The Cross-Section of Household Preferences

Laurent E. Calvet, John Y. Campbell,
Francisco Gomes, and Paolo Sodini*

June 2021

Abstract

This paper estimates the cross-sectional distribution of Epstein-Zin preferences in a large panel of Swedish households. We consider a life-cycle model of saving and portfolio choice, which we match to profiles of wealth accumulation and the risky share. We find some heterogeneity in risk aversion (a standard deviation of 0.47 with a mean/median of 5.24/5.30) and considerable heterogeneity and positive skewness in the time preference rate (standard deviation 6.0% with a mean/median of 6.18/4.08%) and elasticity of intertemporal substitution (standard deviation 0.96 with a mean/median of 0.99/0.42). Risk aversion and the EIS are almost cross-sectionally uncorrelated. We estimate lower risk aversion for households with riskier labor income and higher education, and a higher TPR and lower EIS for households who enter our sample with low wealth.

*Calvet: EDHEC Business School and CEPR, e-mail: laurent.calvet@edhec.edu. Campbell: Harvard University and NBER; e-mail: john_campbell@harvard.edu. Gomes: London Business School, e-mail: fgomes@london.edu. Sodini: Stockholm School of Economics, e-mail: Paolo.Sodini@hhs.se. We acknowledge helpful comments on earlier versions of this paper from Stijn van Nieuwerburgh, Jessica Wachter, Stanley Zin, and seminar participants at Arizona State University, EDHEC Business School, ENSAE-CREST, Harvard University, Imperial College London, Stanford University, the University of Michigan, the Vienna Graduate School of Finance, the 2017 American Economic Association meeting, the UCLA Anderson Fink Center Conference on Financial Markets, the 2019 FMND Workshop, the 2019 NBER Asset Pricing Summer Institute, and the 2021 London Seminar on Micro and Macro Implications of Household Behaviour and Financial Decision-Making. We thank the Sloan Foundation for financial support to John Campbell, and Azar Aliyev, Nikolay Antonov, Huseyin Aytug, and Yapei Zhang for able and dedicated research assistance.

1 Introduction

When households make financial decisions, are their preferences toward time and risk substantially similar, or do they vary cross-sectionally? And if preferences are heterogeneous, how do preference parameters covary in the cross-section with one another and with household attributes such as education and sector of employment? This paper answers these questions using a life-cycle model of saving and portfolio choice fit to high-quality household-level administrative data from Sweden.

Modern financial theory distinguishes at least three parameters that govern savings behavior and financial decisions: the time preference rate (TPR), the coefficient of relative risk aversion (RRA), and the elasticity of intertemporal substitution (EIS). The canonical model of Epstein and Zin (1989) makes all three parameters constant and invariant to wealth for a given household, while breaking the reciprocal relation between relative risk aversion and the elasticity of intertemporal substitution implied by the older power utility model.

We structurally estimate these three preference parameters in the cross-section of Swedish households by embedding Epstein-Zin preferences in a life-cycle model of consumption and portfolio choice in the presence of uninsurable labor income risk and borrowing constraints. In our base case implementation we assume that all agents have common beliefs about income processes and financial returns. As a robustness check, we consider a simple form of heterogeneity in beliefs about risky asset returns.

To mitigate the effects of idiosyncratic events not captured by the model we carry out our estimation on groups of households who share certain observable features, making use of asymptotic properties of our estimation procedure as the size of each group increases. We first group households by their education level, the level of income risk in their sector of employment, and birth cohort. To capture heterogeneity in preferences that is unrelated to these characteristics we further divide households by their initial wealth accumulation in relation to income and by their initial risky portfolio share. This process gives us a sample of 4151 composite households that have data available in each year of our sample from 1999 to 2007.

We allow age-income profiles to vary with education, and the determinants of income risk (the variances of permanent and transitory income shocks) to vary with both education and the household's sector of employment. These assumptions are standard in the life-cycle literature (Carroll and Samwick 1997, Cocco, Gomes, and Maenhout 2005). These life-cycle models more readily match portfolio allocations and wealth accumulation at mid-life than at younger ages or after retirement. Therefore we estimate the preference parameters by matching the time series of wealth and portfolio choice between ages 40 and 60, taking as given the initial level of wealth at the start of each year and the risky asset returns realized during each year. Since we do not observe decisions late in life, we cannot accurately account for bequest motives and instead capture the desire to leave a bequest as a lower rate of time preference.

We measure not only liquid financial wealth, but also defined-contribution retirement assets as well as household entitlements to defined-benefit pension income. However, we confine attention to households who hold some risky financial assets outside retirement accounts, for comparability with previous work and in order to avoid the need to estimate determinants of non-participation in risky financial markets. Our imputation of defined-contribution retirement wealth is an empirical contribution of our paper that extends previous research on Swedish administrative data.

Residential real estate is another important component of household wealth. To handle this, we include real estate in our empirical analysis but map both real estate and risky financial asset holdings into implied holdings of a single composite risky asset. While this is a stylization of reality, the inclusion of real estate wealth is consistent with common practice in life-cycle models (Castaneda, Diaz-Gimenez and Rios-Rull 2003, De Nardi 2004, Gomes and Michaelides 2005).

It is a challenging task to identify all three Epstein-Zin preference parameters. In principle, these parameters play different roles with the TPR affecting only the overall slope of the household's planned consumption path, risk aversion governing the willingness to hold risky financial assets and the strength of the precautionary savings motive, and the EIS affecting both the overall slope of the planned consumption path and the responsiveness of this slope to changes in background risks and investment opportunities. We observe portfolio choice directly, and the slope of the planned consumption path indirectly through its relation with saving and hence wealth accumulation. However, we require time-variation in background risks or investment opportunities in order to identify the EIS separately from the TPR (Kocherlakota 1990, Svensson 1989).

Our model assumes that expected returns on safe and risky assets are constant over time, so we cannot exploit time-variation in the riskless interest rate or the expected risky return to identify the EIS in the manner of Hall (1988) or Yogo (2004). However, the model incorporates time-variation in background risks. Households in the model have a target level of financial wealth that serves both as a buffer stock to smooth consumption in the face of random income variation, and as a means of financing retirement. Households save when their wealth is below the target, and they do so more aggressively when the EIS is high. A related phenomenon is that households with a high level of financial wealth relative to human capital invest more conservatively, which reduces the expected rate of return on their portfolio. In addition, as households age their mortality rates increase, and this alters the effective rate of time discounting. For all these reasons we can identify the EIS from time-variation in the growth rate of wealth within each household group. This identification strategy is a methodological contribution of our paper.

Our main empirical findings are as follows. First, we find considerable heterogeneity in wealth accumulation and portfolio composition across the Swedish population. Average wealth-income ratios increase strongly with the riskiness of income and the level of education while average risky shares do not, but both variables have substantial heterogeneity unrelated to these variables.

Second, we document patterns in wealth and portfolio composition that are broadly consistent with life-cycle financial theory. As households age, they tend to accumulate wealth and reduce their risky portfolio share. The risky portfolio share also declines with the wealth-income ratio after controlling for age. Both patterns are predicted by a life-cycle model in which human capital is safer than risky financial capital.

Third, we estimate heterogeneity in all three preference parameters. The least heterogeneity is in risk aversion, which has a cross-sectional standard deviation of

0.47 around a mean of 5.24. Our other two preference parameters are highly dispersed and right-skewed. The mean TPR is 6.18%, well above the median value of 4.08%, and the standard deviation is 6.03%. The mean EIS is 0.99, well above the median value of 0.42, and the standard deviation is 0.96.

Fourth, our preference parameter estimates are only weakly cross-sectionally correlated. The correlation between risk aversion and the EIS is weakly positive, in contrast with the perfect negative correlation between log risk aversion and the log EIS that we would find if all households had power utility with heterogeneous coefficients. The TPR is negatively correlated with both risk aversion and the EIS, implying a tendency for patient people to be both cautious and willing to substitute intertemporally. The weak correlations across preference parameters imply that Swedish household behavior is heterogeneous in multiple dimensions, not just one. A single source of heterogeneity omitted from our model, such as heterogeneity in household beliefs about the equity premium, cannot explain this pattern.

Fifth, there are some interesting correlations between our parameter estimates, the moments we use for estimation, and exogenous characteristics of households. Risk aversion is lower for households working in risky sectors and for households with higher education. Both these patterns are consistent with the hypothesis that risk-tolerant households choose to acquire education and select risky occupations, but they could also result from households' failure to understand the portfolio choice implications of their income risk exposure and wealth accumulation.

The TPR is negatively correlated with the initial wealth-income ratio of each household group, and positively correlated with the average growth rate of the wealth-income ratio. The symptom of a high TPR in our data is a tendency to accumulate retirement savings later in life, catching up belatedly with those who saved earlier in life. The equivalent correlations for the EIS have the opposite signs, suggesting that households with a high EIS save early in life to reach a target wealth-income ratio, while households with a low EIS save more gradually over time.

Our paper is related to a large literature on household portfolio choice over the life cycle, including Campbell and Viceira (2002), Ameriks and Zeldes (2004), Cocco, Gomes, and Maenhout (2005), and Fagereng, Gottlieb, and Guiso (2017). Our use of comprehensive administrative data from Sweden follows a series of papers by Calvet, Campbell, and Sodini (2007, 2009a, 2009b), Calvet and Sodini (2014), Betermier, Calvet, and Sodini (2017), and Bach, Calvet, and Sodini (2020). A smaller literature on heterogeneity in portfolio choice has recently tried to relate observed household behavior to underlying heterogeneity in preferences and beliefs (Meeuwis, Parker, Schoar, and Simester 2018, Giglio, Maggiori, Stroebel, and Utkus 2019). Relative to this literature, we observe more households over a longer period of time and have more complete data on wealth and portfolio allocation, but we lack data on potentially heterogeneous beliefs.

The organization of the paper is as follows. Section 2 explains how we measure household wealth and its allocation to safe and risky assets, describes the creation of household groups, and reports summary statistics. Section 3 presents the life-cycle model. Section 4 discusses preference parameter identification and develops our estimation methodology. Section 5 reports empirical results and a Monte Carlo analysis of our estimation method. Section 6 concludes. An online appendix provides additional details about our empirical analysis and estimation technique.

2 Measuring Household Wealth and Asset Allocation

Our empirical analysis is based on the Swedish Wealth and Income Registry. This high-quality administrative panel provides the income, wealth, and debt of every Swedish resident. Income data are available at the individual level from 1983 and can be aggregated to the household level from 1991. Wealth data are available from 1999 through 2007. The wealth data include bank account balances, holdings of financial assets, and real estate properties measured at the level of each security or property. We augment the dataset by imputing defined contribution (DC) retirement wealth and entitlements to defined benefit (DB) pension income using income data and the administrative rules governing Swedish pensions.

2.1 The Household Balance Sheet

We measure four components of the household balance sheet: liquid financial wealth, real estate wealth, DC retirement savings, and debt. We define the total net wealth of household h at time t, $W_{h,t}$, as

$$W_{h,t} = LW_{h,t} + RE_{h,t} + DC_{h,t} - D_{h,t}, \tag{1}$$

where $LW_{h,t}$ is liquid financial wealth, $RE_{h,t}$ is real estate wealth, $DC_{h,t}$ is DC retirement wealth, and $D_{h,t}$ is debt.

Liquid financial wealth is the value of the household's bank accounts and holdings of Swedish money market funds, mutual funds, stocks, capital insurance products, derivatives and fixed income securities. Mutual funds include balanced funds and bond funds, as well as equity funds. We subdivide liquid financial wealth into cash, defined as the sum of bank balances and money market funds, and risky assets.

