

TECHNICAL WORKING PAPER SERIES

ESTIMATING EULER EQUATIONS

Orazio P. Attanasio
Hamish Low

Technical Working Paper 253
<http://www.nber.org/papers/T0253>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2000

We would like to thank Chris Carroll, Costas Meghir and Guglielmo Weber for useful discussions. The views expressed herein are those of the authors and are not necessarily those of the National Bureau of Economic Research.

© 2000 by Orazio P. Attanasio and Hamish Low. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Estimating Euler Equations
Orazio P. Attanasio and Hamish Low
NBER Technical Working Paper No. 253
May 2000
JEL No. E2

ABSTRACT

In this paper we consider conditions under which the estimation of a log-linearized Euler equation for consumption yields consistent estimates of the preference parameters. When utility is isoelastic and a sample covering a long time period is available, consistent estimates are obtained from the log-linearized Euler equation when the innovations to the conditional variance of consumption growth are uncorrelated with the instruments typically used in estimation. We perform a Montecarlo experiment, consisting in solving and simulating a simple life cycle model under uncertainty, and show that in most situations, the estimates obtained from the log-linearized equation are not systematically biased. This is true even when we introduce heteroscedasticity in the process generating income. The only exception is when discount rates are very high (47% per year). This problem arises because consumers are nearly always close to the maximum borrowing limit: the estimation bias is unrelated to the linearization. Finally, we plot life cycle profiles for the variance of consumption growth, which, except when the discount factor is very high, is remarkably flat. This implies that claims that demographic variables in log-linearized Euler equations capture changes in the variance of consumption growth are unwarranted.

Orazio P. Attanasio
Department of Economics
University College London
Gower Street
London WC1E 6BT
United Kingdom
and NBER
o.attanasio@ucl.ac.uk

Hamish Low
Faculty of Economics
University of Cambridge
Sidgwick Avenue
Cambridge CB3 9DD
United Kingdom
and Institute for Fiscal Studies
hamish.low@econ.cam.ac.uk

1. Introduction

Much of the empirical literature on consumption has used extensively Euler equations derived from the inter-temporal optimization problem of a generic consumer to estimate structural parameters and test the restrictions implied by the model. Since Hall's (1978) contribution, the Euler equation approach has been applied to both micro and macro data sets to fit a variety of different preference specifications in a variety of different contexts. The main attraction of this approach lies in the fact that it avoids solving explicitly the optimization problem faced by a generic consumer and focuses instead on a specific first order condition implied by that problem. In other words, the Euler equation, by differencing out the fixed effect constituted by the marginal utility of wealth, allows one to avoid its parameterization, while at the same time allowing the estimation of preference parameters. The estimation of these equations, however, is not without problems. Some of these problems have been widely discussed in the literature. Others are less well known.

In this paper, we stress the conditions under which structural preference parameters can be identified using an Euler equation. Our starting point is that to estimate consistently preference parameters and test over-identifying restrictions, aggregation issues require the use of micro data. Moreover, if one does not want to use a representative agent/ complete markets framework, it is crucial to use 'large-T' asymptotics. We then stress the importance of obtaining an equation that is linear in the parameters and with additive residuals. This points to the usefulness of log-linear approximations of Euler equation. Therefore, we discuss the circumstances in which the use of log-linearized Euler equations does yield consistent estimates of the structural parameters of the model.

Recently, the log-linearized Euler equation has received some attention. Ludvigson and Paxson (1999) have considered the approximation error of consumption functions derived from a log-linearized Euler equation and have shown that it can be considerable. The emphasis of this paper is different: we want to establish to what extent a log-linearized Euler equation can be used to estimate structural preference parameters.

Carroll (1998), in a colorfully titled paper, has recently argued that estimates of structural parameters based on log-linearized Euler equations yield seriously biased parameters. The conclusions we reach in this paper are quite different. We provide some Montecarlo evidence on the performance of Instrumental Variable estimation of a log-linearized Euler equation derived from a dynamic optimization problem with intertemporally separable and isoelastic utility function. We show that, if

the sample is long enough, given the variability of the relevant stochastic variables, one obtains well-behaved and relatively efficient estimates from log-linearized Euler equations. The main difference between our approach and Carroll's is in his focus on cross section rather than time series variability.

The main achievement of our paper can be summarized very simply. A log-linearized Euler equation yields consistent estimates of the parameters of interest (provided the Euler equation holds and the sample period is long enough) if the innovations to the conditional second (and possibly higher) moments of consumption growth and interest rates are uncorrelated with the instruments one uses in estimation. As the conditional variance of consumption growth is endogenously determined, the conditions under which this condition holds are difficult to establish analytically. Our Montecarlo simulations establish a set of such conditions. In a standard model and for a wide set of parameter values the simulations indicate no systematic biases arising from estimating structural preference parameters with a log-linearized Euler equation. This result also holds when we introduce heteroscedasticity in the process generating household income. In particular, we notice that the presence of precautionary behaviour does not seem to introduce any bias in the estimation of preference parameters. We also provide estimates of the process of the variance of consumption growth and discuss how under plausible parameter configurations it is unlikely that this 'omitted variable' from the log-linearized Euler equation is captured by demographic variables or introduces any significant bias.

As a by-product of our exercise, we also obtain a number of other results. For example, an increase in the variability of the interest rate (one of the two sources of variability in our model) does not necessarily decrease the precision of our estimates of the elasticity of intertemporal substitution and often increases it. The reason is quite intuitive: such an increase induces more variation in intertemporal prices and this allows more precise estimates of intertemporal preferences. We also show that, as to be expected, the precision of our estimates decreases with the variability of wages (the second source of randomness in our model) and increases with T . Therefore, how large T should be to get reasonably precise estimates is a function of the variability of wages and of the interest rate. Calibrating wage variance is a very challenging task. In our Montecarlo exercise, we calibrate it so that the simulated consumption growth mimics the variance of observed cohort data.

The use of log-linear approximations of the Euler equation is not without problems, however. The most important, in our opinion, is the fact that one loses the ability of identifying the pure discount

factor, which gets buried, together with the mean of the conditional higher moments, in the constant of the log-linearized equation. The interpretation of this constant is therefore problematic.

An issue related to the interpretation of the intercept is that of the interpretation of the parameters on any other variables that are found to be significant in the Euler equation. This is relevant since it turns out, in practice, that to fit micro data it is necessary to allow a number of variables, most notably demographic variables, to enter the equation. While it is in theory possible that such variables capture movements in the conditional variance of consumption, as argued for instance by Gourinchas and Parker (1999), in the final section we claim that a much more plausible interpretation of these variables is as preference shifters.

The rest of the paper is organized as follows. Section 2 outlines the problems associated with Euler equation estimation. Most of this material is well known, but this has not prevented many researchers from using estimators that are only consistent under very restrictive assumptions. It is therefore worthwhile emphasizing the conditions on data availability and on the nature of unobserved heterogeneity that guarantee consistency of various estimators.

Section 3 discusses the effects that log-linearization has on the estimates of Euler equation parameters and section 4 then presents Montecarlo evidence. This evidence on the consistency of log-linearization is important because without an equation that is linear in parameters it is very difficult to deal with measurement error and unobserved heterogeneity in micro economic data.

2. Euler equation estimation

Many studies of inter-temporal consumption use some version of the following equation:

$$(1) \quad E_t \left[\frac{U_c(C_{t+1}^h, z_{t+1}^h, v_{t+1}^h) \beta (1 + R_{t+1})}{U_c(C_t^h, z_t^h, v_t^h)} \right] = 1$$

where subscripts index time and superscripts index individuals. U_c denotes the marginal utility of consumption which is assumed to depend on: C , (non-durable) consumption; z , a vector of observable variables; and v , some random variable meant to capture unobserved heterogeneity across

individuals and/or measurement equation in consumption.¹ R is the real ex-post interest rate on an asset for which the individual consumer is not at a corner and β is the 'pure' discount factor. To estimate the structural parameters in equation (1) one needs to specify the functional form of the utility function, the variables that affect the marginal utility of non durable consumption and the interest rate whose variation one wants to use to estimate some of the parameters of interest.

The isoelastic specification $U(C, z, v) = \frac{C^{1-\gamma}}{1-\gamma} \exp(\theta z + v)$ is often used as it gives rise to relatively tractable specifications. In particular, computing the relevant quantities using such a specification, equation (1) above becomes:

$$(2) \quad E_t \left[\left(\frac{C_{t+1}^h}{C_t^h} \right)^{-\gamma} \exp(\theta \Delta z_{t+1}^h + \Delta v_{t+1}^h) \beta (1 + R_{t+1}) \right] = 1$$

Several considerations are in order if one is interested in estimating the parameters of equation (2).

