Journal of Macroeconomics 79 (2024) 103566

Contents lists available at ScienceDirect

Journal of Macroeconomics

journal homepage: www.elsevier.com/locate/jmacro

Life-cycle wealth accumulation and consumption insurance[☆]

Claudio Campanale[a,*], Marcello Sartarelli[b,1]

[a] Universitá di Torino, CeRP (Collegio Carlo Alberto), Italy
[b] Universidad Complutense de Madrid, ICAE, Spain

ARTICLE INFO

JEL classification:
D15
E21

Keywords:
Precautionary savings
Epstein–Zin
Consumption insurance coefficients
Life-cycle

ABSTRACT

Households appear to smooth consumption in the face of income shocks much more than implied by life-cycle versions of the standard incomplete market model under reference calibrations. In the current paper we explore in detail the role played by the life-cycle profile of wealth accumulation. We show that a standard model parameterized to match the latter can rationalize between 81 and 100 percent of the consumption insurance against permanent earnings shocks empirically estimated by Blundell, Pistaferri and Preston (2008), depending on the tightness of the borrowing limit.

1. Introduction

Microeconomic evidence on individual consumption growth shows a large degree of idiosyncratic volatility, the observable sign of imperfect risk sharing in the data. Macroeconomic models with heterogeneous agents inherently feature imperfect risk-sharing at least qualitatively, but if they are to be credible tools for analyzing economic phenomena and assess economic policies then they need to be able to explain quantitatively the extent to which consumption is insured against income shocks that we observe in the data. The prototype standard incomplete market model (henceforth SIM), arguably the workhorse of heterogeneous agents macroeconomics, when parameterized according to reference calibrations falls significantly short of matching the empirical values of insurance against permanent earnings shocks, as estimated by Blundell et al. (2008), henceforth BPP. In this kind of model wealth, in the form of a single asset, is used to smooth consumption in the face of earnings fluctuations. In the current research we revisit the SIM model, but instead of only constraining average wealth accumulation as it has been done so far, we focus on the role that the exact profile of wealth over working life plays in determining the degree of consumption smoothing.

In order to study the role of the life-cycle pattern of wealth accumulation we modify the baseline self-insurance model by moving from standard expected utility to Epstein–Zin preferences. The key feature of Epstein–Zin preferences is that, contrary to standard expected utility preferences, they permit a complete separation between the elasticity of inter-temporal substitution (eis) and risk-aversion. This allows for a redistribution of wealth over the life-cycle in ways that lead to higher insurance coefficients. The intuition behind this result is that in the Epstein–Zin case, raising risk aversion while keeping the elasticity of inter-temporal substitution high, allows the model to increase early life precautionary savings without concurrently creating a strong motive for holding a very large stock of retirement wealth. As a consequence it becomes possible to match the empirical wealth-to-income ratio with plausibly high values for patience. The combination of high risk aversion, high patience and the willingness to accept inter-temporal redistribution

☆ We thank Luis Franjo, Giovanna Nicodano and participants to the 2018 EEA-ESEM and 2018 ASSET conference and 2019 SMN workshop for comments and suggestions. Financial support from the Spanish Ministerio de Economía y Competitividad, Spain (ECO201343119, ECO2015-65820-P and ECO2016-77200-P) is gratefully acknowledged. All remaining errors are our own.

* Correspondence to: Dipartimento ESOMAS, Universitá di Torino, Cso. Unione Sovietica 218bis, 10127, Torino, Italy.
E-mail addresses: claudiogiovanni.campanale@unito.it (C. Campanale), msartare@ucm.es (M. Sartarelli).

1 Universidad Computense de Madrid, Análisis Económico y Economía cuantitativa and ICAE, Campus de Somosaguas, 28223, Spain.

https://doi.org/10.1016/j.jmacro.2023.103566
Received 11 October 2022; Received in revised form 20 October 2023; Accepted 23 October 2023
Available online 7 November 2023

of consumption away from young ages moves wealth accumulation towards the early part of the life-cycle when it is most needed for insurance purposes and away from middle age, when the accumulation of retirement wealth and the lower effective residual persistence of the shocks makes the latter more easily insurable. For this reason, this mechanism increases the insurance coefficients for permanent shocks in the first part of the life-cycle without affecting those in mid-life. This has the effect of raising the average coefficients and at the same time of making their age profile flatter, hence closer to the flat profile found in the data.

Given that the mechanism that allows the model to generate insurance coefficients in line with the data is the redistribution of wealth across different periods of the life-cycle, it is important that the resulting life-cycle profiles of wealth match the data. For this reason we solve the model calibrating the coefficient of risk aversion so that it minimizes the distance between the model and data wealth accumulation profiles during working life. We find that 81 percent of the insurance coefficients against permanent earnings shocks measured by the BPP index can be rationalized by the model with a zero borrowing limit, compared to 25 percent for the standard expected utility model only matching the average wealth-to-income ratio. This figure raises to virtually 100 percent in the version of the model where borrowing is allowed subject only to the constraint that the household is able to repay for sure. Our analysis suggests that this result is almost entirely due to the calibration of the life-cycle wealth profiles with little specific contribution left for Epstein–Zin preferences. Moreover the result is obtained without adversely affecting the model performance along other dimensions of consumption like the age profile of the insurance coefficients and of the variance of log consumption.

Quantitative and empirically relevant studies of life-cycle consumption behavior date back to at least the work of Gourinchas and Parker (2002), yet a thorough exploration of the insurance properties of the SIM model has lagged behind, mainly because measuring and studying consumption insurance in the data, thus providing an empirical benchmark against which to test the model, has proven challenging. This occurs for two reasons: first, high quality panel data on both consumption and earnings are needed but not directly available and, second, the problem of identifying different shocks from the observable income process must be circumvented. The first problem arises because the two main data sets used to study household behavior in the US, that is the Panel Study on Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX) respectively lack consumption data or the panel dimension. With respect to the first issue, an example of an early effort in this sense is Attanasio and Davis (1996) who used several issues of the CEX to construct synthetic cohorts and study how the evolution of between groups earnings inequality translated into consumption inequality. Strictly speaking though, this does not measure insurance of shocks per se. With respect to the second issue, efforts have been made to distinguish between permanent and temporary shocks by using proxies like disability and short unemployment spells respectively (Dynarski and Gruber, 1997). Alternatively, others like (Krueger and Perri, 2006) have chosen to simply analyze the response of consumption to income shocks without trying to identify the different shocks.

A major step forward was made by Blundell et al. (2008), that used the CEX to estimate a food demand equation and then applied its inverse to PSID data on food consumption, thus obtaining an artificial data set with both a panel dimension and joint data on consumption and income. This, coupled with a suitable strategy to identify shocks, allowed them to come up with a first estimate of insurance coefficients.[2] Kaplan and Violante (2010) first evaluated the standard SIM model against the data to test if it can match BPP estimates of the insurance coefficients. They found that under standard parameterizations this model can explain between 19 and 61 percent of the empirical estimates of insurance coefficients against permanent shocks, depending on the assumption of the zero or natural borrowing constraint respectively.

In the wake of their paper, a few other quantitative papers have been written to extend the basic SIM model to better fit insurance data. These fall into four different strands. A first approach emphasizes less than unitary persistence in the earnings process. This strand includes the works by Karahan and Ozkan (2013) and De Nardi et al. (2020). Karahan and Ozkan (2013) estimate a richer earnings process allowing for age-varying earnings persistence, find that persistence is smaller early in the working life and that this feature increases the insurance coefficients against persistent shocks compared to the standard quantitative model. De Nardi et al. (2020) further extend their model by adding non-linearity and non-normality in the earnings process. This approach confirms the results in Karahan and Ozkan (2013) concerning insurance coefficients while at the same time improving the match of the model with the data on lifetime growth in consumption inequality. A second line of attack to the problem focuses on wages rather than earnings as primitives and studies the extent of insurance against wage shocks emphasizing spousal labor supply as the key mechanism. This line of research includes the works of Blundell et al. (2016) who provided benchmark empirical estimates and Wu and Krueger (2018) who developed the corresponding quantitative model. The third approach is represented by the work of Cerletti and Pijoan-Mas (2012) who extended the prototype SIM model to explore the role of substitutability of durable and non-durable goods in the consumption bundle to explain the observed insurance of non-durable consumption. Finally, the fourth approach is represented by the work of Hryshko and Manovskii (2017) who claim that the observed shortcoming of the SIM model is a consequence of measurement issues in the PSID that according to the authors' analysis is not truly representative of the population.

