

Heterogeneous Earnings Risk in Incomplete Markets

Eva F. Janssens

University of Amsterdam

November 5, 2021

PRELIMINARY AND INCOMPLETE DRAFT, PLEASE DO NOT QUOTE OR
DISTRIBUTE WITHOUT PERMISSION

Abstract

This paper provides a novel characterization of time-varying heterogeneous earnings risk through a Markov process with heterogeneous transition probabilities. The resulting earnings process allows for a richer notion of earnings risk heterogeneity than previously studied by the literature. Assumptions are derived under which a combination of savings and earnings data can be used to identify the earnings process parameters. Alternatively, a narrower interpretation of earnings risk can be adopted, limiting risk heterogeneity to heterogeneous variances of earnings shocks, such that the earnings process is identifiable from earnings data only. This gives rise to two identification strategies. Applying both strategies to the Survey of Income and Program Participation dataset shows that individuals face considerable inequality of earnings risk. High-risk states are found to be temporary, while low-risk states are persistent. Comparing both strategies shows that only allowing for variance heterogeneity is too restrictive, and a rich notion of risk is required to capture the joint dynamics of individuals' savings and earnings.

Keywords: Labor income risk, idiosyncratic shocks, discretized autoregressive conditional heteroskedasticity process, indirect inference estimation

JEL classification codes: C14, C33, D14, D31, D52

* Acknowledgements: This research is funded through the NWO Research Talent Grant *Estimation and Identification of Parameters in Macroeconomic Models with Incomplete Markets* (project number 406.18.514). I thank Frank Kleibergen and Christian Stoltenberg for their valuable discussions and comments, and Robin Lumsdaine for her feedback and questions. I am grateful for the feedback and questions of participants at various seminars and conferences, in particular, Dirk Krueger, Mikkel Plagborg-Møller and Sophocles Mavroeidis at the IAAE (*International Association for Applied Econometrics*) 2021 conference, the CEF (*Computing in Economics and Finance*) 2021 conference, the ESEM (*European Summer meeting of the Econometric Society*) 2021, the UvA Econometrics brown bag seminar and PhD lunch seminar, as well as the brown bag macro seminar at University of Pennsylvania.

Contact information: Eva F. Janssens: University of Amsterdam, Postbus 15867, 1001 NJ Amsterdam, The Netherlands, e-mail: e.f.janssens@uva.nl

1 Introduction

Earnings risk affects the consumption and saving choices individuals make. As such, improving our understanding of earnings risk can help explain wealth and consumption inequality, and the response of individuals to tax or insurance policies. This paper asks whether and to what extent individuals are heterogeneous with respect to their earnings risk, and how this risk changes over time. In answering this question, earnings risk is not limited to a selected number of moments. Rather, individuals are heterogeneous with respect to the entire probability distribution of earnings shocks. To achieve this, this paper proposes an earnings process where individuals can transition between discrete earnings levels through a Markov process with heterogeneous transition probabilities. Given the discrete nature of its state space, the proposed earnings process can readily be incorporated into a general equilibrium heterogeneous agent model with incomplete asset markets, such that the welfare and distributional effects of earnings risk heterogeneity can be evaluated. The paper derives two identification strategies and estimates the parameters of the earnings process through indirect inference, using the Survey of Income and Program Participation (SIPP), a panel data set for earnings and savings of individuals in the United States.¹

The earnings process proposed in this paper is not identifiable from a panel of just earnings data. Therefore, the paper explores two identification strategies: (i) using additional information, and (ii) restricting the earnings process. The first strategy exploits the information contained in the savings decisions of individuals. A savings panel can aid identification because individuals facing different earnings processes also make different saving decisions. I use this idea to non-parametrically assign labels to individuals. These labels reflect individuals' earnings and savings levels relative to individuals with similar observables and previous period's savings levels. Conditional on these observables, individuals with relatively low savings likely face a different earnings risk than individuals with relatively high savings, and, as such, this procedure groups individuals with similar earnings risk together. Next, an auxiliary transition matrix is constructed describing the transitions of individuals between these assigned states over time. I derive conditions under which this non-parametric labelling procedure correctly assigns individuals into the same risk states. These conditions relate to the way earnings risk affects saving decisions, and bear some similarity to the assumptions on instrumental

¹Indirect inference is a simulation-based estimated method by Smith, Jr. (1993) and Gourieroux, Monfort, and Renault (1993).

variables in treatment effect estimation.² If these conditions are satisfied, the auxiliary transition matrix is already a consistent estimator of the transition probability matrix characterizing the earnings process of interest. For better finite-sample properties and less reliance on these conditions, indirect inference is used to find the earnings process that matches the joint dynamics in savings and earnings summarized in this auxiliary transition matrix as closely as possible.

The second identification strategy shows how the earnings process parameters can be identified using only earnings data when the earnings process is restricted. The restriction this paper considers is to assume that individuals are heterogeneous with respect to the variance of the probability distribution from which earnings shocks are drawn, rather than assuming heterogeneity with respect to the entire probability distribution. Individuals then transition between different discrete earnings and variance levels, and I assume these transitions are governed by a (state-space) discretized generalized autoregressive conditional heteroskedasticity (GARCH) process.³ To achieve this, this paper proposes an improved method to discretize ARCH-type processes as a Markov chain with a discrete state space, adapted from Duan and Simonato (2001).⁴ Through the restrictions implied by an ARCH-type process, identification is aided, because only the parameters of this restricted process have to be estimated, rather than all transition probabilities of the Markov process. Similar to the first approach, I estimate the parameters of the earnings process using indirect inference, but, instead of matching the joint dynamics in savings and earnings using the unrestricted earnings process, I only match the dynamics in earnings, using a discretized AR(1)-GARCH(1,1) model.

To get a sense of the restrictiveness of only allowing for heterogeneous variances, rather than unrestricted heterogeneous transition probabilities, the estimates resulting from the two identification strategies (i) and (ii) are compared. I also consider a combination of both identification strategies, where the GARCH-restrictions are imposed on the estimates obtained from the first identification strategy.

In addition to the theoretical contributions described above, this paper shows five main empirical results. The first result relates to the patterns in saving behavior uncovered by the non-parametric analysis of the SIPP data. I find large dispersion in savings levels of

²As first derived in Imbens and Angrist (1994).

³Introduced by Engle (1982) and generalized by Bollerslev (1986).

⁴The method of Duan and Simonato (2001) is applied in an option-pricing setting, and their method generates highly sparse transition matrices. I propose an adaption to their method, such that the resulting Markov process is more likely to be stable, and as such compatible with general equilibrium models requiring a stable Markov process. I furthermore introduce an extension that can accommodate transitory earnings shocks.

individuals, even after controlling for last period's savings levels and a large set of control variables. Furthermore, the incidence of transitions between low- and high-savings states implied by the auxiliary transition matrix of the first identification strategy reveals a strong reversion in saving behavior towards low-savings states. To explain these empirical patterns in saving behavior, a source of time-varying unobservable heterogeneity is needed. This motivates the earnings process considered in this paper, modeling time-varying heterogeneous earnings risk.

Second, the earnings process parameter estimates from both identification strategies reveal how individuals face not only large heterogeneity in earnings levels, but also in earnings risk. The estimates indicate the existence of low and high-risk states, and these states are, among others, distinguishable by their large differences in job-loss probabilities. Individuals in high-risk states are four to ten times more likely to become unemployed next year than individuals in low-risk states.

Third, the estimates from both identification strategies indicate that there is a difference in the dynamics of high- and low-risk states. High-risk states are unstable; in the absence of (large) shocks to their earnings level, individuals are likely to move to a low-risk state next period. In contrast, low-risk states are persistent. This strong reversion towards low-risk states can explain the patterns in saving behavior found in the non-parametric analysis described under the first main result.

The fourth main result relates to the restrictiveness of this discretized GARCH process used in the second identification strategy. Limiting the notion of earnings risk to a time-varying variance by imposing the restrictions of a discretized GARCH process is highly restrictive. This can, for example, be seen by enforcing the restrictions of the second identification strategy on the estimates of the first identification strategy. Doing so leads to a deterioration in the indirect inference objective function of 200%. This underscores the necessity of a rich notion of risk heterogeneity, as proposed by this paper. Interestingly, I find that the GARCH-process estimates are different from what is typically observed in data on financial returns, a common application for GARCH models. While GARCH models for the volatility of financial returns tend to find a low sensitivity to shocks and high persistence of the variance component in order to capture the long periods of high volatility in financial data, the (discretized) GARCH model for earnings risk implies strong mean reversion (low persistence) and a high sensitivity to shocks, in order to match the strong reversion towards the low risk states.⁵

⁵Common parametrizations of GARCH models are discussed in, for example, Zivot (2009).

Finally, heterogeneous earnings risk has implications for the estimated general equilibrium model. A model with earnings risk heterogeneity generates a savings distribution with a larger variance, and attributes a higher probability to the right tail of the distribution than a model without earnings risk heterogeneity. The welfare gain of eradicating earnings risk heterogeneity, computed by comparing the consumption increase needed to make individuals, before they know their risk types, indifferent between an economy with and without earnings risk heterogeneity, equals 0.51%.⁶

In the next subsection, I discuss the literature this paper relates and contributes to. In Section 2, I present the data and the non-parametric analysis that provides evidence for the presence of heterogeneous time-varying earnings risk. Section 3 describes the general equilibrium model, the earnings process with heterogeneous earnings risk, and its relationship to ARCH-type processes and other earnings processes proposed in the literature. In Section 4, I derive results related to the identification strategies of this paper. In Section 5, I describe the econometric approach used by this paper, followed by the results in Section 6. Section 6 also analyzes the general equilibrium effects of heterogeneous earnings risk. Section 7 concludes.

Related literature The earnings process proposed by this paper can subsume the properties of multiple earnings processes previously considered in the literature, but is more general. This includes earnings processes with heterogeneous variances as in Meghir and Pistaferri (2004), Jensen and Shore (2011), Botosaru and Sasaki (2018), Botosaru (2020), and Almuzara (2021), heterogeneous job-loss and job-finding probabilities (Guvenen, Karahan, Ozkan, and Song, 2021; Mueller, Spinnewijn, and Topa, 2021), heterogeneous persistence (Arellano, Blundell, and Bonhomme, 2017; Guvenen et al., 2021), as well as skewness and other non-normalities, as in Busch, Domeij, Guvenen, and Madera (2021) and Guvenen et al. (2021). The process also relates to models of heterogeneous earnings expectations (as in, e.g., Stoltenberg and Singh (2020)), where it is assumed individuals have heterogeneous information about the same earnings process; in contrast, I assume the processes itself are heterogeneous.

Most closely related within this literature are the papers by Guvenen et al. (2021) and Arellano et al. (2017). Guvenen et al. (2021) model earnings dynamics over the lifecycle using a persistent-transitory model with mixtures of normals and state-dependent employment risk. My results confirm the importance of heterogeneity in persistence and

⁶To put this into context, the welfare gains from eliminating earnings risk heterogeneity are similar in size to, for example, the average welfare effects of eliminating business cycles (Krusell, Mukoyama, Şahin, and Smith Jr., 2009) and the welfare effect of the increase in income inequality seen during the period of 1972-2000 (Krueger and Perri, 2004).

heterogeneous job-loss probabilities highlighted in their paper. Arellano et al. (2017) also focuses on heterogeneous persistence of earnings shocks. There are several important differences between their approaches and mine, one being the continuous state space of their earnings processes. Guvenen et al. (2021) suggest using existing discretization methods such as by Tauchen (1986) to discretize the state-space, albeit with an irregular and large grid. Yet, it is unclear whether this discretized earnings process can fully accommodate heterogeneous persistence and heterogeneous job-loss probabilities. In addition, the large number of grid points required to do so increases the computational burden of solving the macro model that would use this earnings process. The discrete Markov chain with heterogeneous transition probabilities in my paper can accommodate heterogeneous persistence and heterogeneous job-loss probabilities even with a small number of grid points. Compared to the approaches of both Guvenen et al. (2021) and Arellano et al. (2017), my earnings process accommodates richer heterogeneity; even when conditioning on previous earnings levels (and age), individuals are heterogeneous with respect to the earnings process they face for next period. Furthermore, my proposed identification strategy also exploits the information that savings decisions contain on the earnings processes individuals face, while these papers only use earnings data to estimate their earnings process.

In fact, the majority of papers studying earnings processes are descriptive and aim to estimate properties of earnings processes using earnings data only. There are a few papers that do look at income data and savings (or consumption) panel data jointly. These include, e.g., Guvenen and Smith (2014), and Alan, Browning, and Ejrnæs (2018), but there are important differences to my approach. The process considered by Guvenen and Smith (2014) restricts how individuals move from one expectation on future earnings to another by assuming individuals update their expectations through Bayesian learning, while my process does not restrict how individuals form expectations. Alan et al. (2018) estimate a flexible earnings process, but, in contrast to this paper, do not consider time-varying sources of unobserved heterogeneity. Furthermore, both papers consider a partial equilibrium setting, while I consider a general equilibrium.

The improved discretization method this paper proposes for GARCH-type processes contributes to the literature on the discretization of continuous-state-space stochastic processes, with earlier contributions of Tauchen (1986) and Rouwenhorst (1995) for the discretization of AR(1) processes. Recently, De Nardi, Fella, and Paz-Pardo (2020) proposed a simulation-based discretization method for the specific earnings process of

Arellano et al. (2017). My proposed earnings process is also in line with papers that aim to directly estimate a discrete process, rather than discretizing a continuous process, such as, e.g., Castaneda, Diaz-Gimenez, and Rios-Rull (2003), and Druedahl and Munk-Nielsen (2020). Both, however, do not consider unobservable heterogeneity.

The identification results I derive for the proposed earnings process relate to the discussion on the identification of latent-variable models such as Hidden Markov Models, see, e.g., Bonhomme, Jochmans, and Robin (2016, 2017). An important distinction between the latent-variable models they consider and the earnings process proposed in my paper is that I assume that earnings levels follow a first-order Markov process, even when conditioning on the latent variable, while Hidden Markov Models assume that conditional on the latent variable, the observed variable is no longer first-order Markov. My proposed process can generate richer dynamics, but it complicates identification, and requires a novel identification strategy, which I provide in this paper.

2 Data and non-parametric analysis

This section presents the data, and a non-parametric analysis of the Survey of Income and Program Participation (SIPP) dataset. In this analysis, I first demonstrate that there is a large role for unobserved heterogeneity in explaining saving choices. Next, I show that part of this unobserved heterogeneity must be time-varying. I do this by constructing an auxiliary transition matrix, describing the dynamics of individuals' earnings and savings levels. The earnings process proposed and estimated by this paper aims to match the dynamics found by this non-parametric analysis.

The SIPP data is discussed in the next subsection, followed by the non-parametric analysis. In the last subsection, some additional data sources are described that are used to estimate the general equilibrium model presented in Section 3 of this paper.

2.1 Survey of Income and Program Participation

The panel data used in this paper is the SIPP dataset of the United States Census Bureau. This survey follows individuals and households for a selected number of years through a number of interviews over this period. This paper uses the 2014 wave, containing four annual interviews and ending in 2017.

Variables of interest The main variables of interest are person-level measures for employment status, earnings, and net worth. The earnings variable is defined as the sum of monthly earnings over all jobs held.⁷ Net worth is the sum of the values of all assets minus the sum of the liabilities for a person.⁸

The analysis also relies on a set of control variables, including data on age, family composition, gender, education, region of residence, having health insurance, ethnicity, living in a metropolitan area and whether the individual owns or rents a house. For more details on the SIPP dataset and the exact variables and transformations used, see Appendix Section A.1.⁹ This appendix section also provides several summary statistics, see Table A1.

Frequency Although the SIPP asks its respondents to report their earnings at a monthly frequency, these data suffer from seam bias, having little variation in earnings reported during one interview, and large changes between measurements across different interviews (Moore, 2008). Furthermore, individuals report their net worth at a yearly frequency, at the value of the last day of the preceding year. Therefore, I aggregate the data to a yearly frequency. For earnings, the average earnings over each year are used.¹⁰ For the unemployment status, an individual is labeled as unemployed in one year if she was without any job at least six months of that year. This is a relatively strict interpretation of unemployment and ignores short unemployment spells.

Sample selection criteria Self-employed individuals are omitted, because they tend to have very different dynamics in savings and earnings (Elkhishin, West, and MacDonald, 2019). I also omit individuals outside of the labor force due to child birth, disease, disability, retirement, study, or other reasons not to work other than not being able to find a job. Only individuals aged between 18 and 67 are considered.

⁷With a maximum of seven different jobs in a specific month.

⁸Assets include retirement accounts, interest-earning assets, other income-generating assets such as stocks, mutual funds or rental properties, and other assets such regular checking accounts, primary residence, other real estates, vehicles, collectibles, other loans owned to the individual, and owned businesses. Liabilities include debts secured by assets such as mortgage debts, vehicle debt, debts on businesses, owned, as well as unsecured debts such as credit card debt, student loans and other debts.

⁹For a complete description, consider the SIPP manual, available at <https://www2.census.gov/programs-surveys/sipp/tech-documentation/methodology/2014-SIPP-Panel-Users-Guide.pdf>.

¹⁰Note that for the analysis below, what is important are the earnings levels relative to other individuals, rather than the absolute values.

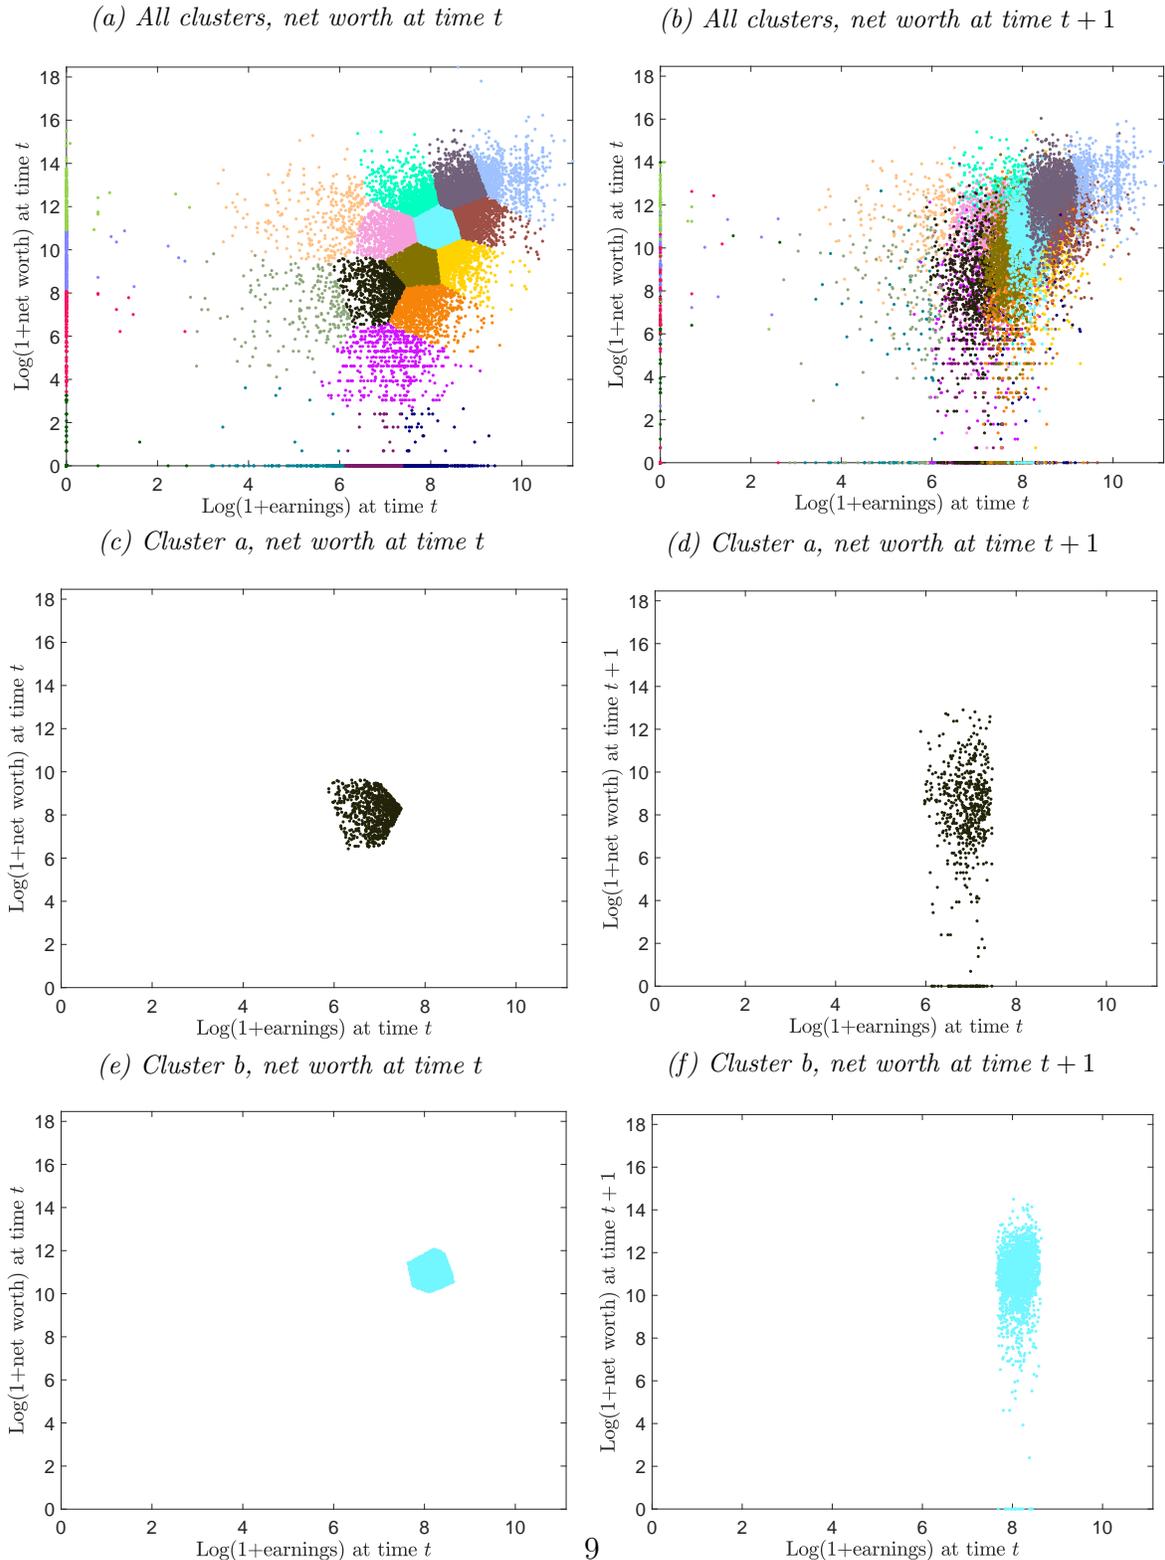
2.2 Non-parametric analysis of the dispersion in saving decisions

Incomplete-markets models with unobserved heterogeneity differ from models without unobserved heterogeneity, such as, for example, the heterogeneous agent models of Aiyagari (1994) or Krusell and Smith (1997) that model uninsurable income shocks, because these models make different predictions on the saving behavior of individuals. In a model with only income heterogeneity (or any other observed heterogeneity), individuals that are homogeneous with respect to the current period's earnings and savings levels will make the same savings choice next period. In a model with unobserved heterogeneity, individuals with similar earnings and savings levels that differ with respect to this unobserved heterogeneity make different saving choices. For example, in the case of earnings risk heterogeneity, conditional on the current earnings and savings levels, individuals with a low probability of becoming unemployed will – for typical model parametrizations – save less compared to individuals with higher job loss probabilities, as the latter will want to maintain a larger buffer for potentially harder times to come.

A non-parametric analysis of the data can empirically motivate the need to incorporate unobserved heterogeneity when explaining saving decisions. For this purpose, I divide all employed individuals in the sample into 20 groups using k-means – a clustering algorithm – based on two variables: their net worth level on December 31st, 2013 and their earnings over 2014. Consequently, each group (cluster) contains individuals that are similar with respect to these two variables. Under the assumption that there is no additional source of heterogeneity, these individuals should have similar net worth levels on December 31st, 2014. The overall level of net worth within these groups can shift up or down, but the dispersion should remain similar to the dispersion of net worth on December 31st, 2013.

