The Cross-Section of Household Preferences

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

July 2019

*Calvet: Department of Finance, EDHEC Business School, 393 Promenade des Anglais, BP 3116, 06202 Nice Cedex 3, France, and CEPR; e-mail: laurent.calvet@edhec.edu. Campbell: Department of Economics, Harvard University, Littauer Center, Cambridge MA 02138, USA, and NBER; e-mail: john_campbell@harvard.edu. Gomes: London Business School, Regent's Park, London NW1 4SA, UK; e-mail: fgomes@london.edu. Sodini: Department of Finance, Stockholm School of Economics, Sveavagen 65, P.O. Box 6501, SE - 113 83 Stockholm, Sweden; e-mail: Paolo.Sodini@hhs.se. We acknowledge helpful comments on a preliminary version of this paper from Stijn van Nieuwerburgh, Stan Zin, and seminar participants at EDHEC Business School, ENSAE-CREST, the University of Michigan, Stanford University, and the 2017 American Economic Association meeting. We thank the Sloan Foundation for financial support to John Campbell, and Nikolay Antonov, Huseyin Aytug, and Yapei Zhang for able and dedicated research assistance.

Abstract

This paper estimates the cross-sectional distribution of preferences in a large administrative panel of Swedish households. We consider a life-cycle model of saving and portfolio choice with Epstein-Zin preferences which incorporates risky labor income, safe and risky financial assets inside and outside retirement accounts, and real estate. We study middle-aged households grouped by education, industry of employment, and birth cohort as well as by their accumulated wealth and risky portfolio shares. We find some heterogeneity in risk aversion (a standard deviation of 0.44 around a mean of 5.06) and considerable heterogeneity in the elasticity of intertemporal substitution (standard deviation 0.99 around a mean of 1.32) and the rate of time preference (standard deviation 4.13% around a mean of 1.57%, a low value which reflects in part the absence of a bequest motive in our model). Our estimated preference parameters are only weakly cross-correlated. We estimate higher time preference for households who enter our sample with low initial wealth, lower time preference for households with higher education, and lower risk aversion for households with riskier labor income.

1 Introduction

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

Modern financial theory distinguishes at least three parameters that govern savings behavior and financial decisions: the rate of time preference, the coefficient of relative risk aversion, and the elasticity of intertemporal substitution (EIS). The canonical model of Epstein and Zin (1989, 1991) makes all three parameters constant and invariant to wealth for a given household, while breaking the reciprocal relation between relative risk aversion and the elasticity of intertemporal substitution implied by the older power utility model.

We structurally estimate these three preference parameters in the cross-section of Swedish households by embedding Epstein-Zin preferences in a life-cycle model of optimal consumption and portfolio choice in the presence of uninsurable labor income risk and borrowing constraints. We assume that all agents have the same beliefs about income processes and financial returns; to the extent that any heterogeneity in beliefs exists, it will be attributed to heterogeneous preferences by our estimation procedure.

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. We first group households by their education level, the level of income risk in their sector of employment, and birth cohort. To capture heterogeneity in preferences that is unrelated to these characteristics we further divide households by their initial wealth accumulation in relation to income and by their initial risky portfolio share. This process gives us a sample of 4468 composite households that have data available in each year of our sample from 1999 to 2007.

We allow households' age-income profiles to vary with education, and the determinants of income risk (the variances of permanent and transitory income shocks) to vary with both education and the household's sector of employment. These assumptions are standard in the life-cycle literature (Carroll and Samwick 1997, Cocco, Gomes, and Maenhout 2005). It is well known that life-cycle models are much better at jointly matching portfolio allocations and wealth accumulation at mid-life than at younger ages or after retirement. Therefore we estimate the preference parameters by matching the age profiles of wealth and portfolio choice between ages 40 and 60, taking as given the initial level of wealth observed at the start of our sample period. Since our model is not designed to capture decisions late in life, we cannot accurately account for bequest motives and instead reflect the desire to leave a bequest as a lower rate of time preference.

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

Residential real estate is another important component of household wealth. To handle this, we include real estate in our empirical analysis but map both real estate and risky financial asset holdings into implied holdings of a single composite risky asset. While this is a stylization of reality, the inclusion of real estate wealth is consistent with common practice in life-cycle models (Hubbard, Skinner and Zeldes 1984, Castaneda, Diaz-Gimenez and Rios-Rull 2003, De Nardi 2004, Gomes and Michaelides 2005).

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

Our model assumes that expected returns on safe and risky assets are constant over time, so we cannot exploit time-variation in the riskless interest rate or the expected risky return to identify the EIS in the manner of Hall (1988). However, the model incorporates time-variation in background risks. As households approach retirement their human capital diminishes relative to their financial wealth, and this alters their desired portfolio composition and hence the rate of return on the portfolio. A secondary effect is that as households age their mortality rates increase, and this alters the incentive to save or equivalently the effective rate of time discounting. These changes alter wealth accumulation in a way that is mediated by the EIS. Accordingly we are able to identify the EIS from time-variation in the growth rate of wealth within each household group. This identification strategy that exploits accelerating or decelerating wealth accumulation is a methodological contribution of our paper.

Our main empirical findings are as follows. First, we estimate reasonable average levels of each Epstein-Zin preference parameter. Average risk aversion is 5.06, and the average EIS is 1.32. The average level of risk aversion is moderate in part because we treat real estate as a risky investment rather than ignoring it or treating it as a safe asset. The average EIS is somewhat above one and far above the reciprocal of average risk aversion, contrary to the restriction of the power utility model. The average rate of time preference is quite low at 1.57%, but this estimate may also be reasonable because we use this parameter to capture bequest motives as well as pure time preference.

Second, we estimate considerable heterogeneity in preference parameters across the Swedish population. The cross-sectional standard deviations are 0.44 for risk aversion, 0.99 for the EIS, and 4.13% for the rate of time preference. There is a debate in the asset pricing literature about whether the EIS is less than one, as estimated by Hall (1988), Yogo (2004) and others in time-series data, or greater than one, as assumed by Bansal and Yaron (2004) and a subsequent literature on long-run risk models. We find that the EIS is less than one for 49% of households, and even less than the reciprocal of risk aversion for 15% of households, while it is greater than one for 51% of households. This much cross-sectional variation suggests that aggregate results are likely to be sensitive to the way in which households are aggregated and are unlikely to be precise, consistent with large standard errors reported by Calvet and Czellar (2015) in a structural estimation exercise using aggregate data.

Third, the preference parameters are only weakly correlated with one another across households. We estimate that risk aversion has a slightly positive cross-sectional correlation of 0.07 with the EIS, contrary to the strong negative correlation implied by the power utility model in which one parameter is the reciprocal of the other. Risk aversion also has a weak positive correlation of 0.18 with the rate of time preference. The strongest cross-sectional correlation is a 0.36 correlation between our estimates of the EIS and the rate of time preference. These weak correlations across preference parameters imply that Swedish household behavior is heterogeneous in multiple dimensions, not just one. A single source of heterogeneity omitted from our model, such as heterogeneity in household beliefs about the equity premium, would not be able to generate this multidimensional heterogeneity in our preference estimates.

Fourth, our parameter estimates have intuitive relations with the moments we use for estimation. Risk aversion has a negative correlation of -0.22 with the average risky share. The rate of time preference is negatively correlated (-0.36) with the initial wealth-income ratio of each household group, and positively correlated (0.49) with the average growth rate of the wealth-income ratio. These patterns reflect the fact that households who enter our sample with a low wealth-income ratio accumulate wealth more rapidly, but not as much more rapidly as they would do if they were as patient as households with a higher initial wealth-income ratio. In other words, the symptom of a high rate of time preference in our data is a tendency to accumulate retirement savings later in life, catching up belatedly with those who saved more earlier in life. We estimate that households with higher education tend to have a lower rate of time preference, in part because they enter the sample with higher wealth-income ratios.

Fifth, we find that riskier labor income is associated with lower risk aversion across household groups. This pattern is consistent with the hypothesis that risk-tolerant households self-select into risky occupations, but could also result from households' failure to understand the investment implications of their income risk exposure.

The organization of the paper is as follows. Section 2 explains how we measure household wealth and its allocation to safe and risky assets, describes the creation of household groups, and reports summary statistics for the wealth-income ratio and the risky share across these groups. Section 3 presents the life-cycle model and household labor income processes. Section 4 discusses preference parameter identification and develops our estimation methodology. Section 5 reports empirical results on the cross-section of household preferences, and section 6 concludes. An online appendix provides additional details about our empirical analysis and estimation technique.

2 Measuring Household Wealth and Asset Allocation

Our empirical analysis is based on the Swedish Wealth and Income Registry, a high-quality administrative panel that has been used in earlier research.[1] The registry 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. The wealth data include bank account balances, holdings of financial assets, and real estate properties measured at the level of each security or property. The registry does not report durable goods, private businesses, or defined-contribution (DC) retirement wealth, but we augment the dataset by imputing DC contributions using income data and the administrative rules governing DC pensions in Sweden. We accumulate these contributions to estimate DC wealth at each point in time, and use a similar procedure to calculate entitlements to DB pension income. All our data series end in 2007.

[1] See, for instance, Bach, Calvet, and Sodini (2018), Betermier, Calvet, and Sodini (2017), Calvet, Campbell and Sodini (2007, 2009a, 2009b), and Calvet and Sodini (2014).

2.1 The Household Balance Sheet

We first aggregate the data to the household level. We define a household as a family living together with the same adults over time. The household head is the adult with the highest average non-financial disposable income; or, if the average income is the same, the oldest; or, if the other criteria fail, the man in the household.

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

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

where $LW_{h,t}$ is liquid financial wealth, $RE_{h,t}$ is real estate wealth, $DC_{h,t}$ is DC retirement wealth, and $D_{h,t}$ is debt. In aggregate Swedish data in 1999, the shares of these four components in total net wealth are 36%, 76%, 13%, and −25%, respectively.

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

Real estate consists of primary and secondary residences, rental, commercial and industrial properties, agricultural properties and forestry. The Wealth and Income Registry provides the holdings at the level of each asset. The pricing of real estate properties is based on market transactions and tax values adjusted by a multiplier, as in Bach, Calvet, and Sodini (2018).

Debt is the sum of all liabilities of the household, including mortgages and other personal liabilities held outside private businesses.[2] Since Swedish household debt is normally floating-rate, we treat debt as equivalent to a negative cash position but paying a borrowing rate that is higher than the safe lending rate.

The hardest balance sheet component to measure is DC retirement wealth. We impute this by reconstructing the details of the Swedish pension system, as we discuss in the next subsection. This detailed pension analysis also enables us to measure each household's entitlement to defined benefit (DB) pension payments in retirement.

[2] Because we do not observe durable goods (such as appliances, cars and boats), the value of household debt can exceed the value of the assets we oberve for some households. To avoid this problem, the debt variable $D_{h,t}$ is defined as the minimum of the total debt and real estate wealth reported in the registry. This approach is consistent with the fact that we proxy the borrowing rate by the average mortgage rate offered by Swedish institutions.

As described here, the household balance sheet excludes durable goods and private businesses, whose values are particularly difficult to measure. Private businesses are an important component of wealth for the wealthiest households in Sweden (Bach, Calvet, and Sodini 2018), but unimportant for most Swedish households. As a robustness check we can exclude self-employed households from the analysis, but our main sample includes them.

2.2 Pension Imputation

The Swedish pension system consists of three pillars: occupational pensions, state pensions, and private pensions. We discuss these in turn. Full details are provided in an online appendix.

Occupational pensions were introduced to Sweden in 1991. They are regulated for the vast majority of Swedish residents by four collective agreements that cover different occupational categories: blue-collar private-sector workers, white-collar private-sector workers, central government employees, and local government employees. These agreements specify workers' monthly pension contributions, the fraction directed to defined benefit (DB) and defined contribution (DC) pension plans, and the DC choices available to workers.

The collective agreements specify DC contributions as a percentage of pension qualifying income. These contributions are invested through insurance companies in either variable annuity products (called TradLiv in Sweden), or in portfolios of mutual funds, chosen by workers from a selection provided by the insurance company. There are also DB contributions which have been declining over time relative to DC contributions under the terms of the agreements, in a gradual transition from a DB to a DC pension system. We are able to impute both DC contributions and DB entitlements at the household level by following the rules of the collective agreements as detailed in the appendix. We can do this accurately because the DB collective pension payouts are a function of at most the last 7 years of pension qualifying income during working life, which can be observed in our income data.

