ρ is fixed an another level.

ing consumption for composition effects can more readily be achieved using quasi-panel data.

Moreover, we show that the intuitively appealing approach of only considering consumption of households without children will still lead to a hump-shaped consumption profile. The reason for this somehow surprising finding is that in the no children' group consists of different types of households at different ages; e.g. for young ages the group contains households who are going to have children later and households who never will have children. Although the consumption of the no children' group cannot directly be used for testing for the life cycle model, we can still derive an adjusted no children' consumption profile on the basis of the predicted ratio of households which will end up with having children. In the paper we present one way of deriving the conditional probability of having children given that the household do not have any at the present. Based on these conditional probabilities we obtain an expression for the total cost of children, which contains foregone income and extra consumption when children are present. Even though the total cost empirically is very imprecisely determined, we provide a framework in which it is possible to estimate the total cost of children, which to our knowledge not have been presented before. To summarize, we use the total sample to estimate the ae scale while the total cost of children can be estimated from the no children' group'.

In more general terms, this excise shows when working with repeated cross sections selecting according to time constant variable raises very different issues that when selecting a time varying variable, e.g. the present of children. This might also be relevant in other contexts, e.g. when the selection is made on the basis of occupation rather than education, where we expect education to be almost constant over time while the occupation may change over a life cycle. This is one of the reasons for why the imputed education is introduced, which allow us to base the cohorts on education even in the time periods where education is not reported. A comparison between consumption paths of actually and imputed reveals a high degree of correspondence, which might make this technique useful for other purposes.

However, one problem which remains unsolved is the selection of married couples. To deal properly with this problem we need to develop a marriage model, which explains the connection between marital status and life time income.

References

[1] Attanasio, Orazio, James Banks, Costas Meghir and Guglielmo Weber (1999), "Humps and bumps in lifetime consumption", *Journal of Economic and Business Statistics*, 17(1), 22-35.

[2] Attanasio, Orazio and Martin Browning (1995), "Consumption over the life cycle and over the business cycle", *American Economic Review*, 85, 1118-1137.

[3] Attanasio, Orazio and Guglielmo Weber (1995), "Is consumption growth consistent with intertemporal optimisation? Evidence from the Consumer Expenditure Survey", *Journal of Political Economy*, 103, 1121-1157.

[4] Attanasio, Orazio and Guglielmo Weber (1993), "Consumption growth, the interest rate and aggregation", *Review of Economic Studies*, 60, 634-49.

[5] Blundell, Richard, Martin Browning and Costas Meghir (1994), "Consumer demand and the lifetime allocation of consumption, *Review of Economic Studies*, 61, 57-80.

[6] Blundell, Richard and Thomas MaCurdy (1999), "Labor Supply: A review of alternative approaches", *Working paper* Institute of Fiscal Studies

[7] Browning, Martin (1999), "Saving in a two person household", *Scandinavian Journal of Economics*, forthcoming.

[8] Browning, Martin, Angus Deaton and Margaret Irish (1985), "A profitable approach to labor supply and commodity demands over the life-cycle", *Econometrica*, 53, 503-44.

[9] Browning, Martin, Lars Peter Hansen and James J. Heckman (1999), "Micro Data and General Equilibrium Models", *Discussion paper*, Institute of Economics, University of Copenhagen.

[10] Calhoun, C.A. and T. S. Espenshade (1988), "Childbearing and wives' foregone earnings", *Population Studies*, 42(1), 5-38.

[11] Carroll, Christoher (1994), "How does future income affect current consumption?", *Quarterly Journal of Economics*, 109, 111-148.

[12] Carroll, Christopher and Lawrence Summers (1991), "Consumption growth parallels income growth:some new evidence", in B.D. Bernheim and J.B. Shoven (eds) *National Saving and Economic Performance*, Chicago University Press for NBER, Chicago, 305-43.

[13] Deaton, Angus (1991), "Saving and liquidity constraints", *Econometrica*, 59, 1221-48.

[14] Gourinchas, Pierre-Olivier and Jonathan Parker (1999), "Consumption over the life cycle", draft, Princeton University.

[15] Heckman, James (1974), "Life cycle consumption and labor supply: an explanation of the relationship between income and consumption over the life cycle", *American Economic Review*, 64, 188-194.

[16] Hoem, Jan M.(1997), "Educational Gradients in Divorce Risks in Sweden in Recent Decades", *Population Studies*, 51, 19-27.

