THE IMPLICATIONS OF RICHER EARNINGS DYNAMICS FOR CONSUMPTION, WEALTH, AND WELFARE

Mariacristina De Nardi
Giulio Fella
Gonzalo Paz Pardo

The Implications of Richer Earnings Dynamics for Consumption, Wealth, and Welfare
Mariacristina De Nardi, Giulio Fella, and Gonzalo Paz Pardo
NBER Working Paper No. 21917
January 2016
JEL No. D14,D31,E21,J31

ABSTRACT

Earnings dynamics are richer than those typically used in macro models. This paper provides multiple contributions. First, it proposes a non-parametric way to model rich earnings dynamics that is easy to use in structural models. Second, it constructs a large, synthetic, data set that matches the earnings dynamics of the U.S. tax earnings. Third, it estimates our non-parametric earnings processes using two data sets: the Panel Study of Income Dynamics and our synthetic tax data. Fourth, it compares the implications of our earnings processes to those of a standard AR(1) in a life cycle structural model of savings and consumption.

Mariacristina De Nardi
Federal Reserve Bank of Chicago
230 South LaSalle St.
Chicago, IL 60604
and University College London
and Institute For Fiscal Studies - IFS
and also NBER
denardim@nber.org

Giulio Fella
Queen Mary University of London
Mile End Road
London E1 4NS
United Kingdom
and and Centre for Macroeconomics - CFM
and Institute for Fiscal Studies - IFS
g.fella@qmul.ac.uk

Gonzalo Paz Pardo
University College London
Drayton House
30 Gordon St
London WC1H 0AX
United Kingdom
gonzalo.pardo.13@ucl.ac.uk

1 Introduction

Earnings risk, consumption, and wealth accumulation are tightly linked. In particular, the magnitude and persistence of earnings shocks determine how saving and consumption adjust to buffer their impact on current and future consumption. Appropriately capturing earnings risk is therefore essential to understand consumption and wealth decisions and their resulting inequalities and welfare implications.

With few exceptions, most of the literature on earnings dynamics focuses on linear models with normal innovations that imply that the persistence of earnings shocks is independent of age and earnings histories and that positive and negative income changes are equally likely. Recent work (e.g., Arellano, Blundell and Bonhomme (2015), and Guvenen, Karahan, Ozkan and Song (2015) take advantage of large data sets or new methodologies to show that household's earnings dynamics feature substantial asymmetries and non-linearities. In particular, earnings changes display substantial skewness and the persistence of shocks depends both on age and current earnings. Guvenen et al. (2015) have documented these features using Social Security Administration tax earnings (W2) data for a large panel of individuals, while Arellano et al. (2015) have shown that similar features hold in the PSID too.

It is well known (De Nardi, 2004; Cagetti and De Nardi, 2006, 2009) that life cycle models with linear earnings process have substantial difficulties matching important features of the empirical wealth distribution, in particular its long right tail, without entrepreneurship or a non-homothetic bequest motive. In contrast, Castañeda, Díaz-Giménez and Ríos-Rull (2003) have shown that introducing a skewed distribution of earnings shock calibrated to match a number of wealth moments dramatically improves the fit of the right tail of the wealth distribution. A key question is to what extent the required degree of skewness is consistent with the empirical earnings distribution. Or, equivalently, what are the implications for consumption, wealth inequality and welfare of an estimated earnings process which matches the richer dynamics discussed above.

The complexity of the recently estimated earnings processes makes it computationally very costly to incorporate them in the rich structural models of household's decision making that are needed needed to study consumption and saving decisions and their implied consumption, wealth, and welfare inequalities, both in the cross-section and over the life cycle. For this reason, they have not been introduced in structural models of household's decision making over the life cycle so far.

This paper fills this gap by proposing a new, non-parametric, way to model a non-linear life-cycle earnings dynamics that is both consistent with the new empirical findings and easy to introduce in structural quantitative models of optimal households' decision making. Our methodology exploits the recent availability to researchers of longitudinal data sets on earnings with a large number of individuals at each point in time. This availability makes it possible to extend to the study of life-cycle earnings dynamics techniques that until now have been used to account for inter-generational income mobility (e.g. Atkinson, Bourguignon and Morrisson (1992)).

Our proposed method has also the advantage of being very simple. It estimates the process for the stochastic component of lifetime earnings (namely the residual of the regression of earnings over a deterministic age profile, time and fixed effects) in the following way. First, given the distribution of (residual) earnings at each age t, it computes the age-specific transition matrix from the percentile rank of the earnings distribution at age t to that at age $t + 1$. Second, it discretizes the distribution of earnings at each age by replacing the (heterogeneous) values of earnings in each rank percentile with their average. The result is a non-parametric representation of the earnings process that follows a Markov chain with an age-dependent transition matrix and a fixed number of age-dependent earnings states.

The method is very flexible and, given a sufficiently fine partition of the earnings distribution at each age, is also very accurate. It imposes no restriction on the earnings process other than the common (Markov) assumption that it has one-period memory. Achieving a

fine partition of the earnings distribution requires that the number of transitions across all rank percentiles at each age is large enough for sampling noise not to be an issue. Therefore, the method requires either the availability of longitudinal data sets with a large number of individuals or the ability to construct large artificial data sets by simulating parametric and semi-parametric processes estimated using large data sets. Our method also lends itself naturally to being used in structural models which usually employ discrete (Markov chain) approximations to earnings processes, but relaxes the assumption of symmetry and of age- and earnings-independent persistence.

We apply our method both to the PSID and to a large synthetic panel generated simulating the parametric earnings processes estimated by Guvenen et al. (2015) on their W2 data.[1] The advantage of the PSID is that it is a well known data set, that includes substantial more information about earners than the W2 data. Conversely, unlike the PSID, the W2 data set is extremely large and does not feature top-coding or under-representation of very high earners. For this reason we apply our non-parametric estimation method to both sets of data and compare their implications. We find that our approximation replicates very well the first four unconditional moments of log earnings and the conditional moments of log earnings growth in both datasets.

We then introduce these two earnings processes in a standard life cycle model of savings and consumption with incomplete markets, and we compare their implications to those of the standard AR(1) earnings process commonly used in the literature (e.g. Huggett, 1996). In sum, we compare the implications of three earnings processes. We name "benchmark" the earnings process derived from discretizing an AR(1) process calibrated to match the US earnings Gini and to approximate the life-cycle profile of earnings inequality, as done by Huggett (1996). We name "non-linear (NL) PSID process" the earnings process obtained by applying our estimation method to the PSID data, while we refer to the "non-linear

---

[1]Our synthetic data set matches, by construction, all of the moments that these authors use in their estimation procedure and allows us to sidestep the problem that accessing the W2 tax data is very involved.

(NL) synthetic W2 process" when using the earnings process resulting from applying our estimation procedure to our synthetic W2 data panel.

Our main findings are the following. First, both our PSID and W2 non-parametric earnings processes imply much less persistence at low and high earnings levels than those implied by the discretized AR(1) process used by Huggett (1996), which is very similar to those commonly used in the quantitative macro literature with heterogeneous agents. Second, when introduced in a standard quantitative life cycle model, both the NL PSID and the NL W2 earnings processes generate a much better fit of the wealth holdings of the bottom 60% of individuals, but fail to improve the fit of the right tail of the wealth distribution relative to the standard benchmark process. In fact, in contrast with the process used by Castañeda et al. (2003)), even the large skewed earnings risk faced by the very high earners under the W2 earnings processes, is not sufficient to generate the large savings of the high wealth and high-earners that leads to the large concentration of wealth that we observe in the data (in contrast with the process used by Castañeda et al. (2003)). This is a frequent drawback of life-cycle models with no entrepreneurship or transmission of bequests (De Nardi, 2004), (Cagetti and De Nardi, 2006), and (Cagetti and De Nardi, 2009) and we show that it still holds even in the presence of reasonable heterogenity in household's patience.[2] Third, everything else equal, a new born worker under the veil of ignorance would require only a one-time transfer equal to 7% of average earnings to switch from the W2 to the AR(1) world, despite the high skewness of the W2 process. Decomposing this result shows, however, that this small ex-ante compensation masks large heterogeneity by earnings levels, with the low earnings households preferring the W2 economy and the high earnings households preferring the AR(1) economy. For instance, a worker starting off at one-quarter of average earnings level would require a compensation of about twice average earnings to move from the W2 to the AR(1) world. This happens because the smaller persistence of earnings at low earnings levels implies that

---

[2]In computations available upon request, we verify that this result is not sensitive to increasing the value of the CRRA coefficient up to a rather high value of 4.

low-earnings households with little or no assets have an easier time to self-insure against earnings shocks under the W2 earnings processes, while high earnings households prefer the higher persistence in the AR(1) environment. These findings on the persistence of the shocks and their implications at different earnings levels are important and likely to have a large effects on the costs and benefits of many government transfers programs, such as unemployment, insurance Supplemental Social Insurance (SSI), and redistributive taxation.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 describes the main features of the data concerning earnings dynamics and inequality in consumption and wealth that our model seeks to match. Section 4 explains how we either calibrate or estimate our three earnings processes. Section 5 highlights the implications of our earnings processes over the life cycle and in the cross-section. Section 6 specifies our model and its calibration. Section 7 discusses the main implications of our earnings process, when introduced in standard quantitative model of savings and consumption over the life cycle. Section 8 reports compensating differentials between different earnings processes to evaluate the household's ability to self-insure under different configurations of earnings risk. Section 9 explores the consequences of allowing for discount rate heterogeneity and Section 10 concludes. Appendix A discusses key features of both the PSID and our synthetic W2 data set. Appendix B shows how well our non-parametric earnings processes match the PSID and W2 data.

2   Related literature

Our goal is to take a standard life cycle model with exogenously incomplete markets (as in Huggett (1996)) and study how the three earnings processes that we consider affect consumption, wealth, and their inequalities, both over the life cycle and in the cross-section, and to uncover whether these earnings processes have different welfare implications. Thus, our work connects with two branches of the literature. First, it relates to the literature on

quantitative models with heterogenous agents, which have been used, for instance, to study inequality and the effects of government policy reforms. Second, it relates to the literature on earnings dynamics and its effects on consumption choices. Both branches of the literature are vast.

The literature on quantitative models with heterogenous agents typically assumes that the logarithm of earnings follow a Gaussian autoregressive process of order one, which is a very parsimonious specification of the earnings process. This Markov Chain process is then typically discretized using the methods described in Tauchen (1986) or Tauchen and Hussey (1991). An important set of applications of the literature on quantitative models with heterogenous agents studies consumption and wealth inequality and their evolution over the life cycle. See Quadrini and Ríos-Rull (2014), Cagetti and De Nardi (2008) and De Nardi (2015) for a discussion of what features of the wealth data various versions of these models are able to match.

It has been argued, however, that a discretized Gaussian AR(1) process might not be an accurate approximation for actual earnings dynamics. Geweke and Keane (2000) point out that log earnings follow a more complex time-series process and that its shocks are not normally distributed. For instance, the evidence shows that the distribution of log earnings shocks is leptokurtic, which means that it has fat tails: there is a small but relevant number of individuals who suffer large short-run earnings shocks. A Gaussian AR(1) fails to capture this feature of the data, thus overestimating persistence of log earnings.

Meghir and Pistaferri (2004) relax the assumption of i.i.d. log earnings innovations and model the conditional variance of log earnings shocks by allowing for an autoregressive conditional heteroschedasticity (ARCH) model of the variance of log earnings. More specifically, they allow the variance of log earnings innovations to depend on year effects, the education of the individual, and unobservable individual factors. In this context, the current realizations of the variance are thus informative about future earnings. They find evidence of ARCH-type

7

variances for both the permanent and transitory shocks and for individual-specific variances.

Browning, Ejrnaes and Álvarez (2010) show that it is necessary to include individual-level heterogeneity in earnings processes to account for many important features of the data. They compare a standard unit-root model for log earnings with a more complex ARMA(1,1) process with large heterogeneity, where individual-specific parameters are derived from three stochastic latent factors. This richer process significantly improves the fit of the earnings process to the data.

Altonji, Smith and Vidangos (2013) consider a multivariate model of earnings and jointly model log wages, job changes, unemployment transitions and hours worked by explicitly allowing these processes to interact (for instance, wages depend on job tenure, while job changes depend on the expected wage should the individual remain in the same job). Their model is not tractable for our purposes, but provides important insights on the evolution of earnings and job tenure. Furthermore, it further stresses the importance of allowing for heterogeneity and state dependence when approximating earnings dynamics.

Arellano et al. (2015) model a log earnings process composed by the sum of a transitory stochastic component and a persistent Markovian process, which allows for significant nonlinearities (capturing e.g. job losses, health shocks or career changes).

Blundell, Graber and Mogstad (2015) use Norwegian panel data and find that a better description of earnings dynamics requires allowing for heterogeneity by education levels and accounting for non-stationarity. Their preferred specification for log earnings includes an idiosyncratic constant term, an idiosyncratic experience profile, an AR(1) persistent shock component and an MA(1) transitory shock component. All components considered are allowed to have a skill-dependent distribution and the variances of shocks are allowed to be age-dependent and hence non-stationary.

In a recent paper Guvenen et al. (2015) study the evolution of male earnings over the life cycle using a large Social Security Administration panel data set. Using non-parametric

methods they find that labor earnings do not conform to the standard assumptions in most of the empirical literature surveyed in Blundell (2014); namely, that the stochastic component of earnings is the sum of linear processes with i.i.d, normal innovations. In particular, they find that earnings shocks display strong negative skewness and very high kurtosis. They also show that cross-sectional moments change with age and previous earnings levels. Finally, they estimate a complex set of parametric processes that match a large number of moments in the data, including higher order moments. Unfortunately, the number of state variables needed to include these kind of processes in a structural model is very large. Even the most simplified model that they consider appropriate for empirical work, includes an heterogeneous lifetime income profile (which implies two state variables, one for the constant and one for the slope), individual-specific variances and a mixture of two AR(1) components. This specification generates both computational and modeling issues because it is not obvious how the discretization of each of these variables should be performed to preserve their relationships and dynamics.

Recent developments in this literature are discussed in Meghir and Pistaferri (2011). The consequences of these richer earnings processes on consumption, savings and welfare remain, however, an important open question and one that we seek to address.

