Consumption Uncertainty and Precautionary Saving [#]

Dimitris Christelis
University of Naples Federico II, CSEF, CFS, CEPAR and Netspar

Dimitris Georgarakos
Deutsche Bundesbank, CFS and University of Leicester

Tullio Jappelli
University of Naples Federico II, CSEF and CEPR

Maarten van Rooij
De Nederlandsche Bank and Netspar

30 November 2015

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, the self-employed and also correlates positively with income risk. We insert these subjective expectations (instead of consumption realizations, as in the existing literature) in an Euler equation for consumption, and estimate the degree of prudence by associating expected consumption risk with expected consumption growth. Both robust OLS and IV estimates indicate a coefficient of relative prudence of around two.

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

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

[#] Tullio Jappelli acknowledges financial support from the Italian Ministry of University and Research (PRIN project) and the FP7[th] EU Framework Program (MOPACT project). We thank Laura Bartiloro, Tom Crossley, Luigi Pistaferri, Ernesto Villanueva, Guglielmo Weber, seminar participants at the 2015 NBER Summer Institute, 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.

1. Introduction

The effect of uncertainty on consumer behavior is a long-standing topic of research in the household saving literature (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 will depend 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 different approaches, like structural models or quasi-experimental approaches. Structural models require far more 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 forecast error, this substitution of expectations for realizations makes the observed consumption growth variability almost surely correlated with the error term of the Euler equation. 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 forecast error) enters in the error term of the estimated Euler equation. As discussed by Hayashi (1987), this expectational error should converge to zero as the time dimension of the data increases, but the same is not true in a short panel. This problem, as Chamberlain (1984) points out, is a serious one because it leads to inconsistent estimates in short panels, such as those that typically contain information on consumption.

The first contribution of the paper is the construction of measures of expected consumption growth and expected consumption risk. These measures are derived from responses of survey participants to questions on expectations about their future consumption. The survey data that 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 that are deduced from these questions take values that are household-specific. Hence, one can correlate them to observable household characteristics. We find that the measures of expected consumption risk are associated with such characteristics in the direction suggested by economic intuition. For example, expected consumption risk is higher for the young and the self-employed. Furthermore, 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. All in all, the survey questions 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 use, for the first time in the literature (to the best of our knowledge), of these expectation measures in order to estimate an Euler equation for consumption. The use of expectations-based variables instead of observed magnitudes eliminates the aforementioned problem of having an expectational error in the disturbance term. Furthermore, expectational variables address the issues regarding the endogeneity of the variable denoting observed consumption growth variability, which are discussed in Carroll (2001) and Bertola et al. (2005).

While avoiding the endogeneity problems present in the existing literature, our estimates still need to account for the possibility that expected consumption risk is correlated with the error term of the Euler equation. In order to address this problem, we use 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 clearly correlated with consumption risk but not with expected consumption growth. We also check the validity of the instruments using an instrumental variable (IV) method recently 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 about two. If the utility function is isoelastic, then this estimate in turn implies that the coefficient of relative risk aversion is about one. These results also hold when we exclude from the sample households that are likely to be liquidity constrained, for which the Euler equation does not apply. The IV estimates are similar to the robust OLS ones. Importantly, when testing for the endogeneity of the expected consumption risk variable, we cannot reject the null hypothesis of no endogeneity. All in all, the results for the strength of the precautionary saving motive and for the measures of household prudence and of the 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 the main empirical tests, with special focus on approaches relying on the Euler equation. Section 4 presents the survey data and Section 5 the empirical results. Section 6 extends the analysis to the presence of liquidity constraints. Section 7 concludes.

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 to the optimal consumption rule along the lines suggested by Blanchard and Mankiw (1988). 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:

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

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

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

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:

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

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 also equal 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 intuitive reason 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] A test of the hypothesis that consumption risk does not affect consumption growth is therefore also a test of the validity of the certainty equivalence model.

---

[1] If one is willing to make specific assumptions about preferences and the probability distribution of future consumption growth, one obtains an explicit solution for the expected growth rate of consumption (Hansen and Singleton, 1991). For instance, assuming an isoelastic utility function, constant interest rate and that the conditional distribution of consumption growth is normally distributed, one obtains the closed form solution $E_t \Delta lnc_{t+1} = \gamma^{-1}(r-\delta) + (\gamma/2) Var_t \Delta lnc_{t+1}$.

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

As we shall see in Section 3, 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 this framework, some models distinguish between movements in hours and in wages. For instance, Abowd and Card (1989) decompose fluctuations in earnings, which is often the standard measure of risk, into exogenous fluctuations in wages and endogenous choice of hours. It is only the first type of fluctuations that represent genuine risk, even though variation in hours has welfare consequences, as 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 the 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 may reflect also uncertainty about other random variables, and may not be related to income risk alone. While for young individuals income risk might be the most important source of consumption risk, at other stages of the life-cycle it might not even be the most important one. Besides income risk, people face a number of other uninsurable or partially uninsurable risks, which can affect intertemporal consumption decisions. Some of these risks (and their associated costs) have received prominent attention in the literature. They include the risk of future liquidity constraints, shocks to asset prices (including house prices), risk of medical and other unexpected expenditures, and the risk of family dissolution (see, e.g., Palumbo, 1999; Voena, 2015).

