Anatomy of Consumption Risk

Luigi Guiso

EIEF and CEPR

Tullio Jappelli

University of Naples Federico II, CSEF and CEPR

April 2, 2025

Abstract

We use quarterly panel data from the Italian Survey of Consumer Expectations to analyze the relative contribution of idiosyncratic and aggregate risks on household consumption uncertainty. Consumption uncertainty is primarily driven by individual shocks rather than macroeconomic fluctuations. Estimating pass-through coefficients, we find that idiosyncratic risks (particularly those related to health and income) account for 75% of consumption risk, while aggregate risks (particularly GDP and house price fluctuations) contribute less than 20%. Young workers with limited cash reserves face greater exposure due to limited insurance options. Using subjective expectations data and an instrumental variables approach to the Euler equation, we estimate a relative prudence coefficient of 2–3. On average, precautionary savings amount to 2.7% of current consumption, three quarters of this driven by idiosyncratic risks. Only one third of these savings buffer labor income shocks, while the remainders is a response to expenditure shocks. These findings suggest that focusing solely on income risk significantly understates precautionary savings. Consistent with theoretical predictions, precautionary savings are highest among younger individuals and decline with age.

Keywords: Consumption Risk; Income Risk; Health Risk; Aggregate Risk; Precautionary Savings.

JEL: D12, D14, D15, C8, C99

The research is funded by the European Union - Next Generation EU, in the framework of the GRINS - Growing Resilient, Inclusive and Sustainable Project (GRINS PE00000018 – CUP E63C22002140007). The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the European Union and the European Union cannot be held responsible for them. We thank seminar participants at the Collegio Carlo Alberto (Torino), Bank of Italy, Einaudi Institute for Economics and Finance, University of Cyprus. We are grateful to Rob Alessie and Orazio Attanasio for comments and discussions, and Alessandro Sciacchetano for excellent research assistance.

1. Introduction

In intertemporal consumption models, expected consumption volatility plays a central role in shaping consumers' decisions and welfare. This volatility is a sufficient statistic for all consumption-relevant sources of uncertainty. However, with some exceptions, standard intertemporal consumption models typically assume that income is the sole source of risk and that markets do not provide insurance against it.[1] While income risk is undoubtedly a major driver of fluctuations in household budget constraints, individuals face many other risks that can significantly impact consumption volatility and financial well-being. These include shocks to healthcare costs, energy prices, asset prices, and business cycle fluctuations. The *relative* and *overall* importance of these risks for consumption volatility and thus for the accumulation of precautionary savings, remains an open empirical question. This paper seeks to address this gap. Understanding which risks disrupt consumption and how households respond to them is essential for reducing vulnerability to shocks and developing mitigation strategies. It is also crucial for improving macroeconomic models that emphasize consumers' uncertainty.

To analyze the anatomy of consumption risk, we use a consumer expectations survey to elicit the probability distribution of future consumption growth, constructing an individual-level measure of expected consumption risk. Because individuals base their responses on available information, this measure cannot be inferred from volatility of realized consumption. Additionally, we gather subjective probability distributions for a broad set of risks that represent key sources of uncertainty—ranging from income and health expenditures to energy price shocks, GDP fluctuations, and other aggregate variables.

This approach allows us to trace expected consumption risk back to both individual-specific risks (such as income, health expenditures, and energy bills) and aggregate risks. By doing so, we estimate how each source of uncertainty affects consumption volatility, quantify their relative importance, and assess their heterogeneous impact across consumers. To our knowledge, this is the first empirical study to systematically examine the anatomy of consumption risk and evaluate the contribution of each risk factor to overall consumption uncertainty.

To elicit subjective expectations, we designed the Italian Survey of Consumer Expectations (ISCE), a panel of 5,000 individuals that has collected quarterly data since October 2023. The survey gathers participants' expected probability distributions for

[1] For recent surveys, see Attanasio and Weber (2010), Jappelli and Pistaferri (2017), and Violante (2024).

consumption growth, income growth, energy prices, health expenditures, and key aggregate variables (GDP growth, inflation, unemployment, house prices and interest rates). Respondents allocate 100 probability points across a set of interval growth rates for each variable over the 12 months following the interview. In many cases, growth rate intervals range symmetrically from negative (below -8%) to positive (above 8%).

For each respondent, we compute moments of these distributions, using variance as our measure of risk. Consumption risk exhibits substantial heterogeneity across individuals and varies over time within the survey period. Subjective consumption risk tends to be lower than labor income risk, consistent with consumption smoothing. Both idiosyncratic and aggregate risk measures correlate with observable characteristics in intuitive ways. For instance, self-employed individuals have a higher income variance and a stronger correlation between expected income and GDP growth, reflecting their greater exposure to business cycle fluctuations. Similarly, young consumers anticipate lower health expenditure growth than older individuals. These, and other validation tests confirm that the elicited probability distribution are meaningful and informative.

To estimate pass-through coefficients, we regress consumption risk on the variances of its underlying risk sources, initially assuming these risks are independent. Extending Banks et al. (2001) to a multi-risk framework, we show that pass-through coefficients reflect both exposure to specific risks and ability to insure against them, either formally or informally. In our preferred specification, we control for individual fixed effects, using only time variation in measured risks for identification.

Our results indicate that these risk sources have strong explanatory power, collectively accounting 57% of the variation in consumption risk. Among individual level risks, personal income, health and energy expenditure risks have the largest impact. Health expenditure risk exhibits the highest pass-though (0.37), surpassing that of income risk (0.21) and energy expenditure risk (0.19). We argue this is because health shocks influence consumption volatility both directly through health expenditures, and indirectly, as health status affects the marginal utility of consumption, as demonstrated by Blundell et al. (2024). Aggregate risks also play a role, particularly the expected variability in GDP and house price growth.

We also find systematic differences across individuals in how source risks translate into consumption uncertainty, reflecting both heterogeneity in exposure and differences in the ability to absorb shocks. Younger individuals (under 35) experience a 50% higher pass-through of income risk than older individuals (60+), likely due to greater labor market volatility and lower financial buffers. The self-employed face a 70% higher pass-through of energy expenditure risk and twice as high of GDP risk compared to employees, as they are more exposed to these factors. Older individuals, with higher homeownership rates (82% vs. 68% for younger groups), are twice as sensitive to housing price fluctuations. Finally, individuals with high liquid cash reserves exhibit lower risk pass-throughs of income, health, and energy risks on consumption volatility, underscoring their role as financial buffers.

In the total sample, most of the expected consumption risk is explained by idiosyncratic sources, which together account for 73% of expected consumption risk, reflecting both the size of the pass-through and the average variance of the source risk. Aggregate risks account for about 17% of consumption risk, suggesting that the uncertainty affecting consumer welfare is largely driven by individual-level shocks rather than shocks associated with business cycle fluctuations. Among macroeconomic risks, GDP and house price risks are the most important contributors to consumption risk, while interest rate, inflation, and unemployment risks together account for only 4% of predicted consumption risk.

Importantly, our conclusions hold even when we relax the assumption of independence between the underlying sources of risk. If the risks are correlated, interactions between the standard deviations of risk pairs might influence expected consumption risk. However, our tests show that only the interaction between the standard deviation of expected health expenditures and income and that between the standard deviations of expected inflation and house prices significantly affects expected consumption uncertainty. Despite this, when evaluated at the sample means, the contributions of the individual sources of risk remain unchanged.

In a second step of our study by using the information on subjective distributions to estimate the structural prudence parameter of the Euler equation for consumption and to bring additional evidence on the relevance of precautionary motives for savings. We find that expected consumption growth is related positively to expected consumption risk in line with the predictions of precautionary saving models. Our estimates imply a coefficient of relative prudence in the plausible 2-3 range. We then combine the Euler equation estimates with the risk measures, to quantify the size of precautionary savings in our sample. We estimate that the average flow of precautionary savings due to the underling risks amounts to 2.7% of consumption. Of this, about 70% is due to idiosyncratic uncertainty, with roughly similar contributions from labor market risk, health and energy expenditures. Thus, focusing only on labor market risk as done by a large part of the literature is bound to grossly understate the size of precautionary savings. Aggregate risks explain the remaining 30% of precautionary saving, with roughly equal contributions of inflation, GDP and house price risk. We also find a marked decline in the importance of precautionary savings over consumers age, as implied by finite horizon models of precautionary wealth accumulation (Caballero, 1990).

This paper contributes to the classic literature on precautionary saving in the presence of non-insurable risks (e.g., Caballero, 1990, 1991; Carroll and Kimball, 1996; Banks et al., 2001) and to research estimating prudence using the Euler equation approach (e.g., Dynan, 1993; Carroll, 2001; Ludvigson and Paxon, 2001; Attanasio and Low, 2004; Bertola et al., 2003; Fagereng et al., 2017; Attanasio et al., 2019). While much of this literature focuses on labor income risk, various studies have expanded the scope of precautionary savings to other sources of uncertainty. Some papers emphasize health shocks alongside income risk. Palumbo (1999) estimates a structural model in which uncertain future medical expenses play a key role in saving decisions. De Nardi and Fella (2017) argue that earnings risk, life expectancy uncertainty, and medical expenditure risk are all crucial in shaping consumption decisions and wealth inequality. More recently, Blundell et al. (2024) analyze how consumption and medical expenses jointly respond to income and health shocks. Other studies highlight the role of aggregate risks. Ryngaert (2022) finds that perceived inflation risk correlates with higher real consumption growth and a greater propensity to purchase durables. Coibion et al. (2024) use European household survey data to examine how exogenous variation in perceived macroeconomic uncertainty influences spending decisions. However, no prior work has jointly examined the full set of idiosyncratic and aggregate risks to assess their impact on consumption volatility and precautionary savings—precisely the focus of this paper.

Our work is also closely related to Fulford and Low (2024), who use survey data from the U.S. Consumer Financial Protection Bureau to document key patterns in expense and income shocks. They find that expense shocks are more frequent and larger than income drops and, except for health expenses, are generally uncorrelated with income declines. While their study focuses on realized shocks, we examine ex ante expected uncertainty and its role in perceived consumption volatility—ultimately what drives precautionary saving behavior. Additionally, we extend the analysis beyond personal income and expenditure shocks to include aggregate uncertainty. Despite these differences, our papers are highly complementary.

Finally, we contribute to the growing literature using elicited subjective expectations to measure risks, including income and unemployment risk (Dominitz and Manski, 1997; Guiso et al., 2002; Jappelli and Pistaferri, 2000), pension risk (Guiso et al., 2013), inflation uncertainty (Crump et al., 2015), consumption uncertainty (Christelis et al., 2020), and business cycle volatility (Georgarakos et al., 2025).