[17] Hubbard, R.G., Jonathan Skinner and Stephen Zeldes (1994), "Precautionary saving and social insurance", *Journal of Political Economy*, 103, 360-399.

[18] Joshi, Heather (1990), "The cash opportunity costs of childbearing: an approach to estimation using British data", *Population Studies*, 44, 41-60.

[19] Lucas, Robert and L. Rapping (1970), "Real wages, employment and inflation", in E.S. Phelps (ed) *Microeconomic Foundations of Employment and Inflation Theory*, New York.

[20] Nagatani, Keizo (1972), "Life cycle saving: theory and fact" *American Economic Review*, 62, 344-53.

[21] National research council (1995), *Measuring Poverty: a new approach*, National Academy Press, Washington DC.

[22] Schultz, Paul T.(1999), "Who is head of the household", presented at the ESPE conference, Turin 1999.

[23] Thurow, Lester (1969), "The optimum lifetime distribution of consumption expenditures" *American Economic Review*, 59, 324-30.

[24] Tobin, James (1967), "Life cycle saving and balanced growth" in W. Fellner et al (eds) *Ten Economic Studies in the Tradition of Irving Fisher*, New York.

[25] Van der Klaauw, Wilbert (1996), "Female Labour Supply and Marital Status Decision" *The Review of Economic Studies*, 63, 199-235.

A Appendix

A.1 Estimating the predicted education

On the basis of the education dummy which is observed from 1978-1995 a logit model is estimated. The explanatory variables are dummy variables for occupation, household tenure, employment status of husband and wife, respectively, birth cohorts, region of residence and a third order polynomium in the husband's year of birth and finally the age difference between the wife and husband in levels and interacted with year of birth. In total we use 58 explanatory variables. The results show that occupation dummies and dummies for household tenure have a large explanatory power. More surprisingly is that the age difference between husband and wife seems to predict the husband's education level. If there husband is much older than the wife the husband is less likely to belong to the more educated group. However the effect vanishes for younger birth cohorts. The estimation is performed on a sample of households with the husband aged from 20 to 64. The sample consists of 64,084 observations and the R^2 is equal to 0.2107.

On the basis of the estimated probability he imputed is constructed as described in section 4. A comparison between the actual and imputed education reveals a high degree of correspondence. Below is shown a cross table of the actual and imputed education.

Table 3: A comparison of actual and imputed education

	Imputed		
Actually	Min. education	Above min education	Total
Min. education	29,128	8,475	37,603
Above min edu.	8,470	18,011	26,481
Total	37,598	26,486	64,084

Pearson goodness of fit test: $\chi^2(1) = 13252.67$ p-value=0.00

A.2 Estimating the conditional distribution of types of households

This section describes how the conditional distributions of different households types are estimated. Let z_{ht} be the number of children present in household h when the wife is aged t. Let z_h be the total number of children household h ever will have. The values for of variables are 0,1, 2 and 3+ (3 and more than 3 children). In the following household of type j will refer to $z_h = j, j = 0, 1, 2, 3+$.

For the correction of consumption of households with no children living in the household, we will need the distribution of types conditioned on being observed with no children:

$$\eta_{j,t} = \Pr(z_h = j | z_{ht} = 0) \qquad j = 0, 1, 2, 3+$$

To be able to identify and estimate these probabilities from the source of data available some assumptions are needed.

ASSUMPTION 1:

At one point in the life of the households the completed fertility of the household can be observed. Here, the total number of children is assumed to be observable when the wife is aged 37 and 38. This means that wife is assumed not to give birth after the age of 37 and on the other hand children are assumed to stay in the household until the wife is aged 37. The assumption can be formalized as

$$\begin{aligned} \Pr(z_h &= j_1 | z_{h37} = j_2) = \begin{cases} 1 & \text{if } j_1 = j_2 \\ 0 & \text{if } j_1 \neq j_2 \end{cases} \\ \Pr(z_h &= j_1 | z_{h38} = j_2) = \begin{cases} 1 & \text{if } j_1 = j_2 \\ 0 & \text{if } j_1 \neq j_2 \end{cases} \end{aligned}$$

Given assumption 1 the unconditional distribution of type of households can be determined as

$$\Pr(z_h = j) = \Pr(z_{h37} = j).$$

In the sample the probability of $\nu_{jt} = \Pr(z_{ht} = j)$ can easily be estimated simply by the fraction of households with the wife aged s and observed with j children to the total number of households with a wife aged s

$$\hat{\nu}_{jt} = \frac{\#\{\text{households with } age = s \text{ and } z_{hs} = j\}}{\#\{\text{households with } age = s\}}$$