The *state pension* system requires each worker in Sweden to contribute 18.5% of its pension qualifying income: 2.5% to a DC system and the remaining 16% to the pay-as-you-go DB system. DC contributions are invested in a default fund, that mirrors the world index during our sample period, unless the worker opts out and chooses a portfolio of at most 5 funds among those offered on the state DC platform (on average around 650 funds from 1999 to 2007). State DB payouts are function of the pension qualifying income earned during the entire working life. Since our individual income data begin in 1983, we cannot observe the full income history for older individuals in our dataset. To handle this, we back-cast their income back to the age of 25 by using real per-capita GDP growth and inflation before 1983, as explained in the appendix. We then use the state DB payout rules to impute state DB pension payments for each individual retiring during our sample period.

Defined contribution *private pensions* have existed in Sweden for a long time but our dataset provides us with individual private pension contributions from 1991. We assume that these contributions are invested in the same way as occupational and state DC contributions. We follow Bach, Calvet and Sodini (2018) and allocate 40% of the aggregate stock of private pension wealth in 1991 to retirees and 60% to workers.[3] Across workers, we allocate pension wealth proportionately to their savings in subsequent years, taking into account both age effects and individual savings propensities, as explained in the appendix.

To calculate DC retirement wealth at each point in time, we accumulate contributions from all three pillars of the Swedish pension system. To do this we assume that all contributions are invested in cash and the MSCI equity world index, without currency hedging, earning the index return less a 70 basis point fee which prevailed during our sample period. This assumption reflects the high degree of international diversification observed in Swedish equity investments (Calvet, Campbell, and Sodini 2007). The equity share in each household's DC retirement portfolio is rebalanced with age following the representative age pattern of life-cycle funds available in Sweden during our sample period.

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

2.3 Household Asset Allocation

Our objective is to match the rich dataset of household income and asset holdings to the predictions of a life-cycle model, which will allow us to estimate household preferences. To accomplish this, we need to map the complex data into a structure that can be related to a life-cycle model with one riskless and one risky asset. We do this in three stages. First we map all individual assets to equivalent holdings of diversified stocks, real estate, or cash. Second, we assume a variance-covariance matrix for the excess returns on stocks and real estate over cash that enables us to compute the volatility of each household portfolio. Third, we assume

[3] This breakdown is obtained from the condition that imputed pension wealth should be roughly the same just before and just after retirement. We find that this property holds when retirees are allocated 40% of private pension wealth.

that all household portfolios earn the same Sharpe ratio so that the volatility of the portfolio determines the expected return on the portfolio. Equivalently, we convert the volatility into a "risky share" held in a single composite risky asset. For ease of interpretation, we normalize that risky asset to have the same volatility as a world equity index.

At the first stage, we treat liquid holdings of individual stocks, equity mutual funds, and hedge funds as diversified holdings of the MSCI world equity index.[4] We treat liquid holdings of balanced funds and bond funds as portfolios of cash and stocks, with the share in stocks given by the beta of each fund with the world index.[5] We assume that unclassifiable positions in capital insurance, derivatives, and fixed income securities are invested in the same mix of cash and stocks as the rest of liquid financial wealth. We treat all real estate holdings as positions in a diversified index of Swedish residential real estate, the FASTPI index. We assume that DC retirement wealth is invested in cash and the MSCI equity world index, with weights calculated from the representative asset allocation of life-cycle funds available in Sweden during our sample period.

This mapping gives us implied portfolio weights in liquid stocks, real estate, and DC stocks for each household. Write these weights as $\omega^h_{S,t}$, $\omega^h_{RE,t}$, and $\omega^h_{DC,t}$ for household h at time t, and define the vector $\omega^h_t = (\omega^h_{S,t}, \omega^h_{RE,t}, \omega^h_{DC,t})'$. The second stage of our analysis is to calculate the variance of the excess return on the household portfolio as

$$\sigma^2(R^e_{h,t+1}) = \omega'_{h,t} \; \Sigma \; \omega_{h,t}, \tag{2}$$

where Σ is the variance-covariance matrix of the excess returns on liquid stocks, real estate, and stocks held in DC plans, $R^e_{t+1} = (R^e_{S,t+1}, R^e_{RE,t+1}, R^e_{DC,t+1})'$.

To estimate the elements of Σ, we assume that cash earns the Swedish one-month risk-free rate net of taxes, that liquid equity earns the MSCI world index return less a 30% long-term capital income tax rate (Du Rietz et al. 2015), that real estate earns the FASTPI index return less a 22% real estate capital gain tax rate, and that stocks held in DC plans earn the pre-tax MSCI world index return before the adjustment of their value to an after-tax basis. Using data from 1984–2007, we estimate the post-tax excess return volatility for stocks at 13.3% and for real estate at 5.5%, with a correlation of 0.27. The pre-tax excess stock return volatility is 19%.

In the third stage of our analysis, we define a numeraire asset, the aggregate Swedish

[4] This reflects the global diversification of Swedish equity portfolios documented by Calvet, Campbell, and Sodini (2007). It abstracts from underdiversification, which the same paper shows is modest for most Swedish households although important for a few. The impact of underdiversification in liquid wealth is further reduced when one takes account of DC retirement wealth as we do in this paper.

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

portfolio of cash, stocks, and real estate scaled to have the same volatility as the after-tax global equity index return:

$$R^e_{N,t+1} = (1+L)(\omega'_{agg,t}\ R^e_{t+1}). \tag{3}$$

Here $R^e_{N,t+1}$ is the return on the numeraire asset and $\omega_{agg,t}$ is the vector containing the weights of equity, real estate and the DC retirement portfolio in the aggregate net wealth of all Swedish households in our sample. The leverage parameter L is chosen so that the volatility of $R^e_{N,t+1}$ is equal to the volatility of the after-tax return in local currency on the global equity index.

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

$$\alpha_{h,t} = \sigma\left(R^e_{h,t+1}\right)/\sigma\left(R^e_{N,t+1}\right), \tag{4}$$

where household volatility is computed from equation (2). This approach implicitly assumes that all households earn the same Sharpe ratio on their risky assets, but guarantees that the standard deviation of a household's wealth return used in our simulations coincides with its empirical value. A value of one for $\alpha_{h,t}$ says that a household's portfolio has the same volatility, 13.3%, as if it invested solely in stocks held outside a retirement account, without borrowing or holding cash.

2.4 Composite Households

The full Swedish Income Registry data set contains almost 41 million household-year observations over the period 1999 to 2007, but we impose several filters on the panel. We exclude observations in which the head is a student, retired before the start of our sample, missing information on education or sector of employment, or outside the set of cohorts we consider. After requiring adequate data availability the resulting panel contains 4.8 million household-year observations or about 532,000 households in an average year.

We classify households by three levels of educational attainment: (i) basic or missing education, (ii) high school education, and (iii) post-high school education. We also classify households by 12 sectors of employment. Within each education level, we rank the sectors by their total income volatility and divide them in three categories. In this way we create a 3×3 grid of 9 large education/sector groups where the sectors of employment are aggregated by income volatility.[6]

[6] The estimation of income volatility is described in section 3, and details of the sectoral income risk classification are provided in table IA.2 in the online appendix.

We subdivide each of these large groups using a two-way sort by deciles of the initial risky share and the initial wealth-income ratio. We use the lowest two and highest two deciles and the middle three quintiles, giving us a 7×7 grid of 49 initial wealth-income and risky share groups.[7] At this stage we have 441 groups, which we further subdivide by 13 cohorts (birth years from 1947 to 1959) to create 5733 groups. After excluding some small groups that do not contain members in each year from 1999 to 2007, our final sample is a balanced panel of 4468 groups.

The median group size varies across years between 81 and 90 households, but the average group size is larger at about 120 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. To minimize the impact of households entering and exiting groups, 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.

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

We now consider the cross-section of the wealth-to-income ratio and risky share, averaging across all years in our sample.

The top panel of Table 1 shows the variation in average risky portfolio shares and wealth-income ratios across groups with each level of education and sectors of employment with each level of income risk, averaging across cohorts and the subdivisions by risky share and wealth-income ratio. For the purpose of computing these summary statistics, households in each group are treated as a single composite household that owns all wealth and receives all income of the group, and groups are weighted by the number of households they contain. Average risky shares vary in a narrow range from 68% to 72%, while wealth-income ratios vary more widely from 3.4 to 5.9. Within each sector, average risky portfolio shares vary little with education but average wealth-income ratios are higher for more educated households, particularly those with post-high school education. Across sectors, the level of income risk

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

has a weak negative effect on the risky portfolio share and a strong positive effect on the wealth-income ratio.

The averages in the top panel of Table 1 conceal a great deal of dispersion across disaggregated groups of households. This is shown by the bottom panel of Table 2, which reports the standard deviations of the risky portfolio share and wealth-income ratio across all the groups with a given education level and working in sectors with a given level of income risk. The standard deviations of the risky share are consistently in the range 21–28%, while the standard deviations of the wealth-income ratio are in the range 3.2–4.0. Across all 4468 groups, the average risky share has a mean of 69.5% with a standard deviation of 24.3% while the average wealth-income ratio has a mean of 4.4 with a standard deviation of 3.8.[8]

Figure 1 plots the cross-sectional distributions of the initial risky share and initial wealth-income ratio across all groups. The risky shares have a cross-sectional distribution that looks approximately normal in the range 0.2 to 1.2, but with a fat right tail including some probability mass above 2 (corresponding to a portfolio volatility above 26.6%). The wealth-income ratios have a strongly right-skewed distribution, with many groups having only a year or two of income accumulated, and a few having well over a decade of income.

The cross-sectional variation in wealth and asset allocation documented in Table 1 and Figure 1 suggests that it will be difficult to account for Swedish household behavior without allowing for cross-sectional variation in preferences. However, we have not yet accounted for cross-sectional variation in the wealth-income ratio at the start of our sample, which may reflect past shocks to income and wealth as well as heterogeneous savings behavior driven by preferences. We now develop a life-cycle model that we can use to estimate preferences from the evolution of wealth and asset allocation during our sample period, taking as given the initial wealth-income ratio and the income and financial returns received in each year of our sample.

3 Income Process and Life-Cycle Model

In this section, we present the labor income process and the life-cycle model of saving and portfolio choice that are used to estimate household preferences.

[8]The aggregation of households into groups naturally reduces the dispersion that is visible in the household-level data. However, the reduction is modest. Across all individual households in our dataset, the average risky share is 75% with a standard deviation of 53%, while the average wealth-income ratio is 4.4 with a standard deviation of 3.8.

3.1 Labor Income Process

We consider the labor income specification of Cocco, Gomes, and Maenhout (2005):

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

where $L_{h,t}$ denotes real income for household h in year t, a_h is a household fixed effect, $x_{h,t}$ 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 Cocco, Gomes, and Maenhout model by distinguishing between shocks that are common to all households in a group and shocks that are specific to each household in the group. To simplify notation, we neglect the group index g in the rest of this section.

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

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

The components ξ_t and $z_{h,t}$ follow independent random walks:

$$\xi_t = \xi_{t-1} + u_t, \tag{7}$$

$$z_{h,t} = z_{h,t-1} + w_{h,t}. \tag{8}$$

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

Finally, we assume that the three income shocks are i.i.d. Gaussian:

$$(u_t, w_{h,t}, \varepsilon_{h,t})' \sim \mathcal{N}(0, \Omega) \tag{9}$$

where Ω is the diagonal matrix with diagonal elements σ_u^2, σ_w^2, and σ_e^2. The model can be estimated by maximizing the likelihood function based on the Kalman filter, as the online appendix explains.

We estimate the income process from consecutive observations of household yearly income data over the period 1991 to 2007, excluding the first and last year of labor income to avoid measuring annual income earned over less than 12 months. In each year, we winsorize non-financial real disposable income to 1000 kronor or about $150. We consider the total income received by all members of the household, but classify households by the head's education level and age. Since the vast majority of Swedish residents retire at 65, we consider two age groups: (i) non-retired households less than 65, and (ii) retired households that are at least 65.

For active households younger than 65, we estimate b by running pooled regressions of equation (5) for each of the three education groups. As in Cocco, Gomes, and Maenhout (2005), the vector of explanatory variables $x_{h,t}$ includes age dummies. We also control for marital status, household size, and whether the head of the household is receiving unemployment benefits. We then regress the estimated age dummies on a third-degree polynomial in age and use the fitted third-degree polynomial in our life-cycle model.

For retired households, we impute the state and occupational after-tax pension benefit of each individual from 1999 to 2007, as explained in the online appendix. We fill forward the imputed pension benefit in real terms until 2007 at individual level, and aggregate income at the household level in each year. The replacement ratio is estimated for each education group as the fraction of the average income of non-retired 64-year-old households to the average income of retired 65-year-old households across the 1999 to 2007 period.

Figure 2 illustrates the estimated age-income profiles for our three education groups. The profiles are strikingly steep compared to profiles estimated in the US. Combined with the replacement ratios we estimate, they imply retirement income that is comparable to labor income received at age 40.

