Economic Quarterly—Volume 106, Number 4—Fourth Quarter 2020—Pages 159–171

A Life-Cycle Model with Individual Volatility Dynamics

Marios Karabarbounis

A large literature has studied how the presence of uninsurable labor-income risk affects the patterns of savings and portfolio allocation over the life cycle. For example, workers in risky companies, occupations, or industries may have a larger incentive to accumulate wealth to insure against adverse events, such as unemployment, and to prepare for retirement. Moreover, they are likely to hold different investment portfolios, e.g., how much they invest in risky assets and how much of their investment is directed toward liquid versus illiquid accounts. In models with heterogeneous agents, income risk is usually represented by a probability distribution over income draws with a constant variance. Nonetheless, there is increasing evidence that labor-income risk is itself idiosyncratic. For example, Meghir and Pistaferri (2004) use income data from the Panel Study of Income Dynamics to show that there is strong support in favor of income dynamics with a time-varying volatility. Guvenen, Karahan, Ozkan, and Song (2015) show that an income process where variance switches stochastically between low and high regimes can match several higher-order of income moments including the high kurtosis of earnings in the U.S. data. Chang, Hong, Karabarbounis, Wang, and Zhang (2020) use administrative data from Statistics Norway to calibrate a life-cycle model with stochastic volatility in earnings and explore its implications for portfolio choice.

The author thanks Grey Gordon, Paul Ho, Colton Lapp, James Lee, and Urvi Neelakantan for valuable comments. Any opinions expressed are those of the author and do not necessarily represent those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

In this article, I incorporate a dynamic income process with time-varying volatility into an otherwise standard life-cycle model of consumption and savings. Volatility dynamics take the form of an autoregressive, conditional, heteroskedasticity (ARCH) model in the spirit of Meghir and Pistaferri (2004). While there can be many other ways to model time-varying volatility, this process is relatively parsimonious in terms of parameters. In the model, in every period, workers draw idiosyncratic productivity shocks from a normal distribution. The variance of this distribution is not constant but depends on (the previous year's) income draw. If workers experienced a sudden high income draw, they also experience a higher subsequent volatility of earnings. If workers experienced a stable income stream, they also experience a low subsequent volatility of earnings. Given the idiosyncratic income and volatility, workers choose how much to consume or save in risk-free bonds. For simplicity, the analysis is kept in partial equilibrium, and it abstracts from the role of volatility on the real interest rate.

I analyze how the ARCH process affects consumption and savings patterns. To analyze the implications of an ARCH process, I perform comparative statics with respect to the key parameter affecting the income dynamics: the "responsiveness" of the variance to the idiosyncratic income draws (denoted ϕ in the model in Sections 1 and 2). I show that the higher ϕ is, the larger the economy-wide risk is. As a result, workers accumulate more assets to maintain a relatively stable path of life-cycle consumption, and consumption inequality decreases.

1. AN ARCH INCOME PROCESS

I explain here the simple mechanics that follow an ARCH income process. Assume y_t represents the income draw at period t given by

$$y_t = \rho y_{t-1} + \varepsilon_t.$$

Income depends on two components. First, income depends on lagged income y_{t-1}, and ρ represents the degree of persistence. ε_t is the innovation (shock) that comes from the following normal distribution

$$\varepsilon_t \sim N(0, \sigma_t^2) \quad \text{with} \quad \sigma_t^2 = a + \phi[\varepsilon_{t-1}]^2.$$

The ARCH income process is characterized by parameter ϕ, which is assumed to be nonnegative. If $\phi = 0$, then $\sigma_t^2 = a$, e.g., income risk is constant across time and individuals. If $\phi > 0$, then the variance

Figure 1 Income Level (Left Panel) and Income Variance (Right Panel)

