Household Wealth and Consumption Variability

Guido Lorenzoni*

April 27, 2001

Abstract

This paper presents evidence of a decreasing relationship between household wealth and consumption variability using PSID data. This relationship offers evidence in favor of models of precautionary savings vs. models of full insurance and vs. standard models of the permanent income hypothesis. Moreover, the presence of this relationship invalidates the identification conditions behind the log-linearized consumption Euler equation. We construct a specification test of the log-linearized Euler equation based on this finding and we obtain a rejection of the assumption needed for log-linearization.

1 Introduction

This paper presents evidence of a decreasing relation between household wealth at a given date and the variability of consumption expenditure in the following year. This evidence bears upon two different issues: first, it provides empirical support to models of precautionary savings, second, it provides a specification test for the log-linearized version of the Euler equation.

Models of precautionary wealth accumulation rely either on prudent preferences or on liquidity constraints or both. A typical feature of these of models is that consumers use accumulated financial assets to self insure against income shocks. An empirical prediction of these models is that household with larger levels of accumulated wealth will face a lower level of consumption variability because they are more willing to use accumulated wealth as a buffer stock against income shocks. This happens because consumers with larger levels of accumulated wealth are either less concerned about hitting

*MIT. I would like to thank Ricardo Caballero, Withney Newey and participants to the MIT Econometrics lunch for helpful suggestions. All remaining errors are mine.

the liquidity constraint in the near future or less risk averse as they reached higher levels of consumption. This prediction distinguishes models of precautionary behavior both from quadratic utility models of the permanent income hipothesis and from models of full insurance. The first objective of this paper is to test this prediction by looking at the variability of consumption as a function of household wealth.

The results obtained can be interpreted as a complement to the test of liquidity constraints by Zeldes (1989): he was concentrating on the conditional first moment of the Euler equation residual, which —in presence of liquidity constraints— displays a discontiuous jump at wealth levels at which the constraint is binding, here instead we study the conditional second moment of the Euler equation residual which —in presence of liquidity constraints and/or prudent preferences— is smoothly decreasing with wealth over an extended range. An advantage is that we do not need any assumption about the maximum level of borrowing allowed. Our results can also be interpreted as a further test of the full insurance model in the vein of Townsend (1990) with the twist that we are testing full insurance against a well defined alternative hypothesis, namely against self-insurance through asset accumulation.

In terms of the consumption Euler equation the conditional variability of consumption growth is associated with the conditional second moment of the Euler equation residual. The log-linearized version of the Euler equation is correctly identified only under the implicit assumption that this second moment is uncorrelated with the instruments used. This happens because the log-linearized identification condition is only an approximation of the true identification condition, and the accuracy of the approximation depends on the higher conditional moments of the residual. Current wealth is clearly correlated with past measures of income and consumption, which are usually included in the instruments set. Thus, a relationship between wealth and the second moment of the residuals hints at a possible misspecification of the log-linearized Euler equation. We setup a specification test along these lines and we obtain a rejection of the hypothesis of a second moment uncorrelated with the instruments. Our results constitute a further warning about the use of the log-linearized Euler equation in empirical work and they complement well the discussion by Carroll (1997) which was based on simulation results. While sharing Carrol's concerns on the log-linearized approach, we maintain a more optimistic view on the general possibility of obtaining useful information in a Euler equation framework. Thus, we try some possible fixes to the identification problem. In presence of heteroskedasticity one can either go back to the nonlinear setup or try and improve the approximation in the traditional setup adding terms to the Taylor expansion. This paper tries the second alternative because, under some additional assumptions, it let us keep linearity in the individual effects, and this can be quite useful from the point of view of panel data estimation.

Section 2 contains a theoretical motivation to interpret the evidence presented in terms of models of precautionary savings. We show simulated series for consumption and wealth for a simple model of optimal consumption with incomplete insurance and borrowing constraints and show that the conditional variance of consumption at time t is decreasing with the level of liquid asset holdings at time $t-1$. Under full insurance or under quadratic preferences the same model displays instead a conditional variance that is constant across wealth levels. In section 3 we discuss the relation between the orthogonality condition derived from consumption theory and the log-linearized version of it often used in the empirical literature and discuss the identification problems that arise in the log-linearized model in presence of heteroskedasticity. In section 4 we describe the data used and the estimation strategy for the log-linearized Euler equation. The empirical results based on the PSID data are reported and discussed in section 5.

2 Theoretical motivation

In this section I use a standard model of consumption behavior with uninsurable income risk to show, with some simple simulations, that the presence of prudence and/or of borrowing constraints generates a negative relation between household wealth and the variability of consumption.

Consider the problem of an infinitely lived consumer that receives a random labor income Y_{it}, can accumulate a risk free asset, and faces a borrowing constraint $A_{it} \geq \underline{A}$.

$$\begin{aligned}
\max \qquad & E_t \sum_{s=t}^{\infty} \beta^t u(C_{is}) \\
s.t. \qquad & C_{is} + A_{is} = (1+r)A_{is-1} + Y_{is} \\
& A_{is} \geq \underline{A}
\end{aligned}$$

Assume that Y_{it} has a permanent and a transitory component, X_{it} and η_{it}. The income process is described by the two equations

$$\begin{aligned}
Y_{it} &= X_{it}\eta_{it} \\
X_{it} &= (1+g)\, X_{it-1}\epsilon_{it}
\end{aligned}$$

where η_{it} and ϵ_{it} are i.i.d. log-normal shocks. Define $c_t = C_t/X_t$ and $a_t = A_t/X_t$.

We distinguish three different cases: (1) full insurance, (2) PIH: only a risk free asset is available, utility is quadratic and $\underline{A} = -\infty$ (3) precautionary savings with CARA utility and borrowing constraints: only a risk free asset is available, utility is CARA and $\underline{A} > -\infty$. We are interested in characterizing the conditional variance of consumption in each case, in particular we will look at $Var_{t-1}(c_t)$ and at $Var_{t-1}(\ln C_t)$.

The case of full insurance requires to embed the consumer's problem above in a general equilibrium framework. We assume that there is no aggregate uncertainty so that the realized cross sectional averages of ϵ_{it} and η_{it} are equal to the mean of the corresponding random variables[1]. In this case for any strictly concave utility function full insurance allows any consumer i to achieve a constant level of consumption across states of the world. In this case both $Var_{t-1}(c_t)$ and $Var_{t-1}(\ln C_t)$ are zero. It is worth noticing that full insurance has a stronger implication: even in presence of aggregate uncertainty, the variability of consumption conditional on the aggregate shock would be zero.

In the case of quadratic utility we need to assume $\beta\frac{1+r}{1+g} = 1$ in order to obtain a stationary process for a_t.[2] In this case the optimal consumption policy is

$$C_t = r\left[(1+r)\,A_{t-1} + \sum_{j=0}^{\infty}(1+r)^{-1}\,E_tY_{t+j}\right]$$

and we can show that both $Var_t(c_t)$ and $Var_t(\ln C_t)$ are constant and do not depend either on a_t or on A_t.