These steep age-income profiles reflect rapid growth of aggregate real labor income in Sweden during our sample period. Sweden experienced much slower growth in real labor income before the start of our sample, and we are not confident that Swedish households expected the income growth they received during the sample period.[9] Accordingly we also consider age-income profiles implied by real income growth that is 1% per year lower than estimated during our sample. These adjusted age-income profiles are also illustrated in Figure 2 and we use them as the basis for the results we report in the paper. We report results using unadjusted age-income profiles in the online appendix. The main effect of the income growth adjustment is to raise the average rate of time preference.

[9] The average growth rate of real labor income in Sweden was 1.77% per year in our sample period 1991–2007, but only 0.27% per year in the previous 15 years 1976–1990. The difference is 1.50%, larger than the adjustment we make in our base case.

3.2 Income Risks Across Groups

Table 2 reports the estimated standard deviations of group-level income shocks (permanent by assumption) and of permanent and transitory idiosyncratic income shocks, across levels of education and sectors of employment sorted into three categories by their total income risk.

Looking across sectors, group-level income volatilities and permanent idiosyncratic income volatilities vary relatively little, but transitory idiosyncratic income volatilities are considerably higher for high-risk sectors. The online appendix reports the underlying sectors that fall in each category. The patterns are intuitive, with relatively little transitory income risk in the public sector and in mining and quarrying, electricity, gas, and water supply, and relatively high transitory income risk in hotels and restaurants, real estate activities, construction for less educated workers, and the financial sector for more educated workers.

Table 2 also shows that educated households face larger transitory income risk, whereas permanent income risk is more evenly distributed across education levels. This pattern is consistent with Low, Meghir, and Pistaferri (2010), but it contrasts with earlier studies showing that in the United States, more educated people have lower transitory income risk and higher persistent income risk, or put slightly differently, that low-education people have "layoff risk" and high-education people have "career risk."

The explanation is likely due to the fact that in Sweden, uneducated workers face lower unemployment risk and enjoy higher replacement ratios than in many other countries, while educated workers face relatively high income losses when they do become unemployed. This results from the following features of the Swedish labor market. First, it is straightforward for companies to downsize divisions, but extremely difficult for them to lay off single individuals unless they have a high managerial position. Second, companies that need to downsize typically restructure their organizations by bargaining with unions. Third, unions are nationwide organizations that span large areas of employment and pay generous unemployment benefits. Fourth, the pay cut due to unemployment is larger for better paid jobs. After an initial grace period, an unemployed person will be required to enter a retraining program or will be assigned a low-paying job by a state agency. All these features imply that unemployment is slightly more likely and entails a more severe proportional income loss for workers with higher levels of education.[10]

We have already noted in discussing Table 1 that average wealth-income ratios tend

[10] See Brown, Fang, and Gomes (2012) for related research on the relation between education and income risk.

to be higher in sectors with riskier transitory income. This pattern is intuitive given that labor income risk encourages precautionary saving, which is particularly valuable to smooth consumption over periods of transitory income loss. However, there is little tendency for risky portfolio shares to be lower in sectors with riskier income.

Table 3 further explores these effects by regressing the average wealth-income ratio and risky share on dummies for high school and post-high school education, on either total income volatility or the separate volatilities of the three income shocks, and on age. All regressions also include fixed time effects.[11]

The first two columns of the table show that average wealth-income ratios increase with education and with income volatility, particularly the volatility of transitory income shocks. This is consistent with the view that wealth is accumulated in part as a buffer stock against such temporary shocks. Unsurprisingly, wealth-income ratios tend to increase with age as households save for retirement.

The third and fourth columns show that education and income risk have little effect on the risky share. However, age has a strong negative effect. This reflects the tendency for older Swedish investors to have less volatile portfolios as they pay down their mortgages and reduce the risk exposure of their financial asset holdings.

The last two columns add the average wealth-income ratio to the regression explaining the average risky share. This also has a strong negative effect, even controlling for age. Households with higher wealth-income ratios typically invest more conservatively. This reflects the fact that such households have less leverage and more cash than households with lower wealth-income ratios.

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 (Campbell 2018, p.309). Older households have fewer earning years remaining so their human capital is lower; and households with a higher wealth-income ratio have more financial capital relative to human capital. In both cases the tilt towards risky assets is reduced.

We work with a richer lifecycle model in which labor income is risky and nontradable, but that model implies a similar pattern of age and wealth effects on the risky share. We

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

will use our model to study the distribution of preferences across households with higher or lower education working in riskier or safer sectors.

3.3 Life-Cycle Model

We consider a standard life-cycle model, very similar to the one in Cocco, Gomes and Maenhout (2005).

Households have finite lives and Epstein-Zin utility over a single consumption good. The utility function V_t is specified by the coefficient of relative risk aversion γ, the elasticity of intertemporal substitution ψ, and the time preference parameter δ. It satisfies the recursion

$$V_t = \left[C_t^{1-1/\psi} + \delta \left(\mathbb{E}_t p_{t,t+1} V_{t+1}^{1-\gamma} \right)^{(1-1/\psi)/(1-\gamma)} \right]^{\frac{1}{1-1/\psi}}, \tag{10}$$

where $p_{t,t+1}$ denotes the probability that a household is alive at age $t+1$ conditional on being alive at age t. Utility, consumption, and the preference parameters γ, ψ, and δ all vary across households but we suppress the household index h in equation (10) for notational simplicity. The age-specific probability of survival, $p_{t,t+1}$, is obtained from Sweden's life table.

Capturing the wealth accumulation of young households poses several problems for life-cycle models which do not include housing purchases, transfers from relatives, investments in education, or changes in family size. In addition it is well-known that such models predict an extremely high equity share at early ages which is hard to reconcile with our data. For this reason, we focus on the stage of the life-cycle during which households have substantial retirement saving. We initialize the model at age 40 and endow households with the same initial wealth level as the one they actually have in the data. We follow the standard notational convention in life-cycle models and let the time index in the model, t, start at 1, so that t is calendar age minus 39. Each period corresponds to one year and agents live for a maximum of $T = 61$ periods (corresponding to age 100).

Matching the behavior of retirees is also hard for simple life-cycle models that do not incorporate health shocks or bequest motives. For this reason, we only consider the model's implications for ages 40 to 60 years. Our model includes no bequest motive, because it would be difficult to separately identify the discount factor and the bequest motive using our sample of households in the 40 to 60 age group, and we prefer not to add one more weakly identified parameter. Our estimates of the time discount factor can be viewed as having an upward bias due to the absence of a bequest motive in the model.

Before retirement households supply labor inelastically. The stochastic process of the household labor income, $L_{h,t}$, is described in Section 3.1. All households retire at age

65, as was typically the case in Sweden during our sample period, and we set retirement earnings equal to a constant replacement ratio of the last working-life permanent income. Consistent with the discussion in Section 2, total wealth, $W_{h,t}$, consists of all the assets held by the household. For tractability, we assume in the model that total wealth is invested every period in a one-period riskless asset (bond) and a composite risky asset. Initial wealth $W_{h,1}$ is calibrated from the data; it is set equal to the average wealth of a household with a 40-year-old head with similar characteristics as h.

The household chooses the consumption level $C_{h,t}$ and risky portfolio share $\alpha_{h,t}$ every period, subject to a constraint that prevents borrowing to finance consumption. We do allow borrowing to finance a risky asset position, that is, we allow $\alpha_{h,t} \geq 1$. We assume that the rate paid on borrowing exceeds the rate earned on cash investments, as discussed in section 2.3. Household wealth satisfies the budget constraint

$$W_{h,t+1} = (R_f + \alpha_{h,t} R^e_{t+1})(W_{h,t} + L_{h,t} - C_{h,t}), \tag{11}$$

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

3.4 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 of permanent group-level, permanent idiosyncratic, and transitory idiosyncratic income shocks reported in Table 2.

In our model we assume that all safe borrowing and lending takes place at a single safe interest rate of 2.0%. This is calibrated as a weighted average of a safe lending rate of 0.8% and the average household borrowing rate of 3.6%, using the cross-sectional average household debt level to construct the average.[12]

We set the volatility of the numeraire risky asset at 13.3%, which is equal to the volatility of post-tax excess stock returns as discussed in section 2.3. We assume that the average excess return on the numeraire asset over the 2.0% 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

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

1984–2007. Putting these assumptions together, we assume a Sharpe ratio of 0.26. Naturally there are alternative assumptions that can be made about the average reward for risk, but changing this parameter would primarily affect the average values of preference parameters (particularly risk aversion), and our main focus is on the cross-sectional dispersion in these parameters.

The remaining parameter that must be calibrated is the correlation between the numeraire risky asset return and group-level income shocks. We estimate this correlation lagging the risky asset return one year, following Campbell, Cocco, Gomes, and Maenhout (2001), to capture a delayed response of income to macroeconomic shocks that move asset prices immediately. Empirically the correlation is very similar across the 9 large education-sector groups, and we set it equal to the average value of 0.44. This is intermediate between a lower value of 0.27 for the correlation estimated using only stock returns, and a higher value of 0.87 for the correlation estimated using only real estate returns.

The correlation between the numeraire risky asset return and individual income growth is much smaller than 0.44, because most individual income risk is idiosyncratic. To illustrate with a representative example, a household with group-level standard deviation of 3%, permanent idiosyncratic standard deviation of 8%, and transitory idiosyncratic standard deviation of 12% would have a correlation with the numeraire risky asset of 0.20 for its permanent income shocks and only 0.11 for its total income shocks. Nonetheless, the group-level income correlation plays an important role in our model, because it helps to choke off household demand for risky assets even at moderate levels of risk aversion.

4 Identification and Estimation

This section explains our procedure for estimating and identifying household preference parameters. Using the calibrated income and asset-return parameters as inputs, we solve the life-cycle model for each of our 4468 household groups on a multi-dimensional grid for the three unknown preference parameters.[13] At each point on the grid, after solving the model, we simulate it conditioning on the initially observed wealth-income ratio and feeding in historically realized income and return shocks. From these simulations we calculate the model's implied path of wealth accumulation and asset allocation ($\{\alpha_{it}\}_{t=2,9}$ and $\left\{\left(\frac{W}{Y}\right)_{it}\right\}_{t=2,9}$), which we then use to estimate the preference parameters using the procedure described below.[14]

[13] As described below, in our estimation procedure we use interpolation methods which allow us to consider solutions that are not in the initial grid.

[14] We could potentially also include the portfolio allocation in the first year (α_{i1}), since that is also an endogenous moment from the simulations. We decided to exclude it in order to have the same number of

We consider a large number of groups (4468) so that we can indeed measure the cross-sectional distribution of preferences, but we do not estimate the model for each individual household for two main reasons. First, by grouping households into bins we hope to eliminate, or at least significantly decrease, the impact of idiosyncratic events that they might face and which we do not capture in our model. Second, the use of multiple households allows us to derive properties for our estimator relying on cross-sectional asymptotics.

4.1 The Identification Challenge

Our goal is to identify three separate preference parameters: the subjective discount factor (δ), the elasticity of intertemporal substitution (EIS, ψ) and the coefficient of relative risk aversion (γ). In a model with incomplete markets all three parameters affect both portfolio shares and wealth accumulation making their identification non-trivial. The main challenge comes from separately identifying the discount factor and the EIS, as we discuss next.

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

$$1 = \mathrm{E}_t \left[\widetilde{\delta}_{t+1} \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{1}{\psi}} \left(\frac{V_{t+1}}{\mu(V_{t+1})} \right)^{\frac{1}{\psi}-\gamma} R^P_{t+1} \right] \tag{12}$$

where $\widetilde{\delta}_{t+1} = \delta p_{t+1}$, $R^P_{t+1} = \alpha R^e_{t+1} + (1-\alpha) R_f$, and $\mu(V_{t+1})$ denotes the certainty equivalent of V_{t+1}.[15]

Taking logs of both sides and making the usual assumption of joint log-normality we obtain

$$\begin{aligned} 0 \;=\;& \log(\widetilde{\delta}_{t+1}) - \frac{1}{\psi}\mathrm{E}_t g_{t+1} + \left(\frac{1}{\psi} - \gamma\right) \mathrm{E}_t \widetilde{v}_{t+1} + \mathrm{E}_t r^P_{t+1} + \frac{1}{2\psi^2}\sigma^2_g \\ &+ \frac{1}{2}\left(\frac{1}{\psi} - \gamma\right)^2 \sigma^2_{\widetilde{v}} + \frac{1}{2}\sigma^2_r + \frac{1}{\psi}\left(\frac{1}{\psi} - \gamma\right)\sigma_{g\widetilde{v}} + \left(\frac{1}{\psi} - \gamma\right)\sigma_{\widetilde{v}r} + \frac{1}{\psi}\sigma_{gr}, \end{aligned} \tag{13}$$

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

---

moments related to the risky share and those related to the wealth accumulation.

