UNIVERSITÀ
DEGLI STUDI
DI TORINO

ALMA UNIVERSITAS
TAURINENSIS

Founded in 1404

DEPARTMENT OF
ECONOMICS AND STATISTICS
WORKING PAPER SERIES

Quaderni del Dipartimento di Scienze
Economico-Sociali e Matematico-Statistiche

ISSN 2279-7114

CONSUMPTION INSURANCE AND EDUCATION: A PUZZLE?

CLAUDIO CAMPANALE

Working paper No. 69 - September 2020

Consumption insurance and education: A puzzle? *

Claudio Campanale

Università di Torino

CeRP (Collegio Carlo Alberto)

September 22, 2020

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 uncover a related puzzle: households with different educational levels show similar insurance against permanent shocks in the model while in the data empirically estimated by Blundell *et al.* (2008) college educated households seem to smooth consumption much more than high school educated households.

Keywords: precautionary savings, consumption insurance coefficients, life-cycle, education

JEL Codes: E21

*Claudio Campanale, Departement of Economics and Statistics (ESOMAS), Università di Torino , Cso. Unione Sovietica 218 bis, Torino, Italy. Phone: +39 0116706072 ext. 3262. E-mail: claudiogiovanni.campanale@unito.it.

All errors and inconsistencies are my own.

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 on the whole population falls significantly short of matching the empirical values of insurance against permanent earnings shocks, as estimated by Blundell *et al.* (2008). In the present research we show that if we break down the population into educational types and calibrate their respective earnings processes separately a new shortcoming emerges: the insurance coefficients against permanent shocks in the model turn out to be roughly equal across groups while in the estimation provided in Blundell *et al.* (2008) high school graduates seems not to smooth permanent shocks at all while college graduates smooth roughly 60 percent of those shocks.

The model analyzed in the present research is a standard life-cycle model with uninsurable earnings risk and borrowing constraints. Agents receive stochastic income during working life and a fixed pension benefit after retirement. There are two different types of agents in the model that are identified as high school graduates and college graduates. The two agents differ in their income process, both its deterministic component and the properties of the stochastic shock process, in particular the innovation to the permanent and to the temporary shock. Preferences are of the expected utility type with standard low risk aversion coefficient. Given the importance of wealth accumulation in determining the possibility of agents to smooth shocks the subjective discount factor is set so as to match a realistic average wealth to income ratio. Also there is a careful modeling of the pension system that mimics the progressive features of the US benefit scheme. This is very important to our goal since, as it was pointed out by Huggett and Ventura (2000) a progressive pension system helps explain why high permanent income households save more than low permanent income households and the income process for the two types implies that college households have higher permanent incomes than high school households.

We solve and simulate the model and find that in the baseline case the insurance coefficient against permanent shocks are about 0,21 for both groups while according to Blundell *et al.* (2008) these same coefficients are 0.06 for high school graduates and 0.58 for college graduates. Analyzing the results we document that the puzzle arises because in the standard estimation of the earnings process college graduates have a steeper profile at young ages, thus they are more inclined to anticipate consumption and reduce savings in the early years of their career when the low level of wealth makes it more difficult to insure shocks. After having established the result we test several simple possible corrections to the model to address the puzzle. These include the possibility that college graduates have better access to credit, in the sense of having a looser natural borrowing constraint and the possibility that college graduates are more patient, a feature that is suggested by the fact that the baseline calibration of the model falls short of matching the data with respect to the ratio of median wealth of college relative to non college households. The latter feature increases the insurance coefficients for them above those for high school graduates but by far too little to match the data. Further, we introduce defined benefit pensions which increase effective replacement ratios for households with higher life-time incomes and finally we play with other preference parameters making high school graduates less elastic inter-temporally. While both of these features move the model towards the data in a qualitative sense they grossly fail quantitatively. These experiments reinforce the case for a new puzzle in the consumption insurance properties of the SIM model.

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. 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. 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. More recently Campanale and Sartarelli (2018) showed that if the whole profile of wealth accumulation during working life is matched, rather than only the average wealth to income ratio the prototype SIM model can explain a large part of the insurance coefficients observed in the data. Finally, 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 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 featuring 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. Agents are divided into two types that are distinct based on their endowment process meant to mimic households headed by a college and high school graduate respectively. A cohort of agents is simulated and the model-generated patterns of consumption insurance are studied. In what follows for the sake of simplifying the notation we describe the model for a general household, being implicit that the household optimization problem must be solved twice, one for each educational type. This can be done without confusion since the problem of the two type of households only differs in the calibrated values of the earnings process and not in the equations that describe it.