The current paper complements the quantitative literature mentioned in the previous paragraph by more carefully modeling the process of accumulation of wealth over working life in an otherwise standard incomplete-markets model with idiosyncratic earnings risks of various persistence. The key applied contribution of our approach is to the explanation of the gap between model and data estimates of the insurance coefficients of permanent shocks that the original work of Kaplan and Violante (2010) showed is harder to bridge. In pursuing this goal the emphasis on the whole profile of wealth accumulation over working life is crucial since the aggregate insurance coefficient is an average of the coefficients by age groups which in turn strongly depend on the groups' wealth, hence targeting average wealth over the whole population alone is not fully adequate to study this problem. While our approach

[2] In a recent paper Christelis et al. (2019) use survey questions to estimate the propensity to consume out of shocks. This approach is very promising but to our knowledge it has been so far applied only to temporary shocks.

provides a perspective on how to reconcile the SIM model with data about consumption insurance that is alternative to the ones put forth in the related literature, it still complements them since no matter what shock process or consumption bundle is specified, wealth accumulation is key in providing insurance and a model that explains the latter should be consistent with the former.

The rest of the paper is organized in the following way. Section 2 is devoted to explaining the model, Section 3 presents the calibration and Section 4 discusses the results. Finally, in Section 5 a brief conclusion is outlined. An online appendix outlines the methodology used to compute the insurance coefficients, presents further analysis of the model to support the intuition in the main text and describes the data and procedure used to estimate the life-cycle profiles of wealth.

2. Model

We consider a standard life-cycle economy where agents are endowed with Epstein–Zin preferences. Agents are meant to represent households and the two terms will be used interchangeably in what follows. The economy features a large number of ex-ante identical agents. Agents have finite lives and go through the two stages of life of working age and retirement. During working life they receive an exogenous stochastic stream of earnings that cannot be insured due to incomplete markets. During retirement they receive a constant pension benefit that depends on the full history of the household's earnings. They have access to a single risk-free asset that they can use to smooth consumption in the face of variable earnings, subject to a borrowing constraint. The model is cast in a partial equilibrium framework and there is no aggregate uncertainty. A cohort of agents is simulated and the model-generated patterns of consumption insurance are studied.

2.1. Demographics and preferences

Time is discrete with model periods of one year length. Agents live for a maximum of $T = 80$ model periods. They enter the model at age 20 and retire after $T^{ret} = 45$ years of work. In each period of life t they face a probability π_{t+1} of surviving one more year. Agents care only about their own consumption and do not value leisure, hence they supply inelastically their unitary endowment of time.

Households value the uncertain stream of future consumption according to the following inter-temporal utility function:

$$V_t(S_t) = \{c_t^{\gamma} + \beta E[\pi_{t+1} V_{t+1}^{\alpha}(S_{t+1})]^{\frac{\gamma}{\alpha}}\}^{\frac{1}{\gamma}} \tag{1}$$

where the variable S_t represents the set of all past histories of shocks up to age t and initial assets that can at each age be summarized into three state variables. As it will become clear in the next few sections, these state variables are cash-on-hand at the beginning of the period, the value of the permanent earnings shock and the average past realizations of gross labor earnings. In the above representation of utility γ is the parameter that controls the elasticity of substitution between current consumption and the certainty equivalent of future utility, the elasticity of substitution being given by $\frac{1}{1-\gamma}$. On the other hand, α is the parameter that controls the curvature of the future utility certainty equivalent function and corresponds to a risk aversion of $1 - \alpha$.[3] Finally, the parameter β determines the weight of future versus current utility and represents the subjective discount factor. In the expression above the expectation E is taken with respect to histories S_{t+1} up to $t + 1$ conditional on history S_t being realized up to age t.

2.2. Income process

During working life agents receive a stochastic flow of net earnings Y_{it} which can be expressed as:

$$logY_{it} = g_t + y_{it} \tag{2}$$

and

$$y_{it} = z_{it} + \varepsilon_{it} \tag{3}$$

where g_t is a deterministic component common to all households and y_{it} is the stochastic component of the labor income. In turn, the stochastic component can be decomposed into a transitory part ε_{it} and a permanent part z_{it} that follows the process:

$$z_{it} = z_{i,t-1} + \zeta_{it} \tag{4}$$

The initial realization of the permanent component is drawn from an initial distribution with mean 0 and variance σ^2_{z0}. The shocks ε_{it} and ζ_{it} are normally distributed with mean 0 and variances σ^2_{ε} and σ^2_{ζ}, are independent of each other, over time and across agents. Retired households receive a fixed pension benefit $P(\overrightarrow{Y}_i)$ where $\overrightarrow{Y}_i$ is the vector collecting all the realizations of gross earnings for agent i, that is, the pension benefit is a function of the history of all past earnings. Agents can save in a single asset. We denote the amount of the asset held by household i at age t with A_{it} and assume that the asset pays a constant return r. We assume that a borrowing constraint $A_{it} \geq \underline{A}$ holds. The household's budget constraint can then be written:

$$c_{it} + A_{i,t+1} = (1 + r)A_{it} + I_{it}Y_{it} + (1 - I_{it})P(\overrightarrow{Y}_i) \tag{5}$$

where I_{it} is an indicator function that takes a value of 1 if $T < T^{ret}$ and 0 otherwise.

[3] Alternatively adopting habit formation preferences would also allow us to control separately the elasticity of inter-temporal substitution and risk aversion. See for example Díaz et al. (2003). However this would come at a substantial extra computational cost, given the need to add the level of habit as a further state variable, without any advantage in terms of economic modeling.

2.3. Household's optimization problem

With the description of the model given above and omitting for simplicity of notation the index i for the household, we can write the optimization problem at each age. This will be described by the Bellman equation:

$$V_t(X_t, z_t, \overline{Y}_t) = \max_{c_t, A_{t+1}} \{c_t^{\gamma} + \beta E[\pi_{t+1} V_{t+1}^{\alpha}(X_{t+1}, z_{t+1}, \overline{Y}_{t+1})]^{\frac{\gamma}{\alpha}}\}^{\frac{1}{\gamma}} \tag{6}$$

where V_t is the value function at age t and the state variables are current cash-on-hand X_t, the realization of the permanent component of the earnings process z_t, and the average of past gross earnings realizations up to age t denoted with $\overline{Y}_t$. The households maximize the CES aggregator of current consumption and the certainty equivalent of future utility with respect to consumption c_t and asset holdings A_{t+1}, that are carried into the next period. The maximization is performed subject to the following constraints:

$$c_t + A_{t+1} \leq X_t \tag{7}$$

$$X_{t+1} = A_{t+1}(1+r) + I_{t+1} Y_{t+1} + (1 - I_{t+1}) P(\overline{Y}_{t+1}) \tag{8}$$

$$\overline{Y}_{t+1} = \begin{cases} \frac{t\overline{Y}_t + \tilde{Y}_{t+1}}{t+1} & \text{if } t < T^{ret} \\ \overline{Y}_{T^{Ret}} & \text{if } t \geq T^{ret} \end{cases}$$

The first inequality is a standard budget constraint that tells us that consumption plus assets carried into the next period cannot exceed current cash-on-hand. The second equality is the law of motion of cash-on-hand. Cash-on-hand in the next period is given by the assets carried into the next period augmented by the net interest rate earned, plus non financial income. It is understood that if the indicator function $I_{t+1} = 1$ then the agent is working and earns net labor income Y_{t+1}, while if $I_{t+1} = 0$ the agent is retired and collects social security benefits $P(\overline{Y}_{t+1})$. The last equation represents the law of motion of average past gross earnings that enter the calculation of the pension benefits. Gross earnings at age t are denoted $\tilde{Y}_t$ and are obtained from net earnings Y_t by way of a suitable tax function $\tau(\tilde{Y}_t)$. Before retirement average past earnings up to t are averaged with the newly received realization of gross earnings $\tilde{Y}_{t+1}$ to update the new value of average past earnings. After retirement the value of average past earnings is fixed at the level matured at retirement time, and denoted with $\overline{Y}_{T^{Ret}}$. Finally, the maximization is subject to the stochastic earnings processes defined in the previous subsection, and to the borrowing constraint $A_{t+1} \geq \underline{A}$.

