

CONSUMPTION UNCERTAINTY AND PRECAUTIONARY SAVING

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

Abstract—Using survey data from a representative sample of Dutch households, we estimate the strength of precautionary saving by eliciting subjective expectations on future consumption. Expected consumption risk is positively correlated with self-employment and income risk and negatively with age. We insert these subjective expectations (rather than 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. Robust OLS and IV estimates indicate a coefficient of relative prudence of around 2. We obtain similar results via partial identification methods using weak assumptions.

I. Introduction

GOING back to Keynes, the effect of uncertainty on consumer behavior is a long-standing topic in research on household saving.¹ Life cycle models of consumption behavior typically imply that higher income uncertainty will increase precautionary saving and consumption growth by lowering current consumption. This increase in saving depends on the third derivative of the utility function and the associated relative prudence coefficient (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. Hence, most tests of precautionary saving use other methods, such as structural models or quasi-experimental approaches. Structural models require a greater number of assumptions than the Euler equation, and so do quasi-experimental methods.

In this paper, we elicit subjective expectations of the distribution of consumption one year ahead. In particular, we

construct measures of expected consumption growth and expected consumption risk using responses to a survey that asked participants about their future consumption. The survey data we use come from the CentER Internet panel, sponsored by the Dutch National Bank and is maintained by CentERdata at Tilburg University and representative of the Dutch population. The measures of expected consumption growth and its variability reflect household-specific subjective expectations that may or may not be fulfilled in future consumption realizations.² In other words, our framework can accommodate deviations from consistent expectations for any number of individuals in our sample, and therefore we do not need to assume that expectation errors in consumption tend to 0 over time for every sample unit.

When we correlate our measure of expected consumption risk with certain demographic and economic characteristics of the sample, we find that expected consumption risk correlates with these characteristics in the direction suggested by both economic theory and intuition. For example, expected consumption risk is higher for the young and the self-employed. Moreover, income risk is positively associated with consumption risk but is not the only determinant of consumption risk. This means that other sources of risk (such as health risk) and institutions (e.g., the pooling of incomes within the family or social insurance programs) are likely to affect consumption risk and the relation between income and consumption risk. Overall, elicited consumption expectations are well aligned with household characteristics that are commonly thought to influence uncertainty about future consumption.

We use these subjective expectations to estimate an Euler equation for consumption. Using expectations-based variables rather than observed magnitudes eliminates the problem of (a) an expectation error in the disturbance term and (b) the endogeneity of the variable denoting observed consumption growth variability, discussed by Carroll (2001) and Bertola et al. (2005). More generally, the paper contributes to a growing literature that uses subjective expectations to elicit individual-level measures of income and unemployment uncertainty (Dominitz & Manski, 1997; Guiso, Jappelli, & Pistaferri, 2002; Jappelli & Pistaferri, 2000), pension uncertainty (Guiso, Jappelli, & Padula, 2013), and interest rate uncertainty (Crump et al., 2015).³ Manski (2004, 2018)

Received for publication May 3, 2017. Revision accepted for publication July 9, 2018. Editor: Yuriy Gorodnichenko.

*Christelis: University of Naples Federico II, CSEF, CFS, and Netspar; Georgarakos: European Central Bank and CFS; Jappelli: University of Naples Federico II, CSEF, and CEPR; van Rooij: De Nederlandsche Bank and Netspar.

We thank the editor and four anonymous referees for helpful comments and suggestions. T.J. acknowledges financial support from the Italian Ministry of University and Research (PRIN project) and the EU Seventh Framework Program (MOPACT project). We thank Efi Adamopoulou, Laura Bartiloro, Tom Crossley, Luigi Pistaferri, Ernesto Villanueva, Gianluca Violante, and Guglielmo Weber; and seminar participants at the 2015 NBER Summer Institute, the ECB Conference on Household Finance and Consumption, the joint Bank de France–Deutsche Bundesbank Spring Conference, the Symposium on Economics and Institutions, the Modena-Netspar Conference on Advances in Household Finance, and CSEF for helpful comments. The opinions expressed in the paper are our own and do not necessarily reflect the views of the European Central Bank or De Nederlandsche Bank.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/rest_a_00819.

¹See, among others, Drèze and Modigliani (1972), Skinner (1988), Deaton (1991), Dynan (1993), Carroll (2001), and Bertola, Guiso, and Pistaferri, (2005).

²Coibion and Gorodnichenko (2012), using survey data from consumers, firms, central bankers, and professional forecasters, show that mean forecasts do not fully adjust to actual shocks and that the implied degrees of informational rigidities are economically large and consistent with significant macroeconomic effects.

³Crump 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 data set includes consumers' expectations of consumption

surveys the literature and discusses the advantages of using measures of expectations in macroeconomics.

When estimating the Euler equation, we find that expected consumption risk is positively associated with expected consumption growth, consistent with intertemporal consumption models with precautionary saving. Using robust OLS regression methods, we find that the implied magnitude of the coefficient of relative prudence is around 2. If the utility function is isoelastic, this implies that the coefficient of relative risk aversion is about 1. These results hold also if we exclude from the sample households likely to be liquidity constrained (to which the Euler equation does not apply) and for various sample splits that allow testing for group differences in the parameters of the utility function.

We also examine the possibility that expected consumption risk is measured with error and correlated with the Euler equation error term. We address this issue by using expected income risk as an instrument, as in Bertola et al. (2005). We find that the IV estimates are similar to the robust OLS ones.

As a final way to check the validity of our results, we relax the assumption of the exogeneity of income risk and allow it to be positively correlated with expected consumption growth. Under this much weaker assumption, we can obtain an identification region for the coefficient of relative prudence using, for the first time in the consumption literature, partial identification methods that are nonparametric and rely on assumptions much weaker than those of OLS and IV. We find that a value of 2 again falls within the identification region of the relative prudence coefficient. Overall, the results for the strength of the precautionary saving motive are robust and plausible.

The paper is organized as follows. Section II surveys the empirical literature on precautionary saving and section III our empirical strategy. Section IV describes the survey data, and section V presents OLS and robust regression results. Section VI extends the analysis to include the presence of liquidity constraints. Section VII reports results from IV estimation, and section VIII discusses partial identification analysis. Section IX concludes.

II. Empirical Tests of Precautionary Saving

Researchers have tested the importance of precautionary saving using three approaches: reduced-form estimation, simulation methods, and estimation of Euler equations for consumption.

A first group of studies attempts to estimate the impact of income risk on consumption or wealth using reduced-form equations. Measures of income risk based on actual earnings are difficult to compute even with long panel data, and in part may reflect a choice (e.g., the choice to work in a risky occupation). Empirical evidence based on this approach is

mixed. Most papers find a positive relation between wealth and income risk, which is consistent with the precautionary saving model. However, the magnitude of the effect varies widely across studies and on net tends to be small (Jappelli & Pistaferri, 2010). This approach provides evidence in favor of or against the size of 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 estimates the paths of consumption and wealth in models with precautionary saving, matching simulated data to observed wealth and consumption distributions. Pioneering this approach, Gourinchas and Parker (2002) use consumption data from the U.S. 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. They estimate, by minimizing the distance between actual consumption and its predicted life cycle profile, a time preference rate of approximately 4% and an elasticity of intertemporal substitution of about 0.5, corresponding to a coefficient of relative risk aversion of about 2. With an isoelastic utility function, the implied coefficient of relative prudence is about 3.

Cagetti (2003), using U.S. data from the PSID and the Survey of Consumer Finances (SCF) and a structural model, finds higher estimates of the time preference rate, a higher coefficient of relative risk aversion (around 4 for the high school sample), and an implied coefficient of relative prudence of around 5. However, structural models deliver estimates of the parameters of the utility function but require that the utility function, the budget constraint, the sources of risks, and the income process be specified.