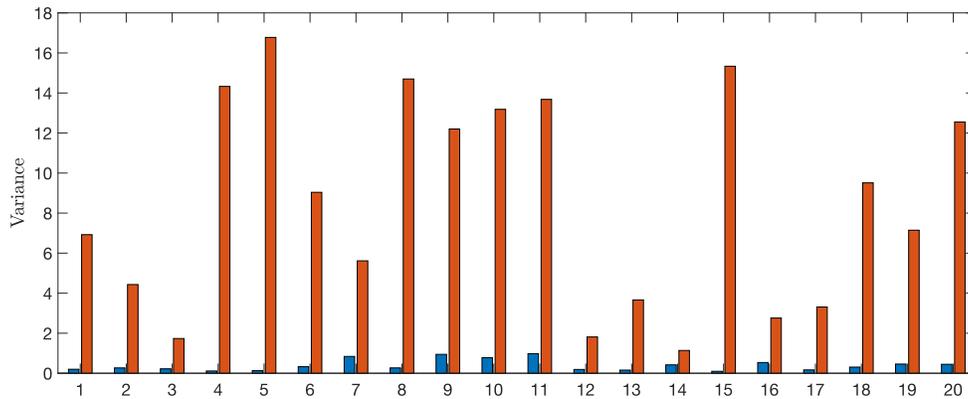
Figure 1a visualizes the 20 clusters identified by the k-means algorithm in a scatter plot, with the net worth levels on December 31st, 2013 on the y -axis and earnings over 2014 on the x -axis. Next, Figure 1b visualizes these same clusters, but instead of using the net worth at December 31st, 2013, it plots the net worth at December 31st, 2014 on the y -axis.

Figure 1: *k*-means analysis with 20 clusters applied to the (standardized and $\log(1+x)$ -transformed) person-level average monthly earnings over 2014 and net worth on Dec-31-2013. The left panel depicts a scatter plot of net worth on Dec-31-2013 (time t) and average monthly earnings in 2014, the right panel uses net worth on Dec-31-2014 (time $t+1$). Clusters *a* and *b* represent two isolated clusters also depicted in Figures 1a-1b. Cluster *a* is cluster 14 in Figure 2, cluster *b* is cluster 18.



Figures 1c-1f focus on two specific clusters and display the large increase in dispersion within these groups going from the end of 2013 to the end of 2014. Figure 1 shows how individuals with similar saving levels at the end of 2013 and earnings over 2014 do not have similar saving levels for next period. This is confirmed by Figure 2; this figure displays the variance of the net worth on December 31st, 2013 and December 31st, 2014, respectively, for the individuals in all 20 clusters and shows that the variance increases are typically large.

Figure 2: Variance of the $\log(1+\text{net worth})$ between individuals in each of the 20 k-mean clusters (labeled on the x-axis) of Figure 1 in Dec-31-2013 (blue) and Dec-31-2014 (red).



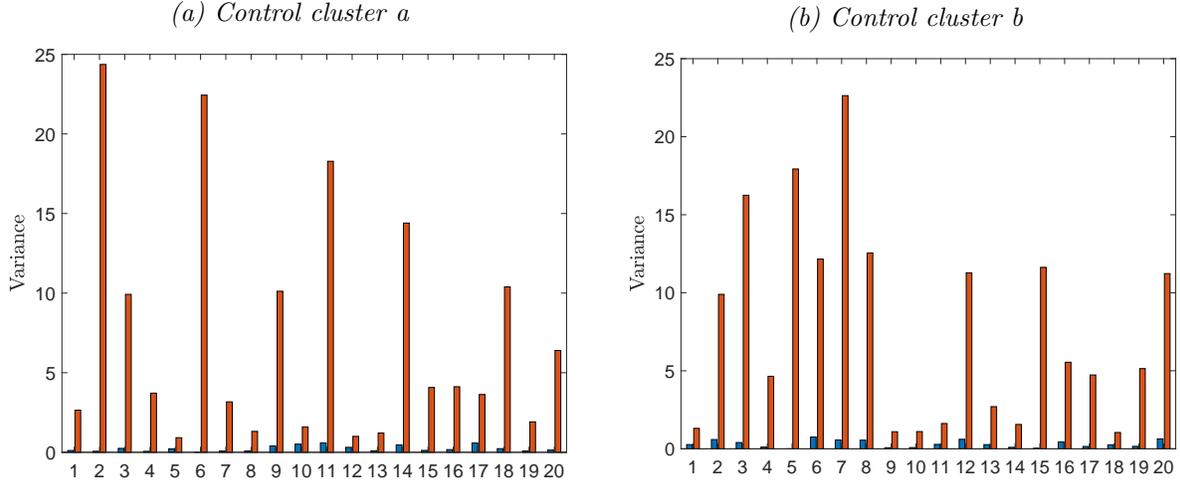
Evidently, various factors can lead to this large increase in the dispersion of net worth levels depicted in Figure 1, such as differences in age, gender, or education level, but also whether people own a house and are hit by a local housing market shock. For that reason, the analysis that generated Figure 2 is repeated, but now accounting for all control variables described in Data section 2.1. If, after controlling for this wide set of observable characteristics, the dispersion in next period’s net worth growth is still large, it suggests that the large differences in net worth are not only due to observable differences between individuals, and that unobserved heterogeneity plays a potentially important role in explaining this dispersion.

To control for these observable characteristics, I use the approach of Bonhomme, Lamadon, and Manresa (2021) to account for individual observable heterogeneity by applying k-means to a large set of control variables. Specifically, I apply k-means clustering to the set of control variables to construct groups of individuals that are similar with respect to a set of observables. I refer to the resulting groups of similar individuals as measured through their observable characteristics as a “control cluster” below. The advantage of this approach, compared to the often-used approach where the variables of interest (income and savings) are first regressed on the control variables, is that no functional

form has to be imposed between the variables of interest and the control variables. Instead, the analysis groups similar individuals together and any subsequent analysis is performed separately for each subgroup.

I use this method to construct 30 control clusters, and, after that, I create 20 subclusters for each control cluster, again based on earnings over 2014 and net worth on December 31th, 2013. These subclusters thus group individuals together that have similar control variables, as well as similar (2014) earnings and (end of 2013) savings levels. Next, Figure 3 shows the variances of the net worth at December 31th, 2013 versus December 31th, 2014 for all 20 subclusters of two of these control clusters. As can be seen, even when controlling for potentially important determinants of heterogeneity in next period’s savings, dispersion in next period’s savings is large within the different clusters, suggesting an important role for unobserved heterogeneity in explaining individuals’ saving choices. The figure description shows that the control clusters are quite specific, implying that the control clusters take care of most of the observed heterogeneity in the sample.

Figure 3: Variance of the $\log(1+\text{net worth})$ between individuals in each of the 20 k-means clusters (labeled on the x-axis) on Dec-31-2013 (blue) and Dec-31-2014 (red), for two different control clusters (see note below for details).



The individuals in control cluster a are all male, on average 54 years old, all have a college degree, 85% identify themselves as white, 98% are non-renter and 77% live in the South or South East of the US, 60% live in a Metropolitan area, and 80% have health insurance. Only 4% still take care of a kid younger than 18. On the other hand, individuals in control cluster b are all females, on average 26 years old, 57% have a college degree or less, 82% identify themselves as white, 99% live in the South, 84% have health insurance, 54% live in a Metropolitan area, 89% are non-renter and 78% have kids of 18 or younger to look after.

Note that these are different clusters than in Figures 1-2, because the methodology for Figure 3 requires constructing control clusters first, after which the methodology used for Figure 2 is applied to each control cluster separately.

2.3 Non-parametric transitions in earnings and savings

Next, a second non-parametric approach is used, to demonstrate that the unobserved heterogeneity driving similar individuals to make different saving choices has a time-varying component. This is done by characterizing the transitions of individuals between groups of high and low savings – relative to individuals with similar controls, earnings and previous period’s savings – and high and low earnings levels – relative to individuals with similar controls.

In formal notation, denote by \hat{S}_{it} the state of an individual i at time t . First, k-means is used to construct a set of control clusters $c \in \mathcal{C} = \{1, \dots, C\}$ with individuals that are comparable with respect to the control variables. C is set to 15. For all individuals that are unemployed at time t , set $\hat{S}_{it} = 0$. Next, group all employed individuals at each period t into two earnings groups $g \in \mathcal{G} = \{1, 2\}$ using the following rule: for individual i in control cluster $c \in \mathcal{C}$, $g = 1$ if his or her earnings at time t are below a threshold $\text{TH}_{c,t}^e$, and individual i of control cluster $c \in \mathcal{C}$ falls in group 2 if his or her earnings are above this threshold. The threshold $\text{TH}_{c,t}^e$ is chosen such that 70% of all individuals in control cluster c at time t fall in group $g = 1$, and all other individuals of this cluster fall in group $g = 2$.

Next, each group $g \in \mathcal{G}$ within a control cluster c is divided into $b \in \mathcal{B} = \{1, \dots, B\}$ savings bins, where B should be chosen such that each bin has a sufficient number of individuals, yet the range of savings within a bin is small. In this case, B is set to 60. The bin-limits are chosen using k-means applied to the current savings level (at the start of year t). This results in groups labeled $(c, g, b) \in (\mathcal{C} \times \mathcal{G} \times \mathcal{B})$, where the composition of these groups can vary at each point in time. These groups have similar control variables (through control cluster c), earnings over period t (through the wage group assignment $g \in \mathcal{G}$) and savings up to time t (through bin assignment $b \in \mathcal{B}$). Then, for each group (c, g, b) at time t , a threshold $\text{TH}_{(c,g,b),t}^s$ is chosen that splits the groups into two based on next period’s savings, that is, savings at the beginning of time $t + 1$. The mean savings of next period of this subgroup (c, g, b) is used as choice for the threshold $\text{TH}_{(c,g,b),t}^s$. Overall, this procedure results in an $N \times (T - 1)$ panel of imputed states \hat{S}_{it} for $t = 2, \dots, T$ and $n = 1, \dots, N$.

In equations:

$$\hat{S}_{it} = \begin{cases} 0 & \text{if } u_{it} = 1, \\ 1 & \text{if } u_{it} = 0, w_{it} < \text{TH}_{c,t}^e, \text{ and } k_{it+1} \leq \text{TH}_{(c,1,b),t}^s, \\ 2 & \text{if } u_{it} = 0, w_{it} < \text{TH}_{c,t}^e, \text{ and } k_{it+1} > \text{TH}_{(c,1,b),t}^s, \\ 3 & \text{if } u_{it} = 0, w_{it} \geq \text{TH}_{c,t}^e, \text{ and } k_{it+1} \leq \text{TH}_{(c,2,b),t}^s, \\ 4 & \text{if } u_{it} = 0, w_{it} \geq \text{TH}_{c,t}^e, \text{ and } k_{it+1} > \text{TH}_{(c,2,b),t}^s, \end{cases} \quad (1)$$

if individual i is in control cluster c and savings bin b at time t . u_{it} is the employment indicator of the individual, where $u_{it} = 1$ if the individual is unemployed according to the definition used by the employment panel, w_{it} are the average monthly earnings of individual i at year t , and k_{it+1} are the savings of the individual at the end of year t . The reason to use cluster-controls in this analysis is because it is control-specific whether one's earnings are considered high or low, and whether savings are high or low. One of the advantages of the approach presented below is that this labeling method is robust to aggregate shocks affecting all individuals' savings levels, as well as to local group-specific shocks (e.g. a shock that hits all home-owners in the South of the US).

Next, one can compute the frequency of transitions between these five states. The elements of the transition matrix are then defined as the total number of times individuals moved from the (imputed) state $\hat{S}_{it} = k$ to $\hat{S}_{i,t+1} = l$, divided by the total number of periods individuals were in (imputed) state $\hat{S}_{it} = k$, for all combinations of k and l :

$$\hat{P}_{kl} = \frac{\sum_{i=1}^N \sum_{t=1}^{T-1} I(\hat{S}_{it} = k, \hat{S}_{i,t+1} = l)}{\sum_{i=1}^N \sum_{t=1}^{T-1} I(\hat{S}_{it} = k)}.$$

This results in the following transition matrix, given in Table 1.

According to this division of households into five groups and introducing a somewhat simplified interpretation, 3% of individuals are unemployed for at least half of the year, whereas 46% are in state 1 (low earnings, low relative savings), 19% in state 2 (low earnings, high relative savings), 22% in state 3 (high earnings, low relative savings) and 10% are in state 4 (high earnings, high relative savings).¹¹ As an example of how to interpret the transition matrix $\hat{P}_{\text{non-param}}$, of the individuals that were unemployed for at least 6 months in the previous year, 41.5% are unemployed for at least 6 months in the next year as well. 50% of the individuals that were unemployed for at least 6 months are relative low earners in the next year,¹² of which 71% is a relative low saver.¹³

¹¹These follow from computing the stationary distribution of the transition matrix in Equation (1).

¹²35.5% + 14.5% = 50%, consider the second and third input of the first row.

¹³35.5%/(50%)=71%.

Table 1: Auxiliary transition matrix describing the relative frequency of transitions between the five states defined in Equation (1).

Transition probabilities (from row/to column)		Unemployed		Low earnings		High earnings					
				Low saver	High saver	Low saver	High saver				
Unemployed		0.415	(0.024)	0.355	(0.027)	0.145	(0.020)	0.067	(0.014)	0.018	(0.008)
Low earnings	Low saver	0.020	(0.002)	0.607	(0.006)	0.280	(0.005)	0.063	(0.004)	0.030	(0.002)
	High saver	0.017	(0.003)	0.663	(0.010)	0.183	(0.009)	0.107	(0.007)	0.030	(0.004)
High earnings	Low saver	0.007	(0.001)	0.159	(0.008)	0.100	(0.005)	0.486	(0.009)	0.248	(0.007)
	High saver	0.008	(0.002)	0.150	(0.012)	0.049	(0.007)	0.602	(0.016)	0.192	(0.012)
Stationary distribution		0.031		0.458		0.193		0.219		0.099	

Bootstrapped standard errors are provided in parentheses, based on 1999 case-sampled bootstrap samples.

The transition matrix in Table 1 shows that these five groups are heterogeneous with respect to their transition probabilities, indicating the existence of five disparate groups. Furthermore, assuming that different saving choices are indicative of the differences in the amount of risk faced by the individuals, the transition matrix contains non-parametric evidence for the existence of time-varying heterogeneous earnings risk. For example, consider the large flows of individuals within the same earnings level states. Conditional on being in the high-earnings-high-savings group in the previous period, 60.2% of individuals move to the high-earnings-low-savings state in the next period. Similarly, for the low-earnings-high-savings individuals, the probability of moving to the low-earnings-low-savings state next year is 66.3%. This shows that, interpreting relative high savings as high risk, individuals that stay within the same earnings category for longer than one year perceive their risk to be lower and save less next period.

Summarizing, applying the non-parametric analysis outlined above on the SIPP dataset provides evidence of unobservable time-varying heterogeneity that affects the saving choices individuals make. The dynamics in savings indicate that non-time-varying unobservable heterogeneity is not sufficient, as studied, for example, in models with heterogeneous preferences, as in, Calvet, Campbell, Gomes, and Sodini (2021), or heterogeneous non-time-varying variances of earnings shocks, as in Botosaru and Sasaki (2018), Botosaru (2020), and Almuzara (2021). Therefore, in the next section, this paper proposes and estimates an earnings process with time-varying unobservable heterogeneity that can match the earnings and saving dynamics of Table 1.

2.4 Other data used for estimation

In addition to the panel data set described above, the estimation of the general equilibrium model that will be described in the next section requires additional data sources. In particular, to estimate the firm-side parameters of this model, I use proxies for aggregate output, aggregate depreciated capital, and aggregate wages. Furthermore, I use data for the interest rate (per annum), the unemployment rate and the average tax rate. Details on the variables used are in Appendix A.2. The proxies for the variables are visualized in Figure A1 for the time-period 1984-2019 in the same Appendix section.

The proxy used for interest rates is the discount rate used by the Federal Reserve. Other proxies for the interest rates can also be considered, but the choice of the proxy is unlikely to affect the overall dynamics in earnings risk found in this paper, and will mainly affect the estimate of the time discount rate (introduced below). Regarding the proxy for the tax rate, the following should be noted. The government that will be presented in the next section only uses taxes for the payment of unemployment benefits, as there are no other government expenditures in the model. For that reason, I approximate the average tax rate by the unemployment insurance expenditures of the government divided by the total earnings of employees. Given the model assumption of a balanced government budget at each point in time, this approximation is appropriate. In real life, governments (should) smooth their taxes over the business cycle.¹⁴

3 Heterogeneous Earnings Risk Model

The model with heterogeneous earnings risk I propose falls into the category of heterogeneous agent models with idiosyncratic risk, in the line of Aiyagari (1994). In the model, ex-ante identical households are ex-post heterogeneous due to incomplete markets in the form of a borrowing constraint, and the idiosyncratic shocks they experience over their lifetime. Whereas models of household heterogeneity as in Aiyagari (1994) focus on idiosyncratic shocks to earnings levels, in the heterogeneous earnings risk model I propose, workers face a second source of idiosyncratic variation. As a result, individuals are heterogeneous both with respect to their earnings levels, and the probabilities of facing certain future earning shocks. A government keeps a balanced budget while tax-

¹⁴Similar to the proxy used for the interest rate, the approximation of the tax rate is, given the identification strategy that the paper uses, not going to affect the dynamics of the earnings risk. It will, however, affect the unemployment benefits in the model, and, consequently, the impact of being unemployed on consumption and saving decisions.

ing labor earnings and providing unemployment insurance, individuals provide labor for a representative firm producing consumption goods. Markets clear, implying a general equilibrium framework.

First, the earnings process proposed by this paper is described, followed by its relationship to other earnings processes proposed in the literature. After that, further details on the general equilibrium model are provided.

3.1 An earnings process with heterogeneous and time-varying earnings risk

A common way to model earnings processes in general equilibrium models is through a discrete-time first-order Markov chain. Let $\mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\}$ denote the set of discretized earnings levels, with L distinct levels. Denote the corresponding random variable at time t as Y_{it} with realizations $y_{it} \in \mathcal{Y}$. Transitions between the different earnings levels are governed by a transition probability matrix Q , with elements q_{ij} , $i, j = 1, \dots, L$, representing the probability of transitioning from current earnings level $Y_{it} = \bar{y}(i)$ to earnings level $Y_{it+1} = \bar{y}(j)$ next period. That is, $q_{ij} = P(Y_{it+1} = \bar{y}(j) | Y_{it} = \bar{y}(i))$.

The recursive equilibrium solution of this general equilibrium model is unique and exists if our earnings process satisfies two conditions: (i) ergodicity and (ii) stationarity (Stokey, Lucas Jr., and Prescott, 1989, Chapter 9). Ergodicity means that each state in \mathcal{Y} should be reachable from any other state in a finite number of steps. Stationarity refers to the existence of a stationary distribution q^{stat} that remains unaffected after applying a matrix multiplication with Q : $q^{\text{stat}} = q^{\text{stat}}Q$. If both conditions are satisfied, the Markov chain (\mathcal{Y}, Q) is called stable.

To obtain heterogeneous transition probabilities for individuals with the same earnings level, I propose to extend the state space with the random variable ξ_t , which has realizations from the set $\mathcal{X} = \{\bar{\xi}(1), \dots, \bar{\xi}(M)\}$, where M is the number of discrete levels. The interpretation of ξ is that this is an additional source of heterogeneity; people differ not only with respect to their earnings levels Y , but also with respect to the variable ξ . The state an individual is in at time t is then characterized by random variables $(Y_{it}, \xi_{it}) = (\bar{y}(j), \bar{\xi}(l)) \in (\mathcal{Y} \times \mathcal{X})$. The two variables Y and ξ can be dependent. The transitions are therefore governed by an extended transition probability matrix P , where $p_{(i,k);(j,l)}$ is the probability of transitioning from state $(\bar{y}(k), \bar{\xi}(m))$ to $(\bar{y}(l), \bar{\xi}(n))$. That is,

$$p_{(k,m);(l,n)} = P(Y_{it+1} = \bar{y}(l), \xi_{it+1} = \bar{\xi}(n) | Y_{it} = \bar{y}(k), \xi_{it} = \bar{\xi}(m)). \quad (2)$$

As above, this Markov chain $((\mathcal{Y} \times \mathcal{X}), P)$ should be stable.

Extending the state space with ξ brings about that two individuals with the same current level of earnings $\bar{y}(i)$ have different transition probabilities to next period's earnings levels $\bar{y}(k) \in \mathcal{Y}$, if these individuals differ with respect to the current value of ξ . This results in a rich notion of earnings risk heterogeneity; individuals differ with respect to the entire probability distribution of future earnings, rather than only a limited set of moments. Furthermore, the process allows for time-varying earnings risk, even in the absence of changes in the earnings levels, as individuals can transition between different values of $\xi_t \in \mathcal{X}$, and each state in \mathcal{X} potentially corresponds to different transition probabilities to next period's earnings levels. The implications of this process are made intuitive through Example 1.

Example 1. Consider a setting where L , the number of discrete earnings levels, equals 2, and M , the number of discrete levels of ξ , also equals 2, so that the total number of states of the extended state space equals 4:

$$(\mathcal{Y} \times \mathcal{X}) = \{(\bar{y}(1), \bar{\xi}(1)), (\bar{y}(1), \bar{\xi}(2)), (\bar{y}(2), \bar{\xi}(1)), (\bar{y}(2), \bar{\xi}(2))\}.$$

Assume $\bar{y}(1) < \bar{y}(2)$. The transition probability matrix of moving from one state to the next, consistent with the notation in Equation (2) is given by P in Equation (3), with a concrete example in P_{e1} :

$$P = \begin{bmatrix} P_{(1,1);(1,1)} & P_{(1,1);(1,2)} & P_{(1,1);(2,1)} & P_{(1,1);(2,2)} \\ P_{(1,2);(1,1)} & P_{(1,2);(1,2)} & P_{(1,2);(2,1)} & P_{(1,2);(2,2)} \\ P_{(2,1);(1,1)} & P_{(2,1);(1,2)} & P_{(2,1);(2,1)} & P_{(2,1);(2,2)} \\ P_{(2,2);(1,1)} & P_{(2,2);(1,2)} & P_{(2,2);(2,1)} & P_{(2,2);(2,2)} \end{bmatrix}, \quad P_{e1} = \begin{bmatrix} 0.90 & 0.01 & 0.08 & 0.01 \\ 0.01 & 0.30 & 0.60 & 0.09 \\ 0.01 & 0.01 & 0.08 & 0.90 \\ 0.50 & 0.01 & 0.09 & 0.40 \end{bmatrix} \quad (3)$$

As can be seen, two individuals i and j with current earnings $y_{it} = y_{jt} = \bar{y}(1)$ can face a different earnings process for next period, if individual i has $\xi_{it} = \bar{\xi}(1)$, while individual j has $\xi_{jt} = \bar{\xi}(2)$. Consider the concrete example in P_{e1} . An individual with $(\bar{y}(1), \bar{\xi}(1))$ faces low upward risk and will remain in the same state with probability 0.9, while an individual with $(\bar{y}(1), \bar{\xi}(2))$ actually has a substantial probability of moving to the higher income state $\bar{y}(2)$ ($0.6+0.09=0.69$). For individuals with $y_{it} = \bar{y}(2)$, individuals with $\bar{\xi}(1)$ have low downward risk, while individuals with $(\bar{y}(2), \bar{\xi}(2))$ have a 0.51 probability of moving to a lower earnings level. \triangleleft

In the next subsections, the relationship of this general earnings process to other processes proposed by the literature is discussed. In particular, I discuss how the earnings process $((\mathcal{Y} \times \mathcal{X}), P)$ can be mapped into (i) a GARCH-type process, (ii) an earnings

process with heterogeneous (non-time-varying) variances, (iii) a model with heterogeneous job-loss and job-finding probabilities, (iv) an earnings process with skewness and other non-normalities, (v) an earnings process with heterogeneous and time-varying persistence, and (vi) a model with heterogeneous signals about the earnings process. The first process will be discussed in more detail, as this process will be estimated and compared to the unrestricted process in the Results section of this paper. In Section 4, the identification of the parameters of this process is discussed.

3.1.1 Relationship to discretized GARCH-type processes

In the analysis below, I want to compare the estimates of the unrestricted Markov process of Equation 2 with the estimates that follow from a more restrictive process, in this case an AR(1)-GARCH(1,1) process. Therefore, I will now describe how one can discretize the state space of an AR(1)-GARCH(1,1) and map the latter process into a discrete-state-space first-order Markov process. The idea behind this mapping is that, instead of describing transitions between discrete levels of earnings y and the unobserved state ξ , the process describes transitions between discrete values of earnings y and σ^2 , where σ^2 is the variance of the earnings shocks. These transitions are governed by the dynamics of an AR(1)-GARCH(1,1) process (Bollerslev, 1986; Engle, 1982). The discretization method I will describe below can also be used for related processes, such as NGARCH, EGARCH, or TGARCH processes.