Real estate consists of primary and secondary residences, rental, commercial and industrial properties, agricultural properties and forestry. The Wealth and Income Registry provides the holdings at the level of each asset. The pricing of real estate properties is based on market transactions and tax values adjusted by a multiplier, as in Bach, Calvet, and Sodini (2020).

Debt is the sum of all liabilities of the household, including mortgages and other personal liabilities held outside private businesses. Since Swedish household debt is normally floating-rate, we treat debt as equivalent to a negative cash position but paying a borrowing rate that is higher than the safe lending rate.

The hardest balance sheet component to measure is DC retirement wealth. We do not measure this directly but impute it by reconstructing the details of the Swedish pension system, as we discuss in the next subsection. This detailed pension analysis also enables us to measure each household's entitlement to defined benefit (DB) pension payments in retirement.

As described here, the household balance sheet excludes durable goods and private businesses, whose values are particularly difficult to measure. Private businesses are an important component of wealth for the wealthiest households in Sweden, but unimportant for most Swedish households (Bach, Calvet, and Sodini 2020).

2.2 Pension Imputation

The Swedish pension system consists of three pillars: state pensions, occupational pensions, and private pensions.

The *state pension* system requires each worker in Sweden to contribute 18.5% of their pension qualifying income: 16% to the pay-as-you-go defined benefit (DB) system and the remaining 2.5% to a defined contribution (DC) system called premiepension system. DC contributions are invested in a default fund, that mirrors the world index during our sample period, unless the worker opts out and chooses a portfolio of funds among those offered on the state DC platform. State DB payouts are a function of the pension qualifying income earned during the entire working life.

Occupational pensions were introduced to Sweden in 1991. They are regulated for the vast majority of Swedish residents by four collective agreements applying to blue-collar private-sector workers, white-collar private-sector workers, central government employees, and local government employees. Since these agreements specify workers' monthly pension contributions, the fraction directed to DB and DC pension plans, and the DC choices available to workers, we are able to impute both DC contributions and DB entitlements at the household level.

The collective agreements specify DC contributions as a percentage of pension qualifying income. These contributions are invested through insurance companies in either variable annuity products (called TradLiv in Sweden), or in portfolios of mutual funds, chosen by workers from a selection provided by the insurance company.

Defined contribution *private pensions* have existed in Sweden for a long time but our dataset provides us with individual private pension contributions from 1991. We assume that they are invested like occupational and state DC contributions.

To calculate DC retirement wealth at each point in time, we accumulate contributions from all three pillars, with appropriate assumptions about the equity share, the investment of equity contributions, and the initial level of private pension wealth in 1991. We describe this procedure in detail in the appendix.

DC retirement wealth accumulates untaxed but is taxed upon withdrawal. To convert pre-tax retirement wealth into after-tax units that are comparable to liquid financial wealth, we assume an average tax rate τ on withdrawals (estimated at 32% which is the average tax rate on nonfinancial income paid by households with retired heads over 65 years old) and multiply pre-tax wealth by $(1-\tau)$. In the remainder of the paper, we always state retirement wealth in after-tax units.

2.3 Household Asset Allocation

Our objective is to match the rich dataset of household income and asset holdings to the predictions of a life-cycle model, which will allow us to estimate household preferences. To accomplish this, we need to map the complex data into a structure that can be related to a life-cycle model with one riskless and one risky asset. We do this in three stages. First we map all individual assets to equivalent holdings of diversified stocks, real estate, or cash. Second, we assume a variance-covariance matrix for the excess returns on stocks and real estate over cash that enables us to compute the volatility of each household portfolio. Third, we assume that all household portfolios earn the same Sharpe ratio so that the volatility of the portfolio determines the expected return on the portfolio. Equivalently, we convert the volatility into a "risky share" held in a single composite risky asset. For ease of interpretation, we normalize that risky asset to have the same volatility as a world equity index.

At the first stage, we treat liquid holdings of individual stocks, equity mutual funds, and hedge funds as diversified holdings of the MSCI world equity index.[1] We treat liquid holdings of balanced funds and bond funds as portfolios of cash and

[1] This reflects the global exposure of Swedish equity portfolios documented by Calvet, Campbell, and Sodini (2007). It abstracts from underdiversification which is documented in the same paper. The impact of underdiversification in liquid wealth is reduced when one takes account of diversified DC retirement wealth as we do in this paper.

stocks, with the share in stocks given by the beta of each fund with the world index.[2] We assume that unclassifiable positions in capital insurance, derivatives, and fixed income securities are invested in the same mix of cash and stocks as the rest of liquid financial wealth. We treat all real estate holdings as positions in a diversified index of Swedish residential real estate, the FASTPI index. We assume that DC retirement wealth is invested in cash and the MSCI equity world index as described in section 2.2.

This mapping gives us implied portfolio weights in liquid stocks, real estate, and DC stocks in the net wealth of each household. For household h at time t, write these weights as $\omega_{S,t}^h$, $\omega_{RE,t}^h$, and $\omega_{DC,t}^h$, and the corresponding excess returns as $R_{S,t+1}^e, R_{RE,t+1}^e, R_{DC,t+1}^e$. The excess return on net wealth for household h is then:

$$R_{h,t+1}^e = \omega_{S,t}^h R_{S,t+1}^e + \omega_{RE,t}^h R_{RE,t+1}^e + \omega_{DC,t}^h R_{DC,t+1}^e + (1 - \omega_{S,t}^h - \omega_{RE,t}^h - \omega_{DC,t}^h) R_{D,t+1}^e,$$

where $R_{D,t+1}^e$ is the borrowing rate in excess of the Swedish T-bill.

The second stage of our analysis is to calculate the variance of the excess return on household net wealth. Since the borrowing rate is deterministic, we only need to consider the vector $\omega_t^h = (\omega_{S,t}^h, \omega_{RE,t}^h, \omega_{DC,t}^h)'$, and we can calculate the variance of $R_{h,t+1}^e$ as $\sigma^2(R_{h,t+1}^e) = \omega_{h,t}' \ \Sigma \ \omega_{h,t}$,where Σ is the variance-covariance matrix of $R_{t+1}^e = (R_{S,t+1}^e, R_{RE,t+1}^e, R_{DC,t+1}^e)'$. To estimate the elements of Σ, we assume that cash earns the Swedish one-month risk-free rate net of taxes, that liquid equity earns the MSCI world index return less a 30% long-term capital income tax rate (Du Rietz et al. 2015), that real estate earns the FASTPI index return less a 22% real estate capital gain tax rate, and that stocks held in DC plans earn the pre-tax MSCI world index return before the adjustment of their value to an after-tax basis. Using data from 1984–2007, we estimate the post-tax excess return volatility for stocks at 13.3% and for real estate at 5.5%, with a correlation of 0.27. The pre-tax excess stock return volatility is 19%.

In the third stage of our analysis, we define a numeraire asset, the aggregate Swedish portfolio of cash, stocks, and real estate scaled to have the same volatility as

[2] We cap the estimated fund beta at 1, and use the cross-sectional average fund beta for funds with less than 24 monthly observations.

the after-tax global equity index return: $R^e_{N,t+1} = (1+L)(\omega'_{agg,t}\ R^e_{t+1})$. Here $R^e_{N,t+1}$ is the return on the numeraire asset and $\omega_{agg,t}$ is the vector containing the weights of equity, real estate and the DC retirement portfolio in the aggregate net wealth of all Swedish households in our sample. The scaling factor L is chosen so that the volatility of $R^e_{N,t+1}$ is equal to the volatility of the after-tax return in local currency on the global equity index.

The empirical risky share $\alpha_{h,t}$ is the ratio of the volatility of household h's portfolio to the volatility of the numeraire asset:

$$\alpha_{h,t} = \sigma\left(R^e_{h,t+1}\right) / \sigma\left(R^e_{N,t+1}\right). \tag{2}$$

This approach implicitly assumes that all households earn the same Sharpe ratio on their risky assets, but guarantees that the standard deviation of a household's wealth return used in our simulations coincides with its empirical value. A unit value for $\alpha_{h,t}$ says that the portfolio has the same volatility, 13.3%, as if it invested solely in stocks held outside a retirement account, without borrowing or holding cash.

2.4 Composite Households

The Swedish Income Registry data set contains data on the entire population of Sweden, but our focus is on middle-aged households aged between 40 and 60 during our sample period from 1999 to 2007. There are 7.7 million household-year observations on the 13 cohorts born between 1947 and 1959, but we impose several filters. We exclude 2.3 million observations on households that do not hold risky financial assets outside retirement accounts. We also exclude households in which the head is a student, working in the agricultural sector, retired before the start of our sample, missing information on education or sector of employment, or missing data in any year between 1999 and 2007. We exclude households that change their employment sector during our sample in such a way as to alter the level of income volatility they are exposed to. To limit the impact of the wealthiest (for whom our measurement procedures may be less adequate), we also exclude households whose financial wealth is above the 99th percentile of the wealth distribution in 1999. These filters exclude another 2.8 million observations, leaving us with a balanced panel containing 2.6 million household-year observations and 291,488 households.

We classify households by three levels of educational attainment: (i) basic or missing education, (ii) high school education, and (iii) post-high school education. We also classify households by 12 sectors of employment. Within each education level, we rank the sectors by their total income volatility and divide them in three categories. In this way we create a 3×3 grid of 9 large education/sector categories where the sectors of employment are aggregated by income volatility. We subdivide each of these categories using a two-way sort by deciles of the initial wealth-income ratio and initial risky share. We use the lowest two and highest two deciles and the middle three quintiles, giving us a 7×7 grid of 49 bins for the initial wealth-income ratio and risky share.[3] We further subdivide by 13 cohorts to create $5733 = 9 \times 49 \times 13$ groups. After excluding small groups that do not contain members in each year from 1999 to 2007, our final sample is a balanced panel of 4151 groups.

The median group size across years is 63 households, but the average group size is larger at about 86 households. The difference reflects a right-skewed distribution of group size, with many small groups and a few much larger ones. The group-level statistics we report in the paper are all size-weighted in order to reflect the underlying distributions of data and preference parameters at the household level.

We treat each group as a composite household, adding up all wealth and income of households within the group. Because we assume scale-independent Epstein-Zin preferences, we scale wealth by income and work with the wealth-income ratio as well as the implied risky share held in our composite numeraire asset.

[3]The wealth-income and risky share breakpoints are set separately in each of the 9 categories. This ensures that across categories we have the same proportion of households at each of the 7 risky share and wealth-income levels. However, the number of households can differ across the 49 bins defined by the two-way sort, to the extent that the wealth-income ratio and the risky share are cross-sectionally correlated.

Table 1: Size-Weighted Wealth-Income Ratios and Risky Shares by Levels of Education and Income Volatility

Panel A. Cross-Sectional Means

	WY				RS			
	No High School	High School	Post-High School	All	No High School	High School	Post-High School	All
Low	3.57	4.03	4.84	4.39	0.685	0.681	0.662	0.672
Medium	4.15	4.42	4.97	4.56	0.671	0.667	0.664	0.666
High	4.72	5.04	6.00	5.34	0.656	0.669	0.667	0.665
All	4.13	4.42	5.14	4.68	0.671	0.672	0.664	0.668

Panel B. Cross-Sectional Standard Deviations

	WY				RS			
	No High School	High School	Post-High School	All	No High School	High School	Post-High School	All
Low	3.04	3.27	3.6	3.45	0.244	0.224	0.208	0.219
Medium	3.62	3.56	3.61	3.59	0.252	0.224	0.208	0.223
High	3.71	3.85	3.71	3.80	0.233	0.219	0.192	0.212
All	3.47	3.55	3.66	3.61	0.242	0.223	0.204	0.218

Panel A reports cross-sectional means of the wealth-income ratio (WY) and risky share (RS) for Swedish household groups with 3 levels of education and working in sectors with 3 levels of income volatility given in Table 2, and for aggregates of these groups. Panel B reports cross-sectional standard deviations of WY and RS across the groups in each of these categories and their aggregates. All statistics weight groups by their size, that is by the number of households they contain, to recover the underlying household-level statistics assuming homogeneity of WY and RS within groups. Summary statistics on group size are reported in the online appendix.

