

The Cross-Section of Household Preferences

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

April 2021

*Calvet: Department of Finance, EDHEC Business School, 16 rue du Quatre-Septembre, 75002 Paris, 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 earlier versions of this paper from Stijn van Nieuwerburgh, Jessica Wachter, Stanley Zin, and seminar participants at EDHEC Business School, ENSAE-CREST, the University of Michigan, Stanford University, the 2017 American Economic Association meeting, the Vienna Graduate School of Finance, Arizona State University, Harvard University, the UCLA Anderson Fink Center Conference on Financial Markets, the 2019 FMND Workshop, Imperial College London, the 2019 NBER Asset Pricing Summer Institute, and the 2021 London Seminar on Micro and Macro Implications of Household Behaviour and Financial Decision-Making. We thank the Sloan Foundation for financial support to John Campbell, and Azar Aliyev, Nikolay Antonov, and Huseyin Aytug for able and dedicated research assistance.

Abstract

This paper estimates the cross-sectional distribution of Epstein-Zin preference parameters in a large administrative panel of Swedish households. We consider a life-cycle model of saving and portfolio choice that incorporates risky labor income, safe and risky financial assets inside and outside retirement accounts, and real estate. We study middle-aged stockowning 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 0.47 around a mean of 5.24 and median of 5.30) and considerable heterogeneity in the time preference rate (standard deviation 6.0% around a mean of 6.2% and median of 4.1%) and elasticity of intertemporal substitution (standard deviation 0.96 around a mean of 0.99 and median of 0.42). Risk aversion and the EIS are almost cross-sectionally uncorrelated, in contrast with the strong negative correlation that we would find if households had power utility with heterogeneous risk aversion. The TPR is weakly negatively correlated with both the other parameters. We estimate lower risk aversion for households with riskier labor income and higher levels of education, and a higher TPR and lower EIS for households who enter our sample with low initial wealth.

1 Introduction

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

Modern financial theory distinguishes at least three parameters that govern savings behavior and financial decisions: the time preference rate (TPR), the coefficient of relative risk aversion (RRA), and the elasticity of intertemporal substitution (EIS). The canonical model of Epstein and Zin (1989, 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. In our base case implementation we assume that all agents have common 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. As a robustness check we do consider a simple form of heterogeneity in beliefs about risky asset returns.

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

We allow 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). These 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 time series of wealth and portfolio choice between ages 40 and 60, taking as given the initial level of wealth at the start of each year and the risky asset returns realized during each year. Since we do not observe decisions late in life, we cannot accurately account for bequest motives and instead capture the desire to leave a bequest as a lower rate of time preference.

We measure not only liquid financial wealth, but also defined-contribution retirement assets as well as household entitlements to defined-benefit pension income. However, we confine attention to households who hold some risky financial assets outside 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 (Castaneda, Diaz-Gimenez and Rios-Rull 2003, De Nardi 2004, Gomes and Michaelides 2005, Hubbard, Skinner and Zeldes 1984).

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

Our model assumes that expected returns on safe and risky assets are constant over time, so we cannot exploit time-variation in the riskless interest rate or the expected risky return to identify the EIS in the manner of Hall (1988) or Yogo (2004). However, the model incorporates time-variation in background risks. Households in the model have a target level of financial wealth that serves both as a buffer stock to smooth consumption in the face of random income variation, and as a means of financing retirement. Households save when their wealth is below the target, and they do so more aggressively when the EIS is high. A related phenomenon is that households with a high level of financial wealth relative to human capital

invest more conservatively, which reduces the expected rate of return on their portfolio. In addition, as households age their mortality rates increase, and this alters the effective rate of time discounting. For all these reasons we can identify the EIS from time-variation in the growth rate of wealth within each household group. This identification strategy that exploits the time pattern of wealth accumulation is a methodological contribution of our paper.

Our main empirical findings are as follows.

