Heterogeneity in returns to wealth and consumption inequality

Claudio Daminato[*]and Luigi Pistaferri[†]

February 2023

Abstract

A recent literature argues that type and scale dependence in wealth returns may play an important role for explaining features of the wealth distribution. Using panel data from the PSID, we first document that a common individual component (which we identify as the stock of endowed cognitive skills) appears to drive persistent heterogeneity in wealth returns and earnings. We then embed return heterogeneity within a life-cycle model of consumer behavior and allow persistent heterogeneity in wealth returns to be correlated with persistent heterogeneity in wages. We show that eliminating persistent return heterogeneity or common factors in both earnings and returns would dramatically understate average returns for people at the top of the wealth distribution as well as the level and rise of consumption inequality over the life cycle.

1 Introduction

A recent literature (Benhabib and Bisin, 2018; Gabaix et al., 2016) argues that features of the distribution of stochastic wealth returns (persistent heterogeneity, or type dependence, and a positive correlation with wealth, or scale dependence) may play an important role for explaining features of the wealth distribution, such as its long left tail, as well as the increase in top shares over time documented in e.g., Saez and Zucman (2016). Several papers show empirically that stochastic wealth returns indeed display type and scale dependence (see Fagereng et al., 2020, and Bach et al., 2020).

In this paper we make two contributions. The first is to embed return heterogeneity within an otherwise standard life-cycle model of consumer behavior. The second is to allow

[*]Lund University.

[†]Stanford University, SIEPR, NBER and CEPR.

persistent heterogeneity in wealth returns to be correlated with (the more frequently documented) persistent heterogeneity in wages. We argue that such correlation may arise from common factors (cognitive and non-cognitive unobserved skills or abilities) driving both.

We estimate the parameters of interest (extent of persistent heterogeneity in wealth returns and earnings, as well as preference parameters) using panel data from the PSID. The data confirm findings from other countries that both type and scale dependence characterize the behavior of returns from assets. The data also show that a common component (which we identify as the stock of endowed cognitive skills) appears to drive persistent heterogeneity in wealth returns and earnings: individuals who do persistently well in labor markets appear to do persistently better in asset markets as well. The model replicates well the rise over the life cycle of the wealth-income ratios and of consumption inequality. In counterfactual exercises, we document that eliminating persistent return heterogeneity or common factors in both earnings and returns would dramatically understate average returns for people at the top of the wealth distribution as well as the level and rise of consumption inequality over the life cycle.

2 Key Facts from Household Panel Data

We start by documenting some key facts about household wealth returns and household earnings using panel data from various waves (1998-2018) of the Panel Study of Income Dynamics (PSID). During this period the survey was conducted bi-annually. Besides demographic information, the PSID routinely collects information on labor earnings of household members (husbands and wives in traditional couples), as well as detailed information on household assets and liabilities as well as sources of capital income and debt payments.

2.1 Returns to Wealth from the PSID

We construct two measures of returns to wealth from the PSID:

1. Returns to net wealth:
$$r_t = \frac{(y_t^c + cg_t - y_t^d)}{(A_{t-1} + 0.5F_t)} \tag{1}$$
where y_t^c are interests and dividends, cg_t the (realized) "capital gains/losses" from business, rents, stocks, real estate, pension/IRA, y_t^d payments on debt, A_{t-1} is total household's net wealth at the beginning of the previous period, and F_t is the net investment flow into businesses, stocks, real estate and pensions (also known as the Dietz' correction, see Dietz, 1968).

2. Returns to gross wealth:

$$r_t^G = \frac{(y_t^c + cg_t)}{(A_{t-1}^G + 0.5F_t)}$$

where $A_{t-1}^G = A_{t-1} + D_{t-1}$ is gross household wealth, with D indicating household debt.

In Figure 1 we plot the estimated return to net wealth (the blue dots) against the percentile of net wealth. The figure show strong evidence for *scale dependence*, at least above the 30-th percentile. Average returns to net wealth increase when we move from the 30-th to the bottom percentile because in this region households have negative net worth and payments on debt exceeding interests and dividends, so that both the numerator and denominator of equation (1) are negative. In the same figure we also plot the return on assets (ROA, the red dots), defined as $r_t^A = \frac{(y_t^c + cg_t - y_t^d)}{(A_{t-1}^G + 0.5F_t)}$, which eliminates this issue: it asks how much net income is generated by one dollar of assets. The degree of scale dependence is substantial: individuals in the top percentile have returns to net wealth (gross wealth) that are approximately 30 (20) percentage points higher than someone with median wealth. There is also an increase in the slope of the relationship (i.e., convexity) as we move to the upper percentiles. In Figure 2 we replicate Figure 1 for the return to gross wealth and find again evidence for scale dependence, at least above the median.