2.5 Cross-Section of Wealth-Income Ratio and Risky Share

We now consider the cross-section of the wealth-income ratio and risky share, averaging across all years in our sample. The top panel of Table 1 shows the variation in average wealth-income ratios and risky portfolio shares across groups with each level of education and sectors of employment with each level of income risk, averaging across cohorts and the subdivisions by initial wealth-income ratio and risky share. For the purpose of computing these summary statistics, households in each group are treated as a single composite household that owns all wealth and receives all income of the group, and groups are weighted by the number of households they contain. Average wealth-income ratios vary widely from 3.6 to 6.0, while average risky shares vary in a narrow range from 65% to 69%. Within each sector, average wealth-income ratios are higher for more educated households, particularly those with post-high school education, but average risky portfolio shares vary little with education. Across sectors, the level of income risk has a strong positive effect on the wealth-income ratio and a weak negative effect on the risky portfolio share.

The category averages in the top panel of Table 1 conceal a great deal of dispersion across disaggregated groups of households. This is shown by the bottom panel of Table 1, which reports the standard deviations of the wealth-income ratio and the risky portfolio share across groups in each of the nine categories of education and sectoral income risk. The standard deviations of the risky share are consistently in the range 19–25%, while the standard deviations of the wealth-income ratio are in the range 3.0–3.9. Across all 4151 groups, the average wealth-income ratio has a mean of 4.7 with a standard deviation of 3.6, while the average risky share has a mean of 67% with a standard deviation of 22%.[4]

The cross-sectional variation in wealth and asset allocation documented in Table 1, even for groups with similar income risks and levels of education, suggests that it will be difficult to account for household behavior without allowing for heterogeneity in preferences. We now develop a life-cycle model that we can use to estimate preferences from the evolution of wealth and asset allocation during our sample period.

[4]Figure A.1 in the online appendix plots the size-weighted group-level distributions. The aggregation of households into groups naturally reduces the dispersion that is visible in the household-level data. Across all individual households in our dataset the average wealth-income ratio is 6.5 with a standard deviation of 185. After winsorizing the underlying household distribution at the 99th percentile, the average falls to 4.6 and the standard deviation falls to 4.0. The average risky share across all individual households is 71% with a standard deviation of 43%.

3 Income Process and Life-Cycle Model

3.1 Income Risk Across Group

We consider the labor income specification used in Carroll and Samwick (1997), Gourinchas and Parker (2002) and Cocco, Gomes, and Maenhout (2005), among others:

$$\log(Y_{h,t}) = a_c + b' x_{h,t} + \nu_{h,t} + \varepsilon_{h,t}, \tag{3}$$

where $Y_{h,t}$ denotes real income for household h in year t, a_c is a fixed effect for the cohort to which the household belongs, $x_{h,t}$ is a vector of characteristics, $\nu_{h,t}$ is a permanent random component of income, and $\varepsilon_{h,t}$ is a transitory component.

We enrich the model above by distinguishing between shocks that are common to all households in a group and shocks that are specific to each household in the group. We assume that the permanent component of income, $\nu_{h,t}$, is the sum of a group-level component, ξ_t, and an idiosyncratic component, $z_{h,t}$:

$$\nu_{h,t} = \xi_t + z_{h,t}. \tag{4}$$

To simplify notation, we do not write an explicit group index but write group-level shocks using a single time index. The components ξ_t and $z_{h,t}$ follow independent random walks: $\xi_t = \xi_{t-1} + u_t$, and $z_{h,t} = z_{h,t-1} + w_{h,t}$.

The transitory component of income, $\varepsilon_{h,t}$, is by contrast purely idiosyncratic. This fits the fact that group average income growth in our Swedish data is slightly positively autocorrelated, whereas it would be negatively autocorrelated if transitory income had a group-level component.

Finally, we assume that the three income shocks impacting household h are i.i.d. Gaussian: $(u_t, w_{h,t}, \varepsilon_{h,t})' \sim \mathcal{N}(0, \Omega_Y)$, where Ω_Y is the diagonal matrix with diagonal elements σ_u^2, σ_w^2, and σ_ε^2.

We estimate the income process (3) using household yearly income data, following a procedure described in detail in the appendix. This gives us estimates of the age-income profile for each education group, which we plot in appendix Figure A.2.

To estimate income risk, we further divide households with the same education level into business sector categories. σ_u^2 is estimated by averaging the regression residuals within each education-business sector category, and by computing the sample variance of the resulting income innovations. We then apply a Carroll and Samwick (1997) decomposition to estimate the permanent and transitory idiosyncratic income risks, σ_w^2 and σ_ε^2, of each education-business sector category.

We proceed in two steps. First, we implement the procedure above on 36 education-business sector categories obtained by dividing households with each of three education levels into the 12 business sectors corresponding to the first digit of the SNI industry code. Equipped with income risk estimates for each of the 36 categories, we aggregate business sectors into three levels of total income risk for each education level.[5] Second, we re-apply the procedure above to estimate income risk for the resulting nine education-business sector categories.

Table 2 reports the estimated standard deviations of group-level income shocks (permanent by assumption) and of permanent and transitory idiosyncratic income shocks, across these nine categories. Group-level income volatilities and permanent idiosyncratic income volatilities vary relatively little across sectors, but transitory idiosyncratic income volatilities are considerably higher for high-risk sectors.

Table 2 also shows that educated households, particularly those with higher education, face higher transitory income risk and lower permanent income risk than less educated households. This pattern is consistent with Low, Meghir, and Pistaferri (2010), but it contrasts with earlier studies showing the opposite pattern in the U.S. A likely explanation is that in Sweden, uneducated workers face lower unemployment risk and lower effects of unemployment on income than in many other countries, while

[5] Appendix Table A.1 reports the number of households and Table A.2 reports the underlying sectors in each category. The patterns are intuitive, with relatively little transitory income risk in the public sector and in mining and quarrying, electricity, gas, and water supply, and relatively high transitory income risk in hotels and restaurants, real estate activities, construction for less educated workers, and the financial sector for more educated workers.

Table 2: Percentage Volatilities of Income Shocks

	Total			Systematic		
	No High School	High School	Post-High School	No High School	High School	Post-High School
Low	13.59	13.42	15.85	2.76	2.88	3.49
Medium	17.20	16.31	16.69	2.86	3.37	3.44
High	20.15	19.83	21.48	3.08	3.27	3.91
	Idiosyncratic permanent			**Idiosyncratic transitory**		
	No High School	High School	Post-High School	No High School	High School	Post-High School
Low	7.48	6.88	5.27	11.01	11.15	14.53
Medium	8.84	7.8	5.64	14.47	13.93	15.33
High	6.95	7.35	5.91	18.66	18.13	20.28

This table reports the standard deviations of income shocks, in percentage points, for Swedish household groups with 3 levels of education and working in sectors with 3 levels of income volatility. The top left panel reports the total standard deviation of income shocks, the top right panel reports the standard deviation of systematic (group-level) permanent income shocks, the bottom left panel reports the standard deviation of idiosyncratic (household-level) permanent income shocks, and the bottom right panel reports the standard deviation of idiosyncratic transitory income shocks.

educated workers face relatively high income losses when they become unemployed.[6]

We have already noted in discussing Table 1 that average wealth-income ratios tend to be higher in sectors with riskier income. This pattern is intuitive given that labor income risk encourages precautionary saving. However, there is little tendency for risky portfolio shares to be lower in sectors with riskier income. Table 3 further explores these effects by regressing the average wealth-income ratio and risky share on age, total income volatility, and dummies for high school and post-high school education. All regressions also include year fixed effects.[7]

[6] This results from the following features of the Swedish labor market. First, it is straightforward for companies to downsize divisions, but extremely difficult for them to lay off single individuals unless they have a high managerial position. Second, companies that need to downsize typically restructure their organizations by bargaining with unions. Third, unions are nationwide organizations that span large areas of employment and pay generous unemployment benefits. Fourth, the pay cut due to unemployment is larger for better paid jobs. After an initial grace period, an unemployed person will be required to enter a retraining program or will be assigned a low-paying job by a state agency. All these features imply that unemployment is slightly more likely and entails a more severe proportional income loss for workers with higher levels of education.

[7] We do not include cohort effects in this table. It is well known that unrestricted time, age, and cohort effects cannot be identified (Ameriks and Zeldes 2004, Fagereng, Gottlieb, and Guiso 2017). Here we use unrestricted time effects, a linear age effect, and exclude cohort effects.

Table 3: Size-Weighted Panel Regressions of Wealth-Income Ratio and Risky Share on Group Characteristics

	(1) WY	(2) RS	(3) RS
Age	0.147***	-0.014***	-0.010***
	(0.016)	(0.001)	(0.001)
Total income volatility	18.375***	-0.193	0.345**
	(2.152)	(0.128)	(0.105)
High school	0.538***	-0.011	0.005
	(0.136)	(0.009)	(0.008)
Post-high school	1.020***	-0.017	0.013
	(0.145)	(0.009)	(0.007)
WY			-0.029***
			(0.001)
Constant	-6.437***	1.515***	1.326***
	(0.817)	(0.056)	(0.046)
Year fixed effects	Yes	Yes	Yes
R^2	0.105	0.180	0.388

This table reports panel regressions of the wealth-income ratio (WY) and risky share (RS) on group characteristics including the age of households in the group, total income volatility (in natural units), and dummies for high-school and post-high-school education. All regressions weight groups by their size, to recover underlying relationships at the household level, and include year fixed effects. Standard errors are reported in parentheses and statistical significance levels are indicated with stars: * denotes 1-5%, ** 0.1-1%, *** less than 0.1% significance. There are 37,359 observations on groups, corresponding to 2,623,392 observations on underlying households.

The first column of the table shows that the average wealth-income ratio increases with age and with income volatility. This is consistent with the view that wealth is accumulated in part to finance retirement, and in part as a buffer stock against temporary shocks to income. In addition, the average wealth-income ratio increases with the level of education.

The second column shows that the average risky share decreases with age, but income risk and education are not significant predictors of the average risky share although the coefficient on income risk is negative as one might expect. The third column adds the average wealth-income ratio as a predictor for the risky share, and finds a negative effect. After controlling for the wealth-income ratio, income risk has a significantly positive effect on the risky share. This finding suggests that households with risky income tend to have lower risk aversion, as Section 5 will confirm.

The negative effects of age and the wealth-income ratio on the risky share are consistent with the predictions of a simple static model in which labor income is safe and tradable, so that human capital is an implicit cash holding that tilts the composition of the financial portfolio towards risky assets (Bodie, Merton, and Samuelson 1992, Campbell and Viceira 2002).[8] We work with a richer lifecycle model in which labor income is risky and nontradable, but that model implies a similar pattern of age and wealth effects on the risky share.