[15] With labor income risk and a utility function that satisfies $u'(0) = -\infty$ the agent will always choose to hold some financial assets and therefore, even in the presence of borrowing constraints, the Euler equation still holds with equality. In our model we also have short-sales constraints on asset holdings but these do not bind for the middle-aged households we are considering.

Solving for $\mathrm{E}_t g_{t+1}$:

$$\begin{aligned}\mathrm{E}_t g_{t+1} &= \psi(\mathrm{E}_t r^P_{t+1} - \log(\widetilde{\delta}_{t+1})) + (1-\gamma\psi)\,\mathrm{E}_t \widetilde{v}_{t+1} \\ &\quad + \frac{1}{2\psi}\sigma^2_g + \frac{\psi}{2}\left[\left(\frac{1}{\psi}-\gamma\right)^2 \sigma^2_{\widetilde{v}} + \sigma^2_r + \left(\frac{1}{\psi}-\gamma\right)\sigma_{\widetilde{v}r}\right] + \left(\frac{1}{\psi}-\gamma\right)\sigma_{g\widetilde{v}} + \sigma_{gr}. \quad (14)\end{aligned}$$

The first term in equation (14) highlights the identification problem. If the expected portfolio return and discount rate are constant over time, then because δ and ψ appear multiplicatively, one can always change δ and ψ in offsetting ways without altering the value of that term.

Equation (14) also illustrates two possible solutions to this problem. One obvious approach is to exploit the remaining terms in the equation. Unfortunately, previous analysis of consumption and portfolio choice with Epstein-Zin preferences shows that the EIS and the discount rate primarily affect savings behavior through the first term, while the remaining terms capture precautionary savings behaviour which is primarily driven by risk aversion.[16]

Our identification is instead based on exploiting time variation in $\mathrm{E}_t r^P_{t+1} - \log(\widetilde{\delta}_{t+1})$. Even though our model has no exogenous variation in expected returns for any individual asset, we have endogenous variation driven by changes in the agent's optimal portfolio as financial wealth accumulates and future labor income declines. We present evidence for such age and wealth effects on portfolio choice in Table 3 below. Furthermore, in our model $\widetilde{\delta}_{t+1}$ is adjusted for survival probabilities which are a function of age. These two sources of variation imply that the profile of the wealth-to-income ratio as a function of age is affected in different ways by the EIS and by the discount factor. Our identification strategy, discussed below, will build on this intuition.

4.2 Identification Strategy

4.2.1 Intuition

We motivate 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 a series of moments from the simulated

[16] Campbell and Viceira (1999) show that the optimal consumption- wealth ratio is, to a first-order, driven by the trade-off between the (endogenous) expected return on invested wealth and the discount rate, exactly the first term in equation (14). Their results are obtained in an infinite horizon model without labor income but Gomes and Michaelides (2005) reach the same conclusion numerically in a life-cycle model that is almost identical to the one we consider in this paper.

data.[17]

We first consider the risk aversion parameter. The portfolio share is an intuitive moment to explore here, so we run the following regression:

$$\gamma_i = k_\gamma^0 + k_\gamma^1 \overline{\alpha}_i + e_i \tag{15}$$

where i denotes a unit of observation in our sample, t denotes calendar time, and

$$\overline{\alpha}_i = \left(\frac{1}{9} \sum_{t=1}^{9} \alpha_{it} \right) \tag{16}$$

Confirming that the average risky share is a very good moment for identifying the risk aversion parameter, we find that the R^2 from this regression is 70%. This is an extremely high number since we are estimating a linear regression and imposing the same coefficients across groups. We know that the true relationship is non-linear and depends on the initial wealth-to-income ratio ($(W/Y)_{i1}$), which varies across groups. In Panel A of Table 4, specification 1, we report results obtained with a fifth-order polynomial in $\overline{\alpha}_i$ and fitting separate regressions to groups with similar initial wealth-to-income ratios. We consider 7 different clusters of $(W/Y)_{i1}$, with the cutoffs chosen to deliver a similar number of observations in each of them.[18] The R^2 is above 80% for 6 out of the 7 regressions (70% for the other), and above 90% for 3 out of 7.

Having found a good moment to identify risk aversion, we now turn our attention to the time discount factor and explore the moments related to wealth accumulation. The first moment that we consider is the growth rate of wealth over the full sample,

$$grWY_i = \left[\left(\frac{W}{Y} \right)_{i9} / \left(\frac{W}{Y} \right)_{i1} \right] \tag{17}$$

We regress the discount factor (δ) on $grWY_i$ and, for the reasons discussed above we consider a fifth-order polynomial and run separate regressions for the same 7 clusters of the initial wealth-to-income ratio. The results are shown in panel A of Table 4, specification 2. The R^2 is between 68% and 75% for the different regressions, confirming that this is a very good moment to identify the time discount factor.

[17] The values for the preference parameters are the same as those considered in the estimation. Likewise, for each set of preference parameters we consider all 4468 observed values of the initial wealth-income ratio, and all other inputs of the model are the same as in our baseline estimation.

[18] The 7 clusters are $]0,1]$, $]1,2]$, $]2,3]$, $]3,5]$, $]5,7]$, $]7,10]$ and > 10. In theory we should get even better results if we consider separate regressions for each initial value of W/Y, but in that case we would have thousands of regression results to report.

Finally we consider the EIS. As explained in the previous sub-section the EIS will create differences in the growth rate of wealth at each age, because $\mathrm{E}_t r^P_{t+1} - Ln(\widetilde{\delta}_{t+1})$ is also changing with age. This suggests that a good moment to identify the EIS would be the curvature (concavity or convexity) of the wealth to income ratio as a function of age:[19]

$$curvWY_i = \frac{\frac{1}{2}\left[\frac{W}{Y}_{i1} + \frac{W}{Y}_{i9}\right]}{\frac{W}{Y}_{i5}} - 1 \tag{18}$$

Given that our estimation will consider all moments and parameters jointly, we now run regressions of the EIS on all three moments ($\overline{\alpha}_i$, $grWY_i$, and $curvWY_i$). As before we consider fifth-order polynomials, and run 7 separate regressions for different clusters of the initial wealth-to-income ratio. The R^2 is between 26% and 35% across the 7 different regressions. These results suggest it is possible to identify the EIS based on these 3 moments, although not as powerfully as the other 2 parameters.

4.2.2 Identification with all moments

In the previous section we illustrated the intuition for the identification strategy. Motivated by those results, in our estimation we use the set of moments previously described: $\{\alpha_{it}\}_{t=2,9}$ and $\left\{\left(\frac{W}{Y}\right)_{it}\right\}_{t=2,9}$. By including the wealth-income ratio in each year as a separate moment we capture the full wealth accumulation profile, namely its growth rate and curvature, the two important features previously discussed. Panel B of Table 4 reports the R^2 from regressions of the three preference parameters on those 16 moments. As in the previous section, we consider a fifth order polynomial and run separate regressions for 7 different clusters of the initial wealth-income ratio. All the risk aversion regressions have an R^2 in excess of 90%, while the time discount factor regressions have an R^2 of 75% or higher. The EIS regressions deliver lower values for the R^2, but now ranging from 38% to 42%. [20]

[19] An alternative variable based on the same intuition would be

$$diffg_i = \left(\frac{W}{Y}\right)_{i9} / \left(\frac{W}{Y}\right)_{i8} - \left(\frac{W}{Y}\right)_{i2} / \left(\frac{W}{Y}\right)_{i1}$$

We have considered this alternative and the results are similar, but slightly weaker.

[20] Although these regressions include a much higher number of variables than those in the previous section, the significantly higher R^2 is not automatic. These are regressions on clean simulated data with more than 1 million observations, so variables with low explanatory power add close to nothing to the R^2. Along these lines we can compare a linear regression of these 16 moments against a fifth-order polynomial in the previous three (15 variables). The R^2 of the former is 83%, 65% and 24%, respectively for risk aversion, discount factor and the EIS. On the other hand, the corresponding R^2s of the second set of regressions are 78%, 58%, and 9%. So, keeping the number of variables in the regression almost the same (16 versus 15), a linear specification of these moments has significantly more explanatory power than a highly non-linear function of the previous three (particularly for the EIS).

Our identification is even more precise than the previous results suggest for two reasons. First, we condition on the exact value of the initial wealth-income ratio, not just on a broad range, and second we are estimating the three preference parameters simultaneously. Running separate regressions for each of 4468 groups in our estimation and each combination of the other two parameter values would result in close to 1 million regressions, and crucially we would not have enough data points in each regression to be able to fit these high-order polynomials. However, to illustrate this important point, in the final two columns of Table 4 we repeat the previous EIS regression conditioning on two alternative values of risk aversion, 3 and 5. Even though we are still using coarse wealth-to-income buckets, the regression R^2 in some cases is already close to 80%.[21]

4.3 Indirect Inference Estimator

Using the parameters from Table 2 as inputs, we solve the life-cycle model for the 4468 different groups of households. For each group g, we compute the wealth-income ratio and the risky share predicted by the model for every year between 1999 and 2007. To make the output from the model comparable with the data, we initialize the simulation by giving each group the same initial wealth-income ratio as they had in the data in 1999. Furthermore, in the simulations the realizations of risky asset returns and group-level income growth are based on the actual returns and income growth observed between 1999 and 2007.

The estimation of the vector of preference parameters, $\theta^g = (\delta^g, \gamma^g, \psi^g)'$, in each group g proceeds by indirect inference (Smith 1993, Gouriéroux Monfort and Renault 1993). This method compares a vector of auxiliary statistics produced by the model to the vector of empirical auxiliary statistics in the group. We denote by $p = 3$ the number of components of θ^g, and by N^g the number of households in the group.

For every $t \in \{2, ..., 9\}$, we consider the following auxiliary statistics: (i) the risky share of the group, defined as the ratio of the group's risky wealth to the group's total wealth:

$$\hat{\mu}^g_{1,t} = \frac{\sum_{i=1}^{N^g} \alpha_{i,t} W_{i,t}}{\sum_{i=1}^{N^g} W_{i,t}}, \tag{19}$$

and (ii) the group's wealth-to-income ratio:

$$\hat{\mu}^g_{2,t} = \frac{\sum_{i=1}^{N^g} W_{i,t}}{\sum_{i=1}^{N^g} Y_{i,t}}. \tag{20}$$

[21] If we were to condition on a specific value of the discount factor as well then the R^2 would be above 99% for all cases. Naturally, for a fixed value of the discount factor there is no identification problem.

We stack these auxiliary statistics into a column vector $\hat{\mu}^g$, which we call the *empirical auxiliary estimator*. By construction, the auxiliary estimator $\hat{\mu}^g$ has $q = 16$ components.

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

For a given parameter vector θ, we compute $\tilde{\mu}^g_S(\theta)$ by simulating the sample paths of S households over the 9 years in our sample and then computing the risky share and wealth-to-income ratios in year t by the same method as above. In practice, we use $S = 10,000$ simulations for each group. As the number of households in the group goes to infinity, the auxiliary estimator $\tilde{\mu}^g_S(\theta)$ converges to the *binding function* $\mu^g(\theta) \in \mathbb{R}^q$ with components

$$\begin{aligned} \mu^g_{1,t}(\theta) &= \frac{E^g_\theta(\alpha_t W_t)}{E_\theta(W_t)}, \\ \mu^g_{2,t}(\theta) &= \frac{E^g_\theta(W_t)}{E^g_\theta(Y_t)}, \end{aligned}$$

where $E^g_\theta(\,\cdot\,)$ denotes the cross-sectional mean of households in the group. The expectation is taken across realizations of idiosyncratic income shocks at the household level, conditional on the initial value of the wealth-income ratio and the realizations of asset returns and group-level income shocks.

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

$$\hat{\theta}^g = \arg\min_\theta \left[\tilde{\mu}^g_S(\theta) - \hat{\mu}^g\right]' \Omega_g \left[\tilde{\mu}^g_S(\theta) - \hat{\mu}^g\right], \tag{21}$$

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

If our model is correctly specified, the preference parameters estimated by this procedure converge to the true preference parameters as the number of households in each group increases provided that we feed into the simulation all group-level shocks, as we now explain. Even though the definition of the empirical auxiliary estimator is not entirely standard, we know that $\hat{\mu}^g$ is asymptotically normal:

$$\sqrt{N^g}\,[\hat{\mu}^g - \mu^g(\theta)] \to N(0, W_g).$$

This result follows from the delta method and the fact that the auxiliary statistics (19) and (20) can be interpreted as ratios of sample moments. For the same reason, the simulated vector $\tilde{\mu}^g_S(\theta)$ is also asymptotically normal.