The third strategy is based on the Euler equation and is the closest to our approach. Dynan (1993) estimates an Euler equation for consumption substituting expected consumption growth and expected consumption risk with their realized counterparts. Substitution of expectations by realizations introduces a forecast error in the Euler equation, which is almost surely correlated with realized consumption growth variability.⁴ To address this endogeneity issue, Dynan uses an IV approach applied to panel data drawn from the CEX. The set of instruments includes education and occupation on the assumption that they are correlated with expected consumption risk and that they affect expected consumption growth only through this channel. Overall, these instruments have low power, and hence the coefficient of relative prudence is imprecisely estimated.

Jappelli and Pistaferri (2000) estimate an Euler equation using Italian panel data from the Survey of Household Income and Wealth (SHIW) that over 1989 to 1993 has measures of subjective income expectations. They find that realized consumption growth is positively correlated with

⁴As discussed in Hayashi (1987), this expectation error should converge to 0 as the time dimension of the data increases, but the same is not true for a short panel. This, as Chamberlain (1984) pointed out, is a serious problem because it leads to inconsistent estimates in short panel surveys such as those that typically contain information on consumption.

growth (but not consumption risk) and inflation, with the latter providing subjective variation in ex ante real interest rates.

the subjective variance of income growth, but their empirical strategy does not deliver an estimate of the degree of prudence.

Dynan's approach was refined by Bertola et al. (2005), who use the subjective variance of income growth available in the SHIW as an instrument for realized consumption variability. As they point out, subjective income risk is a valid instrument because income risk has no direct effect on consumption growth once one conditions on expected consumption risk. In other words, expected consumption risk is a sufficient statistic for expected consumption growth. This implies that all unobservable variables that affect income and its variability are already incorporated into the expected consumption risk term. Examples of such unobservable variables include permanent and transitory income shocks and health and family problems.

The function of expected consumption risk as a sufficient statistic implies that if one conditions on expected consumption risk, one conditions also on all unobservable variables that affect income risk. This makes income risk redundant in the Euler equation and thus a good candidate instrument for consumption risk. Bertola et al. (2005) find that subjective income risk is not only a valid instrument but also one that delivers empirically plausible results. In particular, their coefficient of relative prudence is about 2 and is precisely estimated, thus providing evidence supporting the precautionary saving model.

In this paper, we use an Euler equation approach to estimate the degree of prudence but use measures of expected consumption growth and expected consumption risk based on subjective expectations rather than their realizations. One advantage of our approach is that we do not assume that income risk is the only source of consumption risk, as in other applications. Indeed, in more general models, consumption risk might also reflect uncertainty related to other random variables and thus might not be related only to income risk. Alongside income risk, people face a number of other uninsurable or partially uninsurable risks, which can affect intertemporal consumption decisions. The literature gives prominence to the risks associated with asset price volatility (including house prices), medical and other unexpected expenditures (Palumbo, 1999), family dissolution (Voena, 2015), future liquidity constraints, and the state of the economy.

III. The Empirical Strategy

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

$$u'(c_{i,t}) = \frac{1+r}{1+\delta_{i,t}} E_{it} u'(c_{i,t+1}), \quad (1)$$

as in our empirical specification the discount factor δ depends on demographic characteristics $X_{i,t}$. We assume that the utility function is time separable. In this class of models, the coefficient of relative risk aversion equals the inverse of the EIS.⁵

A second-order Taylor series expansion of $u'(c_{i,t+1})$ around $c_{i,t}$ delivers an expression for the expected growth rate of consumption $E_{it}(g_{i,t+1})$:

$$E_{it}(g_{i,t+1}) = \sigma(c_{i,t}) \left(\frac{r - \delta_{i,t}}{1+r} \right) + \frac{1}{2} p(c_{i,t}) E_{it}(g_{i,t+1}^2) + W_{i,t}, \quad (2)$$

where p denotes Kimball's coefficient of relative prudence, σ denotes the elasticity of intertemporal substitution, and W is a remainder term in the Taylor approximation. The second uncentered moment of the distribution of expected consumption growth $E_{it}(g_{i,t+1}^2)$ is a measure of expected consumption risk.

Equation (2) indicates that an increase in expected consumption risk is associated with higher expected consumption growth. The intuition is that in order to buffer the increase in consumption risk, individuals consume less in period t relative to period $t+1$, and thus increase current saving. Furthermore, the sensitivity of consumption growth to consumption risk is proportional to the coefficient of relative prudence. If utility is quadratic, then $u'''(c_{i,t}) = 0$.⁶ Consequently, a test of the hypothesis that consumption risk does not affect consumption growth is also a test of the validity of the certainty equivalence model. Our parameter of interest is therefore the degree of prudence p .

From an empirical point of view, the main problem related to estimating the Euler equation is that expected consumption growth and expected consumption risk are generally not observable. Were it possible to measure the expectation-related terms in equation (2), it would be also possible to estimate

⁵This assumption is restrictive because there are no good reasons to believe that willingness to substitute consumption across random states of nature should be tightly linked to willingness to substitute consumption deterministically over time. The Epstein-Zin recursive utility model allows greater flexibility with regard to attitudes toward risk and intertemporal substitution, as the EIS is a function of the unobservable continuation value of the future consumption plan. To estimate the Epstein-Zin model, one can rely on aggregate data and exploit the relation between the continuation value and the return on the aggregate wealth portfolio (Campbell, 1996; Chen, Favilukis, & Ludvigson, 2013). On the other hand, in microsurvey data, questions on capital income are very difficult for households to answer, and thus the underlying concepts are typically not measured well, even in well-established surveys like the U.S. Survey of Consumer Finances.

⁶In the presence of 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 & Singleton, 1983).

the coefficient of the term related to expected consumption risk, which would be proportional to relative prudence.

There are two strategies to estimate the degree of prudence. Dynan (1993) and Bertola et al. (2005) replace expectations with their realized counterparts and assume that the elasticity of substitution and the degree of prudence are constant. In this case, equation (2) can be written in a regression framework as

$$g_{i,t+1} = \alpha + \beta g_{i,t+1}^2 + \gamma' X_{i,t} + f_{i,t+1} + \varepsilon_{i,t+1}, \quad (3)$$

where $\gamma' X_{i,t}$ approximates the discount factor with demographic variables, $f_{i,t+1}$ is the difference between realized and actual consumption (the forecast error), and $\varepsilon_{i,t+1}$ is a composite error term that includes higher-order terms of the Taylor expansion—the $W_{i,t}$ term in equation (2)—unobserved heterogeneity, and measurement error of actual spending. The coefficient β is directly related to the strength of the precautionary saving motive.

The forecast error $f_{i,t+1}$ is clearly correlated with $g_{i,t+1}^2$. For instance, if households have positive news about the economy between periods t and $t+1$, they may revise consumption upward in period $t+1$, affecting both the mean and the variance of the (ex post) consumption distribution. Since in general $\text{cov}(f_{i,t+1}, g_{i,t+1}^2) \neq 0$, equation (3) needs to be estimated by IV, relying on instruments that are correlated with consumption risk but not with the forecast error.

An alternative strategy is to measure directly expected consumption growth and expected consumption risk and estimate

$$E_{it}(g_{i,t+1}) = \alpha + \beta E_{it}(g_{i,t+1}^2) + \gamma' X_{i,t} + v_{i,t}. \quad (4)$$

Estimating equation (4) rather than equation (3) has several advantages with respect to previous tests of precautionary saving. First, the error term of equation (3) includes the expectation error of the Euler equation, while by construction, the error term of equation (4) does not. Hence, assuming $\text{cov}[E_{it}(g_{i,t+1}^2), v_{i,t}] = 0$, the error is not correlated with expected consumption risk. We discuss this assumption further below.