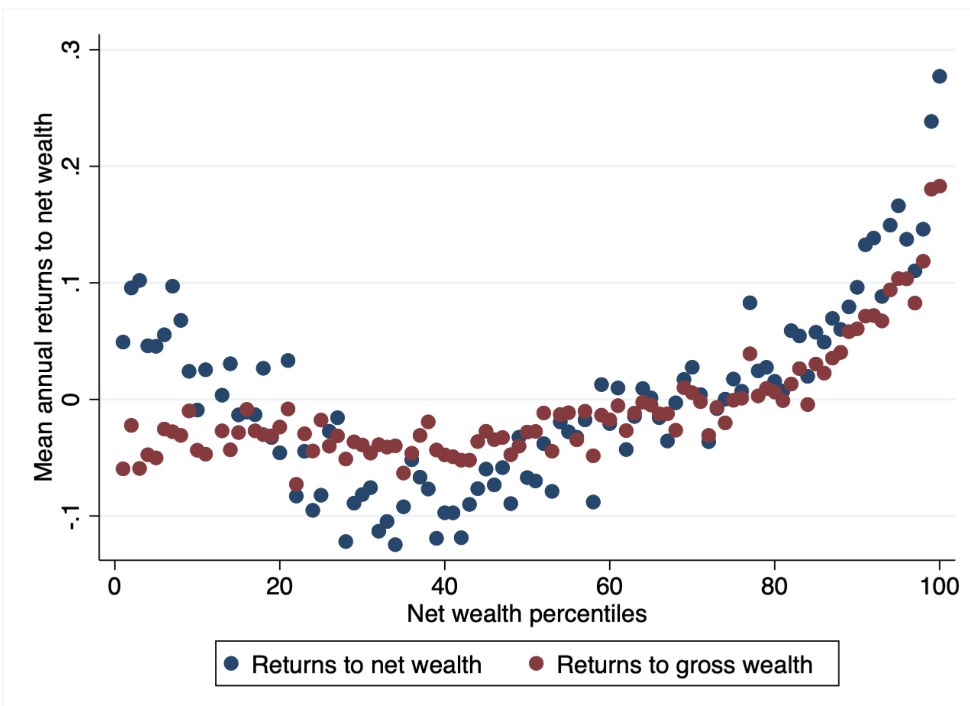

Figure 1: Annualized returns to wealth across the net wealth distribution

To get a gauge of the importance of *type dependence* in wealth returns, in Table 2 we replicate the approach of Fagereng et al. (2020) and regress wealth returns against various controls plus household fixed effects. We do so for two samples: the whole sample and those with non-missing information on intergenerational transfers and the degree of risk aversion. The latter is based on survey questions that elicit people's preferences for risk, see Kimball et al. (2009); however, data on elicited preferences are only available in two waves, 1996 and 2012. The relevance of household fixed effects can be assessed by looking at the increase in the adjusted R^2 when we move from a specification without to one with fixed effects. Depending on the sample, in Table 1 fixed effects increase the explained variation in wealth returns by 20 to 30 percent, controlling for demographics, portfolio composition, and degree of risk aversion.

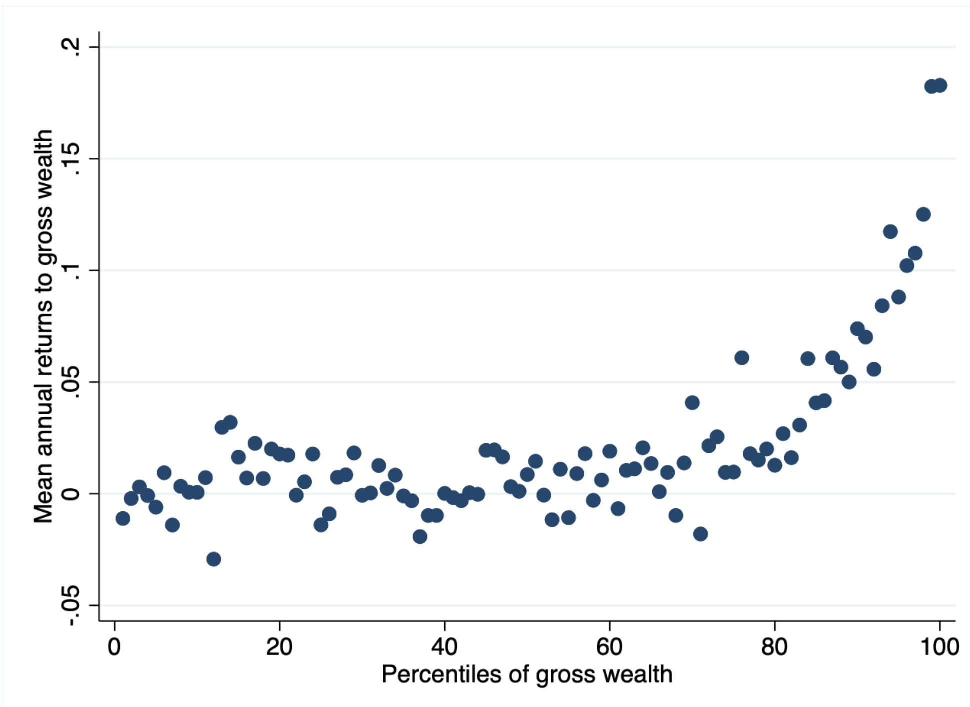

Figure 2: Annualized gross returns to wealth across the gross wealth distribution

Table 1: Fixed effects in returns to wealth

	Non-missing risk avs and transfers				*Whole sample*	
	(1)	(2)	(3)	(4)	(5)	(6)
Baseline controls	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Shares*Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Risk aversion	No	Yes	Yes	No	No	No
Intergenerational transfers	No	No	Yes	No	No	No
Individual FE	No	No	No	Yes	No	Yes
Adj R squared	0.247	0.247	0.248	0.299	0.209	0.272
N	2566	2566	2566	2566	8274	8274

Note: Dependent variable is net returns to wealth. Regressions control for age, education, employment, year and state dummies, share of wealth allocated to different asset classes, leverage of mortgage and other debt, and wealth percentiles.

2.2 The Correlation between Wealth Return Fixed Effects and Wage Fixed Effects

Figure 3 documents the existence of a positive rank correlation between fixed effect in household wealth returns (as estimated in the previous section) and fixed effect in household log earnings. In Figure 4 we replicate the same exercise for log wages (defined as log household earnings divided by log household hours). Both figures are binscatters: they plot the average percentile of the distribution of wealth return fixed effects by percentile of the earnings or wage distribution. Figure 3, for example, shows that a move from the 10-th to the 90-th percentile of the earnings fixed effect distribution would increase the average percentile of the wealth return fixed effect from the 40-th to the 70-th, a rather large effect. The rank correlation is 0.498. The fact that the Figures are quite similar suggest that hour choices are not the ones driving the correlation. For simplicity, in the model below we assume hours are chosen exogenously.

In Table 2 we estimate the rank correlation for different subgroups in the population. It ranges from 0.38 for households whose heads have high school education or less, to 0.542 for households whose heads are non-white.

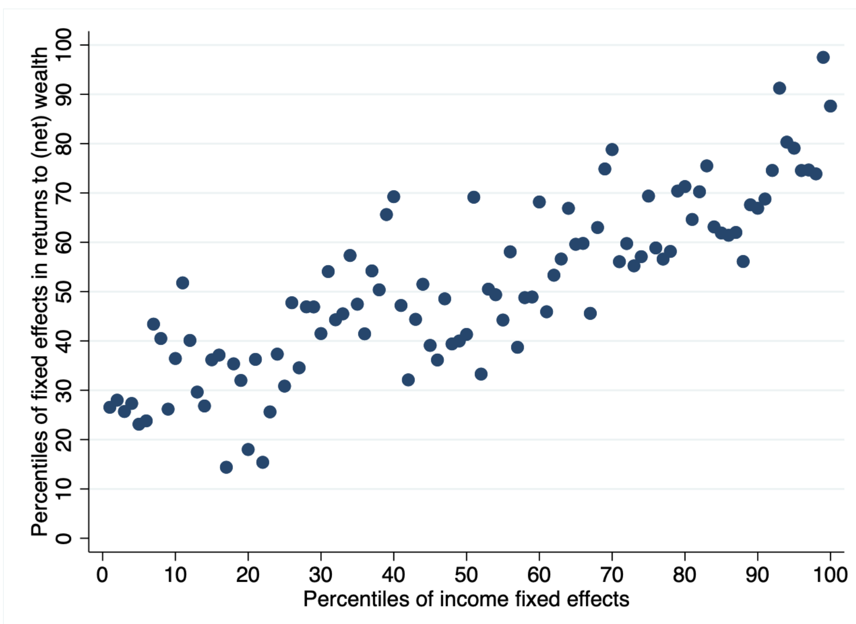

Figure 3: Average percentile of net returns to wealth across the distribution of earnings fixed effects

