Life-cycle wealth accumulation and consumption insurance

(Article begins on next page)

14 October 2023

Life-cycle Wealth Accumulation and Consumption Insurance

Claudio Campanale[*]
Università di Torino
CeRP (Collegio Carlo Alberto)

Marcello Sartarelli[†]
Universidad de Alcalá

July 9, 2019

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 86 and more than 97 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.

Keywords: precautionary savings, Epstein-Zin, consumption insurance coefficients, life-cycle

JEL Codes: D15, E21

---

[*]Corresponding author: Claudio Campanale, Dipartimento ESOMAS, Università di Torino, Cso. Unione Sovietica 218bis, 10127, Torino, Italy. Contact email: claudiogiovanni.campanale@unito.it

[†]Departamento de Economía, Universidad de Alcalá, Plaza de la Victoria 2, 28801 Alcalá de Henares, Spain. Contact email: marcellosartarelli@gmail.com

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 by focussing specifically on the role that life-cycle wealth accumulation plays in determining the degree of consumption smoothing. To preview the results, we find that when the model is calibrated so that it matches the whole empirical profile of wealth accumulation over the working life, a version of the model with the tightest borrowing constraint can match up to 86 percent of the Blundell *et al.* (2008) estimates of insurance against permanent earnings shocks. When the model is solved under the natural borrowing limit it can virtually match the empirical values.

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 con-

currently 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 of consumption away from young ages reshuffles wealth 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. As an example, a calibration with an EIS of 0.8 allows us to match the insurance coefficient for permanent shocks of 0.36 with a risk aversion of 20 while still keeping the discount factor above 0.9 and matching the empirical average wealth-to-income ratio. By contrast, a comparable model with expected utility and risk aversion of 20 would require a discount factor of 0.545 to match a realistic average wealth-to-income ratio and would still fall significantly short of the insurance coefficient estimated in the data.

Given that the main mechanism that allows the model to generate insurance coefficients that are in line with the data is the redistribution of wealth across different periods of the life-cycle, it is important to verify that the resulting life-cycle profiles of wealth match the data. For this reason we solve a preferred calibration of the model where the coefficients of risk aversion and the inter-temporal elasticity of substitution are chosen so as to minimize the distance between the model and data wealth accumulation profiles during working life. In this case we find that 86 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. This figure raises to over 97 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.

Quantitative and empirically relevant studies of life-cycle consumption behaviour 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 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 *et al.*, 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.[1] 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

---

[1]In a very 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 date has been applied only to temporary shocks.

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. Among those, we can cite Cerletti and Pijoan-Mas (2012) who extended the model to explore the role of non-durable goods and the adjustment in the consumption bundle that this allows and Karahan and Ozkan (2013) who estimated an earnings process featuring age-varying persistence and showed that this improves the life-cycle profiles of insurance coefficients of an otherwise standard model. Finally, more recently a parallel line of research that uses wages rather than earnings as primitives and studies the extent of insurance against wage shocks has developed. In this line of research Blundell *et al.* (2016) provided benchmark empirical estimates and Wu and Krueger (2018) developed a first quantitative study to test an extension of the SIM model, featuring households with double earners, against those estimates.

The present paper is most closely related to Kaplan and Violante (2010) in that it also constructs a quantitative SIM model to study its implications for the insurance coefficients using earnings as the primitive shocks. There are three main differences between our study and theirs. First, while Kaplan and Violante (2010) offers a broad theoretic-quantitative analysis of the problem, our research is focused on trying to explain the gap between model and data estimates of insurance coefficients of permanent shocks, that has proven harder to bridge. In this respect we obtain the very important result that indeed a simple SIM model when appropriately parameterized displays empirically plausible insurance properties. Second and key to pursuing that goal we emphasize the role that the whole pattern of wealth accumulation over the working life plays in determining the model-generated insurance coefficients, rather than simply constraining the average wealth-to-income ratio to its empirical counterpart. This is crucial since the total insurance coefficient is an average of the coefficients by age groups which in turn strongly depend on the groups' wealth, hence using average wealth over the whole population alone is not adequate to study this problem. Finally we adopt Epstein-Zin rather than standard expected utility preferences since this gives us more flexibility in matching the wealth profile. With respect to the issue of matching working life profiles of wealth and indirectly of consumption as well, the paper is also

4

related to Cagetti (2003) that builds a life-cycle model and uses a simulated method of moments to estimate preference parameters based on the whole life-cycle wealth profiles obtained from the Survey of Consumer Finance and Gourinchas and Parker (2002) that structurally estimate preferences parameters based on the profile of consumption over the working life. However differently from their work, our values for the preference parameters come from calibration rather than full structural estimation.

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.

2 Model

We consider a standard life-cycle economy where agents are endowed with Epstein-Zin preferences. 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.

5

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 $1-\alpha$.[2] 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

---

[2]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.

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_{z0}^2. The shocks ε_{it} and ζ_{it} are normally distributed with mean 0 and variances σ_ε^2 and σ_ζ^2, 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.

