European Economic Review 44 (2000) 337–358

EUROPEAN ECONOMIC REVIEW

www.elsevier.com/locate/econbase

Using subjective income expectations to test for excess sensitivity of consumption to predicted income growth

Tullio Jappelli[a,c,*], Luigi Pistaferri[b]

[a] *CSEF, University of Salerno, 84084 Fisciaus, Salerno, Italy*
[b] *University College, London, UK*
[c] *CEPR, London, UK*

Received 1 March 1997; accepted 1 May 1998

Abstract

We test for precautionary saving and excess sensitivity of consumption to predicted income growth using a 1989–1993 panel survey of Italian households that has measures of subjective income and inflation expectations. These expectations provide a powerful instrument for predicting income growth. The empirical specification controls for predictable changes in labor supply and allows a fairly general specification for the stochastic structure of the forecast error. We find that consumption growth is positively correlated with the expected variance of income and uncorrelated with predicted income growth. Overall, the results support the precautionary saving model. © 2000 Elsevier Science B.V. All rights reserved.

JEL classification: E21

Keywords: Subjective expectations; Precautionary saving; Excess sensitivity

1. Introduction

An important implication of the permanent income hypothesis is that individual consumption growth should not respond to expected income growth.

* Corresponding author. E-mail: jappelli@synapsis.it.

PII: S0014-2921(98)00069-5

The certainty-equivalence version of the model also suggests that consumers do not react to income risk. But for applied economics the fundamental problem is measuring income risk and predicting future income on the basis of variables that are in individuals' information set and can be observed by the econometrician. In this paper we test the theory of intertemporal consumers choices using data on subjective income expectations to predict realized income growth. The advantage is that no assumption about the process that generates income is required. In the Euler equation we also control explicitly for the potential effect of income risk, predictable changes in households' labor supply, and the nature of the aggregate shocks.

The data are drawn from the 1989–1993 rotating panel of the Bank of Italy Survey of Household Income and Wealth (SHIW). The panel offers unique measures of subjective income and inflation expectations, an annual measure of non-durable consumption that is not affected by seasonality factors, and a wealth of information on financial and real assets. The availability of a good measure of assets is particularly useful to check the possibility of asymmetric response of consumption to predicted income growth.

To date, the panel component of the SHIW has not been extensively exploited for econometric analyses. For the purpose at hand, the main limitation of the panel is that it is relatively short. Even though over long periods of time the forecast error in consumption growth should be zero on average, in our case it may not. In short panels the null hypothesis that the coefficient of predicted income growth in the Euler equation is zero is a joint test of the orthogonality condition implied by the permanent income hypothesis and of the maintained assumptions about the particular stochastic structure of the forecast error. Rejection of the null could be attributed either to a failure of the theory or to the inconsistency of the estimator in short panels. Our test must therefore be designed to tackle this important econometric problem.

In Section 2 we review the literature on excess sensitivity tests, motivate our methodology and describe how it differs from alternative approaches. The construction of subjective income expectation is presented in Section 3. Here we also compare income expectations with income realizations and discuss the validity of expectations as an instrument for predicting realizations. Data and specification issues are discussed in Section 4. Euler equation estimates, reported in Section 5, indicate that consumption growth is positively correlated with the expected variance of income growth, but uncorrelated with predicted income growth. To check for possible asymmetries in the response of consumption to predicted income growth, we also split the sample according to the level of assets (as in Zeldes, 1989) and distinguish between positive and negative expected income growth (as in Shea, 1995). In short, we cannot reject the orthogonality conditions implied by intertemporal optimization, but can reject the certainty equivalence version of the permanent income hypothesis. Section 6 summarizes our main findings and how they can be reconciled with the institutional evidence showing the pervasiveness of borrowing constraints in the Italian economy.

2. Review of the literature and motivation

Several authors have tested the permanent income hypothesis by estimating versions of the following Euler equation with panel data:

$$\begin{aligned}\Delta \ln C_{i,t+1} = {} & \alpha' \Delta F_{i,t+1} + \rho^{-1}(E_{i,t} r_{i,t+1} - \delta) \\ & + \rho/2 \mathrm{var}_{i,t}(\Delta \ln C_{i,t+1} - \rho^{-1} r_{i,t+1}) \\ & + \beta E_{i,t} \Delta \ln Y_{i,t+1} + \varepsilon_{i,t+1}, \end{aligned} \tag{1}$$

where i is a household index, $C_{i,t+1}$ a measure of non-durable consumption, $F_{i,t+1}$ includes predictable indicators of households' preferences (such as age), $r_{i,t+1}$ is the real after-tax rate of interest, ρ^{-1} the intertemporal elasticity of substitution, δ the rate of time preferences, $E_{i,t}$ the expectation operator and $\varepsilon_{i,t+1}$ the forecast error. Eq. (1) can be derived exactly assuming that preferences are of the isoelastic form and that the distribution of the real interest rate and of consumption growth is jointly lognormal. Alternatively, it can be regarded as a second-order approximation to the first-order conditions of the consumer optimization problem.

Predicted income growth is often added to the Euler equation in order to test the orthogonality condition implied by intertemporal optimization, i.e. that $E_{i,t} \Delta \ln Y_{i,t+1}$ should not help in explaining consumption growth ($\beta = 0$). It should be noted that the excess sensitivity test we perform has power against some, but not all, alternative consumption models. For instance, while myopic behavior will lead to excess sensitivity in every period, in a model with prudence and borrowing constraints the orthogonality condition may not be violated most of the time (and even perhaps all the time), as households save in the anticipation of future constraints. Empirically, it is very hard to distinguish between precautionary saving and models with liquidity constraints.

The empirical literature faces several serious problems in testing the restriction $\beta = 0$. First, it is difficult to find viable instruments for income growth that are truly exogenous and yet have good predictive power. Second, the conditional variance of the uncertain components – consumption and the real interest rate – is difficult to observe and is therefore generally omitted from the estimation. Third, excess sensitivity may result from a failure to control properly for non-separability between consumption and leisure. Finally, excess sensitivity may also arise spuriously from the misspecification of the stochastic structure of forecast errors. We address these four problems in turn.

