Income Volatility and Portfolio Choices*

Yongsung Chang
Seoul National University & SIER

Jay H. Hong
Seoul National University & SIER

Marios Karabarbounis
Federal Reserve Bank of Richmond

Yicheng Wang
University of Oslo

March 15, 2020

Abstract

Based on administrative data from Statistics Norway, we find economically significant shifts in households' financial portfolios around structural breaks in income volatility. When the standard deviation of labor-income growth doubles, the share of risky assets decreases by 4 percentage points. We ask whether this estimated marginal effect is consistent with a standard model of portfolio choice with idiosyncratic volatility shocks. The standard model generates a much more aggressive portfolio response than we see in the data. We show that Bayesian learning about the underlying volatility regime can reconcile the gap between the model and the data.

JEL Classification: E2, G1, J3.
Keywords: Income Volatility, Portfolio Choice, Risky Share, Bayesian Learning

*Emails: yongsung.chang@gmail.com, jayhong@snu.ac.kr, marios.karabarbounis@rich.frb.org, and yicheng.wang@econ.uio.no. Any opinions expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. For helpful comments, we thank our discussant José Mustre-del-Río, and also Mark Aguiar, Mark Bils, Marcus Hagedorn, Grey Gordon, Per Krusell, Ben Moll, Luigi Pistaferri, Kjetil Storesletten, Morten Ravn, Alex Wolman and other participants at U of Oslo, Econometric Society China Meeting 2018 in Shanghai, SED 2018, the 12th Nordic Summer Symposium in Macroeconomics, EEA/ESEM 2018 in Cologne, AEA 2019, Rochester, Yonsei, Seoul National U, Pittsburgh, Richmond Fed, SHUFE, Tsinghua PBC and ITAM. Yicheng Wang acknowledges financial support from the European Research Council ((FP7/2007-2013)/ERC grant n.324085, Principal investigator: Kjetil Storesletten at the University of Oslo). We also thank the Frisch Center at the University of Oslo, Andreas Moxnes and Marcus Hagedorn for access to the administrative data of Statistics Norway.

1 Introduction

How do households respond to background risk? A large literature has studied how the presence of uninsurable labor-income risk affects the patterns of savings, consumption, and portfolio allocation over the life cycle. Despite the extensive research done, there is still substantial progress to be made. On one hand, quantitative models analyzing the role of labor-market risk in the allocation of financial assets often report difficulties in matching the empirical patterns of portfolio choice, such as the average stock holdings or the risky share profile across ages (for example, Cocco, Gomes, and Maenhout (2005), Benzoni, Collin-Dufresne, and Goldstein (2011), Huggett and Kaplan (2016), to name only a few). On the other hand, empirical studies—which are much fewer in number—have difficulty in finding an economically sizable effect of labor-income volatility on households' portfolios (for example, Guiso, Jappelli, and Terlizzese, 1996; Palia, Qi, and Wu, 2014).

Our paper contributes to this literature both empirically and quantitatively. First, based on administrative panel data from Statistics Norway, we establish a sizable negative relationship between income risk and the risky share in financial assets. We overcome two problems that have led to small estimated responses in the literature: (i) the lack of high-quality panel data, and (ii) the measurement of *exogenous* variations in income volatility. Second, we build a structural model that successfully reproduces the estimated response of portfolio to income risk. This is a new attempt because the existing studies have been mostly interested in matching the average risky share not the *marginal* effect of income risk. As our analysis illustrates, the marginal effect is useful to (i) understand the nature of earning dynamics and (ii) assess the welfare cost of background risk.

Households in Norway are obliged to report detailed information about their income and wealth to the tax authority every year. As a result, our dataset includes a complete description of households' labor income and financial assets as well as their allocation to safe and risky financial accounts. We merge the households' income and financial data with other data regarding labor market status, demographic characteristics, and, more importantly for our analysis, employer information.

We overcome the challenges in estimating the marginal effect of income risk in two ways. First, we identify the "structural'' breaks in income volatility, which are the periods when an individual worker experiences the largest change in the standard deviation of income growth. By looking at big events, we can potentially avoid noisy variations unrelated to true regime changes. Second, since not all the structural breaks are exogenous (or unpredictable) to households, we use firm-side information as an instrumental variable—an innovative method pioneered by Fagereng, Guiso, and Pistaferri (2017). Specifically, based on a matched

employer-employee register, we use the volatility of firms' sales as an instrument to isolate an orthogonal variation in the individual worker's income volatility.

We find a clear negative relationship between income volatility and risky share. According to our instrumental variable estimation, when the income volatility doubles, the risky share decreases by 4 percentage points over a 4-year horizon. This economically sizable estimate illustrates that our methodology (individual structural breaks combined with a firm-side instrumental variable) is effective.

We then build a structural model of portfolio choice to reproduce the estimated risky share response to income volatility. Our benchmark structural model features: (i) a life-cycle economy with incomplete asset markets, (ii) a portfolio choice between risk-free bonds and risky equity (e.g., Cocco, Gomes, and Maenhout (2005) and Gomes and Michaelides (2005)), (iii) an exogenous borrowing limit, (iv) labor earnings that consist of a mix of heterogeneous-income profiles (Guvenen and Smith, 2014) and uninsurable shocks, and finally (iv) our highlighted new element, idiosyncratic shocks to income volatility.

The structural model is estimated by indirect inference using various moments from the Norwegian panel including the estimated portfolio response to volatility. To identify the (hard-to-observe) persistence of the idiosyncratic volatility shock, we exploit the age profile of the kurtosis of earnings emphasized by Guvenen, Karahan, Ozkan, and Song (2015): the age profile of the cross-sectional kurtosis heavily depends on this persistence of the volatility shocks.[1]

We show that the standard model cannot replicate the portfolio response in the data. It predicts that in response to an increase in (uninsurable) income volatility, a typical household should decrease the risky share much more aggressively than what we find from the Statistics Norway panel. To fill the gap between the model and the data, we introduce Bayesian learning about the volatility of underlying income process. While there are other ways to introduce frictions into the model—e.g., adjustment costs in rebalancing the portfolio, Bayesian learning has a couple of merits. First, it is a parsimonious way of reconciling the model with the data—it does not introduce an additional free parameter. Second, adjustment costs are probably less important for the intensive margin (the adjustment of portfolios for households that already participated in risky investments) on which our analysis is focused.[2]

Finally, according to our structural model, the welfare cost of background risk is fairly big when a worker experiences an increase in income risk in the early stage of life. For example,

[1]This identification strategy mirrors that by Storesletten, Telmer, and Yaron (2004), who show that the persistence of labor-income shocks is important for the slope of the life-cycle profile of the cross-sectional variance of income.

[2]Adjustment costs have been found to play an important role at the extensive margin, the decision to participate in the stock market or not.

a three-fold increase in the standard deviation of income growth at age 22 can generate a welfare cost of 9 percent in consumption-equivalent units. While the majority of welfare losses come from the volatile consumption (lack of insurance), a half-of-a-percentage point welfare loss is due to rebalancing the financial portfolio toward safe assets—missing the opportunity to exploit the profitable but risky investment (the equity premium).

Our contribution to the existing literature on labor-market risk and portfolio choice can be summarized as follows. Most of the literature (based on cross-sectional data) reports that an effect of background risk on risky share is qualitatively consistent with economic theory but quantitatively small. Based on the Italian household survey of expectations about future income, Guiso, Jappelli, and Terlizzese (1996) find a small effect of risk on portfolio. Palia, Qi, and Wu (2014) report that the impact of background risk (such as labor income, house prices, and business income) on the risky share is not sizable: a one standard deviation increase in the labor-income variance (which amounts to a one-and-half times increase in the standard deviation of labor-income growth on average) reduces the risky share by 1.9 percentage points, much smaller than our estimates. A notable exception is Angerer and Lam (2009). Based on the National Longitudinal Survey of Youth for 1979, a 10 percent increase in the standard deviation of (the permanent component of) labor-income growth decreases the risky share by 3.7 percentage points, close to an upper bound in the literature and almost three times larger than our estimates.

We overcome the identification challenges by identifying individual structural breaks (coupled with a firm-side instrumental variable) from a high-quality administrative panel data set from Statistics Norway. Our analysis yields a large and precisely estimated response of risky share to income risk.[3] We closely follow the methodology of Fagereng, Guiso, and Pistaferri (2017), who first used the firm-side information as an instrument to isolate the orthogonal variations in households' income risk. We then expand their method by identifying the individual-specific structural breaks in the underlying income process. Our approach of using structural breaks has its own merit. By further reducing the influence of frequent noisy events, it increases the estimated response by 20 percent.

There has been extensive progress in quantitative analysis of portfolio choice and labor-market risk, represented by Heaton and Lucas (2000), Haliassos and Michaelides (2003), Cocco, Gomes, and Maenhout (2005), Gomes and Michaelides (2005), Benzoni, Collin-

[3]There has been increased use of the administrative data from Statistics Norway. Fagereng, Gottlieb, and Guiso (2017) analyze the portfolio responses to volatility or portfolio allocation over the life cycle. Fagereng and Halvorsen (2015) study household debt and heterogeneity in the marginal propensity to consume. Eika, Mogstad, and Vestad (2017) analyze consumption expenditure using data on income and assets. Fagereng, Holm, Moll, and Natvik (2019) analyze the saving rate across the wealth distribution and highlight the importance of capital gains.

Dufresne, and Goldstein (2011), Athreya, Ionescu, and Neelakantan (2015), Huggett and Kaplan (2016), Fagereng, Gottlieb, and Guiso (2017), Chang, Hong, and Karabarbounis (2018), and Catherine (2019), to name only some. The existing literature has been mostly interested in matching the average risky share, not the *marginal* effect of income risk on which we focus.[4] Our analysis differentiates itself from this literature by (i) introducing idiosyncratic volatility shocks and parsimonious Bayesian learning about the underlying income process and (ii) by structurally estimating and testing the model.

Our paper also contributes to the literature on the estimation of earnings dynamics. Structural models of earnings dynamics typically analyze the level of labor earnings in conjunction with data on consumption choices (for example, Primiceri and van Rens, 2009; Guvenen and Smith, 2014). In contrast, we build a structural model to analyze the individual *volatility* of labor earnings using as information a reliable estimated response of portfolio choice to income risk. The risky share is a suitable alternative target on which to base our estimation, especially since quality panel data on consumption choices are hard to obtain and—as we show in our empirical analysis—households actively use the portfolio margin to insure against larger income fluctuations.

Similarly, our paper is related to the literature analyzing the dynamic process for idiosyncratic volatility. Meghir and Pistaferri (2004) model volatility dynamics using an ARCH specification based on data from the Panel Study of Income Dynamics. Guvenen, Karahan, Ozkan, and Song (2015) use tax-record data to document a series of stylized facts regarding higher-order moments of the earnings distribution. They estimate a flexible specification that allows for i.i.d. volatility shocks. We consider a dynamic process for volatility that we model as an AR(1). We use higher-order moments of labor-income growth to discipline the size and persistence of the volatility shocks.[5]

The rest of the paper is structured as follows. Section 2 describes our empirical specifications and documents the basic patterns regarding the response of portfolio choice to income volatility. Section 3 sets up the structural model. Section 4 describes the estimation results as well as the identification of the structural parameters. Section 5 explores specific aspects of stochastic volatility, such as the workers' information set or its persistence, and also computes welfare losses from volatility fluctuations. Section 6 concludes.

[4]There are a few exceptions to this; Gollier and Pratt (1996) develop a model but without endogenous wealth accumulation. Bertaut and Haliassos (1997) and Viceira (2001) analyze the optimal allocation of assets under various degrees of labor-income risk.

[5]Our paper introduces learning about the volatility of labor income, whereas the existing literature has focused on learning about the level of income (Pischke, 1995; Guvenen, 2007; Guvenen and Smith, 2014).

2 Empirical Analysis

We utilize a wealth of information regarding labor income, asset holdings, and portfolio composition from Statistics Norway to document two facts:

1. A large fraction of workers experience sharp changes in the volatility of their labor-income growth (measured by the standard deviation).

2. The risky share of financial assets significantly decreases (increases) in response to an increase (decrease) in labor-income volatility.

2.1 Data

The Norwegian Registry is a set of comprehensive, relatively measurement-error-free data with detailed information on labor income and household financial assets. Households in Norway are subject to not only an income tax but also a wealth tax. Thus, they are obliged to report their complete income and wealth holdings to the tax authority every year. Employers, banks, brokers, insurance companies, and any other financial intermediaries are also obliged to send information on the value of personal assets to both the individual and the tax authority.[6]

The financial accounts in our data include bank deposits, financial securities, shares in mutual funds, shares in private companies, pension agreements, insurance policies, total debt (loans, credit purchases, mortgages), and others. We also have information on homeownership as well as house values.[7]

We merge our wealth data with other data sets such as: (1) the Income Registry Data, which have detailed information on earned income including cash salary, taxable benefits and sickness and maternity benefits each year, capital investment income, entrepreneurial income, unemployment benefits, and pensions; (2) the Central Population Register, which contains yearly individual demographic information (e.g., gender, date of birth, marital status, number of children, to name a few); (3) the National Educational Database, which has the history and the latest education record for each resident, and finally, (4) the Employer-Employee Register, which provides annual information on workers' labor market status (full-/part-time

[6]Traded financial securities are reported at market value. The value of shares in private companies is reported by individuals as well as private companies to the tax authority. The tax authority will then combine the information from companies' reports with those from individual households and adjust if necessary.

[7]Reliable information on the house value is available only for the period 2010-2014. For earlier years, the housing value for homeowners reported in the tax registry data may not be the true market values, typically underreported, due to self-reporting errors and/or treatment policies of wealth tax. This is well-documented in the literature (see Eika, Mogstad, and Vestad (2017), Fagereng and Halvorsen (2017), among others). For our definition of the risky share, we focus on financial assets and exclude housing and mortgage debt. Nonetheless, in the empirical analysis, we include homeownership dummies in the list of control variables in the regression.

employment, employer ID, beginning/ending time of job, total payments from each employer, industry, occupation, etc.). All data sets are merged using unique personal identifiers assigned to each individual and firm (similar to social security numbers and employment identification numbers in the U.S.). For more details on our data, see Appendix A.

The data are uniquely suitable to address many challenges arising in the empirical analysis of portfolio choices and income volatility. Traditional data sets, which are typically based on surveys, present at least four issues. First, respondents often misreport their labor income or wealth intentionally or unintentionally.[8] In our data, information is directly collected by third parties (employers or financial institutions) for tax purposes, which substantially reduces measurement errors. Second, household surveys are often top-coded. This is problematic when analyzing higher-order moments of earnings that may be driven by top earners (Guvenen, Karahan, Ozkan, and Song, 2015).[9] Third, traditional data with detailed information on households' financial assets (such as the Survey of Consumer Finances) are repeated cross-sections. Thanks to the panel dimension of our data, we can eliminate bias stemming from unobserved fixed heterogeneity (e.g., risk aversion across individuals) in the estimation. Fourth, there is frequent attrition in traditional data, whereas attrition in our data occurs only due to migration or death.

From the whole population with income tax registration records in 1993, we first choose the native Norwegian males, older than 25 in 1993, with no missing records on education and demographic variables. We then randomly select 10 percent from the above sample (amounting to 137,776 individuals and 2,880,970 person-year observations). For each individual male in our sample, we can then obtain his household information (marital status, demographics, total household income and wealth, and so on). For more details on our sample selection and construction, see Appendix A.2.

2.2 Risky Share

Following the standard literature, we classify financial assets into two categories: safe and risky. Safe assets include deposits in Norwegian banks, the cash value of life insurance policies, and debt securities traded in the financial market (mainly government bonds). Risky assets include shares in mutual funds, shares in private companies, and financial securities (mainly stocks and equity certificates traded in financial markets).[10] Total financial assets are the

[8]For example, see the handbook chapter "Measurement Error in Survey Data," by Bound, Brown, and Mathiowetz (2001).

[9]This is important for our analysis as the life-cycle profile of the kurtosis of labor-income growth plays a key role in identifying the persistence of volatility shocks (as we show in Section 4.1).

[10]Since we do not have detailed information on the riskiness of the individual's investment in mutual funds, we use aggregate statistics from Statistics Norway to split the assets in mutual funds into risky and

sum of safe and risky assets at the household level. The risky share of financial assets is the value of risky assets over the value of total financial assets. In our benchmark definition, we do not consider debt and focus on gross savings.

2.3 Individual Structural Break in Income Volatility

We measure labor-market risk based on an *individual* structural break in income volatility.[11] The main idea is to identify an episode of a "large" change in labor-income volatility.[12] More specifically, we look for a year when an individual worker experiences the largest change in terms of the standard deviation in labor-income growth. The algorithm to identify the structural volatility break is as follows:

1. Compute the residual (net of age and time effects) annual labor earnings of individual i at time t: $y_{i,t}$.

2. Construct labor-income growth: $\Delta y_{it} \equiv y_{it} - y_{i,t-1}$. We focus on the changes in income *growth* rather than the *level* to eliminate potential income variations due to heterogeneity in income profiles (which is strongly supported by the data).[13]

3. We then construct the standard deviation before and after τ, $SD(\Delta y)_{i,t<\tau}$ and $SD(\Delta y)_{i,t\geq\tau}$, respectively, for all τ. The change in income volatility for a worker i in year τ is $\Delta SD_{i,\tau} = SD(\Delta y)_{i,t\geq\tau} - SD(\Delta y)_{i,t<\tau}$.

4. Given the sequence of volatility changes for worker i: $\{\Delta SD_{i,\tau}\}$, we identify the structural break period τ^* such that $\tau^* = \text{argmax}_\tau \ \text{abs}(\Delta SD_i)$. The corresponding volatility change in the structural-break year is denoted by $\Delta SD_{i,\tau^*}$.