3. Calibration

The model period is taken to be one year. Agents enter the labor market, hence the model, at age 20, retire at age 65 and die for sure at age 100. Before that age, the probability of survival from one year to the next are taken from the Berkeley Mortality Database. With respect to preference parameters we set β, the subjective discount factor, so that the average wealth-to-income ratio is equal to 2.5. While this value is lower than the one in the aggregate data, in practice it reflects correctly the wealth-to-income ratio in the bottom 95 percent of the earnings distribution in the PSID.[4] This is the part of the population we are interested in given that the empirical estimates of the insurance coefficients are based on the PSID and CEX, which are well known not to represent accurately the top of the distribution. We set exogenously the value of the intertemporal elasticity of substitution. Estimates in the literature vary a lot, from near zero to well over one, including 2 and above. As a compromise value we set it to 1.1 but explore also how the results change when it is set to 0.5 and 1.7.

Risk aversion is set so that the distance between the model and data wealth profile over the working life is minimized. In practice the calibration of the subjective discount factor and risk aversion — the two parameters calibrated endogenously — proceeds in two stages. We first set a grid of values for risk aversion. Then for each value of risk aversion we choose the value of β such that the model matches a wealth-to-income ratio of 2.5, a value that we can denote with $\beta^{KV}(\alpha)$. Finally we pick the value of α such that the distance between model and data wealth is minimized. The metric used for the minimization is the minimum squared distance. Thus if we denote with $W_t(\alpha, \beta)$ the average normalized wealth of age t agents generated by the model as a function of risk aversion and the subjective discount factor, and with $\widetilde{W}_t$ the corresponding empirical values the problem we solve is:

$$\min_{\alpha} \sqrt{\frac{1}{T^{ret}} \sum_{t=1}^{T^{ret}} (W_t(\alpha, \beta) - \widetilde{W}_t)^2} \tag{9}$$

where T^{ret} is the retirement age, subject to the constraint

$$\beta = \beta^{KV}(\alpha) \tag{10}$$

For the deterministic common component of the labor income process we take a third order polynomial in labor market experience, that is, age minus 20, and use coefficients for the average population based on work by Cocco et al. (2005). The process generates a peak to age 25 earning ratio of about 2.1 in line with the process used in Kaplan and Violante (2010), with the peak occurring at the same age. As for the stochastic component of earnings, we have to assign three parameters, that is, the variance

[4] While the best source for data about wealth is the Survey of Consumer Finances (SCF), as pointed out by Bosworth and Anders (2008), the two data sets generate very similar results once the top 5 percent wealthiest households are removed.

of the permanent and temporary shocks ζ and ε and the initial variance of the permanent shock $\sigma^2_{z_0}$. We give σ^2_ζ a value of 0.01 to match the increase in earnings dispersion over the life-cycle observed in PSID data and we assign a value of 0.05 to σ^2_ε based on the point estimate by Blundell et al. (2008). Finally, we set $\sigma^2_{z_0}$ to 0.15 so as to match earnings dispersion at age 25. All the estimates used are based on net of taxes household labor income data, consistent with the model specification.

With respect to assets we set an interest rate of 3 percent. We do not determine the interest rate in equilibrium since the model is not meant to capture the behavior of households in the top of the wealth distribution who hold a disproportionate share of total wealth and, hence, are key in determining the equilibrium value of returns. Assets can be held subject to a borrowing constraint. We experiment with the zero borrowing constraint, with the case where agents may borrow up to the natural borrowing limit and also with an alternative formulation where agents can borrow under the constraint that they can repay for sure their debt but the borrowing rate is higher than the lending rate. The borrowing limit is calculated as the present value of future income and social security benefits assuming at each age the lowest possible realization for the income shocks in every remaining period of the working life. While with log-normal shocks this would be zero, in practice it is not because the model is solved by converting the continuous process into a first order discrete Markov chain using an adaptation of the method presented in Tauchen (1986), hence the lowest possible realization of earnings becomes strictly positive.

We model social security benefits so as to mimic the actual US system. In order to do that, we need to compute the average gross earnings over the lifetime of the agent and then to apply a formula that converts that average into a gross pension benefit. The formula for the US that we apply assigns a 90 percent replacement ratio for earnings up to 18 percent of average, a 32 percent replacement ratio from this bend point to the next one, set at 110 percent of average earnings, and finally a 15 percent replacement ratio for earnings above 110 percent average earnings. Finally, we scale the benefits up so that the replacement ratio for the average earner is 45 percent.[5]

Given that in our model the earnings process is based on net earnings, while in the US social security system the benefit formula is computed based on average gross earnings, we need to back out gross earnings from our model net earnings. To do that we invert the progressive tax function formula estimated by Gouveia and Strauss (1994) and now widely used in macroeconomics. If we denote the tax function with the letter τ and gross earnings of individual i at time t by $\tilde{Y}_{i,t}$ the cited tax function takes the form:

$$\tau(\tilde{Y}_{i,t}) = \tau^b[\tilde{Y}_{i,t} - (\tilde{Y}_{i,t}^{-\tau^\rho} + \tau^s)^{-\frac{1}{\tau^\rho}}] \tag{11}$$

To attribute values to the parameters of this function, we set $\tau^b = 0.258$ and $\tau^\rho = 0.768$ from the original work of Gouveia and Strauss (1994) and then set τ^s so that the ratio of personal income tax receipts to labor income is about 25 percent like in the US. With the tax function fully defined it is possible to recover gross earnings from net earnings by solving the equation: $\tilde{Y}_{i,t} - \tau(\tilde{Y}_{i,t}) = Y_{i,t}$. The tax function described above is then also used on 85 percent of gross social security benefits to get net benefits, following the US rule.

4. Results

In this section we report the results of the quantitative analysis of the model. First, in Section 4.1 we focus on the baseline version of the model with zero borrowing constraints and the specification and calibration described above. We report the main result and provide some intuition. In the following section we present robustness checks based on alternative specifications of the model that include different forms of the borrowing constraint, of the specification of non-financial income during retirement and of the earnings process. In that section we also consider a version of the baseline model where agents start life with positive wealth. Finally we devote a section to present a discussion on the role of the preference specification and parametrization.

The statistics of interest are the values of the insurance coefficients as described by Kaplan and Violante (2010) and used since then. These coefficients are defined by the formula:

$$\phi^x = 1 - \frac{cov(\Delta c_{i,t}, x_{i,t})}{var(x_{i,t})} \tag{12}$$

where $c_{i,t}$ is the log deviation of consumption for household i at age t from its deterministic life-cycle trend and $x_{i,t}$ is a shock to the labor income process against which we measure insurance, in our model the permanent and temporary shock. Shocks to the earnings process are not directly observable in the data, however BPP showed that under suitable identifying restrictions the variances and covariances needed to evaluate the coefficients in Eq. (12) can be estimated from observable variables alone, that is, current and lagged income and consumption. We apply the methodology in BPP to our model simulated data and refer to these coefficients as “Model BPP”.[6] At the same time having a model at hand we can observe the shocks and compute the coefficients directly from Eq. (12). We refer to these coefficients as “Model true”.

In what follows we concentrate mostly on the permanent shock since this is the one that the standard incomplete market model has a hard time to explain. Given our focus on exploring a solution to the inability of the model to match the data, we will focus on the model counterpart of the empirically estimated coefficients, the ones labeled “Model BPP”. In the last subsection we also briefly comment on the “Model true”coefficients to get an idea of the magnitude of the bias introduced by the estimation.

[5] This step function for the replacement ratios is commonly used and can be found for example in Huggett and Ventura (2000).

[6] We refer the reader to the online appendix for the identification strategy and computations performed.

Table 1
Simulated insurance coefficients.