2.1. Predicting income growth

Testing for excess sensitivity requires reliable instruments to predict income growth. However, finding such instruments in panel data has proved to be extremely difficult, particularly in US studies. The Panel Study of Income Dynamics (PSID), which has been extensively used to estimate Euler equations, has relatively good data on income but information on consumption is limited to food expenditures. The Consumer Expenditure Survey (CEX) does give detailed measures of consumption, but the information on income is scanty and suffers from severe measurement error. Three approaches have been proposed to enhance the power of the instruments: out-of-sample information, two-sample instrumental variables techniques, and subjective income expectations.

Shea (1995) isolates a subset of households in the PSID whose heads can be matched to labor union contracts. Information on these contracts is then used to construct a measure of expected nominal wage growth. The latter is found to be strongly correlated with actual wage growth (a coefficient of 0.86 with a t-statistic of 3.8). Inflation expectations, however, are estimated on aggregate data through an autoregressive forecasting model. Shea then estimates an equation similar to Eq. (1) omitting the conditional variance term and replacing the income term with the expected real wage growth of the household head. He finds that expected wage declines affect consumption more strongly than expected wage increases, a result that is not consistent with either myopia or with the hypothesis that excess sensitivity is due to liquidity constraints.[1] There are several problems with this approach. One is that it assumes that the history of past inflation is known to each households in the sample. Another is that Shea ends up with a small sample (647 consumption changes drawn from 285 households), often resulting in poor standard errors, particularly if the sample is split according to the asset-income ratio.[2] Finally, since only food consumption is available in the PSID, he requires an assumption of separability between food and other non-durable expenditures in the household utility function. Yet as Attanasio and Weber (1995) point out, this assumption is rejected in the CEX.

A second possibility is to enhance the power of the test by using two-sample instrumental variable techniques. Lusardi (1996) uses consumption data from the CEX and income data from the PSID, thus overcoming the problem of using just food consumption to estimate the Euler equation. The data are matched by a two-sample instrumental variable estimator. Nonetheless, the adjusted R^2 of

---

[1] Garcia et al. (1997) apply a switching regression model with unknown sample separation to data drawn from the CEX and report a similar finding.

[2] The effect of expected real wage growth is never significantly different from zero in the regressions in Shea (Table 5, p. 195). When Shea splits expected income according to positive and negative expected wage growth he finds an implausible coefficient of 2.242 for expected wage decreases.

the regressions of actual income growth on the instruments (demographic variables, education and occupation dummies) is only about 1% (see Lusardi, 1996, Table 4). Even though Lusardi finds evidence of excess sensitivity to predicted income growth, she does not investigate whether excess sensitivity arises from non-separabilities, myopia, liquidity constraints or other sources.

Hayashi's and Flavin's approach (Hayashi, 1985; Flavin, 1994) is the closest in spirit to the one we use in this paper. Hayashi uses a unique data set of Japanese households reporting subjective expectations for income and consumption on a quarterly basis. He derives the theoretical covariance between the forecast errors in consumption growth and the subjective income expectations, estimates the parameters of the Euler equation by applying a minimum distance estimator and finds some evidence in favor of excess sensitivity. The procedure does not require assumptions about the nature of the aggregate shocks, and is therefore consistent even in short panels.

The 1967–1969 US Survey of Consumer Finances used by Flavin contains a categorical variable about expectations of family income changes. These, in addition to lagged disposable income, are used as an instrument for income growth. Using a robust instrumental variable estimator to control for the presence of influential outliers, Flavin finds evidence of excess sensitivity for both high and low asset households. Evaluating the overall predictive power of Flavin's instrument is not easy, because first-stage results are not fully reported. Data are again problematic in this application. The Survey does not contain a consumption measure, which must therefore be inferred from income and assets. The sample size is small (774 observations), especially when the sample is split by assets.

2.2. The conditional variance of consumption growth

The conditional variance term in Eq. (1) is generally omitted from the estimation.[3] This is correct only under the certainty equivalence version of the model, which implies that households do not react to the expected variance of consumption growth. However, if the utility function exhibits decreasing risk aversion, prudent households react to expected consumption risk by reducing consumption in period t relative to period $t+1$, to an extent that depends on the degree of prudence.[4] The reason the variance term is omitted in actual estimation is not that applied researchers believe in quadratic utility.[5] Rather, it

[3] A notable exception is Dynan (1993).

[4] Kimball (1990) defines absolute prudence as the ratio between the third derivative and the second derivative of the within-period utility function. With isoelastic preferences, relative prudence $-U'''/CU''$ equals 1 plus relative risk aversion.

[5] Research on precautionary saving is in fact steadily growing, see Browning and Lusardi (1996).

is that it has turned out to be extremely difficult to find suitable proxies for the conditional variance.

If the conditional variance term is omitted, one cannot of course test for quadratic preferences or estimate the degree of prudence. But the consequences of this omission could be far more serious. Ludvigson and Paxson (1997) point out that estimating a linearized Euler equation can bias the coefficient of the intertemporal rate of substitution. Furthermore, insofar as the conditional variance of consumption is correlated with $E_{i,t}\Delta \ln Y_{i,t+1}$, the latter will proxy for the omitted effect of consumption risk, generating spurious evidence of excess sensitivity. Carroll (1992) goes one step further, and points out that even Zeldes' (1989) sample splitting approach may produce spurious evidence in favor of liquidity constraints if one does not control properly for expected consumption risk. In fact, Zeldes' test consists in splitting the sample according to the asset–income ratio: if liquidity constraints are at the root of excess sensitivity, one should find no violation of the orthogonality conditions in the high-asset, and excess sensitivity in the low-asset group. But omitting the conditional variance term creates a spurious correlation between consumption growth and income that is stronger for low-wealth households. The reason is that rich households have greater capacity than poor ones to buffer income fluctuations by drawing down their assets, so that a finding of excess sensitivity in the group of poor households only – as in Zeldes – could be rationalized once the assumption of certainty equivalence is dropped by the theory of intertemporal choices.

There are two ways to solve the problem. One would be to estimate the non-linear Euler equation by the generalized method of moments. The second, which is used here, is to introduce explicit proxies for the conditional variance of consumption in the linearized Euler equation. This approach is more directly comparable with previous studies; it also allows us to use standard statistical tools to test if preferences are quadratic or if households react to expected income risk.