Let θ^g denote the true but unknown vector of structural parameters. By Gouriéroux Monfort and Renault (1993), the indirect inference estimator is asymptotically normal:

$$\sqrt{N^g}\left(\hat{\theta}^g - \theta^g\right) \to \mathcal{N}(0, V^g). \tag{22}$$

The asymptotic variance-covariance matrix is given by

$$V^g = \left(1 + s_g^{-1}\right)(D_g \Omega_g D_g')^{-1} D_g \Omega_g W_g \Omega_g D_g' (D_g \Omega_g D_g')^{-1}, \tag{23}$$

where $s_g = S/N_g$ is the ratio of the number of simulated households to the actual group size, and $(D_g)' = \partial \mu^g(\theta^g)/\partial\theta'$ is the Jacobian matrix of the binding function $\mu^g(\,\cdot\,)$ evaluated at the true parameter θ^g. The asymptotic variance-covariance matrix of the indirect inference estimator is consistently estimated by the sample analogue of equation (23).[22]

In the current version of the paper, we set the $q \times q$ weighting matrix Ω_g equal to a diagonal matrix $\Omega = (\Omega_{i,j})$, which is the same for all household groups. For each auxiliary statistic $i \in \{1, ..., q\}$, the diagonal element $\Omega_{i,i}$ is the inverse of the sample variance of the empirical auxiliary statistic $\hat{\mu}_i^g$ across groups. In order to reduce noise, we weigh the empirical auxiliary statistics $\hat{\mu}_i^g$ by the group size N_g in the calculation of the variance.

The asymptotic variance covariance matrix of the auxiliary estimator, $\hat{\mu}^g$, is estimated by the jackknife estimator

$$\frac{\hat{W}^g}{N^g} = \frac{N^g - 1}{N^g} \sum_{i=1}^{N^g} (\hat{\mu}_{[i]}^J - \overline{\mu^J})(\hat{\mu}_{[i]}^J - \overline{\mu^J})', \tag{24}$$

where $\hat{\mu}_{[i]}^J$ is the auxiliary estimator obtained by excluding the i^{th} observation, and $\overline{\mu^J} = (N^g)^{-1} \sum \hat{\mu}_{[i]}^J$. The use of the jackknife allows us estimate W^g even though the auxiliary estimator is not defined as a sample sum or an M-estimator and therefore standard estimators of W^g do not apply. The asymptotic distribution of the indirect inference estimator then follows from (22) and (23).

The numerical optimization of equation (21) proceeds by evaluating the function $\tilde{\mu}_S^g(\,\cdot\,)$ on a grid of preference points θ. Values of $\tilde{\mu}_S^g(\,\cdot\,)$ outside the grid are obtained by spline interpolation. The online appendix gives details of the numerical optimization procedure.

[22] A finite-sample estimator of the variance-covariance matrix V_g is

$$\hat{V}_g = \left(1 + s_g^{-1}\right)\left(\hat{D}_g \Omega_g \hat{D}_g'\right)^{-1} \hat{D}_g \Omega_g V_g^* \Omega_g \hat{D}_g' (\hat{D}_g \Omega_g \hat{D}_g')^{-1},$$

where $\hat{D}_g'$ is the Jacobian matrix of the binding function $\tilde{\mu}_S^g$ evaluated at $\hat{\theta}_g$.

In the next version of the paper, we plan to implement the following efficient indirect inference procedure. As previously, we consider the jackknife estimator (24) of the asymptotic variance covariance matrix of the auxiliary estimator, $\hat{\mu}^g$. We next consider the weighting matrix $\hat{\Omega}^g = (\hat{W}^g)^{-1}$ and compute the *efficient* indirect inference estimator $\hat{\theta}^g$ defined by (21). Note that a single step is required to compute the efficient estimator, which is due to the fact that household choices within the groups are assumed to be uncorrelated. The asymptotic variance-covariance matrix of the indirect inference estimator $\hat{\theta}^g$ then reduces to

$$\left(1+\frac{1}{s_g}\right)\left(\hat{D}_g\hat{\Omega}^g\hat{D}_g'\right)^{-1},$$

where as previously, $(\hat{D}_g)'$ is an estimator of the Jacobian matrix of the binding function at θ^g. The objective function at the optimum satisfies

$$\frac{S}{s_g+1}\left[\tilde{\mu}_S^g(\hat{\theta}^g)-\hat{\mu}^g\right]'\hat{\Omega}^g\left[\tilde{\mu}_S^g(\hat{\theta}^g)-\hat{\mu}^g\right]\to\chi^2(q-p),$$

which provides a convenient overidentification test and therefore allows us to assign a p-value to the optimal value of the objective function obtained for each group.

5 Empirical Results

5.1 The Cross-Sectional Distribution of Preferences

Table 5 summarizes the estimates of our model's three preference parameters. Panel A reports the means and standard deviations of the fitted preference parameters, the initial risky share and wealth-income ratio, and the three observable moments used in estimation. Panel B reports the cross-sectional correlations of the preference parameters, and panel C reports the correlations of the observed group characteristics with each other and with the preference parameters.

The cross-sectional mean of risk aversion is 5.06, with a cross-sectional standard deviation of 0.44. The mean of the EIS is 1.32 with a standard deviation of 0.99. The standard deviation of the EIS is over twice as large in absolute terms as the standard deviation of risk aversion. The difference in standard deviations is even larger in proportional terms, as shown at the right of Table 5, panel A where moments are reported for log risk aversion and the log EIS. The power utility model would imply equal standard deviations for log risk aversion and the log EIS, but the empirical standard deviation is 12 times as large for the log EIS.

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

The mean rate of time preference is 1.57% with a large standard deviation of 4.13%. Since our model omits bequest motives, we interpret the relatively low average estimate as a reflection of households' desire to accumulate wealth for bequests as opposed to retirement consumption. The mean rate of time preference would be even lower if we used the sample average growth rate of real labor income in calibrating our model, as discussed above.

Figure 3 illustrates the cross-sectional distributions of the three preference parameters. The distribution of risk aversion is somewhat left-skewed, with almost no mass above 6. The distribution of the EIS is U-shaped, with probability mass concentrated below 1 and above 2. A significant fraction of groups have EIS estimates at the edge of our parameter space. These are not plotted in the figure but are reported in the figure note: 5% of households have EIS estimated at 0.1 (the lowest value we allow), and 26% have EIS estimated at 2.5 (the highest value we allow). The EIS is less than one for 49% of households, and even less than the reciprocal of risk aversion for 15% of households, while it is greater than one for 51% of households.

The distribution of the rate of time preference has relatively little probability mass around zero, but significant mass at negative values between -1% and -3% and again at positive values between 2% and 4%. There is a long right tail of high time preference estimates, and 6% of households (not plotted in the figure) have time preference of -5%, the lowest value we allow.

Panel B of Table 5 shows that our estimates of preference parameters are only weakly cross-sectionally correlated, indicating that heterogeneity in household preferences is multi-dimensional. The coefficient of risk aversion has a very weak positive correlation of 0.07 with the EIS, and the correlation is equally weak in logs. This is contrary to the strong negative correlation implied by power utility preferences. With power utility, risk aversion is the reciprocal of the EIS so the two parameters would have a perfect negative correlation in logs.

The rate of time preference is weakly positively correlated (0.18) with the coefficient of risk aversion, and somewhat more strongly correlated (0.36) with the EIS.

Panel C reports correlations with observable variables that help to understand these estimates. The initial risky share is almost perfectly correlated with the average risky share, implying that the risky share is a highly persistent household attribute. The initial wealth-income ratio is strongly negatively correlated (-0.80) with the average growth rate of the wealth-income ratio, indicating a tendency for mean-reversion in the wealth-income ratio. The average risky share is negatively correlated (-0.48) with the initial wealth-income ratio.

Our estimate of risk aversion is negatively correlated (-0.22) with the average risky share, an intuitive result that is also consistent with our identification analysis. Risk aversion is also negatively correlated (-0.27) with the initial wealth-income ratio, reflecting the fact that households with more initial wealth have only a weak tendency to lower their risky portfolio shares.

Our estimate of the EIS has only weak correlations with observables, reflecting the fact that the EIS affects wealth accumulation differently for different values of the time discount factor. The strongest correlation (0.30) is with the initial wealth-income ratio.

Our estimate of the rate of time preference is negatively correlated (-0.36) with the initial wealth-income ratio and positively correlated (0.49) with the average growth rate of the wealth-income ratio in our sample period. Mechanically, this is due to the fact that households that enter our sample with low initial wealth accumulate wealth more rapidly than average households, but not as rapidly as they would do if they had an average rate of time preference. Economically, it is intuitive that impatient households accumulate less wealth before age 40 and then belatedly catch up as retirement approaches. The rate of time preference is also positively correlated with the average risky share, reflecting the lower average wealth-income ratio of impatient households that justifies a riskier investment strategy.

Table 6 explores the relation between preference parameters in a different way. Panel A of this table regresses log risk aversion on cohort dummies in the first column, the log EIS in the second column, the log EIS and cohort dummies in the third column, both the log EIS and time preference in the fourth column, and the two preference parameters and cohort dummies in the fifth column. The cohort dummies are generally insignificant and show no clear patterns. The log EIS has a small positive effect on risk aversion when included by itself, but is driven out of the regression when time preference is also included. The explanatory power is low in all these regressions, never exceeding 3%.

Panel B of Table 6 proceeds in a similar fashion with the log EIS as the dependent variable. Again there are no important cohort effects. Time preference has a positive effect and the R^2 statistic when all regressors are included is almost 15%.

Finally, panel C uses the rate of time preference as the dependent variable. Cohorts born earlier have lower rates of time preference, and both risk aversion and the EIS predict time preference positively. The R^2 statistic when all regressors are included is almost 17%.

We noted earlier that some of our parameter estimates are at the edge of the parameter space. This tends to be associated with a poor fit of the model to the data (a high value of the minimized objective function). More generally, some groups have an erratic history of their wealth-income ratios that is difficult for our life-cycle model to fit. A crude way to assess whether this affects our qualitative description of the results is to recompute the results only for those households that are fit sufficiently well by the model. We have done this for the 50% of groups with the lowest minimized objective function, and most cross-sectional patterns are similar. In particular, the means and cross-sectional standard deviations of preference parameters are little different in this subset of the population. The positive cross-sectional correlation between the EIS and the rate of time preference is considerably stronger at 0.64, but risk aversion remains very weakly positively correlated with both the other preference parameters.

5.2 Household Characteristics and Preferences

In Table 7 we explore the relationship between preference estimates and household characteristics: education and labor income volatility, measured either as total income volatility or as separate volatilities for permanent and temporary income shocks. All regressions also include cohort fixed effects, but the cohort coefficients are not reported.

A striking result in the table is that volatile labor income, and particularly volatile permanent components of labor income, are associated with lower estimated risk aversion. Mechanically, this results from the fact documented in Table 3 that income volatility has little effect on the risky share: if risk aversion were the same for safe and for risky occupations, then the risky share should fall with income risk. Economically, the finding suggests that risk-tolerant individuals may self-select into risky occupations, although such a pattern could also result if households fail to fully understand the importance of income risk for optimal investment strategies. The table also shows some evidence that the EIS and the rate of time preference are higher for people with risky incomes.

There are also some interesting correlations between education and preference parameters. Households with higher education tend to have a lower rate of time preference, an intuitive result that is consistent with their greater wealth accumulation. These households also have a lower EIS and, if all components of income risk are controlled for, a lower coefficient of risk aversion. High school education has weaker and less consistent effects but appears similar to higher education at least in its correlations with EIS and risk aversion.

These results, like those reported in Tables 5 and 6, are qualitatively robust to considering only the 50% of groups for which our model achieves the best fit.

6 Conclusion

In this paper we have asked whether the patterns of wealth accumulation and risky investment among Swedish households can be rationalized by a life-cycle model with homogeneous preferences, or whether households appear to have different preferences. By sorting households not only on birth cohort, level of education, and the income risk in their sector of employment, but also on initial wealth accumulation and the risky portfolio share, we create household groups with substantial heterogeneity in their financial behavior.

The maintained assumption throughout the paper is that all households have common expectations about the riskless interest rate and risky asset returns, understand the stochastic processes driving their labor income, and invest rationally given their preferences and information. Under this assumption and with the parameters we calibrate for income and asset returns, our model fits the data with a cross-sectional standard deviation of risk aversion of 0.44 around a mean of 5.06, a standard deviation in the EIS of 0.99 around a mean of 1.32, and a standard deviation in the time preference rate of 4.13% around a low mean of 1.57% that may in part reflect bequest motives that are omitted from our life-cycle model.

