TI 2022-049/VI
Tinbergen Institute Discussion Paper

Consumption Choices and Earnings Expectations: Empirical Evidence and Structural Estimation

Christian Stoltenberg[1]
Arne Uhlendorff[2]

1 MInt, Department of Economics, University of Amsterdam and Tinbergen Institute
2 CREST, CNRS

CONSUMPTION CHOICES AND EARNINGS EXPECTATIONS: EMPIRICAL EVIDENCE AND STRUCTURAL ESTIMATION

Christian A. Stoltenberg*, Arne Uhlendorff†

July 19, 2022

Abstract

In this paper, we document that households' consumption expenditures depend on their expected earnings – even after controlling for realized earnings and wealth. To explain this evidence, we develop and structurally estimate a standard-incomplete markets model in which rational households possess private advance information on their future earnings. We find that households are better informed about their future earnings than an econometrician and that individual expectations are more relevant for the consumption choices of households in the left tail of the wealth distribution. Furthermore, households with advance information prefer less progressive earnings taxes.

JEL classification: E21, D31, D52

Keywords: Private information, household consumption, earnings dynamics, incomplete markets, subjective expectations.

* MInt, Department of Economics, University of Amsterdam and Tinbergen Institute, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands, email: c.a.stoltenberg@uva.nl, tel: +31 20 525 3913.

† CREST, CNRS, 5 avenue Henry Le Chatelier, 91120 Palaiseau, France, email: arne.uhlendorff@ensae.fr.
We thank Vincent Verger for providing excellent research assistance. We would like to thank Stephane Auray, Stephane Bonhomme, Jeanne Commault, Tomaso Duso, Xavier D'Haultfoeuille, David Gomtsyan, Peter Haan, Fabian Kindermann, Winfried Koeniger, Arnaud Maurel, Serdar Ozkan, Laura Pilossoph, Morten Ravn, Harald Uhlig, Gianluca Violante, Hans-Martin von Gaudecker, Pablo Winant, Johannes Wohlfart, and participants at seminars at CREST, University of Amsterdam, LISER, the Berlin Applied Micro Seminar (BAMS), and at the workshop of the population economics section of the German Economic Association for valuable comments. Arne Uhlendorff is grateful to Investissements d'Avenir (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047) for financial support.

1 Introduction

What is households' income uncertainty when they decide about their savings to insure against undesirable fluctuations of their consumption? The answer to this question is of central importance to understanding households' consumption-savings choices; only what households don't know yet constitutes uncertainty they seek to hedge. Typically, households' income uncertainty is measured as the innovation of the estimated income process. Browning, Hansen, and Heckman (1999) emphasize that such a procedure can create a disconnect between the uncertainty as assessed by an *econometrician* and income uncertainty as perceived by households. One reason why the two uncertainties can differ is that households have private advance information on their future income. There is ample evidence that this is the case for earnings.[1]

Our contribution is twofold. First, we provide direct evidence that households' consumption choices depend on their earnings expectations – even after conditioning on realized earnings and wealth. Second, we propose and estimate a consumption-savings model with private advance information to explain the empirical evidence and to quantify households' earnings uncertainty.

The consumption-savings model is a standard incomplete markets model. As the new element here, we extend households' information set by private signals that inform households about their earnings in the next period with certain precision. While the stochastic earnings process constitutes the income uncertainty as assessed by an econometrician, the joint process of signals and earnings represents households' income uncertainty. The difference between the two uncertainties depends on the precision of the signals; the more precise are the signals, the smaller are households' forecast errors for income growth and the lower is households' income uncertainty. Households in our model are rational and with informative signals they know more than an econometrician about their future earnings. The extension of households' information set is motivated by mounting evidence that finds a strong correlation between individual expectations and subsequent realizations, even when other information available to an econometrician such as households' earnings history is taken into account.[2]

Our empirical evidence stems from the Italian Household Survey of Income and Wealth

[1] For early examples see Dominitz and Manski (1997) and Dominitz (1998) who analyze survey data containing subjective expectations about future income of households in the US. Dominitz (1998) links these subjective expectations to future realizations and provides evidence for a significant relationship between expected and subsequent realized earnings while controlling for observed characteristics and current earnings.

[2] The predictive power of subjective expectations for future realizations has not only been demonstrated for earnings but also in other contexts, including the risk of job loss (Campbell, Carruth, Dickerson, and Green, 2007; Hendren, 2017; Stephens, 2004), the duration of unemployment (Mueller, Spinnewijn, and Topa, 2021), and longevity and death (Smith, Taylor, and Sloan, 2001). Manski (2018) provides an overview of this literature.

(SHIW), a representative sample containing detailed information on households' consumption, wealth and income. In some waves, this survey further elicits information on individuals' expected earnings, which makes the dataset particularly suitable for our analysis. We document that households' nondurable consumption expenditures depend positively on their subjective earnings expectations, even after controlling for realized earnings and wealth. Part of the sample is observed in several years, which allows for the estimation of panel data models. It turns out that the cross-sectional correlation between subjective expectations and consumption prevails when we control for unobserved heterogeneity in fixed effects regressions. This suggests that the correlation is not driven by time-constant unobserved characteristics such as over-confidence or permanent income that might jointly determine consumption behavior and expectations about the future.

We follow Guvenen and Smith (2014) and estimate the parameters of our model with indirect inference using a parsimonious auxiliary model. One equation of the auxiliary model captures the dynamics of earnings and the other equation consumption choices. One advantage of our approach is that we directly observe and employ households' earnings expectations in the consumption equation. Allowing for classical measurement error in consumption expenditures, we jointly estimate the key parameters of the model including signal precision, the tightness of the borrowing limit, the discount factor, and the parameters of the earnings process. The estimated structural model fits the data well and all parameters are precisely estimated. Thereby, the identification of signal precision stems from the cross-sectional correlation between consumption expenditures and individual earnings expectations.

Our main findings can be summarized as follows. First, we strongly reject the hypothesis that households do not posses advance information. Thus, households and econometricians do not share the same information set, violating a standard assumption in macroeconomic consumption-savings models. An econometrician who does not take the private information on future earnings into account, over-estimates households' true income uncertainty – as measured by households' conditional earnings variance – by approximately 19%. Second, the consumption-expectations elasticity depends on the wealth position of the household; households' in the first quintile of the wealth distribution exhibit an elasticity that is approximately three times larger than the one of households in the last quintile. Third, households smooth their consumption better than an econometrician would predict; with advance information, households' consumption exhibits an approximately 29% lower sensitivity to earnings changes. Furthermore, an econometrician tends to overestimate the importance of current earnings and to underestimate the importance of expected earnings for consumption choices.

In the next step, we ask whether these observed systematic differences in consumption-savings behavior with and without advance information are policy relevant. We find that social welfare with advance information permanently exceeds welfare in a counterfactual situation without private signals by 2.4%, measured in certainty equivalent consumption. The welfare difference has direct implications for the optimal design of progressive earnings tax schemes. Assuming that the detrimental effects of progressive taxation are independent from households' information on future earnings, our findings imply that households with advance information prefer less progressive earnings taxes.

Related literature Our paper contributes to a growing literature documenting that individual expectations matter for economic decisions.[3] Most closely related to our paper are Wiswall and Zafar (2021) and Arcidiacono, Hotz, Maurel, and Romano (2020). While Wiswall and Zafar (2021) provide evidence that students sort into college majors based on their expected future earnings, Arcidiacono et al. (2020) demonstrate that students' subjective earnings beliefs are key to understand their occupational choices.

The focus of our paper is to rationalize how subjective expectations on future realizations of individual risks affect economic decisions. Our paper complements an increasing number of field experiments documenting that economic decisions depend on expectations about the macroeconomy. For example, Roth and Wohlfart (2020) find that individuals that receive information on GDP forecasts update their expectations about macroeconomic developments and their own economic situation accordingly to eventually adjust their future consumption plans. Coibion, Gorodnichenko, and Weber (2022) provide evidence that individuals revise their inflation expectations as a response to different information treatments, and they also document an impact of these treatments on household spending.[4]

There exists a broad literature on consumption-savings decisions. Typically, macroeconomic consumption-savings models assume that households and econometricians share the same information set. Notable exceptions are Kaplan and Violante (2010), Guvenen and Smith (2014) and

[3] See Fuster and Zafar (2022) for an overview of the recent literature on expectation formation and the impact of expectations on economic behavior with a focus on studies relying on randomized information provision.

[4] It is straightforward to design information treatments and randomly allocate them to some subjects in the context of macroeconomic indicators. Such a design is more difficult - if not impossible - in the context of private information about future earnings, which is why we rely on non-experimental methods.

Pedroni, Singh, and Stoltenberg (2022). Kaplan and Violante (2010) investigate whether advance information can bridge the gap between consumption insurance as estimated in Blundell, Pistaferri, and Preston (2008) to the one stemming from a standard-incomplete market model. Guvenen and Smith (2014) estimate a life-cycle model for the US economy with households that have private advance information on their deterministic earnings profiles. In our paper, households' possess advance information about future shocks to their earnings, which drives a wedge between the perceived income uncertainties of households and econometricians. Using only panel data on consumption and income – but no data on individual income expectations – Pedroni et al. (2022) provide evidence that US households possess advance information on their future income.

Part of the literature on consumption-savings decisions is based on the SHIW dataset.[5] Some of these papers explore the information on earnings expectations. For example, Guiso, Jappelli, and Terlizzese (1992) provide evidence that earnings uncertainty is correlated with precautionary saving, and Guiso, Jappelli, and Terlizzese (1996) find that perceived earnings uncertainty impacts households' decisions on the share of risky assets in their portfolio. Guiso, Jappelli, and Pistaferri (2002) document how individual earnings and employment risk change over the life-cycle and how subjective risk attitudes can impact occupational choices. Kaufmann and Pistaferri (2009) use the data to distinguish anticipated and unanticipated income changes and to estimate consumption insurance. In the same spirit as these authors, we provide direct evidence that subjective earnings expectations matter for households consumption-savings decisions. Relative to the aforementioned papers, we additionally use the direct evidence to structurally estimate a consumption-savings model, which allows us to study the welfare implications of advance information.

Finally, our work is related to Guvenen, Karahan, Ozkan, and Song (2021), Arellano, Blundell, and Bonhomme (2017) and Arellano, Blundell, Bonhomme, and Light (2021) who explore non-linear earnings dynamics. Guvenen et al. (2021) document that the distribution of earnings changes shows non-Gaussian features such as negative skewness. Arellano et al. (2017) argue that earnings shocks are characterized by non-linearities in persistence. Arellano et al. (2021) use the framework provided in Arellano et al. (2017) to study heterogeneous consumption re-

[5] A recent example is Auclert (2019), who estimates cross-sectional statistics on the marginal propensity to consume of Italian households to understand the role of redistribution in the transmission of monetary policy. Fella, Frache, and Koeniger (2020) fit a buffer-stock savings model and conclude that Italian households mainly rely on self-insurance and do not have significant access to other forms of insurance.

sponses to earnings shocks with non-linear persistence. While our earnings process is more conventional, employing expectations data allows us to uncover the heterogeneity of consumption responses to earnings shocks resulting from households with advance information. Overall, our paper complements the work of the aforementioned authors.

This paper proceeds as follows. Section 2 introduces the dataset, details the sample selection criteria and presents some descriptive statistics. Section 3 provides reduced form evidence on the relevance of expected earnings for consumption decisions. In Section 4, we introduce the consumption-savings model with private signals. Section 5 outlines the estimation method. In Section 6, we discuss identification and present the estimation results. Section 7 explores how much advance information Italian households possess and whether the relevance of expectations for consumption choices is heterogenous across the wealth and earnings distribution. The policy implications of households with advance information are studied in Section 8. The last section concludes.

2 Data, Sample and Descriptives

Our analysis is based on the SHIW, a representative sample of the Italian population. The SHIW is a panel data set with household members usually being interviewed every 2 years. The data provide information at the individual level including information about employment, family status, education, earnings and age. Moreover, the SHIW contains detailed information on wealth and consumption at the household level, which is provided by the head of household.[6] In some waves, the interviewees are asked to report their expected annual earnings in the next 12 months. In 1989 and 1991, they are asked about their expected earnings changes. For this, the respondents assign probabilities to 12 different intervals for the expected earnings change. From the surveys in 1995 and 1998, we have individual responses on the expected minimum annual earnings (y_{min}), the expected maximum earnings (y_{max}), and the probability that expected earnings y are lower or equal to the midpoint y_{mid}: $\pi : y \leq (y_{min} + y_{max})/2$. Based on these information, we construct the expected annual earnings (see Appendix A.1 for details).[7]

Our main estimation sample consists of employed household heads for whom we observe

[6] The head of household is defined as the person who is primarily responsible for the household budget.

[7] For 1989 and 1991, the question about earnings expectations is asked unconditional on the future employment status. In 1995 and 1998, the respondents are first asked about their employment probability, and after that they are asked about the expected earnings conditional on being employed. We will later show that our results on the relationship between expected earnings and consumption behavior hold if we estimate models only using the waves 1989/91 or those of 1995/98.