A second, related issue is that one can estimate equation (4) with even a cross-section or a short panel by exploiting the cross-sectional variability in expectations of the consumption distribution. The literature shows that Euler equation estimates derived from panel data may be inconsistent when the time dimension of the panel is short (Chamberlain, 1984; Hayashi, 1987). The reason is precisely that the error term of equation (3) includes a forecasting error. The expectation of a forecasting error, conditional on any information available at t , should be 0 over a long horizon. In other words, the error should not exhibit systematic patterns if the model is correct. Following this logic, the empirical equivalent of $E_{it}(f_{i,t+1})$ in equation (3) is a household-level average taken over many periods. However, panel surveys containing information on consumption are typically short, and hence researchers often

proceed under the assumption that consistency is achieved assuming that forecasting errors average out to 0 in the cross-section. There is no reason to believe that this assumption holds generally. For example, if there are aggregate shocks, households will likely make forecasting errors in the same direction in a given year (Altug & Miller, 1990). In this case, the cross-sectional average of the forecasting error will most likely be different from 0.⁷

Third, given that the Euler equation (4) does not require that consumption expectations be aligned with realizations, systematic deviations of expectations from realizations for some individuals in our sample do not affect the consistency of our estimates. In other words, we do not need to assume that consumption realizations match subjective expectations. This makes our estimates more robust than those based on realizations, which require that expectation errors should average 0 over time for all sample units. While our estimation strategy does not require alignment of expectations and realizations, we report in section IV that subjective expectations have plausible sample distributions and are associated with observables in ways that conform to economic theory and intuition. Importantly, the use of expectations implies that measurement error that arises from differences between reported and actual expenditures is not relevant in our case, as we do not make any use of consumption realizations.⁸

The fact that expected consumption risk is a sufficient statistic for expected consumption growth implies that the former encapsulates the effect of all observable and unobservable factors that determine the various risks a household faces. Therefore, conditioning on expected consumption risk implies that the assumption that $v_{i,t}$ has a 0 conditional expectation in the cross-section is plausible. This implies that one can use OLS to estimate equation (4). Nevertheless, the error term might contain higher-order terms of the Taylor expansion that are correlated with expected consumption risk. Hence, we check the robustness of the results by using an IV estimator that relies on expected income risk as an instrument (section VII), and nonparametric methods that partially identify the relative prudence coefficient using weak assumptions (section VIII).

IV. The Data

We use data from the CentER Internet panel. The baseline survey is conducted once a year via the Internet and collects detailed information on a range of demographics and asset holdings for a representative sample of Dutch-speaking

⁷One way to overcome this problem is to add year dummies to the Euler equation. However, this approach still might fail to deliver consistent estimates if the aggregate shocks are unevenly distributed across consumers, so that the time dummies do not completely absorb their impact.

⁸Measurement error of a different sort can arise if survey respondents do not report their true expectations, although we are not aware of any evidence that this is a common problem. Moreover, there could be some measurement error if, for instance, respondents make mistakes when inputting numbers in the online questionnaire. As discussed in section VII, we use IV to address this issue.

households. In addition to the baseline survey, it is common for households to be asked to participate during the course of a year in special-purpose surveys.

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

To elicit the distribution of expected consumption, we follow a procedure similar to Guiso et al. (2002, 2013), who estimate the subjective distribution of future income and the pension replacement rate, respectively. Specifically, we asked respondents first to report the minimum (y_m) and the maximum (y_M) values of next year's consumption in a typical month, and then to rank on a scale of 0 to 100 the probability that consumption will be higher than the midpoint between the minimum and the maximum, that is, $\pi = \text{Prob}(y \geq (y_m + y_M)/2)$. We provide in online appendix A.1 the questions used to capture expected consumption uncertainty (as well as all other relevant concepts).⁹

To estimate the moments of the subjective distribution of future consumption, we assume that the subjective distribution is either simple triangular (i.e., symmetric around $(y_m + y_M)/2$, assuming $\pi = 0.5$) or split triangular ($\pi \neq 0.5$; see figure A.1. in the online appendix). Based on the elicited values for y_m , y_M (and for π if we assume a split triangular distribution), we compute the household-specific mean and standard deviation of the distribution of expected consumption one year ahead. The formulas of these statistics are reported in online appendix A.2.¹⁰

We set to missing values observations where y_m , y_M , or π is missing or if respondents chose the “do not know” option. The original sample includes 4,323 observations in the three survey waves. Due to missing values, the estimation sample consists of 3,271 household-level observations for the simple triangular distribution and 3,167 observations for the split triangular distribution.

The survey also asked households to directly report the expected change in their spending one year ahead. In particular, they were asked first to think about spending on all goods and services in the coming twelve months and report whether it

would be higher, about the same, or lower than their current spending. Subsequently, they were asked to report the expected change in spending in percentage terms.

Using information on each household's expected consumption growth and minimum and maximum levels of consumption one year ahead, it is straightforward to compute a household-specific expected variance, standard deviation, and expected square of consumption growth. In particular, $E(g_{i,t+1}^2) = f(y_{m,i}, y_{M,i}, \pi_i)$ is the last term that appears in equation (4), and it is the one used in estimation. Results are essentially unchanged between the split and simple triangular distributions, and thus we report in the paper results only for the split triangular, while in the online appendix we report those for the simple triangular. We report results derived using both distributions only in table 5, as in that case they slightly differ across the two distributions.

The survey also asked for information that enables the computation of the moments of the distribution of income one year ahead. Specifically, households provide minimum and maximum values of annual household income net of any taxes for the next twelve months and indicate the probability that income will be higher than the midpoint 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.

The left and right upper panels of figure 1 report the distribution of the expected minimum and maximum levels of consumption twelve months ahead, respectively. For each observation in the sample, the maximum is greater than the minimum. The lower left panel of figure 1 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%” responses but also a sizable number of respondents reporting values larger or smaller than 50%. Notice that the question on this probability, which is arguably more difficult to answer, is not used in the regressions based on the simple triangular distribution.

Table 1 reports cross-sectional statistics of the central tendency and dispersion of the subjective distribution of expected consumption 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 1,400 euro, while the maximum is 1,750 euro (the means are equal to 1,484 euro and 1,882 euro, respectively), and the median probability is 0.5 (average 0.48). Assuming that the distribution is split triangular, we estimate that the sample median of expected consumption growth is 0 (average 1.8%). Since forecasts in the Netherlands indicate that in 2014, consumption expanded slightly (by approximately 0.2%), consumption expectations, as summarized by sample means and medians, seem aligned with aggregate realizations.¹¹

⁹There is always a trade-off between the accuracy of economic concepts and question design. In this case, the question does not distinguish between consumption and spending and refers to a “typical month.” In principle, one could ask more detailed questions, distinguishing between durable and nondurable consumption and referring to specific months of the year. However, we feel that this would increase the complexity of the questionnaire and ultimately reduce the reliability of responses.

¹⁰We assume that for each individual, y_m and y_M , represents the actual minimum and maximum of the distribution. This is potentially a strong assumption. Dominicz and Manski (1997) use the percentage chance format to elicit the subjective income distribution and show that individuals associate low probabilities to the “lowest possible” (and “highest possible”) value.

