# PRELIMINARY AND INCOMPLETE. PLEASE DO NOT QUOTE

## Consumption Uncertainty and Precautionary Saving [#]

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

**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

March 2015

Abstract

Using data from the Dutch CentERdata Internet panel, we estimate the effect of expected consumption uncertainty on precautionary saving. Instead of using realized consumption values, as is the case in existing literature, we ask survey participants expectation questions and construct measures of expected consumption risk and growth that are household-specific. We find that expected consumption risk is associated with demographic and economic characteristics in the direction suggested by economic intuition, i.e. consumption risk is higher for the young, the self-employed and correlates positively also with income risk. We subsequently estimate, using OLS and IV methods, a standard Euler equation for consumption that relates expected consumption growth to expected consumption variability. We find that consumption uncertainty induces precautionary saving, and that the implied estimate of the coefficient of relative prudence is around 2. We obtain similar results when we also estimate, for the first time in the literature, our Euler equation by nonparametric partial identification methods that bound the effect of interest while making weak assumptions about the data.

JEL Classification: D12, D14, D81, E03, E21, C14
Keywords: Consumption Risk, Euler Equation, Prudence, Precautionary Saving, Expectations, Partial Identification.

[#] Tullio Jappelli acknowledges financial support from the Italian Ministry of University and Research (PRIN project) and the FP7 EU Framework Program (Mopact project). We thank seminar participants at CSEF for helpful comments. The opinions expressed in the paper are those of the authors and do not reflect the views of the Deutsche Bundesbank or of the De Nederlandsche Bank.

## 1. Introduction

The effect of uncertainty on consumer behavior has been a long-standing topic of research in the household saving literature. Life-cycle models of consumption behavior typically imply that increased income uncertainty will increase precautionary saving and consumption growth by lowering current consumption (see e.g. Skinner, 1988; Deaton, 1991). This increase in precautionary 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 uncertainty induced by income uncertainty 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. The few empirical tests of the precautionary saving motive using Euler equation estimates substitute for these two magnitudes their observed counterparts: actual consumption growth is regressed on actual consumption variability (Dynan, 1993; Bertola et al., 2005).

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 variability on expected consumption growth becomes problematic. In addition, using realized consumption values implies that the difference between them and their expectations (i.e. the expectational error) enters in the error term of the estimating equation. As discussed by Hayashi (1987), this expectational error should converge to zero as the time dimension increases, but the same is not true in the cross-section. This problem, as Chamberlain (1984) points out, is a serious one because it leads to inconsistent estimates when using cross-sectional data or short panels.

The first contribution of our 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 CentERdata Internet panel, which is sponsored and maintained by the Dutch National Bank. 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 such risk. This means that other risks (such as health risk, divorce risk, etc.), customs and institutions (for instance, the pooling of incomes within the family, or social insurance programs) are likely to affect consumption risk as well as 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 our 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 the variability of expected consumption growth is potentially correlated with the error term of the Euler equation. In order to address this problem, we use, following Bertola et al. (2005), the expected variability of income growth as an instrument. The third contribution that the paper makes to the literature on precautionary saving is the use of instrumental variable (IV) methods, recently introduced by Lewbel (2012), which can serve to check the validity of our instrument.

Finally, the fourth contribution of our paper lies with the use of nonparametric partial identification methods, due to Manski (1990) and Manski and Pepper (2000), to estimate the Euler equation. These methods impose very weak assumptions on the data

(especially on the nature of the instruments needed for the identification of the causal effect of interest) and allow us to bound from below the estimate of the coefficient of relative prudence.

When estimating the Euler equation using both OLS and IV methods, we find that expected consumption variability has a sizeable positive impact on expected consumption growth. In other words, we find strong evidence for the existence of a precautionary saving motive. The implied magnitude of the coefficient of relative prudence is about 2, implying in turn a value of the coefficient of relative risk aversion of about 1 (if one assumes an isoelastic utility function). Our nonparametric estimates corroborate the existence of a strong precautionary saving motive and provide an informative lower bound for the prudence coefficient, which lies between 0.5 and 1.3. Therefore, our 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. Section 5 discusses the estimation methods used. Section 6 presents the empirical results. Section 7 concludes.

## 2. The Euler equation with precautionary saving

The relationship between expected consumption variability 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:

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

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

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

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

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

where $p(c) \equiv u'''(c_t)c_t / u''(c_t)$ is Kimball's coefficient of relative prudence and $EIS \equiv -u'(c_t)/(u''(c_t)c_t)$ is the elasticity of intertemporal of substitution. 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 variability of consumption.

Equation (3) indicates that an increase in the expected variability of consumption induces 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]

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

[1] If one is willing to make specific assumptions about 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 $E_t \Delta lnc_{t+1} = \gamma^{-1}(r - \delta) + (\gamma/2) Var_t \Delta lnc_{t+1}$.

between movements in hours and in wages. Abowd and Card (1989) decompose fluctuations in earnings, which would normally be the preferred measure of risk, into exogenous fluctuations in wages and endogenous choice of hours. It is only the first type of fluctuation that represents genuine risk, although variation in hours has welfare costs 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). Low and Pistaferri (2010) 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. 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. Examples of such risks (and their associated costs) that have received prominent attention in the literature include the risk of future liquidity constraints, shocks to asset prices (including house prices), risk of medical expenditures, and of family dissolution.

## 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 choice (for instance, the choice to work in a risky occupation). Empirical evidence based on this approach is mixed. Most papers find a positive relationship consistent

with the precautionary saving model, but 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).

A second group of studies simulates 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 3 percent, they estimate a rate of time preference 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. Prudence is not directly estimated, but with an isoelastic utility function the implied coefficient of relative prudence is about 3. 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 US 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 4 for the high school sample), and an implied coefficient of relative prudence of around 5.

The most direct strategy to test for precautionary saving, however, is to use survey data and measure the expected variability of future consumption in equation (3) based on respondents' own assessment of uncertainty. This is the approach that we follow in this paper. The approach is useful, because from an empirical point of view, the main problem of estimating the Euler equation is that consumption risk is 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 variability, which is proportional to relative prudence.

The first paper that attempts to identify precautionary saving in an Euler equation is Dynan (1993), who, however, does not have in her data any information related to expectations about consumption. Hence, she has to replace expectations with their

realized counterparts. In this case equation (3) can be written in a regression framework as:

$$\frac{c_{i,t+1} - c_{i,t}}{c_{i,t}} = \alpha + \beta \left( \frac{c_{i,t+1} - c_{i,t}}{c_{i,t}} \right)^2 + \varepsilon_{i,t+1} \qquad (4)$$

where $i$ indexes households $(i = 1, \dots, N)$ and $\varepsilon_{i,t+1}$ reflects innovations to consumption growth, higher order terms of the Taylor expansion and measurement errors. As in equation (3), the coefficient $\beta$ is equal to $(1/2)p(c)$, and thus directly related to the strength of the precautionary saving motive. This substitution of realized consumption changes for their expectations implies that the error term $\varepsilon_{i,t+1}$ in (4) includes the difference between expected and realized magnitudes, and thus is clearly correlated with the term denoting realized consumption variability

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

Dynan's approach has been refined by Bertola et al. (2005), who use the subjective variance of income as an instrument for $((c_{t+1} - c_t)/c_t)^2$. As they point out, subjective income risk should be a very good instrument, as (3) implies that, conditional on expected consumption risk, income risk has no direct effect on consumption growth. In other words, the expected variability of consumption 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 estimate of the coefficient of relative prudence is around 2 and precisely estimated, thus providing evidence in support of the precautionary saving model.

In this paper, and in contrast to Dynan and Bertola et al., we estimate directly equation (3) using subjective expectations of future consumption, rather than relying on realized consumption magnitudes. This has several advantages with respect to previous tests of the precautionary saving motive. First, the error term of the estimated Euler equation does not include expectational errors in consumption, and therefore is not correlated with the ex-ante variability of consumption due to this reason. Second, the fact that we have a short panel at hand is not a problem for the consistency of our estimates. As mentioned in the Introduction, estimates derived using an Euler equation in cross-sections or short panels may be inconsistent when the time dimension of the panel is short (Chamberlain, 1984; Hayashi, 1987). The reason is that the error term of the Euler equation (4) 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 $E_{i,t}(\varepsilon_{i,t+1})$ in (4), which includes the expectation of the forecast error, is a household-level average taken over $T$ periods. Hence, $T \longrightarrow \infty$ is needed to ensure the consistency of the estimated parameters of the Euler equation. However, panel expenditure surveys are typically short ones, and hence researchers 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. The drawback of this simple solution is that it 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, expectational errors do not enter into the error term of our estimating equation, and thus the assumption that this error has a zero conditional expectation in

the cross-section is a reasonable one. Still, it might be the case that that there are unobservables that might be correlated both with expected consumption variability and consumption growth. Hence, we check the robustness of the results using an IV estimator as well as some alternative identification methods that relax the more stringent assumptions invoked in a standard IV framework.

## 4. The data

We use data from the CentERdata 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. It collects detailed information on a range of demographics and asset holdings for a representative sample of Dutch-speaking households.[2] 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 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 the analysis.