Using this methodology we identify the structural break year τ^* for each worker. Each structural break is associated with a positive or a negative change in the standard deviation of labor-income growth. Since we identify the largest change in the worker's income history, each worker has a single structural break.

risk-free components. Our results are quite robust to different splitting rules.

[11] For an application of this approach in the context of neighborhood segregation, see Card, Mas, and Rothstein (2008), and for housing and labor markets, see Charles, Hurst, and Notowidigdo (2018).

[12] In our benchmark specification, the unit of our analysis is the individual. Nonetheless, portfolio responses may be affected by insurance between the primary and the secondary earner. Therefore, we explore in our robustness section the portfolio responses when we also consider the household as the unit of analysis using household-level total disposable income.

[13] According to the labor-income specification in Section 3, labor income for worker i at age j is $y_{ij} = a_i + \beta_i \times j + x_{ij}$. Labor-income growth equals $\Delta y_{ij} = y_{ij} - y_{i,j-1} = \beta_i + \Delta x_{ij}$. Therefore, variability over some periods $Var(\Delta y_{ij})$ will ignore the constant term β_i and only consider the variability in the time-varying component $Var(\Delta x_{ij})$.

Figure 1: Construction of Individual-Specific τ^*

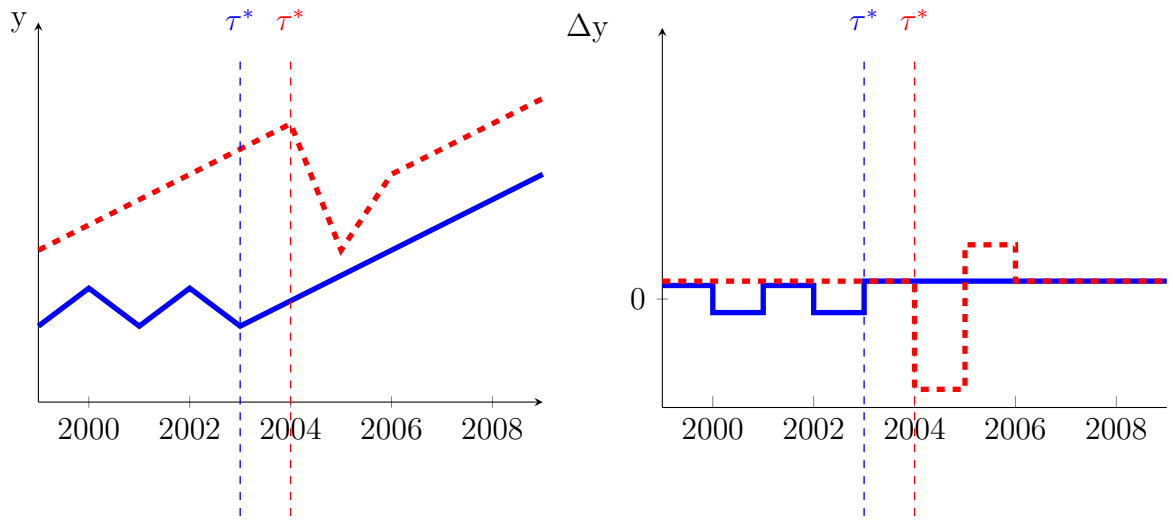

Notes: The left (right) panel shows hypothetical income paths (growth rates) for two workers. For each worker we show the year identified as a structural break by our method.

We demonstrate our methodology to identify the individual structural break in Figure 1. The left panel plots hypothetical labor-income paths, y, for two workers: worker A (solid line) and worker B (dotted line). The right panel shows the growth rate of income, Δy. Worker A's growth rate fluctuates before 2003, and stabilizes after 2003. The structural break τ^* occurs in 2003 and the volatility of income growth decreases after the structural break. Worker B's growth rate is constant up to 2004 and fluctuates thereafter. The structural break τ^* occurs in 2004, and the volatility of income growth increases after the structural break.

For our benchmark, the sample is restricted to workers with at least 18 years of observations of labor earnings. We focus only on employed workers with identified employer IDs, and we do not include self-employed workers. These restrictions decrease the sample to 48,768 individuals. Moreover, we require at least 16 years of positive risky shares, which decreases the sample to 18,156 individuals.[14] We also require total financial assets to be above 50,000 NOK in 2005 (which decreases the sample further to 16,700 individuals).[15] In the robustness section below, we confirm that our results are robust with different sample-selection criteria considered.

[14]Hence, we are basically looking at the response of risky shares conditional on participation.

[15]This is close to the 10th percentile of the cross-sectional distribution of total financial assets in 2005. We impose this restriction because the risky share could be very noisy for workers with a small amount of total financial assets.

2.4 Response of the Risky Share

We examine the household's portfolio choice around the individual volatility break τ^*. In particular, we compute the change in the risky shares with a window of k-years before and after the structural break, $\text{RS}_{i,\tau^*+k} - \text{RS}_{i,\tau^*-k}$, for $k = 1, 2, 3, 4$.

We estimate the response of the risky share with the following regression:

$$\text{RS}_{i,\tau^*+k} - \text{RS}_{i,\tau^*-k} \quad = \quad \beta \Delta SD_{i,\tau^*} + \quad \alpha X_{i,\tau^*} + \delta D_t + \epsilon_{i,\tau^*}. \tag{1}$$

Our benchmark specification includes year dummies (D_t) and a set of individual controls (denoted by X_{i,τ^*}). This includes *changes* (between $\tau^* - k$ and $\tau^* + k$) in log of household's disposable income, log of the household's wealth, total number of children, number of children younger than 10, and number of children younger than 5. We also include dummies for changes in marital status and homeownership. Moreover, we include the *levels* of log household income and wealth as well as age and education dummies at τ^*.

The coefficient β gives the change in the risky share (in percentage points) if the standard deviation of income growth increases by one unit. For example, the median of the standard deviation of income growth in our benchmark sample is 0.25 (see Table 9 in Appendix B). Thus, if income volatility is doubled for the median worker, the impact would be $\beta \times 0.25$ percentage points on the risky share.

Figure 2 plots a simple scatter-bin plot between the change in income volatility (measured as the largest change in income growth ΔSD_{τ^*}) and the change in the risky share over an 8-year window ($k = 4$) without any controls. It shows a clear negative relationship between changes in labor-income volatility and changes in the risky share. The plot confirms that around periods of heightened income volatility, households reduce their exposure to risk in financial investments by decreasing the risky share.

Table 1 reports the estimates of the regression. We show the estimates with and without any controls X_{i,τ^*}. For all windows, the risky share decreases in response to an increase in volatility. Without controls, when the standard deviation of income growth increases by one unit, the risky share decreases by 1.18, 4.36, 5.71, and 7.85 percentage points, for k=1 to 4, respectively. With controls, the risky share decreases by 0.77, 2.91, 2.37, and 2.91 percentage points, for k=1 to 4, respectively. Except for $k = 1$, all estimates are statistically significant at the 1 percent level. The control variables (most notably the level and the changes in income and wealth) absorb a lot of the changes in the risky share and decrease the magnitude of coefficients. Looking at a 4-year window ($k = 2$), if the standard deviation of income growth doubles in size for the median worker, the risky share decreases by 0.72 percentage points (2.91×0.25).

Figure 2: Risky Share and Income Volatility

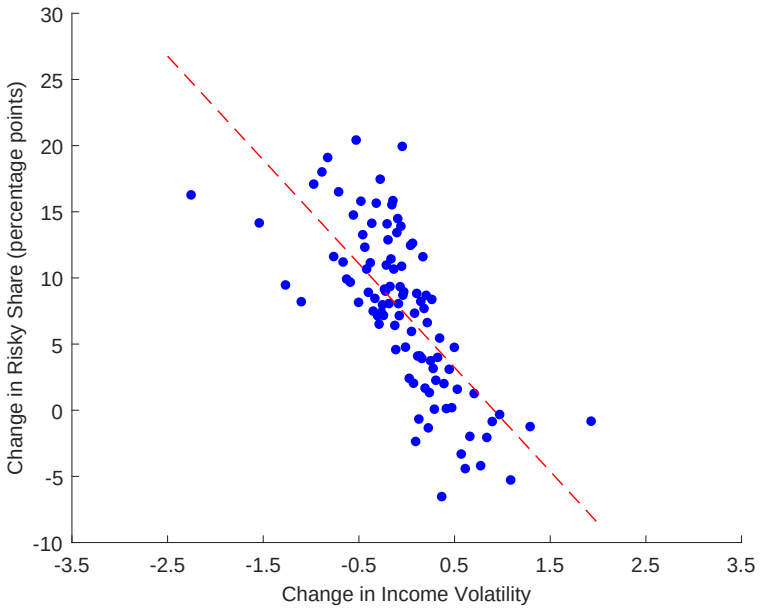

Notes: The figure is a bin-scatter plot. The $y-$axis represents the change in the risky share (ΔRS) and the $x-$axis represents the change in income volatility (ΔSD) over an 8-year ($k = 4$) window around the individual structural break $\Delta SD_{i,\tau^*}$.

In our analysis, we have addressed a major challenge highlighted in the literature estimating the portfolio response to volatility: unobserved heterogeneity. Studies that have estimated the volatility effect on portfolios typically rely on cross-sectional variation (with the exception of Fagereng, Guiso, and Pistaferri, 2017). In our regression, we estimate the change in the risky share to the change in the individual income volatility. Thus, our methodology eliminates unobserved fixed heterogeneity. Nonetheless, our OLS estimates of the risky share response are not substantially larger than those in the studies using cross-sectional variation. Therefore, based on our exercise, unobserved fixed heterogeneity is not necessarily a major driving force behind the small responses of risky share found in the literature.

A second major concern is the measurement of labor-income *risk*. Typically, risk is measured using the observed income volatility. However, risk can be very different from the observed volatility, especially if households already anticipate the changes in income volatility. In fact, a recent literature suggests that a substantial portion of the residual variation in earnings is predictable and reflects individual choices rather than risk (e.g., Primiceri and van Rens, 2009; Guvenen and Smith, 2014). According to Cunha and Heckman (2007), the statistical decomposition of earnings cannot distinguish uncertainty from other sources of income variability. Misinterpretation of labor-income volatility as pure income risk is likely

Table 1: Response of Risky Share to Income Volatility

	Controls	$k=1$	$k=2$	$k=3$	$k=4$
$\Delta SD_{i,\tau^*}$ (OLS)	No	-1.18**	-4.36***	-5.71***	-7.85***
		(0.50)	(0.58)	(0.62)	(0.65)
	Yes	-0.77	-2.91***	-2.37***	-2.91***
		(0.58)	(0.72)	(0.74)	(0.76)
$\widehat{\Delta SD}_{i,\hat{\tau}}$ (IV)	No	-8.45***	-20.70***	-25.67***	-33.67***
		(2.09)	(2.49)	(2.63)	(2.84)
	Yes	-4.96	-16.35***	-20.61***	-24.80***
		(4.12)	(4.86)	(5.23)	(5.43)
J test†		0.26	0.67	0.92	0.66

Notes: OLS estimates report the coefficients from regression (1). The IV estimates also use the first-stage regression (2). The numbers in parentheses are the robust standard errors. The asterisk(s) denote the statistical significance at three p-values: 10%, 5%, and 1%. † denotes the p-value for the over-identifying restriction tests (Hansen's J test). The first-stage F-tests have p-values of 0.00 for all cases and thus are not reported.

to bias the estimated response of portfolio choice toward zero.

In order to identify the marginal effect on portfolio choice of (uninsurable) income risk one needs exogenous variation in the latter. To this end, we employ an instrumental variable estimation based on firm-side variables, a method developed by Fagereng, Guiso, and Pistaferri (2017). The identifying assumption is that an individual worker cannot influence the firm's overall performance. To identify exogenous variations in income volatility, we use as instruments the volatility of growth rates for sales and value added (both scaled by assets) of firm f where a worker i is employed.[16] Using the exact same steps 1-3 described above, we compute the change in the volatility of sales, $\Delta SD^s_{f,t}$, and value added, $\Delta SD^v_{f,t}$, before and after period t. Henceforth, to simplify notation, we bundle both instruments in vector $\Delta SD_{f,t}$. As a first-stage regression, we run the following:

$$\Delta SD_{i,t} = \gamma \Delta SD_{f,t} + X_{i,t} + u_{i,t}, \tag{2}$$

where $X_{i,t}$ is the same set of worker characteristics described in Equation (1) at period t. The coefficient γ can be interpreted as the "pass-through" of firm volatility to workers' earnings

[16]Sales refer to gross revenue minus operating costs, and value added is gross revenue minus operating costs plus wage bills.

volatility.[17] By projecting $\Delta SD_{i,t}$ on $\Delta SD_{f,t}$, we obtain $\widehat{\Delta SD}_{i,t}$, an arguably exogenous component of earnings volatility.[18]

Based on the time series of $\widehat{\Delta SD}_{i,t}$'s, we identify the structural break for each individual. We denote the year of the exogenous structural break as $\hat{\tau}$ and the projected volatility change at that year as $\widehat{\Delta SD}_{i,\hat{\tau}}$. Note that in the benchmark IV regression, we also include workers who change firms around the structural break. Nonetheless, we estimate, as a robustness exercise, the portfolio responses for a subsample of workers who always stay with the same employer (roughly about 80 percent of the total sample) with the empirical estimates being similar in magnitude.

Table 2 compares the summary statistics of $\Delta SD_{i,\tau^*}$'s and $\widehat{\Delta SD}_{i,\hat{\tau}}$'s. Clearly, the exogenous variation of income volatility shows a much smaller dispersion, as the standard deviation of the volatility change decreases from 0.58 to 0.13. This occurs because our raw measure of volatility $\Delta SD_{i,\tau^*}$ is a mix of predictable and unpredictable episodes, while $\widehat{\Delta SD}_{i,\hat{\tau}}$ isolates episodes that are closer to how we think of background risk.

Table 2: Summary Statistics for $\Delta SD_{i,\tau^*}$ and $\widehat{\Delta SD}_{i,\hat{\tau}}$

	Obs.	Mean	S.D.	10th Percentile	90th Percentile
$\Delta SD_{i,\tau^*}$	7,982	-0.04	0.58	-0.65	0.59
$\widehat{\Delta SD}_{i,\hat{\tau}}$	7,982	-0.02	0.13	-0.16	0.16

Notes: Summary statistics for the change in the standard deviation of labor-income growth $\Delta SD_{i,\tau^*}$ and the projected change in the standard deviation of labor-income growth $\widehat{\Delta SD}_{i,\hat{\tau}}$.

The bottom panel of Table 1 reports the IV estimates.[19] Without controls, the risky share decreases by 8.45, 20.70, 25.67, and 33.67 percentage points, for k=1 to 4, respectively. As with the OLS estimation, the control variables mitigate the responses. In particular, with controls, the risky share decreases by 4.96, 16.35, 20.61, and 24.80 percentage points, for k=1 to 4, respectively. Looking at a 4-year window ($k = 2$), if the standard deviation of labor-income growth doubles in size for the median worker, the risky share decreases by 4.09

[17] To exclude outliers from the estimation, we only keep those between the 1st and 99th percentiles of $\Delta SD_{f,t}$'s in each cross-section. The estimate of γ in the first-stage regression is around 3 percent for sales and 0.8 percent for value added, both statistically significant at the 1 percent level. According to the standard test for over-identifying restrictions, both sales and value added are valid instruments. This suggests considerable insurance on behalf of the firms to the workers (see also Guiso, Pistaferri, and Schivardi, 2005 and Fagereng, Guiso, and Pistaferri, 2017).

[18] We also inspect the correlation between current risky share changes, $RS_{i,t+1} - RS_{i,t}$, and future shocks, $\widehat{\Delta SD}_{i,t+k}$, $k \geq 2$. We find basically that there is no significant correlation. That is, households cannot anticipate our identified exogenous shocks in the future.

[19] The OLS estimation also considers the same sample of workers who can be matched with their employer.

Table 3: Using Structural Breaks vs. All Observations

	$k=1$	$k=2$	$k=3$	$k=4$
Using Structural Breaks				
Dependent Variable: $RS_{i,\hat{\tau}+k} - RS_{i,\hat{\tau}-k}$				
Coefficient on $\widehat{\Delta SD}_{i,\hat{\tau}}$ (IV)	-4.96	-16.35***	-20.61***	-24.80***
	(4.12)	(4.86)	(5.23)	(5.43)
Obs.	7,851	7,729	7,600	7,440
Using All Observations				
Dependent Variable: $RS_{i,t+k} - RS_{i,t-k}$				
Coefficient on $\widehat{\Delta SD}_{i,t}$ (IV)	-7.48***	-13.63***	-17.11***	-22.51***
	(1.91)	(2.25)	(2.43)	(2.57)
Obs.	56,225	55,585	54,775	53,783

Notes: Structural breaks are based on the instrumented labor-income growth. When we use all observations we construct successive windows around period t. The numbers in parentheses are the robust standard errors. The asterisk(s) denote the statistical significance at three p-values: 10%, 5%, and 1%. We include the same set of controls as in our benchmark regression.

percentage points (16.35× 0.25). This is substantially larger than the corresponding OLS estimate of 0.72 percentage points.

Finally, Table 3 compares the estimates based on all years—i.e., estimates based on the volatility changes in every t, $\widehat{\Delta SD}_{i,t}$—to those based on the structural breaks. Naturally, using more observations improves the statistical significance, especially for short-term responses. However, our benchmark specification generates larger responses. For example, looking at a 4-year window ($k = 2$), when the standard deviation of labor-income growth doubles for the median worker, the risky share decreases by 3.40 percentage points (13.63 × 0.25), for the case where we use all t, compared with 4.09 percentage points based on the structural breaks (a 20 percent increase).