	P-S	T-S	P-S	T-S	$\%W < 0$	ra	β
Data	0.36	0.95					
	Model BPP		Model True				
Zero borrowing constraint (ZBC)							
Baseline	0.29	0.88	0.40	0.87	0.0	13.5	0.94
DB pensions	0.28	0.88	0.39	0.88	0.0	11.0	0.95
AR(1)	0.32	0.87	0.43	0.86	0.0	11.0	0.94
$\sigma_\xi^2 = 0.02$	0.34	0.87	0.41	0.87	0.0	9.0	0.93
Natural borrowing constraint (NBC)							
Baseline	0.37	0.90	0.47	0.90	0.21	15.0	0.94
DB pensions	0.37	0.91	0.44	0.91	0.19	12.0	0.95
AR(1)	0.43	0.90	0.50	0.90	0.25	12.0	0.94
$\sigma_\xi^2 = 0.02$	0.37	0.89	0.44	0.89	0.16	9.0	0.93
NBC - $r_b > r_l$							
Baseline	0.31	0.88	0.41	0.88	0.06	14.0	0.94
DB pensions	0.31	0.89	0.39	0.89	0.04	11.5	0.95
AR(1)	0.34	0.88	0.43	0.88	0.07	11.0	0.94
$\sigma_\xi^2 = 0.02$	0.35	0.88	0.42	0.88	0.07	9.0	0.93
Non-zero initial wealth (NZIW)							
Baseline with ZBC	0.30	0.88	0.40	0.88	0.0	13.0	0.94

4.1. *Baseline model*

In the current section we present the main results of the research. They refer to the baseline version of the model with zero borrowing constraint, the standard earnings process with permanent plus temporary shocks and the preference parameters calibrated so as to match the profile of wealth accumulation over the working part of the life-cycle in the model with the one in the data. Agents in the current experiment start life with zero wealth. The empirical wealth profiles are based on PSID wealth data for the years 1984, 1989 and 1994, all converted to 2000 dollars. We remove the top 6 percent wealth observations in the data set so that the wealth-to-earnings ratio of the working population is the same in the data that we use for the estimation and in the model. On these data we estimate a pooled cross section with age and cohort dummies.[7] The profiles thus obtained are then re-scaled so as to express them in the same unit as model wealth using average earnings as the re-normalization factor. In order to compare empirical and simulated profiles we use the minimum square distance of model and data wealth over the working part of the life-cycle. We make this restriction on age because shocks are received during working life, hence it is important to have a precise match of wealth during this portion of the life-cycle to make statements about the ability of the model to explain the level of insurance.[8]

Results for this experiment are shown in the first row of the first panel of Table 1. The remaining rows of Table 1 report results for alternative parameterizations of the model and are discussed in the later Section 4.2 titled “extensions”. In each row of the table we report for the different versions of the model the following objects: in the first and second column the model BPP coefficients against permanent and temporary shocks. In the third and fourth column we report the corresponding “model true”coefficients. Finally in the last three columns we report, in the given order, the fraction of agents with negative wealth and the coefficient of relative risk aversion and the subjective discount factor that result from the calibration. The top row just below the table heading reports the insurance coefficients estimated by Blundell et al. (2008) on the data and that represent the coefficients against which to compare the model BPP coefficients. In this and the next section we describe extensively results concerning the “model BPP”coefficients which are directly comparable to the empirical ones and are therefore more relevant for the applied nature of the present research.

What we can see from the first row of the top panel of Table 1 is that for the baseline model with no borrowing the best fit wealth profile is obtained with a risk aversion coefficient of 13.5. The insurance coefficient for the permanent shock is 0.29 which is close to but slightly lower than the empirical value of 0.36. The insurance coefficient for the temporary shock is 0.88, close to the empirical value. The subjective discount factor that allows the model to match the wealth-to-income target is 0.94 basically in line with what is used in macroeconomics.

Next we describe the mechanism that generates our results. We do that with the help of Table 2 which reports the insurance coefficients against permanent shocks for a small selection of values of risk aversion and elasticity of inter-temporal substitution. For each pair it also reports the subjective discount factor that allows the model to comply with the targeted wealth-to-income ratio. The first row shows a standard low risk aversion expected utility parametrization. As it was known since the work of Kaplan and Violante (2010) the insurance coefficient in this case falls quite short of the data, in fact it is only 0.09, compared to the empirical

[7] We also estimated a second equation with age and time dummies but the profiles turned out to be very similar to the ones of the estimation with cohort dummies, hence we do not report separate results for this case. The internet appendix reports a brief description of the wealth profiles estimation.

[8] This restriction is also made in the structural estimates in Gourinchas and Parker (2002) and Cagetti (2003).

Table 2
Permanent insurance coefficients by risk aversion and eis: selected cases.

	Model BPP	Model true	β
ra = 2, eis = 0.5	0.09	0.26	0.98
ra = 10, eis = 0.1	0.16	0.29	0.81
ra = 10, eis = 0.5	0.19	0.31	0.93

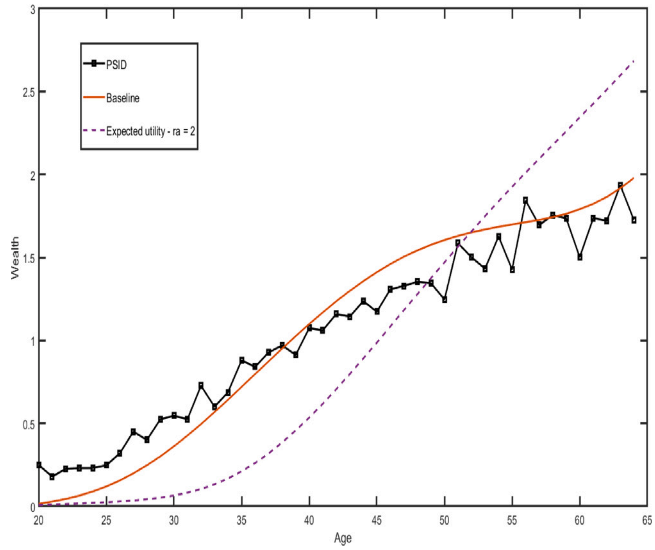

Fig. 1. Life-cycle profiles of wealth.

value of 0.36. As risk aversion increases, households dislike more consumption fluctuations, hence they will save a larger proportion out of positive shocks to use those savings to insulate consumption from negative earnings shocks. This generates both an increase in savings early in the life-cycle and an increase in the observed insurance coefficients. This effect is common to the model with expected utility and to the model with Epstein–Zin preferences. This can be seen in the second and third row of the table where it is shown that raising risk aversion from 2 to 10 increases the insurance coefficients to 0.16 in the expected utility case and to 0.19 in the Epstein–Zin case with unchanged elasticity of inter-temporal substitution. Where the two preference specifications differ is with respect to savings late in the working life. With expected utility, raising risk aversion implies reducing the elasticity of inter-temporal substitution. A lower elasticity of inter-temporal substitution though, leads to higher saving in mid-life because the agents want a flatter consumption profile, hence they need more wealth for the retirement period.

As a consequence, the extra wealth accumulation is more limited in the Epstein–Zin case than in the expected utility case and the subjective discount factor needs not be reduced as much to support the raise in risk aversion and still match the average wealth-to-income target.[9] As the table shows, in fact, the increase of the coefficient of relative risk aversion to 10 leads to a subjective discount factor of 0.81 in the expected utility case but of only 0.93 if we keep constant the elasticity of inter-temporal substitution. Moreover, as it can be seen by comparing the third and second row, the insurance coefficient for the permanent shock increases when the elasticity of inter-temporal substitution increases even in the absence of any increase in risk-aversion. The intuition is similar. With a higher elasticity of inter-temporal substitution the household is willing to accept a more downward sloping consumption profile during retirement, hence it will save less out of late working age income. Given the constant wealth-to-income ratio required by the calibration, this allows the model to accept a higher value of β. In turn, this raises savings early in life. Savings early in life is also increased because the tension between anticipating consumption in the face of an upward sloping earnings profile and delaying it to accumulate precautionary wealth is more easily resolved in favor of the latter if the agent is sufficiently elastic. In summary, having the possibility to keep the elasticity of inter-temporal substitution high implies that wealth is more effectively reshuffled from mid-life, when there is more than enough to insure shocks, to early life when insurance is poor.

Fig. 1 shows this point. The figure reports three lines: the line with square markers represents the data wealth over the working part of the life cycle, the continuous line does the same for the fully calibrated model while the dashed line represents wealth in the model with standard low risk aversion expected utility that only matches the average wealth-to-income ratio. Wealth is expressed

[9] In the appendix we report results for a wider array of values of risk aversion and the elasticity of inter-temporal substitution, showing that while with expected utility a rise of risk aversion to 20 would require a subjective discount factor of 0.53 and still would show insurance coefficients for the permanent shock 0.15 points below the empirical counterpart, a model with an elasticity of inter-temporal substitution of 1.25 can match the insurance coefficient against permanent shocks with risk aversion of 20 and a corresponding subjective discount factor above 0.93.