The remainder of the paper is structured as follows. Section 2 outlines our conceptual framework for measuring consumption risk. Section 3 describes the Italian Survey of Consumer Expectations and the design of the subjective probability elicitations. Section 4 presents the data and validation exercises. Section 5 analyzes consumption risk by estimating risk pass-throughs and quantifying the contribution of different risk sources. Section 6 uses our risk indicators to estimate the Euler equation for consumption and measure precautionary savings. Section 7 concludes.

2. Measuring consumption risk

The literature employs various approaches to measure consumption risk, including ex-post consumption volatility, income volatility, asset pricing models, and subjective expectations, each with its own strengths and limitations. Some studies use realized consumption volatility to proxy for consumption risk in Euler equation estimates (e.g. Dynan, 1993). A limitation of this approach is that expected consumption volatility does not coincide necessarily align with realized volatility. The latter, beyond capturing genuine innovations, also reflects individual choice, making it endogenous.[2] Another approach to measure consumption risk is to use labor income volatility as a proxy, often defined as the standard deviation or variance of realized labor income growth (Bertola et al., 2005). However, while income risk is an important factor, it represents only one of many risks people face, leading to an understatement of overall consumption risk. Additionally, similar to consumption volatility, realized income volatility reflects both genuine risk as well as choices (e.g., decisions to work more or fewer extra hours), making it partially endogenous. Some studies instead employ consumption-based asset pricing models, such as the CAPM, to estimate consumption risk directly (Ludvigson, 2013). These models link consumption growth to asset returns, using financial market data to infer the riskiness of future consumption streams.

Our approach, by contrast, relies on self-reported household spending behavior and

[2] To address endogeneity issues in the Euler equation estimate, Dynan uses education and occupation as instruments, but these instruments have low power.

subjective probability distributions of various risks. The use of subjective expectations to measure economic uncertainty was pioneered by Manski (2004) and Guiso, Jappelli and Terlizzese (1992), and has recently been applied by Caplin et al. (2023), and Arellano et al. (2024), among others. More broadly, Bachmann et al. (2022) and Stancheva (2022) provide numerous applications of survey-based subjective expectations across fields such as education, labor, health, and macro-finance.

2.2. Conceptual framework

Our starting point is that consumption uncertainty is induced by many economic factors. To varying degrees income risk, price variability, health shocks, asset prices volatility all contribute to consumption risk. Some of these risks are common to all consumers, while some are individual specific. By limiting people's ability to cope with income and prices volatility and health shocks, financial and insurance markets imperfections can amplify the impact of these sources of risk on expected consumption uncertainty.

In this paper we document how various sources of risk affect consumption uncertainty using insights from the literature on intertemporal consumption decisions. The first insight is that expected consumption volatility is the relevant measure of the uncertainty affecting consumers' saving response to risk and thus the amount of precautionary savings. The second insight is that expected consumption volatility is a sufficient statistics for all consumption-relevant sources of risk and can be traced back to these sources. The third insight is that the need to engage in precautionary saving in response to a specific risk depends on the level of exposure to this risk. For instance, exposure to energy price shocks depends on the weight of energy costs in the individual budget constraint. The fourth insight is that some risks can be insured, totally or partially, either through formal insurance markets or through informal mechanisms. This attenuates their impact on overall consumption risk depending on available insurance opportunities.

To capture these insights, we build on Banks et al. (2001) and assume that preferences are constant relative risk averse (CRRA) and time separable. We also assume that optimal consumption c_t (net of energy and health expenses) is approximately proportional to individual wealth w_t, $c_t \approx \mu w_t$. When labor income is the only source of risk, Banks et al. (2001) show that, to a second-order approximation in the Euler equation, the expected variance of consumption growth is proportional to the expected variance of income growth, scaled by the ratio of labor income to total wealth. This scaling factor implies that consumption risk (i.e., the expected variance of consumption growth) is more sensitive to income risk for individuals with lower wealth.[3]

We extend this framework to incorporate multiple sources of risk, making simplifying assumptions to derive an explicit solution for the relation between the expected variance of consumption growth and the underlying risks, one that can be implemented empirically. To illustrate the methodology, we focus here on four key sources of risk, for which we elicited subjective probability distributions. Let $\tilde{y}_{t+1}$ denote random labor income in period t+1. Random health expenditures are given by $p_h\tilde{h}_{t+1}$, where p_h is the certain and constant price of health goods and services, and $\tilde{h}_{t+1}$ represents uncertain health status. Energy costs are denoted by $e\tilde{p}_{t+1}$, the product of a known quantity of energy needs e and the uncertain energy price $\tilde{p}_{t+1}$. Finally, $\tilde{r}_{t+1}$ is the uncertain return to wealth, encompassing both financial and real assets. For a consumer i (omitted in notation for simplicity) the budget constraint in period t+1 is:

$$\widetilde{w}_{t+1} = (1 + \tilde{r}_{t+1} + v\,\widetilde{p}_{t+1})w_t + \tilde{y}_{t+1} - c_t - p_h\tilde{h}_{t+1} \qquad (1)$$

where we assume $e = vw_t$, that is, the quantity of energy purchased is a constant share of individual wealth - a simple way to capture the idea that energy consumption is highly price inelastic. In Appendix A we show that a second order Taylor approximation of the expected marginal utility of consumption delivers the following expression:

$$\sigma^2_{\xi,t+1} = \pi^2_y\sigma^2_{y,t+1} + \pi^2_h\sigma^2_{h,t+1} + \pi^2_p\sigma^2_{p,t+1} + \pi^2_r\sigma^2_{r,t+1} \qquad (2)$$

where $\sigma^2_{\xi,t+1}$ is the variance of the innovation of consumption growth, and the four variances on the right-hand-side of equation (2) are, respectively, innovations of income growth, health expenditure, energy prices, and interest rate variances. All these variances are conditional on the information available to consumers in period t and can evolve over time as new information is acquired. In deriving the equation, we assume that all covariances between these risks are zero, though we will later relax this assumption.

The scaling factors π^2_z are proportional to the weight of each risk $(z = y, h, p, r)$ on predictable consumption $(\mu\overline{w}_{t+1})$, and thus measures the consumer's exposure to each of the four risks. We

[3] Banks et al. (2001) also show that transitory income processes have a small impact on consumption risk, and that this impact increases with the persistence of this process.

treat π_z^2 as parameters, potentially varying across consumers. For example, the conditional variance of income growth is more important for individuals whose lifetime wealth is largely composed of human capital (a high π_y^2 term) than for those less dependent on human capital (a relatively low π_y^2). Similarly, individuals close to retirement are less exposed to labor income risk than those just entering the labor market. In a world where these risks cannot be insured or hedged, π_z^2 measures the pass-through of risk z to overall consumption risk.

In practice, consumers can mitigate some of these risks through formal markets, informal networks, public intervention, or accumulated precautionary savings.[4] To allow for partial insurance we let the pass-through coefficient of risk z be equal to $\beta_z = \alpha_z \pi_z^2$, where $0 \leq \alpha_z \leq 1$ is a risk attenuation factor reflecting insurance opportunities vis a vis risk source z. Absence of insurance implies $\alpha_z = 1$, and full insurance $\alpha_z = 0$. We can than re-write equation (2) as:

$$\sigma_{\xi,t+1}^2 = \beta_y \sigma_{y,t+1}^2 + \beta_h \sigma_{h,t+1}^2 + \beta_p \sigma_{p,t+1}^2 + \beta_r \sigma_{r,t+1}^2 + \varepsilon_{t+1} \quad (3)$$

The β_z coefficients in equation (3) measure the pass-through of risk z on consumption risk and reflect both exposure and insurability. In the (unrealistic) case of complete markets, the β_z coefficients are all equal to zero, and the rate of growth of individual consumption has no idiosyncratic volatility. Otherwise, the coefficients reflect the sensitivity of consumption volatility to the underlying risks, due to both exposure and insurance.

Since our aim is to trace consumption risk back to its sources using a regression framework, in equation (3) we also include the term ε. This captures additional determinants of expected consumption volatility not included in our list of source risks, higher order terms omitted from the Taylor expansion, individual specific unobserved factors that contribute to consumption risk such as ability to process information, and measurement error.[5] We model the error term as $\varepsilon_{i,t+1} = \kappa_i + \vartheta_{i,t+1}$, the sum of an individual fixed effect and an error term. Thus,

[4] For example, in Italy some health shocks are fully covered by the National Health System (NHS), while out-of-pocket health expenditures are primarily related to services not covered by the NHS, such as dental care and preventive healthcare. Another example are welfare programs, including unemployment insurance, fiscal transfers, and social safety nets, that help reduce expected consumption volatility for individuals facing income and unemployment risk. Additionally, some income shocks are partially offset by financial support from parents, relatives, or friends (e.g., Fagereng et al., 2024).

[5] Notice that this error is not an expectational error (i.e. the difference between planned and realized consumption) as in the Euler consumption equations which often are estimated with panel data. All terms in the equation are variances computed conditional on the information available to the consumer at the time they predict consumption, income, and the other variables of equation (3).

using cross-sectional data to estimate equation (3) could result in inconsistent estimates of the pass-through parameters due to the presence of unobserved heterogeneity and any other time-invariant characteristic. To address this issue, we estimate model (3) using a fixed effect estimator in panel data. Thus, our identification strategy of the pass-throughs relies on variation over time in expected consumption risk σ_{ξ}^2 and its underlying risk sources σ_z^2.

Equation (3) provides the framework for our decomposition exercise and has interesting implications. It shows that consumption risk reflects all sources of risk that affect the consumer's budget constraint. Additionally, what matters for expected consumption risk is the individual's subjective perception of the risk sources, not the ex-post volatilities calculated using observational data. Finally, since the β_z coefficients reflect risk insurability and exposure (i.e. the α_z and π_z parameters), and since both might vary across individuals, the estimated pass-throughs may also differ across individuals. To account for these differences, and indirectly to validate the decomposition of consumption risk, we also estimate equation (3) for different groups of consumers with supposedly different exposure to specific risk sources or with different capability to buffer consumption-relevant risks.

In some specifications we allow also for covariances among risk sources, and write model (3) as:

$$\sigma_{\xi,t+1}^2 = \beta_y \sigma_{y,t+1}^2 + \beta_h \sigma_{h,t+1}^2 + \beta_p \sigma_{p,t+1}^2 + \beta_r \sigma_{r,t+1}^2 + \sum_z \sum_{x<z} \beta_{zx} \pi_{zx} \sigma_{z,t+1} \sigma_{x,t+1} + \varepsilon_{t+1} \quad (4)$$

where $\beta_{zx} = 2\rho_{zx}\pi_{zx}$, and ρ_{zx} is the correlation coefficient between risk sources *z* and *x* and $\pi_{zx} = \pi_z \pi_x$ the product of the two exposures. Non-zero correlations imply that the interaction terms between the standard deviations of pairs of risk sources might also affect consumption uncertainty. The interaction terms are irrelevant if two risks are uncorrelated or if the consumer is exposed to only one of them, or if one of the two risks is fully insured.