First, our data analysis uncovers considerable heterogeneity in wealth accumulation and portfolio composition across the Swedish population. Average wealth-income ratios vary intuitively with the riskiness of income and the level of education: among groups with the lowest income risk and lowest education, the average wealth-income ratio is 3.6, whereas it is 6.0 among groups with the highest income risk and highest level of education. There is also considerable heterogeneity unrelated to these variables: within each category of income risk and education, the standard deviation of the wealth-income ratio exceeds 3.0 across the groups we consider. Average risky shares vary little with income risk or education, averaging between 0.65 and 0.69 across all income risk and education categories, but within each category the standard deviation of the risky share across groups is substantial, ranging from 0.19 to 0.25.

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

Third, we estimate that the average level of risk aversion across Swedish households is 5.24, close to the median level of 5.30. The cross-sectional standard deviation of risk aversion is 0.47. The ratio of standard deviation to mean is smaller for risk aversion than for the risky portfolio share because households have human capital in our model; if there were no human capital, the usual formula for risky portfolio choice under homogeneous beliefs implies that risk aversion and the risky share must be proportional to one another and therefore must have the same ratio of standard deviation to mean.

Fourth, we estimate the average TPR to be 6.18%, well above the median value of 4.08% because the distribution of the TPR is right-skewed and dispersed, with a standard deviation of 6.03%. We estimate negative rates of time preference for 6.5% of Swedish households; these low values of time preference may in part reflect the existence of bequest motives that are omitted from our life-cycle model.

Fifth, we estimate the average EIS to be very close to one at 0.99. The median EIS however is well below one at 0.42, reflecting a right-skewed distribution with some households estimated to have EIS above two. The cross-sectional standard deviation of the EIS is substantial at 0.96. 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 60% of households, and even less than the reciprocal of risk aversion for 35% of households, while it is greater than one for 40% 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.

Sixth, our estimates of preference parameters are weakly cross-sectionally correlated. The correlation between risk aversion and the EIS is weakly positive at 0.11, in contrast with the perfect negative correlation between log risk aversion and the log EIS that we would find if all households had power utility with heterogeneous coefficients. The TPR is negatively correlated with both risk aversion (-0.28) and the EIS (-0.29), implying a tendency for patient people to be both cautious and willing to substitute intertemporally. The weak correlations across preference parameters imply that Swedish household behavior is heterogeneous in multiple dimensions, not just one. A single source of heterogeneity omitted from our model, such as heterogeneity in household beliefs about the equity premium, would not generate this multidimensional heterogeneity in our preference estimates.

Seventh, there are some interesting correlations between our parameter estimates, the moments we use for estimation, and exogenous characteristics of households. Risk aversion is lower for households working in risky sectors: households in these sectors invest somewhat more conservatively, but not as conservatively as they would do if they understood their income risk and had the same risk aversion as households working in safer sectors. Risk aversion is also lower for households with higher education: these households tend to have higher wealth-income ratios, but similar risky shares. Both these patterns are consistent with the hypothesis that risk-tolerant households choose to acquire education and select risky occupations, but they could also result from households' failure to understand the portfolio choice implications of their income risk exposure and wealth accumulation.

The TPR is negatively correlated (-0.30) with the initial wealth-income ratio of each household group, and positively correlated (0.35) 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 TPR in our data is a tendency to accumulate retirement savings later in life, catching up belatedly with those who saved earlier in life. Conversely, the EIS is positively correlated (0.32) with the initial wealth-income ratio and negatively correlated (-0.18) with the average growth rate of the wealth-income ratio, suggesting that households with a high EIS save early in life to reach a target wealth-income ratio, while households with a low EIS save more gradually over time.

Our paper is related to a large literature on household portfolio choice over the life cycle, including Campbell and Viceira (2002), Ameriks and Zeldes (2004), Cocco, Gomes, and Maenhout (2005), and Fagereng, Gottlieb, and Guiso (2017). Guiso and Sodini (2013) provide a comprehensive survey. Several recent papers have highlighted the role of target date funds as a low-cost way to implement a declining age profile in the risky portfolio share, but have also explained ways in which these funds fall short of the theoretical optimum (Dahlquist, Setty, and Vestman 2018, Parker, Schoar, and Sun 2020). Our use of comprehensive administrative data from Sweden follows a series of papers by Calvet, Campbell, and Sodini (2007, 2009a, 2009b), Calvet and Sodini (2014), Betermier, Calvet, and Sodini (2017), and Bach, Calvet, and Sodini (2020).