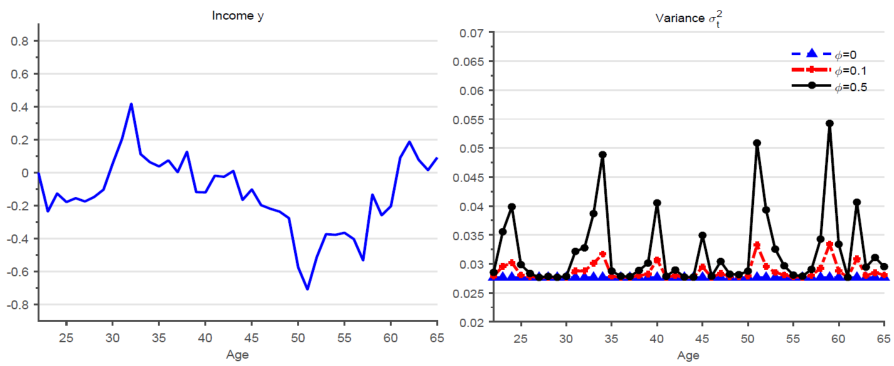

Notes: The figure plots a simulated path of income and volatility for a single worker. Volatility paths are plotted for different values of ϕ.

at time t depends on $\varepsilon_{t-1} = y_{t-1} - \rho y_{t-2}$. If at time $t-1$ the worker draws a y_{t-1} that is very different than the expected draw ρy_{t-2}, then σ_t^2 increases. As a result, the worker will receive an income draw at time t from a relatively wide income distribution. In contrast, if at time $t-1$ the worker draws a y_{t-1} that is similar to the expected draw $\rho y_{t\ \ 2}$, then the worker will receive an income draw at time t from a relatively narrow income distribution.

In Figure 1, I use simulated data and plot a time-series path for log income y (left panel) and the corresponding path for σ_t^2 (right panel). I plot the variance for three different values of $\phi = \{0.0, 0.1, 0.5\}$. Note that a different value of ϕ also implies a different shock ε and thus a different y. However, in all cases, the life-cycle path of income is quite similar, so for simplicity, I plot only one path (the one for $\phi = 0.1$). The values for ρ and a correspond to the parametrization I choose for our main model and discuss in Section 2. When $\phi = 0$, income variations have no bearing on the variance. When $\phi > 0$, during periods of large income variations (e.g., large $\varepsilon_{t\ \ 1}$), the income variance increases. And the larger ϕ is, the larger the change in variance. The increases in the variance are usually short-lived. This occurs because y shocks are relatively persistent (ρ is set to 0.74), and ε does not deviate persistently from zero.

This simple framework gives some idea about the additional state variables we need to solve the model. While in models with constant income variance (e.g., $\phi = 0$), the worker relies on y_{t-1} to form a forecast about y_t, in ARCH the necessary information also includes y_{t-2} in order to also estimate the variance σ_t^2. Therefore, the model requires an additional state variable. Note that a GARCH specification—one of the most popular extensions of ARCH—models the conditional variance as $\sigma_t^2 = a + \rho_\sigma \sigma_{t-1}^2 + \phi[\varepsilon_{t\ \ 1}]^2$. In this case, we would need to add one more state variable (relative to ARCH), namely σ_{t-1}^2, which would significantly increase the computational cost.

2. LIFE-CYCLE MODEL

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:

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

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

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

$$\log Y_{ij} = z_j + y_{ij} \quad \text{with} \quad y_{ij} = a_i + \beta_i \times j + x_{ij}. \tag{2}$$

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:

$$x_{ij} = \rho_x x_{i,j-1} + \nu_{ij} \qquad \text{with} \quad \nu_{ij} \sim \text{ i.i.d. } N(0, \sigma_{ij}^2). \tag{3}$$

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

Matrix for Income Process We define the matrices M_{j-1} and H_j that allow us to define the recursive problem in terms of income y, as follows:

$$\mathbf{M}_{j-1} = \begin{bmatrix} a \\ \beta \\ \rho_x x_{j-1} \end{bmatrix}, \quad \mathbf{H}_j = \begin{bmatrix} 1 \\ j \\ 1 \end{bmatrix}. \tag{4}$$