In case (3) we need to resort to simulations to characterize the stochastic properties of consumption. The budget constraint can be written as: $c_t + a_t = \frac{1+r}{1+g}a_{t-1}/\epsilon_t + \eta_t$. Defining $z_t = \frac{1+r}{1+g}a_{t-1}/\epsilon_t + \eta_t$ we can show that optimal consumption will be characterized by the equation

$$C_t = h(z_t)X_t$$

where the function h satisfies the functional equation

$$u'(h(z)) \geq (1+r)\beta E u'\left(h\left(\frac{1+r}{1+g}(z-h(z))\exp(-\epsilon)+\exp(\eta)\right)\right)$$

[1]This assumption is consisten with the assumption of a fixed r in cases (2) and (3). If we derive r in a general equilibrium framework r can be constant if we reach a stationary distribution of the risk-free asset and there is no aggregate risk (see e.g. Huggett (1993)).

[2]Again, this seems a usful requirement if we want to ember the consumer problem in an equilibrium framework and obtain a constant r.

with $h(z) = z$ in case of inequality.

We calibrate the process for income to match the process for individual income estimated by MaCurdy (1982) using PSID data, (in this we follow Deaton (1991)). We use the following parameters: $r = 2\%$, $\beta = .95$, $\gamma = 2$, $Var(\ln \epsilon) = .1$ and $Var(\ln \eta) = .07$. Many studies have characterized in detail the optimal asset accumulation policy in problems of this type: an agent accumulates the asset in good times and runs down his stock in bad times, there is a stationary transition for asset holdings and a long run stationary distribution of asset holdings with bounded support. To compute the optimal consumption policy we follow a simple iteration method on the Euler equation.

Figures 1 and 2 illustrate the simulated relation between a_{t-1} and $Var(c_t|a_{t-1})$ and the simulated relation between a_{t-1} and $Var(\ln C_t|a_{t-1})$. There is a strong negative relation between lagged wealth holdings and the variability of consumption reflecting the fact that a low variability of the residual is a sign of better ability to insure against adverse income shocks.

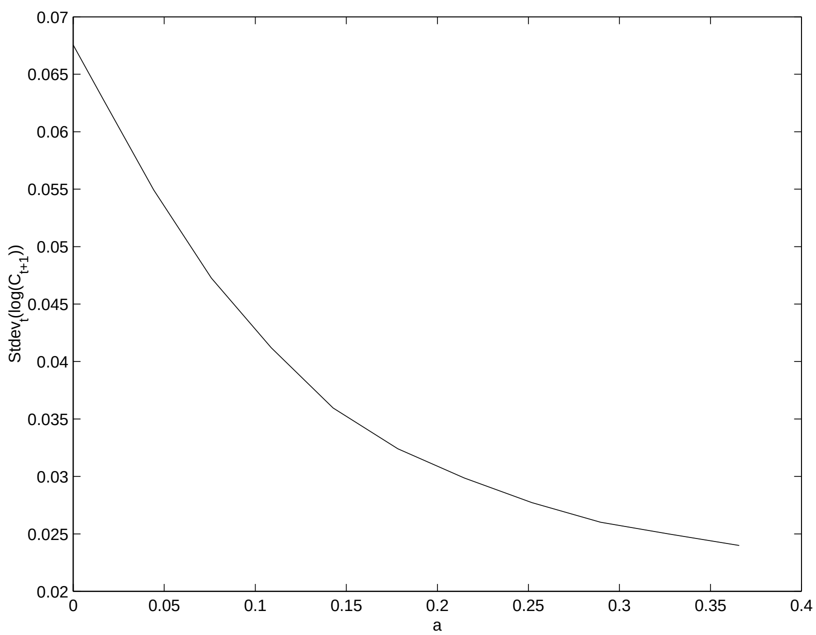

Figure 1: Wealth and variability of consumption (c)

Notice that the relationship above is strictly related to the concavity of the function h (see Carroll and Kimball (1996), and Parker (2001)). Notice

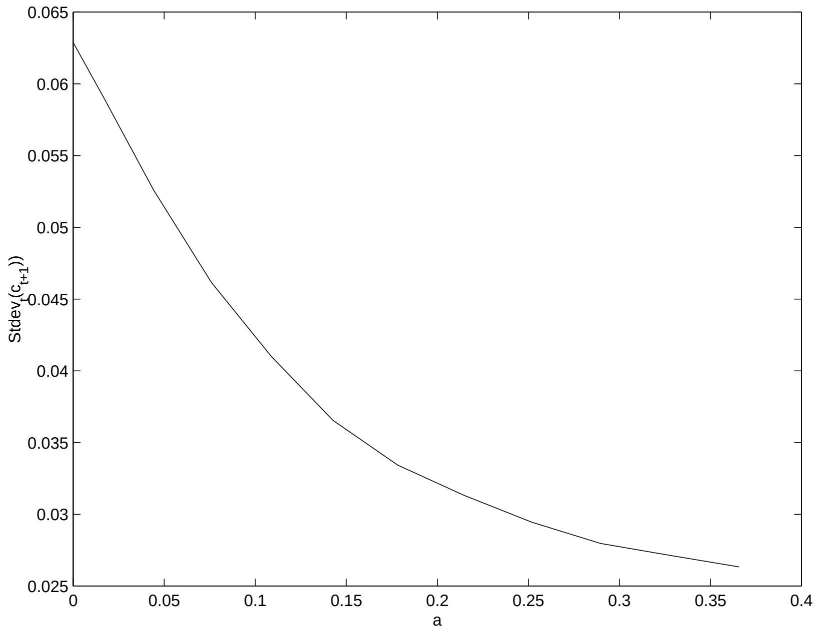

Figure 2: Wealth and variability of consumption ($\ln C$)

that for a given a we have the approximate relation

$$Var(c_{it}|a) \approx (h')^2(a\sigma^2_{1/\epsilon} + \sigma^2_\eta)$$

Two effects are present here, on the one hand there is a simple scale effect, according to which when financial wealth is larger A/X is more volatile, on the other hand the concavity of h makes the term $(h')^2$ larger for low levels of wealth. The figures above show that the second effect seems to dominate for realistic parameter values.

In quantitative terms the effect is sizeable and non-linear in the simulated model. An increase in financial wealth from 0 to 20% of permanent income reduces the standard deviation of consumption from 6.2% to 3.2% percentage points. Further increases in the wealth to permanent income ratio have a much smaller effect on the volatility of consumption. Notice that we have experimented with various values for γ and the results are very similar: the main effect of increases in γ is that the stationary distribution of wealth levels moves to the right but the wealth-variability frontier does not change much.

Summing up, models (1), (2) and (3) differ markedly for their prediction regarding the conditional variability of consumption. In models of full insurance and in quadratic models of the permanent income hypothesis the conditional variability of consumption does not depend on accumualted wealth, in simulated models of precautionary savings models instead household wealth reduces consumption variability especially for low levels of wealth to permanent income ratios. In the empirical section I will use the PSID data on consumption to estimate a nonparametric regression of $Var_{t-1}(\ln C_t)$ on some measures of lagged asset holdings, and I will compare it with the negative relationship depicted in figure 2.