In order to minimize the impact of sampling variation we have smoothen these probabilities by estimating a multinomial logit model for the four alternatives. The explanatory variables are a fourth order polynomium in age and cohort dummies. Based on this calculation the unconditional distribution of types are determined to

type with $z_i = 0$	0.066
type with $z_i = 1$	0.175
type with $z_i = 2$	0.442
type with $z_i = 3+$	0.317

From the source of data available, there are two alternative ways of determining $\mu_{j,t}$. The first method exploits the fact that we know the age of the children living in the household (unfortunately, the age is truncated at 18 years old). From this information it is possible to construct the distribution of the age of first birth for each type by using the distribution at the age of 37.

$$\begin{aligned}\Pr(z_{ht} &= 0|z_{h37} = j) \\ &= \frac{\#\{\text{households with } age = 37, \text{ age of first birth} > t \text{ and } z_{h37} = j\}}{\#\{\text{households with } age = 37\text{and } z_{h37} = j\}} \\ \text{for } t &< 37\end{aligned}$$

Since we know the unconditional distributions one can easily construct $\eta_{j,t}$ for $t < 37$:

$$\Pr(z_h = j|z_{ht} = 0) = \Pr(z_{ht} = 0|z_h = j)\frac{\Pr(z_h = j)}{\Pr(z_{ht} = 0)}.$$

By replacing the actually probabilities with corresponding estimates, an estimate of $\eta_{j,t}$ is obtained.

The disadvantage of this method is that it is not possible to construct estimates of $\eta_{j,t}$ when the age is above 38. Instead of using the distribution of the age of first birth we can use the fact that we know the fraction of household observed with 0,1 2 and 3+ children at different ages. By imposing restrictions on the transitions between different number of children living in the household we obtain identification. The assumption is the following:

ASSUMPTION 2

Only one child is born per year and only one child can leave the home within a year for each household.

Given these assumptions we can construct the transition matrix. Let transition probability be given by

$$\lambda_{jt} = \Pr(z_{ht} = j | z_{ht-1} = j) \; j = 0, 1, 2, 3+.$$

The transition matrix for $t < 37$ is given by

$$\Lambda_t = \begin{pmatrix} \lambda_{0t} & 0 & 0 & 0 \\ (1-\lambda_{0t}) & \lambda_{1t} & 0 & 0 \\ 0 & 1-\lambda_{1t} & \lambda_{2t} & 0 \\ 0 & 0 & (1-\lambda_{2t}) & 1 \end{pmatrix}.$$

For $t = 37$ the transition matrix is $\Lambda_{37} = I$ and $t \geq 38$ the transition matrix is given by

$$\Lambda_t = \begin{pmatrix} 1 & (1-\lambda_1) & 0 & 0 \\ 0 & \lambda_{1t} & (1-\lambda_{2t}) & 0 \\ 0 & 0 & \lambda_{2t} & (1-\lambda_{3+,t}) \\ 0 & 0 & 0 & \lambda_{3+,t} \end{pmatrix}.$$

The parameters of the transition matrix can be identified from the transitions

$$\nu_t = \Lambda_{t-1} v_{t-1}$$

where $\nu_t' = (v_{0,t}, v_{1,t}, v_{2,t}, v_{3,t})$.

From the transition matrices we can construct $\Pr(z_h | z_{it} = 0)$. When $t < 37$, the probability $\Pr(z_h | z_{it} = 0)$ is given by

$$\begin{aligned} \Pr(z_h | z_{ht} &= 0) = \Pr(z_{h37} | z_{ht} = 0) \\ &= \Lambda_{36} \cdots \Lambda_t e, \end{aligned}$$

where $e' = (1, 0, 0, 0)$. For $t > 38$ we use that

$$\begin{aligned} \Pr(z_{ht} &= 0 | z_h = j) = \Pr(z_{ht} = 0 | z_{h38} = j) \\ &= \Lambda_t \Lambda_{t-1} \cdots \Lambda_{38} e_j, \end{aligned}$$

where $e' = (1_{\{j=0\}}, 1_{\{j=1\}}, 1_{\{j=2\}}, 1_{\{j=3+\}})$. By using the following equation:

$$\Pr(z_h = j | z_{ht} = 0) = \Pr(z_{ht} = 0 | z_h = j) \frac{\Pr(z_h = j)}{\Pr(z_{ht} = 0)}$$

we can obtain an expression for $\eta_{j,t}$.