3   Some important facts about earnings and wealth

The earnings shocks experienced by US workers display important deviations from the assumptions of log-normality and independence from age and earnings realizations as documented by Arellano et al. (2015) for the PSID and Guvenen et al. (2015) for their W2 tax data. The top panel of Figure 1 reports our computations for the conditional moments of log earnings growth from the PSID data. The bottom panel reports the same moments estimated on the synthetic panel dataset that we generate by simulating the parametric process estimated by Guvenen et al. (2015) on their W2 data. To capture the non-linearity

Figure 1: Conditional moments of log earnings changes by age cohort (PSID and Synthetic W2 data)

PSID data

W–2 synthetic panel
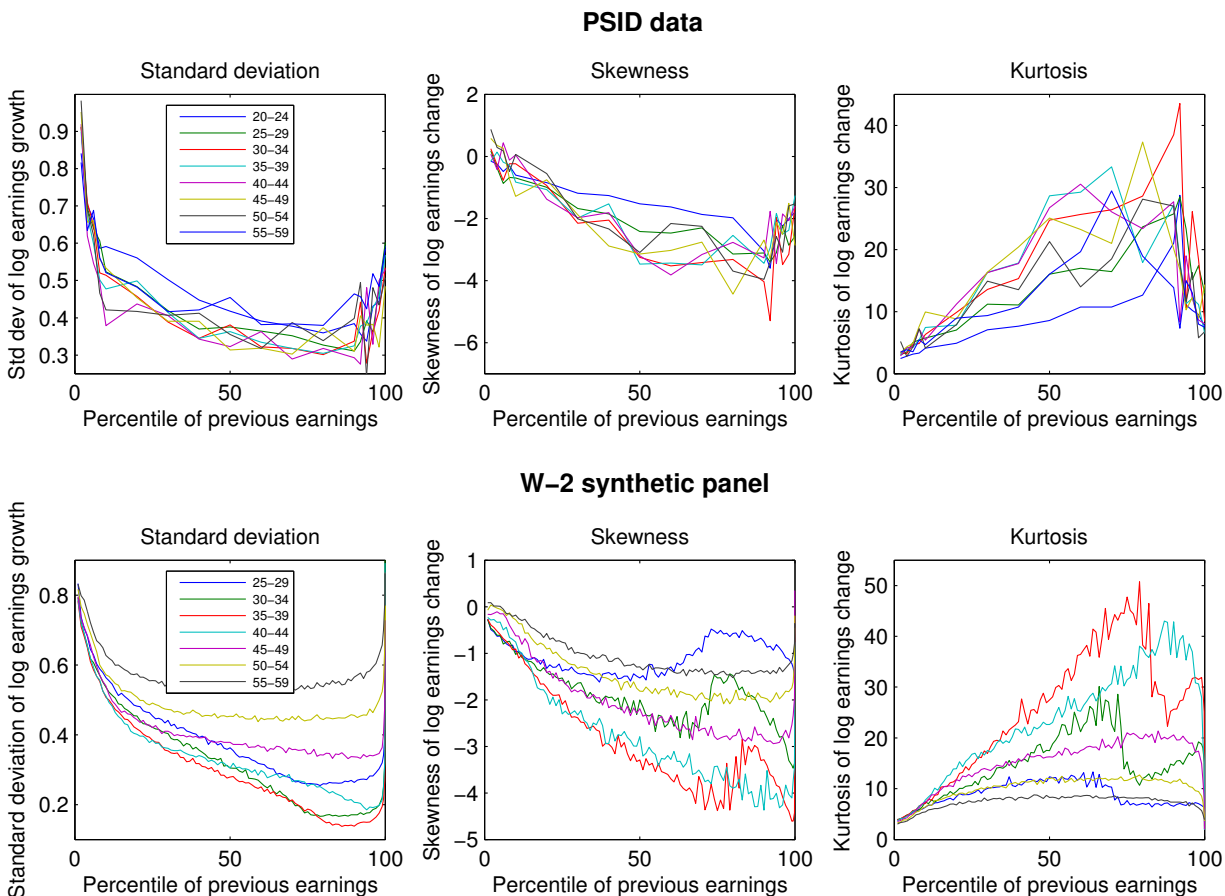
and asymmetry in the data, Guvenen et al. (2015) fit a flexible process consisting of a mixture of AR(1) plus a heterogeneous income profile and estimate it by simulated method of moments.[3]

The figure shows that the conditional variance of log earnings growth is U-shaped across all age groups: individuals with the largest and smallest earnings are the ones that suffer from higher earnings risk. Furthermore, it declines until age 35 years and starts increasing after age 45.

The figure also shows that log earnings growth has strong negative skewness and very high

---

[3]The parameterization is the one denoted "benchmark model" and reported in column (2) in Table III of their paper. Further details about this process, as well as the PSID data, are provided in Appendix A.

kurtosis, and that these moments depend both on age and previous earnings. A negative skewness signals a fat left tail in log earnings changes and means that individuals face a non-negligible risk of a large fall in earnings and smaller probabilities of very large earnings increases. The skewness is more negative for individuals in higher earnings percentiles and for individuals between 35 and 45 years of age. This indicates that individuals face a larger downward risk as they approach middle age.[4]

The kurtosis is a measure of the peakedness of the distribution of log earnings changes. A high kurtosis means that most of the people experience negligible earnings shocks in a given year but that at the same time a small proportion of individuals face very large earnings shocks. The kurtosis of the US earnings shocks distribution, when conditional on previous earnings, gets as large as 30 (compared to 3 in a standard normal distribution). The level of kurtosis is hump-shaped by earnings percentile and increases until age 35-45 to then decrease thereafter. Our graphs thus indicates large deviations between the moments observed in the data and those implied by the standard AR(1) earnings process.

Overall, despite a much smaller sample size, the moments derived from PSID data are both qualitatively and quantitatively very close to those computed from our synthetic W2 panel. However, it should be noted that the W2 synthetic panel is much larger and is not affected by top-coding or differential survey responses, and thus offers a better approximation of the moments for the earnings-rich. This implies, for instance, that even though the variance of earnings is decreasing in earnings until the 95th percentile, it then increases very significantly for the earnings-rich and more so in the W2 data set. Similarly, the negative skewness and the kurtosis are much more pronounced for the high earners in the W2 data set.

Turning to wealth and consumption inequality, the top line of Table 1 summarizes the main statistics of the wealth in the US.[5] Overall, wealth is very unequally distributed, with

---

[4]Graber and Lise (2015) account for this kind of earnings behavior in the context of a search and matching model with a job ladder.

[5]The data on wealth are from Wolff (1987) and come from the 1983 Survey of Consumer Finances. We

a Gini coefficient of 0.72 (the corresponding value is for the earnings distribution is 0.51 (Quadrini and Ríos-Rull, 1997)). Using different time periods yields a slightly more concentrated wealth distribution in the hands of the richest few.

As discussed by Quadrini and Ríos-Rull (2014), Huggett (1996), Cagetti and De Nardi (2008), and De Nardi (2015), the standard life cycle model with incomplete markets and discretized AR(1) earnings shock cannot match the large concentration of wealth in the hands of the richest few and generates too many people at zero (or negative) wealth. An important question is whether a better representation of earnings risk can help match wealth inequality and along which dimensions.

Table 1: Wealth and consumption distribution statistics, U.S. data

		Percentage held by the top						At negative or
	Gini	1%	5%	20%	40%	60%	80%	zero
Wealth	.72	28	49	75	89	96	99	5.8-15.0%
Non-durable consumption	.31	5	15	40	62	79	92	0%

The second line in Table 1 displays the distribution of equivalized consumption of non durables and services for the entire US adult population in 1989. The consumption Gini that we compute is in line with the estimates in the literature (around 0.29 for 1989 in Fisher, Johnson and Smeeding (2013)) and so is the shape of the distribution (our implied 90/10 and 50/10 ratios are 4.23 and 2.12 respectively, which are in line with those in Meyer and Sullivan (2010)). Consumption inequality is thus significant, but much lower in magnitude than both earnings and wealth inequality. Relatedly, the variances of log consumption in the US are lower and increase less over the life-cycle than variances of income and earnings, but they do increase (Deaton and Paxson (1994)).[6]

use the consumption data from the CEX sample derived by Heathcote, Perri and Violante (2010), and we use their sample A and variable definitions.

[6]Some studies have shown consumption inequality to be steadily rising since 1980 (Aguiar and Bils (2015), Heathcote et al. (2010), Blundell, Pistaferri and Preston (2008)), at least until the Great Recession (Attanasio and Pistaferri (2014)), but its evolution in recent years remains controversial and worthy of more study.

4 The discretized earnings processes

The most common earnings process used in quantitative models of consumption and saving is a discrete approximation to an AR(1) process for the logarithm of earnings (e.g. Huggett, 1996). We are going to compare the implications of such a process with those of two alternative earnings processes estimated by applying our non-parametric methodology on the PSID and the synthetic W2 data. In a nutshell, our methodology provides an alternative discretization method which is much more flexible—i.e. it does not impose symmetry and linearity—that the standard discretization methods based on Tauchen (1986) and its variants.

4.1 Benchmark earnings process

The benchmark earnings process is based on Huggett (1996)'s calibration, where the labor endowment process is an AR(1) with persistence parameter $\gamma = 0.96$ and Gaussian shocks with variance $\sigma_\epsilon^2 = 0.045$. The initial endowment of the first cohort of agents is also normally distributed, with variance $\sigma_{y_1}^2 = 0.38$. Huggett's choice of value for the variance of shocks σ_ϵ^2 is based on similar estimates in previous literature (e.g. Lillard and Willis (1978) or Carroll, Hall and Zeldes (1992)). The variance of the initial condition $\sigma_{y_1}^2$ is chosen to match the earnings Gini for young agents (Lillard (1977), Shorrocks (1980)). Given both variances, γ is calibrated to match an overall earnings Gini for US males of 0.42.

We then follow the discretization strategy applied by Huggett (1996), which is based on Tauchen (1986): The state space for log earnings is divided in 18 equidistant points, that range between $-4\sigma_{y_1}$ and $4\sigma_{y_1}$. To better approximate the upper tail of the earnings distribution, a further point is included ad hoc, situated at $6\sigma_{y_1}$. The small number of individuals that are situated in this upper point earn about 40 times median earnings. The rest of the grid ranges between 8% of median earnings and 11 times median earnings. The computation of the transition matrix relies on the fact that, conditional on earnings at age

$h-1$, earnings at age h are a draw from a normal distribution with mean γy_{h-1} and standard deviation σ_ϵ.

There are two main differences between our benchmark earnings process and the original one in Huggett (1996). First, Huggett studies agents from age 20 to 65, while we define working the life as spanning from age 25 to 60 for comparability with our W2 synthetic data. To preserve the match of the earnings process to the US earnings Gini for men to be the same as in Huggett's paper, we slightly increase the persistence parameter γ to 0.963, while keeping both $\sigma_{y_1}^2$ and σ_ϵ^2 constant. Second, we borrow the life-earnings profile from Hansen (1993), since the exact values used by Huggett into his model are not readily available.

4.1.1 Non-parametrically estimated processes for the PSID and W2 earnings data

We estimate two non-parametric earnings processes, one from the PSID data and one from the Social Security W2 synthetic panel that we generate from the empirical processes estimated by Guvenen et al. (2015). Appendix A discusses the PSID data and our synthetic W2 data. The main difference with respect to the discretization method in the previous section is that the alternative discretization we propose is very flexible and therefore capable to match the asymmetries and non-linearities in the PSID and W2 data.

We first purge the original earnings data from time and age effect and then discretize the residual stochastic component of earnings. Let y_{ht}^i denote the logarithm of labor earnings for an individual i at time h and age t. We assume the process for y_{ht}^i takes the following form

$$(1) \qquad\qquad y_{ht}^i = d_h + f(\theta_i, t) + \eta_{ht}^i,$$

where d_h is a yearly dummy and $f(\theta_i, t)$ is a quartic function of age. The term η_{ht}^i captures

the stochastic component of earnings.[7]

We assume that the distribution of the stochastic component of earnings η_{ht}^i is i.i.d. across individuals but do not impose any additional restriction other than assuming that conditionally on age t, η_{ht}^i follows a Markov chain of order one, with age-dependent state space $Z^t = \{\bar{z}_1, \ldots, \bar{z}_N\}$, $t = 1, \ldots, T$ and an age dependent transition matrix Π^t, which has size $(N \times N)$. That is, we assume that the dimension N of the state space is constant across ages but we allow its possible realizations and its transition matrix to be age-dependent.

We determine the points of the state-space and the transition matrices at each age in the following way.

1. We recover the stochastic component of earnings as the residual of running the regression associated with equation (1).

2. Take a number of bins, for each age, N. At each age, we order the realized log earnings residuals by their size and we group them into bins, each of which contains $1/N$ of the number of observations at that age. Let b_n^t denote the interval associated with bin $n = 1, \ldots, N$ at age t.

   Because the PSID and the our synthetic W2 data greatly differ in their sample size, we have chosen N in the two data sets as follows.

   - Due to the limited sample size of the PSID data, we have to strike a balance between a rich approximation of the actual earnings dynamics by earnings level (that is, a large number of bins) and keeping the sample size in each bin sufficiently large. We have thus evaluated many possibilities. In our main specification we report the results for bins representing deciles, with the exception of the top and bottom deciles, that we split in 5. Therefore, bins 1 to 5 and 14 to 18 include

---

[7]We do not need to include yearly dummies when using W2 synthetic data because it is already extracted in the original estimation procedure.

2% of the agents at any given age, while bins $n = 6, \ldots, 13$ include 10 % of the agents at any given age. This implies a total of 18 bins.

- Our synthetic W2 dataset is simulated to contain 18 million observations (500,000 individuals over 36 years), which implies that we are not constrained by issues of insufficient sample size. For this data set, we thus report results for a discretization with 103 gridpoints, which aims at accurately capturing the earnings dynamics of the earnings-rich. The bottom 99 gridpoints correspond to the bottom 99 percentiles of the earnings distribution, while the top 1% is divided into 4 bins. More specifically, we separate in a special bin the top 0.01%, we create another bin for the rest of the top 0.05% and we divide the remaining people if the top percentile in two bins.

3. The points of the state space at each age t are chosen so that point z_t^i is the mean of $\eta_{ht}^i \in b_n^t$. We have considered specifications in which the summary statistic of each bin is the median instead of the mean, with no impact on the results.

4. The initial distribution at model age 0, is the empirical distribution $Pr\{\eta_{it} \in b_n^0\}$, $n = 1, \ldots, N$.

5. The elements π_{mn}^t of the transition matrix Π^t between age t and $t+1$ are the proportion of individuals in bin m at age t that are in bin n at age $t + 1$. The use of transition matrices is well established in the study of income mobility (e.g. Jäntti and Jenkins (2015)). The main difference is that while studies of income mobility are usually concerned about intra- or inter-generational mobility across relative rankings in the earnings distribution, we are interested in capturing mobility across earnings levels.

For this procedure to provide consistent estimates of the population earnings distribution over the life cycle, we need a large enough number of individuals in the sample for every age group. This is not a problem for our W2 synthetic data, that we simulate, but is an issue

for PSID data. To solve this problem, we assume that age t actually includes people aged $t-1$, t and $t+1$. Specifically, we create, for every age t in the sample, a fictitious t-year-old cohort which is formed of all individuals in the sample who are $t-1$, t and $t+1$ years old. We then apply the method described above to this fictitious cohort to derive the state space for agents of age t. Since we repeat this for every age, most observations in the original data base are used three times.

To keep comparability between the benchmark earnings process and our estimated processes, we use the same age-efficiency profile in all cases. Hence, we discard the age-efficiency profile that we extract from the PSID and W2 data and we use those data sets to estimate earnings mobility. Furthermore, average earnings in each of the three economies are normalized to 1 so that the total amount of resources that are exogenously entering the economies are the same.