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 first perform a set of experiments for different

8

values of the elasticity of inter-temporal substitution and risk aversion. We then move to a set of experiments where we search for the values of risk aversion and the elasticity of inter-temporal substitution that minimize the distance between the model and data wealth profile over the working life. In each case 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.[3] 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.

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 the coefficients estimated by Cocco *et al.* (2005). As for the stochastic component of earnings, we have to assign three parameters, that is, the variance of the permanent and temporary shocks ζ and ε and the initial variance of the permanent shock $\sigma_{z_0}^2$. We give σ_ζ^2 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_ε^2 based on the point estimate by Blundell *et al.* (2008). Finally, we set $\sigma_{z_0}^2$ to 0.15 so as to match earnings dispersion at age 25.

With respect to assets we set an interest rate of 3.5 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.

---

[3]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.

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.[4]

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:

$$\tau(\tilde{Y}_{i,t}) = \tau^b[\tilde{Y}_{i,t} - (\tilde{Y}_{i,t}^{-\tau^\rho} + \tau^s)^{-\frac{1}{\tau^\rho}}] \tag{9}$$

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 order to uncover the mechanics of the model we perform

---

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

an extensive exploration of the parameter space. Initially we specialize the Epstein-Zin preferences to the usual expected utility case by setting $\alpha=\gamma$. In this case we consider values of risk aversion of 2, 5, 10, 15 and 20. Then for each risk aversion case we solve again the model for values of the elasticity of inter-temporal substitution of 0.5, 0.8 and 1.25. We report the results both for the model with a zero borrowing constraint and for the opposite case of the natural borrowing limit. Second, in light of the lessons learned with this analysis we report the results of the model solved under our preferred calibration where preference parameters are chosen so as to minimize the distance between the profile of asset accumulation during the working part of the life-cycle in the model and in the data.

We report values of the insurance coefficients of both the permanent and the temporary shock, although our focus will be on the former given the finding of Kaplan and Violante (2010) that these are the ones 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.[5] Before moving to the actual description of the results in the next subsection, we will briefly describe how the insurance coefficients are defined and computed.

4.1 BPP insurance coefficients

Let $y_{i,t}$ and $c_{i,t}$ be the log deviation of net labor income and consumption from their respective deterministic life-cycle trend. In general, the deviation of income can be the result of different shocks that we can generically denote $x_{i,t}$. We can define the insurance coefficient for shock $x_{i,t}$ as:

$$\phi^x = 1 - \frac{cov(\Delta c_{i,t}, x_{i,t})}{var(x_{i,t})} \tag{10}$$

If the received shock translated one-to-one into a change in consumption ϕ^x would be equal to 0, while in the opposite case where consumption did

---

[5]Having a model at hand, one can also compute the true insurance coefficients and use them to study the magnitude of the bias of the BPP estimator, however this is outside the scope of the present research. See Kaplan and Violante (2010) for a discussion of the source of the estimation bias and when it is most likely to be greater.

not react at all to the shock the index would be equal to 1. The index then is a measure of the proportion of the shock that is not translated into consumption growth and, hence, is a measure of the extent to which shocks are insured, with a higher value corresponding to better insurance. Having data on both consumption and the shock, as it happens in a model simulation, one can directly compute the true value of the index. Alternatively, given the earnings process described in the previous sections and used in much of the quantitative macroeconomics literature the coefficients can be estimated from income and consumption data alone provided the following two identifying restrictions are assumed:

$$cov(\Delta c_{i,t}, \zeta_{i,t+1}) = cov(\Delta c_{i,t}, \varepsilon_{i,t+1}) = 0 \qquad (11)$$

and

$$cov(\Delta c_{i,t}, \zeta_{i,t-1}) = cov(\Delta c_{i,t}, \varepsilon_{i,t-2}) = 0 \qquad (12)$$

The two assumptions state that consumption growth can be correlated neither with future nor past shocks.[6] Under these assumptions it can be shown that:

$$- cov(\Delta y_{i,t}, \Delta y_{i,t+1}) = var(\varepsilon_{i,t}) \qquad (13)$$

$$- cov(\Delta c_{i,t}, \Delta y_{i,t+1}) = cov(\Delta c_{i,t}, \varepsilon_{i,t}) \qquad (14)$$

which allows the econometrician to identify ϕ^ε and

$$- cov(\Delta y_{i,t}, \Delta y_{i,t-1} + \Delta y_{i,t} + \Delta y_{i,t+1}) = var(\zeta_{i,t}) \qquad (15)$$

$$- cov(\Delta c_{i,t}, \Delta y_{i,t-1} + \Delta y_{i,t} + \Delta y_{i,t+1}) = cov(\Delta c_{i,t}, \zeta_{i,t}) \qquad (16)$$

that identifies ϕ^ζ, the insurance coefficients for the permanent shock. Failure of the given assumptions leads to biased estimates that can be assessed when using model simulated data.[7]

---

[6]Given this identifying assumption, Epstein-Zin preferences are a possible source of bias, however assuming that the empirical data are generated by household that have these preferences, the application of the BPP estimator to both data and model would introduce the same kind of bias making the comparison legitimate.