A smaller literature on heterogeneity in portfolio choice has recently tried to relate observed household behavior to underlying heterogeneity in preferences and beliefs (Meeuwis et al 2018, Giglio et al 2019). Relative to this literature, we observe more households over a longer period of time and have more complete data on wealth and portfolio allocation, but we lack data on potentially heterogeneous beliefs.

The organization of the paper is as follows. Section 2 explains how we measure household wealth and its allocation to safe and risky assets, describes the creation of household groups, and reports summary statistics 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. Subsection 5.1 summarizes the cross-sectional distribution of preference parameter estimates, and subsection 5.2 relates these estimates to observable household characteristics. Subsection 5.3 discusses parameter uncertainty and subsection 5.4 describes the fit of our model. Subsection 5.5 presents a simple Monte Carlo analysis of our estimation method. Subsection 5.6 demonstrates robustness of our results to some variations of our methodology, in particular allowing for a simple form of heterogeneity in beliefs about the return on risky assets. Section 6 concludes. An online appendix provides additional details about our empirical analysis and estimation technique.

2 Measuring Household Wealth and Asset Allocation

Our empirical analysis is based on the Swedish Wealth and Income Registry. This high-quality administrative panel provides the income, wealth, and debt of every Swedish resident. Income data are available at the individual level from 1983 and can be aggregated to the household level from 1991. Wealth data are available from 1999. 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.

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

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

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 (2020).

Debt is the sum of all liabilities of the household, including mortgages and other personal liabilities held outside private businesses.¹ Since Swedish household debt is normally floating-rate, we treat debt as equivalent to a negative cash position but paying a borrowing rate that is higher than the safe lending rate.

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

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

2.2 Pension Imputation

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

The *state pension* system requires each worker in Sweden to contribute 18.5% of their pension qualifying income: 16% to the pay-as-you-go defined benefit (DB) system and the remaining 2.5% to a defined contribution (DC) system called *premiepension* system or PPM. 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 a 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

¹Because 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 observe 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.

income back to the age of 25 by using real per-capita GDP growth and inflation before 1983. We then use the state DB payout rules to impute state DB pension payments for each individual retiring during our sample period.

Occupational pensions were introduced to Sweden in 1991. They are regulated for the vast majority of Swedish residents by four collective agreements applying to blue-collar private-sector workers, white-collar private-sector workers, central government employees, and local government employees. These agreements specify workers' monthly pension contributions, the fraction directed to DB and 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. We can do this accurately because data on pension qualifying income are available from 1991 (the year occupational pensions were introduced) and because the DB collective pension payouts are a function of at most the last 7 years of pension qualifying income during working life.

Defined contribution *private pensions* have existed in Sweden for a long time but our dataset provides us with individual private pension contributions from 1991.² We assume that these contributions are invested in the same way as occupational and state DC contributions. We follow Bach, Calvet and Sodini (2020) and allocate 58% of the aggregate stock of private pension wealth in 1991 to workers.³ Across workers, we allocate pension wealth proportionately to their private pension contributions in 1991.

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 for the state DC system, we follow the investment policy and cost of the system's default fund and assume that equity contributions are invested in the MSCI equity world index, without currency hedging, and are subject to a fee of 15 basis points. The equity share in the state DC system mirrors the allocation rules of the system's default fund: a 130% levered position in the world index up

²Our dataset contains exact information on private pension contributions from 1994, and only reports a capped version of the variable from 1991. We impute full contributions from 1991 to 1993 taking into account both age effects and individual savings propensities in subsequent years, as explained in the appendix.

³This allocation is chosen to satisfy the condition that imputed pension wealth should be roughly the same just before and just after retirement.

to the age of 55, which is then gradually rebalanced with age to an increasingly conservative portfolio. For occupational and private DC pensions, we assume that equity contributions are invested in the unhedged MSCI equity world index, subject to 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 occupational and private DC retirement portfolio is rebalanced with age following the representative age pattern of life-cycle funds available in Sweden during our sample period. We assume that all DC wealth not invested in equities is invested in cash.