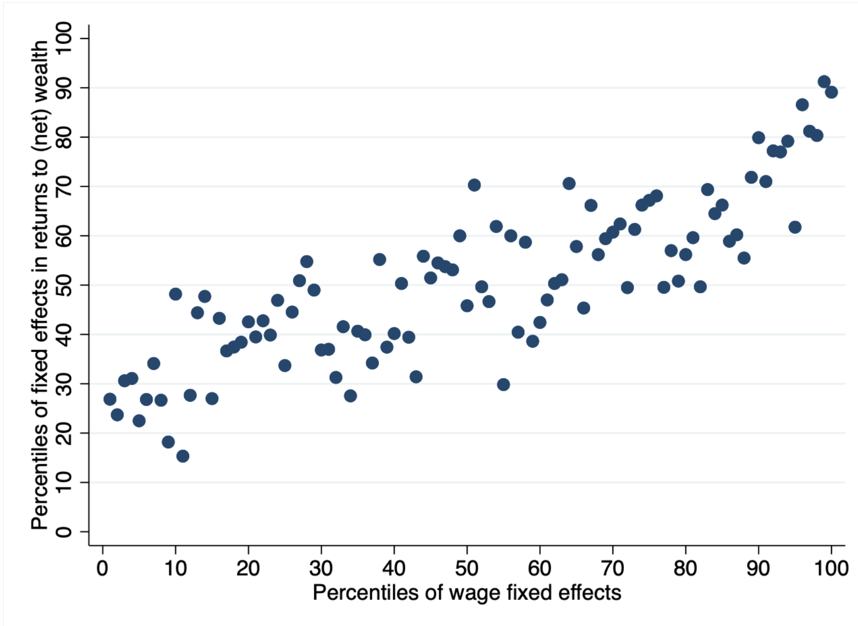

Figure 4: Average percentile of net returns to wealth across the distribution of wage fixed effects

Table 2: Rank correlation between permanent component of wages and returns to wealth

	Education		Race		Total
	High school or below	College	White	Other	
	(1)	(2)	(3)	(4)	(5)
Total labor income					
$\rho(P_i, \psi_i)$	0.380	0.438	0.467	0.542	0.498
N	310	692	807	194	1002

Note: This Table reports the rank correlation between fixed effects in net returns to wealth and fixed effects in total household labor income across sub-groups of the population.

2.3 What explains the correlation between wages and returns fixed effects?

Both cognitive and non-cognitive individual characteristics can potentially explain the positive rank correlation between returns to wealth and wages fixed effects in Figure 4 and Table 2. Suppose we write the wealth return fixed effect (ψ_i) and the earnings fixed effect (P_i) as:

$$\psi_i = f(\Theta_i, \Gamma_i)$$
$$P_i = f(\Theta_i, \Gamma_i)$$

where Γ_i is a set of non-cognitive individual characteristics and Θ_i can be thought of as the individual's ability or genetic endowment. Non-cognitive factors potentially influencing both individual's wages and returns to wealth include, among other things: (i) individual's degree of risk aversion; (ii) cohort effects; (iii) intergenerational transfers. We can test for the importance of individual's risk aversion and intergenerational transfers on the observed correlation between wages and returns fixed effects using data from the 1996 and 2012 waves of the PSID, respectively. The test is carried out by partialling out alternative potential explanations with a regression approach. We estimate the following equations in the cross-section:

$$\psi_i = \alpha_\psi + \gamma\Gamma_i + \eta_i^\psi$$
$$P_i = \alpha_P + \delta\Gamma_i + \eta_i^P$$

where the residuals η_i^ψ and η_i^P can be interpreted as the components of persistence in returns and wages due to individual's cognitive characteristics. We can then use η_i^ψ and η_i^P to test for the degree of rank correlation between wages and returns to wealth fixed effects that is driven by cognitive factors. Results are reported in Table 3. They show that the largest chunk of the rank correlation is due to the unobserved cognitive component (as well as, possibly, non-cognitive components not perfectly captured by the observables). We put some structure on the role of these components in the theoretical model below.

Table 3: Residual rank correlation after partialling out alternative explanations

	Baseline	cohort effects	cohort effects and risk avs	cohort effects and transfers	cohort effects risk avs and transfers
	(1)	(2)	(3)	(4)	(5)
Male earner					
$\rho(\eta^P, \eta^\psi)$	0.498	0.479	0.481	0.412	0.438
N	1,002	1,002	406	908	355
$\rho(P, \psi)$ subsample		0.498	0.532	0.493	0.539

2.4 Persistent Wealth Returns, Earnings, and Consumption

The final key facts connect fixed effects in labor income and in wealth returns to fixed effects in household consumption. Households with persistently high wealth returns or income tend to be households with persistently high consumption, controlling for observables such as age, education, and the like.

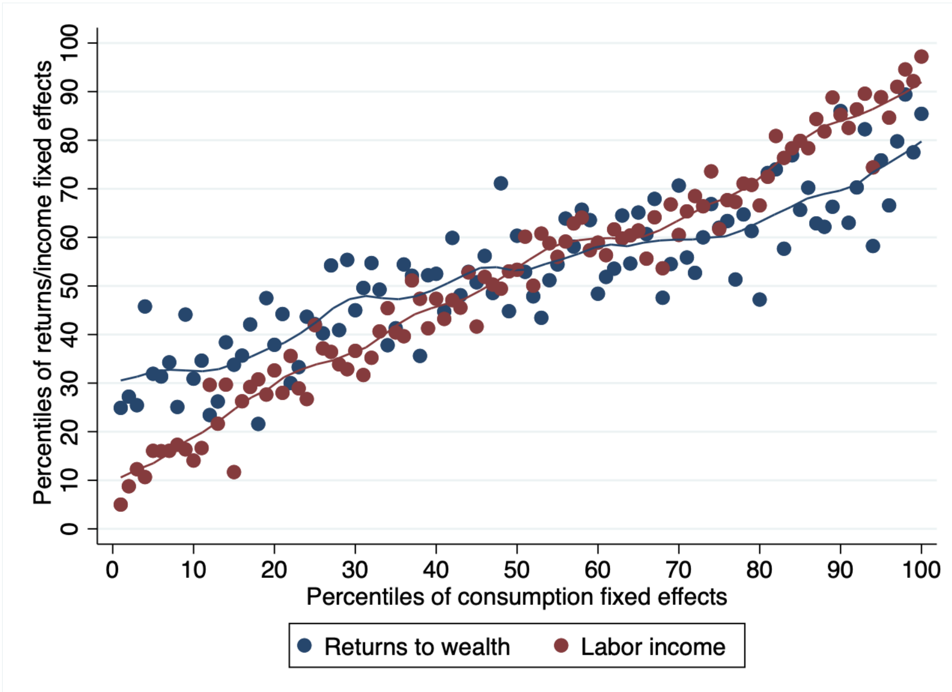

Figure 5: Rank correlation between fixed effect in consumption and returns and labor income fixed effects

3 Model setup and assumptions

Motivated by the empirical evidence presented above, we develop a life-cycle model to explore the effects of persistent wealth returns, and their correlation with the permanent component of wages, on the level of consumption inequality.