Let me now briefly compare this test with other similar tests designed to test the full insurance hypothesis and the precautionary saving model separately. The test is close in spirit to the test of liquidity constraints performed by Zeldes (1989). In that article a low level of wealth is associated with a higher probability of facing a liquidity constraint. The presence of binding liquidity constraints makes the Euler equation invalid for constrained consumers and therefore Zeldes tested wether the conditional moment $E(u_t|a_{t-1})$ was dependent on a_{t-1}. Here we are essentially studying the behavior of $E(u_t^2|a_{t-1})$ as a function of a_{t-1}. The latter conditional moment is decreasing over the whole wealth range and it does so also in cases in which the consumers are facing a binding constraint very rarely. In these cases a direct test of the liquidity constraint may be inconclusive due to the small number of observations where the constraint is actually binding.

This test is also related to tests of the full insurance model based on the first order properties of the residual as the tests performed by Townsend (1990), Cochrane(1991) and Mace (1991). On the one hand, the present test has more power because it tests the full insurance case against a specific alternative. On the other hand it may be less convincing because it relies on the additional assumptions that taste shocks and measurement error are homoskedastic conditional on lagged values of wealth holdings and income shocks.

We can also compare this test with the a wide class of tests that have been used to study the precautionary motive looking at the relation between individual income variability and wealth holdings, on the grounds that the expectation of higher income veriability induces individual to accumulate more wealth *ex ante* (Browning and Lusardi, Lusardi (1997)). These tests assume some (exogenous) heterogeneity in income varibility across individuals, and include this variability in a reduced form savings function. Apart from the use of an explicit saving function these tests rely basically on exogenous cross-sectional heterogeneity in income variability, while the test proposed here exploits the fact that the ex-ante identical consumers may be differently exposed to residual variability depending on their asset holdings. In a sense their approach relies on individual heterogeneity to obtain identification, while the approach taken here relies on having a sufficiently homogeneous population of consumers, so that one can look at the cross sectional dimension to have information about the time series path of income, wealth and consumption.

3 Linearization and Identification

The original form of the Euler equation is a conditional moment restrictions and can be written in the form

$$E[e^u|z] = 1 \tag{1}$$

where z is any variable that is in the individual information set at time $t-1$ —in particular, it may be a lagged value of wealth holdings or income— and u is the traditional log-linearized residual at time t. The expression for u I will use in the following is

$$u^i_{t+1} = \log \beta^i + \log(1 + r^i_{t+1}) + b\Delta d^i_{t+1} + \Delta(\log U'(c^i_{t+1})) + \eta^i_{t+1}$$

where β is the discount factor, r is the real interest rate, d is a vector of household characteristics (i.e. observable taste shocks), U is the istantaneous utility function, c is consumption and η captures the effect of unobservable taste shocks.

The usual estimation strategy is to use the linearized version of (1), that is $E[u|z] = 0$.

This approximation may be more or less accurate in different cases. Consider first the simple case in wich z and u_t are jointly normally distributed. In this case the condition above implies $E(exp(u + z)) = E(exp(z))$ which can be transformed into an expression involving only the moments of u_t and z, that is

$$Eu + \frac{1}{2}(Var(u) + 2Cov(u, z) + Var(z)) = \frac{1}{2}Var(z)$$

Using this expression when z is a constant we derive the restriction $Eu + 1/2Var(u) = 0$ that in turns gives us $Cov(u, z) = 0$. The latter equality is exactly the identification assumption required to estimate the traditional linearized version of the Euler equation, thus under joint normality the traditional approach is fully jusified. [3] Unfortunately, not only

[3]Under normality absence of linear correlation is equivalent to independence, thus the condition $E[u|z] = 0$ follows as well.

the joint normal case is a special case, but is also particularly ill-suited in this context, because it embeds homoschedasticity by construction.

In the general case we can take a Taylor expansion of $e^u - 1$ around $u = 0$. Rearranging terms we have

$$E(u|z) = -E(1/2u^2 + 1/6u^3 + ...|z) \tag{2}$$

and if the right hand side of this condition is constant in z we can derive the identification condition $E(u - Eu|z) = 0$.

Many researchers have encountered some version of equation (2), that is they have recognized that the mean of u will not be zero and that a variance term should appear in the residual of the regression expression. Usually they have dealt with the problem just redefining appropriately the residual and the constant term in the regression, that is assuming –more or less explicitely– that $E(1/2u^2 + 1/6u^3 + ...|z)$ is constant in z. .

This last assumption is crucial for the identification of the linearized model. Notice that the first term appearing in the expectation is $E(1/2u^2|z)$, and this term reflects the heteroskedasticity of u conditional on the instrument z. Zeldes recognises correctly this problem: "the conditional variance of the forecast error could be a function of wealth or disposable income. For example, when household assets are especially low, uncertainty about the growth rate of consumption could be higher. (...) In each of these cases the estimation scheme presented below will be inconsistent." [4]

Therefore when we do our test of homoschedasticity we are also testing the accuracy of the linearized expression, and if we reject $E(u^2|z) = 0$ this implies that our identification assumption is incorrect. In this setup the presence of heteroskedasticity has negative implications for the consistence of our estimator and not only for its efficiency, due to the form of our original identification assumption and the fact that we are using an approximation of it. After rejecting the homoschedastic hypothesis we cannot rely anymore on the estimates obtained by using $E(u|z) = 0$. At this point we can either go back to the non-linear specification or find some fix for the linear identification condition. From the point of view of matching theoretical predictions with identification assumptions the ideal would be to work directly with the nonlinear form, as Hansen and Singleton originally did on aggregate data, and as Runkles did on individual data. But when using panel data there are substantive gains by keeping linearity in the individual effects, so some type of linearization seems useful.

[4]Zeldes (1989), p.319

One possibility is to add a term of the Taylor expansion and use as the identification condition

$$E(u + 1/2u^2|z) = 0 \tag{3}$$

Clearly to obtain linearity in the individual effects from (3) we need to make additional assumptions. The nature of the problem is the following. We don't actually observe u_t^i, instead we observe $v_t^i = u_t^i - \alpha^i - \eta_t^i$. Here α_i is an individual effect constant over time, in our model it corresponds to a different discount factor, that is to $\log \beta^i$ in (3). η_t^i is a time-varying effect, in our model it correspond to the unobserved taste shock part, but it may also include measurement error in consumption. If we rewrite (3) in terms of observables and parameters (the parameters are implicit in v_t^i), we have:

$$E[v_t^i + 1/2(v_t^i)^2 - \alpha^i - \eta_t^i|z] - K(z) = 0 \tag{4}$$

Where $K(z) = E[(\alpha^i + \eta_t^i)^2 + u_t^i(\alpha^i + \eta_t^i)|z]$.

If we are willing to make the additional assumption that $K(z)$ is constant we have an identification assumption which is linear in the individual effects and we can carry out a richer analysis using most of the tools of linear panel data models. Actually, in this paper I simply analyze a random effect specification (that is I will assume $E(\alpha^i|z) = 0$), but I plan to extend the analysis and do the appropriate specification tests in future work.