To elicit the distribution of expected consumption we follow a similar procedure as Guiso et al. (2002) and Guiso et al. (2013), who are interested in approximating the 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 future consumption 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. $Prob(y > (y_m + y_M)/2) = \pi$. We reproduce the wording of the questions in Appendix A.1.

[2] Panel members without Internet access are provided with equipment to participate in the surveys. This group is small, however, as more than 90% of the Dutch population has access to the Internet.

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 a split uniform one. Based on the elicited values of $y_m$, $y_M$ and $\pi$ , we compute the household-specific mean, standard deviation and coefficient of variation. The formulae of these statistics are reported in Appendix A.2.[3]

We set to missing values observations for which $y_m$, $y_M$ or $\pi$ are missing or respondents choose a 'do not know' option. The resulting sample includes about 76% percent of the original sample (2,537 household-level observations).

Moreover, the survey asks household financial respondents to report directly the expected change to their household spending in the 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.[4] Details on the relevant questions survey can be found in Appendix A.1.

In order to get a measure of the square of expected consumption growth, which is needed for estimating the Euler equation in (3), we divide the mean expected consumption level (derived through the split uniform distribution) by 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. Once we obtain the variance of expected consumption growth for each household, we add to it the square of the reported expected

[3] 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.

[4] 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 to obtain the expected square of consumption growth, i.e. our variable of interest.[5]

The survey also asks for information that allows one to compute moments of the distribution of future income in the same way as for future consumption. In particular, households are asked about the minimum and maximum values of the 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 (i.e. we use the split uniform distribution). We obtain 2,331 useful household-level values for expected income risk, i.e. about 70% of the total sample. As we will see, we use expected income variability as an instrument for expected consumption variability in the second part of the paper.

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 or equal 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.

Table 1 reports cross-sectional statistics of the central tendency and dispersion of the subjective distribution of consumption, assuming that the distribution is split uniform, and of the variables that will be used in the estimation (age, household size, marital status). On average, the minimum expected level of consumption is 1,583 euro, the maximum is 2,020 euro (the medians are equal to 1,500 and 1,900 euro, respectively), and the average probability is 0.48 (the median is 0.5). Assuming that the distribution is split uniform, we estimate that the average and median expected levels of consumption are equal to 1,824 and 1,725 euro, respectively.

Cross-sectional averages are useful to describe the subjective distribution of the average household, but hide important heterogeneity across households. Assuming that the distribution is a split uniform one, Figure 4 plots the cross-sectional distribution of

[5] Recall that for any random variable $y$ it is the case that $E(y^2) = Var(y) + (E(y))^2$.

the standard deviations of the 2,537 household-specific distributions of future consumption. The figure highlights considerable heterogeneity in the responses. For instance, for 25% percent of households the standard deviation is less than 16.5 euro, for another 50% between 16.5 and 59.5 euro, and for the top 25% percent more than 59.5 euro. The proportion of respondents that report the minimum expected consumption equal to the maximum (therefore expressing no uncertainty) is 11.4 percent. The coefficient of variation of the cross-sectional distribution of expected consumption is about 1.39.

Turning to consumption growth rates, Table 1 shows that the sample average of expected consumption growth is 1.2 percent (the median is zero). Since forecasts in the Netherlands indicate that in 2014 real consumption expanded slightly (approximately by 0.1 percent), consumption expectations are in aligned with realizations. The sample average of the standard deviation of consumption growth is 2.9 percent. The cross-sectional distribution of the standard deviation of consumption growth is plotted in Figure 5, showing again substantial heterogeneity in perceived consumption risk.

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 (such as age and occupation) that should influence consumption uncertainty. Figures 6 and 7 present kernel-smoothed means of the standard deviation of consumption growth distribution by age and education, respectively.

Figure 6 indicates that consumption risk declines during the life-cycle (the standard deviation of consumption growth falls by about 3 to 2 percent, signalling that younger households perceive more uncertainty than older consumers, in line with the findings of Dominitz and Manski (2006) for income. The age coefficient may capture several effects: first and most importantly, lifetime uncertainty is higher for younger workers; second, people differ by cohort. Unfortunately, our short panel data do not contain enough information to distinguish between these different explanations. Figure 7 shows that households with lower levels of education report smaller consumption uncertainty than those with a high-school or a college degree.

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, education, the standard deviation of expected income growth and inflation (constructed in similar fashion), and union membership (as a further measure of income volatility). 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. Consumption uncertainty strongly correlates with income uncertainty, but the correlation is far less than one-for-one, showing that other factors affect consumption uncertainty, besides income risk. Consumption uncertainty is also strongly correlated with a measure of households' perception of aggregate uncertainty, namely inflation risk (constructed in a similar fashion to consumption risk), and with having at least one self-employed member in the household. 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. On the other hand, age and being retired are negatively associated with consumption uncertainty, which is again as expected given the reduced income uncertainty associated in older age. All in all, the fact that these associations of consumption risk have the expected sign and are also quite strong suggests that our measure of the standard deviation of future consumption is a good indicator of the actual consumption uncertainty faced by the households in our sample.

## 5. Estimation methodology

### 5.1 OLS and IV

We are interested in estimating the impact of expected consumption risk on expected consumption growth, an indicator of precautionary saving. To perform the analysis, we specify equation (3) in Section 2 in a regression framework:

$$E_{i,t}\left(g_{i,t+1}\right) = \alpha + \beta E_{i,t}\left({g_{i,t+1}}^2\right) + \gamma X_{i,t} + v_{i,t+1} \qquad (5)$$

where $g$ indicates consumption growth between $t$ and $t+1$, and $v$ is a term that includes unobservables that affect expected consumption growth, as well as any measurement errors in our variables. Finally, the vector $X$ includes demographic variables such as the age and sex of the household financial respondent, whether (s)he has a partner, the size of the household, as well as an indicator for the survey wave. The demographic variables are included in the specification to capture any additional sources of behaviour heterogeneity (e.g. with respect to the rate of time preference)

Before presenting the econometric results, in Figure 8 we plot $E_{i,t}\left(g_{i,t+1}\right)$ against $E_{i,t}\left({g_{i,t+1}}^2\right)$, using a locally weighted scatter plot smoothing. The two variables are strongly positively correlated, and the slope of the relation between the two is not far to unity, with an implied coefficient of relative prudence of around 2. As we shall see, our estimation results, using conventional and robust OLS methods, IV methods, as well as nonparametric partial identification methods, are consistent with the descriptive evidence.

In order to reduce the influence of outliers, in our mainline specifications we winsorize both $g$ and $g^2$ at the top and bottom 0.5% of observations (i.e. we set the values of those observations equal to those at the $99.5^{th}$ and $0.5^{th}$ quantile, respectively). We estimate equation (5) first by conventional OLS, and then by robust regressions (using the M-estimator in Huber, 1973).

As already discussed, the use of elicited expectations in the estimation of (5) circumvents some of the identification issues inherent in existing studies that base their inference on consumption realizations. In particular, the use of expectations implies that the error term $v$ is not a forecast error, as is usually the case in Euler equation estimates. Nonetheless, there is still the possibility that unobservables in our error term $v$ are correlated with our measure of consumption variability. Furthermore, as described in Section 4, reported expected consumption growth is used to construct the square of expected consumption growth, and thus it is possible that unobservables that affect the former also affect the latter by construction.

Hence, we also estimate equation (5) also using IV methods, so as to take into account these possible endogeneity problems. We use as an instrument our measure of expected income variability (constructed in a similar fashion to that of consumption

variability, as described in Section 4). This is the same variable used in Bertola et al (2005), and, as already discussed in Section 3, it should represent a good instrument choice given that it does not appear in the Euler equation.

We also use an alternative identification approach due to Lewbel (2012), which generates additional instruments by exploiting variation of the higher moments of the errors in a first stage regression of the possibly endogenous variable on the exogenous regressors while excluding the original instrument. In particular, these additional instruments can be generated if two assumptions are fulfilled. First, the errors from a first stage regression of the endogenous covariate on a (subset) of other covariates in the model should be heteroscedastic. In our set-up, this first stage equation is as follows:

$$E_{i,t}\left(g_{i,t+1}{}^2\right) = \theta_1 + \theta_2 W_{i,t} + \omega_{i,t+1} \tag{6}$$

where $\omega$ denotes an error term and $W$ denotes covariates used to estimate (5). In our case, these include the age and gender of the financial respondent, family size, and a couple indicator. Heteroscedasticity in (6) can be tested on the basis of a Breusch-Pagan test and it turns out that is it very strong in our data.

The second condition needed for the generated regressors in Lewbel (2012) to be valid is that $W$ is assumed to satisfy $cov(W, v \cdot \omega) = 0$. This condition holds even when $v$ and $\omega$ share a common unobserved factor component (and are thus correlated), as long as the product of their idiosyncratic error components is uncorrelated with $W$. Given that numerous additional instruments are generated, their exogeneity can be tested using a standard test of overidentifying restrictions, under the assumption that one of them (e.g. age) is exogenous.