The stochastic life-cycle model builds on Deaton (1991), Carroll (1997), Attanasio et al. (1999) and Gourinchas and Parker (2002). We contribute to previous quantitative models of wealth inequality (De Nardi and Fella, 2017) by introducing two sources of individual heterogeneity: (i) initial level of permanent labor income and (ii) cognitive endowment. Further, we allow for correlation between the two components of persistent heterogeneity. In the model, individual cognitive endowment affects both the household's returns to wealth as well as its labor income. Households face uncertainty with respect to financial assets and human capital returns, out-of-pocket medical expenditures as well as length of life. The time unit is one year.

Heterogeneity Besides the ex-post heterogeneity induced by the realization of income and asset returns shocks, we allow for further (ex-ante) heterogeneity with respect to educational attainment. Different education groups face different age profiles of labor income over their working lives.

Preferences The utility function is intertemporally separable. The period utility function is:

$$u(C_t; z_t) = q(z_t)\frac{\widetilde{C}_t^{1-\gamma}}{1-\gamma}$$

where C_t is consumption, $q(z_t)$ is a function of demographic shifters to account for the evolution of households composition over the life-cycle, $\widetilde{C}_t = \frac{C_t}{q(z_t)}$.

Length of life Households live at most until age T, but can die before. Therefore, length of life is uncertain. To model the uncertainty of the length of life, we denote as d_t, the probability that the household is alive in period $t+1$, conditional on being alive in period t.

Bequests If households die at age t, the remaining wealth, A_t, is left to their heirs. As in De Nardi (2004), households value bequests according to the bequest function $b(A_t) = \theta\frac{(A_t+k)^{1-\gamma}}{1-\gamma}$, where θ is the intensity of the bequest motive and k the parameter controlling the curvature of the bequest function.

Social security Households retire exogenously at age 65 and start drawing public pension benefits Y^p. Public pension benefits are computed as a non-linear function of individual's lifetime earnings H.

Cognitive endowments Household i starts her life with an endowment of cognitive skills $\Phi_{i,0}$ which (in the baseline) remains constant over the life-cycle, $\Phi_{i,t+1} = \Phi_{i,t}$ for each $t \in (0, T)$. The endowment of skills is heterogeneous in the population and it has variance $\sigma^2_{\Phi_i}$.

We posit that higher cognitive endowments increase returns to wealth (e.g., by giving access to better investment opportunities or by decreasing the probability of making financial mistakes) and enter the process of human capital formation.

Earnings Each period of their working life, households receive gross labor earnings $Y_{i,t}$. We write the log of real labor income of household i at time t as:

$$\log Y_{i,t} = X'_{i,t}\beta^y_j + P^y_i(\Phi_{i,0}) + u_{i,t} + \sum_{j=1}^{t} v_{i,j} \tag{2}$$

where P^y_i is a permanent individual component (which we assume to be a function of the stock of cognitive endowment $\Phi_{i,0}$), $X_{i,t}$ are observed characteristics of earners in household i that affect wages, and $u_{i,t}$ and $v_{i,t}$ are transitory shocks and permanent i.i.d. shocks to earnings, respectively, with constant variances. The permanent individual component depends on both an individual fixed effect ω^y_i (independent on individual's cognitive abilities), which has mean zero and variance $\sigma^2_{y_i}$, as well as on the individual cognitive endowment. We hence write the permanent individual component of labor income as:

$$P^y_i = \exp(\omega^y_i + \delta^y \Phi_{i,0})$$

Financial assets returns Households allocate their wealth, A_t, between a riskless asset B_t and a risky asset S_t. We assume costly collection and processing of the financial information needed to access the return from the risky asset. Since the access decision is made on a period basis, households pay a per-period fixed cost, κ, to hold the risky assets (as in, e.g., Fagereng et al., 2017). Moreover, we assume borrowing constraints (non-negative share of the riskless asset) and short-sale constraints (non-negative share of risky asset): hence, the share of risky assets, $\alpha^s_t = S_t/A_t$, lies between zero and one. It depends, among other things, on the degree of risk aversion.

The return from a household's portfolio can then be written as:

$$r_{i,t}^{p} = r^{b} + \alpha_{i,t-1}^{s}(r_{i,t}^{s} - r^{b}) \tag{3}$$

We assume that the individual return on the risky asset $r_{i,t}^{s}$ depends on a persistent individual component ψ_i, which in turn is a function of the stock of cognitive endowment:

$$\psi_i = \omega^{r} + \delta^{r}\Phi_{i,0} \tag{4}$$

Since the stock of cognitive endowment enters both the persistent component of earnings and the persistent component of wealth returns, we need to impose a normalization for identification purposes, and hence assume $\delta^{r} = 1$.

Combining (3) and (4), the excess return on the risky asset evolves according to:

$$r_{i,t}^{s} - r^{b} = (\mu_S + \omega^{r}) + \delta^{r}\Phi_{i,0} + \xi_{i,t}^{s} \tag{5}$$

where $\mu_S > 0$ and $\xi_{i,t}^{s}$ are independently and identically distributed according to $\mathcal{N}(0, \sigma_S^2)$. As in Fagereng et al. (2017), we also allow for tail risk in the risky assets return distribution: the return in the tail event is r_{tail} and the probability is p_{tail}.

We assume zero correlation between labor income shocks and shocks to risky returns. This assumption is motivated by previous studies that have found weak evidence regarding the correlation between wages and returns from stocks.[1] In contrast, motivated by the empirical evidence documented above, we allow for correlation between household labor income and returns to wealth through the correlation between their permanent components ($Cov(r_{i,t}^{s}, Y_{i,t}) = Cov(\psi_i, P_i^{y})$).

