



Consumption and Children

Browning, Martin; Ejrnæs, Mette

Publication date:
2002

Document version
Publisher's PDF, also known as Version of record

Citation for published version (APA):
Browning, M., & Ejrnæs, M. (2002). *Consumption and Children*. Cph.: Department of Economics, University of Copenhagen.

Consumption and Children

**Martin Browning
Mette Ejrnæs**

2002-06

Consumption and Children

Martin Browning and Mette Ejrnæs*

August 2002

Abstract

Several recent papers have concluded that precautionary saving motives are needed to reconcile data on lifetime patterns of consumption and income with a standard optimising model. In this paper we contest that we necessarily need a precautionary motive and we show that if we take consumption to take proper account of the number and ages of children, then adjusted consumption does not track income. We do *not* infer from this that children are the sole cause of the rise in consumption in the early part of life but simply that the data we have are not informative enough to convincingly identify the cause of the rise.

JEL: D91, J13, D12

KEYWORDS: Consumption, Children, Precautionary motive, Life-cycle

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 four principal elements to this: allocation between a period of human capital formation and later; allocation within the working life; allocation between the pre-retirement and retirement period and planning for bequests. 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, Lester Thurow (1969), Martin Browning, Angus Deaton and Margaret Irish (1985) and Christopher Carroll and Lawrence Summers (1991)). The same sources also show that household income displays a very similar pattern to consumption over

*Martin Browning: Centre for Applied Microeconometrics and Institute of Economics, University of Copenhagen, Studiestræde 6, DK-1455 Copenhagen K, Denmark, email: Martin.Browning@econ.ku.dk. Mette Ejrnæs: Centre for Applied Microeconometrics and Institute of Economics, University of Copenhagen, Studiestræde 6, DK-1455 Copenhagen K, Denmark, email: Mette.Ejrnæs@econ.ku.dk. We thank the Danish SSF and the Danish National Research Foundation (through its grants to CAM) for support of the research presented here.

the working life. Any contribution to understanding why consumption and income have such a high life-cycle correlation is of critical importance in many policy debates.¹

There have been five broad responses to the observed ‘tracking’ of consumption and income over the life-cycle.² 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, this 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*. Keizo 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. James 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 households increase consumption when children are present then the associated pattern

¹Many issues of cyclical policy turn on the high frequency correlation between consumption and income; we shall not address this in the paper.

²One alternative explanation that has not been explored is that the correlation is an artefact of the sampling procedure. All studies condition their sample selection on being married. 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. To properly deal with the selection bias we need a model of the relationship between the timing of marriage and lifetime income so we leave this to future work.

³This intertemporal substitution explanation argument is closely related to the explanation of cyclical variations in employment being due to cyclical variations in wages; see Robert Lucas and Leonard Rapping (1970) for the original analysis and Richard Blundell and Thomas MaCurdy (1999) and Martin Browning, Lars P. Hansen and James Heckman (1999) for recent surveys of the evidence.

is consistent with a life-cycle model with no liquidity constraints. James 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). Richard Blundell, Martin Browning and Costas Meghir (1994), Orazio Attanasio and Guglielmo Weber (1995) and Orazio Attanasio and Martin 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 the various factors mentioned above 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 Angus Deaton (1991)) whilst remaining agnostic on the importance of demographics. Christopher Carroll (1994) and Glenn R. Hubbard, Jonathan Skinner and Stephen 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 accounting for demographics. Orazio Attanasio, James Banks, Costas Meghir and Guglielmo 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. Finally, Pierre-Olivier Gourinchas and Jonathan Parker (2002) argue, on the basis of a estimated simulation model, 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 a precautionary motive. To do this, we specify a model of intertemporal allocation in which households move resources from periods when they do not have children into periods when they do have children (with households that never have children 'smoothing' real expenditures). Given such a model we show how to construct two counter-factuals from a time series of cross-section family expenditure surveys that allow us to test whether households 'smooth' consumption.

For the first counter-factual we take means over cohorts in each period (this is the conventional construction of quasi-panel data). In any year this sample comprises four groups: those who never have children, those who currently have children, those who will have children in the future and those who have had children who have now left home. We then show how to adjust consumption in households with children to give their consumption before and after chil-

dren. In the second counter-factual 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 two groups: those who never have children and those who have children at some time. 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 for constructing the consequent adjusted paths of consumption.

We thus construct two series of adjusted consumption (in practice, we stratify on education and present results for two education groups separately) under the assumption that the only age effects are due to a trend (for patience or impatience) and children. The identifying assumptions to construct the two adjusted measures are quite different, so that they represent independent evidence. We show that neither of the two adjusted consumption series shows any sign of residual non-linear age effects. In particular, adjusted consumption does not track income at the beginning of the life-cycle. We stress that we do not thus claim that the precautionary motive or liquidity constraints are necessarily unimportant but only that they are not strictly needed to rationalise the observed low frequency seen in quasi-panel data. In fact, our general conclusion from our empirical analysis is that 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.

I Constructing counter-factual consumption.

A The basic model.

The focus of this study is the cause of the non-linear relationship between age and consumption. Our strategy is to assume away any direct non-linear age effects due to liquidity constraints, prudence and direct effects of age on utility and try to capture all of the age effects with children variables. To that end we propose a relatively simple model and then ask how well it fits the facts concerning consumption over the life-cycle. To simplify the exposition we assume that households have at most one child and that the effect of this child on consumption when they are present in the household is independent of age.⁴ Let z_{ht} be a dummy variable denoting the presence of a child in household h at age t . We assume that, conditional on the children variable, the within period utility function is independent of age:

$$(1) \quad u_{ht} = u^h(C_{ht}, z_{ht})$$

⁴In our empirical work below we take account of the numbers and ages of children

where C_{ht} is the consumption of household h at age t . This rules out age utility effects other than those induced by the presence of children. To make progress we need to make further assumptions about functional form. In the consumption literature, most previous investigators have used variants of the following scheme:

$$(2) \quad u^h(C, z) = v_h(Ce^{-\delta_h z}) e^{\delta_h z}$$

where $v_h(\cdot)$ is strictly increasing and strictly concave. To see why we adopt this form⁵, consider the case of a household which has a discount factor equal to the real rate and which does not face any uncertainty. The Euler equation gives:

$$(3) \quad \begin{aligned} u_C^h(C_{ht}, z_{ht}) &= u_C^h(C_{ht+1}, z_{ht+1}) \Rightarrow \\ v_h'(C_t e^{-\delta_h z_{ht}}) &= v_h'(C_{t+1} e^{-\delta_h z_{h,t+1}}) \Rightarrow \\ \Delta c_{ht+1} &= \delta_h \Delta z_{ht+1} \end{aligned}$$

where the lower case c_{ht} denotes the log of C_{ht} . Thus consumption changes are proportional to changes in the child state.

It is very important to emphasise that the (multiplicative) child response function $e^{-\delta_h z}$ is *not* an adult equivalence scale (in the sense of keeping utility constant) nor is it (proportional) expenditures on children, although the latter will be an element of the response. It is the change in consumption that keeps the marginal parental utility of expenditure constant over changes in the child state. It is usual to require that parents increase consumption when children are present ($\delta_h > 0$) but there is no logical necessity for this. It might be, for instance, that some parents choose to spend more when children are *not* present (for example, on going out or having expensive holidays) and then to cut back on expenditures when the child is in the household. As we shall see below, this issue becomes important when we allow for the dependence of the response function on the age of the child.