3.2 Life-Cycle Model

We consider a standard life-cycle model, very similar to the one in Cocco, Gomes and Maenhout (2005) and Gomes and Michaelides (2005). Households have finite lives and Epstein-Zin utility over a single consumption good. The utility function V_t is specified by the coefficient of relative risk aversion γ, the time discount factor δ or equivalently the time preference rate $-\log(\delta)$, and the elasticity of intertemporal substitution ψ. The utility V_t satisfies the recursion

$$V_t = \left[C_t^{1-1/\psi} + \delta \left(\mathbb{E}_t p_{t,t+1} V_{t+1}^{1-\gamma} \right)^{(1-1/\psi)/(1-\gamma)} \right]^{\frac{1}{1-1/\psi}}, \tag{5}$$

where $p_{t,t+1}$ denotes the probability that a household is alive at age $t+1$ conditional on being alive at age t, calibrated from Sweden's life tables. Utility, consumption, and the preference parameters γ, δ, and ψ all vary across households but we suppress the household index h in equation (5) for notational simplicity.

The wealth accumulation of young households is significantly influenced by housing purchases, transfers from relatives, investments in education, or changes in family size, which for tractability we do not include in our model. For this reason, we focus on the stage of the life-cycle during which households have substantial retirement

[8] At first glance the negative effect of the wealth-income ratio on the risky share may appear to contradict evidence that wealthier individuals take more financial risk (Carroll 2002, Wachter and Yogo 2010, Calvet and Sodini 2014). The discrepancy is likely due to several factors. Our sample excludes non-participants in risky financial markets and the wealthiest 1% of Swedish households; we measure the risky portfolio share taking account of housing and leverage through mortgage borrowing; and we predict the risky share using the wealth-income ratio rather than the absolute level of wealth.

saving and initialize our model at age 40. The time index in the model, t, starts at 1, so that t is calendar age minus 39. Each period corresponds to one year and agents live for a maximum of $T = 61$ periods (corresponding to age 100).

Matching the behavior of retirees is also hard for simple life-cycle models that do not incorporate health shocks or bequest motives. For this reason, we only consider the model's implications for ages 40 to 60. Our estimates of the TPR can be viewed as having a downward bias due to the absence of a bequest motive in the model.

Before retirement households supply labor inelastically. The stochastic process of labor income, $Y_{h,t}$, is described in Section 3.1. All households retire at age 65, as was typically the case in Sweden during our sample period, and we set retirement earnings equal to a constant replacement ratio of the last working-life permanent income.

Consistent with the discussion in Section 2, wealth in the model is invested every period in a one-period riskless asset (bond) and a composite risky asset. In each period we recalibrate beginning-of-period wealth to the level observed in the data and use the model to predict end-of-period wealth.

The household chooses its consumption level $C_{h,t}$ and risky portfolio share $\alpha_{h,t}$ subject to a constraint that financial wealth is positive—that is, the household cannot borrow to finance consumption. We do allow borrowing to finance a risky asset position, that is, we allow $\alpha_{h,t} \geq 1$. Household wealth satisfies the budget constraint

$$W_{h,t+1} = (R_f + \alpha_{h,t} R^e_{N,t+1})(W_{h,t} + Y_{h,t} - C_{h,t}), \tag{6}$$

where $R^e_{N,t+1}$ is the return on the composite numeraire asset in excess of the gross risk free rate R_f. The excess return $R^e_{N,t+1}$ is Gaussian $\mathcal{N}(\mu_r, \sigma_r^2)$.

3.3 Calibrated Parameters

The parameters of our life-cycle model can be divided into those describing the income process, and those describing the properties of asset returns. For income, we have age profiles and retirement replacement ratios as illustrated in appendix Figure A.2, and the standard deviations of permanent group-level, permanent idiosyncratic, and transitory idiosyncratic income shocks reported in Table 2.

In our model we assume that all safe borrowing and lending takes place at a single safe interest rate of 2.0%. This is calibrated as a weighted average of a safe lending rate of 0.8% and the average household borrowing rate of 3.6%, using the cross-sectional average household debt level to construct the average.[9]

We set the volatility of the numeraire risky asset at 13.3%, which is equal to the volatility of post-tax excess stock returns as discussed in section 2.3. We assume that the average excess return on the numeraire asset over the safe interest rate is 3.5%, the same as the average post-tax equity premium on the MSCI world index in local currency over the period 1984–2007. Putting these assumptions together, we assume a Sharpe ratio of 0.26. In section 5.5 we discuss robustness of our results to assuming a higher Sharpe ratio of 0.40 or a lower Sharpe ratio of 0.15.

Following Campbell, Cocco, Gomes, and Maenhout (2001), we estimate the correlation between the numeraire risky asset return and group-level income shocks by lagging the risky asset return one year to capture a delayed response of income to macroeconomic shocks that move asset prices immediately. Empirically the correlation is very similar across the 9 education-sector categories, with an average value of 0.27 for stock returns, and 0.87 for real estate returns. A weighted-average of the respective covariances yields a correlation of 0.44 which is the value we used in our model for all groups. It is important to note that the correlation between the numeraire risky asset return and individual income growth is much smaller than 0.44, because most individual income risk is idiosyncratic.

[9] Our model would allow us to assume that households pay a higher rate when they borrow to buy the numeraire asset (that is, when they have a risky share greater than one). However, this assumption would not be a better approximation to reality than the one we make, since households who borrow to buy housing pay the borrowing rate even when their risky share is below one.

4 Identification and Estimation

Our goal is to estimate the three preference parameters of the Epstein-Zin utility model. The main challenge is to separately identify the TPR and the EIS, using data on portfolio risk and wealth accumulation.

The Euler equation for the return on the optimal portfolio is given by

$$1 = \mathrm{E}_t \left[\widetilde{\delta}_{t+1} \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{1}{\psi}} \left(\frac{V_{t+1}}{\mu(V_{t+1})} \right)^{\frac{1}{\psi}-\gamma} R^P_{t+1} \right] \tag{7}$$

where $\widetilde{\delta}_{t+1} = \delta p_{t,t+1}$, $R^P_{t+1} = R_f + \alpha R^e_{t+1}$, and $\mu(V_{t+1})$ denotes the certainty equivalent of V_{t+1}.[10] Under the usual assumption of conditional joint lognormality, we obtain

$$\begin{aligned} \mathrm{E}_t g_{t+1} \quad = \quad & \psi[\mathrm{E}_t r^P_{t+1} - \log(\widetilde{\delta}_{t+1})] + (1-\gamma\psi)\, \mathrm{E}_t \widetilde{v}_{t+1} + \frac{1}{2\psi}\sigma^2_{g,t} + \sigma_{gr,t} \\ & + \frac{\psi}{2} \left[\left(\frac{1}{\psi} - \gamma \right)^2 \sigma^2_{\widetilde{v},t} + \sigma^2_{r,t} + \left(\frac{1}{\psi} - \gamma \right) \sigma_{\widetilde{v}r,t} \right] + \left(\frac{1}{\psi} - \gamma \right) \sigma_{g\widetilde{v},t}, \end{aligned} \tag{8}$$

where lower case letters denote logs of upper case letters, $g_{t+1} = \log(C_{t+1}/C_t)$, and $\widetilde{V}_{t+1} = V_{t+1}/\mu(V_{t+1})$.

Equation (8) highlights the identification problem. If the expected portfolio return, the time discount factor, and the conditional variances are constant over time, then the expected consumption growth rate $\mathrm{E}_t g_{t+1}$ is constant and for any value of ψ there is a corresponding time discount factor δ that delivers the same level of $\mathrm{E}_t g_{t+1}$. Without additional restrictions on δ or ψ these two parameters cannot be separately identified, as shown by Kocherlakota (1990) and Svensson (1989).

Equation (8) also suggests three possible solutions. First, one can exploit time-variation in variance terms, which arises in life-cycle models with undiversifiable risky labor income such as ours. However these changes tend to be more substantial early

[10]This Euler equation holds with equality even though our model has borrowing constraints, because with labor income risk and a Bernoulli utility function that satisfies $u'(0) = \infty$ the agent will always choose to hold some financial assets. Our model also has short-sales constraints on risky asset holdings, but these do not bind for the middle-aged households we are considering.

in life, when households have less wealth to smooth shocks (Gomes and Michaelides 2005). A second channel is time variation in the expected portfolio return. Even though our model has no exogenous variation in expected asset returns, we have endogenous variation driven by changes in the agent's portfolio as a function of age. The third channel is time variation in the effective time discount factor $\widetilde{\delta}_{t+1} = \delta p_{t,t+1}$, driven by the survival probabilities $p_{t,t+1}$ which are also a function of age.

All three sources of variation imply that the profile of the wealth-income ratio is affected in different ways by the TPR and the EIS, at different ages. Our identification strategy builds on this intuition. In the appendix (section 2 and Table A.3), we confirm the good identification of the preference parameters in simulated data.

4.1 Indirect Inference Estimator

The estimation of the vector of preference parameters, $\theta^g = (\delta^g, \gamma^g, \psi^g)'$, in each group g proceeds by indirect inference (Smith 1993, Gouriéroux, Monfort, and Renault 1993). This method compares a vector of auxiliary statistics produced by the lifecycle model to the vector of empirical auxiliary statistics in the group. We denote by $p = 3$ the number of components of θ^g, by N^g the number of households in the group, and by $T = 8$ the number of years in the panel.

For every $t \in \{1, ..., T\}$, we consider the following auxiliary statistics: (i) the wealth-income ratio of the group, defined as the ratio of the group's total wealth to the group's total income:

$$\hat{\mu}^g_{1,t} = \frac{\sum_{h=1}^{N^g} W_{h,t}}{\sum_{h=1}^{N^g} Y_{h,t}}, \tag{9}$$

and (ii) the group's risky share:

$$\hat{\mu}^g_{2,t} = \frac{\sum_{h=1}^{N^g} \alpha_{h,t} W_{h,t}}{\sum_{h=1}^{N^g} W_{h,t}}. \tag{10}$$

The auxiliary statistics $\hat{\mu}^g_{1,t}$ and $\hat{\mu}^g_{2,t}$ provide reliable measures of wealth accumulation and risk-taking based on group aggregates. We note that $\hat{\mu}^g_{1,t}$ and $\hat{\mu}^g_{2,t}$ can be interpreted as ratios of sample moments but are not sample moments themselves, which motivates the use of indirect inference rather than moment-based estimators. We stack these auxiliary statistics into the *empirical auxiliary estimator*

$$\hat{\mu}^g = (\hat{\mu}^g_{1,1}, \ldots, \hat{\mu}^g_{1,T}, \hat{\mu}^g_{2,1}, \ldots, \hat{\mu}^g_{2,T})'.$$

By construction, $\hat{\mu}^g$ has $q = 16$ components.[11]

The data-generating process is based on the policy functions of households with preference parameter vector θ, the return process, and the labor income process defined in earlier sections. As the number of households in the group goes to infinity, the empirical auxiliary estimator $\hat{\mu}^g$ converges to the *binding function* $\mu^g(\theta) \in \mathbb{R}^q$ with components $\mu^g_{1,t}(\theta) = \mathrm{E}^g_\theta(W_t)/\mathrm{E}^g_\theta(Y_t)$ and $\mu^g_{2,t}(\theta) = \mathrm{E}^g_\theta(\alpha_t W_t)/\mathrm{E}_\theta(W_t)$, where $\mathrm{E}^g_\theta(\cdot)$ denotes the cross-sectional mean of households in the group. These expectations are computed under the assumption that all households in the group earn the riskfree rate R_f and the synthetic excess risky return $R^e_{N,t}$ on their risky asset holdings.