(i) Equation (2) is *not* a consumption function, but an equilibrium condition. Such a condition can be used to derive orthogonality conditions that can be used to estimate the parameters of the utility function and test eventual over-identifying restrictions.

(ii) An assumption that is usually made when bringing equation (2) to the data, is that of rational expectations. This means that the conditional expectation operator in (2) coincides with the mathematical expectation of the expression within the brackets and implies that deviations from the actual value of that expression and the left-hand side of (2) are, on average, zero and orthogonal to all information available to the consumer at time t . If we neglect, for the moment, the presence of the 'taste-shock' (or measurement error) v_t , we can in principle use these orthogonality conditions to estimate the parameters of interest, as long as the number of the latter is less or equal than the number of orthogonality conditions. It should be stressed, however, that the average is zero for each consumer *over time*. There is no reason to believe, without additional assumptions, that the cross section mean of these deviations is zero at any point in time. This point was first stressed by Chamberlain (1984) and has been discussed, among others, by Hayashi (1987), Altug and Miller (1990), Deaton (1992), Keane and Runkle (1994) and Attanasio (1998). Its most important implication is the fact that one can use purely cross sectional data to estimate Euler equations only under very special circumstances. In particular, as stressed by Altug and Miller (1990), one needs the assumption of complete markets to guarantee that the expectational errors u_t^h can be decomposed into the sum of an aggregate component η_t and a purely idiosyncratic one ε_t^h . In this case, the aggregate shock is absorbed by the constant of

¹ The consideration of unobserved heterogeneity used here is quite flexible because we are dealing with a single commodity. In the presence of multiple commodities the issues are much more complicated. We are not going to deal with those issues here.

the equation or, if there is more than one time period, by time dummies. In the absence of complete markets, when the expectational errors of different individuals are correlated but not identical, the introduction of time dummies does not solve the problem. This is because the consideration of each instrument would need the introduction of a new set of time dummies to capture the correlation of the instrument with the expectational errors: the model is effectively not identified. Therefore, to estimate an Euler equation in the absence of complete markets one needs a large T : in general, N -asymptotics does not work.

(iii) In the absence of heterogeneity, or under assumptions that guarantee the existence of a representative consumer, and in the absence of measurement error, one can estimate the discount factor and the coefficient of relative risk aversion using aggregate time series data. Hansen and Singleton (1982) were the first to use GMM to estimate a version of equation (2). However, because of the non-linear structure of equation (2), the presence of measurement error in consumption and/or unobserved heterogeneity can have devastating effects on our ability to obtain consistent estimates of the parameters of interest.

(iv) The vector z can include endogenous variables, such as labour supply choices and durable consumption. This would be the case if we modeled preferences *conditional* on the optimal value of some other variables that we do not model explicitly. This is extremely useful in the case of variables for which non-convexities and corners are important, such as durables, labour supply and so on.

(v) There is now a substantial amount of evidence that aggregation issues are important for the estimation and evaluation of Euler equations. In particular, both non-linearities and the role played by certain demographic and labour supply variables can be crucial in explaining the results typically obtained with aggregate time series data. Furthermore, using individual data, one can focus on individuals or groups of individuals, who are more likely to satisfy the assumptions posed to obtain an empirically tractable specification.

An obvious implication of point (v) above is that Euler equations can properly be estimated only using individual data.² On the other hand, point (iii) above stresses the necessity of working with linear (in parameters) models to be able to deal with the unavoidable measurement error present in individual level data.³ This implies that an equation such as (2) is of limited use. A possibly appealing alternative,

² However, it should be stressed that one does not necessarily need panel data. A time series of cross sections, as long as the model does not involve temporal non-separabilities, will be enough. One can form groups with constant membership over time and follow the relevant aggregates for these groups. As the asymptotics is on large T anyway, little is lost by going from the individual to the grouped data. It is crucial, however, that the relevant non-linear transformation of the data is computed.

³ To be more precise: if one believes that measurement error in consumption is multiplicative, one needs to obtain a model where consumption enters in logs, so that measurement error becomes additive.

therefore, is to consider a log-linear version of (2). Following the steps in Hansen and Singleton (1983), for instance, from equation (2) one can derive the following:

$$(3) \quad \ln\left(\frac{C_{t+1}^h}{C_t^h}\right) = \frac{1}{\gamma} (k_t + \theta \Delta z_{t+1}^h + \ln(1 + R_{t+1}) + \Delta v_{t+1}^h) + u_{t+1}^h$$

If the relevant variables are log-normal,

$$k_t = \ln(\beta) + \gamma^2 \text{var}_t\left(\ln\left(\frac{C_{t+1}^h}{C_t^h}\right)\right) + \text{var}_t(\ln(1 + R_{t+1})) - 2\gamma \text{cov}_t\left(\ln\left(\frac{C_{t+1}^h}{C_t^h}\right), \ln(1 + R_{t+1})\right)$$

where the t subscripts on the second moments indicates that these are conditional on the information available at t . When the conditional distribution of the relevant variables is not log-normal, the term k_t will include higher conditional moments as well. As conditional second (and higher) moments are typically unobserved, it is useful to re-write equation (3) as follows:

$$(4) \quad \ln\left(\frac{C_{t+1}^h}{C_t^h}\right) = \frac{1}{\gamma} (\bar{k} + \theta \Delta z_{t+1}^h + \ln(1 + R_{t+1})) + e_{t+1}^h$$

where the term \bar{k} includes the log of the discount factor and the unconditional mean of the second (and higher) moments of consumption growth and real interest rates. The residual e includes expectational errors u , unobserved heterogeneity v and the deviations of k_t from \bar{k} . Notice that, of the parameters in the utility function, only γ and θ can be identified. Equation (4) is typically estimated by Instrumental Variables or GMM, using as instruments variables that are known at time t .

Obviously equations (3) and (4) are only approximations to equation (2). In the next section we discuss the extent to which these approximations can be used to estimate the structural parameters γ and θ . Before that, however, it is worth mentioning two additional points. First, as we stressed above, to estimate Euler equations in the absence of complete markets, the availability of data covering a long time horizon is crucial. Unfortunately, long panel data containing information on consumption are almost non-existent. In the US and in the UK, only time series of cross sections are available.⁴ To estimate any dynamic model, therefore, it is necessary to use synthetic panel techniques, such as those pioneered by Browning, Deaton and Irish (1985). However, to apply such techniques it is necessary to deal with relationships that are linear in parameters, such as equation (3), but not (1). The presence of multiplicative measurement error in consumption, only reinforces the desirability of dealing with an equation such as (4) as opposed to (2). Second, the fact that equation (4) is non-linear in variables (but

⁴ The British FES covers almost 30 years now, while the US CEX is available since 1980 on a continuous basis.

linear in parameters) is a problem for aggregate time series data (as shown by Attanasio and Weber, 1993), but not for synthetic panel techniques, where one controls the aggregation process directly.

3. How useful are log-linear Euler equations?

When studying intertemporal preferences, it is obviously desirable to allow for unobserved heterogeneity in preferences and measurement error in consumption. The discussion of the previous section implies, on the other hand, that it is important to be able to obtain a specification linear in parameters with additive errors. Such a specification should then be estimated on individual data covering a long time period.

One possible strategy, therefore, is to *assume* such a relationship. That is, rather than making assumptions about the utility function, one could start from a specification for the marginal utility of consumption that combines flexibility with linearity in parameters and additive residuals. If necessary, one can derive the implied utility function integrating the assumed expression for the marginal utility. After all, the first order conditions one uses to characterize the inter-temporal allocation of consumption involve its marginal utility. Notice, that precautionary saving would arise from such a modeling strategy if the marginal utility of consumption turns out to be convex.

Alternatively, one can start from a specification for the direct utility function and use a log-linear approximation to the Euler equation derived from it. A relevant question, therefore, is to ask to what extent such an approximation is a good one and, above all, to what extent its use introduces biases in the estimation of the structural parameters of interest. Once again, notice that the use of a log-linear approximation to the Euler equation has no implications for precautionary saving. The latter follows only from the nature of preferences.

It is relatively easy to state a condition under which the use of an IV technique yields consistent estimates of the parameters in equation (4). That is: the instruments one uses should be uncorrelated with the residuals of the equation. As mentioned above, these residuals are made up of three parts: (i) the expectational error, (ii) the terms reflecting unobserved heterogeneity, Δv_{t+1}^h , and (iii) the innovation to the conditional second (and possibly higher) moments of consumption growth and interest rates. The main goal of this paper is to study the effect of the third component.