Medical expenses After retirement, households face uncertainty with respect their out-of-pocket medical expenses m_t. We write the log of medical expenses of household i at time t:

$$\log m_{i,t} = m(H, t) + \zeta_{i,t} \tag{6}$$

where the idiosyncratic component $\zeta_{i,t}$ is modeled using a permanent-transitory decomposition, as suggested in French and Jones (2004). We assume that medical expenses reduce income available for consumption but produce no utility benefits.

[1] In particular, using data from the PSID, Cocco et al. (2005) do not reject the null of no correlation between labor income shocks and stock returns.

Taxes Households face a progressive labor income tax schedule. We write the log of after-tax labor income: $\log Y_t^n = (1-\tau)\log Y_t + \log \nu$. Returns to wealth are taxed at a constant rate τ_c.

Government transfers Following Hubbard et al. (1995), we allow for the presence of a consumption floor $\underline{c}$ through means-tested transfers:

$$T_t = \max(\underline{c} + m_t - R^p A_t - Y_t^n) \tag{7}$$

3.1 The households' problem

Households choose consumption and the portfolio share of risky assets to maximize:

$$E_0 \left\{ \sum_{t=0}^{T} \beta^t \left[d_t u(C_t; z_t) + (1-d_t) b(A_t) \right] \right\}$$

where $\beta < 1$ is the subjective discount factor. Before retirement, the dynamic budget constraint reads as:

$$A_{t+1} = \left(1 + r_{t+1}^p\right) A_t + Y_t^n - C_t + T_t - \kappa \times \mathbb{1}(\alpha_t^s > 0) \tag{8}$$

and after retirement as:

$$A_{t+1} = \left(1 + r_{t+1}^p\right) A_t + Y^p - m_t - C_t + T_t - \kappa \times \mathbb{1}(\alpha_t^s > 0) \tag{9}$$

State variables The dynamic optimization problem of the household is characterized by seven state variables (X): age (t), assets (A), cognitive endowment (Φ), initial labor income (ω^y), history of permanent income shocks ($\sum_{j=1}^{t-1} \upsilon_j$), average lifetime earnings H and permanent medical expense shock.

The recursive optimization problem of the household is:

$$V_t(X_t) = \max_{\{C_t, \alpha_t\}} \left\{ u(C_t; z_t) + \beta E_t \left[d_{t+1} V_{t+1}(X_{t+1}) + (1-d_{t+1}) b(A_{t+1}) \right] \right\}$$

subject to equations (3-9).

The solution algorithm combines continuous and discrete choices based on a modification of the algorithm in Iskhakov et al. (2017). Compared to the standard EGM proposed by Carroll (2006), the employed upper envelop algorithm disregards non-optimal consumption emerging due to the presence of kinks in the value function at the points of the state space where the household is indifferent between alternatives in discrete choice space.

3.2 Model estimation

To estimate the model we adopt a standard two-step strategy (Gourinchas and Parker, 2002). In the first step, we calibrate a set of parameters that do need the usage of the model for identification, either estimating them directly in the data or using previous estimates in the literature. These include the parameters characterising the earnings process and out-of-pocket medical expenses, the risky asset returns distribution, the survival probabilities, the demographic shifters, the curvature of the bequest function, the labor income tax schedule and the pension rules. In the second step we use a Minimum Distance estimator to estimate the remaining structural parameters $(\beta, \gamma, \kappa, \widetilde{\theta}, \underline{c}, \sigma^2_{y_i}, \sigma^2_{\Phi_i}, \delta^y)$, taking the set of first step parameters as given.

3.2.1 Identification

The implications of persistent heterogeneity in returns to wealth and its relation with heterogeneity in returns to human capital on consumption inequality critically depend on the extent of (co-)variation in these measures. The goal is to identify $\sigma^2_{y_i}$, $\sigma^2_{\Phi_i}$ and δ^y given the equations:

$$
\begin{aligned}
P_i^y &= \omega_i^y + \delta^y \Phi_{i,0} + e_i^y \\
\psi_i &= \omega^r + \Phi_{i,0} + e_i^r \\
\widetilde{c}_i &= \lambda^r(\omega^r + \Phi_{i,0}) + \lambda^y(\omega_i^y + \delta^y \Phi_{i,0}) + e_i^c
\end{aligned}
$$

where λ^r and λ^y indicate the effects of fixed effects in returns to wealth and labor income, respectively, on the average level of individual consumption. We allow for measurement error in fixed effects of labor income e_i^y, returns to wealth e_i^r, and average consumption e_i^c. The consumption equation can be interpreted as an auxiliary equation for our simulated method of moment estimation.

Estimating $\sigma^2_{y_i}$ and $\sigma^2_{\Phi_i}$ directly from the data is problematic due to the likely presence of measurement error in both fixed effects in returns to wealth and human capital. To separate $\sigma^2_{y_i}$, $\sigma^2_{\Phi_i}$ from the measurement error variances, we make use of the economic model, use consumption data and the covariance between $\widetilde{c}_i$, ψ_i and P_i^y:

$$Cov(\psi_i, P_i^y) = \delta^y \sigma_{\Phi_i}^2 \quad (10)$$

$$Cov(\widetilde{c}_i, \psi_i) = (\lambda^r + \lambda^y \delta^y)\sigma_{\Phi_i}^2 \quad (11)$$

$$Cov(\widetilde{c}_i, P_i^y) = (\lambda^r \delta^y + \lambda^y (\delta^y)^2)\sigma_{\Phi_i}^2 + \lambda^y \sigma_{y_i}^2 \quad (12)$$

We minimize the distance between the empirical covariances (by education group) and the model-predicted covariances. The education-specific parameters $\sigma_{y_i}^2$, $\sigma_{\Phi_i}^2$ and δ^y are therefore just identified. We further augment the set of target moments to include moment conditions that capture the age profile of wealth-to-income ratio, risky asset market participation, and the age profile of the 10th consumption percentile.

3.2.2 Structural parameters estimation

We minimize the distance between target moments estimated in the data and the corresponding moments simulated by the economic model. We target two sets of moments. The first set of moments describe the median behavior of households with respect to wealth accumulation and risky assets participation over the life cycle: median wealth-to-income ratios and average risky assets participation rates in the age groups 25-34, 35-44, 45-54 and 55-64, separately for households with some tertiary education and upper secondary education or less. We run a median regression of the wealth-to-income ratio (linear probability model of risky assets participation) on a third-order polynomial of the head's age (and its interaction with a dummy for the head having attained a college degree), dummies for the number of adults and kids, an education dummy and year fixed effects. We take the predicted conditional median wealth-to-income ratio (participation) by age group and education level. We take a similar approach to obtain the age group specific 10th consumption percentile. The second set of target moments is given by the empirical pairwise covariances between the fixed effects of consumption, returns to wealth and labor income. We make use of the model to construct the simulated counterpart of the empirical target moments. We simulate the behavior of 10,000 households over the life cycle starting from the age of 25. Initial assets are drawn from the empirical distribution in the PSID data for households aged 24-28. The distribution of cognitive endowments and initial income fixed effects are randomly drawn from a log-normal and normal distribution, respectively. To simulate the behavior over the life-cycle, we take random draws from the education-specific earnings process and risky assets returns, out-of-pocket medical expenses and mortality distributions. In the simulated data, we run the auxiliary model:

$$\widetilde{c}_i = \lambda^r \omega^r + (\lambda^r + \lambda^y \delta^y)\Phi_{i,0} + \lambda^y \omega_i^y + \epsilon_i^c$$

to recover estimates of $\widehat{\lambda}^r$ and $\widehat{\lambda}^y$. We then compute the model-predicted FE covariances using equations (9)-(12).