We generate the simulated auxiliary estimator as follows. Using the parameters from Table 2 as inputs, we solve the life-cycle model for the 4151 different household groups. For each group g and preference parameter θ, we compute the wealth-income ratio and risky share predicted by the model for the years 2000 to 2007. For each year t, the starting point is an information set at the end of year $t-1$ containing: (a) the empirical wealth-income ratio of group g at the end of year $t-1$, and (b) the empirical group-level income shock in year t. Consistent with the life-cycle model, we assume that households have this much advance information about wages and hours. We then feed in year t's empirical risky asset return and group-level income shock and simulate the idiosyncratic permanent and transitory income shocks of each household in the group, which we use to simulate the group's wealth-income ratio and risky share at the end of year t. More specifically, these simulations proceed in four steps.

i. We simulate $S = 10,000$ households/paths in the group over year t. For each simulated unit $i \in \{1..., S\}$, we simulate labor income and permanent income $(\tilde{Y}_{i,t}, \tilde{Y}^P_{i,t})$ in period t. We set wealth at the beginning of period t, $\tilde{W}_{i,t-1}$, equal to

[11] We could also include the risky share in the initial year (α_{i0}), since it is also an endogenous moment from the simulations. We exclude it in order to have an equal number of auxiliary statistics related to the wealth-income ratio and to the risky share.

$\tilde{Y}^P_{i,t}$ times the group's average wealth-income ratio at the end of period $t-1$. Using the lifecycle model's policy functions $\alpha^*_t(\,\cdot\,)$ and $C^*_t(\,\cdot\,)$, we compute the risky share, $\tilde{\alpha}_{i,t-1} = \alpha^*_t(\tilde{Y}_{i,t}, \tilde{W}_{i,t-1}, \tilde{Y}^P_{i,t}; \theta)$, and consumption, $\tilde{C}_{i,t} = C^*_t(\tilde{Y}_{i,t}, \tilde{W}_{i,t-1}, \tilde{Y}^P_{i,t}; \theta)$, of each simulated unit during year t. Consistent with the model, the simulated unit sets both quantities at the end of year $t-1$ and keeps them constant during year t.[12]

ii. We compute the predicted wealth of each simulated unit at the end of year t :

$$\hat{W}_{i,t} = (R_f + \tilde{\alpha}_{i,t-1} R^e_{p,t})(\tilde{W}_{i,t-1} + \tilde{Y}_{i,t} - \tilde{C}_{i,t}), \tag{11}$$

as equation (6) implies. The prediction incorporates empirical financial returns in year t, assumed equal for all groups.

iii. We obtain the group's predicted wealth-income ratio at the end of year t:

$$\widetilde{\mu}^g_{1,t}(\theta) = \frac{\sum_{i=1}^S \hat{W}_{i,t}}{\sum_{i=1}^S \tilde{Y}_{i,t}}. \tag{12}$$

iv. We observe the information set available at the end of year t, we sample S households, and we compute the group's predicted risky share at the end of year t:

$$\widetilde{\mu}^g_{2,t}(\theta) = \frac{\sum_{i=1}^S \tilde{\alpha}_{i,t} \tilde{W}_{i,t}}{\sum_{i=1}^S \tilde{W}_{i,t}}. \tag{13}$$

We stack the resulting values into the column vector $\tilde{\mu}^g_S(\theta)$.

We estimate the vector of preference parameters by minimizing the deviation $\tilde{\mu}^g_S(\theta) - \hat{\mu}^g$ between the lifecycle model and the data:

$$\hat{\theta}^g = \arg\min_\theta \left[\tilde{\mu}^g_S(\theta) - \hat{\mu}^g\right]' \Omega \left[\tilde{\mu}^g_S(\theta) - \hat{\mu}^g\right], \tag{14}$$

where Ω is a weighting matrix. The indirect inference estimator $\hat{\theta}^g$ is overidentified since we employ $q = 16$ auxiliary statistics to estimate $p = 3$ structural parameters.

[12] This methodology exploits the homogeneity of $\alpha^*_t(\,\cdot\,)$ and $C^*_t(\,\cdot\,)$, with respect to $(\tilde{Y}_{i,t}, \tilde{W}_{i,t-1}, \tilde{Y}^P_{i,t})$.

We use a diagonal weighting matrix Ω that is common to all groups. Each diagonal element of Ω is a scale factor that converts the wealth-income ratios and risky shares into comparable units. Specifically, we let $\Omega = diag(\omega_1, ..., \omega_1, \omega_2, ..., \omega_2)$, where $\omega_1 = \left(\frac{1}{GT} \sum_{g=1}^{G} \sum_{t=1}^{T} \hat{\mu}_{1,t}^{g} \right)^{-2}$ and $\omega_2 = \left(\frac{1}{GT} \sum_{g=1}^{G} \sum_{t=1}^{T} \hat{\mu}_{2,t}^{g} \right)^{-2}$. These weights have $(\omega_2/\omega_1)^{1/2} = 7.57$, consistent with an average risky share of around 0.5 and an average wealth-income ratio of 3.5. Using a common weighting matrix Ω implies that the objective function in (14) is comparable across groups. Further details of the numerical implementation of our estimation methodology are given in the online appendix.

4.2 Asymptotic Properties

If our model is correctly specified, the indirect inference estimator obtained by this procedure converges to the true preference parameter vector as the number of households in each group increases, as we now show.

The empirical auxiliary estimator $\hat{\mu}^g$ is asymptotically normal:

$$\sqrt{N^g} \left[\hat{\mu}^g - \mu^g(\theta) \right] \xrightarrow{d} \mathcal{N}(0, W_g) \tag{15}$$

as the group size N^g goes to infinity. This result follows from the delta method and the fact that the auxiliary statistics (9) and (10) can be interpreted as ratios of sample moments. We estimate the asymptotic variance covariance matrix of $\hat{\mu}^g$ by the jackknife estimator

$$\frac{\hat{W}_g}{N^g} = \frac{N^g - 1}{N^g} \sum_{i=1}^{N^g} (\hat{\mu}_{[i]}^J - \overline{\mu^J})(\hat{\mu}_{[i]}^J - \overline{\mu^J})', \tag{16}$$

where $\hat{\mu}_{[i]}^J$ is the auxiliary estimator obtained by excluding the i^{th} household, and $\overline{\mu^J} = (N^g)^{-1} \sum \hat{\mu}_{[i]}^J$.

The indirect inference estimator is asymptotically normal:

$$\sqrt{N^g} \left(\hat{\theta}^g - \theta^g \right) \to^d \mathcal{N}(0, V^g), \tag{17}$$

as Gouriéroux, Monfort, and Renault (1993) show. Furthermore, the asymptotic variance-covariance matrix is given by

$$V^g = (1 + s_g^{-1})\,(D_g\,\Omega\,D_g')^{-1}\,D_g\,\Omega\,W_g\,\Omega\,D_g'(D_g\,\Omega\,D_g')^{-1}, \tag{18}$$

where the ratio $s_g = S/N_g$ accounts for simulation noise and $(D_g)' = \partial\mu^g(\theta^g)/\partial\theta'$ is the Jacobian matrix of the binding function $\mu^g(\,\cdot\,)$ evaluated at the true parameter θ^g. In practice, we estimate the asymptotic variance-covariance matrix of V_g by its sample equivalent $\hat{V}^g = (1+s_g^{-1})\,(\hat{D}_g\,\Omega\,\hat{D}_g')^{-1}\,\hat{D}_g\,\Omega\,\hat{W}_g\,\Omega\,\hat{D}_g'(\hat{D}_g\,\Omega\,\hat{D}_g')^{-1}$, where $\hat{D}_g$ is a finite-difference approximation of D_g that we discuss in the appendix.

When the size of each group g is large, we could achieve efficient estimation by setting the second-stage weighting matrix equal to the inverse of the jackknife estimator: $\Omega^{(2)} = \hat{W}_g^{-1}$, and then solving the optimization problem (14). Efficient estimation, however, is not feasible in our sample because most groups are too small to obtain a reliable estimator of W_g^{-1}. The median group size is 63, while the symmetric matrix W_g contains 136 ($= 16 \times 17/2$) distinct elements. A related problem is that in many groups, the weighting matrix $\Omega^{(2)} = \hat{W}_g^{-1}$ assigns almost all the weight to the risky share, while the wealth-income ratio plays essentially no role in estimation. Efficient estimation is therefore unsatisfactory in our sample on statistical and economic grounds.[13] For these reasons, we henceforth focus on one-step estimation based on the diagonal weighting matrix Ω defined in Section 4.1. Since this approach does not provide global specification tests based on the value of the objective function (14), we focus on measures of fit based on root mean squared error or economic significance.

[13]These difficulties are consistent with the finite-sample inaccuracy of two-step generalized method of moments studied in Hwang and Sun (2018).

5 Empirical Results

5.1 The Cross-Sectional Distribution of Preference Estimates

Tables 4 and 5 and Figure 1 summarize the cross-sectional distributions of our estimated preference parameters. Table 4 reports the cross-sectional means, medians, and standard deviations of the estimated parameters along with summary statistics of the data: the average risky share, the initial wealth-income ratio, and the growth rate of the wealth-income ratio. Table 5 reports the cross-sectional correlations of the estimated parameters and summary statistics. A number of interesting patterns are visible in these tables.

Table 4 reports a moderate mean RRA of 5.24, close to the median estimate of 5.30 and in the middle of the range from 1 to 10 defined as reasonable by Mehra and Prescott (1985). The cross-sectional standard deviation of estimated RRA is modest at 0.47, less than a tenth of the mean and median estimates.

It may at first seem puzzling that the cross-sectional standard deviation of RRA is lower in proportional terms than the cross-sectional standard deviation of the risky portfolio share, which was shown in Table 1 to be almost one-third of its mean. In a simple one-period portfolio choice model without labor income, the risky portfolio share and RRA are inversely proportional to one another so they must have equal proportional standard deviations; and the same is true in a model where labor income is safe and all investors have the same wealth-income ratio. Two features of our model help to account for this finding. First, there is variation across groups in their wealth-income ratios which helps to account for some of the cross-sectional variation in risky shares as illustrated in Table 3. Second, we estimate that labor income risk is correlated with financial risk; this increases the change in the risky financial share that is needed to generate a given change in a household's overall risk exposure.

The other two preference parameters have much greater cross-sectional variation, and both are strongly right-skewed. The median TPR is 4.08%, considerably lower than the mean of 6.18%, and the cross-sectional standard deviation of TPR is 6.03%.

Table 4: Size-Weighted Cross-Sectional Distributions of Estimated Preference Parameters and Group Financial Characteristics

	Mean	Median	Std. Dev.	10%	25%	75%	90%
RRA	5.24	5.30	0.47	4.60	4.90	5.50	5.90
TPR (%)	6.18	4.08	6.03	1.01	3.15	7.42	18.39
EIS	0.99	0.42	0.96	0.10	0.10	1.97	2.50
Average RS	0.65	0.62	0.17	0.45	0.52	0.75	0.89
Initial WY	4.13	2.88	3.78	0.96	1.69	4.71	7.67
Growth of WY	1.08	1.07	0.05	1.03	1.05	1.10	1.14

This table reports the mean, median, standard deviation, and 10th, 25th, 75th, and 90th percentiles of estimated preference parameters and group financial characteristics. All statistics weight groups by their size to recover the underlying cross-sectional distributions at the household level. There are 4,151 groups containing 291,488 households.