Table 3
Insurance coefficients by age groups (Permanent shock - “Model BPP”).

Age group	27–31	57–61
ra = 2, eis = 0.5	−0.44	0.63
ra = 10, eis = 0.1	−0.25	0.63
ra = 10, eis = 0.5	−0.13	0.62
ra = 13.5, eis = 1.1	0.12	0.59
ra = 13, eis = 1.1, NZIW	0.13	0.59

in model units. There are two messages that are conveyed by the figure. First, supporting the intuition described above, the fully calibrated model with Epstein–Zin preferences shows substantially more wealth accumulation in the early part of the life-cycle compared to the model with expected utility, while in the last portion of working life the reverse holds true. Second the baseline model with best calibration does indeed do a good job at matching the wealth accumulation profile over working life. The model wealth profile slightly underestimate the profile in the data up to about age 30 and then follows it closely until retirement except for some small point deviations for few years in mid working life.

4.1.1. *Consumption insurance and other consumption statistics over the life-cycle*

The consequences of this wealth redistribution over working life for the insurance coefficients are explored in Table 3, where we report the insurance coefficients of the permanent shocks for two sample age groups for the same parameterizations represented in Table 2 and for the baseline case. The reported coefficients are the “model BPP” ones. As it can be seen, the redistribution of wealth towards the early part of the working life raises the insurance coefficients at younger ages which increases from a low of −0.44 in the standard low risk aversion expected utility environment to −0.25 in the expected utility high risk aversion case and to 0.12 in the baseline, best fit model featuring a risk aversion of 13.5 and eis of 1.1. On the other hand the coefficient late in the working life remains confined in a narrow range between approximately 0.59 and 0.63. The above analysis, beside providing insights into why Epstein–Zin preferences allow the model to get closer to matching empirical coefficients then also points to another benefit of this choice. In fact, according to Blundell et al. (2008) estimates, the insurance coefficients for permanent shocks do not show any clear trend with age. The analysis conducted in this paper shows that as we move towards parameterizations that increase early life wealth accumulation we take a step in the correct direction by flattening the life-cycle profile of the coefficients.

To further check the validity of the approach based on matching the age profile of wealth accumulation we report in Fig. 2 the pattern of the variance of log consumption over the life-cycle. The continuous line represents the pattern of the given variable in the model with expected utility and risk aversion of 2, while the dashed line refers to the baseline model. We can see that the two models generate a very similar rise in consumption inequality over the life-cycle with an increase of between 0.21 and 0.24. The increase is slightly larger for the baseline model, which on top exhibits an initially convex profile. This convex part is due to the fact that in the baseline model the agents are more risk averse and elastic hence at young ages they will save more out of positive earnings shocks, which compresses the initial consumption inequality. The two profiles are then basically overlapping after age 35. For both models, the increase in the variance of log consumption exceeds the one in the data as estimated for example in Heathcote et al. (2010). At the same time the cited paper reports a profile that at the beginning of the life-cycle shows a convex stretch. In this respect then the baseline model moves closer to the data at least qualitatively.

Fig. 3 reports the life-cycle profiles of average consumption for the baseline and the low risk aversion expected utility case. The life-cycle consumption profiles show the familiar hump-shape and the excursion in the level of consumption for the low risk aversion expected utility case is in line with the one reported in Kaplan and Violante (2010). Not surprisingly in our baseline model the life cycle profile of consumption for the baseline model tracks the one of earnings less closely than the one for the expected utility low risk aversion case. This is a consequence of increased precautionary savings early in life in the baseline model. In the data, as reported for example in Fernández-Villaverde and Krueger (2003), the excursion in consumption expenditures is less pronounced, moreover the peak is reached somewhat before age 50. In this respect the baseline model, which produces a slightly earlier peak in consumption moves in the correct direction at least qualitatively, compared to the standard low risk aversion expected utility model. Overall the evidence presented in this subsection suggests that the baseline model, calibrated to match the whole working life profile of wealth accumulation on top of average wealth, not only matches a significant part of the insurance coefficients estimated in the data but also better fits a number consumption statistics over the life cycle at least qualitatively.

4.2. *Extensions*

In this section we present results of versions of the model that differ from the baseline to check how the results are robust to at least some alternative settings. In particular we report results for a model with defined benefit pensions on top of social security, a model that assumes that the persistent component of earnings follows an AR(1) process with the parameters taken from Guvenen (2009) and a model with increased variance of the shock to the permanent component of the earnings process. For all of the three versions we present both results with the zero borrowing constraint and for the model with the natural borrow limit. We also add an alternative scenario where borrowing is allowed but with a wedge between the borrowing and the lending rate. Finally, we consider a version of the baseline model with no borrowing where agents start life with positive wealth. We report both the “model BPP”and “model true”coefficients but the discussion will focus on the former.

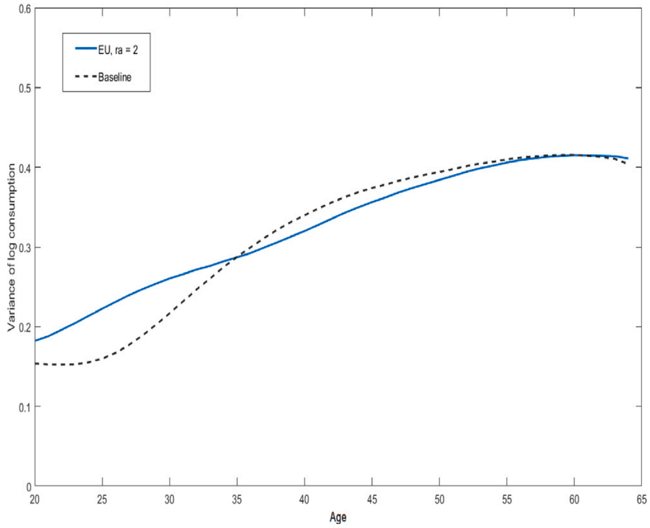

Fig. 2. Life-cycle profiles of variance of log consumption.

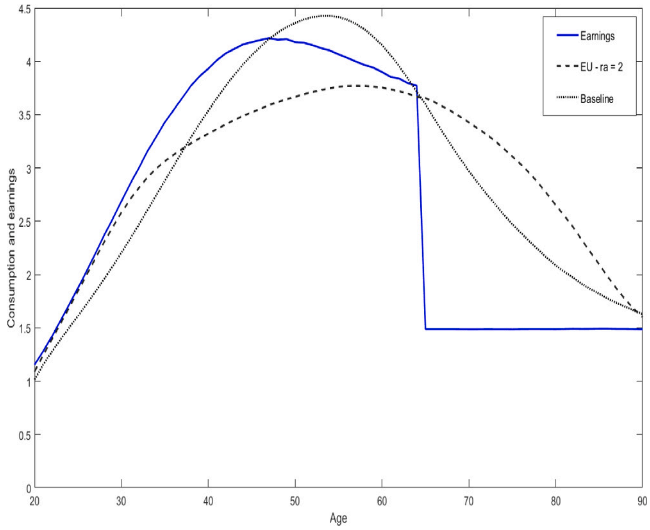

Fig. 3. Life-cycle profiles of average consumption.