In the empirical implementation of equations (3) and (4) we account also for indicators of subjective aggregate risk, demographic variables, time and individual fixed effects. We also test for the significance and influence of the covariance terms, and whether their presence affects our decomposition of consumption risk.

3. The Italian Survey of Consumer Expectations

The Italian Survey of Consumer Expectations (ISCE) is a quarterly rotating panel. We use the first five waves collected in October 2023, January 2024, April 2024, July 2024 and October 2024.[6] The survey collects data on demographic variables, household resources (income and wealth components), consumption, individual expectations of future distribution of individual-level variables (consumption, income, energy expenditures and health expenditures), and macroeconomic variables (GDP growth, unemployment, inflation, nominal interest rate, and house prices growth).

The survey builds upon two international online, high-frequency surveys: the New York Federal Reserve Survey of Consumer Expectations, which collects information on consumers' views and expectations regarding inflation, employment, income, and household finances, and the European Central Bank Consumer Expectation Survey which collects monthly data on households' expectations in the euro area economies.

ISCE targets Italian resident population aged between 18 and 75. The survey includes approximately 5,000 interviews in each of the five waves, and a total of 25,026 interviews conducted during the first 7-15 days of the reference month.[7] The average response rate across waves (ratio of completed interviews to invitations) is around 40%. The average retention rate (percentage of individuals interviewed in two consecutive waves) is around 80%.[8] To maintain population representativeness the sample is refreshed at quarterly intervals. The ISCE Statistical Bulletin provides detailed information on the survey, see Guiso and Jappelli (2025).

The ISCE sampling scheme is modeled after that used in the Bank of Italy's Survey of Household Income and Wealth (SHIW). The sample is based on a stratification of the Italian resident population according to several criteria: area of residence (North-East, North-West, Central and South Italy), age group (18-34, 35-44, 45-54, 55-64, and over 65), gender, education (college degree, high school degree, and less than high school), and employment status (working, not working). We apply sample weights to ensure that the descriptive statistics are representative of the population.

Appendix B provides further details on the survey and compares the sample means of selected ISCE variables with those from the 2022 SHIW, the most recent available wave. The

---

[6] We have chosen this quarterly sequence to avoid interviewing consumers during the seasonal months of August and December.

[7] All interviews are conducted using the Computer Assisted Web Interviewing (CAWI) method, with an average duration of 18 minutes.

[8] For example, of the 5,011 individuals interviewed in the fifth wave, 2,978 had participated since wave 1, 422 since wave 2, 416 since wave 3, 635 since wave 4 and 560 were interviewed for the first time in wave 5.

.

gender, age, family composition and education levels in both samples are similar (see Table B1), which is noteworthy, given that the ISCE sample likely includes individuals to have internet access and to respond to online questionnaires. Consumption medians are generally comparable between the two surveys, but income appears to be underestimated in ISCE relative to SHIW (see Table B2). A possible explanation is that SHIW collects much more detailed information on disposable income than ISCE. Homeownership is reported at 72% in both surveys, while ISCE respondents report higher financial wealth and greater participation in financial markets (bonds, stocks, private pensions and life insurance).

ISCE elicits subjective probability distributions of consumption, income and the other risk sources over the 12 months following the interview. In the case of income, respondents are asked to assign probabilities to 11 possible income growth intervals, ranging from less than -8% to more than 8%.[9] To construct the moments of the subjective income growth distribution, we take the midpoint of each growth interval. For the lowest (less than -8%) and highest (more than 8%) open intervals, we assume values of -10% and 10%, respectively. The variance of the individual distributions serves as our measure of income risk, which we use to analyze the determinants of consumption risk. ISCE elicits respondents' subjective distributions of individual-level and aggregate variables in a similar way. The intervals are the same across variables and respondent are asked to make sure that probabilities sum up to 100.

4. Descriptive analysis and validation

Table 1 presents summary statistics for the pooled sample across five survey waves conducted from October 2023 to October 2024. The top panel reports the means and standard deviations of key demographic variables. The sample is evenly split between males and females, with an average age of 48 years. Additionally, 33% of respondents hold a college degree. Homeownership is prevalent, with 75% of participants owning their homes. Approximately 9% are self-employed, while 45% are employees. On average, cash-on-hand amounts to 29,844 euros, with a standard deviation of a similar magnitude.

The second panel displays the sample means and standard deviations for the first moments of subjective distributions of consumption growth and eight sources of uncertainty. Average growth rates vary across variables, reflecting considerable diversity in expectations, as indicated by standard deviations that generally exceed the means. On average, respondents

[9] Except for unemployment (intervals ranging from zero to 14%) and the nominal interest rate (from 0 to 8%).

anticipate a 0.7% increase in consumption, with a standard deviation of 3.9%. They are less optimistic about disposable income growth, projecting an average decline of approximately 1% across the five waves. Over the sample period, health costs are expected to rise by 0.8%, while energy costs are projected to increase by 1.3%, both with a standard deviation of around 3.5%.

The lower part of Table 1 presents the sample means of the variances (scaled by a factor of 100) and the second moments of the expected distributions of consumption growth and its underlying sources of risk. The variances of income, health, and energy expenditure growth rates are similar in magnitude and display substantial dispersion. The variances of consumption and income growth are also comparable, though the income variance is larger, suggesting some degree of smoothing. Appendix C (Figure C1) provides additional insights into the average values of our risk measures over time.

Crucially, expectations vary both over time for the same individual and across respondents. A simple variance decomposition reveals that 57% of the total variance of σ_{ξ}^2 is due to cross-sectional differences, while the remaining 43% results from time variation within individuals. Similarly, the variance decomposition of our risk indicators shows that within-individual variance accounts for approximately 40% of total variance.

4.1. Validation of uncertainty measures

We validate our measures of subjective uncertainty using two strategies. First, we examine whether the moments of our risk indicators correlate with observable characteristics — such as age, occupation, and cash-on-hand — in ways consistent with theoretical expectations. We find that all perceived risks, particularly consumption and income risk, are higher for younger individuals, who experience greater career uncertainty and more precarious employment (see Appendix C, Figures C2 and C3). Additionally, we observe a strong negative relationship between cash-on-hand and consumption risk, which aligns with precautionary savings models predicting that the variance of innovations to consumption growth should decrease monotonically as cash-on-hand increases (Carroll, 2001).

Further, when comparing self-employed individuals to employees, we find that the self-employed exhibit greater income variance and a 40% higher correlation between expected income and GDP growth, consistent with their heightened exposure to business cycle fluctuations. Similarly, comparing younger respondents (under 35) with older respondents (60+), we find that the younger group anticipates a significantly lower rate of health expense growth (0.4% vs. 1.6%) and half the expected variance of income growth.

These findings reinforce the validity of our elicited probability distributions, suggesting that respondents meaningfully assess and incorporate the uncertainties they face, an essential condition for accurately accounting for the relevant risk sources when expressing their consumption risk.

Our second strategy for validating our measures of uncertainty is to assess the consistency between the elicited risk sources and consumption risk as implied by equation (4), Section 2.[10] To do this, we define a dummy variable, $I(\sigma_{\xi}^{2})$, which equals one if the variance of expected consumption growth is positive and zero otherwise. Similarly, we define another dummy variable, $I(\sigma_{s}^{2})$, which equals one when at least one source of risk has a positive variance and zero if all sources of risk have zero variance. By interacting these two dummies, we obtain a 2×2 matrix with four possible outcomes.

If our risk sources are exhaustive, equation (3) implies that individuals who report no source risks should also report no consumption risk, that is, $I(\sigma_{\xi}^{2}) = 0$ when $I(\sigma_{s}^{2}) = 0$. This condition holds for 6,982 cases (27.9% of the sample, Table 2, Panel A). A second interesting case arises when consumption variance is positive, $I(\sigma_{\xi}^{2}) = 1$ despite no reported source risks, $I(\sigma_{s}^{2}) = 0$. This occurs in only 1% of the sample (247 observations), and may suggest either that there are other risk sources not captured by our questions or an inconsistent response (people expect consumption volatility despite facing no risk). But such cases are rare. We interpret this as evidence that our list of source risks is fairly comprehensive for the vast majority of the sample.

The third case involves respondents who perceive some risk, $I(\sigma_{s}^{2}) = 1$, and expect their consumption to be volatile, $I(\sigma_{\xi}^{2}) = 1$. This should be the typical scenario unless individuals can fully insure against all reported risks. Indeed, this group represents the majority of observations (50.86%).

The fourth case, comprising 20.1% of the sample (5,058 observations), requires more explanation. In this scenario, respondents expect no consumption risk, $I(\sigma_{\xi}^{2}) = 0$ despite reporting that at least one positive variance of the risk source, $I(\sigma_{s}^{2}) = 1$. This could indicate that some consumers effectively insure against these risks.. Alternatively, it may suggest that they fail to pass reported risks to consumption risk, due to inattentiveness during the interview,

[10] Recall that equation (4) shows that the variance of consumption innovation is a weighted sum of source risks.

measurement issues, or the insignificance of these risks (i.e., small variances are simply ignored).

Panel B of Table 2 compares the means variances of different risk sources among respondents who report positive risk factors, distinguishing between those who expect no consumption risk, $I(\sigma_{\xi}^2) = 0$, and those who do, $I(\sigma_{\xi}^2) = 1$. The results show that the variance of each individual risk factor is much lower among those reporting no consumption risk. The last row reveals that the average variance of these risks is 0.014 in the group expecting no consumption risk, and 5.6 time larger among those reporting positive consumption risk (0.079). Moreover, the former group reports an average of 2.6 risk sources with positive variance (out of 8), compared to 6.0 in the latter group.

Because consumption risk is additive in source risks (see equation 3), this evidence suggests that when risks are either small or few in number, individuals are more likely to insure against them or simply do not perceive them as a threat to consumption stability. To explore this further, we examine the relation between relate the probability of reporting zero consumption risk and two factors: the sum of the variances of the risk sources ($sum\_\boldsymbol{\sigma_s^2}$)[11] and the number of risk sources with positive variance, $num\_I(\boldsymbol{\sigma_s^2})$. Using a linear probability model, we estimate this regression among respondent with at least one positive-variance source risk ($I(\sigma_s^2) = 1$) using the panel and controlling for individual fixed effects. The results confirm a significant inverse relationship between the probability of reporting zero consumption risks and both $sum\_\boldsymbol{\sigma_s^2}$ and $num\_I(\boldsymbol{\sigma_s^2})$ (Table 4, panel C).[12]

4.2. Potential correlations among responses to different risks

In Appendix D, we examine how correlations between different risks impact our analysis. We find that the variances of various risks are positively correlated in both the cross-section and over time (Table D1). For example, the cross-sectional correlation between income risk and health expenditure risk is 0.71, while the correlation between changes in these risks over time is 0.44. More broadly, correlations in risk levels are around 0.7, and correlations in first differences are approximately 0.4. This suggests that individuals who anticipate high risk in one domain tend to expect higher risks in others. Similarly, those who report an increase in risk in

[11] The risk sources are weighted by the pass-through coefficients estimated in Table 3.