Table 5: Size-Weighted Cross-Sectional Correlations of Estimated Preference Parameters and Group Financial Characteristics

	RRA	TPR	EIS	Average RS	Initial WY	Growth of WY
RRA	1.000					
TPR	-0.275***	1.000				
EIS	0.113***	-0.289***	1.000			
Average RS	-0.412***	0.472***	-0.155***	1.000		
Initial WY	-0.088***	-0.296***	0.320***	-0.524***	1.000	
Growth of WY	0.003	0.352***	-0.179***	0.621***	-0.685***	1.000

This table reports the cross-sectional correlations across estimated preference parameters and group financial characteristics. Correlations weight groups by their size to recover the underlying cross-sectional correlations at the household level. Statistical significance levels of correlation coefficients are indicated with stars: * denotes 1-5%, ** 0.1-1%, *** less than 0.1% significance. There are 4,151 groups containing 291,488 households.

Similarly, the median EIS is 0.42, considerably lower than the mean of 0.99, and the cross-sectional standard deviation of the EIS is 0.96. This cross-sectional standard deviation is twice as large for the EIS as for RRA in absolute terms, and over 20 times as large in proportional terms; this contrasts with the prediction of a power utility model, which would imply equal proportional standard deviations for RRA and the EIS since one parameter is the reciprocal of the other.[14]

Table 5 shows that our estimates of preference parameters are only weakly cross-sectionally correlated. RRA and the EIS have a low positive correlation of 0.11, a

[14] Appendix Figure A.3 plots the univariate distributions of all three preference parameters.

Figure 1: Joint Distribution of TPR and EIS

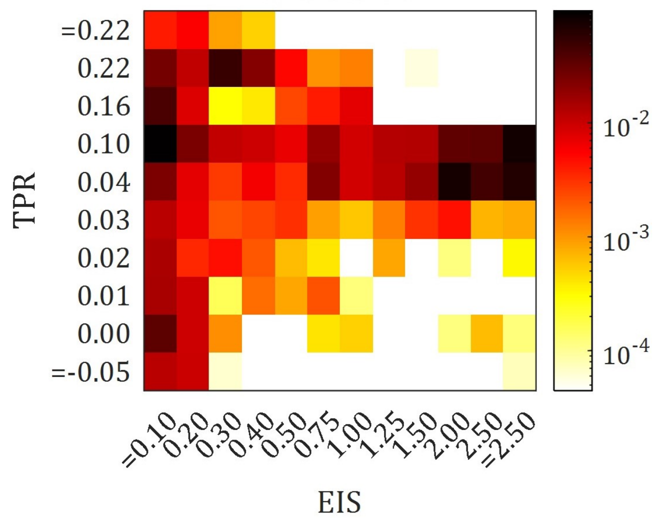

This figure presents bivariate heat map for estimates of TPR and EIS across 4,151 groups of Swedish households, size-weighted to recover the underlying distribution across households under the assumption that preferences are homogeneous within groups. Each axis label shows the upper cutoff value of the corresponding bin, except for labels beginning with = which indicate that the bin contains only estimates of the exact value indicated by the label. The logarithmic color scheme indicates the fraction of the sample in each bin. This fraction is 9.6% for the darkest color and 0.0% for the brightest color.

finding that contrasts with the perfect negative correlation between the logs of RRA and the EIS implied by power utility. Both RRA and the EIS are negatively correlated with the rate of time preference, but the correlations are modest at -0.28 and -0.29 respectively. These relatively weak correlations imply that heterogeneity in household preferences is multi-dimensional and cannot be explained by any single factor missing from our model such as heterogeneity in beliefs about the equity premium.

Figure 1 is a heat map of the bivariate distribution of the TPR and EIS. The distribution of the EIS is U-shaped, with probability mass concentrated below 1 and near the upper edge of our parameter space which we set to 2.5. The distribution of the TPR is also dispersed, but the figure shows that more extreme values of the TPR are associated with low values of the EIS. This makes sense since equation (8) implies that a low EIS reduces the impact of the TPR on observable savings decisions.

5.2 Preference Estimates and Household Characteristics

Our parameter estimates have interesting correlations with observable variables that help us understand what drives the estimates. The lower portion of Table 5 explores these correlations. Looking first at the correlations among observables, the initial wealth-income ratio has a correlation of -0.52 with the average risky share and a correlation of -0.69 with the average growth rate of the wealth-income ratio. These correlations are consistent with the predictions of our life-cycle model that the risky share declines with the level of financial wealth in relation to human capital, and that households that enter the sample with low financial wealth have a strong motive to accumulate wealth to finance retirement. Unsurprisingly given these patterns, the average risky share and the average growth rate of the wealth-income ratio have a strong positive correlation of 0.62.

Our estimate of RRA is negatively correlated (-0.41) with the average risky share, an intuitive result that is also consistent with our identification analysis. RRA is also weakly negatively correlated (-0.09) with the initial wealth-income ratio. Mechanically, this reflects the fact that households who enter the sample with high wealth have risky shares that are insufficiently lower than the risky shares of other households to be consistent with the same level of RRA.

Our estimate of the TPR is negatively correlated (-0.30) with the initial wealth-income ratio and positively correlated (0.35) with the average growth rate of the wealth-income ratio in our sample period. Mechanically, this is due to the fact that households that enter our sample with low initial wealth accumulate wealth more rapidly than average households, but not as rapidly as they would do if they had an average TPR. It is intuitive that impatient households accumulate less wealth before age 40 and then belatedly catch up as retirement approaches. The TPR is also positively correlated (0.47) with the average risky share, reflecting the lower wealth-income ratio of impatient households that justifies a riskier investment strategy.

Our estimate of the EIS is positively correlated (0.32) with the initial wealth-income ratio and negatively correlated (-0.18) with the average growth rate of the wealth-income ratio. Economically, this suggests that households with a high EIS save for retirement early in life, before our sample begins; such households have a high willingness to adjust consumption to reach their target wealth-income ratio, whereas households with a low EIS save more gradually over time.[15]

We next ask how our estimates are related to households' income risk and education. Table 6 regresses preference estimates on labor income volatility, the level of education, and cohort fixed effects. RRA is most strongly related to these observables. Households with riskier labor income tend to have lower risk aversion. Mechanically, this results from the fact documented in Table 3 that income volatility has little effect on the risky share: if risk aversion were the same for safe and for risky occupations, then the risky share should fall with income risk. The finding suggests that risk-tolerant individuals self-select into risky occupations, or that households fail to fully understand the importance of income risk for investment. Beyond this, more educated people also tend to have lower RRA. The R^2 of the regression for risk aversion is almost 38%. Households with high income risk also tend to have a lower TPR, but the explanatory power of this regression is less than 7%. Finally, the EIS has almost no relation to income risk, education, or cohort fixed effects.

5.3 Parameter Uncertainty

The discussion in the previous subsections treats our point estimates of parameters as if they are equivalent to the parameters themselves. In this subsection we use our asymptotic standard errors to take parameter uncertainty into account.

In Table 7 we report hypothesis tests based on our asymptotic standard errors and using 5% significance levels. The top panel focuses on hypotheses about the TPR. The first row of the table reports that 6.5% of households are in groups estimated to have a negative TPR. The second row of the table reports that we can reject the null of a positive TPR at the conventional 5% significance level for only 2.1% of households; thus, a significantly negative TPR is a rare occurrence in our sample. The third row of the table shows that we can reject the null of a negative TPR at the

[15]Appendix Table A.4 presents similar results using multiple regressions rather than univariate correlations.

Table 6: Size-Weighted Cross-Sectional Regressions of Preference Parameters on Income Volatility and Education

	(1) RRA	(2) TPR	(3) EIS
Total income volatility	-6.752***	-0.334***	1.595*
	(0.239)	(0.042)	(0.721)
High school	-0.462***	0.023***	-0.123**
	(0.018)	(0.003)	(0.046)
Post-high school	-0.570***	0.009***	-0.051
	(0.015)	(0.003)	(0.046)
Constant	6.599***	0.132***	0.701***
	(0.053)	(0.010)	(0.156)
Cohort dummies	Yes	Yes	Yes
R^2	0.379	0.068	0.006

This table reports the cross-sectional regression coefficients across estimated preference parameters and group characteristics including the total income volatility (in natural units), and dummies for high-school and post-high-school education. All regressions weight groups by their size, to recover the underlying cross-sectional relationships at the household level. Standard errors are reported in parentheses and statistical significance levels are indicated with stars: * denotes 1-5%, ** 0.1-1%, *** less than 0.1% significance. There are 4,151 groups containing 291,488 households.

5% level for 61.1% of households, and the fourth row shows that we can reject the null of a zero TPR using a two-sided test for 58.7% of households. Thus the TPR is significantly positive for a majority of Swedes.

The next panel of the table focuses on the EIS. The first row reports that 59.7% of households are in groups with estimated EIS less than one. The following rows show that we can reject the null of an EIS greater than one for 49.0% of households, and can reject the null of an EIS less than one for only 17.2% of households. The asymmetry reflects the fact, illustrated in appendix Figure A.4, that the asymptotic standard error of the EIS is positively correlated with the level of the estimated EIS. Thus it is far more common for Swedish households to have an EIS significantly below one than an EIS significantly above one.

Turning to power utility in the third panel, 34.7% of households have an estimated EIS that is lower than the reciprocal of risk aversion. The table shows that we can

Table 7: Size-Weighted Statistical Test Results for Estimated Preference Parameters

Condition	Fraction of the Population
TPR < 0	0.065
Reject null that TPR > 0	0.021
Reject null that TPR < 0	0.611
Reject null that TPR = 0	0.587
EIS < 1	0.597
Reject null that EIS > 1	0.490
Reject null that EIS < 1	0.172
Reject null that EIS = 1	0.633
EIS < 1/RRA	0.347
Reject null that EIS > 1/RRA	0.054
Reject null that EIS < 1/RRA	0.331
Reject null that EIS = 1/RRA	0.340
Reject null that RRA = mean(RRA)	0.891
Reject null that TPR = mean(TPR)	0.472
Reject null that EIS = mean(EIS)	0.629
Reject the joint null of the above three rows	0.976
Reject null that RRA = median(RRA)	0.888
Reject null that TPR = median(TPR)	0.204
Reject null that EIS = median(EIS)	0.501

This table reports the size-weighted fraction of Swedish household groups, or equivalently the fraction of Swedish households, for which each condition stated in the row label applies. All hypothesis test rejections are at the 5% significance level. Hypothesis tests in the bottom panel treat the cross-sectional median preference parameter as known rather than estimated. There are 4,151 groups containing 291,488 households.

reject the null that EIS exceeds the reciprocal of RRA for 5.4% of households, and can reject the null that the EIS is lower than the reciprocal of RRA for 33.1% of households. We can reject the power utility null using a two-sided test for 34.0% of households.

The bottom two panels of the table test hypotheses about the cross-sectional dispersion of preferences. We report the fraction of households that are in groups for which we can reject the null that the group preference estimate equals the cross-sectional mean, taking account of statistical uncertainty about that mean. This is the case for 89.1% of the RRA estimates, 47.2% of the TPR estimates, and 62.9% of the EIS estimates. We can reject the null that all three parameters equal their cross-sectional means for 97.6% of households. Results are similar in the bottom panel where we test whether group preference estimates equal the cross-sectional median estimates, treating the medians as known for simplicity. Overall, the table presents strong statistical evidence against homogeneity of preferences within our framework.

A second use of our asymptotic standard errors is to adjust our estimates of the heterogeneity in true preference parameters. Table 4 and Figure 1 describe the cross-sectional distribution of our parameter estimates, but this is increased by noise in the estimation procedure. Since our asymptotic standard errors estimate the noise for each group, in principle we can correct for the effect of noise on the estimated cross-sectional variance of parameters by subtracting the cross-sectional average squared standard error from the cross-sectional variance of our estimates.