[7]The introduction in the text is a basic description of the parameters of interest that we compute in the model experiments. For a thorough introduction to the estimation of the insurance coefficients see Blundell *et al.* (2008) and Kaplan and Violante (2010) upon which our treatment of the issue is based.

Table 1: Insurance coefficients by risk aversion (Expected utility)

	Permanent shock	Transitory shock	β
Data	0.36	0.95	
ra = 2	0.167	0.885	0.975
ra = 5	0.170	0.892	0.936
ra = 10	0.212	0.904	0.822
ra = 15	0.238	0.902	0.675
ra = 20	0.266	0.900	0.545

4.2 The expected utility case

We report in Table 1 the insurance coefficients obtained from the model simulation when preferences are assumed to be of the standard expected utility form and the coefficient of relative risk aversion ranges from 2 to 20. Looking at the first row, we see that the empirical estimates, taken from Blundell *et al.* (2008), are 0.36 for permanent shocks and 0.95 for transitory shocks. In the case of transitory shocks, the estimates based on simulated data are always very close to their data counterparts, ranging from 0.885 in the case of risk aversion of 2, to about 0.90 for risk aversion of 10 or more. The picture changes radically as far as the insurance coefficients for permanent shocks are concerned. The estimated coefficient is 0.167 for risk aversion of 2 and it rises up to 0.266 for risk aversion of 20. While this increase is substantial, it still leaves the coefficient 0.1 points below the empirical value. The reason for the increase is that more risk-averse agents dislike consumption volatility more and, hence, in the face of positive shocks they will save a larger part of them in order to finance consumption when the bad shock hits. This increase in savings also has another effect that can be seen in the last column of Table 1: in order for the model to still match the targeted level of the wealth-to-income ratio, it is necessary to reduce the subjective discount factor from 0.975 to 0.545, clearly a value that is well outside of what is accepted in macroeconomic practice.

Table 2: Insurance coefficients by risk aversion (Epstein-Zin)

	Permanent shock	Transitory shock	β
Data	0.36	0.95	
EIS = 0.5			
ra = 2	0.167	0.885	0.975
ra = 5	0.186	0.887	0.947
ra = 10	0.267	0.902	0.925
ra = 15	0.314	0.897	0.91
ra = 20	0.340	0.889	0.897

4.3   The Epstein-Zin Case

It is well known that Epstein-Zin preferences allow the model to fully disentangle risk aversion from inter-temporal substitution. In this subsection we thus exploit this increased freedom in choosing preference parameters to check if it is possible to improve the ability of the model to match the empirical insurance coefficients. We proceed in two steps. First, in Table 2 we keep the elasticity of inter-temporal substitution fixed at 0.5 and consider the usual values of risk aversion in the range 2 to 20. As can be seen in the second column, the estimated insurance coefficient for the permanent shock raises from 0.167 when risk aversion is 2 to 0.34 when risk aversion is 20. The latter value is already quite close to 0.36, the measure found in the data. As it can be seen in the last column of Table 2, this can be obtained with a substantially smaller decrease in the value of the subjective discount factor. The value of β that is needed to match to targeted wealth-to-income ratio is 0.975 when risk aversion is 2 and declines only to 0.897 when risk aversion is 20, a value that while still smaller, it is not far from what is accepted and used in macroeconomic modelling.

In the second step, we alternatively proceed by fixing risk aversion at a value of 10 and checking how results change when the elasticity of inter-temporal substitution is raised from the corresponding expected utility value of 0.1 to 1.25. Looking at the second column in Table 3, we can notice that even for constant risk aversion an increase in the elasticity of inter-temporal substitution brings about a substantial increase in the esti-

Table 3: Insurance coefficients by EIS (Baseline)

	Permanent shock	Transitory shock	β
Data	0.36	0.95	
ra=10			
EIS = 0.1	0.212	0.904	0.822
EIS = 0.5	0.267	0.902	0.925
EIS = 0.8	0.287	0.901	0.935
EIS = 1.25	0.317	0.899	0.942

mated coefficient for permanent shocks, from 0.212 to 0.317. Also, looking at the last column of Table 3 we see that this is obtained with a contemporaneous increase in the required value of the subjective discount factor, from 0.822 when the elasticity of inter-temporal substitution is 0.1 to 0.942 when it is 1.25. The latter is a value that is already close to the values used in macroeconomic modelling.

Finally, in Table 4 we put together the insights obtained in the previous analysis and consider a broad range of parameter values including for risk aversion the values of 2, 5, 10, 15 and 20 and for each of them setting the elasticity of inter-temporal substitution at 0.5, 0.8 and 1.25. What the table shows is that for certain combinations of the elasticity of inter-temporal substitution and risk aversion the model goes a long way towards rationalizing the observed empirical values of the insurance coefficients of permanent shocks. For example, for the combination of EIS of 0.8 and risk aversion of 20 the insurance coefficient is 0.366 and the targeted wealth-to-income ratio is obtained for β set to 0.917, while if we are willing to accept a value of the EIS of 1.25 we can get an insurance coefficient of 0.369 for risk aversion equal to 15 and a subjective discount factor of 0.934.