3. Empirical tests of precautionary saving

Tests of the importance of precautionary saving follow several research strategies. 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 drawn from actual earnings are not easy to compute even with long panel data and may reflect in part 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 a great deal across studies and on net tends to be on the small side (Jappelli and Pistaferri, 2010). Furthermore, this approach 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 path of consumption and wealth in models with precautionary saving, matching simulated data to the observed distribution of wealth and consumption. 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 the actual and the predicted life-cycle profile of consumption. Setting the real interest rate at three percent, they estimate a rate of time preference of approximately four percent and an elasticity of intertemporal substitution of about 0.5, corresponding to a coefficient of relative risk aversion of about two. With an isoelastic utility function the implied coefficient of relative prudence is about three. Cagetti (2003) estimates the same preference parameters as in Gourinchas and Parker (2002) by matching simulated and actual median wealth profiles over the life cycle using U.S. data from the PSID and the Survey of Consumer Finances (SCF). He finds higher estimates of the rate of time preference, a higher coefficient of relative risk aversion (around four for the high school sample), and an implied coefficient of relative prudence of around five. Structural estimations deliver estimates of the parameters of the utility function, but require specifying not only the utility function, but also the budget constraint, the sources of risks, and the income process.

An alternative, direct strategy to estimate the coefficient of relative prudence is to use survey data and measure expected consumption growth and expected consumption risk in equation (3) based on respondents' own assessments of these variables. This is the approach that we follow in this paper.

From an empirical point of view, the main problem of estimating the Euler equation is that expected consumption growth and expected consumption risk are generally not observable. If one could measure the expectation-related terms in equation (3), then it would be possible to estimate the equation by OLS and identify the coefficient of the term related to expected consumption risk, which is proportional to relative prudence. The first paper that attempts to identify this coefficient in an Euler equation framework is Dynan (1993), who, however, does not have in her data any information related to expectations about consumption. Hence, she replaces expectations with their realized counterparts. In this case, indexing households by i (i = 1, …, N) , equation (3) can be written in a regression framework as:

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

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 errors, and possibly heterogeneity in preferences. In particular, the substitution of realized consumption changes for their expectations implies that $\varepsilon_{i,t+1}$ in equation (6) 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 instrumental variables approach applied on panel data drawn from the CEX. The set of instruments includes education and occupation, assuming 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.

Dynan's approach has been refined by Bertola et al. (2005). Using Italian data from the Survey of Household Income and Wealth, they use the subjective variance of income one year ahead as an instrument for $\left((c_{i,t+1} - c_{i,t})/c_{i,t}\right)^2$. As they point out, subjective income risk should be a valid instrument, as (3) implies that, conditional on expected consumption risk, income risk has no direct effect on consumption growth. In other words, expected consumption risk is a sufficient statistic for expected consumption growth. In addition, subjective expectations of income risk result in a household-specific subjective distribution of income uncertainty, and therefore in a measure of income risk that takes different values across households. 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 around two and precisely estimated, thus providing evidence in support of the precautionary saving model.

In this paper, and in contrast to Dynan (1993) and Bertola et al. (2005), we estimate directly equation (3) using subjective expectations of future consumption, rather than relying on realized consumption magnitudes. Let's rewrite equation (3) as:

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

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 (6) has two advantages with respect to previous tests of precautionary saving. First and most importantly, the error term of equation (6), $\varepsilon_{i,t+1}$, includes the expectational error of the Euler equation, while the error term of equation (5), $v_{i,t+1}$, by construction does not. Hence, it is not correlated with expected consumption risk.

A second, related issue is that one can estimate equation (5) even with a cross-section or with a short panel, exploiting the cross-sectional variability in expectations of the consumption distribution. The literature shows that Euler equation estimates in short panels 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 (6) includes a forecast error. The life-cycle and permanent income models imply that the expectation of the forecast error, conditional on any information available at t should be zero over a long horizon, i.e. the error should not exhibit systematic patterns, if the model is correct. Following this logic, the empirical equivalent of $E_{i,t}\left(\varepsilon_{i,t+1}\right)$ in (6), which includes the expectation of the forecast error, is a household-level average taken over T periods. Importantly, $T \longrightarrow \infty$ is needed to ensure the consistency of the estimated parameters of the Euler equation. However, panel surveys with information on consumption are typically short, and hence researchers often proceed under the assumption that consistency is achieved with $N \longrightarrow \infty$, i.e. assuming that forecast errors average out to zero in the cross-section.

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

As already discussed, given that we use expectations of consumption instead of realizations, the assumption that $v_{i,t+1}$ has a zero conditional expectation in the cross-section is quite reasonable, implying that one can use OLS to estimate equation (5). Still, it might be the case that that there are unobservable variables that are correlated with expected consumption risk as well as with expected consumption growth, or that the error term contains higher order terms 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, as discussed in the Introduction (see also Section 5 below).