The macroeconomic profession is well-acquainted with methods to discretize AR(1) processes, relying on the work of Tauchen (1986), Rouwenhorst (1995) and others, and practitioners apply their discretization methods when incorporating AR(1) processes in macroeconomic models, but discretized AR-GARCH processes have not yet been used in macroeconomic models. Given that the solution methods used for heterogeneous agent models require a discrete state space, discretization methods are necessary when one wants to incorporate continuous-state-space processes in macroeconomic models.

In an option pricing application, Duan and Simonato (2001) show that GARCH-type processes can be approximated by a Markov chain and they prove convergence to the underlying GARCH model as the number of Markov states increases. The method of Duan and Simonato (2001) is suitable for option pricing applications, but for a macroeconomic application, it has certain disadvantages. In particular, the approximation they propose leads to sparse transition matrices, and as such the implied Markov chain may be unstable. I will first present the method of Duan and Simonato (2001), after which I will propose an adjustment to their method that improves the approximation.

Assume the following AR(1)-GARCH(1,1) income process for the (log) earnings levels of individual i at time t , $Y_{i,t}$:

$$Y_{i,t+1} = a + b \cdot Y_{it} + \sigma_{i,t+1} \cdot u_{i,t+1}, \quad b < 1 \quad (4)$$

$$\sigma_{i,t+1}^2 = c + d \cdot \sigma_{it}^2 + f \cdot u_{it}^2, \quad u_{it} \sim N(0, 1), \quad d + f < 1. \quad (5)$$

Now approximate this by a Markov chain with L discrete income levels $(\bar{y}(i), i = 1, \dots, L)$ and n variance levels $(\bar{\sigma}^2(j), j = 1, \dots, M)$. The transition probability matrix P^G of the Markov process then contains the probabilities of moving from one combination $(\bar{y}(i), \bar{\sigma}^2(j))$ to another pair $(\bar{y}(k), \bar{\sigma}^2(l))$ for all pairs (i, j) and (k, l) . Once we know how to compute the elements of the transition probability matrix P^G , it is evident that the discretized AR(1)-GARCH(1,1) model is a restricted version of the earnings process of Equation (2), where σ_{it}^2 takes the role of the second source of heterogeneity ξ_{it} and the transition dynamics in transition probability matrix P are now instead governed by the model Equations (4)-(5).

Next, following Duan and Simonato (2001), I introduce some general notation for the cells H that correspond to the L discretized income levels $\bar{y}(i)$: $H(i) = [h(i), h(i+1)]$ for $i = 1, \dots, L$ with $h(1) = -\infty$, $h(i) = (\bar{y}(i) + \bar{y}(i-1))/2$ for $i = 2, \dots, L$ and $h(L+1) = +\infty$. For the M levels of the conditional variance, similarly, define cells G , where $G(i) = [g(i), g(i+1)]$ for $i = 1, \dots, M$, where $g(1) = -\infty$, $g(i) = (\bar{\sigma}^2(i) + \bar{\sigma}^2(i-1))/2$ for $i = 2, \dots, M$ and $g(M+1) = +\infty$. Defining function

$$R(\sigma_{i,t+1}^2, y_{i,t+1}, y_{it}) = c + d \cdot \sigma_{i,t+1}^2 + f \cdot \left(\frac{y_{i,t+1} - a - b \cdot y_{it}}{\sigma_{i,t+1}} \right)^2,$$

Duan and Simonato (2001) approximate the transition probabilities

$$p_{(i,j);(k,l)}^G = P(Y_{i,t+1} \in H(k), \sigma_{i,t+2}^2 \in G(l) | Y_{it} \in H(i), \sigma_{i,t+1}^2 \in G(j)) \quad (6)$$

by:

$$\begin{aligned} & P(Y_{i,t+1} \in H(k), \sigma_{i,t+2}^2 \in G(l) | Y_{it} = \bar{y}(i), \sigma_{i,t+1} = \bar{\sigma}^2(j)) \\ & \approx \begin{cases} P(Y_{i,t+1} \in H(k) | Y_{it} = \bar{y}(i), \sigma_{i,t+1}^2 = \bar{\sigma}^2(j)) & \text{if } R(\bar{\sigma}^2(j), \bar{y}(k), \bar{y}(i)) \in G(l) \\ 0 & \text{else,} \end{cases} \quad (7) \end{aligned}$$

where $P(Y_{i,t+1} \in H(k) | Y_{it} = \bar{y}(i), \sigma_{i,t+1}^2 = \bar{\sigma}^2(j)) = P(L_{ij}(k) \leq Z < L_{ij}(k+1))$ with $L_{ij}(k) = \frac{h(k) - a - b\bar{y}(i)}{\bar{\sigma}(j)}$ and $Z \sim N(0, 1)$. Reading Equations (6) and (7) carefully, one sees

that this computation relies on two consecutive approximation steps. First, going from Equation (6) to the first line in (7) requires the following approximation:

$$\begin{aligned} P(Y_{i,t+1} \in H(k), \sigma_{i,t+2}^2 \in G(l) | Y_{it} \in H(i), \sigma_{i,t+1}^2 \in G(j)) &\approx \\ P(Y_{i,t+1} \in H(k), \sigma_{i,t+2}^2 \in G(l) | Y_{it} = \bar{y}(i), \sigma_{i,t+1}^2 = \bar{\sigma}^2(j)), &\end{aligned} \quad (8)$$

that is, instead of conditioning on the previous values of y and σ^2 to fall inside their respective intervals, the approximation assumes the previous values of y and σ^2 are equal to the midpoints of these intervals. Next, Equation (7) also relies on an approximation. Instead of evaluating whether $\sigma_{i,t+2}^2$ falls into interval $D(l)$ for the different possible realizations of $Y_{i,t+1} \in H(k)$, it only evaluates whether the $\sigma_{i,t+2}^2$ that corresponds to $\bar{y}(k)$, i.e., the midpoint of the interval $H(k)$, falls into $G(l)$.

This second approximation step has some serious limitations for the application considered in this paper. As articulated before, it is necessary that the Markov process used is stable, otherwise a solution to the general equilibrium model may not exist or may not be unique. The approximation method of Duan and Simonato (2001) generates highly sparse transition probability matrices and therefore can generate Markov chains that fail the ergodicity condition; the stationary distribution of the Markov chain is not reachable from any initial distribution. To make this issue less likely to occur, I suggest replacing Equation (7) by the following:¹⁵

$$\begin{aligned} p_{(i,j);(k,l)}^G &\approx P(Y_{i,t+1} \in H(k), \sigma_{i,t+2}^2 \in G(l) | Y_{it} = \bar{y}(i), \sigma_{i,t+1}^2 = \bar{\sigma}^2(j)) \\ &= \int_{L_{ij}(k)}^{L_{ij}(k+1)} I\{(c + d \cdot \bar{\sigma}^2(j) + f \cdot u^2) \in G(l)\} \phi(u) du, \end{aligned} \quad (9)$$

where $\phi(u)$ denotes the probability density function of the standard normal distribution and $I\{\cdot\}$ is an indicator function. As can be seen, I still use the approximation in Equation (8), but the computation of this probability is then exact.

This issue and how my proposed computation of the transition probabilities resolves this issue are visualized in Figure 4. Figure 4 also displays how the approximation of Duan and Simonato (2001) can generate undesirable sparsity, as it is not possible to reach a high-variance state from the low-variance state when using their proposed Equation

¹⁵Whether or not the Markov chain is ergodic will still depend on the underlying parameters of the GARCH process, for example, if the GARCH process indicates a constant value for the conditional variance, because the GARCH parameters are not identified, and transitions between different variance states can never occur.

(7), violating the ergodicity condition. This is not the case when using Equation (9), because, as one can see in the lower part of Figure 4, it is also possible to reach a high-variance state, albeit with a low probability.

Accommodating a transitory component in the income process

Another commonly encountered earnings process relies on the assumption that (log) earnings Y_{it} are composed of both a transitory and a permanent component. The Markov transition matrix can also be mapped into such an earnings process of the form

$$Y_{it} = \eta_{it} + \nu_{it}, \tag{12}$$

where η_{it} follows the AR(1) process above of Equation (4), and ν_{it} is a white-noise transitory shock. This model is often referred to as the canonical model of earnings. As an extension, assume that η_{it} follows an AR(1)-GARCH(1,1) process. The additivity of the transitory-permanent income process allows using the discretization technique proposed above for η_{it} , and in a second step computing the final discretized Markov process for Y_{it} . Details on the discretization method of the process in Equation (12) are provided in Appendix Section B.1, where it is shown that a mapping can be made from this model to the general earnings process $((\mathcal{Y} \times \mathcal{X}), P)$.

Accommodating a leverage effect in the GARCH process

As indicated above, this discretization method is compatible with the larger class of GARCH processes. To illustrate this, I show in Appendix B.2 how the discretization method proposed above changes when one is interested in the discretization of a TGARCH process. Unlike the traditional GARCH process, a TGARCH allows for a leverage effect, where negative shocks have a different effect on the variance than positive shocks.

3.1.2 Relationship to other earnings processes

In Appendix Sections B.3-B.7, I discuss how the general earnings process of Equation (2) is related to other earnings processes studied in the literature. This includes earnings processes with heterogeneous variances as studied by Botosaru and Sasaki (2018), Botosaru (2020), Almuzara (2021), and others, as discussed in Section B.3, models with heterogeneous job-loss and job-finding probabilities, as in Guvenen et al. (2021) and Mueller et al. (2021), discussed in Section B.4, and models featuring skewness and other higher-order moments, as shown to be important by Busch et al. (2021), and discussed

in Section B.5. The relationship to models with heterogeneous persistence as in Arellano et al. (2017) is discussed in Section B.6, and the mapping to models with heterogeneous signals about earnings as in Stoltenberg and Singh (2020) is discussed in Section B.7.

3.2 The general equilibrium model

This subsection describes the general equilibrium model in which the earnings process of this paper will be incorporated.

Workers A unit mass of ex ante identical workers optimizes its discounted future stream of utilities, by allocating their incomes between savings and consumption. Workers face earnings shocks consistent with the earnings process $((\mathcal{Y} \times \mathcal{X}), P)$ in Equation (2), for which, specifically, the following discretization will be assumed:

$$(y_{it}, \xi_{it}) \in \begin{pmatrix} y = \text{unemployed}, \xi = (\bar{\xi}(1) \text{ or } \bar{\xi}(2)) \\ y = \bar{y}(1), \xi = \bar{\xi}(1) \\ y = \bar{y}(1), \xi = \bar{\xi}(2) \\ y = \bar{y}(2), \xi = \bar{\xi}(1) \\ y = \bar{y}(2), \xi = \bar{\xi}(2) \end{pmatrix}. \quad (13)$$

This means workers can be in one of five states: (i) unemployed (unemployed, $\xi = (\bar{\xi}(1) \text{ or } \bar{\xi}(2))$), (ii) employed with low earnings and low earnings risk ($y = \bar{y}(1), \xi = \bar{\xi}(1)$), (iii) employed with low earnings and high earnings risk ($y = \bar{y}(1), \xi = \bar{\xi}(2)$), (iv) employed with high earnings and low earnings risk ($y = \bar{y}(2), \xi = \bar{\xi}(1)$), and (v) employed with high earnings and high earnings risk ($y = \bar{y}(2), \xi = \bar{\xi}(2)$). For simplicity, $\bar{\xi}(1)$ is interpreted as low risk and $\bar{\xi}(2)$ as high risk, but no restrictions are yet imposed to achieve such an ordering. Note that the unemployed state is an aggregation of two states, namely (unemployed, $\bar{\xi}(1)$) and (unemployed, $\bar{\xi}(2)$). Denote the five states by $\epsilon = (y, \xi) \in \{0, 1, 2, 3, 4\}$.

If a worker is unemployed, she receives unemployment benefits μw_t , where w_t denotes the aggregate wage rate and μ is a parameter capturing the fraction of average wages received as unemployment benefits. Employed workers face heterogeneity with respect to their wage rate and transition probabilities to next period's earnings state. In the low-wage state, a worker earns $w_t(1 - v_l)$, while in the high-wage state, a worker earns $w_t(1 + v_h)$. v denotes the wage differential between the two wage groups. v_h is chosen such that, given the proportion of workers in the low-wage state and high-wage state

respectively, the average wage over these groups is equal to w_t . This implies:¹⁶

$$v_h = v_l \frac{(p_1 + p_2)}{(p_3 + p_4)}.$$

where p_ϵ denotes the stationary probability of being in state ϵ .

Employed workers' earnings are taxed with rate τ_t . All workers can save in capital k , with a rate of return equal to r_t . Workers cannot borrow, that is, $k_{t+1}^i \geq 0$. This results in the following optimization problem for workers:

$$\max_{\{c_t^i, k_{t+1}^i\}_{i=0}^\infty} \mathbb{E} \sum_{t=0}^{\infty} \beta^t \frac{(c_t^i)^{1-\gamma} - 1}{1-\gamma} \quad (14)$$

$$\text{s.t. } c_t^i + k_{t+1}^i = (1 + r_t)k_t^i + \begin{cases} \mu w_t & \text{if } \epsilon_{it} = 0 \\ (1 - \tau_t)w_t \bar{l}(1 - v_l) & \text{if } \epsilon_{it} = 1, 2 \\ (1 - \tau_t)w_t \bar{l}(1 + v_h) & \text{if } \epsilon_{it} = 3, 4 \end{cases} \quad (15)$$

$$k_{t+1}^i \geq 0, \quad c_t^i > 0 \quad (16)$$

where k_t^i is the beginning-of-period level of capital of the workers, c_t^i is the individual consumption level, \bar{l} is the number of hours an employed individual works, and k_{-1} is given.

Firm There is a representative firm, characterized by a Cobb-Douglas production function. K_t is capital per capita, depreciating at rate δ and L_t is the employment rate. Then, per capita output is given by:

$$Y_t = K_t^\alpha (\bar{l}L_t)^{1-\alpha} \quad (17)$$

Market competition implies that prices are given by

$$w_t = (1 - \alpha)\bar{l} \left(\frac{K_t}{\bar{l}L_t} \right)^\alpha \quad \text{and} \quad r_t + \delta = \alpha \left(\frac{K_t}{\bar{l}L_t} \right)^{\alpha-1} \quad (18)$$

Government The government taxes employed agents in order to pay the unemployment benefits to the unemployed agents. Assuming a balanced government budget

¹⁶Solve $w_t = ((p_1 + p_2)w_t(1 - v_l) + (p_3 + p_4)w_t(1 + v_h))/(p_1 + p_2 + p_3 + p_4)$ for v_h .

implies

$$(\bar{l}L_t)\tau_t = \mu u_t \tag{19}$$

where $u_t = 1 - L_t$ is the unemployment rate at time t .

Recursive competitive equilibrium A stationary recursive competitive equilibrium consists of a value function $V(k, \epsilon)$, aggregate prices r and w , an allocation K (aggregate capital), a consumption decision rule $c(k, \epsilon)$, a savings decision rule $k'(k, \epsilon)$, and a joint probability measure of the individual capital choices k and states ϵ , $\Phi_{k,\epsilon}$, such that:

- (i) $V(k, \epsilon)$ is attained by decision rules $c(k, \epsilon)$, $k'(k, \epsilon)$, given prices r and w
- (ii) Prices are consistent with Equation (18), taxes follow Equation (19).
- (iii) The distribution over assets and states $\Phi_{k,\epsilon}$ is stationary
- (iv) The asset market clears: $K = \int \sum_{\epsilon} k'(k, \epsilon) d\Phi_{k,\epsilon}$.

Note that in the stationary recursive competitive equilibrium, prices, the distribution of assets and states and the aggregate variables are constant. There is, however, variation in the choices that individuals make; workers move within the stationary distribution of assets and states.

4 Identification

This section analyzes the identification of the parameters of the earnings process, as well as the other structural parameters of the general equilibrium. To simplify the identification analysis, the paper will make use of the two-step identification approach in Janssens (2019), implying that identification of subsets of parameters can be analyzed separately, when the moments used to identify these parameters do not depend on the other parameters in the model. First, the two-step method and its advantages are briefly described. Next, the paper will discuss identification of the earnings process parameters. The discussion of the identification of the other parameters follows.

4.1 Two-step identification and estimation

This section briefly describes how the two-step identification and estimation procedure of Janssens (2019) can be used for the heterogeneous earnings risk model presented in the

previous section. As described by Janssens (2019) in more detail, the idea of two-step identification and estimation is that the parameter vector of interest, denoted by θ , can be subdivided into two subsets $\theta = (\theta_1, \theta_2)$, where for the first subset of parameters, θ_1 , a set of moment conditions (or for indirect inference, auxiliary parameters) exists that is satisfied for any value of θ_2 . Then, $\theta_2 = \theta \setminus \theta_1$, for which another set of moment conditions is defined with no additional restrictions except that they globally and locally identify θ_2 , conditional on θ_1 . Under the choice of an efficient weighting matrix, the second step estimates of θ_2 are asymptotically efficient. These results hold both for Generalized Method of Moments, as well as indirect inference.

This two-step approach offers several advantages, of which the following two are most relevant for the application in this paper. First of all, two-step estimation is computationally efficient. In macroeconomic applications, and the application in this paper, the first-step parameters and moments can be chosen such that they do not require the full model solution, which is computationally intensive to obtain. By estimating a subset of parameters in the first step, fewer parameters have to be estimated in the step where the objective function does rely on the evaluation of the full model. A second advantage is that the identification analysis is simplified. The first-step identification analysis can often-times be performed analytically, see also Janssens (2019). Furthermore, in the second-step, identification analysis can be done conditional on the first-step estimates.

This two-step estimation method will be used to estimate the parameters of interest in this application. Specifically, in the first step, the parameters are estimated that can be identified from moments that characterize the firm-side and government-side of the problem. The first-step parameters in this application are (i) α (elasticity of capital), (ii) δ (depreciation rate), and (iii) μ (the unemployment benefits). For more details on why these first-step parameters are independent of the other parameters of this model, see Appendix C. In the second step, the other parameters will be estimated, being (i) P (the transition probability matrix), (ii) γ (worker's risk aversion), (iii) β (the discount rate of the worker), and (iv) ν_l (the income differential between high and low earners).¹⁷ The consequence of this division of parameters is that the analysis only needs to show joint identification of two groups of parameters, (α, δ, μ) , and $(P, \gamma, \nu_l, \beta)$, but not all parameters jointly.

¹⁷ \bar{l} (hours worked by an employed household) is used to normalize the average total labor supply of the economy to one and therefore not estimated.

4.2 Identification of the earnings process parameters

As the earnings process is the main object of interest in this paper, this subsection analyzes the identification of the parameters of the earnings process of Section 3.1. This set of parameters consists of the transition probabilities matrix of Equation (2), as well as the discretized earnings levels. The proofs are based on the general process, while numerical results are based on the process in Equation (13), or a smaller toy example used for ease of exposition.

Whether a parameter is identified depends on how this parameter enters the structural model, but also on the data that are available for estimation. For this discussion, a distinction will be made between different sources of identification coming from four potential data types: (i) cross-sectional earnings data, (ii) an earnings panel data set, (iii) cross-sectional net worth data, and (iv) a panel with net worth data and earnings data. Data types (i) and (iii) contain information on the cross-sectional distribution of individuals, while (ii) and (iv) consider a panel dimension.

For the results on data types (i)-(iii), the focus is on local first-order identification. Informally, a set of parameters is locally first-order identified if the Jacobian of the set of moments or auxiliary parameters used to estimate these parameters has the same rank as the number of parameters. If it is possible to change one of the parameters, and adjust the other parameters such that the moments are unaffected, these parameters are not jointly locally identified, and in that case, the rank of the Jacobian will be lower than the number of parameters.

For data type (iv), a different identification argument is used. I derive a set of assumptions under which one can perfectly retrieve the underlying earnings states of the individuals in the panel. These assumptions bear some similarity to the relevance and monotonicity assumption on instrumental variables used when estimating treatment effects. Under monotonicity, it is assumed that the chosen treatment is monotonically increasing/decreasing in the instrumental variable. Instead, in this setting, I need to make assumptions on the effect that the earnings state has on the saving decisions of the individual. Furthermore, I need to assume that the effect of the earnings state on the savings decision of individuals is sufficiently large, similar to the relevance condition for instruments, which requires that instruments have a significant effect on the treatment choice.

4.2.1 Cross-sectional data on earnings

The following proposition can be derived regarding the identification of the earnings process of Section 3.1 when using moments from the cross-sectional distribution of earnings levels.

Proposition 1. *Consider the first-order Markov process with earnings states $\{\bar{y}(1), \dots, \bar{y}(L)\} \times \{\bar{\xi}(1), \dots, \bar{\xi}(M)\}$ and transition probability matrix P . Any set of moments derived from an $N \times T$ earnings (repeated) cross section, with $T \geq 1$ and observations on earnings levels $y \in \mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\}$ can generate a Jacobian of the moments with respect to the transition probabilities P with a rank at most L .*

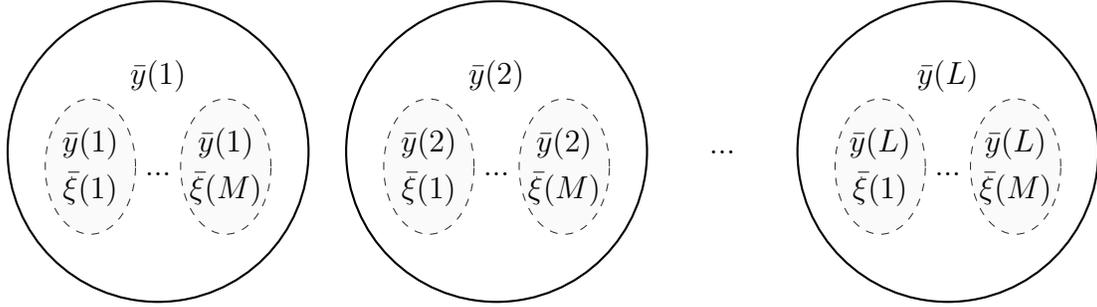
Proof In Appendix D.1.

Given that the estimation of the transition probability matrix P requires the estimation of $(L \times M) \times (L \times M - 1)$ parameters, it follows from Proposition 1 that the parameters of the earnings process are severely under-identified when using cross-sectional data on earnings. The intuition for this result is as follows. Earnings data only contain information about the observed component of the earnings state: the earnings levels $y \in \mathcal{Y}$. This characterizes an “observed” earnings process, with states $y \in \mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\}$ and transition probability matrix Q^{obs} given by

$$Q^{\text{obs}} = \begin{bmatrix} q_{11} & \dots & q_{1L} \\ \vdots & \ddots & \vdots \\ q_{L1} & \dots & q_{LL} \end{bmatrix} \quad (20)$$

where $q_{ij} = P(y_{t+1} = \bar{y}(j) | y_t = \bar{y}(i))$. The difference between the observed and unobserved earnings process is visualized in Figure 5. Denote by q_1, \dots, q_L the stationary probabilities of the observed income process. These are functions of the transition probabilities q_{ij} in Q^{obs} , which are, in return, also functions of the transition probabilities of the unobserved earnings process in matrix P . It follows that any moment derived from an observed cross section of earnings levels will be a function of the stationary probabilities q_1, \dots, q_L , which has dimensionality L , and, as such, the Jacobian of any set of moments with respect to the transition probabilities in P can have a rank of at most L . In general, moments derived from the stationary distribution of a Markov process are never sufficient to identify its (unrestricted) transition probability matrix, but in this case, the identification issue is more stringent, because we only observe the stationary distribution of the smaller Markov chain (y, Q^{obs}) , rather than of the Markov chain $((y, \xi), P)$.

Figure 5: Observed earnings levels (outer circles, with black solid outline) versus hidden earnings states (grey ellipses within the outer circles, with dotted outline).



While cross-sectional data on earnings and employment does not suffice to identify the transition probabilities in matrix P , moments derived from these data are able to identify all the levels of the discrete earnings states in \mathcal{Y} .

4.2.2 Earnings panel

Next, it is assessed whether an earnings panel can be used to identify and estimate the earnings process introduced in Section 3.1. The conclusion that follows from this analysis is negative, and is summarized in Proposition 2.

Proposition 2. *Consider the first-order Markov process with earnings states $\{\bar{y}(1), \dots, \bar{y}(L)\} \times \{\bar{\xi}(1), \dots, \bar{\xi}(M)\}$ and transition probability matrix P . Any set of moments derived from an $N \times T$ earnings panel with observations on earnings levels $y \in \mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\}$ can generate a Jacobian of the moments with respect to the transition probabilities P with a rank at most $L \times (L - 1)$.*

Proof. In Appendix D.2.