2.5 Robustness

We further examine whether our benchmark IV findings are robust with respect to different specifications and measurements. We discuss briefly the estimates of each robustness check and report the estimates in Table 10 in Appendix B. By and large, our baseline results are robust with respect to the following variations: employer changes, exclusion of small firms and managers from the sample, definitions of risky share, sample-selection criteria, additional

control variables (such as high-order polynomials of income and wealth, mortgage debt, and capital income), alternative measures of income volatility, and different subsamples for those always being married or being single, households' disposable income (as opposed to individual income), and others.

Staying with Same Employer: Switching firms is an endogenous choice, and it might be influenced by some unobserved factors that also drive portfolio adjustment. Therefore, we also estimate the portfolio responses for a subsample of workers who always stay with the same employer (roughly about 80 percent of the total sample). The empirical estimates are similar in magnitude and in some cases more pronounced relative to the benchmark.

Excluding Small Firms and Managers: One concern about the instrument's validity is that (i) in smaller firms individual workers can have some impact on the firm's performance, and (ii) some workers may have managerial positions that allow them to directly influence firms' decisions. Thus, as a robustness check, we exclude from our sample first, firms in the bottom quartile of the employment distribution and second, workers who are managers. The IV estimates retain their size and significance and in some cases become even more pronounced.

Alternative Measures of Firm Volatility: We also consider using value added/assets or sales/assets separately as instrumental variables. For both specifications we have negative estimates in the same range as in the benchmark specification, but the estimates are more noisy. We also experiment with alternative ways to measure volatility at the firm level (such as growth rate for net sales but not scaled by firm assets). The results are still consistent.

Alternative Control Variables for Nonlinear Wealth Effects: We also consider using high-order polynomials in income and wealth and their changes to possibly control some nonlinear wealth effects.

Industry/Occupation Controls: Also, as pointed our in the literature (e.g., Heaton and Lucas, 2000; Campbell, Cocco, Gomes, and Maenhout, 2001; Angerer and Lam, 2009), individual income (labor income or total income) may be correlated with aggregate stock returns, and this correlation may be different across industries and/or occupations. To see whether this has a big impact on our estimated response of risky share to income volatility, we additionally include two-digit industry and/or occupation dummies in our benchmark regression.

Further Sample Restrictions Another concern is that the responses may be driven by workers who have too little assets (or risky assets) or who experience a temporary unemployment spell. In the next robustness check, we exclude workers (i) who collect unemployment benefits, (ii) with financial assets in the bottom quintile, or (iii) with a risky share of less than 0.05. All of the restricted samples generate coefficients of similar magnitude

and statistical significance as in the benchmark specification.

Alternative Definition for $\hat{\tau}$: In the benchmark specification, we constructed $\widehat{\Delta SD}_{i,t}$, an arguably exogenous variation of earnings volatility. Based on the time series of $\widehat{\Delta SD}_{i,t}$'s, we identified the structural break for each individual, $\hat{\tau}$, and denoted the projected volatility change around $\hat{\tau}$ as $\widehat{\Delta SD}_{i,\hat{\tau}}$. It is worth analyzing the empirical responses if we still use the firm's volatility as an instrument but look around the original worker's structural break τ^*. Therefore, in this specification we use $\widehat{\Delta SD}_{i,\tau^*}$ instead of $\widehat{\Delta SD}_{i,\hat{\tau}}$. In this case, the coefficients are even more pronounced than in the benchmark, but more noisy for shorter horizons ($k < 3$).

Household Disposable Income: In our benchmark specification, we constructed income volatility using the individual's earnings growth. However, the decision to hold risky assets may be influenced by the volatility of household income, not only individual income. For example, if there is intra-household insurance, individual income volatility may not necessarily affect household savings. Therefore, we apply the structural-break approach based on total household-level disposable income instead of individual labor earnings. Once more, the estimates become more pronounced relative to the benchmark specification.

Marriage and Disposable Income: We also separately look at those individuals who are always married (from $-k$ to $+k$) and those who are always single to control for some potential influence from marital status. Results are robust for married workers but less statistically significant for singles.

2.6 Heterogeneity across Groups

We examine the response of the risky share to income volatility (IV estimation) across different groups (by age, education, and income growth). The results are collected in Table 11 in Appendix B.

By Age: Given that individual earnings and risky share exhibit clear life-cycle patterns (Guvenen, 2007 and Chang, Hong, and Karabarbounis, 2018), we ask whether our main finding is distinctive for a particular age group. The sample is classified into three groups: the young (younger than 40 years old), middle-aged (40 to 55), and old (older than 55) based on a worker's age in 2005. This splits the sample into 18 percent, 56 percent, and 26 percent, respectively.[20] Table 11 shows that for most horizons the young and the middle-aged exhibit the most significant and distinctive responses.

By Education: Workers with different levels of education may face different labor markets—e.g., stability of jobs, career paths, etc. Table 11 shows the estimated response

[20] In our sample for the benchmark regression, the mean age for the young group (middle-aged, old group) is about 36 (47, 58).

of the risky share separately for college graduates and high-school only graduates. While both groups show statistically significant responses, the college graduates exhibit sharper responses. We also examine whether knowledge about financial markets matters (Lusardi and Mitchell, 2007). Based on the Norwegian Education Registry, we classify the college graduates into those who majored in "Econ/Finance" (economics, business, and other related business/management majors) and the rest. Those who majored in economics/finance-related subjects do show stronger responses.

By Income Growth: We estimate the responses for different groups based on average income growth during the period (from $-k$ to $+k$). Based on income growth before and after the individual structural break year τ^* for a given time horizon, we classify workers into three groups: low growth (bottom 25 percent), high growth (top 25 percent) and the rest. In Table 11 we confirm that different groups show similar responses of the risky share to changes in income volatility.

3 Life-Cycle Model

3.1 Economic Environment

Demographics The economy is populated by a continuum of workers with total measure of one. A worker enters the labor market at age $j = 1$, retires at age j_R, and lives until age J. The decision to retire is exogenous. During each period the worker faces a probability of surviving s_j.

Preferences Each worker maximizes the time-separable discounted lifetime utility:

$$U = E \sum_{j=1}^{J} \delta^{j-1} (\mathbf{\Pi}_{t=1}^{j} s_t) \frac{{c_j}^{1-\gamma}}{1-\gamma}, \tag{3}$$

where δ is the discount factor, c_j is consumption in period j, and γ is the relative risk aversion.[21] For simplicity, we abstract from the labor effort choice and assume that labor supply is exogenous.

Labor-Income Profile We assume that the log earnings of a worker i with age j, $\log Y_{ij}$, is:

$$\log Y_{ij} = z_j + y_{ij} \quad \text{with} \quad y_{ij} = a_i + \beta_i \times j + x_{ij}. \tag{4}$$

[21] Alternative preferences have also been proposed in the literature analyzing portfolio choice. For example, Gomes and Michaelides (2005) use Epstein-Zin preferences with heterogeneity in both risk aversion and inter-temporal elasticity of substitution. Wachter and Yogo (2010) use non-homothetic preferences.

Log earnings consist of common (z_j) and individual-specific (y_{ij}) components. The common component, z_j, represents the average age-earnings profile, which is assumed to be the same across workers. The idiosyncratic component, y_{ij}, consists of an individual-specific profile, $a_i + \beta_i \times j$, which is constant along the life cycle, and stochastic shocks, x_{ij}, which follow an AR(1) process:

$$x_{ij} = \rho_x x_{i,j-1} + \nu_{ij}, \qquad \text{with} \quad \nu_{ij} \sim \text{ i.i.d. } N(0, \sigma^2_{ij}). \tag{5}$$

Note that the volatility of income shocks, σ^2_{ij}, is also idiosyncratic, and its stochastic process is described below.

Variance of Labor Income The idiosyncratic labor-income volatility is assumed to follow an AR(1) process:

$$\log(\sigma^2_{ij}) = (1 - \rho_\sigma) \log(\sigma^2_\nu) + \rho_\sigma \log(\sigma^2_{i,j-1}) + \zeta_{ij}, \qquad \text{with} \quad \zeta_{ij} \sim \text{ i.i.d. } N(0, \sigma^2_\zeta). \tag{6}$$

We use a log specification to ensure that income volatility is positive. Three parameters govern its dynamics: (i) σ^2_ν, which is the average variance of x, (ii) σ^2_ζ, which is the variance of the volatility shocks, and (iii) ρ_σ, which governs their persistence. We approximate the autoregressive process for the volatility shock using a Markov chain. In particular, we assume that the labor-income volatility takes N possible values (regimes): $\sigma^2 = \{\sigma^2_1, ..., \sigma^2_N\}$. The Markov chain is defined as $\Gamma(\sigma^2_j | \sigma^2_{j-1})$.

Savings There are two types of assets for savings: a risk-free bond, b, (paying a gross return of R in consumption units) and a stock, s, (paying $R_s = R + \mu + \eta$), where μ (> 0) represents the risk premium and η is the stochastic rate of return.[22] We denote the probability distribution of the stock realization by $\chi(\eta)$. Workers save for insuring themselves against labor-income volatility (precautionary savings) as well as for retirement (life-cycle savings). We allow workers to borrow using the risk-free bond ($b' \geq \underline{b}$), where $\underline{b}$ is the credit limit.

Tax System and Social Security The government performs two functions in the model. First, it taxes individual earnings Y_{ij} using the tax function $T(Y_{ij})$. We specify a flexible tax function based on Heathcote, Storesletten, and Violante (2014) that allows for transfers (see Section 4.1). Second, it runs a social security system. When a worker retires from the labor market at age j_R, the worker receives a social security benefit. To avoid the computational

[22] For simplicity, we abstract from the general equilibrium aspect by assuming exogenous average rates of return to both stocks and bonds.

complexity of tracking one more state variable (history of earnings), we make the social security benefit dependent on earnings received in the last working year before the exogenous retirement (Guvenen, 2007). The social security benefit of worker i is denoted by $ss(Y_{J_R-1})$, which is financed by the social security tax rate τ_{ss}.

Matrix for Income Process We define the matrices $\mathbf{M}_{j-1}$ and $\mathbf{H}_j$ that allow us to define the recursive problem in terms of income y, as follows:

$$\mathbf{M}_{j-1} = \begin{bmatrix} a \\ \beta \\ \rho_x x_{j-1} \end{bmatrix}, \quad \mathbf{H}_j = \begin{bmatrix} 1 \\ j \\ 1 \end{bmatrix}. \tag{7}$$

The following period's $\mathbf{M}_j$ is:

$$\mathbf{M}_j = \mathbf{R} \begin{bmatrix} \mathbf{M}_{j-1} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} (y_j - \mathbf{H}_j' \mathbf{M}_{j-1}) \end{bmatrix} \tag{8}$$

with $\mathbf{R}$ denoting a (3×3) matrix whose diagonal elements are $(1, 1, \rho_x)$. Note that, $\mathbf{H}_j' \mathbf{M}_{j-1}$ is the conditional expectation of period j's labor income as of age $j-1$. Moreover, $y_j - \mathbf{H}_j' \mathbf{M}_{j-1} = x_j - \rho_x x_{j-1} = \nu_j$, is the innovation of the shock to x. When the worker enters period j, log labor earnings y_j are drawn from a normal distribution F with mean $\mathbf{H}_j' \mathbf{M}_{j-1}$ and variance σ_j^2 (denoted as $F(y_j \mid \mathbf{H}_j' \mathbf{M}_{j-1}, \sigma_j^2)$).

Benchmark: Perfect Information about Volatility Shocks In the benchmark model, workers have perfect information about their individual labor-income volatility (perfect information model or PIM, henceforth). In particular, workers enter age j observing volatility σ_j^2 and forming expectations about the next period's volatility based on the law of motion $\Gamma(\sigma_{j+1}^2 | \sigma_j^2)$.

Imperfect Information about Volatility Shocks The second case is one where workers have imperfect information about their individual income volatility (imperfect information model or IIM, henceforth). In this case, workers enter age j with a prior probability $\boldsymbol{\pi}_{j|j-1} = \{\pi_{j|j-1}^g\}_{g=1}^N$ for each possible regime g with $\sum_g \pi_{j|j-1}^g = 1$ (in the PIM, the prior is a just degenerate at the true regime). They form a posterior belief for each regime $\boldsymbol{\pi}_{j|j} = \{\pi_{j|j}^g\}_{g=1}^N$ based on the Bayes rule. In particular, workers compute the probability that a particular regime g is currently active given the available information $\{y, \mathbf{M}_{j-1}\}$. As a result, the posterior beliefs

are given by:

$$\pi_{j|j}(\sigma_g^2 \mid y_j, \mathbf{H}_j'\mathbf{M}_{j-1}) = \frac{F(y_j \mid \mathbf{H}_j'\mathbf{M}_{j-1}, \sigma_g^2) \times \pi_{j|j-1}^g(\sigma_g^2)}{\sum_{h=1}^N F(y_j \mid \mathbf{H}_j'\mathbf{M}_{j-1}, \sigma_h^2) \times \pi_{j|j-1}^h(\sigma_h^2)}, \tag{9}$$

where $F(y_j \mid \mathbf{H}_j'\mathbf{M}_{j-1}, \sigma_g^2)$ is the probability that labor-income realization, y_j, is observed given that the last year's labor income is $\mathbf{H}_j'\mathbf{M}_{j-1}$ and that the current volatility regime is σ_g^2. If the absolute value of $y_j - \mathbf{H}_j'\mathbf{M}_{j-1}$ (the innovation ν_j) is small, the worker places a larger probability on the low-volatility regimes and vice versa. Given the posterior probabilities, the worker forms the next period priors:

$$\pi_{j+1}(\sigma_g^2) = \sum_{h=1}^{N} \Gamma(\sigma_g^2 \mid \sigma_h^2) \times \pi_{j|j}(\sigma_h^2). \tag{10}$$

Note that in both perfect and imperfect information models, workers know the law of motion (transition probability) for the volatility regime, Γ. What is different in the two cases is the initial regime. Under perfect information, workers know the true regime, while under imperfect information, workers have a probability distribution over the possible regimes.

Value Functions We collapse financial wealth into one variable, "total financial wealth," $W = bR + sR_s$. Then, the state variables include workers' wealth (W), current income (y_j), the expected income ($\mathbf{M}_{j-1}$), and the prior probability about the current volatility regime, $\boldsymbol{\pi}_{j|j-1}$. The value function of a worker at age j is:

$$V_j(W, y_j, \mathbf{M}_{j-1}, \boldsymbol{\pi}_{j|j-1}) = \max_{c,s',b'} \left\{ \frac{c_j^{1-\gamma}}{1-\gamma} + \delta s_j \sum_g \int_{\eta'} \int_{y_{j+1}} \pi_{j+1}(\sigma_g^2) V_{j+1}(W', y_{j+1}, \mathbf{M}_j, \boldsymbol{\pi}_{j+1|j}) dF(y_{j+1}|\mathbf{H}_{j+1}'\mathbf{M}_j, \sigma_g^2) d\chi(\eta') \right\}$$

s.t. $c + s' + b' = [(1-\tau_{ss})Y_j - T(Y_j)] \times \mathbf{1}\{j < j_R\} + ss(Y_{J_R-1}) \times \mathbf{1}\{j \geq j_R\} + W$

PIM: $\pi_{j+1}(\sigma_g^2)$ is based on law of motion $\Gamma(\sigma_g^2|\sigma_j^2)$

IIM: $\pi_{j+1}(\sigma_g^2)$ is given by Equations (9) and (10)

$\mathbf{M}_j$ is given by Equation (8)

$F(y_{j+1}|\mathbf{H}_{j+1}'\mathbf{M}_j, \sigma_g^2)$ is the prob. distribution for the next period income given $\mathbf{M}_j, \sigma_g^2$

$b' \geq \underline{b}$ and $s' \geq 0$,

where $\mathbf{1}\{\cdot\}$ is an indicator function and total labor income is $Y_j = e^{z_j + y_j}$.

4 Estimation

There are several sets of parameters to pin down: (i) life-cycle parameters $\{j_R, J, s_j\}$, (ii) preferences $\{\gamma, \delta\}$, (iii) asset-market parameters $\{R, \mu, \sigma_\eta^2, \underline{b}\}$, (iv) labor-income process $\{z_j, \sigma_a^2, \sigma_\beta^2, \rho_x, \rho_\sigma, \sigma_\nu^2, \sigma_\zeta^2\}$, and (v) tax and transfers $\{\tau_1, \tau_2, \tau^*, \tau_{ss}, ss\}$. One set of parameters is calibrated directly from the data or the existing literature. The remaining parameters are estimated using indirect inference.

4.1 Calibrated Parameters

Table 4 gives the list of calibrated parameters. The model period is a year. Workers are born and enter the labor market at $j = 1$ and live for 80 periods, $J = 80$. This life cycle corresponds to ages 21 to 100. Workers retire at $j_R = 45$ (age 65) when they start receiving the social security benefit. We estimate the survival probability $\{s_j\}$ at each age using the data on mortality rates from Statistics Norway.