About 85% of households are estimated to have risk aversion that is higher than the reciprocal of the EIS, violating the parameter restriction of power utility. Furthermore, the cross-sectional association between risk aversion and the EIS is very weakly positive which contradicts the qualitative prediction of a power utility model that these two parameters should be strongly negatively correlated. We find that the rate of time preference is weakly positively correlated with risk aversion and more strongly positively correlated with the EIS.

There is a negative correlation between income risk and risk aversion. This is consistent with a model in which risk-tolerant households self-select into risky occupations, but could also reflect rules of thumb for asset allocation that do not adapt appropriately to income risk.

Households with higher education have lower rates of time preference on average, as well as lower EIS and risk aversion when one controls for the income risk in their sectors of employment.

Our results shed light on a number of issues in asset pricing and household finance.

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

Even if individual households have constant preference parameters, cross-sectional heterogeneity in these parameters can break the relation between household preferences and the implied preferences of a representative agent. In a representative-agent economy, preferences with habit formation are needed to generate countercyclical variation in the price of risk (Constantinides 1990, Campbell and Cochrane 1999), but in heterogeneous-agent economies, countercyclical risk premia can arise from time-variation in the distribution of wealth across agents with different but constant risk preferences (Dumas 1989, Chan and Kogan 2002, Guvenen 2009). Gomes and Michaelides (2005 and 2008) illustrate the importance of preference heterogeneity for simultaneously matching the wealth accumulation and portfolio decisions of households. Our empirical evidence can be used to discipline these modeling efforts.

Importantly, we estimate multi-dimensional heterogeneity in preferences: the correlations among our estimated preference parameters are relatively low. This implies that a single factor omitted from our model, such as heterogeneity in expected stock returns of the sort documented in survey data by Vissing-Jørgensen (2003), Dominitz and Manski (2011), Amromin and Sharpe (2013), and Giglio, Maggiori, Stroebel, and Utkus (2019), is unlikely to reconcile the data with homogeneous underlying preferences.

In household finance, there is considerable interest in estimating risk aversion at the individual level and measuring its effects on household financial decisions. This has sometimes been attempted using direct or indirect questions in surveys such as the Health and Retirement Study (Barsky et al 1997, Koijen et al 2014), the Survey of Consumer Finances (Bertaut and Starr-McCluer 2000, Vissing-Jørgensen 2002 *b* , Curcuru et al 2010, Ranish

2014), and similar panels overseas (Guiso and Paiella 2006, Bonin et al 2007). One difficulty with these attempts is that even if risk aversion is correctly measured through surveys, its effects on household decisions will be mismeasured if other preference parameters or the properties of labor income covary with risk aversion. Our estimates suggest that this should indeed be a concern.

Similarly, there is interest in measuring the effects of labor income risk on households' willingness to take financial risk (Calvet and Sodini 2014, Guiso, Jappelli, and Terlizzese 1996, Heaton and Lucas 2000). Models such as those of Campbell et al (2001), Viceira (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 (Ranish 2014). Our estimates suggest that this too is a serious empirical issue.

Our findings may also contribute to an ongoing policy debate over approaches to consumer financial protection. If all households have very similar preference parameters, strict regulation of admissible financial products should do little harm to households that optimize correctly, while protecting less sophisticated households from making financial mistakes. To the extent that households are heterogeneous, however, such a stringent approach is likely to harm some households by eliminating financial products that they prefer (Campbell et al 2011, Campbell 2016, Calvet et al 2017).

Our model omits some features of the household decision problem that may potentially be important and deserve further research. We assume that preference parameters do not vary with wealth at the household level, contrary to evidence that relative risk aversion declines with wealth (Carroll 2000, 2002, Wachter and Yogo 2010, Calvet and Sodini 2014). We treat labor income as exogenous and do not consider the possibility that the household can endogenously vary its labor supply (Bodie, Merton, and Samuelson 1992, Gomes, Kotlikoff and Viceira 2008). We ignore the possibility that some components of consumption involve precommitments or generate habits that make them costly to adjust (Gomes and Michaelides 2003, Chetty and Szeidl 2007, 2010). We also do not model fixed costs of stock market participation (Haliassos and Bertaut 1995, Vissing-Jørgensen 2002, Ameriks and Zeldes 2004, Gomes and Michaelides 2005, Fagereng, Gottlieb, and Guiso 2017) because we restrict our sample to middle-aged households who have already decided to hold risky assets outside their retirement accounts and pay any one-time costs of doing so. We do not model homeownership jointly with other financial decisions as in Cocco (2005). Household-level data on asset allocation and wealth accumulation provide a laboratory in which these aspects of household financial decision-making can be explored.

References

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

Amromin, Gene, and Steven A. Sharpe, 2013, From the horse's mouth: Economic conditions and investor expectations of risk and return, *Management Science* 60, 845–866.

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini, 2018, Rich pickings? Risk, return, and skill in household wealth, unpublished paper, ESSEC Business School, EDHEC Business School, and Stockholm School of Economics.

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.

Bertaut, Carol, and Martha Starr-McCluer, 2002, Household portfolios in the United States, in Luigi Guiso, Michael Haliassos and Tullio Jappelli, eds.: *Household Portfolios* (MIT Press).

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

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

Bonin, Holger, Thomas Dohmen, Armin Falk, David Huffman, and Uwe Sunde, 2007, Cross-sectional earnings risk and occupational sorting: The role of risk attitudes, *Labour Economics* 14, 926–937.

Brown, Jeffrey, Chichun Fang, and Francisco Gomes, 2012, Risk and returns to education, NBER Working Paper 18300.

Calvet, Laurent E., John Y. Campbell, Francisco Gomes, and Paolo Sodini, 2019, Online appendix to the cross-section of household preferences, unpublished paper, EDHEC Business School, Harvard University, London Business School, and Stockholm School of Economics.

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

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

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

Calvet, Laurent E., Claire Célérier, Paolo Sodini, and Boris Vallée, 2017, Can financial innovation solve household reluctance to take risk?, unpublished paper, EDHEC Business School, Rotman School of Management, Stockholm School of Economics, and Harvard Business School.

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

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

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

Campbell, John Y., 2018, *Financial Decisions and Markets: A Course in Asset Pricing* (Princeton University Press).

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

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

Campbell, John, Y. and Luis Viceira, 1999, Consumption and portfolio decisions when expected returns are time varying, *Quarterly Journal of Economics*, 114, 433–495.

Campbell, John Y., Howell Jackson, Brigitte Madrian, and Peter Tufano, 2011, Consumer financial protection, *Journal of Economic Perspectives* 25(1), 91–114.

Carroll, Christopher D., 2000, Why do the rich save so much?, in Joel B. Slemrod, ed.: *Does Atlas Shrug? The Economic Consequences of Taxing the Rich* (Harvard University Press).

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

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

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

Chan, Yeung Lewis and Leonid Kogan, 2002, Catching up with the Joneses: Heterogeneous preferences and the dynamics of asset prices, *Journal of Political Economy* 110, 1255–1285.

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

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

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

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

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

Curcuru, Stephanie, John Heaton, Deborah Lucas, and Damien Moore, 2010, Heterogeneity and portfolio choice: Theory and evidence, in Yacine Ait-Sahalia and Lars Peter Hansen eds. *Handbook of Financial Econometrics Vol. 2*, 337–382.

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

Dominitz, Jeff and Charles F. Manski, 2011, Measuring and interpreting expectations of equity returns, *Journal of Applied Econometrics* 26, 352–370.

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

Du Rietz, Gunnar, Dan Johansson, and Mikael Stenkula, 2015, Swedish capital income taxation (1862–2013), in Magnus Henrekson and Mikael Stenkula eds. *Swedish Taxation: Developments since 1862* (Palgrave Macmillan), 123–178.

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

Epstein, Larry G. and Stanley E. Zin, 1991, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis, *Journal of Political Economy* 99, 263–286.

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

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2019, Five facts about beliefs and portfolios, unpublished paper, Yale University, Harvard University, New York University, and Vanguard Group.

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

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

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

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

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

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

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

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

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

Hall, Robert E., 1988, Intertemporal substitution in consumption, *Journal of Political Economy* 96, 221–273.

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

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, 1984, Expanding the life-cycle model: Precautionary saving and public policy, *American Economic Review* 84, 174–179.

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.

Koijen, Ralph S. J., Stijn van Nieuwerburgh, and Motohiro Yogo, 2016, Health and mortality delta: Assessing the welfare cost of household insurance choice, *Journal of Finance* 71, 957–1010.

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

Ranish, Benjamin, 2014, Why do households with risky labor income take greater financial risks?, Chapter 1 in *Essays on Stock Investing and Investor Behavior*, Ph.D. dissertation, Harvard University.

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

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

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

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

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

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

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

Table 1
Summary Statistics on Risky Share and Wealth-Income Ratio

This table reports the cross-sectional mean (Panel A) and the cross-sectional standard deviation (Panel B) of the risky share and the wealth-income ratio for three groups of households sorted by education groups and groups of employment sectors sorted by volatility.

Panel A: Cross-Sectional Mean

	No High School Degree		High School Degree		Post-High School Education	
	Risky Share	Wealth-Income Ratio	Risky Share	Wealth-Income Ratio	Risky Share	Wealth-Income Ratio
Sectors						
- High Volatility	68.8%	4.42	69.7%	4.70	68.5%	5.88
- Medium Volatility	70.5%	3.81	69.6%	4.02	68.7%	4.81
- Low Volatility	71.9%	3.36	71.2%	3.74	68.4%	4.67
Aggregate	70.4%	3.86	70.2%	4.09	68.5%	5.00

Panel B: Cross-Sectional Standard Deviation

	No High School Degree		High School Degree		Post-High School Education	
	Risky Share	Wealth-Income Ratio	Risky Share	Wealth-Income Ratio	Risky Share	Wealth-Income Ratio
Sectors						
- High Volatility	25.8%	3.86	24.2%	4.02	21.2%	4.00
- Medium Volatility	27.7%	3.68	25.1%	3.62	23.2%	3.83
- Low Volatility	26.9%	3.18	25.1%	3.34	23.0%	3.80
Aggregate	26.8%	3.60	24.9%	3.65	22.6%	3.89

Table 2
Labor Income Risk

This table reports the characteristics of labor income risk for three groups of households sorted by educational attainment. The risk characteristics reported for each group are the standard deviation of systematic permanent shocks, the standard deviation of idiosyncratic permanent shocks, and the standard deviation of idiosyncratic transitory shocks.

	No High School Degree				High School Degree				Post-High School Education			
	Systematic	Idiosyncratic Risk			Systematic	Idiosyncratic Risk			Systematic	Idiosyncratic Risk		
	Risk	Permanent	Transitory	Size	Risk	Permanent	Transitory	Size	Risk	Permanent	Transitory	Size
Sectors												
- High Volatility	3.30%	6.77%	15.91%	277,932	3.18%	7.75%	15.39%	526,070	3.77%	7.81%	16.95%	472,809
- Medium Volatility	3.05%	8.26%	11.62%	186,065	3.11%	7.63%	11.59%	855,215	3.40%	6.42%	12.76%	454,663
- Low Volatility	2.74%	6.93%	8.68%	296,619	2.85%	6.64%	9.13%	712,517	3.19%	6.21%	11.63%	1,012,289

Table 3
Risky Share, Wealth-Income Ratio, and Household Characteristics

This table reports regressions of the wealth-income ratio (columns 1 and 2) and the risky share (columns 3 to 6) on household characteristics. The explanatory variables include dummy variables for educational attainment, measures of labor income risk, the age of the household head, and time fixed effects. In columns 5 and 6, we also control for the average level of the wealth-to-income ratio.