We assume that the sub-utility function $v(\cdot)$ takes the iso-elastic form:

$$(4) \quad v_h(Ce^{-\delta_h z}) = \frac{(Ce^{-\delta_h z})^{(1-\theta_h)}}{(1-\theta_h)}$$

Assuming a constant real rate r and a discount factor β_h and no liquidity constraint, the exact Euler equation is:

$$(5) \quad (C_{ht} e^{-\delta_h z_{ht}})^{-\theta_h} = \beta_h (1+r) E_t \left[(C_{ht+1} e^{-\delta_h z_{ht+1}})^{-\theta_h} \right]$$

We now adopt the linearised version of this exact Euler equation. Although the linearised Euler equation is held in disrepute as a vehicle for estimating the elasticity of intertemporal

⁵In the demand literature more sophisticated schemes are used to incorporate the effects of children. The restriction used here is usually referred to as ‘demographic separability’.

substitution (see, for example, Sydney Ludvigson and Christina Paxson (2001) and Christopher Carroll (2001)), it captures exactly our hypothesis that the precautionary motive can be ignored. Specifically, suppose that the precautionary motive is more significant in the early part of the life-cycle. If we exclude the conditional variance of consumption and higher conditional moments this will lead to significant (but spurious) ‘age’ effects on consumption. In the case considered here, consumption growth will be higher in earlier years even if the household does not have children. The linearised Euler equation is given by:

$$(6) \quad \Delta c_{ht+1} = \alpha_h + \delta_h \Delta z_{ht+1} + \varepsilon_{ht+1} \text{ with } E_t [\varepsilon_{ht+1}] = 0$$

where $\alpha_h = -\theta_h^{-1} (\ln(\beta_h (1+r)))$. Our ‘null’ hypothesis is thus that the linearised Euler equation does not display any age effects except for those associated with changes in demographics, z . We now consider how to test this using data drawn from a time series of cross-sections. We present two alternative methods that rely on quite different identifying assumptions.

B An unconditional sampling scheme.

In our empirical work below we adopt two different schemes for sampling from a time-series of cross-sections of married couples. In the first we take means across the whole sample for different birth-education cohorts in each period. This is the conventional sampling scheme for time series of cross-sections; we refer to it as the ‘unconditional mean’ sampling scheme. The composition of this cohort should be stable over time, except for marriage, divorce, death, emigration etc.⁶ We denote the (unconditional) population mean of variable y_{ht} for cohort d at age t by $E^d [y_{ht}]$. Taking means through (6) we have:

$$(7) \quad \Delta E^d [c_{ht+1}] = E^d [\alpha_h] + \Delta E^d [\delta_h z_{ht+1}] + E^d [\varepsilon_{ht+1}]$$

We assume that δ_h and Δz_{ht+1} are uncorrelated across households at every age so that:

$$(8) \quad \begin{aligned} \Delta E^d [\delta_h z_{ht+1}] &= E^d [\delta_h] E^d [\Delta z_{ht+1}] \\ &= \delta_d \Delta p_{dt+1} \end{aligned}$$

where $\delta_d = E^d [\delta_h]$ is the age invariant cohort mean child effect and p_{dt} is the proportion of households in cohort d at age t who have a child present. We then re-write the cohort Euler equation as:

$$(9) \quad \Delta c_{d,t+1} = \alpha_d + \delta_d \Delta p_{d,t+1} + \varepsilon_{d,t+1}$$

⁶In practice, the composition of the empirical group also changes over time since we only take samples from the cohort population. Since the issues that arise from this are well known (see Angus Deaton (1985)), we shall ignore this in our exposition.

where $c_{d,t+1} = E^d[c_{ht+1}]$ and similarly for the other terms. In our empirical work below we assume that δ_d is the same for all cohorts⁷ ($= \delta$), so that we can pool across cohorts. Given a series of cohorts and a long time series for each cohort we can estimate δ using conventional GMM techniques. Given an estimate $\hat{\delta}$ we can then construct adjusted consumption paths using:

$$(10) \quad \hat{c}_{d,t+1} = c_{d,t+1} - \hat{\delta}z_{d,t+1} \Rightarrow E(\hat{c}_{d,t+1}) = \alpha_d t$$

Under our hypothesis that we can capture all non-linear age effects by children, adjusted consumption, $\hat{c}_{d,t+1}$, should not display anything other than linear effects in age. We shall test this using statistical and graphical techniques.

C A ‘no children present’ sampling scheme.

The sampling scheme just presented is conventional and yields quasi-panels; that is, the construction of groups with compositions that are invariant over time. Our second sampling scheme is novel and presents analysis using groups with time varying membership. This sampling scheme is motivated by the observation that under our null hypothesis, the consumption path of those who never have children has no non-linear age effects. Unfortunately in our data we do not know which currently childless households never have a child so that we cannot draw such a sample. Instead we take means for the sub-sample of households in a given cohort-period who do not currently have any children present; we refer to this as ‘no child present’ sampling scheme. The composition of this sub-sample changes radically over time since at early and late ages the sample includes almost everyone whereas around ‘middle’ age the sample consists mostly of households that never have children. This will introduce spurious age effects; below we present identifying assumptions that allow us to remove these effects and to construct an estimate of the consumption path of households that never have children.

The proportion of households who have a child present is p_{dt} ; the proportion of the cohort who have a child at some time is denoted π_d . We have $0 \leq p_{dt} \leq \pi_d$. At any age the ‘no children present’ sample is composed of two groups. The first, group, I , is those who never have children. The proportion of such households in the ‘no child present’ sample is $(1 - \pi_d) / (1 - p_{dt})$. The second group, II , do not currently have a child present but have either had a child or will have one; the proportion of this group is $(\pi_d - p_{dt}) / (1 - p_{dt}) = m_{dt}$. Denote the age t mean, conditional on never having a child by where $E_I^d[c_{ht}]$ of cohort d , with $E_{II}^d[c_{ht}]$ as the corresponding cohort mean for those who do have a child at some time but not currently. In general $E_{II}^d[c_{ht}]$ will not be equal to $E_I^d[c_{ht}]$ for at least two reasons. First, lifetime resources and the propensity to have children may be correlated either because of correlated

⁷In our empirical work we always stratify on the husband’s education, so that we actually allow that the mean child effect depends on education.

heterogeneity or because having children lowers lifetime income.⁸ Second, if potential parents have lower consumption before they have children, then even if lifetime resources are the same across the two groups, the group who will have children at some time have lower consumption when the children are not present. This suggests that $E_I^d [c_{ht}] > E_{II}^d [c_{ht}]$ which will induce a spurious inverted-U age effect in the consumption path of this sample. This follows since in the middle age ranges (the wife is aged about 40), the sample is mostly those who never have children (group I) and consequently mean consumption is higher at those ages.