Before delving into a discussion of the relevant problems, however, it is worth discussing briefly the issue of unobserved heterogeneity. The v term in equation (2) (and subsequent equations) could capture heterogeneity in discount factors and/or multiplicative measurement error in consumption. If one were to use panel data, provided a long enough panel existed, consistency of the estimator would depend on the stochastic properties of such a term. If one was willing to assume that such a term is a constant or a pure random walk, its presence would not jeopardize consistency. In the presence of persistent taste shocks, however, problems would arise if one used individual specific instruments. This would be the case, for instance, if different individuals were characterized by differences in discount factors. When dealing with grouped data, however, the issue is greatly simplified. The reason is, once again, that one relies on large T to obtain consistency.⁵ Therefore one can in principle control for heterogeneity between groups by inserting group specific intercepts. This constitutes yet another strong argument (beside the availability of relatively long time series) for the use of grouped data, such as synthetic cohorts.

We now discuss the possible problems caused by the presence of the innovation to the second (and higher) moments of consumption growth and the interest rate in the residuals of equation (4) arising from the log-linearization of equation (2). Going beyond the statement that to obtain consistency these innovations should be uncorrelated with the instruments used in estimation, however, is very hard, as these conditional moments are endogenously determined. Without a closed form solution for consumption specifying its relation to the interest rate and its other determinants, it is difficult to establish how the conditional variance of consumption growth should evolve. Even if one is willing to assume the absence of heteroscedasticity in the interest rate and in the other determinants of income (such as income innovations), one cannot rule out the possibility that consumption growth exhibits heteroscedasticity of a very persistent nature. If this were the case, the use of lagged instruments to estimate equation (4) would induce important biases in the estimation of the parameter of interest, such as the elasticity of intertemporal substitution.

To investigate this issue we perform a Montecarlo exercise. Given some assumptions on preferences and on the processes generating income and interest rates, we solve numerically for the optimal consumption of a generic individual. With the consumption function so obtained, we simulate consumption and interest rate paths for groups of individuals for a relatively large number of time periods. We then use the generated data to estimate versions of equation (4) both for individual consumers and for different individuals grouped together. We repeat this exercise many times to

⁵ In addition, averaging across individuals removes part of the heterogeneity.

establish the properties of the estimator of the log-linearized Euler equation. The details of the exercise are described in the next section.

4. Montecarlo evidence

In this section we specify a life-cycle model in which a generic consumer maximizes expected utility in the face of wage and interest rate uncertainty. We use numerical solutions of this problem to simulate the behaviour of a large number of individuals. The simulated data are then used to estimate the log-linearized Euler equation. The aim of the exercise is to investigate whether the estimates of the structural parameters obtained with typical time-series regressions are systematically biased. In particular, we report the results obtained under a variety of scenarios, in terms of preferences and stochastic environments faced by the individual consumers and information available to the econometrician. Given the considerations discussed above, we will not investigate the results one would obtain estimating the Euler equation from a single (or a few) cross-sections.

4.1. The life-cycle problem and its numerical solution

We begin by outlining the optimization problem of a generic consumer and describing the method used to obtain consumption functions. As mention above, we assume expected utility maximization. In addition, we assume that utility is a function of a homogeneous consumption good, it is intertemporally separable and is homothetic. In empirical estimation researchers have typically used non-durable consumption, and so we are implicitly assuming separability between this component of consumption and durables, leisure and so on. We are also neglecting the role that demographic variables play in the intertemporal allocation of resources. Therefore, the problem of individual consumer h can be written as:

$$(5) \quad \begin{aligned} \text{Max } E_t \left[\sum_{j=0}^{S-t} \frac{(C_{t+j}^h)^{1-\gamma}}{1-\gamma} \beta^j \right] \\ \text{s.t. } A_{t+j+1}^h = (1 + R_{t+j}^h) A_{t+j}^h + Y_{t+j}^h - C_{t+j}^h \end{aligned}$$

where C denotes, consumption, A assets, Y labour income and R the rate of interest. The length of life S is assumed to be certain. All variables are in real terms. The discount factor β is assumed to be between zero and unity and the coefficient of risk aversion γ to be greater than zero. We assume that the consumer does not have a bequest motive and that she is born with no assets. These assumptions imply that $A_1 = 0$ and $A_{S+1} = 0$. To complete the description of the problem faced by consumer h , we need to specify the processes that generate income and the interest rate. We consider the following income process:

$$(6) \quad Y_t^h = Y_{t-1}^h \varepsilon_t^h;$$

We assume that the innovations to income ε are i.i.d. over time and across individuals. We also assume that these innovations are bounded below from minus infinity and that their variance is constant. While we start with this simple process for income, we also investigate the possibility that the variance of income shocks changes over time. The assumption that the innovations to income are independently distributed across individuals is equivalent to the absence of aggregate shocks to income. Furthermore, we assume that individuals can use only asset A to smooth consumption over time. That is, we rule out the presence of complete insurance markets that could be used to smooth out idiosyncratic shocks. We assume that individuals can borrow at most an amount that they will be able to repay with certainty. In this respect the assumption that income is bounded away from zero is crucial. If income could be zero with positive probability at each point in time, consumers never borrow. The maximum amount they borrow in our model is the present discounted value of the minimum income realization over the remainder of their life-time. This constraint will never bind exactly because the marginal utility of zero consumption is infinite and so consumers ensure strictly positive consumption in all remaining periods. However, when the discount factor is high relative to the current interest rate or when future income is expected to grow quickly, this constraint will be close to binding. In other words, the individual will consume very close to the maximum possible and this means that changes in consumption due to intertemporal substitution and variation in intertemporal prices will be dominated by changes in current income. The interest rate will still affect consumption because it affects the maximum that the individual can borrow. This variation, however, is not useful in identifying the elasticity of intertemporal substitution.

The interest rate is generated by an AR(1) process:

$$(7) \quad r_t^h = \rho r_{t-1}^h + u_t^h$$

where $r = \log(1+R)$ and the innovations to the interest rate are assumed to be uncorrelated over time. We consider both the case in which innovations to the interest rate are idiosyncratic and the case in which innovations are common across subgroups. For simplicity, we assume that innovations are homoscedastic. Notice that to be able to estimate the elasticity of intertemporal substitution with a log-linear Euler equation, it is necessary to have an interest rate that varies over time and does so in a way that is at least partly predictable. The latter requirement is necessary if one wants to use instrumental variable techniques of the kind typically used in the empirical analysis of Euler equations. This assumption is not unrealistic.

The consumption data we use in our Montecarlo exercise are generated by solving numerically for the consumption function implied by the stochastic processes and underlying preferences just described, and then simulating actual consumption behaviour using this solution and particular realisations of incomes and interest rates.

The numerical solution to this problem is by now standard. Following Deaton (1991) the model is solved recursively from the terminal condition to give consumption in each period as a function of two state variables: the ratio of cash-in-hand to permanent income and the interest rate. A few points need to be noted. First, given the assumption about finite lives, the problem is inherently non-stationary. For this reason we have to derive S consumption functions. Second, an additional element of non-stationarity is introduced by the random walk assumption for income. For this reason, we reparameterize the problem in (5) by dividing by current income. Third, even when income innovations are an i.i.d. process, we have two state variables as the current interest rate gives information about future interest rates. When we introduce heteroscedasticity (and persistence in the variance of income), we have an additional state variable.

The variable cash-in-hand is continuous, but the numerical solution is found explicitly only at a finite set of nodes in the state space. Values of the solution in between these nodes are found by local interpolation. The range of the grid for cash-in-hand is different at different points of the life cycle and is chosen such that no extrapolation is used. Further, the grid is most dense near to the lower bound and 200 points are used in the discretisation. In equation (7) the interest rate is a continuous variable, but in the numerical solution we assume that the interest rate takes one of 12 distinct values. Following Tauchen and Hussey (1991), we form a first-order Markov process, giving transition probabilities from r_{t-1} to r_t , which mimics the underlying continuous AR(1) process. A similar technique is used to describe the evolution of the variance of the shock to income. The underlying process is that

$$(8) \quad \sigma_{w,t}^h = \lambda \sigma_{w,t-1}^h + \eta_t^h$$

where σ_w is the variance of income and η is the shock to the variance. In the numerical solution, σ_w can either be high or low and the probability of observing $\sigma_{w,t}$ high is greater if $\sigma_{w,t-1}$ is high. Finally, in solving the Euler equation for the consumption function, it is necessary to calculate conditional expectations over each stochastic variable. When the stochastic variable is discretised, as with r , the transition probabilities are explicitly given. When the stochastic variable is continuous, as with Y , the probabilities are calculated using Gauss quadrature methods. Since these also involve discretisation, the processes are in effect very similar.