2.1 Demographics and preferences

Time is discrete with model periods of one year length. The model is populated by a continuum of households. 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.

Most of the simulations will be carried out using expected utility preferences, however for the sake of generality and since some experiments will be based on Epstein-Zin preferences we use the latter to write the description of the individual decision problem. 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 $1 - \alpha$.[1] 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 of a given educational level 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:

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

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 η_{it} are normally distributed with mean

[1] As it is well known we can obtain the standard expected utility formulation by setting $\gamma = \alpha$.

0 and variances σ_ε^2 and σ_η^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.[2] With respect to preference parameters we set risk aversion to 2 and the elasticity of inter-temporal substitution to 0.5, thus considering the standard case of expected utility. We set β, the subjective discount factor, so that the average wealth-to-income ratio is equal to 2.5. While at first sight 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

[2]The database does not provide separate survival probabilities by education hence we work with common values for both types even though it is reasonable to expect that in reality those probabilities may differ.

[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.

estimated by Cocco *et al.* (2005) separately for households with a high school or college graduate head. As for the stochastic component of earnings, for each educational type we have to assign three parameters, that is, the variance of the permanent and temporary shocks η and ε and the initial variance of the permanent shock $\sigma^2_{z_0}$. We give σ^2_η a value of 0.018 for high school and a value of 0.023 for college graduates and we assign σ^2_ε the value of 0.0318 for high school and 0.0434 for college graduates. All the given values are taken from the point estimates in Blundell *et al.* (2008). Finally we set $\sigma^2_{z_0}$ to 0.15 so as to approximately match earnings dispersion at age 25.[4]

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 both with the zero borrowing constraint and 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.

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 next one, set at 110 percent of average earnings, and finally a 15 percent replacement ratio for earnings above 110 percent average earnings.[5] Finally, we scale the benefits up so that the replacement ratio for the average earner is 45 percent.

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

[4]We assign the same value of the initial dispersion of earnings to both groups and match the initial dispersion of earnings in the overall population. While one could assign a different value to each of the two groups, in practice this parameter has a negligible impact on the object of our study.

[5]This formula is by now a standard in comparable quantitative macro literature and was first used by Huggett and Ventura (2000).

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 follow Kaplan and Violante (2010) and 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.

4 Results

In this section we report the results of the quantitative analysis of the model. In the first subsection we document the puzzle that arises when insurance coefficients are computed separately by educational type and briefly interpret the mechanism that gives rise to the puzzle. In the following subsections we propose some simple modifications to the baseline model to try to solve the puzzle. The first modifications that we propose include different access to credit by the two groups and a more careful matching of wealth differences, obtained through different degrees of patience by educational types. Next we turn to the introduction of defined benefit pensions which effectively increase the replacement ratios for higher life-time income families and finally we consider more preference heterogeneity. As we will see, the proposed extensions to the baseline model all move the model results in the correct direction with respect to the data but they all fall short by a substantial amount quantitatively. There are two shocks to the earnings process in the model: permanent and temporary. The results reported below will focus only on the permanent shock. The reason is that the SIM model does a very good job at reproducing the insurance coefficient for temporary shocks. This was known for the general population since the work of Kaplan and Violante (2010) and remains true when separate consideration is given to groups with different educational levels.[6]