Appendix B extensively discusses how well our non-parametric estimation method matches the observed moments both in the PSID and the synthetic W2 data, including for the same number of earnings bin. The main conclusion from that comparison is that our method does a very good job of matching the vast majority of moments in the data, and especially so for the earnings process with 103 bins.

5 What are the implications of our earnings processes?

5.1 Moments across the life-cycle

Figure 2 shows the average earnings profile that we assume for all three processes, calibrated using the age-efficiency profile in Hansen (1993). The earnings process is calibrated such that average earnings for everyone in the economy (including those who have no labor earnings because they are retired) is 1.

Table 2 reports the cross-sectional Gini coefficient of earnings for the overall working

17

Figure 2: Average earnings by age
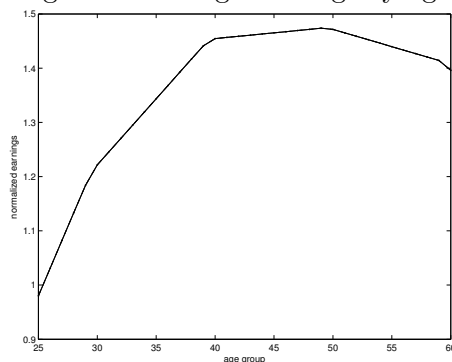
population, and the Gini coefficients for the first and last cohort of the simulated processes (25 and 60 years old).

Table 2: Earnings process statistics

Process	Overall Gini	Gini at 25	Gini at 60
Benchmark	0.4160	0.3371	0.4398
NL PSID Process	0.3324	0.2649	0.3890
NL Synthetic W2 Process	0.4120	0.3456	0.4980

The table shows that, compared to the benchmark economy, whose earnings process is chosen to match the level and approximate the rise of the Gini coefficient by age, the NL earnings process from the PSID process generates a significantly lower Gini coefficient both in the whole cross-section of working ages and at younger and older ages. This is likely due to the features of the PSID data we discussed earlier (small sample size, lack of over-sampling of high earners, lower response rate of high earners, etc.). The NL synthetic W2 process, instead, is much more successful in terms of replicating the actual level of overall inequality and actually generates earnings inequality around the age of retirement that is higher than that generated by the other two processes.

Figure 3 reports the variances of log earnings by age cohort for the 25-60 years old workers. This figure confirms the findings implied by the Gini coefficients. The NL PSID process generates less inequality than the benchmark, but displays, consistently with the

18

data, a significant variance increase across the life-cycle. It should be noted, in fact, that in the benchmark, the variance of shocks during the lifetime is constant and that the realized increase in the variance of the earnings process over the life cycle is driven by the fact that the variance of initial conditions is assumed to be lower than the variance of the process itself. The NL synthetic W2 process generates a higher amount of inequality at every age and a very significant variance growth towards later ages (which is consistent with the one in our synthetic W2 data, see Appendix B). This explains the large Gini coefficient of earnings at age 60 that we observed earlier.

Figure 3: Variance of the earnings processes by age
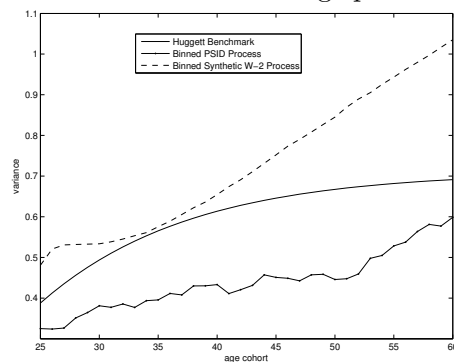
5.2 Earnings mobility

When compared to the benchmark, both of our empirical processes (NL PSID and NL synthetic W2) generate larger earnings mobility, particularly at earlier ages. This can be seen in Table 3, that summarizes earnings mobility between the age of 25 and 30. In particular, the NL synthetic W2 process, despite showing a level of overall earnings inequality similar to the benchmark, implies more mobility. This implies that individuals are not very likely to be stuck at bad earnings realizations if they begin their lives with a negative earnings shock. However, individuals are also more likely to fall from a higher to a lower earnings state.

Table 3: Earnings mobility between earnings quintiles - 25 to 30 years old

Process	Quintile at 25	Quintile at 30				
		1st	2nd	3rd	4th	5th
Benchmark	1st	0.54	0.34	0.10	0.01	0
	2nd	0.25	0.41	0.26	0.07	0.01
	3rd	0.07	0.24	0.41	0.22	0.06
	4th	0.01	0.08	0.28	0.4	0.22
	5th	0	0.02	0.12	0.36	0.51
NL PSID Process	1st	0.28	0.25	0.20	0.15	0.12
	2nd	0.25	0.24	0.21	0.17	0.13
	3rd	0.20	0.21	0.21	0.20	0.18
	4th	0.16	0.17	0.20	0.23	0.24
	5th	0.12	0.13	0.17	0.25	0.33
NL Synthetic W2 Process	1st	0.37	0.28	0.18	0.10	0.06
	2nd	0.27	0.26	0.21	0.15	0.1
	3rd	0.19	0.21	0.22	0.21	0.16
	4th	0.11	0.15	0.21	0.26	0.26
	5th	0.06	0.09	0.16	0.27	0.42

6   The model

The model is based on Huggett (1996)'s paper. There is an infinitely lived government and overlapping generations of individuals who are equal at birth but receive idiosyncratic shocks to labor income throughout their working lives. The model period is one year.

6.1   Demographics

Each year, a positive measure of agents is born. Agents start working life at age 25 with no assets and a random productivity draw. At 30, each agent begets $(1+n)$. Working life ends at 60, when agents retire.

Agents face a positive probability of dying throughout their lifetimes. This probability grows with age and is 1 at age 85; i.e., agents die for sure before turning 86 years old.

We restrict attention to stationary equilibria, hence variables are only indexed by age t.

6.2 Preferences and technology

Preferences are time separable, with a constant discount factor. The intra-period utility is CRRA: $u(c_t) = c_t^{1-\sigma}/(1-\sigma)$.

Agents are endowed with one indivisible unit of labor which they supply inelastically at zero disutility. The efficiency of their labor supply is subject to random shocks and follows a Markov chain of order 1 with (possibly) age-dependent state spaces and transition matrices.

We assume a closed economy with an aggregate Cobb-Douglas production function $F(K, L) = AK^\alpha L^{1-\alpha}$, where K is aggregate capital and L is aggregate labor. Capital depreciates at rate δ.

6.3 Markets and the government

There are competitive markets for the two factors of production. The wage per efficiency unit of labor is denoted by w and r is the rental price for undepreciated capital.

Asset markets are incomplete. Agents can only invest in physical capital and cannot borrow. Since, there are no annuity markets to insure against the risk of premature death, there is a positive flow of accidental bequests in each period. These are distributed in equal amount b among all agents alive in the economy.

The government taxes labor earnings and capital income to finance an exogenous stream of public expenditure and a pension system. Income from capital and labor income are taxed at flat rates τ_k and τ_l respectively. Retired agents receive a lump-sum pension p from the government until they die.

6.4 The household's problem

For any given period, a t-year old agent chooses consumption c and risk-free asset holdings for the next period a', as a function of the state variables $x = (t, a, y)$. The term t indicates the agent's age, a indicates current asset holdings of the agent, and y stands for the earnings

process realization. For given prices, the optimal decision rules are functions $(c(x)$ and $a'(x))$ that solve the dynamic programming problem described below.

(i) From age $t = 1$ to age $t = 36$ (from 25 to 60 years of age), the agent is working and has a probability of dying before the next period $1 - s_t$. The problem he solves is:

(2)
$$V(t, a, y) = \max_{c,a'}\left\{u(c) + \beta s_t E_t V(t + 1, a', y')\right\}$$
$$\text{s.t. } a' = \left[1 + r\left(1 - \tau_k\right)\right] a - c + (1 - \tau_l)wy + b, \quad a' \geq 0.$$

The evolution of y' follows the stochastic process described above. At every age t, y can lie in an age-specific grid y_t and its evolution towards y_{t+1} is determined by an age-specific transition matrix Q_{y_t}.

(ii) From t_r to T (from 61 to 86) agents no longer work and live off pensions and interest. Their value function satisfies:

(3)
$$W(t, a) = \max_{c,a'}\left\{u(c) + s_t \beta W(t + 1, a')\right\}$$
$$\text{s.t. } a' = \left[1 + r\left(1 - \tau_a\right)\right] a - c + p + b, \quad a' \geq 0.$$

The terminal period value function $W(T + 1, a)$ is set to equal 0 (agents do not derive utility from bequeathing).

The definition of equilibrium is in Appendix C.

6.5 The model calibration

A model period is one year. The coefficient of relative risk aversion is set to 1.5. The share of capital that goes to output is set to .36 and depreciation to 6%. All of these values are commonly used in the literature.

The discount rate β is calibrated to match a K/Y ratio of 3. and equals 0.955 in our benchmark, 0.954 in the NL PSID process, and 0.950 in the NL synthetic W2 process. We

also compute the results for the case in which all discount factors are the same and are equal to the average of the three (that is 0.953). We report the results for the case of equal betas in Appendix D. They show that our conclusions hold in this case as well.

The technology parameters generate an interest rate of 6% and a wage of 1 when the capital/output ratio equals 3.

The population growth rate n is set to 1.2% per year. The survival probabilities s_t are from Bell, Wade and Goss (1992). Government spending is 19% of GDP (g) (as in the the Council of Economic Advisors (1998) data), while capital taxes G (τ_a) are taken from Kotlikoff, Smetters and Walliser (1999), (see Table 4). The labor tax rate adjust to balance the government budget constraint.

Table 4: Calibration parameters

σ	A	α	δ	n	g	τ_k
1.5	0.895844	0.36	0.06	0.012	0.19	0.2

The social security pension benefit p equals 40% of the average earnings of a person in the economy (as in De Nardi (2004)).

7   The results from the model

7.1   The consumption implications

The model generates lifetime patterns of average consumption which are hump-shaped (Figure 4), as in the data for the US (Carroll and Summers, 1991) and the UK (Attanasio and Weber, 2010).

Note that given that we have imposed a common average income profile for the three earnings processes, any differences in the average consumption profiles stem from differences

Figure 4: Average of log consumption profiles
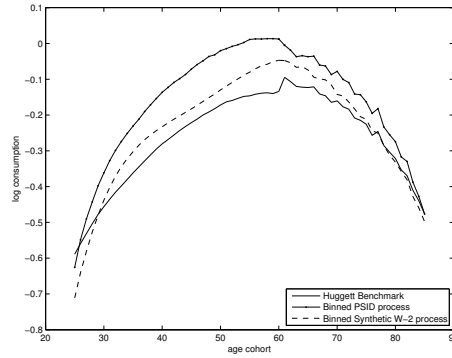
in precautionary saving due to differences in the riskiness of respective procesess. The average consumption profile has a higher initial level, grows at a slower rate and peaks at a lower level for the benchmark process, reflecting the lower riskiness of the process and therefore lower precautionary saving. Precautionary saving in mid-life (between age 45 and 60) seems also more important for the W2 process, reflecting the much higher level and rate of growth of earnings uncertainty it implies over those ages (see 3).

None of the earnings processes generates a drop in consumption at retirement which is consistent with the data (for instance, Bernheim, Skinner and Weinberg (2001) report an average consumption drop of 0.14 log points in the two years after retirement).[8] Instead, the benchmark displays a counterfactual increase in log consumption at retirement, which is related to the existence of a flat social security benefit that implies that low-earnings low-asset individuals experience an important rise in income after retiring. The empirical earnings processes do generate a drop in consumption, but it is very small with respect to the data (around 0.02 log points for W2 and 0.03 for PSID).

Figures 5, 6 and 7 plot the age profile of the variance of log earnings, of consumption and their covariance for, respectively, the benchmark, the NL PSID process, and the NL synthetic W2 process. The joint evolution of the three profiles is informative on the persistence of

---

[8]As is well known, accounting for the drop in consumption at retirement in life cycle models requires complementarities between work and consumption, Banks, Blundell and Tanner (1998), or changes in shopping time and food production, Aguiar and Hurst (2005)), or unexpected events which we do not model here.

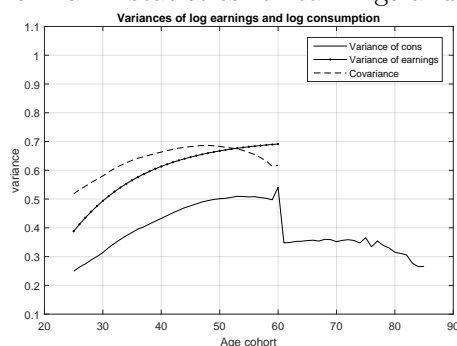Figure 5: Benchmark: statistics for earnings and consumption

Figure 6: NL PSID process: statistics for earnings and consumption

earnings risk at different stages of the life cycle. Intuitively, the growth in the variance of log earnings as individuals age may be reflecting increases in either permanent or transitory uncertainty, or both. Blundell and Preston (1998) show that consumption data can help disentangle the relative contribution of these two effects. For a given cohort at a given age, they identify the variance of the permanent shocks by the increase in the variance of log consumption. The variance of permanent shocks can also be identified by the increase in the covariance between log income and log consumption, which provides an over-identifying restriction. On the other hand, the increase in the variance of transitory shocks is identified by the difference between the increase in the variances of log income and log consumption for a given cohort at a certain age. If the variance of log income increases more than the variance of log consumption, the variance of transitory shocks is also increasing.

The life-cycle profiles depicted in Figures 5, 6 and 7 have the same qualitative character-

Figure 7: NL synthetic W2 process: statistics for earnings and consumption
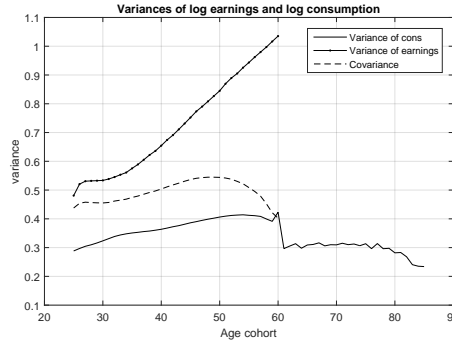
istics as the data reported in Blundell and Preston (1998). Since the variance of log earnings increases much more than the variance of log consumption, the variance of transitory shocks increases as cohorts age.

By construction, the benchmark earnings process has constant persistence and constant innovation variance from age 26 till retirement. The fanning out of the cross sectional earnings and consumption distribution in Figure 5 reflects the fact that the initial earnings draw at age 25 is drawn from a distribution which has a smaller variance than the limit stationary distribution of earnings that obtains when age diverges to infinity. Since the earnings process is highly persistent, individual consumption can only partially smooth the shocks over the lifetime and the covariance of log-consumption and income increases until age 50. From age 50 onwards, even highly persistent shocks are effectively transitory as retirement approaches the covariance falls with age.

The other two earnings processes do not impose constant persistence and innovation variance over the working lifetime. For this reason, the differential evolution of the consumption and income variance and covariance profiles is informative about the changing nature of earnings risk over the life cycle.