In order to briefly conclude this section, we also want to point out at the results concerning the estimated insurance coefficients for temporary shocks. With the exception of the cases with risk aversion set to 2 and with an elasticity of inter-temporal substitution of 0.8 or 1.25, where the insurance coefficients declines to 0.863 and 0.790 respectively, the insurance coefficients for temporary shocks remain in the narrow range between 0.89 and 0.90, values that are very close to the empirical estimates.

Table 4: Insurance coefficients by risk aversion and EIS (Baseline)

	Permanent shock	Transitory shock	β
Data	0.36	0.95	
EIS = 0.5			
ra = 2	0.167	0.885	0.975
ra = 5	0.186	0.887	0.947
ra = 10	0.267	0.902	0.925
ra = 15	0.314	0.897	0.91
ra = 20	0.340	0.889	0.897
EIS = 0.8			
ra = 2	0.135	0.863	0.963
ra = 5	0.201	0.890	0.950
ra = 10	0.287	0.901	0.935
ra = 15	0.337	0.894	0.925
ra = 20	0.366	0.887	0.917
EIS = 1.25			
ra = 2	0.027	0.790	0.951
ra = 5	0.212	0.890	0.950
ra = 10	0.317	0.899	0.942
ra = 15	0.369	0.891	0.934
ra = 20	0.394	0.881	0.928

4.4 Interpretation

In this section we describe the mechanism that generates our results. Models like the one considered here exhibit precautionary savings, which is the dominant factor for accumulating wealth in the initial part of the life-cycle.[8] 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. As a consequence, we observe both an increase in savings early in the life-cycle and an increase in the observed insurance coefficients. This effect is common both to the model with expected utility and to the model with

---

[8]See Gourinchas and Parker (2002).

Epstein-Zin preferences. The increase in wealth accumulation for precautionary reasons, given the calibration constraint on the wealth-to-income ratio, implies the need to reduce the value of the subjective discount factor. 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, which is connected to the former by an inverse relationship. 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 to unrealistically low values to support a substantial raise in risk aversion and still match the average wealth-to-income target.

Moreover, as it can be seen in Table 3, 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 late in life, 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 reshuffled from mid-life, when there is more than enough to insure shocks, to early life when insurance is poor.

This is confirmed by Figure 1, which reports the life-cycle profile of wealth for parameterizations of the model with a constant risk aversion, set equal to 10, and increasing values of the elasticity of inter-temporal substitution, ranging from 0.1 (expected utility case) up to 1.25. Figure 1 shows that the age profile of wealth during working life is convex shaped in the expected utility case. When the elasticity of inter-temporal substitu-

17

Figure 1: Life-cycle profiles of wealth.

tion is progressively raised, it changes to a convex-concave shape that gives it a distinct hunchbacked profile. In the expected utility case wealth after 20 model periods is roughly one third of peak wealth around retirement age, while in the model with an elasticity of substitution of 1.25 it is about 53 percent of peak wealth. The consequences for the insurance coefficients are explored in Table 5, where we report the insurance coefficients of the permanent shocks for the same parameterizations represented in Figure 1. As it can be seen, the substantial reduction in peak wealth caused by the increase in the elasticity of inter-temporal substitution barely affects insurance coefficients near the end of the working life, which remain confined in a narrow range between 0.615 and 0.627. On the other hand, the larger wealth accumulation early in life raises the insurance coefficients in a substantial way for the age group between 27 and 31: the coefficient increases from 0.029 when the elasticity of inter-temporal substitution is 0.1 to 0.198 when it is 1.25.

The above analysis, beside providing insights into why Epstein-Zin preferences allow the model to get closer to matching empirical coefficients, also points to another benefit of this choice. In fact, according to Blundell

Table 5: Estimated coefficients by age groups (Permanent shock)

Age group	27-31	57-61
ra=10		
EIS = 0.1	0.029	0.615
EIS = 0.5	0.175	0.627
EIS = 0.8	0.162	0.627
EIS = 1.25	0.198	0.617

et al. (2008) estimates, the insurance coefficients for permanent shocks do not show any trend with age, while as shown in Kaplan and Violante (2010) the standard expected utility model generates strongly increasing and convex insurance coefficients by age.[9] The analysis conducted in this paper though, shows that increasing the elasticity of inter-temporal substitution makes one step in the correct direction by flattening the life-cycle profile of the coefficients. Results are even starker for the case of a risk aversion of 20, the highest value considered here. In that case, combined with an elasticity of inter-temporal substitution of 1.25, we get an estimated insurance coefficient for the permanent shock that is 0.342 for the age group 27 to 31 and 0.595 for the age group 57 to 61.[10]

4.5 The natural borrowing limit case

The models solved so far have assumed that agents cannot borrow. In this subsection we repeat the same experiments that we performed before but in a version of the model where households can borrow, subject only to the constraint that they can repay for sure their debt by the time they reach the maximum possible age, that is, subject to the so called natural borrowing limit.

Results are reported in Table 6 for all the risk aversion and elasticity of inter-temporal substitution values that we considered in the zero borrowing limit case. Looking at the first column of the table, we see that introducing

---

[9]We get their same results when using expected utility with low risk aversion, which we do not report for the sake of brevity.