4 Estimation results

Table 4 reports estimates of the structural parameters. The estimates for the discount factor and the CRRA are in the ballpark of what estimated in previous papers. The financial market participation cost is approximately \$1,000, while the consumption floor is \$3267. Finally, the MPB (marginal propensity to bequeath) is 0.78.

In the rest of Table we report, separately by education, estimates of the parameters capturing variation in initial earnings, the stock of cognitive endowment, and the effect of the latter on permanent income. Cognitive skills have a larger impact on the permanent component of the highly educated. Low-educated individuals have lower dispersion in initial earnings and a slightly lower dispersion in cognitive skill endowments.

In Figure 6 we plot selected target moments (the median wealth/income ratio by age group). There is a positive age gradient - wealth grows faster than income over the life cycle.

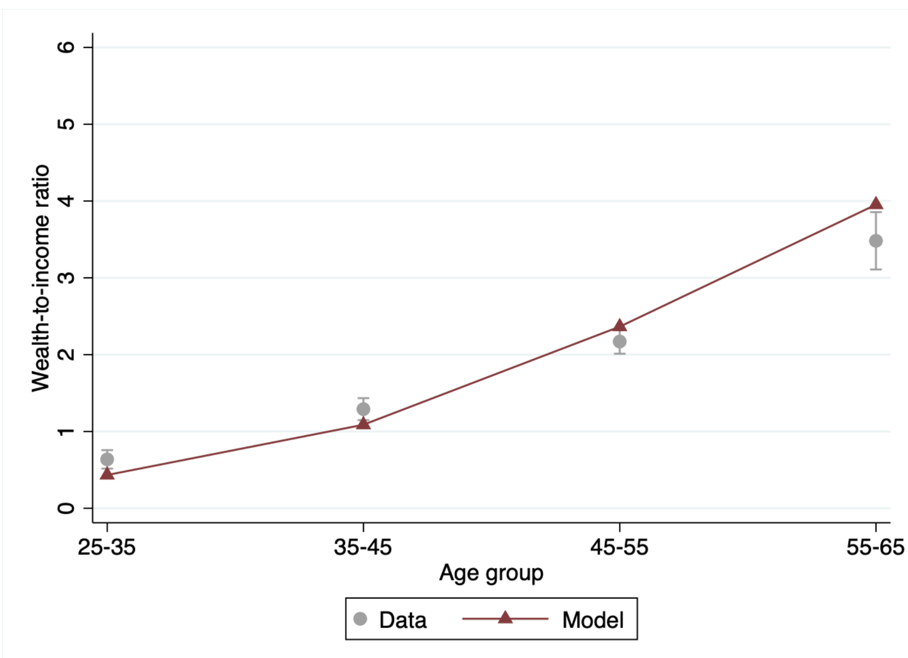

Figure 6: Median wealth to income ratio by age - Model vs. Data

Table 4: Estimated structural parameters

Parameter		Value
Time discount factor	β	0.9787
Coefficient of relative risk aversion	γ	3.325
Financial markets participation cost	κ	1005.98
Consumption floor	$\underline{c}$	3267
Marginal propensity to bequeath	$\widetilde{\theta}$	0.7846
Upper secondary education		
Variance skills FEs	$\sigma^2_{\Phi_i}$	0.0091
Variance initial earnings	$\sigma^2_{y_i}$	0.0275
Effect FEs skills on earnings	δ^y	2.723
Some college degree		
Variance skills FEs	$\sigma^2_{\Phi_i}$	0.0102
Variance initial earnings	$\sigma^2_{y_i}$	0.0231
Effect FEs skills on earnings	δ^y	5.1944

Notes: The estimates are obtained using a minimum distance approach. We minimize the weighted distance between moments of actual and simulated data using the inverse of the diagonal of the bootstrapped variance-covariance matrix of the moments as a weighting matrix. The cost of risky assets participation and the consumption floor are expressed in 2014 dollars.

Model validation Figure 7 is a first validation of the model. We plot median log consumption over the life cycle in the data and as predicted by the model. The model replicates the concave shape, although it tends to underpredict consumption early in life perhaps because borrowing constraints may not be, in practice, as extreme as we have imposed.

Figure 8 moves from explaining averages from explaining inequality in log consumption over the life cycle (as measured by the 75-th/25-th percentile difference). In the data, inequality increases over most of the working life cycle before slowing down in the years before retirement. In the model the increase is well captured, and smoother. The lack of fit at the very end of the life cycle could be explained by early retirement episodes not allowed by the model (where no retirement is allowed before age 65). There is less income inequality (and hence less consumption inequality) at/after retirement since pension income is less volatile than labor income.