[12] Effects are large. A one standard deviation increase in $sum\_\boldsymbol{\sigma_s^2}$ lowers the probability of fully insuring consumption by 23 percentage points. Holding the number of risks constant, an increase in $num\_I(\boldsymbol{\sigma_s^2})$ by one standard deviation lowers the probability of fully insuring consumption by 3.5 percentage points.

one area are likely to report rising risks in other domains as well.[13]

There are at least two possible explanations for these correlations. One possibility is behavioral: individuals who have experienced economic instability may develop a heightened perception of uncertainty across multiple domains (Ben-David et al., 2018). Alternatively, individuals' general perception of uncertainty—and their ability to interpret probabilities—may influence how they assess risks in different areas. Since our goal is to understand how underlying risks translate into consumption uncertainty, we must account for the possibility that respondents answer similar questions in a repetitive or mechanical manner.

To investigate this, we analyze how respondents assign probabilities to different growth intervals across seven risk indicators, all measured on the same scale (from "less than -8%" to "more than 8%"). If respondents were simply applying the same probability distribution to each risk mechanically—e.g., assigning 20% to a given interval across all risks—it would artificially induce correlations across risk measures. If this kind of automatic response were applied to all intervals, it would result in identical variances (and other statistical moments) across all risks for the same respondent. In that case, consumption risk would appear to be correlated with all risk sources, even those that do not actually contribute to consumption uncertainty.

We address this issue in Appendix D. First, we present the distribution (Figure D1), sample means (Table D2), and regressions (Table D3) of the probability weights assigned to each growth interval. We then run a formal test of the hypothesis that the values assigned to each interval across risks are similar, controlling for individual, time and question fixed effects. This hypothesis is overwhelmingly rejected for all growth intervals. Moreover, individual fixed effects account for only 10% to 30% of the total variance in assigned probabilities across intervals. Overall, our test shows that there is genuine individual variability in responses across questions, implying that our estimated relation between consumption risk and source risks is unlikely to reflect equal answers to different questions.

5. Anatomy of consumption risk

5.1 Pass-through estimates

Table 3 presents our main findings. Column (1) reports the panel estimates for a variant

[13] Notice that here we refer to the cross-sectional correlation of say, the variance of income risk and the variance of health costs risk and not to the correlation between the subjective *distribution* of income growth and health costs growth. Since our survey elicits the subjective marginal probability distributions of the two variables not their joint distribution, we cannot make any claims about their correlation.

of equation (3).[14] Since some individuals exit the panel, the sample size decreases from 25,026 to 23,117 individual-wave observations. Because the regression includes fixed effects, the results can be interpreted as indicating that changes in quarterly consumption risk between October 2023 and October 2024 are partly explained by fluctuations in perceived income, energy, and health risks, as well as by changes in aggregate uncertainty. The regression also controls for demographic factors, including age, gender, family size, education, region of residence, occupation, cash-on-hand, and homeownership. However, given the high-frequency nature of our panel, these effects are largely absorbed by the fixed effects, so their coefficients are not reported in Table 3.

As already explained, the pass-through coefficients of the risk indictors reflect the weight of human capital in total wealth (π), and expected formal and informal insurance opportunities (α). All pass-through coefficients are positive, implying that – on average - consumers do not expect to be able to insure completely, either formally or informally against the risks they face.

In particular, the coefficient of income risk is positive and precisely estimated. The size of the coefficients indicates that 20% of the expected income variance is transmitted to the consumption variance. Also, the other coefficients of microeconomic risk (health and energy expenditures) are positive, precisely estimated, and significantly below one.

The pass through of health expenditure risk (0.37) is higher than that of income and energy expenditure risks. This is consistent with the evidence in Blundell et al. (2024) who find that the pass-through of transitory health shocks to consumption is exceeds that of transitory income shocks. The reason is that health shocks affect consumption in two ways: (i) directly, because shocks to health affect health related expenses with a very large pass-through, as they cannot typically be deferred; (ii) indirectly, because health status affects the marginal utility of consumption, discouraging expenditures when bad health hits, and vice versa. Energy expenditure shocks, like health shocks, are difficult to defer but have no direct effect on the marginal utility of consumption. Their pass-through is of the same magnitude as that of income risk.

Among the macro risks, only the coefficients of GDP and house price risks are positive and statistically different from zero, while interest rate risk, inflation risk, and unemployment risk have a minor influence. Table D4 in Appendix D shows that results are stable across the

[14] In the Appendix we regress the variance of consumption on the variances of the other risk indicators, separately for each of the five ISCE waves (Table C2). In these OLS regressions we exploit the cross-sectional variability of the risk indicators.

five waves for the income risk and health risk coefficients, and that in each of the subsamples the regression explains more than 70% of consumption volatility.

One potential concern is that the positive correlation between consumption risk and underlying risk factors may be affected by the individuals who report positive consumption variance but no variance in the source risks.[15] These responses are either inconsistent or arise from uncertainty sources not captured in our data. Excluding this group leaves the results unchanged (column 2, Table 3). As explained in Section 4.2, a second, larger group of respondents (27.9%) reports zero variance for each of the risk factors, that is, $I(\sigma_\xi^2) = 0$ and $I(\sigma_s^2) = 0$. In column 3 we drop these observations. Also in this restricted sample estimated pass-through coefficients very close to those in column (1).

In column (4) we further drop all observations reporting zero consumption risk even though they report positive values for one or more risk factors, that is, $I(\sigma_\xi^2) = 0$ and $I(\sigma_s^2) = 1$. In this restricted sample with 11,575 observations, the pass-through coefficient of income risk increases to 0.24 (from 0.21 in column 1), while the pass-through of health and energy risks are slightly lower than in column 1 (0.31 and 0.16, respectively). But qualitatively results are similar across sample selections.

Table 4 presents the panel fixed-effects estimates of equation (4), incorporating potential interactions between different risks—specifically, the product of the standard deviations of the subjective distributions, as specified in equation (4). We first estimate a model that includes the full set of covariances among the eight risk factors. We then refine the model by retaining only the two covariances with statistically significant coefficients.

Column 1 shows that the covariance between expected health expenditure growth and expected income growth has a positive coefficient, aligning with findings from Blundell et al. (2024) and Fulford and Low (2024) that suggest a positive correlation between health and income shocks. In contrast, the covariance between inflation and house price growth has a negative coefficient, consistent with the idea that housing serves as a hedge against inflation. We then report in column 2 of Table 4 the marginal effects and associated standard errors of the risk indicators, evaluated at the sample means of the expected variances of the risks

[15] In Section 4.2. we identify 1% of the cases for which $I(\sigma_\xi^2) > 0$ when $I(\sigma_s^2) = 0$.

considered.[16] Comparing the results in Tables 2 and 3, we conclude that the pass-through of the risk factors are quite stable even allowing for potential covariance effects.

5.2 Pass-through heterogeneity

Recall that the pass-through coefficients in equation (3) should vary across population groups with different risk exposure and access to insurance. While our data do not allow us to disentangle these two factors – meaning we cannot separately estimate the α_z and π_z coefficients, we can still assess their relevance by focusing on those groups where we would expect, a priori, we expect in both risk exposure (π coefficients) and insurance (α coefficients).

To identify these groups, the descriptive analysis in Section 4 is useful. For instance, individuals with high levels of human capital relative to wealth (the young with more career uncertainty) should be more exposed to income risk (have higher π_y^2 in equation 3) whereas those individuals with high levels of wealth relative to income (older and retired individuals) have a larger buffer and are less exposed to income risk. Similarly, compared to the self-employed, employees should be less exposed to business cycle fluctuations, as their wages tend to be more stable—a pattern reflected in the elicited expectations discussed in Section 4..

Table 5 presents separate pass-through regressions for different consumer groups based on age (≤35 and >60), occupation (employee vs. self-employed), and sector of activity. For the latter, we classify workers into "high-risk sectors," which include agriculture, manufacturing, and construction, and "low-risk sectors," which encompass services and public administration.

The pass-through coefficients align with a priori expectations regarding differences in risk exposure across groups. For instance, compared to older individuals (age ≥60), the pass-through of income risk is higher for younger individuals (age ≤35). Similarly, while the pass-through of unemployment risk is small and not statistically significant in the overall sample (Table 4), it becomes larger and significant among younger individuals but remains small and insignificant among older individuals. Additionally, the pass-through of house price risk is higher for the older group, which is consistent with their greater exposure to house price

[16] For instance, the marginal effect of income risk is evaluated as $\hat{\beta}_y + 2\hat{\beta}_{yh}\left(\frac{\sqrt{\bar{\sigma}_h^2}}{\sqrt{\bar{\sigma}_y^2}}\right)^2$, where $\bar{\sigma}_h^2$ and $\bar{\sigma}_h^2$ are, respectively, the average variances of the health and income distributions.

fluctuations due to higher homeownership rates. Employees exhibit a significantly lower pass-through of energy expenditure risk than the self-employed (0.165 vs. 0.271), reflecting the fact that self-employed individuals are more vulnerable to energy cost fluctuations, as energy serves as both a production input and a consumption good. Furthermore, the self-employed show a pass-through of GDP growth risk that is twice as large as that of employees. Consistent with these findings, individuals employed in high-risk sectors exhibit a higher pass-through of income risk compared to those in low-risk sectors (0.33 vs. 0.22).

Table 6 presents estimates based on cash-on-hand as a measure of heterogeneity in access to self-insurance. Individuals with low cash-on-hand (less than €10,000) exhibit a 35% higher pass-through of income risk to consumption volatility compared to those with high cash-on-hand (above €30,000). Additionally, their pass-through of health and energy expense risks is 9% and 12% higher, respectively. These findings support the hypothesis that cash-on-hand provides a form of self-insurance against income and expenditure shocks, particularly income shocks. Overall, the evidence in Tables 5 and 6 reinforces our decomposition exercise and suggests that both risk exposure and access to insurance play key roles in explaining the determinants of consumption risk.

The last two columns of Table 6 highlight that the pass-through of income risk is significantly higher for younger individuals with low cash-on-hand (0.295). This indicates that income risk has a greater impact when financial buffers are limited and individuals are in the early stages of their careers, where income volatility is typically higher. Conversely, the pass-through of income risk declines considerably for older individuals with greater cash reserves (0.127). Furthermore, macroeconomic risks have a much lower impact on consumption volatility among younger individuals with low cash-on-hand compared to older individuals with substantial financial buffers.