The profiles for the W2 earnings process in Figure 7 are qualitatively similar to those for the benchmark process above. The main difference is the dramatically higher rate of increase for the variance of earnings and the significantly lower rate for the variance of

26

consumption. This indicates the main share of the increase in income uncertainty over the lifetime reflect an increase in transitory uncertainty. This is confirmed by the much lower level, and rate of growth, in the consumption-earnings covariance. Furthermore, this covariance falls substantially more from age 50 until retirement suggesting an additional fall in the persistence of shocks beyond that stemming from approaching retirement. This reflects the much more transitory nature of earnings shocks in the W2 data. Comparing the NL synthetic W2 process with the NL PSID process, the variance of log consumtion at all ages is higher in the first case, but both show a similar profile. The covariance between log earnings and log consumption follows a similar pattern, while the variance of log earnings rises very significantly as cohorts age.

Finally, the profiles for the PSID earnings process in Figure 6 are quantitatively, and to some extent qualitatively, qsubstantially different from those for the other two processes. In particular, both the very low levels and slope of the profiles for the consumption variance and the consumption-earnings covariance indicate, in line with the mobility indicators in Tables 3, that the process implies highly transitory earning shocks.

Figure 8 depicts the variances of consumption implied by our three earnings processes and structural model, compared with some estimates from the literature. Namely, we include data from Heathcote, Storesletten and Violante (2014) (HSV) and data derived from Heathcote et al. (2010) (HPV). Both use CEX data and cohort fixed effects but differ in methodology, particularly in the equivalization procedure (family composition dummies in the former and OECD equivalence scale in the latter). We also include CEX data extracted from Aguiar and Hurst (2013) (AH), who use family composition dummies and control directly for cohort fixed effects and indirectly for year fixed effects (through a set of adjusted year dummies).

There are multiple estimates in the literature of the age profile for the variance of consumption (Deaton and Paxson (1994),Blundell and Preston (1998), Storesletten, Telmer and Yaron (2004), etc.), usually derived following different procedures. For instance,Storesletten

et al. (2004) reports variances once education has been controlled for. As Attanasio and Weber (2010) point out, the identification of age effects is complicated because it is not possible to separately identify age, time, and cohort effects from the data. Heathcote et al. (2010) provide results controlling first for year effects and then for cohort effects. The variance of log consumption increases more along the life-cycle in the latter case.
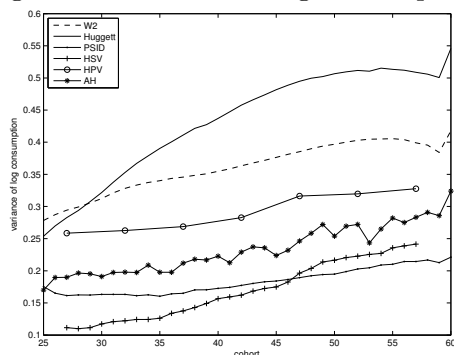
The benchmark earnings process fails to generate a realistic profile in the variances of log consumption. Its level is too high with respect to all the sources of data we report, and increases over the life-cycle by more than most of the estimates in the literature.

The NL PSID process is closer to the level of the variance of log consumption reported in the graph or, for example, to Storesletten et al. (2004). However, it fails to generate a sufficient increase in variance along the life-cycle.

The NL synthetic W2 process overestimates the variance of log consumption with respect to most of these data sources, but its slope is reasonably close to the profile in Heathcote et al. (2010) (HPV). In addition, even though the variance for younger agents is too high, the synthetic W2 process accurately replicates the life-cycle increase in variances of log consumption. It matches exactly the increase in the data reported in Heathcote et al. (2014) (HSV) and is very similar to the one in the other sources. A potential reason for this discrepancy is that, given that we focus on males earnings, we do not acknowledge the important insurance role of family labor supply (Blundell, Pistaferri and Saporta-Eksten (2015)).

In addition, the limitations of the survey data and the concerns about the quality of the CEX data, especially in recent years, raise the question of how close measured consumption is to the actual data. For instance, the CEX has been shown to display serious problems of non-classical error (Attanasio, Hurst and Pistaferri (2014)). When studying consumption patterns close to retirement, other surveys like the Health and Retirement Study (HRS) provide additional high quality data.

Figure 8: Variances of log consumption

7.2 Wealth distribution

We now compare the key statistics of the wealth distributions generated by our benchmark earnings process with those of our non-parametrically estimated earnings processeses[9].

Table 5: Wealth distribution statistics

Capital-output ratio	Bequests-output ratio	Wealth Gini	Percentage wealth in the top						Percentage with negative or zero wealth
			1%	5%	20%	40%	60%	80%	
U.S. data									
3.0	2.6%	.72	28	49	75	89	96	99	5.8–15.0
Benchmark									
3.0	2.8%	.72	11	35	74	93	98.98	100	19
NL PSID Process									
3.0	2.8%	.63	9	28	63	86	96	99.52	6
NL Synthetic W2 Process									
3.0	2.7%	.67	10	31	68	89	97	99.74	8

The benchmark calibration suffers from the limitations that have already been analyzed in the literature. Although it matches the wealth Gini coefficient we observe in the data, it does so by generating too many people with zero assets (the left tail is too fat) and not enough concentration at the top (the right tail is too thin). For instance, the number of

---

[9]Source for US data: De Nardi (2004), except for bequest-output ratio, which is taken from Villanueva (2005)

people with zero assets[10] is (19%), which is even higher than the highest estimate in the data, while the richest 1% only hold 11% of total wealth, compared with 28% in the data.

Both the NL PSID process and the NL synthetic W2 better match the left tail of the distribution. They generate a substantially smaller amount of agents with zero assets (6% and 8% respectively)) and better fit the holdings of the poorest 60% of people: in the distribution generated by the NL W2 process, the poorest 60% hold 11% of total net worth, which matches the actual data, and in the distribution derived from the NL PSID process they hold 14%, which is still much closer to the 11% observed in the data, compared to only 7% implied by the benchmark AR(1) earnings process. This is due to the two forces. First, a higher earnings mobility at low earnings levels than in the benchmark (as seen in Section 5.2) implies that individuals are able to hold a positive stock of wealth at some point without sacrificing too much consumption. Second, the high likelihood of falling out of a higher earnings state generates an extra savings motive.

The result of these forces can be seen in the left panel of Figure 9. As one might have expected, a large fraction of the zero-wealth individuals in the benchmark process economy are young agents that haven't had the opportunity to build up any assets. However, there is still a significant amount of individuals at the borrowing constraint at age 40, for example (see right panel). Instead, in the W2 and PSID processes most agents, even young, manage to build up a small amount of wealth. These striking implications of the earnings processes can be seen in the left panel of Figure 10 (which depicts the point of highest density in the wealth distribution by age) and the right panel in the same figure (median wealth by age). Under the benchmark earnings process, the mode of wealth is 0 for all ages and the agent with median wealth doesn't reach 50.000$ until he is 45 years old.

Even the more flexible and more realistic earnings processes do not help to better match the right tail of the distribution. Although Castañeda et al. (2003) have shown that a highly

---

[10]Given that all young agents are born at zero wealth by assumption, we do not include in any of the model computations of the fraction of agents at zero or nagative wealth.
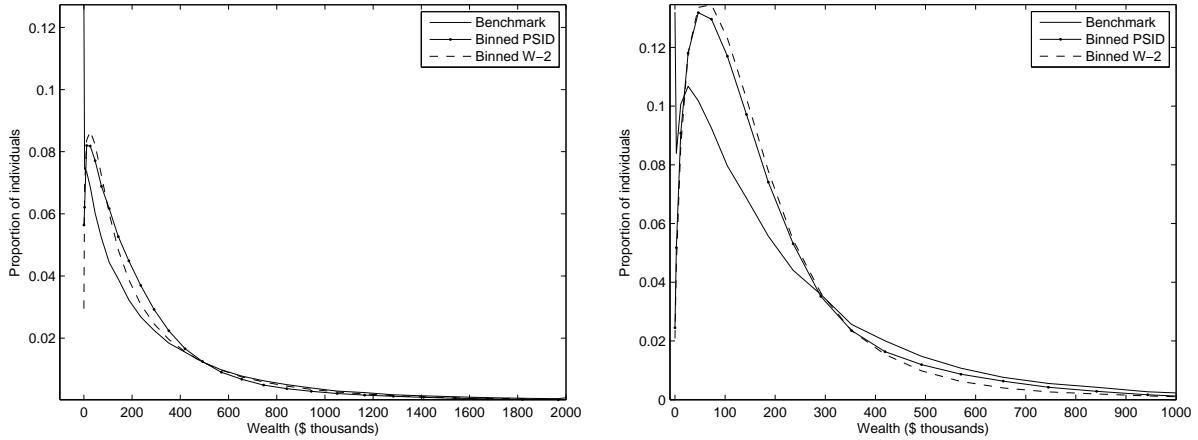
Figure 9: Left panel: Wealth distribution by earnings process. Right panel: Wealth distribution by earnings process: 40-year-olds
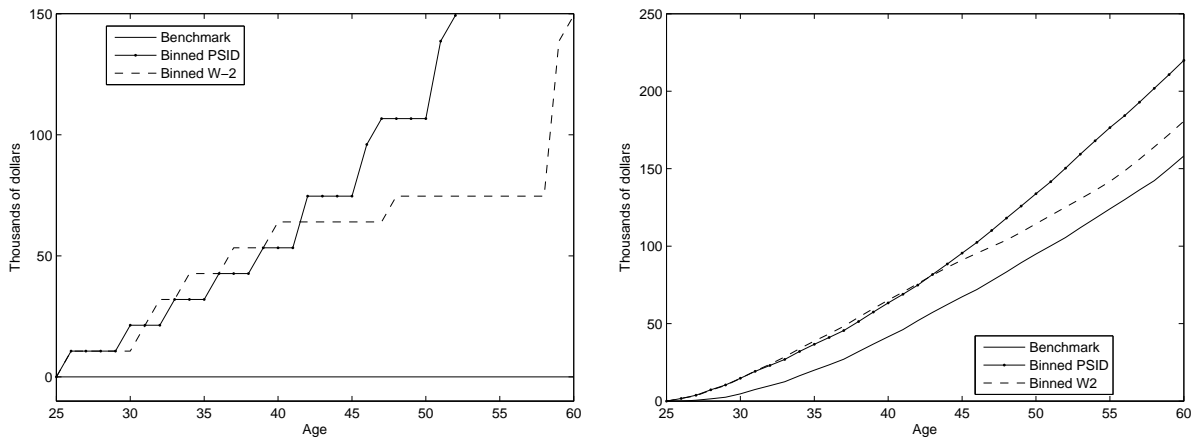

Figure 10: Left panel: Wealth mode. Right panel: Wealth median

skewed earnings process can generate a long right tail for the wealth distribution, it turns out that even the NL synthetic W2 process, which as shown earlier has a higher degree of earnings inequality and negative skewness than our NL PSID process, generates too small a wealth concentration amongst top earners.

On the positive side our empirical method, applied to synthetic W2 data, performs approximately as well as the benchmark in the right tail, outperforming it in the left tail. This is remarkable given that neither of our empirical processes have been calibrated to match any particular moments of the wealth distribution. The inclusion of intergenerational links and bequest transmission, as in De Nardi (2004), is likely to take these results closer to the

actual data, particularly in what respects to the right tail of the distribution. It should also be considered that we are not taking inter-vivos transfers into account, which may represent an additional source of wealth, particularly at earlier ages.

Figure 11: Wealth Ginis by cohort
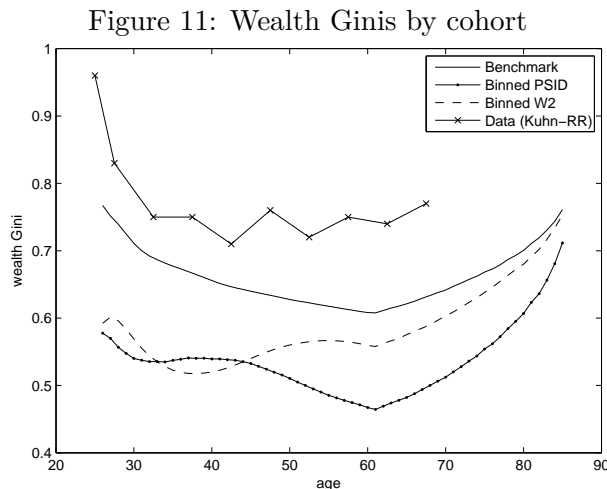
With respect to the Gini coefficient by age, the Huggett model delivers the largest wealth Gini at every age (Figure 11 - the source for the data (year 1989) is Kuhn and Ríos-Rull (2015)). However, as argued earlier, this is generated by pushing too many agents towards zero. Besides, it implies a very large drop in wealth Gini from ages 20 to 60. On the other hand, the W2 process generates a qualitatively reasonable wealth Gini profile, although levels are too low with respect to the data. This, too, can be due to many important mechanisms that our model abstracts from (bequests, entrepreneurship, inter-vivos transfers, etc.).

7.3 The wealth profiles over the life-cycle

The wealth profiles generated by this model (see Figure 12) display the same limitations in Huggett (1996): since agents do not have a bequest motive, nor face medical expenses, they run down their assets fast during retirement (see De Nardi, French and Jones (2010) for a discussion of the actual data). In the benchmark process, the 10 percent quantile barely accumulates savings at the time of retirement (just 3% of average yearly earnings). The

NL PSID process shows less persistence than the benchmark processes (as we have argued earlier in Section 5.2): the agents in the 10 percent quantile are able to save a positive amount (though 25 times lower than the top 5%). The amount they save is around one year's worth of average earnings in the economy. In the NL synthetic W2 process, there is a larger level of earnings inequality. This translates into more unequal wealth profiles, where the top 5% accumulates about 80 times more wealth on average than the 10 percent quantile, and where to find individuals with peak savings equivalent to one year of average earnings we need to get to the 25th percentile.

Figure 12: Wealth profiles for the 0.1, 0.3, 0.5, 0.8 and 0.9 quantiles
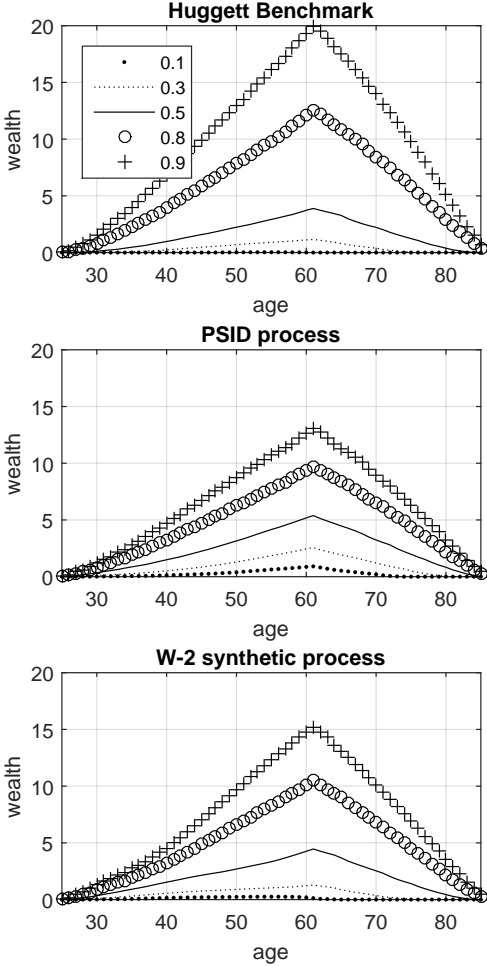
8 Compensating differentials

To evaluate how well individuals can self-insure in the Huggett AR(1) economy compared to the W2 economy, we compute compensating differentials.[11] More specifically, we compute λ by comparing the value functions at age 25 at the same points of the state space for the two economies as $V^H(25, \lambda, y_{25}) = V^{W2}(25, 0, y_{25})$, where y_{25} is earnings realization upon entering the labor market. We measure λ in units of average yearly earnings. Hence, λ is the one time asset transfer (expressed as a fraction of average yearly earnings) that we would have to give an individual in the Huggett world, with given initial current earnings and given age, to make him indifferent between the Huggett and W2 worlds.

