



WORKING PAPER NO. 421

Consumption Uncertainty and Precautionary Saving

Dimitris Christelis, Dimitris Georgarakos, Tullio Jappelli and Maarten van Rooij

December 2015

This versione January 2016



University of Naples Federico II



University of Salerno



Bocconi University, Milan

WORKING PAPER NO. 421

Consumption Uncertainty and Precautionary Saving

Dimitris Christelis^{*}, Dimitris Georgarakos^{},
Tullio Jappelli^{***} and Maarten van Rooij^{****}**

Abstract

Using survey data from a representative sample of Dutch households, we estimate the strength of the precautionary saving motive by eliciting subjective expectations on future consumption. We find that expected consumption risk is higher for the young and the self-employed, and is correlated positively with income risk. We insert these subjective expectations (rather than consumption realizations, as in the existing literature) in a Euler equation for consumption, and estimate the degree of prudence by associating expected consumption risk with expected consumption growth. Robust OLS and IV estimates both indicate a coefficient of relative prudence of around 2.

Keywords: Consumption Risk, Euler Equation, Prudence, Precautionary Saving, Subjective Expectations.

JEL Classification: D12, D14, D81, E03, E21, C14

Acknowledgements: Tullio Jappelli acknowledges financial support from the Italian Ministry of University and Research (PRIN project) and the EC 7th Framework Program (MOPACT project). We thank Efi Adamopoulou, Laura Bartiloro, Tom Crossley, Luigi Pistaferri, Ernesto Villanueva, Gianluca Violante, Guglielmo Weber, seminar participants at the 2015 NBER Summer Institute, the ECB Conference on Household Finance and Consumption, the joint Bank de France - Deutsche Bundesbank Spring Conference, the Symposium on Economics and Institutions, the Modena-Netspar Conference on Advances in Household Finance, and CSEF for helpful comments. The opinions expressed in the paper are those of the authors and do not necessarily reflect the views of the Deutsche Bundesbank or of De Nederlandsche Bank.

* Università di Napoli Federico II, CSEF, CFS, CEPAR and Netspar

** Deutsche Bundesbank, CFS and University of Leicester

*** Università di Napoli Federico II, CSEF and CEPR

**** De Nederlandsche Bank and Netspar

Table of contents

1. Introduction

2. The Euler equation with precautionary saving

3. Empirical tests of precautionary saving

4. The data

5. Empirical results

6. Liquidity constraints

7. Conclusions

References

Figures, Tables and Appendix

1. Introduction

The effect of uncertainty on consumer behavior is a long-standing topic in research on household saving (see e.g. Skinner, 1988; Deaton, 1991; Dynan, 1993; Bertola et al., 2005). Life-cycle models of consumption behavior typically imply that increased income uncertainty will increase precautionary saving and consumption growth by lowering current consumption. This increase in saving depends on the third derivative of the utility function and the associated coefficient of prudence (Kimball, 1990) which in the case of isoelastic utility, is proportional to relative risk aversion.

In a standard Euler equation framework, expected consumption risk induced by income risk or other sources of risk (such as health risk) raises expected consumption growth. However, neither expected consumption growth nor its variability are typically observed in household surveys. For this reason, most tests of precautionary saving use other methods such as structural models or quasi-experimental approaches. Structural models require a greater number of assumptions than the Euler equation; quasi-experimental estimates do not deliver estimates of the structural parameters of the utility function.

The few empirical tests of precautionary saving using the Euler equation substitute for expected consumption growth and consumption risk their observed counterparts: actual consumption growth is regressed on actual consumption risk (Dynan, 1993; Bertola et al., 2005). Since the Euler equation typically includes a forecasting error, this substitution of expectations for realizations renders the observed consumption growth variability almost surely correlated with the Euler equation error term. As a result, the identification of the effect of consumption risk on expected consumption growth becomes very problematic. Indeed, using realized consumption growth and risk instead of expected consumption growth and risk implies that the difference between them (i.e. the forecasting error) enters the error term of the estimated Euler equation. As discussed in Hayashi (1987), this expectation error should converge to zero as the time dimension of the data increases but the same is not true for a short panel. This, as Chamberlain (1984) pointed out, is a serious problem because it leads to inconsistent estimates in short panel surveys such as those that typically contain information on consumption.

The first contribution of the present paper is that it constructs measures of expected consumption growth and expected consumption risk. These measures are derived from responses to a survey that asked participants about their future consumption. The survey data we use come from the CentER Internet panel which is sponsored by the Dutch National Bank and maintained by CentERdata at Tilburg University, and is representative of the Dutch population. The measures of expected consumption growth and its variability deduced from these questions take household-specific values which can be correlated to observable household characteristics. We find that our measures of expected consumption risk are associated with these characteristics, in the direction suggested by economic intuition. For example, expected consumption risk is higher for the young, and for the self-employed. Moreover, income risk is positively associated with consumption risk but is not the only determinant of consumption risk. This means that other sources of risk (such as health risk), and institutions (for instance, the pooling of incomes within the family, or social insurance programs) are likely to affect consumption risk and the relation between income and consumption risk. Overall, the survey questions responses seem to be good indicators of the uncertainty about future consumption experienced by the households in the sample.

The second contribution of the paper is the first (to our knowledge) use of these expectation measures to estimate an Euler equation for consumption. Using expectations-based variables rather than observed magnitudes eliminates the aforementioned problem of an expectation error in the disturbance term. It also addresses the issue of endogeneity of the variable for observed consumption growth variability, discussed by Carroll (2001) and Bertola et al. (2005). More generally, the paper contributes to a growing strand of literature, pioneered by Manski (2004), that uses subjective expectations to elicit individual-level measures of income and unemployment uncertainty (Dominitz and Manski, 1997; Guiso et al., 2002; Jappelli and Pistaferri, 2000), pension uncertainty (Guiso et al., 2013), and interest rate uncertainty (Crump et al., 2015).

However, while avoiding the endogeneity problems present in the existing literature, we also examine the possibility that expected consumption risk is correlated with the Euler equation error term. We address this issue by using expected income risk as an instrument, as in Bertola et al. (2005). Since in the Euler equation expected consumption

risk is a sufficient statistic for expected consumption growth, income risk is a suitable instrument: it is correlated with consumption risk but not with expected consumption growth. We also check the validity of the instruments using the instrumental variable (IV) method introduced by Lewbel (2012).

When estimating the Euler equation, we find that expected consumption risk is positively associated with expected consumption growth, consistent with intertemporal consumption models with precautionary saving. Using robust OLS regression methods, we find that the implied magnitude of the coefficient of relative prudence is around 2. If the utility function is isoelastic, then in turn, this estimate implies that the coefficient of relative risk aversion is about 1. These results hold also if we exclude from the sample those households that are likely to be liquidity constrained, to which the Euler equation does not apply. The IV estimates and the robust OLS estimates are similar. Importantly, when testing for endogeneity of the expected consumption risk variable, we cannot reject the null hypothesis of no endogeneity. Overall, the results for the strength of the precautionary saving motive and for the measures of household prudence and curvature of the utility function are empirically robust and economically plausible.