5.3 Contribution of risk sources to consumption uncertainty

To evaluate the economic significance of different sources of risk, Table 7 presents the contribution of each risk factor to predicted consumption risk. For the full sample, we use the pass-through coefficients estimated in Table 3 (column 1) and calculate contributions based on the mean values of the source risk variables. Specifically, for each risk source, its contribution to predicted consumption risk is obtained by multiplying the estimated pass-through coefficient by the sample mean of that risk's variance and dividing by the sample average of consumption risk. We differentiate between idiosyncratic and aggregate risks and provide estimates for the full sample as well as various consumer sub-groups classified by risk exposure (young vs. old; employed vs. self-employed; high-risk vs. low-risk sectors) and self-insurance capacity (low vs. high cash-on-hand).

Panel A of Table 7 shows that, for the full sample, idiosyncratic risks account for the majority of consumption risk (72.8%), while aggregate risks contribute only 17.4%, and demographic factors and time effects together explain the remaining 9.8%. Notably, health expenditure risk emerges as the largest contributor to expected consumption risk (33.6%). Since the average variance of health risk is lower than that of income risk (see Table 1), its large contribution is driven by the higher pass-through coefficients documented in Tables 4 to 6, consistent with findings in Blundell et al. (2024). Similarly, while energy risk has a lower average variance than income risk, its substantial pass-through coefficients highlight the significance of expenditure shocks, supporting the results of Fulford and Low (2024).

In the remaining columns of Panel A in Table 7, we further decompose consumption risk across different groups, highlighting heterogeneity in risk exposure. The share of consumption risk attributable to individual-level risks varies significantly, ranging from a low of 61% among individuals aged 60 and above to a high of 77% for those working in high-risk sectors. We also observe relatively high contributions of micro risks for younger individuals (77%) and the self-employed (75%). In contrast, aggregate risks play a larger role for older individuals (25%) but are less significant for younger individuals and those in high-risk sectors (16%). For other groups, aggregate risks contribute between 15% and 18% of total consumption risk. Demographic factors and time effects account for approximately 10% of consumption risk.

Panel B shifts the focus to differences in self-insurance capacity. Here, income risk is shown to be more important for individuals with low cash-on-hand, while macroeconomic risks are substantially more significant for wealthier individuals with greater financial buffers.

The findings in Table 7 indicate that consumption uncertainty affects all population segments, though the relative importance of different risk sources varies significantly across groups. This variation underscores the role of precautionary savings, though the motivations for accumulating such savings differ based on life stage, occupation, employment sector, and individual capacity to manage financial risks.

5.4. Willingness to pay to fully insure residual consumption risk

As a final exercise, we leverage data on consumption uncertainty and its sources to estimate individuals' willingness to pay for insurance against consumption risk. Following Lucas (2003), this can be computed as $\left(\frac{1}{2}RRA \times \sigma_c^2\right)$, representing the risk premium a consumer with risk aversion *RRA* would be willing to pay to insure against consumption growth volatility of size σ_c^2.

Since we observe σ_c^2 , we can back up the contribution from business cycle fluctuations (Lucas focus) and from idiosyncratic uncertainty. However, it is important to note that our measure of σ_c^2 reflects *residual* consumption risk, i.e., the portion that remains after consumers have utilized available insurance mechanisms to mitigate the pass-through of source risks onto consumption uncertainty. This distinction underscores that the estimated willingness to pay represents compensation for risks that individuals are unable to fully insure against through existing financial or informal risk-sharing arrangements.

From Table 7, we observe that σ_c^2 =0.057%. In Section 6, our Euler equation estimates suggest a relative risk aversion (*RRA*) of 1.4. Therefore, the willingness to pay to avoid residual consumption uncertainty for the average consumer in our sample is $0.5 \times 1.4 \times 0.057 = 0.04\%$ of consumption. The contribution to consumption uncertainty from micro and macro risks are 0.728 and 0.174, respectively. Thus, the willingness to pay to insure micro risks amounts to 0.029% of consumption, while the willingness to pay to insure aggregate fluctuation to 0.011%. This constitutes a small welfare cost, confirming Lucas' (2003) conclusion that the willingness to pay for insurance against consumption uncertainty is relatively low.

6. Euler equation estimates

The first step of our analysis has shown that the expected volatility of consumption can be attributed to a variety of risks, with pass-through coefficients consistently below 1, reflecting a combination of risk exposure and insurance opportunities. However, the impact of consumption risk on expected consumption growth remains to be explored. In this section, we use our data on consumption risk and its determinants to estimate an Euler equation for consumption. The goal is to assess whether expected consumption volatility influences the expected rate of consumption growth and to measure the strength of the precautionary savings motive.

Following Blanchard and Mankiw (1988), we approximate the Euler equation using a second-order Taylor expansion of the marginal utility of consumption $u'(c_{t+1})$ around c_t. Solving for the expected growth rate of consumption we obtain:

$$E_t\left(\frac{c_{t+1}-c_t}{c_t}\right) \cong EIS\left(\frac{r-\delta}{1+r}\right) + \frac{1}{2}p(c)E_t\left(\frac{c_{t+1}-c_t}{c_t}\right)^2 + \psi_{it+1} \quad (5)$$

where $p(c) = -\frac{u'''(c_t)c_t}{u''(c_t)}$ is Kimball's coefficient of relative prudence, $EIS = -\frac{u'(c_t)}{u''(c_t)c_t}$ is the elasticity of intertemporal substitution which in this framework is also equal to the inverse of the coefficient of relative risk aversion, and ψ is a remainder term in the Taylor approximation. The second moment of the conditional distribution of consumption growth $E_t\left(\frac{c_{t+1}-c_t}{c_t}\right)^2$ is a measure of the expected variability of consumption.

Equation (5) shows that the anticipated volatility of consumption growth is linked to a higher growth rate of consumption. Consumption uncertainty prompts consumers to reduce current consumption and increase savings, with the extent of this adjustment depending on their degree of prudence.[17] There two ways to estimate equation (5). One approach is to use realized consumption data, as pioneered by Dynan (1993). However, this method faces significant identification challenges due to the endogeneity of realized consumption and the difficulty in finding valid instruments (see, for example, Carroll, 2001). Our proposed alternative, following Christelis et al. (2020), is to rely on data regarding consumers' subjective expectations of the distribution of consumption growth. This allows us to obtain measures of the key variables in equation (5) – the expected rate of consumption growth and its expected second moment, both conditional on the information available to consumers. We assume that prudence is a parameter invariant to wealth (as implied by CRRA preferences), and express the equation in a regression framework:

$$E_{it}(g_{i,t+1}) = \gamma_0 + \omega E_{it}(g_{i,t+1}^2) + f_i + d_{t+1} + v_{i,t+1} \quad (6)$$

[17] In the certainty equivalence model consumers do not respond to uncertainty and $p(c) = 0$. Caballero (1991) demonstrates that if utility is exponential and income follows a random walk, then expected consumption growth equals the product of the coefficient of relative prudence and the variance of income, normalized by current income. A closed-form solution can also be derived if utility is isoelastic and consumption growth is normally distributed.

where $E_{it}(g_{i,t+1})$ and $E_{it}(g_{i,t+1}^2)$ represent, respectively, the first and second moment of the distribution of expected consumption growth in period t+1, $(g_{i,t+1})$; f_i and d_{t+1} are individual and time fixed effects, and $v_{i,t+1}$ is an error term. The individual fixed effects capture systematic differences in returns to wealth, as documented in Fagereng et al. (2020). The parameter ω equals ½ of the degree of relative prudence and thus identifies the strength of the precautionary saving motive. In some specification we allow $\omega = \omega(z_i)$ to vary depending on observables z_i that the literature has shown to correlate with individual risk aversion. We do this to increase confidence that we are identifying a structural preference parameter.

Notice that the error term $v_{i,t+1}$ of equation (6) is not a forecast error, which allows equation (6) to avoid endogeneity problems that typically affect estimates based on ex-post consumption realizations. Additionally, systematic deviations of expectations from realizations do not undermine the consistency of our estimates, even with short panels.[18] However, there remains the possibility that unobservable variables included in the error term are correlated with consumption risk (this could arise, for example, from omitted higher-order terms in the Taylor expansion). Furthermore, subjective expectations might be influenced by measurement error, which could also affect our estimates in Table 8 presents the estimates of two versions of equation (6): (i) OLS without and with fixed-effects panel estimation; (ii) IV fixed-effects panel estimation, using the second moments of risk factors as instruments. The IV specification is our preferred one, as it allows us to trace the effect of consumption volatility back to the underlying risk sources. Column (1) shows a basic OLS regression The coefficient of consumption risk is precisely estimated at 1.137, implying a coefficient of relative prudence of 2.27. This value is in line with previous evidence based on different methodologies and datasets.[19]

The Euler equation (6) is derived under the assumption of perfect capital markets, but this

[18] If one uses realized consumption data, one can write equation (6) in a regression framework as: $g_{i,t+1} = \gamma_0 + +\omega g_{i,t+1}^2 + f_i + d_t + v_{i,t+1} + e_{i,t}$. But in this case the term $v_{i,t+1}$ is the difference between realized and expected consumption (the forecast error), which is likely to be correlated with $g_{i,t+1}^2$. For instance, if households receive positive news, they may revise consumption upward, affecting both the mean and the variance of the (ex post) consumption distribution. Consequently, any regression of realized consumption growth on realized consumption volatility produces inconsistent estimates of the strength of the precautionary motive (Carroll, 2001). Another consequence of the presence of the expectational error is that estimates derived from short panel data may be inconsistent (Chamberlain, 1984).

[19] Dynan (1993) estimates the Euler equation using realized consumption data and an instrumental variable approach. Bertola et al. (2005) use the subjective variance of income growth from the SHIW as an instrument for realized consumption variability, finding a coefficient of relative prudence of approximately 2. Similarly, Fagereng et al (2017) estimate a prudence coefficient of 2 by instrumenting income variance with the variance of shocks to the firm the consumer is employed. Christelis et al. (2020) regress expected consumption growth on expected consumption risk in Dutch data and also obtain a relative prudence estimate of 2.

equation may not hold in the presence of liquidity constraints or myopic consumers. To capture these individuals, we assume that some consumers set consumption equal to income in each period. Column 2 nests the two models and adds expected income growth to the regressors. The coefficient of consumption risk is 2.06 (implying relative prudence of 4.12), and the coefficient of income growth is 0.29, consistent with a large body of empirical evidence showing that consumption growth is sensitive to expected income growth.[20]

In column 3 of Table 8 we add individual fixed effects, which capture heterogeneity in returns on wealth across individuals. The estimate on consumption risk, (-1.76), is of the same order of magnitude as in the OLS estimates, yielding a measure of prudence of 3.52.