Under the two aforementioned assumptions, Lewbel shows that a set of valid instruments for estimating (5) can be generated as $(W_i' - \bar{W}_i')\widehat{\omega_i}$, where $\bar{W}$ is the mean of $W$ and $\widehat{\omega}$ are estimated residuals from (6). The advantage of using generated instruments in our context is twofold: (i) it allows us to compare our estimates from the standard IV approach to those obtained using the generated instruments from the Lewbel method, because the generated instruments create a number of overidentifying restrictions, that, if true, can in turn render our estimates more efficient; (ii) given that

our standard IV model is exactly identified, the additional instruments allow us to test the validity of our main instrument (expected income risk) using again a standard test of overidentifying restrictions.

### 5.2. Partial Identification[6]

As already mentioned, expected consumption variability could be endogenous to expected consumption growth, and thus one would need to use IV estimation methods. However, in the presence of heterogeneity of the effect of consumption variability on consumption growth, IV methods identify only the local average treatment effect (LATE, see Imbens and Angrist, 1994), i.e. the effect for those who change their perception of consumption risk because of a corresponding change in the value of the instrument, i.e. their perceived income risk. This group of respondents cannot be identified in the data, and in any case we would prefer to estimate, if possible, the effect of interest in the whole population.

In order to address this issue, we use partial identification methods that identify the causal effect of interest for the whole population. These methods, introduced by Manski (1989, 1990, 1994), are nonparametric and produce bounds on the average treatment effect (ATE henceforth). In other words, they locate the ATE in an identification region instead of calculating a point estimate. Importantly, partial identification methods use assumptions that are much weaker than those used in OLS and IV estimation methods.

As in Manski (1997), for each individual $i$ there is a response function $y_i(\bullet): D \to Y$ that maps mutually exclusive and exhaustive treatments $d \in D$ into outcomes $y_i(d) \in Y$. Importantly, these response functions $y_i(\bullet)$ can differ across individuals in arbitrary ways, thus allowing for unlimited response heterogeneity. Let $w_i$ denote the realized treatment received by $i$, and $y_i \equiv y_i(w_i)$ denote the associated observed outcome. On the other hand, $y_i(d)$ is a latent potential outcome when $d \neq w_i$. In our context, the outcome is expected consumption growth, while the treatment variable is expected consumption variability. In order to use partial identification

[6] This section is based on Christelis and Dobrescu (2014).

methods we need to discretize our treatment variable. We thus assign to it the values of 0 and 1 if a household exhibits expected consumption uncertainty below and above the sample median, respectively. Consequently, $y_i(0)$ and $y_i(1)$ are the two possible values of the outcome for household $i$. We would like to estimate the ATE of expected consumption risk on expected consumption growth, i.e.

$$ATE = E[y(1)] - E[y(0)] \tag{7}$$

Note that the ATE in (7) represents the difference in the two mean outcomes, which are both evaluated using all population units while keeping all other observable and unobservable variables fixed at their realized values (Manski 1997, p. 1322). In our context, the ATE is equal to the difference in expected outcomes when every household experiences high consumption uncertainty as opposed to when every household experiences low uncertainty. By the law of iterated expectations, and given that $E[y(d)|w = d] = E[y|w = d]$, the expected outcome as a function of $d$ is equal to

$$E[y(d)] = E[y|w = d]P(w = d) + E[y(d)|w \neq d]P(w \neq d) \tag{8}$$

where $P(w = d)$ denotes the probability that $w = d$. The term $E[y(d)|w \neq d]$ in the right hand side of (8) is a counterfactual one because it denotes the expectation of the outcome as a function of $d$ when the treatment actually received is different from $d$. On the other hand, the remaining three terms in the right hand side of (8) have sample analogues that are observed in the data. Given that $E[y(d)|w \neq d]$ is unobserved, the unconditional expectation $E[y(d)]$ is also unobserved, i.e. it represents a potential outcome. Hence the ATE in (7) is equal to the difference between two average potential outcomes.

If one assumes that the counterfactual conditional expectation $E[y(d)|w \neq d]$ is equal to the observed one when the treatment actually received is equal to $d$, i.e. if

$$E[y(d)|w \neq d] = E[y|w = d] \tag{9}$$

then from (8) it follows that

$$E[y(d)] = E[y|w = d] \quad (10)$$

Equation (10) states that the unobserved potential outcome under $d$ is equal to the mean outcome when the treatment actually received is equal to $d$. As the sample analogue of the latter is observed in the data, one can estimate the unobserved potential outcome $E[y(d)]$, and then in turn the ATE defined in equation (7), which is equal to

$$ATE = E[y|w = 1] - E[y|w = 0] \quad (11)$$

We refer to the estimate of the ATE in (11) as the one under exogenous treatment selection (ETS henceforth) because it is derived under the assumption that (9) holds, which in turn implies that households receiving different treatments (i.e. experiencing different levels of expected consumption uncertainty) are not systematically different from each another. In other words, (9) implies that selection into treatment is exogenous.

Equation (9) is likely to be true in the case of a randomized control trial, in which treatment assignment is indeed exogenous. In observational data, however, (9) might not hold because treatment assignment is not random, especially when the treatment variable reflects an attitude or expectation experienced by households, as is the case in our context. Hence, the expected potential outcome $E[y(d)]$ might differ, for any $d$, between these households actually experiencing different levels of the treatment,. In other words, the fact that equation (9) may not to hold in our data is due to the endogeneity of the expectation about consumption risk. Such endogeneity can have any source, as all sources eventually lead to non-random treatment assignment, i.e. to the violation of (9).

Once one rules out the application of (9), the problem of estimating the unobservable potential outcome $E[y(d)]$ arises. As a solution, Manski (1989) suggested bounding this outcome from above and below. Let us denote the lower and upper bounds on $E[y(d)]$, computed using a particular method $M$, as $LB^M[d]$ and $UB^M[d]$, respectively.[7] Given that $LB^M[d] \leq E[y(d)] \leq UB^M[d]$, Manski (1990)

[7] It is important to note that one calculates the bounds of the potential outcome $E[y(d)]$ for each value $d$ of the treatment independently from all other treatment values.

points out that equation (7) in turn implies that one can bound the ATE using method $M$ as follows:

$$LB^M[1] - UB^M[0] \leq ATE \leq UB^M[1] - LB^M[0] \quad (12)$$

The interval between the lower and the upper bound on the ATE is its identification region, and since it is an interval the ATE is partly identified.

When calculating the upper and lower bounds on $E[y(d)]$, a natural starting point is to assume that, for any value $d$ of the treatment, the outcome space $Y$ is bounded below and above by finite values $Y_{min}$ and $Y_{max}$, respectively. In our application, we use the observed sample minima and maxima as $Y_{min}$ and $Y_{max}$.

As in Manski (1990), by using equation (8) and after replacing the unobserved term $E[y(d)|w \neq d]$ by $Y_{min}$ and $Y_{max}$, one can bound $E[y(d)]$ from below and above as follows:

$$\begin{aligned} &E[y|w = d]P(w = d) + Y_{min}\, P(w \neq d) \\ &\leq E[y(d)] \leq \\ &E[y|w = d]P(w = d) + Y_{max}P(w \neq d) \end{aligned} \quad (13)$$

The bounds in (13) are obtained without imposing any assumptions on the data, other than the existence of finite $Y_{min}$ and $Y_{max}$. We thus denote them as the no assumptions (NA) bounds. Moreover, the NA bounds can be readily calculated using their sample analogues, as these are observed in the data. As Manski (1989) points out, taking sample averages leads to consistent estimates of $E[y|w = d]$, $P(w = d)$ and $P(w \neq d)$.

It follows from (13) that the distance between the NA upper and lower bound of $E[y(d)]$ (i.e. the length of its identification interval) is equal to $(Y_{max} - Y_{min})P(w \neq d)$. Moreover, the distance between the upper and the lower bound of ATE computed using the NA method is equal to $Y_{max} - Y_{min}$ (Manski, 1990).

The NA identification region for the ATE is typically very wide, and always includes zero. Hence, one has to make additional assumptions in order to make it narrower. The first such assumption is that of monotone treatment response (MTR

henceforth; see Manski, 1997). The MTR assumption states that for all sample units $i$, and for any two treatment values $d_1 \in D$ and $d_2 \in D$ such that $d_2 > d_1$,

$$y_i(d_2) \geq y_i(d_1) \tag{14}$$

In our context, the MTR assumption implies that, for all households in our sample, being more uncertain about future consumption does not decrease consumption growth, i.e. does not reduce precautionary saving. Importantly, (14) holds irrespective of the treatment actually received. Given that at each point in time one observes only one outcome for each household in the sample, one cannot test for the validity of (14) in isolation using the data at hand. However, it is difficult to think of a mechanism through which higher uncertainty would reduce a household's precautionary saving. Moreover, and to the best of our knowledge, there is no empirical evidence for such a detrimental effect. Importantly, the MTR assumption fully allows for the possibility that there is no effect of expected consumption uncertainty on expected consumption growth; therefore, we believe that the MTR assumption is a reasonable one to make.