The paper is organized as follows. Section 2 surveys the empirical literature on precautionary saving and Section 3 reviews the main empirical tests, focusing especially on approaches relying on the Euler equation. Section 4 describes the survey data and Section 5 presents the empirical results. Section 6 extends the analysis to include the presence of liquidity constraints. Section 7 reports results from IV estimation, while Section 8 concludes the paper.

2. The Euler equation with precautionary saving

The relationship between expected consumption risk and expected consumption growth can be described using a second-order approximation of the optimal consumption rule along the lines suggested by Dynan (1993). With a constant interest rate, the Euler equation for consumption states the marginal utility of consumption in period t is proportional to the expected marginal utility:

$$u'(c_t) = \frac{1+r}{1+\delta} E_t u'(c_{t+1}) \quad (1)$$

A second-order Taylor series expansion of $u'(c_{t+1})$ around c_t of equation (1) yields:

$$u'(c_t) = \frac{1+r}{1+\delta} E_t \left[u'(c_t) + u''(c_t)(c_{t+1} - c_t) + \frac{1}{2} u'''(c_t)(c_{t+1} - c_t)^2 + n_t \right] \quad (2)$$

where n_t is a remainder with third and higher order terms in the approximation.

Dividing equation (2) by c_t^2 , and solving for the expected growth rate of consumption one obtains:

$$E_t \left(\frac{c_{t+1} - c_t}{c_t} \right) = EIS \left(\frac{r - \delta}{1 + r} \right) + \frac{1}{2} p(c) E_t \left(\frac{c_{t+1} - c_t}{c_t} \right)^2 + R_t \quad (3)$$

where $p(c) \equiv u'''(c_t)c_t/u''(c_t)$ denotes Kimball's coefficient of relative prudence, $EIS \equiv -u'(c_t)/(u''(c_t)c_t)$ is the elasticity of intertemporal of substitution, and R_t is a remainder term due to the second-order approximation. Note that, with standard preferences, the EIS is equal also to the inverse of the coefficient of relative risk aversion. The second uncentered moment of the distribution of expected consumption growth $E_t[((c_{t+1} - c_t)/c_t)^2]$ is a measure of the expected consumption risk.

Equation (3) indicates that an increase in the expected consumption risk is associated with higher expected consumption growth. The intuition is that in order to buffer the increase in consumption risk individuals consume less in period t relative to period $t + 1$, and thus increase current saving. Furthermore, the sensitivity of consumption growth to consumption risk is proportional to the coefficient of relative prudence. If utility is quadratic, then $u'''(.) = 0$; therefore, expected consumption risk does not affect expected consumption growth, and the consumption profile depends only on the elasticity of intertemporal substitution, the interest rate, and the rate of time preference.¹ Therefore, a test of the hypothesis that consumption risk does not affect consumption growth is also a test of the validity of the certainty equivalence model.

¹ However, in the presence of specific assumptions about preferences and the probability distribution of future consumption growth, an explicit solution for the expected growth rate of consumption is obtained

Equation (3) is as an equilibrium condition that contains a parameter of interest $p(c)$. It can also be thought of as an equation describing the effect of a change in expected consumption risk (induced by underlying income, health, family, or other risks) on expected consumption growth, and thus also on precautionary saving.

As discussed in Section 3 below, in most applications consumption risk is assumed to depend only on income risk because the variability of future earnings is assumed to be the only source of uncertainty. Within that framework, some models distinguish between movements in hours and in wages. For instance, Abowd and Card (1989) decompose fluctuations in earnings, often taken as the standard measure of risk, into exogenous fluctuations in wages and endogenous choice of hours. However, only the first type of fluctuation represents genuine risk, even though variation in hours has welfare consequences since people value leisure. Low et al. (2010) model labor supply and job mobility in a search and matching framework. Their approach distinguishes between shocks, and responses to shocks, and between employment risk (such as exogenous job destruction and lack of offers when unemployed), and productivity shocks (such as health shocks or poor matches in the labor market). Low and Pistaferri (2015) model productivity risk by distinguishing between health shocks (in the form of risk of disability), and shocks to the price and quantity of skills (such as those related to skill biased technological changes).

In more general models, consumption risk might also reflect uncertainty related to other random variables, and thus might not be related only to income risk. While income risk might be the most important source of consumption risk for young individuals, this might not be the case at other stages in the life-cycle. Alongside income risk, people face a number of other uninsurable or partially uninsurable risks which can affect intertemporal consumption decisions. The literature gives prominence to some of these risks (and their associated costs). They include risk of future liquidity constraints, shocks to asset prices (including house prices), risk of medical and other unexpected expenditures, and risk of family dissolution (see, e.g., Palumbo, 1999; Voena, 2015).

(Hansen and Singleton, 1983). For instance, assuming an isoelastic utility function, constant interest rate, and a normally distributed conditional distribution of consumption growth, one obtains the closed form solution $E_t \Delta \ln c_{t+1} = \gamma^{-1}(r - \delta) + (\gamma/2) \text{Var}_t \Delta \ln c_{t+1}$.

3. Empirical tests of precautionary saving

Several research strategies have been used to test the importance of precautionary saving. A first group of studies attempts to estimate the impact of income risk on the reduced forms of consumption or wealth. Measures of income risk based on actual earnings are difficult to compute even with long panel data, and in part, may reflect a choice (for instance, the choice to work in a risky occupation). Empirical evidence based on this approach is mixed. Most papers find a positive relation between wealth and income risk which is consistent with the precautionary saving model. However, the magnitude of the effect varies widely across studies, and on net tends to be small (Jappelli and Pistaferri, 2010). This approach also provides evidence in favor or against precautionary saving but does not deliver estimates of the parameters of the utility function (such as the coefficient of relative prudence).

A second group of studies estimate the paths of consumption and wealth in models with precautionary saving, matching simulated data to observed wealth and consumption distributions. Pioneering this approach, Gourinchas and Parker (2002) use consumption data from the US Consumer Expenditure Survey (CEX) and income data from the Panel Study of Income Dynamics (PSID) to estimate the rate of time preference and risk aversion. Their estimation methodology minimizes the distance between actual consumption and its predicted life-cycle profile. Setting the real interest rate at 3 percent, Gourinchas and Parker estimate a time preference rate of approximately 4 percent, and an elasticity of intertemporal substitution of about 0.5, corresponding to a coefficient of relative risk aversion of about 2. With an isoelastic utility function the implied coefficient of relative prudence is about 3. Cagetti (2003) estimates the same preference parameters as Gourinchas and Parker (2002), by matching simulated and actual median wealth profiles over the life cycle using US data from the PSID and the Survey of Consumer Finances (SCF). Cagetti finds higher estimates of the time preference rate, a higher coefficient of relative risk aversion (around 4 for the high school sample), and an implied coefficient of relative prudence of around 5. Structural estimations deliver estimates of

the parameters of the utility function but require the utility function, the budget constraint, the sources of risks, and the income process to be specified.