Column 4 presents IV fixed-effects estimates, using the second moments of the expected distributions of our microeconomic risk factors (expected income growth, expected healthcare expenditure growth, and expected growth of energy costs) as instruments. Column 5 also includes in the instrument list macroeconomic sources of risk, such as GDP, inflation, interest rate, unemployment, and house price risks. [21] In both cases, the effect of consumption risk is precisely estimated, with implied prudence ranging between 2.8 and 2.5. We take the latter as our reference estimate. Excess sensitivity remains nearly unchanged across both specifications and is also precisely estimated.[22]

Panel B of Table 8 exploits heterogeneity in risk aversion based on observable characteristics to strengthen the structural interpretation of the estimated prudence parameter. Under CRRA preferences, relative prudence is directly related to relative risk aversion, following the relation: prudence = 1+ RRA. To explore this relationship, we split the sample in two ways. First, columns (1) and (2) distinguish between younger individuals (age<=35) and older ones (age>60). Second, columns (3) and (4>) split the sample by occupational risk, given that risk aversion is positively related to and age, and influences job selection – more risk averse individuals are more likely to choose safer jobs, such as positions in the public sector (e.g. Albert Duffy, 2012). The Euler equation estimates reveal a clear pattern: prudence is significantly lower among the young (approximately 2.2) compared to the old (5.8), consistent with the idea that older individuals tend to be more risk averse. Similarly, prudence is lower for

[20] For further references, see Jappelli and Pistaferri (2017), Havranek and Sokolova (2020), Crawley and Theloudis (2024).
[21] Table C3 in Appendix C reports the first stage regressions corresponding to columns 1 and 2 in Table 8.
[22] As robustness check, we exclude observations where consumption risk, risk factors, or both equal to zero, following the approach in Table 3. This exclusion does not affect the results, particularly the coefficient of prudence.

individuals employed in low-risk sectors (around 0.6, though with a high standard error) than for those employed in high-risk sectors (approximately 5.2).

7. Size and drivers of precautionary saving

Using our estimate of the Euler equation, along with information on subjective distributions of consumption and the underlying risks, we assess how precautionary savings influence consumption and its sensitivity to various risks that drive expected consumption volatility. The detailed methodology and step-by-step approach are provided in Appendix A2. Here, we outline the key intuition and discuss the results.

For an individual of age a, the Euler equation estimates the expected consumption growth rate, which consists of two main components: (i) a precautionary saving component, $\omega\sigma_a^2$, which corresponds to the second term on the right-hand side of equation (6). The parameter ω is estimated in Table 8, while the second moment of expected consumption growth σ_a^2 is derived from the elicited distributions of future consumption growth. The second component ($\alpha = \gamma_0 + f_i$) captures other motives for savings, including intertemporal substitution and heterogeneity in returns to capital, which is accounted for by the individual fixed effects. Our assumption is that this component remains unchanged when the precautionary motive is removed.

Following Caballero (1990, 1991) we assume that uncertainty influences the volatility of future resources but does not alter their expected value over the remaining lifetime. Under this assumption, uncertainty affects only the slope of the age profile of expected consumption from age a onwards, while leaving its total lifetime value unchanged. Since $\omega > 0$ current consumption under uncertainty, C_a^u, will be lower than current consumption under certainty, C_a. In the appendix we derive the ratio of these two consumption levels:

$$\frac{C_a^u}{C_a} = \frac{\sum_{k=0}^{T_a-a}[\frac{1+\alpha_a}{1+r}]^k}{\sum_{k=0}^{T_a-a}[(1+\alpha_a+\beta\sigma_a^2)/(1+r)]^k} = \mu_a < 1$$

Here r represents the real rate of interest used to discount expected consumption over the life cycle. Thus $(1-\mu_a)$ is the level of precautionary savings relative to consumption under certainty. Following the prediction in Caballero (1991), we can test twhether precautionary savings decreases as an individual's horizon shortens.[23]

In Table 9 we report the average values of $(1-\mu_a)$ for the entire sample and for three age groups: young (ages 18-24), middle-aged (40-45), and pre-retirement (60-65). To estimate the consumer's horizon, we use life expectancy tables (separately for males and females) and set the discount rate at 3%. We obtain the relevant parameters from the first stage estimates of regression (5) in Table 8 (column 5) and evaluate $(1-\mu_a)$ at the sample means of the risk factors used as instruments. On average, precautionary savings accounts for 2.74% of consumption, with individual-level risks contributing nearly 2 percentage points and aggregate risk accounting for 0.8 percentage points. These results remain consistent when performing the calculations separately for males and females, or when assuming that the discount rate is 2 or 4%.

For the 18-24 age group, $(1-\mu_a)=4.2\%$, higher than the sample average. In terms of the sources of precautionary saving, 69% stems from microeconomic risks, while the remaining 31% from aggregate risks. For the middle-aged and the pre-retirement groups precautionary savings decline notably with age (3.3% and 1.7%, respectively), in line with predictions from precautionary savings models. There is also a mild age-related shift in the composition of these savings. While idiosyncratic risk accounts for 69% of precautionary savings among the young, this share rises to 74% among older individuals, largely due to the increasing importance of energy and health expenditure risks in later life.

Finally, we can contrast the welfare cost of these fluctuations estimated in Section 5 with the amount of precautionary savings that consumers are willing to pile up to cushion business cycle fluctuations. From Table 9 this amounts to 0.78% of consumption, much higher than the willingness to pay to insure against aggregate fluctuations (0.007%). This is what it should be, since paying to avoid the business cycle entails a sure loss of consumption while precautionary savings entail only procrastination of consumption.

7. Conclusions

Intertemporal consumption models suggest that consumers' decisions depend on the

[23] In Caballero (1990) the statistic $(1-\mu_a)$ has a closed-form solution given by $\frac{p}{2}\sigma^2(T-a)$ where p represents the degree of relative prudence, σ^2 is the squared coefficient of variation of the innovation to the income process, assumed to be normally distributed, and $T-a$ denotes the remaining horizon for an individual aged a.

expected variability of consumption, which, in turn, reflects all relevant sources of uncertainty they face. However, most of the literature has primarily focused on income risk and the strategies consumers use to mitigate it. While income uncertainty is a major concern, it is not the only source of risk affecting consumption. Focusing solely on income risk overlooks other important uncertainties. This paper fills that gap by eliciting the probability distribution of expected consumption risk, along with a comprehensive set of underlying risk sources. This approach allows us to uncover several novel findings that enhance our understanding of the anatomy of consumption risk:

First, we find significant heterogeneity in both consumption risk and its underlying sources. About 28% of consumers report no uncertainty from risk sources and thus experience no consumption risk. Another 20% also report no consumption risk, possibly due to their ability to insure against the risks they face—which typically happens when risks are limited in numbers and small. However, the majority of consumers experience multiple sources of risk that, in turn, influence their expected consumption volatility.

Second, the extent to which risk sources affect consumption varies between zero and one for all risks, but is significantly higher for individual-specific risks than for aggregate risks. Third, the pass-through of expenditure risks arising from health and energy price shocks is at least as large as that of income risk. Fourth, when considering both the size of risks and their pass-through rates, we find that individual-specific risks account for most of consumption uncertainty, whereas aggregate risks (including business cycle fluctuations) contribute less than a quarter of the overall consumption uncertainty. Finally, we document substantial heterogeneity in both pass-through rates and the relative importance of different risk sources, reflecting differences in risk exposure (e.g., age and employment sector) and access to insurance (e.g. liquid resources available for self-insurance).

Using our data, we estimate the structural parameter of the Euler equation for consumption and leverage it to quantify precautionary savings flows, examining how different risks contribute to these savings and how they vary across consumers. Our findings suggest that precautionary savings account for slightly less than 3% of consumption, with about 2 percentage points allocated to buffer individual-level risks and 1 percentage point in response to aggregate risk. Precautionary savings are higher among younger consumers (~4%), a group more exposed to microeconomic risks.

We see this study as a first step in a broader research agenda aimed at evaluating the welfare costs of limited insurability when consumers face multiple sources of risk. This requires a full-fledged structural life-cycle model that accounts for heterogeneous exposure and insurability across multiple risks. Such a model, which will be calibrated using our empirical findings, will be developed in a separate study.

References

Albert, Steven M, John Duffy (2012), "Differences in Risk Aversion Between Young and Older Adults", *Neuroscience and Neuroeconomics* 2012, 3–93.

Armantier, Olivier, Giorgio Topa, Wilbert van der Klaauw, Basit Zafar (2016), "An Overview of the Survey of Consumer Expectations," Federal Reserve Bank of New York Staff Reports No. 800.

Attanasio, Orazio P., Guglielmo Weber (2010), "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy", *Journal of Economic Literature* 48, 693-751.

Attanasio, Orazio P., Hamish Low (2004), "Estimating Euler Equations", *Review of Economic Dynamics* 7, 406-435.

Attanasio, Orazio, Agnes Kovacs, Krisztina Molnar (2019), "Euler Equations, Subjective Expectations and Income Shocks", *Economica* 87, 406-441.

Arellano, Manuel, Orazio Attanasio, Margherita Borella, Gonzalo Paz-Pardo (2024), "Subjective Earnings and Employment Dynamics", mimeo.

Bachmann, Ruediger, Giorgio Topa, Wilbert van der Klaauw (2022), *Handbook of Economic Expectations*, Elsevier.

Banks, James, Richard Blundell, Agar Brugiavini (2001), "Risk Pooling, Precautionary Saving and Consumption Growth," *The Review of Economic Studies* 68, 757–779.

Ben-David, Itzhak, Elyas Fermand, Camelia M. Kuhnen, Geng Li (2018), "Expectations Uncertainty and Household Economic Behavior," NBER Working Papers 25336.

Bertola, Giuseppe, Luigi Guiso, Luigi Pistaferri (2005), "Uncertainty and Consumer Durables Adjustment," *Review of Economic Studies* 72, 973–1007.

Blanchard, Olivier and Gregory Mankiw (1988), "Consumption: Beyond Certainty Equivalence," *The American Economic Review P&P* 78, 173-177.

Blundell, Richard, Margherita Borella, Jeanne Commault, Mariacristina De Nardi (2024), "Old Age Risks, Consumption, and Insurance," *American Economic Review* 114, 575-613.

Boguth Oliver and Lors-Alexander Kuehn (2013), "Consumption Volatility Risk", *The Journal of Finance*, 2589-2615.

Caballero, Ricardo (1990), "Consumption Puzzles and Precautionary Savings," *Journal of Monetary Economics* 25, 113-136.

Caballero, Ricardo (1991), "Earnings Uncertainty and Aggregate Wealth Accumulation," *American Economic Review* 81, 859-71.

Caplin, Andrew, Victoria Gregory, Eungik Lee, Søren Leth-Petersen and Johan Sæverud (2023) "Subjective Earnings Risk", NBER Working Paper 31019.

Carroll, Christopher D. (2001), "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)," *Topics in Macroeconomics* 1, Article 6.