4.1 The puzzle

We set the puzzle by way of table 1. The first panel of the table reports as a benchmark the results from the model when it is assumed that there is a single type of agent in the economy. This agents in particular features a deterministic component and variances of income shocks that are based on population wide estimates. The insurance coefficient against permanent shocks generated by the model for the whole population is in this case equal to 0.208, which falls short of the value of 0.36 estimated in BPP. This inability of the model to match the insurance coefficients under a standard low risk aversion expected utility calibration is well known since the work of Kaplan and Violante (2010). What we want to stress here though is another shortcoming of the model. This is done in the second panel of the table which reports the results for the model when agents are divided into college and non-college graduates, labelled "high school", each endowed by a specific earnings process that differ both in its deterministic component and in the shock process. In this case the insurance coefficient for the overall population remains almost constant and takes a value of 0.213. When we look at the insurance coefficients by educational type we see that they are almost identical: the one for college graduates is 0.214 and the one for high school graduates is 0.207. If we compare these values with the empirical ones, reported in the top line of the table, we see that in the data the insurance coefficient against permanent shocks for college graduates is 0.58, showing that college graduates can attain a large degree of consumption smoothing, while the coefficient for high school graduates is only 0.06, showing almost no smoothing at all. This is a further failure of the standard model. The last panel of the table reports results of an alternative experiment where the deterministic component of the earnings profile is different by education but the shocks embodied in the stochastic process are restricted to have the same variance, equal to the one estimated on the whole population. Under this assumptions it is even the case that the insurance coefficient against permanent shocks, at a value of 0.195, is lower for college graduates than for high school graduates for whom it is 0.224. The interpretation of the results can be explained with the help of figure 1 which reports the simulated average life-cycle profile of earnings for the two groups. The

[6]Results concerning temporary shocks are nonetheless available upon request.

Table 1: Simulated insurance coefficients: Baseline ZBC

	All	College	High School
Data (P-S)	0.36	0.58	0.06
Model: Common			
Model IP	0.208		
β	0.967		
$\%W < 0$	0.0		
W/Y	2.55		
Model: Heterogeneous			
Model IP	0.213	0.214	0.207
β		0.967	0.967
$\%W < 0$	0.0	0.0	0.0
W/Y	2.55	2.87	2.28
Model: Heterogeneous - Restricted			
Model IP	0.224	0.195	0.224
β		0.966	0.966
$\%W < 0$	0.0	0.0	0.0
W/Y	2.55	2.58	2.37

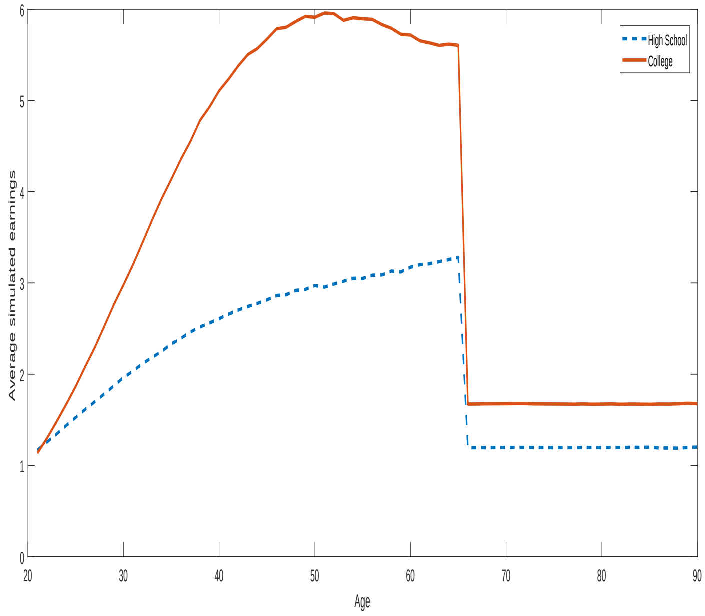

Figure 1: Life-cycle earnings profiles by education.

two profiles start at the same level but as it is well known the one for college graduates peaks at a substantially higher level, hence it is steeper. For this reason college graduates have a stronger incentive to anticipate consumption and everything else equal would save less early in the life cycle. This explains why the insurance coefficient is lower for this group. Once the variance of the shocks is allowed to be different, the result is overturned, albeit to a trivial amount. The reason is that the insurance coefficient is a relative metric since it is normalized by the variance of the shock itself, hence there are no reasons to expect that a more volatile shock negatively affects the coefficient. However a higher variance of the shock makes the earnings process more risky, inducing more savings, hence better possibilities to self-insure. The variance of the permanent shock is higher for the college graduates, hence when we move from the common variance specification to the separately calibrated volatility, the insurance coefficient increases for them. This said the amount of the variation of the insurance coefficient across the two experiments is small compared to their deviation from their empirical counterpart.