Figure 1: Distribution of earnings and expected earnings

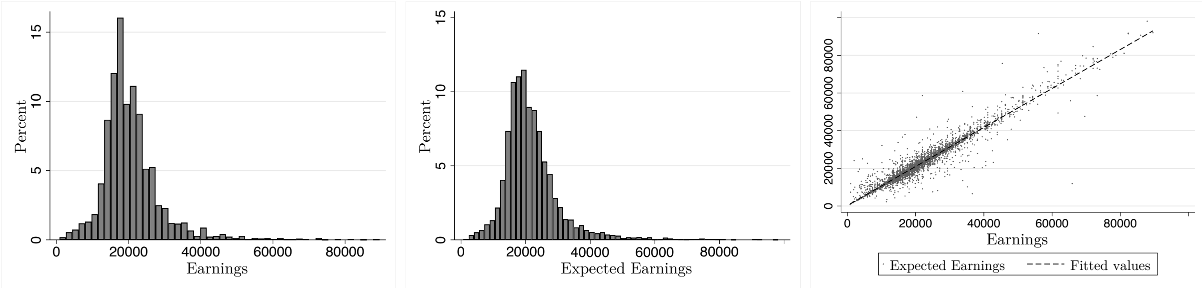

Notes: The left panel displays the distribution of annual earnings, the middle panel the distribution of expected annual earnings, excluding observations above 100,000€. The right panel displays the joint distribution (n=7,653). In Appendix A.2, we report the earnings distributions separately for the periods before and after 1993 (Figures A.1 and A.2). Real terms in Euro of the year 2010.

their expected earnings. We exclude self-employed workers and individuals who are younger than 18 and older than 65. Our final sample consists of 7,659 observations for 6,501 individuals. The vast majority of household heads are men, and around 83% are married. Table A.1 presents descriptive statistics. The average annual earnings amount to 20,880€, expressed in real terms with 2010 as the base year. The interviews of the SHIW usually take place in May, with the respondents reporting their income and consumption expenditures of the previous calendar year. At this moment, individuals expect to earn on average 21,940€ in the next 12 months. Both earnings measures refer to after-tax earnings. We observe similar distributions of the expected and realized earnings variables (see Figure 1). The joint distribution of the expected and the current earnings shows a strong positive relationship between these two variables. The average annual consumption per household member amounts to 12,915€ per year, and the households have on average a net wealth of 66,267€ per household member at the beginning of the previous calendar year.[8]

For part of our sample, we can link the expected earnings with actual earnings in the future. This allows us to investigate whether subjective expectations are informative about future earnings. For this, we regress the log of future earnings for year $t + 2$ on the log of earnings in year t and the log of subjective expectations about future earnings observed at the moment of the interview in year $t + 1$. Figure 2 provides an overview of the data structure based on the example of the surveys in 1998 and 2000.[9] In a model in which we control for

[8] The net wealth corresponds to the sum of financial and real assets, subtracting the financial liabilities. We observe the net wealth at the end of the previous calendar year and the amount of savings for the whole year. To calculate the net wealth at the beginning of the calendar year, we subtract the savings from the net wealth observed at the end of the year. The distributions of the net wealth and the annual consumption of non-durable goods are reported in Appendix A, Figure A.3.

[9] While we usually observe actual earnings every two years, the gap between the surveys 1995 and 1998 is

Figure 2: Timing of realizations and expectations for the 1998 and 2000 surveys

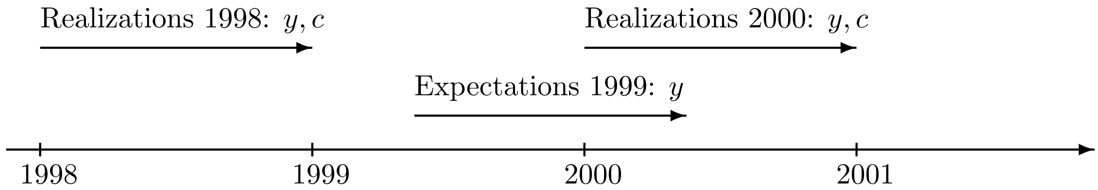

basic socio-demographic characteristics, we find a strong positive relationship between the log of expected earnings and the log of actual earnings in the future (see Table A.2 in the Appendix). The estimated coefficient drops from 0.646 to around 0.401 when we control additionally for annual earnings in t, but it is still significantly positive at the 1% level. These results suggest that the individuals in our sample have private information about their future earnings.

While our results are in line with the findings in the literature on subjective expectations about future economic outcomes, the timing of the different measurements of our main variables is not perfectly designed for analyzing this type of questions. On the one hand, there is a time lag between the realized earnings in $t-1$ and the date of the interview. Part of the subjective expectations about future earnings might therefore reflect changes in the earnings between the end of the calendar year and the time of the interview. On the other hand, our realizations are only partially overlapping with the expectations. While this calls for a cautious interpretation of the findings, we are confident that the evidence suggests that individuals possess some private information which goes beyond their current earnings.

3 Expected Earnings and Consumption: Empirical evidence

We are interested in the role of expected earnings for current consumption behavior. We observe the consumption of non-durable goods as well as the annual earnings for the calendar year t and the expectations about the earnings in the next 12 months at the moment of the interview in year $t+1$. Figure 3 displays the bivariate distributions of the log of non-durable consumption and the log of annual earnings (left panel) and the log of expected annual earnings (right panel). We observe a positive relationship between actual and expected earnings with consumption expenditures.

In a next step, we estimate linear regressions controlling for observed characteristics includ-

three years.

Figure 3: Log (expected) earnings and log consumption

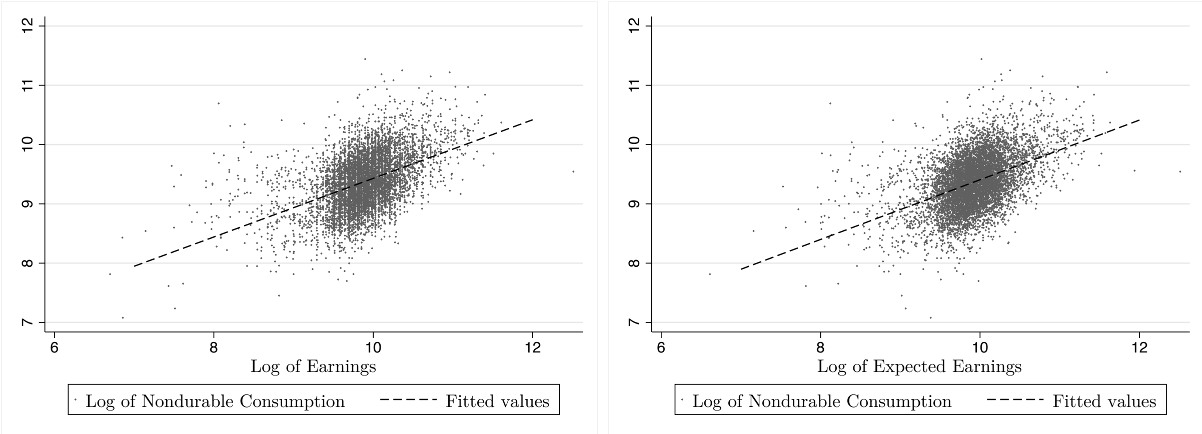

Notes: The left panel displays the joint distribution of the log of annual nondurable consumption and the log of annual earnings, the right panel the joint distribution of the log of annual nondurable consumption and the log of expected annual earnings.

ing net wealth measured at the beginning of period t. The estimations are based on pooled cross sections controlling for year and regional fixed effects. We estimate the following equation:

$$\log(c_{it}) = \varepsilon_{c,y} \log(y_{it}) + \varepsilon_{c,\mathrm{E}\,y} \log(\mathrm{E}_{it+1}\, y_{it+2}) + \gamma_1 X_{1it} + \eta_{1,it}, \tag{1}$$

The corresponding OLS results are reported in Table 1. The coefficients $\varepsilon_{c,y}$ and $\varepsilon_{c,\mathrm{E}\,y}$ can be interpreted in terms of estimated elasticities. In a regression without controlling for current earnings, we estimate a coefficient of 0.321 for the consumption elasticity with respect to expected earnings measured in $t+1$. Once we include earnings in period t, we estimate a consumption elasticity of 0.133 for the expected earnings and of 0.201 for current earnings. Both coefficients are statistically significant at the 1% level. This evidence clearly suggests that households adjust their consumption choices to their expectations about their future earnings.

The evidence also indicates advance information as a potential mechanism to explain the correlation between consumption expenditures and earnings expectations. Households' expectations are conditional on information in $t+1$ which includes their future earnings in period $t+1$. If households have advance information on their future earnings, their consumption in t reacts to their earnings in $t+1$, explaining the correlation between consumption and expectations.[10]

These results are robust to the selection criteria for our sample, to controlling more flexibly for earnings and to the specification used for calculating the expected earnings. For part of our

[10] This argument resembles the logic of the tests to detect advance information in Cunha et al. (2005), Blundell et al. (2008) and Pedroni et al. (2022). Using data from the Panel Study of Income Dynamics (PSID), Pedroni et al. (2022) provide evidence that US households possess advance information on their future income.

Table 1: Consumption elasticities: pooled cross sections

	(1)	(2)
Log of expected earnings	0.321***	0.133***
	(0.0142)	(0.0438)
Log of earnings	-	0.201***
	-	(0.0457)

Notes: Standard errors are clustered at the individual level and reported in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. n = 7,659. We control for net wealth, age, gender, education, marital status, year and region. Earnings and consumption are measured in real terms in Euro of the year 2010. The full estimation results are reported in Table A.3 in the Appendix A.4.

observation period, we observe individuals' inflation expectations. It turns out that adjusting expected earnings using individual inflation expectations does not affect our results. We present the different sensitivity exercises in Appendix A.5.

One concern might be that unobserved time-constant characteristics which are correlated with subjective expectations and consumption behaviour could explain these findings. For example, over-confidence could be positively correlated with expectations about future earnings and consumption expenditures. For evaluating this, we estimate the following unobserved effects models:

$$\log(c_{it}) = \varepsilon_{c,y} \log(y_{it}) + \varepsilon_{c,\mathrm{E}\,y} \log(\mathrm{E}_{it+1}\, y_{it+2}) + \gamma_1 X_{1it} + \alpha_i + \epsilon_{1,it},$$

In this model, α_i captures the unobserved time-constant individual characteristics. The fixed effects estimations are based on 930 individuals having more than one observation in our data set. In contrast to the model specified in equation (1), the identification of the elasticities is solely based on the variation of (expected) earnings and consumption expenditures within individuals over time. The results are reported in Table 2. It turns out that they are very similar to the OLS results. While the point estimates for the elasticities are slightly larger (0.161 for the expected earnings and 0.228 for the actual earnings), they are within the confidence intervals of the estimates reported in Table 1. This suggests that our results are not driven by unobserved time-constant characteristics like over-confidence, permanent income, cognitive ability or time

Table 2: Elasticity of consumption: fixed effects estimations

	(1)	(2)
Log of expected earnings	0.349***	0.161***
	(0.0343)	(0.0554)
Log of earnings	-	0.228***
	-	(0.0528)

Notes: Standard errors in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. n = 2,087. We control for net wealth age, gender, education, marital status, year and region. Earnings and consumption are measured in real terms in Euro of the year 2010. The full estimation results are reported in Table A.4 in the Appendix A.4.

preferences.

While we observe the consumption of non-durable goods and the annual earnings for the calendar year t, the expectations about future earnings are measured at the interview date in year $t+1$ (see Figure 2). The expectations might have been updated between t and $t+1$, and the expectations in year $t+1$ should be only a proxy for or a function of the expectations in t. Therefore, we tend to underestimate the true underlying elasticity of consumption with respect to earnings expectations, and our results are a conservative estimate of this relationship.

We additionally investigate to what extent households react differently if they receive rather positive or rather negative signals about their future earnings. For this, we allow for heterogeneous consumption elasticities depending on the expected change in earnings. We split the sample (i) at the median expected change and (ii) at the terciles of the expected change and allow the expectation elasticities to depend on the position in the distribution of expected changes. For both regressions, we do not find any evidence for heterogenous responses (see Tables A.10 and A.11 in the Appendix). This finding does not support the idea that households might react differently to positive and negative signals.

4 A structural consumption-savings model

In this section, we describe an economy in which risk-averse households face uninsurable idiosyncratic income shocks as in Aiyagari (1994). As the main novel feature, households receive private signals that inform on their endowment shock realization in the next period.

Preferences and endowments Consider an economy with a continuum of households indexed by i. Time is discrete and indexed by t from zero onward. Households have preferences over consumption streams and evaluate them conditional on the information available at $t = 0$

$$U\left(\{c_t^i\}_{t=0}^{\infty}\right) = (1-\beta)\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t^i), \tag{2}$$

where the instantaneous utility function $u : \mathbb{R}_+ \to \mathbb{R}$ is strictly increasing, strictly concave and satisfies the Inada conditions, and $0 < \beta < 1$ is the discount factor.

Household i's disposable labor income in period t is given by $w_t y_{it}$, where w_t is the real wage per unit of effective labor and y_{it} are individual effective labor unit endowments. Effective labor unit endowments are generated by a stochastic process $\{y_{it}\}_{t=0}^{\infty}$, where the set of possible realizations in each period is time-invariant and finite $y_{it} \in Y \equiv \{y_1, ..., y_N\} \subseteq \mathbb{R}_{++}$, ordered. The history $(y_0, ...y_t)$ is denoted by y^t. Effective labor units are independent across households and evolve across time according to a first-order Markov chain with time-invariant transition matrix P whose elements $\pi(y' = y_k|y = y_j)$ for all j, k are the conditional probabilities of next period's endowment y_k given current period endowment y_j. There is no aggregate risk, and the Markov chain induces a unique invariant distribution of endowments $\pi(y)$ such that the aggregate labor endowment is constant and equal to $L_t = \bar{y} = \sum_y y\pi(y)$. In the following, all relevant transition probabilities are time-invariant, which is why we employ a recursive notation such that $x(x')$ denotes the value of a generic variable x in the current (future) period.

Information Our information structure resembles the one developed in Stoltenberg and Singh (2020). Except for observing past and current endowment shocks, household i receives in each period $t \geq 0$ a private signal $k_{it} \in Y$ that informs about endowment realizations in the next period. The signal has as many realizations as endowment states and its precision κ is captured by the time-invariant conditional probability that signal and future endowment coincide, $\kappa = \pi(y' = y_j|k = y_j)$, $\kappa \in [1/N, 1]$. Uninformative signals are characterized by precision $\kappa = 1/N$, perfectly informative signals by $\kappa = 1$. Hence, at each point in time the agents can find themselves in one of the states $s_t = (y_t, k_t)$, $s_t \in S$, where S is the Cartesian product $Y \times Y$ and $s^t = (y^t, k^t) = (s_0, ...s_t)$ is the history of the state.

Using the recursive notation, the conditional probabilities of future endowments y' conditional on today's state $s = (y, k)$ are denoted by $\pi(y'|s)$. The latter probabilities are given by

$$\pi\left(y' = y_j | k = y_m, y = y_i\right) = \frac{\pi_{ij}\kappa^{\mathbf{1}_{j=m}}\left(\frac{1-\kappa}{N-1}\right)^{1-\mathbf{1}_{j=m}}}{\sum_{z=1}^{N}\pi_{iz}\kappa^{\mathbf{1}_{z=m}}\left(\frac{1-\kappa}{N-1}\right)^{1-\mathbf{1}_{z=m}}} \quad , \tag{3}$$

where tomorrow's endowment is $y' = y_j$, today's endowment is $y = y_i$ and today's signal indicates endowment state y_m in the future, $k = y_m$; $\mathbf{1}_{j=m}$ is an indicator function that equals one if the signal and the actual realization of the endowment coincide. The formula resembles a "hit-or-miss" specification and its logic follows from Bayes' theorem. There are two independent "signals" on future endowment realizations, current endowments and the private signal. Both signals are weighted with their precision, endowments with transition probability π_{jk} and signals with precision κ. Intuitively, the signal informs about future endowment shock realizations by implicitly providing advance information on future innovations to endowments.