Carroll, Christopher D., and Miles S. Kimball, (1996), "On the Concavity of the Consumption Function," *Econometrica* 64, 981–92.

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, Matthew White (2017), "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics* 8, 977-1020.

Chamberlain, Gary (1984), "Panel Data," in Zvi Griliches and Michael. D. Intriligator, eds. *Handbook of Econometrics*, vol. 2. Amsterdam: Elsevier Science.

Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Maarten van Rooij (2020), "Consumption Uncertainty and Precautionary Saving," *The Review of Economics and Statistics* 102, 148–161.

Coibion, Olivier, Yuriy Gorodnichenko, Geoff Kenny, Michael Weber (2024), "The Effect of Macroeconomic Uncertainty on Household Spending," *American Economic Review* 114, 645-77.

Crawley, Edmund, Theloudis Alexandros (2024), "Income Shocks and their Transmission into Consumption," CentER Discussion Paper 2024-012, forthcoming in the Encyclopedia of Consumption.

De Nardi, Maria Cristina, Giulio Fella (2017), "Saving and Wealth Inequality," *Review of Economic Dynamics* 26, 280-300.

Dominitz, Jeff, Charles F. Manski (1997), "Using Expectations Data to Study Subjective Income Expectations," *Journal of the American Statistical Association* 92, 855–867.

Dynan, Karen, (1993) "How Prudent Are Consumers?" *Journal of Political Economy* 101, 1104–1113.

ECB (2021), "ECB Consumer Expectations Survey: An Overview and First Evaluation," Occasional Paper Series, No. 287, December.

Fagereng, Andreas, Luigi Guiso and Luigi Pistaferri (2017), "Firm-Related Risks and Precautionary Savings Response", *American Economic Review P&P* 107, 393-97.

Fagereng, Andreas, Luigi Guiso and Luigi Pistaferri (2017), "Heterogeneity and Persistence in Returns to Wealth", *Econometrica* 88, 115-170.

Fagereng, Andreas, Luigi Guiso, Luigi Pistaferri, Marius Ring (2024), "Insuring Labor Income Shocks: The Role of the Dynasty", EIEF Working Paper.

Fulford, Scott L., David Low (2024), "Expense Shocks Matter", Consumer Financial Protection Bureau Office of Research Working Paper No. 2024-08

Georgarakos, Dimitris, Kwang Hwan Kim, Olivier Coibion, Myungkyu Shim, Myunghwan Andrew Lee, Yuriy Gorodnichenko, Geoff Kenny, Seowoo Han, Michael Weber (2025), How Costly Are Business Cycle Volatility and Inflation? A Vox Populi Approach, NBER Working Paper 33476.

Gourinchas, Pierre-Olivier, Jonathan A. Parker (2002), "Consumption over the Life Cycle," *Econometrica* 70, 47–89.

Guiso, Luigi, Tullio Jappelli (2025), "The Italian Survey of Consumer Expectations: Statistical Bulletin," https://isrlab.it/italian-survey-consumer-expectations.

Guiso, Luigi, Tullio Jappelli, Daniele Terlizzese (1992), "Earnings Uncertainty and Precautionary saving," *Journal of Monetary Economics* 30, 307-337.

Guiso, Luigi, Tullio Jappelli, Mario Padula (2013), "Pension Wealth Uncertainty," *Journal of Risk and Insurance* 80, 1057–1085.

Guiso, Luigi, Tullio Jappelli, Luigi Pistaferri (2022), "An Empirical Analysis of Earnings and Employment Risk," *Journal of Business and Economic Statistics* 20, 241–253.

Havranek, Tomas, Anna Sokolova (2020), "Do Consumers Really Follow a Rule of Thumb? Three Thousand Estimates from 144 Studies Say Probably Not," *Review of Economic Dynamics* 35, 97-122.

Jappelli, Tullio, Luigi Pistaferri (2000), "Using Subjective Income Expectations to Test for Excess Sensitivity of Consumption to Predicted Income Growth," *European Economic Review* 44, 337–358.

Jappelli, Tullio, Luigi Pistaferri (2017), *The Economics of Consumption*. Oxford: Oxford University Press.

Kaplan, Greg, Giovanni L. Violante, Justin Weidner (2014), "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity* 48 (2014),77–153.

Lucas, Robert, E Jr. (2003), "Macroeconomic Priorities," *American Economic Review* 93, 1–14.

Ludvigson Sydney C. (2013), "Advances in Consumption-Based Asset Pricing: Empirical Tests," in *Handbook of the Economics of Finance*, edited by George M. Constantinides, Milton Harris and Rene M. Stulz, vol. 2, pp. 799-906. Elsevier Science B.V., North Holland, Amsterdam.

Ludvigson, Sydney, Christina Paxon (2001), "Approximation Bias in Linearized Euler Equations", *The Review of Economics and Statistics* 83, 242-256.

Manski, Charles F. (2004), "Measuring Expectations," *Econometrica* 72, 1329–1376.

Palumbo, Michael G. (1999), "Uncertain Medical Expenses and Precautionary Saving near the End of the Life Cycle," *Review of Economic Studies* 66, 395–421.

Ryngaert, Jane M. (2022), "Inflation Disasters and Consumption," *Journal of Monetary Economics* 129, S67-S81.

Stantcheva, Stefanie (2023), "How to Run Surveys: A Guide to Creating your Own Identifying Variation and Revealing the Invisible," *Annual Review of Economics* 15, 205-234.

Violante, Gianluca (2024), "Consumption," in *Macroeconomics: A Comprehensive Textbook for First-year Ph.D. Courses in Macroeconomics*, edited by Marina Azzimonti, Per Krusell, Alisdair McKay, and Toshihiko Mukoyama. Unpublished (https://phdmacrobook.org).

Table 1. Descriptive statistics

	Mean	Standard deviation
Male	.493	.5
Age	48.379	14.299
Family size	2.784	1.133
College degree	.222	.415
North	.464	.499
Centre	.195	.396
South	.341	.474
Employed	.449	.497
Self-employed	.088	.283
Cash-on-hand	29.844	23.185
Homeowner	.758	.428
Consumption growth	.007	.039
Income growth	-.01	.035
Health expenditures growth	.008	.035
Energy expenditures growth	.013	.036
Nominal interest rate	.031	.023
GDP growth	-.013	.036
Inflation	.013	.038
Unemployment rate	.09	.039
House price growth	.001	.032
	Variance of	**Second moment of**
Consumption growth	.057	.211
Income growth	.064	.195
Health expenditures growth	.051	.182
Energy expenditures growth	.047	.195
Nominal interest rate	.012	.157
GDP growth	.048	.194
Inflation	.045	.206
Unemployment rate	.033	.999
House price growth	.043	.143
Observations	25,026	25,026

Note. The table reports the sample means and standard deviations of selected variables from the ISCE. The average growth rates and their standard deviations are presented as the original numbers, while the variances are multiplied by 100. All statistics are computed using sample weights from the pooled 2023-24 ISCE.

Table 2. Distribution of consumption risk, by risk factors

Panel A. Cross-tabulation of $I\left(\sigma_\xi^2\right)$ and $I(\sigma_s^2)$

$I\left(\sigma_\xi^2\right)$	$I(\sigma_s^2)$		Total
	0	1	
0	6,982 (27.90)	5058 (20.21)	12,040 (48.1)1
1	257 (1.03)	12,729 (50.86)	12,986 (51.89)
Total	7,239 (28.93)	17,787 (71.07)	25,026 (100)

Panel B. Average risk factors, by consumption risk groups

	$I(\sigma_s^2) = 1$ and $I\left(\sigma_\xi^2\right) = 0$	$I(\sigma_s^2) = 1$ and $I\left(\sigma_\xi^2\right) = 1$
Income risk	.039	.111
Health risk	.008	.098
Energy risk	.007	.089
Interest rate risk	.006	.022
GDP risk	.017	.087
Inflation risk	.013	.082
Unemployment risk	.015	.058
House price risk	.005	.083
Average risk	.014	.079
Average number of positive risks	2.59	6.04
N	5,058	12,729

Panel C. Probability of reporting zero consumption risk in the sample with some positive source risk

	Coefficient	Standard error
$sum\_\sigma_s^2$	-0.504	(0.059)***
$num\_I(\sigma_s^2) = 1$	-0.087	(0.002)***
R^2	0.533	
N	16,251	

Note. Panel A presents a cross-tabulation of two dummy variables: $I(\sigma_s^2)$ equals one if the variance of at least one risk source is positive, and zero otherwise; $I\left(\sigma_\xi^2\right)$ equals one if the variance of expected consumption growth is positive, and zero otherwise. Percentage are reported in parenthesis. Panel B shows the averages of risk sources and the average number of positive risk factors when the variance of consumption growth is zero or positive. Statistics are computed sample weights. Panel C reports results from a linear regression of the probability of zero consumption risk on the weighted sum of the variances of source risk and on the number of non-zero-variance source risks.

Table 3. Determinants of consumption risk

	Whole sample	Excluding $I(\sigma_\xi^2) = 0$ when $I(\sigma_s^2) = 1$	Excluding also $I(\sigma_\xi^2) = 0$ when $I(\sigma_s^2) = 0$	Excluding also $I(\sigma_\xi^2) = 0$ when $I(\sigma_s^2) = 1$
	(1)	(2)	(3)	(4)
Income risk	0.206	0.207	0.191	0.240
	(0.013)***	(0.013)***	(0.014)***	(0.019)***
Health risk	0.373	0.373	0.360	0.314
	(0.022)***	(0.022)***	(0.022)***	(0.025)***
Energy risk	0.195	0.196	0.202	0.157
	(0.023)***	(0.023)***	(0.024)***	(0.024)***
Interest rate risk	0.063	0.065	0.052	0.053
	(0.056)	(0.056)	(0.060)	(0.074)
GDP risk	0.074	0.075	0.070	0.065
	(0.018)***	(0.018)***	(0.019)***	(0.022)***
Inflation risk	0.008	0.008	0.016	0.010
	(0.020)	(0.020)	(0.021)	(0.024)
Unemployment risk	0.024	0.024	0.020	0.014
	(0.021)	(0.021)	(0.023)	(0.029)
House price risk	0.104	0.103	0.105	0.084
	(0.022)***	(0.022)***	(0.023)***	(0.025)***
R^2	0.57	0.58	0.53	0.45
N	23,117	22,885	16,251	11,575

Note. The variable $I(\sigma_\xi^2)$ equals one for observations where the variance of expected consumption growth is positive, and zero otherwise. The variable $I(\sigma_s^2)$ equals one when at least one source of risk has a positive variance, and zero if all sources of risk have zero variance. All regressions are panel fixed effects estimates and include time fixed effects and demographic variables. Standard errors in parentheses. * significance at 10%, ** significance at 5%, *** significance at 1%.