Table 4: Calibrated Parameters

Parameter	Variable	Value	Target / Source
Life Cycle	J	80	–
Retirement Age	j_R	45	–
Risk-free Rate	$R-1$	1.43%	Klovland (2004)
Equity-Risk Premium	μ	3.14%	Dimson, Marsh, and Staunton (2008)
Stock-Return Volatility	σ_η	23.8%	Dimson, Marsh, and Staunton (2008)
Social Security Benefit	ss	–	Statistics Norway
Tax Function	τ_1	0.73	Statistics Norway
Tax Function	τ_2	0.16	Statistics Norway
Tax Function	τ^*	0.85	Statistics Norway
Tax Function	Y^*	1.7	Statistics Norway
Survival Probability	$\{s_j\}$	–	Statistics Norway
Common Age-Earnings Profile	$\{z_j\}$	–	OECD
Number of Volatility Regimes	N	7	–

According to Dimson, Marsh, and Staunton (2008), the annualized real returns to equity for Norway from 1900-2005 were 4.28 percent. We follow Fagereng, Gottlieb, and Guiso (2017) and adjust the returns to reflect an 80 percent bias of Norwegian investors toward domestic over foreign stocks. Since the world average returns were 5.75 percent, according

to Dimson, Marsh, and Staunton (2008) for the same period, we set the rate of returns to equity at 4.57 percent. Using the estimates from Klovland (2004), we calibrate a real return for the risk-free rate to 1.43 percent. Therefore, the equity premium in our model μ is 3.14 percent. The standard deviation of the innovations to the rate of return to stocks, σ_η, is 23.8 percent, also computed using a weighted average of the standard deviation of Norwegian stocks and of foreign stocks, which are 26 percent and 17 percent, respectively, based on Dimson, Marsh, and Staunton (2008). We assume that the stock returns are orthogonal to labor-income risks.[23]

To compute the amount of tax and transfers, we use the following specification:

$$T(Y) = Y - \tau_1 Y^{1-\tau_2} + \mathbf{1}_{\{Y^* > Y\}} \tau^* (Y - Y^*).$$

A version of this type of tax function has recently been used to analyze tax and transfers in the U.S. (Heathcote, Storesletten, and Violante, 2017). In particular, parameter τ_1 captures the average tax rate in the economy and parameter τ_2 the degree of progressivity of the schedule. As seen in the left panel of Figure 3, the tax system in Norway becomes very progressive for income levels around twice the average labor income. To capture the high progressivity of the Norwegian tax system, we add the term $\mathbf{1}_{\{Y^* > Y\}} \tau^* (Y - Y^*)$. With our detailed administrative data, we can calibrate all parameters using information on before- and after-tax labor earnings. The before-tax earnings are cash salary, while after-tax earnings are before-tax earnings net of taxes and transfers. Transfers include unemployment benefits, sickness benefits, money received in government activity programs, and disability benefits. The left panel of Figure 3 shows that our model matches well the relationship between before- and after-tax individual labor income.

The social security benefit is calibrated to replicate the average benefit for each labor-income decile we observe in the data (right panel of Figure 3). As mentioned, in the model we condition the social security benefit on the earnings received in the last working year before retirement. For consistency, in the data we find the relationship between social security benefits and labor income during ages 60 to 65. We find that a worker with the mean labor income during ages 60 to 65 receives a benefit equal to 36 percent of his/her pre-retirement labor income. A worker with twice the mean labor income during ages 60 to 65 receives around 55 percent of pre-retirement labor income.

We calibrate the common age profile of income (z_j) based on the age profile of real wages

[23] In our data, the correlation between stock market return and average real wage (using aggregate data from national accounts) is small and equal to -0.08 with a standard deviation of 0.16. These numbers are similar to the numbers reported in Heaton and Lucas (2000) for the U.S. In other studies that have used U.S. data, Davis and Willen (2000) find a small, positive correlation, while Campbell, Cocco, Gomes, and Maenhout (2001) find a positive correlation only for specific population groups.

in Norway from the OECD. The real wages for 30-, 40-, and 50-year old workers are on average approximately 5, 15, and 20 percent higher, respectively, than those of 25-year-old workers. Finally, we assume that there are 7 regimes for income volatility: $N = 7$.

Figure 3: Tax and Social Security System: Model vs. Data

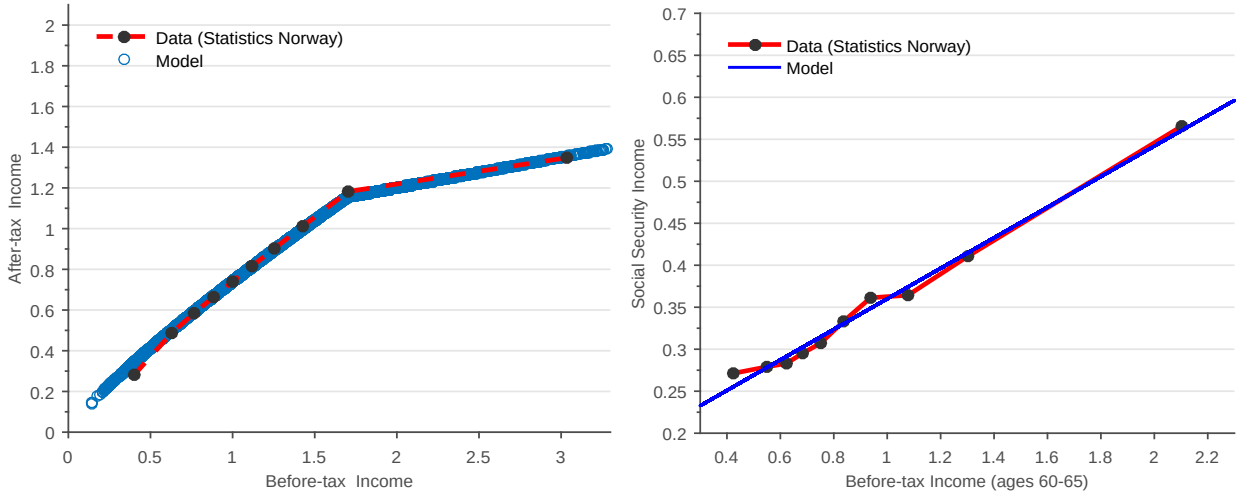

Notes: The left panel shows the relationship between before- and after-tax labor income for the model and the data. The right panel shows the relationship between before-tax labor income for ages 60 to 65 and the social security benefit in the model and in the data. We normalize labor-income data by the average earnings, 553,414 NOK. Data are from Statistics Norway and authors' calculations.

4.2 Structural Estimation

We use indirect inference to estimate the remaining parameters by minimizing the distance between the model statistics and their empirical counterparts.[24] We need to estimate a total of 9 parameters represented by the vector:

$$\Theta = [\delta, \gamma, \underline{b}, \sigma_a^2, \sigma_\beta^2, \rho_x, \rho_\sigma, \sigma_\nu^2, \sigma_\zeta^2].$$

Let $M^m(\Theta)$ denote the set of model-generated moments and M^d their empirical counterparts. The moments used in the estimation are

1. Average assets-income ratio.

2. Average risky assets-total assets ratio.

[24] According to Guvenen and Smith (2014), when the estimation uses any type of statistical association as information, even if that requires specifying an auxiliary model (as the IV estimation in our case), the estimation falls under the classification of indirect inference. In contrast, with simulated method of moments, the statistics used in the estimation are directly computed from the data.

3. Average consumer debt-income ratio.

4. Variance-covariance matrix of log labor income across ages (a total of 441 moments).

5. Life-cycle profile of the cross-sectional kurtosis of labor-income growth.

6. Dispersion of volatility change (projected): $\widehat{\Delta SD}$.

7. Life-cycle profile of the cross-sectional variance of log consumption.[25]

8. Response of risky share to volatility (IV estimate for $k = 2$ in Table 1).

In total we have 511 moments. The estimator minimizes the loss function:

$$\min L_\Theta = [M^d - M^m(\Theta)]' W [M^d - M^m(\Theta)]. \tag{11}$$

We pick a weighting matrix W (freely) to focus on some specific dimensions of the data we view as more important to match. In particular, we place more weight on the two moments that describe the portfolio behavior: the average risky share and the response of the risky share to volatility. We also place more weight on moments that describe the dynamics of labor income, such as the cross-sectional variance of log labor income over the life cycle and the dispersion in volatility changes. Appendix D provides the weights used in the estimation.

It is useful to connect some parameters with the moments that are the most informative about their values. For example, the discount factor δ is identified primarily through the average financial assets to income ratio. In our data, the average financial assets for ages 25 to 60 are 1,212,000 NOK, and the average household income is 553,414 NOK. Therefore, we estimate the discount factor δ to match an asset-to-income ratio of 2.19. The risk aversion γ is identified primarily through the risky to total assets ratio. Since the average financial assets in risky accounts are 690,830 NOK, we target an aggregate risky share of 0.57.[26] According to the data from the Bank of Norway, the credit card debt accounts for 3 percent of total debt (which averages at 908,587 NOK). Therefore, the average credit card debt to income ratio is $3\% \times 908,587/553,414 = 4.9\%$. The debt-to-income ratio is crucial to the value of borrowing constraint $\underline{b}$.

[25] We thank Martin Holm from the University of Oslo for providing statistics for this moment. In our data set, we could not have reliable estimates on individual consumption since we do not have access to those data required for constructing consumption (transaction data on housing prices; housing physical characteristics; detailed information on capital gains; and so on). See Eika, Mogstad, and Vestad (2017) among others.

[26] The average risky share in the Norwegian panel is a much lower value of 0.31. As is well-known, it is hard to match such low values of the average risky share unless we resort to unrealistically high degrees of risk aversion or highly risky events such as a stock market crash or long-term unemployment spell. Instead, we base our estimation on a more feasible target, the aggregate risky share.

For the estimation of parameters $\{\sigma_a^2, \sigma_\beta^2, \rho_x, \sigma_\nu^2\}$ that govern the income process, we use the variance-covariance matrix of log labor earnings similar to Guvenen (2009). For the estimation of the stochastic process of volatility, $\{\rho_\sigma, \sigma_\zeta^2\}$, we use the life-cycle profile of the cross-sectional kurtosis of income growth and the dispersion of the volatility change $\widehat{\Delta SD}$. (See Appendix C for the construction of the life-cycle profile of kurtosis.) We discuss the identification of all income parameters in more detail in Section 4.4, below.

Our estimation takes into account important economic choices. The first regards the variance of log consumption along the life cycle. According to Guvenen and Smith (2014), consumption inequality is informative about (i) the presence of borrowing constraints and (ii) the lack of prior information about the income profile. Here, for simplicity, we choose not to assume imperfect information about the *level* component of the income profile because (i) it simplifies our analysis, which is mainly focused on learning about the volatility shock, and (ii) the imperfect information on the growth component does not necessarily improve the empirical fit of the model.[27] This is in contrast to the imperfect information about the *volatility*, which, as we show below, is necessary to match the data better.

Second, the response of portfolio choice to income volatility is a key moment. More specifically, the response coefficient β in Equation (1) is a part of our target moments (the IV estimates in Table 1). Obviously, the portfolio response depends on several parameters, such as risk aversion (γ), the borrowing constraints ($\underline{b}$), and the persistence of the volatility shocks (ρ_σ). More importantly, we show that portfolio choice is useful to sharply discriminate between alternative models of information about the income volatility (PIM and IIM).

4.3 Estimation Results and Model Fit

We estimate two versions of the model: the benchmark case of perfect information about the volatility regime (PIM) and imperfect information (IIM). We also estimate the parameters of the income process separately using income moments only (in which case the distinction between PIM and IIM does not matter). Table 5 shows the estimation results, while Figure 4 and Table 6 show the models' fit to the empirical targets.

For PIM, the variance of the fixed effect component is $\sigma_a^2 = 0.050$, the variance of the slope $\sigma_\beta^2 = 0.0089\%$, and the average variance of the idiosyncratic shocks $\sigma_\nu^2 = 0.026$. The

[27] Guvenen and Smith (2014) estimate the amount of prior information about the growth component of the income profile on labor market entry using the dispersion of labor income and consumption based on U.S data. The imperfect information about the level component helps the model to match the fast-increasing variance of (log) consumption over the life cycle in the U.S. data (10 log points). But in the Norwegian panel, the profile of the variance of consumption is much flatter (an increase of 4 log points). As a result, imperfect information about the level component is not crucial for matching the data. Moreover, introducing imperfect information on both components of the income profile makes the model highly complicated.

Table 5: Parameter Estimates

Moments Used in Estimation		All Moments		Income Only
Model Specification		PIM	IIM	–
Variance of Fixed Component	σ_a^2	0.050	0.054	0.053
		(0.003)	(0.004)	(0.004)
Variance of Growth Component	$\sigma_\beta^2 \times 100$	0.0089	0.0085	0.0081
		(0.0005)	(0.0006)	(0.0005)
Persistence of Level Shocks	ρ_x	0.783	0.788	0.783
		(0.012)	(0.016)	(0.015)
Variance of Level Shocks	σ_ν^2	0.026	0.023	0.023
		(0.002)	(0.003)	(0.003)
Persistence of Volatility Shocks	ρ_σ	0.936	0.931	0.933
		(0.006)	(0.006)	(0.005)
Variance of Volatility Shocks	σ_ζ^2	0.071	0.071	0.070
		(0.011)	(0.010)	(0.010)
Discount Factor	δ	0.919	0.923	
		(0.033)	(0.025)	
Risk Aversion	γ	5.44	4.86	
		(1.21)	(1.15)	
Credit Limit	$\underline{b}$	-0.144	-0.114	
		(0.052)	(0.039)	

Notes: Estimation results for model parameters. "All Moments" uses all available moments (income, financial wealth, consumption, and portfolio). "Income Only" uses income moments only. The standard errors for each parameter are reported in parentheses.

persistence of the idiosyncratic shocks is $\rho_x = 0.783$.

The upper left panel of Figure 4 plots the cross-sectional variance of log earnings across ages net of cohort effects. In the data, the variance of income initially decreases during ages 25 to 30, a pattern we cannot generate from the model. After age 30, we see a familiar increasing variance—also well documented in the U.S. data (see, for example, Storesletten, Telmer, and Yaron, 2004; Guvenen and Smith, 2014; Heathcote, Storesletten, and Violante, 2014). In Norway, it increases by about 10 log points over ages 30 to 55, while in the U.S. the increase is about 20-25 log points. The PIM captures this increasing profile fairly well after age 30.

In both our estimation and Guvenen and Smith (2014), who estimated these parameters using the U.S. data, the idiosyncratic growth component is fairly large, implying a mildly

Figure 4: Model Fit: Income, Consumption and Portfolio Response

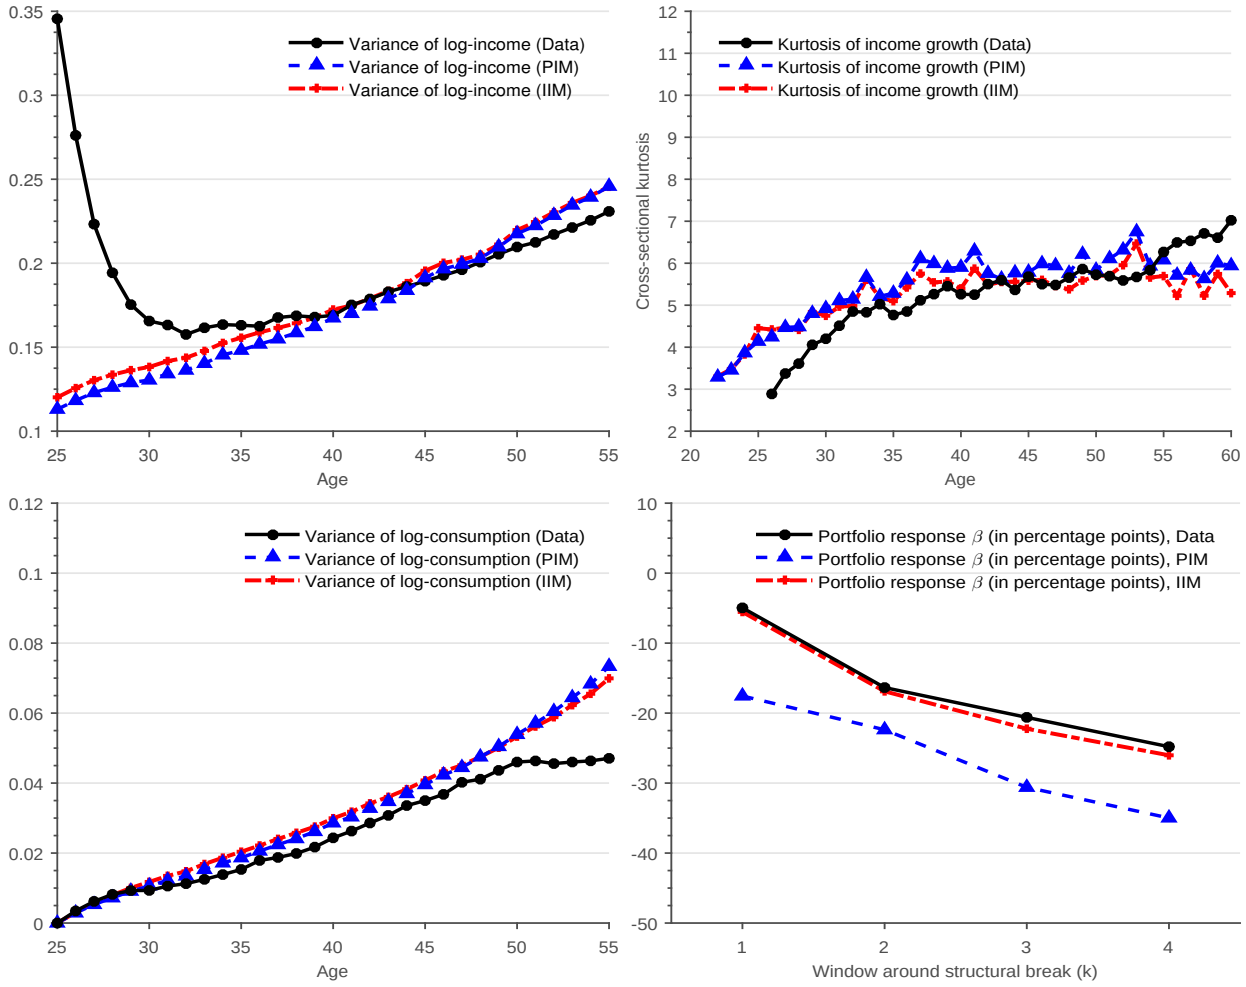

Notes: Upper-left panel shows the variance of log earnings over the life cycle. Upper-right panel is the cross-sectional kurtosis of income growth. Lower-left panel is the variance of log consumption. Lower-right panel shows the regression coefficients of response of the risky share in Equation (1) for $k = 1$ to 4.

persistent process for the shock to the income level. In particular, we find an almost identical value for ρ_x (0.783 versus 0.789 in the U.S. data). One difference is that in our Norwegian data, the variance of the σ_ν^2 is half of what Guvenen and Smith (2014) estimate for the U.S. data, which reflects the sharper increase in the variance of log labor income over the life cycle, in the U.S. relative to Norway.