4.2 Debt and patience

We next propose two possible amendments to the model to check if they can help realigning it to the data. The results are presented in table 2. We first check the possibility that allowing for debt might solve the problem. We assume that households can take debt with the constraint that they are able to repay it even in case they face a string of worst possible shocks for the remaining part of the life-cycle, that is, the so called natural borrowing limit. However we assume that there is a wedge between lending and borrowing rates, with the latter fixed at 8 percent based on the work of Davis *et al.* (2006). In principle this might lead to higher coefficients of insurance for college graduates than for high school graduates, since the former have higher worst case realizations of earnings shocks and hence will have a looser constraint on borrowing.[7] More borrowing in turn promotes better insurance. As the first panel of the table shows it turns out that this is not the case. The insurance coefficient is 0.226 for high school graduates and 0.215 for college graduates, hence in this case the consumption smoothing is even stronger for the high school group although the difference is again quantitatively minor. Notice also that this formulation of the borrowing constraint does at least a reasonable job at capturing households' behavior with respect to debt, since it generates a percentage of agents with negative wealth of 16.7, close to the values reported in Huggett (1996) which range from 5 to 15 percent depending on whether cars are included or not in the definition of net worth.[8]

As a second possible explanation we fix our attention on wealth. As it was described before the model features a realistic progressive social security system which as it was pointed out in Huggett and Ventura (2000) is one important factor in determining why high income households save proportionally more than low income household. The results in table 1 and in the first panel of 2 show this clearly: the wealth-to-income ratio is around 2.9 for the college graduates and it is only between 2.2 and 2.3 for high school graduates. However a

[7]Strictly speaking shocks are log-normal hence their lowest possible realization is zero and so should be the natural borrowing limit. A negative limit then arises because the discretization of the stochastic processes moves the lowest realization away from zero. The criterion used to turn the two continuous processes for earnings discrete is the same for both, insuring consistency in defining the natural borrowing limit for the two groups.

[8]This justifies our choice to add a wedge between borrowing and lending rates, whose absence would have led to a grossly overestimated fraction of agents with negative wealth.

Table 2: Simulated insurance coefficients: Debt and patience

	All	College	High School
Data (P-S)	0.36	0.58	0.06
Model: baseline NBC			
Model IP	0.225	0.215	0.226
β	0.968		
$\%W < 0$	0.167	0.14	0.18
W/Y	2.55	2.89	2.22
Model: Heterogeneous β			
Model IP	0.207	0.231	0.179
β		0.972	0.962
$\%W < 0$	0.0	0.0	0.0
W/Y	2.55	3.16	1.99

closer comparison to the data shows that quantitatively this mechanism does not generate a sufficient difference in wealth accumulation by the two groups. For example, based on the Survey of Consumer Finances, Bucks *et al.* (2006) and Kennickell *et al.* (2000), find that the ratio of median net worth of households with a college degree to that of households with a high school degree varied between 2 in 1995 and 3.67 in 2001, the year of the dot.com bubble, averaging 2.76 over the surveys conducted between 1983 and 2004. The ratio of college to high school median wealth in the baseline model with zero borrowing constraint is 2.04 almost coinciding with the lowest values reported in the surveys. Hence we next propose a solution by which college and non college households are characterized by different subjective discount factors with the college households being more patient. This assumption may be justified on the ground that a more patient agent will have a larger propensity to invest, as young in human capital, that is, education, and later in capital through savings. Knowles and Postlewaite (2004) construct a simple model along these lines and find some empirical support for this idea. Alternatively it has been proposed that education may foster more patience. Perez-Arce (2017) tests this idea on a survey of college students in Mexico City and finds a positive result. For calibration we choose the difference in discount factors so that the ratio of median wealth among college graduates is 3 times that of high school graduates, somewhat larger than the average value for the period 1983 to 2004 reported above. We make this choice to give a stronger chance to the model to match the data. As the second panel of table 2 shows this still falls largely short of what is needed to close the gap between the model and the data. The table shows that we need a difference in the subjective discount factor of 0.01 to match the ratio of median wealth of the two groups. The insurance coefficient against permanent shocks raises to 0.231 for college households and drops to 0.179 for high school households, hence a significant difference in favor of the former emerges. Even so, the insurance coefficient for college graduates is substantially lower than the one in the data while the reverse holds true for high school graduates.