Table 4. Determinants of consumption risk, with interaction terms

	Coefficients	Marginal effects
	(1)	(2)
Income risk	0.143	0.247
	(0.017)***	(0.017) ***
Health risk	0.249	0.378
	(0.036)***	(0.222)***
Energy risk	0.184	0.184
	(0.023)***	(0.023)***
Interest rate risk	0.065	0.065
	(0.055)	(0.055)
GDP risk	0.070	0.070
	(0.017)***	(0.017)***
Inflation risk	0.063	-0.004
	(0.028)**	(0.020)
Unemployment risk	0.026	0.026
	(0.021)	(0.021)
House price risk	0.155	0.085
	(0.033)***	(0.021)***
s.d.(income growth)×s.d(health exp. growth)	0.231	
	(0.048)***	
s.d.(inflation)×s.d.(house price growth)	-0.137	
	(0.050)***	
R^2	0.58	
N	23,117	

Note. Coefficients in column (1) represents panel fixed effects estimates. Column (2) reports marginal effects and their associated standard errors, evaluated at the sample means of the expected variances of the risks considered. All regressions include demographic variables. Standard errors in parentheses. * significance at 10%, ** significance at 5%, *** significance at 1%.

Table 5. Pass-through to consumption volatility. Sample splits by employment and age

	Age<=35	Age>60	Employees	Self-employed	Low risk sector	High risk sector
Income risk	0.186	0.143	0.216	0.194	0.220	0.326
	(0.022)***	(0.026)***	(0.020)***	(0.047)***	(0.029)***	(0.055)***
Health risk	0.388	0.372	0.374	0.351	0.352	0.272
	(0.036)***	(0.054)***	(0.032)***	(0.058)***	(0.040)***	(0.061)***
Energy risk	0.231	0.168	0.165	0.271	0.142	0.129
	(0.043)***	(0.050)***	(0.034)***	(0.057)***	(0.034)***	(0.060)**
Interest rate risk	-0.010	-0.016	0.072	0.190	0.015	-0.123
	(0.106)	(0.103)	(0.088)	(0.205)	(0.124)	(0.170)
GDP risk	0.030	0.099	0.061	0.126	0.057	0.123
	(0.036)	(0.032)***	(0.024)**	(0.063)**	(0.029)**	(0.065)*
Inflation risk	-0.016	0.031	0.014	-0.026	0.034	0.002
	(0.039)	(0.032)	(0.028)	(0.066)	(0.035)	(0.076)
Unemployment risk	0.106	0.028	0.032	-0.044	0.007	0.013
	(0.047)**	(0.032)	(0.036)	(0.080)	(0.049)	(0.084)
House price risk	0.096	0.198	0.117	0.084	0.078	0.074
	(0.042)**	(0.056)***	(0.031)***	(0.064)	(0.043)*	(0.052)
R^2	0.56	0.60	0.55	0.58	0.45	0.44
N	5,592	5,779	10,493	2,041	4,601	1,804

Note. All regressions use the panel fixed effects estimator and include demographic variables and time effects. High-risk sectors includes respondents working in agriculture, manufacturing and construction, while low-risk sectors include services and public administration. Standard errors in parentheses. *significance at 10%, **significance at 5%, ***significance at 1%
.

Table 6. Pass-through to consumption volatility. Sample splits by cash-on-hand

	Cash-on-hand<=10	10<Cash-on-hand<=30	Cash-on-hand>30	Cash-on-hand<=10 and Age<=35	Cash-on-hand>30 and Age>60
Income risk	0.265	0.202	0.197	0.295	0.127
	(0.034)***	(0.020)***	(0.028)***	(0.062)***	(0.038)***
Health risk	0.327	0.446	0.300	0.288	0.221
	(0.050)***	(0.034)***	(0.032)***	(0.103)***	(0.052)***
Energy risk	0.245	0.175	0.218	0.386	0.194
	(0.067)***	(0.036)***	(0.040)***	(0.137)***	(0.070)***
Interest rate risk	0.049	0.095	-0.022	-0.040	0.201
	(0.142)	(0.090)	(0.098)	(0.216)	(0.152)
GDP risk	0.024	0.125	-0.000	0.025	0.103
	(0.036)	(0.028)***	(0.031)	(0.076)	(0.051)**
Inflation risk	-0.023	-0.011	0.052	0.072	-0.009
	(0.050)	(0.030)	(0.040)	(0.087)	(0.062)
Unemployment risk	0.048	0.005	0.033	0.051	-0.026
	(0.052)	(0.031)	(0.043)	(0.098)	(0.056)
House price risk	0.105	0.055	0.140	-0.004	0.242
	(0.045)**	(0.033)*	(0.051)***	(0.086)	(0.108)**
R^2	0.56	0.61	0.52	0.61	0.48
N	4,874	10,294	7,949	1,110	2,728

Note. All regressions use the panel fixed effects estimator and include demographic variables and time effects. Cash-on-hand is defined as financial assets plus monthly income in euros. Standard errors in parentheses. * significance at 10%, ** significance at 5%, *** significance at 1%

Table 7. Anatomy of consumption risk

Panel A. Whole sample and risk exposure groups	Total	Age <=35	Age >60	Employed	Self-Employed	Low risk sectors	High risk sectors
	(1)	(2)	(3)	(4)	(5)	(7)	(8)
Income risk	.232	.223	.144	.246	.226	.231	.279
Health risk	.336	.346	.335	.335	.314	.344	.316
Energy risk	.161	.194	.136	.137	.21	.136	.178
Micro risks	**.728**	**.764**	**.616**	**.718**	**.75**	**.711**	**.773**
GDP risk	.062	.029	.071	.049	.099	.047	.09
House price	.079	.074	.139	.09	.07	.088	.074
Interest rate risk	.013	0	-.003	.016	.037	.013	-.017
Inflation risk	.006	-.017	.026	.012	-.018	.02	-.015
Unemployment risk	.014	.06	.019	.019	-.026	.012	.024
Macro risks	**.174**	**.146**	**.252**	**.188**	**.163**	**.18**	**.156**
Demographics and time	**.098**	**.09**	**.133**	**.094**	**.087**	**.109**	**.071**
Total	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Consumption risk	.057	.069	.035	.059	.065	.060	.063

Panel B. Self-insurance groups	Cash<10	10<Cash<30	Cash>30
	(1)	(2)	(3)
Income risk	.304	.227	.219
Health risk	.289	.399	.278
Energy risk	.202	.145	.175
Micro risks	**.795**	**.772**	**.671**
GDP risk	.021	.105	-.003
House price	.08	.043	.1
Interest rate risk	.01	.02	-.004
Inflation risk	-.019	-.008	.04
Unemployment risk	.027	.003	.019
Macro risks	**.12**	**.163**	**.152**
Demographics and time	**.084**	**.065**	**.177**
Total	1	1	1
Consumption risk	.065	.066	.040

Note. The table reports the contribution of microeconomic and aggregate risks (evaluated at the sample means) to consumption risk for the entire sample and for different groups. Column (1) in Panel A uses the regression coefficients of Table 2 column (1), while the other columns use the corresponding regression coefficients from Table 5. Panel B uses the regressions coefficients from Table 6.

Table 8. Euler equation estimates

Panel A	OLS	OLS	FE	IV FE	IV FE
	(1)	(2)	(3)	(4)	(5)
Consumption risk	1.137	2.061	1.764	1.408	1.250
	(0.088)***	(0.088)***	(0.106)***	(0.243)***	(0.232)***
Wave 2	0.002	0.002	0.002	0.002	0.002
	(0.001)***	(0.001)***	(0.001)***	(0.001)***	(0.001)***
Wave 3	0.003	0.002	0.002	0.001	0.001
	(0.001)***	(0.001)***	(0.001)**	(0.001)**	(0.001)*
Wave 4	0.001	0.001	0.001	0.000	0.000
	(0.001)*	(0.001)	(0.001)	(0.001)	(0.001)
Wave 5	0.002	0.002	0.002	0.001	0.001
	(0.001)***	(0.001)***	(0.001)**	(0.001)	(0.001)
Expected income growth		0.288	0.252	0.248	0.246
		(0.007)***	(0.009)***	(0.009)***	(0.009)***
N	23,117	23,117	23,117	23,117	23,117

Panel B	IV FE Age<=35	IV FE Age>60	IV FE Low risk	IV FE High risk
	(1)	(2)	(3)	(4)
Consumption risk	1.097	2.887	0.281	2.516
	(0.422)***	(0.572)***	(0.380)	(0.659)***
Wave 2	0.001	0.004	0.001	0.001
	(0.001)	(0.001)***	(0.001)	(0.002)
Wave 3	-0.001	0.004	-0.000	0.001
	(0.001)	(0.001)***	(0.001)	(0.002)
Wave 4	-0.001	0.003	-0.002	-0.002
	(0.001)	(0.002)**	(0.001)	(0.002)
Wave 5	0.001	0.005	-0.002	0.001
	(0.002)	(0.002)***	(0.001)	(0.002)
Expected income growth	0.285	0.245	0.246	0.300
	(0.018)***	(0.020)***	(0.015)***	(0.026)***
N	5,592	5,779	8,529	3,208

Note. The dependent variable is expected consumption growth. Consumption risk is the second conditional moment of the distribution of expected consumption growth. The upper panel presents full sample estimates: OLS estimates in columns (1) and (2), fixed effects in column (3), and IV with fixed effects in columns (4) and (5). In column (4), the instruments are the microeconomic risks, while in column (5) the instruments include both micro and macro risks. The lower panel reports fixed effects IV estimates using micro and macro risks. High-risk sectors includes the sample of respondents working in agriculture, manufacturing and construction, while low-risk sectors include services and public administration Standard errors are in parentheses. * significance at 10%, ** significance at 5%, *** significance at 1%.

Table 9. Contribution to precautionary saving

	Total sample	18-24 sample	40-45 sample	60-65 sample
Income risk	0.64	0.99	0.81	0.37
Health risk	0.76	1.13	0.85	0.53
Energy risk	0.56	0.81	0.68	0.37
Micro risks	**1.96**	**2.93**	**2.34**	**1.27**
House price risk	0.24	0.39	0.31	0.13
GDP risk	0.19	0.34	0.26	0.10
Interest rate risk	0.01	0.01	0.01	0.01
Inflation risk	0.27	0.45	0.35	0.16
Unemployment risk	0.07	0.10	0.07	0.04
Macro risks	**0.78**	**1.29**	**1.00**	**0.44**
Excess saving	**2.74**	**4.22**	**3.34**	**1.71**

Note. The contribution to precautionary saving is computed using the Euler equation estimates from column (5) of Table 8. All values are expressed as a percentages. Expected life is based on the 2023 ISTAT life tables, separately for males and females, and averaged over the 18-24 age interval. Column 1 predicts contribution using the sample of individuals 18-24 years old. Column 2 excludes individuals reporting no risk. All statistics are presented as percentages.