We next turn to the estimated parameters that govern the process of income volatility. For PIM, the estimated variance is $\sigma_\zeta^2 = 0.071$. This parameter is identified mainly by the cross-sectional dispersion of changes in income volatility ΔSD. The estimated volatility

shocks are highly persistent: $\rho_\sigma = 0.936$.[28] The highly persistent process reflects the life-cycle profile of kurtosis in the data. Looking at the upper-right panel of Figure 4, the model captures the increasing (and slightly concave) shape of kurtosis over the life cycle. Both the size and the persistence of the volatility shocks are precisely estimated with t-statistics higher than 10. This suggests that the moments we picked to identify the volatility process are very informative. We give more details about the link between the empirical moments and the parameters of the volatility process in Section 4.4.

The IIM generates broadly similar estimates for the income parameters. This occurs because income parameters are primarily identified from the income data (the variance-covariance matrix of realized incomes, the cross-sectional kurtosis, etc.) and are not heavily influenced by economic choices (financial wealth, consumption inequality, portfolio, etc.). We verify this by estimating the parameters of income process using the income-related moments only (the last column of Table 5). The estimated parameters are fairly similar to the case where we use all moments.

Table 6: Model Fit: Selected Statistics

	Data	PIM	IIM
Financial Assets / Income	2.19	2.21	2.22
Risky Assets / Financial Assets	0.57	0.57	0.56
Credit Card Debt / Income	4.9%	6.2%	5.4%
Standard Deviation of ΔSD	0.13	0.12	0.12
Response of Risky Share $\hat{\beta}$ $(k = 2)$	-16.3	-22.3	-16.8

Notes: We report selected statistics based on two models (PIM and IIM) and corresponding empirical targets based on Statistics Norway.

The discount factor in the PIM is $\delta = 0.919$, while in the IIM, it is 0.923. The estimated credit limit ($\underline{b}$) is -14.4 percent and -11.4 percent of annual income, in PIM and IIM, respectively. The estimate of credit limits is influenced by two sets of statistics from the data. First, the average debt-to-income ratio. Second, the increase in the cross-sectional variance of log consumption. Intuitively, if borrowing constraints are tight, consumption inequality increases in parallel to income inequality. In the data, the variance of log consumption increases by about 4 log-points (from ages 25 to 50), which is matched fairly well by both models (lower-left panel of Figure 4).

[28] The average number of volatility changes implied by our estimated values for ρ_σ and σ_ζ^2 (and given that we set $N = 7$ grid points) is 7.8 in our simulation.

Where the two models (PIM and IIM) diverge substantially is with respect to the estimate of risk aversion (5.44 versus 4.86 in PIM and IIM, respectively) as well as their implications for portfolio choice. The lower-right panel of Figure 4 plots the regression coefficients of the risky share to volatility in Equation (1) using model-generated data.[29] The model with perfect information generates responses much larger (in absolute terms) than those we find in the data. For example, if the standard deviation of labor-income growth increases by one unit, the risky share decreases by 22 and 34 percentage points, for $k = 2$ and $k = 4$, respectively. The corresponding estimates in the data are 16 and 24 percentage points, respectively. Therefore, even though parameter estimation in PIM can match well most of the income and financial moments (see Table 6), it fails to match the marginal response of the portfolio to income volatility.

On the other hand, parameter estimation in IIM can naturally deliver the marginal response of the portfolio while still matching the other moments. When the standard deviation of income growth increases by one unit, the risky share decreases by 12 and 26 percentage points, for $k = 2$ and $k = 4$, respectively, which aligns well with the data. The key mechanism is a gradual learning about the change in volatility regime that mitigates the response of portfolio choice. The discrepancy is more visible at shorter time horizons because workers gradually learn about the change in the volatility regime. To confirm that imperfect information can naturally match the portfolio response, we also consider an estimation where the marginal response is *not* targeted. Even in this case, the model matches the portfolio response very closely, suggesting that the improvement in the fit does not come at the expense of other moments. One might think that a low value of risk aversion in IIM is crucial for the different estimated response of the risky share. To explore this possibility, we simulate PIM under the same risk aversion as the estimate from IIM ($\gamma = 4.86$). See Appendix E for the estimated responses. The response of the risky share in PIM continues to substantially overpredict the response in the data.

4.4 Identification of the Income Process

We discuss the estimation of the stochastic process of labor income in detail. The discussion is concentrated on the parameters regarding the volatility (i.e., the second moment) shocks $\{\sigma_\zeta^2, \rho_\sigma\}$. The other parameters are discussed briefly because their estimation has been extensively analyzed in the literature (e.g., Storesletten, Telmer, and Yaron, 2004; Guvenen,

[29] Similar to the empirical analysis, the regression includes a set of individual controls, namely, changes (between $\tau^* - k$ and $\tau^* + k$) in income and wealth, as well as the levels of income, wealth, and age at τ^*. We exclude outliers by keeping observations between the 1st and 99th percentiles of the distribution of the risky share.

Figure 5: Identification: Variance of Volatility Shocks σ_ζ^2

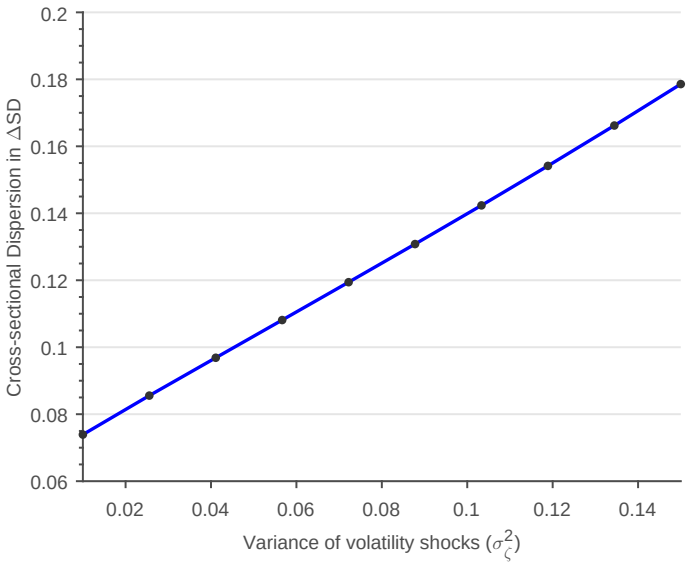

Notes: The x-axis is the variance of volatility shocks (σ_ζ^2), and the y-axis is the cross-sectional standard deviation of the volatility changes ΔSD from the model-generated data.

2009; Guvenen, Karahan, Ozkan, and Song, 2015).

Variance of Volatility Shocks σ_ζ^2 The variance of the volatility shocks is identified mainly by the cross-sectional dispersion of exogenous changes in volatility (i.e., instrumented ΔSD). When σ_ζ^2 increases, the dispersion of ΔSD increases as more workers draw income shocks from a wider distribution. Figure 5 plots the model-generated dispersion of ΔSD for a wide range of σ_ζ^2. The model produces a monotonically increasing relationship between σ_ζ^2 and ΔSD. Our estimation exploits this relationship to identify the variance of volatility shocks σ_ζ^2.

Persistence of Volatility Shocks ρ_σ The persistence of volatility shocks is identified by the life-cycle profile of the cross-sectional kurtosis of earnings growth. The transitory process of volatility implies that the kurtosis quickly approaches its long-run value (thus a flat life-cycle profile except for the beginning), whereas highly persistent shocks to volatility imply that the kurtosis gradually converges to its long-run value over time (an increasing profile).

We explain the identification of ρ_σ by setting up a simple example. We assume there is no profile heterogeneity ($\sigma_a^2 = 0, \sigma_\beta^2 = 0$). Therefore, labor-income variations are associated with the stochastic component x only. We further assume that income volatility can take

Figure 6: Identification: Persistence in Volatility Shocks: ρ_σ

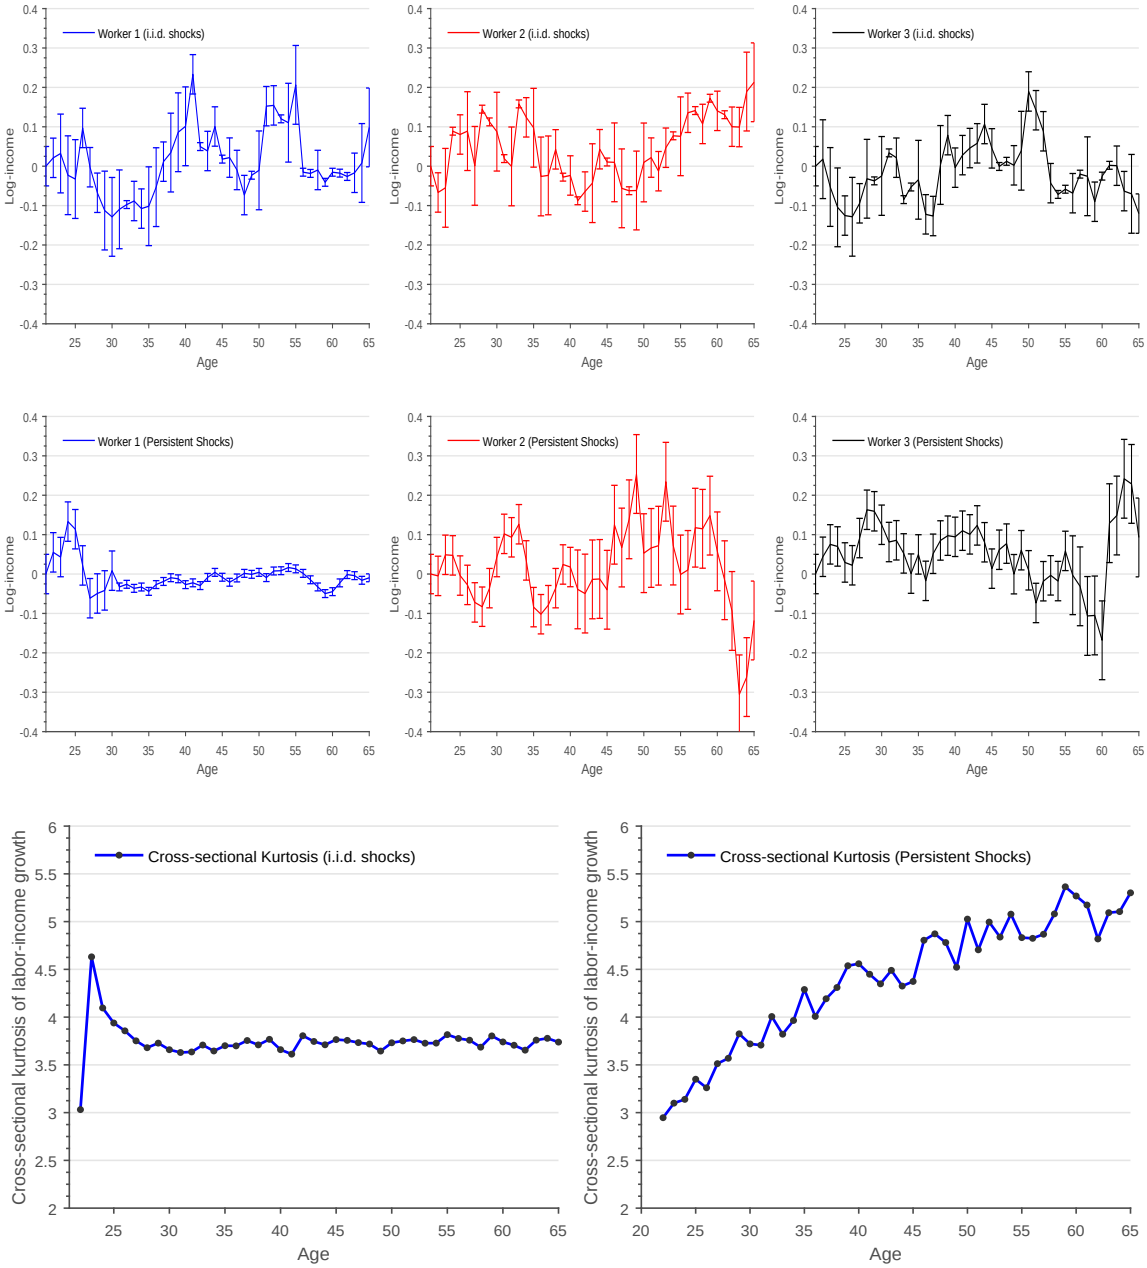

Notes: The upper and middle panels show the evolution of log labor income for three workers from the model-generated data. The upper panels correspond to i.i.d. shocks to income volatility, whereas the middle panels correspond to persistent volatility shocks. The bottom panels show the life-cycle profile of the cross-sectional kurtosis of income growth when shocks are i.i.d (left panel) and persistent (right panel), respectively.

three values $\sigma_\zeta^2 = \{0.1\%, 5\%, 10\%\}$. Let's compare two extreme cases: (i) an i.i.d. volatility regime where the transition matrix $\Gamma_{ij} = 1/3, \forall i, j$ and (ii) a highly persistent stochastic process of volatility where $\Gamma_{ij} = 0.98$ if $i = j$, and $\Gamma_{ij} = 0.01$ otherwise.

The upper panels of Figure 6 show the log labor income (as well as the volatility bands) for three workers with i.i.d. volatility shocks. In each period, workers receive a shock that is drawn from a distribution whose variance is independent of the last period's volatility. As a result, starting from the beginning of the life cycle, workers jump around the distributions. The middle panels of Figure 6 show log labor income for three workers when the volatility follows a highly persistent stochastic process described above. At the beginning of the life cycle, workers draw shocks from the same middle distribution ($\sigma_\zeta^2 = 5\%$). As the life cycle advances, the volatility of income decreases for worker 1 at age 30. Consequently, labor income becomes almost constant after the break. For workers 2 and 3, the volatility increases at ages 40 and 55, respectively, resulting in a more volatile labor-income path.

The bottom panels of Figure 6 show the cross-sectional kurtosis of labor-income growth for both cases of the stochastic regime changes: (i) i.i.d. in the left panel and (ii) persistent regime changes in the right panel. In both cases, the average kurtosis is higher than 3 (which is the value of kurtosis if the distribution is normal) implying a leptokurtic cross-sectional distribution of income growth. Indeed, Guvenen, Karahan, Ozkan, and Song (2015) show that the labor-income process with volatility shocks gives rise to a leptokurtic distribution. Intuitively speaking, earnings are not drawn from a single normal distribution. Instead earnings are drawn from a mixture of normals generating a cross-sectional distribution that looks leptokurtic. We argue that it is not only the average kurtosis but also the life-cycle profile of kurtosis that conveys useful information about the stochastic process of the income's volatility regime—especially with respect to its persistence.

When volatility shocks are i.i.d., workers are distributed across the different normal distributions from the beginning of the life cycle. Although workers jump around distributions every period, the measure of workers allocated to each distribution remains invariant. As a result, the life-cycle profile of the cross-sectional kurtosis in labor-income growth is constant (left panel of Figure 6). When volatility shocks are persistent, kurtosis experiences a gradual and linear increase over time (right panel of Figure 6). Workers start their life cycle drawing shocks from the median distribution. As the life cycle advances, some workers draw from a narrower distribution (decrease in volatility) and some draw from a wider distribution (increase in volatility). Due to these transitions, income growth becomes increasingly leptokurtic. The increasing (and slightly concave) age profile of kurtosis in the data results in a persistent process of volatility.

Note that our identification strategy for the persistence of volatility, ρ_σ, mirrors the

identification strategy for the persistence of labor-income shocks, ρ_x. As Storesletten, Telmer, and Yaron (2004) show, higher ρ_x leads to a linearly increasing profile in the *second* moment of labor income. In our case, higher persistence in the volatility shocks ρ_σ generates a linearly increasing profile in the *fourth* moment of labor income, e.g. the kurtosis.

Remaining Parameters for the Income Process $\{\sigma_a^2, \sigma_\beta^2, \rho_x, \sigma_\nu^2\}$ The identification of the remaining parameters for the income process is similar to that in Guvenen (2009). We discriminate between the heterogeneous income profile ($\sigma_\beta^2 > 0$) and the homogeneous income profile ($\sigma_\beta^2 = 0$). In the latter case, the age profile of income inequality should be accounted for by the stochastic shocks only; thus the estimation yields a substantially more persistent stochastic process of x.