2.3. Non-separability between consumption and leisure

If leisure is an argument of the utility function, and if consumption and leisure are non-separable, today's consumption decisions will be affected by predictable changes in households' labor supply. This implies that consumption growth is positively correlated with predictable growth in hours of work. Since predicted growth in hours will almost surely correlate with predicted income growth, failure to control for labor supply indicators may lead to spurious evidence of excess sensitivity (that is, it could bias the estimated β coefficient upwards). This point has been forcefully made by Attanasio and Weber (1995) and Meghir and Weber (1996) with CEX data. But the same authors also indicate a way out to this problem. Following their suggestions, we augment Eq. (1) with labor supply indicators.

2.4. The stochastic structure of the forecast errors

The disturbance term $\varepsilon_{i,t+1}$ in Eq. (1) is a forecast error, the difference between realized and expected consumption growth. According to the permanent income hypothesis with rational expectations, the conditional expectation of a forecast error must be zero, i.e. $E_{i,t}(\varepsilon_{i,t+1}) = 0$. The empirical analog of this expectation is an average taken over long periods of time, not across a large number of households. In fact, as pointed out by Chamberlain (1984), there is no guarantee that the cross-sectional average of forecast errors will converge to zero as the dimension of the cross-section gets large. For instance, if the forecast error is the sum of an aggregate and of an idiosyncratic shock, then in a short panel the orthogonality condition fails even if the permanent income model is true: aggregate shocks induce a cross-sectional correlation between expected consumption growth and predicted income growth. The problem is sometimes handled by including time dummies in the Euler equation. This approach is restrictive, because it rules out that aggregate shocks are not evenly distributed in the population.

For this reason, excess sensitivity tests performed on short panels are in fact joint tests of the null hypothesis that $\beta = 0$ *and* the stochastic structure of the forecast error has a known form (so that the distance between the true forecast error and its empirical analog can be suitably adjusted). Rejection of the null need not be interpreted as the failure of the theory, but could also be attributed to misspecification of the stochastic structure of the forecast error.[6] Distinguishing between the two alternatives is difficult, unless the true structure of the forecast error is known. Yet, as will be seen, subjective expectations provide a guide to modeling the stochastic structure of the forecast error, thereby diminishing the problems one faces when testing for excess sensitivity with short panels.

3. Predicting income growth and consumption risk with subjective expectations

We estimate the Euler equation using the 1989–1993 panel section of the Bank of Italy Survey of Household Income and Wealth (SHIW). An Appendix available upon request details the sample design, response rates, timing of the interviews, wording of the questions, and definitions of the variables. Here we describe only the subjective expectations used to predict income growth and to proxy for consumption risk.[7] Several surveys contain subjective income

[6] Deaton (1992, pp. 147–148) provides an example with non-additive aggregate shocks leading to spurious evidence in favor of excess sensitivity.

[7] Guiso et al. (1992) used the same SHIW questions to study the effect of earnings risk on 1989 saving and households' wealth. They also discuss the pros and cons of using subjective income expectations.

expectations, but vary considerably as to the way expectations are elicited.[8] In the case of the SHIW, in 1989 and 1991 each labor income and pension recipient interviewed was asked to attribute probability weights, summing to 100, to given intervals of inflation and nominal income increases one year ahead.

3.1. Expected income growth

Let $E_{i,t}z_{i,+1}$ denote the expected growth rate of nominal earnings or pension income, $E_{i,t}\pi_{i,t+1}$ the expected rate of inflation and $g^e_{i,t} = E_{i,t}z_{i,+1} - E_{i,t}\pi_{i,t+1}$ the expected growth rate of real earnings. This is the instrument we use for $\Delta \ln Y_{it+1}$, the actual growth rate of earnings of the household head. Although each labor income recipient is asked to answer the survey questions, we rely only on the information provided by the head of the household. The reason is that in most cases information on income recipients other than the head is lacking. As we explain below, subjective expectations are also used to construct a measure of income risk, and use of data on income recipients other than the head would require making difficult assumptions about risk sharing arrangements within the household.

Table 1 compares nominal earnings expectations with realizations by demographic and household-income groups. In comparing expectations with realizations, it must be stressed that respondents report forecasts for the 12 months following the day of the interview. Interviews were taken between May and July of 1990 for the 1989 survey, and between May and October 1992 for the 1991 survey,[9] whereas income realizations refer to the calendar years 1989, 1991 and 1993. Thus we use as instrument the one-year forecast of income growth given in May–July 1990 for the growth rate of earnings between 1989 and 1991 and the one-year forecast of earnings given in May–October 1992 for the growth rate of earnings between 1991 and 1993. This implies that expectations and realizations do not coincide in time, and are not immediately comparable.

In an instrumental variables context, this is not a concern. All that is needed is that the expectation be correlated with actual income growth and uncorrelated over time with the innovation term of the Euler equation (1). Under the null hypothesis of the permanent income model, the latter condition is met. Our

---

[8] The 1982 Japanese Survey of Family Consumption contains information about consumption and income expectations. The Dutch VSB Panel, the 1967 US Survey of Consumer Finances, and the US Survey of Economic Expectations (SEE) contain information on income prospects, but not on expected or actual consumption. Das and Van Soest (1997) and Dominitz and Manski (1997) using the VSB and the SEE, respectively, compare income expectations with realizations.

[9] SHIW interviews usually start in May, with households asked about their income, assets and consumption of the previous calendar year. The reason is that previous experience has shown that people report income more accurately when filing the income tax forms, which must be returned by May 31.

Table 1
Comparing expectations and realizations of nominal income growth

	1990–1991 Expectation (1)	1990–1991 Average realization (2)	1992–1993 Expectation (3)	1992–1993 Average realization (4)
Age group				
< 35	0.0758	0.0588	0.0521	0.0361
35–55	0.0640	0.0475	0.0399	0.0178
> 55	0.0426	0.0543	0.0306	0.0248
Education				
Junior high-school or less	0.0498	0.0559	0.0333	0.0101
High school	0.0667	0.0428	0.0439	0.0319
University degree or more	0.0754	0.0493	0.0446	0.0716
Occupation				
Employed	0.0565	0.0623	0.0371	0.0446
Self-employed	0.0607	0.0043	0.0345	− 0.1207
Region of residence				
North	0.0576	0.0435	0.0342	0.0246
Center	0.0541	0.0869	0.0306	0.0469
South	0.0582	0.0443	0.0429	− 0.0049
Household income				
I quartile	0.0455	0.1048	0.0381	0.0422
II quartile	0.0620	0.0689	0.0351	0.0113
III quartile	0.0548	0.0310	0.0377	0.0171
IV quartile	0.0637	0.0129	0.0360	0.0111
Total sample	0.0573	0.0565	0.0367	0.0201

Notes: The table compares expectations and realizations of nominal income growth. The realization is the average growth rate over the two years. Expectations are given in May–July of 1990 (column (1)) and May–October 1992 (column (3)) for the subsequent 12 months. Income is defined as after-tax earnings and pension benefits of the household head.

approach is valid even if individuals underestimate or overpredict future income: all we need is that expected income growth helps predicting actual income growth. In the next section we show that income expectations are indeed strongly correlated with realizations. Here we limit ourselves to a descriptive analysis.