4.3 Defined benefit pensions

The social security system provides a key insurance mechanism and this is more so the higher the replacement ratio. Since replacement ratios are declining in life-time income the social security system is more effective at providing insurance to the poorer agents in the economy hence, if we split the population into high school and college graduate households it will tend to reduce the insurance coefficients for the latter group relative to the former. For their nature pension payments in the form of defined benefit plans may play a similar insurance role. However in general pension payments are more concentrated towards the higher end of the earnings distribution. As a consequence while social security provides a mechanism that tends to reduce the differences in insurance opportunities across the distribution of life-time earnings private pensions operate in the opposite direction. In this section then we check if the existence of defined benefit pensions may give an important contribution towards aligning the insurance coefficients by education groups that the model generates to the one in the data.

In order to calibrate pensions we use 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.[9] 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, avoids the computational burden of adding a further state variable with potentially as many realizations as there are working years, while still allowing us to capture the median replacement ratio for defined benefit pensions and the fact that since these are increasing in lifetime earnings they tend to undo the progressive insurance element intrinsic in social security.[10]

We perform two experiments. In the first one we simply solve the model with defined benefit pensions assuming a common discount factor for both type of agents, in the second one we again allow college graduates to be more patient and calibrate the two subjective discount factors so that the model can match the ratio of median wealth of the two groups in the data. Results are reported in table 3. The table is divided into two panels. The top panel shows the results of the model solved under the assumption of equal β in the two groups, the bottom panel reports the results for the model with heterogeneous subjective discount factors. The top panel shows that simply adding defined benefit pensions leads to the counterfactual result that the insurance coefficients is higher for the high school graduates. This result is at first sight counter-intuitive since replacement ratios for defined benefit pensions are higher for households with higher life-time incomes that in turn are more concentrated among the college graduates. This apparent contradiction though is easily explained by taking into account that higher replacement ratios also discourage savings which in turn has a countervailing effect on the insurance coefficients. This can be seen by comparing the ratio of median wealth in the two groups which is 2.04 in the baseline model but plunges to 1.64 in the model with defined benefit pensions.

When the subjective discount factors are allowed to differ so that the ratio of median wealth of college and non-college households of about 3 in the data is matched, the insurance

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

[10]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 3: Simulated insurance coefficients: DB pensions

	All	College	High School
Data (P-S)	0.36	0.58	0.06
Model: Common β			
Model IP	0.235	0.228	0.237
β		0.974	0.974
W/Y	2.54	2.77	2.33
Median wealth ratio	1.64		
Model: Heterogeneous β			
Model IP	0.221	0.255	0.193
β		0.981	0.962
W/Y	2.54	3.21	1.87
Median wealth ratio	3.03		

coefficient for college graduates rises to about 0.26 while the one for high school graduates falls to about 0.19, a difference in favor of the former group of 0.07. While this difference is non negligible it is still very small in comparison to the 0.5 that we see in the data. Moreover the difference between the two groups is little higher than the one in the model with no pensions. Hence we conclude that even if qualitatively the introduction of defined benefit pensions goes in the right direction towards explaining the differences in insurance coefficients between college and high school households, quantitatively the impact is largely insufficient.

4.4 Preference specification

All the models in the previous sections were solved by assuming standard expected utility preferences with a quite standard value of risk aversion of 2. As in previous research like Kaplan and Violante (2010) and Campanale and Sartarelli (2018) this generates an aggregate insurance coefficient against permanent shocks that is well below the one measured in the data. However Campanale and Sartarelli (2018) showed that a model with Epstein-Zin preferences featuring a relatively high level of the inter-temporal elasticity of substitution and moderate risk aversion could reconcile the model and the data by changing the shape of the life-cycle profile of wealth in the sense of moving wealth towards the early portion of