The following period's M_j is:

$$\mathbf{M}_j = \mathbf{R} \left[\mathbf{M}_{j-1} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} (y_j - \mathbf{H}_j' \mathbf{M}_{j-1}) \right] \tag{5}$$

with R denoting a (3×3) matrix whose diagonal elements are $(1, 1, \rho_x)$. Note that, $H_j' M_{j-1}$ is the conditional expectation of period j's labor income as of age $j-1$. Moreover, $y_j - H_j' M_{j\ 1} = x_j - \rho_x x_{j\ 1} = \nu_j$, 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 $H_j' M_{j-1}$ and variance σ_j^2 (denoted as $F(y_j \mid H_j' M_{j-1}, \sigma_j^2)$).

Variance of Labor Income The idiosyncratic labor-income volatility is assumed to follow an ARCH process (which is also used in Meghir and Pistaferri, 2004). The individual variance of labor-income growth is:

$$\sigma_{i,j+1}^2 = \sigma_\nu^2 + \phi(y_j - \mathbf{H}_j' \mathbf{M}_{j-1})^2. \tag{6}$$

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

Savings There is a single asset used for savings: a risk-free bond, b, (paying a gross return of $1 + r$ in consumption units). For simplicity, I abstract from the general equilibrium aspect by assuming exogenous average rates of return. 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 3). Second, it runs a Social Security system. After a worker retires from the labor market at age j_R, the worker receives a Social Security benefit. To avoid the computational complexity of tracking one more state variable (history of earnings), I 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}.

Value Functions The state variables include workers' bonds (b), current income (y_j), and the expected income (M_{j-1}). The value function of a worker at age j is:

$$
\begin{aligned}
V_j(b, y_j, \mathbf{M}_{j-1}) = \max_{c,b'} & \left\{ \frac{c_j^{1-\gamma}}{1-\gamma} \right. \\
& \left. + \ \delta s_j \int_{y_{j+1}} V_{j+1}(b', y_{j+1}, \mathbf{M}_j) dF(y_{j+1} | \mathbf{H}'_{j+1}\mathbf{M}_j, \sigma^2_{j+1}) \right\} \\
\text{s.t.} \quad c + b' & = [(1-\tau_{ss})Y_j - T(Y_j)] \times \mathbf{1}\{j < j_R\} \\
& + ss(Y_{J_R-1}) \times \mathbf{1}\{j \geq j_R\} + \ (1+r)b \\
& \mathbf{M}_j \text{ is given by Equation (5)} \\
& F(y_{j+1} | \mathbf{H}'_{j+1}\mathbf{M}_j, \sigma^2_{j+1}) \text{ is the probability distribution} \\
& \text{for the next period income given } \mathbf{M}_j, \sigma^2_{j+1} \\
& \sigma^2_{i,j+1} = \sigma^2_\nu + \phi(y_j - \mathbf{H}'_j\mathbf{M}_{j-1})^2 \\
& b' \geq \underline{b},
\end{aligned}
$$

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

Discussion of State Variables We need to track the following state variables: b, the bond holdings, y_j, income, and $M_{j-1} = [a, \beta, \rho_x x_{j-1}]$. The combination of y_j and $M_{j\ 1}$ gives us σ^2_{j+1} (see Equation (6)). The additional state variable that is required from the ARCH specification is x_{j-1}. This suggests that we keep a one-period history of earnings. The combination of y_j and $x_{j\ 1}$ (as well as the worker's type a, β) reveals the innovation at age j, ν_j, and thus, the income variance at age $j+1$. If we had a constant volatility environment ($\phi = 0$), then we wouldn't need to track x_{j-1}.