Only if incomes grew steadily over the two-year span one would expect subjective predictions to mirror half of the actual income growth rate. The last row of Table 1 suggests that while in 1990–1991 expectations are quite close to realizations (5.73 against 5.65%), in 1992–1993 expectations overpredict realizations (3.7 against 2%). Subjective expectations can be criticized because respondents may not fully understand the survey questions: households with better education might therefore give more accurate income forecasts simply because they understand the survey questions better. However, individuals with less education do not appear to answer the survey questions less accurately than those with more. For instance, in 1989 individuals with junior high school or less report an average expectation of 5 percent (vis-à-vis a realization of 5.6%), while individuals with college degrees overpredict income growth (7.5% vis-à-vis 5%). In 1991 it is the group with higher education that makes better forecasts. One explanation of the discrepancy between expectations and realizations is the sharp and largely unanticipated 1993 recession. The explanation usually offered for the recession was strong fiscal contraction and pension reform enacted by the Government in the Fall of 1992 (after the survey was completed), raising taxes, cutting pension benefits and increasing contributions. The recession had different effects for various population groups, hitting particularly the self-employed and the residents of the South. As will be seen, we will exploit knowledge of the groups that suffered mostly from the recession in modeling the structure of the forecast error.

The pattern of expectations and realizations by population groups are also of interest. The young expect earnings to grow faster than the middle-aged and the elderly. Also employees predict their earnings growth more accurately than self-employed in both surveys. In part this is due to the fact that the self-employed have greater income volatility. Yet, comparison between subjective expectations and realizations for the self-employed is difficult, because this group experienced an income decline of 12% in 1992–1993, due to the 1993 recession and tax increases. Finally, expectations by income quartile do not indicate that rich households predict earnings better than poor ones.

Table 2 displays inflation expectations. In both surveys, average expected inflation is roughly 7%, quite close to the forecasts in 1990 (for 1991) and 1992 (for 1993) of sophisticated econometric models and international institutions. Respondents' average expectation for 1990–1991 (7.2%) comes closer to the realized value of 6.8% than OECD's forecast for June 1990–June 1991 (5.4%). Results are reversed for the June 1992–June 1993 period: OECD projections are closer to realizations (4.2 and 4.8%, respectively), while individuals overestimate the actual rate (with average expectations of 7.2%).[10] An interesting feature is that these average subjective inflation expectations do not in fact mask a great number of implausible extreme values. More than 50% of the sample bunches the entire probability distribution for inflation between 5 and 7%. Finally, there is no clear pattern of subjective expectations by region, age, education or income.

[10] One possibility for the larger gap between expectations and realizations in 1992 is that individuals were surprised by the implementation of income policies in July of 1992. These income policies are generally thought to have been effective in reducing the actual inflation rate. An alternative possibility is that consumers form adaptive expectations (in both 1989 and 1991 the inflation rate was 6.3%).

Table 2
Inflation expectations

	1990–1991 (1)	1992–1993 (2)
Age group		
< 35	0.0719	0.0704
35–55	0.0722	0.0747
> 55	0.0715	0.0698
Education		
Junior high-school or less	0.0732	0.0717
High school	0.0693	0.0742
University degree or more	0.0714	0.0712
Occupation		
Employed	0.0720	0.0720
Self-employed	0.0700	0.0734
Region of residence		
North	0.0698	0.0749
Center	0.0663	0.0704
South	0.0760	0.0700
Household income		
I quartile	0.0745	0.0744
II quartile	0.0751	0.0730
III quartile	0.0679	0.0705
IV quartile	0.0708	0.0711
Total sample	0.0719	0.0722
OECD Projection (Consumer prices)	0.0540	0.0415
Realization (Consumer prices)	0.0680	0.0480

Notes: Inflation expectations are given in May–July of 1990 (column (1)) and May–October 1992 (column (2)) for the subsequent 12 months. OECD inflation projections are 12-month forecasts given in June 1990 (column (1)) and June 1992 (column (2)). Inflation projections and realizations refer to the same time periods. Source for inflation projections: *OECD Economic Outlook*, June 1990, vol. 46, pp. 125–127; and June 1992, vol. 51, pp. 127–129. Source for inflation realizations: *OECD Economic Outlook*, June 1994, vol. 55, pp. 68–70.

3.2. Income risk

In the Euler equation it is the term $\mathrm{var}_{i,t}(\Delta \ln C_{i,t+1} - \rho^{-1} r_{i,t+1})$ that affects consumption growth. We assume that the only non-insurable risk faced by individuals is income risk, thus neglecting such other possibilities as rate of return and health risks. The subjective variance of the growth rate of real earnings is $\sigma^2_{i,g} = \sigma^2_{i,z} + \sigma^2_{i,\pi} - 2\phi\sigma_{i,z\pi}$. We have data on the marginal distributions of z and π, but lack information on ϕ, the correlation coefficient between nominal earnings shocks and inflation shocks. Thus in the empirical analysis we rely mainly on the subjective variance of the growth rate of nominal earnings ($\sigma^2_{i,z}$) as our preferred proxy for expected consumption risk. One justification for this choice is that it avoids arbitrary assumptions about the value of ϕ; furthermore, indexation clauses in labor contracts often provide insurance against inflation increases.