	Wealth-Income Ratio (1)		Wealth-Income Ratio (2)		Risky Share (3)		Risky Share (4)		Risky Share (5)		Risky Share (6)	
	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat
Average wealth-income ratio									-0.031	-39.87	-0.031	-39.87
Education dummies												
- High school	0.396	2.96	0.373	2.75	-0.013	-1.34	-0.013	-1.35	-0.001	-0.11	-0.002	-0.20
- Post-high school	1.092	7.35	1.007	2.81	-0.026	-2.66	-0.017	-0.79	0.008	1.00	0.014	0.78
Labor income risk												
- Total volatility	19.090	7.83			-0.235	-1.55			0.357	2.82		
- Systematic income shocks			27.860	0.31			-1.739	-0.32			-0.870	-0.19
- Idiosyncratic permanent shocks			8.733	0.71			0.398	0.48			0.661	0.94
- Idiosyncratic transitory shocks			14.418	2.22			-0.159	-0.40			0.289	0.87
Age	0.155	9.49	0.155	9.49	-0.015	-14.24	-0.015	-14.24	-0.011	-11.82	-0.011	-11.82
Year dummies												
- 2000	-0.329	-15.64	-0.329	-15.64	-0.036	-21.07	-0.036	-21.06	-0.043	-24.97	-0.043	-24.97
- 2001	-0.583	-14.85	-0.583	-14.85	-0.076	-27.65	-0.076	-27.65	-0.087	-32.19	-0.087	-32.18
- 2002	-0.955	-16.37	-0.955	-16.37	-0.146	-37.71	-0.146	-37.71	-0.162	-42.66	-0.162	-42.65
- 2003	-0.728	-10.23	-0.728	-10.23	-0.116	-25.43	-0.116	-25.43	-0.135	-32.10	-0.135	-32.10
- 2004	-0.413	-4.83	-0.413	-4.83	-0.121	-22.00	-0.121	-22.00	-0.145	-28.84	-0.145	-28.84
- 2005	0.070	0.70	0.071	0.70	-0.109	-16.94	-0.109	-16.94	-0.138	-23.72	-0.138	-23.72
- 2006	0.368	3.17	0.368	3.17	-0.114	-15.36	-0.114	-15.36	-0.147	-22.44	-0.147	-22.44
- 2007	0.321	2.41	0.322	2.41	-0.130	-15.42	-0.130	-15.42	-0.168	-22.80	-0.168	-22.79
Constant	-6.518	-7.82	-6.963	-3.67	1.611	27.06	1.620	13.92	1.421	28.45	1.416	14.61
Adjusted R^2	0.090		0.090		0.160		0.160		0.365		0.365	
Number of groups	4,468		4,468		4,468		4,468		4,468		4,468	

Table 4
Identification of the Preference Parameters

This table reports the R^2 coefficients of regressions of preference parameters on characteristics of the wealth-income profile and asset allocation obtained from simulations, as explained in Section 4.2 of the main text. In both panels, the regressions are based on fifth-order polynomials of the explanatory variables.

Panel A: First Set of Identification Regressions

	R^2 Coefficients		
	Relative Risk Aversion γ	Discount Factor δ	Elasticity of Intertemporal Substitution ψ
	(1)	(2)	(3)
Explanatory Variables			
Mean risky share	Yes	No	Yes
Ratio of W/Y at t=9 to W/Y at t=1	No	Yes	Yes
Curvature of W/Y	No	No	Yes
Initial W/Y			
]0,1]	0.700	0.684	0.270
]1,2]	0.797	0.734	0.343
]2,3]	0.857	0.738	0.333
]3,5]	0.897	0.744	0.317
]5,7]	0.923	0.734	0.287
]7,10]	0.925	0.731	0.286
>10	0.951	0.705	0.261

Table 4
Identification of the Preference Parameters - *Continued*

Panel B: Second Set of Identification Regressions					
	R^2 Coefficients				
	Relative Risk Aversion γ	Discount Factor δ	Elasticity of Intertemporal Substitution ψ	Elasticity of Intertemporal Substitution ψ	Elasticity of Intertemporal Substitution ψ
	(1)	(2)	(3)	(4)	(5)
Restriction				γ=3	γ=5
Explanatory Variables					
W/Y at dates t=2 to 9	Yes	Yes	Yes	Yes	Yes
Risky share at dates t=2 to 9	Yes	Yes	Yes	Yes	Yes
Initial W/Y					
]0,1]	0.905	0.784	0.380	0.716	0.416
]1,2]	0.925	0.818	0.397	0.724	0.467
]2,3]	0.934	0.812	0.398	0.768	0.467
]3,5]	0.950	0.809	0.407	0.780	0.449
]5,7]	0.967	0.807	0.424	0.761	0.462
]7,10]	0.968	0.802	0.417	0.770	0.451
>10	0.977	0.783	0.378	0.633	0.391

Table 5
Cross-Sectional Properties of Preference Parameters and Financial Characteristics

This table reports the cross-sectional moments of preference parameters and observed financial characteristics. Panel A tabulates cross-sectional means and standard deviations. Panel B reports the cross-sectional correlation of preference parameters, both in levels and in logs. Panel C reports the cross-sectional correlations of observed financial characteristics, as well as their correlations with preference parameters.

Panel A. Cross-Sectional Mean and Standard Deviation

	In Levels		In Logs	
	Mean	Std. Deviation	Mean	Std. Deviation
Preference parameters				
- Relative risk aversion γ	5.06	0.44	1.62	0.09
- Elasticity of intertemporal substitution ψ	1.32	0.99	-0.19	1.11
- Time preference rate -log(δ)	1.57%	4.13%		
Observed financial characteristics				
- Initial risky share	0.85	0.38		
- Initial wealth-income ratio	4.03	4.26		
- Average risky share	0.70	0.22		
- Concavity of W/Y	0.23	0.09		
- Average growth rate of W/Y	1.08	0.05		

Panel B. Cross-Sectional Correlation of Preference Parameters

	In Levels			In Logs	
	Relative Risk Aversion γ	Elasticity of Intertemporal Substitution ψ	Time Pref. Rate -log(δ)	Relative Risk Aversion γ	Elasticity of Intertemporal Substitution ψ
Preference parameters					
- Relative risk aversion γ	1.000			1.000	
- Elasticity of intertemporal substitution ψ	0.068	1.000		0.066	1.000
- Time preference rate -log(δ)	0.175	0.356	1.000		

Table 5
Cross-Sectional Properties of Preference Parameters and Financial Characteristics - *Continued*

Panel C. Cross-Sectional Correlation of Preferences and Financial Characteristics								
	Preference Parameters			Observed Financial Characteristics				
	Relative Risk Aversion γ	Elasticity of Intertemporal Substitution ψ	Time Pref. Rate $-\log(\delta)$	Initial Risky Share	Initial Wealth-Income Ratio	Average Risky Share	Concavity of W/Y	Average Growth Rate of W/Y
Observed financial characteristics								
- Initial risky share	-0.228	0.060	0.219	1.000				
- Initial wealth-income ratio	-0.269	0.299	-0.355	-0.436	1.000			
- Average risky share	-0.215	0.051	0.276	0.982	-0.480	1.000		
- Concavity of W/Y	-0.070	0.027	0.127	0.077	-0.002	0.142	1.000	
- Average growth rate of W/Y	0.232	-0.026	0.494	0.637	-0.800	0.657	0.112	1.000

Table 6
Preference Parameters and Cohort Fixed Effects

This table reports regressions of the coefficient of relative risk aversion (Panel A), the elasticity of intertemporal substitution (Panel B), and the rate of time preference (Panel C) on the other two preference parameters. Cohort fixed effects are also included in columns 1, 3, and 5 of each panel.

Panel A. Dependent Variable: Log of Risk Aversion Coefficient

	(1)		(2)		(3)		(4)		(5)	
	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat
Preference Parameters										
- Log of elasticity intertemporal substitution, log(ψ)			0.005	3.17	0.005	3.34	0.000	0.13	0.000	0.16
- Time preference rate, -log(δ)							0.366	6.27	0.368	6.40
Cohort Dummies										
- Birth year 1958	-0.001	-0.78			-0.002	-1.12			0.001	0.23
- Birth year 1957	0.000	0.00			0.000	-0.19			0.001	0.43
- Birth year 1956	-0.002	-0.69			-0.001	-0.68			0.002	0.92
- Birth year 1955	-0.003	-1.26			-0.003	-1.21			0.001	0.44
- Birth year 1954	-0.002	-0.61			-0.002	-0.53			0.002	0.36
- Birth year 1953	-0.001	-0.21			-0.001	-0.14			0.003	0.73
- Birth year 1952	-0.002	-0.50			-0.003	-0.60			0.001	0.20
- Birth year 1951	-0.001	-0.26			-0.002	-0.31			0.003	0.58
- Birth year 1950	-0.003	-0.53			-0.003	-0.50			0.002	0.30
- Birth year 1949	-0.003	-0.42			-0.003	-0.43			0.001	0.13
- Birth year 1948	-0.002	-0.24			-0.001	-0.14			0.003	0.37
- Birth year 1947	-0.001	-0.15			0.000	0.03			0.005	0.61
Intercept	1.620	60.25	1.619	68.35	1.621	59.98	1.613	69.17	1.611	60.09
Adjusted R^2	0.000		0.004		0.005		0.029		0.029	
Number of groups	4,468		4,468		4,468		4,468		4,468	

Table 6
Preference Parameters and Cohort Fixed Effects - *Continued*

Panel B. Dependent Variable: Log of Elasticity of Intertemporal Substitution										
	(1)		(2)		(3)		(4)		(5)	
	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat
Preference Parameters										
- Log of coefficient of relative risk aversion, $-\log(\gamma)$			0.819	2.42	0.818	2.51	0.029	0.13	0.036	0.16
- Time preference rate, $-\log(\delta)$							10.070	3.87	9.988	3.77
Cohort Dummies										
- Birth year 1958	0.108	0.94			0.109	0.94			0.151	1.58
- Birth year 1957	0.077	0.70			0.077	0.70			0.113	1.40
- Birth year 1956	-0.052	-0.37			-0.051	-0.37			0.050	0.58
- Birth year 1955	-0.072	-0.64			-0.070	-0.62			0.038	0.49
- Birth year 1954	-0.050	-0.45			-0.048	-0.43			0.052	0.79
- Birth year 1953	-0.046	-0.32			-0.046	-0.31			0.058	0.76
- Birth year 1952	0.089	1.04			0.091	1.07			0.172	2.51
- Birth year 1951	0.055	0.63			0.056	0.64			0.164	2.63
- Birth year 1950	-0.035	-0.38			-0.033	-0.35			0.096	1.46
- Birth year 1949	0.012	0.15			0.015	0.17			0.109	1.67
- Birth year 1948	-0.140	-1.79			-0.139	-1.76			-0.016	-0.24
- Birth year 1947	-0.271	-3.21			-0.270	-3.14			-0.101	-1.46
Intercept	-0.162	-1.94	-1.518	-2.65	-1.486	-2.57	-0.397	-1.04	-0.473	-1.22
Adjusted R^2	0.008		0.004		0.013		0.141		0.146	
Number of groups	4,468		4,468		4,468		4,468		4,468	

Table 6
Preference Parameters and Cohort Fixed Effects - *Continued*

Panel C. Dependent Variable: Time Preference Rate										
	(1)		(2)		(3)		(4)		(5)	
	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat
Preference Parameters										
- Log of coefficient of relative risk aversion, -log(γ)			0.078	4.56	0.078	4.65	0.067	4.45	0.067	4.47
- Log of elasticity intertemporal substitution, log(ψ)							0.014	7.43	0.014	7.19
Cohort Dummies										
- Birth year 1958	-0.004	-0.98			-0.004	-0.95			-0.006	-1.43
- Birth year 1957	-0.004	-0.67			-0.004	-0.67			-0.005	-1.03
- Birth year 1956	-0.010	-2.03			-0.010	-2.02			-0.009	-2.47
- Birth year 1955	-0.011	-2.94			-0.011	-2.85			-0.010	-3.29
- Birth year 1954	-0.010	-2.11			-0.010	-2.03			-0.009	-2.37
- Birth year 1953	-0.010	-1.69			-0.010	-1.67			-0.010	-2.25
- Birth year 1952	-0.008	-4.27			-0.008	-3.98			-0.009	-6.43
- Birth year 1951	-0.011	-2.81			-0.011	-2.78			-0.012	-4.09
- Birth year 1950	-0.013	-2.42			-0.013	-2.32			-0.012	-2.76
- Birth year 1949	-0.010	-3.02			-0.009	-2.74			-0.010	-3.56
- Birth year 1948	-0.012	-3.76			-0.012	-3.50			-0.010	-3.66
- Birth year 1947	-0.017	-4.27			-0.017	-3.93			-0.013	-3.46
Intercept	0.025	5.35	-0.111	-4.03	-0.101	-3.69	-0.091	-3.82	-0.081	-3.51
Adjusted R^2	0.011		0.029		0.039		0.162		0.169	
Number of groups	4,468		4,468		4,468		4,468		4,468	

Table 7
Preference Parameters and Household Characteristics

This table reports regressions of the log of risk aversion (columns 1 and 2), the elasticity of intertemporal substitution (columns 3 and 4), and the rate of time preference (columns 5 and 6) on dummy variables for educational attainment and measures of labor income volatility.