Just for the purposes of presentation, we assume that there is no stochastic volatility. As outlined in Section 3, the (log) earnings of worker i with age j is:

$$y_{ij} = a_i + \beta_i \times j + x_{ij},$$

where a_i is the fixed effect, β_i is the individual-specific slope, and $x_{ij} = \rho_x x_{i,j-1} + \nu_{ij}$ is the stochastic component that follows an AR(1) process with $\nu_j \sim N(0, \sigma^2)$. For a worker who is h years older:

$$y_{i,j+h} = a_i + \beta_i(j+h) + x_{i,j+h}.$$

The variance and covariance of earnings across ages help us to uncover the underlying process. Specifically,

$$\begin{aligned} Cov(y_j, y_j) &= \sigma_a^2 + 2\sigma_{a\beta} j + \sigma_\beta^2 j^2 + V(x_j) \\ Cov(y_j, y_{j+h}) &= \sigma_a^2 + 2\sigma_{a\beta}[2j+h] + \sigma_\beta^2 j(j+h) + \rho_x^h V(x_j) \end{aligned}$$

with

$$V(x_j) = \begin{cases} 0, & j = 1 \\ \sum_{h=0}^{j-1} \rho_x^{2h} \sigma_\nu^2, & j > 1. \end{cases}$$

The variance of earnings at age j is driven by two components: (i) $\sigma_\beta^2 j^2$, which is convex in age j, and (ii) $V(x_j) = \sigma_\nu^2 \sum_{h=0}^{j-1} \rho^{2h}$, which is concave in age j (as long as $\rho < 1$). The shape of the cross-sectional variance along the life cycle is informative about which component is stronger. The autocovariance of labor income at age j also depends on two components: (i) $\sigma_\beta^2 j(j+h)$ is increasing in the time horizon h, and (ii) the autocovariance should decay monotonically at rate ρ_x due to component $\rho_x^h V(x_j)$. Absent the profile heterogeneity, the autocovariances should decay monotonically from age j onward. In sum, the variance-

covariance matrix of labor income from our administrative data identifies $\{\sigma_a^2, \sigma_\beta^2, \rho_x, \sigma_\nu^2\}$. Note that a common challenge in this identification method is the lack of a long panel, a challenge we do not face because we have about 20 years of observations.

5 Discussion on Stochastic Volatility

In this section, we illustrate how the marginal response of the risky share to income volatility is useful to understand: (i) the role of imperfect information, (ii) the welfare cost of income volatility, (iii) the persistence of the volatility shocks, and (iv) to distinguish our benchmark model from other well-known models of stochastic volatility such as ARCH and GARCH.

Role of Imperfect Information Our structural model exhibits significantly different portfolio responses between PIM and IIM. To illustrate the importance of imperfect information, we consider a simple example. Suppose that in our model economy all workers start with the median volatility and experience a single structural change in the volatility regime at age 43 (i.e., $\tau^* = 43$). For one-half of the workers, the volatility increases; for the rest of the workers, the volatility decreases. The volatility change is assumed to be unexpected and permanent (until the workers retire from the labor market).

Based on our simulated data, we run the following regression (used in our empirical analysis) for $t = \{\tau^* - 3, .., \tau^* + 4\}$:

$$\mathrm{RS}_{i,t} - \mathrm{RS}_{i,\tau^*-4} \;=\; a \;+\beta_t \Delta SD_{i,\tau^*} \;+X_{i,t} + \;\epsilon_{i,t}, \tag{12}$$

where $X_{i,t}$ is a list of control variables (age, wealth, etc.). We plot the marginal responses (coefficients β_t) in Figure 7. Before the structural break, the risky share differs only due to time-varying worker characteristics. Therefore, controlling for characteristics, the marginal responses are approximately equal to zero. The marginal responses turn negative upon the structural break, reflecting a decrease in the risky share for workers who experience an increase in income volatility. The negative relationship between the risky share and income volatility confirms the standard intuition that labor-income risk crowds out financial risk. With imperfect information, the marginal responses are also negative. But the magnitude is smaller and slowly increases as workers gradually learn about the shift in the regime.

Welfare Cost of Income Volatility One advantage of the structural model is that it allows us to examine the welfare cost of background income risk. Here, we study two types of welfare cost of volatile income. First, consumption becomes more volatile because of the

Figure 7: Response of Risky Share: Perfect vs. Imperfect Information

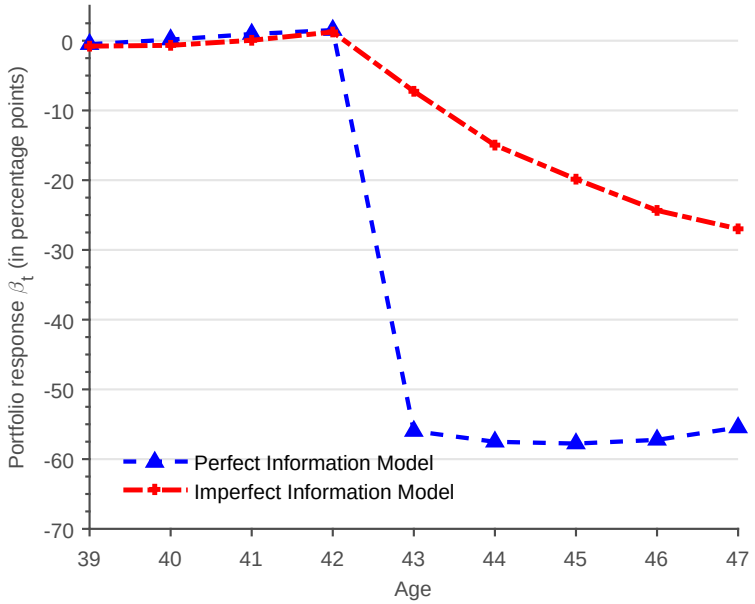

Notes: The figure shows the estimated response—the regression coefficients from (12)—based on a simple example where one half of workers experience an increase in income volatility at age 43, while the other half experiences a decrease.

inability to insure against fluctuations in income streams (referred to as the "volatility effect"). Second, faced with higher risk in the labor market, households would like to reduce their exposure to risky investment in the financial market and, as a result, forgo the opportunity to exploit the risk premium (the "portfolio effect").

Suppose that some workers are exogenously switched to a high-volatility income regime in the early stage of the life cycle. More specifically, at age 22 ($\tau^* = 22$), one year after the labor-market entry in our model, these workers experience a sharp increase in income volatility from the middle to the highest volatility regimes ($\sigma_\zeta = 2.5\%$ to 20%). This increase in income volatility is equivalent to the three-fold increase in the standard deviation of income growth: 0.24 to 0.74. We denote these workers as "H" (high volatility). This change in income volatility is assumed to be unexpected and permanent. We compare their welfare to that of those who remain in the middle volatility regime, denoted by "M" (middle volatility).

Table 7 reports the volatility of consumption and the average rate of return from investment (during the working period, ages 21 to 65) for both groups. The workers in the H regime experience 18 percent more volatile consumption (the standard deviation of log consumption of 0.40 versus 0.34 in M). This results in a welfare cost of 8.4 percent in consumption-equivalent units. At the same time, they receive a lower average rate of return

of 3.0 percent (compared with 3.3 percent in the M group) as they invest their savings in a more conservative way. This results in a welfare cost of 0.5 percent in consumption-equivalent units. (See Appendix F for the details of computing the two types of welfare cost.)

Table 7: Welfare Costs of Volatile Income

Volatility Regime	"M"	"H"	Welfare Cost	
SD of Log Consumption	0.34	0.40	Volatility Effect	8.4%
Average Rate of Return	3.3%	3.0%	Portfolio Effect	0.5%

Persistence of Volatility Shocks The indirect inference requires a highly persistent volatility process in order to match the increasing age profile of kurtosis in the data. We examine the role of the persistence of volatility shocks on portfolio choice. Specifically, we vary the persistence and estimate the marginal effect of income volatility on the risky share (i.e., $\hat{\beta}$ in Equation (1), for $k = 4$, using model-simulated data). In each case, we adjust (i) σ_ζ^2 to match the same standard deviation of ΔSD as in the benchmark and (ii) the discount factor, δ, to match the same assets-income ratio as in the benchmark. Figure 8 plots the marginal effect (regression coefficients) separately for the PIM and IIM models.

As the volatility shocks become more persistent, the response of the risky share ($\hat{\beta}$) becomes larger: workers adjust the risky share aggressively when the volatility change is long-lived. For small values of ρ_σ, the responses are similar between PIM and IIM. When the shocks are transitory, the realized income is not so informative about the regime. As the volatility shocks become persistent, the response of the risky share diverges between PIM and IIM since the current outcome is more informative about the future regime.

Alternative Models of Stochastic Volatility We next examine the portfolio choices from the models under the alternative specification of stochastic volatility such as ARCH/GARCH (which are commonly used in the empirical analysis, e.g., Engle, 1982).

First, suppose that the income process follows an ARCH (which is also used in Meghir and Pistaferri (2004)). The individual variance of labor-income growth is:

$$\sigma_{i,j+1}^2 = \sigma_\nu^2 + \phi(y_j - \mathbf{H}_j'\mathbf{M}_{j-1})^2.$$

Now, the variance of income growth is a constant term (the median variance σ_ν^2) plus a term that depends on the squared innovation ν_j $(= y_j - \mathbf{H}_j'\mathbf{M}_{j-1})$. In this specification, future variations in income volatility are linked to realizations of innovations in earnings.

Figure 8: Persistence of Volatility Shocks and Response of Risky Share

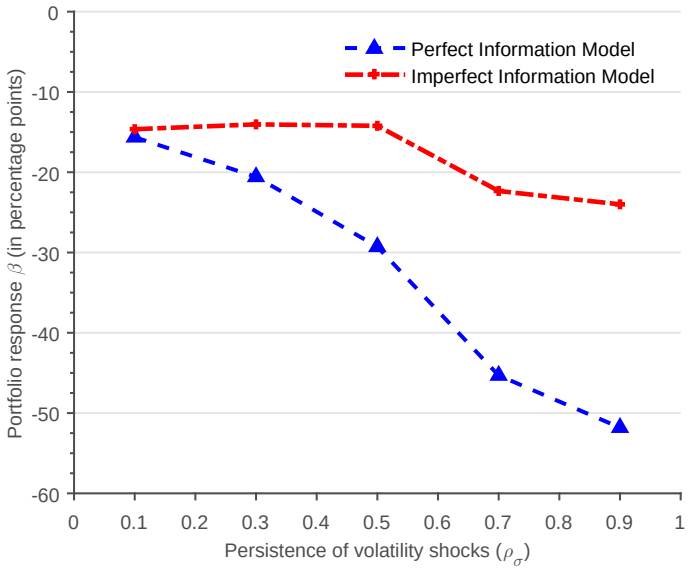

Notes: The x-axis is the persistence of volatility shocks, ρ_σ, and the y-axis is the estimated response of the risky share, $\hat{\beta}$, with $k = 4$, for PIM and IIM.

Under GARCH, the variance of income growth can be written as:

$$\sigma^2_{i,j+1} = (1 - \rho_\sigma)\sigma^2_\nu + \rho_\sigma \sigma^2_{i,j} + \phi(y_j - \mathbf{H}'_j \mathbf{M}_{j-1})^2.$$

This explicitly allows for persistence in the volatility regime. However, note that none of these specifications can accommodate imperfect information about the regime. In ARCH/GARCH, the next period's (expected) volatility depends on the innovation of the current labor income, which is assumed to be perfectly observed.

For both ARCH and GARCH, we calibrate the parameters of the model in a similar way to our benchmark. Specifically, ϕ is calibrated to match the cross-sectional dispersion of the projected (instrumented) change in volatility, $\widehat{\Delta SD}$. For GARCH, the persistence, ρ_σ, is calibrated to match the average kurtosis over the life cycle. We calibrate the discount factor, the risk aversion, and the credit limit to match the same financial moments. Finally, we keep the other labor-income parameters (including σ_ν) equal to the values estimated from the benchmark case. Figure 9 illustrates the implications of each stochastic volatility model for (a) the variance of log income, (b) the kurtosis of labor-income growth, (c) the variance of log consumption, and (d) the estimated portfolio response—the regression coefficient $\hat{\beta}$.

While the ARCH specification generates a realistic variance of labor income along the life cycle, it fails to generate an increasing kurtosis profile as there is no built-in persistence in the

Figure 9: Alternative Models: Evaluation

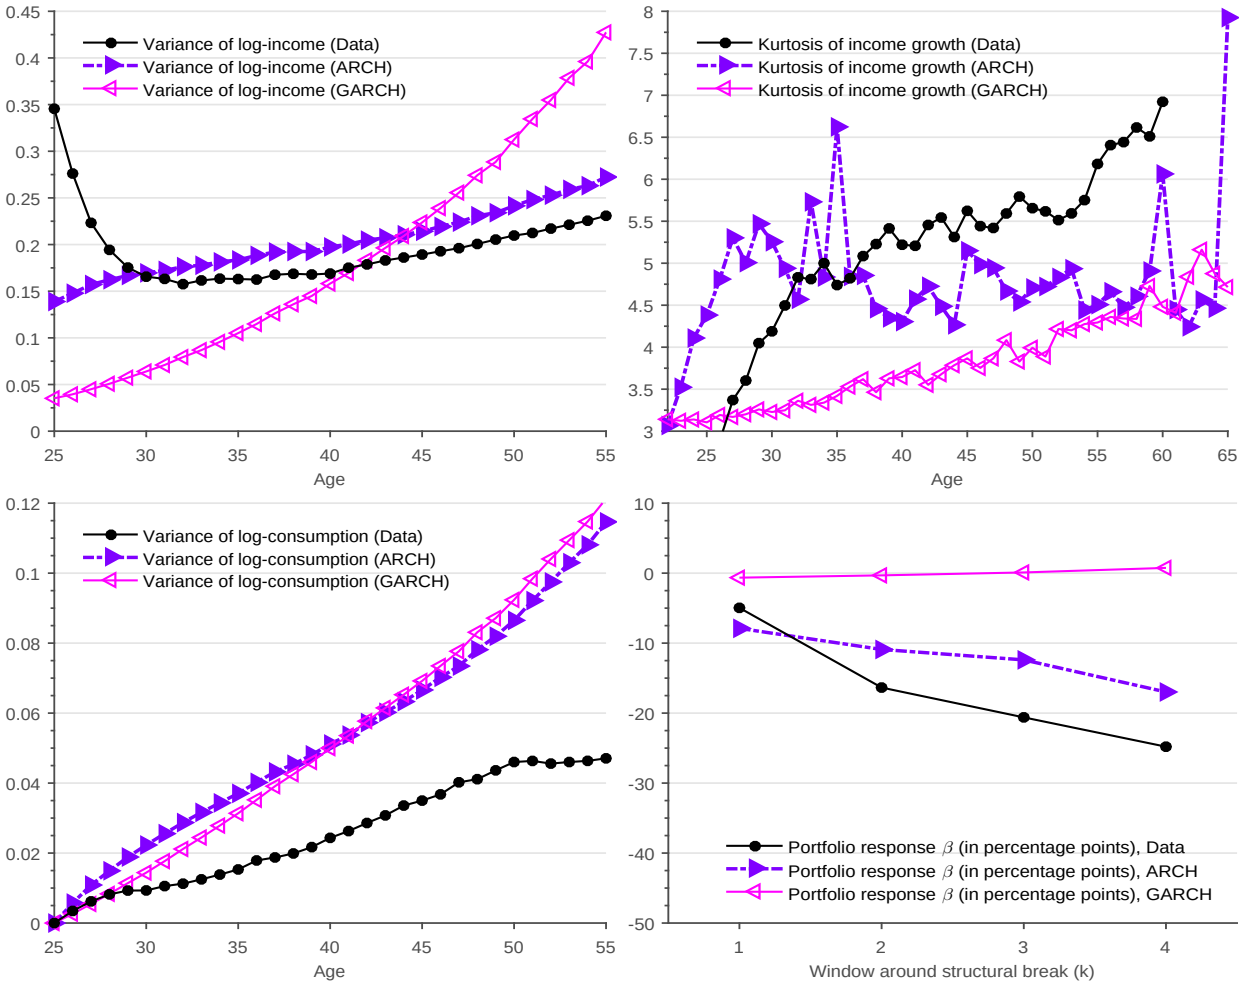

Notes: Model fit for ARCH and GARCH. Upper panels show the variance of log labor income (left) and kurtosis of labor-income growth (right), respectively. Lower-left panel shows the increase in the variance of log consumption along the life cycle. Lower-right panel shows the coefficients from a model-generated regression of risky share on volatility.

volatility process (as we explained in Section 4.4). Moreover, ARCH fails to generate a realistic profile for consumption inequality because the process is characterized by a strong mean reversion. A high income today increases the next period's volatility, inducing high-income workers to save more for precautionary reasons. A low income today also increases the next period's income volatility, inducing low-income workers to borrow because escaping from the low-income state is more likely when the income process is highly volatile. As a result, the ARCH model generates a large inequality in wealth that leads to a larger inequality in consumption. Finally, the ARCH model generates portfolio responses smaller than what we find in the data because there is no persistence in the volatility.

One characteristic of GARCH is the lack of symmetry in the evolution of income volatility. On average, volatility increases for most workers along the life cycle as the income process is persistent. This is not the case in our our benchmark model because workers can experience both increases and decreases in income volatility, which lead to a nearly constant average profile over the life cycle. As a result, the GARCH specification overestimates—by a wide margin—the income inequality along the life cycle. Consumption inequality rises even faster relative to ARCH because of the rapidly increasing income inequality. The GARCH can deliver an increasing kurtosis profile, which confirms that persistence in income volatility is a key element for the realistic age profile of kurtosis. However, we could not find a combination of parameters (ϕ, ρ_σ) to match both the dispersion of ΔSD and the average kurtosis. Finally, GARCH cannot deliver a realistic response of portfolio choices due to a lack of symmetry in the changes in volatility.

Overall, judging from the joint dynamics of labor income and portfolio choices (as well as consumption), we argue that our model of persistent stochastic volatility is the best representation of the data.

6 Conclusion

Households' portfolio decisions depend on the background risk in the labor market. Based on detailed administrative panel data from Statistics Norway, we find a statistically and economically significant shift in the risky share around the structural break in income volatility: if the standard deviation of labor-income growth doubles in size, the worker decreases the risky share by 4 percentage points on average. We find substantially larger estimates compared to many previous studies due to our identification strategy: individual-specific structural breaks of income volatility combined with a firm-side instrumental variable.

