NBER WORKING PAPER SERIES

THE CROSS-SECTION OF HOUSEHOLD PREFERENCES

Laurent E. Calvet
John Y. Campbell
Francisco Gomes
Paolo Sodini

Working Paper 28788
http://www.nber.org/papers/w28788

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2021, Revised May 2025

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 2019FMNDWorkshop, the 2019NBERAsset Pricing Summer Institute, the 2021 London Seminar on Micro and Macro Implications of Household Behaviour and Financial Decision-Making, and the 2021 NBER Big Data Conference. We thank the Sloan Foundation for financial support to John Campbell, and Nikolay Antonov, Huseyin Aytug, Andreas Karlsson, Yuhuang Sun, Yapei Zhang, and especially Azar Aliyev for able and dedicated research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w28788

The Cross-Section of Household Preferences
Laurent E. Calvet, John Y. Campbell, Francisco Gomes, and Paolo Sodini
NBER Working Paper No. 28788
May 2021, Revised May 2025
JEL No. E21, G51

ABSTRACT

This paper estimates the cross-sectional distribution of Epstein-Zin preferences using the wealth and risky portfolio shares of a large panel of Swedish households. We find modestly heterogeneous risk aversion (a standard deviation of 0.97 with a median of 7.50) and a meaningfully heterogeneous and right-skewed time preference rate (standard deviation 7.31% with a median of 4.08%) and elasticity of intertemporal substitution (standard deviation 3.17 with a median of 0.70). Risk aversion and the EIS are only very weakly negatively correlated. We estimate lower risk aversion for households with riskier labor income, and a higher TPR and lower EIS for households who enter our sample with low wealth.

Laurent E. Calvet
EDHEC Business School
Department of Finance
laurentecalvet@gmail.com

John Y. Campbell
Harvard University
Department of Economics
and NBER
john_campbell@harvard.edu

Francisco Gomes
London Business School
Finance Department
fgomes@london.edu

Paolo Sodini
Stockholm School of Economics
Department of Finance
Paolo.Sodini@hhs.se

A data appendix is available at http://www.nber.org/data-appendix/w28788

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.

The canonical model of Epstein and Zin (1989) distinguishes three parameters that govern financial decisions: the time preference rate (TPR), the coefficient of relative risk aversion (RRA), and the elasticity of intertemporal substitution (EIS). We structurally estimate these 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. Our baseline implementation assumes that all agents have common beliefs about income processes and financial returns, but we also consider heterogeneity in beliefs about expected returns on risky assets.

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 group households by their initial wealth in relation to income and by their initial risky portfolio share. This process gives us a sample of 4264 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 to vary with both education and the household's sector of employment. These assumptions are common in classic life-cycle models (Carroll and Samwick 1997, Cocco, Gomes, and Maenhout 2005). Moreover, these 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 wealth-income ratio at the start of each year as well as realized group-level income shocks and risky asset returns during the year.

Our measure of wealth includes liquid financial wealth, real estate, defined-contribution retirement assets, and household entitlements to defined-benefit pension income. Our imputation of defined-contribution retirement wealth is an empirical contribution that extends previous research on Swedish administrative data. 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. To reduce the dimensionality of the model, we map both real estate and risky financial asset holdings into implied holdings of a single composite risky asset.

We address the challenge of identifying 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.

Identifying the EIS separately from the TPR requires time-variation in background risks or investment opportunities (Kocherlakota 1990, Svensson 1989). Our model is identified because it generates endogenous variation in household risk exposures. Households in the model have an age-specific target level of wealth that serves to buffer income shocks and finance retirement. [1] Households with a higher EIS adjust their savings behavior more aggressively: they save more if their wealth is below their target level and dissave more if it is above. Relatedly, households with high financial wealth relative to human capital invest more conservatively, which reduces the expected return on their financial wealth. 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 wealth accumulation profiles, although identification is weak in some cases which we handle by imposing a very small shrinkage penalty for EIS values distant from one. Our identification strategy is a methodological contribution of the paper.

We develop an indirect inference estimator of the preference parameters in the household population, which we define as follows. For each of the 4264 household groups, the indirect inference estimator is the vector of preference parameters under which the life-cycle model matches most closely the empirical time series of the group's wealth-income ratio and risky share, taking account of a shrinkage penalty. We obtain this optimum by conducting a grid search over 2760 combinations of the preference parameters.

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 in-

[1] The existence of a wealth target is a general property of life-cycle models with random income, as emphasized by Carroll and Toche (2009).

come 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 financial theory. As households age, they accumulate wealth and reduce their risky portfolio share. The risky portfolio share also declines with the wealth-income ratio conditional on 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 0.97 around a median of 7.74 and a mean of 7.50. Our other two preference parameters are highly dispersed and right-skewed. The median TPR is 4.08%, well below the mean value of 6.81%, and the standard deviation is 7.31%. The median EIS is 0.70, well below the mean value of 2.01, and the standard deviation is 3.17.

Fourth, our preference parameter estimates are only weakly cross-sectionally correlated. The correlation between risk aversion and the EIS is very weakly negative (-0.091), 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 weakly positively correlated with risk aversion (0.096) and negatively correlated with the EIS (-0.253), implying a tendency for impatient people to be both cautious and unwilling 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 cannot explain this pattern.

Fifth, we document notable 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. This pattern is consistent with the hypothesis that risk-tolerant households select risky occupations. In addition, 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. 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.

Sixth, when we allow for heterogeneity in beliefs about the expected return on the risky asset, we find that belief heterogeneity has little effect on the fit of our model and does not reduce the cross-sectional dispersion of estimated preference parameters. The dispersion in estimated risk aversion actually increases, because our model uses heterogeneous beliefs to fit savings behavior and adjusts risk aversion to avoid counterfactual implications for risky portfolio shares. Moreover, the fit of the model deteriorates drastically when we allow for heterogeneity in beliefs about the Sharpe ratio but restrict preferences to be homogeneous across households, which confirms the importance of the preference heterogeneity we estimate.

To the best of our knowledge, our paper is the first to estimate the Epstein-Zin preference parameters of a life-cycle model using micro data in a context where rates of return on financial assets and liabilities are constant over time and across households.[2] Since the estimation of recursive preferences in this context is both new to the literature and numerically intensive, we choose to focus on a widely used

[2]Many papers estimate Epstein-Zin preference parameters using time-variation in expected asset returns; see for example Vissing-Jørgensen and Attanasio (2003), Yogo (2004), and Chen, Favilukis, and Ludvigson (2013). Best, Cloyne, Ilzetski, and Kleven (2020) instead exploit variation in mortgage rates across borrowers with different loan-to-value ratios.

specification of a life-cycle model. In future work, our estimation approach could be readily extended to richer settings, involving for instance more complex labor income processes or preferences, at the cost of greater computational burden.

Our paper is related to a large literature on portfolio choice over the life cycle,[3] and a series of papers using the Swedish administrative data.[4] We contribute to the literature by reporting micro-level preference estimates of these models. In addition, our results provide useful inputs for models investigating the impact of heterogeneous agents on financial and macroeconomic outcomes (e.g., Guvenen 2009, Kaplan and Violante 2022).

A small and growing literature on heterogeneity in portfolio choice has recently tried to relate observed household behavior to underlying heterogeneity in preferences and beliefs (Giglio, Maggiori, Stroebel, and Utkus 2021, Meeuwis, Parker, Schoar, and Simester 2021). Relative to this literature, we observe more households over a longer period of time and have more complete data on wealth and portfolio allocation, although we lack data to measure potentially heterogeneous beliefs.

The organization of the paper is as follows. Section 1 explains how we measure household wealth and reports summary statistics. Section 2 presents the life-cycle model. Section 3 discusses the identification of the preference parameters and develops our estimation methodology. Section 4 reports empirical results. Section 5 concludes. An Internet Appendix provides additional results and details about our empirical analysis and estimation technique.

[3]See for instance Campbell and Viceira (2002), Cocco, Gomes, and Maenhout (2005), and Fagereng, Gottlieb, and Guiso (2017).

[4]See, for example, Calvet, Campbell, and Sodini (2007, 2009), Calvet and Sodini (2014), Betermier, Calvet, and Sodini (2017), and Bach, Calvet, and Sodini (2020).