The preference specification and the stochastic environment in which we put our consumers may sound very restrictive. In particular, it has been shown that to fit actual micro data it is necessary to allow preferences to depend on demographic and possibly labour supply variables. Furthermore, it is certainly unrealistic to assume the absence of aggregate shocks to wages as well as deterministic trends (and age effects) in earning profiles. However, the reason for these simplifications is that we want to focus on the performance of standard IV estimators.

The effect of introducing demographic variables would be equivalent to the introduction of a varying discount factor, and would therefore change the incentives consumers have to anticipate or postpone consumption for any interest rate. As we discuss the effect of different discount factors on our results, it is relatively straightforward to guess the results that one would obtain allowing for the effect of demographic variables. Analogously, one could guess the effect of deterministic earning trends by recognising that they will be relevant if they make it more likely that liquidity constraints become binding. In other words, the effect on our estimators of most complications that can be translated into a larger (or smaller) probability of liquidity constraints being binding, can be guessed by considering our results for different discount factors.⁶

Finally, the case of a single consumer, which we consider, could be interpreted as a case in which all shocks are aggregate ones. As we will see, the effect of grouping completely idiosyncratic shocks is that of reducing the variance of wage shocks. The more realistic case of shocks that are partly idiosyncratic and partly aggregate, as far the properties of our estimator are concerned, is therefore likely to be in the middle of the two extreme assumptions we consider.

4.2. Generating artificial data

For each of several parameter specifications, we run 10,000 simulations of individual behaviour to generate consumption paths which differ ex-post despite being generated by the same preferences and stochastic processes. We then perform IV estimation of the log-linearized Euler equation using the simulated consumption and interest rate time series both for a single consumer and for groups of consumers. In the latter case, we assume all income risk is idiosyncratic and that interest rate shocks are common across sub-groups of individuals. When we consider groups of individuals, we construct consumption growth averages across the whole group and use these averages in the estimation.⁷ We

⁶ Of course the interpretation of the estimates one gets when allowing for demographic or other variables is a different issue that we discuss briefly in the conclusions.

⁷ This differs slightly from synthetic panel techniques, proposed by Browning, Deaton and Irish (1985), where only a subset of the group would be used in the average for each t .

return to the choice of instruments below. The time span for the estimation begins at a random date after $t = 10$, and ends before $t = 70$, thus reducing the impact of terminal value conditions.

We report the parameter values used in the baseline specification and in some alternatives in Table 1. Baseline parameters are shown in bold, alternative parameters are given in brackets. In terms of preferences, we fix the elasticity of intertemporal substitution at the plausible value of 0.67 and experiment with different assumptions on the discount factor, the relevant stochastic processes and T , the number of time periods used in estimation. As we will see, in order to obtain efficient estimates, it is crucial that T is big enough. How big, in practice, depends mainly on the variability of income innovations, which we discuss below. We interpret a period as being one quarter.

Changing the true value of the discount factor is particularly important. This is because with higher values of δ , impatient consumers will borrow close to the maximum possible. It is in this case that it has been argued that precautionary saving behaviour becomes most relevant. Notice that in our baseline case, the discount factor is higher than the average interest rate. We experiment with two values in which it is lower than the baseline (0.01 and 0.015), one of which coincides with the mean interest rate. We also experiment with two cases in which the discount factor is very high: if we take into account that our period is calibrated to be one quarter, the values of 0.04 and 0.1 in Table 1 correspond to 16% and 47% discount factors on an annual basis.

Table 1: Parameter Values

Parameter Name	Value
E.I.S. ($1/\gamma$)	0.67
Mean r	0.015
AR(1) coefficient on r (ρ)	0.6
S (length of life)	80
T (number of periods in estimation)	20, 40 and 60
$\delta = 1/\beta - 1$ (discount rate)	0.02 (0.01, 0.015, 0.04, 0.1)
σ_w (standard deviation of income shocks)	0.02 (0.1, 0.05)
σ_r (standard deviation of interest rate shocks)	0.011 (0.1, 0.033, 0.0033)

The interest rate process was estimated from time series data for 3-months T-bill real returns in the UK over the period 1960-1996. The AR(1) model provides a reasonable representation of the time series process associated with the annual data, but it is less good for quarterly data. Using an AR(4) would add an additional 3 state variables and so is less practical. However, we would expect a higher order AR process to yield better instruments for the interest rate and so improve the efficiency of the estimates.

Calibrating the standard deviation of the innovation to the income process is more complicated. Even if one were to ignore the difficulty in distinguishing uncertainty from unobserved heterogeneity in the data, it is not obvious that wages or earnings are the right variable to use in such a calibration, as the income innovation represents all the uncertainty, besides that on interest rates, faced by consumers. The implications of different levels of uncertainty have to be interpreted with great caution. The approach we use is to calibrate the uncertainty in the income process so that the variability of the resulting consumption growth time series is matched to the available evidence. Even this criterion, however, is difficult to implement as very little exists on the *time series* variability of consumption growth. We use as a rough guide the estimates presented in Attanasio (1998) obtained from synthetic panels for the US and the UK.

Before performing our Montecarlo experiment, we use the simulations to check the properties of our

numerical solution method. In particular, we check that $E \left[\left(\frac{C_{t+1}^h}{C_t^h} \right)^{-\gamma} (1 + R_{t+1}) \right] = \frac{1}{\beta}$. In the first three

columns of Table 2, we report the average values for $1/\beta$ implied by the simulated values of the non-linear Euler equation with $T=20, 40$ and 60 . The average is computed over time and across simulations. The standard errors reported in parentheses are computed across simulations. The first row refers to the baseline case, in which the true value is 1.02. In the following rows we change the variance of interest rate and income innovations. We then introduce heteroscedasticity in the innovations to income. Finally, in the last four rows, we consider different assumptions about discount factors. All results in Table 2 use individual data.

Notice that, with the partial exception of the case in which the discount factor is very high (10%), the discounted ratio of marginal utilities estimates the discount factor very well. This is not surprising at all, but constitutes an important check on the validity of the numerical algorithms that generate the time series used in estimation. When the discount factor is very high, individuals consume close to the maximum and so consumption in any given period is very sensitive to shocks in that period, leading to

a high variance in consumption. This is shown in the behaviour of our simple estimator in the last row. Notice that an increase in T induces, as expected, a reduction in the standard error of the estimator. In the case where $\delta = 0.1$, however, our estimate of the discount factor is poor, particularly when $T = 60$: individuals are consuming close to the maximum possible and this makes it hard to identify parameters with much precision.

In the last column of the Table, we consider the standard deviation of consumption growth implied by each preference specification. This is particularly useful to judge how appropriate are the assumptions on the variability of earning innovations, which we discussed above. As mentioned above, we set the variance of earnings in the baseline case so that the variability of consumption growth matches that estimated in Attanasio (1998). Notice also that, as expected, consumption growth is much more variable in the last row where high discounting means individuals consumer close to their current income.

Table 2: Checking the solution method

Parameter Specification	$E\left[\left(\frac{C_{t+1}^h}{C_t^h}\right)^{-\gamma} (1+R_{t+1})\right]$			Standard deviation of $\Delta \log(C_{t+1}^h)$
	T=20	T=40	T=60	
Baseline $\sigma_w = 0.02$; $\sigma_r = 0.011$	1.020 (0.010)	1.020 (0.007)	1.020 (0.006)	0.025 (0.002)
$\sigma_r = 0.0033$	1.020 (0.008)	1.020 (0.006)	1.020 (0.005)	0.022 (0.002)
$\sigma_r = 0.033$	1.020 (0.020)	1.021 (0.014)	1.020 (0.012)	0.041 (0.004)
$\sigma_r = 0.1$	1.020 (0.044)	1.020 (0.030)	1.020 (0.026)	0.079 (0.013)
$\sigma_w = 0.05$	1.020 (0.019)	1.020 (0.013)	1.020 (0.011)	0.051 (0.005)
$\sigma_w = 0.1$	1.020 (0.029)	1.020 (0.020)	1.020 (0.017)	0.082 (0.009)
Heteroscedasticity in income Innovations: ⁸ $P(\sigma_{w_t} = \sigma_{w_{t-1}}) = 0.8$	1.023 (0.028)	1.027 (0.021)		0.086 (0.032)
$\delta = 0.01$	1.010 (0.008)	1.010 (0.006)	1.010 (0.005)	0.020 (0.002)
$\delta = 0.015$	1.0150 (0.009)	1.0150 (0.006)	1.0150 (0.005)	0.023 (0.002)
$\delta = 0.04$	1.035 (0.016)	1.037 (0.010)	1.038 (0.008)	0.045 (0.014)
$\delta = 0.1$	1.084 (0.020)	1.090 (0.012)	1.076 (0.010)	0.11 (0.021)

⁸ For $S > 40$, computational errors made the solution unreliable, and so $S = 40$ in this row.