Only if utility is exponential and income is a random walk there is a one-to-one correspondence between income risk and consumption risk in the Euler equation. Otherwise, the relation between the two is non-linear, depending on the utility function and the income process. For this reason one cannot give a structural interpretation of the estimated coefficients, i.e. in terms of prudence or underlying preference parameters. We are also aware that our measure of income risk is open to criticism.[11] For instance, we rule out the potential effect of other non-insurable risks faced by households. And yet if income risk is poorly measured, or if income risk is only poorly correlated with consumption risk, one should find no statistical relation between consumption growth and the subjective variance of income.[12]

4. Sample and specification issues

The panel component of the SHIW includes 1,137 households interviewed in 1989 and 1991, 2,420 households interviewed in 1991 and 1993, and 1,050 households interviewed in 1989, 1991, and 1993. Defining an 'observation' as two years of data, this corresponds to 5,657 potential observations (2,187 in the 1989–1991 panel, and 3,470 in the 1991–1993 panel). We drop cases in which the household head changed (355 observations); those with inconsistent data on age, sex, or education (515 observations); those lacking data on subjective expectations (1,123 observations); and those lacking data for other variables used in the empirical analysis (130 observations). The final sample therefore includes 3534 'observations' (1102 for 1989–1991, and 2432 for 1991–1993). Since in most cases we have only one observation per household, we test primarily if the cross-sectional variation in consumption growth is explained by the cross-sectional variation in predicted income growth. We explain below how we deal with this problem.

[11] Given the wording of the questions, the probability of low income states, such as unemployment, may not be reported.

[12] If the variance of income growth is measured with error, one should use an instrumental variables approach. However, it is hard to find suitable instruments to predict the cross-sectional evolution of income risk. In the experiments we performed, the effect of instrumenting the variance of income was to lower the precision of the estimates, possibly because of the weakness of the instruments used.

As in previous studies, we control for individual preferences with age and change in family size.[13] Testing for non-separabilities in the utility function is interesting in its own right and ensures that excess sensitivity does not arise from preference misspecification. Given that in our sample virtually no head is unemployed, we introduce in the Euler equation the change in the employment status of the spouse. As mentioned, omitting labor supply indicators can bias upward the coefficient of expected income growth of the household head. The problem is not as serious than if we had total household earnings (employment is almost surely positively correlated with predicted income growth). However, the earnings of the head may still be correlated with the working spouse dummy because common macroeconomic shocks affect the probability of working and income prospects in the same direction. Other labor supply indicators – such as the change in the number of income recipients – were either not significantly different from zero or did not alter the results.[14]

As mentioned, one should control for the structure of aggregate shocks, particularly in short panels. Even though forecast errors in consumption are unobservable, we do observe the cross-sectional pattern of forecast errors in income. This can be used to extract potentially useful information about the structure of forecast errors in consumption, which depends on the income innovations. For instance, in the absence of common shocks, time dummies should not explain the forecast error. If instead macroeconomic shocks are important, time dummies will be correlated with the forecast errors in income and in consumption, and therefore cannot be used as instruments to predict income. Rather, one should allow for time effects in the Euler equation.

Preliminary analysis indicates that the forecast errors in income ($\Delta \ln Y_{i,t+1} - g^e_{i,t}$) is correlated not only with time dummies, but also with education, and dummies for occupation and region. Given the characteristics of the recessionary episode of 1993, we find it plausible to assume that the forecast error contains an aggregate component which is unevenly distributed across population groups and an idiosyncratic component that averages out in the cross-section.[15] Tax increases for the self-employed or a stronger effect of the

[13] We also tried changes in other demographic variables, such as the number of adults or the number of children. In no case were the main results affected.

[14] Estimating the elasticity of intertemporal substitution has proven to be extremely difficult with panel data. Even in long panels – such as the PSID – the coefficient of the real interest rate is often poorly determined or implausible. Initially, we constructed a measure of the household-specific real interest rate, subtracting inflation expectations from the nominal rate on Treasury bills. However, the coefficient of the elasticity of intertemporal substitution thus obtained was not significantly different from zero and theoretically implausible. In the end, we decided to drop the interest rate from the regressions: using two-year consumption changes with one-period ahead inflation expectations, it is simply impossible to get the timing of the interest rate right.

[15] In the empirical specification we thus assume that the forecast error in consumption growth can be decomposed as $\varepsilon_{i,t+1} = \theta_j \mu_{t+1} + v_{i,t+1}$, where $v_{i,t+1}$ denotes the idiosyncratic component.

1993 recession in the South would have such an effect (see also Miniaci and Weber, 1996). This implies that group dummies (such as region and employment status) should not be used as excluded instruments to predict actual income growth.

Table 3 reports the first-stage coefficients obtained by regressing actual income growth on expected income growth, time dummies, education, regional

Table 3
Predicting actual income growth

	Total sample (1)	Excluding self-employed and farmers (2)
Expected income growth	0.5003*	0.6660*
	(0.1099)	(0.1133)
Education*1991	– 0.0040	– 0.0020
	(0.0029)	(0.0027)
Education*1993	0.0104*	0.0104*
	(0.0021)	(0.0019)
Resident in the South*1991	– 0.0813**	– 0.1053*
	(0.0370)	(0.0345)
Resident in the South*1993	– 0.1110*	– 0.0842*
	(0.0248)	(0.0231)
Resident in the North*1991	– 0.0840**	– 0.1048*
	(0.0372)	(0.0340)
Resident in the North*1993	– 0.0404	– 0.0331*
	(0.0239)	(0.0220)
Self-employed*1991	– 0.1242*	
	(0.0334)	
Self-employed*1993	– 0.3372*	
	(0.0248)	
Farmer*1991	0.0842	
	(0.0696)	
Farmer*1993	0.0291	
	(0.0349)	
Working spouse	0.0156	0.0027
	(0.0171)	(0.0160)
Sample size	3534	2680
Adj.-R^2	0.0708	0.0270
Adj.-R^2 on excluded instruments	0.0055	0.0109
F-test	10.78	15.78
(degrees of freedom)	(2; 3531)	(2; 2677)

Notes: The dependent variable is the growth rate of real after-tax earnings and pensions of the household head. Standard errors are reported in parenthesis. Each regression also includes a constant term, a time-dummy, age, change in family size and the variance of income growth. Column 2 excludes farmers and the self-employed. One and two stars indicate that the variable is statistically different from zero at the 1% and 5% level, respectively.

dummies and employment status interacted with year dummies, lagged employment status of the spouse, age, family size, and income risk. Overall, the first stage regression has good predictive power (the adjusted R^2 statistics is 0.07). The coefficient of expected income growth is 0.5 and significantly different from zero at the 1% level.[16] A conventional F-test on the excluded instruments (expected income and lagged employment status of the spouse) yields a p-value below 1%, confirming the validity of the instruments.