It follows that an earnings panel is insufficient to estimate the parameters of the transition probability matrix P , which has $(L \times M) \times (L \times M - 1)$ parameters that need to be pinned down. The intuition behind Proposition 2 is similar to that of Proposition 1. A panel with earnings data only gives information about the observed component of the earnings state: $y \in \mathcal{Y}$, and therefore, is only informative about the observed Markov chain with transition probability matrix Q^{obs} given in Equation (20). Any set of moments that can be derived from an earnings and employment panel will be a function of these transition probabilities in Q , which has rank $L \times (L - 1)$, while the object of interest, the transition probability matrix P of the unobserved Markov chain, has rank $(L \times M) \times (L \times M - 1)$.

Identifying assumptions To solve the identification problem posed in Proposition 2, identifying assumptions can be introduced. For example, one can assume that certain transition probabilities are zero, or equal to other transition probabilities. Alternatively, one can consider the restrictions imposed by an AR(1)-GARCH(1,1) model. Recall that the AR(1)-GARCH(1,1) can be mapped into the general earnings process proposed by this paper, as was shown in Section 3.1.1. The AR(1)-GARCH(1,1) specification only depends on five parameters, while the earnings process of interest in the application of this paper, introduced in Equation (13), depends on twenty parameters. Following Proposition 2, a vector of moments derived from an earnings and employment panel will have a rank of at most six when trying to identify the parameters of the earnings process in Equation (13). This means that under the restrictions of a GARCH process, the earnings process of Equation (13) is identifiable from an earnings panel.

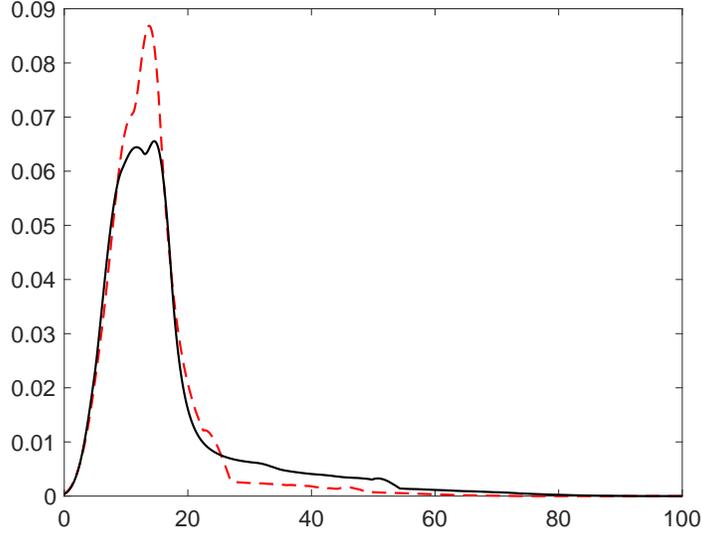
4.2.3 Cross-sectional data on net worth

The next step to analyze is whether moments from the cross-sectional distribution of net worth can be used to identify and estimate the parameters of this earnings process. At first thought, one may wonder how the cross sectional distribution of net worth can identify the parameters of the earnings process if the cross sectional distribution of earnings can not. If individuals facing different earnings risk make different saving choices, however, the cross sectional distribution contains additional information that is not reflected by the earnings distribution. Unlike for the analysis of the employment and wage distribution, when analyzing the cross-sectional distribution of net worth, one needs to rely on numerical approximations to the recursive competitive equilibrium of the general equilibrium model, as a closed-form analytical solution for the cross-sectional net worth distribution is not available.

As follows from the proof of Proposition 1, the moments of the earnings distribution are a function of the stationary probabilities of the earnings process, implying that two earnings processes with the same stationary distribution will generate the same earnings distribution. This is not the case for the savings distribution, as can be illustrated through Example 3 and corresponding Figure 6.

Example 3. In this example, heterogeneity with respect to earnings levels is ignored, and the following earnings states are considered: (unemployed, $\xi(1)$ or $\xi(2)$), (employed, $\xi(1)$), (employed, $\xi(2)$). Specifically, two processes with the following transition matrices

Figure 6: Comparing two savings distributions, computed from two different transition matrices, but the same stationary distribution as in Equation (21).



To generate this figure, the other structural parameters of the general equilibrium model are set to:
 $\beta = 0.95$, $\gamma = 2$, $\alpha = 0.36$, $\delta = 0.025$, $\mu = 0.15$.

are considered:

$$P^{\text{red}} = \begin{bmatrix} 0.5 & 0.4 & 0.1 \\ 0.2 & 0.6 & 0.2 \\ 0.05 & 0 & 0.95 \end{bmatrix}, \quad P^{\text{black}} = \begin{bmatrix} 0.380 & 0.286 & 0.334 \\ 0.200 & 0.714 & 0.086 \\ 0.070 & 0 & 0.930 \end{bmatrix}$$

where both distributions have stationary distribution $P^{\text{stationary}} = (0.125, 0.125, 0.750)$ (21)

As can be seen in Figure 6, these two earnings processes imply two different saving distributions as a solution to the general equilibrium model, even though these earnings processes generate the same earnings and employment distribution. The difference between the saving distributions the two earnings processes generate can be measured, for example, through the variances of these savings distributions: 22.9 (corresponding to the red-dotted distribution line) versus 28.0 (solid black). <

To analyze (local) identification, one can numerically evaluate the Jacobian to assess the identifying power of the cross-sectional distribution of net worth. In principle, one has to compute the Jacobian for many possible parameter values P and assess its rank. As this will not be the source of identification used to estimate the transition probabilities in this paper, the Jacobian will only be computed for two examples.

Example 4. Consider the same earnings process as in Example 3, with earnings states (unemployed, $\bar{\xi}(1)$ or $\bar{\xi}(2)$), (employed, $\bar{\xi}(1)$), (employed, $\bar{\xi}(2)$). and the following transition matrix:

$$P = \begin{bmatrix} 0.5 & 0.4 & 0.1 \\ 0.2 & 0.6 & 0.2 \\ 0.05 & 0 & 0.95 \end{bmatrix} \quad (22)$$

The other structural parameters of the general equilibrium model are set to: $\beta = 0.95$, $\gamma = 2$, $\alpha = 0.36$, $\delta = 0.025$, $\mu = 0.15$.

Next, consider the following choice of moments: $(\frac{1}{100}\text{Var}(k_i), \frac{1}{10}\text{Skewness}(k_i), \frac{1}{10}\text{Percentile}_{95}(k_i), \frac{1}{10}\text{Percentile}_{90}(k_i)/\text{Percentile}_{10}(k_i), \frac{1}{10}\text{Percentile}_5(k_i), \frac{1}{5}\text{Percentile}_{70}(k_i)/\text{Percentile}_{30}(k_i))$. Numerically computing the Jacobian of this moment vector with respect to the transition probability matrix in Equation (22) gives the following Jacobian:

$$J = \begin{bmatrix} 0.991 & -0.239 & 0.080 & 0.000 & 4.008 & 0.000 \\ -0.879 & -0.534 & -0.387 & 0.000 & 0.000 & 0.000 \\ -0.773 & -0.869 & 0.080 & 0.000 & 4.080 & 0.000 \\ -0.515 & -0.650 & 0.003 & 0.000 & 0.000 & 0.000 \\ 2.750 & 2.041 & 3.546 & -4.099 & 4.008 & -0.654 \\ -0.887 & 0.594 & 0.805 & -0.541 & 4.008 & 0.755 \end{bmatrix}$$

Computing the rank of this Jacobian gives a value of 6, indicating that observing these specific moments of the wealth distribution locally identifies the full transition matrix. The condition number of this Jacobian, a statistic used to characterize the identification strength of the moments and parameters, see Iskrev (2010), is approximately 40.22, which suggests that identification at the current parametrization is fairly strong. \triangleleft

Example 4 deals with the identification of six transition probabilities, which requires the computation of six moments of the net worth distribution.¹⁸ In the case where one wants to estimate all transition probabilities of the earnings process of Equation (13) using the cross sectional distribution of net worth, this requires the computation of twenty moments. Although in principle this can lead to a Jacobian matrix that is full rank, it is probable that this identification will be weak: when computing twenty moments from the same distribution, it is likely that the way these moments depend on the underlying transition probabilities will be strongly correlated. This is confirmed by Example 5.

¹⁸The other three transition probabilities follow from the first six.

Example 5 Next, consider the earnings process of Equation (13), with earnings states (unemployed, $\bar{\xi}(1)$ or $\bar{\xi}(2)$), $(y(1), \bar{\xi}(1))$, $((y(1), \bar{\xi}(2))$, $(y(2), \bar{\xi}(1))$, $(y(2), \bar{\xi}(2))$, evaluated at the following transition probability matrix:

$$P = \begin{bmatrix} 0.60 & 0.10 & 0.10 & 0.10 & 0.10 \\ 0.02 & 0.92 & 0.02 & 0.02 & 0.02 \\ 0.25 & 0.11 & 0.60 & 0.02 & 0.02 \\ 0.02 & 0.02 & 0.02 & 0.92 & 0.02 \\ 0.25 & 0.05 & 0.05 & 0.05 & 0.60 \end{bmatrix}$$

A vector of twenty moments is defined, consisting of the variance, skewness and several ratios of percentiles of the net worth distribution, the variance of the earnings distribution, and the average unemployment rate. Computing the Jacobian evaluated at P given above and the same choice of structural parameters as in Example 4, and $v_l = 0.5$, a rank of 20 is obtained. However, the condition number is equal to $3.7e5$, confirming the conjecture of weak identification. \triangleleft

4.2.4 A panel with individuals' net worth levels and earnings levels

This last subsection analyzes how a panel of individuals' net worth levels, supplemented with earnings levels, can be used to identify the parameters of the earnings process proposed in this paper. Under a set of assumptions, a panel on the net worth holdings and the earnings levels of individuals can be shown to provide a direct mapping to the underlying risk states ξ (up to label swapping). The assumptions are summarized in Assumptions 1 and 2, and the final result is described in Proposition 3.

Assumption 1: Ordering *The optimal saving rule of the households is denoted by $k'(k, (y, \xi))$ with domain $y \in \mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\}$, $\xi \in \mathcal{X} = \{\bar{\xi}(1), \dots, \bar{\xi}(M)\}$ and $k \geq 0$. There exists an ordering of the elements in $\epsilon_{ij} = (\bar{y}(i), \bar{\xi}(j)) \in \mathcal{Y} \times \mathcal{X}$ over j , where the ordered elements are denoted by $\tilde{\epsilon}_{ij}$ such that*

$$k'(k, \tilde{\epsilon}_{i1}) > k'(k, \tilde{\epsilon}_{i2}) > \dots > k'(k, \tilde{\epsilon}_{iM}) \quad \forall k \geq 0.$$

This ordering exists for each $i = 1, \dots, L$ but is allowed to differ for each i .

Assumption 1 indicates that there should be a fixed ordering in the saving choices over the risk states ξ . That is, if a certain combination of $(\bar{y}(i), \bar{\xi}(j))$ implies a higher next period's savings level k' than at another risk level ξ but the same earnings level y , it

should do so at each level of previous savings k .

Assumption 2a: Locally Lipschitz continuous *The domain of k can, for each value of $\bar{y}(j) \in \mathcal{Y}$ be subdivided into K_N^j subdomains $[k_1^j, k_2^j], \dots, [k_i^j, k_{i+1}^j], \dots, [k_{K_N^j}^j, k_{K_N^j+1}^j]$, such that for each subdomain $[k_i^j, k_{i+1}^j]$, the saving function $k'(\cdot, (\bar{y}(j), \bar{\xi}(i)))$ satisfies*

$$d(k'(k_i^j, (\bar{y}(j), \bar{\xi}(l))), k'(k_{i+1}^j, (\bar{y}(j), \bar{\xi}(l)))) < Z_{ijl} \cdot d(k_i^j, k_{i+1}^j), \quad \forall l = 1, \dots, M-1, \quad \forall \bar{y}_j \in \mathcal{Y}$$

where $Z_{ijl} = d(k'(k_i^j, (\bar{y}(j), \bar{\xi}(l))), k'(k_i^j, (\bar{y}(j), \bar{\xi}(l+1))))$.

Assumption 2b: Minimum number of observations *Each subdomain specified in Assumption 2a contains at least M observations with earnings level $\bar{y}(j)$ that can be labeled (without loss of generality) $(k_{1,t}, k_{1,t+1}), (k_{2,t}, k_{2,t+1}), \dots, (k_{M,t}, k_{M,t+1})$ for which $k_{i,t} > k_{i+1,t}$ and $k_{i,t+1} < k_{i+1,t+1}$ for $i = 1, \dots, M-1$.*

Assumption 2a comes down to assuming that the savings function k' is locally Lipschitz continuous, in other words, is not too steep. Whether this assumption is satisfied depends both on the (local) Lipschitz coefficient, in this case defined as $Z_{ijl} = d(k'(k_i, (\bar{y}(j), \bar{\xi}(l))), k'(k_i, (\bar{y}(j), \bar{\xi}(l+1))))$, that is, the difference between next period's net worth of an individual with risk state $\bar{\xi}(l)$ and an individual with risk state $\bar{\xi}(l+1)$, as well as on the steepness of $k'(\cdot, (\bar{y}(i), \bar{\xi}(l)))$ in the subdomain. Simply put, individuals with different risk states should differ enough in their saving choices, and how different this should be depends on the steepness of the savings function $k'(k, (\bar{y}(j), \bar{\xi}(l)))$ on the subdomain. Assumption 2b asserts that there are enough observations within these subdomains such that the different risk states can actually be identified.

Proposition 3 *Consider the first-order Markov process with earnings states $\{\bar{y}(1), \dots, \bar{y}(L)\} \times \{\bar{\xi}(1), \dots, \bar{\xi}(M)\}$ and transition probability matrix P . Assume Assumptions 1 and 2a-2b hold. A panel of individual net worth k_{it} , and earnings/employment status y_{it} , with $t = 1, \dots, T$ and $i = 1, \dots, N$, $T \geq 2$, can be used to perfectly retrieve a panel of (reordered) unobserved states ξ_{it} for $t = 2, \dots, T$ and $i = 1, \dots, N$.*

Proof In Appendix D.3.

The intuition behind this result is as follows. The conditions in Assumptions 1 and 2a-b assert that given current earnings level y , previous savings k and end-of-year savings k' , one can deduce which observations have the same risk state ξ . For this, it is necessary that people that face different levels of risk ξ make sufficiently different saving decisions

k' . This is enforced by the assumptions. If so, then conditional on their previous savings k and current earnings level y , their saving choice for next period reveals the amount of risk they perceive. An $N \times T$ panel of net worth k_{it} and earnings y_{it} can then be used to “reverse-engineer” an $N \times (T - 1)$ panel of (reordered) earnings states $\tilde{\epsilon}$. Next, this reverse-engineered panel can be used to identify and estimate the parameters of transition probability matrix P .

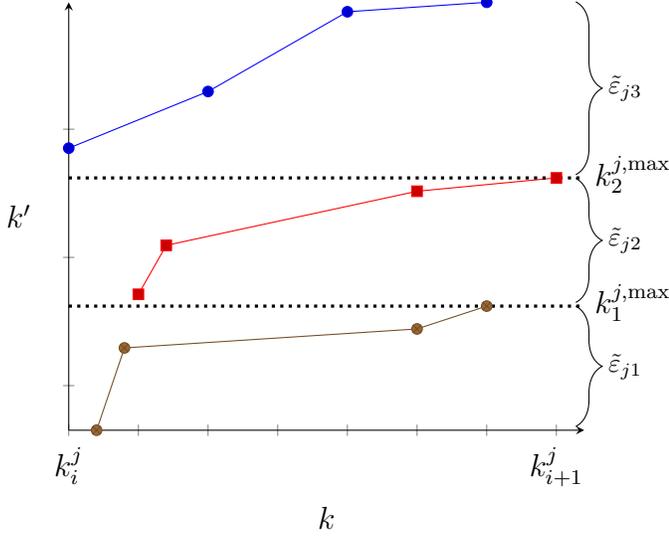
Assumptions 1 and 2a-b, and their implication for the possibility to reverse-engineer the risk states ξ are visualized in Figure 7a for an example with three different risk states, that is, $M = 3$. Assumption 1 imposes that the three lines in this figure, that correspond to sets of (k, k') observations from three different risk states but the same income level $\bar{y}(j)$, do not intersect on the observable domain, and have a fixed ordering of realizations k' across all values of k . Furthermore, assumption 1 imposes a labeling of these re-ordered risk states, being $\tilde{\epsilon}_{j1}$, $\tilde{\epsilon}_{j2}$ and $\tilde{\epsilon}_{j3}$ respectively. Next, assumption 2a ensures that the interval $[k_i^j, k_{i+1}^j]$ is such that the three lines on which the observations lie fall inside three non-overlapping subranges of k' , with range boundaries $[0, k_1^{j,\max}]$, $[k_1^{j,\max}, k_2^{j,\max}]$ and $[k_2^{j,\max}, \infty)$. At last, assumption 2b ensures that there exist at least three observations that are such that the range-boundaries $k_1^{j,\max}$ and $k_2^{j,\max}$ can be defined. $k_2^{j,\max}$ follows from being the largest value of k' and most right (largest k) point for which there does not exist another point that has a lower value of k yet a higher value of k' . Next, all observations that have a value of k' above $k_2^{j,\max}$ are labeled as $\tilde{\epsilon}_{j3}$, and are removed from the sample. Then, $k_1^{j,\max}$ is defined in the same way and results in the labeling of observations with risk state $\tilde{\epsilon}_{j2}$. The remaining observations are labeled as $\tilde{\epsilon}_{j1}$.

Some analogies can be drawn between the assumptions on instrumental variables in treatment effect estimation, and the assumptions given above. Assumption 1 is somewhat analogous to the monotonicity assumption used in the estimation of treatment effects using instrumental variables, assuming that the instrumental variable has a monotonic effect on the treatment choice (Imbens and Angrist, 1994). The monotonicity assumption is also sometimes called the no-defier assumption, that is, assuming the monotonic effect is positive, there are no individuals that decide not to take the treatment after the instrumental variable increases, while they would have taken the treatment for a lower value of the instrumental variable. Assumption 2a can be compared to the idea of a relevance assumption on instrumental variables when estimating treatment effects, which requires that the instrumental variable sufficiently affects the treatment choice of the individuals. Here, the assumption is that the saving choices of individuals with

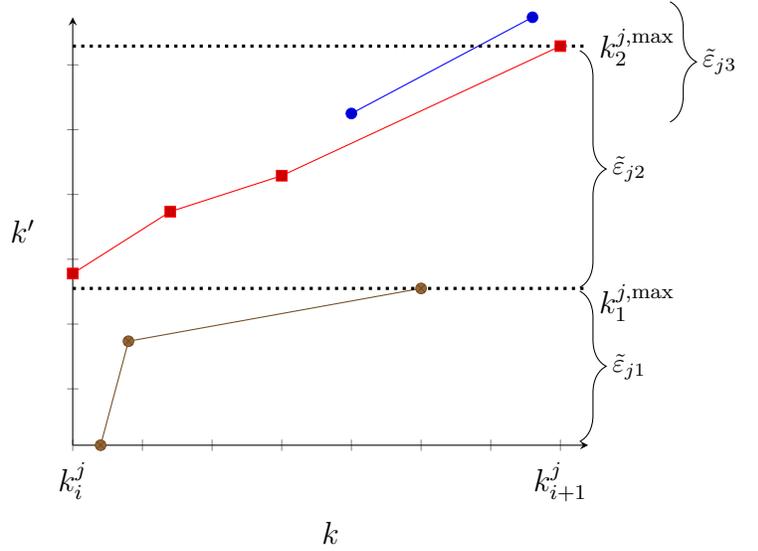
different earnings states are sufficiently different. Assumption 2b can be compared to the assumptions that there are compliers in the sample: individuals that only take the treatment for a sufficiently high value of the instrumental variable.

Figure 7: A visualization of Assumptions 1, 2a-2b and Proposition 3. Dots represent observations of (k, k') with the same earnings level $y = \bar{y}(j)$. Dots in the same color have the same risk state ξ .

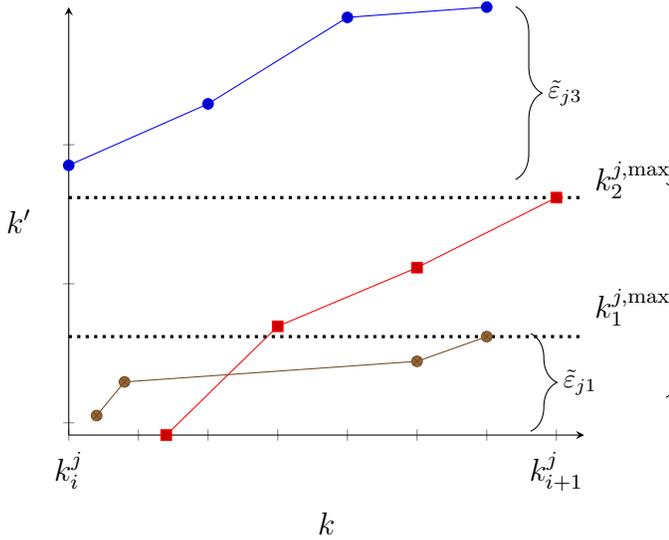
(a) This example satisfies all assumptions.



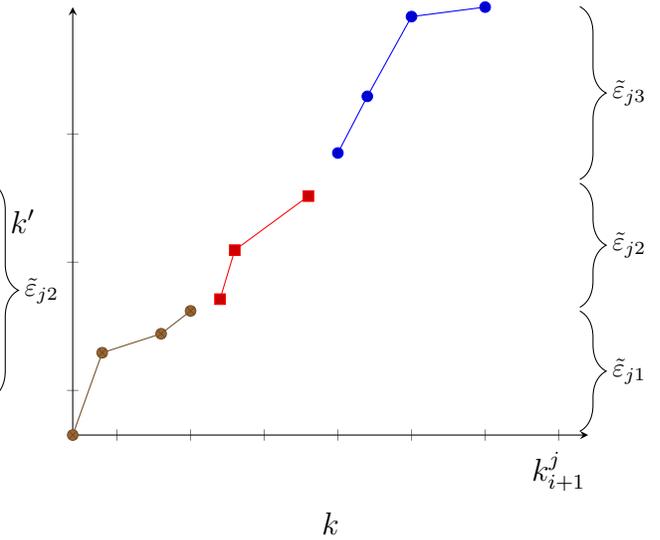
(b) This example fails Assumption 2a.



(c) This example fails Assumption 1.



(d) This example fails Assumption 2b.



Next, it is of interest to understand the implications of these assumptions better, and, furthermore, to understand the implication of their potential violations. If Proposition 3 holds, then it means that the risk states (y, ξ) are indirectly observed. It follows that the transition probability matrix P characterizing the earnings process can then

be estimated using maximum likelihood, as it simply boils down to the estimation of a transition probability matrix of an observed Markov chain. Proposition 3 thus leads to a non-parametric identification method of the transition probabilities.

However, the assumptions under which Proposition 3 hold may be too stringent. Consider Assumption 2a. A visualization is presented in Figure 7b of a case where this assumption fails to hold. In this example, it is not possible to create an interval $[k_i^j, k_{i+1}^j]$ such that $k'(k, \tilde{\varepsilon}_{j3}) > k'(k, \tilde{\varepsilon}_{j2})$ for all observations in the interval, while simultaneously having observations in this interval for which $k_i > k_{i+1}$ and $k'_i < k'_{i+1}$ for $i = 1, \dots, L - 1$ (as required by assumption 2b). The function $k'(k, \tilde{\varepsilon}_{j2})$ is too steep. In words, this means that the saving choices of individuals with risk states $\tilde{\varepsilon}_{j2}$ and $\tilde{\varepsilon}_{j3}$ are not different enough, and are not distinguishable. The example in Figure 7d visualizes what happens if Assumption 2b is violated. In this case, it is not possible to construct the range-boundaries $k_1^{j,\max}$ and $k_2^{j,\max}$, because there are no observations within the interval that allow us to draw conclusions on the risk states individuals are in. Only if we have at least two observations where one individual has a higher amount of current savings, but chooses a lower value of future savings than the other individual, can we conclude that these individuals are in different risk states. This is not the case in Figure 7d, implying violation of Assumption 2b.

In both cases above where Assumption 2a and 2b are violated, the issue is likely to disappear as $N \rightarrow \infty$, because as the number of observations increases, the number of intervals K_N^j can be increased and consequently the width of the intervals will decrease, such that it is more likely that both assumption 2a and 2b will be simultaneously satisfied. The probability of misclassification of certain observations due to violation of Assumptions 2a-b thus decreases to zero as $N \rightarrow \infty$. The assumptions do rule out risk neutrality, or near-risk neutrality, as under such parametrizations, the saving decisions of households under the different risk states are too similar, and no intervals can be made that simultaneously satisfy Assumptions 2a and 2b. Similarly, the probability distributions the individuals face under the different risk states should be sufficiently different. This likely is satisfied in the application, which only considers two different risk states. The assumption is also more likely to be violated at very high savings levels, because income risk then becomes less important, and the saving decisions between individuals with different saving levels become more similar.