We start by reporting the ex ante compensating variation λ^{ex} satisfying $\mathbb{E}_{y_{25}} V^H(25, \lambda^{ex}, y_{25}) = \mathbb{E}_{y_{25}} V^{W2}(25, 0, y_{25})$, than would make an individual indifferent between being born in either economy, under the veil of ignorance; i.e., before knowing the initial earnings realization. In our base case, the value that households attribute to between being born in the W2 world rather than in the AR(1) Huggett world, corresponds to a one-time asset compensation equal to one year of average earnings. This result, however, is due to the fact that we recalibrate the discount factors to match the same aggregate wealth and when the discount factor is lower, the value function, which is negative, is higher. Instead, when we hold the discount factor constant across earnings processes the value that agents attribute to being born in the W2 rather than in the Huggett world is equivalent a one time compensations of just 7% of average yearly earnings.

To better understand the effects of earnings risks across the earnings distribution, Figure 13 reports the compensating variation as a function of the initial earnings realization both the common β (left panel) and different β cases (right panel). In both cases, the lower the initial level of earnings the more individuals prefer to be born in the W2 economy, which features a substantially lower persistence of earnings shocks than the benchmark. In both

---

[11] For reasons of space constraints, the comparisons with the PSID economy are available from the authors upon request.

cases, very high-earnings individuals are better off under Huggett's earnings process due to the fact that the Huggett process displays a much higher persistence (the graph reports only positive compensations, all other earners would be better off in the AR(1) world). This is because the AR(1) Huggett process shows significantly less mobility, particularly at younger ages, and young people start out with little to no assets. Thus, self-insurance is more difficult for people with low earnings and low assets in presence of higher earnings persistence. What changes across the two panels is the cutoff level of earnings at which a household prefers to be born in the W2 economy. In the same β case, only individuals with less than average earnings prefer to be born in the W2 economy, while in the case with recalibrated patience, even households with two times average earnings would prefer to be born in the W2 economy.
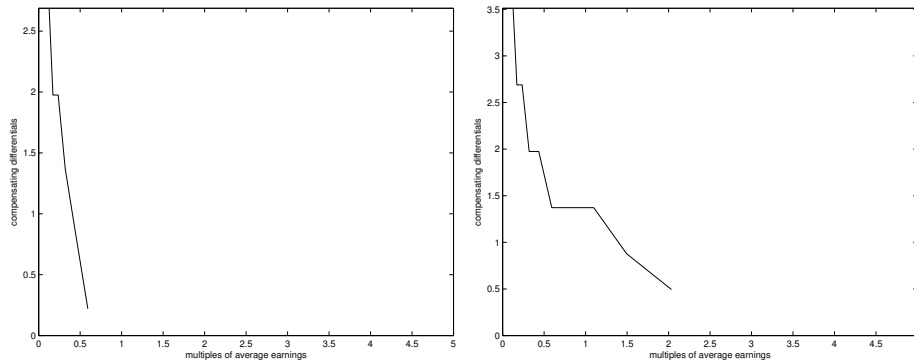
Figure 13: Compensating differentials. Left, same β; right, different β.

9 Heterogeneity in discount factors

Heterogeneity in time preference has been proposed as a factor that helps explain the large levels of wealth inequality observed in the data, together with the observed dispersion in wealth at retirement given lifetime earnings (Venti and Wise (2000)). Cagetti (2003) shows that matching mean wealth profiles in a life cycle model implies that both discount factors and coefficients of risk aversion differ significantly for different educational groups.

Krusell and Smith (1998) show that a limited degree of time preference heterogeneity dramatically improves the ability of a dynastic Bewley model to match the right tail of the wealth distribution. Conversely, Hendricks (2007) has shown that even a substantial degree of preference heterogeneity does improve significantly the fit of the right tail of the wealth distribution in a life-cycle model with persistent earnings shocks. Comparing his finding to those of Krusell and Smith (1998) he points out that in his enviroment the elasticity of saving to impatience is limited due to the fact that the bulk of saving reflects life-cycle considerations and self-insurance against large and persistent earnings shocks. Viceversa, in (Krusell and Smith, 1998) environment there is no life-cycle motive and there is little precautionary saving as earnings shocks are small and not very persistent. Since the main motive for saving is the difference between the rate of return and the time preference rate, time preference heterogeneity has quantitatively important effects on the wealth distribution. In the light of (Hendricks, 2007) findings that the size and persistence of earnings shocks determines the response of top wealth inequality to the introduction of time preference heterogeneity, it seems relevant to investigate its interaction with our W2 earnings process.

To account for some preference heterogeneity in the context of the W2 earnings process economy, we also study several configurations of preference heterogeneity (or distributions of discount factors), which we take from Hendricks (2007) and Krusell and Smith (1998). Table 6 reports the results. In the Hendricks' case, discount rates are placed on a grid

$$\bar{\beta} \cdot [0.94, 0.98, 1, 1.02, 1.06]$$

with stationary distribution across gridpoints (0.431, 0.081, 0.199, 0.184, 0.105). In the Krusell-Smith case, discount rates lie on a grid

$$\bar{\beta} \cdot [0.997, 1, 1.003]$$

with stationary distribution (0.1, 0.8, 0.1).

Table 6: Wealth distribution statistics (preference heterogeneity)

Capital-output ratio	Wealth Gini	Percentage wealth in the top						Percentage with negative or zero wealth
		1%	5%	20%	40%	60%	80%	
U.S. data								
3.0	.72	28	49	75	89	96	99	5.8–15.0
NL PSID Process - no heterogeneity)								
3.0	.63	9	28	63	86	96	99.52	6
NL PSID Process with heterogeneity (Hendricks)								
3.0	.72	11	35	75	93	98.56	99.91	12.6
NL PSID Process with heterogeneity (Krusell and Smith)								
3.0	.63	9	29	64	86	96	99.52	6
NL Synthetic W2 Process (no heterogeneity)								
3.0	.67	10	31	68	89	97	99.74	8
NL Synthetic W2 Process (Hendricks)								
3.0	.74	12	38	77	94	98.85	99.92	13.3
NL Synthetic W2 Process with heterogeneity (Krusell and Smith)								
3.0	.66	9	31	68	89	97	99.74	7.5

Both distributions of discount factors have little to no effect on the wealth holding of the richest, a result that is consistent with the one in Hendricks (2007) in presence of an AR(1) earnings process.

10   Conclusions

Our main findings are the following. First, both our PSID and W2 non-parametric earnings processes imply much less persistence at low and high earnings levels than those implied by the discretized AR(1) process used by Huggett (1996), which is very similar to those commonly used in the quantitative macro literature with heterogeneous agents. Second, both the PSID and the W2 earnings processes, when introduced in a standard quantitative model of consumption and savings over the life cycle, generate a much better fit of the wealth holdings of the bottom 60% of people, but vastly underestimate the level of wealth concentration at the top of the wealth distribution, just as the the model with the discretized AR(1) earn-

ings process. This is a frequent drawback of life-cycle models with no entrepreneurship or transmission of bequests (De Nardi, 2004), (Cagetti and De Nardi, 2006), and (Cagetti and De Nardi, 2009) and we show that it still holds in presence of reasonable heterogenity in household's patience. Third, but very importantly, the smaller persistence of earnings at low earnings levels implies that low-earnings, low-wealth households, and especially young ones, have an easier time to self-insure against earnings shocks under the PSID and W2 earnings processes. In contrast, older households with high earnings but low wealth are more exposed to earnings risk because they also face higher earnings mobility. These findings on the persistence of the shocks and their implications at different ages, earnings, and asset levels are important and likely to have a large effects on the costs and benefits of many government transfers programs, such as unemployment, insurance Supplemental Social Insurance (SSI), and redistributive taxation.

References

Aguiar, Mark and Bils, Mark (2015), 'Has consumption inequality mirrored income inequality?', *American Economic Review* **105**(9), 2725–2756.

Aguiar, Mark and Hurst, Erik (2005), 'Consumption versus expenditure', *Journal of Political Economy* **113**(5), 919–948.

Aguiar, Mark and Hurst, Erik (2013), 'Deconstructing life cycle expenditure', *Journal of Political Economy* **121**(3), 437–492.

Altonji, Joseph G., Smith, Anthony A. and Vidangos, Ivan (2013), 'Modeling earnings dynamics', *Econometrica* **81**(4), 1395–1454.

Arellano, Manuel, Blundell, Richard and Bonhomme, Stéphane (2015), Earnings and consumption dynamics: A non-linear panel data framework. unpublished manuscript.

Atkinson, Anthony Barnes, Bourguignon, François and Morrisson, Christian (1992), *Empirical studies of earnings mobility*, Hardwood, Chur.

Attanasio, Orazio and Pistaferri, Luigi (2014), 'Consumption inequality over the last half century: some evidence using the new psid consumption measure', *The American Economic Review* **104**(5), 122–126.

Attanasio, Orazio, Hurst, Erik and Pistaferri, Luigi (2014), The evolution of income, consumption, and leisure inequality in the us, 1980-2010, *in* C.Carroll, T.Crossley and J.Sabelhaus, eds, 'Improving the Measurement of Consumer Expenditures', University of Chicago Press.

Attanasio, Orazio P. and Weber, Guglielmo (2010), 'Consumption and saving: Models of intertemporal allocation and their implications for public policy', *Journal of Economic Literature* **48**, 693–751.

Banks, James, Blundell, Richard and Tanner, Sarah (1998), 'Is there a retirement-savings puzzle?', *The American Economic Review* **88**(4), 769–788.

Bell, Felicitie C., Wade, Alice H. and Goss, Stephen C. (1992), 'Life tables for the United States Social Security Area: 1900-2080', (Social Security Administration, Office of the Actuary).

Bernheim, B. Douglas, Skinner, Jonathan and Weinberg, Steven (2001), 'What accounts for the variation in retirement wealth among U.S. households?', *The American Economic Review* **91**(4), 832–857.

Blundell, Richard (2014), 'Income Dynamics and Life-cycle Inequality: Mechanisms and Controversies', *The Economic Journal* **124**(576), 289–318.

Blundell, Richard and Preston, Ian (1998), 'Consumption inequality and income uncertainty', *The Quarterly Journal of Economics* **113**(2), 603–640.

Blundell, Richard, Graber, Michael and Mogstad, Magne (2015), 'Labor income dynamics and the insurance from taxes, transfers and the family', *Journal of Public Economics* **127**, 58–73.

Blundell, Richard, Pistaferri, Luigi and Preston, Ian (2008), 'Consumption inequality and partial insurance', *The American Economic Review* pp. 1887–1921.

Blundell, Richard, Pistaferri, Luigi and Saporta-Eksten, Itay (2015), 'Consumption inequality and family labor supply', *American Economic Review* p. Forthcoming.

Browning, Martin, Ejrnaes, Mette and Álvarez, Javier (2010), 'Modelling income processes with lots of heterogeneity', *Review of Economic Studies* **77**, 1353–1381.

Budria, Santiago, Diaz-Gimenez, Javier, Quadrini, Vincenzo and Rios-Rull, Jose-Victor

(2002), 'New facts on the distribution of earnings, income and wealth in the U.S.', *Federal Reserve Bank of Minneapolis Quarterly Review* **26**, 2–35.

Cagetti, Marco (2003), 'Wealth accumulation over the life cycle and precautionary savings', *Journal of Business and Economic Statistics* **21**(3), 339–353.

Cagetti, Marco and De Nardi, Mariacristina (2006), 'Entrepreneurship, frictions and wealth', *Journal of Political Economy* **114**(5), 835–870.

Cagetti, Marco and De Nardi, Mariacristina (2008), 'Wealth inequality: Data and models', *Macroeconomic Dynamics* **12**(S2), 285–313.

Cagetti, Marco and De Nardi, Mariacristina (2009), 'Estate taxation, entrepreneurship, and wealth', *American Economic Review* **99**(1), 85–111.

Carroll, Christopher D. and Summers, Lawrence H. (1991), Consumption growth parallels income growth: Some new evidence, *in* B. DouglasBernheim and John B.Shoven, eds, 'National Saving and Economic Performance', University of Chicago Press, pp. 305–348.

Carroll, Christopher D, Hall, Robert E and Zeldes, Stephen P (1992), 'The buffer-stock theory of saving: Some macroeconomic evidence', *Brookings papers on economic activity* pp. 61–156.

Castañeda, Ana, Díaz-Giménez, Javier and Ríos-Rull, José-Víctor (2003), 'Accounting for the U.S. Earnings and Wealth Inequality', *Journal of Political Economy* **111**(4), pp. 818–857.

Council of Economic Advisors (1998), 'Economic Report of the President', (Washington: United States Government Printing Office).

De Nardi, Mariacristina (2004), 'Wealth inequality and intergenerational links', *Review of Economic Studies* **71**, 743–768.

De Nardi, Mariacristina (2015), 'Quantitative models of wealth inequality: A survey', National Bureau of Economic Research working paper no. 21106.

Deaton, Angus and Paxson, Christina (1994), 'Intertemporal choice and inequality', *Journal of Political Economy* **102**(3), 437–67.

De Nardi, Mariacristina, French, Eric and Jones, John B. (2010), 'Why do the elderly save? The role of medical expenses', *Journal of Political Economy* **118**(1), 39–75.

Fisher, Jonathan D, Johnson, David S and Smeeding, Timothy M (2013), 'Measuring the trends in inequality of individuals and families: Income and consumption', *The American Economic Review* **103**(3), 184–188.

Geweke, John and Keane, Michael (2000), 'An empirical analysis of earnings dynamics among men in the psid: 1968-1989', *Journal of Econometrics* **96**, 293–356.

Graber, Michael and Lise, Jeremy (2015), 'Labor market frictions, human capital accumulation and consumption inequality', mimeo, University College London.

Guvenen, Fatih, Karahan, Fatih, Ozkan, Serdar and Song, Jae (2015), What do data on millions of U.S. workers reveal about life-cycle earnings risk? Working paper, University of Minnesota.

Haider, S. J. (2001), 'Earnings instability and earnings inequality of males in the united states: 1967-1991', *Journal of Labor Economics* **19**(4), 799–836.