A practical difficulty in doing this is that some groups have extremely high standard errors. Although these high standard errors are not pervasive enough to undermine our ability to reject homogeneous preferences for most households in the group-specific tests reported in Table 7, they do have a strong influence on the cross-sectional average of squared standard errors. In fact, if we do not limit the influence of outliers the average squared standard error is higher than the cross-sectional variance of estimates for TPR and EIS, implying a negative cross-sectional variance for true TPR and EIS. We obtain more reasonable results if we winsorize the group-specific standard errors at the 90th percentile of the cross-sectional distribution. This procedure implies a cross-sectional standard deviation of 0.46 for RRA, 3.68% for the TPR, and 0.54 for the EIS, as compared with the cross-sectional standard deviations of estimates reported in Table 5 which are 0.47, 6.03%, and 0.96 respectively. The appendix (section 4 and Table A.5) provides further details.

5.4 Model Fit

In this subsection we consider measures of model fit more directly. We begin by describing the cross-sectional distribution of the errors our model makes in fitting the 16 auxiliary statistics that are the target of our estimation procedure. We take the 8 wealth-income ratios and the 8 risky shares, and for each of these variables

Table 8: Size-Weighted Cross-Sectional Distributions of Model Fit Indicators

	Mean	Median	Std. Dev.	10%	25%	75%	90%
WY RMSE	30.81	19.92	36.69	8.67	12.60	36.21	66.71
RS RMSE	4.66	4.14	2.13	2.48	3.13	5.69	7.43
Scaled WY RMSE	6.25	4.04	7.45	1.76	2.56	7.35	13.54
Scaled RS RMSE	7.16	6.37	3.27	3.81	4.82	8.75	11.42
RMSE-scaled OF	6.67	5.92	5.42	2.42	4.31	7.99	11.59

This table reports the mean, median, standard deviation, and 10th, 25th, 75th, and 90th percentiles of several measures of model fit. All statistics weight groups by their size to recover the underlying cross-sectional distributions at the household level. WY (RS) RMSE is the root mean squared error of the 8 WY (RS) moments used in estimation, multiplied by 100 so that the units are percentage points of income or wealth. Scaled WY RMSE divides by the cross-sectional mean of WY, 4.93, to express the WY RMSE in proportional percentage units. Scaled RS RMSE divides by the cross-sectional mean of RS, 0.65, to express the RS RMSE in proportional percentage units. RMSE-scaled OF (objective function) is the square root of the objective function divided by 4 and multiplied by 100 to express it in RMSE-equivalent percentage units. It differs slightly from the average of scaled WY and scaled RS RMSE because of interpolation in our estimation procedure. There are 4,151 groups containing 291,488 households.

we calculate the root mean squared error (RMSE), the square root of the average squared deviation of the model-fitted variable from the observed variable. The results are reported in percentage points in the first two rows of Table 8. The mean RMSE across all groups is 30.8% for the wealth-income ratio and 4.7% for the risky share. In other words, the average error in fitted wealth is just under 4 months of income and the average error in the risky share is just below 5% of wealth. The RMSE distribution is somewhat right-skewed as indicated by the fact that the median RMSEs are below the mean RMSEs at 19.9% and 4.1% respectively.

To interpret these numbers, we note that the standard deviation of the change in the wealth-income ratio, around a mean of zero, has an average across groups of 33.1% and a median of 30.5%. Thus our model has a slightly better mean performance and a much better median performance than an atheoretical random walk model for WY. The standard deviation of the risky share around its group-specific time-series mean has an average across groups of 6.2% and a median of 5.2%. Thus our model, which captures variation in the risky share with age and wealth accumulation, fits asset allocation better than an atheoretical model that simply captures the mean risky share for each group.

Our estimation procedure takes into account that the wealth-income ratio and the risky share have different units, and scales them in proportion to their grand cross-sectional means. The next two rows of Table 8 similarly divide the RMSEs for the wealth-income ratio and risky share by their grand means of 4.93 and 0.65, respectively, to express them in proportional units. The mean proportional RMSE is 6.3% for the wealth-income ratio and 7.2% for the risky share.

Finally, we report a transformation of the objective function that is rescaled to express it in RMSE-equivalent units. The objective function is the sum of squared proportional errors, so we divide by the number of auxiliary statistics (16) and take the square root, then multiply by 100 to express the RMSE-scaled objective function in percentage points. Group by group, the result is not exactly the average of the proportional errors for the wealth-income ratio and the risky share because of the interpolation method we use in estimation; and the quantiles of the cross-sectional distribution also may refer to different groups. Nonetheless, the bottom row of Table 8 is similar to an average of the previous two rows. The mean RMSE-scaled objective function is 6.7%, with a moderately right-skewed distribution.

In the appendix Table A.6, we show how model fit deteriorates if we suppress heterogeneity in preferences. The mean RMSE-scaled objective function more than doubles to 15.8% if we fix RRA at its cross-sectional mean. Fixing TPR at its cross-sectional mean produces a mean RMSE-scaled objective function of 7.9%. The smallest impact comes from restricting the EIS to its cross-sectional mean, which delivers a mean RMSE-scaled objective function of 7.6%. Both the TPR and the EIS have a strong impact on savings, so a model that fixes one of these can compensate to some degree by varying the other. Unsurprisingly, if we restrict both these parameters simultaneously to their cross-sectional means, the model fit deteriorates further and we get a mean RMSE-scaled objective function of 10.2%. Finally, fixing all parameters at their cross-sectional means is disastrous in the sense that it increases the mean RMSE-scaled objective function to 26.6%. A life-cycle model with homogeneous preferences, under our maintained assumption of homogeneous rational beliefs, delivers an extremely poor fit to the cross-section of household financial behavior.

5.5 Monte Carlo Analysis and Robustness

We evaluate the finite-sample performance of our procedure by a simple Monte Carlo exercise. For each group in our sample, we simulate our model under the group's initial conditions and the preference parameters we estimated for the group. We combine simulated households into hypothetical groups each containing N_g^* households, where N_g^* is a measure of the effective empirical group size. We repeat this procedure to obtain 1,000 hypothetical groups and calculate the mean parameter estimate. A comparison of this mean with the preference parameters under which the model was simulated allows us to assess finite-sample bias in our estimation method.

This Monte Carlo analysis does not fully capture the heterogeneity in household-level data, even under the assumption that our model holds without error at the household level and that all households in each group have identical preferences. This is because we simulate each household in the group assuming that the household has the group average wealth-income ratio at the start of the period. In the data, by contrast, and in the ergodic distribution of wealth-income ratios implied by the model, different households have different income and wealth levels at each point of time reflecting the influence of past idiosyncratic income shocks. Hence, the group average wealth-income ratio is more strongly influenced by those households with higher wealth. To partially capture this effect, we adjust our simulations to set the effective group size N_g^* equal to the reciprocal of the sum of squared wealth shares of individual households in the group, rather than the number of households in the group N_g. We find that N_g^* is on average about 3/4 of N_g, with relatively little variation in this ratio across groups.

In the appendix Table A.7 we report regression coefficients of Monte Carlo mean parameter estimates on the parameter estimates that were used to generate the simulated data ("true" parameters for the purpose of this exercise). The results are very good for RRA, which has a slope coefficient of 1.006, insignificantly different from one, and an R^2 statistic of 93%. The regression for TPR has a slope coefficient of 0.930 and an R^2 statistic of 87%. Results are not quite as good for EIS, which has a slope coefficient of 0.704 and an R^2 statistic of 61%. This regression places most of its weight on the high EIS estimates, which are noisy; but results are similar for the log of the EIS. An important lesson of these results is that small-sample bias cannot explain the substantial cross-sectional heterogeneity in our preference parameter estimates. There is almost no small-sample bias for RRA, and minimal bias for the TPR; and while there is some bias in our EIS estimates, a bias correction would have little effect on the cross-sectional dispersion of the EIS.

The appendix also reports several robustness checks. As a first check (Tables A.8-A.9), we show that when we exclude households in the top decile of initial wealth-income ratio, model fit improves considerably, while the distribution of preference parameters changes only slightly. As a second check (Tables A.10-A.11) we increase the assumed Sharpe ratio of the composite risky asset from 0.26 to 0.40. Unsurprisingly, this change increases the RRA needed to rationalize household portfolio choices. Our EIS estimates decline slightly, with the mean falling from 0.99 to 0.67, while the mean TPR increases from 6.2% to 8.7%. As a third check (Tables A.12-A.13), we consider a simple form of heterogeneity in beliefs by considering three alternative assumptions about the Sharpe ratio: the base value of 0.26, the high value of 0.40, and a low value of 0.15. Then, for each group we pick the assumption and preference parameters that minimize the objective function. The base case Sharpe ratio is selected for groups representing 49% of households, the low Sharpe ratio for 28% of households, and the high Sharpe ratio for 23% of households. Allowing for heterogeneity in household beliefs improves the fit of our model modestly. The cross-sectional standard deviation of RRA is almost four times larger, at 1.91, in the heterogeneous-beliefs model compared to our baseline, and the cross-sectional standard deviations of TPR and EIS are very slightly larger. The explanation is that the model uses heterogeneous beliefs to better fit wealth accumulation, and offsets belief heterogeneity with RRA heterogeneity to avoid counterfactual heterogeneity in the risky share.

6 Conclusion

In this paper, we have considered a life-cycle model of consumption-portfolio choice, which we have estimated on a panel of Swedish households. Our estimates of the RRA and EIS are almost uncorrelated across households, which contradicts the predictions of power utility. The TPR is negatively correlated with both risk aversion and the EIS. We estimate a negative correlation between income volatility and risk aversion. Income volatility is also negatively related to the TPR and weakly positively related to the EIS. More educated households tend to have lower risk aversion, controlling for income volatility, but we do not find that educated households are more patient; if anything, the relationship is the opposite when we control for income volatility.

Our results shed light on a number of issues in asset pricing and household finance. In general equilibrium asset pricing models, Epstein-Zin preferences are popular because they are scale-independent and therefore accommodate economic growth without generating trends in interest rates or risk premia. For this reason Epstein-Zin preferences have been assumed for a representative agent in many recent asset pricing papers. In particular, the long-run risk literature following Bansal and Yaron (2004) has argued that many asset pricing patterns are explained by a moderately high coefficient of relative risk aversion (typically around 10) and an EIS around 1.5. We estimate a lower cross-sectional average risk aversion around 5 and a cross-sectional average EIS considerably lower than that assumed in the long-run risk literature. We also estimate a cross-sectionally dispersed EIS such that relatively few households have an EIS between 1 and 2.

Even if individual households have constant preference parameters, cross-sectional heterogeneity in these parameters can break the relation between household preferences and the implied preferences of a representative agent. In a representative-agent economy, preferences with habit formation are needed to generate countercyclical variation in the price of risk (Constantinides 1990, Campbell and Cochrane 1999), but in heterogeneous-agent economies, countercyclical risk premia can arise from time-variation in the distribution of wealth across agents with different but constant risk preferences (Dumas 1989, Guvenen 2009). Gomes and Michaelides (2005 and 2008) illustrate the importance of preference heterogeneity for simultaneously matching the wealth accumulation and portfolio decisions of households. Our empirical evidence can be used to discipline these modeling efforts.

Importantly, we estimate multi-dimensional heterogeneity in preferences: the correlations among our estimated preference parameters are relatively low. This implies that a single factor omitted from our model, such as heterogeneity in expected stock returns of the sort documented in survey data by Vissing-Jørgensen (2003), Meeuwis, Parker, Schoar, and Simester (2018), and Giglio, Maggiori, Stroebel, and Utkus (2019) is unlikely to reconcile the data with homogeneous underlying preferences.