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

2.3 Household Asset Allocation

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

At the first stage, we treat liquid holdings of individual stocks, equity mutual funds, and hedge funds as diversified holdings of the MSCI world equity index.⁴ We treat liquid holdings of balanced funds and bond funds as portfolios of cash and stocks, with the share in

⁴This reflects the global exposure of Swedish equity portfolios documented by Calvet, Campbell, and Sodini (2007). It abstracts from underdiversification which is documented in the same paper. The impact of underdiversification in liquid wealth is reduced when one takes account of diversified DC retirement wealth as we do in this paper.

stocks 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 as described in section 2.2.

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

$$R_{h,t+1}^e = \omega_{S,t}^h R_{S,t+1}^e + \omega_{RE,t}^h R_{RE,t+1}^e + \omega_{DC,t}^h R_{DC,t+1}^e + (1 - \omega_{S,t}^h - \omega_{RE,t}^h - \omega_{DC,t}^h) R_{D,t+1}^e, \quad (2)$$

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

The second stage of our analysis is to calculate the variance of the excess return on household net wealth. Since the borrowing rate is deterministic, we only need to consider the vector $\omega_t^h = (\omega_{S,t}^h, \omega_{RE,t}^h, \omega_{DC,t}^h)'$, and we can calculate the variance of $R_{h,t+1}^e$ as

$$\sigma^2(R_{h,t+1}^e) = \omega_{h,t}' \Sigma \omega_{h,t}, \quad (3)$$

where Σ is the variance-covariance matrix of $R_{t+1}^e = (R_{S,t+1}^e, R_{RE,t+1}^e, R_{DC,t+1}^e)'$.

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

In the third stage of our analysis, we define a numeraire asset, the aggregate Swedish portfolio of cash, stocks, and real estate scaled to have the same volatility as the after-tax global equity index return:

$$R_{N,t+1}^e = (1 + L)(\omega_{agg,t}' R_{t+1}^e). \quad (4)$$

Here $R_{N,t+1}^e$ 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

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

all Swedish households in our sample. The scaling factor L is chosen so that the volatility of $R_{N,t+1}^e$ 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:

$$\alpha_{h,t} = \sigma(R_{h,t+1}^e) / \sigma(R_{N,t+1}^e), \quad (5)$$

where household portfolio volatility is computed from equation (3). 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 Swedish Income Registry data set contains data on the entire population of Sweden, but our focus is on middle-aged households aged between 40 and 60 during our sample period from 1999 to 2007. There are 7.7 million household-year observations on the 13 cohorts born between 1947 and 1959, but we impose several filters on the panel. We exclude 2.3 million observations on households that do not participate in risky financial markets outside retirement accounts. We also exclude households in which the head is a student, working in the agricultural sector, retired before the start of our sample, missing information on education or sector of employment, or missing data in any of the years between 1999 and 2007. We exclude households that change their sector of employment during our sample in such a way as to alter the level of income volatility they are exposed to. To limit the impact of the wealthiest households (for whom our measurement procedures may be less adequate), we also exclude households whose financial wealth is above the 99th percentile of the wealth distribution in 1999. These filters exclude another 2.8 million observations, leaving us with a balanced panel containing 2.6 million household-year observations and 291,488 households.

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

⁶The estimation of income volatility is described in section 3, and details of the sectoral income risk

We subdivide each of these categories using a two-way sort by deciles of the initial wealth-income ratio and the initial risky share. We use the lowest two and highest two deciles and the middle three quintiles, giving us a 7×7 grid of 49 bins for the initial wealth-income ratio and the risky share.⁷ We further subdivide by 13 cohorts to create $5733 = 9 \times 49 \times 13$ 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 4151 groups.

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

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

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

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

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

classification are provided in table IA.2 in the online appendix.

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

negative effect on the risky portfolio share.

The category averages in the top panel of Table 1 conceal a great deal of dispersion across disaggregated groups of households. This is shown by the bottom panel of Table 1, which reports the standard deviations of the wealth-income ratio and the risky portfolio share across groups in each of the nine categories of education and sectoral income risk. The standard deviations of the risky share are consistently in the range 19–25%, while the standard deviations of the wealth-income ratio are in the range 3.0–3.9. Across all 4151 groups, the average wealth-income ratio has a mean of 4.7 with a standard deviation of 3.6, while the average risky share has a mean of 67% with a standard deviation of 22%.⁸

Figure 1 plots the cross-sectional distributions of the time-series average wealth-income ratio and risky share across all groups, size-weighting the groups to recover the underlying household distribution. The wealth-income ratios have a strongly right-skewed distribution, with many households having only a year or two of income accumulated, and a few having well over a decade of income. The risky shares also have a right-skewed cross-sectional distribution including some probability mass above 1 (corresponding to a portfolio volatility above 13.3%).