Assumption 1 requires a strict ordering of saving decisions over risk states, but this ordering is allowed to vary across income levels. Figure 7c visualizes a potential violation of this assumption. However, violations of this assumption do not seem economically

relevant. If individuals in a certain risk state save more at low current levels of saving, they will also do so at high current levels of savings. The difference between the saving decisions k' does become smaller at higher levels of k , as also alluded to in the previous paragraph, because at high savings levels, income risk is less important, but the ordering is not expected to change.

To be less reliant on the assumptions discussed above, instead of using the non-parametric classification method that follows from Proposition 3, and estimating the transition probability matrix of the reverse-engineered states using maximum likelihood estimation and presenting these as the final estimates, a transition probability matrix computed based on the idea of Proposition 3 is used as the auxiliary model parameters in an indirect inference procedure instead. The advantage of such an approach is that the auxiliary model is not required to be fully correctly specified, but, as the auxiliary model grows closer to the true model (that is, if assumptions 1 and 2 are satisfied), the indirect inference estimator of these parameters will converge to be fully efficient. The indirect inference approach will be further discussed in Section 5.

4.3 Identification of the other structural parameters

Identification of the other structural parameters is discussed in detail in Appendix C. The source of identification of the first-step parameters α (elasticity of capital), δ (depreciation rate) and μ (unemployment benefits) are moments derived from the firm-side and government-side of the general equilibrium model. It follows that these parameters can be estimated independently of all other parameters, using only the means of several aggregate macro variables. Risk aversion can, as in Janssens (2019), be identified from the cross-sectional distribution of net worth. As the earnings process will be identified from the panel dimension and thus rely on a different source of identification, these parameters will most likely be jointly identified, although this can only be verified numerically. The discount factor β is not estimated, but follows as a solution to the model for a given interest rate.

5 Econometric Approach

This section discusses the estimation procedure. It uses the two-step indirect inference of Janssens (2019), briefly described in Section 4.1. The estimation procedure for the first-step parameters α , δ and μ will first be described. After performing this step, these parameters will be fixed to their estimates, and will be used in the second step, which is described in larger detail in the second subsection below.

5.1 First-step estimation of α , δ and μ

As derived in Appendix C, the following moment conditions can be used to consistently estimate α , δ and μ :

$$\begin{aligned}\mathbb{E}\left(\frac{w_t^{\text{obs}}(1 - u_t^{\text{obs}})}{Y_t^{\text{obs}}} - (1 - \alpha)\right) &= 0 \\ \mathbb{E}\left(\frac{Y_t^{\text{obs}}}{(\delta K_t)^{\text{obs}}} - \frac{r_t^{\text{obs}} + \delta}{\delta\alpha}\right) &= 0 \\ \mathbb{E}\left(\frac{\tau_t^{\text{obs}}}{u_t^{\text{obs}}} - \mu\right) &= 0,\end{aligned}$$

under the assumption that the sample means of the observed counterparts of the variables, denoted τ_t^{obs} , (the average tax rate), u_t^{obs} , (the unemployment rate), w_t^{obs} (the average wage rate), Y_t^{obs} (output), r_t^{obs} (interest rate) and $(\delta K_t)^{\text{obs}}$ (depreciation of capital) converge to their long-run values.¹⁹ Method of moments can be used to obtain consistent estimates of these three parameters.

5.2 Second-step estimation of P , v_l and γ

After estimating the first-step parameters, the other parameters of the structural model can be estimated conditional on the estimated values of the first-step parameters. Of the second-step parameters, the earnings process parameters are of main interest. Based on the identification analysis in Section 4.2, two different identification strategies arise that can be used to estimate the parameters of the earnings process:

1. Estimate the earnings process parameters using a panel data set on earnings, employment and net worth data using the identification strategy implied by Proposition 3.
2. Estimate the earnings process parameters using a panel data set on earnings and employment data, and restrict the parameters of the earnings process, for example, using the mapping into an AR(1)-GARCH(1,1) process from Section 3.1.1.

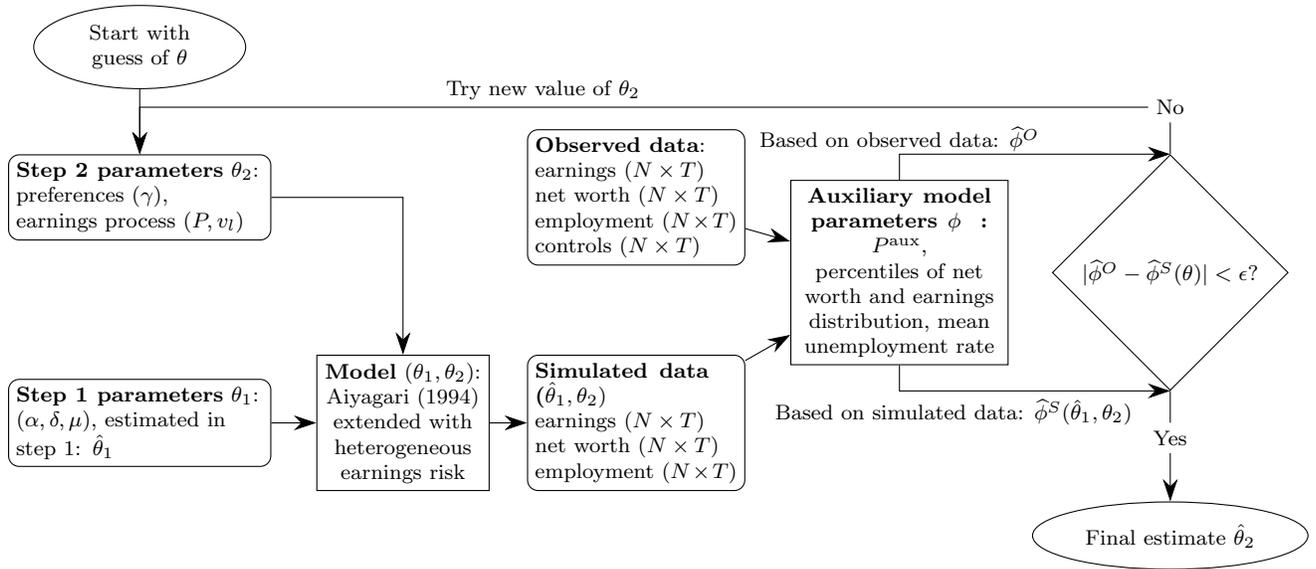
In addition, I will also consider a combination of these two identification strategies, that is, I will use a panel on earnings, employment and net worth data, and estimate

¹⁹ $(\delta K_t)^{\text{obs}}$ is used because this variable is directly observed in the data, see also the Data section. A similar condition can be derived using K_t^{obs} instead.

the AR(1)-GARCH(1,1)-restricted earnings process. The estimation procedures corresponding to these three approaches are visualized in three flow charts, in Figures 8-10, and will be discussed in more detail below.

Identification strategy 1 The first identification strategy uses a panel on net worth, earnings and employment status (as well as other observable characteristics of individuals) to estimate the unrestricted earnings process. The estimation procedure is visualized in Figure 8. In this procedure, an auxiliary model with parameters ϕ is estimated for both the observed data and simulated data. The optimization procedure then aims to find a parameter combination $\theta = (P, v_l, \gamma)$ that minimizes the distance between the auxiliary model parameters of the observed data $\hat{\phi}^O$ and the simulated data, $\hat{\phi}^S(\theta)$. Simulating data from the general equilibrium model requires solving the recursive competitive equilibrium, and is therefore computationally intensive.

Figure 8: Flowchart describing estimation of the unrestricted earnings process, using a structural model and net worth data.



The auxiliary model parameters comprise the following elements. First of all, recall the auxiliary transition probability matrix presented in the Data section in Table 1. The method to compute this transition probability matrix actually uses the intuition of the labeling procedure implied by Proposition 3. As such, if Assumptions 1-2 hold as the number of observations $N \rightarrow \infty$ and the number of clusters grows large, the auxiliary transition matrix of Table 1 is already a consistent estimator of the transition probability matrix P . Here, the matrix is used as one of the parameters of the auxiliary model, to

be more robust against violations to the assumptions imposed in smaller samples.²⁰ The other auxiliary model parameters include various moments of the net worth and earnings distribution, as well as the average unemployment rate. The full objective function is given below in Equation (23).

Objective function for second step of the unrestricted model:

$$\max_{P, \gamma, v_l} \left(\hat{\phi}^O - \hat{\phi}^S(\theta) \right)' W \left(\hat{\phi}^O - \hat{\phi}^O(\theta) \right) \quad (23)$$

$$\text{where } \hat{\phi} = \left(\text{vec}(\hat{P}_{\text{non-param}}), \text{Prctl}_{90}(k_{it})/\text{Prctl}_{15}(k_{it}), \text{Prctl}_{99}(k_{it})/\text{Prctl}_{90}(k_{it}), \dots \right. \\ \left. \dots \text{Prctl}_{50}(k_{it})/\text{Prctl}_{15}(k_{it}), \text{Prctl}_{90}(w_{it})/\text{Prctl}_{30}(w_{it}), \bar{u}_t \right) \quad (24)$$

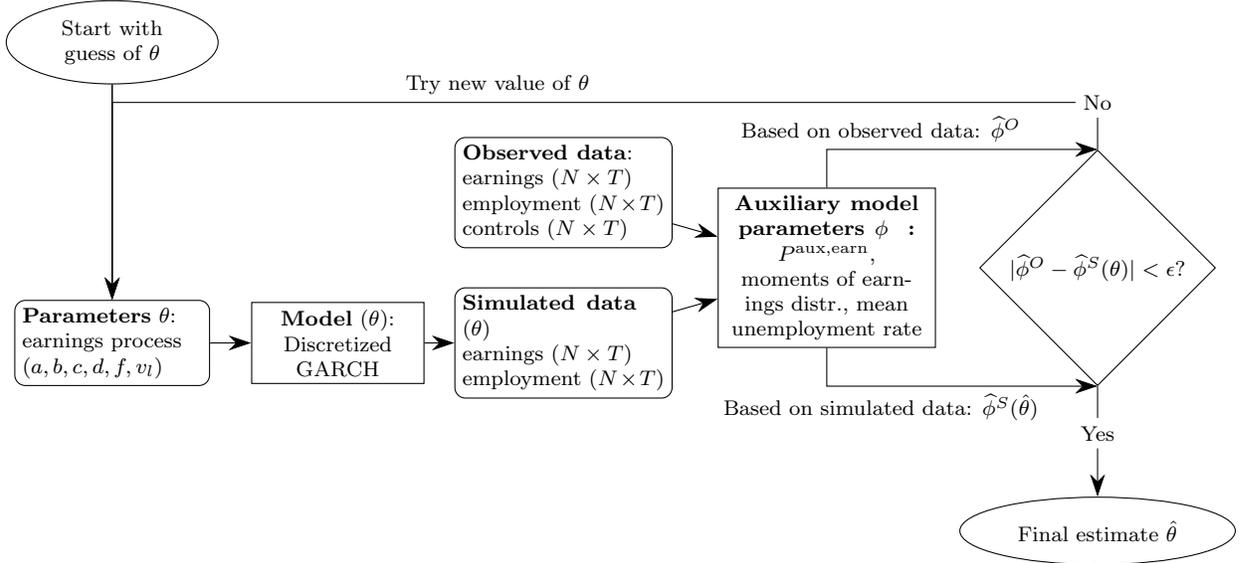
Here \bar{u}_t denotes the sample mean of the unemployment rate. $\hat{P}_{\text{non-param}}$ is described in Equations (1) and given in Table 1. W is a weighting matrix, currently not chosen as the optimal efficient matrix, but chosen to ensure all elements of \hat{m}^{obs} have a similar scale, by letting all weighted moments ($\sqrt{W_{ii}}\hat{m}_i^{\text{obs}}$) have values between zero and one. A diagonal matrix is used, where all weights related to $\hat{P}_{\text{non-param}}$ receive a weight of 1. The other moment weights satisfy (following the order of Equation (24)): $\sqrt{W_{ii}} = (\frac{1}{20000}, \frac{1}{100}, \frac{1}{1000}, \frac{1}{2000}, 1)$. The notation $\text{Prctl}_x(k_{it})$ denotes the x^{th} percentile of the net worth distribution.

Identification strategy 2 Consider Figure 9, visualizing the flow chart explaining the estimation procedure for the second identification strategy. This strategy does not require solving the structural model for the estimation of the earnings process parameters P and v_l , as the earnings process is assumed to be evolving exogenously, and, in this case, can be identified fully from earnings data only. Thus, in the indirect inference procedure, one does not need to solve the recursive competitive equilibrium to simulate data from the general equilibrium for each combination of parameters, but instead one simulates from the discretized GARCH model, which is analogous to simulating data from a first-order Markov process. Next, the estimation procedure for P and v_l aims to find the parameters that minimize the distance between the auxiliary model parameters computed for the simulated earnings and employment data and the observed data.

For the auxiliary model, an auxiliary transition matrix $P^{\text{aux, earn}}$ is computed that measures the (empirical) probability of transitions between unemployment, low (below av-

²⁰A slightly different auxiliary transition probability matrix is used in the indirect inference procedure, based on a different choice for the number of saving bins and control groups. This auxiliary transition probability matrix is given in Appendix E.

Figure 9: Flow chart describing estimation of earnings process restricted using GARCH process, without structural model and without net worth data.



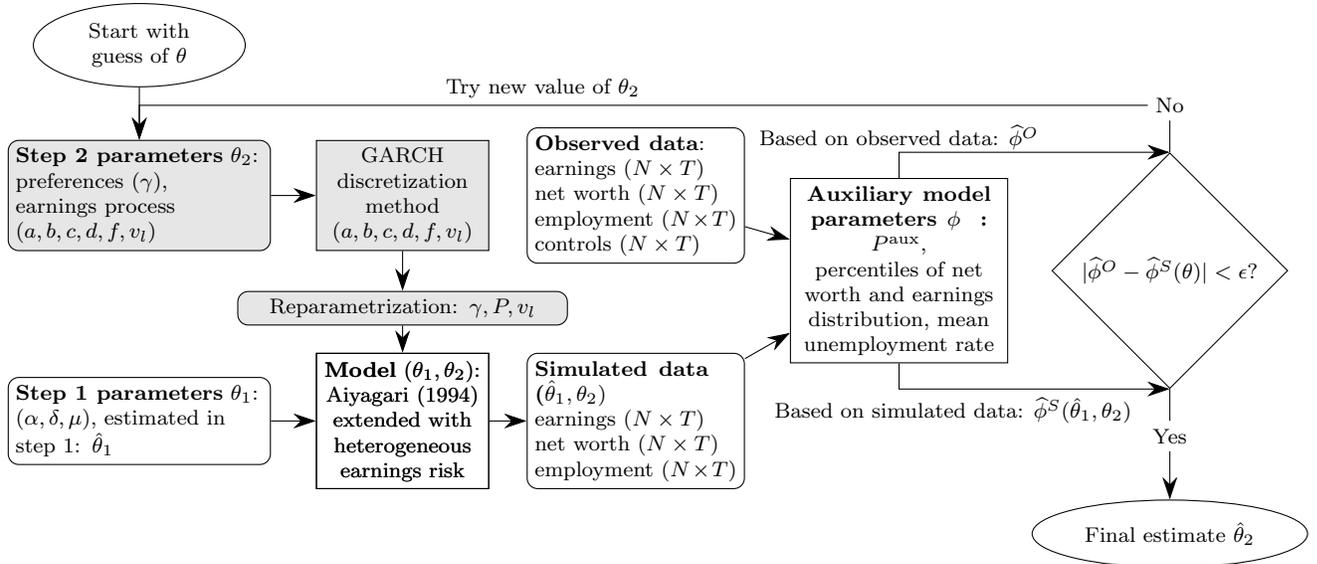
erage) earnings and high (above average) earnings. For the observed data, this division into high- and low-earnings groups is done relative to its control cluster. For more details on how these control clusters are constructed, see also Sections 2.2 and 2.3. This auxiliary transition matrix is supplemented with some additional moments: the average earnings of individuals with below average earnings, the average earnings of individuals that have above-average earnings, the unemployment rate in the sample, the average over all non-zero earnings and the variance over all non-zero earnings. As in Figure 9, denote the auxiliary model parameters computed for the observed dataset by $\hat{\phi}^O$ and the auxiliary model parameters for the simulated data by $\hat{\phi}^S(\theta)$, where θ in this case are the parameters of the earnings process, P and v_l . The indirect inference procedure stops when the distance between $\hat{\phi}^O$ and $\hat{\phi}^S(\theta)$ is minimized. An identity matrix is used for weighting.

The other structural parameters of the model, being risk aversion γ and β , can afterwards be obtained in a third step, in case one also has at least one higher-order moment (at least second-order) of the net worth distribution. This does require solving the recursive competitive equilibrium of the general equilibrium, and simulating net worth data from the model. This paper will, however, focus on the estimates of the earnings process parameters.

Combining identification strategy 1 and 2 The third approach combines both identification strategies. The corresponding estimation procedure is visualized in Figure

10, where the differences with the estimation procedure of identification strategy 1, visualized in Figure 8, are highlighted in grey. This approach uses a panel of net worth, earnings and employment status (as well as various individual-level control variables) to estimate the earnings process, where the earnings process parameters are restricted by imposing that it behaves according to a discretized AR(1)-GARCH(1,1) process. The reason to consider this combination of strategies is because it can be used to evaluate the restrictiveness of the restrictions posed by the AR(1)-GARCH(1,1) process. In essence, the AR(1)-GARCH(1,1) model limits the notion of risk in the model, as it no longer relates to the entire probability distribution of future shocks, but is limited to a time-varying variance. The most important difference between the two approaches is the amount of parameters to be estimated: due to the restrictions of the discretized AR(1)-GARCH(1,1) process, the earnings process is characterized by only six parameters, instead of 21 as in the unrestricted process. Apart from these restrictions, the procedure is essentially the same as identification strategy 1 and uses the same auxiliary model.

Figure 10: Flowchart describing estimation of earnings process restricted using GARCH process, with structural model and with net worth data. Shaded elements indicate the elements that are different from the flowchart in Figure 8.



Optimization All three estimation procedures outlined above require the use of an optimization algorithm. For the GARCH-restricted estimates, the following procedure is used. First, a grid search is performed over the range of feasible AR(1)-GARCH(1,1) parameters. Next, the ten lowest objective function evaluations and their corresponding parameters are used as a starting point of a pattern search optimization algorithm.²¹

²¹For an overview of pattern search algorithms, see Torczon (1997). This paper relies on the imple-

The estimate corresponding to the lowest objective is reported. For the unrestricted estimates, grid search is not feasible due to the high dimensionality of the parameter vector. However, the GARCH-restricted estimates provide good starting points for the unrestricted estimates. Thus, ten random starting points are generated around the GARCH-restricted estimates that are obtained from the procedure that combines both identification strategies. Again, pattern search is used on these ten starting points, and the estimate corresponding to the lowest objective is reported.

6 Estimates of the Heterogeneous Earnings Risk Model

This section summarizes the estimation results from the heterogeneous earnings risk model. The first subsection discusses the first-step parameters. The next subsection discusses the results obtained on the parameters of the earnings process, and how these vary for the three approaches outlined in the Econometric Approach section. In the third subsection, I explore the implications of heterogeneous earnings risk in an incomplete markets model, by comparing the recursive competitive equilibrium solution with that of a model without heterogeneous earnings risk.

6.1 First-step estimates of α (elasticity of capital), δ (depreciation rate), μ (unemployment benefits)

Table 2: Estimates for α (elasticity of capital), δ (depreciation rate), and μ (unemployment benefits). Standard errors obtained using block-bootstrap with blocks of 3-years width.

Parameter	Estimate	Standard errors
α	0.373	0.003
δ	0.026	0.004
μ	0.132	0.011

This subsection presents the estimates obtained for the parameters α (elasticity of capital), δ (depreciation rate), and μ (unemployment benefits). Recall that these parameters are identified and estimable from a subset of moments that is independent of all the other parameters, and they do not require the full solution of the stationary equilibrium. They only require the solution to the firm-side problem (for α and δ) and the government budget constraint (for μ). The obtained estimates are reported in Table 2. As can be seen, the estimates are similar to those commonly used in calibrations, e.g., in Den Haan

mentation of pattern search in MATLAB, see <https://nl.mathworks.com/help/gads/patternsearch.html>.

(2010), and are estimated with high precision.

6.2 Estimation of earnings process parameters

This subsection presents and discusses the estimation results for the earnings process parameters, and, when applicable, the other structural parameters of the general equilibrium model. Results are presented for the three procedures outlined in Section 5.

6.2.1 Estimation results of identification strategy 1: using net worth, earnings and employment panel

This subsection presents the estimates obtained from the first identification strategy. This strategy enables the identification and estimation of the earnings process, without imposing additional restrictions on the earnings process. These unrestricted estimates are presented in Table 3. In Appendix E, the auxiliary model parameters are presented as simulated from the model, and compared to the auxiliary model parameters as estimated from the observed data.

Table 3: Estimates of unrestricted earnings process, using identification strategy 1.

Transition probabilities (from/to)		Unemployed	Low earnings		High earnings	
			Low risk	High risk	Low risk	High risk
Unemployed		0.318	0.482	0.052	0.009	0.139
Low earnings	Low risk	0.012	0.736	0.004	0.236	0.012
	High risk	0.081	0.367	0.448	0.043	0.061
High earnings	Low risk	0.021	0.106	0.043	0.792	0.038
	High risk	0.110	0.145	0.302	0.143	0.300
Stationary distribution		0.036	0.377	0.068	0.474	0.045

Note: Other parameter estimates: ν_l : 0.68, ν_h (implied): 0.58, β (implied): 0.974, γ (risk aversion): 5.34. Standard errors are (currently) unavailable due to the high computational demands imposed by the model in combination with the large number of parameters.

Several observations can be made about the estimation results. First of all, in terms of interpretation, starting by comparing the low and high-income states, the estimates of ν_l and ν_h show that individuals in the low-income state earn 32% below average wage, while the high income individuals earn 58% above average wage. Individuals face not only income inequality, but also income risk inequality. The second and fourth row of \hat{P} can be unambiguously interpreted as states with low risk. The probability of becoming unemployed is fairly low, and the probability of staying in the same state

next period is relatively large: 0.736 for the low-income-low-risk state, and 0.792 for the high-income-low-risk state. On the other hand, the third and fifth row represent the transition probabilities for the high-risk states. Comparing the low-income-low-risk and low-income-high-risk states (rows 2 and 3), the probability of becoming unemployed is about seven times as large in the high-risk state. Similarly, the probability of becoming unemployed is around five times as large in the high-income-high-risk state than in the high-income-low-risk state.

Risk encompasses more than simply job-loss probabilities. Comparing the second and third rows, one can see that individuals in the low-income-high-risk state (third row), having a higher probability of becoming unemployed, also have a considerably lower probability of moving into the high-income state ($0.043+0.061=0.104$), compared to the low-income-low-risk individuals (second row), who face a probability of 0.248 ($0.236+0.012$) of moving to a high-income state next period. This supports the notion that earnings risk is a multidimensional object, and, as the results below will show, a variance-based risk measure will not be able to capture this, as having a larger variance would also imply a larger probability for positive income shocks. The estimates show that the high-income-high-risk state in row five has a higher probability of moving to the low-income level than the high-income-low-risk individuals in row four.

Another interesting result is the difference in stability of the high and low-risk states. Comparing rows two and three, and rows four and five, it can be seen that low-risk states (row two and four) are more stable than the high-risk states (row three and five), in the sense that the probabilities of remaining in the same state next period (the diagonal elements) are much larger in the low-risk than in the high-risk states.

6.2.2 Estimation results of identification strategy 2: using earnings and employment panel and GARCH restrictions

The results of the second identification strategy are summarized in Table 4, where the AR(1)-GARCH(1,1) parameters are presented, as well as the corresponding discretization into a first-order Markov process, denoted in the Table by $\hat{P}_{\text{earnings}}^{\text{GARCH}}$. P^{stat} denotes the corresponding stationary distribution.

There are several important differences between and similarities to the estimates that follow from the first identification strategy, presented in the previous subsection. A notable similarity is the difference in stability between high- and low-risk states; the estimates of the second identification strategy confirm that individuals in high-risk states

Table 4: GARCH estimates using earnings data only. Estimated using indirect inference. AR(1)-GARCH(1,1) is given by: $y_{it+1} = a + by_{it} + \sigma_{it+1}u_{it+1}$, where $\sigma_{it+1} = c + d\sigma_{it}^2 + fu_{it}^2$. y is log earnings.