Hansen, Gary D. (1993), 'The cyclical and secular behaviour of labour input: Comparing efficiency units and hours worked', *Journal of Applied Econometrics* **8**(1), 71–80.

Heathcote, Jonathan, Perri, Fabrizio and Violante, Giovanni L (2010), 'Unequal we stand: An empirical analysis of economic inequality in the united states, 1967-2006', *Review of Economic dynamics* **13**(1), 15–51.

Heathcote, Jonathan, Storesletten, Kjetil and Violante, Giovanni L (2014), 'Consumption and labor supply with partial insurance: An analytical framework', *American Economic Review* **104**(7), 2075–2126.

Hendricks, Lutz (2007), 'How important is discount rate heterogeneity for wealth inequality?', *Journal of Economic Dynamics and Control* **31**(9), 3042–3068.

Huggett, Mark (1996), 'Wealth distribution in life-cycle economies', *Journal of Monetary Economics* **38**(3), 469–494.

Jäntti, Markus and Jenkins, Stephen P. (2015), Income mobility, *in* Anthony B.Atkinson and FrançoisBourguignon, eds, 'Handbook of Income Distribution', Vol. 2 of *Handbook of Income Distribution*, Elsevier, pp. 807–935.

Kotlikoff, Laurence J., Smetters, Kent and Walliser, Jan (1999), 'Privatizing social security in the united states - comparing the options', *Review of Economic Dynamics* **2**(3), 532–574.

Krusell, Per and Smith, Anthony A. (1998), 'Income and wealth heterogeneity in the macroeconomy', *Journal of Political Economy* **106**(5), 867–896.

Kuhn, Moritz and Ríos-Rull, José-Víctor (2015), '2013 update on the us earnings, income and wealth distributional facts: a view from macroeconomics', Working paper.

Lillard, Lee A. (1977), 'Inequality: earnings vs. human wealth', *American Economic Review* **67**, 42–53.

Lillard, Lee A. and Willis, Robert J. (1978), 'Dynamic aspects of earning mobility', *Econometrica* **46**(5), 985–1012.

Meghir, Costas and Pistaferri, Luigi (2004), 'Income variance dynamics and heterogeneity', *Econometrica* **72**(1), 1–32.

Meghir, Costas and Pistaferri, Luigi (2011), 'Earnings, consumption and life cycle choices', *Handbook of Labor Economics* **4**, 773–854.

Meyer, Bruce D and Sullivan, James X (2010), 'Consumption and income inequality in the us since the 1960s', University of Chicago manuscript.

Quadrini, Vincenzo and Ríos-Rull, José-Víctor (1997), 'Understanding the US distribution of wealth', *Federal Reserve Bank of Minneapolis Quarterly Review* **21**(2), 22–36.

Quadrini, Vincenzo and Ríos-Rull, José-Víctor (2014), Inequality in macroeconomics, *in* A.B.Atkinson and F.Bourguignon, eds, 'Handbook of Income Distribution, vol. 2B', Elsevier, chapter 15.

Shorrocks, Antony F. (1980), Income stability in the United States, *in* N.Klevmarken and J.Lybeck, eds, 'The statics and dynamics of income', Tieto Ltd.

Storesletten, Kjetil, Telmer, Chris I. and Yaron, Amir (2004), 'Cyclical dynamics in idiosyncratic labor-market risk', *Journal of Political Economy* **112**, 695–717.

Tauchen, George (1986), 'Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions', *Economics Letters* **20**(2), 177–181.

Tauchen, George and Hussey, Robert (1991), 'Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models', *Econometrica* **59**(2), 371–96.

Venti, Steven F. and Wise, David A. (2000), 'Choice, chance, and wealth dispersion at retirement', No. w7521. National Bureau of Economic Research.

Villanueva, Ernesto (2005), 'Inter vivos transfers and bequests in three OECD countries.', *Economic Policy* **20**(43), 506–565.

Wolff, Edward N. (1987), 'Estimates of household wealth inequality in the US, 1962-1983', *Review of Income and Wealth* **33**(3), 231–256.

A   Appendix for online publication only: The PSID and synthetic W2 data

A.0.1   PSID data

The Panel Study of Income Dynamics (PSID) follows a large number of US households over time and reports information about their demographic characteristics and sources of income. The PSID was initially composed of two major subsamples. The first of them, the SRC (Survey Research Center) or core subsample, was designed to be representative of the US population and is a random sample itself. The second, the SEO (Survey of Economic Opportunity) subsample, was created to study the characteristics of the most deprived households. Later, Immigrant and Latino subsamples were also added to the PSID.

From 1968 to 1997, the survey was yearly. After 1997, it started having a biennial structure. Since the model period is one year, we restrict ourselves to the yearly part of the survey (1968-1997). We only consider the SRC or core subsample, since the SEO oversamples the poor. Obtaining a more accurate picture of the wealth or income-poor is not particularly helpful for our purposes, since the poor account for a small proportion of the overall wealth. Furthermore, after dropping the SEO and Latino samples we are left with a random sample, which makes computations simpler since weights are not needed Haider (2001) [12].

Because our binning strategy requires a large amount of data per age group, we drop individuals still working after age 60 and we assume that people retire at 60. For comparability with the W2 synthetic dataset that we derive from the econometric processes in Guvenen et al. (2015), that start age age 25, we also drop individuals below age 25.

In line with previous literature, we restrict our study to male heads of household. A valid observation is formed of two values for individual labor income: this year's and next

---

[12]In our case, we are creating a dataset of pairs of income observations from the PSID panel. It must be taken into account that the weighting of our final dataset can be affected by attrition and by the fact that we are neglecting observations of yearly income under 900 $

year's. Neither of the two income values can be missing. We deflate labor income by CPI-U (taken from FRED economic data, where index 100=1982-1984). We do not consider earnings observations with values under 900$. This avoids the problem of taking logs of very small numbers and is in line with usual practice in previous literature (De Nardi (2004) or Guvenen et al. (2015), for instance).

Labor income does not include unemployment or workers' compensation, nor earnings of the spouse, because adding them would reduce the sample period to 1977-1992, due to data availability issues.

The major shortcomings the PSID data are that the sample size is not very large (and the earnings-rich are not oversampled), that they are less likely to respond to surveys, and that earnings are top-coded. As a result, our approximation of the right tail of the income distribution is not well estimated (Budria, Diaz-Gimenez, Quadrini and Rios-Rull (2002)). This is an important limitation because the earnings risk faced by the high earners is potentially an important determinant of their saving behavior and wealth concentration in the upper tail (Castañeda et al. (2003)).

A.0.2   W2 synthetic data

The original data used by Guvenen et al. (2015) come from the Master Earnings File (MEF) of the U.S. Social Security Administration (SSA) records and contain information for every individual in the US who has ever been issued a Social Security number. They also include some demographic information (sex, race, date of birth...) and earnings from the employee's W2 forms, that U.S. employers are legally required to report to the SSA since 1978. Labor earnings are measured annually, and they include all wages, salaries, bonuses and exercised stock options. The data are uncapped. Therefore, due to their very nature, the W2 tax data, in contrast to the PSID data, does not suffer from small sample size, under-sampling of the rich, differential survey response, or top-coding.

Unfortunately, the original W2 data are not publicly available, but Guvenen et al. (2015) have estimated a flexible, and rather complex, parametric process using the simulated method of moments and shown that it fits the data well. We use their estimates to simulate a synthetic W2 data set. More specifically, the parameterization that we simulate is their "benchmark model" reported in column (2) of Table III in their paper. The model assumes that the stochastic component of earnings is the sum of: (i) an individual-specific linear time trend (heterogeneous income profile); (ii) a mixture of AR(1) processes where each component receives an innovation in a given year with age-dependent probability; and (iii) an i.i.d. transitory shock.

Let $\tilde{y}_t^i$ denote the stochastic component of earnings for individual i at age $t = 1, 2, \ldots$. The specification of the parametric model is:

$$\tilde{y}_t^i = (\alpha^i + \beta^i t) + y_t^i$$

$$y_t^i = z_t^i + x_t^i + \nu_t^i + \varepsilon_t^i$$

$$z_t^i = .259 z_{t-1}^i + \eta_{zt}^i, \quad z_0^i \sim N(0, .133)$$

$$x_t^i = .425 x_{t-1}^i + \eta_{xt}^i, \quad x_0^i \sim N(0, .089)$$

$$\nu_t^i = \nu_{t-1}^i + \eta_{\nu t}^i, \qquad \nu_0^i = 0$$

$$\varepsilon_t^i \sim N(0, .029).$$

The vector (α, β) has a bivariate normal distribution with zero mean and the covariance matrix implied by $\sigma_\alpha = .552$, $\sigma_\beta \times 10 = .13$ and $\mathrm{corr}_{\alpha\beta} \times 100 = -.49$. The innovations $\eta_{jt}^i, j = z, x, v$ to the three AR(1) components are mutually exclusive with respective probabilities p_j. More formally, $\eta_{jt}^i = \eta_{jt}^{*i} \times \mathbb{I}\{s_{it} \in I_{p_j}\}$ with s_{it} a standard uniform random variable. The indicator function $\mathbb{I}\{s_{it} \in I_{p_j}\}$ takes value one if s_{it} falls in the interval I_{p_j}, with $I_{p_z} = [0, p_z]$, $I_{p_x} = (p_z, p_z + p_x]$, $I_{p_\nu} = (p_z + p_x, 1]$, where $p_z + p_x \leq 1$ and $p_\nu = 1 - p_z - p_x$. The (mixing) probabilities depend on age and on the previous realization of the stochastic component of earnings according to

$$p_j(y_{t-1}^i, t) = .067 + .038 \times \frac{t}{10} - .203 \times y_{t-1}^i + .071 \times \frac{t}{10} \times y_{t-1}^i$$

and are truncated when they fall outside the $[0, 1]$ interval.

Finally, the innovations to the AR(1) components have distribution $\eta_{jt}^{*i} \sim N(\mu_j, \sigma_j^i)$, with individual-specific standard deviation distributed according to $\log \sigma_j^i \sim N(\bar{\sigma}_j - \frac{\sigma_{jj}^2}{2}, \sigma_{jj}), j = z, x$ and $\sigma_\nu^i = \sigma_\nu$. The parameter values are $\mu_z = -.426$, $\mu_x = .021$, $\mu_\nu = .060$, $\bar{\sigma}_z = .847$, $\bar{\sigma}_x = .361$, $\sigma_{zz} = .107$, $\sigma_{xx} = 1.94$ and $\sigma_\nu = .087$.

Our synthetic W2 panel is obtained by simulating the above process to generate a panel of 500,000 individuals histories between ages 25 and 60.

B   Appendix for online publication only: How well does our non-parametric estimation of earnings processes fit the PSID and W2 data?

Because we allowed for a different number of earnings bins for the two data sets (18 for the PSID and 103 for the synthetic W2 data set), here we also report results for the case in which we allow for a smaller bin size in the synthetic W2 data set (18 points) so that the results across the two data sets can be compared keeping the coarseness of the earnings binning constant.

We first show results for the unconditional moments of log earnings by age and then for conditional moments of earnings growth, also by age.

We do not report results for the unconditional first moment because, since the data is detrended, it is an average of OLS residuals and therefore mean zero by definition.

Both the 18-gridpoint and 103-gridpoint discretizations successfully replicate unconditional variances over the life cycle and generate a negative skew and high kurtosis for log earn-

Figure 14: W2: Unconditional moments of log earnings
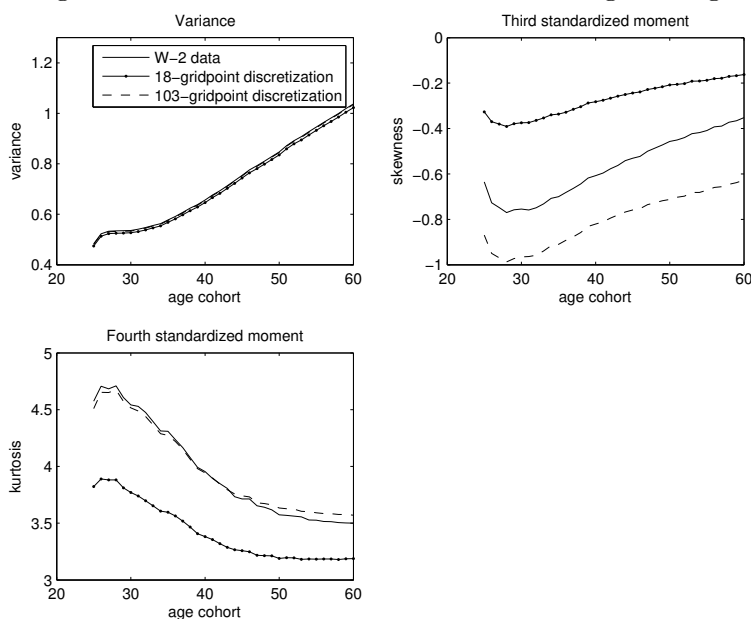
ings (Figure 14). The 18-gridpoint discretization understimates both the negative skew and the kurtosis of the log earnings distribution, while the 103-gridpoint discretization matches very closely the kurtosis profile, but overestimates the negative skew. The comparison of our NL PSID process with PSID data yields very similar results (see Figure 15).

In order to analyze the conditional moments of log earnings and for comparability with Guvenen et al. (2015), we group individuals in 5-year cohorts. Still, our graphs are not fully comparable to Section 3 of their paper, since we condition on the earnings of the previous year instead of the earnings of the previous 5 years. As in their paper, log earnings growth is defined as the difference between detrended earnings in two consecutive years.

Now the division in quantiles now serves two purposes. As earlier, it defines the size of the bins, but now it also determines the conditioning variable for the computation of the moments. In the "103-gridpoint discretization", we compute all moments conditional on each of the 103 bins in which we divide the overall population, while in the "18-gridpoint discretization" they are conditional on the 18 bins.

Figure 15: PSID: unconditional moments of log earnings
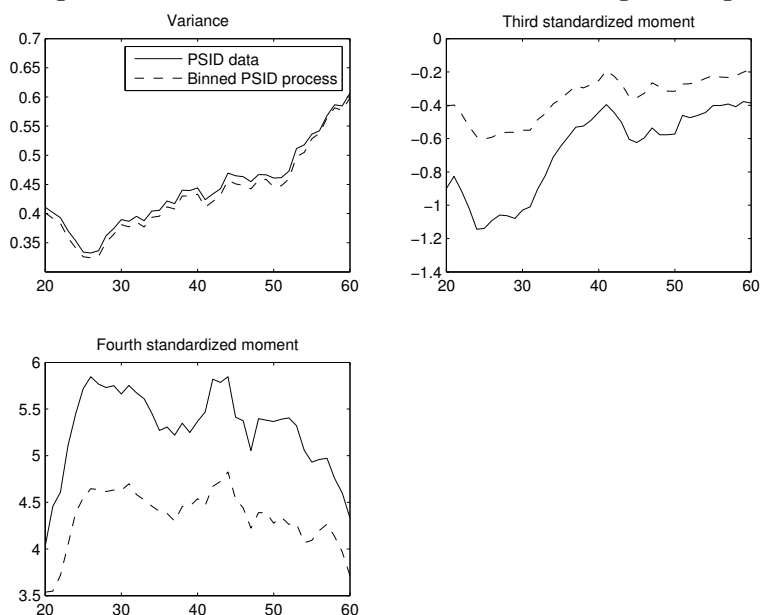
We first consider the conditional mean (Figures 16, 17 and 18). Both discretizations replicate the fact that individuals who are situated in a low earnings quintile can expect larger earnings increases.