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 every 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, households are frequently asked to participate during the course of a year in special purpose surveys.

We have designed such a survey that contains questions aiming to measure individual uncertainty about future consumption and income as well as expected household consumption growth. We first asked these questions to participants in the Internet panel in June 2014. In January 2015 and in June 2015 we asked again the same set of questions in order to see if there is any strong seasonal pattern in responses and to increase the sample size used in our analysis. In each household, we target the financial respondent, i.e. the person who is responsible for the household finances.

In a recent paper that is related to ours, 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 of 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 EIS is around 0.8.

To elicit the distribution of expected consumption we follow a similar procedure as Guiso et al. (2002; 2013), who are interested in approximating the subjective distribution of future income and pension replacement rate, respectively. In particular, we first ask respondents to report the minimum (y_m) and the maximum (y_M) values of next year's consumption in a typical month, and subsequently 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. $\pi = Prob(y \geq (y_m + y_M)/2)$. We reproduce the wording of the questions in Appendix A.1.

To estimate the moments of the subjective distribution of future consumption we rely on the same 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 $(y_m + y_M)/2$ by 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 when a split triangular distribution is assumed) 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 for which y_m, y_M or π are missing or respondents choose a '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.

Moreover, the survey asks household financial respondents to report directly the expected change to their household spending one year ahead. Respondents are first asked to think about their household spending on all goods and services in the next twelve months and report whether they think that will be higher, about the same, or lower than their current spending. Subsequently, they are asked to report in percentages the expected change in spending.[3] Details on the relevant questions survey can be found 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. Subsequently, we use information for each household on expected consumption growth, and on the minimum and maximum level of consumption one year ahead. Then, it is straightforward to compute the household-specific expected variance, standard deviation, and square of expected

[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 (respectively, high) probability.

[3] Similar questions aiming to measure expected changes in household consumption in the year ahead or realized changes in consumption in the last year have been asked in US surveys conducted in the aftermath of the Great Recession. A number of studies has used 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 US Health and Retirement Study, while Shapiro (2009), uses data from the 2009 Cognitive Economics Study.

consumption growth. The latter is the term that appears in equation (5) and is used in the estimation.

The survey also asks for information that allows one to compute moments of the distribution of income one year ahead in the same way as for future consumption. In particular, households are asked about the minimum and maximum values of annual family income, gross of any taxes, during the next twelve months, and 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 (π). 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 that use 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 and 1,971 euro, respectively), and the average probability is 0.5 (the average is 0.48). Assuming that the distribution is simple triangular, we estimate that the sample median of expected consumption growth is zero (the average is 1.4 percent), while the median (mean) standard deviation of the distribution of expected consumption risk is about 4.2 (5) percent. Since forecasts in the Netherlands indicated that in 2014 consumption expanded slightly (approximately by 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 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. The figure 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% between 2.1 and 7 percentage points, and for the top 25% percent 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%.

The next step of 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) that 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, as 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 of Dominitz and Manski (2006) for the subjective distribution of income uncertainty. Notice that the age gradient may capture also cohort effects, so that Figure 5 may signal that younger cohorts face higher uncertainty, regardless of age. Unfortunately, our survey does not contain enough information to distinguish between these two different explanations.

To provide further insights on the reliability of our measures of subjective expectations, Table 2 reports associations of the standard deviation of expected consumption growth with age, the standard deviation of 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 strongly correlates with expected income risk, but the correlation is far less than one-for-one, showing that other factors affect consumption risk, besides income risk. A further 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 being self-employed. The direction of this latter 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 is again as expected, given the reduced income uncertainty associated with older age. Being a union member likely implies 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. All in all, the fact that these associations of consumption risk have the expected sign, are sizeable and also 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:

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

where the vector X includes demographic variables, in particular: age and gender of the household financial respondent, whether (s)he has a partner, the size of the household, as well indicators of the survey wave and regional dummies. The demographic variables are included in the specification to capture any additional sources of expected consumption growth heterogeneity.

[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.

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 a bit larger than one, with an implied coefficient of relative prudence of a bit larger than two. 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% of observations; that is, we set the values of those observations equal to those at the 99.5th and 0.5th percentile, respectively. We also use Huber-White robust standard errors.

We estimate equation (6) first by conventional OLS. In Table 3, columns 1-3, we report OLS results of equation (6), using the simple triangular distribution for expected consumption risk. The estimated coefficient of consumption risk is 0.64 with a standard error of 0.122, implying a prudence coefficient of about 1.28, and highly statistically significant (p-value<.01). The coefficients of age, female householder 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 columns 4-6 in Table 3 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 two. As this estimate of prudence is robust to outliers, we consider it as more reliable than the OLS one.