In the following section we thus present instrumental variable estimates of the following Euler equation:

$$\begin{aligned}\Delta \ln C_{i,t+1} = {} & \alpha_1 age_{i,t+1} + \alpha_2 \Delta \ln FS_{i,t+1} \\ & + \eta \sigma^2_{i,z,t} + \gamma \Delta ww_{i,t+1} + \beta \Delta \ln Y_{i,t+1} \\ & + \theta_j \mu_{t+1} + v_{i,t+1}, \end{aligned} \tag{2}$$

where $FS_{i,t+1}$ denotes family size, $\Delta ww_{i,t+1}$ is the change in a dummy for spouse working full-time, $\sigma^2_{i,z,t}$ denotes the expected variance as of time t of nominal income growth, j the population groups affected by macroeconomic shocks, and θ_j captures the effect of unevenly distributed aggregate shocks μ_{t+1} on the forecast error in consumption.[17] In the empirical application we will also present estimates replacing predicted income growth $E_{it}\Delta \ln Y_{it+1}$ with the subjective expectation of income growth $g^e_{i,t}$.

5. Euler equation estimates

The results of estimating Eq. (2) are reported in column 1 of Table 4. The coefficients of the demographic variables are well determined and have the 'right' sign. The positive and significant coefficient of the change in the spouse's employment status indicates that expecting to work more in the future reduces current consumption. This will indeed be the case if leisure and consumption are non-separable. The coefficients of the group dummies are not reported for brevity.

The proxy for consumption risk is positive and significantly different from zero at the 1 percent level, and supports the theory of precautionary saving. Since what we measure is not the expected variance of consumption but the expected variance of income growth, the coefficient has no structural interpretation. Nevertheless, its size (5.67) is most suggestive. With isoelastic utility,

[16] Our instrument predicts well both income increases and income decreases. The first stage coefficients of expected income growth are, respectively, 0.45 and 0.64 in the samples expecting positive and negative income growth.

[17] Our identifying assumption is therefore $p\lim_{N\to\infty} N^{-1}\sum_i v_{i,t+1} = 0$, where N is the number of households.

Table 4
Euler equation estimates

	Baseline specification		Splitting the sample by the wealth–income ratio		Using expected income as a regressor	
	Total sample (1)	Excluding farmers and self-employed (2)	Low-wealth (3)	High-wealth (4)	Expected income replaces actual income (5)	Asymmetry test (6)
Age	0.0009**	0.0013**	(0.0007	0.0011	0.0013**	0.0009**
	(0.0005)	(0.0006)	(0.0007)	(0.0007)	(0.0005)	(0.0005)
Δ ln (Family size)	0.3405*	0.3334*	0.3706*	0.3375*	0.3442*	0.3441*
	(0.0533)	(0.0583)	(0.0810)	(0.0656)	(0.0510)	(0.0510)
Δ (Working spouse)	0.3391*	0.4156*	0.2977*	0.3314*	0.3447*	0.3444*
	(0.0693)	(0.0814)	(0.1028)	(0.0902)	(0.0667)	(0.0667)
Variance of income growth	5.6719*	5.9123*	9.1301	5.3033*	5.6442*	5.3578*
	(1.9744)	(1.7715)	(20.4515)	(1.8285)	(1.9273)	(1.9566)
$\Delta \ln Y_{i,t+1}$	− 0.0835	− 0.0514	0.2309	− 0.034		
	(0.1928)	(0.1469)	(0.7714)	(0.1532)		
Expected income growth					− 0.0418	
					(0.0924)	
Expected income increase						0.0522
						(0.2563)
Expected income decline						− 0.0687
						(0.1051)
Sample size	3,534	2,680	1,108	2,426	3,534	3,534

Notes: The dependent variable is the growth rate of non-durable consumption expenditures. $\Delta \ln Y_{i,t+1}$ is the after-tax real growth rate of earnings and pensions of the household head. Each regression also includes time dummies, interaction of education with year, and interactions of year and dummies (dated t) for region, self-employed and farmer (omitted in column (2)). In columns (1)–(4) the instruments used are expected income growth and the lagged employment status of the spouse. In columns (3) and (4) an observation is included in the low-asset group (high-asset group) if, at the beginning of the period, the wealth-income ratio is smaller (greater) than 2 (wealth is real estate plus financial assets less household debt). Standard errors corrected for heteroscedasticity of unknown form are reported in parenthesis. One and two stars indicate that the variable is statistically different from zero at the 1% and 5% level, respectively.

prudence equals one plus relative risk aversion, and reasonable values for risk aversion vary between 1 and 10.

It is important to note that ignoring the group dummies induces a correlation between the cross-sectional variation in consumption growth and the cross-sectional variation in income growth leading to spurious evidence in favor of excess sensitivity. In fact, if one assumes that the forecast errors can be decomposed into an aggregate shock and an idiosyncratic shock, as in most of the literature (though we know it cannot from the pattern of the forecast error in income), introducing time dummies in the Euler equation should provide consistent estimates. If education and dummies for region and occupation, in addition to expected income, are then used as instruments for income growth, one does find excess sensitivity (a coefficient of 0.32 with a t-statistics of 5). However, when the time dummies *and* their interactions with group dummies are added to the Euler equation (thus controlling for the structure of the forecast error) such evidence vanishes, as in Table 4. Note also that excluding the dummy for working wife and the variance of income growth does not affect the excess sensitivity coefficient. Thus, in our sample there is no excess sensitivity even when the Euler equation is misspecified.

How should one interpret the role of group dummies and education in the Euler equation? Even though they were introduced as a device to eliminate the inconsistency of IV estimates in short panels, at least two other interpretations are possible. First, group dummies may account for preference shifters and for this reason should not be omitted from the Euler equation, otherwise income growth will simply proxy for the omitted variables (absent group dummies, excess sensitivity is just a signal of misspecified preferences). The second possibility is that there is a subtler form of excess sensitivity, arising not from the correlation between consumption and income, but from the correlation between consumption and income predictors. To clarify this point, suppose that (low) education, residence in the South and self-employment are predictors of the probability of being liquidity constrained in period t. If so, one may expect them to predict higher consumption growth between period t and $t+1$. However, in the regressions of Table 4 the dummies for South and self-employment are negative, while the coefficient of education is positive (with the exception of the dummy for South in 1993, the other interaction terms are not statistically significant). While alternative explanations for the effect of group dummies are therefore possible, we find it more plausible to attribute their role to the effect of unexpected aggregate shocks.