An alternative direct strategy to estimate the coefficient of relative prudence is to measure expected consumption growth and expected consumption risk in equation (3) using survey data that record respondents' own assessments of these variables. This is the approach adopted in the present paper.

From an empirical point of view, the main problem related to estimating the Euler equation is that expected consumption growth and expected consumption risk are generally not observable. Were it possible to measure the expectation-related terms in equation (3), then it would be possible also to estimate the equation using OLS, and to identify the coefficient of the term related to expected consumption risk, which would be proportional to relative prudence. The first paper to attempt to identify this coefficient in the Euler equation framework was Dynan (1993); however, her data do not include any information on expectations about consumption. As a result, she replaces expectations with their realized counterparts. In this case, indexing households by i ($i = 1, \dots, N$), equation (3) can be written in a regression framework as:

$$g_{i,t+1} = \alpha + \beta g_{i,t+1}^2 + \varepsilon_{i,t+1} \quad (4)$$

where $g_{i,t+1} = (c_{i,t+1} - c_{i,t})/c_{i,t}$, and the coefficient β equals $(1/2)p(c)$, and thus is directly related to the strength of the precautionary saving motive. The term $\varepsilon_{i,t+1}$ is a composite error term reflecting innovations to consumption growth, higher order terms of the Taylor expansion, measurement error, and possibly heterogeneity in preferences. In particular, the substitution of realized consumption changes for their expectations implies that $\varepsilon_{i,t+1}$ in equation (4) includes the difference between expected and realized magnitudes, and thus is clearly correlated with the term denoting realized consumption risk. For instance, if households have positive news about the economy between periods t and $t+1$, they may revise consumption upwards in period $t+1$, affecting both the mean and the variance of the (ex-post) consumption distribution.

To address this endogeneity issue, Dynan uses an IV approach applied to panel data drawn from the CEX. The set of instruments includes education and occupation on the assumption that these characteristics are correlated with the expected consumption risk,

and that they affect expected consumption growth only through this channel. Overall, these instruments have low power, and hence the coefficient of relative prudence is imprecisely estimated.

Jappelli and Pistaferri (2000) test for precautionary saving using Italian panel data from the Survey of Household Income and Wealth (SHIW) that in 1989-93 has measures of subjective income expectations. These expectations of income risk generate a household-specific subjective distribution of income uncertainty, and therefore in a measure of income risk that takes different values across households. They find that realized consumption growth is positively correlated with the subjective variance of income growth.

Dynan's approach was refined by Bertola et al. (2005) who use the subjective variance of income growth available in the SHIW as an instrument for $((c_{i,t+1} - c_{i,t})/c_{i,t})^2$. As they point out, subjective income risk should be a valid instrument because, as (3) implies, income risk has no direct effect on consumption growth once one conditions on expected consumption risk. In other words, expected consumption risk is a sufficient statistic for expected consumption growth. Bertola et al. (2005) find that subjective income risk is not only a powerful instrument but also one that delivers empirically plausible results. In particular, their coefficient of relative prudence is about 2 and is precisely estimated, thus providing evidence supporting the precautionary saving model.

In the present paper, and in contrast to Dynan (1993), Jappelli and Pistaferri (2000) and Bertola et al. (2005), we estimate equation (3) directly using subjective expectations of future consumption, rather than relying on realized consumption magnitudes. We can rewrite equation (3) as:

$$E_{i,t}(g_{i,t+1}) = \alpha + \beta E_{i,t}(g_{i,t+1}^2) + v_{i,t+1} \quad (5)$$

where $v_{i,t+1}$ is a composite error term that includes higher order terms of the Taylor expansion, measurement error, and possibly other unobservable variables that affect expected consumption growth.

Estimating equation (5) rather than equation (4) has two advantages with respect to previous tests of precautionary saving. First and most importantly, the error term of

equation (4), $\varepsilon_{i,t+1}$, includes the expectation error of the Euler equation, while by construction, the error term of equation (5), $v_{i,t+1}$, does not. Hence, it is not correlated with expected consumption risk.

A second, related issue is that equation (5) can be estimated even with a cross-section or with a short panel, by exploiting the cross-sectional variability in expectations of the consumption distribution. The literature shows that Euler equation estimates derived from panel data may be inconsistent when the time dimension of the panel is short (Chamberlain, 1984; Hayashi, 1987). The reason is precisely that the error term of the Euler equation (4) includes a forecasting error. Life-cycle and permanent income models imply that the expectation of a forecasting error, conditional on any information available at t should be zero over a long horizon. In other words, the error should not exhibit systematic patterns if the model is correct. Following this logic, the empirical equivalent of $E_{i,t}(\varepsilon_{i,t+1})$ in (4), which includes the expectation of the forecasting error, is a household-level average taken over T periods. Importantly, $T \rightarrow \infty$ is needed to ensure consistency of the estimated parameters of the Euler equation. However, panel surveys containing information on consumption are typically short, and hence researchers often proceed under the assumption that consistency is achieved with $N \rightarrow \infty$, i.e. assuming that forecasting errors average out to zero in the cross-section.

There is no reason to believe that this assumption holds generally. For example, if there are aggregate shocks, households will likely make forecasting errors in the same direction in a given year (Altug and Miller, 1990). In this case, the cross-sectional average of the forecasting error will most likely be different from zero. One way to overcome this problem is to add year dummies to the Euler equation. However, this approach still might fail to deliver consistent estimates if the aggregate shocks are unevenly distributed across consumers, so that the time dummies do not completely absorb their impact.

As already discussed, given that we use consumption expectations rather than realizations, the assumption that $v_{i,t+1}$ has a zero conditional expectation in the cross-section is quite reasonable, implying that OLS can be used to estimate equation (5). Nevertheless, there might be unobservable variables that are correlated with both the expected consumption risk and the expected consumption growth, or the error term might

contain higher order terms of the Taylor expansion that are correlated with expected consumption risk. Hence, we check the robustness of the results using an IV estimator, relying on expected income risk as an instrument (see Section 7).

4. The data

We use data from the CentER Internet panel, which is sponsored by the Dutch National Bank and maintained by CentERdata at Tilburg University. The baseline survey is conducted once a year via the Internet, and collects detailed information on a range of demographics and asset holdings for a representative sample of Dutch-speaking households. In addition to the baseline survey, it is common for households to be asked to participate during the course of a year in special purpose surveys.

We designed such a survey containing questions aimed at measuring individual uncertainty about future consumption and income, and expected household consumption growth. We administered our survey first to Internet panel participants, in June 2014. We repeated the survey in January 2015 and June 2015 to check for a seasonal pattern in responses, and to increase the sample size used in our analysis. We targeted the financial respondent in each household, i.e. the person responsible for household finances.