1 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.

1.1 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

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

where $LW_{h,t}$ is liquid financial wealth, $DC_{h,t}$ is DC retirement wealth, $RE_{h,t}$ is real estate wealth, and $D_{h,t}$ is debt. In aggregate Swedish data in 1999, liquid financial wealth, DC retirement wealth, and real estate net of debt respectively account for 44.9%, 15.4%, and 39.7% of aggregate net wealth. Non-cash net wealth is

$$NCW_{h,t} = LW^S_{h,t} + DC^S_{h,t} + RE_{h,t} - D_{h,t}, \tag{2}$$

where $LW^S_{h,t}$ and $DC^S_{h,t}$ are the risky components of liquid financial wealth and DC wealth, respectively.

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 (*Kapitalförsäkring*), 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.

We impute DC retirement wealth by reconstructing the contribution rules of several types of Swedish DC pensions. We accumulate these contributions since 1991, with appropriate assumptions about asset allocation and the initial level of DC pension wealth in 1991.[5] We describe this procedure in detail in Sections I.B and II of the Internet 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.

Real estate consists of primary and secondary residences, rental, commercial and industrial properties, agricultural properties and forestry. As in Bach, Calvet,

[5]We can accurately impute DC contributions in Sweden because all companies are either part of a collective agreement or benchmarked against one, and employees cannot opt out of a DC pension scheme provided by their employer. For each employee, we compute DC contributions as the sum of (i) the mandatory contributions stipulated by the relevant collective agreement and (ii) additional private contributions, which we also observe. These extra pension contributions were relatively uncommon during our sample period due to a monthly cap of 1,000 SEK on the amount eligible for tax deferral.

and Sodini (2020), we value real estate properties using Statistics Sweden data.[6]

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.

As described here, the household balance sheet excludes durables 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). The balance sheet also excludes entitlements to DB pension income, which we treat as a form of post-retirement labor income rather than as a capital asset.

1.2 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. 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. This mapping proceeds in three stages.

At the first stage, we map all individual assets to equivalent holdings of diversified stocks, real estate, or cash. We treat liquid holdings of individual stocks, equity mutual funds, and hedge funds as diversified holdings of the MSCI world equity

[6]Real estate prices are compiled by Statistics Sweden from two main sources. Every 3 to 7 years, tax authorities assess the tax value of properties using detailed property characteristics and hedonic pricing. In addition, Statistics Sweden continuously collects data on every real estate transaction in the country, which permits the construction of sales-to-tax-value multipliers for different geographic locations and property types. The transaction data are also used to value apartments at the level of each residential building.

index.[7] We treat liquid holdings of balanced funds and bond funds as portfolios of cash and stocks, where cash pays the Swedish Treasury bill rate and where the share in stocks is given by the beta of each fund with the world index.[8] We assume that DC retirement wealth is invested in cash and the MSCI equity world index, as section I.B of the Internet Appendix explains.[9] We treat all real estate holdings as positions in a diversified index of Swedish residential real estate, the FASTPI index. Moreover, 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.

For each household h at time t, this mapping gives us the implied weights of liquid stocks, $\omega_{h,t}^{S}$, DC stocks, $\omega_{h,t}^{DCS}$, real estate, $\omega_{h,t}^{RE}$, and debt, $\omega_{h,t}^{D}$, in the household's non-cash net wealth. The excess return on *non-cash* net wealth is then:

$$R_{NCW,h,t+1}^{e} = \omega_{h,t}^{S} R_{S,t+1}^{e} + \omega_{h,t}^{DCS} R_{DCS,t+1}^{e} + \omega_{h,t}^{RE} R_{RE,t+1}^{e} - \omega_{h,t}^{D} R_{D,t+1}^{e}. \tag{3}$$

where $R_{S,t+1}^{e}$, $R_{DCS,t+1}^{e}$, and $R_{RE,t+1}^{e}$ denote the excess return over cash on risky liquid wealth, risky DC wealth, and real estate, respectively, and $R_{D,t+1}^{e}$ is the household borrowing rate over cash.

The second stage of our analysis is to calculate the conditional variance of $R_{NCW,h,t+1}^{e}$. Since the borrowing rate is conditionally deterministic, we only need to consider the vector $\omega_{h,t} = (\omega_{h,t}^{S}, \omega_{h,t}^{DCS}, \omega_{h,t}^{RE})'$ and the variance-covariance ma-

[7] 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.

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

[9] In Sweden, DC retirement wealth is highly diversified and invested either in variable annuity products (*traditionell försäkring*) or in pension funds chosen from a menu available on pension saving platforms provided by insurance companies (*fondförsäkring*).

trix Σ of $R^e_{t+1} = (R^e_{S,t+1}, R^e_{DCS,t+1}, R^e_{RE,t+1})'$. The variance of $R^e_{NCW,h,t+1}$ is then $\sigma^2(R^e_{NCW,h,t+1}) = \omega'_{h,t}\ \Sigma\ \omega_{h,t}$. To estimate Σ, we calculate excess returns on risky assets over cash, assuming that cash earns the Swedish one-month risk-free rate net of taxes, that liquid equity earns the MSCI world index return net of a 30% long-term capital income tax rate (Du Rietz et al. 2015), that real estate earns the FASTPI index return net of 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.

In the third stage, we convert the volatility into a *risky share* held in a single composite risky asset. The composite asset, also called numeraire, is the aggregate portfolio of Swedish households, scaled to have the same volatility as the after-tax MSCI world index: $R^e_{N,t+1} = (1 + L)(\omega'_{agg,t} R^e_{agg,t+1})$. Here $R^e_{N,t+1}$ is the return on the numeraire and $\omega_{agg,t}$ is the vector containing the weights of equity, real estate and risky DC wealth in the aggregate non-cash 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.

Total net wealth earns the excess return $R^e_{h,t+1} = (NCW_{h,t}/W_{h,t}) R^e_{NCW,h,t+1}$. The empirical risky share $\alpha_{h,t}$ is the ratio of the standard deviation of household h's overall portfolio to the standard deviation of the numeraire asset:

$$\alpha_{h,t} = \frac{\sigma(R^e_{h,t+1})}{\sigma(R^e_{N,t+1})} = \left(\frac{NCW_{h,t}}{W_{h,t}}\right) \frac{\sigma(R^e_{NCW,h,t+1})}{\sigma(R^e_{N,t+1})}. \tag{4}$$

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 is invested solely in the MSCI world stock index outside a retirement account.

1.3 Composite Households

We consider Swedish households that are aged between 40 and 60 during the 1999 to 2007 period and hold risky financial assets outside retirement accounts. This corresponds to 5.4 million household-year observations on the 13 cohorts born between 1947 and 1959, but we impose several filters. We exclude households in which the head is a student, working in the agricultural sector, retired before 1999, missing information on education or sector of employment, or missing data in any year. 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. Since our measurement procedures may be less adequate for the wealthiest, we also exclude households whose financial wealth is above the 99th percentile of the wealth distribution in 1999. These filters exclude 2.7 million observations, leaving us with a balanced panel containing 2.7 million household-year observations and 298,646 households.

We classify households by three levels of educational attainment: (i) basic or missing, (ii) high school, and (iii) post-high school. 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. We obtain a 3×3 grid of 9 large education/sector categories where the sectors of employment are aggregated by income volatility.

Because our goal is to measure preference heterogeneity, we further subdivide each of these categories using a two-way sort by deciles of the initial wealth-income ratio and initial risky share. The wealth-income ratio reflects saving that has already taken place before the start of our sample, so it and the risky share should reflect any preference heterogeneity that exists in the Swedish population. To capture the tails of the cross-sectional distribution of these variables, we work with the lowest two and highest two deciles and aggregate the remaining six deciles into three quintiles, giving us a 7×7 grid of 49 bins for the initial wealth-income ratio and risky share.[10] Finally, we again subdivide by 13 cohorts to create $5733 = 9 \times 49 \times 13$ groups. After excluding groups with less than 10 members or a wealth-income ratio higher than 25 in each year from 1999 to 2007, our final sample is a balanced panel of 4264 groups.

The median group size across years is 54 households, but the average group size is larger at about 70 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.