4 Data and Estimation Strategy

The estimates are computed using a sample of 2,350 families from the Panel Study of Income Dynamics (PSID) covering the years 1976 to 1985. The advantages and disadvantages of using this data set for studying intertemporal consumption have been discussed at length in the literature[5]. One specific disadvantage in our context is the recognized presence of large measurement error in the consumption data. Apart from the additional assumptions that need to be done in presence of measurement error to get identification, its presence here means also that we have little information about the total variability of the residual. For this reason we are limited in our ability to make statements about the importance of wealth accumulation in reducing variability in proportional terms, even if we have good estimates of its absolute effect.

[5] Studies on consumption using the PSID include: Hall and Mishkin (1982), Zeldes (1989), Keane and Runkle (1992), Gruber and Dynarski (1997).

The variables used are: food consumption, the real rate of interest, household disposable income, family wealth and two demographic variables.

The food consumption variable is constructed adding the value of food stamps to the food at home variable and then adding food at home and food out of home, each deflated by the appropriate consumption price index. The rate of interest is computed by applying the household marginal tax rate to the rate of return on 3-months Treasury bills, and the real rate of interest is obtained by subtracting the rate of inflation for consumption goods. Disposable income is computed subtracting taxes and social security payments from the total income of the family unit, disposable income is deflated using the price index

The first measure of family wealth available is housing wealth, and is computed subtracting the value of the outstanding mortgage from the value of the house. The second measure is actually a flow measure: the income from interests and dividends. An approximate stock measure can be obtained dividing this measure by an appropriate rate of return. In the estimation of the euler equation wealth levels are used only as instruments and in that case I have preferred to use the two wealth variables separately, keeping the second one as a flow variable. For the nonparametric regressions of u^2 instead I have computed an aggregate measure of wealth, converting the second measure to a flow using the rate of return on treasury bills. [6] The demographic variables used are the age of the head of the household and the size of the family unit.

We can use the definition of the observable residual v in section 2, assume CRRA utility ($\frac{c^{1-\gamma}}{1-\gamma}$) and rewrite the expression for the linearized residual in (3) as:

$$v^i_{t+1} = b_0 + \log R^i_{t+1} + b\Delta d^i_{t+1} - \gamma\Delta(\log c^i_{t+1}) \tag{5}$$

To estimate the parameters of this expression I use the GMM estimator based on the condition $E(Z_t' v_t) = 0$ where Z_t is a vector of instrumental variables that are uncorrelated with v_t. The demographic variables used are: the age of the head, the age of the head squared, and the change in family size. To take into account aggregate forecast error I have added time dummies to the right hand side of (5). Z_t includes lagged values of the wealth variables, lagged values of disposable income (2 lags), and the lagged value of the marginal tax rate.

[6] In the construction of the variables above I have followed closely Zeldes (1989), therefore I refer to that article's appendix for a careful discussion of this construction.

In principle one could use a different set of instruments for every time period (more lags are available for larger t), and one could make efficient use of the period-by-period distinct conditions $plim_{N\to\infty}(Z_t'v_t)/N = 0$ for $t = 2, 3, ..., T$, where N is the number of families and T the number of time periods. Here, for simplicity, I have used the same set of instruments for each period and I have used the time aggregate condition $plim_{N\to\infty}\sum_t(Z_t'v_t)/NT = 0$, that is the weighting of the different time period orthogonality conditions is not efficient.[7]

Therefore, my estimation is reduced to a simple efficient instrumental variable regression on the model $\Delta\log c_t^i = bx_t^i + \epsilon_t^i$, where each (i, t) is treated as a separate observation, x_t^i includes all variables on the r.h.s. of (5) except consumption, and $\epsilon_t^i = -(1/\gamma)v_t^i$.

I am not differencing the v_t to eliminate the individual effect, that is I am considering a random effects estimator. In making this choice I rely on results by Keane and Runkle (1992). They adapt the Hausman and Taylor (1981) specification test to the case of not strictly exogenous instruments and do not reject the hypothesis of individual effects uncorrelated with the instruments. Therefore, in the following I use the random effects specification.

Still, the panel dimension was taken account of in computing a constistent estimate of the variance matrix $\Omega = E(Z'\epsilon\epsilon'Z)$ needed for efficient weighting. In doing that I have considered the presence of the individual component in v_t^i and I have allowed for non-zero values of $E(\epsilon_r^i\epsilon_s^i)$ for all r, s. The expression below is the appropriate extension of White's (1980) heteroskedasticity consistent covariance estimator, assuming $E(\epsilon_r^i\epsilon_s^j) = 0$ for any r, s, i, j such that $i \neq j$.[8]

$$\hat{\Omega}/N = (1/N)\sum_{t=1}^{T}\sum_{i=1}^{N}(\hat{\epsilon}_{i,t})^2 z_{i,t}z_{i,t}' + 2(1/N)\sum_{l=1}^{T-1}\sum_{t=l+1}^{T}\sum_{i=1}^{N}\hat{\epsilon}_{i,t}\hat{\epsilon}_{i,t-l}z_{i,t}z_{i,t-l}' \quad (6)$$

The heteroskedasticity test is simply based on the regression of $\epsilon_{i,t}^2$ on the vector of instrumental variables in Z. In that regression I have also tested the joint significance of the four variables immediately related to availability

[7]The efficient use of period-by-period orthogonality conditions separately is computationally more demanding but would allow one to construct a sequence of GMM estimators with asymptotic variances approaching the information bound as the set of instruments is expanded, as shown in Chamberlain (1992).

[8]Compare it with (3.12) in Holtz-Eakin, Newey and Rosen and notice that here period-by-period orthogonality conditions are aggreg5ated in a fixed (inefficient) way.

of liquid assets, namely the two measures of wealth and the two lagged values of disposable income.

A second set of parameter estimates is obtained from the following orthogonality condition, adjusted by a quadratic term,

$$E(Z_t'(\epsilon_t - 1/2\gamma\epsilon_t^2) = 0. \tag{7}$$

which is derived from (4) and, as argued in section 2, is more reliable in case of heteroskedasticity of the residual.

The GMM estimation strategy is analogous to that for the linearized model. Letting $h = Z'(\epsilon + \gamma 1/2\epsilon^2)$, I have first obtained an estimate of the parameters, minimizing $h'Wh$ with $W = I$. Then, using the residuals $\hat{\epsilon}$, I have computed the efficient weighting matrix $W = \Omega^{-1}$ using an expression analogous to (6) with the adjusted residual $\hat{\epsilon} - \hat{\gamma}1/2\hat{\epsilon}^2$ replacing $\hat{\epsilon}$. Using this efficient weigthing matrix I have obtained the parameter estimates.