In practice, we use a weaker, and thus more conservative, version of the MTR assumption than the one in (14). This weaker version states that for any treatment value $d \in D$, and any two values $d_1 \in D$ and $d_2 \in D$ such that $d_2 > d_1$,

$$E[y(d_2)|w = d] \geq E[y(d_1)|w = d] \tag{15}$$

Equation (15) implies that expected consumption variability has a weakly positive effect on expected consumption growth on average, but not necessarily for every individual in the sample. Furthermore, this average weak monotonicity holds for all subsamples that are defined by the treatment actually received.[8] Clearly, (14) implies (15), but the converse is not necessarily true.

As Manski (1997) shows the MTR assumption implies that the bounds on $E[y(d)]$ can be expressed as follows:

[8] Given that (15) holds for all values $d$ of the observed treatment $w$, it is clearly the case that the weak monotonicity in (15) applies also to the unconditional expectation, i.e. (15) implies that $E[y(d_2)] \geq E[y(d_1)]$. However, the converse need not be true.

$$E[y|w<d]P(w<d)+E[y|w=d]P(w=d)+Y_{min}\,P(w>d) \leq E[y(d)] \leq Y_{max}P(w<d)+E[y|w=d]P(w=d)+E[y|w>d]P(w>d) \quad (16)$$

The narrower identification region of $E[y(d)]$ under MTR yields in turn a narrower identification region of ATE resulting from (16), which is now bounded below by zero (Manski, 1997). This is to be expected, as the MTR assumption in (15) rules out the possibility that a higher level of the treatment induces a lower mean outcome, while allowing for the possibility of a zero effect. Furthermore, the upper bound on ATE derived using the MTR assumption is equal to the NA one.

One can further tighten the identification region of $E[y(d)]$ and the ATE by combining the MTR assumption with the use of a variable $Z$ that is exogenous to $E[y(d)]$, in the sense that for all values $z$ of $Z$ the following condition holds:

$$E[y(d)|Z=z]=E[y(d)] \quad (17)$$

While equation (17) requires that $Z$ be exogenous to the outcome, Manski and Pepper (2000, MP henceforth) point out that it still needs to be correlated with the treatment variable, as typically required in IV estimation. In addition, given that (17) is an assumption on the unobserved potential outcome $E[y(d)]$, the exogeneity of $Z$ cannot be tested without using further assumptions. Obviously, this difficulty is not particular to bounds-based estimation, as it is present also in any conventional IV setup.

If an exogenous $Z$ can be found, then one can choose various values in its domain, and (17) implies that each of these values can be used to obtain an estimate of $E[y(d)]$. Hence, after calculating the lower and upper bounds of $E[y(d)]$ under MTR for each value of $Z$, one can calculate the maximum of the lower bounds of $E[y(d)]$ across all these values, and use this maximum as an estimate of the lower bound of $E[y(d)]$. Correspondingly, one could compute the minimum across all the values of Z of all upper bounds of $E[y(d)]$ under MTR, and use this minimum as the estimate of the upper bound of $E[y(d)]$. As a result, the bounds on $E[y(d)]$ under the combination of MTR with an exogenous instrument can be computed as follows:

$$\max_z LB^{MTR}[d|Z=z] \leq E[y(d)] \leq \min_z UB^{MTR}[d|Z=z] \quad (18)$$

The optimization operations in (18) can produce a larger lower bound and a smaller upper bound on $E[y(d)]$ compared to their MTR counterparts in (16).

In our context, and as already discussed, we use the expected consumption variability as an exogenous instrument, and in order to compute the minima and maxima over its values in (18) we divide it in quintiles.

A considerably weaker kind of instrumental variable is the monotone one (MIV henceforth), introduced by MP, which satisfies the following requirement for any pair of values $z_1, z_2$ of $Z$ such that $z_2 > z_1$,

$$E[y(d)|Z=z_2] \geq E[y(d)|Z=z_1] \quad (19)$$

Equation (19) states that the MIV can influence the outcome in a given direction, while the possibility of no influence whatsoever is also allowed for. Hence, the assumption of a monotone instrumental variable is a much weaker than the assumption of an exogenous one. It is important to note that (19) captures only a positive association of $Z$ with $Y$; a causal relationship is neither implied nor required.

Once one adds the MIV assumption to MTR, then for a given value $z$ of $Z$ one can calculate the maximum of the MTR lower bounds over all values of $Z$ that are smaller or equal to $z$. By (19), this maximum lower bound cannot be larger than $E[y(d)|Z=z]$. Similarly, one can calculate the minimum MTR upper bound over all values of $Z$ larger or equal to $z$, and (19) implies that this minimum cannot be smaller than $E[y(d)|Z=z]$. Hence the MTR +MIV assumption implies that

$$\max_{z_1 \leq z} LB^{MTR}[d|Z=z_1] \leq E(y(d)|Z=z) \leq \min_{z \geq z_2} UB^{MTR}[d|Z=z_2] \quad (20)$$

Once the bounds in (20) have been computed for all $z$, one can take their weighted average over all $z$ and bound the potential outcome $E\big(Y(d)\big)$ as follows:

$$\sum_{z} P(Z=z) \max_{z_1 \leq z} LB^{MTR}[d|Z=z_1]$$

$$\leq \sum_{z} P(Z=z) E[y(d)|Z=z] = E[y(d)] \leq \tag{21}$$

$$\sum_{z} P(Z=z) \min_{z \geq z_2} UB^{MTR}[d|Z=z_2]$$

Hence, by integrating $Z$ out of the conditional expectation in (20) one can obtain bounds on $E[y(d)]$. In contrast, this is not needed when using exogenous instruments because such instruments do not affect $E[y(d)]$ by assumption, as can be seen in (17).

We choose as our monotone instruments two variables denoting uncertainty experienced by households. The first one is uncertainty about expected income growth (in quintiles), i.e. the variable already used as an exogenous instrument. Using it as a monotone instrument is a much weaker assumption than that of exogeneity, as it allows uncertainty about income to have a direct effect on the outcome of interest, i.e. to (weakly) increase expected consumption growth. In the same vein, we use as an additional monotone instrument uncertainty about expected inflation (computed in analogous way as consumption and income uncertainty), also in quintiles. The reasoning behind the use of this variable as a monotone instrument is that aggregate uncertainty also has a weakly positive effect on precautionary saving.

As is the case with exogenous instruments, these weak monotonicity assumptions are imposed on the unobserved potential outcome $E[y(d)]$; hence, they cannot be tested using the observed data without imposing further assumptions. On the other hand, we also note that the weak inequality of the MIV condition in (19) accommodates the possibility that the monotone instruments have no effects on the outcome.

In order to conduct inferences on $E[y(d)]$ and the ATE we compute confidence intervals (CIs henceforth) that cover these magnitudes with 95% probability by using 1,000 bootstrap replications, as described in Imbens and Manski (2004). Given that both the XIV and the MIV bounds involve optimization operations, the bootstrapped estimates of the bounds can be biased. Therefore, we apply the bias correction

procedure suggested by Kreider and Pepper (2007). We found, however, that this bias correction made essentially no difference to our results. As we have repeated observations for a large proportion of households in our sample, we draw the bootstrap samples after clustering by households.

The features of the bounds-based estimation method presented in this Section provide many reasons why one would consider it a valid alternative to more commonly used methods, e.g. OLS-, IV- or panel data-based ones, when trying to estimate the causal effect of interest. First, bounds-based estimation is completely nonparametric, as it only involves taking simple sample averages of the outcome and the treatment. Hence, it is not affected by the problem of estimates resulting from local minima and maxima. Second, bounds-based estimation leads to estimates of the ATE across all sample units, and not of the LATE as is the case with IV estimation when the treatment effect is heterogeneous across the population. Third, it allows for arbitrary forms of heterogeneity of the treatment effect because the ATE is just an average magnitude across sample units. Hence, the treatment effect for each sample unit can depend on any other variable in a fully flexible way. Such unlimited heterogeneity of the treatment effect is not typically allowed for, as in most estimation methods one makes particular assumptions about how the treatment variable enters into the specification. Moreover, if one is interested in studying the heterogeneity of the treatment effect in particular dimensions in the context of partial identification, then one can simply restrict estimation to subsamples defined by particular combinations of values of control variables. Fourth, as already discussed, the computation of the ATE takes the distribution of all observables and unobservables in the sample as given. Hence, one does not need to worry about: i) which variables to add in the empirical specification; ii) the manner in which they appear; iii) whether they are endogenous or not. Fifth, bounds-based estimation accommodates any form of endogeneity (e.g. due to both time-varying and time-invariant unobservables or selectivity), as it allows for non-random selection into treatment in any form. This also implies that one does not need to assume particular properties of an error term, as is the case with regression methods. Sixth, partial identification methods use very few and quite mild assumptions to narrow the identification region of the estimates. Importantly, they are completely

transparent about how the addition of each assumption affects the identification region. In contrast, most commonly used estimation methods typically impose many assumptions on the empirical model at the same time, and thus it is typically not clear how each assumption affects estimates. Seventh, partial identification methods allow the use of both exogenous and monotone IVs, which can tighten identification regions. As is the case with standard IV estimation, the assumptions behind these IVs cannot be tested without making further assumptions. However, MIVs, which cannot be used in standard IV estimation, have properties that are much less restrictive than those of exogenous IVs. Finally, in bounds-based estimation one uses the data as a cross-section, and thus panel data are not required. One can accommodate any dependencies among sample units (e.g. due to repeated observation or features of the sampling process) through the appropriate clustering and stratification when bootstrapping standard errors.