In household finance, there is considerable interest in estimating risk aversion at the individual level and measuring its effects on household financial decisions. This has sometimes been attempted using direct or indirect questions in surveys such as the Health and Retirement Study (Barsky, Juster, Kimball, and Shapiro 1997), the Survey of Consumer Finances (Vissing-Jørgensen 2002b), and similar panels overseas (Guiso and Paiella 2006). One difficulty with these attempts is that even if risk aversion is correctly measured through surveys, its effects on household decisions will be mismeasured if other preference parameters or the properties of labor income covary with risk aversion. Our estimates suggest that this should indeed be a concern.

Similarly, there is interest in measuring the effects of labor income risk on households' willingness to take financial risk (Calvet and Sodini 2014, Guiso, Jappelli, and Terlizzese 1996, Heaton and Lucas 2000). Models such as those of Campbell, Cocco, Gomes, and Maenhout (2001) and Cocco, Gomes, and Maenhout (2005) show the partial effect of labor income risk for fixed preference parameters, which will be misleading if risk aversion or other parameters vary with labor income risk. Our estimates suggest that this too is a serious empirical issue.

Our findings may also contribute to an ongoing policy debate over approaches to consumer financial protection. If all households have very similar preference parameters, strict regulation of admissible financial products should do little harm to households that optimize correctly, while protecting less sophisticated households from making financial mistakes. To the extent that households are heterogeneous, however, such a stringent approach is likely to harm some households by eliminating financial products that they prefer (Campbell 2016).

In this paper we have focused on households that participate in risky asset markets outside their retirement accounts. An interesting extension of our work will be to estimate preference parameters for non-participants, although such an analysis may require taking into account fixed costs of entering risky financial markets of the sort considered by Haliassos and Bertaut (1995), Vissing-Jørgensen (2002), Ameriks and Zeldes (2004), Gomes and Michaelides (2005), and Fagereng, Gottlieb, and Guiso (2017). A related exercise will be to estimate the effect of self-employment (private business ownership) on financial decisions.

Our model omits some other features of the household decision problem that may potentially be important and deserve further research. We assume that financial market returns are iid rather than time-varying, ignoring intertemporal effects on asset allocation discussed by Campbell and Viceira (2002). In our base case we assume homogeneous beliefs about financial market returns. As a robustness check we have allowed for a simple form of belief heterogeneity, but certainly it would be possible to go further in this direction, for example along the lines suggested by Calvet, Célérier, Sodini, and Vallée (2020). We fix preference parameters for each household and do not allow them to vary with wealth at the household level, contrary to evidence that relative risk aversion declines with wealth (Carroll 2002, Wachter and Yogo 2010, Calvet and Sodini 2014). We model labor income risk using normally distributed shocks rather than the skewed distributions estimated by Guvenen, Ozkan, and Song (2014) and others. We treat labor income as exogenous and do not consider the possibility that the household can endogenously vary its labor supply (Bodie, Merton, and Samuelson 1992, Gomes, Kotlikoff and Viceira 2008). We ignore the possibility that some components of consumption involve precommitments or generate habits that make them costly to adjust (Gomes and Michaelides 2003, Chetty and Szeidl 2007, 2010). We do not model homeownership jointly with other financial decisions as in Cocco (2005). Household-level data on asset allocation and wealth accumulation and our structural approach to estimation of a life-cycle model provide a laboratory in which these aspects of household financial decision-making can be explored.

References

Ameriks, John, and Stephen Zeldes, 2004, How do household portfolio shares vary with age?, unpublished paper, Columbia University.

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini, 2020, Rich pickings? Risk, return, and skill in household wealth, *American Economic Review* 110, 2703–2747.

Bansal, Ravi and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Barsky, Robert, Thomas Juster, Miles Kimball, and Matthew Shapiro, 1997, Preference parameters and behavioral heterogeneity: An experimental approach in the Health and Retirement Study, *Quarterly Journal of Economics* 112, 537–579.

Betermier, Sebastien, Laurent E. Calvet, and Paolo Sodini, 2017, Who are the value and growth investors? *Journal of Finance* 72, 5–46.

Bodie, Zvi, Robert C. Merton, and William Samuelson, 1992, Labor supply flexibility and portfolio choice in a life cycle model, *Journal of Economic Dynamics and Control* 16, 427–49.

Calvet, Laurent E., John Y. Campbell, Francisco Gomes, and Paolo Sodini, 2021, Online appendix to the cross-section of household preferences, unpublished paper.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, 2007, Down or out: Assessing the welfare costs of household investment mistakes, *Journal of Political Economy* 115, 707–747.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, 2009a, Fight or flight? Portfolio rebalancing by individual investors, *Quarterly Journal of Economics* 124, 301–348.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, 2009b, Measuring the financial sophistication of households, *American Economic Review Papers and Proceedings* 99, 393–98.

Calvet, Laurent E., Claire Célérier, Paolo Sodini, and Boris Vallée, 2020, Can security design foster household risk-taking?, unpublished paper.

Calvet, Laurent E., and Veronika Czellar, 2015, Through the looking glass: Indirect inference via auxiliary equilibria, *Journal of Econometrics* 185, 343–358.

Calvet, Laurent E. and Paolo Sodini, 2014, Twin picks: Disentangling the determinants of risk taking in household portfolios, *Journal of Finance* 69, 867–906.

Campbell, John Y., 2016, Restoring rational choice: The challenge of consumer financial regulation, Ely Lecture, *American Economic Review Papers & Proceedings* 106(5), 1–30.

Campbell, John Y., Joao Cocco, Francisco Gomes, and Pascal Maenhout, 2001, Investing retirement wealth: A life-cycle model, in John Y. Campbell and Martin Feldstein eds., *Risk Aspects of Investment-Based Social Security Reform* (University of Chicago Press), 439–473.

Campbell, John Y., and John Cochrane, 1999, By force of habit: A consumption based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Campbell, John Y. and Luis Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors* (Oxford University Press).

Carroll, Christopher D., 2002, Portfolios of the rich, in Luigi Guiso, Michael Haliassos and Tullio Jappelli, eds. *Household Portfolios* (MIT Press).

Carroll, Christopher D., and Andrew Samwick, 1997, The nature of precautionary wealth, *Journal of Monetary Economics* 40, 41–71.

Castaneda, Ana, Javier Diaz-Gimenez, and Jose-Victor Rios-Rull, 2003, Accounting for the U.S. earnings and wealth inequality, *Journal of Political Economy*, 111, 818–857.

Chetty, Raj, and Adam Szeidl, 2007, Consumption commitments and risk preferences, *Quarterly Journal of Economics* 122, 831–877.

Chetty, Raj, and Adam Szeidl, 2016, Consumption commitments and habit formation, *Econometrica* 84, 855–890.

Cocco, Joao, 2005, Portfolio choice in the presence of housing, *Review of Financial Studies* 18, 535–567.

Cocco, Joao, Francisco Gomes, and Pascal Maenhout, 2005, Consumption and portfolio choice over the life cycle, *Review of Financial Studies* 18, 491–533.

Constantinides, George, 1990, Habit formation: A resolution of the equity premium puzzle, *Journal of Political Economy* 98, 519–543.

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

Dumas, Bernard, 1989, Two-person dynamic equilibrium in the capital markets, *Review of Financial Studies*, 2, 157–188.

Du Rietz, Gunnar, Dan Johansson, and Mikael Stenkula, 2015, Swedish capital income taxation (1862–2013), in Magnus Henrekson and Mikael Stenkula eds. *Swedish*

Taxation: Developments since 1862 (Palgrave Macmillan), 123–178.

Epstein, Larry G. and Stanley E. Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso, 2017, Asset market participation and portfolio choice over the life-cycle, *Journal of Finance* 72, 705–750.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2019, Five facts about beliefs and portfolios, unpublished paper.

Gomes, Francisco, and Alexander Michaelides, 2003, Portfolio choice with internal habit formation: A life-cycle model with uninsurable labor income risk, *Review of Economic Dynamics* 6, 729–766.

Gomes, Francisco, and Alexander Michaelides, 2005, Optimal life-cycle asset allocation: Understanding the empirical evidence, *Journal of Finance* 60, 869–904.

Gomes, Francisco, and Alexander Michaelides, 2008, Asset pricing with limited risk sharing and heterogeneous agents, *Review of Financial Studies* 21, 415–448.

Gomes, Francisco, Laurence Kotlikoff, and Luis Viceira, 2008, Optimal life-cycle investing with flexible labor supply: A welfare analysis of life-cycle funds, *American Economic Review* 98, 297-303.

Gouriéroux, Christian, Alain Monfort, and Eric Renault, 1993, Indirect inference, *Journal of Applied Econometrics* 8, S85–S118.

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese, 1996, Income risk, borrowing constraints and portfolio choice, *American Economic Review* 86, 158–171.

Guiso, Luigi and Monica Paiella, 2006, The role of risk aversion in predicting individual behavior, in Pierre-Andre Chiappori and Christian Gollier eds. *Insurance: Theoretical Analysis and Policy Implications* (MIT Press).

Guvenen, Fatih, 2009, A parsimonious macroeconomic model for asset pricing, *Econometrica* 77, 1711–1750.

Guvenen, Fatih, Serdar Ozkan, and Jae Song, 2014, The nature of countercyclical income risk, *Journal of Political Economy* 122, 621–660.

Haliassos, Michael, and Carol Bertaut, 1995, Why do so few hold stocks?, *Economic Journal* 105, 1110–1129.

Hall, Robert E., 1988, Intertemporal substitution in consumption, *Journal of Political*

Economy 96, 221–273.

Heaton, John, and Deborah Lucas, 2000, Portfolio choice and asset prices: The importance of entrepreneurial risk, *Journal of Finance* 55, 1163–1198.

Hwang, Jungbin, and Yixiao Sun, 2018, Should we go one step further? An accurate comparison of one-step and two-step procedures in a generalized method of moments framework, *Journal of Econometrics* 207, 381–405.

Kocherlakota, Narayana R., 1990, Disentangling the coefficient of relative risk aversion from the elasticity of intertemporal substitution: An irrelevance result, *Journal of Finance* 45, 175–190.

Low, Hamish, Costas Meghir, and Luigi Pistaferri, 2010, Wage risk and employment risk over the life cycle, *American Economic Review* 100, 1432–1467.

Meeuwis, Maarten, Jonathan A. Parker, Antoinette Schoar, and Duncan I. Simester, 2018, Belief disagreement and portfolio choice, NBER working paper 25108.

Mehra, Rajnish and Edward Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.

Smith, Anthony, 1993, Estimating nonlinear time series models using simulated vector autoregressions, *Journal of Applied Econometrics* 8, S63–S84.

Svensson, Lars E.O., 1989, Portfolio choice with non-expected utility in continuous time, *Economics Letters* 30, 313–317.

Viceira, Luis, 2001, Optimal portfolio choice for long-horizon investors with nontradable labor income, *Journal of Finance* 56, 433–470.

Vissing-Jørgensen, Annette, 2002 , Limited asset market participation and the elasticity of intertemporal substitution, *Journal of Political Economy* 100, 825–853.

Vissing-Jørgensen, Annette, 2003, Perspectives on behavioral finance: Does "irrationality" disappear with wealth? Evidence from expectations and actions, *NBER Macroeconomics Annual* 18, 139–194.

Wachter, Jessica, and Motohiro Yogo, 2010, Why do household portfolio shares rise in wealth?, *Review of Financial Studies* 2, 3929–3965.

Yogo, Motohiro, 2004, Estimating the elasticity of intertemporal substitution when instruments are weak, *Review of Economics and Statistics* 86, 797–810.