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 υ is not a forecast error, as is usually the case in Euler equation estimates. Nonetheless, there is still the possibility that unobservable variables in the error term υ are correlated with expected consumption risk. Hence, we estimate equation (6) also using IV methods, so as to take into account possible endogeneity problems and measurement error. We use as an instrument expected income risk (constructed in similar fashion to expected consumption risk, as described in Section 4). This is the same variable used by Bertola et

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

al. (2005) as an instrument for realized consumption volatility. It represents a good instrument choice given that it does not appear in the Euler equation (6) when expected consumption risk is included, and given that it is positively correlated with consumption risk.

Results from IV estimation are shown in columns 7-9 of Table 3. The estimated effect of expected consumption risk on expected consumption growth is 0.89 and strongly significant (p-value<.01). Moreover, it is very similar in magnitude to the estimate from the robust regression. The first-stage regression confirms that consumption risk correlates positively with income risk. Nevertheless, the corresponding F-statistic is about 3.29, and thus below the rule of thumb threshold of 10 that is generally recommended in order to be able to make dependable inferences. One can also test the endogeneity of consumption risk using a standard Hausman test. The test statistic has a p-value equal to 0.27, indicating that the null hypothesis of exogeneity cannot be rejected. Thus, IV estimation is, on this ground, not needed.

Given that the F-statistic from the first stage regression is rather weak, we provide additional evidence by using an IV procedure proposed by Lewbel (2012). In this procedure, additional instruments are generated 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.[6] To provide some intuition for these additional generated instruments, let's consider an unobservable variable that is contained in w (e.g. a demand shock in a particular industry), that affects workers differentially, depending on, say, age and residential location. This could be the case when older workers in that industry may find it difficult to find a new job, and thus their consumption uncertainty may be higher than that of younger workers. In addition, those living in an area in which the local economy performs well may also experience reduced consumption uncertainty. 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 be heteroskedastic with respect to X. If this condition holds, then

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

Lewbel (2012) shows that one can generate additional valid instruments, equal to the product of the demeaned regressors in *X* with the residual *w*.

Another important benefit of these generated instruments is that they produce additional over-identifying restrictions, which allow one to test for the validity of the original instrument (expected income risk in our case). Furthermore, more 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, as well as time and regional dummies, and we use all of them to form the additional instruments.

A Breusch-Pagan test gives very strong evidence of heteroskedasticity in the residuals of the regression of expected consumption risk on *X*, as shown in column 10 of Table 3 (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. Importantly, its magnitude is remarkably close to those of the corresponding robust regression and standard IV estimates. It is also the case that our results remain essentially unchanged when income risk is excluded from the set of instruments and only generated instruments are used.

The combined F-statistic of our original instrument and the generated instruments is about 26, while the F-statistic of only the generated instruments is about 19. Hence, inference derived from the specification using the generated instruments should be more reliable than the one derived from the specification using only the original instrument.

One can perform a test of overidentifying restrictions when using only the additional generated instruments, as well as when using both these instruments and our original instrument. In both cases, as can be seen in Table 3, the p-value of the test clearly indicates the null of exogeneity cannot be rejected. These results suggest that both expected income risk and the additional generated instruments are valid instruments. Finally when we perform again a Hausman test for the endogeneity of the consumption risk variable, the result (p-value: 0.27) once more indicates that the null cannot be refuted. Hence, once again we conclude that one can rely on the robust regression estimates, as there is no indication that expected consumption risk (our regressor of interest) is affected by endogeneity problems.

Furthermore, we perform the same estimation using the split triangular distribution instead of the simple one. Our results are shown in Table 4, and we note that they are essentially unchanged in both size and statistical significance, regardless of the estimation method used.

6. Liquidity constraints

The Euler equation that we estimate in Section 5 is derived assuming perfect capital markets. However, in the presence of liquidity constraints or myopic consumers, the equation fails. Consider, for instance, a simple alternative model, where consumption equals income in each period. In that model, expected consumption growth will equal expected income growth in each period. This means that our estimates might be contaminated by the presence of some households that do not necessarily engage in precautionary saving. From an econometric point of view, we face an omitted variable problem, which might bias the coefficient of interest, i.e. the sensitivity of expected consumption growth with respect to expected consumption variability.

In order to address this issue, we present in Table 5 results from robust regressions that exclude from the estimation sample households that are possibly 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 using the split triangular are similar).

Results shown in columns 1-3 are derived using a sample from which we drop households that report that they have been turned down for credit in the past 12 months, or that were discouraged from borrowing; that is, they reported that they did not apply for credit because they thought that they would 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 that do not appear in the baseline survey this information will not be available. After excluding these households, we find 97 households (about 3.6% of the sample) that have been denied (or discouraged from applying for) credit.

For deriving the results shown in columns 4-6, we exclude those who are liquidity constrained according to another measure: when asked what they would do if they were given a windfall sum equal to one month's income, they answer that they would spend at least 90% of it on non-durables and durables.[7] The number of households designated as liquidity constrained using this definition is equal to 83 (about 2.5% of the estimation sample), and to 20 if one excludes the purchase of durable goods from spending.