¹¹It should be noted that aggregate consumption growth may differ from the cross-sectional average of households' consumption growth rates (means or medians). This difference depends, among other things, on the consumption distribution. Still, we find reassuring that the survey average

FIGURE 1.—THE DISTRIBUTION OF EXPECTED CONSUMPTION




TABLE 1.—DESCRIPTIVE STATISTICS

Variable	Mean	Median	Std. dev.
Minimum expected consumption level	1,484.2	1,400.0	912.7
Maximum expected consumption level	1,882.2	1,750.0	1,147.8
Probability that the expected consumption level is above the average of the expected minimum and maximum values	0.476	0.500	0.228
Expected consumption growth	0.018	0.000	0.096
Std. dev. of expected consumption growth	0.049	0.040	0.041
Expected square of consumption growth	0.016	0.003	0.071
Std. dev. of expected income growth	0.028	0.016	0.043
Age	52.0	52.0	16.1
Female householder	0.41	0.00	0.49
Household size	2.20	2.00	1.25
Has a partner	0.60	1.00	0.49
Number of observations	3,167		

Zarnowitz and Lambros (1987) and Manski (2018) point out that the cross-sectional dispersion of point predictions provides little information about individual uncertainty. We provide further simulation-based evidence to that effect in online appendix A.3, where we show that for values of the minimum and maximum consumption close to the median, the standard deviation of the cross-sectional distribution of realized consumption is about 39% (24%) larger than the

and median growth rates are close to the aggregate one. Notice also that one should not expect the variance of the actual cross-sectional consumption distribution to match the variance of the subjective distribution, because dispersion of realized outcomes does not necessarily correlate with the risk perceived by respondents (see Manski, 2018).

average standard deviation of the household-level expected consumption distributions, when the latter are simple triangular (uniform).

The above implies that comparing the distribution of individual consumption or income uncertainty to the cross-sectional distribution of actual income and consumption is not appropriate.¹² On the other hand, a good way to validate the expectation data in our survey is to compare them with subjective expectation data from other surveys. Although measures of consumption risk are not available in the literature, for the Netherlands we can compare our measures of income risk with those of Das and Donkers (2005). They estimate income risk using the same method as we do, and their estimated distribution of income risk is quite close to ours.¹³

Notice also that the median (mean) standard deviation of the distribution of expected consumption growth is 4% (4.9%). The median (mean) expected square of consumption growth is 0.3% (1.6%). This is the variable that appears on the right-hand side of equation (4).

Cross-sectional averages are useful to describe the subjective consumption distribution of a typical household, but they

¹²If one had long panel data, one could calculate the observed consumption or income variability for each household, which could in turn be compared to that household's corresponding expected variability.

¹³Das and Donkers (2005, table 4) document that 18% of respondents report a 0 coefficient of variation of income risk, another 40% greater than 0 and less than 0.025, 24% between 0.025 and 0.065, and 18% above 0.065. These figures are close to the corresponding ones in our survey, which are as follows for the respective coefficient of variation intervals: 18%, 39%, 22%, and 21%.

FIGURE 2.—EXPECTED SQUARE OF CONSUMPTION GROWTH, BY AGE GROUPS




TABLE 2.—CORRELATIONS OF MEASURES OF CONSUMPTION UNCERTAINTY WITH OTHER MAGNITUDES

Variable	(1)	(3) Std. Dev. of Expected Consumption Growth			(5)	(7) Expected Square of Consumption Growth		
	Coefficient	SE	p-Value	Standardized Change	Coefficient	SE	p-Value	Standardized change
Std. dev. of expected income growth	0.2623	0.0363	0.0000	0.2772	0.0377	0.0062	0.0000	0.0228
Age/100	-0.0375	0.0044	0.0000	-0.0092	-0.0039	0.0006	0.0000	-0.0005
Self-employed	0.0062	0.0037	0.0909	0.1527	0.0013	0.0005	0.0112	0.0184
Retired	-0.0116	0.0014	0.0000	-0.2859	-0.0011	0.0001	0.0000	-0.0157
Is a union member	-0.0056	0.0025	0.0241	-0.1378	-0.0010	0.0003	0.0032	-0.0137
Household size	0.0017	0.0006	0.0048	0.0415	0.0002	0.0001	0.0318	0.0022

Columns 1–3 (5–7) report results from robust regressions using Huber's (1973) M-estimator in which the dependent variable is the standard deviation of expected consumption growth (expected square of consumption growth) and the only regressor is the variable shown in each line. The standardized change shown in column 4 (8) refers to the change (in standard deviations) in the standard deviation of expected consumption growth (in the expected square of consumption growth) induced by the following change in the regressors: (a) the standard deviation of expected income growth increases by 1 standard deviation, and (b) the remaining variables increase by one unit (age increases by one year). Standard errors are computed using 1,000 bootstrap replications.

hide important heterogeneity across households. The lower right panel of figure 1 plots the histogram of the distribution of the 3,167 household-specific measures of consumption risk—namely, the expected square of future consumption growth. We note the considerable heterogeneity in the responses. For instance, for 25% of households, the measure of consumption risk is less than 0.1 square percentage points; for another 50%, it is between 0.1 and 0.9 square percentage points; and for the top 25% it is more than 0.9 square percentage points. The proportion of households for which the expected square of consumption growth is 0 is 6.6%, while the corresponding proportion for which the standard deviation of expected consumption growth is 0 is 13.5%.

As discussed, a novel feature of our analysis is that it makes use of directly elicited expectations on consumption growth and associated uncertainty to estimate an Euler equation in its original form. Subjective expectations are not necessarily aligned with consumption realizations, but one could still check how expected consumption risk correlates with socioeconomic variables.


Figure 2 plots the median standard deviation of the expected square of consumption growth distribution by ten-year age bands. It indicates that consumption risk declines during

the life cycle; the expected square of consumption growth falls by about 2 percentage points (the pattern is very similar if we plot the standard deviation of expected consumption growth against age). This finding suggests that younger households perceive a higher uncertainty than older ones, in line with the findings in Dominitz and Manski (1997) for the subjective distribution of income uncertainty.¹⁴ Notice that the age gradient might also capture cohort effects, so figure 2 might signal that younger cohorts face higher uncertainty regardless of age. Unfortunately, our survey does not provide enough information to distinguish between these two explanations.

Table 2 reports associations of two measures of consumption uncertainty: the standard deviation of expected consumption growth and the expected square of consumption growth (i.e., the term $E_{it}(g_{i,t+1}^2)$ in equation [4]), with age, expected income growth (constructed in similar fashion), self-employment, retirement status, union membership (as a further measure of income volatility), and household size.

¹⁴ Another possible explanation is that the young report higher consumption uncertainty because they have less knowledge and experience in evaluating the likelihood of future economic events.

FIGURE 3.—EXPECTED CONSUMPTION GROWTH, BY GROUPED VALUES OF EXPECTED SQUARE OF CONSUMPTION GROWTH



These associations are derived from robust regressions (using the M-estimator in Huber, 1973) of the two measures of consumption uncertainty on each of the variables. Since the expected square of consumption growth is a generated variable, all standard errors in the descriptive statistics and in the regressions reported in the paper are computed using 1,000 bootstrap replications.

Both measures of consumption uncertainty are strongly correlated with expected income risk but at much less than one-to-one, showing that other factors besides income risk affect consumption risk. Another reason to expect that consumption and income risk are not perfectly correlated is that under the permanent income hypothesis, consumption risk should reflect only permanent, not transitory, income risk. Consumption uncertainty is also correlated with self-employment, especially in the case of the expected square of consumption growth. The direction of this association is as expected, given that the self-employed typically face a higher-than-average income risk, which should lead to higher consumption uncertainty.¹⁵

Age and being retired are negatively associated with consumption uncertainty, as we expect given the reduced income uncertainty associated with older age. Being a union member may imply more predictable wages and thus lowers consumption uncertainty. Finally, consumption uncertainty increases with household size, possibly because larger households are exposed to larger expenditure shocks.