For example, with uninformative signals ($\kappa = 1/N$) the conditional probability of endowment y_j tomorrow given today's endowment y_i and given any signal k today can be computed as

$$\pi\left(y' = y_j | k, y = y_i\right) = \frac{\pi_{ij}\frac{1}{N}}{\frac{1}{N}\sum_{z=1}^{N}\pi_{iz}} = \pi_{ij}.$$

To derive the transition probabilities of the state $\pi(s'|s)$, we assume that signals follow an exogenous first-order Markov process with time-invariant transition probabilities $\pi(k'|k)$. Combining this assumption with (3) yields a time-invariant Markov transition matrix P_s with conditional probabilities $\pi(s'|s)$ as elements

$$\begin{aligned}\pi(s'|s) &= \pi\left(y' = y_j, k' = y_l | k = y_m, y = y_i\right) \\ &= \pi(k' = y_l | k = y_m)\pi\left(y' = y_j | k = y_m, y = y_i\right). \end{aligned} \tag{4}$$

The Markov chain induces a unique invariant distribution of the state denoted by $\pi(s)$. The formula (4) applies to any first-order Markov process for signals. To discipline the Markov signal process, we apply a "reverse-engineering" procedure to choose the Markov process for signals such that households' perceived transitions of the endowment shock y equal the actual transitions of the shock (household rationality). Appendix B.1 provides details on the derivation of the formulas for the joint distribution of endowments and signals. Appendix B.2 describes the reverse-engineering procedure.

Production A representative firm hires labor L_t and capital K_t at rental rates w_t and r_t to maximize profits. Capital depreciates at rate δ and the production of consumption goods Y_t takes place via a linear homogenous production function

$$Y_t = AF(L_t, K_t),$$

with A as a productivity parameter that is constant in the stationary equilibria that we focus on in the following. Aggregate labor endowments L_t are normalized to unity.

Household problem Households can only trade in a single non-state contingent bond with gross return R and face an exogenous borrowing limit $\underline{a}$. We focus directly on stationary allocations. Given asset holdings a, state $s = (y, k)$, and an interest rate R, households' problem can be written recursively as

$$V(a,s) \quad = \quad \max_{c,a'} \left[(1-\beta)u(c) + \beta \sum_{s'} \pi(s'|s) V(a', s') \right]$$

subject to a budget and a borrowing constraint

$$\begin{aligned} c + a' \quad &\leq \quad wy + Ra \\ a' \quad &\geq \quad \underline{a}. \end{aligned}$$

Here, households differ with respect to initial asset holdings and initial shocks where the heterogeneity is captured by the probability measure $\Psi_{a,s}$. The state space is given by $M = A \times S$, where $A = [\underline{a}, \infty)$.

The stationary recursive competitive equilibrium is summarized in the following definition.

Definition 1 *A stationary recursive competitive equilibrium in the standard incomplete markets economy comprises a value function $V(a,s)$, prices R, w, an allocation $c(a,s), a'(a,s), K$ a joint probability measure of assets and the state $\Psi_{a,s}$, and an exogenous borrowing limit $\underline{a}$ such that* *(i)* $V(a,s)$ *is attained by the decision rules* $c(a,s)$*,* $a'(a,s)$ *given* R

(ii) *The joint distribution of assets and state* $\Psi_{a,s}$ *induced by* $a'(a,s)$ *and* P_s *is stationary.*

(iii) *Factor prices satisfy*

$$R - 1 = AF_K(1, K) - \delta$$

$$w = AF_L(1, K)$$

(iv) *The bond market clears*

$$\int a'(a,s)\,\mathrm{d}\,\Psi_{a,s} = K.$$

5 Quantitative exercise

In this section, we provide functional forms for the earnings process, households' preferences and firms' production technology. Afterwards, we describe how we choose the structural parameters. One part of the parameter vector is preset to match salient features of the Italian economy post 1980. The other part of the parameter vector – including the precision of private signals – is structurally estimated by indirect inference.

5.1 Earnings process, preferences and technology

Typically, the log of household income is modelled as the sum of persistent and orthogonal transitory shocks, that is, there are two innovation terms. With just one signal but two innovation terms, it remains unclear on which future innovation the signal is informative. For this reason, we employ the results provided in Ejrnæs and Browning (2014) to model log earnings of household i as an ARMA(1,1)-process with a single innovation term that is equivalent to a persistent–transitory specification for $\theta \leq \rho$

$$\log(y_{it}) = \rho \log(y_{it-1}) - \theta u_{it-1} + u_{it}, \quad u_{it} \overset{iid}{\sim} \mathcal{N}(0, \sigma_u^2), \tag{5}$$

with unconditional variance

$$\mathrm{var}\left[\log(y_{it})\right] = \frac{1 + \theta^2 - 2\rho\theta}{1 - \rho^2} \sigma_u^2. \tag{6}$$

Table 3: Calibrated and preset parameters

	Parameter	Value
R	Gross interest rate	1.0414
α	Production elasticity of capital	0.3061
δ	Depreciation of capital	0.0870
A	Technology parameter, production function	0.9876
σ	Relative risk aversion	4
ρ	Autoregressive coefficient, ARMA(1,1)	0.9989

Further, we consider CRRA-preferences with relative risk aversion parameter σ and a Cobb-Douglas production function with elasticity of capital α such that the vector of structural parameters is $(R, \alpha, \delta, A, \sigma, \rho, \kappa, \underline{a}, \theta, \sigma_u, \beta)$.

5.2 Calibrated and preset parameters

Our calibration is designed to capture some salient features of the post-1980 Italian economy on an annual frequency. The return of the risk-free asset is $R = 1.0414$ which is the average risk-free return post 1980 found in Jordà, Knoll, Kuvshinov, Schularick, and Taylor (2019). The capital-production elasticity is set to $\alpha = 0.3061$ as estimated in Piketty and Zucman (2014). Given R and α, we choose the depreciation of the capital stock δ and the technology parameter A to yield a real wage rate of unity and a wealth-to-income ratio (based on capital demand) of 3.0091, estimated by Caprara, de Bonis, and Infante (2018) based on OECD data. With a wage rate of unity, labor income is $wy = y$, and we use the terms individual endowment, labor income and earnings interchangeably. Further, we choose a standard value for relative risk aversion, $\sigma = 4$.[11]

For the earnings process, we choose $\rho = 0.9989$, which is the value originally estimated in Storesletten, Telmer, and Yaron (2004). Given a particular value of the persistence parameter, we approximate the ARMA(1,1) as a finite-state first-order Markov process with six distinct states. We normalize the value of all earnings states such that mean earnings (or aggregate labor endowment) is equal to unity. For each of the six earnings states, there are therefore six signals such that the joint earnings-signals state S is approximated by 36 states. Table 3 summarizes the calibrated and preset parameter values.

[11] In Table B.2 in the Appendix, we provide estimation results for alternative degrees of risk aversion.

5.3 Parameters estimated with indirect inference

We use indirect inference to estimate the parameter vector $\Theta = (\kappa, \underline{a}, \theta, \sigma_u, \beta, \sigma_c)$, with σ_c as the standard deviation of a classical measurement error in consumption expenditures. The parameters θ, σ_u govern the dynamics of the ARMA(1,1) earnings process. For a given earnings process, the precision of private signals κ, the borrowing limit $\underline{a}$ and the discount factor β are the key parameters for households' consumption-savings decisions, which eventually shape the wealth distribution.

The economic model is given by the structural consumption-savings model described in Section 4. As Guvenen and Smith (2014), we employ a partial equilibrium version of the structural model in the estimation and do not impose the clearing of the bond market – as formalized in Definition 1 part (iv).[12]

Indirect inference relies on a parsimonious auxiliary model to connect the consumption-savings decisions of the households in the data to the decisions in the economic model. The methodology is similar to Guvenen and Smith (2014). One important difference is that we additionally observe and employ data on households' beliefs in the auxiliary model. The auxiliary model comprises a consumption equation, and an additional equation describing earnings dynamics. The first equation captures how current consumption of household i depends on current earnings and expected earnings one period ahead:

$$\log(c_{it}) = \varepsilon_{c,y} \log(y_{it}) + \varepsilon_{c,\mathrm{E}\,y} \log(\mathrm{E}_{it+1}\, y_{it+2}) + \gamma_1 X_{cit} + \eta_{c,it}, \tag{7}$$

where $\varepsilon_{c,y}, \varepsilon_{c,\mathrm{E}\,y}$ are the elasticities of consumption with respect to current earnings and future expected earnings, X_{cit} is the vector of control variables including a constant, with residual $\eta_{c,it} \sim (0, \sigma^2_{\eta,c})$.[13] For the consumption-savings model, beginning-of-period net wealth a_{it-1} is included in X_{cit}. In the data, $\mathrm{E}_{it+1}\, y_{it+2}$ is the mean of the expected-earnings distribution of a particular household constructed from her survey response. In the model, the expectation operator is given by the conditional expectation $\mathrm{E}_{it+1}\, y_{it+2} = \mathrm{E}_{t+1}(y_{it+2}|y_{it+1}, k_{it+1})$.

The second equation represents the earnings-dynamics for the two-year frequency available in the Italian households data:

[12] In Section 8, we study the policy relevance of advance information in a general equilibrium version of the model. There, we adjust the interest rate R to satisfy bond-market clearing.

[13] The control variables in case of the survey data are described in Section 2.

$$\log(y_{it}) = a_1 \log(y_{it-2}) + \gamma_2 X_{2,it} + \eta_{y,it}, \tag{8}$$

with $\eta_{y,it} \sim (0, \sigma^2_{\eta,y})$ and control vector $X_{2,it}$, which is empty in case of the consumption-savings model. In the SHIW, we estimate $a_1 = 0.6359$ with a standard error of 0.0136 and $\sigma_{\eta,y} = 0.3229$.[14]

With indirect inference, the parameter vector Θ is chosen to minimize the relative distance between the coefficients of the auxiliary model in Equations (7)-(8) in the data and in the model. The coefficients in the auxiliary model include the two consumption elasticities in the survey data, $\widehat{\varepsilon}_c = (\widehat{\varepsilon}_{c,y}, \widehat{\varepsilon}_{c,\mathrm{E}\,y})$, and in the structural model, $\varepsilon_c(\Theta)$. Further, the auxiliary model comprises the auto-regression coefficient $\hat{a}_1, a_1(\Theta)$ and the two residual standard deviations, $\hat{\sigma}_\eta = (\hat{\sigma}_{\eta,c}, \hat{\sigma}_{\eta,y}), \sigma_\eta$. To capture consumption insurance, we further target the consumption response, $\log(c_{it}) - \log(c_{it-2})$, to unexpected earnings changes, $\log(y_{it}) - \log(\mathrm{E}_{it-1}\, y_{it})$, by the corresponding regression coefficient, $\widehat{\beta}_{INS} = 0.1982$.[15] The wealth distribution is taken into account by targeting the median, $\widehat{\mathrm{Med}}_{WIR} = 2.6503$, and the 10th percentile of the wealth-to-income ratio distribution in the survey data, $\widehat{10\mathrm{th}}_{WIR} = 0.0132$. The vector of percentage deviations is

$$err(\Theta) = \left[\begin{pmatrix} \varepsilon_c(\Theta)' \\ a_1(\Theta) \\ \sigma_\eta(\Theta) \\ \beta_{INS}(\Theta) \\ \mathrm{Med}_{WIR}(\Theta) \\ 10\mathrm{th}_{WIR}(\Theta) \end{pmatrix} - \begin{pmatrix} \widehat{\varepsilon}'_c \\ \widehat{a}_1 \\ \widehat{\sigma}'_\eta \\ \widehat{\beta}_{INS} \\ \widehat{\mathrm{Med}}_{WIR} \\ \widehat{10\mathrm{th}}_{WIR} \end{pmatrix} \right] ./ \begin{pmatrix} \hat{\varepsilon}'_c \\ \widehat{a}_1 \\ \widehat{\sigma}'_\eta \\ \widehat{\beta}_{INS} \\ \widehat{\mathrm{Med}}_{WIR} \\ \widehat{10\mathrm{th}}_{WIR}(\end{pmatrix} = [m(\Theta) - \widehat{m}] ./ \widehat{m},$$

with ./ as point-wise division. In total, there are eight targets for six structural parameters in Θ. For estimation, we specify the following objective function

$$\Theta_{\min} = \underset{\Theta}{\arg\min} \;\; f(\Theta) = err(\Theta)' * \mathbf{I}_8 * err(\Theta), \tag{9}$$

[14] We control for observed characteristics including age, gender, marital status, education and region. For our sample period 1989–2000, we end up with 4,771 observations. We do not use pairs of earnings observed in 1995 and 1998, because here the gap between two observations is three years.

[15] Using SHIW data from 1989–1993, Duso (1999) reports point estimates of β_{INS} between 0.207 to 0.236.

for $1/N \leq \kappa \leq 1$, $\sigma \geq 0$, $0 \leq \theta \leq \rho$, $\sigma_u \geq 0$ and $0 < \beta < 1$, and $\mathbf{I}_8$ denotes an eight-dimensional identity matrix. The minimization problem is highly non-linear in the structural parameters. In Appendix B.3, we describe our estimation procedure.

6 Estimation results

In this section, we present our estimation results. We begin with a discussion on identification. Afterwards, we present the parameter estimates and continue with inspecting the auxiliary model to analyze the fit of the economic model. Eventually, we use the economic model to answer several questions that cannot be addressed with the survey data. First, we quantify how much advance information Italian households possess on their future earnings. Second, we ask how households' consumption decisions are affected by current earnings expectations instead of earnings expectations formed in the future. Finally, we study the heterogeneity of consumption-expectations responses across the wealth and earnings distribution.