4.3. Consistency of Estimates derived from a log-linearized Euler equation

In our Montecarlo experiment, we estimate equations similar to (4) above. In particular, we focus on the estimation of the elasticity of intertemporal substitution (EIS) $1/\gamma$. In addition to the version of (4) corresponding to our model, we also focus on two other specifications. First, instead of considering a first order Taylor expansion to perform the log-linearization, we consider a second order expansion. It has been claimed that such a method might solve the possible problems that the omission of the second (and higher) moments of consumption growth and interest rates might create. A second-order expansion has been used, for instance, by Dynan (1993), albeit in a cross section study. Second, we add to the log-linearized Euler equation the rate of growth of income. It has been argued (Carroll, 1997) that 'excess sensitivity' of consumption growth to income, might be reflecting the presence of 'precautionary saving'. According to this line of reasoning, income growth would capture the missing variance term.

To summarize, for each simulation we estimate by Instrumental Variables the following three equations:

$$(9) \quad \ln\left(\frac{C_{t+1}^h}{C_t^h}\right) = \bar{k} + \frac{1}{\gamma}r_{t+1} + e_{t+1}^h$$

$$(10) \quad \ln\left(\frac{C_{t+1}^h}{C_t^h}\right) = \bar{k} + \phi \ln\left(\frac{C_{t+1}^h}{C_t^h}\right)^2 + \frac{1}{\gamma}r_{t+1} + e_{t+1}^h$$

$$(11) \quad \ln\left(\frac{C_{t+1}^h}{C_t^h}\right) = \bar{k} + \lambda \ln\left(\frac{Y_{t+1}^h}{Y_t^h}\right) + \frac{1}{\gamma}r_{t+1} + e_{t+1}^h$$

The estimation uses lagged interest rates, lagged income and consumption growth and (in the case of the second order approximation) lagged consumption growth squared as instruments. These are the instruments typically used in the empirical analysis. It should be stressed, however, that the lack of persistence in innovations in income (and the limited persistence in interest rates) makes the model harder to estimate because of the scarcity of informative instruments.

As mentioned above, we perform two different types of exercise. First, we use data generated for individual consumers (and the process of interest rates) to estimate the log-linearized Euler equation. This exercise corresponds to the hypothetical situation in which the econometrician follows a single individual for a long time period. Second, we generate data for groups of individuals who receive

idiosyncratic and completely uninsurable income shocks, but who receive interest rate shocks common within the group. In other words, we follow all members of a group facing a common interest rate across time. If, instead, we sampled only a random subset of the group at each point in time, we would have a source of MA(1) in the residuals and would need to use variables lagged two periods as instruments. As the focus of this exercise is on the possibility of obtaining *consistent* estimates, we follow the same individuals over time, therefore not considering the additional source of noise that one would have if the sample is renewed each period.

We report results for the case in which interest rates are idiosyncratic (ungrouped data) and when interest rates are common across sub-groups (grouped data). Wage shocks are always idiosyncratic. Grouping averages out heterogeneity due to differences in realized wage rates and this reduces the standard errors of the first and second order regression equations. Grouping makes little difference to the estimating equation that is already augmented with a wage term.

In Table 3, we report results on the estimation of the elasticity of intertemporal substitution as well as the Sargan test of overidentifying restrictions for our baseline parameters for both individual and grouped data. The means and standard errors of the estimates of $1/\gamma$ (whose 'true' value is 0.67) are computed as averages over 10,000 replications. Similarly, the Sargan test statistics are the averages of the Sargan tests for each individual regression, with the percentage of regressions rejecting at 5% given in brackets. Under the null hypothesis, there would be a 5% rejection rate.

The first point to make is that a large T is required to achieve unbiased estimates and low standard errors. As discussed in section 2, consistency depends on T asymptotics and hence we need a large T . Even when $T = 40$, standard errors are large even though grouping removes a large part of unobserved heterogeneity. This need for a large T is reflected in the large number of rejections of the Sargan test when $T = 20$.

The second point is that grouping can be quite effective in reducing the variability of the estimates. This is apparent if we compare the last two rows of the table. With groups of size 100 (which is smaller than what is typically used in the average cohort analysis), the standard error of the estimates is reduced from 0.80 to 0.55 when $T = 40$, and from 0.58 to 0.40 when $T = 60$. The increase in the percentage of rejections of the Sargan test with grouping reflects the greater power of the test when the variance of the errors falls.

Third, the introduction of the second order term does not change the result dramatically and the average of the point estimates are close to the true value of the elasticity of intertemporal substitution

when the sample size is 40 or 60. However, the introduction of the squared consumption term increases the standard errors dramatically. This is likely to be a consequence of the difficulty in identifying good instruments for this term.

Finally, the introduction of expected wage growth does not change the results much. In this model, expected wage growth does not improve the performance of the estimator for low T and only increases the variability of the estimates for the case in which the estimation horizon is longer. This term does not seem to be correlated with the variance of consumption growth.

Table 3: Estimates for Different T

Parameter Specification	Log-linearized Equation (9)		Second-Order Equation (10)		Augmented Equation (11)	
	Estimate	Sargan Test	Estimate	Sargan Test	Estimate	Sargan Test
T = 20, ungrouped	0.35 (1.55)	2.54 (9.4)	0.23 (4.49)	0.84 (4.1)	0.36 (2.78)	1.39 (10.7)
T = 20, group size 100	0.37 (1.01)	4.14 (24.8)	0.28 (1.94)	1.17 (7.8)	0.37 (2.43)	1.34 (10.0)
T = 40, ungrouped	0.71 (0.80)	2.34 (8.0)	0.60 (2.30)	0.74 (2.9)	0.69 (1.35)	0.95 (5.4)
T = 40, group size 100	0.69 (0.55)	2.88 (12.2)	0.53 (1.30)	0.86 (4.4)	0.70 (1.33)	0.92 (5.4)
T = 60, ungrouped	0.72 (0.58)	1.90 (5.0)	0.66 (1.23)	0.65 (2.3)	0.71 (0.92)	0.68 (2.8)
T = 60, group size 100	0.72 (0.40)	2.44 (9.2)	0.58 (1.01)	0.81 (3.9)	0.71 (0.88)	0.79 (3.9)

Note: For estimates, standard errors in parentheses. Baseline specification. True value = 0.67. Sargan test performed for each individual: Mean Sargan Statistic reported, percentage of rejections at 5% in parentheses. Critical values are (5.99, 3.84, 3.84) respectively for the 3 equations.

To conclude the discussion of the baseline case, this first set of simulations shows that the log-linearization of the Euler equation does not introduce systematic biases in the estimation of the elasticity of inter-temporal substitution when T is sufficiently large. Furthermore, the estimates that one gets in some specifications, for example with grouped data when T=60, indicate that such a parameter can be estimated with reasonable precision. On the negative size, however, we notice that when we get more precise estimates (for instance when we use grouped data), the Sargan statistic shows some sign of misspecification of the model. This would occur, for instance, if the omitted

variances are correlated with the instruments we use. It is apparent, however, that this misspecification does not show up in a systematic bias of the estimates of the structural parameter.

In Table 4 and Table 6 we experiment with changes in the discount factor. The two tables refer to results obtained with ungrouped and grouped data respectively. In particular, we consider values of δ ranging from 0.01 to 0.1. Table 5 reports Sargan statistics for estimates using ungrouped data. The results on the tests of overidentifying restrictions for the grouped estimates are not particularly different from those in Table 5 and are available upon request.