3. CALIBRATION

The model is calibrated using information from Chang, Hong, Karabarbounis, Wang, and Zhang (2020), who rely on administrative data from

Statistics Norway. Earnings data of U.S. households are typically available from household surveys, so they are subject to measurement error. Also it is more difficult to analyze higher-order dynamics of income—which are usually driven by high earners—since surveys are often top-coded (Guvenen, Karahan, Ozkan, and Song, 2015).

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 $\{r, \underline{b}\}$, (iv) labor-income process $\{z_j, \sigma_a^2, \sigma_\beta^2\}$, and (v) tax and transfers $\{\tau_1, \tau_2, \tau^*, \tau_{ss}, ss\}$. For the key parameter ϕ, I perform comparative statics in the range of $\{0.0, 0.1, 0.5\}$.

Table 1 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 eighty periods, $J = 80$. This life cycle corresponds to ages twenty-one to one hundred. Workers retire at $j_R = 45$ (age sixty-five) when they start receiving the Social Security benefit. I estimate the survival probability $\{s_j\}$ at each age using the data on mortality rates from Statistics Norway. Using the estimates from Klovland (2004), I calibrate a real return for the risk-free rate to 1.43 percent.

Table 1 Calibrated Parameters

Parameter	Variable	Value
Life Cycle	J	80
Retirement Age	j_R	45
Risk-Free Rate	$R-1$	1.43%
Tax Function	τ_1	0.73
Tax Function	τ_2	0.16
Tax Function	τ^*	0.85
Tax Function	Y^*	1.7
Survival Probability	$\{s_j\}$	Chang, Hong et al. (2020)
Common Age-Earnings Profile	$\{z_j\}$	Chang, Hong et al. (2020)
Number of Volatility Regimes	N	7
Variance of Fixed Component	σ_a^2	0.057
Variance of Growth Component	$\sigma_\beta^2 \times 100$	0.0088
Persistence of Level Shocks	ρ_x	0.74
Variance of Level Shocks	σ_ν^2	0.027
ARCH Coefficient	σ_ν^2	0.38
Discount Factor	δ	0.95
Risk Aversion	γ	2
Credit Limit	$\underline{b}$	-0.10

To compute the amount of tax and transfers, I 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 United States (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. The tax system in Norway is very progressive. To capture the high progressivity of the Norwegian tax system, I add the term $1_{\{Y^*>Y\}}\tau^*(Y-Y^*)$. We show in Chang, Hong, Karabarbounis, Wang, and Zhang (2020) that with our detailed administrative data, we can calibrate all parameters based on information on before- and after-tax labor earnings.

The Social Security benefit is calibrated to replicate the average benefit for each labor-income decile in the data. As mentioned, in the model, I condition the Social Security benefit on the earnings received in the last working year before retirement. For consistency in the data, I find the relationship between Social Security benefits and labor income during ages sixty to sixty-five. I find that a worker with the mean labor income during ages sixty to sixty-five receives a benefit equal to 36 percent of his or her preretirement labor income. A worker with twice the mean labor income during ages sixty to sixty-five receives around 55 percent of preretirement labor income.

The common age profile of income (z_j) is calibrated based on the age profile of real wages in Norway from the OECD. The real wages for thirty-, forty-, and fifty-year-old workers are on average approximately 5, 15, and 20 percent higher, respectively, than those of twenty–five-year-old workers. Finally, I assume that there are seven regimes for income volatility: $N = 7$.