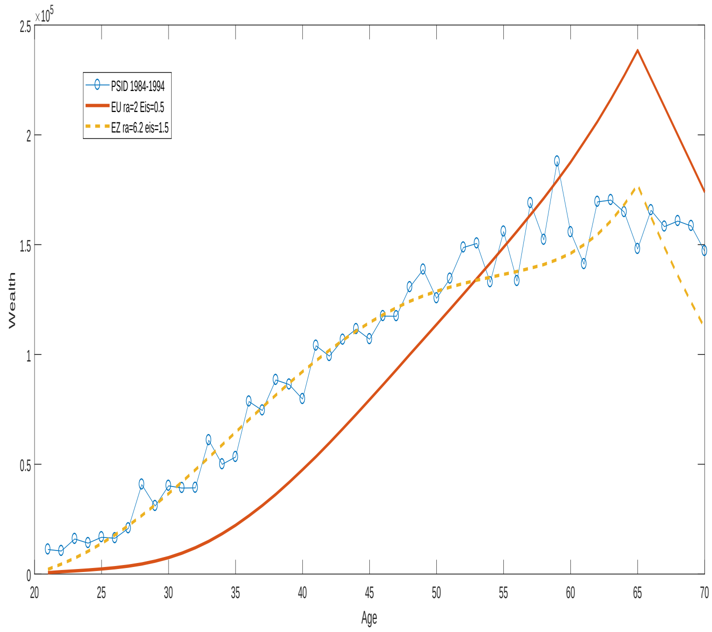

Figure 2: Life-cycle wealth profiles.

the life-cycle. In this section we check if this mechanism can also help explaining the large differences in insurance coefficients between the two educational types.

We propose a specification of Epstein-Zin preferences with an elasticity of inter-temporal substitution of 1.5 and a relative risk aversion of 6.2. To support the intuition provided in the previous paragraph figure 2 reports the profile of wealth accumulation during working life for the baseline model with expected utility and risk aversion of 2, for the model with Epstein-Zin preferences that we consider in this section and for comparison purposes for the data. 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. Model wealth is converted to dollar values by normalization with average earnings. As it can be seen from the continuous line the profile is convex for the expected utility case, with little wealth accumulation occurring in the first 20 years of life, while it has a sort of convex-concave shape for the Epstein-Zin model, represented by the dashed line, with substantial wealth accumulation having occurred already by age 35. The latter profile is also more consistent with the data as it is shown by the line with circle markers largely overlapping with the dashed line.

Results are reported in table 4. The model is solved by allowing heterogeneity in the subjective discount factor and matching the ratio of median wealth between the two groups. No borrowing is allowed in this case. The insurance coefficient for the overall population is 0.33 which is very close to the one in the data. In the late working years the shorter horizon reduces the effective persistence of the permanent shock, moreover the reduction in peak wealth in the Epstein-Zin model is not dramatic both implying that the ability to insure the permanent shock does not change in this part of the life-cycle. The higher wealth accumulation though improves the ability to insure early in the life-cycle. This mechanism explains the increase in the overall insurance coefficient.[11] When we check the results for the two groups separately the difference in the insurance coefficient, at almost 10 percentage points becomes non trivial. This figure is greater than the difference observed in the model with standard expected utility which was only six percentage points, hence even if risk aversion and the elasticity of substitution remain equal for the two groups the change in their value acts asymmetrically leading to better insurance for the college graduates, which is qualitatively consistent with the empirical evidence. The interpretation for this result is that college graduates face a much steeper profile in earnings early in the working life, hence they have a stronger incentive to anticipate consumption and delay the accumulation of wealth. Increasing the elasticity of inter-temporal substitution makes agents more willing to accept uneven consumption profiles, hence to postpone consumption when income is indeed increasing. This force acts more strongly on those, that is the college graduates whose income profile is steeper and for this reason would otherwise lead to less early savings. Even if this parametrization of preferences moves the model insurance in the right direction, still from a quantitative point of view this is far short of what is needed to match the data.

The last experiment suggests that it is possible to play with the elasticity of inter-temporal substitution to further increase the difference between the insurance coefficients of the two groups. In particular if we let high school graduates be substantially less elastic than college graduates wealth accumulation for the former group would occur relatively later in life, thus based on the previous experiment this would lead to a further reduction in their insurance coefficient. Some support for this strategy comes from empirical evidence

[11] See Campanale and Sartarelli (2018)for a more detailed exploration of this mechanism.

Table 4: Simulated insurance coefficients: Epstein-Zin

	All	College	High School
Data (P-S)	0.36	0.58	0.06
Model: Equal eis			
Model IP	0.331	0.386	0.292
β		0.942	0.935
W/Y	2.56	3.01	1.86
Median wealth ratio	3.02		
Model: Heterogeneous eis			
Model IP	0.287	0.373	0.227
β		0.94	0.794
W/Y	2.54	3.05	2.06
Median wealth ratio	3.02		

that suggests that the elasticity of inter-temporal substitution monotonically increases with income. Also it has been found that stockholders, who are on average wealthier tend to show higher values for this parameter.[12] Although the categories of higher income and stockholding households do not fully correspond to the one of college graduates it is surely true that college graduates have on average higher incomes and at the same time they also show higher participation rates in the stock market.[13] Based on this evidence we consider a version of the model where we let the risk aversion be equal to 6.2, like in the previous experiment, for both groups. We then let the elasticity of inter-temporal substitution be different and fix it at 0.1 for the group of high school households and 1.5 for the group of college households. The first is based on the value used by Guvenen (2006) for the non-stockholders, while we fix the one for college graduates arbitrarily with the goal of maximizing the potential effect on generating differences in the insurance coefficients for the two groups.[14] The subjective discount factors are then adjusted so as to keep the same

[12]See on this point Guvenen (2006) and the references cited therein.