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

In each of the three sets of estimation results, the coefficient of expected consumption risk is around one, confirming the baseline results of Table 3 for the whole sample. We thus conclude that our baseline estimates are unlikely to be affected by the presence of liquidity constrained households in our sample.

All in all, 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 for a precautionary saving motive among the households in our sample. Our estimates imply a coefficient of relative prudence of around two, which lies within the range of values that 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 about one.

[7] Respondents are asked to given the percentages of the use of the windfall gain corresponding to the following four alternatives: i) save for future expenses; ii) repay debt; iii) purchase within 12 months of durable goods (cars, home improvement, furniture, jewelry, other durable good) that they would otherwise would not have purchased or that they 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 that normally you would not have purchased. Similar questions have been used in a different context by Shapiro and Slemrod (1995, 2003), Jappelli and Pistaferri (2014).

7. Conclusions

The goal of the paper is to investigate the existence of a precautionary saving motive that affects household saving behaviour. To that effect, we estimate an Euler equation for consumption using subjective expectations on consumption, which conform better to the original formulation of the Euler equation than the ex-post consumption realizations that have been used up to now in the related literature. One additional motivation for the use of expectational data is that they allow researchers to circumvent problems related to inconsistency and endogeneity that plague ex-post realizations. In order to get the expectational data, we designed a questionnaire that asks households about expectations of their future consumption distribution, and used it on a representative sample of Dutch households.

Using these expectational data, we estimate 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 use a variety of estimation methods, namely OLS, robust regression and IV, and get consistent results across them that clearly point to the existence of a strong precautionary motive in the saving behaviour of the households in our sample. The estimated relative prudence coefficient is around two, in line with existing results in the literature.

Given that the expectational data correlate with observable household characteristics in accordance with both theory and intuition, they likely provide good measures of the underlying uncertainty experienced by households. We thus believe that questions on households' expectations are very useful ones to ask when one is interested in such uncertainty. More generally, these questions can provide valuable input in estimating economic relationships in which households' expectations play a role. Hence, we believe that including them in the questionnaires of household surveys is an advisable practice.

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., and A. Tambalotti, Giorgio 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.

Gollier C., and J. W. Pratt (1996), "Risk Vulnerability and the Tempering Effect of Background Risk," *Econometrica* 4, 1109-1124.

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.

Guiso, L., Jappelli, T., and D. Terlizzese (1992), "Earnings Uncertainty and Precautionary Saving," *Journal of Monetary Economics* 30, 307-337.

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. (2010a), "The Effects of the Economic Crisis on the Older Population," Michigan Retirement Research Center Working Paper No. 2010-231.

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, Arthur (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.

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) Coeff.	(2) Std. Error	(3) P value	(4) Change (in std. deviations) in the std. deviation of consumption growth when the variable changes by one standard deviation	(5) Change (in std. deviations) in the std. deviation of consumption growth when the variable changes by one unit
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 member of a union	-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. Columns 4-5 reports the implied change in the standard deviation of expected consumption growth (in terms of its own standard deviation) when the regressors change as specified.

Table 3. Euler equation estimates, simple triangular distribution

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	OLS			Robust Regression			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	0.640	0.122	0.000	0.963	0.008	0.000	0.888	0.165	0.000	0.994	0.129	0.000
Age	0.000	0.000	0.322	0.000	0.000	0.953	0.000	0.000	0.704	0.000	0.000	0.807
Female Householder	0.002	0.003	0.473	0.000	0.001	0.911	0.000	0.004	0.963	0.000	0.004	0.991
Household Size	0.001	0.002	0.595	0.000	0.000	0.453	0.001	0.002	0.723	0.000	0.002	0.783
Couple	-0.004	0.004	0.330	-0.001	0.001	0.599	-0.002	0.005	0.642	-0.002	0.005	0.728
Constant	0.014	0.009	0.099	0.006	0.003	0.014	0.003	0.010	0.741	0.001	0.008	0.950
Number of observations	3,271			3,271			2,890			2,890		
F-test	-..-			-..-			3.290			26.147		
F-test of generated instruments	-..-			-..-			-..-			19.244		
Breusch-Pagan test for heteroskedasticity - p-value	-..-			-..-			-..-			0.000		
Test of overidentifying restrictions - p-value	-..-			-..-			-..-			0.381		
Test of overidentifying restrictions of generated instruments - p-value	-..-			-..-			-..-			0.460		
Test of endogeneity of the treatment variable (consumption uncertainty) - p-value	-..-			-..-			0.273			0.274		