In a recent paper, Crump et al. (2015) estimate the elasticity of expected consumption growth with respect to variation in the expected real interest rate using the Federal Reserve Bank of New York Survey of Consumer Expectations (SCE). This dataset includes consumers' expectations of consumption growth (but not consumption risk) and inflation, with the latter providing subjective variation in ex ante real interest rates. The Euler equation estimates which omit the conditional variance term of the Euler equation indicate that the elasticity of intertemporal of substitution is around 0.8.

To elicit the distribution of expected consumption we follow a similar procedure to Guiso et al. (2002, 2013) whose interest is in approximating the subjective distribution of future income and the pension replacement rate, respectively. Specifically, we asked respondents first to report the minimum (y_m) and the maximum (y_M) values of next year's consumption in a typical month, and then to rank on a 0-100 scale the probability

that consumption will be higher than the mid-point between the minimum and the maximum, i.e. $\pi = \text{Prob}(y \geq (y_m + y_M)/2)$. The questions used are provided in Appendix A.1.

To estimate the moments of the subjective distribution of future consumption we rely on the assumptions and methods used by Guiso et al. (2002) for the subjective distribution of future income. We assume that the subjective distribution is either simple triangular (i.e. symmetric around $(y_m + y_M)/2$, assuming $\pi = 0.5$), or split triangular ($\pi \neq 0.5$; see Fig. A.1. in the Appendix). Based on the elicited values of y_m , y_M (and of π if we assume a split triangular distribution) we compute the household-specific mean and standard deviation of the distribution of expected consumption one year ahead. The formulae of these statistics are reported in Appendix A.2.²

We set to missing values observations where y_m , y_M or π are missing, or if respondents chose the ‘do not know’ option. The original sample includes 5,034 observations in the three survey waves. Due to missing values, the estimation sample includes 3,271 household-level observations for the simple triangular distribution, and 3,167 observations for the split triangular distribution.

The survey also asked households' financial respondents to report directly expected change to their household spending one year ahead. Respondents were asked first to think about their household spending on all goods and services in the coming 12 months, and to report whether they thought it would be higher, about the same, or lower than their current spending. They were then asked to report expected change in spending as a percentage.³ The corresponding survey questions are provided in Appendix A.1.

As described in Appendix A.1 and A.2, we assume that the consumption distribution is simple triangular or split triangular. We then use information on each household's expected consumption growth, and minimum and maximum levels of

² We assume that y_m and y_M represent the actual minimum and maximum of the distribution. This is potentially a strong assumption. Dominitz and Manski (1997) use the percentage chance format to elicit the subjective income distribution, and show that individuals associate the “lowest possible” (and “highest possible”) with low (resp. high) probability.

³ U.S. surveys administered in the aftermath of the Great Recession include similar questions aimed at measuring expected changes in household consumption in the year ahead, or realized changes in consumption in the previous 12 months. A number of studies use this information to explore household consumption adjustments in response to the financial crisis. Christelis et al. (2015) and Hurd and Rohwedder (2010), use data from the 2009 Internet panel of the U.S. Health and Retirement Study, while Shapiro (2010) uses data from the 2009 Cognitive Economics Study.

consumption one year ahead. On this basis, it is straightforward to compute a household-specific expected variance, standard deviation, and square of expected consumption growth. The last is the term that appears in equation (5) and is used in the estimation.

The survey also asked for information that enables computation of the moments of the distribution of income one year ahead (similar to information to estimate future consumption). Specifically, households were asked to provide minimum and maximum values of annual family income gross of any taxes, for the coming 12 months, and to indicate the probability that income will be higher than the mid-point between the minimum and the maximum reported values. This allows us to compute expected income and expected income risk making the same distributional assumptions as for future consumption (either triangular or split triangular).

Figures 1 and 2 report the distribution of the expected minimum and maximum levels of consumption 12 months ahead. For each observation in the sample, the maximum is greater than the minimum. Figure 3 reports the distribution of the probability that the expected consumption is above the average of the expected minimum and maximum values. There is a prevalence of “50 percent” responses but also a sizable number of respondents reporting values larger or smaller than 50 percent. Notice that the question on this probability, which is arguably more difficult to answer, is not used in the regressions based on the simple triangular distribution.

Table 1 reports cross-sectional statistics of the central tendency and dispersion of the subjective distribution of consumption, assuming that the distribution is a simple (i.e. symmetric) triangular, and of the variables that will be used in the estimation (age, household size, marital status). At the median, the minimum expected level of consumption is 1,500 euro, while the maximum is 1,900 euro (the means are equal to 1,561 euro and 1,971 euro, respectively), and the median probability is 0.5 (average 0.48). Assuming that the distribution is simple triangular, we estimate that the sample median of expected consumption growth is zero (average 1.4 percent), while the median (mean) standard deviation of the distribution of expected consumption risk is about 4.2 (5.0) percent. Since forecasts in the Netherlands indicate that in 2014 consumption expanded slightly (by approximately 0.2 percent), consumption expectations are aligned with realizations.

Cross-sectional averages are useful to describe the subjective consumption distribution of a typical household but they hide important heterogeneity across households. Assuming that the distribution is simple triangular, Figure 4 plots the histogram of the standard deviations of the 3,271 household-specific distributions of future consumption growth. Figure 4 highlights considerable heterogeneity in the responses. For instance, for 25 percent of households the standard deviation is less than 2.1 percentage points, for another 50 percent it is between 2.1 and 7 percentage points, and for the top 25 percent it is more than 7 percentage points. The proportion of respondents for which the standard deviation is zero (i.e. they report no future consumption risk) is 14 percent.

The next step in the analysis is to relate consumption expectations to household characteristics. We are particularly interested in studying how the subjective expectations of consumption risk correlate with characteristics (e.g. age and occupation) which should influence consumption uncertainty. Figure 5 plots the median standard deviation of the expected consumption growth distribution by ten-year age bands. Figure 5 indicates that consumption risk declines during the life-cycle; the standard deviation of expected consumption growth falls by about 2 percentage points. This finding suggests that younger households perceive more uncertainty than older consumers, in line with the findings in Dominitz and Manski (1997) for the subjective distribution of income uncertainty. Notice that the age gradient might also capture cohort effects, so Figure 5 might signal that younger cohorts face higher uncertainty regardless of age. Unfortunately, our survey does not provide enough information to distinguish between these two explanations.

To obtain further insights on the reliability of our measures of subjective expectations, Table 2 reports associations of the standard deviations of expected consumption growth with age, expected income growth (constructed in similar fashion), self-employment, retirement status, union membership (as a further measure of income volatility), and household size. These associations are derived from robust regressions (using the M-estimator in Huber, 1973) of the standard deviation of expected consumption growth on each of the aforementioned variables.