With respect to the conditional variance (Figures 19, 20 and 22), the empirical process is also successful. For clearer illustration, we have included how our processes replicate the standard deviation we observe for individuals at ages 35-40 (Figure 21).

With respect to the third standardized moment, all discretizations closely match the negative skew that we observe in the data and the fact that it is larger for individual with higher previous earnings (Figures 23 24, 26). This can also be seen in the example in Figure 25.

Finally, both synthetic processes (see Figures 27, 28, 29, 30) match the high kurtosis in the data, with just a slight underestimation in the 18-gridpoint case.

Therefore, we conclude that our procedure generates a reasonably good approximation up to the fourth moment of the actual earnings process, particularly for discretizations with a large number of gridpoints, when applied to a sufficiently large dataset. This represents a

50

Figure 16: 18-gridpoint discretization: conditional mean of earnings growth
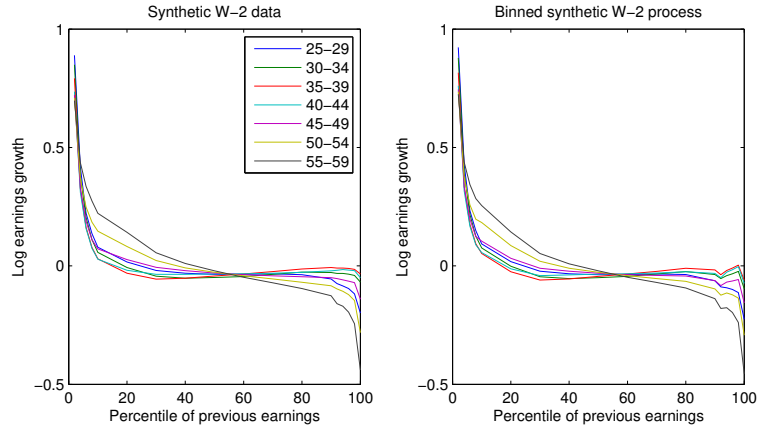
Figure 17: 103-gridpoint discretization: conditional mean of earnings growth
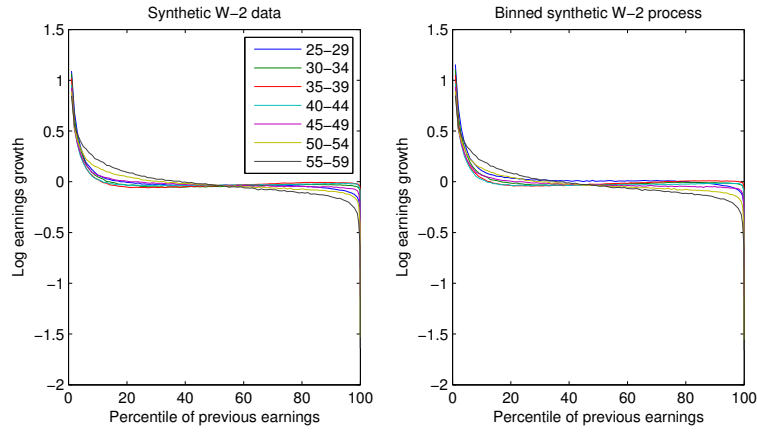
Figure 18: PSID: conditional mean of earnings growth

51
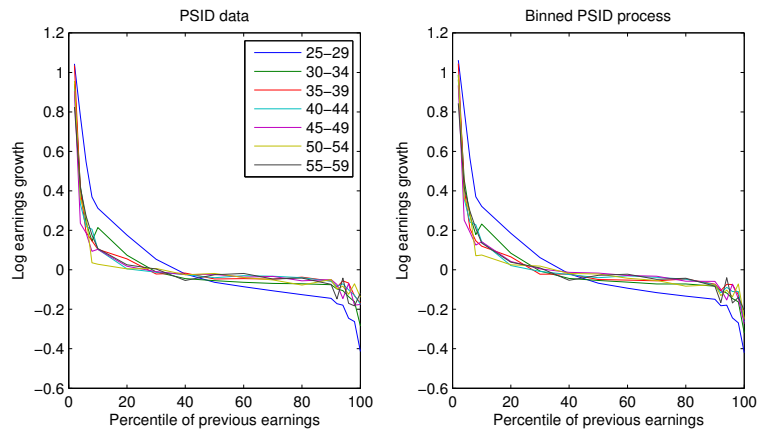
Figure 19: 18-gridpoint discretization: conditional standard deviation of earnings growth
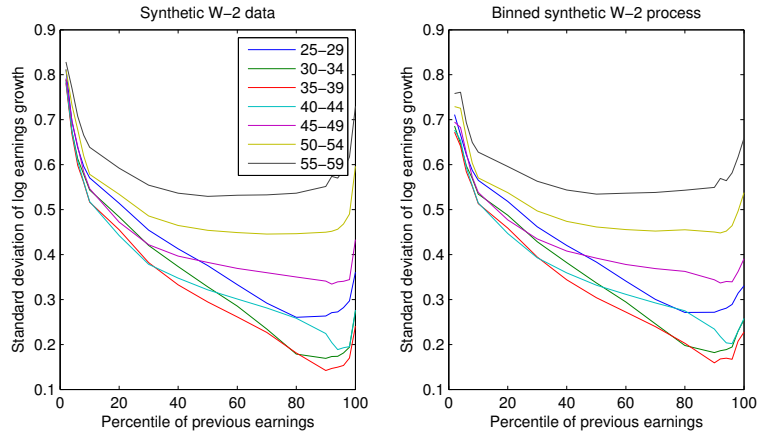
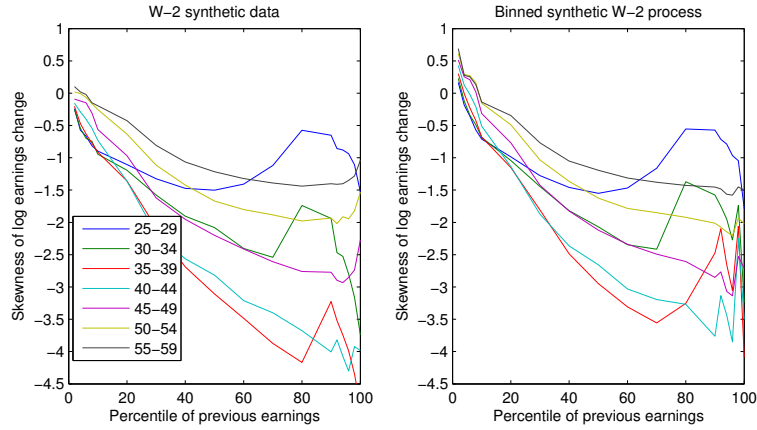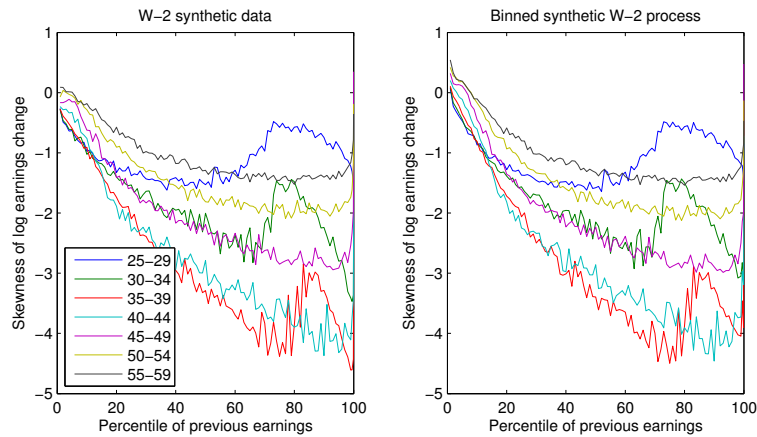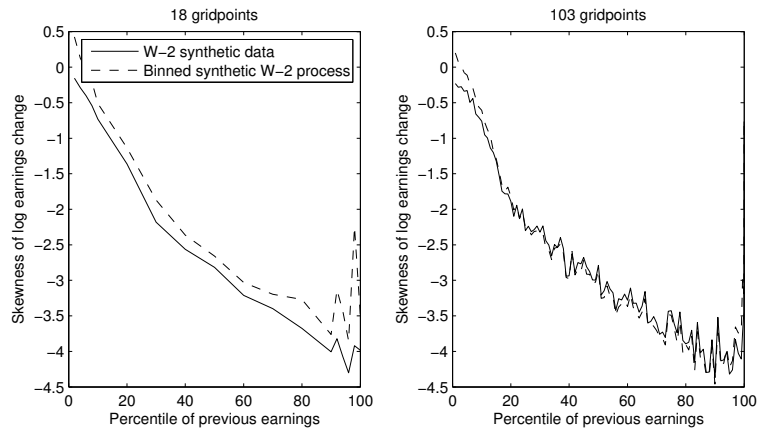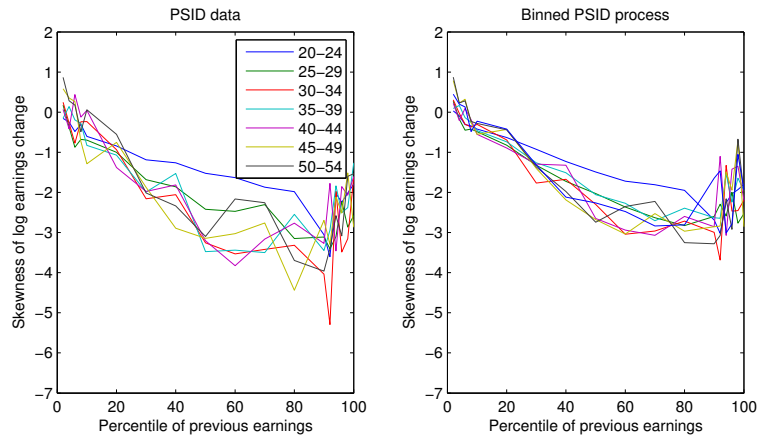Figure 20: 103-gridpoint discretization: conditional standard deviation of earnings growth
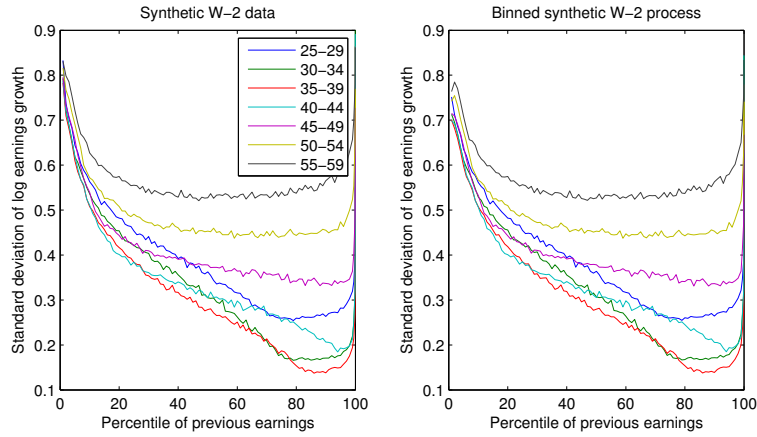
Figure 21: Conditional standard deviation for 35-40 year-old cohort
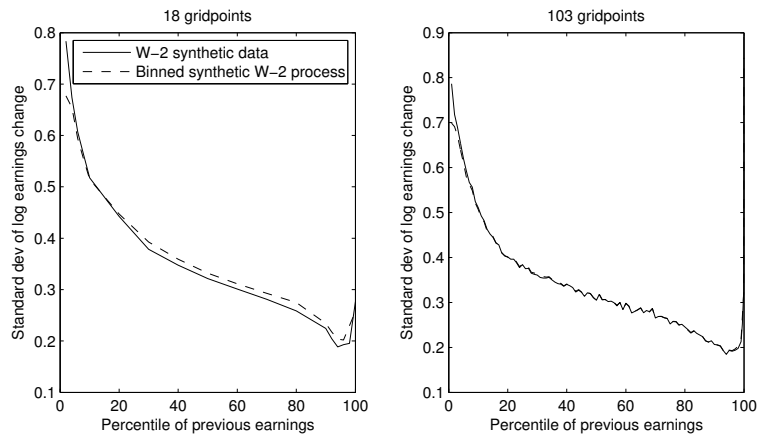
Figure 22: PSID: conditional standard deviation of earnings growth
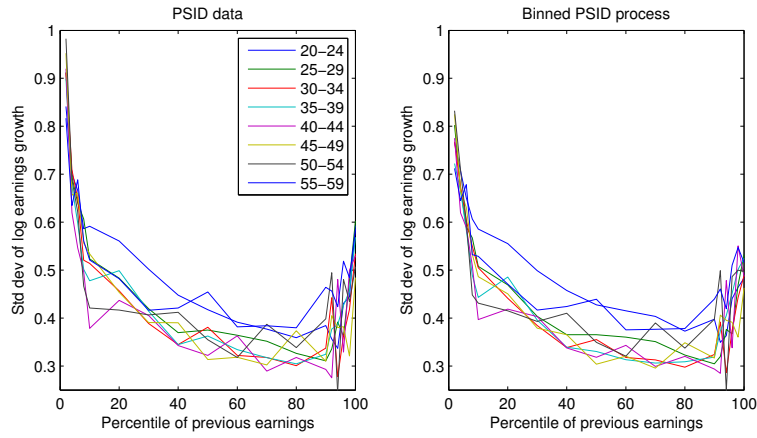
Figure 23: 18-gridpoint discretization: conditional third moment of earnings growth

Figure 24: 103-gridpoint discretization: conditional third moment of earnings growth

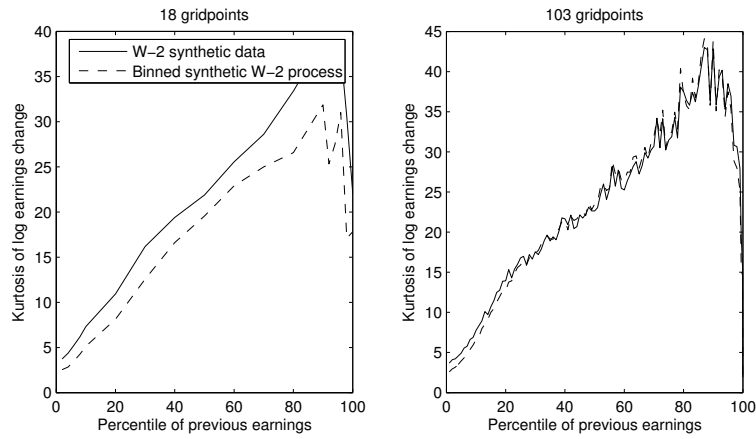Figure 25: Conditional third moment for 35-40 year-old cohort