We then ask whether our estimated marginal effect is consistent with a standard model of portfolio choice with idiosyncratic shocks to income volatility. We structurally estimate the model using various moments on income, consumption, and financial portfolio choices from the Norwegian panel. The standard model fails to replicate the portfolio choice we see in the data. It predicts that in response to a sudden change in income volatility, a typical household should decrease the risky share much more aggressively than we see in the data. We show that imperfect information (Bayesian learning) about the income volatility can fill the gap between the model and the data by mitigating the response of the household's portfolio choice to the changes in income volatility. We also show that the structural model is useful for identifying the higher moments of earning dynamics as well as assessing the welfare cost of background risk.

References

Angerer, Z. and P.S. Lam. 2009. "Income risk and portfolio choice." *Journal of Finance* 64 (2):1037–1055.

Athreya, K., F. Ionescu, and U. Neelakantan. 2015. "Stock market participation: the role of human capital." Working paper, Federal Reserve Bank of Richmond.

Benzoni, L., P. Collin-Dufresne, and S. R. Goldstein. 2011. "Portfolio choice over the life-cycle when the stock and labor markets are cointegrated." *Journal of Finance* 62 (5):2123–2167.

Bertaut, C.C. and M. Haliassos. 1997. "Precautionary portfolio behavior from a life-cycle perspective." *Journal of Economic Dynamics and Control* 21 (8-9):1511–1542.

Bound, J., C. Brown, and N. Mathiowetz. 2001. "Measurement error in survey data." *Handbook of Econometrics* 5:3705–3843.

Campbell, J., J. Cocco, F. Gomes, and P. Maenhout. 2001. "Investing retirement wealth: a life-cycle model." In *Risk Aspects of Investment-Based Social Security Reform*, edited by M. Feldstein and J. Campbell. NBER.

Card, D., A. Mas, and J. Rothstein. 2008. "Tipping and the dynamics of segregation." *Quarterly Journal of Economics* 123 (1):177–218.

Catherine, S. 2019. "Countercyclical labor income risk and portfolio choices over the life cycle." Working paper, University of Pennsylvania.

Chang, Y., J. Hong, and M. Karabarbounis. 2018. "Labor-market uncertainty and portfolio choice puzzles." *American Economic Journal: Macroeconomics* 10 (2):222–262.

Charles, K.K., E. Hurst, and M.J. Notowidigdo. 2018. "Housing booms and busts, labor market opportunities, and college attendance." *American Economic Review* 10 (108):2947–2994.

Cocco, F. J., J. F. Gomes, and J. P. Maenhout. 2005. "Consumption and portfolio choice over the life-cycle." *Review of Financial Studies* 18 (2):491–533.

Cunha, F. and J. J. Heckman. 2007. "The evolution of inequality, heterogeneity and uncertainty in labor earnings in the US economy." Working paper, National Bureau of Economic Research.

Davis, S. J. and P. Willen. 2000. "Occupation-level income shocks and asset returns: their covariance and implications for portfolio choice." Working paper, National Bureau of Economic Research.

Deaton, A. and C. Paxson. 1994. "Intertemporal choice and inequality." *Journal of Political Economy* 102 (3):437–467.

Dimson, E., P. Marsh, and M. Staunton. 2008. "The worldwide equity premium: A smaller puzzle." In *Handbook of the Equity Risk Premium*, edited by Rajnish Mehra. Elsevier.

Eika, L., M. Mogstad, and O. L. Vestad. 2017. "What can we learn about household consumption expenditure from data on income and assets?" Working paper, University of Chicago.

Engle, F.R. 1982. "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation." *Econometrica* 50 (4):987–1008.

Fagereng, A., C. Gottlieb, and L. Guiso. 2017. "Asset market participation and portfolio choice over the life-cycle." *Journal of Finance* 72 (2):705–750.

Fagereng, A., L. Guiso, and L. Pistaferri. 2017. "Portfolio choices, firm shocks, and uninsurable wage risk." *Review of Economic Studies* 85 (1):437–474.

Fagereng, A. and E. Halvorsen. 2015. "Debt and household level responses." Working paper, Norges Bank.

———. 2017. "Imputing consumption from Norwegian income and wealth registry data." *Journal of Economic and Social Measurement* 42 (1):67–100.

Fagereng, A., M.B. Holm, B. Moll, and G. Natvik. 2019. "Saving behavior across the wealth distribution: The importance of capital gains." Working paper, National Bureau of Economic Research.

Floden, M. 2001. "The effectiveness of government debt and transfers as insurance." *Journal of Monetary Economics* 48 (1):81–108.

Gollier, C. and J. W. Pratt. 1996. "Risk vulnerability and tempering effect of background risk." *Econometrica* 64 (5):1109–1123.

Gomes, J. F. and A. Michaelides. 2005. "Optimal life-cycle asset allocation: Understanding the empirical evidence." *Journal of Finance* 60 (2):869–904.

Guiso, L., T. Jappelli, and D. Terlizzese. 1996. "Income risk, borrowing constraints and portfolio choice." *American Economic Review* 86 (1):158–172.

Guiso, L., L. Pistaferri, and F. Schivardi. 2005. "Insurance within the firm." *Journal of Political Economy* 113 (5):1054–1087.

Guvenen, F. 2007. "Learning your earning: Are labor income shocks really very persistent?" *American Economic Review* 97 (3):687–712.

———. 2009. "An empirical investigation of labor income processes." *Review of Economic Dynamics* 12 (1):58–79.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song. 2015. "What do data on millions of U.S. workers reveal about life-cycle earnings risk." Working paper, National Bureau of Economic Research.

Guvenen, F. and T. Smith. 2014. "Inferring labor income risk and partial insurance from economic choices." *Econometrica* 82 (6):2085–2129.

Haliassos, M. and A. Michaelides. 2003. "Portfolio choice and liquidity constraints." *International Economic Review* 44 (1):143–177.

Heathcote, J., K. Storesletten, and G.L. Violante. 2014. "Consumption and labor supply with partial insurance: An analytical framework." *American Economic Review* 104 (7):2075–2126.

———. 2017. "Optimal tax progressivity: An analytical framework." *Quarterly Journal of Economics* 132 (4):1693–1754.

Heaton, J. and D. Lucas. 2000. "Portfolio choice in the presence of background risk." *Economic Journal* 110 (460):1–26.

Huggett, M. and G. Kaplan. 2016. "How large is the stock component of human capital?" *Review of Economic Dynamics* 22:21–51.

Klovland, J.T. 2004. "Bond markets and bond yields in Norway 1819-2003." In *Historical Monetary*

Statistics for Norway 1819-2003, edited by Oyvind Eitrheim, J.T. Klovland, and J.F. Qvigstad. Norges Bank.

Lusardi, A. and O.S. Mitchell. 2007. "Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth." *Journal of Monetary Economics* 72 (1):1–32.

Meghir, C. and L. Pistaferri. 2004. "Income variance dynamics and heterogeneity." *Econometrica* 72 (1):1–32.

Palia, D., Y. Qi, and Y. Wu. 2014. "Heterogeneous background risks and portfolio choice: Evidence from micro-level data." *Journal of Money, Credit and Banking* 8 (46):1687–1720.

Pischke, J. 1995. "Individual income, incomplete information, and aggregate consumption." *Econometrica* 63 (4):805–840.

Primiceri, G. and T. van Rens. 2009. "Heterogeneous life-cycle profiles, income risk and consumption inequality." *Journal of Monetary Economics* 56 (1):20–39.

Storesletten, K., I. C. Telmer, and A. Yaron. 2004. "Consumption and risk sharing over the life cycle." *Journal of Monetary Economics* 51 (3):609–633.

Viceira, L. M. 2001. "Optimal portfolio choice for long-horizon investors with non-tradable labor income." *Journal of Finance* 56 (2):433–470.

Wachter, A. J. and M. Yogo. 2010. "Why do household portfolio shares rise in wealth?" *Review of Financial Studies* 23 (11):3929–3965.

Appendix

A Data

A.1 Data Sources

Based on Statistics Norway, we combined the several data sets using the unique personal identifiers. The details of the data are as follows.

Central Population Register: The data contain individual demographic information for all Norwegian residents from 1992 to 2014. This includes personal variables (country of birth, first stay date, immigration category, country background, gender, date of birth) as well as time-varying characteristics (marital status, spouse's ID if married). The family identifiers can be used to link spouses and cohabiting couples with common children. Family structure and family types variables are also available: total number of persons in the family, the age of youngest child, the age of youngest and oldest person in the family, the number of children under the ages of 18, 16, 11, and 6, family type, father ID and mother ID at the time of birth.) Some of these variables are missing for several years (e.g., family type).

National Educational Database: All individual statistics on education have been gathered in the National Education Database (NUDB) since 1970. Educational attainment is reported by the educational establishment directly to Statistics Norway. By October 1 of each year, the completed education from the previous school/academic year is updated and the information about the highest attained level of education for the whole population is updated as well.

Administrative Tax and Income Records: All households in Norway are subject to an income tax and a wealth tax, and they are obliged to report their complete income and wealth holdings to the tax authority every year. Also, employers, banks, brokers, insurance companies and any other financial intermediaries are obliged to send the information on the value of the assets owned by the individual to the individual and to the tax authority as well. Traded financial securities are reported at market value; value of shares in private companies is reported by individuals as well as private companies to the tax authority. The tax authority will combine the information from companies' reports on net worth with individuals', and adjust if necessary. For more details, see annual reports from the tax authority (http://www.skatteetaten.no) as well as other literature (e.g., Fagereng and Halvorsen (2017))

Income Registry: From the income registry data, we gathered the following items:

- Earned income includes cash salary, taxable benefits and sickness and maternity benefits during the calendar year.
- Net entrepreneurial income includes the income from land and forestry, fishing and hunting, income from other business activities and sickness benefits in employment during the calendar year.

- Capital investment income includes interest income, dividends, realized gains and other investment income:
  - Interest income (from bank deposits and accounts receivable during the calendar year)
  - Dividends received during the calendar year
  - Realization gains: Taxable gains on the sale of real estate and financial securities during the calendar year.
  - Deductible losses on the sale of real estate and financial securities during the calendar year.
  - Other income: Net income from the rental of real property outside the industry, return on "spare part" of life insurance, income from abroad and other unspecified investment income during the calendar year.
- Unemployment benefits paid to wage earners and self-employed.
- Pensions:
  - Pensions from the scheme (ftryg): includes own pensions and national insurance, including a spouse's allowance and child benefit for children aged 16 or younger.
  - Own pensions: Occupational pensions include the payment to working conditions including contractual pension (AFP). It also includes payments from individual pension agreements (IPA), annuities and maternity council benefits in agriculture and forestry.
- Transfers:
  - Tax-free transfers include child allowance, housing allowance, study grants, social assistance, basic and auxiliary disability compensations.
  - Taxable transfers include pensions from the National Insurance, pension, unemployment benefit and received contributions, as well as other taxable transfers.
- Miscellaneous items in the income tax record:
  - Alimony and annuities outside employment
  - Sum of income and wealth taxes and negative transfers
  - Sum of interest payments (interest on debt to Norwegian and foreign creditors) and residential income (imputed residential income and leisure property and shareholder's share of income from housing companies).

Wealth Registry Record: For persons who are older than 17 years, we can obtain the wealth data from the tax authority every year. Here are the descriptions quoted from the tax administration.[30]

[30]For detailed information, see the corresponding 2015 tax form from the Norwegian Tax Administration: http://www.skatteetaten.no/en/person/Tax-Return/Find-item/#&del1=1&del2=1&del3=1&del4=1&del5=1.

- **Bank Deposits** *Deposits in Norwegian banks (entry 4.1.1 in the 2015 tax form)*

  "This item shows what deposits you and your children who are under 17 years of age at the end of the income year have in Norwegian banks as of 31 December. Deposits belonging to children under 17 will be split with half being assigned to each of the parents if they live together. The amount will normally be pre-filled with the amount that has been reported."

- **Value of Shares in Mutual Funds** *(entry 4.1.4)*

  "The item is pre-filled with information concerning capital in Norwegian unit trusts and certain foreign unit trusts which the Tax Administration has received information about from the management companies concerned."

- **Value of Financial Securities** *(entry 4.1.7)*

  "This item shows the value of bonds and shares in the Norwegian Central Securities Depository (VPS) as of 31 December. The amount will normally be pre-filled with what has been reported."

- **Value of Shares in Private Companies** *(entry 4.1.8)*

  "This item shows the capital value of shares and other securities not registered with VPS."

- **Tax Value of Housing and Other Real Property** *(entry 4.3)*

  Capital such as dwellings, holiday homes, forest property, farms, agricultural property, plots and commercial property, etc.

- **Value of Home Owned** *(entry 4.3.2)*

  "If you own a home as of 31 December, the tax value must be entered under this item. The tax value is determined on the basis of factors such as location, size and year of construction. If the tax value exceeds 30 percent of the home's market value, you can change the value if you are able to document the market value. This concerns your primary home."

- **Premium Funds, Individual Pension Agreements** *(entry 4.5.1)*

  "This item shows your capital in the form of premium funds as of 31 December. The amount will normally be pre-completed with the amount that has been reported by the company you have entered into a pension agreement (IPA) with."

- **Value of Life Insurance Policies** *(entry 4.5.2)*

  "This item shows the surrender value of your annuities as of 31 December. The amount will normally be pre-filled with the amount that has been reported by the insurance company/companies or employer who has made deposits on your behalf."

- **Other Capital** *(entry 4.5.4)*

  "If you have other taxable capital as of 31 December, the value must be entered under this item. 'Other taxable capital' means for example assets in the form of capitalised ground rent, rights linked to forest/uncultivated land, share of company assessed as partnerships (RF-1221) and/or NOKUS (RF-1246)."

- **Total Debt** *(entry 4.8)*

  "The items under '4.8 Debt' concern negative capital such as loans, credit purchases, underpaid tax, private loans, debts in housing cooperatives/jointly owned properties, debts abroad and deductions for leasehold land, etc."

- **Total Net Worth** *(entry 4.9)*

  "Amounts specified under this item form the basis for the calculation of wealth tax payable to municipalities and the state. The basis for this is the sum of your wealth with a reduction for any debt."

Employer-Employee Register: Statistics Norway combines the required report from each firm that hires workers and the tax record from individuals. The data include detailed labor market information for every worker each year (worker ID, employers' ID, job starting date with each employer, job ending date with each employer, total payments to workers from each employer, industry, occupation, actual and expected working hours, total number of days worked, indicator for full-time/part-time employment, etc.).

Register of Social Assistance Received: For each person from 1992, the register records monthly the details of social assistance received. This includes any unemployment benefit, rehabilitation/medical rehabilitation, maternity, temporary disability insurance as well as the benefit for sickness.

A.2 Sample Selection and Variable Definitions

Sample Selection Criteria

1. From the whole population with income tax registration records in 1993 (who are older than 17 years old), we choose male, Norwegian native, older than 25 in 1993, with no missing records on education and other demographic variables.

2. We randomly select 10 percent and construct a panel for the period of 1993-2014.

3. Starting in 1994 and onward, for each year, we add a random sample of 10 percent who became 25 years old to the panel data.

4. Using a unique person ID, we link the labor market information (industry, sector, occupation, working days, and so on) with income tax record and wealth tax record. The occupation data are only available starting in 2003.

5. Using a unique person ID, we link the social assistance (including unemployment benefits, number of months being paid).

Linking Household Information

1. From the Central Population Register data, we obtain marital status and link to his spouse's ID.

2. From the Central Population Register data, we obtain the person's father's ID and mother's ID at the time of his/her birth; we can also link to each of his children's ID. This provides the information on family size, number of children, and number of young children.

3. Based on the person ID, spouse ID, father ID, mother ID, and children's IDs', we construct information on the spouse's income and wealth, father's income and wealth, and each child's income and wealth if older than 17. This yields household-level income and wealth. (The administration registry does not keep track of household-level income and wealth.)

Multiple Jobs and Main Job: In our benchmark sample, approximately 10 percent of workers have multiple jobs within a year. Some of them are associated with multiple employers (different establishment ID). Also, some workers with the same employer (establishment ID) have records of earnings with different starting and/or ending dates. The latter cases reflect, for most workers, changes in titles, job requirements, or new contracts with the same employer. Thus, we simply add up the earnings under the same establishment within a year. Following the literature, we define the main job, possibly among different employers, as the one with the largest earnings within a year. As a result, this will reduce the person-year observations of earnings slightly, from 2,880,970 to 2,864,084.

Labor Market Status: Since the respondent does not report labor market status, we classify them as follows:

- Employed: those with positive earnings and non-missing establishment ID.
- Unemployed: those with positive unemployment benefit during the year and the received unemployment benefit greater than received earnings (if any), and not receiving pensions (either from own pensions or other pension plans).
- Self-employed: those not associated with any employers, not receiving any unemployment benefits during the year, and with positive reported business income.
- Retired: those not associated with any employers, not receiving any unemployment benefits during the year, and with positive pensions (either from own pensions or other pensions plans)

Homeownership: In the wealth registry, we observe the reported value of home. One natural way to define homeowners is just to find those with a positive value of home. However, according to this definition, it turns out the homeownership rate is very low and not consistent with the report from Statistics Norway (e.g., see the link "https://www.ssb.no/en/bo/"). Since some households may have positive value of other real estate—which is either occupied by themselves or rented out—we prefer to add up those values with the value of home. As a result, the homeownership rate is close to 80 percent in the sample, very close to the official statistics.

Financial Assets and Risky Shares:

- **Total Financial Assets**: Value of life insurance + Value of Shares in mutual funds + Bank Deposits + Value of Shares in private companies + Value of Financial Securities + Value in premium fund (mostly pension funds)
- **Risky Assets**: Value of Shares in mutual funds + Value of Shares in private companies + Value of Financial Securities

Same Employer Dummy: If the worker works with the same establishment ID this period and last period—we require both establishment IDs and earnings records to not be missing.