Expected consumption risk is correlated strongly with expected income risk but at much less than one-to-one, showing that other factors besides income risk affect consumption risk. Another reason to expect that consumption and income risk are not perfectly correlated is that, under the permanent income hypothesis, consumption risk should reflect only permanent but not transitory income risk. Consumption uncertainty is also strongly correlated with self-employment. The direction of this association is as expected given that the self-employed typically face a higher than average income risk which should lead in turn, to higher consumption uncertainty.⁴

On the other hand, age and being retired are negatively associated with consumption risk which again, is as expected given the reduced income uncertainty associated with older age. Being a union member may imply more predictable wages, and thus lowers consumption risk. Finally, consumption risk increases with family size, possibly because larger families are exposed to larger expenditure shocks. Overall, the fact that these associations of consumption risk have the expected sign, are sizeable, and also are statistically significant suggests that the survey measure of subjective expected consumption risk is a good indicator of the actual consumption uncertainty faced by the households in our sample.

5. Empirical results

We estimate the relation between expected consumption risk and expected consumption growth, and thus augment equation (5) as follows:

$$E_{i,t}(g_{i,t+1}) = \alpha + \beta E_{i,t}(g_{i,t+1}^2) + \gamma X_{i,t} + v_{i,t+1} \quad (6)$$

where the vector X includes demographic variables, in particular: age and gender of the household financial respondent, whether (s)he has a partner, size of the household, and indicators of survey wave, and regional dummies. The demographic variables are

⁴ Dillon (2015), using data from the PSID and the CPS (Current Population Survey), and controlling for occupational mobility and endogenous labor supply, estimates that the self-employed face substantially higher lifetime earnings risk.

included in the specification to capture any additional sources of expected consumption growth heterogeneity.

Before presenting the econometric results, in Figure 6 we plot $E_{i,t}(g_{i,t+1})$ against binned values of $E_{i,t}(g_{i,t+1}^2)$.⁵ The two variables are strongly positively correlated, and the slope of the relation between the two is slightly more than 1, with an implied coefficient of relative prudence of slightly higher than 2. As we shall see, our estimation results are consistent with this descriptive evidence.

In order to reduce the influence of outliers, we winsorize both g and g^2 at the top and bottom 0.5 percent of the observations; i.e., we set the values of those observations equal to those at the 99.5th and 0.5th percentiles, respectively. We also use Huber-White robust standard errors.

We first estimate equation (6) by conventional OLS. Table 3 columns 1-3, report the OLS results of equation (6), using the simple triangular distribution for expected consumption risk. The estimated coefficient of consumption risk is 0.64 and highly statistically significant (p-value<.01), implying a prudence coefficient of about 1.28. The coefficients of age, female financial respondent, and household size are positive but imprecisely estimated.

In order to check the sensitivity of the OLS results to possible outliers we rely on robust regressions, using Huber's (1973) M-estimator. Results from this estimation are shown in Table 3 columns 4-6 and yield an estimated coefficient of expected consumption risk of 0.96 (p-value<.01). The derived estimate is larger than the corresponding OLS, and implies a prudence coefficient of about 2. As this estimate of prudence is robust to outliers, we consider it more reliable than the OLS estimate.

We performed the same estimation using the split instead of the simple triangular distribution. The results are presented in Table 3, columns 7-12; their size and statistical significance remain essentially the same, regardless of the estimation method used.

⁵ The bins are defined using the deciles of the distribution of the expected square of consumption growth.

6. Liquidity constraints

The Euler equation estimated in Section 5 is derived assuming perfect capital markets. However, the equation fails in the presence of liquidity constraints or myopic consumers. Let us consider a simple alternative model, where consumption equals income in each period. Then, expected consumption growth equals expected income growth in each period. This suggests that our estimates might be contaminated by the presence of some households that may not engage in precautionary saving. From an econometric point of view, this is an omitted variable problem which might bias the coefficient of interest, i.e. the sensitivity of expected consumption growth to expected consumption variability.

In order to address this, Table 4 presents the results of robust regressions that exclude from the estimation sample households which possibly are liquidity constrained, and thus, less likely to engage in precautionary saving. We distinguish liquidity constrained households based on three different measures. Expected consumption growth and consumption risk are calculated using the simple triangular distribution (results for the split triangular are similar).

Table 4 columns 1-3 present the results for the sample excluding households that reported being turned down for credit in the previous 12 months, or households that were discouraged from borrowing; i.e., households reporting that they did not apply for credit because they expected to be turned down. The relevant questions come from the 2014 and 2015 waves of the baseline DNB household survey, and thus, for households in our special purpose survey which were not part of the baseline survey, this information is not available. This resulted in 97 households (about 3.6% of the sample) that had been denied (or discouraged from applying for) credit.

The results in Table 4 columns 4-6 are based on the sample excluding liquidity constrained households based on whether when asked what they would do if they received a windfall sum equal to one month's income, they said they would spend at least 90% of it on non-durables and durables.⁶ The number of households classified as

⁶ Respondents were asked to indicate in percentages use of a windfall gain, according to four alternatives: i) save for future expenses; ii) repay debt; iii) spend within 12 months on durable goods (cars, home

liquidity constrained based on this measure is 83 (about 2.5% of the estimation sample), or 20 if purchase of durable goods is excluded.

Finally, we exclude from the sample households whose household head is unemployed, and those in the bottom quintile of the disposable income distribution (resulting in 671 households, or about 20.5% of the estimation sample being dropped; results are shown in Table 4 columns 7-9).

The coefficient of expected consumption risk is around 1 for all three sets of estimation results, confirming the baseline results in Table 3 for the whole sample. All the results discussed in the current Section are derived using the simple triangular distribution. When we use the split triangular distribution results are essentially identical. Thus, we can conclude that our baseline estimates are unlikely to be affected by the presence in our sample of liquidity constrained households.

Overall, the results from all our estimation methods and different specifications suggest that there is a positive and economically relevant association between expected consumption risk and expected consumption growth. This finding provides strong evidence of a precautionary saving motive among the households in our sample. Our estimates imply a coefficient of relative prudence of around 2, which is within the range of values the literature considers plausible. If one is willing to assume that the utility is isoelastic, then this value implies a coefficient of relative risk aversion as well as an intertemporal elasticity of substitution of around 1.

7. IV estimation

As already discussed, the use of elicited expectations in the estimation of (6) circumvents serious econometric issues affecting existing studies that base inference on consumption realizations. In particular, the use of expectations implies that the error term

improvements, furniture, jewelry, other durable goods) that they would otherwise would not have purchased or would have purchased later; iv) purchase within 12 months of non-durable goods and services that do not last in time (food, clothes, travel, vacation, etc.) and would normally not have been purchased. Similar questions were used in a different context by Shapiro and Slemrod (1995, 2003) and Jappelli and Pistaferri (2014).

v is not a forecasting error, as is usually the case in Euler equation estimates. Nonetheless, there is still the possibility that unobservable variables in the error term v (e.g. higher order terms of the Taylor expansion) are correlated with expected consumption risk. Hence, as a robustness check, we estimate equation (6) also using IV methods to take account of possible endogeneity problems and measurement error. Our instrument is expected income risk (constructed similar to expected consumption risk, as described in Section 4). This variable is used by Bertola et al. (2005) as an instrument for realized consumption volatility. It represents a good instrument choice in an Euler equation framework, given that it does not appear in equation (6) when expected consumption risk is included. Moreover, it is positively correlated with consumption risk.