[13]Bucks *et al.* (2009), based on the Survey of Consumer Finance report that in 2004 the percentage of households holding stocks directly was 31.5 for those headed by a college graduate and 9.3 percent for those headed by a high school graduate. The percentage for mutual funds were 21.4 and 5.8 percent respectively.

[14]Even so, the value is not chosen purely arbitrarily and still reflects estimates for stockholders. See for example Campbell (2017) which reports median estimates of the elasticity of inter-temporal substitution

ratio of median wealth between the two groups. Results are reported in the second panel of table 4. We first note that the ratio of median wealth between the two groups is about 3 as targeted but that in order to obtain this the subjective discount factor of the two groups is substantially more different than in the previous case: In fact the subjective discount factor is 0.94 for the college graduates and it is only 0.794 for the high school graduates. The reason is that now high school graduates are very inelastic hence they would like to save a large amount of wealth in the later part of their working life to keep a flat profile of consumption past the retirement age. In order to satisfy the calibration target concerning their median wealth it is then needed to make them more impatient. In turn being more inelastic make them less willing to postpone consumption early in the life-cycle which together with their impatience delays wealth accumulation. This can be seen in figure 3 that represents with the dashed line the life cycle profile of wealth for the high school graduates and with the continuous line the life cycle profile of wealth for the college graduates. To highlight the comparison of the relative shape of the two, the latter is re-normalized so that peak wealth is the same for the two groups. As it can be seen, apart from the very first years of life, the profile for the college graduates stands above the one for the high school graduates signaling earlier accumulation of wealth. Since it is in the first years of life when wealth is low that insurance coefficients are low this translates into relatively higher coefficients for the college graduates. The first line of the bottom panel of table 4 confirms this: the insurance coefficient against permanent shocks is 0.373 for the college graduates and it is 0.227 for the high school graduates. This difference of 0.15 is substantial but still quite short of the difference of 0.52 of the data.

The difference in the elasticity of inter-temporal substitution between the two groups used in the latter experiment is most likely an upper bound of what can be thought of as reasonable, given that it was based on the most favorable empirical estimates for stockholders versus non stockholders or estimates based on income level but that there is not a full overlap between the above categories and the classification of the population in this research between high school and college graduates. Since the difference in the insurance coefficient against permanent shocks remains quite below the one in the data we can then conclude that heterogeneity in preferences beyond the one in discount factors needed to

for the population of Swedish stockholding households of 1.5.

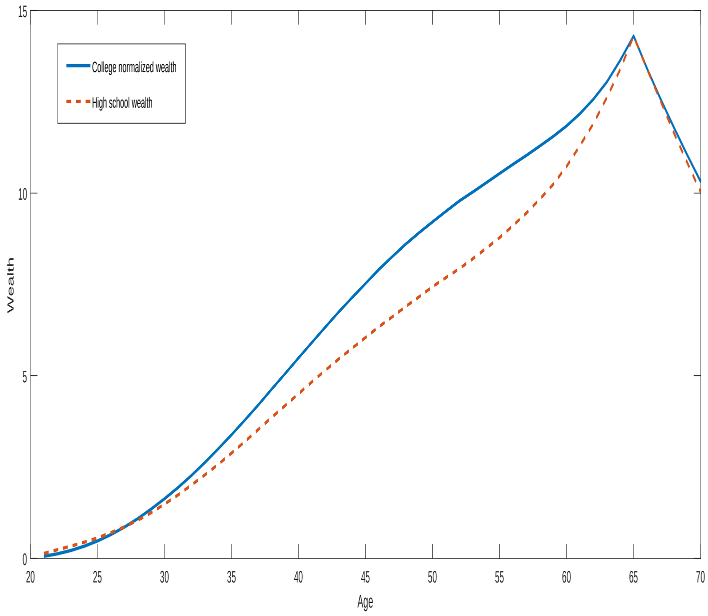

Figure 3: Normalized life-cycle wealth profiles by type.