Before discussing the results we describe the procedure used to calibrate defined benefit pensions. The procedure uses data reported in Scholz et al. (2006). The authors report data on median earnings and on median defined benefit wealth by deciles of the life-time earnings distribution in their sample from the Health and Retirement Study. Using our average past earnings distribution at retirement age we similarly partition it into deciles. We then attribute to each cell a pension benefit such that the ratio of its expected present value at retirement to median earnings in the model matches the data in the above mentioned paper. This calibration is a simplification for several reasons. First, in partitioning agents at retirement, the concept of average past earnings although very similar is not exactly the same as that of present value of earnings.[10] Second, the only uncertainty about whether an agent will be assigned a defined pension and its level is related to the unfolding of the earnings realizations over the life-cycle. In reality agents may move through different jobs that may or may not offer defined benefit pension plans independently of the earnings shock. Our

[10] The two may differ because of the distribution of shocks over the life-cycle, however the correlation of the two measures is very high.

approach though, beside avoiding the computational burden of adding a further state variable with potentially as many realizations as there are working years, allows us to capture the median replacement ratio for defined benefit pensions and the fact that since these replacement ratios are increasing in lifetime earnings they tend to undo the insurance element intrinsic to social security.[11]

Results are reported in Table 1. The first row of the second sub-panel of the table shows that when borrowing is allowed in the form of the natural borrowing limit the estimated insurance coefficient rises to 0.37 which slightly overestimates the empirical value. This is obtained with a value of risk aversion of 15. The increase in the insurance coefficients is directly explained by the fact that when debt is allowed agents can run wealth into negative territory to insure shocks. At the same time a similar level of wealth in the early part of the life-cycle must be kept in order to match the empirical wealth profile, something that the model achieves through a higher level of risk aversion. The percentage of agents with negative wealth is 21, somewhat higher than the values reported by Huggett (1996).[12]

Interestingly, when defined benefit pensions are introduced there is no improvement in the insurance coefficient generated by the model. The minimizing value of risk aversion falls to 11 for the model with no borrowing and to 12 for the model with the natural borrowing limit. The subjective discount factor rises to 0.95. The interpretation is that with defined benefit pensions the effective replacement ratio of income at retirement would increase, leading to lower savings on average, hence to a higher discount factor to match the average wealth-to-income ratio. However, this also leads to relatively higher wealth accumulation early in life. Since the whole profile of wealth is constrained the calibration reduces risk aversion so as to reduce precautionary savings, which takes places early in the life-cycle, to compensate. This in turn implies that the introduction of defined benefit pensions is virtually neutral with respect to the value taken by the insurance coefficients. This stands in contrast with what one would expect from the fact that a higher position in a payment that depends on the whole history of earnings should improve insurance.[13]

Next we discuss the model that uses the persistent rather than permanent earnings process. To complete the calibration we need to fix the autocorrelation coefficient in the AR(1) process for earnings. Guvenen (2009) estimated two versions of this process that he calls restricted (RIP) and heterogeneous (HIP) income profile. Given that we do not assume heterogeneity in the deterministic part of the income process we use his estimates of the RIP model which comes out at 0.988. The resulting insurance coefficients against the persistent shocks are higher than in the baseline model.[14] In the version with the no borrowing constraint it is 0.32 and in the case with the natural borrowing constraint it is 0.43, even higher than the empirical target. This is consistent with the fact that shocks that have lower persistence are easier to insure. Since the value of the autocorrelation coefficient for the persistent shock is in this case 0.988, this result also shows that even a modest deviation from fully permanent shocks would allow the SIM model to closely match the insurance coefficients for the persistent/permanent component of the shocks in the data. The risk aversion coefficients that minimize the distance between the model and data life-cycle wealth pattern are 11 and 12 respectively and in the case of the model with the natural borrowing limit, the fraction of agents with negative wealth is 25 percent.

The case in which the variance of the shock to the permanent component of earnings is doubled is reported in the last line of each panel. Under this scenario and the zero borrowing constraint the insurance coefficient against the permanent shocks increases to 0.34 from 0.29 in the baseline case. Higher variance of the shocks to permanent earnings translates into higher volatility of consumption so that both the numerator and denominator of the insurance coefficient formula increase. However, with more risky labor income households will tend to save more for precautionary reasons leading to better ex-post insurance. This mechanism is also reflected in the fact that to match the wealth accumulation displayed by the data a lower risk aversion is needed, that is a coefficient of 9 is obtained versus 13.5 in the baseline case. In the case with the natural borrowing limit the insurance coefficients against the permanent earnings shocks is 0.37 and the fraction of agents with negative wealth is 16 percent, in line with the upper bound estimated on the data by the work of Huggett and Ventura (2000).

The third panel of Table 1 reports the results for the version of the model with the alternative borrowing constraint. More specifically in this case borrowing is still allowed subject to the restriction that the household must be able to repay debt for sure, but it is assumed that the borrowing rate is higher than the lending rate. This effectively reduces the maximum amount that can be borrowed and also reduces the incentive to take on debt. As for the calibration the lending rate is left at the baseline value while the borrowing rate is set at 8 percent, a value taken from Davis et al. (2008). Not surprisingly the results are intermediate between the zero and the pure natural borrowing limit. The insurance coefficient against permanent shocks is 0.31 in the baseline model and in the model with defined benefit pensions, 0.34 in the model with the earnings process estimated by Guvenen (2009) and 0.35 in the model with higher variance of the shock to the permanent component of earnings. The associated values of risk aversion are 14 in the baseline and 11.5 in the model with defined benefit pensions, 11 in the model with persistent shocks and 9 in the model with high variance of permanent shocks. Also the subjective discount factor turns out to be very similar to the other specifications

[11] Based on Scholz et al. (2006) data in fact, our pensions are zero in the bottom three deciles of the average past earnings distribution and then they show a monotonically increasing replacement ratio in the remaining ones. For a more detailed modeling of defined benefit pensions one can see Zhou and MacGee (2014).

[12] Huggett (1996) reports two values for the fraction of households with negative wealth: 5.8 and 15 percent. The smaller figure refers to a measure of net worth that includes certain durable goods like cars, while the larger figure does not include them.

[13] The result that increasing the replacement ratio increases the insurance coefficient was obtained by Kaplan and Violante (2010), using only the average wealth-to-income ratio as a calibration target. These authors considered an increase in the replacement ratio that is proportional at all levels of life-cycle earnings rather than a defined benefit pension scheme.

[14] As it was explained in Kaplan and Violante (2010), when shocks are persistent rather than permanent an additional source of bias is introduced in the estimates of the insurance coefficients. However applying the BPP procedure to model simulated data introduces the same bias that is introduced in the data if the data are actually generated by an AR(1) process, hence applying the BPP estimated coefficients to model simulated data and comparing them to the empirical one is still correct.

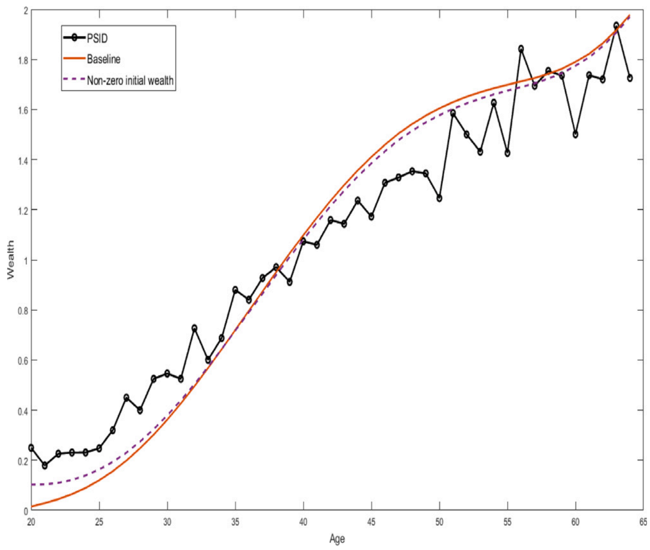

Fig. 4. Life-cycle profiles of wealth.

of the debt limit. The main difference between this and the pure natural borrowing limit version of the model concerns the fraction of agents with negative wealth. In the baseline model this amounts to 6 percent, in the model with defined benefit pensions it is 4 percent and in the model with Guvenen (2009) earnings process and the one with higher variance to the permanent component of the earnings shock it is somewhat higher, standing at 7 percent. These values are close to the lowest of the two estimates provided in Huggett and Ventura (2000).

The last line of Tables 1 and 3 report the results for the baseline model with zero borrowing constraint when we assume that agents start life with positive wealth. For the purpose of calibration we assume that agents receive a certain amount of wealth upon entering the labor market that follows a log-normal distribution and we take the mean and standard deviation of that distribution from the PSID data for households aged 20 to 25. The last line of Table 1 shows that the model BPP insurance coefficient against permanent shocks slightly increases to 0.30 and that the value of risk aversion that best matches the empirical profile of wealth accumulation is 13.[15] The last line of Table 3 shows that the increase in the overall coefficient is obtained through a slight increase in insurance in the early part of the life-cycle, since the coefficient for the age group 27–31 increases to 0.13. We can thus say that adding some positive initial wealth moves the model results in the correct direction although quantitatively the effects are small. There are two reasons to explain this: first the amount of initial wealth must be small to be consistent with the empirical level. Second agents face a steeply increasing earnings profile early in the life-cycle hence they will tend to quickly consume the received wealth. As Fig. 4 shows, this implies that the wealth profile with positive initial wealth will converge to the one without in just about five years leaving little room for increasing the insurance coefficient.