The results of IV estimation are shown in Table 5, columns 1-3. The estimated effect of expected consumption risk on expected consumption growth is 0.89 and is strongly significant ($p\text{-value} < .01$). Moreover, it is very similar in magnitude to the robust regression estimate. The first-stage regression confirms that consumption risk is correlated positively with income risk. Nevertheless, the corresponding F-statistic is about 3.29, and thus below the rule of thumb threshold of 10 generally recommended to allow dependable inferences. The endogeneity of consumption risk can also be tested using a standard Hausman test. The test statistic has a $p\text{-value}$ equal to 0.27, indicating that the null hypothesis of exogeneity cannot be rejected. Thus, on this basis, IV estimation is unnecessary.

Given that the F-statistic from the first stage regression is rather weak, we provide additional evidence by using the IV procedure proposed by Lewbel (2012). This procedure generates additional instruments by interacting (after demeaning) all X variables in equation (6) with the residuals w from a regression of expected consumption risk on all the demeaned X variables.⁷ To provide some intuition for these additional generated instruments, let us consider an unobservable variable that is contained in w (e.g. a demand shock in a particular industry) which affects workers differentially depending on say, age and residential location. For example, this would apply to the case where older workers in that industry find it difficult to get a new job, thus making their

⁷ Additional instruments generated using the Lewbel (2012) method have been used in a number of empirical studies as an alternative to the standard IV approach, e.g. Emran and Hou (2013), and Chowdhury et al. (2014).

consumption uncertainty higher than for younger workers. In addition, residents in an area where the local economy is performing well may also experience reduced consumption uncertainty associated with this unobservable variable. Econometrically, as Lewbel (2012) shows, a necessary condition for the existence of this differential impact of the error term in the regression of consumption risk on the X variables is that it is heteroskedastic with respect to X . If this condition holds, then Lewbel (2012) shows that additional valid instruments can be generated, equal to the product of the demeaned regressors in X times the residual w .

Another important benefit of these generated instruments is that they produce additional over-identifying restrictions which allow the validity of the original instrument (expected income risk in our case) to be tested. Moreover, additional instruments generally lead to more efficient estimates. As already mentioned, the X variables include age and gender of the financial respondent, family size, a dummy for couples, and time and regional dummies, all of which we use to generate additional instruments.

The Breusch-Pagan test provides very strong evidence of heteroskedasticity among the residuals in the regression for expected consumption risk on X , as indicated by the p -value of the test shown in Table 5 column 4 (the value of the test statistic is about 652). The estimated coefficient of expected consumption risk in Lewbel's IV regression is equal to 0.99 and strongly significant. Notably, its magnitude is close to the magnitudes in the corresponding robust regression and standard IV estimates. Our results are also generally unchanged when income risk is excluded from the set of instruments and only generated instruments are used.

The F -statistic when using both our original instrument and the generated instruments is about 26, while it is equal to about 19 when using only the generated instruments. Hence, inferences based on the specification using the generated instruments should be more reliable than that based on the specification using only the original instrument.

The test of overidentifying restrictions can be performed for the additional generated instruments as well as these instruments combined with our original instrument. In both cases (see Table 5 column 4), the p -value strongly indicates that the null of instrument exogeneity cannot be rejected. These results suggest that both expected

income risk and the additional generated instruments are valid instruments. Finally, if we re-run the Hausman test for endogeneity of the consumption risk variable, the result (p-value: 0.27) again indicates that the null cannot be refuted. Thus, we can conclude that the robust regression estimates reported in Section 5 are reliable, as there is no indication that expected consumption risk (i.e. our covariate of interest) is affected by endogeneity problems.

As when using OLS and robust regression, we redid our estimation using the split instead of the simple triangular distribution. The results are presented in Table 5, columns 7-12, and we note that they are essentially identical to those obtained using the simple triangular distribution.

8. Conclusions

Our goal in this paper was to investigate the existence of a precautionary saving motive affecting household saving behavior. We estimated an Euler equation for consumption using subjective expectations of consumption, which conform better to the original Euler equation formulation than the ex-post consumption realizations used in the related literature so far. An additional motivation for using expectation data is that they allow researchers to circumvent problems related to inconsistency and endogeneity which affect ex-post realizations. To obtain expectation data, we designed a questionnaire that asked households about their expectations of future consumption distribution, and administered it to a representative sample of Dutch households.

Using these expectation data, we estimated the Euler equation to get an indication of the existence and strength of the precautionary saving motive (through the magnitude of the associated prudence coefficient). We used a variety of estimation methods, namely OLS, robust regression, and IV, and obtained consistent results pointing clearly to the existence of a strong precautionary motive in the saving behavior of the households in our sample. The estimated relative prudence coefficient is around 2, in line with the existing literature.

Since the expectation data are correlated with observable household characteristics in a manner that conforms to theory and intuition, these data are likely to provide a good measure of the underlying uncertainty experienced by households. This suggests the advantages of asking about households' expectations in order to investigate this uncertainty. More generally, the responses to such questions are valuable for estimating economic relationships in which households' expectations play a key role. Thus, we recommend their inclusion in household surveys.

References

- Abowd, J. M., and D. Card (1989), “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 57, 411-45.
- Altug S. and Miller R. (1990), “Household Choices in Equilibrium”, *Econometrica*, 58, 543—70.
- Bertola, G., Guiso, L., and L. Pistaferri (2005), “Uncertainty and Consumer Durables Adjustment,” *Review of Economic Studies*, 72, 973-1007.
- Cagetti, M. (2003), “Wealth Accumulation over the Life Cycle and Precautionary Savings,” *Journal of Business and Economic Statistics*, 21, 339-353.
- Carroll, C. (2001), “Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation).” *Advances in Macroeconomics*, 1(1), 1003-1003.
- Chamberlain, G. (1984), “Panel Data,” in *Handbook of Econometrics*, vol. 2, Z. Griliches and M. D. Intriligator (eds.). Amsterdam: Elsevier Science.
- Chowdhury, M., Bhattacharya, P.S., Mallick, D., and M.A. Ulubasoglou (2014), “An Empirical Inquiry into the Role of Sectoral Diversification in Exchange Rate Regime Choice,” *European Economic Review*, 67, 210-227.
- Christelis, D., Georgarakos, D., and T. Jappelli (2015), “Wealth Shocks, Unemployment Shocks and Consumption in the Wake of the Great Recession,” *Journal of Monetary Economics*, 72, 21-41.
- Crump, R. K., Eusepi, S., Tambalotti, A., and G. Topa (2015), “Subjective Intertemporal Substitution,” Federal Reserve Bank of New York, Staff Report No. 734.
- Deaton, A. S. (1991), “Saving and Liquidity Constraints,” *Econometrica*, 59, 1221-48.
- Dillon, E. (2015), “Risk and Return Tradeoffs in Lifetime Earnings,” mimeo, Arizona State University.
- Dominitz, J. and C. F. Manski. (1997), “Using Expectations Data to Study Subjective Income Expectations,” *Journal of the American Statistical Association*, 92, 855-67.