Calibration of parameters $\{\sigma_a^2, \sigma_\beta^2\}$ follows the estimation technique described in Chang, Hong, Karabarbounis, Wang, and Zhang (2020), which closely follows Guvenen (2009). The variance of the fixed-effect component is $\sigma_a^2 = 0.057$; the variance of the slope is $\sigma_\beta^2 = 0.0088$ percent, and the average variance of the idiosyncratic shocks is $\sigma_\nu^2 = 0.027$. The persistence of the idiosyncratic shocks is $\rho_x = 0.74$. Guvenen and Smith (2014) estimated these parameters using the U.S. data. Similarly to the values here, they find a fairly large idiosyncratic growth component, implying a mildly persistent process for the shock to the income level. But in the 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 United States relative to Norway. I also set the discount factor $\delta = 0.95$, the risk aversion $\gamma = 2$, and the value of borrowing constraint $\underline{b} = -0.10$, which suggests that households can borrow up to 10 percent of their annual income. Finally, Meghir and Pistaferri (2004) report a value of 0.38 for ϕ in a sample of college graduates.[1]

Quantitative Results

To solve the model, I simulate 20,000 individuals and track their lifetime decisions. For every age group, I calculate the following statistics: the variance of log income, the variance of log consumption, the mean assets, and the mean consumption. Figure 2 plots these cross-sectional statistics. The upper left panel plots variance of log earnings across ages. The model generates a familiar increasing profile. This occurs due to (i) persistent shocks accumulating (component x) and (ii) the profile heterogeneity β. The increase in the cross-sectional variance of log income is consistent with data from the United States (Guvenen and Smith, 2014) and Norway (Chang, Hong, Karabarbounis, Wang, and Zhang, 2020). The upper right panel of Figure 2 plots the cross-sectional variance of log consumption across ages. The increase in consumption variance is around 10 log points lower than the increase in income variance. This arises because (i) the government is taxing part of earnings and (ii) households can to some extent self-insure by accumulating assets or borrowing when a negative shock hits. Asset accumulation is depicted in the lower left panel. Households accumulate assets to insure against negative income shocks (precautionary savings) and to prepare for retirement. The assets-to-income ratio is around 2, which means that I exclude some components of wealth, such as housing, in our measurement. Finally, the consumption path (lower right panel) is increasing and concave during working years.

Comparative Statics

I analyze how the life-cycle profiles vary with the ARCH coefficient ϕ. Higher ϕ increases the variance almost in a parallel way. Thus, incorporating an ARCH specification affects more the overall variance of log earnings but less the increase between ages twenty-five and fifty-five. The upper right panel of Figure 3 plots the cross-sectional variance of log consumption across ages. Although income variance increases with respect to ϕ, consumption variance decreases with respect to ϕ. This occurs due to asset accumulation (lower left panel). A higher income risk generates a larger incentive to accumulate assets. As a result, work-

[1] For high school dropouts, ϕ is estimated to 0.19 and for high school graduates to 0.66. They also estimate conditional variance coefficients for permanent shocks that are not part of our specification.