Figure 26: PSID: conditional third moment of earnings growth

Figure 27: 18-gridpoint discretization: conditional fourth moment of earnings growth

54
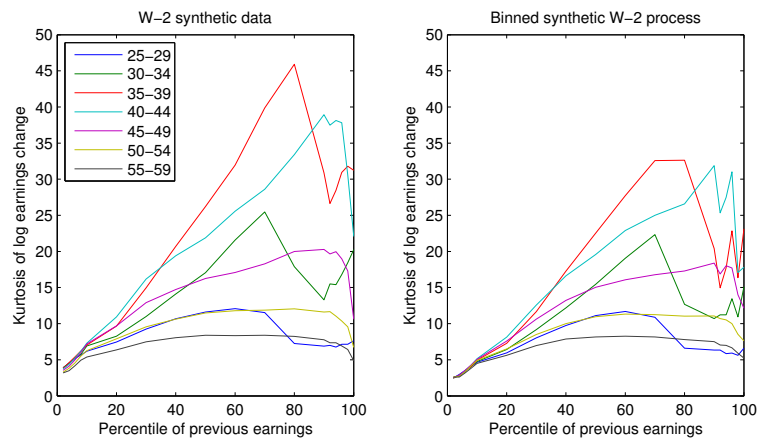
Figure 28: 103-gridpoint discretization: conditional fourth moment of earnings growth
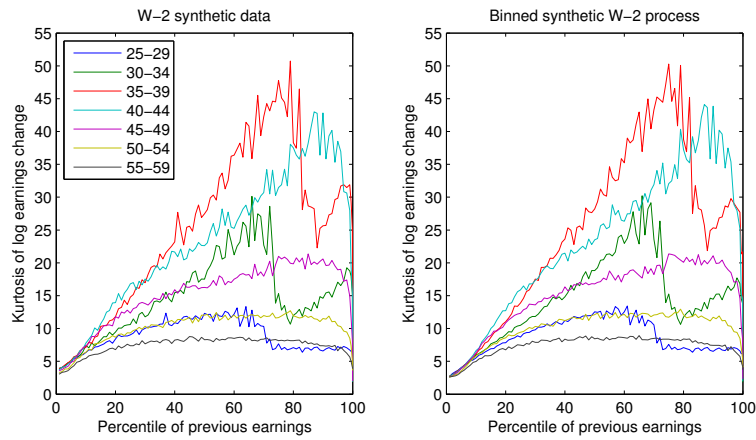
Figure 29: Conditional fourth moment for 35-40 year-old cohort

Figure 30: PSID: conditional fourth moment of earnings growth
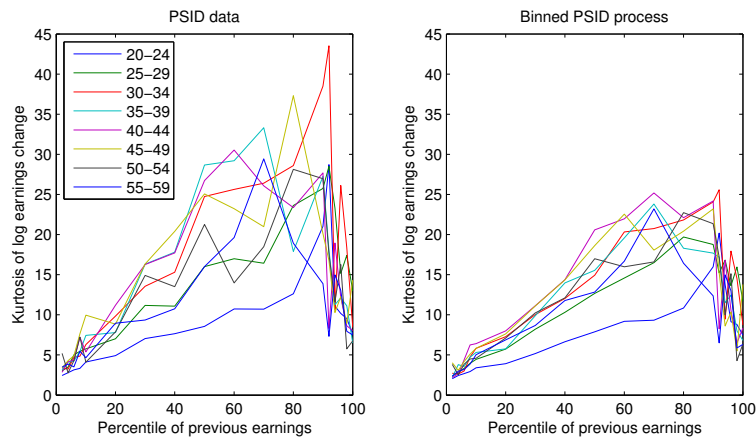
significant improvement with respect to the use of a Gaussian AR(1) process, and is one of the key contributions of this empirical method.

When we refer to "NL synthetic W2 process", unless specifically mentioned, we imply the 103-gridpoint discretization. Using the 18-gridpoint discretization instead has no significant impact on the results (not reported).

C   Appendix for online publication only: Definition of Stationary Equilibrium

A stationary equilibrium is composed of

$$
\begin{cases}
\text{an interest rate } r \text{ and wage rate } w, \\
\text{allocations } c(x), a'(x), \\
\text{government tax rates and transfers } (\tau_a, \tau_l, p), \\
\text{and a distribution } \psi_t(x) \text{ over state variables } x \text{ for every age } t
\end{cases}
$$

such that the following hold (let μ_t be the proportion of agents of age t):

(i) Given the interest rate, the wage and government tax rates and transfers, the functions $c(x)$ and $a'(x)$ solve the described maximization problem for a household with state variables x.

(ii) The tax rate τ_l is chosen so that the government budget constraint balances at every period:

$$
g = \sum_t \mu_t \int_X \left[\tau_a\, r\, a + \tau_l y I_{t<t_r} - p\, I_{t \geq t_r} \right] d\psi_t(x).
$$

(iii) Distributions are consistent with individual behavior and the exogenous process for labor productivity (the transition function P is described in detail below):

$$
(4) \qquad \psi_t(B) = \int_X P(x, t-1, B) d\psi_{t-1}(x) \qquad \forall B \in B(X).
$$

(iv) Aggregate capital K is given by $\sum_t \mu_t \int_X a \, d\psi_t(x)$. Aggregate effective labor is denoted by L. In equilibrium the price of each factor is equal to its marginal product.

The probability measure ψ_t, together with $(X, B(X))$, forms a probability space. $X = [0, \inf) \times Y$ is the state space (of assets and exogenous labor productivity) and $B(X)$ is the Borel σ-algebra on X. [13]

The function $P(x, t, B)$ is a transition function which gives the probability that an age t agent transits to B next period given that his current state is x. This transition function is determined by the decision rule on asset holding and the exogenous transition probabilities on the labor productivity shock. Survival probabilities are exogenous and orthogonal with respect to decision rules and the exogenous labor productivity process, so they are not included in the transition function as it is described here.

The distribution over states of age 1 agents is determined by the exogenous initial distribution of labor productivity, since agents start life with no assets.

D   Appendix for online publication only: The results computed using the same discount factor across earnings processes

The main results calibrated the discount factors to obtain a capital/income ratio of 3 to match that in the U.S. economy. However, this procedure implies that the discount factors are slightly different across the three earnings processes and that this could be inducing some differences in the results which are not due to the characteristics of the earnings processes. To account for this, we have also re-run all results for all three earnings processes when keeping the discount factor β constant at a level equivalent to the average of the three discount

---

[13]Therefore, for a given set $B(X)$, $\psi_t(B)$ is the proportion of individuals of age t that lie in B. $\mu_t \psi_t(B)$ indicates the fraction of total population that agents of age t in B represent.

factors in the main results (0.953). This new calibration implies that agents in the W-2 world are more patient, while agents in the Huggett and PSID world are more impatient.

D.1   Consumption

Average consumption and the variances of consumption don't show many important changes. In the common β case, individuals in the W-2 world are more patient, hence they consume more after retirement than in the main results (Figure 31). The variances of log consumption are very similar in both cases (Figure 32).
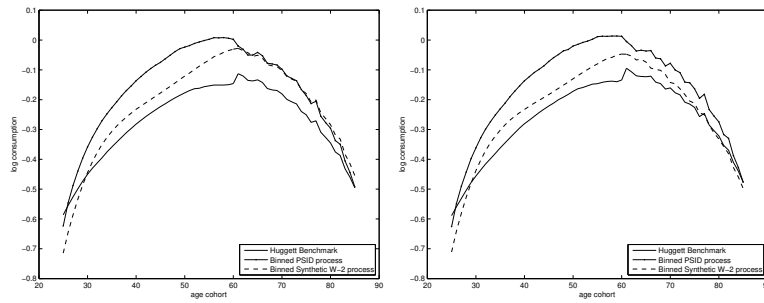

Figure 31: Average of log consumption per age group. Left panel: same β; right panel: different β.
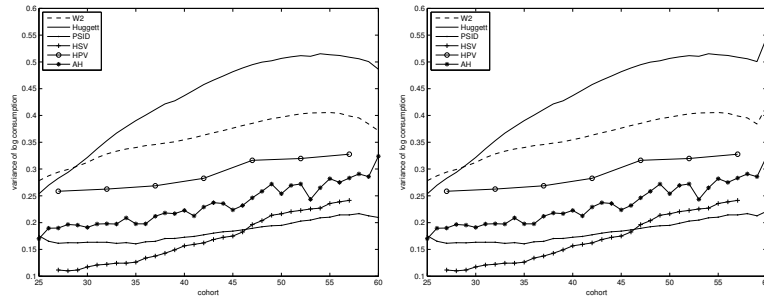

Figure 32: Variances of log consumption. Left panel: same β; right panel: different β.

D.2   Wealth distribution

The results with respect to the wealth distribution (Table 7) are only marginally affected by this change. All economies remain in interval $K/Y = 3 \pm 0.15$. All other results with respect

to wealth (Figures 33, 34, 35, 36, 37, 38) show the expected movements after a small change in patience.

Table 7: Wealth distribution statistics

| Capital-output ratio	Wealth Gini	Percentage wealth in the top						Percentage with negative or zero wealth
		1%	5%	20%	40%	60%	80%	
Same β								
U.S. data								
3.0	.72	28	49	75	89	96	99	5.8–15.0
Huggett Benchmark								
2.9	.73	12	36	75	93	99.15	100	20
NL PSID Process								
2.93	.63	9	28	64	86	96	99.54	6
NL Synthetic W-2 Process								
3.15	.66	10	31	68	88	97	99.69	7.5
Different β								
U.S. data								
3.0	.72	28	49	75	89	96	99	5.8–15.0
Benchmark								
3.0	.72	11	35	74	93	98.98	100	19
NL PSID Process								
3.0	.63	9	28	63	86	96	99.52	6
NL Synthetic W-2 Process								
3.0	.67	10	31	68	89	97	99.74	8
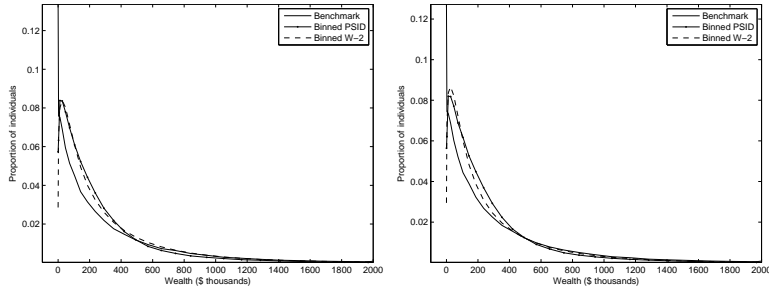
Figure 33: Wealth distribution by earnings process. Left panel: same β; right panel: different β.
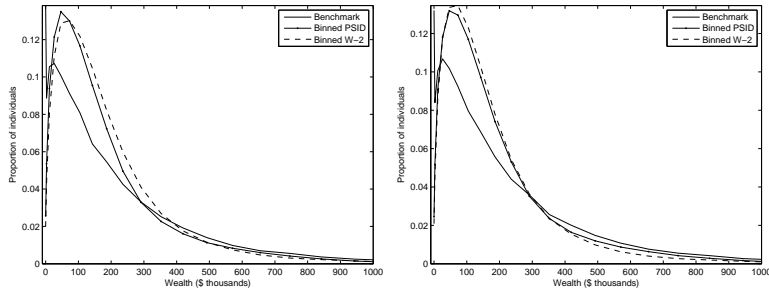

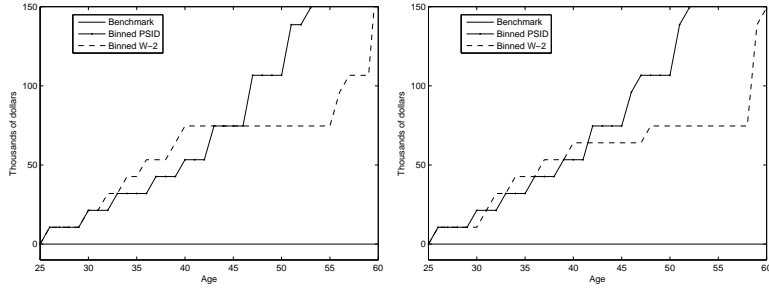Figure 34: Wealth distribution by earnings process: 40-year-olds. Left panel: same β; right panel: different β.

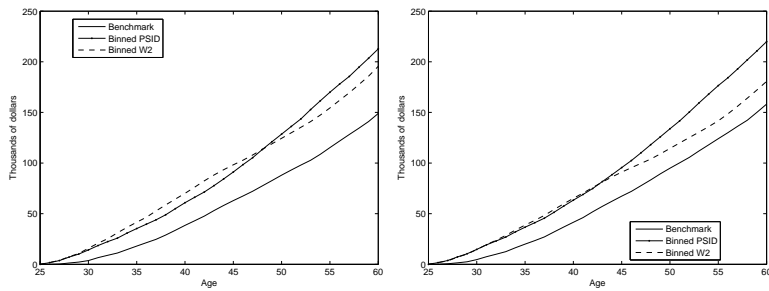Figure 35: Wealth mode. Left panel: same β; right panel: different β.

Figure 36: Wealth median. Left panel: same β; right panel: different β.
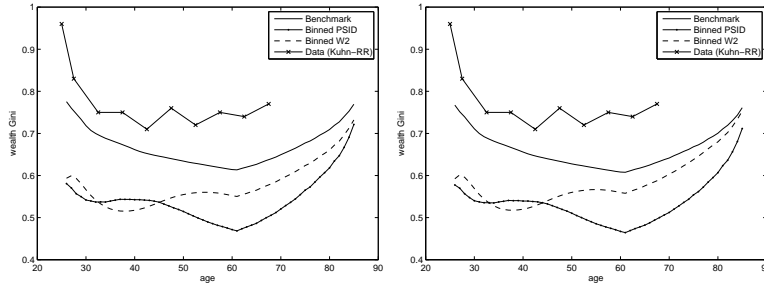
60

Figure 37: Wealth Ginis by age. Top panel: same β; bottom panel: different β.
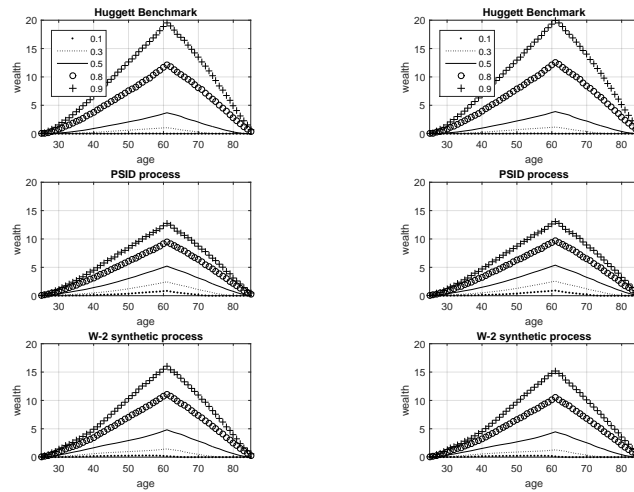

Figure 38: Wealth profiles by wealth quantile. Left panel: same β; right panel: different β.

61