- Dynan, K. (1993), "How Prudent are Consumers?," *Journal of Political Economy* 101, 1104-13.
- Emran S., and Z. Hou (2013), "Access to Markets and Rural Poverty: Evidence from Household Consumption in China," *Review of Economics and Statistics*, 95, 682-697.
- Gourinchas, P., and J. A. Parker (2002), "Consumption over the Life Cycle," *Econometrica*, 70, 47-89.
- Guiso, L., Jappelli, T., and L. Pistaferri (2002), "An Empirical Analysis of Earnings and Employment Risk, *Journal of Business and Economic Statistics*, 20, 241-253
- Guiso, L., Jappelli, T., and M. Padula (2013), "Pension Wealth Uncertainty," *Journal of Risk and Insurance*, 80, 1057-1085.
- Hansen, L.P., and K. J. Singleton (1983), "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy*, 91, 249-65.
- Hayashi, F. (1987), "Tests for Liquidity Constraints: A Critical Survey and Some New Observations," in T.F. Bewley (ed.) *Advances in Econometrics: Fifth World Congress*, Vol II, 91-120.
- Huber, P. (1973), "Robust Regression: Asymptotics, Conjectures and Monte Carlo," *The Annals of Statistics*, 1, 799-821.
- Hurd, M., Rohwedder, S. (2010), "The Effects of the Economic Crisis on the Older Population," Michigan Retirement Research Center Working Paper No. 2010-231.
- Jappelli, T., and L. Pistaferri (2000), "Using Subjective Income Expectations to test for Excess Sensitivity of Consumption to Predicted Income Growth," *European Economic Review*, 44, 337-358
- Jappelli, T., and L. Pistaferri (2010), "The consumption response to income changes," *Annual Review of Economics*, 2, 479-506
- Jappelli, T. and L. Pistaferri (2014), "Fiscal Policy and MPC Heterogeneity," *American Economic Journal: Macroeconomics*, 6(4), 107-136.

- Kimball, M. S. (1990), "Precautionary Saving in the Small and in the Large," *Econometrica*, 58, 53-73.
- Lewbel, A. (2012), "Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models," *Journal of Business and Economic Statistics*, 30, 67-80.
- Low, H., and L. Pistaferri (2015), "Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off," *American Economic Review*, 105(10), 2986-3029.
- Low, H., Meghir, C., and L. Pistaferri (2010), "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review*, 100, 1432-67.
- Manski, Charles F. (2004), "Measuring expectations," *Econometrica* 72, 1329-76.
- Palumbo, M. G. (1999), "Uncertain medical expenses and precautionary saving near the end of the life cycle," *Review of Economic Studies*, 66, 395-421.
- Shapiro, M. D., and J. Slemrod (1995), "Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding," *American Economic Review*, 85, 274-283.
- Shapiro, M. D., and J. Slemrod (2003), "Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys," in *Tax Policy and the Economy*. Vol. 17, ed. James M. Poterba, Cambridge, MA: MIT Press.
- Shapiro, M. D. (2010), "The Effects of the Financial Crisis on the Well-Being of Older Americans: Evidence from the Cognitive Economics Study," Michigan Retirement Research Center Working Paper No. 2010-228.
- Skinner, J. (1988), "Risky Income, Life Cycle Consumption, and Precautionary Savings," *Journal of Monetary Economics*, 22, 237-255.
- Voena, A. (2015), "Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?," *American Economic Review*, 105(8), 2295-2332.

Table 1. Descriptive statistics

Variable	Mean	Median	Std. Deviation
Minimum Expected Consumption Level	1,561.3	1,500.0	905.0
Maximum Expected Consumption Level	1,970.7	1,900.0	1,117.5
Probability that the Expected Consumption Level is above the Average of the Expected Minimum and Maximum Values	0.477	0.500	0.233
Expected Consumption Growth	0.014	0.000	0.089
Standard Deviation of Expected Consumption Growth	0.050	0.042	0.041
Standard Deviation of Expected Income Growth	0.028	0.015	0.050
Age	55.8	58.0	15.5
Female Householder	0.38	0.00	0.49
Household Size	2.23	2.00	1.20
Has a Partner	0.66	1.00	0.47
Number of Observations	3,271		

Table 2. Correlations of the standard deviation of expected consumption growth with other magnitudes

Variable	(1)	(2)	(3)	(4)
	Coeff.	Std. Error	p value	Standardized Change in the Std. Deviation of Consumption Growth due to a Change in the Variable
Standard Deviation of Expected Income Growth	0.199	0.034	0.000	0.238
Age	0.000	0.000	0.000	-0.009
Self-employed	0.008	0.003	0.018	0.191
Retired	-0.012	0.001	0.000	-0.285
Is a Union Member	-0.006	0.003	0.026	-0.142
Household Size	0.002	0.001	0.001	0.044

Notes: Columns 1-3 report results from robust regressions in which the dependent variable is the standard deviation of consumption growth and the only regressor is the variable shown in each line. Column 4 reports the implied change in the standard deviation of expected consumption growth (in terms of its own standard deviation) when the regressors change as follows: i) the standard deviation of expected income growth increases by one standard deviation; ii) the remaining variables increase by one unit..

Table 3. Euler equation estimates, OLS and robust regression estimation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Variable	Simple Triangular Distribution						Split Triangular Distribution					
	OLS			Robust Regression			OLS			Robust Regression		
	Coeff.	Std. Error	p-value	Coeff.	Std. Error	p-value	Coeff.	Std. Error	p-value	Coeff.	Std. Error	p-value
Consumption Uncertainty	0.640	0.122	0.000	0.963	0.008	0.000	0.642	0.123	0.000	0.966	0.008	0.000
Age	0.000	0.000	0.322	0.000	0.000	0.953	0.000	0.000	0.364	0.000	0.000	0.889
Female Householder	0.002	0.003	0.473	0.000	0.001	0.911	0.002	0.003	0.557	0.000	0.001	0.964
Household Size	0.001	0.002	0.595	0.000	0.000	0.453	0.001	0.002	0.493	0.001	0.000	0.257
Couple	-0.004	0.004	0.330	-0.001	0.001	0.599	-0.005	0.004	0.294	-0.001	0.001	0.450
Constant	0.014	0.009	0.099	0.006	0.003	0.014	0.015	0.009	0.095	0.006	0.003	0.011
Number of Observations	3,271			3,271			3,167			3,167		

Notes: In addition to the variables shown, all specifications include regional and survey wave dummies.