On the other hand, partial identification methods can lead to identification regions that are wide, and thus do not allow one to draw strong conclusions about the effect of interest in all directions. This is the price that sometimes one has to pay for imposing very few and weak assumptions on the data. As Manski (1994) notes, the point identification obtained by more commonly used estimation methods may give one a false certainty about results, as the reduction in uncertainty is obtained through assumptions that are not testable, and that might not hold in the data. This discrepancy between point identification through strong assumptions and partial identification through mild assumptions shows up very clearly in our results.

## 6. Estimation results

In Table 3, columns 1-5, we show OLS results from a specification that correlates expected consumption growth with expected consumption growth squared and the demographic variables described in Section 5.1. The estimated coefficient of interest is .42 with a p-value of .025, implying a prudence coefficient of about 0.84.

Second, in order to check the sensitivity of the baseline results from OLS to possible outliers we perform robust regressions, using Huber's (1973) M-estimator. Results from this estimation are shown in columns 6-10 in Table 3 and yield an estimated coefficient of 0.95 on the expected consumption variability (p-value<.01). The derived estimate that is resilient to outliers is twice as large as the corresponding OLS one, which in turn implies a twice as large prudence coefficient of about 2.

We then perform a standard IV estimation using expected income variability as an instrument for expected consumption variability. The validity of this approach requires that our measure of expected income variability correlates with that of consumption variability and moreover the former is orthogonal to the error term in (5). Results from this standard IV approach are shown in Table 4, columns 1-5. The implied effect of expected consumption variability on expected consumption growth is 1.18 and precisely estimated (p-value<.01). First-stage regression results confirm our prior that consumption variability correlates positively with income variability. Nevertheless the corresponding F-statistic is below the rule of thumb threshold of 10. Finally, one can test the endogeneity of the expected consumption variability using a Hausman test, and the result (p-value = 0.02) indicates that the null of exogeneity is not valid, and thus IV estimation is indeed needed.

Turning to IV estimation using generated regressors, a Breuch-Pagan test ($X^2$ = 656, p-value = 0.000) indicates that there is strong heteroscedasticity in the first-stage regression shown in (6). Following this result, we generate four instruments by taking the products of residuals from (6) with age, gender, household size and the couple indicator, centered at their respective sample means. Columns 6-10 in Table 4 show results when using these generated instruments. The coefficient of expected consumption variability is .79, precisely estimated (p-value<.01), and implies a prudence coefficient of about 1.6, i.e. slightly lower than but still quite comparable with the one derived using robust OLS.

The generated instruments meet the exogeneity assumption, as Hansen's J statistic fails to reject the null of exogeneity with high confidence (p-value: .78). One can also test the exogeneity of the original instrument only, using a generalized method of moments-based test (C-test) of exogeneity, as described in Hayashi (2000, pp. 218-

221).[9] The result of the test (p-value: 0.223) do not refute the null hypothesis that expected income variability is a valid instrument.

The combined F-statistic of our original instrument and the generated instruments is about 14.2, while the F-statistic of the generated instruments only is about 16. 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. Finally when we perform again a Hausman test for the endogeneity of the consumption risk variable, the result (p-value: 0.03) again indicates that the null can be refuted.

Finally, we turn to the partial identification results, shown in Panel A of Table 5. The ETS estimate implies that the average difference in expected consumption growth between households who exhibit consumption risk above and below the median is about 1.5 percentage points and precisely estimated. The ATE under NA has a predictably wide and thus uninformative identification region, ranging from -64.2 percentage points to 65.8 points. Adding the MTR assumption makes the lower bound of the ATE equal to zero, while leaving the upper bound unchanged.

Turning now to methods using instruments, when we combine the MTR assumption with the use of expected income risk as an exogenous instrument the lower bound of the ATE is 2.9 percentage points, while the upper bound is a very large and uninformative 49.3 points. The lower bound of the 95% CI is also above zero. Hence, the use of an exogenous instrument results in a non-trivial estimated effect of consumption risk on consumption growth, and points to the existence of a precautionary saving motive.

When we use expected income risk as a monotone instead of an exogenous instrument, and in combination with expected inflation variability, the estimated lower bound of the ATE becomes smaller (1.1 percentage points), but it is still economically relevant and statistically significant at 5%. Hence, even the lowest estimates of the ATEs that are consistent with the data imply the existence of a strong precautionary

[9] The C-test of exogeneity uses the difference in the Sargan-Hansen J statistic between the model using only the generated instruments and the model using both the original and the generated instruments.

saving motive. On the other hand, the price to be paid for making weak assumptions is that we are not able to put an informative upper bound on how strong this motive is.

Given that we use a discrete version of our treatment variable, the scale of the partial identification results is not comparable to that of the OLS and IV results. In order to facilitate scale comparability, we divide all estimates in Panel A of Table 4 with the difference in the mean values of the original continuous treatment variable across the two levels of the discretized treatment.

The adjusted results are shown in Panel B of Table 5 and imply that the ATE under ETS is 0.6, which is close to our preferred IV estimate using generated regressors. The upper bounds are in all cases uninformative, while the lower bound under MTR+XIV is 1.3, i.e. quite close to the robust OLS estimates. On the other hand, the lower bound under MTR+MIV is 0.5, which is a bit smaller than the IV estimates. While the MTR+XIV method implies a lower bound of the prudence coefficient of slightly above 2.6, the MTR+MIV method implies a lower bound of about 1.

All in all, the results from all our estimation methods suggest that there is a positive and economically relevant effect of consumption variability on expected consumption growth. The variation in the magnitude of the implied prudence coefficient is proportional to the strength of the assumptions used to generate the estimates: stronger assumptions like the existence of exogenous instruments result in higher estimates, while weaker (but possibly more credible) assumptions suggest smaller ones. Even the weakest results, however, are economically relevant and significant, and allow us to conclude that the evidence for the existence of a precautionary saving motive is quite robust.

## 7. Conclusions

The goal of the paper is to estimate an Euler equation for consumption using expectational data, which conform much better to the standard formulation of the Euler equation than the realized consumption data that have been used up to now in the

related literature, or reduced form estimates linking saving, consumption or wealth to proxies for income risk. One key 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 to report their expectations of future consumption (as well as income), and used it on a panel of Dutch households.

Using these expectational data we then estimate the Euler equation to get an indication of the existence and strength of the precautionary saving motive. The key estimated parameter is the prudence coefficient, which is a measure of the convexity of the marginal utility of consumption. We use a variety of estimation methods, namely OLS, IV and partial identification methods, 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 magnitude of the prudence coefficient depends on the strength of the assumptions that we make, but even when using very weak assumptions we can establish the existence of precautionary saving.

Given that expectational data about future consumption and its variability correlate as theory and intuition indicates with observable household characteristics, they seem to provide good measures of the underlying uncertainty experienced by households. As a result, we find that questions on households' expectations are very useful for studying household decisions under uncertainty, and that their inclusion in household surveys' questionnaires 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.

Blanchard, Olivier Jean, and N.G. Mankiw (1988), "Consumption: Beyond Certainty Equivalence," *American Economic Review* 78, 173-77.

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.

Christelis, D., Dobrescu, I. (2014), "The Cognitive and Economic Impact of Social Activities in Older Age: Evidence from 17 Countries". SSRN Working Paper No. 2141386. <http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2141386>

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*, dx.doi.org/10.1016/j.jmoneco.2015.01.003.

Deaton, A. S. (1991), "Saving and Liquidity Constraints," *Econometrica* 59, 1221-48.

Dominitz, J. and C. F. Manski C. F. (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.

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.

Hayashi, F. (2000). *Econometric*s. Princeton, NJ: Princeton University Press.

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

Imbens, G. W., Angrist, J. (1994), "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62, 467–476

Imbens, G. W., Manski, C. F. (2004), "Confidence intervals for partially identified parameters," *Econometrica*, 72, 1845–57.

Kimball, M. S. (1990), "Precautionary Saving in the Small and in the Large," *Econometrica* 58, 53-73.

Kreider, B., Pepper, J. V. (2007), "Disability and employment: Reevaluating the evidence in light of reporting errors," *Journal of the American Statistical Association*, 102, 432–41.

Lewbel, Arthur (2012), "Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models," *Journal of Business & Economic Statistics* 30, 67-80.

Low, H., and L. Pistaferri (2010), "Disability Risk, Disability Insurance and Life Cycle Behavior," NBER Working Paper 15962

Low, H., Meghir, C., and L. Pistaferri (2010), "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review* 100, 1432-67.

Manski, C. F. (1989), "Anatomy of the selection problem," *Journal of Human Resources* 24, 343–60.