Figure 2 Life-Cycle Paths: Benchmark Calibration

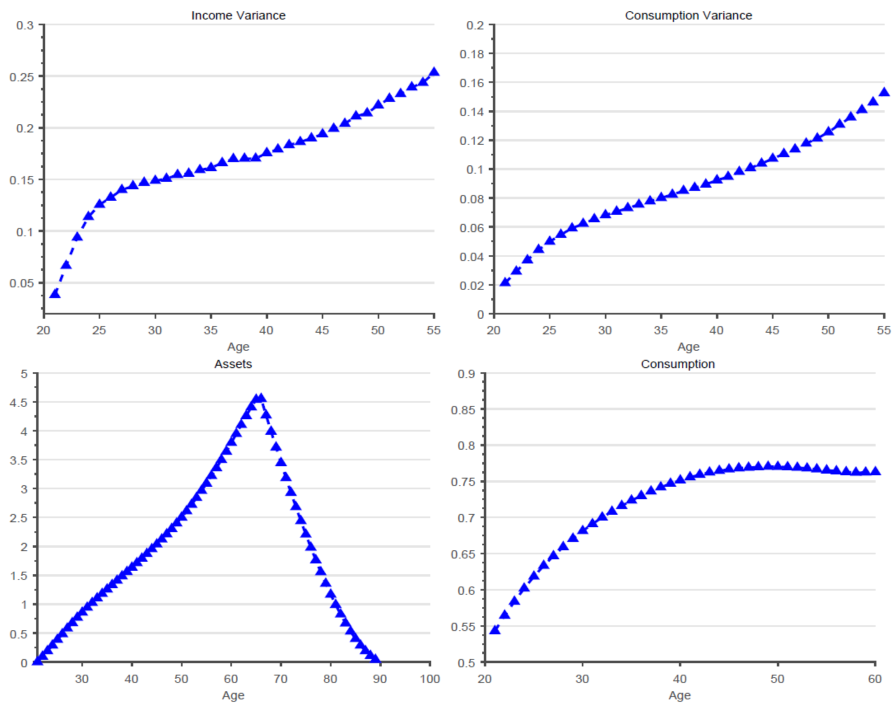

Notes: The upper left panel shows the cross-sectional variance of log earnings and the upper right panel of log consumption. The lower left panel shows average assets by age and the lower right panel average consumption by age.

ers can insure themselves better against income fluctuations, resulting in a lower consumption variance. The consumption path (lower right panel) also confirms that with higher ϕ, workers save more and consume less in the beginning of the life cycle. In sum, higher ϕ results in a high-risk economy with higher income inequality but not necessarily consumption inequality due to a higher degree of assets.

In the above experiments, when we changed ϕ, the overall income variance increased. Here, we set $\phi = 0$, but we increase σ_ν^2 from 2.7 percent to 4.2 percent so that the income variance profile remains unaffected as in the benchmark calibration (with $\phi = 0.38$). Figure 4 plots the income variance across the two calibrations (left panel), the consumption variance (middle panel), and the asset holdings. Even though the parameters σ_ν and ϕ differ between the models, the life-

Figure 3 Life-Cycle Paths: Different Values for ϕ

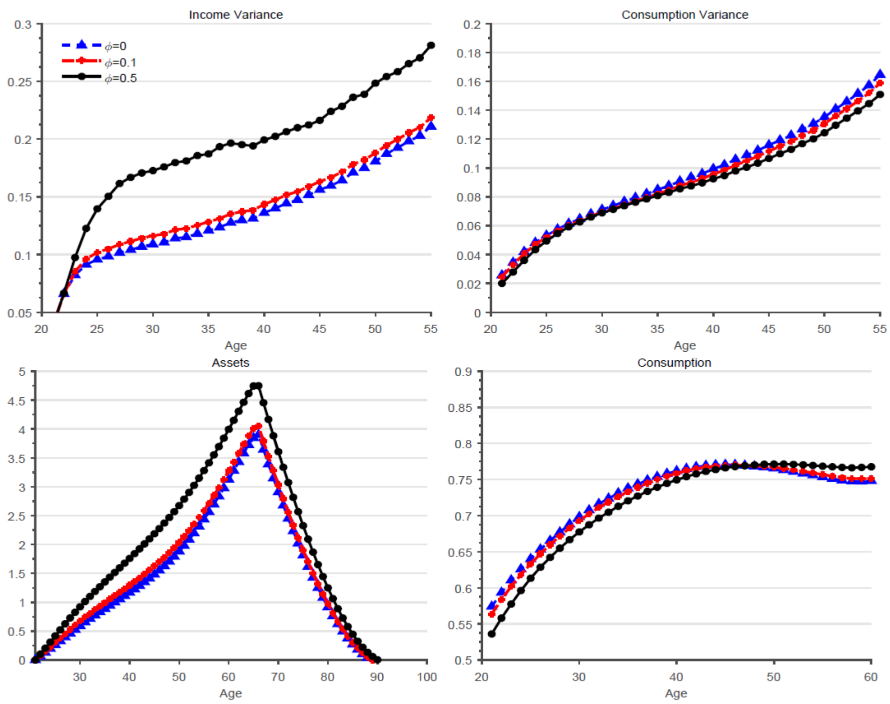

Notes: The upper left panel shows the cross-sectional variance of log earnings and the upper right panel of log consumption. The lower left panel shows average assets by age and the lower right panel average consumption by age.

cycle profiles are broadly similar. So it turns out the household does not care about which component is driving the conditional variance but only about its overall level.