[10]The wealth-income and risky share breakpoints are set separately in each of the 9 education/sector categories. This ensures that across categories we have the same proportion of households in each of the 7 risky share and wealth-income brackets. However, the number of households can differ across the 49 bins defined by the two-way sort.

Table 1: Wealth-Income Ratio and Risky Share by Education and Income Risk

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	3.70	4.13	5.08	4.47	0.692	0.686	0.675	0.682
Medium	4.47	4.50	4.99	4.69	0.669	0.672	0.664	0.669
High	4.78	5.10	6.07	5.46	0.662	0.675	0.668	0.670
All	4.23	4.51	5.27	4.78	0.677	0.678	0.669	0.674

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	3.06	3.27	3.58	3.42	0.249	0.227	0.203	0.221
Medium	3.64	3.56	3.76	3.66	0.245	0.228	0.209	0.223
High	3.80	3.85	3.95	3.92	0.239	0.219	0.191	0.211
All	3.49	3.55	3.76	3.66	0.246	0.226	0.203	0.220

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 Internet Appendix Tables IA.3 and IA.4 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 are based on the 1999 to 2007 period and 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 Internet Appendix Table IA.2.

1.4 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, averaging across cohorts and the subdivisions by initial wealth-income ratio and risky share. 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 3.7 to 6.1, while average risky shares vary in a narrow range from 66% 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 shares vary little with education. Across sectors, income risk has a strong positive effect on the wealth-income ratio and a weak effect on the risky share.

The bottom panel of Table 1 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 3.0–4.0. Across all 4264 groups, the average wealth-income ratio has a mean of 4.8 with a standard deviation of 3.7, while the average risky share has a mean of 67% with a standard deviation of 22%.[11] Figure 1 plots the distribution of wealth-income ratios and risky shares across Swedish households.

The cross-sectional variation in wealth and asset allocation documented in Table 1 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.

[11]Our aggregation procedure into groups, while designed to preserve as much cross-sectional dispersion as possible, does slightly compress the underlying distribution of wealth-income ratios and risky shares at the household level. Table IA.1 in the Internet Appendix reports that the standard deviation of the initial wealth-income ratio at the household level is 4.3 (reflecting some extremely high ratios driven by temporarily low household income), and the standard deviation of the initial risky share is 31%.

Figure 1: Distribution of Wealth-Income Ratio and Risky Share Across Swedish Households

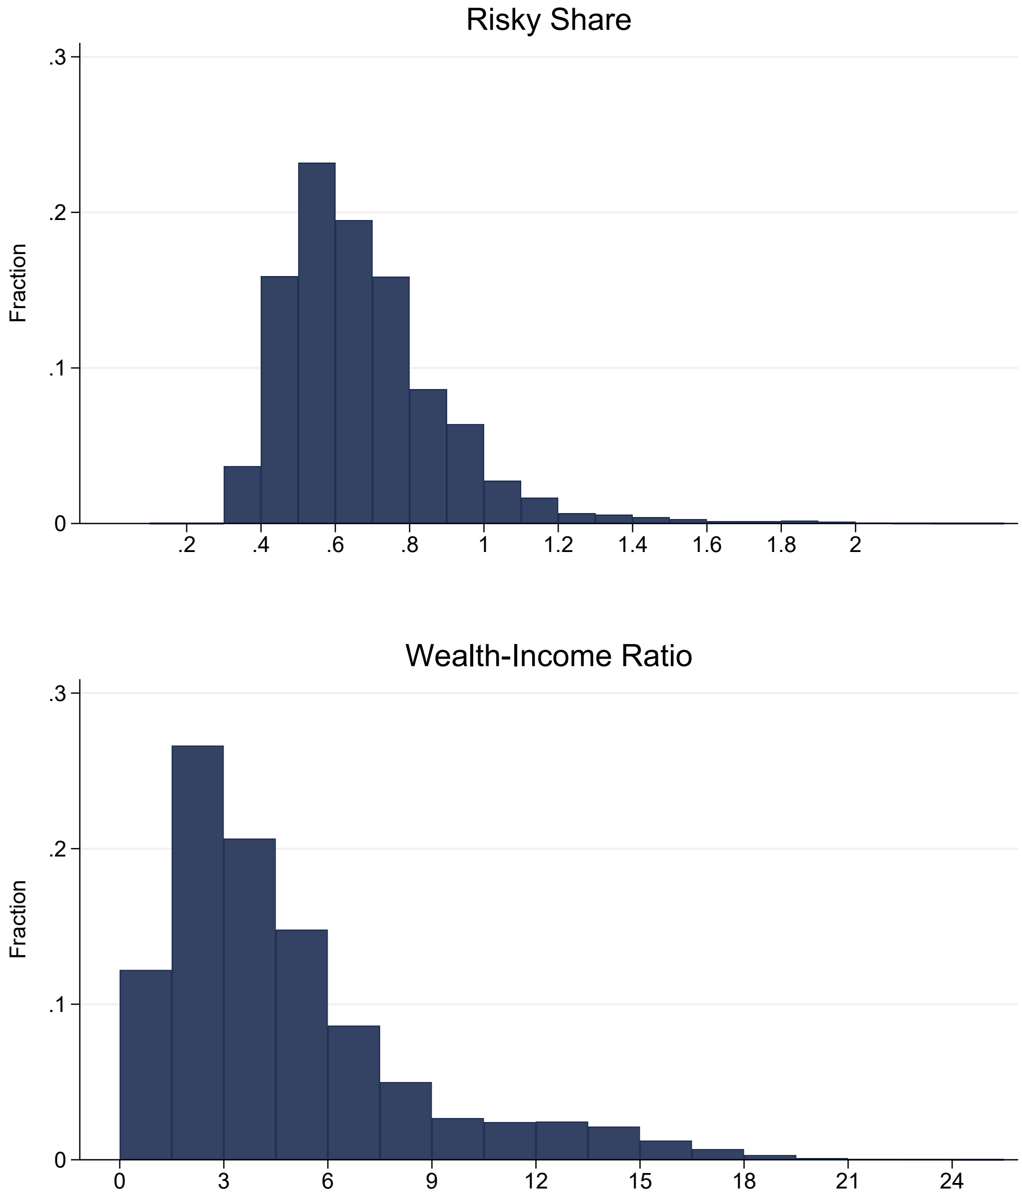

This figure presents histograms for the wealth-income ratio (WY) and risky share (RS) across 4,264 groups of Swedish households, size-weighted to recover the underlying distribution across households under the assumption that WY and RS are homogeneous within groups. Each bin is labeled on the horizontal axis with the upper cutoff value of WY or RS at the right edge of the bin. The vertical axis shows the size-weighted fraction of the sample in each bin.

2 Income Process and Life-Cycle Model

2.1 Measuring Income Risk

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

$$\log(Y_{h,t}) = a_c + b' x_{h,t} + \nu_{h,t} + \varepsilon_{h,t}, \tag{5}$$

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. The distinction is important because we estimate preference parameters that best fit group-level wealth dynamics conditional on realized group-level income shocks and simulated idiosyncratic household-level income shocks.

We model the permanent component of income, $\nu_{h,t}$, as the sum of a group-level component, ξ_t, and an idiosyncratic component, $z_{h,t}$:

$$\nu_{h,t} = \xi_t + z_{h,t}. \tag{6}$$

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 $z_{h,t} = z_{h,t-1} + w_{h,t}$.