[10]For the sake of space we do not report the full set of tables with the insurance coefficients by age groups for all parameterizations but they are available upon request.

Table 6: Insurance coefficients for permanent shock (Model with debt)

	Permanent shock	Transitory shock	%$W < 0$	β
EIS $= 0.5$				
ra $= 2$	0.261	0.943	0.274	0.979
ra $= 5$	0.274	0.931	0.212	0.950
ra $= 10$	0.326	0.915	0.147	0.928
ra $= 15$	0.354	0.906	0.108	0.913
ra $= 20$	0.382	0.899	0.115	0.902
EIS $= 0.8$				
ra $= 2$	0.253	0.94	0.342	0.966
ra $= 5$	0.279	0.928	0.194	0.951
ra $= 10$	0.333	0.913	0.097	0.937
ra $= 15$	0.373	0.903	0.112	0.927
ra $= 20$	0.405	0.896	0.124	0.920
EIS $= 1.25$				
ra $= 2$	0.184	0.909	0.551	0.955
ra $= 5$	0.288	0.924	0.175	0.952
ra $= 10$	0.355	0.908	0.098	0.943
ra $= 15$	0.403	0.898	0.120	0.936
ra $= 20$	0.438	0.891	0.139	0.930

debt further increases the estimated insurance coefficients for permanent shocks. As it can be seen in the top panel, in this case it is possible to reach an insurance coefficient of 0.354, already very close to the empirical one, at a value of risk-aversion of 15 in the case where the elasticity of inter-temporal substitution is 0.5. When the elasticity of substitution is raised to 0.8, the empirical value of the coefficient is reached somewhere between risk aversion of 10 and 15, while looking at the bottom panel we see that when the elasticity of inter-temporal substitution is 1.25 a coefficient of risk-aversion of 10 generates an estimated coefficient of 0.355. For this parametrization a value of the subjective discount factor of 0.943 supports a wealth-to-income ratio of 2.5. The insurance coefficients for temporary shocks increase as well, albeit to a lesser extent. Looking at the third column of the table, we can also see that the fraction of agents with negative wealth ranges from a maximum of 0.551 to a minimum of 0.097. For the parameterizations whose associated insurance coefficients

are consistent with the empirical ones though, the fraction of agents with negative wealth is always close to 10 percent. This value falls in the range of 5.8 percent to 15 percent reported by Huggett (1996) from the Survey of Consumer Finances.[11] Overall, we can then say that adding debt further improves the fit of the model to the data on insurance of permanent shocks. The reason is that when households may hold debt they can drive wealth into negative territory when faced with negative shocks, thus improving their insurance opportunities.

4.6 Optimized wealth profiles

The results so far obtained, showed that the flexibility of Epstein-Zin preferences allows the model to match the average estimated insurance coefficient for permanent shocks with values of the subjective discount factor that are in line with what is normally used in economics. The inspection of the mechanism also showed that this is obtained by shifting wealth from pre-retirement age, when agents are already well-insured, to young age when this is not the case. Whether this is a substantial contribution to the related literature depends on the extent to which the shape of the wealth profiles obtained under the parameterizations of the Epstein-Zin preferences that allowed the model to match the insurance coefficients are at the same time the ones that are closer to their empirical counterpart.

In the current section we want to further follow this line of reasoning and search for the preference parameter combinations that minimize the distance between model and empirical life-cycle wealth profiles. We then check the insurance coefficients under these minimum distance parameterizations. The empirical wealth profiles are based on PSID wealth data obtained by averaging the profiles for the years 1984, 1989 and 1994, all converted to 2000 dollars and removing the top 5 percent observations so that they reflect the same subset of the population that the model is meant to capture. 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

---

[11]The smallest figure refers to a measure of net worth that includes durable goods like cars, while the largest figure does not include them.

we use the minimum square distance of model and data wealth over the working part of the life-cycle. We make this restriction 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.[12] Thus if we denote with $W_t(\gamma, \alpha)$ the average normalized wealth of age t agents generated by the model simulation as a function of the two preference parameters and with $\widetilde{W}_t$ the corresponding empirical values the problem we solve is:

$$\min_{\gamma, \alpha} \sqrt{\frac{1}{T^{ret}} \sum_{t=1}^{T^{ret}} (W_t(\gamma, \alpha) - \widetilde{W}_t)^2} \tag{17}$$

where T^{ret} is the retirement age. Based on the results obtained thus far we restrict the search on the interval of risk-aversion coefficients between 6 and 18. We also restrict the elasticity of inter-temporal substitution to values between 0.3 and 1.45. In all cases, as before we use β to match the aggregate wealth-to-income ratio. We solve several versions of the model beside the baseline models considered in the previous sections. In particular we report results also for a model with defined benefit pensions and a model that assumes that the persistent component of earnings follows an AR(1) process with the parameters taken from Guvenen (2009).[13] 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 tighter constraint than the natural borrowing limit.