Transition probabilities (from/to)		Unemployed	Low earnings		High earnings	
			Low risk	High risk	Low risk	High risk
Unemployed		0.574 (0.004)	0.302 (0.003)	0.107 (0.001)	0.000 (0.000)	0.017 (0.001)
Low earnings	Low risk	0.030 (0.001)	0.443 (0.003)	0.074 (0.001)	0.351 (0.001)	0.103 (0.002)
	High risk	0.038 (0.001)	0.400 (0.003)	0.106 (0.001)	0.313 (0.002)	0.144 (0.002)
High earnings	Low risk	0.005 (0.000)	0.166 (0.000)	0.099 (0.001)	0.628 (0.003)	0.103 (0.002)
	High risk	0.007 (0.000)	0.137 (0.001)	0.137 (0.002)	0.575 (0.003)	0.144 (0.002)
Stationary distribution		0.035 (0.001)	0.263 (0.001)	0.097 (0.001)	0.497 (0.003)	0.109 (0.002)
		<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>f</i>
AR(1)-GARCH(1,1) estimates		2.206 (0.018)	0.702 (0.002)	2.511 (0.008)	0.146 (0.001)	0.114 (0.001)

Note: Discretization is done at the same σ -levels as the earnings process estimated using earnings and savings data that will be presented below, for sake of comparability. The other parameter estimates are: ν_l : 0.6297, ν_h (implied): 0.8868. Standard errors based on 1999 case-bootstrap samples.

have lower probabilities of remaining in the same state next year than individuals in low-risk states. The amount of income inequality is similar, but slightly larger than the estimates found in the previous subsection: low-income earners have a wage 37% below average, while the high-income earners have a wage 89% above average. Note that these values are the same for the different risk states. Also, low-risk states have a smaller probability of becoming unemployed than high-risk states, although this difference is smaller than under the first identification strategy. Overall, the main conclusions are similar.

There are also important differences. The first identification strategy does a better job at matching the probability of remaining unemployed. The probability of remaining unemployed next year given by the non-parametric auxiliary matrix in the Data Section in Equation (1) is around 0.4. The GARCH-restricted estimates imply a probability around 0.57. This is because in the GARCH-restricted estimates, unemployment is interpreted as having a low income level, and the transition probabilities follow from the GARCH model, but these cannot match this feature of the data well.

Another issue with interpreting risk as variance, as the GARCH-restricted model does, is that the results in Table 4 suggest that the low-income-high-risk state has fairly similar upward risk as the low-income-low-risk state, while the results in Table 3 indicate that the low-income-high-risk individuals actually have a considerably lower probability of

moving to the high-income state than low-income-low-risk individuals. Again, this is driven by restricting risk to be a variance: when doing so, a high risk level corresponds to a high variance, and consequently, an increase in both the probabilities of upward and downward changes in income. These issues support the idea of this paper that a rich notion of income risk heterogeneity should be studied and modeled.

A last notable difference between the two identification strategies is the following. Table 4 suggests that the low-income states have considerably larger job-loss probabilities than the high-income states. In the discretized GARCH framework of this paper, unemployment is modeled as having a *very* low income, which is a smaller – and thus more likely – shock when coming from the low-income state than from the high-income state. Therefore, the GARCH restrictions suggest that low-income states are considerably more risky in terms of their probability of moving into unemployment, while this inverse relationship between income and risk is not supported by the unrestricted estimates of the first identification strategy. However, the non-parametric results in Table 1 do support an inverse relationship between earnings level and unemployment probabilities, suggesting that it might not purely be the discretized AR-GARCH restrictions that are driving this result. Instead, it seems plausible that the estimates of the first identification strategy are affected by some of the other parameters in the auxiliary model. For example, assigning higher unemployment probabilities to the high-earnings states may help to generate more wealth inequality, at the cost of a worse fit for some of the other auxiliary parameters.

6.2.3 Estimation results using a combination of identification strategy 1 and 2: net worth, earnings and employment panel, and GARCH restrictions

This section considers the combination of both identification strategies. For these estimates, essentially the transition probability matrix obtained from identification strategy 1 is restricted to be generated by a discretized AR(1)-GARCH(1,1) process. The results of this procedure are summarized in Table 5. The reason to consider these estimates is because they can be used to further analyze the restrictiveness of the GARCH model. As both models minimize the same objective function, a direct comparison gives us a good indication of this restrictiveness. Whereas the GARCH-restricted estimates of Table 5 have an objective function of 0.7642, the unrestricted estimates of Table 3 have an objective function as low as 0.2632, amounting to a deterioration of approximately 200% caused by imposing these restrictions.

A comparison between the results of Table 5 and the unrestricted estimates of Table 3 confirms some of the results above on the restrictiveness of the AR(1)-GARCH(1,1) process. The probability of remaining unemployed in table 5 is 0.6, which is larger than the empirically measured probability of remaining unemployed next year, which is around 0.4 (see the results in Equation (1) in the Data Section), confirming that the discretized GARCH model fails to match this dimension of the data. The estimation results again imply a relationship between income inequality and income risk inequality, as found in both the auxiliary transition matrix of Table 1, and in the estimates of the second identification strategy. Furthermore, the results also confirm the considerable risk inequality faced by the individuals. The low-risk state has a variance of shocks to the log earnings process of 2.079, while the high-risk state has a variance of 4.216, which means the variance of the shocks to log-income doubles.

Table 5: GARCH estimates using an earnings, net worth and employment panel. $y_{it+1} = a + by_{it} + \sigma_{it+1}u_{it+1}$, where $\sigma_{it+1} = c + d\sigma_{it}^2 + fu_{it}^2$. y is log earnings. Process is discretized at the following two variance levels: $c/(1-d)$ and $c + d(c/(1-d) + f((a/(1-b))(lv_l + lv_h))^2)$. lv_l and lv_h are the log-correspondences of the income differences v_l and v_h . The discretized variance levels are: 2.079 and 4.216.

Transition probabilities (from/to)		Unemployed	Low earnings		High earnings	
			Low risk	High risk	Low risk	High risk
Unemployed		0.601	0.299	0.076	0.000	0.025
Low earnings	Low risk	0.024	0.433	0.105	0.311	0.128
	High risk	0.082	0.461	0.000	0.389	0.069
High earnings	Low risk	0.003	0.129	0.126	0.614	0.128
	High risk	0.025	0.255	0.044	0.607	0.069
Stationary distribution		0.046	0.263	0.097	0.484	0.111
		<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>f</i>
AR(1)-GARCH(1,1) estimates		1.850	0.722	2.817	-0.355	0.829

Other parameters: ν_1 : 0.699, ν_2 (implied): 0.498. β : 0.979 and γ (risk aversion) 5.62.

Another example of how the GARCH model is put to the test is the following. The discretized GARCH model has a hard time simultaneously matching the large stability of the low-risk states, and strong reversion to the low-risk states from the high-risk states. The model tries to match these features by choosing a negative parameter d , where d is the autoregressive component in the variance equation, and the f parameter close to one, where f is the relation between next period's variance and the current earnings shock. This combination achieves a large variance after having a sizeable earnings shock, but a

strong reversion towards a low-variance state next period. Note that this is an atypical parametrization for a GARCH model. GARCH models are often used for financial applications, and financial returns are typically characterized by a low sensitivity to shocks (f close to zero) and very high persistence (d close to one). Note, furthermore, that in financial applications, d is typically restricted to be positive, because else, the model might predict negative variances. This is not necessary in this application, because the state space is discretized at positive variance levels.

6.2.4 Conclusions shared by all three estimation procedures

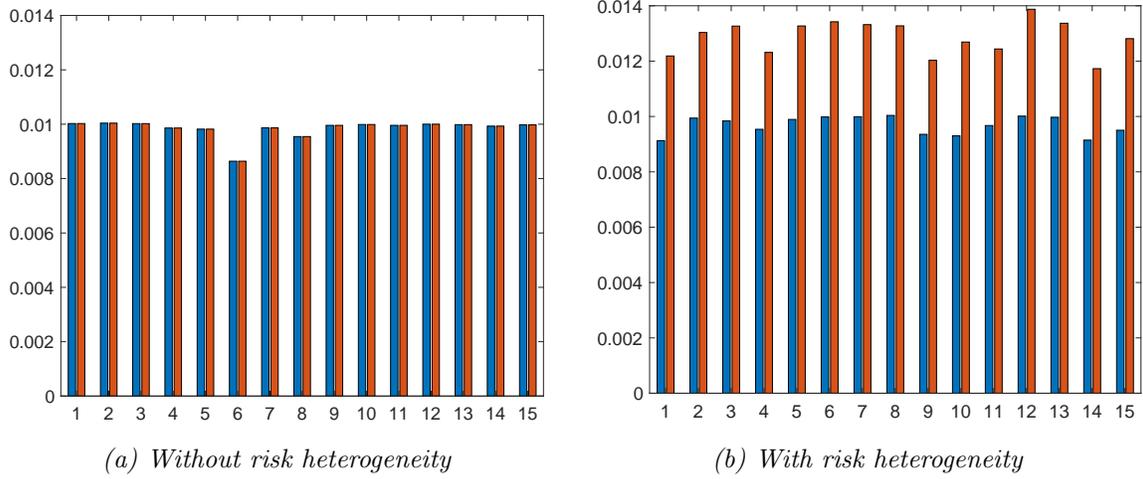
Although the results point at the fact that the GARCH model is overly restrictive, and earnings risk is a richer concept than only the variance of earnings shocks, several results are consistent across all three approaches and are therefore highlighted here. Income inequality is present and sizeable: low-income individuals earn around 30-33% below average wages, while high-income individuals earn 49 to 88% above average. Similarly, income risk inequality is found to be sizeable, and high-risk individuals have probabilities of becoming unemployed four to ten times as large as low-risk individuals. Thirdly, the estimates support that high-risk states are typically temporary, and the probability of remaining in a high-risk state next period is fairly small. On the other hand, the low-risk states are stable, and the probability of remaining in a low-risk state next period is large, with estimates varying from 0.4 to 0.8.

6.3 Comparing an economy with and without heterogeneous earnings risk

In this subsection, I evaluate the implications of heterogeneous time-varying earnings risk in the estimated general equilibrium model. I compare the model implications with those of a general equilibrium with a reduced Markov chain that aggregates the unobservable earnings risk, such that individuals only differ with respect to earnings levels. These individuals thus still experience earnings risk, but there is no earnings risk inequality. This analysis is based on the estimates of Table 3.

I first confirm that the model with risk heterogeneity can generate dispersion in saving behavior, whereas this is not the case for the model without risk heterogeneity. Individuals in high-risk states save more than individuals in low-risk states, conditional on their earnings and previous savings. This is shown in Figure 11. The difference in saving choices between high- and low-risk individuals is larger at lower savings levels.

Figure 11: Variance of savings between individuals in 15 (out of 250 in total) clusters in period t (blue) and period $t+1$ (red) in the model with and without risk heterogeneity, clustered based on their earnings and period- t savings.



In the model without risk heterogeneity, the same transition probability matrix is used as in the model with risk heterogeneity, but states are aggregated such that individuals are only heterogeneous with respect to earnings levels and not transition probabilities. Based on simulations, using the estimates of the unrestricted model in identification strategy 1.

The model with earnings risk heterogeneity implies a savings distribution with a larger variance, and attributes more probability to the right tail of the distribution than the model without earnings risk heterogeneity. Also, aggregate savings go up by approximately 5%. This, and other properties of the savings distribution are summarized in Table 6, and visualized in Figure 12. The differences between the savings distributions of the models with and without earnings risk heterogeneity are mostly driven by the low-wage individuals, for whom the pre-cautionary motive is strong. This can be seen from the left bottom panel in Figure 12. The savings distribution of the low-income-high-risk individuals lies almost strictly to the right of the distribution of the low-income-low-risk individuals. The opposite happens for high-income individuals; the savings distribution of high-risk individuals lies to the left of the savings distribution of low-risk individuals. This can likely be explained as follows. Even though at a given savings level at time t , the high-risk individuals will save more than high-income-low-risk individuals, on average these high-risk individuals will have accumulated lower savings levels, such that the overall distribution of low and high risk individuals with high wages is fairly similar.

Figure 12: Savings distributions for the model with heterogeneous risk (abbreviated as *het. risk*) and without earnings risk heterogeneity (abbreviated as *hom. risk*). In the model without risk heterogeneity, the same transition probability matrix is used as in the model with risk heterogeneity, but states are aggregated such that individuals are only heterogeneous with respect to earnings levels and not transition probabilities. ‘All’ refers to the savings distribution of all individuals jointly. ‘Low wage’ refers to individuals in the low income state, while ‘high wage’ refers to the high income individuals.

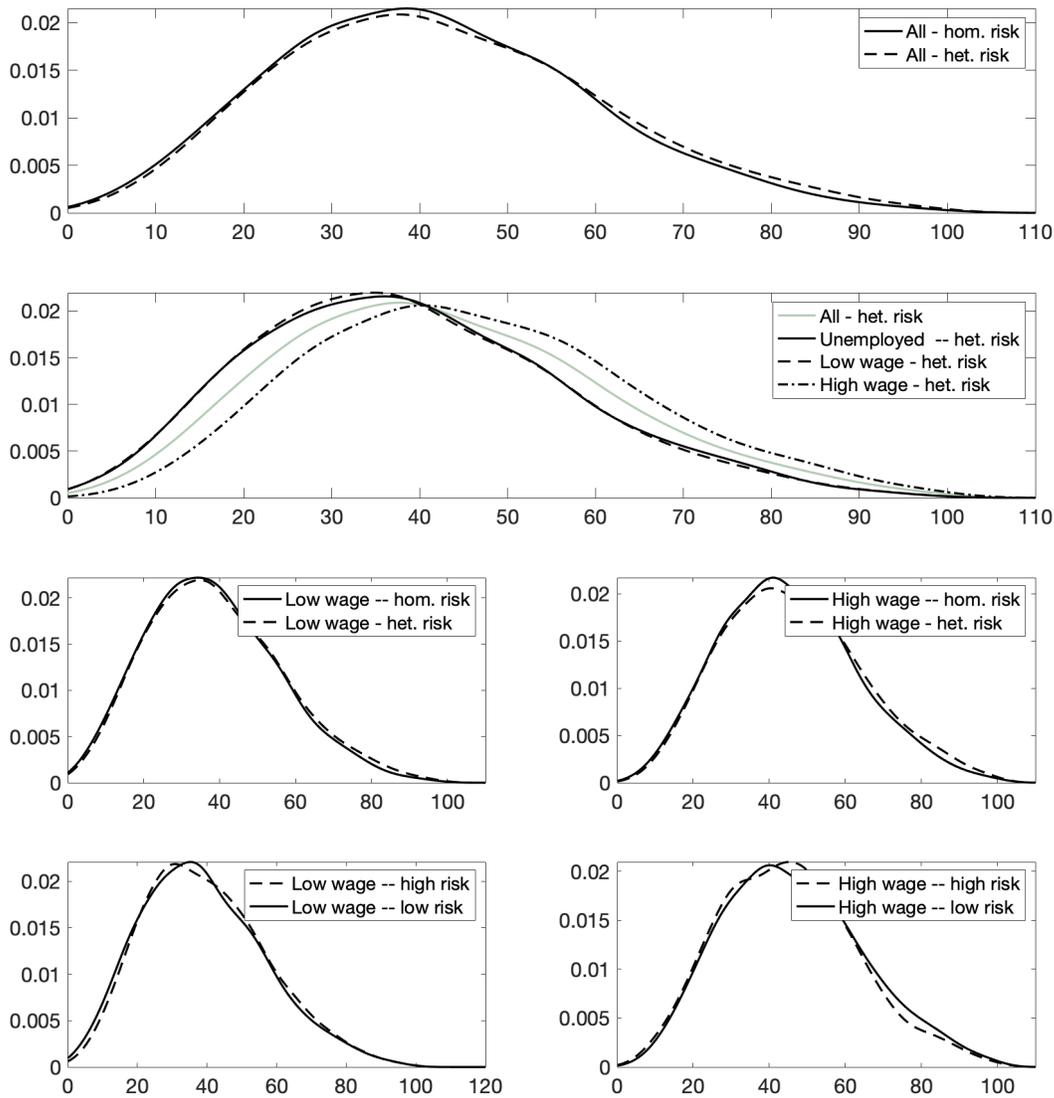


Table 6: Properties of the savings distributions for the model with heterogeneous risk and without earnings risk heterogeneity. In the model without risk heterogeneity, the same transition probability matrix is used as in the model with risk heterogeneity, but states are aggregated such that individuals are only heterogeneous with respect to earnings levels and not transition probabilities.

Model/Statistic	Homogeneous risk	Heterogeneous risk
Mean(k_{it})	40.076	41.041
$\hat{\beta}$ (Implied time discount factor)	0.9637	0.9643
Variance(k_{it})	356.07	372.50
Skewness(k_{it})	0.4296	0.4301
Kurtosis(k_{it})	2.808	2.754
Percentile ₉₅ (k_{it})/Percentile ₉₀ (k_{it})	1.1216	1.1215
Percentile ₉₅ (k_{it})/Percentile ₇₅ (k_{it})	1.4094	1.4124
Percentile ₇₅ (k_{it})/Percentile ₅₀ (k_{it})	1.3636	1.3761
Percentile ₅₀ (k_{it})/Percentile ₂₀ (k_{it})	1.6552	1.6525
Percentile ₂₀ (k_{it})/Percentile ₁ (k_{it})	4.4615	4.2908

To assess the welfare implications of earnings risk heterogeneity, one can compute the percentage increase in consumption, denoted by $g(k, \epsilon)$, needed to make an individual with current savings k and earnings state ϵ indifferent between an economy with heterogeneous and homogeneous earnings risk. Denote the value function of an individual in an economy with heterogeneous earnings risk by $V^{\text{het. risk}}(k, \epsilon)$, and the value function of an individual in an economy with homogeneous earnings risk by $V^{\text{hom. risk}}(k, \epsilon^{\text{hom}})$. Note that ϵ^{hom} aggregates states 1 and 2, and states 3 and 4. Comparing the two economies in equilibrium, g is given by (Krueger and Perri, 2004):

$$g(k, \epsilon) = 100 \cdot \left(\left[\frac{V^{\text{hom. risk}}(k, \epsilon)}{V^{\text{het. risk}}(k, \epsilon)} \right]^{\frac{1}{1-\gamma}} - 1 \right),$$

where, recall, relative risk aversion of the worker is denoted by γ . Conditional on the joint probability measure of the individual capital choices k and states ϵ , denoted $\Phi_{k, \epsilon}^{\text{het. risk}}$, an aggregate measure for welfare is the welfare gain comparing the steady states of both economies, given by

$$g^{ss} = \left(\frac{\int_{k, y} V^{\text{hom. risk}}(k, y) d\Phi_{k, y}^{\text{hom. risk}}}{\int_{k, y, \xi} V^{\text{het. risk}}(k, y, \xi) d\Phi_{k, y, \xi}^{\text{het. risk}}} \right)^{\frac{1}{1-\gamma}} - 1 \quad (25)$$

Computing Equation (25) for the estimates in Table 3 and the entire distribution implies a welfare gain of 0.51%. This means that on average individuals are better off when they all face the earnings process with homogeneous risk.

7 Conclusion

This paper proposes a novel earnings process that models heterogeneity in earnings risk by allowing for heterogeneous transition probabilities between high and low earning levels. This process generates a richer notion of earnings risk heterogeneity than previously considered in the literature. The paper specifies and analyzes a general equilibrium heterogeneous agent model accommodating this heterogeneous idiosyncratic earnings process for household earnings. Theoretical identification results are derived on the parameters of this earnings process, and these results show that a panel with savings levels can be used to identify and estimate the parameters of this process. Alternatively, a more restrictive variant of the idiosyncratic shock process proposed by this paper can be mapped into a discretized (generalized) autoregressive conditional heteroskedasticity model, such that these parameters can be estimated using only a panel of earnings levels. To do so, this paper proposes an improved discretization method for GARCH-type processes. These gives rise to two identification strategies, and these are used to estimate the relevant parameters of the earnings process using indirect inference and data from the Survey of Income and Program Participation.

The results show that earnings risk heterogeneity is empirically relevant, and individuals face large differences in their earnings risk, as evidenced, for example, by the differences in the probabilities of becoming unemployed. Heterogeneous earnings risk also has implications in general equilibrium, for example, the precautionary savings of the low-income-high-risk individuals moves the savings distribution to the right. Comparing the two different identification strategies shows that a (discretized) GARCH model is overly restrictive when modelling earnings risk. This supports the notion that earnings risk is a rich concept, that should be modelled more flexibly, for example, by using the earnings process proposed in this paper. Earnings risk heterogeneity also has welfare effects, and individuals would on average be willing to give up 0.51% of their consumption to live in an economy without earnings risk heterogeneity.

For further research, it would be interesting to analyze the heterogeneous earnings risk model in a setting with aggregate risk, to assess whether heterogeneous earnings risk has other important implications. Such a model can then also be used to analyze whether and how earnings risk inequality varies over the business cycles, and to study the implications of this potentially important source of business cycle variation for tax and monetary policy.

References

- Aiyagari, S. R. (1994). “Uninsured Idiosyncratic Risk and Aggregate Saving”. *The Quarterly Journal of Economics*, 109(3), 659–684.
- Alan, S., Browning, M., and Ejrnæs, M. (2018). “Income and Consumption: A Micro Semistructural Analysis with Pervasive Heterogeneity”. *Journal of Political Economy*, 126(5), 1827–1864.
- Almuzara, M. (2021). “*Heterogeneity in Transitory Income Risk*”. Retrieved from https://martinalmuzara.com/files/transitory_risk.pdf.
- Arellano, M., Blundell, R., and Bonhomme, S. (2017). “Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework”. *Econometrica: Journal of the Econometric Society*, 85(3), 693–734.
- Bollerslev, T. (1986). “Generalized Autoregressive Conditional Heteroskedasticity”. *Journal of Econometrics*, 31(3), 307–327.
- Bonhomme, S., Jochmans, K., and Robin, J.-M. (2016). “Estimating Multivariate Latent-Structure Models”. *The Annals of Statistics*, 44(2), 540–563.
- Bonhomme, S., Jochmans, K., and Robin, J.-M. (2017). “Nonparametric Estimation of Non-Exchangeable Latent-Variable Models”. *Journal of Econometrics*, 201(2), 237–248.
- Bonhomme, S., Lamadon, T., and Manresa, E. (2021). “*Discretizing Unobserved Heterogeneity*”. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3333452.
- Botosaru, I. (2020). “*A Panel Model with Heteroskedasticity, with an Application to Earnings Risk*”. Retrieved from <https://ideas.repec.org/p/sfu/sfudps/dp17-11.html>.
- Botosaru, I., and Sasaki, Y. (2018). “Nonparametric Heteroskedasticity in Persistent Panel Processes: An Application to Earnings Dynamics”. *Journal of Econometrics*, 203(2), 283–296.
- Busch, C., Domeij, D., Guvenen, F., and Madera, R. (2021). “Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance”. *American Economic Journal: Macroeconomics*.
- Calvet, L. E., Campbell, J. Y., Gomes, F., and Sodini, P. (2021). *The cross-section of household preferences* (Working Paper No. 28788). National Bureau of Economic Research. Retrieved from <http://www.nber.org/papers/w28788>.
- Castaneda, A., Diaz-Gimenez, J., and Rios-Rull, J.-V. (2003). “Accounting for the US Earnings and Wealth Inequality”. *Journal of Political Economy*, 111(4), 818–857.