A.3 Figures

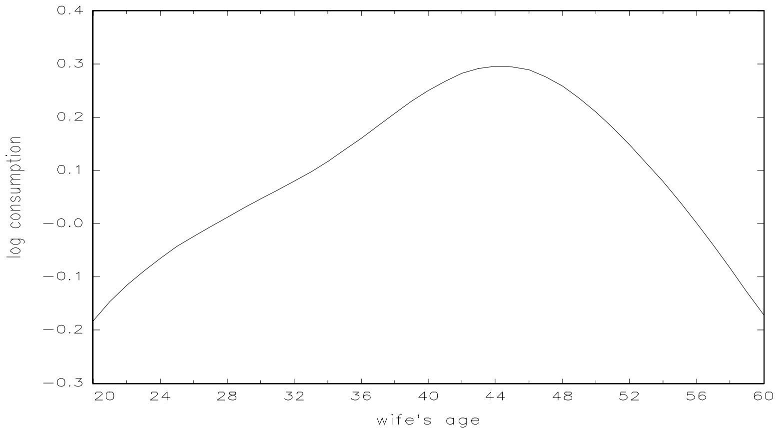

Figure1: Smoothed consumption

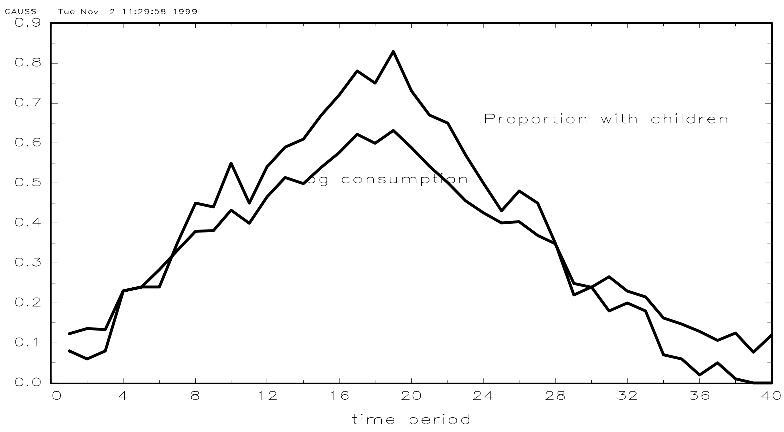

Figure 2: Simulated consumption and children

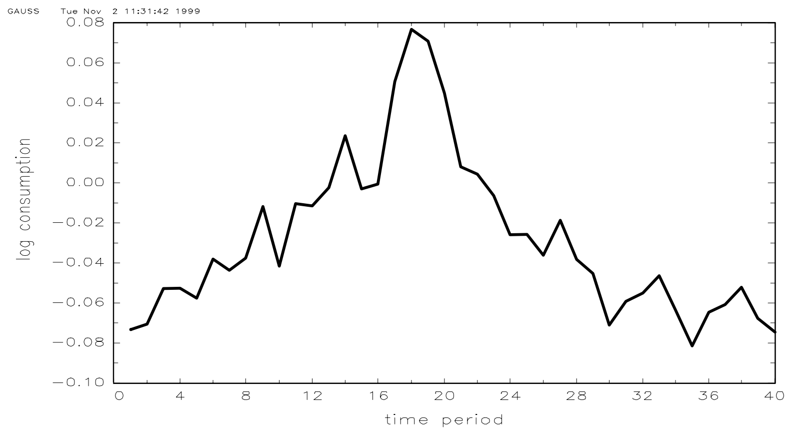

Figure 3: Simulated consumption for no children sample

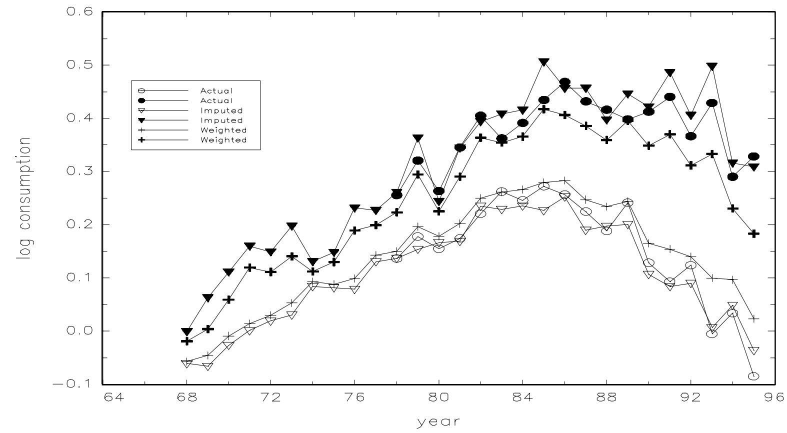

Figure 4: Consumption profiles based on imputed and actual education (birth cohort 1940-1944)