We assume by contrast that the transitory component of income, $\varepsilon_{h,t}$, is purely

Figure 2: Estimated Age-Income Profiles

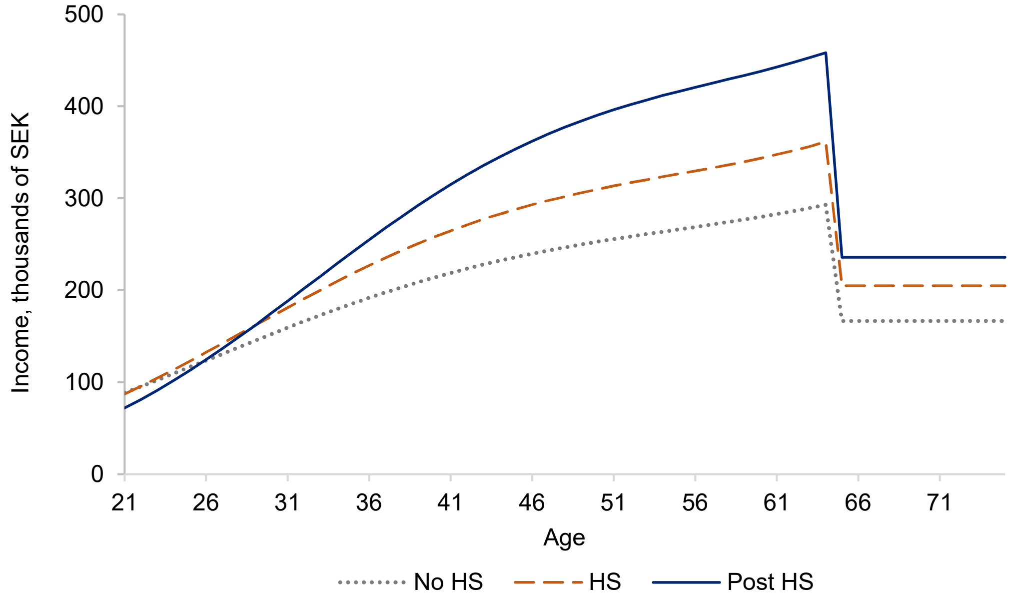

This figure presents estimated age-income profiles, including replacement ratios in retirement, for Swedish households with three levels of education: no high school (No HS), high school (HS), and post-high-school (Post HS). The estimates are based on a labor income process specified in equations (5)-(6).

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 (5) using household yearly income data, following a procedure described in Section I.C of the Internet Appendix. This gives us estimates of the age-income profile for each education group, which we plot in Figure 2. The profiles during working life are steeper than profiles estimated in the US.[12] The flat profiles after retirement reflect the entitlements of Swedish households to DB pension income.

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.[13] Second, we re-apply the procedure above to estimate income risk for the resulting nine education-business sector categories.

Internet Appendix Table IA.4 reports the standard deviations estimated for these nine categories. Permanent (systematic and idiosyncratic) income volatilities vary relatively little across sectors, but transitory idiosyncratic income volatilities are considerably higher for high-risk sectors. The table also shows that educated households, particularly those with higher education, face higher transitory income risk

[12]Dahlquist, Setty, and Vestman (2018) estimate income profiles for Sweden with a pronounced hump shape and lower income towards the end of working life. They use a model that excludes cohort effects, thereby estimating the age-income profile in part by comparing the incomes of households of different ages at a point in time. This procedure is biased if different cohorts receive different lifetime income on average. We obtain similar estimates when we exclude cohort effects from our model of income.

[13]Internet Appendix Table IA.2 reports the number of households and Table IA.3 reports the underlying sectors in each category.

and lower idiosyncratic 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 US. 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 educated workers face relatively high income losses when they become unemployed.[14]

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 2 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. 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 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

[14]This results from institutional features of the Swedish labor market which we explain in Section I.C of the Internet Appendix.

Table 2: Panel Regressions of Wealth-Income Ratio and Risky Share on Group Characteristics

	(1)	(2)	(3)
	WY	RS	RS
Age	0.156***	-0.014***	-0.010***
	(0.016)	(0.001)	(0.001)
Total income volatility	14.816***	-0.191	0.222*
	(1.847)	(0.109)	(0.090)
High school	0.503***	-0.011	0.003
	(0.136)	(0.009)	(0.007)
Post-high school	1.084***	-0.016	0.014*
	(0.137)	(0.008)	(0.007)
WY			-0.028***
			(0.001)
Constant	-6.219***	1.541***	1.368***
	(0.783)	(0.052)	(0.043)
Year fixed effects	Yes	Yes	Yes
R^2	0.103	0.190	0.383

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 38,376 observations on groups, corresponding to 2,687,814 observations on underlying households.

with risky income tend to have lower risk aversion, as Section 4.2 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).[15] We work with a richer lifecycle

[15]The negative effect of the wealth-income ratio on the risky share appears 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 in 1999; we measure the risky portfolio share taking account of housing and leverage through mortgage borrowing; and we

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.

The results of this section are obtained under the maintained assumption that real estate is a risky asset that earns the FASTPI index return net of a 22% real estate capital gains tax. In the Internet Appendix we consider how these findings are modified when real estate is treated as riskless, or as having the same risk as each group's financial portfolio so that the risky share can be measured from the financial portfolio alone. Tables IA.7 and IA.9 in the Internet Appendix report that the cross-sectional average risky share falls from 67% in the base case to 44% if we treat real estate as riskless, and rises to 81% if we consider financial wealth alone. The cross-sectional heterogeneity of the risky share is similar in all these cases, and regressions of the risky share on group characteristics, reported in Tables IA.8 and IA.10, are also similar. As we explain in Section 2.3, the correlation between risky assets and real estate is substantial. As a result, excluding real estate risk from the model reduces the correlation between risky assets and labor income, which would require a large increase in risk aversion to fit observed risktaking behavior; since we do not regard risk aversion estimates above 10 as plausible, we do not consider these alternative assumptions further.

2.2 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). We consider a relatively tractable framework because the estimation of heterogeneous preference parameters

---

predict the risky share using the wealth-income ratio rather than the absolute level of wealth.

in the household population is numerically costly.[16] Households have finite lives and Epstein-Zin utility over a single consumption good. The utility function V_t is specified by the RRA coefficient γ, the time discount factor δ or equivalently the TPR $-\log(\delta)$, and the EIS ψ. The utility V_t satisfies the recursion

$$V_t = \left[C_t^{1-1/\psi} + \delta \left(\mathbb{E}_t p_{t,t+1} V_{t+1}^{1-\gamma} \right)^{(1-1/\psi)/(1-\gamma)} \right]^{\frac{1}{1-1/\psi}}, \tag{7}$$

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. Preference parameters vary across households but we suppress the household index in (7) for 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. Similarly, matching the behavior of retirees is also hard for simple life-cycle models that do not incorporate health shocks or bequest motives.[17] For these reasons, we only consider the model's implications for ages 40 to 60.

We initialize the 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). Before retirement households supply labor inelastically. The stochastic process of labor income, $Y_{h,t}$, is described in Section 2.1. All households retire at age 65, as was typically the case in Sweden during our sample period, and retirement earnings are set to a constant replacement ratio of the last working-life permanent income. Consistent

[16] The estimation of the chosen specification proceeds as follows. For each of the 4264 household groups, we conduct a grid search over 2760 combinations of the preference parameters, which overall requires us to solve the life-cycle model more than 10 million times.

[17] Since we do not observe decisions late in life, we do not include an explicit bequest motive and instead capture the desire to leave a bequest as a lower TPR.

with Section 1, wealth in the model is invested every period in a one-period riskless asset (bond) and a composite risky asset.

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

$$W_{h,t+1} = (R_f + \alpha_{h,t} R^e_{N,t+1})(W_{h,t} + Y_{h,t} - C_{h,t}), \tag{8}$$

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 $\mathcal{N}(\mu_r, \sigma_r^2)$.

2.3 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 Figure 2, and the standard deviations σ_u, σ_w, and σ_ε in Table IA.4 in the Internet Appendix.

We assume that all safe borrowing and lending takes place at a single safe interest rate of 2.0%.[18] 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 1.2. We

[18]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. Our model abstracts from the wedge between these two rates. It would be feasible to assume that households pay a higher rate 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. We do not need to model the cost of unsecured borrowing against future labor income given our focus on middle-aged Swedish households who are saving for retirement.

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. This value is intermediate between the high risk premium on DC wealth and the low risk premium on real estate. Putting these assumptions together, we assume a Sharpe ratio of 0.26. In section 4.5 we discuss robustness of our results to assuming alternative Sharpe ratios.

Following Campbell, Cocco, Gomes, and Maenhout (2001), we estimate the correlation between the numeraire risky asset return and group-level systematic income shocks by lagging the stock return one year to capture a delayed response of income to macroeconomic shocks that move asset prices immediately. Empirically the correlation has an average value across the nine education-sector categories of 0.08 for stock returns, 0.37 for real estate returns, and 0.26 for the composite risky asset.[19] Table IA.5 in the Internet Appendix reports the separate correlations for each of the nine categories that we use in our model.