5 Results

5.1 Some preliminary results

In order to estimate $Var(\ln C_{t+1}|w_{it})$ we need to make assumptions about individual heterogeneity and about aggregate risk. If aggregate risk is absent, as in the model simulated in section 2, then the cross sectional variability of $\ln C_{it}$ provides information about the perceived variability of consumption for a single consumer. Let us assume, for the moment that aggregate risk has negligible effects, and let us focus on the problem of individual heterogeneity. Denote with θ_{it} all sources of individual heterogeneity except the wealth level, in the model in section 2 the only element of heterogeneity, aside from the wealth to permanent income ratio, was the level of permanent income and $\theta_{it} = X_{it}$. Then we can write

$$Var_i(\ln C_{it+1}|w_{it}) = E_i(Var(\ln C_{it+1}|w_{it}, \theta_{it})|w_{it}) + Var_i(E(\ln C_{it+1}|w_{it}, \theta_{it})|w_{it})$$

Where with Var_i we denote a cross sectional variance. Clearly, we are interested in characterizing only the first term on the right hand side. In the model in section 2 $Var(\ln C_{it+1}|w_{it}, \theta_{it})$ is actually constant in θ_{it}, and equal to $Var(\ln C_{it+1}|w_{it})$, therefore if we could isolate this first element on the right hand side we could immediately test the predicted decreasing shape of $Var(\ln C_{it+1}|w_{it})$. Isolating this effect requires a number of identification assumptions. As a preliminary exercise we adopt a rough approach to eliminate the 'level' effect generated by the second term, by just taking differences. If we compute

$$Var_i(\Delta \ln C_{it+1}|w_{it}) = E_i(Var(\Delta \ln C_{it+1}|w_{it}, \theta_{it})|w_{it}) + Var_i(E(\Delta \ln C_{it+1}|w_{it}, \theta_{it})|w_{it})$$

and include C_{it} in the list of the θ_{it} variables, we can exploit the fact that $E(\Delta \ln C_{it+1}|w_{it}, \theta_{it})$ is approximately zero if consumption is close to a random walk. In the context of the simulated model in section 2 this procedure is correct for any w_{it} greater than zero if an approximate Euler equation holds.

Therefore as a preliminary result I have estimated the regression of $(\Delta log(c_{t+1}))^2$ on the wealth level w_t, computed summing the two measures of housing and financial wealth. That is, I have estimated the statistical relation between lagged wealth holdings and variability of consumption. The kernel estimate of this regression appear in figure 3 and 4. [9] The picture seems to support the idea that larger wealth holdings are associated with lower consumption variability. The change in log consumption is an imperfect estimate of the Euler equation residual, for example it does not take into account that changes in demographics may make a given consumption change have different impact on the houshold marginal utility of income. Therefore we turn now to the estimation of the parameters of a fully specified Euler equation in order to recover the residual from this estimation. After that we will go back to the nonparametric analysis of the conditional behavior of the residual.

5.2 Euler equation estimates and tests of heteroskedasticity

The results of the estimation of the linearized model are reported in the first half of table 1. Column II displays the coefficients for an alternative specification that tests for excess sensitivity by adding current income to the group of explanatory variables.

The coefficients of the demographic variables are all significantly different from zero. The coefficient on the interest rate is considerably larger than that usually found in the literature (usually smaller than 1 and often around 0.5 corrsponding to an estimated value of γ around 2). Actually the estimates of this coefficient have pretty large standard errors in all the specifications considered, maybe more efficient estimation procedures (as those discussed in footnote 9) would give results in line with the literature. Observing the

[9] A simple quadratic kernel was used and various bandwidth values where tried. Figure 3 and 4 correspond to a bandwidth value of 0.15.

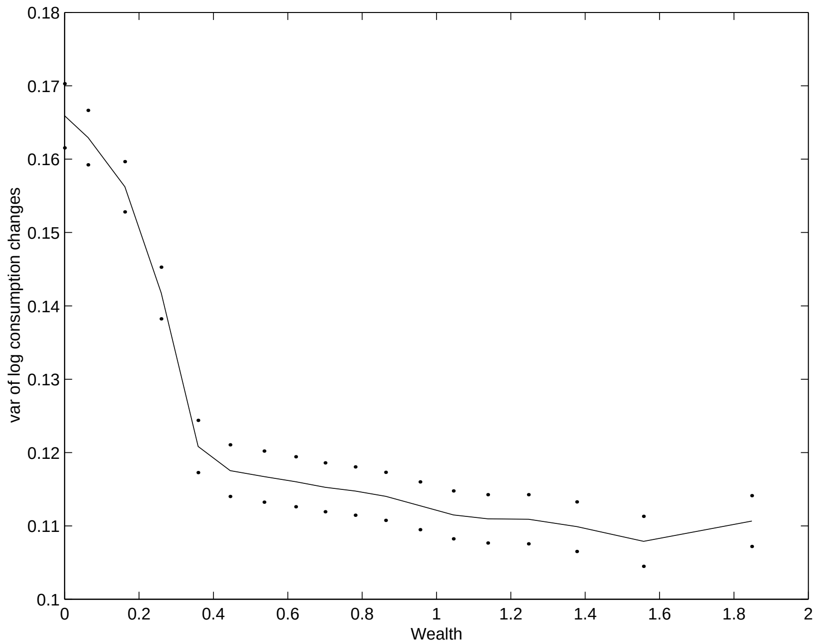

Figure 3: Wealth and consumption variability (w=0 included, bandwidth=.15)

coefficient on disposable income in column II we see no evidence of excess sensitivity in the data, in accordance with the findings of Keane and Runkle and of other studies that allow for a certain richness in the specification of the demographic part (see Attanasio and Browning (1995)). The table reports also the Sargan test of overidentifying restrictions, in both cases the test does not reject the specification used.

The second half of table 1 reports the results of the heteroskedasticity regression, the value of the χ^2 statistic is reported for the general test of homoskedasticity (all coefficients in the regression equal zero), and for the specific test of homoskedasticity conditional on past income and wealth levels (the relative degrees of freedom are reported in parenthesis). Both tests strongly reject the hypothesis of homoskedasticity.

We can take this as *prima facie* evidence in favor of the insurance role of accumulated wealth. Given the correlation of the income and wealth variables and the fact that they are all included to account for the same factor (availability of liquid assets at time $t+1$) we do not expect to get negative and significant coefficients on all variables. The lagged income term carries

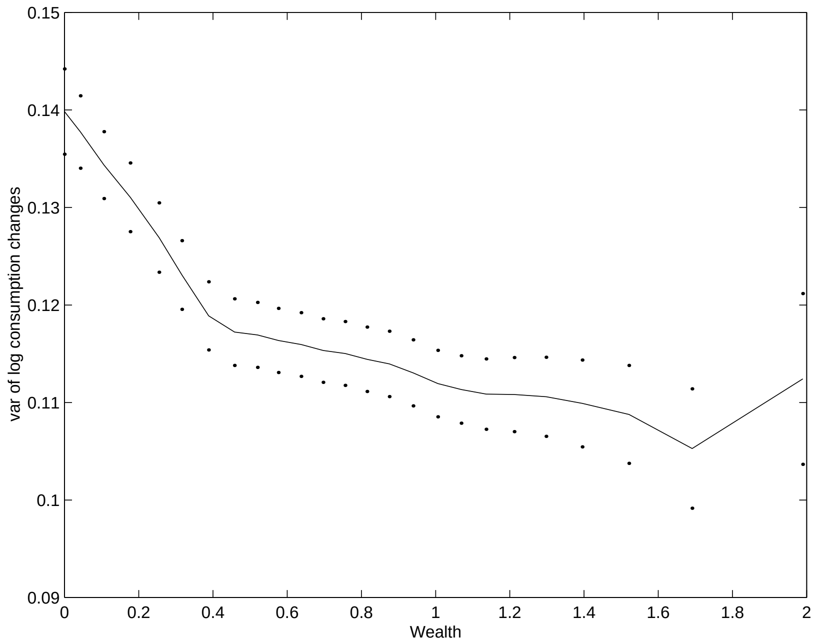

Figure 4: Wealth and consumption variability ($w > 0$, bandwidth=.15)