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

⁸The aggregation of households into groups naturally reduces the dispersion that is visible in the household-level data. Across all individual households in our dataset the average wealth-income ratio is 6.5 with a standard deviation of 185. After winsorizing the underlying household distribution at the 99th percentile, the average falls to 4.6 and the standard deviation falls to 4.0. The average risky share across all individual households is 71% with a standard deviation of 43%.

3.1 Labor Income Process

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

$$\log(Y_{h,t}) = a_c + b'x_{h,t} + \nu_{h,t} + \varepsilon_{h,t}, \quad (6)$$

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

We enrich the model above by distinguishing between shocks that are common to all households in a group and shocks that are specific to each household in the group. To simplify notation, we do not write an explicit group index but write group-level shocks using a single time index.

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, $z_{h,t}$:

$$\nu_{h,t} = \xi_t + z_{h,t}. \quad (7)$$

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

$$\xi_t = \xi_{t-1} + u_t, \quad (8)$$

$$z_{h,t} = z_{h,t-1} + w_{h,t}. \quad (9)$$

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

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

$$(u_t, w_{h,t}, \varepsilon_{h,t})' \sim \mathcal{N}(0, \Omega) \quad (10)$$

where Ω is the diagonal matrix with diagonal elements σ_u^2 , σ_w^2 , and σ_ε^2 .

We estimate the income process from consecutive observations of household yearly income data over the period 1992 to 2007, excluding the first and last year of labor income to avoid measuring annual income earned over less than 12 months.⁹ We consider the total income

⁹In each year, we winsorize non-financial real disposable income to a minimum level of 1000 kronor or about \$150. We also winsorize the pooled data from above at the 0.01% level to take care of extreme outliers at the top of the income distribution.

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 older than 19, and (ii) retired households that are at least 65.

For active households younger than 65, we estimate b by running pooled regressions of equation (6) for each of the three education levels. The vector of explanatory variables $x_{h,t}$ includes age dummies, which we then regress on a third-degree polynomial in age and use the fitted third-degree polynomial in our life-cycle model.

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

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

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 steeper than profiles estimated in the US.¹⁰

¹⁰Dahlquist, Setty, and Vestman (2018) estimate income profiles for Sweden with a pronounced hump shape and lower income towards the end of working life. They use a model that excludes cohort effects, thereby estimating the age-income profile in part by comparing the incomes of households of different ages at a point in time. This procedure is biased if different cohorts receive different lifetime income on average. We obtain similar estimates when we exclude cohort effects from our model of income.

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 the nine categories defined by three 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, particularly those with higher education, face higher transitory income risk and lower permanent income risk than less educated households. This pattern is consistent with Low, Meghir, and Pistaferri (2010), but it contrasts with earlier studies showing the opposite pattern in the United States. The explanation is likely due to the fact that in Sweden, uneducated workers face lower unemployment risk and lower effects of unemployment on income than in many other countries, while educated workers face relatively high income losses when they 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.¹¹

We have already noted in discussing Table 1 that average wealth-income ratios tend to be higher in sectors with riskier income. This pattern is intuitive given that labor income risk encourages precautionary saving. However, there is little tendency for risky portfolio shares to be lower in sectors with riskier income. Table 3 further explores these effects by

¹¹See Brown, Fang, and Gomes (2012) for related research on the relation between education and income risk.

regressing the average wealth-income ratio and risky share on age, total income volatility, and dummies for high school and post-high school education. All regressions also include year fixed effects.¹²

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

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

The negative effects of age and the wealth-income ratio on the risky share are consistent with the predictions of a simple static model in which labor income is safe and tradable, so that human capital is an implicit cash holding that tilts the composition of the financial portfolio towards risky assets (Bodie, Merton, and Samuelson 1992, Campbell and Viceira 2002 p.163, Campbell 2018 p.309). Older households have fewer earning years remaining so their human capital is lower; and at any given age, 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.¹³

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

¹²We 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.