6.1 Identification

In the economic model, all elements of Θ impact the three auxiliary parameters in Equation (7) and the wealth-to-income ratio distribution. The coefficients $a_1, \sigma_{\eta,y}$ that describe the earnings dynamics in the auxiliary model equation (8) are only affected by θ, σ_u. Thus, the identification of θ, σ_u stems from these two parameters of the auxiliary model. Thereby, the auxiliary model parameter a_1 is exclusively linked to θ according to the following quadratic equation

$$\begin{aligned}
a_1(\theta) &= \frac{\operatorname{cov}\left[\log(y_{it}), \log(y_{it-2})\right]}{\operatorname{var}\left[\log(y_{it-2})\right]} \\
&= \frac{\operatorname{cov}\left[\rho^2 \log(y_{it-2}) - \theta\rho u_{it-2}, \log(y_{it-2})\right]}{\operatorname{var}\left[\log(y_{it-2})\right]} \\
&= \frac{\rho^2 \operatorname{var}\left[\log(y_{it-2})\right] - \rho\theta\sigma_u^2}{\operatorname{var}\left[\log(y_{it-2})\right]} \\
&= \frac{\rho^2(1+\theta^2-\rho\theta) - \rho\theta}{1+\theta^2-2\rho\theta},
\end{aligned} \tag{10}$$

where the first line uses the definition of the ARMA(1,1), the second one uses that $\operatorname{cov}\left[\log(y_{it}), u_{it}\right] = \sigma_u^2$ (and weak stationarity), and the last line the variance formula of an ARMA(1,1) given in (6). There are two solutions to this equation, but for $0 < a_1 < 1$, only one satisfies $\theta \leq \rho$, the necessary condition that allows to represent an ARMA(1,1) as a persistent-transitory earnings specification. Thus, with this additional restriction, $\hat{a}_1$ uniquely identifies θ. With θ identified by $\hat{a}_1$, $\hat{\sigma}_{\eta,y}$ and $\hat{a}_1$ identify σ_u via the unconditional variance of log earnings

$$\begin{aligned} \text{var}\left[\log(y_{it-2})\right] &= \frac{\hat{\sigma}^2_{\eta,y}}{1-\hat{a}_1{}^2} = \frac{1+\theta^2-2\rho\theta}{1-\rho^2}\sigma_u^2 \\ \Leftrightarrow \sigma_u^2 &= \frac{\hat{\sigma}^2_{\eta,y}}{1-\hat{a}_1^2}\frac{1-\rho^2}{1+\theta^2-2\rho\theta}. \end{aligned} \tag{11}$$

The parameters $\kappa, \underline{a}, \beta$ jointly impact the consumption elasticities, the residual standard deviation in Equation (7) as well as the moments of the wealth-to-income distribution in the simulated structural consumption-savings model. For example, the consumption elasticity with respect to future expected earnings, $\varepsilon_{c,\mathrm{E}\,y}$, depends not only on the informativeness of signals, but also on patience and the severeness of the borrowing limit. When households are very patient, their current consumption only weakly responds to future earnings expectations, even when these expectations are informative. Alternatively, for households that are at their borrowing limit and live hand-to mouth, only current earnings but not expected earnings matter, independent from the precision of signals; the wealth-to-income ratio is subject to similar arguments. Thus, each of the three parameters can not be separately identified with one particular auxiliary model parameter. Instead, these parameters can only be jointly identified from the consumption elasticities and the moments of the wealth-to-income ratio distribution. Given $\kappa, \underline{a}, \beta$, the parameter σ_c is identified by the remainder of the residual standard deviation $\sigma_{\eta,c}$.

6.2 Structural parameters and model-data comparison

In Table 4, we provide the parameter estimates and report information on the model fit.[16]

Parameters of the ARMA(1,1)-earnings process: θ, σ_u. Given $\rho = 0.999$, we estimate that log earnings depend on past innovations with a coefficient of $\theta = 0.854$. The innovations are found to have a standard deviations of $\sigma_u = 0.127$.[17] Both parameters are precisely estimated.

Parameters identified from the consumption elasticities, consumption insurance and wealth distribution: $\kappa, \underline{a}, \beta$. We estimate informative signals with a precision of $\kappa = 0.569$ and a standard error of 0.060. Uninformative signals exhibit $\kappa = 1/6$. Taking this into account, signal precision has a t-value of $(0.569 - 1/6)/0.060$, which equals approximately 7, safely

[16] All parameter estimates are robust to variations in the preset value for relative risk aversion (see Table B.2 in Appendix B.4 for the details).

[17] In Appendix B.3, we provide details on the computation of standard errors.

rejecting the hypothesis that households do not possess advance information on their future earnings. Further, we estimate $\underline{a} = -0.420$, with standard error of 0.10. Average net earnings are normalized to one, which implies that households can borrow up to 42% of it every year. The estimated discount factor is $\beta = 0.925$, a standard value for an annual frequency, precisely estimated.

The identification of $\kappa, \underline{a}, \beta$ is further illustrated in Figure 4 that displays absolute deviations of the auxiliary model parameters in the model compared to their counterpart in the survey data along two dimensions, conditional on the remaining structural parameters being at Θ_{min}. The parameter regions with the lowest deviation from target is marked in dark blue and the stars mark the minimizer. Moving from left to right, the consumption elasticity with respect to expected earnings informs about κ while the consumption elasticity with respect to current earnings pins down the discount factor β. The 10th percentile of the wealth-to-income ratio distribution is informative about the borrowing limit $\underline{a}$, but not about signal precision.

Table 4: Parameter estimates and auxiliary model

I. Parameter estimates, Θ_{min}

κ	$\underline{a}$	θ	σ_u	β	σ_c
0.569 (0.060)	−0.420 (0.096)	0.854 (0.011)	0.127 (0.009)	0.925 (0.003)	0.299 (0.003)

II. Auxiliary model

	$\varepsilon_{c,y}$	$\varepsilon_{c,\mathrm{E}y}$	$\sigma_{\eta,c}$	a_1	$\sigma_{\eta,y}$	β_{INS}	Med$_{WIR}$	10th$_{WIR}$
Data	0.201	0.133	0.313	0.636	0.323	0.198	2.650	0.013
Model	0.203	0.133	0.313	0.634	0.319	0.207	2.664	0.013

Notes: asymptotic standard errors (SEs) computed from a parametric bootstrap with 200 repetitions in parentheses.

Measurement error of consumption: σ_c**.** The auxiliary model captures the household consumption stemming from the consumption-savings model very well. As a consequence, the residual standard deviation from Equation (7) is too low compared to the data. The void is filled by a consumption measurement error with $\sigma_c = 0.299$, also precisely estimated. For comparison, Guvenen and Smith (2014) find a measurement error of $\sigma_c = 0.36$ using US data.

Model-data comparison: inspecting the auxiliary model Despite over-identification, the estimated coefficients from the auxiliary model in the economic model and in the Italian household data are well aligned (see the last two rows of Table 4). The average percentage deviation from target between data and simulated model amounts to 1.59%, the average absolute deviation equals 0.59%.

No advance information To further investigate the relevance of advance information for understanding the correlation between consumption expenditures and expected earnings, we estimate a model with uninformative signals by restricting signal precision to $\kappa = 1/N$. Table B.3 in Appendix B.4 displays the estimation results. The main message is that without informative signals, the model largely underestimates the consumption elasticity with respect to future earnings expectations, confirming that advance information is indeed the relevant mechanism to explain the elasticity.

7 Implications of the estimated structural model

In this section, we use the structural model to address several questions that can't be answered using survey data alone. We start with quantifying how much advance information Italian households possess on their future earnings. Afterwards, we study how consumption expenditures react to earnings expectations not in the future but today. Finally, we investigate whether the relevance of individual expectations for consumption choices differs across the wealth distribution.

7.1 How much advance information do Italian households possess?

In the Italian household data, earnings expectations and subsequent earnings realizations are not well synchronized, which does not allow to accurately quantify households' advance information. For this reason, we use the economic model to ask how important private signals are

Figure 4: Identification of $\kappa, \underline{a}, \beta$

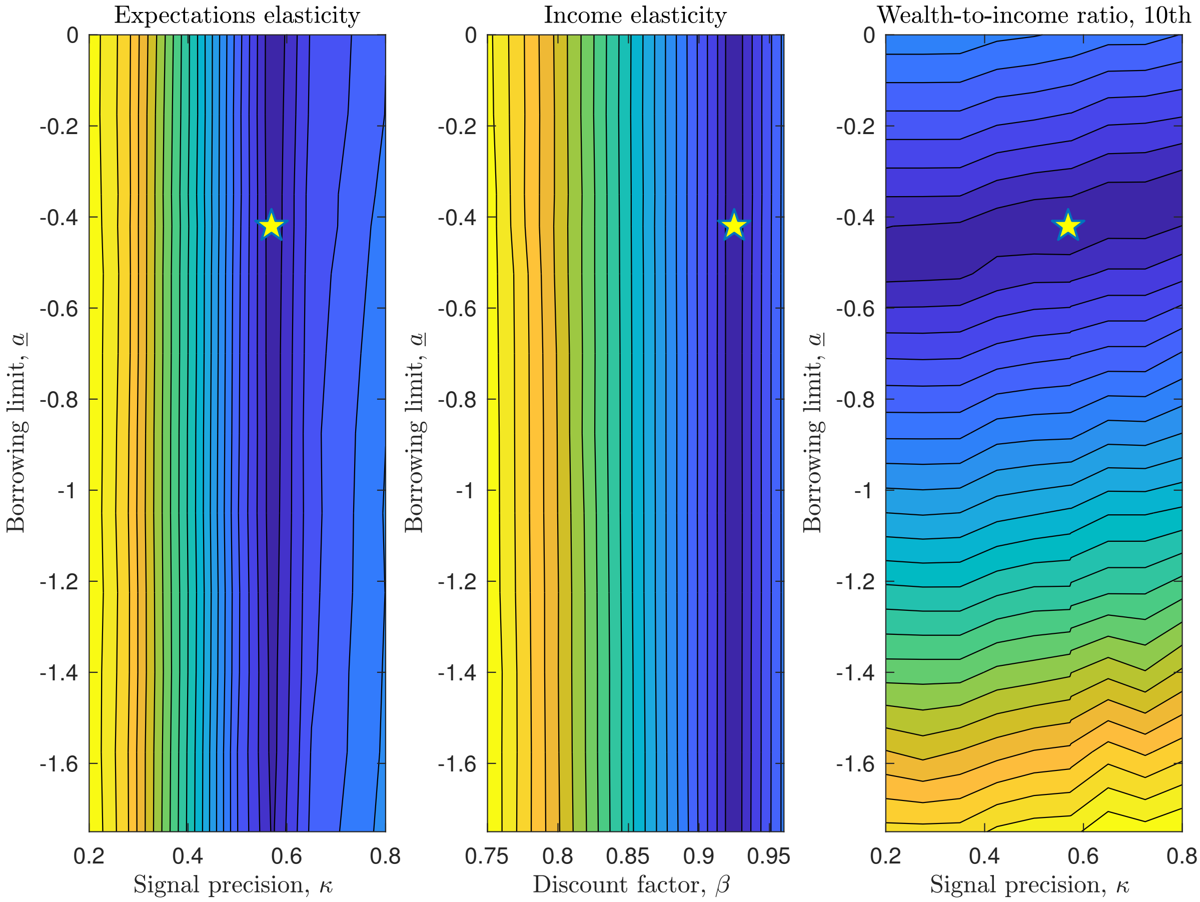

Notes: Absolute deviations of auxiliary model parameter estimates in the structural model from survey data estimates as functions of $\kappa, \underline{a}, \beta$*. All other parameters fixed at their* Θ_{min} *values. Dark blue (yellow) indicates the lowest (highest) deviation.*

for predicting future earnings. The answer to this question also offers an interpretation of the estimated signal precision of $\kappa = 0.569$.

To serve this goal, we compute the reduction in the conditional variance (the mean-squared forecast error) as a result of additionally conditioning expected earnings on private signals. More formally, we compute

$$\tilde{\kappa}(\kappa) = \frac{\mathrm{MSFE}_y - \mathrm{MSFE}_s(\kappa)}{\mathrm{MSFE}_y}, \qquad 0 \leq \tilde{\kappa}(\kappa) \leq 1 \tag{12}$$

with

$$\mathrm{MSFE}_y = \sum_y \pi(y) \sum_{y'} \pi(y'|y) \left\{ \log(y') - \mathrm{E}\left[\log(y')|y\right] \right\}^2$$
$$\mathrm{MSFE}_s(\kappa) = \sum_s \pi(s) \sum_{y'} \pi(y'|s) \left\{ \log(y') - \mathrm{E}\left[\log(y')|s\right] \right\}^2 \leq \mathrm{MSFE}_y,$$

Table 5: Consumption and expectations formed in $t+1$ and t

	Period-$t+1$ expectations	Period-t expectations
$\log(y_{it})$	0.205 (0.029)	0.059 (0.024)
$\log(\mathrm{E}_{it+1}\, y_{it+2})$	0.132 (0.013)	–
$\log(\mathrm{E}_{it}\, y_{it+1})$	–	0.302 (0.016)

Notes: Regression results for the simulated economic model at $\Theta_{\min}$ with earnings expectations conditional on information available in period-$t+1$ (as in the data) and in period t. Consumption without ME. Standard errors (S.E.) computed from drawing 200 times from the asymptotic distribution of Θ in parentheses.

$\pi(s)$ is the joint invariant distribution of earnings and signals induced by P_s, and $\mathrm{E}\,[\log(y')|y]$, $\mathrm{E}\,[\log(y')|s]$ are the expected log earnings conditional on earnings only and jointly on earnings and signals, respectively. Thus, $\tilde{\kappa}$ captures the difference in earnings uncertainty as measured by an econometrician – ignoring the private information on future shocks on the household level – and the uncertainty as perceived by households stemming from their subjective expectations. Quantitatively, we find that private signals reduce the conditional variance by $\tilde{\kappa}(0.569) = 19.18\%$, estimated with a standard error of 0.066.[18]

Dominitz (1998) provides direct evidence on the predictive power of subjective expectations using US data stemming from a specifically designed module in the Survey of Economic Expectations (SEE) in 1993–1994. In the spring and fall of 1993, survey participants are asked to report their current earnings and their subjective earnings expectations are elicited. In the spring of 1994, the respondents are asked again for their realized earnings. Using this data, he runs two types of linear predictor regressions, one that uses current earnings as predictor for future earnings, and one that additionally employs the mean of respondents' subjective earnings expectations for predicting future earnings. He finds that reported subjective expectations yield predictive value for the 6 and 12 months ahead forecasts (see his Table 7 on page 385), reducing the MSFE by 21.35% for the fall prediction (6 months ahead) and by 11.78% for the spring prediction (12 months ahead), respectively. Our estimates are consistent with this direct evidence.