The mean for cohort d at age t over all those with no child currently present is:

$$\begin{aligned}
 E_0^d [c_{ht}] &= \frac{(1 - \pi_d)}{(1 - p_{dt})} E_I^d [c_{ht}] + \frac{(\pi_d - p_{dt})}{(1 - p_{dt})} E_{II}^d [c_{ht}] \\
 &= E_I^d [c_{ht}] + \frac{(\pi_d - p_{dt})}{(1 - p_{dt})} (E_{II}^d [c_{ht}] - E_I^d [c_{ht}]) \\
 (11) \qquad &= E_I^d [c_{ht}] + m_{d,t} (E_{II}^d [c_{ht}] - E_I^d [c_{ht}])
 \end{aligned}$$

The counter-factual consumption we wish to construct is the mean of consumption by those who never have children:

$$(12) \qquad E_I^d [c_{ht}] = E_0^d [c_{ht}] - m_{d,t} (E_{II}^d [c_{ht}] - E_I^d [c_{ht}])$$

Under our null hypothesis, this consumption path should not display any non-linear age effects. The first term on the right hand side is directly observable. The first part of the second term, $m_{d,t}$, is the proportion of the ‘no children present’ sample at age t who will have a child at some time. This can be constructed if we can estimate π_d (since $p_{d,t}$, the proportion of households with a child, can also be observed); we show how implement an estimator of π_d in the next section. Thus to construct our counter-factual consumption level we need a measure of $(E_{II}^d [c_{ht}] - E_I^d [c_{ht}])$. This is a measure of the difference between the cohort mean consumption levels of those who have children but not currently and those who never have children. We define:

$$(13) \qquad E_{II}^d [c_{ht}] - E_I^d [c_{ht}] = \tau_d + \eta_{d,t}$$

where $\eta_{d,t}$ is zero mean. Following equation (12), we define adjusted consumption to be:

$$(14) \qquad \tilde{c}_{d,t} = E_0^d [c_{ht}] - m_{d,t} \tau_d$$

Thus we need to estimate τ_d to calculate the adjusted consumption that corresponds to $E_I^d [c_{ht}]$. In our estimation procedure we assume τ_d is the same for all cohorts. By using equation (11)

⁸Estimates of the impact of childbearing on women’s earnings in the literature (see, for example, Charles A. Calhoun and Thomas S. Espenshade (1988) for the U.S. and Heather Joshi (1990) for the U.K.) suggest that this may be the case.

and (13) we have:

$$(15) \quad E_0^d [c_{ht+1}] = E_I^d [c_{ht+1}] + m_{d,t} (\tau + \eta_{d,t}).$$

From the Euler equation (6) on the time invariant group I we have

$$(16) \quad E_I^d [c_{ht+1}] = \alpha_d t + \gamma_d + \omega_{d,t}.$$

The estimation equation is then

$$(17) \quad c_{d,t}^0 = \alpha_d t + \gamma_d + m_{d,t} \tau + m_{d,1} \eta_{d,t} + \omega_{d,t},$$

where $c_{d,t}^0 = E_0^d [c_{ht+1}]$. We shall use this equation as the basis for estimating τ and hence constructing adjusted consumption $\tilde{c}_{d,t}$. In order to obtain unbiased estimates of τ we will use predicted values of $m_{d,t}$ to ensure that these are uncorrelated with the error terms. Once we have this, we shall once again test for non-linear age effects.

II Econometric issues

A Estimating child response functions.

This sub-section discusses how to estimate the child response function from observed repeated cross-section data on consumption. One scheme that is often used to ‘control’ for children is to first run a regression of log consumption on children in the cross-section data and then to work with the residuals from this regression (see, for example Gourinchas and Parker (2002)). We begin by showing that this does *not* identify the child response function. To do this we consider the levels of log consumption from the Euler equation (6) with the assumption that all households have the same mean parameters (α_h, δ_h):

$$(18) \quad c_{ht} = \alpha \cdot t + \delta z_{ht} + \lambda_{ht}$$

where λ_{ht} is the marginal utility of money and t is a trend. 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:

$$(19) \quad \begin{aligned} E_t(\ln c_{ht} | z_{ht} = 1) - E_t(\ln c_{ht} | z_{ht} = 0) \\ = \delta + E_t(\lambda_{ht} | z_{ht} = 1) - E_t(\lambda_{ht} | z_{ht} = 0) \end{aligned}$$

In general there is no reason to believe that $E_t(\lambda_{ht} | z_{ht} = 1)$ and $E_t(\lambda_{ht} | z_{ht} = 0)$ have the same time means so that the estimator that pools cross-sections will be a biased estimator of the

child effect δ . Moreover, the bias varies with age in a non-linear way because of the dependence on the presence of children, so that there is no obvious adjustment for the bias.

Consistent estimation of the child response function is based on equation (9). This indicates that a simple regression of cohort mean consumption growth on the change in the proportion of households with children will identify the child response function as the coefficient on the latter:

$$(20) \quad \Delta c_{d,t+1} = \alpha_d + \delta \Delta p_{d,t+1} + \varepsilon_{d,t+1}$$

In this estimation we have to take account of the fact that using a time series of cross-sections effectively introduces measurement error. We also have to take account of the fact that there is uncertainty and fertility plans are not certain so that the cohort specific ‘news’ term $\varepsilon_{d,t+1}$ may contain new information on the fertility variable $p_{d,t+1}$. Both of these considerations suggests that we need instruments for the changes in demographics. The obvious instruments are twice or more lagged demographics. These should be strong instruments once we take account of the numbers and ages of children. Moreover, the measurement error induced by the sampling procedure and expectations errors will be uncorrelated with the instruments.

B Estimating completed fertility

Turning to our second sampling scheme, we see from equation (12) that we need estimates of the proportion of currently childless households in cohort d that never have children, $m_{d,t}$. Since $p_{d,t}$ can be observed, this requires an estimate of π_d , the proportion of households that have children at some point. Using the data to hand⁹ we have to impose some assumptions to be able to identify π_d . 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. Given a consistent estimate of π_d and hence of $m_{d,t}$ we see from equation (17) that a regression on a sample of currently childless households of the first difference of mean log consumption on the first difference of $m_{d,t}$ has a coefficient which can be interpreted as the mean over time difference between $E_{II}^d [c_{ht}]$ and $E_I^d [c_{ht}]$. 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.

⁹An alternative would be to use outside information, such as a retrospective fertility survey on the same population.

III The data.

A 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 regularly since 1968, allowing us to follow some birth cohorts for more than 30 years. The sample used in this paper covers the period 1968 to 1999. As usual in studies of household consumption we limit ourselves to households consisting of cohabiting couples.¹⁰ The measure of consumption used in this paper is the nominal expenditure on all goods purchased (excluding housing expenditures) during a two-week period, deflated by the consumer price index for total expenditure excluding housing prices. Furthermore, to remove year effects consumption is adjusted by using the residuals from a regression of household log consumption on year dummies.

B Constructing cohorts

The cohorts we use are based on year of birth and education but contrary to most other studies we use wife's birth year and the husband's education. 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.¹¹ For constructing birth cohorts we use five-year bands where the wife in the eldest cohort is born between 1915-1919 and the youngest cohort is born in 1970-1974.

Another important aspect in constructing homogenous groups is educational attainment. As is well known there are significant differences in the number and timing of children across education groups. Educational attainment also seems to be important 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

¹⁰Extended families (which means families with other adults than grown up children) are excluded.