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

Table 4. Euler equation estimates, split triangular distribution

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	OLS			Robust Regression			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	0.642	0.123	0.000	0.966	0.008	0.000	0.833	0.181	0.000	0.979	0.144	0.000
Age	0.000	0.000	0.364	0.000	0.000	0.889	0.000	0.000	0.812	0.000	0.000	0.950
Female Householder	0.002	0.003	0.557	0.000	0.001	0.964	0.000	0.004	0.998	0.000	0.004	0.955
Household Size	0.001	0.002	0.493	0.001	0.000	0.257	0.001	0.002	0.551	0.001	0.002	0.652
Couple	-0.005	0.004	0.294	-0.001	0.001	0.450	-0.003	0.005	0.546	-0.002	0.005	0.664
Constant	0.015	0.009	0.095	0.006	0.003	0.011	0.003	0.010	0.742	0.000	0.008	0.976
Number of observations	3,167			3,167			2,791			2,791		
F-test	-..-			-..-			3.709			21.782		
F-test of generated instruments	-..-			-..-			-..-			14.517		
Breusch-Pagan test for heteroskedasticity - p-value	-..-			-..-			-..-			0.000		
Test of overidentifying restrictions - p-value	-..-			-..-			-..-			0.280		
Test of overidentifying restrictions of generated instruments - p-value	-..-			-..-			-..-			0.419		
Test of endogeneity of the treatment variable (consumption uncertainty) - p-value	-..-			-..-			0.343			0.472		

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

Table 5. Euler equation estimates, excluding liquidity-constrained households

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Has been denied Credit			Marginal Propensity to Consume ≥ 0.90			Household income above the 20th 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.

Fig. 1. Histogram of the minimum expected consumption level

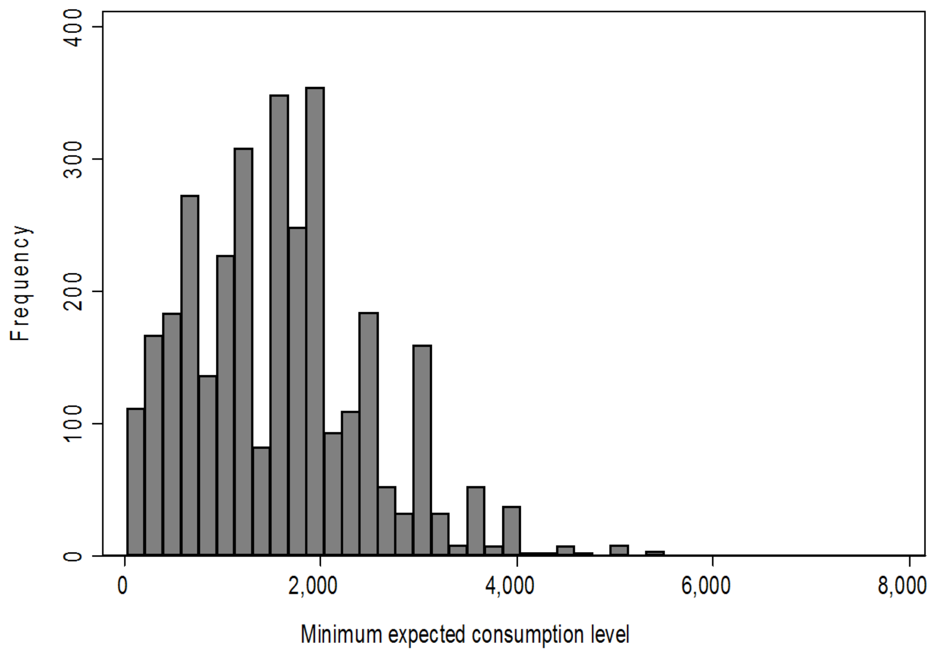

Fig. 2. Histogram of the maximum expected consumption level

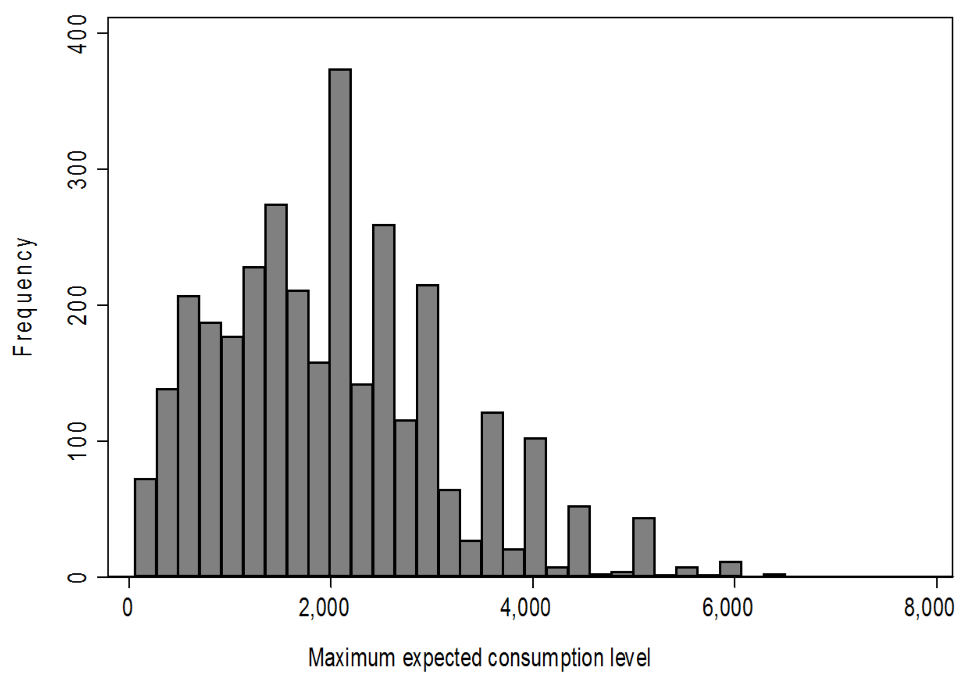

Fig. 3. Histogram of the probability that expected consumption is above the average of the expected minimum and maximum values