7.2 How does consumption respond to current earnings expectations?

In the Italian data, there is a mismatch between the moment when consumption decisions are made and when households are asked for their expected earnings. While consumption expenditures are reported for the year t, expected earnings refer to period $t+2$ based on

[18] All standard errors in this section are based on drawing 200 times from the asymptotic distribution of Θ.

Figure 5: Consumption and advance information at the borrowing limit

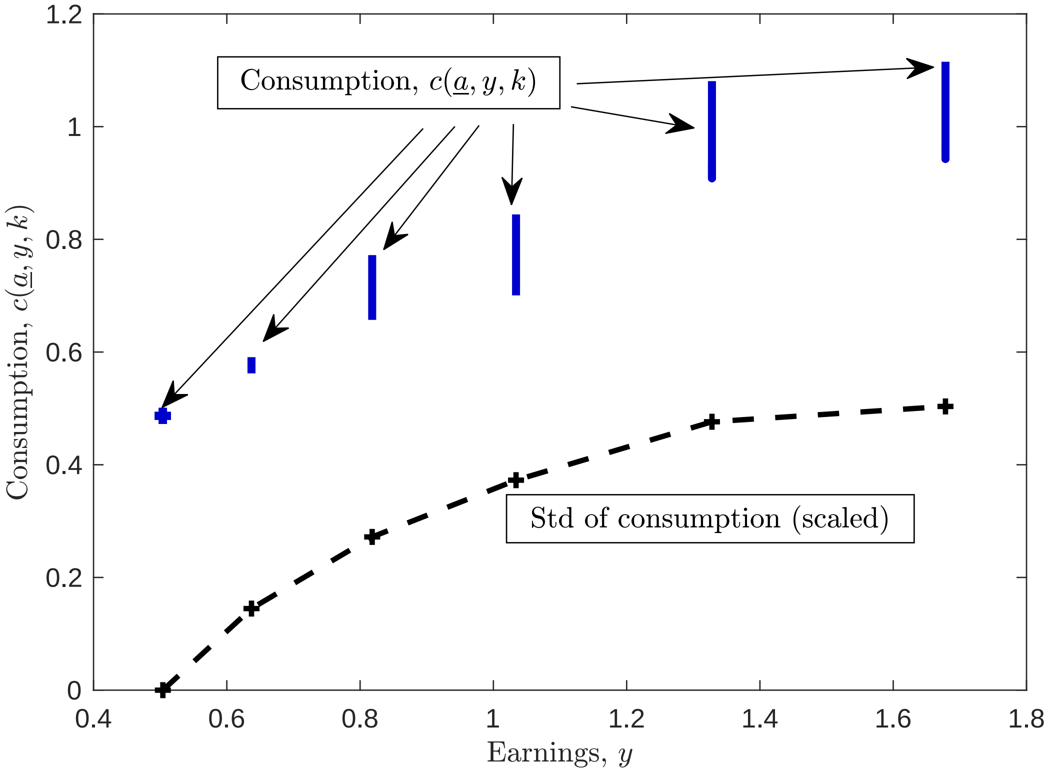

Notes: Consumption of households at the borrowing limit, $c(\underline{a}, y, k)$. Upper end consumption interval: $c(\underline{a}, y, k = \max\{y'\})$. Lower end consumption interval: $c(\underline{a}, y, k = \min\{y'\})$. Standard deviation of consumption conditional on earnings (scaled), $7 \times \mathrm{std}_k(c|y)$.

information at the beginning of year $t + 1$. The economic model can be used to resolve this mismatch and to investigate how current consumption decisions c_{it} react to changes in current earnings expectations $\mathrm{E}_{it}\, y_{it+1}$. The results of this exercise can be found in Table 5.

Compared to the auxiliary model in Equation (7), the consumption-earnings elasticity drops from 0.205 to 0.059 when current expected earnings are taken into account (see Row 1, Columns 1 and 2). Without period-t expectations, current earnings capture two aspects relevant for households' consumption decisions. It determines the resources that are available for consumption today but also informs about available resources for tomorrow because earnings are persistent. When current expected earnings are taken into account, current earnings do not contain any additional information relevant for future resources. As a consequence, the earnings elasticity decreases and the expectations elasticity increases from 0.132 to 0.302.

Table 6: Consumption-expectations elasticities: wealth quintiles

	Households: $\mathrm{E}\,(y_{it+1}\vert y_{it}, k_{it})$	Econometricians: $\mathrm{E}\,(y_{it+1}\vert y_{it})$
$\log(y_{it})$	0.059 (0.023)	0.092 (0.017)
$\log(\mathrm{E}_{it}\, y_{it+1})_{a\in(0,20th]}$	0.516 (0.032)	0.509 (0.045)
$\log(\mathrm{E}_{it}\, y_{it+1})_{a\in(20,40th]}$	0.317 (0.022)	0.273 (0.030)
$\log(\mathrm{E}_{it}\, y_{it+1})_{a\in(40,60th]}$	0.247 (0.012)	0.198 (0.018)
$\log(\mathrm{E}_{it}\, y_{it+1})_{a\in(60,80th]}$	0.204 (0.009)	0.153 (0.015)
$\log(\mathrm{E}_{it}\, y_{it+1})_{a\in(80,100th]}$	0.163 (0.008)	0.106 (0.008)

Notes: Regression results for the simulated economic model at $\Theta_{\min}$ with period-t earnings expectations depending on the wealth-quintile at the beginning of the period. Column (1): expectations of the households, $\mathrm{E}\,(y_{it+1}|y_{it}, k_{it})$. Column (2): expectations of econometricians, $\mathrm{E}\,(y_{it+1}|y_{it})$. No ME in consumption. Standard errors (S.E.) computed from drawing 200 times from the asymptotic distribution of Θ in parentheses.

7.3 Are consumption responses to expected earnings heterogeneous?

The economic model allows us to zoom in on how consumption choices of a particular type of households are affected by signals and expected earnings.

In Figure 5, we illustrate how consumption choices of households at the borrowing constraint are affected by the signals. As displayed in the upper left part of the figure, households with a low earnings realization do not adjust their consumption according to the particular signal realization. Instead, they behave hand-to-mouth and consume only their earnings diminished by their interest payments on their debt independent from their particular signal about future earnings. This picture changes when earnings increase. Instead of being given by a single point, consumption now differs across different signal realizations, with the upper end of the consumption interval in case of a bright future, $k = \max\{y'\}$, and the lower end for $k = \min\{y'\}$.

Thereby, the strength of the consumption response to the signals increases with earnings as indicated by the monotonically increasing earnings-conditional standard deviation of consumption in the lower panel of the figure. The private signals are relevant information in households' decisions on current and future savings. Households that are borrowing constrained and receive a low earnings shock lack the means to save at all. As earnings increase, households have the possibility to save for the future and then the private signals can have a stronger impact on their consumption-savings decisions.

As a next step, we study whether the relevance of subjective expectations for consumption choices differs systematically depending on the wealth position of the household. To investigate this, we estimate consumption elasticities depending on the wealth position at the beginning of the period that are displayed in Table 6. The main message from Column (1) is that subjective expectations play a more important role for households at the left tail of the wealth distribution; the consumption elasticity of households in the first quintile is approximately three times larger than the elasticity of households in the fifth quintile. Thereby, the consumption elasticities decrease monotonically as the wealth position of the household improves.

Intuitively, the better is the wealth position, the less relevant become earnings expectations because the higher wealth allows for better consumption smoothing. At the left tail of the distribution, consumption choices are motivated by precautionary savings considerations to prevent being borrowing constrained in future periods, and earnings expectations are an important factor in these considerations. At the top of the wealth distribution, avoiding to hit the borrowing limit in the future is less of a concern and earnings expectations matter mainly because of their impact on the permanent income of the household.

In Column 2 of Table 6, we replace households' subjective expectations (conditional on current earnings and signals) with the expectations conditional only on current earnings. Comparing the results in both columns, the pattern of the expectations elasticities is similar across the wealth distribution, but there are two main differences. First, the elasticities are smaller over the whole wealth distribution. Second, the estimated earnings elasticity is more than 50% larger than the actual elasticity of households. These results imply that if one ignores the information contained in expected earnings, one tends to over-estimate the importance of current earnings and to under-estimate the importance of expected earnings for consumption choices.

8 Is advance information policy relevant?

In this section, we study whether advance information is relevant for tax policy. First, we contrast estimated Italian economies with and without advance information with respect to consumption inequality and consumption smoothing.[19] Second, we ask whether these differences have quantitatively important implications for the design of an optimal progressive tax scheme.

For comparability, we consider both economies in general equilibrium, that is, given all other parameters, we adjust the interest rate to exactly satisfy bond market clearing – as formalized in Definition 1 part (iv).

[19] The parameter estimates without advance information can be found in Appendix B.4 in Table B.3.

Consumption inequality and smoothing As an inequality measure, we compute the risk-sharing ratio RS, which is defined as follows

$$RS = 1 - \frac{\text{std}_c}{\text{std}_y}, \tag{13}$$

with $\text{std}_x = \text{std}[\log(x)]$ as the cross-sectional standard deviation of log earnings and consumption. As one extreme, if $\text{std}_c = \text{std}_y$, then $RS = 0$, and there is no private risk sharing against fluctuations in after-tax earnings. On the other hand, if $\text{std}[\ln(c)] = 0$, then $RS = 1$, implying full risk sharing with respect to income shocks and the absence of consumption inequality.

With advance information, the risk-sharing ratio equals 0.335, which is very similar to the one without the signals (see Panel I of Table 7). On the first glance, such small differences appear to be surprising because with advance information households can make better informed consumption-saving decisions, which should result in less unequal consumption. On the other hand, however, consumption with advance informations spreads out to incorporate the information about future earnings provided by the signals, which increases consumption inequality.

The two information environments have different implications with respect to consumption smoothing as measured by the sensitivity of consumption growth with respect to earnings growth, which is given by the coefficient $\beta_{\Delta y}$ in the following regression equation

$$\Delta c_{it} = \psi + \beta_{\Delta y} \Delta y_{it} + \nu_{it}, \tag{14}$$

where ψ is a constant, and ν_{it} a residual; Δc_{it} and Δy_{it} are the growth rates of consumption and earnings of household i in period t. When the coefficient $\beta_{\Delta y}$ is zero, then consumption growth is perfectly smooth. The higher the coefficient, the less smoothing is achieved. In Table 7, we display the different estimates of $\beta_{\Delta y}$ with and without advance information, unconditional and conditional on the wealth position of the household.

With advance information, the unconditional sensitivity coefficient is with a value of 0.140 about 30% lower than without the signals, indicating a lower pass-through of earnings changes to consumption with informative signals on average (see upper panel). The reason is that with informative signals part of the earnings changes are predicted in the period before, which allows households to prepare beforehand for the earnings change by adjusting their savings choice. The pass-through is also lower when conditioning on the wealth position of a household; households above and below median wealth smooth consumption better with advance information.

Table 7: Consumption inequality, smoothing and welfare

	Advance information	No advance information
I. Risk sharing, RS	0.335	0.332
II. Consumption smoothing, $\beta_{\Delta y}$		
Unconditional	0.140	0.196
Conditional, below median wealth	0.194	0.255
Conditional, above median wealth	0.085	0.135
III. Welfare and tax progressivity		
Certainty equivalent consumption	0.968	0.945
Equivalent tax progressivity, T	-0.043	0.045

Notes: General equilibrium. Simulated data of the estimated models at $\Theta_{\min}$, with and without advance information. Risk-sharing ratio, RS, consumption sensitivity, $\beta_{\Delta y}$, ex-ante welfare measured in certainty equivalent consumption, and the change in tax progressivity that yields the same welfare with and without advance information.

What are the implications for progressive taxation in Italy? The classical argument in favor for progressive income taxes is that private risk sharing is imperfect with the consequence that income shocks result in undesirable fluctuations in household consumption. On the other hand, when tax progressivity becomes too high, it results in distortions in households' decisions to supply labor and to invest into their human capital.[20] The numbers in the upper two panels of Table 7 indicate roughly the same degree of risk sharing, but better consumption smoothing with advance information. In the following, we ask whether the systematic differences between the two information environments have implications for the optimal design of progressive income taxation.

In the economic model, neither labor supply nor investment in human capital is endogenous. Further, it is not directly clear in what direction advance information affects these costs of progressive taxation relative to the case without signals. In the following, we therefore assume that these costs are identical with and without advance information and focus on the benefits of progressive taxation. Consider the following simple tax function, linking before-tax income y and after-tax income y_{disp} as follows[21]

$$y_{disp} = (1 - T)y + T.$$

In our model, y denotes after-tax earnings. Thus, y_{disp} captures the possibility of more ($T > 0$)

[20] Exemplary papers on optimal tax progressivity are Cochrane (1993) and Heathcote, Storesletten, and Violante (2017).

[21] The functional form is proposed in Krueger and Perri (2011). It is progressive (regressive) for $T > 0$ ($T < 0$) because average taxes increase (decrease) in T, for $T > 0$ ($T < 0$).

or less progressivity ($T < 0$) relative to the status quo. To address the question on the policy relevance of advance information, we first compute social welfare with and without informative signals. As welfare measure, we employ households' ex-ante utility in the invariant distribution, that is, before the realization of any shock, expressed in certainty equivalent consumption.

We find that welfare with advance information permanently exceeds welfare without informative signals by 2.43% (see Panel III of Table 7). Based on the welfare difference, we compute the change in tax progressivity necessary such that welfare is identical in both information environments.

The main message is that taxes with advance information can be less progressive; a value of $T = -0.044$ in the economy with signals results in the same welfare than without signals (see Panel III).[22] This implies a variance of log after-tax earnings that is 10.17% higher than the status quo. Alternatively, an increase in tax progressivity of $T = 0.045$ without signals produces the same welfare as in case of informative signals, which corresponds to a reduction of after-tax earnings risk by 9.97%. Given the assumed constant costs of progressive taxation, we conclude that advance information is policy relevant; households with advance information prefer less progressive income taxes.