So far, our sample has included farmers and the self-employed (854 observations). There are several reasons why it may be desirable to test the robustness of the results when these observations are excluded: reported income for the self-employed income is severely under-estimated (Brandolini and Cannari, 1994); some individuals may have chosen self-employment, a more risky occupation, because they are less risk averse than the rest of the population, inducing sample selection; for farmers it is not easy to measure income or to distinguish it from consumption. The first-stage regression excluding farmers and the self-employed is reported in column (2) of Table 3. The coefficient of expected income growth increases to 0.67, indicating again that this variable is a powerful instrument to predict actual income growth. Column (2) of Table 4 replicates the regressions of column (1) using the restricted sample. There is again no evidence of excess sensitivity (column (2)), and the other coefficients are only marginally affected.

An excess sensitivity coefficient of zero may hide possible asymmetric responses of consumption growth to predicted income growth. The well-known approach of Zeldes (1989) is to split the sample according to the asset–income ratio. If liquidity constraints are the only source of failure of the model, one would find excess sensitivity in the low-asset but not in the high-asset group, in that affluent households can always overcome borrowing constraints by drawing on assets, while the less wealthy cannot. In Table 4 households are defined as 'poor' if total net worth (including real estate wealth) does not exceed twice annual income. The sample split thus places about 30% of the sample in the low-asset group and 70% in the high-asset group. It is apparent that we find no evidence of excess sensitivity in either group (two insignificant coefficients of 0.23 and -0.03 in the low-asset and high-asset groups, respectively).[18]

Under liquidity constraints the response of consumption to predictable income growth should be asymmetric (Altonji and Siow, 1987). If consumers expect their income to increase, they would like to borrow but are prevented from doing so: consumption growth will then respond to predicted income growth. If instead consumers expect income to fall, they will save, not borrow: in this case the liquidity constraint is not binding, and one should not find a violation of the orthogonality conditions.

Our instrument for income growth offers an opportunity to test for the potential asymmetric response of consumption to expected income growth. For comparison with previous estimates, in column (5) of Table 4 we replace (instrumented) actual income growth with expected income growth. Given the endogeneity of $\Delta ww_{i,t+1}$ the equation is again estimated by instrumental variables, and the previous results are confirmed.[19] We then capture the potential

[18] Results are qualitatively unaffected if we split the sample according to the ratio of financial assets to income or if we vary (upwards or downwards) the threshold used to split the sample. In all cases the low-asset group tends to be younger, less educated, with fewer self-employed and lower income than the high-wealth group. Given that reducing the threshold used to split the sample reduces the group of low-asset households, the estimated coefficients tend to be less precisely estimated.

[19] Since expectations are available only about bands of possible income and inflation values, our measure of income risk will entail a certain amount of measurement error. We replicate regression 5 in Table 4 by OLS, omitting the change in the employment status of the wife, with results basically unaffected. Since in an OLS context measurement error in an independent variable tends to bias the coefficients towards zero, we take this as an indication that measurement error cannot explain, alone, a significant coefficient of income risk. For the same reason, we cannot rule out that measurement error in expected income biases the excess sensitivity coefficient towards zero in columns (1)–(4).

non-linear effect of expected income growth estimating

$$\Delta \ln C_{i,t+1} = \alpha_1 age_{i,t+1} + \alpha_2 \Delta \ln FS_{i,t+1} + \eta\sigma^2_{i,z,t} + \gamma \Delta ww_{i,t+1} + \beta_1 g^{e+}_{i,t} + \beta_2 g^{e-}_{i,t} + \theta_j \mu_{t+1} + v_{i,t+1}, \quad (3)$$

where $g^{e+}_{i,t}$ denotes positive (or zero) expected income growth, and $g^{e-}_{i,t}$ denotes negative expected income growth.[20] In column (6) of Table 4 we do not find evidence of asymmetric effects: the coefficients of positive and negative expected income growth are 0.07 and -0.06, respectively, and are not significantly different from zero or from each other. The asymmetry test was replicated also splitting the sample by assets. Under liquidity constraints one should find excess sensitivity mainly in the group of poor households that expect an increase in income. However, even in this case we cannot reject the null hypothesis of no asymmetric effects (whether or not the self-employed are included in the sample).

We performed several tests to check the robustness of the results. Here we briefly comment on higher moments of the expected income growth variable, sample selection arising from non-responses, the definition of the sample, and alternative instruments to predict income growth.[21] The survey questions allow us to estimate higher moments of the conditional distribution of expected income growth, not just the variance, which is only a valid indicator of risk under restrictive assumptions. For instance, households may react more strongly to the risk of low income realizations. We thus introduced an index of asymmetry of the distribution of income growth and dummies for households that expected with relatively high probability (more than 20%) a large decline in income (more than 5%). These variables were not significantly different from zero.

Our estimates may be criticized on the ground that the respondents reporting expectations presumably understand the survey questions better than those who do not. A formal test of this hypothesis can be made by controlling explicitly for selection bias arising from non-responses. We thus run a probit regression for the probability of response, assuming that the probability is related to demographic and economic variables (income, education, age, occupation, industry, and region of residence). The implied Mills ratio was then added as a regressor to the Euler equation. The ratio was not significantly different from zero and results were again similar to those reported in the basic specification, suggesting that this effect is not important. We also checked the stability of the coefficients with respect to several sample exclusions: individuals older than 40 or 50, households with more than two income recipients, and households whose head

[20] Those who expect their income to decline are less wealthy, less educated, and more likely to be near to the retirement (or already retired).

[21] For brevity these results are not reported, but are available on request.

is a pension income recipient. In no case did the pattern of results change appreciably.

Finally, our conclusions are qualitatively unchanged if we use lagged income growth, rather than expected income growth, to predict actual income growth. For this purpose we must use the sub-sample of households surveyed in 1989, 1991, and 1993. Here we find again evidence for excess sensitivity if we do not control for the stochastic structure of the forecast error (a coefficient of 0.19 with a t-statistics of 2.1), but no excess sensitivity when education and group dummies (interacted with time) are introduced as additional regressors to the Euler equation (a statistically insignificant coefficient of -0.01). The problem with using lagged income growth is that if income is measured with error, the first lag of income growth is not a valid instrument, as measurement error violates the orthogonality conditions. The advantage of using expected income growth is that the instrument is valid whether or not income is measured with errors.