most of the explanatory power and it is the only one which has a coefficient significantly different from zero. The coefficient on lagged income is negative and —according to our interpretation— it implies that a 10% increase in disposable income in period t allows the consumer to reduce residual consumption variability in the next period by 2.6%. This elasticity is computed using the mean value of the measure of consumption variability, which is 0.1253. Here we encounter the problem mentioned above regarding measurement error. Measurement error affects total observed variability and therefore inflates the value of the average residual variability. At the same time measurement error should not affect our coefficient estimates (at least we hope so, otherwise our identification assumption would be incorrect). As a consequence elasticities of the type just reported will be in general underestimated. Under heteroskedastic errors we know that the identification condition is inaccurate, so I have used the identification condition (7) to get new parameter estimates. The results are reported in Table 2. The heteroskedasticity regression has been performed also for this model, because of its economic interpretation. The first four coefficients do not show striking differences with the simple linearized specification.

Table 1: Log Linearized Model

	I		II	
Indep. Variable	Coeff.	S.E.	Coeff.	S.E.
Age	0.01430	0.00166	0.01458	0.00171
Age^2	-0.00017	0.00002	-0.00017	0.00002
Δ(Family size)	0.11487	0.00537	0.11511	0.00538
Real interst rate	1.65004	0.21259	1.29191	0.54788
Disposable income			-0.00647	0.00917
Overidentif. Restr.	0.93720	(4)	0.42004	(3)
Regression of squared residuals				
Indep. Variable	Coeff.	S.E.	Coeff.	S.E.
Lagged disp. income	-0.03284	0.00482	-0.03301	0.00482
2 Lagged disp. income	-0.00740	0.00392	-0.00747	0.00392
Housing wealth	-0.00341	0.00312	-0.00343	0.00312
Non-housing wealth	0.00186	0.00237	0.00187	0.00237
χ^2	490.610	(16)	492.288	(16)
χ^2 income/wealth vars.	106.610	(4)	107.816	(4)

The coefficient of $\hat{\epsilon}^2$ is an additional parameter which is directly estimated in this case and corresponds to γ (see (7)). Also this estimated γ is much smaller than the estimates obtained in most empirical studies and it is actually not significantly different from zero. In this specification we have also another available estimate for γ, that is obtained as the inverse of the coefficient of the interest rate, which corresponds to 0.7468. We can test the difference of the two by using a delta method to derive the variance of the difference. Even though both estimates have high standard errors we reject the hypothesis that they are equal[10]

The result that the variability term has a coefficient close to zero is puzzling. This result bears some resemblance with a result obtained by Dynan (1993). She uses a different approxiamte expression for the Euler equation that allows for a general form of the utility function, and she obtains a condition analogous to (7) with a variability term represented by $(\Delta logc)^2$ instead of ϵ^2. She shows that the coefficient of the variability term is approximately equal to cu'''/u'', a measure of consumer 'prudence'. Notice that the CRRA specification imposes $cu''/u' = cu'''/u''$, so our test of equality between the two estimated γ can be considered —in Dynan setup— as

[10] The difference is 0.5088, and it is asymptotically normal with standard error 0.2145.

Table 2: Model with quadratic term correction

Indep. Variable	Coeff.	S.E.
Age	0.01578	0.00254
Age^2	-0.00019	0.00003
Δ(Family size)	0.11555	0.00550
Real interst rate	1.33907	0.47128
$1/2\hat{\epsilon}^2$	0.23798	0.15664
Overidentif. Restr.	0.39969	(4)
N. obs.	18880	
N. Indep vars	12	
N. instr.	16	Regression of squared residuals

Regression of squared residuals

Indep. Variable	Coeff.	S.E.
Lagged disp. income	-0.03284	0.00483
2 Lagged disp. income	-0.00759	0.00392
Housing wealth	-0.00359	0.00313
Non-housing wealth	0.00185	0.00237
χ^2	493.83874	(16)
χ^2 income/wealth vars.	107.75742	(4)

a test of the CRRA specification, and our result would imply a rejection of this specification. Secondly the fact that the coefficient on the squared residual is not significantly different from zero would imply a rejection of the prudence assumption. Starting from the approximate Euler equation Dynan follows a different strategy than ours: first aggregates across time periods and then uses the individual observations obtained in this way to test the significance of the prudence coefficient. She reaches a similar result as that obtained here, namely that the coefficient on the variability term is not significantly different from zero. The approach of aggregating along the time dimension and then using the individual cross section of aggregate data can be criticised because the identification of the Euler equation is basically a time series property, and the identification assumptions needed for cross-sectional analyisis may be unwarranted. Carrol(1997) (pp.22-23) contains a detailed critique of Dynan's test along this lines. In the present setup we are exploiting the time-series dimension therefore the result seems more disturbing. At the same time, the time series dimension of the panel is so short that would be ridicolous to claim asymptotic validity in that direction. It is outside the scope of this paper to analyse the delicate question of the use of a time-series identifying assumption with a 'short' panel dataset[11], let me just notice that an implicit assumption that justifies the standard approach is the assumption that the shocks realized in the population reflect the shocks experienced by a typical individual across his lifetime. This implicit assumption is perfeclty valid for example in the simple model simulated in section 2 where only idiosyncratic shocks are present. In presence of aggregate shocks the question is if these aggregate shocks are 'small' enough and if the time dimension available is large enough to average them away. In this sense it would be interesting to redo the estimation above exploiting the full length of the PSID that covers now 25 years.

Incidentally, it is interesting to observe that a recent article by Attanasio and Browning (1995) seems to bring some indirect comfort to the estimation approach based on condition (7). They use coohort data, and in this way they have a much longer time dimension to exploit. Thus, we hope that they are better protected from Chamberlain's critique. In their IV regressions they include the growth of squared consumption among the explanatory variables and the coefficient they obtain is always significant and positive. This is not the same as including the squared residual but the two variables are likely to be highly collinear. Actually, their interpretation of

[11]The point was first raised by Chamberlain(1984). See also Hayashi (1992) and the reply by Keane and Runkle (1992).

that additional term is quite different from the one given here: they stick to the standard linearized Euler equation, but they use a generalized version of the CRRA utility function which, in the end, leads them to write the squared term among the explanatory variables. It would be very interesting to try the estimation strategy outlined above on that dataset, and see which conclusion we obtain regarding consumers' prudence.