4. CONCLUSION

Typical models with heterogeneous agents treat income risk as constant. In this article, I incorporate a dynamic income process with idiosyncratic, time-varying volatility into an otherwise standard life-cycle model of consumption and savings. This is consistent with empirical evidence on earning dynamics (Meghir and Pistaferri, 2004; Guvenen, Karahan, Ozkan, and Song, 2015). Idiosyncratic volatility is modeled as an ARCH specification. I show that this specification can be in-

Figure 4 Life-Cycle Paths: Conditional Variance Parameters

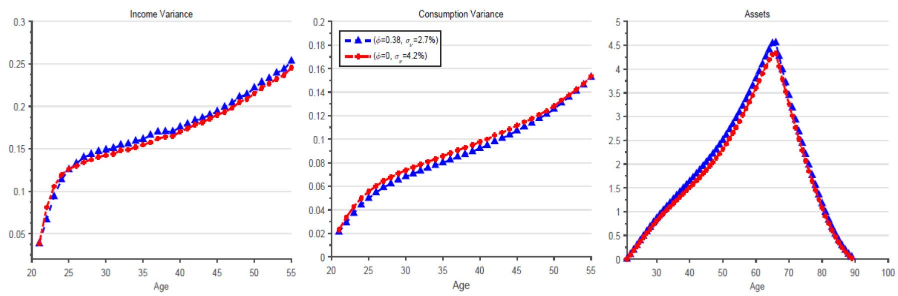

Notes: The left panel shows the cross-sectional variance of log earnings; the middle panel shows log consumption, and the right panel shows the average assets over the life cycle.

tegrated naturally in a standard life-cycle model with the addition of a state variable. I analyze how the ARCH process affects the aggregate consumption and savings patterns. I show that when volatility is more sensitive to income fluctuations, the economy-wide risk increases. As a result, workers accumulate more assets to maintain a relatively stable path of life-cycle consumption. This suggests that the ARCH process affects not only individual dynamics, but also has aggregate implications.

REFERENCES

Chang, Yongsung, Jay H. Hong, Marios Karabarbounis, Yicheng Wang, and Toa Zhang. 2020. “Income Volatility and Portfolio Choices.” Federal Reserve Bank of Richmond Working Paper 20-01, revised September 2020.

Guvenen, Fatih. 2007. “Learning Your Earning: Are Labor Income Shocks Really Very Persistent?” *American Economic Review* 97 (June): 687–712.

Guvenen, Fatih. 2009. “An Empirical Investigation of Labor Income Processes.” *Review of Economic Dynamics* 12 (January): 58–79.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2015. “What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk.” National Bureau of Economic Research Working Paper 20913.

Guvenen, Fatih, and Anthony A. Smith Jr. 2014. “Inferring Labor Income Risk and Partial Insurance from Economic Choices.” *Econometrica* 82 (November): 2085–129.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2014. “Consumption and Labor Supply with Partial Insurance: An Analytical Framework.” *American Economic Review* 104 (July): 2075–126.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2017. “Optimal Tax Progressivity: An Analytical Framework.” *Quarterly Journal of Economics* 132 (November): 1693–754.

Klovland, Jan Tore. 2004. “Bond Markets and Bond Yields in Norway 1819–2003.” In *Historical Monetary Statistics for Norway 1819–2003*, edited by Øyvind Eitrheim, Jan Tore Klovland, and Jan F. Qvigstad. Oslo, Norway: Norges Bank.

Meghir, Costas, and Luigi Pistaferri. 2004. “Income Variance Dynamics and Heterogeneity.” *Econometrica* 72 (January): 1–32.