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 ra-

---

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

[13]We also considered a model where agents start life with non-zero wealth, however the results do not add any new insight so we do not report them.

tio 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.[14] 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 cycle through different jobs that may or may not offer defined benefit pension plans independently of the earnings shock. Our 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.[15]

Results for this experiment are shown in Table 7, where for the different versions of the model, we report the estimated permanent and temporary insurance coefficient in the first and second column respectively, the percentage of agents with negative wealth in the third column and the values of risk aversion and the elasticity of inter-temporal substitution that minimize the distance of model and data wealth over working life in the fourth and fifth columns. Finally, the last column reports the value of the subjective discount factor that allows each model to match the targeted wealth-to-income ratio of 2.5. What we can see from Table 7 is that for the baseline model with no borrowing the best fit wealth profile is obtained for the elasticity of inter-temporal substitution of 1.25 and a risk aversion coefficient of 9.5. The insurance coefficient for the permanent shock is 0.31 which is close to but slightly lower than the empirical value of 0.36. The insurance coefficient for the temporary shock is 0.9, close to the empirical value. The subjective discount factor that allows the model to match

---

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

[15]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 modelling of defined benefit pensions one can see Zhou and MacGee (2014).

Table 7: Estimated Insurance coefficients: Minimum distance parameters

	P-S	T-S	$\%W < 0$	ra	EIS	β
Data	0.36	0.95				
Zero borrowing constraint (ZBC)						
Baseline	0.31	0.90	0.0	9.5	1.25	0.942
DB pensions	0.31	0.91	0.0	7	1.45	0.950
AR(1)	0.35	0.89	0.0	7	1.45	0.945
Natural borrowing constraint (NBC)						
Baseline	0.35	0.91	9.5	9	1.45	0.945
DB pensions	0.35	0.92	8.0	7.5	1.45	0.950
AR(1)	0.41	0.91	12.6	7.5	1.45	0.945
NBC - $r_b > r_l$						
Baseline	0.32	0.90	2.9	10	1.08	0.940
DB pensions	0.32	0.91	1.7	7	1.45	0.951
AR(1)	0.35	0.90	2.1	7	1.45	0.945

the wealth-to-income target is 0.942 basically in line with what is used in macroeconomics. When borrowing is allowed in the form of the natural borrowing limit the estimated insurance coefficient rises to 0.35 which virtually matches the empirical value. This is obtained with a slightly lower value of risk aversion of 9. 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 mix of higher patience and higher elasticity of substitution. The percentage of agents with negative wealth is 9.5, within the values reported by Huggett (1996) and cited in previous sections of this work.

Interestingly, when defined benefit pensions are introduced there is no improvement in the insurance coefficient generated by the model. The minimizing value of the elasticity of inter-temporal substitution is 1.45 while that of risk aversion falls to 7 for the model with no borrowing and to 7.5 for the model with the natural borrowing limit. The subjective discount

factor rises to 0.950. 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 now 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 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.[16]

Finally, in the model that uses the earnings process in Guvenen (2009) the insurance coefficients against permanent shocks are higher than in the baseline model.[17] In the version with the no borrowing constraint it is 0.35 and in the case with the natural borrowing constraint it is 0.41, 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 values of the elasticity of inter-temporal substitution that minimize the distance between the model and data life-cycle wealth pattern is again 1.45 in both the model with and without borrowing. The risk aversion coefficients are 7 and 7.5 respectively and in the case of the model with the natural borrowing

---

[16]We obtain a positive effect of introducing defined benefit pensions when, like in section 4.3 we only use the average wealth-to-income ratio as a target. A similar positive result is obtained by Kaplan and Violante (2010) who also use the average wealth-to-income ratio as a target. These authors though consider an increase in the replacement ratio that is proportional at all levels of life-cycle earnings rather than a defined benefit pension scheme.

[17]As it was explained in Kaplan and Violante (2010), when shocks are persistent 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.

limit, the fraction of agents with negative wealth is 12.6 percent, close to the larger value reported in Huggett (1996).

The third panel of Table 7 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.* (2006). Not surprisingly the results are intermediate between the zero and the pure natural borrowing limit. The insurance coefficient against permanent shocks is 0.32 in the baseline model and in the model with defined benefit pensions and 0.35 in the model with the earnings process estimated by Guvenen (2009). The insurance coefficients against temporary shocks again hover around 0.9. These results are obtained under a value of the elasticity of substitution of 1.08 in the baseline version of the model and 1.45 in the versions with defined benefit pensions and with the persistent rather than permanent shock. The associated values of risk aversion are 10 in the baseline and 7 in the other two versions of the model. Also the subjective discount factor turns out to be very similar to the other specifications 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 2.9 percent, in the model with defined benefit pensions it is 1.7 percent and in the model with Guvenen (2009) earnings process it is 2.1 percent. These values are below their empirical counterpart. It must be said though that the choice of the borrowing rate can be considered as an extreme case given that the figures reported in Davis *et al.* (2006) refer only to unsecured borrowing, while in the data households have access to collateralized borrowing through home equity lines which is in general cheaper.

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 is consistent with the estimates of insurance coefficients provided in Blundell *et al.* (2008): the model with the zero bor-

rowing constraint can explain 86 percent of the consumption smoothing observed in the data, while the model with the natural borrowing limit can explain up to 97.2 percent. A model where agents must repay debt for sure but with a substantially higher borrowing cost would still allow the model to match 90 percent of the empirically observed consumption smoothing.