3 Identification and Estimation

3.1 Identification Strategy

3.1.1 Intuition

Our goal is to estimate the three preference parameters of the Epstein-Zin utility model. The main challenge is that the TPR and the EIS are not separately identified if consumption growth and the portfolio return are independent and identically

[19]The correlation between the numeraire risky asset return and individual income growth is much smaller because most individual income risk is idiosyncratic.

distributed, as Kocherlakota (1990) and Svensson (1989) explain. To identify these parameters we therefore need to exploit time-variation in consumption growth and the portfolio return arising from the life-cycle structure of our model.

To explain the intuition underlying our identification strategy, we consider an Epstein-Zin investor who can trade a riskless asset and a risky asset every period. The Euler equation for the return on the optimal portfolio is given by

$$1 = \mathrm{E}_t \left[\widetilde{\delta}_{t+1} \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{1}{\psi}} \left(\frac{V_{t+1}}{\mu(V_{t+1})} \right)^{\frac{1}{\psi}-\gamma} R^P_{t+1} \right], \tag{9}$$

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}.[20] Under the usual assumption of conditional joint lognormality, we obtain

$$\begin{aligned} \mathrm{E}_t g_{t+1} \quad = \quad & \psi \left[\mathrm{E}_t r^P_{t+1} + \log(\widetilde{\delta}_{t+1}) \right] + (1-\gamma\psi)\, \mathrm{E}_t \widetilde{v}_{t+1} + \frac{1}{2\psi} \sigma^2_{g,t} - \sigma_{gr,t} \\ & + \frac{\psi}{2} \left[\left(\frac{1}{\psi} - \gamma \right)^2 \sigma^2_{\widetilde{v},t} + \sigma^2_{r,t} \right] + (1-\gamma\psi)\, \sigma_{\widetilde{v}r,t} - \left(\frac{1}{\psi} - \gamma \right) \sigma_{g\widetilde{v},t}, \end{aligned} \tag{10}$$

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 (10) highlights the identification problem. If the expected portfolio return, the time discount factor, and the conditional second moments are all constant, then for any value of ψ there is a corresponding time discount factor δ that delivers the same level of $\mathrm{E}_t g_{t+1}$ (which is also constant). Without additional restrictions on δ or ψ these two parameters cannot be separately identified.

[20]This Euler equation holds with equality even though our model has borrowing constraints, because with labor income risk and a 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.

Equation (10) also shows three effects that can deliver identification. First, the variance terms are time varying in life-cycle models with undiversifiable risky labor income such as ours. However such variation tends to be more substantial early in life, when households have less wealth to smooth shocks (Gomes and Michaelides 2005). Second, the expected portfolio return varies over time because the risky share in the agent's financial portfolio is a function of age and accumulated financial wealth. Third, the effective time discount factor $\widetilde{\delta}_{t+1} = \delta p_{t,t+1}$ varies with age, driven by the age-dependent survival probabilities $p_{t,t+1}$.

All three types of time variation imply that the profile of the wealth-income ratio is affected differently by the TPR and the EIS at different ages. Our identification strategy builds on this intuition. One important caveat, visible from equation (10), is that identification will be poor if we have $\mathrm{E}_t r^P_{t+1} + \log(\widetilde{\delta}_{t+1})$, the difference between the expected portfolio return and the TPR $-\log(\widetilde{\delta}_{t+1})$, close to zero. In that case the first term, which is multiplied by the EIS parameter, drops out, and the EIS only enters the equation through the variance terms. Therefore, for groups for which the true preference parameters imply $\mathrm{E}_t r^P_{t+1}$ very close to the TPR, we are likely to encounter an identification problem.

3.1.2 Regressions on Simulated Data

We illustrate the promise of our identification strategy by running a series of regressions based on simulated data from the model. More specifically we regress the underlying preference parameters that were used to generate those simulations against moments from the simulated data. We consider the same 16 moments that we use in estimation: the average risky share and the average wealth-income ratio in

Table 3: Regressions of Preference Parameters on Simulated Moments

Panel A. Unconditional Regressions.

Parameter	RRA	Disc. Factor	EIS
Adjusted R^2	0.909	0.619	0.024

Panel B. EIS Regressions controlling for initial W/Y.

WY range	≤ 1	$(1,2]$	$(2,3]$	$(3,5]$	$(5,7]$	$(7,10]$	> 10
R^2	0.604	0.640	0.634	0.534	0.463	0.343	0.172

This table reports the adjusted R^2 statistics of regressions in simulated data using preference parameters on a grid containing 12 values of relative risk aversion (RRA) ranging from 3 to 12, 11 values of the time preference rate (TPR) ranging from -0.05 to 0.22, and 14 values of the elasticity of intertemporal substitution (EIS) ranging from 0.1 to 2.5. For each of the 1,848 combinations of preference parameters we consider all initial levels of the wealth-income ratio (WY) observed among Swedish household groups. We regress the preference parameters on 16 simulated moments: 8 values of RS and 8 values of W/Y. The dependent variables in the Panel A regressions are RRA (column 2), the discount factor (column 3), and the EIS (column 4). Panel D report results of separate EIS regressions for simulated groups within different ranges of the initial W/Y (indicated in the columns).

every year ($\{\alpha_{it}\}_{t=1,8}$ and $\{(W/Y)_{it}\}_{t=1,8}$).[21] The values of the preference parameters and initial wealth-income ratios considered in these simulations are very similar to those we use in estimation.

The results in Panel A of Table 3 show that risk aversion is extremely well identified from these moments, with an R^2 of 90.9%. This is an extremely high value since we are estimating a linear regression and imposing the same coefficients across groups, even though we know that the true relationship is non-linear and also depends on the initial wealth-income ratio. For the discount factor regression we also obtain a high R^2 of 61.9%. For the EIS, however, the fit is very poor with an R^2 of only 2.4%.

In Panel B, we show how the identification of the EIS can be improved by partially relaxing the linearity assumption and estimating regressions within seven different

[21] Simulated moments are obtained by averaging 10,000 simulations. In this exercise, unlike our empirical analysis, we use the observed wealth-income ratio only in the first year, and take wealth-income ratios in subsequent years from the simulated data rather than from the observed data.

ranges of values of the initial wealth-income ratio. The adjusted R^2 statistics now vary between 17.2% and 64.0%, with higher values corresponding to lower initial wealth-income ratios. Hence, although the results are not as strong as for the other preference parameters given the identification problem that arises when $\mathrm{E}_t r_{t+1}^P$ is close to the TPR, the EIS can be identified within our framework – particularly for households with less wealth relative to income – by conditioning our estimates on the wealth-income ratio.

3.2 Structural Estimation

We estimate the vector of preference parameters, $\theta^g = (\delta^g, \gamma^g, \psi^g)'$, in each group g using indirect inference (Smith 1993, Gouriéroux, Monfort, and Renault 1993). This method compares a vector of auxiliary statistics produced by the lifecycle model to the corresponding vector of empirical auxiliary statistics in the group. To address the weak identification of the EIS in some regions of the parameter space, we employ a shrinkage version of indirect inference to improve accuracy in finite samples, as in Blasques and Duplinskiy (2018).

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 its total income:

$$\hat{\mu}_{1,t}^g = \frac{\sum_{h=1}^{N^g} W_{h,t}}{\sum_{h=1}^{N^g} Y_{h,t}}, \tag{11}$$

and (ii) the group's risky share:

$$\hat{\mu}^g_{2,t} = \frac{\sum_{h=1}^{N^g} \alpha_{h,t} W_{h,t}}{\sum_{h=1}^{N^g} W_{h,t}}. \tag{12}$$

These statistics provide reliable measures of wealth accumulation and risk-taking based on group aggregates. We note that $\hat{\mu}^g_{1,t}$ and $\hat{\mu}^g_{2,t}$ are ratios of sample moments but are not sample moments themselves, which motivates the use of indirect inference rather than moment-based estimators. We stack the 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.[22]

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 earn the riskfree rate R_f and the synthetic excess risky return $R^e_{N,t}$ on their risky asset holdings.