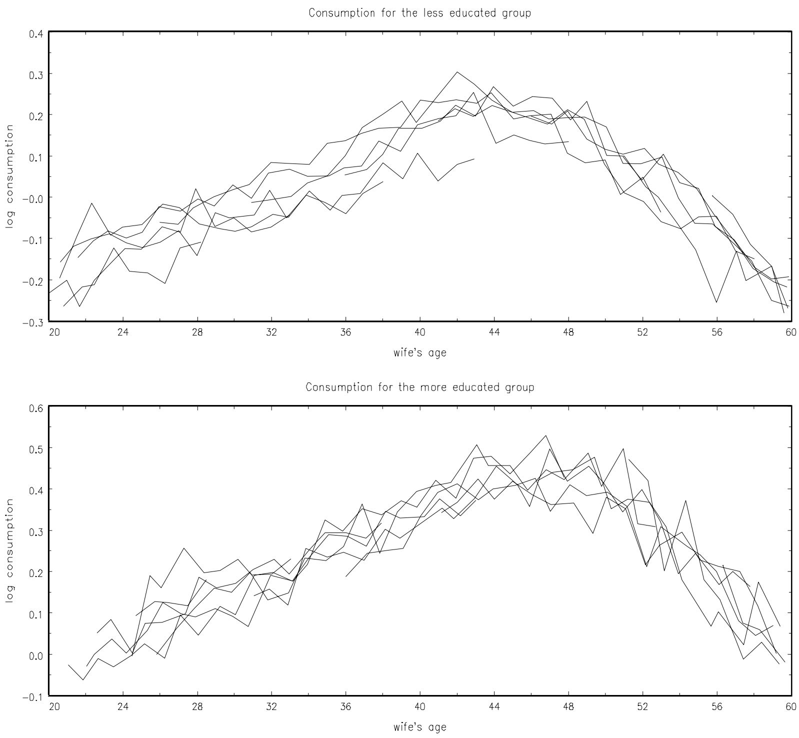

Figure 5: Consumption profiles for the full sample

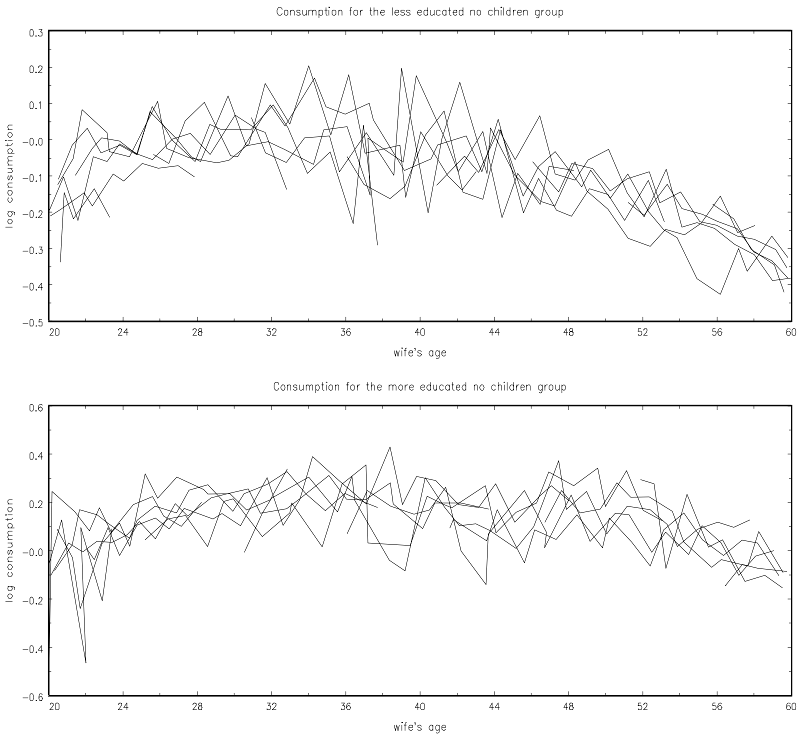

Figure 6: Consumption profiles of 'no children' sample

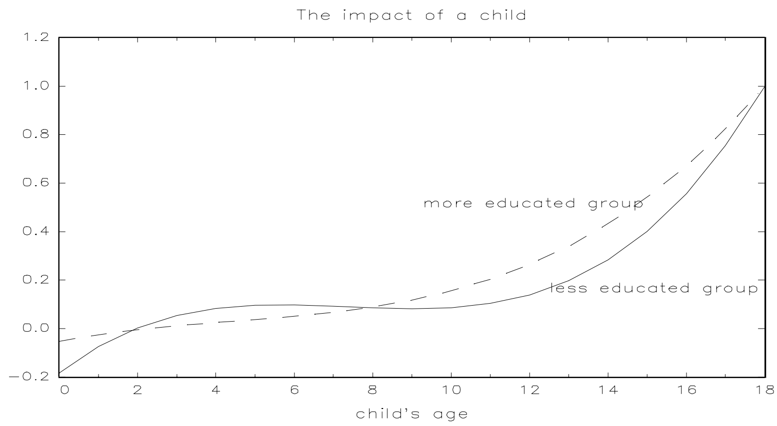

Figure 7: The equivalent adult (adult=1)