4.7 Discussion

In this section we discuss some issues related to the properties of the minimum distance surface over the parameter space and about the choice of the latter in relation to the empirical literature on preference parameters. Starting with the first issue, Figure 2 reports the minimum squared distance between model and data wealth over the working life for the baseline model with the zero borrowing constraint for the set of risk aversion and the elasticity of inter-temporal substitution that were searched over.
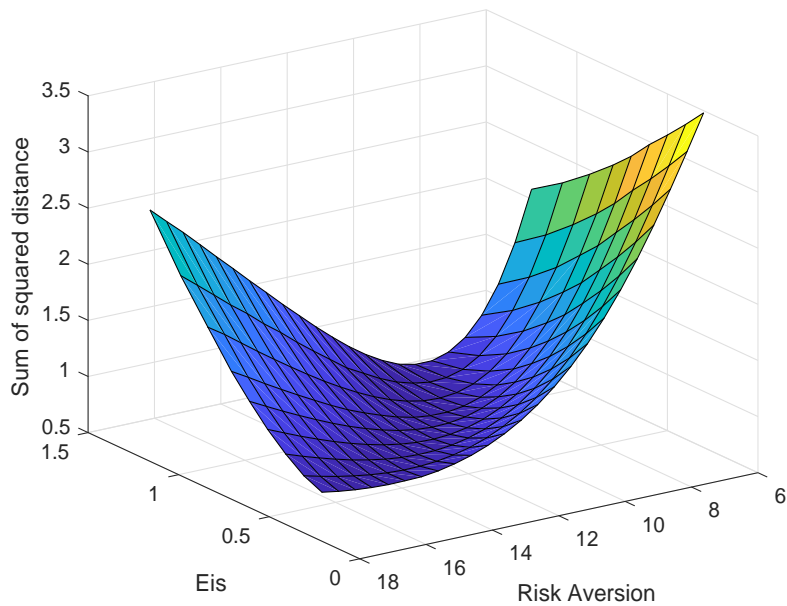

Figure 2: Life-cycle wealth squared distances.

There are three main features of the minimal distance surface in Figure 2 that deserve mention. First, for each value of the elasticity of inter-temporal substitution, the distance metrics shows a U-shaped pattern with

respect to risk aversion. Second, the minimizing value of risk aversion falls as the elasticity of inter-temporal substitution increases, so that minima are reached along a "valley" with a negative slope in the space of the two preference parameters. This reflects the fact that a more elastic agent is more willing to accept variation of consumption over time, hence a lower risk aversion is required to foster early life precautionary savings to levels consistent with the empirical evidence on life-cycle wealth profiles. Third, the bottom of the valley thus obtained is rather flat. This implies that the exact pair of risk aversion and elasticity of substitution that minimizes the distance between model and data working life wealth is poorly identified.

Table 8: Minimum squared distance by EIS (Baseline model)

	ra	P-S	msqd
EIS = 0.3	15.5	0.297	0.776
EIS = 0.4	14.5	0.301	0.767
EIS = 0.5	13.5	0.302	0.766
EIS = 0.6	12.5	0.301	0.765
EIS = 0.7	12.0	0.304	0.764
EIS = 0.8	11.0	0.302	0.760
EIS = 0.93	10.5	0.305	0.756
EIS = 1.08	10.0	0.309	0.753
EIS = 1.25	9.5	0.313	0.752
EIS = 1.45	8.5	0.309	0.753

The values of the insurance coefficients at the bottom of the valley though are very flat as well, hence the exact pair of risk aversion and the elasticity of inter-temporal substitution is not crucial for the result about insurance coefficients. We illustrate all this by way of Table 8. The table reports for each value of the elasticity of inter-temporal substitution, the corresponding value of risk aversion that minimizes the sum of squared deviations between model and data working life wealth in the second column. In the third column it reports the corresponding insurance coefficient against permanent shocks and finally in the last column it reports the value of the mean squared deviations of model and data wealth. This last column illustrates the fact that for any value of the elasticity of substitution it is possible to set risk aversion so that the model profile of wealth accumula-

28

tion matches the empirical one almost as well as in the distance minimizing solution.

For example, when the elasticity of inter-temporal substitution is 0.3 the minimum sum of squared distance is 0.776 which compared to the absolute minimum attained of 0.752 is only 3 percent off, a trivial amount compared to the overall variation of the distance measure that we can observe in Figure 2. The second column shows the trade-off between higher elasticity of inter-temporal substitution and lower risk aversion in matching working life profiles of wealth accumulation. The third column delivers the key message that the model ability to match the insurance coefficients is robust to the exact specification of preference parameters. As it can be seen the insurance coefficient against permanent shocks ranges from 0.297 when the elasticity of inter-temporal substitution is 0.3 to 0.313 in the case where the absolute minimum of the distance function is reached at a value of the elasticity of inter-temporal substitution of 1.25. This excursion is very narrow and represents a variation of between 82.5 and 86.9 percent of the empirical measure of the insurance coefficient against permanent shocks. For this reason the results in terms of the ability of the model to match the insurance coefficients in the data can be seen as quite robust and depending essentially on a correct specification of the profile of wealth accumulation over the working part of the life-cycle: provided the latter condition is satisfied it is not important for which exact pair of preference parameters this is obtained. While we do not report it in the paper for the sake of brevity the same analysis holds true for all the other versions of the model considered in Table 7.