Taken together the results in this section suggest that given the empirically observed pattern of wealth accumulation over working life, a standard incomplete market model comes near the estimates of insurance coefficients provided in Blundell et al. (2008): the baseline model with the zero borrowing constraint can explain 81 percent of the consumption smoothing observed in the data, while the model with the natural borrowing limit can explain virtually all of it. A model where agents must repay debt for sure but with a substantially higher borrowing cost would still allow the model to match 86 percent of the empirically observed consumption smoothing. Raising the variance of the shocks or reducing slightly the persistence to levels consistent with available estimates for AR(1) processes further aligns the model output with the data in the zero borrowing constraint case and in the case with a wedge between the borrowing and lending rate.

4.3. Discussion

In this section we discuss several issues related to the model. First we briefly describe the results concerning the true insurance coefficients computed from simulated data and labeled in the tables “Model true”. Second we explore the role of the elasticity of intertemporal substitution and related to that of the assumption of Epstein–Zin preferences. Finally we briefly discuss the consistency of the values of preference parameters used in the current study with respect to the literature.

[15] Indeed the minimum squared distance decreases as well as it could be expected given that in the data wealth is small but positive even at the very beginning of working life, so that adding some initial wealth makes it easier to match the wealth profile in this early part of the life-cycle.

Table 4
Insurance coefficients: Alternative preference specification.

	P-S	T-S	P-S	T-S	ra	β
Data	0.36	0.95				
	Model BPP		Model True			
Zero borrowing constraint (ZBC)						
Baseline (eis = 1.1)	0.29	0.88	0.40	0.87	13.5	0.94
eis = 1.7	0.30	0.88	0.40	0.87	10.5	0.95
eis = 0.5	0.28	0.87	0.40	0.87	19.0	0.90
Expected utility	0.27	0.87	0.40	0.87	28.0	0.33

4.3.1. Model true coefficients

Given the applied nature of the current work we have so far devoted our attention to the exposition of the coefficients obtained by applying BPP methodology to the simulated data. In a simulated model though the researcher has access to the true earnings shocks, hence can compute the true insurance coefficients as well and by comparing the two gauge the bias introduced by the failure of the model to comply with the conditions for identification required by the BPP identification strategy. It was already shown by Kaplan and Violante (2010) that when agents are close to the borrowing constraint the covariance between consumption and two periods lagged permanent shocks is non zero leading to a bias in the BPP estimator. In our model the use of Epstein–Zin preferences may introduce another source of bias since the BPP procedure is derived from the assumption of expected utility.

Looking at Tables 1 and 2 we can first observe that there is virtually no difference between the “model BPP”and the “model true”coefficients for the temporary shock. As far as the permanent shocks are concerned we observe a substantial difference between the two coefficients: the estimated one is about 0.17 points below the true one in the low risk aversion case of the baseline specification of the model. This difference declines to 0.14 in the expected utility case with risk aversion of 10 and to 0.11 in the baseline case. In the model with expect utility and risk aversion set so that it can match the life-cycle wealth profile we can see from Table 4 that the difference between the two coefficients is 0.13. These results are in line with the findings in Kaplan and Violante (2010): the conditions for the estimator to be unbiased are mainly violated in the case of the permanent shock due to the failure of the correlation between consumption and lagged shock to be zero that occurs near the borrowing constraint. For this reason the bias in the permanent shock coefficients is attenuated for parameterizations that lead to faster accumulation of wealth early in life, hence a less frequently binding borrowing constraint. The fact that in our model with Epstein–Zin preferences the patterns of the bias highlighted in Kaplan and Violante (2010) are confirmed and that the numerical values found are similar suggest that Epstein–Zin preferences do not add in a significant way to the bias. A more rigorous assessment of the bias introduced by Epstein–Zin preferences would require deriving an analogous to the BPP coefficients for this kind of preference and then using it to assess how the omission of the extra terms introduced by these preferences would impact the insurance coefficients. This, however is not a trivial task and as such falls outside the scope of the current research.

4.3.2. The role of the eis and of Epstein–Zin preferences

In the main calibration exercise we fixed endogenously the subjective discount factor and the coefficient of relative risk-aversion but picked exogenously, based on existing literature, the value of the elasticity of intertemporal substitution. Unfortunately available estimates for the latter vary wildly. In this section we check how results would have changed if we had used different values for the eis. At the same time and related to that we check what happens if we calibrate endogenously the coefficient of relative risk aversion and the subjective discount factor according to the strategy adopted in the main analysis but in the context of standard expected utility preferences. The latter experiment allows us to gauge if there is any special role played by Epstein–Zin preferences or if the results in the main section were entirely driven by matching the life cycle profile of wealth accumulation. For sake of clarity we report these results only for the zero borrowing constraint case and the baseline version of the model.[16]

Results are reported in Table 4 and Fig. 5. Looking first to the role of the elasticity of intertemporal substitution, the first three rows of the main panel of the table show that as the elasticity of intertemporal substitution rises three things happen. First a lower risk aversion is needed to obtain the best match with the empirical life-cycle wealth profiles: this value is 13.5 in the baseline case, rises to 19 when the eis is reduced to 0.5 and decreases to 10.5 when the eis increases to 1.7. Second the calibrated value of the subjective discount factor decreases, from 0.95 when the eis is 1.7 to 0.9 when it is 0.5. Third and most importantly the insurance coefficients are not much affected. The greatest changes are observed for the Model BPP coefficients against permanent shocks. Even in this case though differences across values of the eis are small, ranging from 0.28 in the case of an eis of 0.5 to 0.30 when the eis is 1.7. Fig. 5 also shows that the life-cycle profiles of wealth for the three cases are very close to each other with only a trivially improved ability to get close to the empirical one when moving from an eis of 0.5 to an eis of 1.7. This is shown by the comparison of the position of the dashed line, representing the latter case with the position of the line with markers representing the former, compared to the data line. Summarizing, a higher elasticity of intertemporal substitution can be traded off for lower risk aversion. This neither affects the ability of the model to match the empirical wealth profiles, nor leads to any substantial difference in the corresponding values of the insurance coefficients.

[16] Results for the other specifications of the model are similar and are available from the authors upon request.

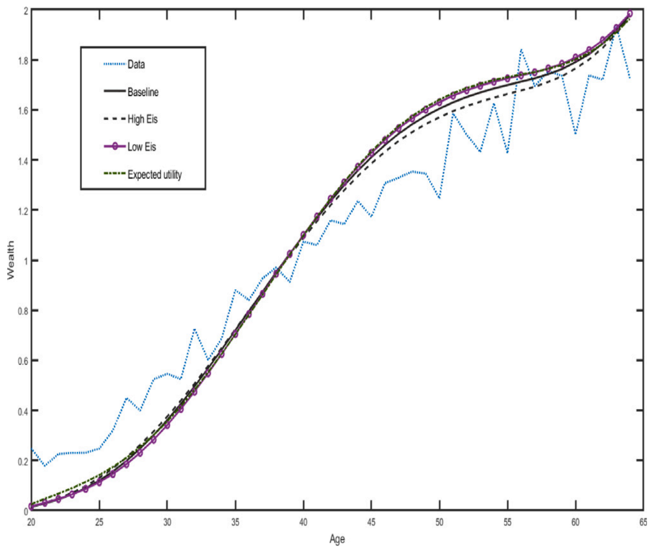

Fig. 5. Life-cycle profiles of wealth.

The last row of Table 4 shows the results for the baseline specification of the model with zero borrowing constraints and using expected utility rather than Epstein–Zin preferences. In this case the minimum distance between the model and data wealth profiles over the working life is reached for a value of risk aversion of 28, corresponding to an eis of 0.0357. The value of the subjective discount factor would then be equal to 0.33. As for the insurance coefficients compared to our baseline they are almost unchanged. For example the largest change observed is a reduction of 0.02 for the model BPP coefficient against permanent shocks. This change would raise to 0.03 if one accepts an eis of 1.7 but would fall to a puny 0.01 if one prefers a lower value of the eis of 0.5. A look at the dashed dot line in Fig. 5 also shows that the ability to match the data of life-cycle wealth profiles is only marginally affected by the use of expected utility rather than Epstein–Zin preferences. Overall this discussion shows that the improved ability to match insurance coefficients delivered by our model is almost entirely due to the fact that by a suitable choice of preference parameters we match the life-cycle profile of wealth, with the added contribution of Epstein–Zin preferences being negligible. This said one should not forget that in the case of expected utility the value of all preference parameters would fall quite outside what is normally accepted in macroeconomics and at a low value of 0.33 this is especially true for the subjective discount factor.