5.3 Non-parametric regressions

We now return to the attempt to estimate non-parametrically $Var(\ln C_{it+1}|w_{it}, \theta_{it})$. Now we will use the decomposition

$$Var_i(\Delta \ln C_{it+1}|w_{it}) = E_i(Var(u_{it+1}|w_{it}, \theta_{it})|w_{it}) + Var_i(E(u_{it+1}|w_{it}, \theta_{it})|w_{it}) \tag{8}$$

This decomposition is very convenient, because if the approximate Euler equation holds, then the second term will have a negligible effect. Unfortunately, the results in the previous section warn us that this cannot be the case, actually this cannot be the case *exactly* when the decreasing relation between w and the variability of consumption are correlated. So in a sense our exercise is self defeating, because is leads to a specification test that invalidates our estimation. On the other hand the Dynan result, obtained also in the previous section, could be taken to indicate that, the second order terms are quantitatively small, so that $Var_i(\Delta \ln C_{it+1}|w_{it})$ mostly capture the behavior of $E_i(Var(u_{it+1}|w_{it}, \theta_{it})|w_{it})$. Here we take this approach and we simply look at estimates of the left hand side of 8. The alternative would be to fully spell out a set of identification assumptions that would allows us to estimate the first term in isolation.

Notice that the results reported below are robust to the use of residual from the standard linearized specification and using the residuals from the adjusted specification.

We report the results for the last specification for brevity and because we still hope it provides a better approximation to the behavior of the actual residual. Figures 5 to 7 show the results of a kernel estimate of the nonlinear regression of the squared residual on w. For this estimates we have used a simple quadratic kernel. Given the skewness of the distribution of w we found more convenient to transform the wealth variable according to $w' = log(1 + w)$ and to use the transformed variable as the explanatory variable in computing the nonparametric regressions[12]. Notice though that

[12] I also computed some regression without transforming the wealth variable, the results were not very different but the choice of the bandwidth was difficult: increasing the band-

the figures display the estimated conditional expectation as a function of the original wealth variable. Wealth is expressed in terms of the average income.

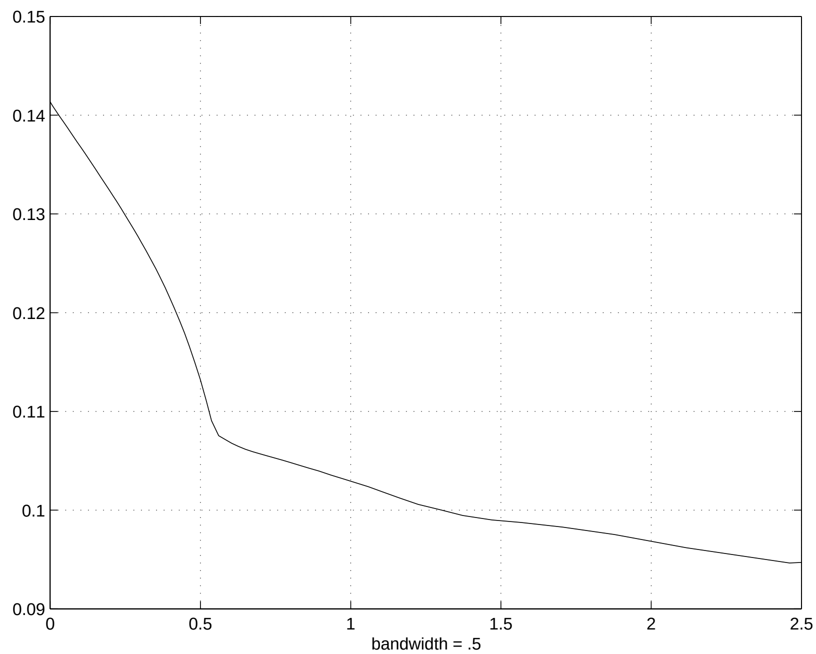

Figure 5: Wealth and residual variability ($w >= 0$)

Figure 5 confirms the preliminary result obtained in figure 3. Even after controlling appropriately for demographic variables and the real interest rate the residual variability is clearly influenced by the wealth level of the household. From the picture it appears that most of the reduction in variability is achieved at the lower levels of wealth holdings. The analysis is made more difficult because the distribution of w is clearly not continuous, since 30% of the data correspond to zero wealth holdings. Moreover at zero wealth holdings the conditional variabilty spikes at 0.1651. This accounts for the steep portion of the regression for low values of wealth, and for the fact that decreasing the value of the bandwidth the steep part shifts left. At this point it is interesting to reestimate the regression admitting a discontinuity at wealth zero. Table 6 displays the results of an estimation using only data with $w > 0$. To get an idea of the jump at zero recall that the mean variabil-

width implied the use of a lot of observations for low levels of wealth and and decreasing it implied the use of too few observations at high levels of wealth.

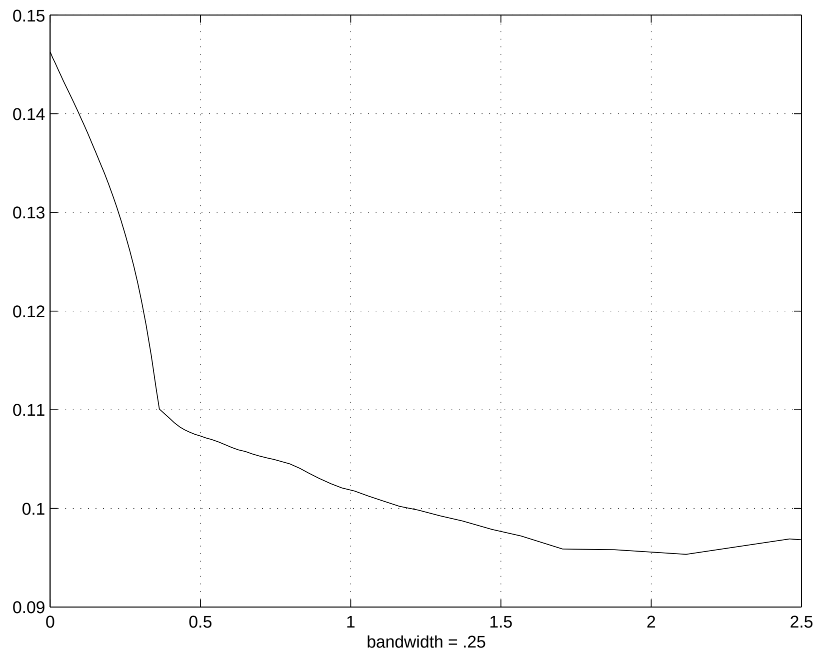

Figure 6: Wealth and residual variability ($w >= 0$)

ity at $w = 0$ is 0.1651. In any event the general result is quite clear. There is a very large reduction in variability passing from zero wealth to some limited amount of wealth holdings, for higher values of wealth holdings the reduction is slower but, interstingly enough, it continues all over the wealth range (the 98% percentile of the distribution is 2.26). Another interesting piece of evidence is contained in the joint distribution of the residual and the wealth variable. The picture of the estimated joint density does not seem very informative, but if we compute the conditional density of the residual at different levels of wealth we can obtain the plot in Figure 8 that confirms the results above, at higher wealth levels corresponds a distribution of the residual that is more concentrated around zero.