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 7 to 10. 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 between 5 and 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 were searched over to find the best match between the model working life profile of wealth with the data fall within the limits and explore a large part of the region suggested by 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 86 percent of the value of the BPP insurance coefficients against permanent earnings shocks in the zero borrowing constraint case. A similar model with benchmark parameterizations can explain less than 40 percent of those coefficients. In the case of the natural borrowing limit our model can explain up to 97 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 reshuffling 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 mainly reflects 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.

References

ATTANASIO, O. P. and DAVIS, S. J. (1996). Relative Wage Movements and the Distribution of Consumption. *Journal of Political Economy*, **104** (6), 1227–1262.

ATTANASIO, O. P. and WEBER, G. (1993). Consumption Growth, the Interest Rate and Aggregation. *Review of Economic Studies*, **60** (3), 631–649.

BARSKY, R. B., JUSTER, F. T., KIMBALL, M. S. and SHAPIRO, M. D. (1997). Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study. *Quarterly Journal of Economics*, **112** (2), 537–579.

BLUNDELL, R., PISTAFERRI, L. and PRESTON, I. (2008). Consumption Inequality and Partial Insurance. *American Economic Review*, **98** (5), 1887–1921.

BLUNDELL, R., PISTAFERRI, L. and SAPORTA-EKSTEN, I. (2016). Consumption Inequality and Family Labor Supply. *American Economic Review*, **106** (2), 387–435.

BOSWORTH, B. and 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. *Journal of Business and Economic Statistics*, **21** (3), 339–353.

CERLETTI, E. and 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. and VAN ROOIJ, M. (2019). Asymmetric Consumption Effects of Transitory Income Shocks. *Economic Journal*, forthcoming.

COCCO, J. F., GOMES, F. J. and MAENHOUT, P. J. (2005). Consumption and Portfolio Choice over the Life Cycle. *The Review of Financial Studies*, **18** (2), 491–533.

DAVIS, S.J., KUBLER, F. and WILLEN, P. (2008). Borrowing Costs and the Demand for Equity over the Life-Cycle. *The Review of Economics and Statistics*, **88** (2), 348–362.

DÍAZ, A., PIJOAN–MAS, J. and RÍOS–RULL, J.V. (2003). Precautionary Savings and Wealth Distribution under Habit Formation Preferences. *Journal of Monetary Economics*, **50** (6), 1257–1291.

DYNARSKI, S. and GRUBER, J. (1997). Can Families Smooth Variable Earnings? *Brookings Papers on Economic Activity*, **28** (1), 229–303.

GOURINCHAS, P.-O. and PARKER, J. A. (2002). Consumption Over the Life Cycle. *Econometrica*, **70** (1), 47–89.

GOUVEIA, M. and STRAUSS, R. P. (1994). Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis. *National Tax Journal*, **47** (2), 317–339.

GRUBER, J. (2013). A Tax-Based Estimate of the Elasticity of Intertemporal Substitution. *Quarterly Journal of Finance*, **03** (1). 1350001-1 – 1350001-20

GUVENEN, F. (2009). An Empirical Investigation of Labor Income Processes. *Review of Economic Dynamics*, **12** (1), 58–79.

HUGGETT, M. (1996). Wealth Distribution in Life-cycle Economies. *Journal of Monetary Economics*, **38** (3), 469–494.

HUGGETT, M. and VENTURA, G. (2000). Understanding Why High Income Households Save More than Low Income Households *Journal of Monetary Economics*, **45** (2), 361–397.

KAPLAN, G. and VIOLANTE, G. L. (2010). How Much Consumption Insurance beyond Self-Insurance? *American Economic Journal: Macroeconomics*, **2** (4), 53–87.

KARAHAN, F. and OZKAN, S. (2013). On the Persistence of Income Shocks over the Life Cycle: Evidence, Theory and Implications. *Review of Economic Dynamics*, **16** (3), 452–476.

KRUEGER, D. and PERRI, F. (2006). Does Income Inequality Lead to Consumption Inequality? Evidence and Theory. *Review of Economic Studies*, **73** (1), 163–193.

SCHOLZ, J. K., SESHADRI, A. and KHITATRAKUN, S. (2006). Are Americans Saving "Optimally" for Retirement?. *Journal of Political Economy*, **114** (4), 607–643.

VISSING-JØRGENSEN, A. and ATTANASIO, O. P. (2003). Stock-Market Participation, Intertemporal Substitution, and Risk-Aversion. *American Economic Review*, **93** (2), 383–391.

WU, C. and 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. and MACGEE, J. (2014). Private Pensions, Retirement Wealth and Lifetime Earnings. *2014 Meeting Papers 181*, Society for Economic Dynamics.