9 Conclusions

Consumption-saving models are usually based on the assumption that households and econometricians share the same information set. This implies that individuals' expected future income and their corresponding income uncertainty are fully captured by the estimated income process.

In this paper, we have investigated the role of individual expectations for consumption-savings decisions. Based on Italian survey data, we have provided evidence that individual earnings expectations affect consumption expenditures. These results are robust to the way we control for current earnings and also hold when we exploit the panel dimension of the data set and control for time-invariant unobserved heterogeneity. In a second step, we have developed and estimated an incomplete markets model with rational households that receive private signals on their future earnings realizations.

Our results show that households receive informative signals and are therefore better in-

[22] The computed reduction in tax progressivity is a conservative estimate because we do not take the detrimental effects of progressive taxation into account; with signals, any decrease T will also reduce the detrimental effects of progressive taxation, requiring additional reductions in T to render the specifications with and without signals equivalent.

formed about their future earnings than an econometrician. Ignoring agents' advance information leads to an overestimation of households' earnings uncertainty by around 19%. Advance information on their future earnings allows households to better smooth their consumption compared to a situation without private signals, and social welfare with advance information permanently exceeds social welfare in a counterfactual situation without private signals by 2.4%. This implies that households with advance information prefer less progressive earnings taxes.

The main contribution of this paper is to provide direct evidence on the relevance of earnings expectations for consumption decisions, and to demonstrate the economic relevance of advance information within the context of a structural consumption-savings model. Recent evidence based on rich panel data suggest non-Gaussian features of the earnings changes (Guvenen et al., 2021) and non-linearities in the earnings process (Arellano et al., 2017). Combining these more flexible earnings processes with a structural model incorporating advance information would be an interesting topic for future research.

References

Aiyagari, S. R. (1994): "Uninsured idiosyncratic risk and aggregate saving," *Quarterly Journal of Economics*, 109, 659–684.

Arcidiacono, P., V. J. Hotz, A. Maurel, and T. Romano (2020): "Ex ante returns and occupational choice," *Journal of Political Economy*, 128, 4475–4522.

Arellano, M., R. Blundell, and S. Bonhomme (2017): "Earnings and consumption dynamics: a nonlinear panel data framework," *Econometrica*, 85, 693–734.

Arellano, M., R. Blundell, S. Bonhomme, and J. Light (2021): "Heterogeneity of Consumption Responses to Income Shocks in the Presence of Nonlinear Persistence," Mimeo, University of Chicago.

Auclert, A. (2019): "Monetary policy and the redistribution channel," *American Economic Review*, 109, 2333–67.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption inequality and partial insurance," *American Economic Review*, 98, 1887–1921.

Browning, M., L. P. Hansen, and J. J. Heckman (1999): "Micro data and general equilibrium models," in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier, vol. 1 of *Handbook of Macroeconomics*, chap. 8, 543–633.

Campbell, D., A. Carruth, A. Dickerson, and F. Green (2007): "Job insecurity and wages," *The Economic Journal*, 117, 544–566.

Caprara, D., R. de Bonis, and L. Infante (2018): "La ricchezza delle famiglie in sintesi: l'Italia e il confronto internazionale," *Banca d'Italia Occasional Papers*, 470.

Cochrane, J. (1993): "A simple test for consumption insurance," *Journal of Political Economy*, 99, 957–976.

Coibion, O., Y. Gorodnichenko, and M. Weber (2022): "Monetary policy communications and their effects on household inflation expectations," *Journal of Political Economy*, 130, 1537–1584.

Cunha, F., J. Heckman, and S. Navarro (2005): "Separating uncertainty from heterogeneity in life cycle earnings," *Oxford Economic Papers*, 57, 191–261.

Dominitz, J. (1998): "Earnings expectations, revisions, and realizations," *The Review of Economics and Statistics*, 80, 374–388.

Dominitz, J. and C. F. Manski (1997): "Using expectations data to study subjective income expectations," *Journal of the American Statistical Association*, 92, 855–867.

Duso, T. (1999): "Complete markets in Italy: an analysis on micro data," *Giornale degli Economisti e Annali di Economia*, 58 (Anno 112), 1–25.

Ejrnæs, M. and M. Browning (2014): "The persistent–transitory representation for earnings processes," *Quantitative Economics*, 5, 555–581.

Fella, G., S. Frache, and W. Koeniger (2020): "Buffer-stock saving and households' response to income shocks," *International Economic Review*, 61, 1359–1382.

Fuster, A. and B. Zafar (2022): "Survey experiments on economic expectations," in *Handbook of Economic Expectations*, Elsevier.

Gourieroux, C., A. Monfort, and E. Renault (1993): "Indirect inference," *Journal of Applied Econometrics*, 8, S85–S118.

Guiso, L., T. Jappelli, and L. Pistaferri (2002): "An empirical analysis of earnings and employment risk," *Journal of Business Economics and Statistics*, 20, 241–253.

Guiso, L., T. Jappelli, and D. Terlizzese (1992): "Earnings uncertainty and precautionary saving," *Journal of Monetary Economics*, 30, 307–337.

——— (1996): "Income risk, borrowing constraints, and portfolio choice," *American Economic Review*, 86, 158–172.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021): "What do data on millions of U.S. workers reveal about lifecycle earnings dynamics?" *Econometrica*, 89, 2303–2339.

Guvenen, F. and A. A. Smith (2014): "Inferring labor income risk and partial insurance from economic choices," *Econometrica*, 82, 2085–2129.

Heathcote, J., K. Storesletten, and G. L. Violante (2017): "Optimal tax progressivity: an analytical framework," *The Quarterly Journal of Economics*, 132, 1693–1754.

Hendren, N. (2017): "Knowledge of future job loss and implications for unemployment insurance," *American Economic Review*, 107, 1778–1823.

Jordà, O., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019): "The rate of return on everything, 1870–2015," *Quarterly Journal of Economics*, 134, 1225–1298.

Kaplan, G. and G. L. Violante (2010): "How much consumption insurance beyond self-insurance?" *American Economic Journal: Macroeconomics*, 2, 53–87.

Kaufmann, K. and L. Pistaferri (2009): "Disentangling Insurance and Information in Intertemporal Consumption Choices," *American Economic Review P&P*, 99, 387–92.

Krueger, D. and F. Perri (2011): "Public versus private risk sharing," *Journal of Economic Theory*, 146, 920–956.

Manski, C. F. (2018): "Survey measurement of probabilistic macroeconomic expectations: progress and promise," *NBER Macroeconomics Annual*, 32, 411–471.

Mueller, A. I., J. Spinnewijn, and G. Topa (2021): "Job seekers' perceptions and employment prospects: heterogeneity, duration dependence, and bias," *American Economic Review*, 111, 324–63.

Pedroni, M., S. Singh, and C. A. Stoltenberg (2022): "Advance information and consumption insurance: evidence from panel data," Tinbergen Discussion Paper 22-032/VI, available at SSRN: https://ssrn.com/abstract=4035762.

Piketty, T. and G. Zucman (2014): "Capital is back: wealth-income ratios in rich countries, 1700-2010," *Quarterly Journal of Economics*, 129, 1255–1310.

Roth, C. and J. Wohlfart (2020): "How do expectations about the macroeconomy affect personal expectations and behavior?" *The Review of Economics and Statistics*, 102, 731–748.

Smith, V. K., D. H. Taylor, and F. A. Sloan (2001): "Longevity expectations and death: Can people predict their own demise?" *American Economic Review*, 91, 1126–1134.

Stephens, M. (2004): "Job loss expectations, realizations, and household consumption behavior," *The Review of Economics and Statistics*, 86, 253–269.

Stoltenberg, C. A. and S. Singh (2020): "Consumption insurance with advance information," *Quantitative Economics*, 11, 671–711.

Storesletten, K., C. I. Telmer, and A. Yaron (2004): "Cyclical dynamics in idiosyncratic labor market risk," *Journal of Political Economy*, 112, 695–717.

Wiswall, M. J. and B. Zafar (2021): "Human capital investments and expectations about career and family," *Journal of Political Economy*, 129, 1361–1424.

Appendix A Data and reduced form evidence

A.1 Measurement of subjective expectations

Subjective expectations about future earnings have been differently elicited before and after 1993. In the following, we document the way the interviewees have been asked about their expectations in the different survey waves and how we use these questions to measure expected earnings.

A.1.1 1989 and 1991 surveys

In the 1989 and 1991 surveys, respondents are asked about their expected evolution of their income from work in the next 12 months. For this they are shown the following intervals for potential income changes:

- More than 25% ___
- Between 20% and 25% ___
- Between 15% and 20% ___
- Between 13% and 15% ___
- Between 10% and 13% ___
- Between 8% and 10% ___
- Between 7% and 8% ___
- Between 6% and 7% ___
- Between 5% and 6% ___
- Between 3% and 5% ___
- Between 0% and 3% ___
- Less than 0% ___

  By how much? ___ %

In a first step, they are asked which intervals they completely exclude. In a second step, they are asked to distribute 100 points to the remaining intervals. We measure the expected change in earnings as the weighted sum of the different categories. For this, we use the mid-points of the intervals. The weights correspond to the number of points allocated by the interviewees to the different intervals divided by 100. For the highest interval we only have the lower bound. For our main specification, we have chosen a length of 5 percentage points for this interval. In some robustness analysis, we use alternative specifications for this interval. First, we have chosen an interval length of 10 percentage points, which results in a midpoint for the highest interval of 30%. Second, we have used the lower bound of this interval (25%). It turns out that these alternative choices do not affect our main results (see Table A.8). In case the interviewees expect with a positive probability a decrease in earnings, they are asked by how much they think the earnings might decrease. We use this answer and weight it with the points the individuals allocate to this potential event.

A.1.2 1995 and 1998 surveys

From the surveys in 1995 and 1998, we have individual responses on the expected minimum annual earnings (y_{min}), the expected maximum earnings (y_{max}) and the probability that expected earnings y are lower or equal to the midpoint y_{mid}: $\pi : y \leq (y_{min} + y_{max})/2$.

Triangular Distribution: In our main specification, we assume that the density $f(y)$ is composed of two triangle distributions, $y_{min} \leq y < y_{mid}$ with mass π = PROBLTX and $y_{mid} \leq y \leq y_{max}$, with mass $1 - \pi = 1 -$ PROBLTX. We get the following subjective density function:[23]

$$f(y) = \begin{cases} 8\frac{\pi(y-y_{min})}{(y_{max}-y_{min})^2} & y_{min} \leq y < y_{mid} \\ 8\frac{(1-\pi)(y_{max}-y)}{(y_{max}-y_{min})^2} & y_{mid} \leq y \leq y_{max} \\ 0 & else. \end{cases} \tag{15}$$

The CDF is

$$F(y) = \begin{cases} 0 & y \leq y_{min} \\ 4\frac{\pi(y-y_{min})^2}{(y_{max}-y_{min})^2} & y_{min} < y \leq y_{mid} \\ 1 - 4\frac{(1-\pi)(y_{max}-y)^2}{(y_{max}-y_{min})^2} & y_{mid} \leq y < y_{max} \\ 1 & y \geq y_{max} \end{cases} \tag{16}$$

The subjective mean of the earnings based on two triangular distributions is

$$\mathrm{E}(y) = \frac{\pi(2y_{min} + y_{max})}{3} + \frac{(1-\pi)(y_{min} + 2y_{max})}{3} \tag{17}$$

Uniform distribution: In some robustness analysis, we assume that the density $f(y)$ is composed of two uniform distributions, $y_{min} \leq y < y_{mid}$ with mass π = PROBLTX and $y_{mid} \leq y \leq y_{max}$, with mass $1 - \pi = 1 -$ PROBLTX. We get the following subjective density

[23] In the 1995 survey, PROBLTX is the survey response to earn a future income less than the midpoint, in 1998 it is more than the midpoint. We adjust the survey response PROBLTX for 1998, accordingly.

function:

$$f(y) = \begin{cases} \frac{2\pi}{y_{max}-y_{min}} & y_{min} \leq y < y_{mid} \\ \frac{2(1-\pi)}{y_{max}-y_{min}} & y_{mid} \leq y \leq y_{max} \\ 0 & else. \end{cases} \tag{18}$$

The CDF is

$$F(y) = \begin{cases} 0 & y \leq y_{min} \\ \frac{2\pi(y-y_{min})}{y_{max}-y_{min}} & y_{min} < y \leq y_{mid} \\ 1 - \frac{2(1-\pi)(y-\frac{y_{max}+y_{min}}{2})}{y_{max}-y_{min}} & y_{mid} \leq y < y_{max} \\ 1 & y \geq y_{max} \end{cases} \tag{19}$$

The subjective mean of the earnings based on two uniform distribution is

$$\mathrm{E}(y) = \frac{\pi(3y_{min} + y_{max})}{4} + \frac{(1-\pi)(3y_{max} + y_{min})}{4}.$$

It turns out that this alternative choice does not affect our main results (see Table A.8).

A.2 Descriptive Figures and Tables

Figure A.1: Distribution of earnings and expected earnings: 1989 and 1991

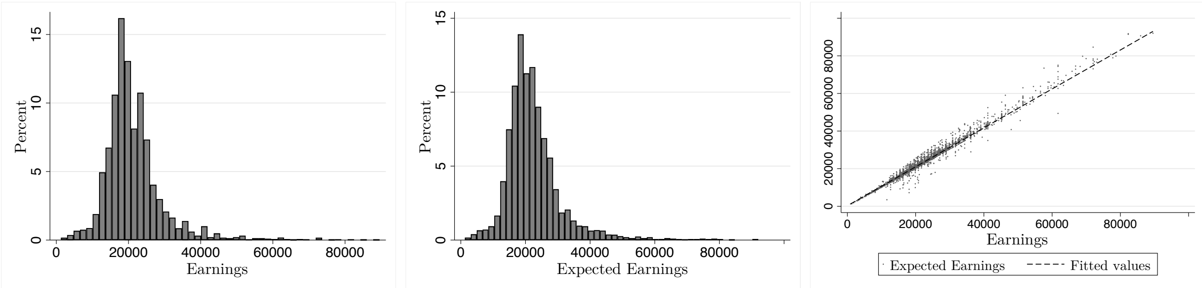

Notes: The left panel displays the distribution of annual earnings, the middle panel the distribution of expected annual earnings for the years 1989 and 1991, excluding observations above 100,000€ (n=5,387). The right panel displays the joint distribution.