Manski, C. F. (1990), "Nonparametric Bounds on Treatment Effects," *American Economic Review Papers and Proceedings*, 80, 319-323.

Manski, C. F. (1994), "The Selection Problem," in *Advances in Econometrics: Sixth World Congress*, ed. by C. Sims. Cambridge, UK: Cambridge University Press, 143–170.

Manski, C. F. (1997), "Monotone Treatment Response," *Econometrica*, 65, 1311-1334.

Manski, C. F. (2004), "Measuring Expectations," *Econometrica* 72, 1329-1376.

Manski, C. F. , and John V. Pepper (2000), "Monotone Instrumental Variables, with an Application to the Returns to Schooling," *Econometrica* 68, 997-1012.

Shapiro, M. (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.

**Fig. 1. Histogram of the minimum expected consumption level**

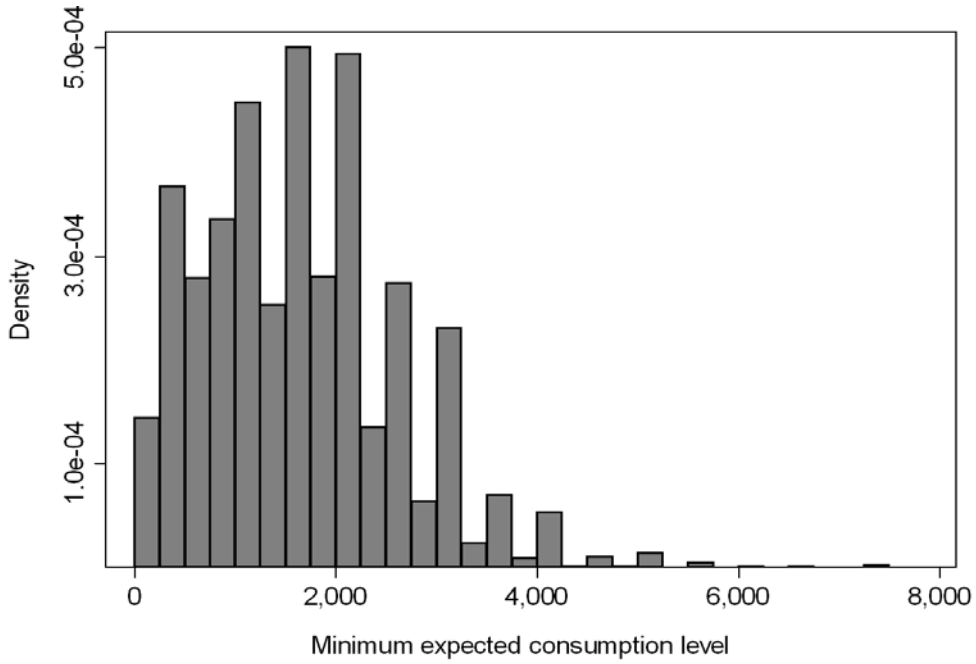



**Fig. 2. Histogram of the maximum expected consumption level**

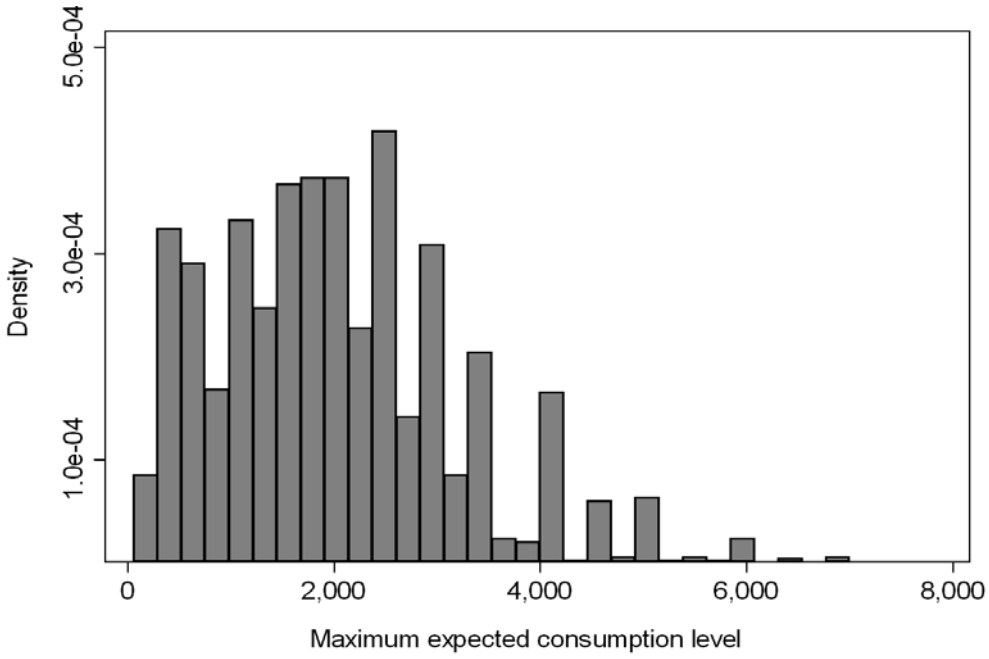


**Fig. 3. Histogram of the probability that expected consumption is above the average of the expected minimum and maximum values**