match average wealth ratios does not look like a feasible candidate to solve the puzzle.[15]

5 Conclusions

In the current paper we have constructed a life-cycle version of the SIM model that is the workhorse of current heterogeneous agents macroeconomics. We have simulated it for two different types of agents calibrated to mimic the earnings process of college and high school educated households. It was shown by Kaplan and Violante (2010) that under reference calibrations the model falls quite short of the data with respect to the value of the insurance coefficients against permanent shocks. The current work uncovered another shortcoming of the model when tested against the data. This arises because the model generates insurance

[15] As a further experiment we also allowed risk aversion to be different between the types, in particular letting high school graduates be less risk averse which would tend to reduce further their wealth accumulation early in life. We fixed risk aversion at two for high school households and left it at 6.2 for the college graduates. This experiment however did not further increase the difference in the insurance coefficients between the two types. The intuition is that median wealth ratios between the two groups are empirically constrained and force the model to make the high school graduates more patient to respect the constraint. Results are available from the author upon request.

coefficients against permanent earnings shocks that are virtually equal for the two types of households while in the estimates provided by Blundell *et al.* (2008) high school households shows almost no consumption smoothing at all while college educated households appear to smooth about 60 percent of permanent shocks. We explored a number of simple possible solutions to this puzzle. These included differential access to credit, more patient college households to match the ratio of median wealth of the two groups in the data, the addition of defined benefit pensions and heterogeneous elasticity of substitution. All of them proved insufficient to align the model with the data. There are several more explanations that we could consider. We briefly discuss one of them here, that is, intergenerational insurance between members of a dynasty. Low insurance arises mainly because of very limited consumption smoothing for young workers given that we always assumed that all agents start with zero wealth. Since it is plausible that college graduates come from more affluent parents they may have better access to intergenerational transfers in the initial years of work when wealth accumulation is still modest hence self-insurance very poor. This and others are surely interesting avenues to pursue to investigate this crucial issue for the quantitative macroeconomics with heterogeneous agents, but given that they represent substantial departures from the basic framework that we used to pose the puzzle we leave them for future research.

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.

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).

Bucks, B., Kennickell, A., Mach, T. and Moore, K. (2009). Changes in U.S. Family

Finances from 2004 to 2007: Evidence from the Survey of Consumer Finances. *Federal Reserve Bulletin*, **92**, 1–56.

Bucks, B., Kennickell, A. and Moore, K. (2006). Recent Changes in U.S. Family Finances: Evidence from the 2001 and 2004 Survey of Consumer Finances. *Federal Reserve Bulletin*, **92**, 1–38.

Campanale, C. and Sartarelli, M. (2018). Life-cycle wealth accumulation and consumption insurance. Working Paper 055, Department of Economics and Statistics, University of Torino.

Campbell, Y.J. (2017). The cross-section of household preferences. Arrow Lecture 2, Stanford University. April 19 .

Cerletti, E. and Pijoan-Mas, J. (2012). Durable Goods, Borrowing Constraints and Consumption Insurance. Discussion Paper DP9035, Centre for Economic Policy Research.

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.

Dynarski, S. and Gruber, J. (1997). Can Families Smooth Variable Earnings? *Brookings Papers on Economic Activity*, **1997** (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*, pp. 317–339.

Guvenen, F. (2006). Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective. *Journal of Monetary Economics*, **53**, pp. 1451–1472.

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.

Kennickell, A., Starr-McCluer, M. and Surette, B. (2000). Recent Changes in U.S. Family Finances: Results from the 1998 Survey of Consumer Finances. *Federal Reserve Bulletin*, **86**, 1–29.

Knowles, J. and Postlewaite, A. (2004). Do Children Learn to Save from Their Parents?. *Working Paper*.

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

Perez-Arce, F. (2017). The Effect of Education on Time Preferences. *Economics of Education Review*, **56**, 52–64.

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

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.

DEPARTMENT OF ECONOMICS AND STATISTICS
UNIVERSITY OF TORINO
Corso Unione Sovietica 218 bis - 10134 Torino (ITALY)
Web page: http://esomas.econ.unito.it/