¹¹A 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.

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 quasi-panel of cohort means. In order to be able to use the data from the period 1968-1977 we use an imputation method; details are given in the Appendix.

C Consumption profiles

The cohort sample means are constructed by averaging over households belonging to the same birth cohort and education group observed in a given year. We construct 12 birth cohorts crossed by two education groups. 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. When estimating the parameters of the model 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 (to avoid problems with retirement). For the more educated group a serious problem is the selection in to marriage/cohabitation. For that reason we have chosen to only look at households where the average cohort age of the wife is above 25. Finally we drop any cells that have less than 75 households. The final sample consists of 93,921 households of whom 60% belong to ‘the less education group’. The synthetic panel contains 450 cohort-year observations.

The cohort sample means are plotted against the average age of the wife in figure 1 (the unconditional sampling scheme of section I.B). As can be seen, 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 its 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 2, we have present the consumption profiles for households where there are no children (or other adults) currently living in the household (the no children sampling scheme of section I.C). 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 do not have at least one child present. As discussed in section I these figures would display an inverted-U shape even if everyone has a flat profile. As can be seen, the nonlinearities are relatively modest as compared to the unconditional sampling scheme.¹² For example, for the more educated the unconditional mean of log consumption rises and then falls by about 0.3, whereas for the no children sample the rise and fall is of the order of 0.08. For the less educated group the consumption profile is still hump-shaped although the picture is not that clear. A closer look at the graph reveals

¹²Although modest, the non-linearities are significant according to tests for a quadratic term in a simple fit against age




Figure 1: The consumption of the full sample

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 consists mostly of households who never have children. When examining the ‘more educated group’ the graph is very noisy due to the very low number of households in each cell.

IV Modeling the demographic effects

A Estimating child response functions.

In this section we estimate the effects of the demographics on consumption paths. This sub-section concerns estimating the child response function, following the strategy outlined in the econometric section. Compared to the simple one-child model we extend the model by allowing that the child response function depends on the age of the child and for economies of scale.¹³ In the second sub-section we estimate the completed fertility from the currently observed fertility (the π_d variables of sections I.C and II.B).

Whilst the authors of 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 child response function and the age of children. 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 way. First, we assume that consumption depends on the age of the children in a continuous monotonic way. 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.¹⁴ 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: Martin Browning (1999) analyses a theoretical non-unitary model of intertemporal allocation in a two person household and Paul T. Schultz (1999) discusses 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.

¹³We do not take account of spacing considerations that suggest that the child response functions may depend on the closeness in age of siblings.

¹⁴Or 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.




Figure 2: Consumption of the "no children" sample

We define the child response function in two steps. In the first step we adjust for the ages of the individuals in the household and sum to give the number of equivalent adults. In the second step we take account of economies of scale. For the first step we define a function that depends on the age of the each individual in the household:

$$(21) \quad f(\text{age}_j) = \mu_0 + \mu_1 \frac{\text{age}_j}{20} + \mu_2 \left(\frac{\text{age}_j}{20} \right)^2 + \mu_3 \left(\frac{\text{age}_j}{20} \right)^3$$

where age_j is equal to the maximum of the individual's age and 20. In estimation we impose:

$$\mu_3 = 1 - \mu_0 - \mu_1 - \mu_2$$

so that the function is continuous with $f(0) = \mu_0$ and $f(20) = 1$. Thus the parameter μ_0 is the impact of an infant ($\text{age}_j = 0$) child and all households members aged 20 or more have a value of unity. Denote the number of adults living in household h in period t by N_{ht} and define the number of *equivalent adults*, n_{ht} , by:

$$(22) \quad n_{ht} = \sum_{j=1}^{N_{ht}} (f(\text{age}_j))$$

Clearly this value is bounded below by 2 (the value for a married couple with no else currently present in the household). Some authors use restricted versions of this formulation. For example, Attanasio et al (1999) do not allow for age effects for children but do allow that adults and children can have different effects:

$$(23) \quad n_{ht} = n_{ht}^a + \mu n_{ht}^c$$

where n_{ht}^a and n_{ht}^c are, respectively, the numbers of adults and children in the household. The authors of National Research Council (1995) make the same suggestion with a value of μ equal to 0.7. Gourinchas and Parker (2002) further restrict age composition effects and use only family size which is equivalent to setting μ equal to unity. As we shall show below, these assumptions have a substantial impact on the inferences concerning the path of consumption.

Denoting the consumption of a household with n equivalent adults by $C(n)$ we set:

$$(24) \quad C(n) = e^{k(n)} C(1)$$

To allow for economies of scale effects we take the following parameterisation:

$$(25) \quad k_{ht} = \delta_1 \ln(n_{ht}) \Rightarrow C(n) = n^{\delta_1} C(1)$$

The degree of economies of scale is captured by the parameter δ_1 . If the parameter δ_1 is equal to zero then there is no change in consumption consequent on an anticipated change in the

composition of the household.¹⁵ On the other hand if δ_1 equals one there are no economies of scale ($C(n) = nC(1)$). Thus the values the scale parameter is bounded between zero and unity. Not all specifications used in the literature respect these bounds. For example, specifications of the form $\ln c = \beta_0 + \beta_1 n^a + \beta_2 n^c$ are widely used (see, for example, Attanasio and Browning (1995)) but these imply diseconomies of scale. The authors of the National Research Council (1995) proposal suggest taking a value equivalent to approximately 0.6. Gourinchas and Parker (2002) 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 (see equation (20)). That is, we run a regression of first differences in cohort mean log consumption on first differences of the mean scale. :

$$(26) \quad \begin{aligned} \Delta E_{\chi t}(\ln c_{ht}) &= \delta_0 + \Delta E_{\chi t}(k_{ht}) + \varepsilon_{\chi t}. \\ &= \delta_0 + \delta_1 \Delta E_{\chi t}(\ln(n_{ht})) + \varepsilon_{\chi t} \end{aligned}$$

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 child response function enter in a nonlinear way we have to perform the estimation in two steps. Given the parameters of the child response function function, the child response function for each household can be calculated and then aggregated into cohort sample means. On the basis of cohort sample means, the parameters δ_0 and δ_1 can be estimated and the cohort residuals can be determined. To estimate the parameters we use a GMM estimator. The instruments used in this article are based the number of children in different age categories: 0, 1,2, 3-4, 5-10, 11-16, 17-18,19-21 22-25 and 26-30. The variables are aggregated into cohort means and first differences variables are constructed. The instruments consist of these variables in first differences lagged one and two times and a constant. This leaves us with 21 instruments denoted by $Z_{\chi,t}$. The parameters of the child response function are then determined such that the cohort residuals projected on the instruments are minimized. Let $\theta = \{\delta_0, \delta_1, \mu_0, \mu_1, \mu_2\}$ and the criterion function is then given by

$$K(\theta) = \left(\sum_{\chi,t} (\varepsilon_{\chi t})' Z_{\chi,t} \right) V(\theta)^{-1} \left(\sum_{\chi,t} (\varepsilon_{\chi t})' Z_{\chi,t} \right)'$$

The estimation is performed by using a continuous-updating GMM estimator (see Lars P. Hansen, John Heaton and Amir Yaron, 1996). The covariance matrix of the parameters is

¹⁵Once again we emphasise that this restriction does not mean that ‘two can live as cheaply as one’; the thought experiment here holds the marginal utility of money constant but not the utility level.

estimated by:

$$V(\theta) = D'(\varepsilon M \varepsilon')^{-1} D$$

$$D = \frac{\partial(Z\varepsilon)}{\partial\theta}$$

In the estimation we assume that (ε_{xt}) is following an MA(1)-process.
The estimation results are reported in Table 1.