Overall, the fact that these associations of consumption uncertainty have the expected sign, are sizable and are also statistically significant suggests that the survey measures of subjective expected consumption uncertainty represent rea-

sonably well the actual consumption uncertainty that the households in our sample face. The associations are generally stronger for the standard deviation of expected consumption growth than for the expected square of consumption growth, likely because the former is the more comprehensive measure of consumption uncertainty.

While these findings are reassuring for the informational quality of our expectation measures, we note that what is sufficient for our estimation is a measure of the expected consumption uncertainty perceived by households, which may or may not be consistent with the actual consumption uncertainty they face.

V. OLS and Robust Regression Estimation

In this section, we estimate equation (4) using data on expected consumption growth and expected consumption risk. The vector X includes age and gender of the household financial respondent, whether she or he has a partner, household size, indicators of survey wave, and regional dummies. The demographic variables are included in the specification to capture preference heterogeneity. Note that β may depend on demographic variables. We explore this issue by estimating the Euler equation on sample splits defined by different socioeconomic groups.

Before presenting econometric results, in figure 3 we plot $E_{it}(g_{i,t+1})$ against binned values of $E_{it}(g_{i,t+1}^2)$.¹⁶ The two variables are strongly positively correlated, and the slope of the relation between the two is slightly more than 1, with an implied coefficient of relative prudence of slightly higher than 2. As we shall see, our estimates are consistent with this descriptive evidence. A similar pattern emerges when plotting $E_{it}(g_{i,t+1})$ against the expected square of income

¹⁵Dillon (2018), 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.

¹⁶The bins are defined using the deciles of the distribution of the expected square of consumption growth.

TABLE 3.—EULER EQUATION ESTIMATES: OLS AND ROBUST REGRESSION, SPLIT TRIANGULAR DISTRIBUTION

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	OLS			Robust regression								
	Whole Sample			Whole Sample			Household Income above the 20th Quantile and Employed			Net Financial Assets Value Larger Than Two Months of Income		
Variable	Coefficient	SE	p-value	Coefficient	SE	p-value	Coefficient	SE	p-value	Coefficient	SE	p-value
Consumption risk	0.642	0.127	0.000	0.966	0.050	0.000	0.961	0.204	0.000	0.959	0.181	0.000
Age	0.000	0.000	0.403	0.000	0.000	0.880	0.000	0.000	0.896	0.000	0.000	0.295
Female householder	0.002	0.003	0.594	0.000	0.001	0.645	−0.001	0.001	0.234	−0.002	0.001	0.147
Household size	0.001	0.002	0.424	0.001	0.001	0.272	0.001	0.001	0.232	0.001	0.001	0.199
Couple	−0.005	0.004	0.250	−0.001	0.001	0.412	0.001	0.002	0.622	−0.002	0.002	0.280
Constant	0.014	0.008	0.083	0.007	0.003	0.021	0.006	0.003	0.086	0.004	0.004	0.336
Number of observations	3,167			3,167			2,515			1,919		

Robust regressions use Huber's (1973) M-estimator. Standard errors are computed using 1,000 bootstrap replications. In addition to the variables shown, all specifications include regional and survey wave dummies. We report the corresponding results using the simple triangular distribution in table A.1 in the online appendix.

growth (see figure A.2 in the online appendix), that is, another, less comprehensive, measure of household uncertainty.

In order to reduce the influence of outliers, we winsorize both g and g^2 at the top and bottom 0.5% of the observations; that is, we set the values of those observations equal to those at the 99.5th and 0.5th percentiles, respectively. We also use Huber-White robust standard errors and bootstrap standard errors in all estimation methods. As a number of households participate in two or three survey waves, we use a clustered bootstrap by resampling household identification numbers.

We first estimate equation (4) by conventional OLS. We proxy the discount factor—the X matrix in equation (4)—with standard demographic variables (age, family size, and female householder). We omit from the baseline specifications labor supply variables (such as retirement status and self-employment) because they are potentially correlated with the error term of the equation, as Attanasio and Weber (1995) argue.

Columns 1 to 3 in table 3 report the OLS results using the split triangular distribution for expected consumption risk. The estimated coefficient of consumption risk is 0.64 and highly statistically significant (p -value $< .01$), implying a prudence coefficient of about 1.28. The coefficients of age, female householder, and household size are positive but imprecisely estimated.

A common alternative to OLS that limits the influence of outliers is Huber's (1973) M-estimator. Results from this estimation method are shown in columns 4 to 6 in table 3. The estimated coefficient of expected consumption risk is 0.97 (p -value $< .01$), a value that is larger than the corresponding OLS one and implies a prudence coefficient of about 2. As this estimate of prudence is robust to outliers, we consider it more reliable than the OLS estimate.¹⁷

Past precautionary behavior may bias the estimate of prudence upward. Households with large amounts of cash on hand might have already engaged in precautionary saving

and thus might exhibit lower expected consumption risk and higher expected consumption growth. However, if we introduce wealth as an additional regressor, results do not change.

To check whether our estimate of the effect of consumption risk on expected consumption growth differs by household characteristics, we split our estimation sample between those below and above 50 years old, the retired and nonretired, singles and couples, and those with and without college education. In all cases, the coefficient of consumption risk is close to 1, that is, essentially the same as in the full sample.

Finally, as discussed in section IV, about 14% of respondents report no expected consumption uncertainty. Such responses could be legitimate but could also be due to misperceptions of the survey questions. As a robustness check, we estimate the Euler equation excluding respondents with no consumption risk. Results are not affected.

VI. Liquidity Constraints

The Euler equation estimated in section V is derived assuming perfect capital markets. However, the equation fails in the presence of liquidity constraints or myopic consumers. Let us consider a simple alternative model, where consumption equals income in each period. Then, expected consumption growth equals expected income growth in each period, and consumption risk plays no role. This suggests that our estimates might be contaminated by the presence of some households that may not engage in precautionary saving. From an econometric point of view, this is an omitted variable problem that might bias the coefficient of interest, that is, the sensitivity of expected consumption growth to expected consumption variability.

In order to address this, we present in table 3 the results of robust regressions that exclude from the estimation sample households that possibly are liquidity constrained and thus less likely to engage in precautionary saving. We distinguish liquidity-constrained households based on two different measures.

Columns 7 to 9 of table 3 present results when we exclude from the sample households whose household head is

¹⁷Trimming the sample at 0.5% or 1% delivers OLS estimates of the expected consumption risk coefficient that are close to the estimate that one obtains using the M-estimator. Results using the simple triangular distribution are similar and are reported in table A.1, columns 1–6 in the online appendix.

unemployed and those in the bottom quintile of the disposable income distribution (resulting in 671 households, or 20.5% of the estimation sample being dropped). In addition, in columns 10 to 12, we show results from a sample of households whose net financial assets are larger than two months of disposable income.¹⁸

The coefficient of expected consumption risk is around 1 for both sets of estimation results, confirming the baseline results in columns 1 to 6 for the whole sample.¹⁹ Thus, we can conclude that our baseline estimates are unlikely to be affected by the presence of liquidity-constrained households in our sample.

Note that we are not suggesting that low income or wealth is necessarily associated with liquidity constraints, but rather that high-income or high-wealth households are unlikely to be liquidity constrained. What is important for our purposes is that the estimation sample (the high income/wealth sample) is not contaminated by the presence of liquidity-constrained households, not that the sample of low-income or low-wealth households does not contain unconstrained ones. In fact, it may well be the case that there are many unconstrained households in the low-income, low-wealth samples (e.g., they have both access to credit markets and low wealth because they choose to borrow). As a result, the effect of consumption risk in these samples might be similar to the one in the whole sample. In other words, we do not claim that we perform a test of liquidity constraints, and we do not specify an alternative model that should hold in their presence.