- De Nardi, M., Fella, G., and Paz-Pardo, G. (2020). “Nonlinear Household Earnings Dynamics, Self-insurance, and Welfare”. *Journal of the European Economic Association*, 18(2), 890–926.
- Den Haan, W. J. (2010). “Comparison of Solutions to the Incomplete Markets Model with Aggregate Uncertainty”. *Journal of Economic Dynamics and Control*, 34(1), 4–27.
- Druedahl, J., and Munk-Nielsen, A. (2020). “Higher-order Income Dynamics with Linked Regression Trees”. *The Econometrics Journal*, 23(3), S25–S58.
- Duan, J.-C., and Simonato, J.-G. (2001). “American Option Pricing under GARCH by a Markov Chain Approximation”. *Journal of Economic Dynamics and Control*, 25(11), 1689–1718.
- Elkhishin, A., West, T., and MacDonald, K. (2019). “Understanding the Retirement Savings of Self-Employed Tradespeople in Australia”. *Australian Journal of Labour Economics*, 22(2), 75–92.
- Engle, R. F. (1982). “Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation”. *Econometrica: Journal of the Econometric Society*, 987–1007.
- Gourieroux, C., Monfort, A., and Renault, E. (1993). “Indirect Inference”. *Journal of Applied Econometrics*, 8(S1), S85–S118.
- Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). “What Do Data on Millions of US Workers Reveal about Life-Cycle Earnings Risk?”. *Econometrica: Journal of the Econometric Society* (forthcoming).
- Guvenen, F., and Smith, A. A. (2014). “Inferring Labor Income Risk and Partial Insurance from Economic Choices”. *Econometrica: Journal of the Econometric Society*, 82(6), 2085–2129.
- Imbens, G. W., and Angrist, J. D. (1994). “Identification and Estimation of Local Average Treatment Effects”. *Econometrica: Journal of the Econometric Society*, 62(2), 467–475.
- Iskrev, N. (2010). “Local Identification in DSGE Models”. *Journal of Monetary Economics*, 57(2), 189–202.
- Janssens, E. (2019). “*Identification in Heterogeneous Agent Models*”. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3463736.
- Jensen, S. T., and Shore, S. H. (2011). “Semiparametric Bayesian Modeling of Income Volatility Heterogeneity”. *Journal of the American Statistical Association*, 106(496), 1280–1290.
- Krueger, D., and Perri, F. (2004). “On the Welfare Consequences of the Increase

- in Inequality in the United States”. In M. Gertler and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2003, Volume 18* (p. 83-138). Cambridge, Massachusetts: The MIT Press.
- Krusell, P., Mukoyama, T., Şahin, A., and Smith Jr., A. A. (2009). “Revisiting the Welfare Effects of Eliminating Business Cycles”. *Review of Economic Dynamics*, 12(3), 393–404.
- Krusell, P., and Smith, A. A. (1997). “Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset returns”. *Macroeconomic Dynamics*, 1(2), 387–422.
- Meghir, C., and Pistaferri, L. (2004). “Income Variance Dynamics and Heterogeneity”. *Econometrica: Journal of the Econometric Society*, 72(1), 1–32.
- Moore, J. C. (2008). “Seam Bias in the 2004 SIPP Panel: Much Improved, but Much Bias Still Remains”. *US Census Bureau Statistical Research Division Survey Methodology Research Report Series*, 3, 2008.
- Mueller, A. I., Spinnewijn, J., and Topa, G. (2021). “Job Seekers’ Perceptions and Employment Prospects: Heterogeneity, Duration Dependence, and Bias”. *American Economic Review*, 111(1), 324–63.
- Rouwenhorst, K. (1995). *Frontiers of Business Cycle Research*. In T. Cooley (Ed.), *Frontiers of Business Cycle Research* (p. 294-330). Princeton, NJ: Princeton University Press.
- Smith, Jr., A. A. (1993). “Estimating Nonlinear Time-series Models using Simulated Vector Autoregressions”. *Journal of Applied Econometrics*, 8(S1), S63–S84.
- Stokey, N. L., Lucas Jr., R. E., and Prescott, E. C. (1989). *Recursive Methods in Economic Dynamics*. Cambridge, Massachusetts: Harvard University Press.
- Stoltenberg, C. A., and Singh, S. (2020). “Consumption Insurance with Advance Information”. *Quantitative Economics*, 11(2), 671–711.
- Tauchen, G. (1986). “Finite State Markov-chain Approximations to Univariate and Vector Autoregressions”. *Economics Letters*, 20(2), 177–181.
- Torczon, V. (1997). “On the Convergence of Pattern Search Algorithms”. *SIAM Journal on Optimization*, 7(1), 1–25.
- Zivot, E. (2009). “Practical Issues in the Analysis of Univariate GARCH Models”. In *Handbook of Financial Time Series* (pp. 113–155). Berlin, Heidelberg: Springer.

A Data

A.1 Individual-level variables

The individual-level variables used in this paper are obtained from the Survey of Income and Program Participation (SIPP) dataset of the United States Census Bureau. The analysis in Section 2.2 uses the panel of 2014. A detailed description of the variables employed from the 2014 panel is now given, together with the sample transformations used.

First of all, it should be noted that only person-level variables are used. I exclude individuals that are self-employed (using the variable `SELFEMPL`), and remove individuals younger than 18 and older than 67 (using `TAGE`) at the start of the panel. Furthermore, I remove the subset of individuals that are not in the labor force, for example because of a disability or health condition, study, care-taking, pregnancy or retirement (indicated by the variable `EN_NOWRKi`).

Variables used in the analysis are: person-level net worth (`TNETWORTH`), sum of earnings of all jobs (sum of `TJBi_MSUM` for $i = 1, \dots, 7$), and a code for the monthly employment status (`RMESR`). It is also possible for jobs to be linked in the panel, such that the employment status of an individual at one specific job can be assessed, but in this paper the focus is on transitions between un- and employment, not in between jobs.

Apart from these variables, various controls are used to ensure that the measured effect is due to the mechanism of interest: heterogeneity of earnings risk, rather than other variables that could explain heterogeneity in saving decisions. The following variables are used: `TAGE` (age of the individual), `ESEX` (indicating gender), `EEDUC` (highest level of school completed or highest degree received by end of reference year), `ERACE` (what race does the respondent consider him/herself to be) or the more detailed variable `TRACE` (these demographic variables can be extended with variable `EORIGIN` indicating whether the respondent is Spanish, Latino or Hispanic), `TEHC_ST` (state of residence), `TEHC_METRO` (living in a metropolitan area or not).

Some other variables that could be important for the perception of risk an individual faces that is unrelated to earnings risk is whether the individual has a health insurance or not (`RHLTHMTH`), whether the parent takes care of children under 18 or not (`ERP`), and coverage of housing (`ETENURE`, i.e. is the house rented, owned or occupied without paying rent).

Table A1: Selected summary statistics of the Survey of Income and Program Participation dataset for 2014-2018 waves after sample selection. Selected standard errors given in parentheses.

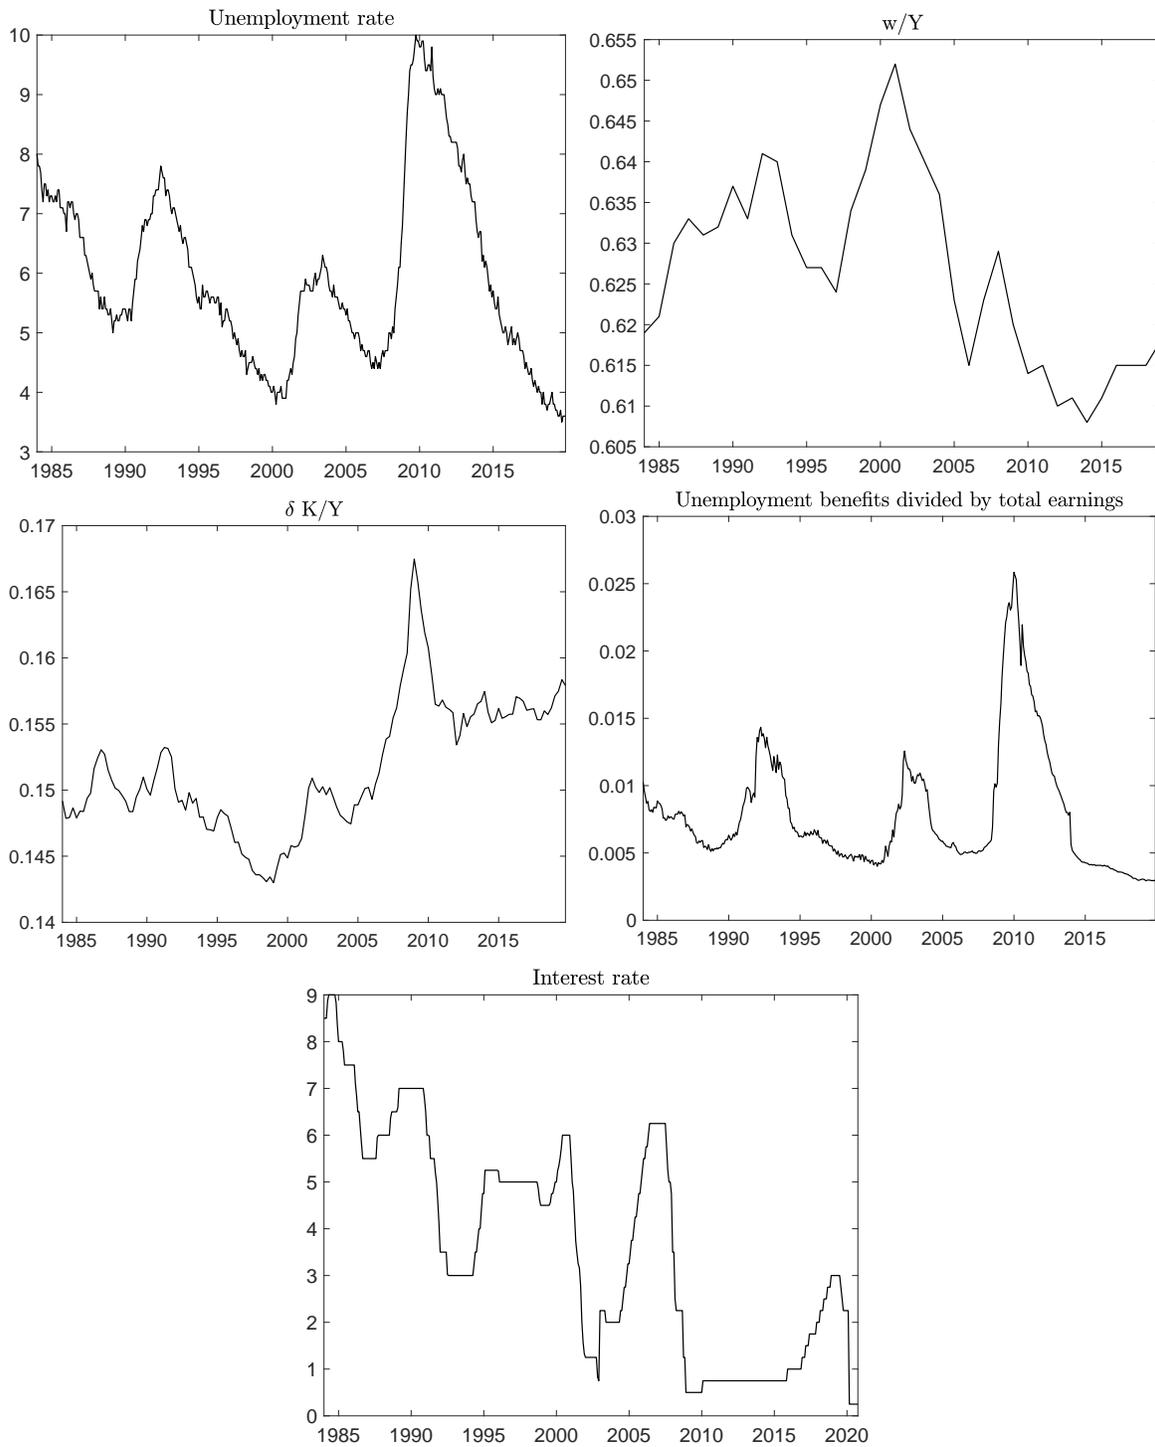
	Sample average	Average for males	Average for females
Age (years)	42	41.44	42.59
Female	38%	–	–
High school	30.94%	29.03%	34.01%
College	23.70%	18.24%	32.36%
University BSc or up	15.61%	12.10%	21.25%
Identifies as black	10.97%	0.08%	15.83%
Not white, not black	6.16%	5.46%	7.29%
Big city	61.34%	51.26%	77.50%
Health insurance	63.46%	51.65%	82.41%
Renter	25.88%	21.66%	32.64%
House owner	51.76%	43.14%	65.61%
Kids	16.42%	2.71%	38.41%
Average monthly earnings	3525.6 (3879.1)	3546.8 (3860.7)	3491.2 (3908.7)
Net worth	155940 (981400)	164190 (1226400)	142880 (325530)
Unemployed (>6 months)	5.95%	5.76%	6.25%
Total number of observations	25124	15479	9645

A.2 Aggregate variables

Table A2: Data for the United States, all data except income tax revenues are retrieved from FRED, Federal Reserve Bank of St. Louis, January 10, 2021 and can be accessed using the codes given in parentheses via <https://fred.stlouisfed.org/>.

Data for estimation parameters α, δ, ν		
Variable	Data proxy/description	Source
Y	Gross national income (A023RC1Q027SBEA)	U.S. Bureau of Economic Analysis
	The sum of:	
	Shares of gross domestic income: Compensation of employees, paid (A4002E1A156NBEA)	
w/Y	Shares of gross domestic income: Net operating surplus: Private enterprises: Business current transfer payments (net) (B029RE1A156NBEA)	U.S. Bureau of Economic Analysis
	Shares of gross domestic income: Net operating surplus: Private enterprises: Proprietors' income with inventory valuation and capital consumption adjustments (A041RE1A156NBEA)	
δK	Consumption of Fixed Capital (COFC)	U.S. Bureau of Economic Analysis
r	Interest Rates, Discount Rate for United States (percent per annum) (INTDSRUSM193N)	International Monetary Fund
u	Unemployment Rate (UNRATE)	U.S. Bureau of Labor Statistics
	Personal current transfer receipts: Government social benefits to persons: Unemployment insurance (W825RC1)	
τ	Divided by: Compensation of Employees, Received: Wage and Salary Disbursements (A576RC1)	U.S. Bureau of Economic Analysis

Figure A1: Aggregate data used for estimation of the first-step parameters



Data sources are described in more detail in Appendix A.2.

B Additional details on the relationship to other earnings processes

B.1 Discretizing a canonical earnings model with a GARCH component

It is common in the literature on earnings processes to assume that (log) earnings Y_{it} are composed of both a transitory and a permanent component, as opposed to only the permanent component that the AR(1) process in Section 3.1.1 of the main paper assumes. The Markov transition matrix can also be mapped into such an earnings process of the form

$$Y_{it} = \eta_{it} + \nu_{it}, \quad (\text{B.1})$$

where η_{it} follows the AR(1)-GARCH(1,1) process of Equation (4) in the main paper, and ν_{it} is a white-noise transitory shock. This model is often referred to as the canonical model of earnings.²² The additivity of the transitory-permanent income process allows using the discretization technique proposed in Section 3.1.1 for η_{it} , and in a second step computing the final discretized Markov process for Y_{it} .

To make this more concrete, assume that the permanent component η_{it} is discretized in N levels $\bar{\eta}(i)$, $i = 1, \dots, N$, and $\sigma_{i,t+1}^2$, the conditional variance of the shocks to η , is discretized into M levels $\bar{\sigma}^2(j)$, $j = 1, \dots, M$. Last, assume that (log) earnings Y_{it} is discretized into L levels, $\bar{y}(k)$, $k = 1, \dots, L$. Retain the definitions for cells H and G introduced in Section 3.1.1 of the main paper, where H is now used to denote the cells of the discretized permanent component, instead of the income level. For the income levels, introduce cells E corresponding to the L income levels $\bar{y}(k)$: $E(k) = [e(k), e(k+1)]$ for $k = 1, \dots, L$, with $e(1) = -\infty$, $e(k) = (\bar{y}(k) + \bar{y}(k-1))/2$, for $k = 2, \dots, L$ and $e(L+1) = +\infty$. Denote the transition probability

$$P(Y_{i,t+1} \in E(l), \eta_{i,t+1} \in H(p), \sigma_{i,t+2}^2 \in G(q) | Y_{i,t} \in E(i), \eta_{i,t} \in H(j), \sigma_{i,t+1}^2 \in G(k))$$

with short-hand notation $\pi((i, j, k); (l, p, q))$. As before, we approximate this probability by $P(Y_{i,t+1} \in E(l), \eta_{i,t+1} \in H(p), \sigma_{i,t+2}^2 \in G(q) | Y_{i,t} = \bar{y}(i), \eta_{i,t} = \bar{\eta}(i), \bar{\sigma}_i^2 = \bar{\sigma}^2(k))$. Next, this probability can be approximated by:

$$p_{((i,j,k);(l,p,q))}^{CG} = \int_{B_{\text{lower}}(p,l)}^{B_{\text{upper}}(p,l)} p_{(p,q;j,k)}^G p(\nu) d\nu, \quad (\text{B.2})$$

²²Alternatively, it is often assumed η_{it} is a unit-root process, but this case is not considered here.

where $p(\nu)$ is the probability density function of the transitory shock ν_{it} , $B_{\text{lower}}(p, l)$ is defined as $e(l) - \bar{\eta}(p)$ and $B_{\text{upper}}(p, l) = e(l + 1) - \bar{\eta}(p)$, that is, the highest and lowest values of the transitory component ν that can correspond with the intervals of the income levels and the permanent component, where the permanent component η is fixed at the value middle of its interval. The definition of $p_{(p,q);(j,k)}^G$ is given in Equation (9). CG is short for canonical-GARCH.

Alternatively, an exact expression for the probability can also be derived, but this is computationally more intensive.

$$p_{((i,j,k);(l,p,q))}^{CG} = \int_{L_{ij}(k)}^{L_{ij}(k+1)} \int_{e(l)-\eta}^{e(l+1)-\eta} p(\nu) d\nu I\{(c + d \cdot \bar{\sigma}^2(j) + f \cdot u^2) \in G(l)\} \phi\left(\frac{\eta - a - b \cdot \bar{\eta}(i)}{\bar{\sigma}(j)}\right) d\eta. \quad (\text{B.3})$$

To compute the transition matrix of interest, that is, transition probabilities $P(Y_{i,t+1} \in E(l), \sigma_{i,t+2}^2 \in G(n) | Y_{i,t} \in E(i), \sigma_{i,t+1}^2 \in G(k))$, one needs to aggregate over the states $\bar{\eta}$:

$$P(Y_{i,t+1} \in E(l), \sigma_{i,t+2}^2 \in G(q) | Y_{i,t} \in E(i), \sigma_{i,t+1}^2 \in G(k)) = \sum_{p=1}^N \sum_{j=1}^N \pi(l, p, q; i, j, k) \frac{P(Y_{i,t+1} \in E(i), \eta_{i,t} \in H(j), \sigma_{i,t+1}^2 \in G(k))}{\sum_{h=1}^m P(Y_{i,t+1} \in E(i), \eta_{i,t} \in H(h), \sigma_{i,t+1}^2 \in G(k))}, \quad (\text{B.4})$$

where the expression in the fraction can be computed directly from the stationary distribution implied by the transition matrix in Equation (B.2). The attractiveness in allowing for a transitory component ν_{it} is that this component can affect earnings temporarily, without any lasting effects on future earnings or on future earnings risk.

B.2 Discretization of a TGARCH model

Unlike the traditional GARCH process, a TGARCH allows for a leverage effect, where negative shocks have a different effect on the variance than positive shocks. This is achieved in a TGARCH(1,1) model by assuming the following process for the variance:

$$\sigma_{i,t+1}^2 = c + d \cdot \sigma_{it}^2 + (f + gI\{u_{it} < 0\}) \cdot u_{it}^2. \quad (\text{B.5})$$

The effect of positive and negative shocks on the variance is asymmetric if $g \neq 0$. If $g > 0$, there is a leverage effect. In financial applications, the typical finding is that g is indeed positive and significant. However, for earnings, this might not necessarily

be the case. In fact, a large negative shock, such as a fall into unemployment does not make a large consecutive (positive or negative) shock more likely. It seems more likely that unemployment will be followed by either another period of unemployment, or a lower-wage job. $g < 0$ is thus also a possible outcome in earnings processes. Adjusting the discretization method to accommodate this asymmetric effect only requires replacing Equation (9) by (B.6) below.

$$\begin{aligned} p_{(i,j;k,l)}^{TG} &= P(y_{i,t+1} \in H(k), \sigma_{i,t+2}^2 \in G(l) | y_{it} = \bar{y}(i), \sigma_{i,t+1} = \bar{\sigma}^2(j)) \\ &= \int_{L_{ij}(k)}^{L_{ij}(k+1)} I\{(c + d \cdot \bar{\sigma}^2(j) + (f + g \cdot I\{\varepsilon < 0\})\varepsilon^2) \in G(l)\} \phi(u) du, \end{aligned} \quad (\text{B.6})$$

with $L_{ij}(k)$, \bar{y} , $\bar{\sigma}^2$, $H(k)$ and $G(l)$ defined above, and $\phi(u)$ the probability density function of the normal distribution. TG is short for TGARCH.

B.3 Discretization of a canonical earnings model with heterogeneous (non-time varying) variances

The general income process introduced in Section 3.1 can also be related to earnings processes with ex-ante non-timevarying heterogeneous variances, as for example studied by Almuzara (2021); Botosaru (2020); Botosaru and Sasaki (2018). Consistent with the previous introduced notation, consider the following process for log earnings:

$$Y_{it} = \eta_{it} + \nu_{it}, \quad \nu_{it} \sim (0, (\sigma_i^\nu)^2) \quad (\text{B.7})$$

$$\eta_{it+1} = a + b\eta_{it} + \sigma_i^u u_{i,t+1}, \quad |b| < 1, \quad u_{it} \sim (0, 1) \quad (\text{B.8})$$

where again η_{it} is the permanent component of earnings and ν_{it} the transitory component. Consider a discretization of M permanent component variance levels ($\bar{\sigma}^u(i)$, $i = 1, \dots, M$), and N transitory component variance levels ($\bar{\sigma}^\nu(j)$, $j = 1, \dots, N$). Also consider L discrete income levels ($\bar{y}(l)$, $l = 1, \dots, L$). The discretized canonical earnings process with heterogeneous non-time-varying variances of Equations (B.7)-(B.8) can be mapped into the transition probabilities of the general process in Equation (2) by defining states as a combination of $(\bar{y}(l), (\bar{\sigma}^u(i), \bar{\sigma}^\nu(j))) = (\bar{y}(l), \bar{\xi}(k))$. In this case, \mathcal{X} from the general model in Equation (2) corresponds to $\{\bar{\sigma}^u(i), i = 1, \dots, M\} \times \{\bar{\sigma}^\nu(j), j = 1, \dots, N\}$. Due to the non-time-varying nature of the variances, transitions between different realizations of $\bar{\xi}(i) \in \mathcal{X}$ are ruled out, and thus have probability zero.

To compute the non-zero transition probabilities, a similar strategy as in the previous subsection can be used. Assume that also the permanent component is discretized in K

levels $\bar{\eta}(k), k = 1, \dots, K$. For the income levels y , introduce again cells E corresponding to the L income levels $\bar{y}(l) : l = 1, \dots, L$ as $E(l) = [e(l), e(l+1)]$ for $l = 1, \dots, L$, with $e(1) = -\infty$, $e(l) = (\bar{y}(l) + \bar{y}(l+1))/2$ for $l = 2, \dots, L$ and $e(L+1) = +\infty$. Denote similarly by H the cells of the discretized permanent component $\bar{\eta}(k)$. Denote the transition probability

$$P(Y_{i,t+1} \in E(m), \eta_{i,t+1} \in H(n), \sigma_i^u = \bar{\sigma}^u(p), \sigma_i^\nu = \bar{\sigma}^\nu(q) | \dots \\ Y_{i,t} \in E(i), \eta_{i,t} \in H(j), \sigma_i^u = \bar{\sigma}^u(k), \sigma_i^\nu = \bar{\sigma}^\nu(l))$$

with short-hand notation $p_{(i,j,k,l),(m,n,p,q)}$. As before, approximate this probability by:

$$P(Y_{i,t+1} \in E(m), \eta_{i,t+1} \in H(n), \sigma_i^u = \bar{\sigma}^u(p), \sigma_i^\nu = \bar{\sigma}^\nu(q) | Y_{i,t} = \bar{y}(i), \eta_{i,t} = \bar{\eta}(j), \sigma_i^u = \bar{\sigma}^u(k), \sigma_i^\nu = \bar{\sigma}^\nu(l)).$$

Next, this probability can be computed as

$$p_{(i,j,k,l),(m,n,p,q)} = \begin{cases} 0 & \text{if } k \neq p, l \neq q \\ \int_{L_j(n+1)}^{L_j(n)} \int_{e(m+1)-\eta}^{e(m)-\eta} p(\nu, \sigma^\nu) d\nu \phi\left(\frac{\eta - a - b\bar{\eta}(i)}{\sigma^u}\right) d\eta & \text{else} \end{cases}, \quad (\text{B.9})$$

where $L_j(n)$ is given by $\frac{h(n) - a - b\bar{\eta}(j)}{\sigma^u}$. ϕ denotes the probability density function of u_{it} . At last, aggregate the transition probabilities similar as in Equation (B.4) to obtain the transition probability matrix for the states $(\bar{y}(l), (\bar{\sigma}^u(i), \bar{\sigma}^\nu(j)))$.