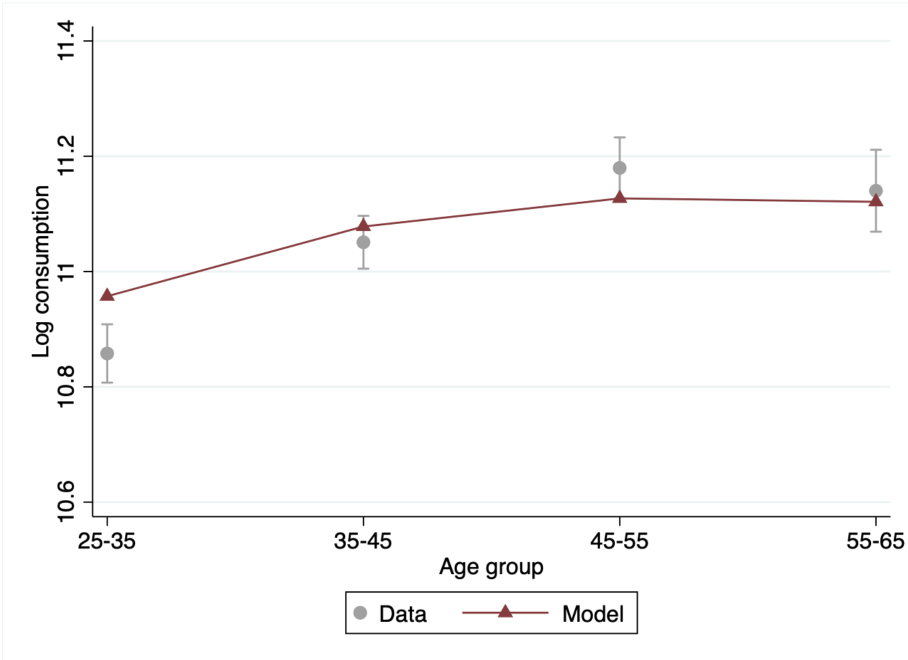

Figure 7: Median consumption by age - Model vs. Data

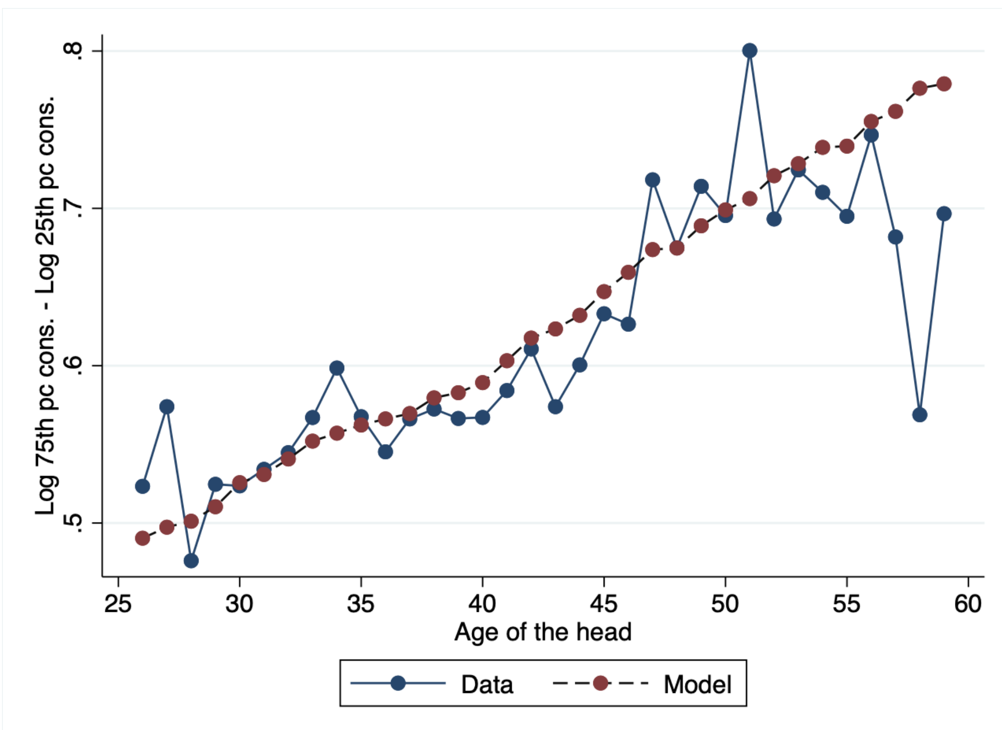

Figure 8: 75th-25th percentile log consumption difference - Model vs. Data

5 Implications

One of the most intriguing questions that the literature on return heterogeneity is confronting is how to explain the degree of scale dependence visible in the data. A traditional explanation is risk taking: because of the insurance provided by their wealth or because they can more easily overcome the cost of participating in stock markets, people at the top of the wealth distribution have an asset portfolio that is more geared towards risky assets, which helps explaining their higher wealth returns. A different, not necessarily alternative explanation, is financial education or sophistication: wealthy individuals may have better analytical skills or knowledge of financial products (here captured by the stock of cognitive endowment) or have the wealth to acquire the services of financial advisors, which may produce higher returns. In Figure 9 we plot the difference between the actual return to wealth and the return to wealth that would be generated in a counterfactual world in which we shut down the contribution of the stock of cognitive endowment. If scale dependence was entirely attributable to risk taking, the figure would be flat at 0. Clearly, that is not the case; more importantly, cognitive factors appear more important at the top - the top percentile of the wealth distribution have wealth returns that are almost 30 percentage points higher than would be predicted by a model in which portfolio choice were driven by portfolio allocation choices between risky and riskless assets.

In Figure 10 we present a different counterfactual. We plot consumption inequality over the life cycle (again, measured by the difference between the 75-th and 25-th percentile of the log consumption distribution) under three different scenarios: (a) the baseline (blue dots), which allows for the stock of cognitive endowment to affect both the permanent components of wealth returns and wages, and –as shown in Figure 8 replicates the data well; (b) a counterfactual in which we let the permanent component of wealth returns and the permanent component of earnings to be uncorrelated (the red squares) – this is obtained setting $\delta^y = 0$; and (c) a counterfactual case in which we shut down persistent returns to wealth heterogeneity (the green triangles) – this is obtained setting $\sigma_{\Phi_i} = 0$. These counterfactual exercises highlight the importance of heterogeneity in wealth returns to understand different consumption inequality dynamics. Without correlation in return and earnings FE's, there will be too little consumption inequality over the life cycle, and it will grow at a slower rate. Further shutting down heterogeneity in returns would exacerbate both effects.