Table 1: Child response parameters.		
	Less educated	More educated
μ_0	-0.266 (0.093)	-0.091 (0.066)
μ_1	2.847 (0.686)	2.469 (0.394)
μ_2	-6.119 (2.084)	-5.730 (1.337)
δ_0	-0.012 (0.001)	0.002 (0.002)
δ_1	0.799 (0.162)	0.921 (0.224)
Sargan (d.f.)	10.997(16)	7.991(16)
p-value	0.809	0.949
N	56322	37599
cohort obs	229	221

All the parameter in the models, except δ_0 (which is discussed in the next section), relate to the impact of children. Since the parameters are somewhat difficult to interpret, we present some illustrative examples in Table 2. For comparison we normalise the consumption of a two adult household to unity.

Number of children	Age of children			Less educated		More educated	
				n	C	n	C
0				2	1	2	1
1	0			1.73	0.89	1.91	0.96
2	0	1		1.59	0.83	1.93	0.97
1	10			2.19	1.07	2.25	1.11
2	10	11		2.39	1.15	2.51	1.23
3	10	11	12	2.62	1.24	2.78	1.36
1	15			2.34	1.13	2.37	1.17
2	15	16		2.76	1.29	2.82	1.37
3	15	16	17	3.28	1.48	3.36	1.61
1	20			3	1.38	3	1.45
2	20	20		4	1.73	4	1.89
3	20	20	20	5	2.08	5	2.33

n is the number of equivalent adults.

C is the level of consumption.

Consider first the scale parameter, δ_1 . Our estimates are somewhat above the value of 0.6 suggested in National Research Council (1995) so that we find weaker economies of scale. We conjecture that this is because we take account of age effects. For the less educated group, consumption increases by 38 per cent if the family increases from two adults to three adults while for the more educated group the similar number is 45 per cent.

We also find a significant impact of the age of the children. In figure 3 we graph the estimated function defined in equation (21) for the two education groups. The estimates imply that the impact of a child is increasing in age (which was not imposed in estimation) and that an infant actually has a negative impact on the total consumption. For example, we find that having one small child actually lowers consumption by about 11 per cent for the less educated group and 4 per cent for the more educated group. The reason for this somehow surprising results is that household with very small children change the composition of consumption. We are not claiming that parents do not spent money on small children, but this increase in child related consumption is outset by a decrease in other consumption components such as entertainment, tobacco, food in restaurants etc.. Since this is a somewhat surprising finding, in Appendix B we give some details of expenditures on different consumption items. From the age of 3 for the less educated and for the age of 1 for the more educated an extra child increases total consumption. These estimates indicate that both the age of the child and the economies of scale are important when considering child response functions.

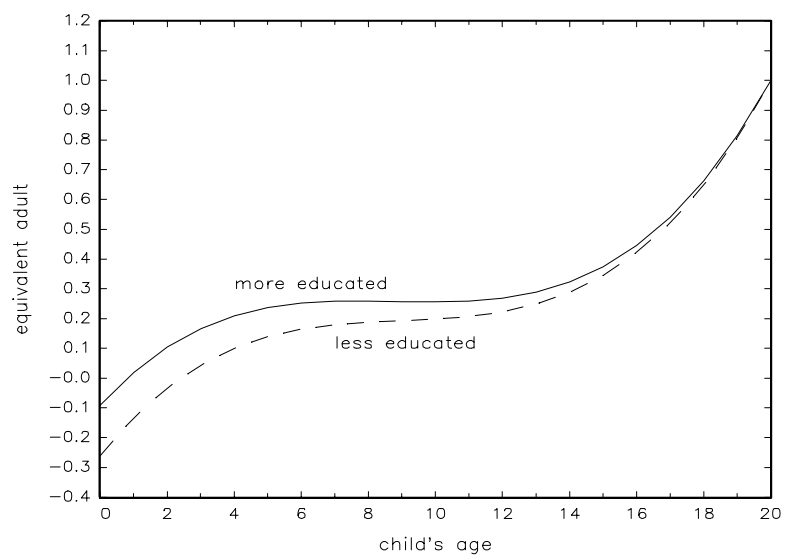


Figure 3: Child response function

B Completed fertility

In the econometrics 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. First, we assume that at one age during the life cycle, completed fertility is observed. 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 of course a very restrictive assumption because women can have children later or some might have had children very early who already have left home.¹⁶ The data indicate 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 the use of this assumption. The second assumption we make is that the 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 4, 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.

¹⁶In Appendix A4 we examine the sensitivity of our results to this assumption. To do this we use an alternative data set which contains information on completed fertility for the populations considered here. As detailed in Appendix A4, we use estimates from these data to adjust the current procedure. The conclusion of this investigation is that assuming that we observe completed fertility at age 37 or 38 does not have a significant impact on the final results.

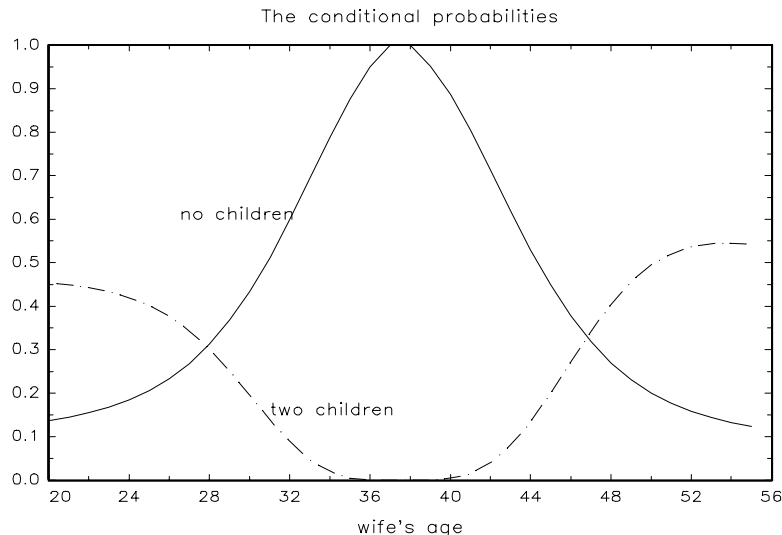


Figure 4: The probability of completed fertility conditioned on no children in the household

By construction the proportion of households which never have children should be unity at the age 37 and 38. Since these ratios are used to construct the counterfactual 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 internally with the FES data.