6 Concluding remarks

In this paper I have analysed the relationship between household wealth and consumption variability. The relationship has been analysed for its bearing on two different questions: (1) the comparison of a precautionary savings model against models of full insurance or quadratic models of the perma-

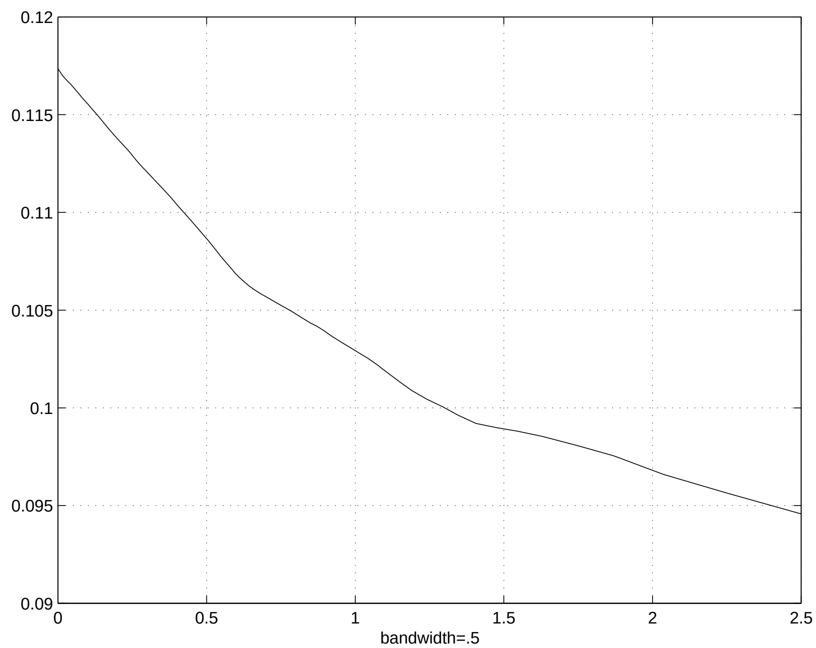

Figure 7: Wealth and residual variability ($w > 0$)

nent income hypothesis, and (2) the econometric use of the approximate consumer Euler equation. Regarding the first question the results seems to support the insurance role of wealth displayng a strong negative relationship between wealth holdings and residual variability. As a clarification is important to notice the objective of this study was not to determine the direction of causation between the two variables, that is we did not try to estimate the direct effect of wealth holdings in reducing consumption variabilty or the effect of greater ex-ante income variability on wealth accumulation. Rather we estimated the equilibrium (reduced form) relationship between the variability and wealth, and we have observed that the shape of this equilibrium relationship seems consistent with the predictions of precautionary savings models.

From the econometric point of view the results show that the use of the traditional linearized Euler equation may lead to inconsisten estimates, being the identification condition an inaccurate approximation of the original condition. At the same time, trying to improve the traditional setup has led to puzzling results similars to those of Dynan (1993). These results may be due to a classical weakness of the Euler equation approach with

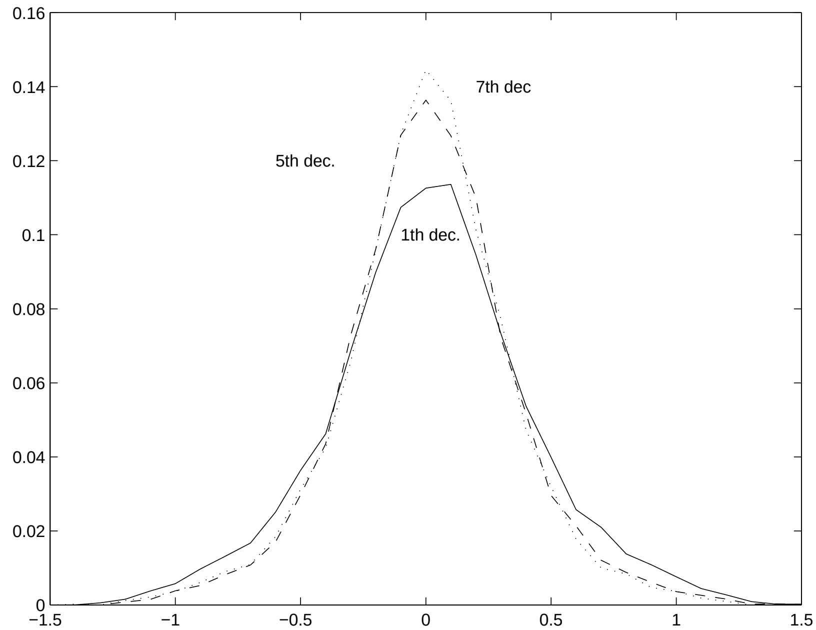

Figure 8: Conditional density of the residual

short panel data, that is, the unwarranted use of a time series identification condition with a small number of time periods. Or they may be due to more fundamental flaws in the specification of consumer preferences.

When the approximate Euler equation holds we can estimate in a simple way the non-parametric relation between wealth and consumption variability, even in presence of individual heterogeneity. When the approximate Euler equation fails, though, this direct estimates are no longer reliable. So in a sense our specification test is self defeating. If our specification test rejects the approximate Euler equation, we can no longer use it as a starting point to study the conditional variance of consumption. Thus, in future work we need to attack directly the problem of identifying the relation between wealth levels and consumption variability without the help of the linearized Euler equation.

References

[1] Attanasio, Orazio and Martin Browning, 1995, "Consumption over the Life Cycle and over the Business Cycle." *American Economic Review*, 85, 1118-1137.

[2] Bewley, Truman, 1977, "The Permanent Income Hypothesis: A Theoretical Formulation." *Journal of Economic Theory*, 16, 252-292.

[3] Carrol, Christopher D., "Death"

[4] Carrol, Christopher D. 1997, "Buffer-stock Saving and the Life Cycle/ Permanent Income Hypothesis." *Quarterly Journal of Economics*, 112, 1-55.

[5] Carrol, Christopher D. and Miles S. Kimball, 1996, "On the Concavity of the Consumption Function." *Econometrica*, 64, 981-92.

[6] Chamberlain, Gary, 1992, "Comment: Sequential Moment Restrictions in Panel Data." *Journal of Business and Economics Statistics*, 10, 20-26.

[7] Cochrane, John H., 1991, "A Simple Test of Consumption Insurance." *Journal of Political Economy*, 99, 957-76.

[8] Deaton, 1991, "Saving and Liquidity Constraints", *Econometrica*, 59, 1221-48.

[9] Dynan, Karen E., 1993, "How Prudent Are Consumers?" *Journal of Political Economy*, 101, 1104-13.

[10] Dynarski, Susan and Jonathan Gruber, 1997, "Can Families Smooth Variable Earnings?" *Brookings Papers on Economic Activity*, 1997, 1, 229-303.

[11] Hansen, Lars Peter and Kenneth J. Singleton, 1982, "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica*, 50, 1269-86.

[12] Holtz-Eakin, Douglas, Withney Newey, and Harvey S. Rosen, 1988, "Estimating Vector Autoregressions with Panel Data." *Econometrica*, 56, 1371-95.

[13] Huggett, Mark , 1993, "The Risk-free Rate in Heterogeneous-agent Incomplete-insurance Economies." *Journal of Economic Dynamics and Control*.

[14] Keane, Micheal P. and David E. Runkle, 1992, "On the estimation of Panel-Data Models With Serial Correlation When Instruments Are Not Strictly Exogenous." *Journal of Business and Economics Statistics*, 10, 1-9.

[15] White, H., 1980, "A Heteroskedasticity-consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica*.

[16] Zeldes, Stephen P., 1989, "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy.*