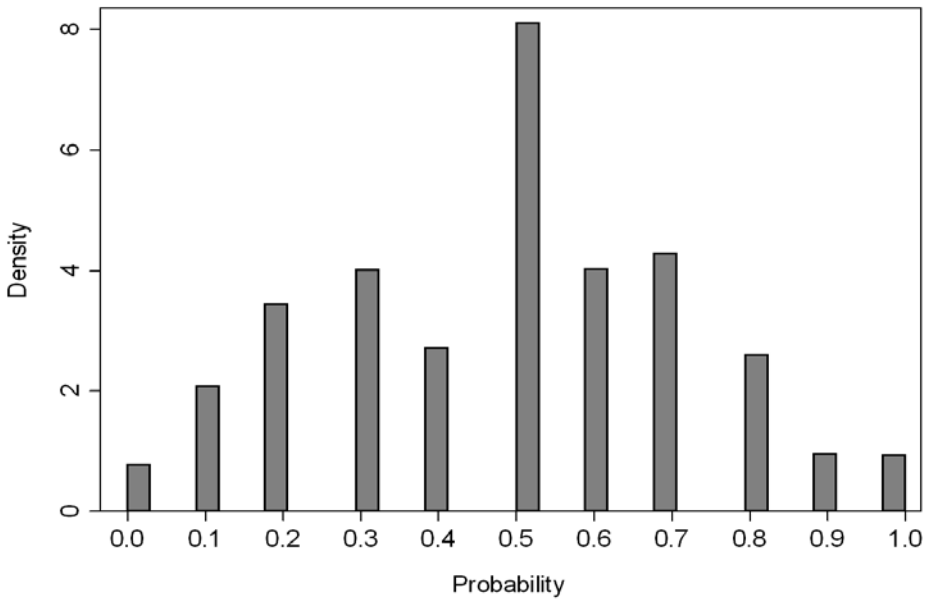



**Fig. 4. Histogram of the standard deviation of the expected level of consumption**

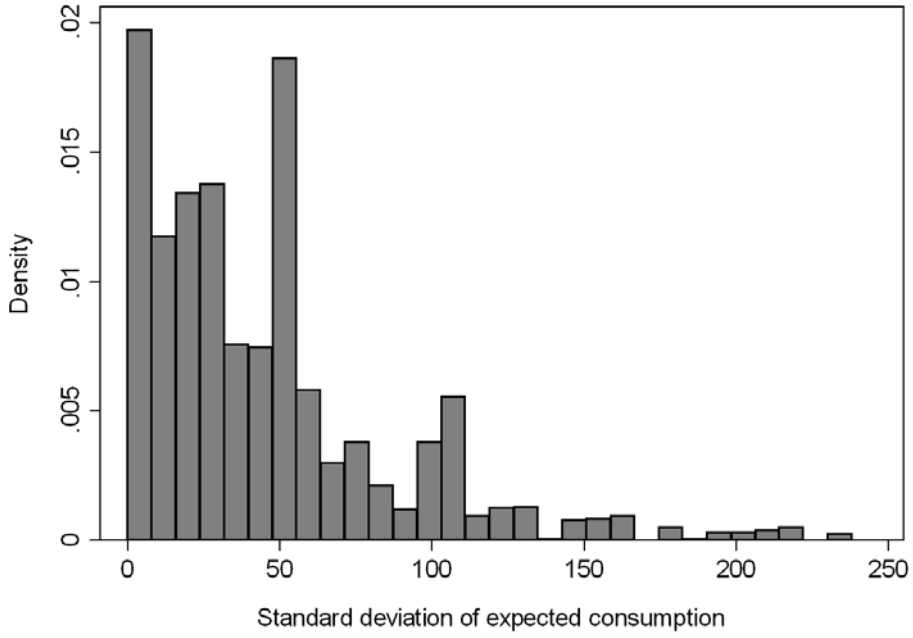


**Fig. 5. Histogram of the standard deviation of expected consumption growth**

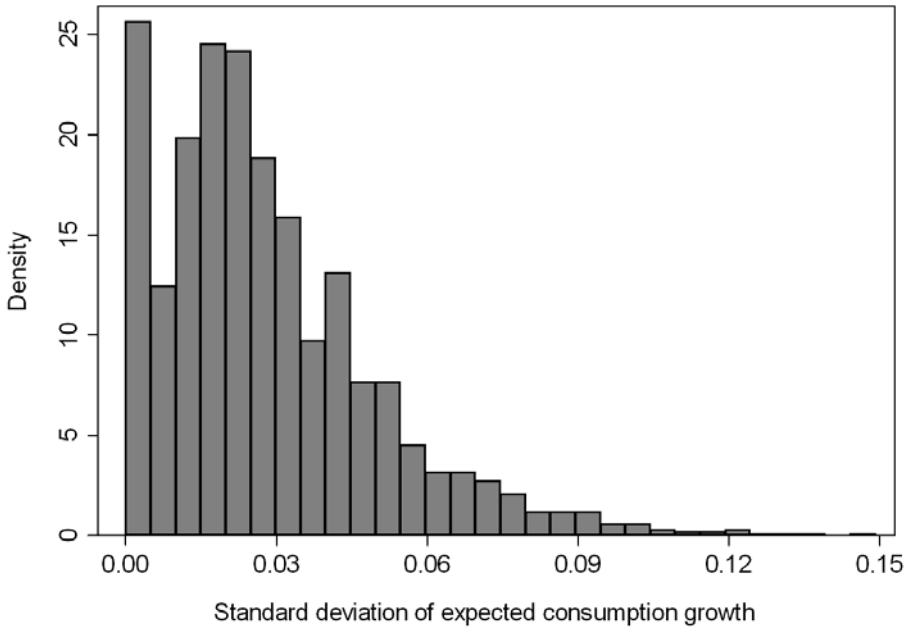



**Fig. 6. Standard deviation of expected consumption growth, by age groups**

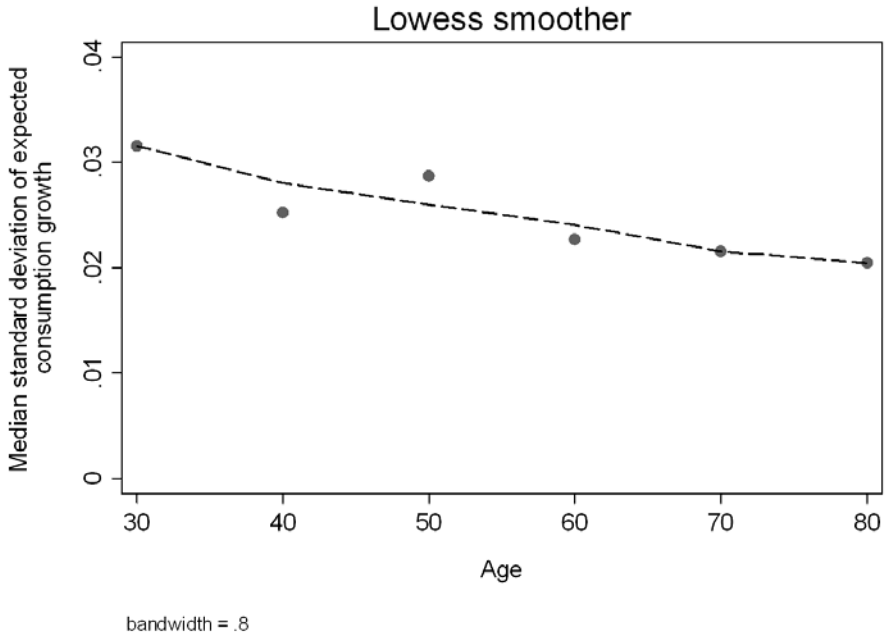


**Fig. 7. Standard deviation of expected consumption growth, by education level**

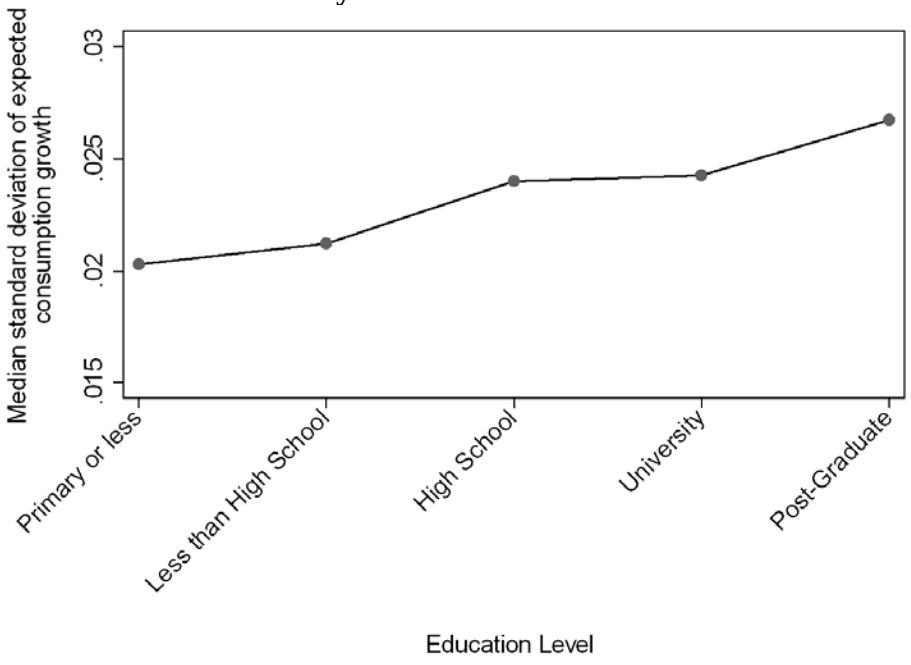
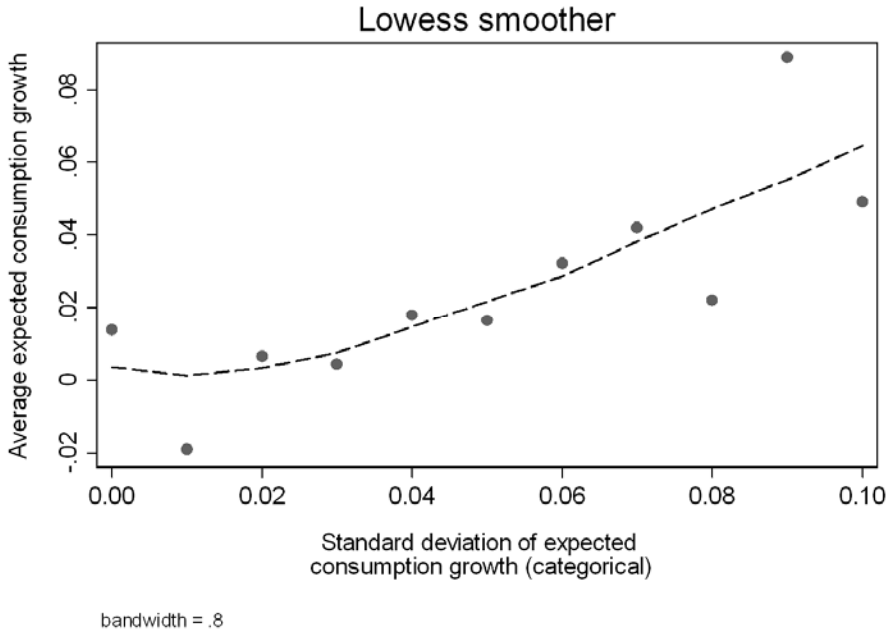



**Fig. 8. Average expected consumption growth, by levels of the standard deviation of consumption growth**



## Table 1. Descriptive Statistics

| Variable | Mean | Median | Std. Deviation |
|---|---|---|---|
| **Minimum expected consumption level** | 1,583.3 | 1,500.0 | 979.2 |
| **Maximum expected consumption level** | 2,020.5 | 1,900.0 | 1,234.9 |
| **Probability that the expected consumption level is above the average of the expected minimum and maximum values** | 0.484 | 0.500 | 0.235 |
| **Expected consumption growth** | 0.012 | 0.000 | 0.088 |
| **Standard deviation of expected consumption growth** | 0.029 | 0.023 | 0.023 |
| **Standard deviation of expected income growth** | 0.015 | 0.008 | 0.028 |
| **Age** | 54.6 | 55.0 | 15.7 |
| **Household size** | 2.4 | 2.0 | 1.2 |
| **Has a partner** | 0.70 | 1.00 | 0.46 |
| **Appears in wave 2** | 0.51 | 1.00 | 0.50 |

**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** | 0.193 | 0.046 | 0.000 | 0.234 | -..- |
| **Standard deviation of expected inflation** | 2.203 | 0.391 | 0.000 | 0.137 | -..- |
| **Age** | 0.000 | 0.000 | 0.000 | -..- | -0.009 |
| **Self-employed** | 0.005 | 0.002 | 0.007 | -..- | 0.232 |
| **Retired** | -0.006 | 0.001 | 0.000 | -..- | -0.275 |
| **Belongs to a union** | -0.003 | 0.001 | 0.042 | -..- | -0.127 |