Based on the estimated conditional distribution of completed fertility we can estimate the relative expenditures of ‘childless’ and ‘children’ households from the consumption of the ‘no children’ group. The estimation is based on equation (17), which in the framework with four types of households can be generalised to

$$E_{\chi_{0,t}}[c_{h,t}] = \gamma_{\chi_0} + \alpha * age + \tau_1 m_t^1 + \tau_2 m_t^2 + \tau_3 m_t^3 + \omega_{\chi_{0,t}},$$

where $E_{\chi_{0,t}}$ denotes the cohort mean conditioned on no children present in the household. The variable m_t^j is the conditional probability of having a completed fertility of j children but at

time t none of the children are living in the household. Since we use the predicted values of m_t^1 , m_t^2 and m_t^3 we ensure that these are independent of the sampling error $\omega_{\chi^0,t}$. However, we may expect imprecise estimates due to the fact that we often have small cell sizes. The estimates are reported in Table 3. In both cases the effects are monotone and negative except for more educated households that have only one child. The estimates indicate that a currently childless household which plans to have two children should lower consumption by 3 per cent for the low education group and 15% for the educated group in order to save for the higher consumption and lower income when the child is present. The results further show a particularly strong and significant effect for the 3+ groups. Within the framework considered here, these effects reflect both the consumption needs of children and the earnings responses of households to planned fertility. For example, a positive correlation for wives between the tastes for going to work and for having children will lead to a negative relationship between completed fertility and lifetime earnings of the couple, even if children have no direct effect on consumption.

Table 3: Estimates for the ‘no children’ sample		
	Less educated	More educated
m_t^1	-0.6 (5.9)	15.3 (9.3)
m_t^2	-3.0 (4.8)	-14.8 (4.8)
m_t^{3+}	-41.1 (8.9)	-30.7 (6.7)
<i>age</i>	-0.3 (0.7)	0.1 (0.09)
Number of cohorts	232	227

In the estimation we control for cohort effects.
All values multiplied by 100. The values in the bracket are standard errors

V Adjusted consumption.

A 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 linear trend and a constant, which may depend on the cohort. This means that the adjusted cohort consumption should be a linear trend over the life cycle. The counter-factual consumption can be constructed for both the unconditional sample or from the no children sample. In the following adjusted consumption is established on the basis of the correction outlined in the previous sections.

When adjusting the unconditional sample, the estimates of the child response function given in Table 1 are used (see equation (10)).¹⁷ The unadjusted and adjusted consumption paths are shown in figure 5. The consumption paths of both education groups are close to being linear for each cohort (formal tests are presented below). However, the trend in adjusted consumption is different for the two education groups. For the less educated adjusted consumption is declining at about 1.2% per year whereas for the more educated consumption is increasing by about 0.2% per year (up to age 55). Since we have taken out common time effects by the first round regression on time dummies, only the differences between these mean much. The differences in the trends can be related to differences in discount factor for the two education groups ($\delta_0 = E[\alpha_h] = E[-\theta^{-1}(\ln(\beta_h(1+r)))]$). Our results indicate that more educated have a higher discount rate than less educated, under the assumption that the two groups have the same θ and the same real rate.

For both groups there is evidence that younger cohorts have higher consumption but the effect is much stronger for the more educated. The differences between the cohort effects for the two education groups are consistent with recent findings on the divergence between the wages of different education groups.

To test if there is any non-linear age pattern once we correct for the presence of children we propose the following test. For the unconditional sample we correct the consumption for each household for presence of children by using the estimates given in Table 1. The adjusted consumption on household level is then aggregated into cohort means (cohort and year observations). In order to eliminate cohort effects we use first differences on the cohort data. The following regression equation is estimated:

$$\Delta E_{\chi t}(\ln c_{ht}) - \hat{\delta}_1 \Delta E_{\chi t}(\ln(n_{ht})) = \rho_0 + \rho_1 \Delta E_{\chi t}(age_{ht}) + w_{\chi t}$$

Subsequently we test the hypothesis: $\rho_1 = 0$, which corresponds to that the corrected consumption only contains a linear trend in age. The test is performed as a t-test. For comparison we also perform the test on the unadjusted consumption. The estimation and test results are reported in the Table 4. The results clearly show that there are significant nonlinearities in age for the unadjusted data but not for the adjusted paths.¹⁸

Table 4: Tests for nonlinearities (unconditional sample).				
	Unadjusted consumption		Adjusted consumption	
	less educated	more educated	less educated	more educated
ρ_1	-0.012	-0.0016	-0.0005	-0.0005
t-test	-3.65	-3.65	-1.70	-1.26
p-value	0.000	0.000	0.090	0.208

¹⁷Notice, that we do not correct for the trend in consumption, which corresponds to the estimate of δ_0 in equation (26).

¹⁸Different specifications of the age-pattern have been tried but in all cases they were not significant in the adjusted consumption.




Figure 5: The adjusted and unadjusted consumption for the full sample

B The ‘no children’ sample

Our second sampling scheme uses the ‘no children’ sample. because of the noisiness of the data, the figures of the adjusted and unadjusted paths are not very illuminating so we present only formal tests. Note first that from Table 3 we see that for the more educated group there is an upwards trend in the adjusted consumption, while for the less educated the trend is downwards sloping. To formally test for an age pattern in the adjusted and unadjusted consumption we use a similar idea as for the full sample. However, for this exercise we use the levels of log consumption (see equation (17)):

$$E_{\chi_{0t}}(\ln c_{ht}) - \hat{\tau}_1 m_{\chi_{0t}}^1 - \hat{\tau}_2 m_{\chi_{0t}}^2 - \hat{\tau}_3 m_{\chi_{0t}}^{3+} = \gamma D_{\chi_{0t}} + \rho_0 + \rho_1 E_{\chi_{0t}}(age_{ht}) + \rho_1 E_{\chi_{0t}}(age_{ht}^2) + w_{\chi t}$$

where $D_{\chi_{0t}}$ is a set of cohort dummies. The test is a test of the hypothesis: $\rho_2 = 0$. Table 5 reports the results of the tests.

Table 5: Tests for nonlinearities (no children sample).				
	Unadjusted consumption		Adjusted consumption	
	less educated	more educated	less educated	more educated ^a
ρ_2	-0.0004	-0.0002	-0.00002	-0.0000
t-test	-9.06	-9.68	-0.48	-0.35
p-value	0.000	0.000	0.633	0.728

a: The test for the more educated group is performed for the age group 25-55

The results show that the no linear age pattern in the adjusted consumption is removed for both groups.