Table 4: Estimates varying the Discount Rate, Ungrouped

Parameter Specification	Log-linearized Equation (9)			Second-Order Equation (10)			Augmented Equation (11)		
	T=20	T=40	T=60	T=20	T=40	T=60	T=20	T=40	T=60
$\delta = 0.01$	0.42 (1.23)	0.70 (0.64)	0.71 (0.47)	0.38 (2.21)	0.59 (1.44)	0.62 (0.95)	0.42 (1.86)	0.69 (0.98)	0.70 (0.77)
$\delta = 0.015$	0.39 (1.37)	0.70 (0.72)	0.72 (0.52)	0.33 (2.41)	0.62 (1.63)	0.66 (1.08)	0.39 (2.34)	0.69 (1.12)	0.71 (1.07)
<i>Baseline: $\delta = 0.02$</i>	0.35 (1.55)	0.71 (0.80)	0.72 (0.58)	0.23 (4.49)	0.60 (2.30)	0.66 (1.23)	0.36 (2.78)	0.69 (1.35)	0.71 (0.92)
$\delta = 0.04$	0.34 (1.85)	0.72 (0.93)	0.97 (0.99)	0.23 (2.89)	0.41 (1.76)	0.67 (2.02)	0.32 (2.48)	0.70 (1.37)	1.05 (4.07)
$\delta = 0.1$	0.86 (5.49)	1.45 (2.49)	2.50 (2.47)	0.44 (10.84)	0.92 (6.97)	1.58 (4.51)	1.02 (11.21)	1.40 (6.92)	2.61 (10.27)

Note: Standard errors in parentheses. True value = 0.67.

The evidence that emerges from these tables is that values of the discount rate, δ , have little affect upon the bias in the estimates of the elasticity of substitution, unless δ becomes extremely large (0.1). When δ is large, individuals will consume close to the maximum possible and consumption will be strongly affected by current shocks. Changes to the interest rate impact on behaviour primarily through changing the maximum amount of feasible borrowing, and this leads to very imprecise estimates of the EIS. High values of δ can be interpreted as being analogous to explicit liquidity constraints and the poor results and frequent rejections of the model in this case show that we can obtain reliable estimates only in the absence of such constraints. Notice, however, that a very high discount factor

Table 5: Sargan Tests for Estimation varying the Discount Rate, Ungrouped

Parameter Specification	Log-linearized Equation (9)			Second-Order Equation (10)			Augmented Equation (11)		
	T=20	T=40	T=60	T=20	T=40	T=60	T=20	T=40	T=60
$\delta = 0.01$	2.44 (8.4)	2.24 (7.2)	2.15 (6.5)	0.81 (3.7)	0.73 (2.9)	0.72 (2.9)	1.37 (10.4)	0.90 (5.0)	0.80 (3.9)
$\delta = 0.015$	2.48 (8.8)	2.25 (7.4)	2.15 (6.4)	0.82 (3.8)	0.74 (2.7)	0.72 (2.6)	1.37 (10.3)	0.90 (4.9)	0.80 (3.8)
<i>Baseline: $\delta = 0.02$</i>	2.54 (9.4)	2.34 (8.0)	1.90 (5.0)	0.84 (4.1)	0.74 (2.9)	0.65 (2.3)	1.39 (10.7)	0.95 (5.4)	0.68 (2.8)
$\delta = 0.04$	2.48 (9.2)	2.25 (7.3)	2.94 (13.6)	0.78 (3.2)	0.74 (3.1)	1.00 (5.7)	1.31 (9.7)	0.91 (5.2)	1.00 (6.0)
$\delta = 0.1$	2.94 (13.2)	3.31 (17.0)	5.22 (34.3)	1.13 (7.9)	1.36 (10.0)	1.16 (7.2)	1.22 (9.2)	1.15 (7.9)	1.70 (13.8)

Note: Sargan test performed for each individual: Mean Sargan Statistic reported, percentage of rejections at 5% in parentheses. Critical values are (5.99, 3.84, 3.84) respectively for the 3 equations.

Table 6: Estimates varying the Discount Rate, Grouped Data

Parameter Specification	Log-linearized Equation (9)			Second-Order Equation (10)			Augmented Equation (11)		
	T=20	T=40	T=60	T=20	T=40	T=60	T=20	T=40	T=60
$\delta = 0.01$	0.48 (0.65)	0.68 (0.35)	0.71 (0.31)	0.23 (1.20)	0.30 (0.76)	0.42 (1.08)	0.48 (1.43)	0.68 (1.07)	0.68 (0.68)
$\delta = 0.015$	0.45 (0.75)	0.68 (0.40)	0.69 (0.30)	0.37 (3.69)	0.61 (1.57)	0.64 (0.88)	0.46 (1.63)	0.69 (0.99)	0.69 (0.74)
<i>Baseline: $\delta = 0.02$</i>	0.37 (1.01)	0.69 (0.55)	0.72 (0.40)	0.28 (1.94)	0.53 (1.30)	0.58 (1.01)	0.37 (2.43)	0.70 (1.33)	0.71 (0.88)
$\delta = 0.04$	0.31 (1.17)	0.70 (0.61)	0.96 (0.70)	0.08 (1.03)	0.10 (0.84)	0.73 (2.06)	0.32 (2.64)	0.73 (1.55)	1.00 (3.35)
$\delta = 0.1$	0.71 (3.72)	0.44 (1.78)	2.47 (1.77)	0.35 (9.04)	0.79 (3.91)	2.37 (5.09)	0.69 (8.30)	1.52 (5.68)	2.55 (4.02)

Note: Standard errors in parentheses. True value = 0.67

created problems even in the non-linear case discussed in Table 2. The log-linearization, therefore is not responsible for these biases. It is also worth noting that with $T=20$ and grouped data, rejections of the Sargan test occur over 20% of the time for equation (9).

Next, we study the effect of changing the level of uncertainty faced by the consumers. Table 7 and Table 8 consider changes in the variance of the interest rate, σ_r , for ungrouped and grouped data, and Table 9 changes in the variance of wages, σ_w . In Table 7 and Table 8, we observe that, while the average of the estimates does not change much with the variance of the interest rate, the estimates become more precise as the variance increases. This is because the estimator is able to exploit the larger variability in intertemporal prices to estimate the elasticity of intertemporal substitution more precisely. Notice that the variance of the interest rate is one of the terms in the 'constant' of the equation that is present because of the log-linearization.

Table 7: Estimates varying σ_r , Ungrouped Data

Parameter Specification	Log-linearized Equation (9)			Second-Order Equation (10)			Augmented Equation (11)		
	T=20	T=40	T=60	T=20	T=40	T=60	T=20	T=40	T=60
$\sigma_r = 0.1$	0.44 (0.70)	0.66 (0.38)	0.68 (0.31)	0.32 (1.48)	0.55 (0.99)	0.60 (0.79)	0.44 (1.66)	0.67 (0.90)	0.67 (0.69)
$\sigma_r = 0.033$	0.37 (0.93)	0.71 (0.61)	0.73 (0.49)	0.21 (3.03)	0.63 (1.12)	0.65 (0.94)	0.31 (2.41)	0.72 (1.29)	0.72 (0.84)
<i>Baseline:</i> $\sigma_r = 0.011$	0.35 (1.55)	0.71 (0.80)	0.72 (0.58)	0.23 (4.49)	0.60 (2.30)	0.66 (1.23)	0.36 (2.78)	0.69 (1.35)	0.71 (0.92)
$\sigma_r = 0.0033$	0.26 (3.75)	0.71 (2.26)	0.72 (1.69)	0.35 (6.43)	0.37 (4.45)	0.50 (2.94)	0.29 (2.27)	0.73 (1.30)	0.68 (1.13)

Note: Standard errors in parentheses. True value = 0.67

Table 8: Estimates varying σ_r , Grouped Data

Parameter Specification	Log-linearized Equation (9)			Second-Order Equation (10)			Augmented Equation (11)		
	T=20	T=40	T=60	T=20	T=40	T=60	T=20	T=40	T=60
$\sigma_r = 0.1$	0.44 (0.68)	0.66 (0.41)	0.67 (0.29)	0.31 (1.91)	0.56 (1.09)	0.59 (0.83)	0.45 (1.81)	0.66 (1.10)	0.67 (0.61)
$\sigma_r = 0.033$	0.38 (0.94)	0.69 (0.52)	0.71 (0.37)	0.29 (2.17)	0.62 (1.45)	0.65 (1.46)	0.39 (2.20)	0.69 (1.64)	0.70 (0.80)
<i>Baseline:</i> $\sigma_r = 0.011$	0.37 (1.01)	0.69 (0.55)	0.72 (0.40)	0.28 (1.94)	0.53 (1.30)	0.58 (1.01)	0.37 (2.43)	0.70 (1.33)	0.71 (0.88)
$\sigma_r = 0.0033$	0.36 (1.10)	0.69 (0.58)	0.72 (0.42)	0.19 (1.45)	0.25 (0.94)	0.25 (0.95)	0.38 (2.33)	0.71 (1.33)	0.71 (0.88)

Note: Standard errors in parentheses. Grouped data. True value = 0.67.