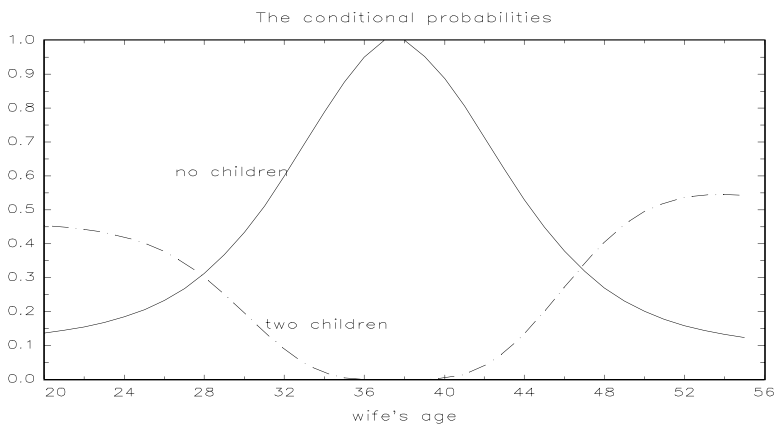

Figure 8: Estimated conditional probabilities given no children are present in the household

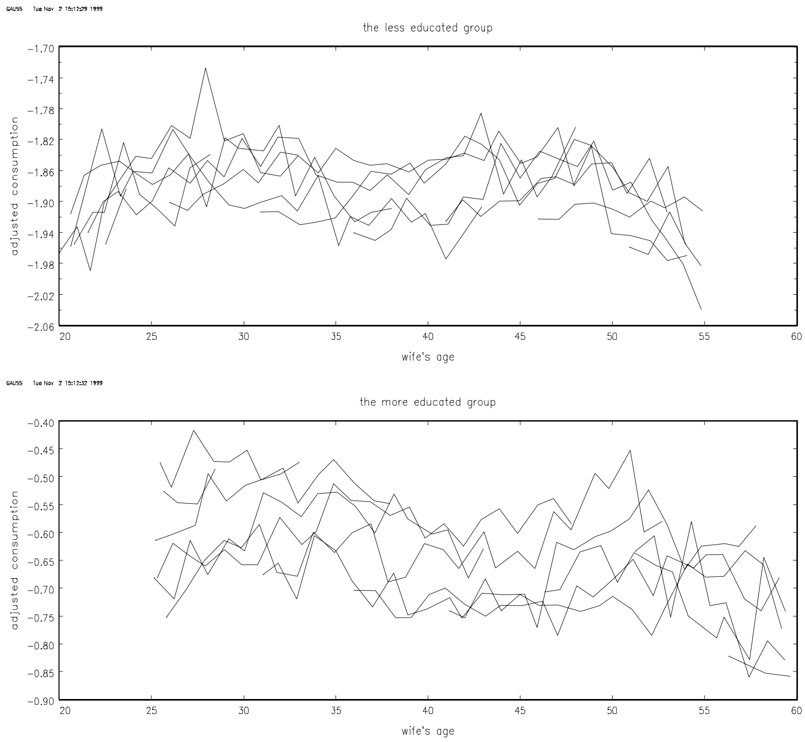

Figure 9: The consumption adjusted for household size and the age of the children (full sample)