VI 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 if we take proper account of the numbers and age of children, family composition can ‘explain’ completely the hump-shape in consumption. This conclusion is based on two largely independent sets of tests, one using a conventional unconditional sampling scheme and the other using on a sample of households with no children present. Our analysis is based on a model in which forward looking households make provision for the future consumption needs and possible falls in income consequent on having children in the household by adjusting consumption down when they are not present. Our finding is in line with the evidence presented in Adriaan Kalwij (1999) who finds that for a sample of Dutch households savings rise before the birth of a child and decline after. We also have the following supplementary conclusions:

- Households that hold the marginal utility of expenditure constant *reduce* consumption consequent on the arrival of an infant in the household (see Tables 1 and 2). A detailed analysis of the associated demand effects is given in Appendix B. As we have been at pains to emphasise throughout this paper, this cannot be interpreted as a ‘utility constant’ change and it does not imply anything about the value of the adult equivalence scale that would be required to keep the parents’ utility constant.
- Consumption is increasing in the age of children with a relatively flat portion of about 0.25 equivalent adults between the age of five and fourteen; see figure 3. The sharp increase in consumption from age 16 onwards may be associated with the increasing earnings of the child and a consequent failure of the unitary model.
- Consumption is increasing in the number of equivalent adults in the household, and there are modest economies of scale (see Table 2). For example, less educated households with 3, 4 and 5 equivalent adults respectively spend 38%, 73% and 108% more than a two adult household.
- Due to the combination of age effects and scale effects consumption in households with children aged about 11 are not greatly higher than before the children are born. For example, each extra child of about age 11 increases consumption by about 8% for less educated households and by about 12% for more educated households.
- We find that higher parity households have much lower consumption when they do not have children present (see Table 3). For example for more educated households, there is no significant difference for one child but for two children or three or more children, consumption falls by 15% and 31% respectively. Given the data to hand we cannot determine whether this is due to the consumption costs of children or to their impact on lifetime earnings.
- Lower educated households have a lower trend in consumption than more educated households. This could be interpreted as evidence that lower educated households have a higher discount rate.
- We find positive cohort effects for consumption with stronger effects for more educated households. This is in line with recent research on the diverging paths of earnings for different education groups.
- On the modelling side, our analysis shows that when working with repeated cross sections selected according to time varying variable, such as the presence of children, raises very different issues arise than when selecting on time invariant variable such as year of birth. This might also be relevant in other contexts. For example, if the selection is on the basis of occupation rather than education.

- We also develop a method to impute education to allow us to use the longer time series for which education is not always available.

As always our conclusions are subject to many reservations concerning the specification. The most important assumptions we make are that we can ignore the selection into and out of the sample (marrying and divorcing); we assume separability between non-housing expenditures and housing; we have a very specific functional form for child effects; we ignore intra-household issues which may be particularly important if there are adult children in the household; we assume that all common effects (including time variation in the real rate of interest) can be captured by an initial regression on time dummies. Subject to these caveats, we find that we can capture consumption movements within the working and married life without recourse to a precautionary motive.

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] Blundell, Richard, Martin Browning and Costas Meghir (1994), "Consumer demand and the lifetime allocation of consumption", *Review of Economic Studies*, 61, 57-80.
- [5] Blundell, Richard and Thomas MaCurdy (1999), "Labor Supply: A review of alternative approaches", *Working paper* Institute of Fiscal Studies
- [6] Browning, Martin (1999), "Saving in a two person household", *Scandinavian Journal of Economics*, forthcoming.
- [7] 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.
- [8] Browning, Martin, Lars Peter Hansen and James J. Heckman (1999), "Micro Data and General Equilibrium Models", *Discussion paper*, Institute of Economics, University of Copenhagen.
- [9] Calhoun, Charles A. and Thomas S. Espenshade (1988), "Childbearing and wives' foregone earnings", *Population Studies*, 42(1), 5-38.
- [10] Carroll, Christoher (1994), "How does future income affect current consumption?", *Quarterly Journal of Economics*, 109, 111-148.
- [11] Carroll, Christoher (2001), "Death to the Log-Linearized Consumption Euler Equation!", *Advances in Macroeconomics*, .
- [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 (1985), "Panel Data from Time Series of Cross-Sections", *Journal of Econometrics*; 30(1-2), 109-26.

- [14] Deaton, Angus (1991), "Saving and liquidity constraints", *Econometrica*, 59, 1221-48.
- [15] Gourinchas, Pierre-Olivier and Jonathan Parker (2002), "Consumption over the life cycle", *Econometrica*, 70(1), 47-90.
- [16] Hansen, Lars Peter, John Heaton and Amir Yaron (1996), "Finite-Sample Properties of Some Alternative GMM Estimators", *Journal of Business & Economics Statistics*, 14, 262-280.
- [17] 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.
- [18] Hubbard, Gleen R., Jonathan Skinner and Stephen Zeldes (1994), "Precautionary saving and social insurance", *Journal of Political Economy*, 103, 360-399.
- [19] Joshi, Heather (1990), "The cash opportunity costs of childbearing: an approach to estimation using British data", *Population Studies*, 44, 41-60.
- [20] Kalwij, Adriaan (1999), "Female employment and fertility decisions: an analysis of the effects of household financial wealth", mimeo, Oxford University.
- [21] Ludvigson, Sydney and Christina Paxson, (2001) "Approximation Bias in Linearized Euler Equations", *Journal of Economics and Statistics*, vol. 83, no:2, 242-256.
- [22] Lucas, Robert and Leonard A. Rapping (1970), "Real wages, employment and inflation", in E.S. Phelps (ed) *Microeconomic Foundations of Employment and Inflation Theory*, New York.
- [23] Nagatani, Keizo (1972), "Life cycle saving: theory and fact" *American Economic Review*, 62, 344-53.
- [24] National research council (1995), *Measuring Poverty: a new approach*, National Academy Press, Washington DC.
- [25] Schultz, Paul T.(1999), "Who is head of the household", presented at the ESPE conference, Turin 1999.
- [26] Thurow, Lester (1969), "The optimum lifetime distribution of consumption expenditures" *American Economic Review*, 59, 324-30.
- [27] 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.

A Appendix

A Imputing 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 1978-1999 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:

$$(27) \quad \begin{aligned} \bar{x}_{h,t}^0 &= \frac{\sum_{h \in \chi_t} (1 - e_h) x_{h,t}}{\sum_{h \in \chi_t} 1 - e_h} \\ \bar{x}_{c,t}^1 &= \frac{\sum_{h \in \chi} e_h x_{h,t}}{\sum_{h \in \chi} e_h} \end{aligned}$$

where the superscripts denote the education group. The alternative way of constructing cohort means is simply by replacing the imputed education, e_h , in (27) 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 table A1 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. 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 over-predicts the difference between the education groups while using the weighting procedure under-predicts. On the basis of pictures for 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. On the basis of the education dummy which is observed from 1978-1999 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 polynomial 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 65 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 71,971 observations and the R^2 is equal to 0.2003.

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.

Actually	Imputed		Total
	less educated	more educated	
Less education	30,971	9,658	40,637
More educated.	9,666	21,676	31,334
Total	40,629	31,342	71,971

Pearson's goodness of fit test: $\chi^2(1) = 14827.66$ p-value= 0.00

B Comparison of expenditures for couples without children and with an infant

In the estimation of child effects we found that the presence of small children in the household actually lowers consumption. To provide more evidence on this somehow surprising result, we have looked at the expenditures on different consumption items. Using the FES data from 88-94 we select a sample of married couples with the wife aged between 25 and 30. Furthermore we restrict the sample to only including couples with no children or one child aged less than one. To make the analyses comparable to our previous analyses we divided the sample according to the educational attainment of the husband. In Table A2 we compare the group of couples with no children to the group that have one baby for each education group.