We estimate the binding function $\mu^g(\theta)$ by simulation of the life-cycle model as follows. For each preference parameter θ, we compute the wealth-income ratio and risky share predicted by the model for the years 2000 to 2007 using the parameters from Table IA.4 in the Internet Appendix as inputs. For each year t, we simulate $S = 10{,}000$ paths in the group, taking as a starting point an information set $\mathcal{I}_t$ containing the empirical wealth-income ratio of group g at the end of year $t-1$.

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

In each simulation, we feed the realized return on the risky asset and the realized empirical group-level income shocks, during the year. Consistent with the life-cycle model, we assume that households have this much advance information about wages and hours. We simulate the idiosyncratic permanent and transitory income shocks of each household, which we combine to $\mathcal{I}_t$ to obtain the group's wealth-income ratio and risky share at the end of year t. We denote the averages of the S simulated values as $\widetilde{\mu}^g_{1,t,S}(\theta)$ and $\widetilde{\mu}^g_{2,t,S}(\theta)$, respectively, and stack them into a column vector $\tilde{\mu}^g_S(\theta)$. Section III.B of the Internet Appendix explains the simulation procedure in detail.

We estimate the vector of preference parameters, θ^g, by the shrinkage indirect inference estimator,

$$\hat{\theta}^g = \arg\min_{\theta} Q^g(\theta), \tag{13}$$

where the criterion function is given by

$$Q^g(\theta) = \left[\tilde{\mu}^g_S(\theta) - \hat{\mu}^g\right]' \Omega \left[\tilde{\mu}^g_S(\theta) - \hat{\mu}^g\right] + \frac{\lambda}{N} (\psi - 1)^2. \tag{14}$$

In this definition, Ω is a positive-definite weighting matrix, $\lambda \geq 0$ is a nonnegative constant, and $N = N^1 + \cdots + N^G$ denotes the total number of households in the panel.

The criterion $Q^g(\theta)$ in equation (14) consists of two terms. The first term, $\left[\tilde{\mu}^g_S(\theta) - \hat{\mu}^g\right]' \Omega \left[\tilde{\mu}^g_S(\theta) - \hat{\mu}^g\right]$, quantifies the distance between the auxiliary statistics produced by the lifecycle model and their empirical counterparts. The second term, $\lambda\,(\psi - 1)^2/N$, is a Tikhonov penalty aimed at addressing the weak identification of the EIS. It is flat near $\psi = 1$, reflecting the ongoing debate about whether the EIS is above or below unity.

We use a diagonal weighting matrix Ω and a parameter λ common to all groups, which ensures that our measure of fit in equation (14) is consistently scaled across groups. More specifically, each diagonal element of Ω is a scale factor that converts the wealth-income ratios and risky shares into comparable units. We define

$$\Omega = \text{diag}(\omega_1, ..., \omega_1, \omega_2, ..., \omega_2),$$

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 ensure comparability between wealth-income ratios and risky shares. In particular, their ratio, $(\omega_2/\omega_1)^{1/2} = 7.78$, aligns with a grand average risky share of 0.65 and an average wealth-income ratio of 5.06 over the 2000-2007 period. Moreover, we set the regularization parameter λ/N to 10^{-4} and verify in Sections V.A and V.B of the Internet Appendix that our baseline results are robust to alternative choices of λ.

When the number of households goes to infinity, the shrinkage indirect inference estimator, $\hat{\theta}^g$, is consistent and asymptotically normal. Consistency follows from the fact that the penalty in equation (14) vanishes at rate N^{-1}, as explained by Blasques and Duplinskiy (2018). We refer the reader to Section III of the Internet Appendix for the closed-form expression of the variance-covariance matrix of our estimator and additional details on its properties.

4 Empirical Results

4.1 The Cross-Sectional Distribution of Preference Estimates

Tables 4 and 5 and Figure 3 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. 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 median RRA of 7.50, close to the mean of 7.74 and in the range considered as reasonable by Mehra and Prescott (1985). The cross-sectional standard deviation of estimated RRA is modest at 0.97, less than 15% of the mean and median estimates.

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 can be borrowed against 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 2. 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

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

	Mean	Median	Std. Dev.	10%	25%	75%	90%
RRA	7.74	7.50	0.97	6.50	7.00	8.00	9.00
TPR (%)	6.81	4.08	7.31	1.01	2.02	6.19	22.31
EIS	2.01	0.70	3.17	0.02	0.20	2.50	5.00
Average RS	0.65	0.63	0.17	0.45	0.53	0.75	0.90
Initial WY	4.28	3.08	3.89	0.83	1.63	5.25	9.18
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 4,264 groups containing 298,646 households.

risk exposure.

The other two preference parameters have much greater cross-sectional variation relative to their means, and both are strongly right-skewed. The median TPR is 4.08%, considerably lower than the mean of 6.81%, and the cross-sectional standard deviation of the TPR is 7.31%. Similarly, the median EIS is 0.70, considerably lower than the mean of 2.01, and the cross-sectional standard deviation of the EIS is 3.17. This cross-sectional standard deviation is over 12 times as large for the EIS as for RRA 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.[23]

Table 5 shows that preference parameter estimates are only weakly cross-sectionally correlated. RRA and the EIS have a weak negative correlation of −0.09, a finding that contrasts with the perfect negative correlation between the logs of RRA and the EIS under power utility. The TPR is positively correlated with RRA

[23]Figure IA.1 in the Internet Appendix plots the univariate distributions of all three preference parameters.

Figure 3: Joint Distribution of TPR and EIS

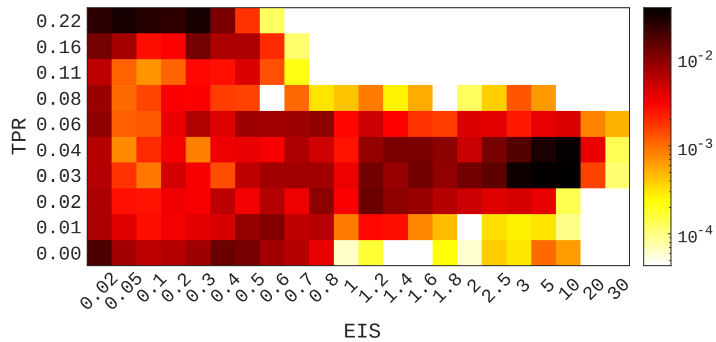

This figure presents bivariate heat map for estimates of TPR and EIS across 4,264 groups of Swedish households, size-weighted to recover the underlying distribution across households under the assumption that preferences are homogeneous within groups. Each axis label indicates the exact value estimates. The logarithmic color scheme indicates the fraction of the sample in each bin. This fraction is 3.7% for the darkest color and 0.0% for the brightest color.

and negatively correlated with the EIS, but the correlations are modest at 0.10 and −0.25 respectively. These 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 3 is a heat map of the bivariate distribution of the TPR and EIS. The negative correlation between the two parameters is visible in the figure, with high TPR values more common when the EIS is estimated to be low, but the main pattern illustrated in the figure is that the TPR estimates are much more dispersed when the EIS is low. This makes sense since a low EIS strengthens the desire to smooth consumption over time and thus reduces the impact of the TPR on observable savings decisions, as equation (10) shows.

Section V.B of the Internet Appendix conducts a Monte Carlo analysis of our procedure. A key lesson is that small-sample bias cannot explain the substantial

Table 5: 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	0.096***	1.000				
EIS	-0.091***	-0.253***	1.000			
Average RS	-0.103***	0.382***	-0.048***	1.000		
Initial WY	-0.548***	-0.377***	0.394***	-0.502***	1.000	
Growth of WY	0.395***	0.358***	-0.082***	0.603***	-0.716***	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 4,264 groups containing 298,646 households.

cross-sectional heterogeneity in our preference estimates. There is almost no bias for the RRA, minimal bias for the TPR, and some bias for the EIS, but correcting this bias has little effect on cross-sectional parameter dispersion.

4.2 Preference Estimates and Household Characteristics

The lower portion of Table 5 explores correlation patterns among preference parameters and observables. The initial wealth-income ratio has a correlation of -0.50 with the average risky share and a correlation of -0.72 with the cumulative growth rate of the wealth-income ratio over our sample period. 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. Correspondingly, the average risky share and the cumulative growth rate of the wealth-income ratio have a positive correlation of 0.60.