Table 4. Euler equation estimates, excluding liquidity-constrained households

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Variable	Has not been denied Credit			Marginal Propensity to Consume < 0.90			Household Income above the 20 th Quantile and no Unemployment		
	Coeff.	Std. Error	p-value	Coeff.	Std. Error	p-value	Coeff.	Std. Error	p-value
Consumption Uncertainty	0.954	0.013	0.000	0.962	0.009	0.000	0.957	0.012	0.000
Age	0.000	0.000	0.608	0.000	0.000	0.857	0.000	0.000	0.615
Female Householder	0.001	0.001	0.347	0.000	0.001	0.912	-0.001	0.001	0.279
Household Size	0.000	0.001	0.647	0.000	0.000	0.335	0.000	0.000	0.345
Couple	0.000	0.001	0.742	0.000	0.001	0.692	0.001	0.001	0.288
Constant	0.007	0.003	0.016	0.007	0.003	0.011	0.006	0.003	0.034
Number of Observations	2,642			3,188			2,600		

Notes: In addition to the variables shown, all specifications include regional and survey wave dummies. The variable denoting consumption uncertainty is derived using the simple triangular distribution.

Table 5. Euler equation estimates, IV estimation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Variable	Simple Triangular Distribution						Split Triangular Distribution					
	Standard IV			IV with Generated Instruments			Standard IV			IV with Generated Instruments		
	Coeff.	Std. Error	p-value	Coeff.	Std. Error	p-value	Coeff.	Std. Error	p-value	Coeff.	Std. Error	p-value
Consumption Uncertainty	0.888	0.165	0.000	0.994	0.129	0.000	0.833	0.181	0.000	0.979	0.144	0.000
Age	0.000	0.000	0.704	0.000	0.000	0.807	0.000	0.000	0.812	0.000	0.000	0.950
Female Householder	0.000	0.004	0.963	0.000	0.004	0.991	0.000	0.004	0.998	0.000	0.004	0.955
Household Size	0.001	0.002	0.723	0.000	0.002	0.783	0.001	0.002	0.551	0.001	0.002	0.652
Couple	-0.002	0.005	0.642	-0.002	0.005	0.728	-0.003	0.005	0.546	-0.002	0.005	0.664
Constant	0.003	0.010	0.741	0.001	0.008	0.950	0.003	0.010	0.742	0.000	0.008	0.976
Number of Observations	2,890			2,890			2,791			2,791		
F-test	3.290			26.147			3.709			21.782		
F-test of Generated Instruments	-..-			19.244			-..-			14.517		
Breusch-Pagan Test for Heteroskedasticity - p-value	-..-			0.000			-..-			0.000		
Test of Overidentifying Restrictions - p-value	-..-			0.381			-..-			0.280		
Test of Overidentifying Restrictions of Generated Instruments - p-value	-..-			0.460			-..-			0.419		
Test of Endogeneity of the Treatment Variable (Consumption Uncertainty) - p-value	0.273			0.274			0.343			0.472		

Notes: In addition to the variables shown, all specifications include regional and survey wave dummies.

Fig. 1. Histogram of the minimum expected consumption level




Fig. 2. Histogram of the maximum expected consumption level




Fig. 3. Histogram of the probability that expected consumption is above the average expected minimum and maximum values




Fig. 4. Histogram of the standard deviation of expected consumption growth




Fig. 5. Standard deviation of expected consumption growth, by age group





Fig. 6. Average expected consumption growth, by levels of the expected square of consumption growth



Appendix

A.1. Questions on future spending.

Respondents to the household survey were asked about future levels of spending as follows:

Thinking ahead about your household spending during the next 12 months, what do you expect to be the value of such future spending in a typical month? Please provide the future monthly expenditure.

- (a) Please give the minimum value: €... (y_m)
- (b) Please give the maximum value: €... (y_M)
- (c) What is the probability that the household spending value is greater than X? (where X is automatically computed as $(y_m + y_M)/2$ and appears to the respondent's screen)

0	10	20	30	40	50	60	70	80	90	100
Absolutely no chance that household spending will be greater than X								Absolutely certain that household spending will be greater than X		

Questions on future spending were addressed to the financial respondent of each household. In particular, there was an introductory screen instructing financial respondents as follows:

‘The next questions are about your household's spending on all goods and services. Please count the spending of everyone who is living with you.’

Then these respondents were asked the following question:

‘Thinking ahead to 12 months from now, how do you expect your household spending on all goods and services at that time to compare to your spending today?’

The possible answers were:

- a) Higher than now
- b) About the same
- c) Lower than now
- d) Do not know

Respondents who anticipated an increase or decrease in their household spending in a year's time were then asked to estimate this change as a percentage:

‘How much (percentage-wise) do you expect that your household spending on all goods and services is [higher/lower] 12 months from now?’

Respondents could answer from 0 to 100 in multiples of 10, or report 'don't know' in which case they were directed to a set of options as follows:

Less than 5%

5-10%

11-15%

16-20%

21-25%

26-30%

More than 30%

Don't Know

Respondents who indicated that they expected their spending to be ‘about the same’ as currently, were asked to be more precise:

‘You have indicated that you expect that your household spending 12 months from now will be about the same as now. This could mean that the change equals zero percent of

that the percent change is small. Please estimate using the categories listed below what situation best describes your situation?’

The possible options vary by 1% starting from -10% up to +10%, and include an ‘exactly the same as now’ and a ‘do not know’ option.

In order to get a measure of the square of expected consumption growth we divide the mean expected consumption level (derived through the triangular distribution) by 1 plus the reported expected growth rate to obtain an estimate of the current level of consumption. We then divide the reported expected consumption extrema (i.e. y_m , y_M) and their average by this estimated current consumption level, and deduce the distribution of expected consumption growth using the procedures followed for the distribution of the expected consumption level. Given that $E(y^2) = Var(y) + E(y)^2$, once the variance of expected consumption growth is obtained for each household, it can be added to the square of the reported expected consumption growth to obtain the expected square of consumption growth, i.e. our variable of interest.

A.2. The subjective distribution of consumption

Let $f(y)$ denote the distribution of future consumption for each individual. The survey provides information on the support of the distribution $[y_m, y_M]$ and on the probability mass to the right of the mid-point of the support $\pi = Prob(y > (y_m + y_M)/2)$. Knowing the support of the distribution, the expected value and variance of y can be expressed as:

$$E(y) = \int_{y_m}^{y_M} yf(y)dy \quad (A.1)$$

$$Var(y) = \int_{y_m}^{y_M} y^2 f(y)dy - \left(\int_{y_m}^{y_M} yf(y)dy \right)^2 \quad (A.2)$$

We assume that the distribution $f(y)$ is triangular over each of the two intervals $[y_m, (y_m + y_M)/2]$ and $[(y_m + y_M)/2, y_M]$, as shown in Figure A.1. If $\pi = 0.5$ the

distribution collapses to a simple triangular distribution over the interval $[y_m, y_M]$. Note that $E(y)$ and $Var(y)$ depend only on the three known parameters (y_m , y_M , and π). The triangular distribution is a plausible description of the probability distribution of consumption growth because outcomes farther from the mid-point receive less weight.

Figure A1
The split triangular distribution