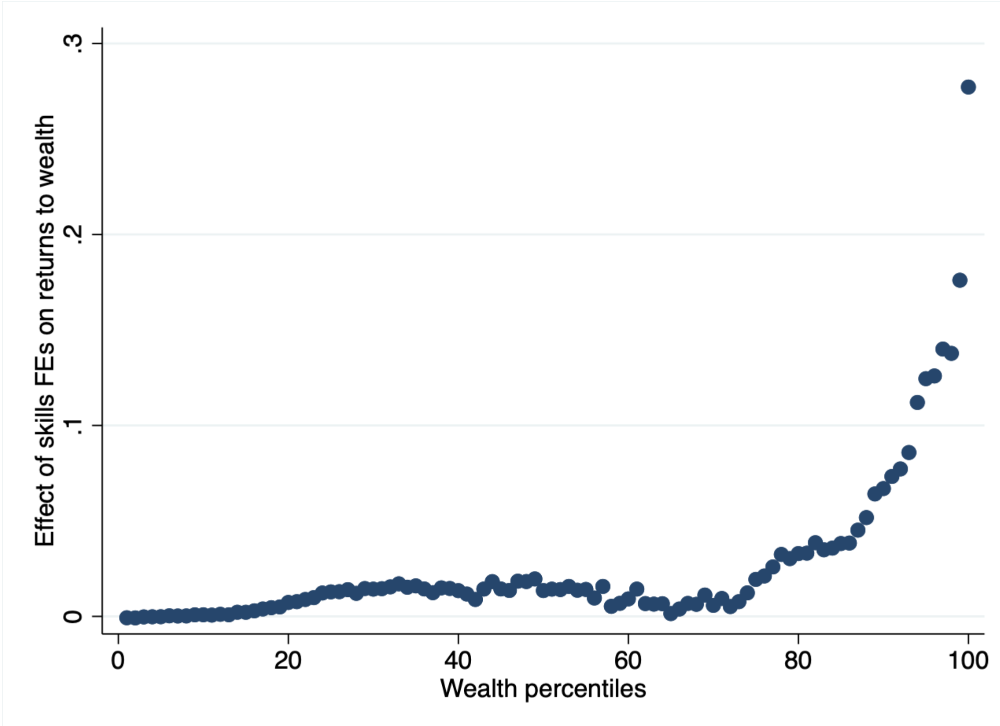

Figure 9: Persistent returns to wealth heterogeneity and average returns to wealth across the wealth distribution

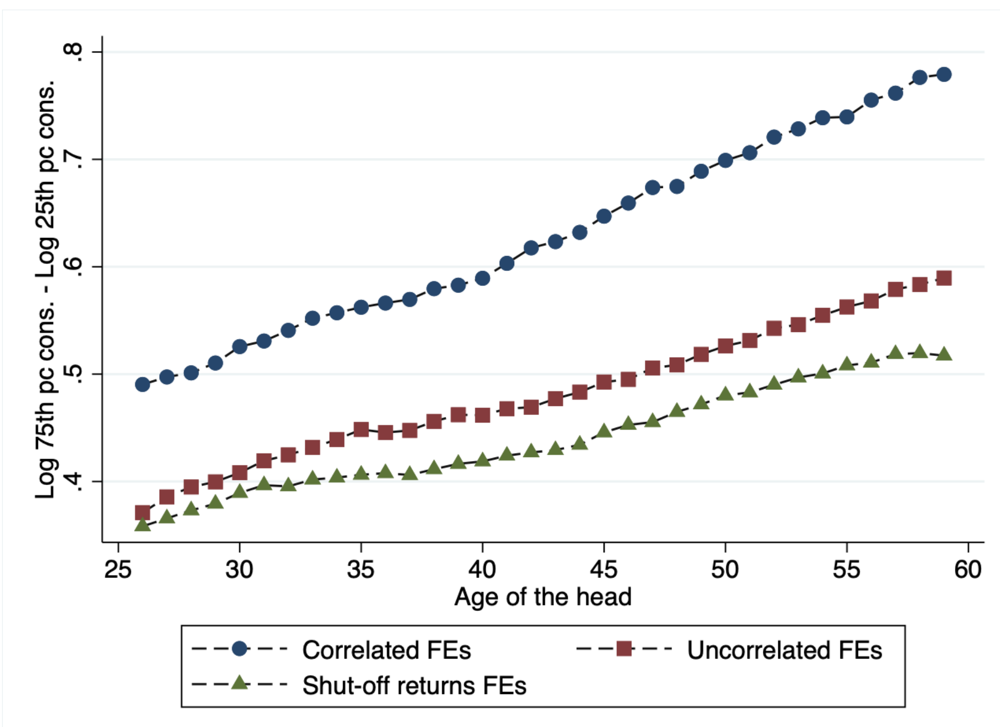

Figure 10: The role of cognitive skill heterogeneity on consumption inequality over the life-cycle

6 Conclusions

In this paper we take a standard life cycle model where household income varies because of i.i.d. transitory and permanent innovations. Unlike previous papers in the literature, we put more focus on the initial component of permanent income and assume it is a function of the stock of cognitive skills an individual is endowed with at birth (or at the point of entry in the labor market) and other fixed components independent of it. In keeping with evidence from a recent literature, we also let returns from assets to include a persistent heterogeneity component, which we also model as a function of the stock of endowed cognitive skills. Hence, the persistent component of wealth returns and the persistent component of earnings are potentially correlated. We find that these variants are important. First, consumption inequality over the life cycle would be lower and grow more slowly if persistent heterogeneity in returns was ignored or assumed independent of persistent heterogeneity in earnings. Second, the model would have a hard time explaining excess wealth returns at the top of the wealth distribution if these variants were ignored.

References

Attanasio, O. P., J. Banks, C. Meghir, and G. Weber (1999): "Humps and Bumps in Lifetime Consumption," *Journal of Business & Economic Statistics*, 17, 22–35.

Bach, L., L. E. Calvet, and P. Sodini (2020): "Rich Pickings? Risk, Return, and Skill in Household Wealth," *American Economic Review*, 110, 2703–47.

Benhabib, J. and A. Bisin (2018): "Skewed Wealth Distributions: Theory and Empirics," *Journal of Economic Literature*, 56, 1261–91.

Carroll, C. D. (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, 112, 1–55.

——— (2006): "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics Letters*, 91, 312–320.

Cocco, J., F. Gomes, and P. Maenhout (2005): "Consumption and Portfolio Choice over the Life Cycle," *The Review of Financial Studies*, 18, 491–533.

De Nardi, M. (2004): "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 71, 743–768.

De Nardi, M. and G. Fella (2017): "Saving and wealth inequality," *Review of Economic Dynamics*, 26, 280–300.

Deaton, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59, 1221–48.

Dietz, P. O. (1968): "Components of a Measurement Model: Rate of Return, Risk, and Timing," *The Journal of Finance*, 23, 267–275.

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

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2020): "Heterogeneity and Persistence in Returns to Wealth," *Econometrica*, 88, 115–170.

French, E. and J. B. Jones (2004): "On the distribution and dynamics of health care costs," *Journal of Applied Econometrics*, 19, 705–721.

Gabaix, X., J.-M. Lasry, P.-L. Lions, and B. Moll (2016): "The Dynamics of Inequality," *Econometrica*, 84, 2071–2111.

Gourinchas, P.-O. and J. A. Parker (2002): "Consumption Over the Life Cycle," *Econometrica*, 70, 47–89.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995): "Precautionary saving and social insurance," *Journal of political Economy*, 103, 360–399.

Iskhakov, F., T. H. Jørgensen, J. Rust, and B. Schjerning (2017): "The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks," *Quantitative Economics*, 8, 317–365.

Kimball, M. S., C. R. Sahm, and M. D. Shapiro (2009): "Risk Preferences in the PSID: Individual Imputations and Family Covariation," *American Economic Review*, 99, 363–68.

Saez, E. and G. Zucman (2016): "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *The Quarterly Journal of Economics*, 131, 519–578.