Table 9: Estimates varying σ_w , Ungrouped Data

Parameter Specification	Log-linearized Equation (9)			Second-Order Equation (10)			Augmented Equation (11)		
	T=20	T=40	T=60	T=20	T=40	T=60	T=20	T=40	T=60
$\sigma_w = 0.1$	0.45 (4.53)	0.72 (2.22)	0.71 (1.61)	0.39 (10.14)	0.68 (5.69)	0.55 (6.26)	0.43 (1.65)	0.68 (0.79)	0.70 (0.62)
$\sigma_w = 0.05$	0.33 (2.69)	0.71 (1.45)	0.72 (1.05)	0.26 (5.62)	0.68 (4.17)	0.63 (2.88)	0.35 (2.27)	0.69 (1.10)	0.71 (0.79)
<i>Baseline:</i> $\sigma_w = 0.02$	0.35 (1.55)	0.71 (0.80)	0.72 (0.58)	0.23 (4.49)	0.60 (2.30)	0.66 (1.23)	0.36 (2.78)	0.69 (1.35)	0.71 (0.92)

Note: Standard errors in parentheses. True value = 0.67

In Table 9, we consider changes in the variability of wages.⁹ Increases in the variance of wages now reduce the precision of our estimates. This is to be expected, given the results we obtained on grouping. It is interesting to notice, however, that changes in the variance of wages, which increase the motive for precautionary saving, have practically no effect on the average of the estimates. This shows,

once again, that precautionary behaviour has no direct effect on the estimates of the log-linearized Euler equation and their bias. This is because any bias depends only on the correlation between the *innovations* to the variance of consumption growth and the instruments used in the estimation.

Finally, we turn to the consideration of heteroscedasticity in the wage process. Given the considerations above, this is the most likely situation in which the use of a log-linearized Euler equation could introduce biases in the estimation of the structural parameters. The particular form of heteroscedasticity considered makes the variance follow a Markov process with two values of the variance. Thus, the variance in t is the same as in $t-1$ with probability ρ , and this is symmetric between the high and low values. The value $\rho = 0.8$, high $\sigma_\varepsilon = 0.1$ and low $\sigma_\varepsilon = 0.005$. We report the results of this experiment in Table 10. As is apparent from the table, this particular form of heteroscedasticity in the wage process does not affect the average estimates of the elasticity of intertemporal substitution. However, the high proportion of rejections of the Sargan test indicates that there might be some misspecification in our estimation.

Table 10: Estimates when Wage Variance depends on Past Variance

Parameter Specification	Log-linearized Equation (9)		Second-Order Equation (10)		Augmented Equation (11)	
	Estimate	Sargan Test	Estimate	Sargan Test	Estimate	Sargan Test
Heteroscedasticity, ungrouped	0.77 (2.90)	3.03 (15.3)	0.79 (4.93)	0.71 (3.3)	0.73 (7.16)	1.38 (11.6)
Heteroscedasticity, group 100	0.76 (0.98)	4.21 (28.3)	0.93 (1.30)	0.79 (3.3)	0.72 (2.80)	1.28 (9.9)
Baseline, ungrouped	0.71 (0.80)	2.34 (8.0)	0.60 (2.30)	0.74 (2.9)	0.69 (1.35)	0.95 (5.4)
Baseline, group 100	0.69 (0.55)	2.88 (12.2)	0.53 (1.30)	0.86 (4.4)	0.70 (1.33)	0.92 (5.4)

Note: For estimates, standard errors in parentheses. Baseline specification. $T=40$. True value = 0.67. Sargan test performed for each individual: Mean Sargan Statistic reported, percentage of rejections at 5% in parentheses. Critical values are (5.99, 3.84, 3.84) respectively for the 3 equations. The value $\rho = 0.8$, high $\sigma_w = 0.1$ and low $\sigma_w = 0.005$.

⁹ We only report ungrouped data: grouping across individuals who face common interest rate shocks has the same effect as reducing the variance of the shock to wages.

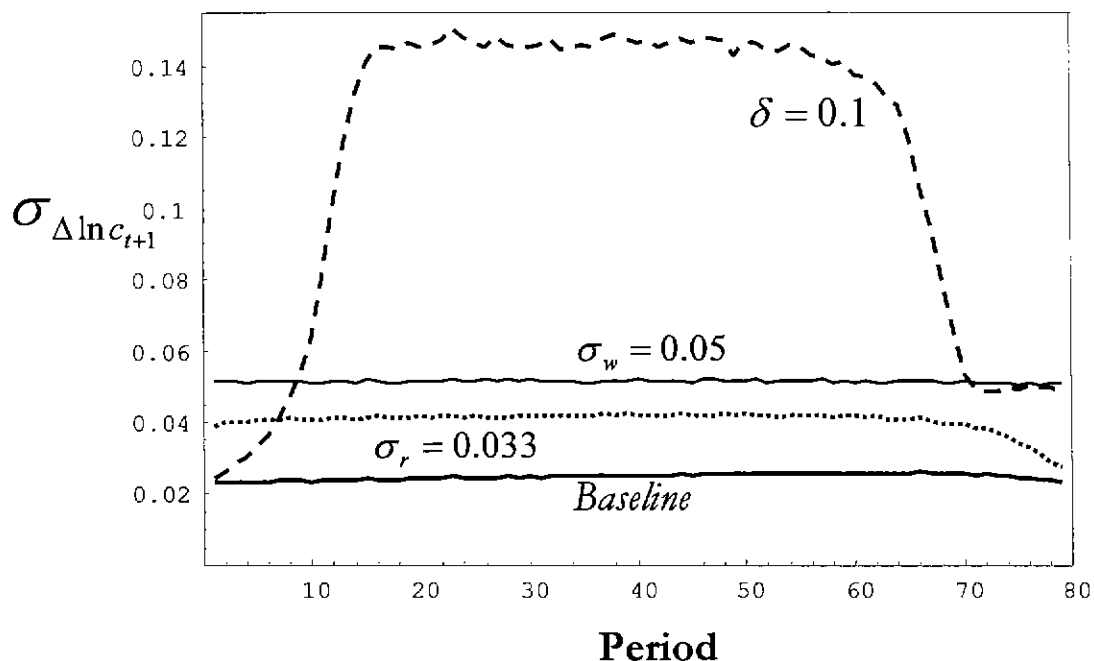
The key to why we are able to obtain consistent estimates of the EIS in certain circumstances is the lack of correlation between the innovations to the variance of consumption growth and the instruments. Cases where the value of δ is very high are particularly problematic, as shown by the high number of rejections of the Sargan test in Table 5.

As we argued above, our main aim was to establish under what conditions innovations to the variance of consumption growth are uncorrelated with the instruments typically used in estimation. Our results show that under a wide set of circumstances this lack of correlation holds true. It is interesting, however, to consider the shape of the life cycle profile for the variance of consumption growth implied by our model. This is particularly interesting for the cases in which we encountered problems, such as that of a very high discount factor.

In Figure 1, we plot the standard deviation of consumption growth obtained computing the standard deviation of actual consumption growth across simulations for individuals who were ex-ante identical but who experience different realization of the wage and interest rate process. The figure contains the plots for different values of the discount factors. The profile for the baseline case is very mildly hump shaped, and this arises because in mid-life-cycle, differential asset accumulation leads to greater variance in consumption growth across individuals who are ex-ante identical. The hump is slightly steeper with high variance in the interest rate because the greater uncertainty induces asset accumulation for precautionary reasons and so there is greater variability in mid-life. However, the most striking feature of Figure 1 is the very strong hump shape of the variance of consumption growth when the discount factor is very high ($\delta=0.1$). In mid-life, individuals consume close to the maximum and are strongly affected by current income. This induces the high variance in consumption growth. This also explains why estimates with larger T appears to give worse estimates for high values of δ : when T is large, individuals spend longer away from the terminal points and so they have more periods with high variability in consumption growth, worsening estimates.

The fact that the variance of consumption growth is relatively flat over the life cycle for most sensible parameter configurations is also informative about the possible interpretation of the role of demographics in log-linearized Euler equation. In particular, the claim made by Gourinchas and Parker (1999) that these variables capture the effect of precautionary saving behaviour seems unwarranted.

Figure 1 Standard Deviation of Consumption Growth across the Life-cycle



5. Conclusions

It is well known that to obtain consistent estimates of structural preference parameters from an Euler equation, in the absence of complete markets, it is necessary to use 'large-T' asymptotics. In the presence of measurement error and unobserved heterogeneity it is advisable to work with specifications that are linear in the parameters and with additive residuals. For this reason, the log-linearization of the non-linear Euler equations typically derived from homothetic preferences seems particularly attractive. The issue is then whether variation in the conditional second (or higher) moments that are typically included in the constant of the log-linearized Euler equation induce any bias in the estimation of the elasticity of intertemporal substitution. We have presented Montecarlo evidence that shows that under a variety of assumptions, the approximated Euler equation yields consistent estimates of the EIS. We have used the isoelastic specification and the instruments often used in the applied literature. Our results show that, at least for the cases that we considered, the endogenously determined conditional variance of consumption growth (and of the interest rates) is uncorrelated with the instruments we use in estimation. Interestingly, our results hold even in the case in which we consider heteroscedastic innovations to the wage processes. We have also shown that the use of grouped data can improve the efficiency of the estimates.