	Log of RRA (1)		Log of RRA (2)		Log of EIS (3)		Log of EIS (4)		TPR (5)		TPR (6)	
	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat
Education dummies												
- High school	-0.026	-0.69	-0.009	-5.10	-0.045	-1.32	-0.064	-5.73	0.001	0.24	-0.001	-0.52
- Post-high school	0.005	0.14	-0.021	-6.54	-0.033	-1.96	-0.186	-5.25	-0.008	-2.89	-0.008	-2.38
Labor income volatility												
- Total volatility	-2.313	-4.94			0.795	1.21			0.004	0.08		
- Systematic income shocks			-7.418	-8.00			50.653	5.93			1.743	1.81
- Idiosyncratic permanent shocks			-7.513	-52.76			-1.560	-1.23			0.536	3.75
- Idiosyncratic transitory shocks			-0.299	-4.61			-2.896	-4.39			-0.226	-3.06
Constant	1.971	31.62	2.437	110.56	-0.244	-3.53	-1.198	-5.83	0.027	3.42	-0.036	-1.72
Cohort fixed effects	Yes		Yes		Yes		Yes		Yes		Yes	
Adjusted R^2	0.363		0.485		0.009		0.010		0.021		0.027	
Number of groups	5,434		5,434		5,434		5,434		5,434		5,434	

Figure 1
Histogram of the Initial Risky Share and Wealth-Income Ratio

This figure reports the cross-sectional distribution of the initial risky share (Panel A) and initial wealth-income ratio across household groups.

A. Initial Risky Share

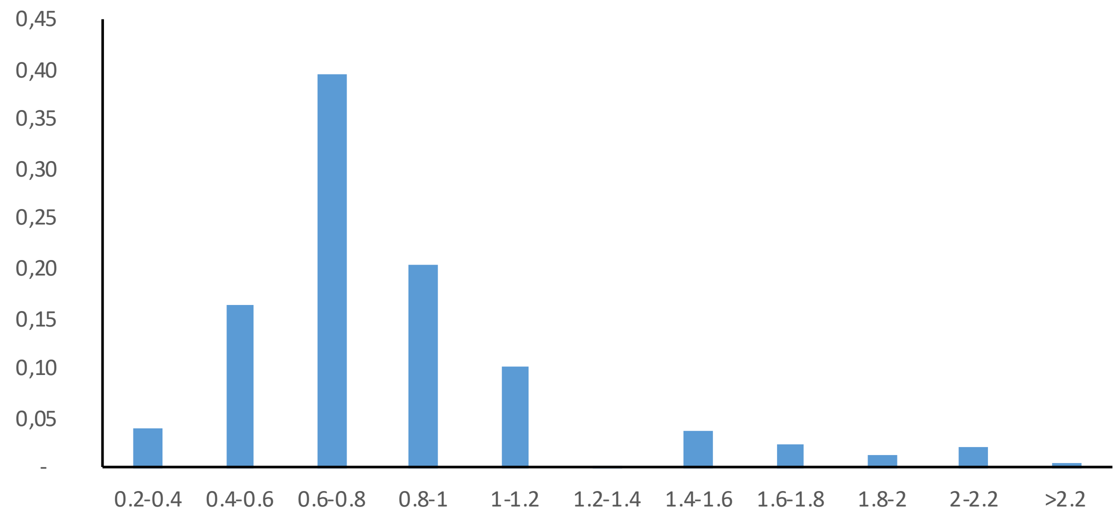

B. Initial Wealth-Income Ratio

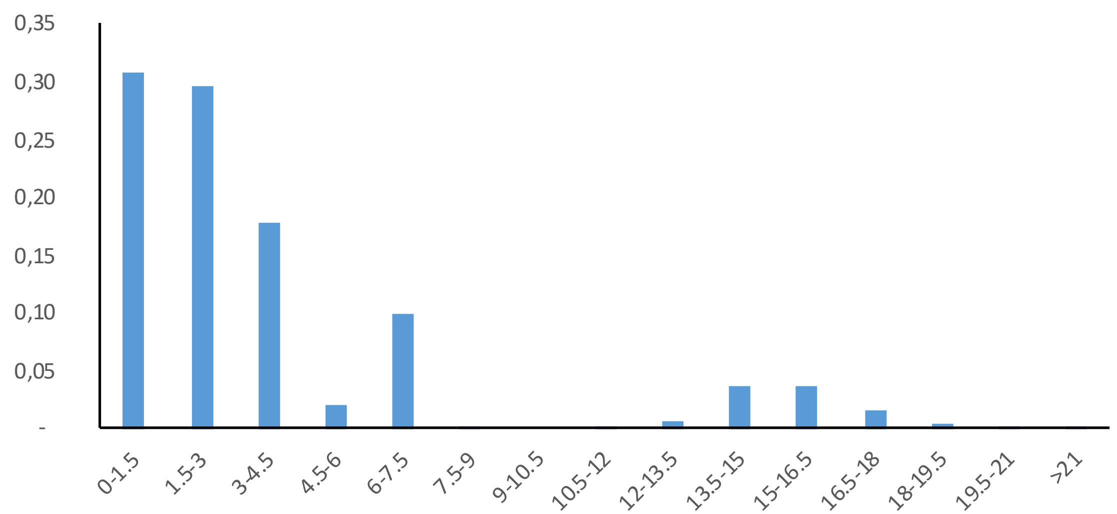

Figure 2
Age-Income Profiles

This figure reports the age-income profiles of households with (1) no high school degree, (2) a high school degree, and (3) post-high school education. For each education group and year, we report both the average empirical value and its fitted value.

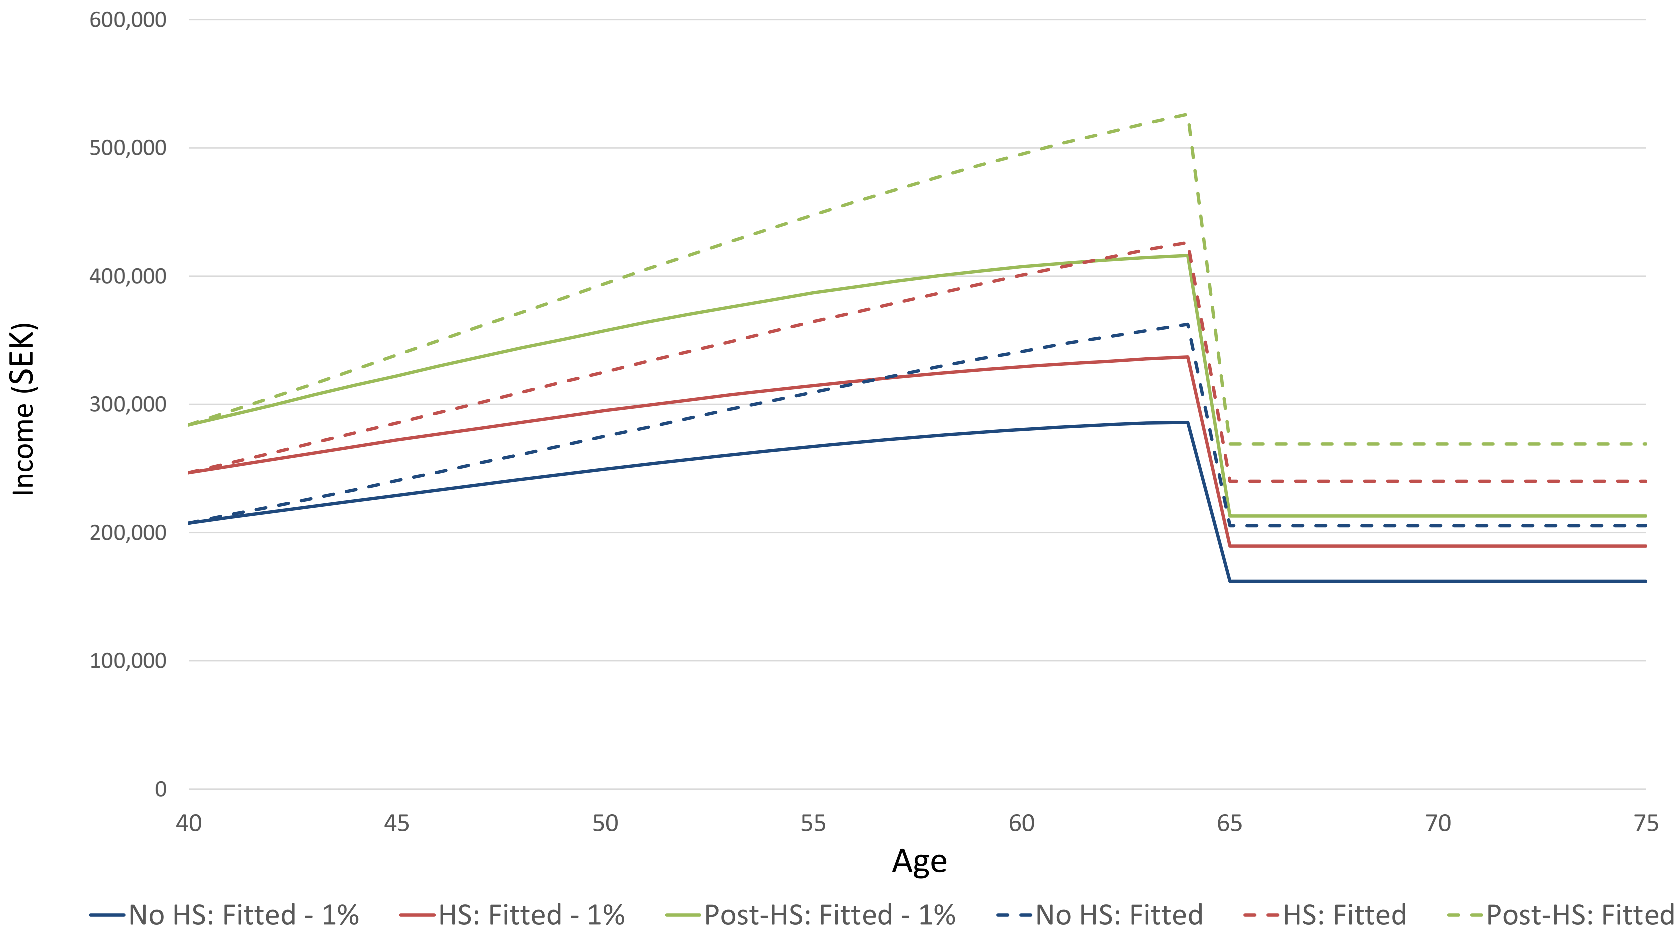

Figure 3
Cross-Sectional Distribution of Preference Parameters

This figure reports the cross-sectional distribution of the coefficient of relative risk aversion α (Panel A), the elasticity of intertemporal substitution ψ (Panel B), and the discount factor δ (Panel C). In Panel B and C the first and the last bins are omitted. In Panel B the omitted bins contain 5.4% and 25.6% of the observations, respectively. In panel C, 6.3% and 0.6%. All preference parameters are obtained from the indirect inference procedure developed in the main text.

A. Risk Aversion

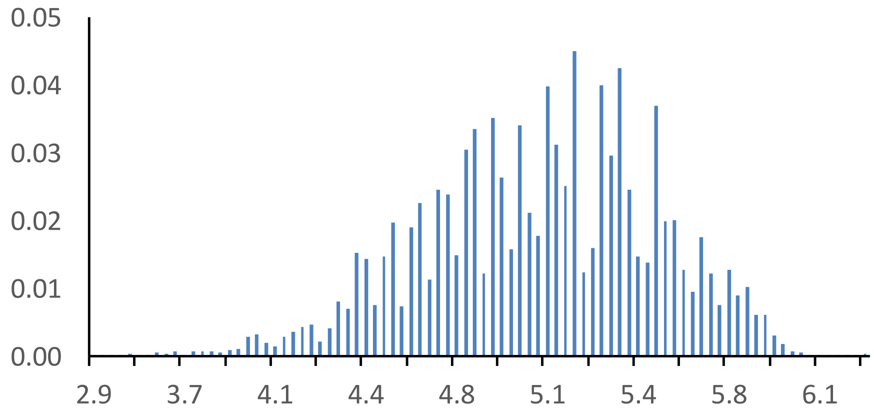

B. Elasticity of Intertemporal Substitution

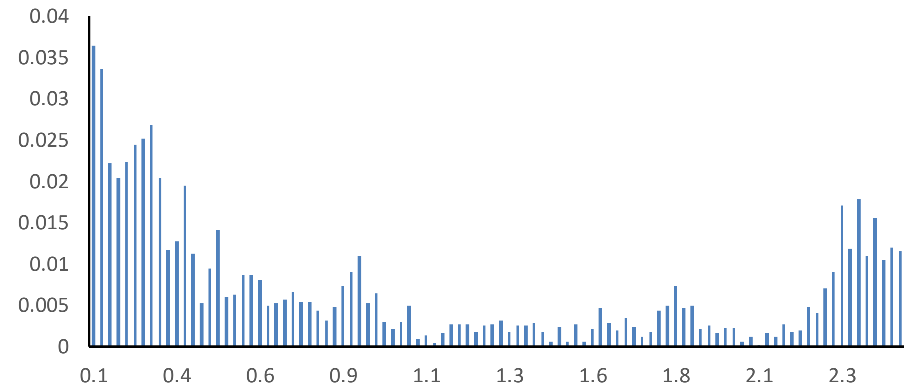

C. Time Preference Rate