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

VII. IV Estimation

The use of elicited expectations in the estimation of equation (4) circumvents serious econometric issues affecting existing studies that base inference on consumption realiza-

¹⁸We obtain information on financial assets and liabilities from the main wave of the DNB survey, conducted in April 2014 and 2015. Some households in the DNB Internet Panel Survey did not participate in the main DNB survey, and thus the sample for which financial assets information is available is smaller than that used in our main specifications. We obtain similar results if we use a one- or three-month income threshold for net financial assets to define the sample split. The one-month income threshold is used by Kaplan, Violante, and Weidner (2014) to distinguish between wealthy and poor hand-to-mouth households.

¹⁹Expected consumption growth and consumption risk are calculated using the split triangular distribution. Results do not change when using the simple triangular distribution (see table A.1, columns 7–12 in the online appendix).

TABLE 4.—EULER EQUATION ESTIMATES—IV ESTIMATION, SPLIT TRIANGULAR DISTRIBUTION

Variable	(1) Coefficient	(2) SE	(3) <i>p</i> -Value
Consumption risk	0.837	0.463	0.071
Age	0.000	0.000	0.915
Female householder	0.000	0.004	0.959
Household size	0.001	0.002	0.534
Couple	−0.003	0.006	0.544
Constant	0.002	0.013	0.855
Number of observations	2,791		
<i>F</i> -test	3.967		
Anderson-Rubin test of significance of the treatment variable (consumption risk)— <i>p</i> -value	0.016		
Test of endogeneity of the treatment variable (consumption risk)— <i>p</i> -value	0.323		

In addition to the variables shown, all specifications include regional and survey wave dummies. The Anderson-Rubin test is a test of the significance of consumption risk. The endogeneity test is a Hausman test of endogeneity of consumption risk. Standard errors are computed using 1,000 bootstrap replications. We report the corresponding results using the simple triangular distribution in table A.2 in the online appendix.

tions. In particular, the use of expectations implies that the error term ε is not a forecasting error, as is usually the case in Euler equation estimates. Nonetheless, there is still the possibility that unobservable variables in the error term ε (e.g., higher-order terms of the Taylor expansion) are correlated with expected consumption risk or that individuals make errors when completing the online questionnaire. Hence, as a robustness check, we estimate equation (4) also using IV methods to take account of possible endogeneity problems and measurement error.

Our instrument is expected income risk (constructed similar to expected consumption risk). This variable is used by Bertola et al. (2005) as an instrument for realized consumption volatility. As discussed in section III, it is a good instrument in an Euler equation framework, given that it does not appear in equation (4) when expected consumption risk is included. In other words, it is not correlated with unobservable variables that are incorporated in consumption risk (including income shocks and personality traits) once the latter is included in the specification. Moreover, income risk should be positively correlated with consumption risk.

Importantly, IV estimation allows for preference heterogeneity. In the presence of homogeneous preferences, the estimated IV coefficient is equal to the population prudence coefficient. When preferences are heterogeneous, however, the IV coefficient equals the local average treatment effect. Hence, it captures the average of the heterogeneous prudence coefficients among those whose expected consumption risk changes when their expected income risk changes, that is, among “compliers” (see Angrist, Imbens, & Rubin, 1996). Compliers should form a large part of the population, because many households experience a change in expected consumption risk when income risk changes.

We show IV estimation results in columns 1 to 3 of table 4, noting that the sample is about 12% smaller than the one used in the conventional OLS and robust regressions shown in table 3 due to missing values of the income risk variable.

The estimated effect of expected consumption risk on expected consumption growth is 0.84 and strongly significant (p -value = 0.071). Moreover, it is very similar to the robust regression estimate. The first-stage regression confirms that expected consumption risk is positively correlated with expected income risk. Nevertheless, the F -statistic is about 3.97, and thus below the rule-of-thumb threshold of 10 generally recommended for dependable inference. The endogeneity of consumption risk can be tested using a standard Hausman test, and the resulting test statistic has a p -value equal to 0.35. This implies that the expected consumption risk (i.e., the covariate of interest) is unlikely to be affected by endogeneity problems, with the caveat that the test might not be reliable given the weak instrument problem.

Given that the F -statistic value from the first-stage regression is rather weak, we test for the significance of the consumption risk term by using the Anderson-Rubin (1949) statistic that is robust to weak instruments (see Stock, Wright, & Yogo, 2002).²⁰ We find that the null of no statistical significance can be rejected at the 5% level.

As when using OLS and robust regressions, we redo our estimation using the simple triangular distribution. Results are presented in table A.2 in the online appendix and are essentially identical to those obtained using the split triangular distribution.

VIII. Partial Identification

Theory suggests that consumption risk encapsulates all relevant factors affecting expected consumption growth. Consequently, income uncertainty should not affect the latter, conditional on consumption risk. In this section, however, we check whether OLS and IV results are still plausible when consumption risk is endogenous and income risk weakly positively correlated (given consumption risk) with expected consumption growth.

To estimate the strength of the precautionary saving motive under these two weaker assumptions we use the partial identification (PI) method introduced by Manski (1990, 1994). PI is nonparametric and produces bounds on the average treatment effect (ATE); it locates it in an identification region instead of producing a point estimate. Furthermore, it has various important advantages over OLS and IV methods. In what follows, we give a brief overview of the use of PI methods in our context and provide additional details in online appendix A.4.

PI methods apply bounds to the counterfactual, and thus unobservable, average potential outcomes across sample units. These outcomes are obtained under a value of the treatment other than the observed one. To put bounds around these unobserved outcomes, PI uses some weak assumptions compared to OLS and IV methods.

The first assumption is that of monotone treatment response (MTR; see Manski, 1997), which states that expected consumption growth is weakly increasing in consumption risk on average. This is a reasonable assumption, as increasing consumption risk is unlikely to have negatively affected precautionary saving on average. In fact, theory predicts the opposite. At most, consumption risk may have no effect on precautionary saving on average, and this is fully allowed by the MTR assumption.

The second assumption is the monotone instrumental variable (MIV) one, introduced by Manski and Pepper (2000), which serves to narrow the identification region of the ATE. Under this assumption, instrumental variables are assumed to be only weakly monotonically correlated with expected consumption growth (i.e., the outcome variable). We use as MIV expected income risk, which, as a source of consumption risk, is expected to be positively associated with precautionary saving. This weak positive correlation of expected income risk with precautionary saving significantly relaxes the exogeneity assumption that lies behind IV estimation, even if this latter assumption cannot be rejected, as discussed in section VII. The use of a MIV involves no assumptions about the correlation between the MIV and the treatment (i.e., consumption risk).

The MIV assumption refers to the correlation of expected income risk with the average outcome if all sample units received a given value of the treatment (i.e., consumption risk). As sample units actually get different treatment values, the MIV assumption cannot be tested, just like the assumptions underlying standard IV estimation. In our sample, income risk is positively correlated with expected consumption growth (the correlation coefficient is 0.0987, significant at 1%). This correlation is not a proof of the validity of the MIV because it is computed using observed outcomes rather than potential ones. However, it does suggest that the MIV assumption is reasonable.

We contrast the PI estimates with those obtained under exogenous treatment selection (ETS), which posits that respondents receiving different treatments are not systematically different from one another. In other words, ETS implies that income risk is essentially randomly assigned across households.