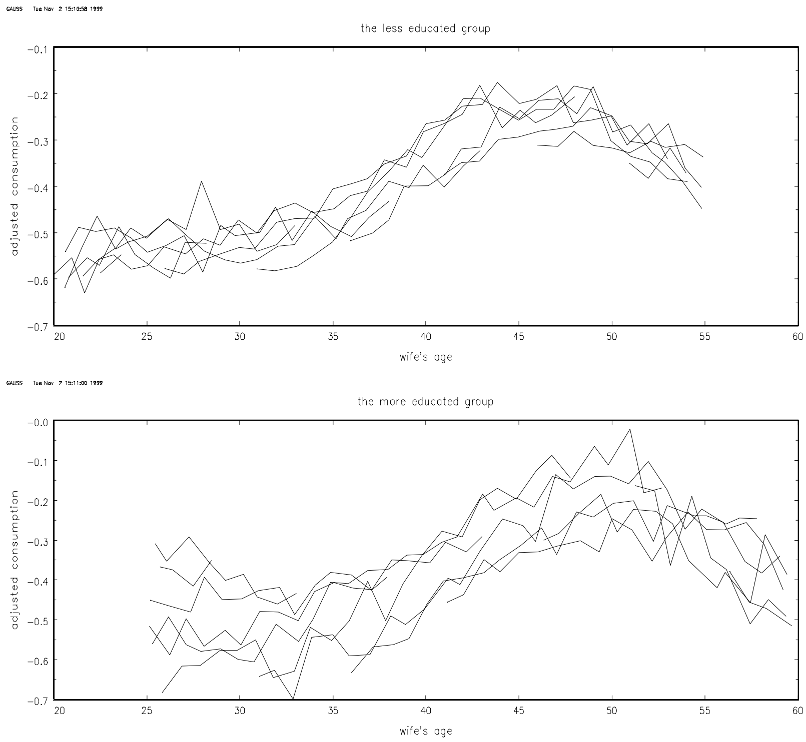

Figure 10: The consumption adjusted for household size (full sample)

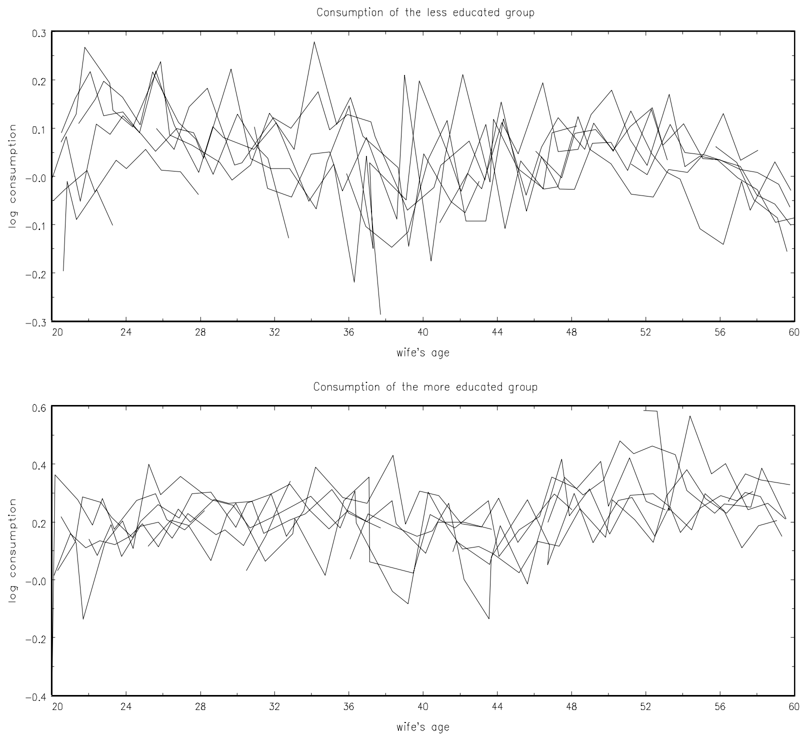

Figure 11: The adjusted consumption of the 'no children sample'