Our estimate of RRA is weakly negatively correlated (-0.10) with the average risky share, an intuitive result that is consistent with our identification analysis. RRA is more strongly negatively correlated (-0.55) 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 (-0.38) with the initial wealth-income ratio and positively correlated (0.36) with the cumulative 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 (0.38) 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 (0.39) with the initial wealth-income ratio and weakly negatively correlated (-0.08) 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.[24]

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

[24]Table IA.6 in the Internet Appendix reports multiple regressions rather than univariate correlations. Most patterns are similar, but controlling for the initial wealth-income ratio, the growth of wealth-income predicts the EIS positively rather than negatively.

Table 6: Education, Income Risk and Preferences

	(1) RRA	(2) TPR	(3) EIS
Total income volatility	-8.223***	-0.148***	3.292
	(0.426)	(0.043)	(1.809)
High school	1.208***	0.034***	0.111
	(0.033)	(0.003)	(0.136)
Post-high school	0.562***	0.025***	0.081
	(0.030)	(0.003)	(0.117)
Constant	8.735***	0.102***	1.449***
	(0.090)	(0.010)	(0.389)
Cohort dummies	Yes	Yes	Yes
R^2	0.265	0.061	0.002

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 4,264 groups containing 298,646 households.

observables. Households with riskier labor income tend to have lower risk aversion. Mechanically, this results from the fact documented in Table 2 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. Controlling for income risk, more educated people tend to have slightly higher RRA. The R^2 of the regression for risk aversion is 27%. Households with high income risk also tend to have a lower TPR and higher EIS, but the explanatory power of the TPR regression is only 6% and that of the EIS regression is less than 1%. We do not find that educated households are more patient; in fact, they tend to have slightly higher TPR controlling for their income risk.

4.3 Parameter Uncertainty

The discussion in the previous subsections treats our point estimates of parameters as if they are equivalent to the parameters themselves. In Sections III.C and V.C of the Internet Appendix, we develop asymptotic standard errors of the preference parameters that take parameter uncertainty into account. Reassuringly, the cross-sectional standard deviations of the RRA, TPR, and EIS, which we have reported in Section 4.1, are only weakly affected by estimation error, as Table IA.24 of the Internet Appendix shows.

In Table 7 we report hypothesis tests based on our asymptotic standard errors (developed in Sections III.C and III.D of the Internet Appendix), using a 5% significance level. We begin by testing hypotheses about the heterogeneity of preferences. We report that 86% of households are in groups for which we can reject the null that the group RRA equals the cross-sectional mean RRA, taking account of statistical uncertainty about that mean. Similarly, we reject equality to the mean TPR for 62% of households, and equality to the mean EIS for 50% of households. We can reject the null that all three parameters equal their cross-sectional means for 99% of households. We also test whether group preference estimates equal the cross-sectional medians, treating the medians as known for simplicity, and reject this hypothesis for risk aversion in 73% of cases, for the TPR in 32% of cases, and for the EIS in 34% of cases. Overall, the table presents strong statistical evidence against homogeneity of preferences within our framework. This evidence is particularly strong for risk aversion, despite the lower heterogeneity of our point estimates, because this parameter is relatively precisely estimated.

Next we turn to hypotheses about the level of our preference parameters. We reject the hypothesis that the TPR is zero for 59% of households. While 57% of

Table 7: Statistical Test Results for Estimated Preference Parameters

	% of Pop.
Reject RRA = mean(RRA)	85.7
Reject TPR = mean(TPR)	62.0
Reject EIS = mean(EIS)	50.3
Reject joint equality to mean	99.1
Reject RRA = median(RRA)	72.9
Reject TPR = median(TPR)	32.4
Reject EIS = median(EIS)	34.4
Reject TPR = 0	59.4
EIS < 1	57.0
Reject EIS > 1	36.7
Reject EIS < 1	4.7
Reject EIS = 1	38.1
EIS < 1/RRA	20.4
Reject EIS > 1/RRA	6.1
Reject EIS < 1/RRA	23.5
Reject EIS = 1/RRA	23.7

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 treat the cross-sectional median preference parameter as known rather than estimated. There are 4,264 groups containing 298,646 households.

households are in groups with point estimates of the EIS less than one, we can reject the null of an EIS greater than one for 37% of households, and can reject the null of an EIS less than one for only 5% of households.[25] Thus seven times as many Swedish households have an EIS significantly below one than have an EIS significantly above one.

Turning to power utility, only 20% of households have an estimated EIS that is lower than the reciprocal of RRA. We reject the null that EIS exceeds 1/RRA for only 6% of households, and reject the null that the EIS is lower than 1/RRA for 24%

[25]The asymmetry reflects the fact, illustrated in Figure IA.2 in the Internet Appendix, that the asymptotic standard error of the EIS is positively correlated with the level of the estimated EIS.

of households. A two-sided test rejects the power utility null for 24% of households. These findings reflect the fact that some households have relatively low EIS and high risk aversion, consistent with the reciprocal relationship implied by power utility; but about one quarter of households have EIS estimates that are significantly greater than the reciprocal of RRA.

4.4 Model Fit

As Figure IA.3 in the Internet Appendix shows, our model captures well the average variation of the risky share and wealth-income over the life-cycle, the usual target for life-cycle models. In this subsection we consider group-level measures of model fit. 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 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 median RMSE across all groups is 26.1% for the wealth-income ratio and 5.4% for the risky share. In other words, the median error in fitted wealth is just over 3 months of income and the median error in the risky share is about 5% of wealth. The RMSE distribution is somewhat right-skewed, particularly for the wealth-income ratio, as indicated by the fact that the mean RMSEs are above the medians at 35.3% and 5.6% respectively.

To interpret these numbers, we note that the RMSE of an atheoretical random walk model for WY has a median across groups of 30.4% and a mean of 33.0%. Thus

Table 8: Cross-Sectional Distributions of Model Fit Indicators

	Mean	Median	Std. Dev.	10%	25%	75%	90%
WY RMSE	35.29	26.08	30.48	12.55	17.96	41.07	69.19
RS RMSE	5.55	5.35	2.15	2.97	4.11	6.75	8.10
Scaled WY RMSE	6.97	5.15	6.02	2.48	3.55	8.11	13.67
Scaled RS RMSE	8.49	8.18	3.29	4.54	6.28	10.32	12.38
Scaled total RMSE	8.27	7.53	3.94	4.82	5.83	9.39	12.62

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, 5.06, 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. The scaled total RMSE is the square root of the objective function, without the penalty component, divided by 4 to create an RMSE, and multiplied by 100 to express it in annual percentage units. It is equivalent to the average of scaled WY and scaled RS RMSE. The cross-sectional means of WY and RS are computed over the 2000-2007 period. There are 4,264 groups containing 298,646 households.

our model has a better median performance and a slightly worse mean performance than a random walk. The standard deviation of the risky share around its group-specific time-series mean has a median across groups of 5.5% and a mean of 6.4%. Thus our model, which captures variation in the risky share with age and wealth accumulation, fits asset allocation better than an atheoretical model that treats the mean risky share for each group as a free parameter.

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 over the 2000-2007 period (5.06 and 0.65) to express them in proportional units. The median proportional RMSE is 5.2% for the wealth-income ratio and 8.2% for the risky share.

Finally, we report a transformation of the objective function, without the shrinkage penalty, that is rescaled to express it in RMSE-equivalent units. The objective function without the shrinkage penalty 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 scaled total RMSE of the model in percentage points.[26] The median scaled total RMSE is 7.5%, and the bottom row of Table 8 is similar to an average of the previous two rows, with a moderately right-skewed distribution.

Table IA.11 in the Internet Appendix shows how model fit deteriorates under homogeneous preferences. The median scaled total RMSE increases to 12.2% if we fix RRA at its cross-sectional mean. Fixing TPR at its cross-sectional mean produces a median total scaled RMSE of 8.2%, and similarly restricting the EIS delivers a median total scaled RMSE of 8.0%. Fixing all parameters at their cross-sectional means is disastrous in the sense that it increases the median total scaled RMSE to 21.9%. 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 behavior.