6. Conclusions

After more than a decade of studies testing the theory of households' intertemporal choices on panel data, the evidence is mixed (Browning and Lusardi, 1996). In this paper we test for excess sensitivity using a 1989–1993 panel of Italian households that provides measures of income and inflation expectations and income risk. The expectations are used as an instrument for predicting income growth. Controlling for income risk, predictable changes in employment status of household members, and for aggregate shocks that affect differently population groups, we find that consumption growth is uncorrelated with the expected earnings growth of the household head. We also find that predictable proxies of changes in labor supply and expected income risk affects positively consumption growth. To the extent that income risk is correlated with expected consumption risk, this finding supports the theory of precautionary saving.

Our results are robust to a variety of experiments such as asymmetric response of consumption to positive or negative expected income growth and sample splits by assets. It is worth stressing that our result of no excess sensitivity depends on the validity of subjective income expectations to predict income growth. The correlation between the two is statistically significant, but the instrument may not be powerful enough to capture small departures from the permanent income hypothesis.

Given the severe imperfections of the Italian credit markets by the standards of other industrialized countries and the pervasiveness of various liquidity constraints, particularly in the mortgage market (Guiso et al., 1994), the fact that we do not find excess sensitivity may come as a surprise, since often excess sensitivity has been linked to liquidity constraints. But it is precisely for this reason that Italian households are high savers, and even at young ages have accumulated considerable assets to buffer income fluctuations. This indicates that excess sensitivity tests have limited power against models in which borrowing constraints play an important role. For instance, prudent consumers will save in anticipation of future constraints, and may never exhibit excess sensitivity to predicted income growth. Consumers who are saving to purchase a house are globally constrained because they must meet a downpayment, but the orthogonality condition does not fail, except perhaps at the time of the purchase. Thus our results should not be viewed as a contradiction that borrowing constraints play an important role in the Italian economy; rather, as evidence confirming how difficult it is to detect liquidity constraints in structural models of intertemporal choices by conventional excess sensitivity tests.

Acknowledgements

We thank two anonymous referees, the editor, Chris Carroll, Angus Deaton, Sydney Ludvigson and seminar participants at Dartmouth College, the New York Fed and Princeton University for helpful comments. We retain responsibility for any errors. This paper is part of a research project on 'Structural Analysis of Household Savings and Wealth Positions over the Life-Cycle'. Financial support has been provided by The Training and Mobility of Researchers Network Program (TMR) of the European Commission DGXII and by the Italian National Research Council (CNR).

References

Altonji, J.G., Siow, A., 1987. Testing the response of consumption to income with (noisy) panel data. Quarterly Journal of Economics 102, 293–328.

Attanasio, O.P., Weber, G., 1995. Is consumption growth consistent with intertemporal optimization? Evidence from the Consumer Expenditure Survey. Journal of Political Economy 103, 1121–1157.

Browning, M., Lusardi, A., 1996. Household saving: Micro theories and micro facts. Journal of Economic Literature 34, 1797–1855.

Brandolini, A., Cannari, L., 1994. Methodological appendix: The bank of Italy survey of household income and wealth. In: Ando, A., Guiso, L., Visco, I. (Eds.), Saving and the accumulation of wealth. Essays on Italian households and government behavior. Cambridge University Press, Cambridge.

Carroll, C.D., 1992. Buffer stock saving and the life cycle/permanent income hypothesis. Brookings Papers on Economic Activity, 2, 61–156.

Chamberlain, G., 1984. Panel data. In: Griliches, Z., Intriligator, M.D. (Eds.), Handbook of Econometrics, vol. 2, North-Holland, Amsterdam.

Das, M., Van Soest, A., 1997. Expected and realized income changes: Evidence from the Dutch socio-economic panel. Journal of Economic Behavior and Organization 32, 137–154.

Deaton, A., 1992. Understanding consumption. Oxford University Press, Oxford.

Dominitz, J., Manski, C.F., 1997. Using expectation data to study subjective income expectations. Journal of the American Statistical Association 92, 855–867.

Dynan, K.E., 1993. How prudent are consumers? Journal of Political Economy 101, 1104–1113.
Flavin, M., 1994. The joint consumption/asset demand decision: A case study in robust estimation. Mimeo. University of California, San Diego.
Garcia, R., Lusardi, A., Ng, S., 1997. Excess sensitivity and asymmetries in consumption: An empirical investigation. Journal of Money, Credit, and Banking 29, 154–176.
Guiso, L., Jappelli, T., Terlizzese, D., 1992. Earnings uncertainty and precautionary saving. Journal of Monetary Economics 30, 307–337.
Guiso, L., Jappelli, T., Terlizzese, D., 1994. Why is Italy's saving rate so high? In: Ando, A., Guiso, L., Visco, I. (Eds.), Saving and the accumulation of wealth: Essays on Italian households and government behavior. Cambridge University Press, Cambridge.
Hayashi, F., 1985. The permanent income hypothesis and consumption durability: Analysis based on Japanese panel data. Quarterly Journal of Economics 100, 1083–1113.
Kimball, M., 1990. Precautionary saving in the small and in the large. Econometrica 58, 53–73.
Ludvigson, S., Paxson, C., 1997. Approximation bias in linearized Euler equations. Mimeo. Princeton University.
Lusardi, A., 1996. Permanent income, current income and consumption: Evidence from two panel data sets. Journal of Business and Economics Statistics 14, 81–90.
Meghir, C., Weber, G., 1996. Intertemporal non-separability or borrowing restrictions? A disaggregate analysis using a US consumption panel. Econometrica 64, 1151–1182.
Miniaci, R., Weber, G., 1996. The Italian recession of 1993: Aggregate implications of microeconomic evidence. Mimeo. University of Padua, Padua.
Shea, J., 1995. Union contracts and the life-cycle permanent income hypothesis. American Economic Review 85, 186–200.
Zeldes, S.P., 1989. Consumption and liquidity constraints: An empirical investigation. Journal of Political Economy 97, 305–346.