**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, OLS and robust OLS**

| Variable | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
|---|---|---|---|---|---|---|---|---|---|---|
| | Standard OLS | | | | | Robust OLS | | | | |
| | Coeff. | Std. Error | P value | Lower 95% CI | Upper 95% CI | Coeff. | Std. Error | P value | Lower 95% CI | Upper 95% CI |
| Expected square of consumption growth | 0.419 | 0.188 | 0.026 | 0.051 | 0.787 | 0.951 | 0.021 | 0.000 | 0.909 | 0.994 |
| Age | 0.000 | 0.000 | 0.016 | -0.001 | 0.000 | 0.000 | 0.000 | 0.017 | 0.000 | 0.000 |
| Female | 0.004 | 0.004 | 0.355 | -0.004 | 0.012 | 0.000 | 0.001 | 0.831 | -0.002 | 0.002 |
| Household Size | 0.000 | 0.002 | 0.957 | -0.004 | 0.004 | 0.000 | 0.001 | 0.916 | -0.001 | 0.001 |
| Couple | -0.002 | 0.005 | 0.664 | -0.013 | 0.008 | -0.001 | 0.001 | 0.525 | -0.004 | 0.002 |
| Wave 2 | -0.007 | 0.003 | 0.039 | -0.013 | 0.000 | -0.001 | 0.001 | 0.385 | -0.003 | 0.001 |
| Constant | 0.031 | 0.010 | 0.003 | 0.011 | 0.051 | 0.012 | 0.003 | 0.000 | 0.007 | 0.018 |
| Number of obs. | 2,060 | | | | | 2,060 | | | | |

**Table 4. Euler equation estimates, standard IV and IV using generated instruments**

| Variable | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
|---|---|---|---|---|---|---|---|---|---|---|
| | Standard IV | | | | | IV with generated instruments | | | | |
| | Coeff. | Std. Error | P value | Lower 95% CI | Upper 95% CI | Coeff. | Std. Error | P value | Lower 95% CI | Upper 95% CI |
| Expected square of consumption growth | 1.179 | 0.270 | 0.000 | 0.650 | 1.709 | 0.792 | 0.199 | 0.000 | 0.401 | 1.183 |
| Age | 0.000 | 0.000 | 0.372 | -0.001 | 0.000 | 0.000 | 0.000 | 0.122 | -0.001 | 0.000 |
| Female | -0.002 | 0.006 | 0.752 | -0.013 | 0.009 | 0.000 | 0.005 | 0.949 | -0.010 | 0.009 |
| Household Size | 0.000 | 0.002 | 0.905 | -0.005 | 0.004 | 0.000 | 0.002 | 0.972 | -0.005 | 0.004 |
| Couple | 0.003 | 0.008 | 0.706 | -0.012 | 0.018 | 0.001 | 0.006 | 0.924 | -0.011 | 0.013 |
| Wave 2 | -0.005 | 0.004 | 0.194 | -0.013 | 0.003 | -0.005 | 0.004 | 0.159 | -0.012 | 0.002 |
| Constant | 0.008 | 0.014 | 0.549 | -0.019 | 0.035 | 0.017 | 0.010 | 0.074 | -0.002 | 0.036 |
| Number of obs. | 1,804 | | | | | 1,804 | | | | |
| F-test | 3.643 | | | | | 14.224 | | | | |
| F-test of generated instruments | -..- | | | | | 16.086 | | | | |
| Test of overidentifying restrictions (p value) | -..- | | | | | 0.776 | | | | |
| Test of endogeneity of the expected variability of income (p value) | -..- | | | | | 0.223 | | | | |
| Test of endogeneity of the expectation of the square of consumption growth (p value) | 0.020 | | | | | 0.030 | | | | |

**Table 5. Estimates using partial identification methods**

| Method | Expected Consumption Growth | | | |
|---|---|---|---|---|
| | Lower Bound | Upper Bound | Low 95% CI | High 95% CI |
| **Panel A. Original Estimates** | | | | |
| **Exogenous Treatment Selection** | 0.015 | | 0.007 | 0.023 |
| **No Assumptions Bounds** | -0.642 | 0.658 | -0.648 | 0.663 |
| **MTR** | 0.000 | 0.658 | 0.000 | 0.663 |
| **MTR + XIV (1 instrument)** | 0.029 | 0.492 | 0.016 | 0.531 |
| **MTR + MIV (2 instruments)** | 0.011 | 0.483 | 0.001 | 0.530 |
| **Panel B. Adjusted Estimates** | | | | |
| **Exogenous Treatment Selection** | 0.659 | | 0.293 | 1.024 |
| **No Assumptions Bounds** | -28.591 | 29.261 | -28.843 | 29.514 |
| **MTR** | 0.000 | 29.261 | 0.000 | 29.514 |
| **MTR + XIV (1 instrument)** | 1.283 | 21.903 | 0.692 | 23.640 |
| **MTR + MIV (2 instruments)** | 0.496 | 21.507 | 0.030 | 23.585 |
| **Number of observations (w/o instr.)** | | 2,060 | | |
| **Number of observations (with 1 XIV instr.)** | | 1,804 | | |
| **Number of observations (with 2 MIV instr.)** | | 1,647 | | |

## Appendix

### A.1. Questions on future spending

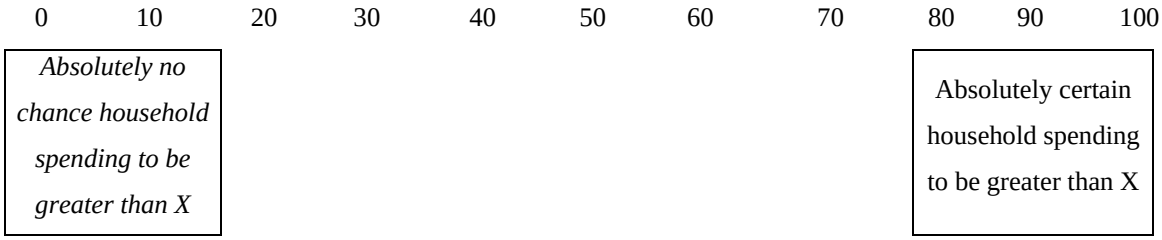
Questions on the future level of spending were asked to the survey respondent in the household as follows:

*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)*




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:

a) Less than 5%
b) 5 - 10%
c) 11 - 15%
d) 16 - 20%
e) 21 - 25%
f) 26 - 30%
g) More than 30%
h) 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 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.

### A.2. The subjective probability distribution of future consumption growth

Let $f(y)$ denote the distribution of 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:

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

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

We consider two assumptions concerning $f(y)$. The first is that *y* is uniformly distributed over each of the two intervals $[y_m, (y_m + y_M)/2]$ and$[(y_m + y_M)/2, y_M]$. If $\pi = 0.5$ the distribution collapses to a single uniform distribution defined in the interval $[y_m, y_M]$. A second possibility is to assume that the distribution is triangular over the same two intervals; if $\pi = 0.5$ the distribution again collapses to a single triangular distribution over the interval $[y_m, y_M]$. Note that in both cases $E(y)$ and $Var(y)$ depend only on the three known parameters ($y_m$, $y_M$, and $\pi$). The triangular distribution is a more plausible description of the probability distribution of consumption growth, because outcomes further away from the mid-point receive less weight. Figures A.1 and A.2 show the p.d.f. under the uniform and the triangular assumptions.

**Figure 1: The uniform distribution**

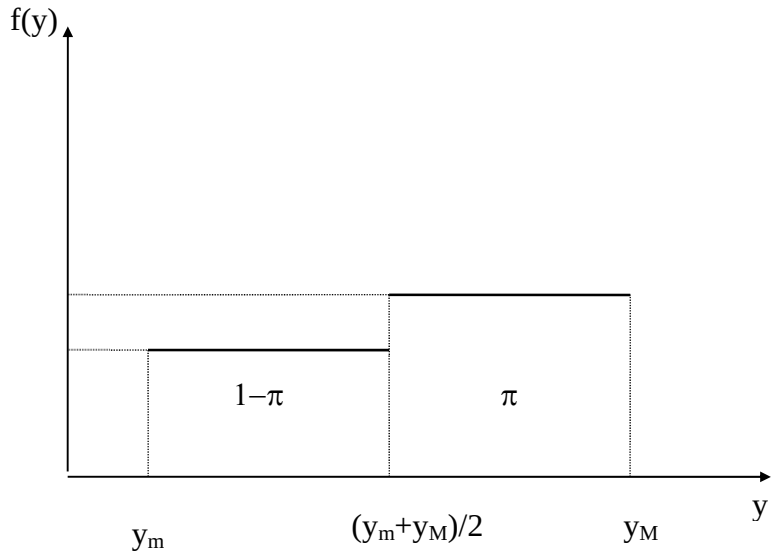



**Figure 2: The triangular distribution**