To implement PI, we discretize the treatment variable, that is, the square of expected consumption growth, so as to bound the unobserved average potential outcome when evaluated at different treatment values. We thus divide our treatment variable into terciles and evaluate the bounds of the average potential outcomes within each tercile. It turns out that the combination of the MTR and MIV assumptions provides an informative lower bound on the ATE of a change in the treatment from its first to its third tercile because in this case, all counterfactual terms can be replaced by observed terms, and thus uncertainty is reduced. On the other hand, upper bounds remain uninformative. We discuss this issue further in online appendix A.4.

²⁰The Anderson-Rubin test is implemented using the wild efficient restricted bootstrap, following Davidson and MacKinnon (2010).

TABLE 5.—PARTIAL IDENTIFICATION RESULTS

Assumptions	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Split Triangular Distribution						Simple Triangular Distribution					
	Lower Bound	Upper Bound	Low 95% CI	High 95% CI	Low 90% CI	High 90% CI	Lower Bound	Upper Bound	Low 95% CI	High 95% CI	Low 90% CI	High 90% CI
Panel A. ATE of a change from the first to the second tercile of expected consumption uncertainty												
Exogenous treatment selection	0.0037		0.0013	0.0059	0.0016	0.0055		0.0057	0.0032	0.0078	0.0036	0.0076
MTR + MIV (without bias correction)	0.0000	0.4469	0.0000	0.4716	0.0000	0.4685	0.0000	0.4382	0.0000	0.4703	0.0000	0.4691
MTR + MIV (with bias correction)	0.0000	0.4314	0.0000	0.4872	0.0000	0.4841	0.0000	0.4255	0.0000	0.4830	0.0000	0.4818
Panel B. ATE of a change from the first to the third tercile of expected consumption uncertainty												
Exogenous treatment selection	0.0283		0.0158	0.0420	0.0174	0.0401		0.0297	0.0160	0.0448	0.0177	0.0428
MTR + MIV (without bias correction)	0.0098	0.6451	0.0013	0.7134	0.0013	0.7031	0.0136	0.6485	0.0063	0.7240	0.0068	0.7094
MTR + MIV (with bias correction)	0.0062	0.6500	0.0000	0.7184	0.0000	0.7080	0.0113	0.6535	0.0040	0.7290	0.0046	0.7143
Panel C. ATE of a change from the second to the third tercile of expected consumption uncertainty												
Exogenous treatment selection	0.0246		0.0117	0.0387	0.0136	0.0367		0.0240	0.0099	0.0394	0.0120	0.0376
MTR + MIV (without bias correction)	0.0000	0.4831	0.0000	0.5477	0.0000	0.5398	0.0000	0.4929	0.0000	0.5648	0.0000	0.5509
MTR + MIV (with bias correction)	0.0000	0.4854	0.0000	0.5500	0.0000	0.5421	0.0000	0.4944	0.0000	0.5663	0.0000	0.5524
Number of observations	2,791											

MTR: monotone treatment response; MIV: monotone instrumental variable. The MIV used are deciles of the variance of expected income growth (the second decile is merged into the first one due to insufficient number of observations). Confidence intervals are computed using the bias-corrected percentile method, implemented by bootstrapping 1,000 times the estimation sample.

To conduct inferences on the ATE, we compute, as in de Haan (2012), 95% and 90% bias-corrected percentile confidence intervals using 1,000 bootstrap replications. Given that the bounds obtained under MTR and MIV involve optimization operations applied to nonparametric estimation, the estimates of the bounds can be biased. Therefore, in one of our estimations, we apply the bias correction procedure that Kreider and Pepper (2007) suggested. On the other hand, the bias estimate can be highly variable (see the discussion in de Haan, 2012, note (6)). We thus report results both with and without the bias correction.

The advantages of PI are considerable. First, it uses fairly weak assumptions (in our case, MTR and MIV). Second, it is nonparametric, and thus results do not depend on the functional form or the inclusion or exclusion of any control variables, the distributions of which are taken as given. Hence, in our context, PI accommodates unlimited heterogeneity in preferences, interest rate, and any other variable, observable or not, affecting expected consumption growth. Third, PI provides estimates of the ATE (and not the LATE) and allows for its full heterogeneity across sample units.

PI can also sometimes lead to identification regions that are wide. As Manski (1994) notes, however, the point identification obtained using strong assumptions may give a false certainty about results, as the reduction in uncertainty is obtained through strong untestable assumptions that might not hold.

Table 5 shows the PI estimates of the ATEs in three different panels that capture a change in expected consumption risk from first to second tercile (panel A), from first to third (panel B), and from second to third (panel C). For every estimation method, we report the lower and upper bounds on the ATE (or, in the case of ETS, the point estimate), as well as the 95% and 90% CIs. Columns 1 to 6 show the results for the split triangular distribution and columns 7 to 12 those for the simple triangular distribution.

Starting with the split triangular distribution, we first examine the ETS estimates, which in practice are equal to those obtained by running a weighted OLS regression on a constant and two dummy variables denoting the second and third terciles of the square of expected consumption growth. The ETS results imply that consumption risk has a strong, positive effect on expected consumption growth, with point estimates ranging from 0.4 to 2.8 percentage points. In addition, the CIs around the ETS estimates are quite narrow.

PI relaxes the assumption of exogeneity of consumption risk that underlies the ETS estimates. When using the MIV together with the MTR assumption, the identification regions have lower bounds equal to 0 when expected consumption uncertainty changes from its first to its second tercile, and the same holds for a change from its second to its third tercile. When, however, expected consumption uncertainty changes from its first to its third tercile, the lower bound on the ATE is equal to 0.98 percentage points when not using the bias correction (significant at 5%, as evidenced by the lower bounds of the 95% and the 90% CIs) and to 0.6 percentage points (not statistically significant) when not using the bias correction. The upper bounds on the ATEs are uninformative in all cases.

To gauge what a lower bound of 0.98 percentage points in expected consumption growth implies for the strength of precautionary saving (and thus for the coefficient of relative prudence), one needs to adjust this estimate by the change in expected consumption uncertainty (i.e., the square of expected consumption growth) from its first to its third tercile. This calculation produces an implied coefficient of relative prudence of 0.92, a value that is close to the estimates of about 1 that we get from robust OLS and IV estimation.²¹

²¹The median consumption uncertainty is equal to 0.0004697 and 0.0110797 in its first and third terciles, respectively. Hence, the implied estimate of half of the coefficient of relative prudence is $(0.0098/(0.0110797 - 0.0004697)) = 0.92$.

The corresponding estimate of the lower bound of half of the coefficient of relative prudence using the bias correction is 0.58, smaller than the robust OLS and IV estimates and not statistically significant. On the other hand, when using the simple triangular distribution, the lower bounds of the implied half of the coefficient of relative prudence without and with bias correction are equal to 1.11 and .92, respectively (the results are in both cases significant at 5%). These estimates are very similar to the robust OLS and IV ones.

While the upper bound of the MTR + MIV identification region remains high, one could, however, consider the ETS value as a plausible upper bound value. This is so because this value is just the difference in mean expected consumption growth between the first and third terciles of consumption risk, and thus is likely to overestimate the true effect of the latter on precautionary saving. Given that the ETS upper bound using the split triangular distribution is 0.0283 and using the same values for median consumption, the implied upper bound of half of the coefficient of relative prudence is 2.67, which in turn implies a coefficient of relative prudence equal to 5.3 and a coefficient of relative risk aversion of 4.3 (assuming isoelastic utility).

Overall, we conclude that even under the weak assumptions used in PI estimation and the resulting larger estimation uncertainty, households in our sample have preferences that imply a precautionary saving motive.