Table IA.12 conducts a similar exercise for a model that imposes the power utility restriction that the EIS equals the reciprocal of RRA. This restriction increases the median scaled total RMSE to 8.1%, roughly comparable to models with homogeneous TPR or EIS. The increase in RMSE is strongly right-skewed across households, reflecting the fact that power utility is a reasonable assumption for many Swedish households, but fits very poorly in some cases.

[26]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.

4.5 Heterogeneous Beliefs

We have shown that heterogeneous preferences are essential to fit household behavior if beliefs are homogeneous. It is natural to ask to what extent this finding is driven by our restriction on beliefs. In Tables IA.13–IA.15 in the Internet Appendix, we address this question by considering a simple form of heterogeneity in beliefs that allows three possible values of the Sharpe ratio: the base value of 0.26, a high value of 0.40, and a low value of 0.15. Then, for each group we pick the Sharpe ratio and preference parameters that minimize the objective function. The base case Sharpe ratio is selected for groups representing 54% of households, while the low Sharpe ratio and the high Sharpe ratio are each selected for 23% of households.

Allowing for heterogeneity in household beliefs has only a modest impact on the average preference parameters we estimate. The median preference parameters remain unchanged, and the mean preference parameters change only slightly relative to the homogeneous-beliefs case. However, the cross-sectional standard deviation of risk aversion is over twice as large at 2.27. 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.

Heterogeneous beliefs necessarily improve the fit of our model by adding free parameters, but the improvement is modest. Table IA.15 shows that the median total scaled RMSE declines only from 7.53% in the homogeneous-beliefs case to 6.95% in the heterogeneous-beliefs case. Importantly, the median total scaled RMSE is a disastrous 21.61% when we try to explain household behavior with heterogeneous beliefs alone, imposing homogeneous preferences.

Our results relate to Giglio et al.'s (2021) finding that portfolio choices respond less strongly to investors' self-reported beliefs than a simple Merton model would predict. A possible explanation is that optimistic households tend to also have high risk aversion. Our estimates display this positive correlation between beliefs and risk aversion; however, Table IA.15 shows that a restricted model that imposes homogeneous preferences for groups that share the same beliefs continues to fit very poorly. Overall, these estimates confirm our message that substantial preference heterogeneity is required to fit household financial decisions.

4.6 Preference Variation across Wealth Groups and Implications for Representative-Agent Modeling

In Section 4.1, we have reported size-weighted averages of preference parameters across groups, corresponding to equally weighted averages across households. While this is the natural weighting scheme in household finance applications, wealthier households have a greater influence on equilibrium asset prices and so asset pricing economists may be interested in wealth-weighted average preference parameters of households. In Table IA.16 in the Internet Appendix, we weight groups by their average wealth during the sample period rather than by their size. We find a similar median risk aversion of 7.50, a much lower median time preference rate of 3.05%, and a somewhat higher median EIS of 1.20. Wealth-weighting shifts the means of these parameters in a similar fashion, but has little effect on their cross-sectional standard deviations.

We also investigate how the cross-sectional dispersion of preference parameter estimates varies across wealth brackets. Table IA.17 in the Internet Appendix presents the standard deviation of preference parameters within wealth quintiles. Preference parameters, and particularly the TPR, are more dispersed among households with lower wealth. The standard deviation of the TPR declines from 7.7% in the lowest wealth quintile to 4.6% in the top quintile. Households with low wealth can have very different TPR and therefore very different willingness to save.

In Table IA.18 in the Internet Appendix, we verify that the wealth-weighted mean risk aversion, mean TPR and mean EIS are extremely stable over time. The time-series standard deviation is 0.03 for risk aversion, 0.11% for the TPR, and 0.05 for the EIS. Hence, despite shifts in the wealth distribution across households with heterogeneous preferences, a representative agent model with a stable set of Epstein-Zin preference parameters may be a reasonable approximation for the Swedish household sector as a whole.

5 Conclusion

In this paper, we have estimated a life-cycle model of consumption-portfolio choice on a panel of Swedish households. Our estimates of the RRA and EIS are only weakly negatively correlated across households, which contradicts the predictions of power utility. The TPR is weakly positively correlated with RRA and negatively correlated with the EIS. We estimate a negative correlation between income volatility and risk aversion. Households that enter our sample with low wealth, having saved relatively little earlier in life, tend to have a high TPR and a low EIS. We show that our results are unlikely to be driven by heterogeneous beliefs about the Sharpe ratio of the aggregate portfolio.

Our work sheds 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. 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 RRA (typically around 10) and an EIS around 1.5. We estimate a somewhat lower cross-sectional median RRA of 7.5, and a lower median EIS of 0.7. Weighting households by their wealth, as may be appropriate in an asset pricing model, leaves median RRA unchanged but increases the median EIS to 1.2. Wealth weights are sufficiently stable that these values might be a good approximation to the preferences of a representative agent even though underlying household preferences are heterogeneous.

In household finance, there is considerable interest in estimating risk aversion at the individual level and measuring its effects on financial decisions. This has sometimes been attempted using questions in surveys (Barsky, Juster, Kimball, and Shapiro 1997, Vissing-Jørgensen 2003). One difficulty with these attempts is that even if risk aversion is correctly measured, 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 financial risk-taking (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 (Campbell 2016, Jackson and Rothstein 2019). 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 can harm some households by eliminating financial products they prefer, as in the structured product example considered by Calvet, Celerier, Sodini and Vallee (2023).

Future research on household savings and portfolio choice can enrich our model in a number of interesting ways. We have studied middle-aged households who already have substantial liquid wealth and participate in the stock market; therefore we do not model hand-to-mouth behavior or nonparticipation in risky asset markets, two phenomena that can be explored using a broader sample of households. It would also be interesting to model real estate as a separate asset class from risky financial assets, to allow for flexibility in working hours, and to study the effects of household structure on financial decisions by comparing the decisions of single people, couples, and families with children. The interaction between private business ownership and financial risktaking is also potentially important, but difficult to measure given the very limited data available on the values of private businesses.

References

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.

Best, Michael Carlos, James S. Cloyne, Ethan Ilzetzki, and Henrik J. Kleven, 2020, Estimating the elasticity of intertemporal substitution using mortgage notches, *Review of Economic Studies* 87, 656–690.

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

Blasques, Francisco, and Artem Duplinskiy, Penalized indirect inference, *Journal of Econometrics* 205, 34–54.

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, 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, 2009, Fight or flight?

Portfolio rebalancing by individual investors, *Quarterly Journal of Economics* 124, 301–348.

Calvet, Laurent E., Claire Célérier, Paolo Sodini, and Boris Vallée, 2023, Can security design foster household risk-taking?, *Journal of Finance* 78, 1917–1966.

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 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.

Carroll, Christopher D., and Patrick Toche, 2009, A tractable model of buffer-stock saving, NBER Working Paper 15265.

Chen, Xiaohong, Jack Favilukis, and Sydney C. Ludvigson, 2013, An estimation of economic models with recursive preferences, *Quantitative Economics*, 4, 39–83.

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

Dahlquist, Magnus, Ofer Setty, and Roine Vestman, 2018, On the asset allocation of a default pension fund, *Journal of Finance* 73, 1893–1936.

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, 2021, Five facts about beliefs and portfolios, *American Economic Review* 111, 1481–1522

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

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.

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

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

Jackson, Howell, and Paul Rothstein, 2019, The analysis of benefits in consumer protection regulations, *Harvard Business Law Review* 9, 197–321.

Kaplan, Greg, and Giovanni Violante, 2022, The marginal propensity to consume in heterogeneous agent models, *Annual Reviews of Economics* 14, 747–775.

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, 2021, Belief disagreement and portfolio choice, *Journal of Finance*, forthcoming.

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.

Vissing-Jørgensen, Annette, 2003, Perspectives on behavioral finance: Does "irrationality" disappear with wealth? Evidence from expectations and actions, *NBER*

Macroeconomics Annual 18, 139–194.

Vissing-Jørgensen, Annette and Orazio P. Attanasio, 2003, "Stock-market participation, intertemporal substitution, and risk-aversion", *American Economic Review* 93, 383–391.

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.