Evidently, this Markov chain is not ergodic with respect to the variance states, as there is no randomness in those states. This means that it is not possible to transition from each state to every other state in the state space. Yet, the process can still be stable with respect to the earnings level states, because all earnings levels should still be reachable from each other earnings level state within a finite number of transitions. The variance levels of each individual should in this case rather be considered as ex-ante heterogeneous, drawn from a set of possible discrete realizations.

B.4 Heterogeneous job-loss and job-finding probabilities

The earnings process of Section 3.1 in the main paper can easily accommodate heterogeneous job-loss and job-finding probabilities as in Guvenen et al. (2021) and Mueller et al. (2021). To achieve this, when defining \mathcal{Y} , the set of discretized earnings levels, one should also include an unemployed state u , that is, $\mathcal{Y} = \{u, \bar{y}(1), \dots, \bar{y}(L)\}$. Next, for any choice for the set \mathcal{X} with cardinality of two or higher, the earnings process generates het-

erogeneous job-finding probabilities for different unemployed individuals. Furthermore, individuals with the same current earnings level have different probabilities of losing their job if they have a different realization of ξ_t .

B.5 Skewness and higher-order moments

The general earnings process proposed by this paper can generate skewness and other non-normalities, as shown to be important by Busch et al. (2021) and others. Furthermore, the amount of skewness in the earnings shock distribution can be heterogeneous and time-varying.

Denote the current state by $(\bar{y}(i), \bar{\xi}(j))$. In that case, the probability distribution of next period's earnings is characterized by the probabilities $p_{(i,j);(k,l)}$ for $k = 1, \dots, L$ and $l = 1, \dots, M$ corresponding to the outcomes $(y_{t+1}, \xi_{t+1}) \in (\mathcal{Y} \times \mathcal{X})$. The expectation of next period's earnings level conditional on the current state can be computed as:

$$\mu_{(i,j)} = \mathbb{E}\left[Y_{t+1} \mid (Y_t, \xi_t) = (\bar{y}(i), \bar{\xi}(j))\right] = \sum_{k=1}^L \sum_{l=1}^M \bar{y}(k) p_{(i,j);(k,l)}. \quad (\text{B.10})$$

The skewness, conditional on current state (i, j) can be computed as:

$$\gamma_{ij} = \frac{\sum_{k=1}^L \sum_{l=1}^M (\bar{y}(k) - \mu_{(i,j)})^3 p_{(i,j);(k,l)}}{\left(\sum_{k=1}^L \sum_{l=1}^M (\bar{y}(k) - \mu_{(i,j)})^2 p_{(i,j);(k,l)}\right)^{\frac{3}{2}}}, \quad (\text{B.11})$$

with $\mu_{(i,j)}$ given in Equation (B.10). If the next period's earnings level distribution is asymmetric, the conditional skewness will always be non-zero. Furthermore, individuals with different realizations of ξ_t can have different values of skewness, generating time-varying heterogeneous skewness. The same holds for other higher-order moments.

B.6 Heterogeneous persistence

In the approach of Arellano et al. (2017), the persistence of earnings shocks is time-varying and heterogeneous, as it depends on the sign and size of the previous earnings shock. This statistical property can also be achieved in the earnings process I propose, as illustrated through Example 2.

Example 2. Assume $\mathcal{Y} \in \{\bar{y}(1), \bar{y}(2), \bar{y}(3)\}$, where $|\bar{y}(1) - \bar{y}(2)| < b$, with b some scalar, $|\bar{y}(2) - \bar{y}(3)| < b$, but $|\bar{y}(1) - \bar{y}(3)| > b$. Also, $\bar{y}(1) > \bar{y}(2) > \bar{y}(3)$. Next, restrict ξ as representing the sign and size of the shock, where the shock is here for simplicity

defined as the difference between $y_{i,t+1} - y_{it}$. That is, take

$$\xi_{t+1} = \begin{cases} \bar{\xi}(1) & \text{if } Y_{i,t+1} - Y_{it} < -b \\ \bar{\xi}(2) & \text{if } -b \leq Y_{i,t+1} - Y_{it} \leq b \\ \bar{\xi}(3) & \text{if } Y_{i,t+1} - Y_{it} > b \end{cases}$$

Note that these restrictions imply that the following combinations of states do not exist: $(\bar{y}(1), \bar{\xi}(3)), (\bar{y}(2), \bar{\xi}(1)), (\bar{y}(2), \bar{\xi}(3)), (\bar{y}(3), \bar{\xi}(1))$. This means that the transition probability matrix would look like:

$$(\mathcal{Y} \times \mathcal{X}) = \begin{bmatrix} (\bar{y}(1), \bar{\xi}(1)) \\ (\bar{y}(1), \bar{\xi}(2)) \\ (\bar{y}(2), \bar{\xi}(2)) \\ (\bar{y}(3), \bar{\xi}(2)) \\ (\bar{y}(3), \bar{\xi}(3)) \end{bmatrix}, P = \begin{bmatrix} 0 & p_{(1,1);(1,2)} & p_{(1,1);(2,2)} & 0 & p_{(1,1);(3,3)} \\ 0 & p_{(1,2);(1,2)} & p_{(1,2);(2,2)} & 0 & p_{(1,2);(3,3)} \\ 0 & p_{(2,2);(1,2)} & p_{(2,2);(2,2)} & p_{(2,2);(3,2)} & 0 \\ p_{(3,2);(1,1)} & 0 & p_{(3,2);(2,2)} & p_{(3,2);(3,2)} & 0 \\ p_{(3,3);(1,1)} & 0 & p_{(3,3);(2,2)} & p_{(3,3);(3,2)} & 0 \end{bmatrix}.$$

If we want $\bar{\xi}(1)$ to be a more persistent state than state $\bar{\xi}(2)$, then the probability of staying at current earnings level $\bar{y}(k)$ should be larger when in $\bar{\xi}(1)$. For example, $p_{(1,1);(1,2)}$ should be larger than $p_{(1,2);(1,2)}$. Similarly, $\bar{\xi}(3)$ is a more persistent state than $\bar{\xi}(2)$ if $p_{(3,3);(3,2)} > p_{(3,2);(3,2)}$. \triangleleft

B.7 Heterogeneous signals

The general earnings process can also be mapped into an earnings process with heterogeneous expectations that arise due to signals received by the individual as in Stoltenberg and Singh (2020). For this mapping, the interpretation of unobserved state ξ is that it is the signal received by the individual, S_{it} , which takes realizations from the set \mathcal{Y} , so $(Y_{it}, S_{it}) \in (\mathcal{Y} \times \mathcal{Y})$. The signal has a certain level of informativeness, represented by $\kappa = P(Y_{it+1} = \bar{y}(j) | S_{it} = \bar{y}(j))$, i.e. the probability that the signal S_{it} coincides with next periods earnings realization.

For the transition probability matrix of states (Y_{it}, S_{it}) , the following holds:

$$\begin{aligned} & P(Y_{it+1} = \bar{y}(k), S_{it+1} = \bar{y}(m) | Y_{it} = \bar{y}(j), S_{it} = \bar{y}(l)) \\ & = P(S_{it+1} = \bar{y}(m) | S_{it} = \bar{y}(l)) \cdot P(Y_{it+1} = \bar{y}(k) | Y_{it} = \bar{y}(j), S_{it} = \bar{y}(l)). \end{aligned}$$

$P(k_{it+1} = \bar{y}(m) | k_{it} = \bar{y}(l))$ depends on the process that is assumed for the signals. Given

the assumptions on the signal's informativeness,

$$P(Y_{it+1} = \bar{y}(k) | Y_{it} = \bar{y}(j), S_{it} = \bar{y}(l)) = \frac{p_{jk} \kappa^{I[l=k]} \left(\frac{1-\kappa}{1-L}\right)^{1-I[l=k]}}{\sum_{z=1}^L p_{jz} \kappa^{I[l=z]} \left(\frac{1-\kappa}{N-1}\right)^{1-I[l=z]}},$$

where p_{ij} is the probability $P(Y_{it+1} = \bar{y}(j) | Y_{it} = \bar{y}(i))$. Recall that L is the number of possible earnings realizations, that is, the cardinality of \mathcal{Y} . If the process governing the signal transitions, $P(S_{it+1} = \bar{y}(m) | S_{it} = \bar{y}(l))$, is left unrestricted, this implies a total of $L \times (L - 1)$ transition probabilities for the signals, $L \times (L - 1)$ transition probabilities for the earnings, and the parameter κ , which is less than the number of parameters of the general earnings process, with $(L \times M) \times (L \times M - 1)$ parameters, for $M \geq 2$ if $L \geq 1$. Here M is the cardinality of the set of unobserved states \mathcal{X} .

C Identification of α , δ , μ , γ and β

First, consider the identification of the first-step parameters α , δ and μ . A set of moment conditions can be derived that identifies parameters α (elasticity of capital), δ (depreciation rate), and μ (unemployment benefits), such that these moment conditions do not depend on the other structural parameters in the model.

The following four equations follow from the government and firm-side problem:

$$\begin{aligned} Y_t &= K_t^\alpha (\bar{l} L_t)^{1-\alpha}, \\ w_t &= (1 - \alpha) \bar{l} \left(\frac{K_t}{\bar{l} L_t} \right)^\alpha, \\ r_t + \delta &= \alpha \left(\frac{K_t}{\bar{l} L_t} \right)^{\alpha-1}, \\ (\bar{l} L_t) \tau_t &= \mu u_t. \end{aligned}$$

\bar{l} is set to $\frac{1}{\bar{L}}$, where \bar{L} denotes the long-run average of L_t . Also denote the long-run values of the other variables as \bar{Y} , \bar{K} , \bar{w} , \bar{r} , $\bar{\tau}$ and \bar{u} respectively. Also, $\bar{L} = 1 - \bar{u}$. The following relations then hold:

$$\begin{aligned} \bar{Y} &= \bar{K}^\alpha, \\ \bar{w} &= (1 - \alpha) \left(\frac{1}{1 - \bar{u}} \right) \bar{K}^\alpha, \\ \bar{r} + \delta &= \alpha \bar{K}^{\alpha-1}, \\ \bar{\tau} &= \mu \bar{u}. \end{aligned}$$

The following moment conditions can be used to consistently estimate α , δ and μ :

$$\begin{aligned} \mathbb{E} \left(\frac{w_t^{\text{obs}} (1 - u_t^{\text{obs}})}{Y_t^{\text{obs}}} - (1 - \alpha) \right) &= 0, \\ \mathbb{E} \left(\frac{Y_t^{\text{obs}}}{(\delta K_t)^{\text{obs}}} - \frac{r_t^{\text{obs}} + \delta}{\delta \alpha} \right) &= 0, \\ \mathbb{E} \left(\frac{\tau_t^{\text{obs}}}{u_t^{\text{obs}}} - \mu \right) &= 0, \end{aligned}$$

under the assumption that the sample means of the variables τ_t^{obs} , u_t^{obs} , w_t^{obs} , Y_t^{obs} , r_t^{obs} and $(\delta K_t)^{\text{obs}}$ converge to their long-run values.²³

²³ $(\delta K_t)^{\text{obs}}$ is used because this variable is directly observed in the data, see also the Data section.

Local (first-order) identification can be shown by deriving the Jacobian of these moment conditions with respect to the parameters α , δ and μ , and by showing that this Jacobian is full rank. This is easily verified, as the first moment condition does not depend on δ and μ , but does depend on α , the second condition does not depend on μ , but does depend on δ and α , and the third moment condition does not depend on δ and α , but does depend on μ . It follows that, for values of δ , α and μ unequal to zero, and α unequal to one, the Jacobian will be full rank.

Janssens (2019) shows that given an idiosyncratic shock distribution, one can identify the risk aversion parameter from the shape of the wealth distribution. However, we do need to verify if it is possible to jointly identify the risk aversion parameter and the transition probabilities of the idiosyncratic shock process. This can only be explored numerically.

Example 6 Consider the setting described in Example 3, where there is no wage heterogeneity ($L = 1$), but there is risk heterogeneity ($M = 2$). Using the same set of moments, with the additional moment $(1/5)\text{Percentile}_{50}(k_i)/\text{Percentile}_{20}(k_i)$, the Jacobian of the moments with respect to the transition probability matrix is evaluated, around risk aversion parameter $\gamma = 1$. This results in a rank of 7 (full) and a condition number of 39.77. ◁

This indicates numerically that the parameters of the earnings process and the risk aversion parameter can be jointly identified, even when both are using the same source of information (moments of a cross-section of net worth data). However, the identification strategy that will be used below will pin down the transition probabilities from a savings and earnings panel, while the identification of the risk aversion comes from moments of the wealth distribution, meaning that identification will come from distinct sources.

The discount factor β is not estimated, but follows as a solution from the general equilibrium model for a given choice of the interest rate level r .

A similar condition can be derived using K_t^{obs} instead.

D Propositions and proofs on the identification of the heterogeneous earnings risk process

D.1 Proof of Proposition 1

Proposition 1. *Consider an first-order Markov process with earnings states $\{\bar{y}(1), \dots, \bar{y}(L)\} \times \{\bar{\xi}(1), \dots, \bar{\xi}(M)\}$ and transition probability matrix P . Any set of moments derived from an $N \times T$ earnings (repeated) cross section, with $T \geq 1$ and observations on earnings levels $y \in \mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\}$ can generate a Jacobian of the moments with respect to the transition probabilities P with a rank at most L .*

The process is characterized by an observed Markov chain with levels $y \in \mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\}$, and transition probability matrix

$$Q = \begin{bmatrix} q_{1;1} & \cdots & q_{1;L} \\ q_{2;1} & \cdots & q_{2;L} \\ \vdots & \ddots & \vdots \\ q_{L;1} & \cdots & q_{L;L} \end{bmatrix}$$

Denote the L stationary probabilities corresponding to this observed Markov chain by q_1, q_2, \dots, q_L . On the other hand, the underlying earnings process has states $\mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\} \times \{\bar{\xi}(1), \dots, \bar{\xi}(M)\}$ and transition probability matrix P given by

$$P = \begin{bmatrix} P_{(1,1);(1,1)} & P_{(1,1);(1,2)} & \cdots & P_{(1,1);(1,M)} & P_{(1,1);(2,1)} & \cdots & P_{(1,1);(L,M)} \\ P_{(1,2);(1,1)} & P_{(1,2);(1,2)} & \cdots & P_{(1,2);(1,M)} & P_{(1,2);(2,1)} & \cdots & P_{(1,2);(L,M)} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ P_{(L,M);(1,1)} & P_{(L,M);(1,2)} & \cdots & P_{(L,M);(1,M)} & P_{(L,M);(2,1)} & \cdots & P_{(L,M);(L,M)} \end{bmatrix}$$

Denote the $L \times M$ stationary probabilities of this Markov chain by $p_{1,1}, \dots, p_{L,M}$. The following mapping exists between the stationary probabilities of the observed Markov chain and the underlying Markov chain:

$$q_i = \sum_{j=1}^M p_{i,j}.$$

The stationary probabilities are functions of the underlying transition probabilities. When deriving moments from the cross-sectional distribution of this earnings process, these moments will be functions of the levels and the stationary probabilities that cor-

respond to these levels. It follows that the Jacobian of any set of moments derived from a cross section of earnings and employment data with respect to the transition probabilities can have rank at most L . \square

D.2 Proof of Proposition 2

Proposition 2. *Consider the first-order Markov process with earnings states $\{\bar{y}(1), \dots, \bar{y}(L)\} \times \{\bar{\xi}(1), \dots, \bar{\xi}(M)\}$ and transition probability matrix P . Any set of moments derived from an $N \times T$ earnings and employment panel with observations on earnings levels $y \in \mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\}$ can generate a Jacobian of the moments with respect to the transition probabilities P with a rank at most $L \times (L - 1)$.*

Proof The process is characterized by an observed Markov chain with levels $y \in \mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\}$, and transition probability matrix

$$Q = \begin{bmatrix} q_{1;1} & \dots & q_{1;L} \\ q_{2;1} & \dots & q_{2;L} \\ \vdots & \ddots & \vdots \\ q_{L;1} & \dots & q_{L;L} \end{bmatrix}$$

whereas the underlying earnings process has states $\mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\} \times \{\bar{\xi}(1), \dots, \bar{\xi}(M)\}$ and transition probability matrix P given by

$$P = \begin{bmatrix} P_{(1,1);(1,1)} & P_{(1,1);(1,2)} & \dots & P_{(1,1);(1,M)} & P_{(1,1);(2,1)} & \dots & P_{(1,1);(L,M)} \\ P_{(1,2);(1,1)} & P_{(1,2);(1,2)} & \dots & P_{(1,2);(1,M)} & P_{(1,2);(2,1)} & \dots & P_{(1,2);(L,M)} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ P_{(L,M);(1,1)} & P_{(L,M);(1,2)} & \dots & P_{(L,M);(1,M)} & P_{(L,M);(2,1)} & \dots & P_{(L,M);(L,M)} \end{bmatrix}$$

One can establish a mapping from P to Q :

$$q_{m;n} = \sum_{j=1}^M \left(P(\eta = \bar{\eta}(j) | y = \bar{y}(m)) \cdot \left(\sum_{i=1}^M P_{(m,j);(n,i)} \right) \right),$$

where $P(\eta = \bar{\eta}(j) | y = \bar{y}(m))$ can be found by obtaining the eigenvector of P corresponding to the eigenvalue of 1 and compute the conditional probability from $P_{\text{stationary}}$.

Deriving moments based on the observed panel data with states $\mathcal{Y} = \{\bar{y}(1), \dots, \bar{y}(L)\}$ and transition matrix Q shows that these moments can be expressed as a function of

the $L \times (L - 1)$ linearly independent transition probabilities in Q , and the levels of \mathcal{Y} . It follows that therefore the Jacobian of these moments with respect to the underlying transition probabilities in P will be of reduced rank, as P has $(L \times M) \times (L \times M - 1)$ moments to be identified, which is larger than $L \times (L - 1)$ for $M \geq 2$. \square

D.3 Proof of Proposition 3

Proposition 3 *Consider the first-order Markov process with earnings states $\{\bar{y}(1), \dots, \bar{y}(L)\} \times \{\bar{\xi}(1), \dots, \bar{\xi}(M)\}$ and transition probability matrix P . Assume Assumptions 1 and 2a-2b hold. A panel of individual net worth k_{it} , and earnings/employment status y_{it} , with $t = 1, \dots, T$ and $i = 1, \dots, N$, $T \geq 2$, can be used to perfectly retrieve a panel of (reordered) unobserved states ξ_{it} for $t = 2, \dots, T$ and $i = 1, \dots, N$.*

Proof By assumption 2a, for each (discrete) level of $y = \bar{y}(j)$, $j = 1, \dots, L$, the domain of k can be subdivided into K_N^j subdomains $[k_1^j, k_2^j], \dots, [k_i^j, k_{i+1}^j], \dots, [k_{K_N^j}^j, k_{K_N^j+1}^j]$. For each subdomain $[k_i^j, k_{i+1}^j]$ where all individuals have a value of $y = \bar{y}(j)$, a rule can be designed such that each observation is classified under the correct risk state ξ . Here, it should be noted that the exact labeling of an observation having risk state $\bar{\xi}(l)$ is not important, as this value does not have an interpretation itself, but the idea is rather that there are M risk states for each of the L income states, and that the observations that belong to the same risk state should be correctly categorized as belonging to one of these $L \times M$ distinct groups, and should not be grouped with observations from another risk state. This is where assumption 1 comes in, which relabels the risk states based on the saving levels to which they correspond, and denotes them by $\tilde{\varepsilon}_{j1}, \dots, \tilde{\varepsilon}_{jM}$.

First, for each subdomain $[k_i^j, k_{i+1}^j]$, the following thresholds should be determined: $k_1^{j,\max}, \dots, k_{M-1}^{j,\max}$. Consider all observations within this subdomain, and sort them in increasing order of k as $(k_0, k'_0), (k_1, k'_1), \dots, (k_n, k'_n)$ where $n + 1$ is the number of observations in the subdomain, k_i denotes the net worth of an individual at time t , $y = \bar{y}(j)$ is the income of the individual at time t and k'_i is the choice of net worth for next period ($t + 1$). So $k_0 \leq k_1 \leq \dots \leq k_n$.

Next, $k_{M-1}^{j,\max}$ is defined as:

$$\begin{aligned} & \max q \in \{0, \dots, n\} \text{ s.t.} \\ & (\exists p : (k_p < k_q \wedge k'_p > k'_q)) \wedge (\nexists m : k'_m > k'_p \wedge (\exists n : (k_n < k_m \wedge k'_n > k'_m))) \\ & \Rightarrow k_1^{M-1,\max} := k'_q \end{aligned}$$

Now, label all observations in this subdomain for which $k' > k_{M-1}^{j,\max}$ as $\tilde{\varepsilon}_{jM}$. Then, remove all observations with label $\tilde{\varepsilon}_{jM}$ from the ordered list. Repeat the same procedure to define the other thresholds $k_m^{j,\max}$ for $m = M - 3, \dots, 1$ and iteratively assign the labels $\tilde{\varepsilon}_{j(M-1)}, \dots, \tilde{\varepsilon}_{j2}$. After having assigned the label $\tilde{\varepsilon}_{j2}$, and having removed those observations from the ordered list, the remaining observations should be labeled as $\tilde{\varepsilon}_{j1}$. Note that the points that define these thresholds always exists through Assumption 2b.

Next, it is left to show that these labels are assigned correctly, that is, no observation is categorized together with an observation that has a different risk state. This follows from the second part of Assumption 2a, that is, in the interval $[k_i^j, k_{i+1}^j]$, if an observation (k_p, k'_p) has $k'_p > k_i^{j,\max}$, then it must be that the risk state of this observation $\tilde{\varepsilon}_j \neq \tilde{\varepsilon}_{jl}$. Else, it would imply that the local Lipschnitz continuity of Assumption 2a is violated and the local Lipschnitz coefficient is larger than $Z_{ijl} = d(k'(k_i^j, \tilde{\varepsilon}_{jl}), k'(k_i^j, \tilde{\varepsilon}_{jl+1}))$. \square

E Auxiliary parameter estimates

Table E1: Auxiliary transition probability matrix computed for observed data, based on 40 control groups and 5 saving bins per group, using Equation (1).

Transition probabilities (from/to)		Unemployed	Low earnings		High earnings	
			Low savings	High savings	Low savings	High savings
Unemployed		0.415	0.285	0.122	0.138	0.041
Low earnings	Low savings	0.019	0.562	0.264	0.118	0.037
	High savings	0.017	0.563	0.197	0.177	0.047
High earnings	Low savings	0.008	0.120	0.071	0.562	0.239
	High savings	0.005	0.086	0.043	0.601	0.265

Table E2: Auxiliary transition probability matrix computed for data simulated from the model with estimates from the first identification strategy, based on a single control cluster and 200 saving bins for the low wage group, and 20 saving bins for the high wage group, using Equation (1).

Transition probabilities (from/to)		Unemployed	Low earnings		High earnings	
			Low savings	High savings	Low savings	High savings
Unemployed		0.332	0.353	0.168	0.077	0.071
Low earnings	Low savings	0.020	0.491	0.247	0.177	0.066
	High savings	0.031	0.488	0.269	0.140	0.073
High earnings	Low savings	0.030	0.102	0.071	0.550	0.248
	High savings	0.030	0.096	0.068	0.556	0.250

Table E3: Auxiliary transition probability matrix computed for data simulated from the model with estimates from the combination of the first and second identification strategy, based on a single control cluster and 200 saving bins for the low wage group, and 20 saving bins for the high wage group, using Equation (1).

Transition probabilities (from/to)		Unemployed	Low earnings		High earnings	
			Low savings	High savings	Low savings	High savings
Unemployed		0.591	0.208	0.174	0.022	0.006
Low earnings	Low savings	0.042	0.274	0.233	0.276	0.176
	High savings	0.042	0.297	0.212	0.290	0.159
High earnings	Low savings	0.007	0.153	0.106	0.448	0.286
	High savings	0.006	0.146	0.120	0.506	0.222

Table E4: Other moments used as auxiliary parameters: observed, as simulated from the model using the estimates from identification strategy one, and as simulated from the model using the estimates from the combination of identification strategy one and two.

Moments	Observed	Identification strategy 1& 2	Identification strategy 1
$\text{Prctl}_{90}(k_{it})/\text{Prctl}_{15}(k_{it})$	320	2.507	2.849
$\text{Prctl}_{99}(k_{it})/\text{Prctl}_{90}(k_{it})$	4.056	1.285	1.319
$\text{Prctl}_{50}(k_{it})/\text{Prctl}_{15}(k_{it})$	35.9	1.566	1.728
$\text{Prctl}_{90}(w_{it})/\text{Prctl}_{30}(w_{it})$	4.367	4.726	5.007
Unrate (>6 months)	0.045%	0.046%	0.036%