Figure A.2: Distribution of earnings and expected earnings: 1995 and 1998

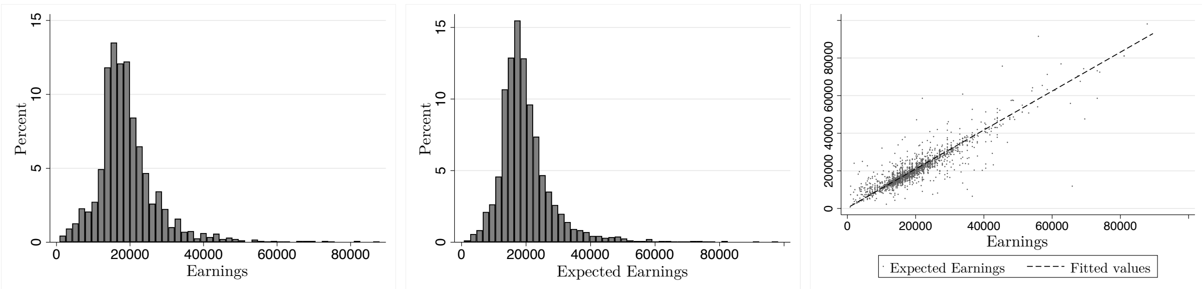

Notes: The left panel displays the distribution of annual earnings, the middle panel the distribution of expected annual earnings for the years 1995 and 1998, excluding observations above 100,000€ (n=2,266). The right panel displays the joint distribution.

Figure A.3: Distribution of nondurable consumption and wealth

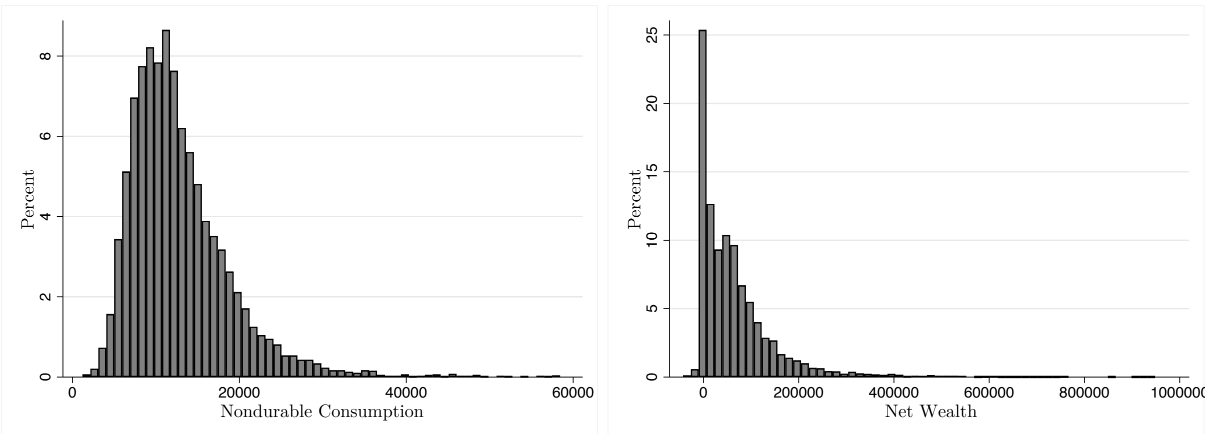

Notes: The left panel displays the distribution of annual nondurable consumption, excluding observations above 60,000€ (n=7,652). The right panel displays the distribution of net wealth at the beginning of the period, excluding observations below -50,000€ and above 1,000,000€ (n=7,651).

Table A.1: Descriptive statistics

Variables	Mean	Std. dev.
Age	42.902	(9.28)
Female	0.115	
Married	0.833	
Education		
None	0.014	
Elementary school	0.184	
Middle school	0.342	
High school	0.338	
Bachelor's degree and post-graduate qual.	0.121	
Region		
North-West	0.264	
North-East	0.172	
Centre	0.185	
South	0.267	
Islands	0.112	
Year		
1989	0.300	
1991	0.404	
1995	0.159	
1998	0.137	
Earnings	20,880	(9,163.86)
Expected earnings	21,940	(9,934.45)
Household's consumption of non-durable goods	12,915	(6,387.71)
Household's net wealth	66,267	(94,834.31)

Sample size : 7,659 observations. Economic variables are expressed in euro and in real terms. The earnings and consumption data refer to the previous calendar year. The household's net wealth refers to the wealth at the beginning of the previous calendar year. Household's consumption expenditures and net wealth are adjusted for the number of household members, applying the OECD equivalence scale.

Table A.2: Regressions of log earnings in $t+2$ on log expected earnings

	(1)	(2)
Log of expected earnings	0.646***	0.401***
	(0.0292)	(0.0745)
Log of lagged earnings	-	0.260***
	-	(0.0745)

Notes: Standard errors in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. n = 2,676. We control for age, gender, education, marital status, year and region. Real terms in Euro of year 2010.

A.3 Bivariate distributions of (expected) earnings and consumption

Figure A.4: Log (expected) earnings and log consumption: 1989 and 1991

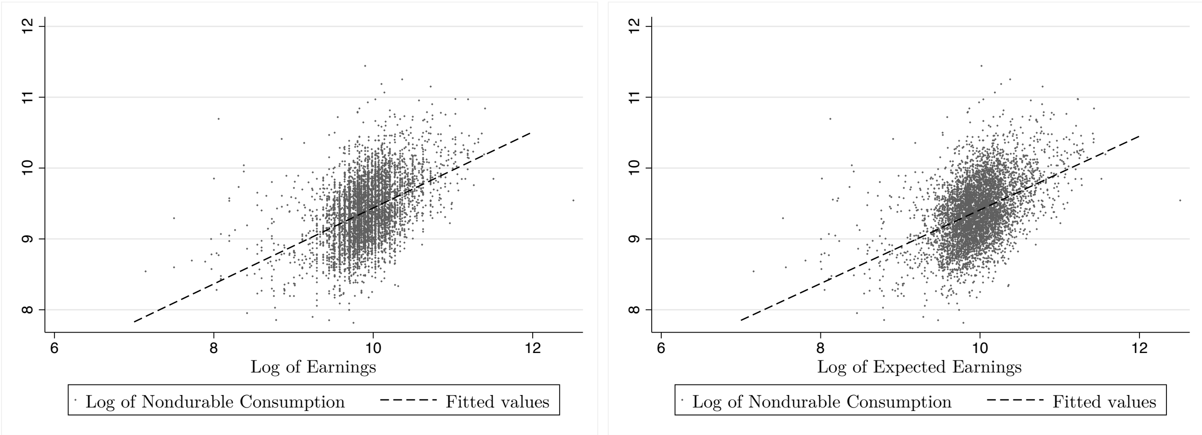

Notes: The left panel displays the joint distribution of the log of annual nondurable consumption and the log of annual earnings, the right panel the joint distribution of the log of annual nondurable consumption and the log of expected annual earnings for the years 1989 and 1991 (n = 5,390).

Figure A.5: Log (expected) earnings and log consumption: 1995 and 1998

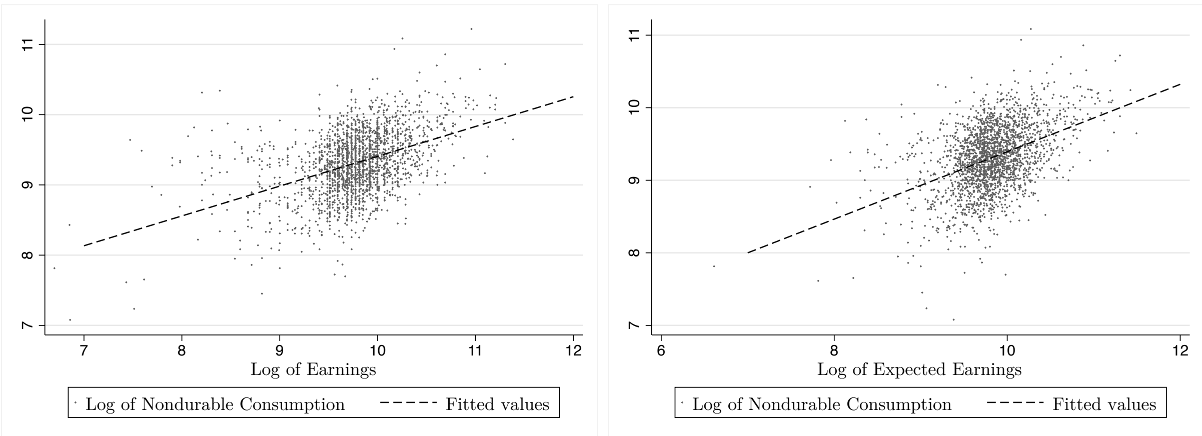

Notes: The left panel displays the joint distribution of the log of annual nondurable consumption and the log of annual earnings, the right panel the joint distribution of the log of annual nondurable consumption and the log of expected annual earnings for the years 1995 and 1998 (n=2,269).

A.4 Expected earnings and consumption: full estimation results

Table A.3: Consumption elasticities: pooled cross sections, full results

	(1)	(2)
Log of expected earnings	0.321***	0.133***
	(0.0142)	(0.0438)
Log of earnings		0.201***
		(0.0457)
Net Wealth	1.57e-06***	1.57e-06***
	(1.64e-07)	(1.64e-07)
Age	-0.0225***	-0.0227***
	(0.00353)	(0.00352)
Age squared	0.000246***	0.000246***
	(4.10e-05)	(4.10e-05)
Female	0.0390**	0.0432**
	(0.0169)	(0.0170)
Elementary School	0.0182	0.0151
	(0.0383)	(0.0380)
Middle School	0.0883**	0.0826**
	(0.0384)	(0.0381)
High School	0.189***	0.182***
	(0.0392)	(0.0390)
Bachelor's degree / post-graduate level	0.284***	0.276***
	(0.0424)	(0.0422)
1991	-0.0335***	-0.0377***
	(0.00803)	(0.00806)
1995	-0.0470***	-0.0468***
	(0.0113)	(0.0112)
1998	-0.117***	-0.118***
	(0.0122)	(0.0121)
North-East	-0.0261**	-0.0263**
	(0.0120)	(0.0120)
Center	-0.0582***	-0.0587***
	(0.0113)	(0.0113)
South	-0.257***	-0.253***
	(0.0108)	(0.0108)
Islands	-0.228***	-0.223***
	(0.0136)	(0.0136)
Married	-0.189***	-0.192***
	(0.0141)	(0.0139)
Constant	6.733***	6.633***
	(0.150)	(0.152)
R-squared	0.501	0.505

Notes: Standard errors are clustered at the individual level and reported in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. n = 7,659. Real terms in Euro of year 2010.

Table A.4: Consumption elasticities: fixed effects estimations, full results

	(1)	(2)
Log of expected earnings	0.349***	0.161***
	(0.0343)	(0.0554)
Log of earnings		0.228***
		(0.0528)
Net Wealth	5.54e-07***	5.54e-07***
	(8.76e-08)	(8.69e-08)
Age	0.00566	0.00905
	(0.0248)	(0.0246)
Age squared	-4.36e-05	-8.42e-05
	(0.000168)	(0.000167)
Elementary School	-0.112	-0.137
	(0.121)	(0.121)
Middle School	-0.0402	-0.0684
	(0.128)	(0.127)
High School	0.0188	-0.00834
	(0.132)	(0.131)
Bachelor's degree / post-graduate level	0.0324	-0.000117
	(0.145)	(0.144)
1991	-0.0305	-0.0344
	(0.0437)	(0.0434)
1995	-0.0405	-0.0432
	(0.126)	(0.125)
1998	-0.141	-0.147
	(0.187)	(0.186)
Married	0.00889	0.00505
	(0.0664)	(0.0659)
R-squared	0.143	0.157

Notes: Standard errors in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. n = 2,087. Earnings, net wealth and consumption are measured in real terms in Euro of the year 2010.

A.5 Expected earnings and consumption: sensitivity anaysis

Table A.5: Consumption elasticities: sensitivity depending on sample selection

	Baseline specification	Excluding managers	Excluding fem. househ. heads	Excluding outliers
Log of expected earnings	0.133***	0.128***	0.101**	0.101***
	(0.0438)	(0.0438)	(0.0488)	(0.0340)
Log of earnings	0.201***	0.205***	0.248***	0.237***
	(0.0457)	(0.0458)	(0.0512)	(0.0348)

Notes: Standard errors in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. For the baseline specification: n=7,659. For the sample without managers: n = 7,450. For the sample without female household heads: n=6,779. For the sample excluding the highest percentiles in (expected) earnings, nondurable consumption and net wealth: n=7,351. We control for net wealth, age, gender, education, marital status, year and region. Real terms in Euro of year 2010.

Table A.6: Consumption elasticities: pooled cross sections, separately for the periods 1989 / 1991 and 1995 / 1998

	All years	1989 / 1991	1995 / 1998
Log of expected earnings	0.133***	0.155*	0.133***
	(0.0438)	(0.0789)	(0.0498)
Log of earnings	0.201***	0.188**	0.199***
	(0.0457)	(0.0822)	(0.0524)

Notes: Standard errors in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. $n_{pooled} = 7,659$. $n_{89/91} = 5,390$, $n_{95/98} = 2,269$. We control for net wealth, age, gender, education, marital status, year and region. Real terms in Euro of year 2010.

Table A.7: Consumption elasticities: flexible control for earnings

	(1)	(2)	(3)	(4)
Log of expected earnings	0.115***	0.155***	0.130***	0.129***
	(0.0398)	(0.0314)	(0.0343)	(0.0372)
Higher order polynomials	yes	no	no	no
25 position indicators for earnings distr.	no	yes	no	no
50 position indicators for earnings distr.	no	no	yes	no
100 position indicators for earnings distr.	no	no	no	yes

Notes: Standard errors in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. n = 7,659. We control for net wealth, age, gender, education, marital status, year and region. Real terms in Euro of year 2010.