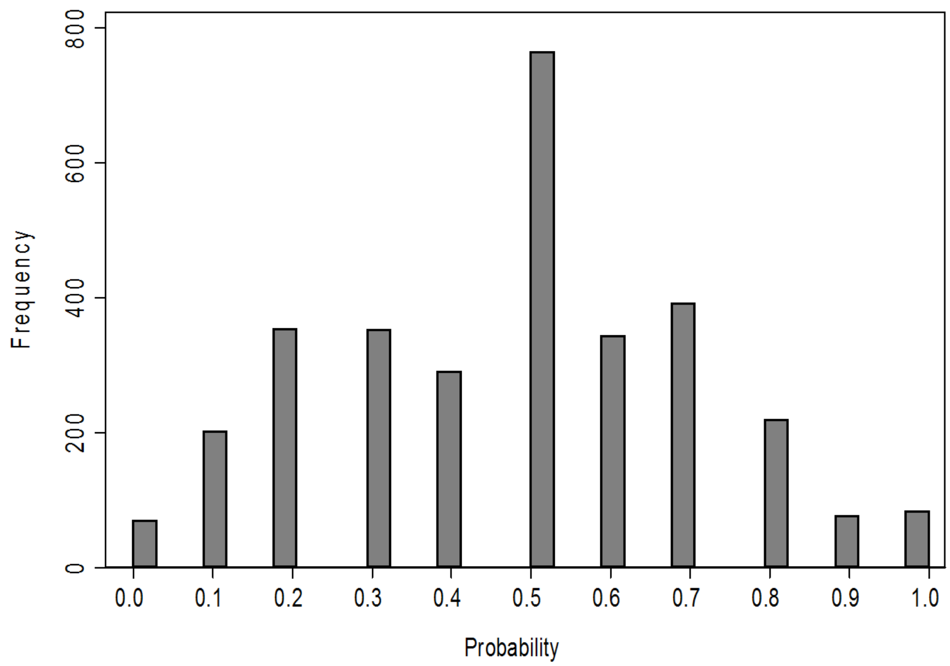

Fig. 4. Histogram of the standard deviation of expected consumption growth

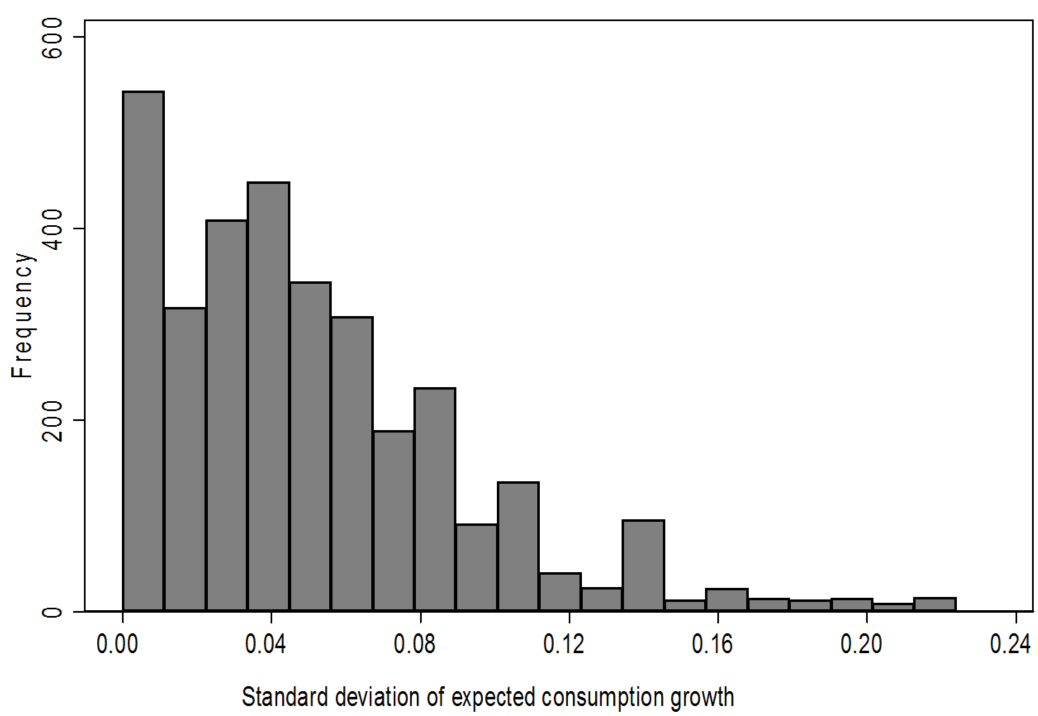

Fig. 5. Standard deviation of expected consumption growth, by age group

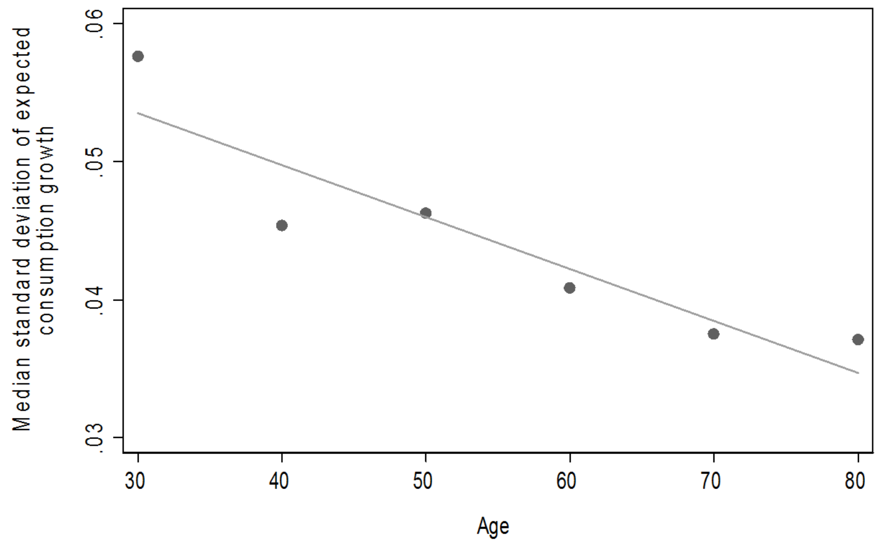

Fig. 6. Average expected consumption growth, by levels of the expected square of consumption growth

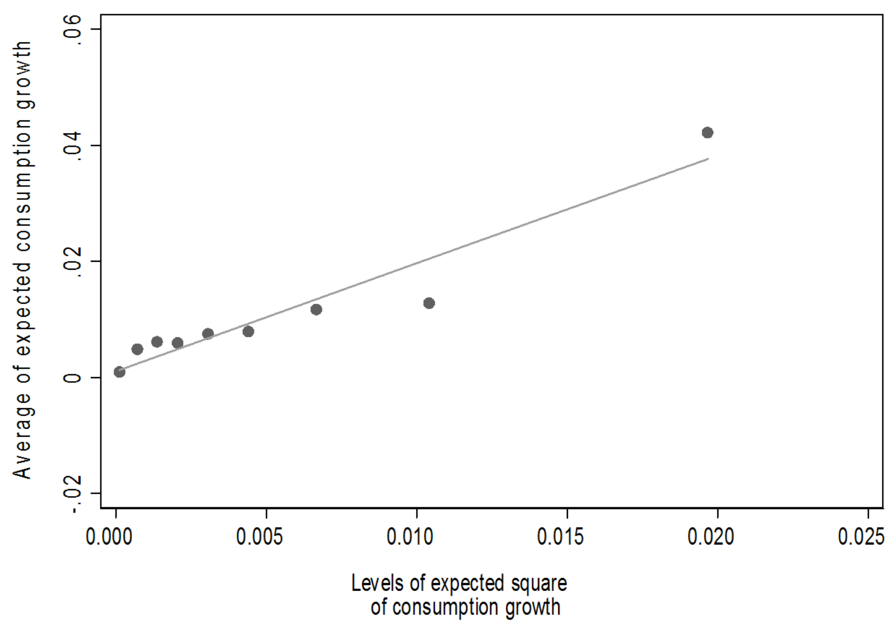

Appendix

A.1. Questions on future spending.

Questions on the future level of spending were asked to the survey respondent in the household 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 monthly future expenditures.

(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 household spending to be greater than X								Absolutely certain household spending to 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.'

Afterwards, respondents were asked the following introductory 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 time were subsequently asked to provide an estimate in percent of this change:

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

Respondents can give any answer from 0 to 100 in multiples of 10, or report that they don’t know the answer. In the latter case they are routed to a series of unfolding brackets questions with the following options:

Less than 5%

5 - 10%

11 - 15%

16 - 20%

21 - 25%

26 - 30%

More than 30%

Don't Know

As regards respondents who indicated in the introductory question above that they expect their spending to be ‘about the same’ to their spending today, they are asked to specify what do they actually mean:

‘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 one percent 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 split uniform distribution) by one plus the reported expected growth rate, and thus 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 same procedures as with the distribution of the expected consumption level. Given that $E(y^2) = Var(y) + E(y)^2$, once we obtain the variance of expected consumption growth for each household, we add to it the square of the reported expected consumption growth so as 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, we can express the expected value and variance of y as:

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

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

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 further away from the mid-point receive less weight.

Figure A1
The split triangular distribution