The only case in which the Euler equation yields imprecise estimates of the parameter of interest is when the discount factor is very large. This, however, is not a consequence of the linearization: when consumers are very impatient, they will consume close to the maximum possible and consumption will be strongly affected by current income, and this leads to poor estimates of the EIS using the interest rate. Estimation of the Euler equation, log-linearized or otherwise, would be imprecise.

The evidence presented here might sound surprising in the light of the recent literature on precautionary saving. It should be stressed, however, that the use of a log-linearized Euler equation for estimation purposes does not imply the absence of a precautionary motive, if one means by that a statement about the curvature of the marginal utility of consumption. The variance of consumption growth (along with the variance of the interest rate, their covariance and, possibly, higher moments) is incorporated in the intercept of our equation. As one uses the variation over time to estimate the parameters, this constant will take care of the average growth of consumption (in excess of the difference between the discount factor and the average interest rate). While this implies that the log-linearized Euler equation cannot be used to identify the discount factor, no bias is introduced in the estimation of the EIS or in the tests of the over-identifying restrictions as long as the difference between the conditional and unconditional second moments is uncorrelated with the instruments used in estimation. However, if by precautionary saving one means the presence of very impatient consumers for whom a liquidity constraint is often binding, then, for these consumers, the Euler equation cannot be used to estimate preference parameters. Once again, however, this failure has nothing to do with the log-linearization. A non-linearized Euler equation would perform as badly. The evidence we have presented is, at face value, a bit disappointing in terms of the precision of the estimates one gets even when using sample periods as long as 60 quarters. It should be remembered, however, that our model is very simple and that we have used a very limited sets of instruments. When interest rates and consumption growth are better fitted by a wider set of instruments, the precision of our estimates could improve considerably, even if one would need to fit additional parameters (for instance to capture the role of demographics). Moreover, it might be worthwhile investigating possible efficiency gains that one would get from fitting the equation for several consumers (or groups of consumer) simultaneously. This would be interesting as it is the standard practice in studies that use time series of repeated cross-sections.

The preferences and the stochastic environment we have considered are very simple. This does not mean that we believe that this representation is a realistic one. Indeed, there is much evidence indicating that, for the model to fit the data, it is necessary to incorporate the effect of demographics

and possibly labour supply on the marginal utility of consumption. The main point of this paper was more easily made within a simple framework. The effect of changes in demographics, likely to reflect changes in needs, can be guessed if one interprets the effect of the variables z in equation (2) as shifts in the discount factor. The effects on the estimators are at least two. On the one hand there is a greater amount of variability across individuals that might improve the precision of the estimates of the curvature parameter. On the other hand, changes in the size of the discount factor might accentuate or alleviate the liquidity constraint problem discussed above. That is, if needs are correlated with the amount of resources available in a given time period, it is less likely that liquidity constraints will be binding (and vice-versa). Attanasio, Banks, Meghir and Weber (1999) present evidence that this is indeed the case. In other words, once one corrects for needs, the life cycle profile of consumption appears (at least until retirement) quite flat and therefore consistent with the theory.

Our evidence, and in particular, the plots of the standard deviation of consumption growth in our simple model have implications for the interpretation of the parameters on demographic variables in a log-linearized Euler equation. Gourinchas and Parker (1999) have recently argued that the demographic variables that Attanasio et al. (1999) consider in their model capture movements in the conditional variance of earnings over the life cycle. The fact that in our model the life cycle profile of the standard deviation of consumption growth (the omitted variable in the log-linearized Euler equation) is remarkably flat, casts doubts on this interpretation. Obviously it is conceptually possible to construct a process for earnings whose conditional variance moves over the life cycle in a way that is consistent with the observed hump-shape in consumption. However, the fact that in a simple model the pattern of the variance is far from being hump-shaped, together with the fact that per-capita or per-adult equivalent consumption does not exhibit much of a hump,¹⁰ makes such an interpretation quite contrived. Once again, this does not mean that the precautionary motive is not important.

Generalizing our results to consider a more realistic or more general stochastic environment is likely to be hard. First of all, it is difficult to determine what are all the aspects of uncertainty that are relevant for a consumer. Moreover, it is almost impossible without direct information on individuals' expectations, to distinguish uncertainty from unobserved heterogeneity. Finally, even if one had complete information on all the stochastic processes of relevance, the implied problem would be

¹⁰ Attanasio et al. (1999) show that preferences estimated from a log-linearized Euler equation that allows for the effect of family size on marginal utility are able to reproduce the differences in the shape of life cycle profiles across different education groups. This point is important because there is nothing in the estimation procedure, which is conducted aggregating across education groups, that fits this particular aspect of the data. However, Attanasio et al. (1999) point out that to match the differences in the ages at which consumption peaks, one needs to assume reasonable differences

extremely hard to solve numerically. After all, the main advantage of Euler equation estimation is precisely that, by considering the changes in the marginal utility of wealth, it avoids the need of fully specifying the stochastic environment faced by individual consumers. We believe, however, that our results are indicative of the sort of conditions that would have to hold for the linearization process to be problematic.

in discount factors (which are not identified by the Euler equation). Browning and Eijmaes (2000) go further and claim that a proper modeling of children explains *all* the shape in life cycle consumption profiles.

6. References

- Altug S. and Miller R. (1990): "Household Choices in Equilibrium", *Econometrica*, 58, 543—70.
- Attanasio, O. (1998): "Consumption", in Taylor, J.B. and M Woodford (eds) *Handbook of Macroeconomics*, North Holland.
- Attanasio, O. and G. Weber, (1993): "Consumption Growth, the Interest Rate and Aggregation," *Review of Economics Studies*, 60, 631-49.
- Attanasio, O., Banks, J., Meghir, C. and Weber, G. (1999): "Hump and Bumps in Lifetime Consumption", *Journal of Business Economics and Statistics*.
- Browning, M., Deaton, A. and M. Irish, (1985): "A Profitable Approach to Labor Supply and Commodity Demand over the Life Cycle" , *Econometrica*, 53, 503-44.
- Browning, M. and M. Ejrnaes (2000), "Consumption and Children", Mimeo, Univ. of Copenhagen.
- Carroll, C. (1997): "Buffer Stock Saving and the Life Cycle -Permanent Income Hypothesis", *Quarterly Journal of Economics*, 112, 1-55.
- Carroll, C. (1998): "Death to the log-linearized Euler equation!", Mimeo, Johns Hopkins University.
- Chamberlain, G. (1984): "Panel Data", in Z. Griliches and M.D. Intriligator (eds.), *Handbook of Econometrics*, Vol. II, North Holland, Amsterdam, 1247-318.
- Deaton, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59, 1221-48.
- Deaton, A. (1992): *Understanding Consumption*, Oxford University Press, Oxford.
- Dynan, K.E. (1993): "How Prudent Are Consumers?", *Journal of Political Economy*, 101, 1104-1113.
- Gourinchas, J.O. and J. Parker (1999): "Life Time Consumption", Mimeo, Princeton University
- Hall, R.E., (1978): "The Stochastic Implications of the Life Cycle- Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86, 971-87.
- Hansen, L. P. and K.J. Singleton (1982): "Generalized instrumental variables estimation of nonlinear rational expectation models," *Econometrica*, 50, 1269-86.
- Hansen, L.P. and K.J. Singleton (1983): "Stochastic Consumption, Risk Aversion and the Temporal Behavior of Asset Returns", *Journal of Political Economy*, 91, 249-265.
- Hayashi, F.(1987): "Tests for Liquidity Constraints: A Critical Survey and Some New Observations," in T.F. Bewley (ed.), *Advances in Econometrics: Fifth World Congress*, Vol II, 91-120.
- Keane, M. and D. Runkle (1992): "On the Estimation of Panel-Data Models with Serial Correlation when Instruments are not Strictly Exogenous," *Journal of Business and Economic Statistics*, 10, 1-9.
- Ludvigson, S. and Paxson C. (1999) : "Approximation Bias in Linearized Euler Equations" NBER Working Paper No. T0236
- Tauchen, G. and Hussey, R. (1991) "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models", *Econometrica*, 59, 371-396