IX. Conclusions

We investigate the existence of a precautionary saving motive affecting household saving behavior. We estimate an Euler equation for consumption using subjective expectations of consumption, which conform better to the original Euler equation formulation than the ex post consumption realizations used in the related literature so far. Furthermore, expectation data circumvent problems related to inconsistency and endogeneity that affect ex post realizations. To obtain expectation data, we design a questionnaire that asks households about their expectations of future consumption and administer it to a representative sample of Dutch households.

Using these expectation data, we estimate the Euler equation and the strength of the precautionary saving motive through the magnitude of the prudence coefficient. We use a variety of estimation methods—OLS, robust regression, IV, and PI—and obtain consistent results pointing clearly to the existence of a precautionary motive in the saving behavior of the households in our sample. The estimated relative prudence coefficient is around 2, in line with the existing literature.

Since expectation data are correlated with observable household characteristics in a manner that conforms to theory and intuition, these data are likely to provide a reasonably good measure of the underlying uncertainty experienced by households. This points to the advantages of asking about households' expectations to investigate this uncertainty. More generally, the responses to such questions

are valuable for estimating economic relationships in which households' expectations play a key role. Thus, we recommend the inclusion of questions about expectations in household surveys.

REFERENCES

- Altug, Sumru, and Robert A. Miller, "Household Choices in Equilibrium," *Econometrica* 58 (1990), 543–570.
- Anderson, T. W., and Herman Rubin, "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *Annals of Mathematical Statistics* 20 (1949), 46–63.
- Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin, "Identification of Causal Effects Using Instrumental Variables," *Journal of the American Statistical Association* 91 (1996), 444–472.
- Attanasio, Orazio P., and Guglielmo Weber, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy* 103 (1995), 1121–1157.
- Bertola, Giuseppe, Luigi Guiso, and Luigi Pistaferri, "Uncertainty and Consumer Durables Adjustment," *Review of Economic Studies* 72 (2005), 973–1007.
- Cagetti, Marco, "Wealth Accumulation over the Life Cycle and Precautionary Savings," *Journal of Business and Economic Statistics* 21 (2003), 339–353.
- Campbell, John Y., "Understanding Risk and Return," *Journal of Political Economy* 104 (1996), 298–345.
- Carroll, Christopher D., "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)," *Advances in Macroeconomics* 1 (2001), 1003–1003.
- Chamberlain, Gary, "Panel Data," in Zvi Griliches and Michael D. Intriligator, eds., *Handbook of Econometrics*, vol. 2 (Amsterdam: Elsevier Science, 1984).
- Chen, Xiaohong, Jack Favilukis, and Sydney C. Ludvigson, "An Estimation of Economic Models with Recursive Preferences," *Quantitative Economics* 4 (2013), 39–83.
- Coibion, Olivier, and Yuriy Gorodnichenko, "What Can Survey Forecasts Tell Us about Informational Rigidities?" *Journal of Political Economy* 120 (2012), 116–159.
- Crump, Richard K., Stefano Eusepi, Andrea Tambalotti, and Giorgio Topa, "Subjective Intertemporal Substitution," Federal Reserve Bank of New York staff report 734 (2015).
- Das, Marcel, and Das Donkers, "How Certain Are Dutch Households about Future Income? An Empirical Analysis," *Review of Income and Wealth* 45 (2005), 325–338.
- Davidson, Russel, and James G. MacKinnon, "Wild Bootstrap Tests for IV Regression," *Journal of Business and Economic Statistics* 28 (2010), 128–144.
- Deaton, Angus S., "Saving and Liquidity Constraints," *Econometrica* 59 (1991), 1221–1248.
- de Haan, Monique, "The Effect of Parents' Schooling on Child's Schooling: A Nonparametric Bounds Analysis," *Journal of Labor Economics* 29 (2012), 859–892.
- Dillon, Eleanor W., "Risk and Return Tradeoffs in Lifetime Earnings," *Journal of Labor Economics* 36 (2018), 981–1021.
- Dominitz, Jeff, and Charles F. Manski, "Using Expectations Data to Study Subjective Income Expectations," *Journal of the American Statistical Association* 92 (1997), 855–867.
- Drèze, Jacques H., and Franco Modigliani, "Consumption Decisions under Uncertainty," *Journal of Economic Theory* 5 (1972), 308–335.
- Dynan, Karen, "How Prudent Are Consumers?" *Journal of Political Economy* 101 (1993), 1104–1113.
- Gourinchas, Pierre-Olivier, and Jonathan A. Parker, "Consumption over the Life Cycle," *Econometrica* 70 (2002), 47–89.
- Guiso, Luigi, Tullio Jappelli, and Mario Padula, "Pension Wealth Uncertainty," *Journal of Risk and Insurance* 80 (2013), 1057–1085.
- Guiso, Luigi, Tullio Jappelli, and Luigi Pistaferri, "An Empirical Analysis of Earnings and Employment Risk," *Journal of Business and Economic Statistics* 20 (2002), 241–253.
- Hansen, Lars P., and Kenneth J. Singleton, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy* 91 (1983), 249–265.

- Hayashi, Fumio, "Tests for Liquidity Constraints: A Critical Survey and Some New Observations" (pp. 91–120), in T. F. Bewley, ed., *Advances in Econometrics: Fifth World Congress*, vol. 2 (Cambridge: Cambridge University Press, 1987).
- Huber, Peter J., "Robust Regression: Asymptotics, Conjectures and Monte Carlo," *Annals of Statistics* 1 (1973), 799–821.
- Jappelli, Tullio, and Luigi Pistaferri, "Using Subjective Income Expectations to Test for Excess Sensitivity of Consumption to Predicted Income Growth," *European Economic Review* 44 (2000), 337–358.
- , "The Consumption Response to Income Changes," *Annual Review of Economics* 2 (2010), 479–506.
- Kaplan, Greg, Giovanni L. Violante, and Justin Weidner, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity* 48 (2014), 77–153.
- Kimball, Miles S., "Precautionary Saving in the Small and in the Large," *Econometrica* 58 (1990), 53–73.
- Kreider, Brent, and John V. Pepper, "Disability and Employment: Reevaluating the Evidence in Light of Reporting Errors," *Journal of the American Statistical Association* 102 (2007), 432–441.
- Manski, Charles F., "Nonparametric Bounds on Treatment Effects," *American Economic Review Papers and Proceedings* 80 (1990), 319–323.
- , "The Selection Problem," in C. Sims, ed., *Advances in Econometrics: Sixth World Congress* (Cambridge: Cambridge University Press, 1994).
- , "Monotone Treatment Response," *Econometrica* 65 (1997), 1311–1334.
- , "Measuring Expectations," *Econometrica* 72 (2004), 1329–1376.
- , "Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise," in Michael Woodford and Jonathan A. Parker, eds., *NBER Macroeconomics Annual 2017* (Cambridge, MA: MIT Press, 2018).
- Manski, Charles F., and John V. Pepper, "Monotone Instrumental Variables, with an Application to the Returns to Schooling," *Econometrica* 68 (2000), 997–1012.
- Palumbo, Michael G., "Uncertain Medical Expenses and Precautionary Saving near the End of the Life Cycle," *Review of Economic Studies* 66 (1999), 395–421.
- Skinner, Jonathan, "Risky Income, Life Cycle Consumption, and Precautionary Savings," *Journal of Monetary Economics* 22 (1988), 237–255.
- Stock, James H., Jonathan H. Wright, Motohiro Yogo, "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments," *Journal of Business and Economic Statistics* 20 (2002), 518–529.
- Voena, Alessandra, "Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?" *American Economic Review* 105 (2015), 2295–2332.
- Zarnowitz, Zarnowitz, Victor, and Louis Lambros, "Consensus and Uncertainty in Economic Prediction," *Journal of Political Economy* 95 (1987), 591–621.