Table A.8: Consumption elasticities: alternative calculation of the expected earnings

	Baseline Specification	Model (1)	Model (2)	Model (3)
Log of expected earnings	0.133***	0.134***	0.133***	0.134***
	(0.0438)	(0.0496)	(0.0436)	(0.0439)
Log of earnings	0.201***	0.199***	0.201***	0.201***
	(0.0457)	(0.0524)	(0.0456)	(0.0458)

Notes: Standard errors in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. n = 7,659. Model (1) is based on a uniform distribution for the waves 1995 and 1998. Model (2) is based on the assumption that the highest interval for the expected growth as a width of 10% points, while in model (3) we use 25% for the highest interval. We control for net wealth, age, gender, education, marital status, year and region. Real terms in Euro of year 2010.

Table A.9: Consumption elasticities: adjusting for subjective inflation expectations using waves 1989 and 1991

	(1)	(2)
Log of expected earnings adjusted	0.333***	0.147*
for inflation expectations	(0.0171)	(0.0787)
Log of earnings		0.195**
		(0.0825)

Notes: Standard errors in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. n = 5,261. We control for net wealth, age, gender, education, marital status, year and region. The expected earnings are deflated at the individual level using the subjective inflation expectations, observed in the waves 1989 and 1991.

A.6 Heterogenous consumption elasticities

Table A.10: Consumption elasticities: good and bad news, median split

	(1)	(2)
Log of expected earnings	0.114**	
	(0.0503)	
Log of expected earnings, expected change below median		0.111**
		(0.0505)
Log of expected earnings, expected change above median		0.123**
		(0.0533)
Log of earnings	0.218***	0.216***
	(0.0513)	(0.0517)

Notes: Heterogeneous consumption-expectation elasticities with respect to good and bad news, median split. Standard errors in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. n = 7,659. We control for net wealth, age, gender, education, marital status, year and region. Real terms in Euro of year 2010.

Table A.11: Consumption elasticities: good and bad news, tertiles

	(1)	(2)
Log of expected earnings	0.113**	
	(0.0535)	
Log of expected earnings, low expected change		0.0932*
		(0.0535)
Log of expected earnings, medium expected change		0.141**
		(0.0586)
Log of expected earnings, high expected change		0.138**
		(0.0577)
Log of earnings	0.219***	0.209***
	(0.0542)	(0.0548)

Notes: Heterogeneous consumption-expectation elasticities with respect to good and bad news, tertile split. Standard errors in parentheses. ***, **, * indicate significance at 1%, 5% and 10% respectively. n = 7,659. We control for net wealth, age, gender, education, marital status, year and region. Real terms in Euro of year 2010.

Appendix B Model specification and estimation

B.1 Joint distribution of endowments and signals: formulas

In this subsection, we explain how to derive the formulas (3) and (4) stated in the main text.

We start with the derivation of the conditional probability of future endowments. Using the general formula for calculating conditional probabilities, we receive

$$\pi\left(y' = y_j | k = y_m, y = y_i\right) = \frac{\pi\left(y' = y_j, k = y_m, y = y_i\right)}{\pi\left(k = y_m, y = y_i\right)}.$$

The conditional probability can be simplified using the identity

$$\sum_{z=1}^{N} \pi\left(y' = y_z | k = y_m, y = y_i\right) = 1$$

to replace the denominator with the following expression

$$\pi\left(k = y_m, y = y_i\right) = \sum_{z=1}^{N} \pi\left(y' = y_z, k = y_m, y = y_i\right).$$

The joint probability in the numerator is

$$\pi\left(y' = y_j, k = y_m, y = y_i\right) = \pi_{ij} \kappa^{\mathbf{1}_{m=j}} \left(\frac{1-\kappa}{N-1}\right)^{1-\mathbf{1}_{m=j}},$$

where π_{ij} is the Markov transition probability for moving from endowment i to endowment z. For all endowment states that are not indicated by the signal, $j \neq m$, we assume here that their probability of occurrence conditional on the signal is identical and therefore equals $(1-\kappa)/(N-1)$. For the conditional probability of endowments, the general formula can then be written as

$$\pi\left(y' = y_j | k = y_m, y = y_i\right) = \frac{\pi_{ij} \kappa^{\mathbf{1}_{m=j}} \left(\frac{1-\kappa}{N-1}\right)^{1-\mathbf{1}_{m=j}}}{\sum_{z=1}^{N} \pi_{iz} \kappa^{\mathbf{1}_{m=z}} \left(\frac{1-\kappa}{N-1}\right)^{1-\mathbf{1}_{m=z}}} \tag{20}$$

which resembles (3) in the main text. For example, with two equally likely persistent endowment states, the conditional probability of receiving a low endowment y_l in the future conditional on a high signal $k = y_h$ and a low endowment today is given according to (20) by

$$\pi\left(y' = y_l | k = y_h, y = y_l\right) = \frac{(1-\kappa)\pi_{11}}{(1-\kappa)\pi_{11} + (1-\pi_{11})\kappa}.$$

The joint transition probability $\pi(s'|s) = \pi(y', k'|k, y)$ can be computed by combining the conditional probability of earnings with an assumption on the signal process. With signals following an exogenous first-order Markov process, the conditional probability $\pi(y', k'|k, y)$ is given by

$$\pi\left(y' = y_j, k' = y_l | k = y_m, y = y_i\right) = \pi_{ml} \frac{\pi_{ij}\kappa^{\mathbf{1}_{m=j}}\left(\frac{1-\kappa}{N-1}\right)^{1-\mathbf{1}_{m=j}}}{\sum_{z=1}^{N} \pi_{iz}\kappa^{\mathbf{1}_{m=z}}\left(\frac{1-\kappa}{N-1}\right)^{1-\mathbf{1}_{m=z}}} \quad \forall k', \tag{21}$$

with $\pi(k' = y_l|k = y_m)$ as the Markov signal transition probabilities that are consistent with household rationality as further explained in the following section.

B.2 Choosing the signal process

In this section, we explain how we reverse engineer internally consistent Markov signal processes (household rationality). We define consistent signal processes as follows.

Definition 2 (Consistent signal processes) *Consider the conditional probabilities $\pi(y'|s)$ and $\pi(s'|s)$ as defined in Equations* (3) *and* (4) *for a given Markov signal transition matrix P_k whose elements are the transition probabilities $\pi(k'|k)$. A Markov signal process is consistent if the following two consistency requirements are satisfied*

- ***Consistency requirement I****: The marginal distribution of the joint invariant distribution of endowments and signals, $\pi(s) = \pi(y, k)$, with respect to endowments equals the invariant distribution of endowments $\pi(y)$*

$$\hat{\pi}(y) = \sum_{k \in Y} \pi(s) = \sum_{k \in Y} \pi(y, k) \doteq \pi(y).$$

- ***Consistency requirement II****: The conditional distribution of endowments $\pi(y'|y)$ follows from integrating $\pi(y'|s) = \pi(y'|y, k)$ with respect to signals*

$$\hat{\pi}(y'|y) = \sum_{k \in Y} \pi(y'|s)\pi(k|y) = \sum_{k \in Y} \pi(y'|y, k)\pi(k|y) \doteq \pi(y'|y),$$

with $\pi(k|y)$ as the probability of signal k conditional on endowment y

$$\pi(k|y) = \frac{\pi(k, y)}{\sum_k \pi(k, y)}$$

Essentially, the two consistency requirements demand that households' subjective endowment transitions equal the actual endowment transitions. One can show that if Consistency requirement II is satisfied, so is Consistency requirement I, but not vice versa. To find consistent Markov processes, we apply a numerical procedure. Let N be the number of distinct states for log earnings y as result of a first order finite state Markov approximation of the earnings process estimated from Italian household data, that is, $y_t^i \in Y \equiv \{y_1, ..., y_N\} \subseteq \mathbb{R}_{++}$ and earnings transition probabilities $p_{kl} = \pi(y' = y_l | y = y_k)$. For each κ, we use the $N^2 - N$ restrictions imposed by Consistency Requirement II to solve for the signal transition probabilities $p_{ij}^k = \pi(k' = k^j | k = k_i)$.

B.3 Econometrical procedure

Numerical solution For each realization of Θ, the numerical procedure follows the following steps:

1. Given θ, σ_u (and ρ), the idiosyncratic earnings process is approximated as a first order Markov process with six earnings states.

2. Given the discretized earnings process from the first step and given κ, the Markov transition probabilities $\pi(k'|k)$ are chosen to satisfy the consistency requirements outlined in Definition 2, resulting in P_s.

3. The household problem is solved using value function iteration with 200 asset grid points.

4. For a fixed seed, the model is simulated for 700,000 time periods. To ensure stationarity, the first 100,000 draws are discarded.

5. Based on the remaining 600,000 periods, the auxiliary-model parameters $m(\Theta)$ are estimated and the objective function (9) is evaluated.

Estimation with global methods The estimation procedure can be summarized as follows.

1. We start with computing the values of the objective function and all auxiliary-model parameters $m(\Theta)$ for an equally spaced 6-dimensional grid of Θ for 235,008 points in total.

2. As one starting point, we choose the minimizer of the objective function over the six-dimensional grid. An alternative starting point is found by evoking a local solver starting

from 100 randomly selected points. The values of the objective function result from separately interpolating the auxiliary model parameters on the grid, using a piecewise cubic hermite interpolating polynomial. From the resulting 100 minimizer candidates, we pick the minimizer that results in the lowest value of the objective function and use this as the second starting value.

3. For all starting values, we continue with employing a two-step procedure

   (a) As first step, we use a derivative-free pattern search algorithm that starts with a search step in which points in the whole parameter range are evaluated. The points are generated by the latin hypercubes method. In the next step, the pattern-search algorithm uses polling to improve on the outcome of the search step until convergence is achieved.

   (b) Taking the optimizer from the first step as starting point, we conduct a second pattern-search round for which the derivative-free Nelder–Mead downhill algorithm is used in the search step. The search step is followed again by polling until convergence.

4. Eventually, we compare the objective function values for all alternative starting values and pick the minimizer associated with the lowest objective function value.

Parametric bootstrap To generate the parametric bootstrap, we set $\Theta = \Theta_{min}$ and simulate the economic model as many times as the given length of the bootstrap, $BT = 200$. For each simulation b, we randomly draw a seed. After discarding the first 100,000 draws to ensure stationarity, we estimate the auxiliary model on exactly as many observations as in the data. We collect the resulting auxiliary-model parameters in vector $\hat{m}_b$, for $b = 1, ...BT$.

Non parametric bootstrap To generate the non parametric bootstrap with clustering at the individual level, we draw for each bootstrap sample n individuals with replacement from the estimation sample. For each bootstrap sample, we estimate the auxiliary model and collect the corresponding auxiliary-model parameters in vector $\hat{m}_b$, for $b = 1, ...BT$, with $BT = 200$.

Asymptotic standard errors Here we follow the procedure described in Gourieroux, Monfort, and Renault (1993), which can be summarized as follows. First, we calculate the variance-covariance matrix of auxiliary model parameters $\text{cov}(\hat{m}_b)$, either from the parametric or non-parametric bootstrap. As next step, we compute the numerical derivative of the vector $m(\Theta)$

at Θ_{min}. Standard errors are then given as the square roots of diagonal elements of

$$\left[\frac{\partial m(\Theta)}{\partial \Theta} \times \text{cov}(\hat{m})^{-1} \times \frac{\partial m(\Theta)'}{\partial \Theta}\right]^{-1}.$$

B.4 Robustness exercises

In Table B.1, we report the asymptotic standard errors stemming from a parametric and non-parametric bootstrap. In Table B.2, we display the parameter estimates for different degrees of relative risk aversion. The estimation results for uninformative signals can be found in Table B.3.

Table B.1: Asymptotic standard errors: parametric versus non-parametric bootstrap

	Parameter estimates, Θ_{min}					
Bootstrap	κ	$\underline{a}$	θ	σ_u	β	σ_c
Parametric, baseline	0.569 (0.060)	−0.420 (0.096)	0.854 (0.011)	0.127 (0.009)	0.925 (0.003)	0.299 (0.003)
Non-parametric	0.569 (0.101)	−0.420 (0.117)	0.854 (0.034)	0.127 (0.024)	0.925 (0.003)	0.299 (0.005)

Asymptotic standard errors (SE) computed from a parametric (first row) and a non-parametric (second row) bootstrap with 200 repetitions in parentheses.

Table B.2: Parameter estimates and risk aversion

	Parameter estimates, Θ_{min}					
	κ	$\underline{a}$	θ	σ_u	β	σ_c
Model, $\sigma = 3$	0.568 (0.057)	−0.358 (0.305)	0.854 (0.013)	0.127 (0.010)	0.938 (0.003)	0.297 (0.003)
Baseline, $\sigma = 4$	0.569 (0.060)	−0.420 (0.096)	0.854 (0.011)	0.127 (0.009)	0.925 (0.003)	0.299 (0.003)
Model, $\sigma = 5$	0.568 (0.042)	−0.484 (0.092)	0.854 (0.010)	0.127 (0.013)	0.911 (0.003)	0.300 (0.003)

Notes: estimation results for relative risk aversion, $\sigma = (3, 4, 5)$. Baseline in boldface. Asymptotic standard errors (SEs) computed from a parametric bootstrap with 200 repetitions in parentheses.

Table B.3: Uninformative signals: parameter estimates and auxiliary model

I. Parameter estimates, Θ_{min}

$\underline{a}$	θ	σ_u	β	σ_c
−0.557 (0.103)	0.845 (0.012)	0.122 (0.008)	0.926 (0.002)	0.306 (0.004)

II. Auxiliary model

	$\varepsilon_{c,y}$	$\varepsilon_{c,\mathrm{E}\,y}$	$\sigma_{\eta,c}$	a_1	$\sigma_{\eta,y}$	β_{INS}	Med_{WIR}	$10\mathrm{th}_{WIR}$
Data	0.201	0.133	0.313	0.636	0.323	0.198	2.650	0.013
Model	0.237	0.048	0.320	0.636	0.323	0.227	3.094	0.013

Notes: Uninformative signals, $\kappa = 1/6$. **earnings**-process parameters set to match the auxiliary model parameters in Equation (8), $\theta = 0.845$ and $\sigma_u = 0.122$. Estimated parameters $\Theta = (\underline{a}, \beta, \sigma_c)$. Asymptotic standard errors (SEs) computed from a parametric bootstrap with 200 repetitions in parentheses.