Tenure for Current Job: This is calculated from the job starting date, which is reported in the Employer-Employee Register. Note that for each worker and each year, if he has multiple job records within the same establishment (e.g., someone works based on monthly contracts), we use the summation of all earnings as the earnings record and use the earliest date as the job starting date. If he has several records across different establishments, we then select the one with the largest earnings.

Number of Job Changes: Total number of times changing employers in his history (in our sample period) up to date.

Industry Stayer Dummy: If the worker works in the same industry (at the 2-digit level) in year $t-1$ and t.

Number of Industry Changes: The total number of times of changing industries up to date.

Occupation Stayer Dummy: If the worker works at the same occupation (at the 2-digit level) in $t-1$ and t.

Number of Occupation Changes: The number of times changing occupations up to date.

Number of Occupations Experienced: The number of different occupations up to date.

Years of Experience in Occupation: The total years of experience with the current occupation up to date.

B Summary Statistics and Robustness

Table 8: Sample Selection

	Number of Individuals	Person-Year Obs.	%
10% of registered Norwegian males	137,776	2,880,970	100.0
+ Earnings $\geq$ 18 years	48,768	1,014,882	35.4
+ Positive RS $\geq$16 years	18,156	379,204	13.1
+ Financial Assets $\geq$ 50K NOK	16,700	348,918	12.1
+ Identified Structural Breaks	16,041	337,899	11.6

Notes: Earnings are deflated by the CPI. "10% of registered men" includes 10% random sample of registered Norwegian males, older than 25, with no missing records on birth year, education, and income and wealth tax.

Table 9: Summary Statistics for Income Growth and Changes in Volatility

	Mean	S.D.	Percentiles						
			5%	10%	25%	50%	75%	90%	95%
y_i	0.096	0.369	-0.492	-0.316	-0.105	0.097	0.308	0.527	0.678
$\Delta y_{i,t}$	-0.002	0.398	-0.499	-0.240	-0.065	0.004	0.069	0.230	0.476
$SD_i[\Delta y_{it}]$	0.319	0.260	0.056	0.075	0.136	0.247	0.419	0.659	0.831
$\Delta SD_{i,\tau^*}[\Delta y_{it}]$	-0.098	0.583	-1.031	-0.722	-0.352	-0.072	0.193	0.498	0.775

Table 10: Robustness for IV Estimates

Dependent Variable: $RS_{i,\hat{\tau}+k} - RS_{i,\hat{\tau}-k}$	$k=1$	$k=2$	$k=3$	$k=4$
Excluding firms <25 pct.	-13.68*	-32.95***	-24.79**	-21.57**
Excluding managers within HHs	-4.85	-21.82***	-24.20***	-30.75***
Using value added/assets only	13.46	-23.24*	-38.11***	-27.36*
Using net sales/assets only	-4.46	-10.95	-4.88	-28.78
Excluding HHs with unemp. spells	-7.11	-25.36***	-23.15***	-25.23***
Financial assets >20 p.p.	-8.41	-31.44***	-30.17***	-28.32***
Risky share >5 p.p.	-7.27	-23.12***	-26.11***	-28.90***
Using τ^* defined from workers' ΔSD	-14.45	-13.86	-32.61**	-36.15***
With higher-order income & wealth	-5.12	-16.98***	-20.02***	-24.92***
With industry dummies	-3.00	-15.37***	-18.94***	-23.00***
With industry & occ. dummies	-3.19	-21.28***	-27.02***	-28.68***
Staying with the same employer	-4.50	-22.67***	-22.44***	-23.07***
Always married from $-k$ to $+k$	-3.85	-15.61***	-20.14***	-24.64***
Always single from $-k$ to $+k$	-20.26	-29.20*	-14.08	-45.41**
Using volatility of disposable income	-14.81**	-34.53***	-39.31***	-39.88***
Controlling mortgage debt & housing	-4.281	-16.37***	-20.08***	-24.16***

Table 11: Responses by Group: IV Estimates

	$k=1$	$k=2$	$k=3$	$k=4$
Age				
Young ($age < 40$ in 2005)				
$\Delta SD_{i,\tau^\star(i)}$	-5.30	-27.02**	-26.22*	-25.28
	(10.71)	(13.50)	(14.76)	(15.85)
Middle ($40 \leq age \leq 55$)				
$\Delta SD_{i,\tau^\star(i)}$	-8.46	-16.83**	-13.55*	-22.02***
	(5.800)	(6.610)	(7.041)	(7.329)
Old ($age > 55$ in 2005)				
$\Delta SD_{i,\tau^\star(i)}$	-3.45	-8.813	-29.48**	-15.19
	(9.144)	(12.54)	(12.14)	(12.70)
Education				
College	8.49	-23.77***	-29.25***	-36.66***
	(7.131)	(8.664)	(9.146)	(10.08)
High School	-11.21*	-17.79**	-20.88***	-25.26***
	(5.936)	(7.089)	(7.505)	(7.661)
Financial Knowledge				
Econ/Finance Major	-3.11*	-4.10*	-6.64***	-5.71**
	(1.761)	(2.376)	(2.221)	(2.450)
Other Majors	-6.34	-17.94***	-24.93***	-28.32***
	(4.455)	(5.267)	(5.656)	(5.862)
Income Growth				
Top 25%	-10.03	-18.25*	-23.80**	-30.69***
	(8.371)	(9.76)	(10.47)	(10.98)
Middle	-1.87	-13.30*	-23.61***	-23.56***
	(6.48)	(7.53)	(7.93)	(8.03)
Bottom 25%	-5.64	-20.86**	-14.30	-25.52**
	(6.98)	(8.49)	(9.34)	(9.98)

C Life-Cycle Kurtosis

We construct the life-cycle profile of kurtosis controlling for cohort effects. Following the methodology of Deaton and Paxson (1994) we first compute kurtosis in labor income growth by age and cohort group and second, regress the computed kurtosis on cohort and age dummies. The left panel of Figure 10 shows the kurtosis profile for selected cohorts, while the right panel plots the age dummies. The labor income growth of Norwegian workers exhibits substantial kurtosis. The average kurtosis is 8.3, which suggests considerable deviations from a normal distribution that has a kurtosis equal to 3.

The high average kurtosis reflects the large dispersion in volatility breaks, ΔSD. As mentioned, we have filtered these changes to capture plausibly exogenous variations in volatility using firm-side volatility. It is thus more appropriate to estimate our model based on a scaled kurtosis profile that is closer to the instrumented worker volatility change $\widehat{\Delta SD}$. To construct the scaled profile we apply the following process. First, we fit a linear trend to the profile of kurtosis and generate the trend coefficients $\{t_j\}$ and the deviations from the linear trend $\{a_j\}$ where j stands for age. Second, we construct a scaled version of the linear trend $\{s_j\}$ by setting $s_1 = t_1$ (both profiles start at the same point at age 25) and $s_{60} = xt_{60}$ (where x is the scale factor).

Our motivation for the scaled kurtosis profile is that it reflects $sd(\widehat{\Delta SD}) = 0.55$, e.g., the raw measure of volatility breaks, and not $sd(\widehat{\Delta SD}) = 0.128$, e.g., the exogenous measure by which we discipline the model. To calculate x we construct a relationship between the standard deviation of $\widehat{\Delta SD}$ and kurtosis. We use model-simulated data from a process without persistence to capture only the relationship between average kurtosis and the standard deviation of $\widehat{\Delta SD}$. We find that if the standard deviation of ΔSD decreases by 0.1, the average kurtosis decreases by 1.25. Since we adjusted downward the dispersion of ΔSD by 0.43 (0.55-0.12), we need to re-scale (decrease) kurtosis by $0.43 \times 1.25/0.1 = 5.37$, which implies that $x = 0.55$. Finally, we add the deviations $\{a_j\}$ to the scaled linear trend to obtain the scaled kurtosis profile (right panel of Figure 10).

Figure 10: Cross-Sectional Kurtosis across Age Groups

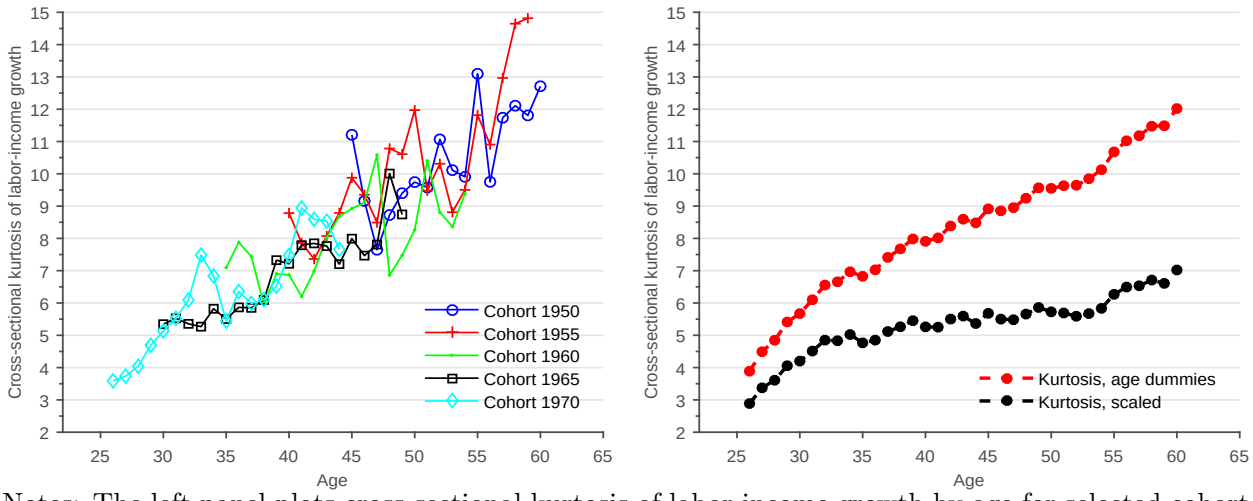

Notes: The left panel plots cross-sectional kurtosis of labor income growth by age for selected cohort groups. The left panel plots the cross-sectional kurtosis of labor income growth net of cohort effects and the scaled profile that takes into account the firm-induced variability.

D Indirect inference

We estimate the model using indirect inference. We need to estimate a total of $P = 9$ parameters represented by the vector

$$\Theta = [\delta, \gamma, \underline{b}, \sigma_a^2, \sigma_\beta^2, \rho_x, \rho_\sigma, \sigma_\nu^2, \sigma_\zeta^2].$$

We estimate the parameters by minimizing the distance between a set of model-generated moments $M^m(\Theta)$ and their empirical counterpart $M^d(\Theta)$. As explained in the main text, the moments used in the estimation are

1. The average assets-income ratio.
2. The average risky assets-total assets ratio.
3. The average consumer debt-income ratio.
4. The variance-covariance matrix of log labor income across ages (a total of 441 moments).
5. The dispersion in the firm-induced volatility change $\widehat{\Delta SD}$.
6. The life-cycle profile in the cross-sectional kurtosis of labor-income growth.
7. The life-cycle profile in the cross-sectional variance of log-consumption.
8. The response of the risky share to volatility (over a 4-year horizon).

In total we have $M = 511$ moments. The estimator minimizes the loss function

$$\min L_\Theta = [M^d - M^m(\Theta)]' W [M^d - M^m(\Theta)]. \tag{D.1}$$

W denotes the weighting matrix of the moments. We first re-normalize the units by setting weight on the diagonal element i equal to $W_i = 1/M_i^d$, e.g., the inverse of the empirical moment we target. We then distribute more weights on moments we think are important to match closely in our estimation. In particular, we multiply 500 times the weight associated with the average assets-income ratio, 50 times for both the weights associated with the average risky share and the response of the risky share to volatility, 500 times the weight associated with the variance of log-earnings at three stages of the life cycle (ages 25, 45, 60). Finally, we multiply by 7,000 the weight associated with the dispersion in volatility. Θ^* is the parameter vector that minimizes the distance between the model and empirical moments given our choice of the weighting matrix. We compute the standard errors as follows.

1. Given Θ^* we generate $N = 100$ different draws of shocks. For each $n = \{1, .., N\}$ we compute the model-generated moments $M_n^m(\Theta^*)$.
2. We compute the variance-covariance matrix of the moments $Cov(M_n^m)$, which is an $M \times M$ matrix. We construct the $M \times M$ matrix S, which includes only the diagonal elements of $S = diag(Cov(M_n^m))$.

3. For each parameter $\theta = \{1, .., P\}$ we compute the numerical derivative with respect to each moment m. We do so by altering slightly the value of the parameter relative to its estimated value and computing the difference in the resulting value of the model-generated moment. We end up with a $P \times M$ matrix $\frac{\Delta m_n}{\Delta \theta}$.

4. We compute the standard errors using the diagonal elements of the inverse of the following matrix

$$V = \frac{1}{N}(\frac{\Delta m_n}{\Delta \theta} S^{-1} \frac{\Delta m_n}{\Delta \theta}').$$

E Robustness: Small Risk Aversion

According to our model, the risky share overshoots in PIM because workers immediately realize the regime change. The estimated relative risk aversion is 5.4 in PIM and 4.8 in IIM. One might be concerned that the smaller risk aversion may drive this different response between PIM and IIM. To explore this possibility, we solve the PIM with the risk aversion of 4.8 the estimate in the IIM (i.e., 4.86). The responses are reported in Figure 11. Even with a smaller risk aversion, the response of the risky share in PIM substantially over-predicts the response of the risky share. Notably, the discrepancy is more visible at shorter time horizons. Workers gradually learn about the regime change so the difference between PIM and IIM remains substantial for shorter time horizons.

Figure 11: Portfolio Response in PIM with Smaller Risk Aversion

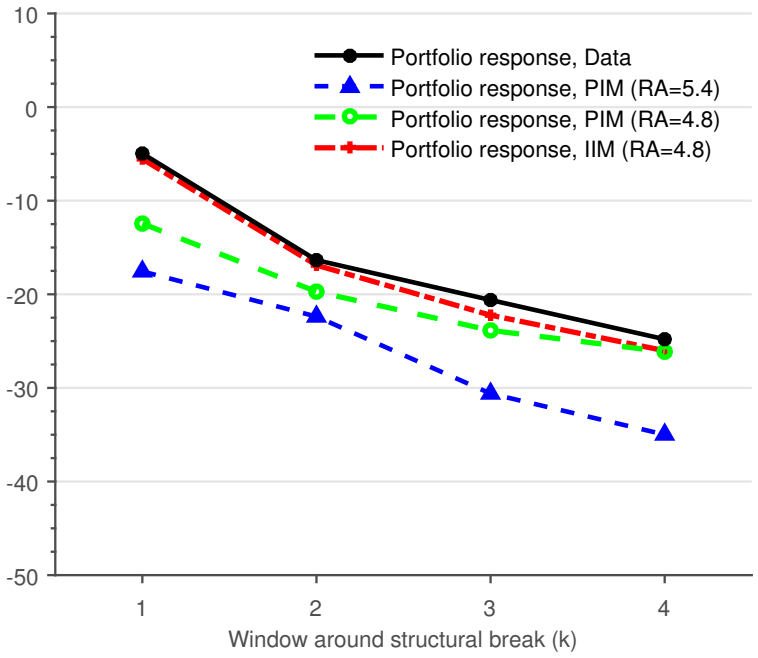

F Calculation of Welfare Costs

To isolate each channel of welfare loss we adopt the decomposition proposed by Floden (2001). The life-time utility of each group $p = \{H, M\}$ is given by:

$$V_p = \sum_{j=1}^{J} \sum_{i \in p} \delta^{j-1} (\mathbf{\Pi}_{t=1}^{j} s_t) \frac{{c_{ij}}^{1-\gamma}}{1-\gamma}.$$

If for each age, each worker had consumption equal to the average consumption within each age group, then welfare is:

$$\bar{V}_p = \sum_{j=1}^{J} \delta^{j-1} (\mathbf{\Pi}_{t=1}^{j} s_t) \frac{c_j^{p\,1-\gamma}}{1-\gamma},$$

where $c_j^p = \sum_i c_{ij}^p / N$ and N is the total number of workers.

To find the volatility effect, we first compute the compensating differential (x_p) for no dispersion in consumption for each group:

$$\sum_{j=1}^{J} \delta^{j-1} (\mathbf{\Pi}_{t=1}^{j} s_t) \frac{c_j^{p\,1-\gamma}}{1-\gamma} = \sum_{j=1}^{J} \sum_{i \in p} \delta^{j-1} (\mathbf{\Pi}_{t=1}^{j} s_t) \frac{(1+x_p) {c_{ij}}^{1-\gamma}}{1-\gamma}.$$

This equals $\bar{V}_p = (1+x_p)^{1-\gamma} V_p$ and $x_p = (\frac{\bar{V}_p}{V_p})^{\frac{1}{1-\gamma}} - 1$. According to our calibrated model, workers in the M group are willing to give up 16 percent of annual consumption to be in an economy with equal consumption within age group, while workers in the H group are willing to give up 26 percent of their consumption. The welfare loss due to an increased volatility in H relative to that in M is $\omega^{\text{v}} = \frac{1+x_H}{1+x_M} - 1$.

Finally, in response to an increased background risk, workers move their portfolio toward safer assets, missing out on the equity premium. To estimate the portfolio effect, we compute life-time utility for each group in a counterfactual economy, where H workers experience on average the same financial returns as the M group in the benchmark example. We find that the equity premium has to increase to $\mu = 3.94$ relative to $\mu = 3.14$ for the H group to receive on average return of 3.3%. We then compute the welfare from being in the counterfactual economy (high volatility and high returns) $V_H^{\mu_H}$. The welfare loss from missing on the equity premium in H group is $\omega^{\mu} = (\frac{V_H^{\mu_H}}{V_H})^{\frac{1}{1-\gamma}} - 1$.