Table A2: Expenditures on different consumption items

Table A2: Expenditures on different consumption items				
	Less educated		More educated	
	No child	One infant	No child	One infant
Number of obs	943	273	1193	265
Age of the wife	27.18	27.20	27.27	27.77
Age of the husband	30.15	29.70	29.49	30.48
Total expenditures	267.43	234.23	312.94	280.60
Selected items:				
Alcohol	12.16	8.49	11.96	9.20
Tobacco	4.28	3.77	2.34	2.03
Food out	13.30	9.03	17.16	9.49
Domestic ser.	0.94	2.21	1.51	5.79
Leisure	18.25	9.47	23.84	11.82

First we see that these two groups are similar in terms of the age of the husband and wife except. The comparison of the total expenditures of the two groups shows that for both education groups, couples with a small child spent less.¹⁹ The reduction in consumption for the less educated group is about 13 percent and for the more educated group about 10 percent. If we look at selected consumption items we find that expenditures on alcohol, food out and leisure are much smaller for couples with a small child.²⁰ In particular the expenditures on leisure for couples with a small child is only half of that for the no child couples. This supports our hypothesis that consumption is actually lower in households with small children. One offsetting category is domestic services, which includes child care; this is higher for couples with a child.

¹⁹All the expenditure are deflated by the consumer price index.

²⁰Leisure includes television licence, entertainment and holidays.

C 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:

$$m_t^j = \Pr(z_h = j | z_{ht} = 0) \quad j = 0, 1, 2, 3 +$$

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

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

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

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 polynomial in age and cohort dummies. Based on this calculation the unconditional distribution of types are determined to:

Table A3: Imputed fertility pattern for women aged 37 and 38 (FES data)		
Completed fertility	less educated group	more educated group
no children	7.93	9.51
one child	16.51	16.14
two children	40.84	47.40
more than two children	34.70	26.92

From the source of data available, there are two alternative ways of determining m_t^j . The first method exploits the fact that we know the age of the children living in the household. 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 m_t^j 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 m_t^j is obtained.

The disadvantage of this method is that it is not possible to construct estimates of m_t^j 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

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

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

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

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 m_t^j .

D The alternative data set

To complement our analyses with the FES data we use an alternative data set: the UK version of the European Community Household Panel (ECHP). This data set is used to examine the assumptions we made on the fertility pattern of the women in the FES data set. The advantage

of this data as compared to the FES is that it contains information on completed fertility. In the survey, all households are asked how many children they had have, the year of birth of their children and how many children who are living outside the household. This information is very valuable in our context since it allows us to investigate some of our critical assumptions concerning the fertility pattern used for the FES data set. Unfortunately since the ECHP has limited duration we are not able to examine how fertility patterns change over time.

In the analyses of the FES data we impose the assumption that the completed fertility is revealed when the wife is aged 37 and 38. This implies that women do not have children after the age of 37 and that no children have left home before the age of 37 of the mother. These assumptions can be examined with ECHP data. In the ECHP we select households consisting of married or cohabiting British couples (extended families are excluded in order to make it comparable to the FES data). In table A4 an FES sample of women aged 37 to 38 in 1994 is compared with the ECHP. In both data sets we show how many children who were living in the household.

Completed fertility	FES	ECHP
no children	10.45	11.69
one child	10.45	15.58
two children	50.25	41.56
more than two children	28.86	31.17
N	201	154

From the table it is seen that there are more households with two children in the FES data set than in the ECHP. However a formal test shows no significant differences in distribution of number of children for the two data sets. To make this as comparable as possible to the FES data we divide the households into two groups according educational attainment of the husband.

To investigate the assumption that women do not have children after the age of 37, we select a sample of households where women are 50 and older in 1994.²¹ In the table A5, the completed fertility of the two education groups is shown. In order to compare these numbers with the FES data we select the same group in the FES data (women aged more than 50 in 1994) and use the method described above to calculate the completed fertility. The "actual" fertility is based on the actual number of children observed, while the imputed fertility is estimated on the basis of a multilogit model. In table A6 the results are shown. A comparison of the completed fertility shows that we find more no children households in the FES data. In the FES sample this fraction is estimated to around 7 percent (the imputed fertility), while we get around 5 percent with the ECHP. This means that we might over-estimate the group of no children in

²¹We are assuming that women do not have children after the age of 50.

the FES data, especially for the more educated group. The other figures seem very much in accordance with the FES data set, except that the fraction of the two children households are under-estimated for the educated group. Looking at the fraction of women having their first child after the age of 37 this fraction is very small. It seems more seriously that about ten percent of the less educated households and about six percent of the more educated households are having children after the age 37. However, further analyses shows that the majority of these households end up having more than two children. Since we do not distinguish between three children and more than three children, this means that it is not very likely that the completed fertility is unduly affected by these households.

	Less educated	More educated
Completed fertility	group	group
no children	5.13	4.62
one child	16.50	15.13
two children	36.55	50.00
more than two children	42.82	30.25
fraction having first birth after 37 (%)	1.96	1.26
fraction having last birth after 37 (%)	9.90	6.30
N	818	238

	Less educated		More educated	
Completed fertility	group		group	
	Actual	Imputed	Actual	Imputed
no children	6.33	7.12	9.26	7.85
one child	15.69	16.23	15.11	15.59
two children	37.19	38.11	41.70	45.29
more than two children	40.71	38.53	33.94	31.28
N	1753		940	

To investigate the second part of the assumption we select a sample of household where the women are between 25 and 37. Here we examine how many of these households where some of the children already have left home. In table A7 the fractions are shown. The table show that in about 15 per cent of the households with a mother below 37 at least one child is not staying in the household. A closer look at the children who are not living in the household shows that more than half of the children are below 10 years old. Most of these children where associated with households where the mother was divorced or separated, but living together with a new

man.²² The fraction of households where older children (above 15) have left home is about 3 per cent for less educated households and two per cent for the more educated households. This indicates that there are two aspects which we have not dealt with in the present approach. Firstly that some children might already have left home and secondly that some children might not live in the same household as their mother (e.g. divorce). However, again if we take a look at these households we find that many of these households are having more than two children. So we conclude that this might not affect the distribution of the completed fertility.

Number of children who are not in the household	less educated group	more educated group
no children	84.62	89.47
one child	10.99	5.26
two children	1.10	5.26
more than two children	3.30	0
Fraction of households with children above 15 who have left home		
	3.29	2.63
N	91	38

In order to check the sensitivity of the assumptions imposed, we have tried to correct the estimated fertility pattern of FES by using the information of the ECHP. We use a very simple correction. For the very specific group (household where the wife is aged more than 50 in 1994) we know the completed fertility. By generalizing this results to the whole sample we are able to adjust the estimated fertility pattern of the FES data. On the adjusted sample we repeat the estimation on the total cost of children. In table A8 the new results are shown. For the both groups the results are very similar to the previous obtained.

²²For the remaining part of the young children living in another household than their mother, it could be because of divorce. If the mother had remarried we are not able to distinguish between the biological father and a step-father.

Table A8: Estimates of 'no children' sample, corrected by ECHP.

	Less educated	More educated
m_t^1	0.004 (0.055)	0.128 (0.083)
m_t^2	-0.033 (0.045)	-0.122 (0.041)
m_t^{3+}	-0.430 (0.091)	-0.317 (0.071)
age	-0.003 (0.007)	0.007 (0.003)
Number of cohort obs	232	227

In the estimation we control for cohort effects

The numbers in the bracket are standard errors