4.3.3. Comments on the values of preference parameters

The remaining part of the discussion concerns preference parameters, that is, the coefficient of relative risk aversion and the elasticity of inter-temporal substitution. With respect to risk aversion, minimum distances between model and data wealth profiles are attained with values that range from 9 to 15. These values are somewhat higher than what is normally assumed in macroeconomic models, but it must be said that the key reason for assuming a low risk aversion is that a reasonable behavior of macroeconomic quantities hinges upon a relatively high elasticity of inter-temporal substitution which under expected utility is linked to the former by an inverse relationship. Such a link is not present in the case of Epstein–Zin preferences. Estimates for risk aversion in an Epstein–Zin setting presented by Vissing-Jørgensen and Attanasio (2003) suggest that a risk aversion of up to 10 for the presumably more risk tolerant stock holders is plausible. If on the other hand one looks at the experimental evidence for example, Barsky et al. (1997) find that about two thirds of their sample shows a risk aversion coefficient of 15, with the rest of the sample equally split between risk aversions of 7, 6 and 4. With respect to the elasticity of inter-temporal substitution, microeconomic estimates vary substantially. For example, using British data (Attanasio and Weber, 1993) find values between 0.3 and 0.7. It is also true that the values tend to increase with wealth: for example Vissing-Jørgensen and Attanasio (2003) estimate an interval ranging from 1 to 1.4 for the population of stockholders, which is wealthier than average. Some estimates are even higher: Gruber (2013) using data from the CEX and exploiting exogenous cross individual differences in after tax real interest rates, finds a value above 2 although admittedly with large standard errors. Overall then, the values of risk aversion and the elasticity of inter-temporal substitution that are used in our calibration to find the best match between the model working life profile of wealth with the data fall within the limits of the available empirical evidence.

5. Conclusions

In this research we revisited the ability of the SIM model to explain the extent of consumption smoothing observed in the data. We focussed specifically on the role played by a careful specification of wealth accumulation over the working part of the life-cycle. Using the insurance coefficients estimated by Blundell et al. (2008) as a benchmark measure for consumption smoothing we found that a standard SIM model that is parameterized to match the working life profile of wealth accumulation can explain 81 percent of the value of the BPP insurance coefficients against permanent earnings shocks in the zero borrowing constraint case. A similar model with expected utility and benchmark low risk aversion parametrization can explain 25 percent of those coefficients. In the case of the natural borrowing limit our model can explain 100 percent of the empirical coefficients. We obtained the match of the working life wealth accumulation profiles by using Epstein–Zin preferences with parameters that fall within the empirical and experimental evidence. Since matching the life-cycle wealth profile implies redistributing wealth from mid-age to young age when compared with standard models, we can conclude that the failure of the baseline model with standard expected utility and benchmark parameterizations to match the empirical insurance coefficient for permanent shocks reflects in an important way its under-prediction of wealth accumulation early in the working-life. Also the increase in wealth accumulation early in life leads to a flatter profile of the insurance coefficients against permanent shocks with respect to age. In this case the result while being qualitatively more consistent with the evidence it still falls short of it quantitatively.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jmacro.2023.103566.

References

Attanasio, O.P., Davis, S.J., 1996. Relative wage movements and the distribution of consumption. J. Polit. Econ. 104 (6), 1227–1262.
Attanasio, O.P., Weber, G., 1993. Consumption growth, the interest rate and aggregation. Rev. Econ. Stud. 60 (3), 631–649.
Barsky, R.B., Juster, F.T., Kimball, M.S., Shapiro, M.D., 1997. Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study. Q. J. Econ. 112 (2), 537–579.
Blundell, R., Pistaferri, L., Preston, I., 2008. Consumption inequality and partial insurance. Amer. Econ. Rev. 98 (5), 1887–1921.
Blundell, R., Pistaferri, L., Saporta-Eksten, I., 2016. Consumption inequality and family labor supply. Amer. Econ. Rev. 106 (2), 387–435.
Bosworth, B., Anders, S., 2008. Saving and Wealth Accumulation in the PSID, 1984-2005. Working Paper 2, Center for Retirement Research (Boston College).
Cagetti, M., 2003. Wealth accumulation over the life cycle and precautionary savings. J. Bus. Econom. Statist. 21 (3), 339–353.
Cerletti, E., Pijoan-Mas, J., 2012. Durable Goods, Borrowing Constraints and Consumption Insurance. Discussion Paper DP9035, Centre for Economic Policy Research.
Christelis, G., Georgarakos, G., Jappelli, T., Pistaferri, L., Van Rooij, M., 2019. Asymmetric consumption effects of transitory income shocks. Econom. J. 129 (622), 2322–2341.
Cocco, J.F., Gomes, F.J., Maenhout, P.J., 2005. Consumption and portfolio choice over the life cycle. Rev. Financ. Stud. 18 (2), 491–533.
Davis, S.J., Kubler, F., Willen, P., 2008. Borrowing costs and the demand for equity over the life-cycle. Rev. Econ. Stat. 88 (2), 348–362.
De Nardi, M., Fella, G., Paz-Pardo, G., 2020. Non linear household earnings dynamics, self-insurance, and welfare. J. Eur. Econ. Assoc. 18 (2), 890–926.
Díaz, A., Pijoan-Mas, J., Ríos-Rull, J.V., 2003. Precautionary savings and wealth distribution under habit formation preferences. J. Monetary Econ. 50 (6), 1257–1291.
Dynarski, S., Gruber, J., 1997. Can families smooth variable earnings? Brookings Pap. Econ. Activity 28 (1), 229–303.
Fernández-Villaverde, J., Krueger, D., 2003. Consumption over the life cycle: Facts from the consumer expenditure survey data. Rev. Econ. Stat. 89 (3), 552–565.
Gourinchas, P.-O., Parker, J.A., 2002. Consumption over the life cycle. Econometrica 70 (1), 47–89.
Gouveia, M., Strauss, R.P., 1994. Effective federal individual income tax functions: An exploratory empirical analysis. Natl. Tax J. 47 (2), 317–339.
Gruber, J., 2013. A tax-based estimate of the elasticity of intertemporal substitution. Q. J. Finance 03 (1), 1350001–1–1350001–20.
Guvenen, F., 2009. An empirical investigation of labor income processes. Rev. Econ. Dyn. 12 (1), 58–79.
Heathcote, J., Perri, F., Perri, F., 2010. Unequal we stand: An empirical analysis of income inequality in the United States, 1967–2006. Rev. Econ. Dyn. 13, 15–51.
Hryshko, D., Manovskii, I., 2017. How much consumption insurance in the U.S.? Unpublished manuscript.
Huggett, M., 1996. Wealth distribution in life-cycle economies. J. Monetary Econ. 38 (3), 469–494.
Huggett, M., Ventura, G., 2000. Understanding why high income households save more than low income households. J. Monetary Econ. 45 (2), 361–397.
Kaplan, G., Violante, G.L., 2010. How much consumption insurance beyond self-insurance? Am. Econ. J.: Macroecon. 2 (4), 53–87.
Karahan, F., Ozkan, S., 2013. On the persistence of income shocks over the life cycle: Evidence, theory and implications. Rev. Econ. Dyn. 16 (3), 452–476.
Krueger, D., Perri, F., 2006. Does income inequality lead to consumption inequality? Evidence and theory. Rev. Econom. Stud. 73 (1), 163–193.
Scholz, J.K., Seshadri, A., Khitatrakun, S., 2006. Are Americans saving optimally for retirement? J. Polit. Econ. 114 (4), 607–643.
Tauchen, G., 1986. Finite state Markov-chain approximations to univariate and vector autoregressions. Econ. Lett. 20 (2), 177–181.
Vissing-Jørgensen, A., Attanasio, O.P., 2003. Stock-market participation, intertemporal substitution, and risk-aversion. Am. Econ. Rev. 93 (2), 383–391.
Wu, C., Krueger, D., 2018. How Much Consumption Insurance in Bewley Models with Endogenous Family Labor Supply? Working Paper 24472, National Bureau of Economic Research.
Zhou, J., MacGee, J., 2014. Private pensions, retirement wealth and lifetime earnings. In: 2014 Meeting Papers 181. Society for Economic Dynamics.