¹³At first glance the negative effect of the wealth-income ratio on the risky share may appear to contradict evidence that wealthier individuals take more financial risk (Carroll 2000, 2002, Wachter and Yogo 2010, Calvet and Sodini 2014). The discrepancy is likely due to several factors. Our sample excludes non-participants in risky financial markets and the wealthiest 1% of Swedish households; we measure the risky portfolio share taking account of housing and leverage through mortgage borrowing; and we predict the risky share using the wealth-income ratio rather than the absolute level of wealth.

3.3 Life-Cycle Model

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

Households have finite lives and Epstein-Zin utility over a single consumption good. The utility function V_t is specified by the coefficient of relative risk aversion (RRA) γ , the time discount factor δ or equivalently the time preference rate (TPR) $-\log(\delta)$, and the elasticity of intertemporal substitution (EIS) ψ . V_t satisfies the recursion

$$V_t = \left[C_t^{1-1/\psi} + \delta \left(\mathbb{E}_t p_{t,t+1} V_{t+1}^{1-\gamma} \right)^{(1-1/\psi)/(1-\gamma)} \right]^{\frac{1}{1-1/\psi}}, \quad (11)$$

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 , δ , and ψ all vary across households but we suppress the household index h in equation (11) 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. For this reason, we focus on the stage of the life-cycle during which households have substantial retirement saving and initialize our model at age 40. 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 preference rate can be viewed as having a downward bias due to the absence of a bequest motive in the model.

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

Consistent with the discussion in Section 2, 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. In each period

we recalibrate beginning-of-period wealth to the level observed in the data and use the model to predict end-of-period wealth.

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

$$W_{h,t+1} = (R_f + \alpha_{h,t}R_{N,t+1}^e)(W_{h,t} + Y_{h,t} - C_{h,t}), \quad (12)$$

where $R_{N,t+1}^e$ is the return on the composite numeraire asset in excess of the risk-free rate R_f . The excess return $R_{N,t+1}^e$ 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.¹⁴

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 1984–2007. Putting these assumptions together, we assume a Sharpe ratio of 0.26. In section 5.6 we discuss robustness of our results to assuming a higher Sharpe ratio of 0.40 or a lower Sharpe ratio of 0.15.

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

¹⁴Our 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.

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 education-sector categories, 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 describes our procedure for estimating household preference parameters. We explain why parameter identification for Epstein-Zin preferences is challenging; we motivate our identification strategy using simulations of our model; and finally we present our indirect inference estimator and discuss its asymptotic properties as the number of households in a group increases.

4.1 The Identification Challenge

Our goal is to identify three separate preference parameters: the coefficient of relative risk aversion (RRA) γ , the time discount factor δ or equivalently the time preference rate (TPR) $-\log(\delta)$, and the elasticity of intertemporal substitution (EIS) ψ . It is a nontrivial task to identify these three parameters, and particularly to separately identify the TPR and the EIS, using data on portfolio risk and wealth accumulation.

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

$$1 = \text{E}_t \left[\tilde{\delta}_{t+1} \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{1}{\psi}} \left(\frac{V_{t+1}}{\mu(V_{t+1})} \right)^{\frac{1}{\psi}-\gamma} R_{t+1}^P \right] \quad (13)$$

where $\tilde{\delta}_{t+1} = \delta p_{t,t+1}$, $R_{t+1}^P = R_f + \alpha R_{t+1}^e$, and $\mu(V_{t+1})$ denotes the certainty equivalent of V_{t+1} . This Euler equation holds with equality even though our model has borrowing constraints, because with labor income risk and a utility function that satisfies $u'(0) = \infty$ the agent will always choose to hold some financial assets.¹⁵

Taking logs of both sides and making the usual assumption of conditional joint lognormality, we obtain

$$\begin{aligned} 0 = & \log(\tilde{\delta}_{t+1}) - \frac{1}{\psi} \mathbf{E}_t g_{t+1} + \left(\frac{1}{\psi} - \gamma \right) \mathbf{E}_t \tilde{v}_{t+1} + \mathbf{E}_t r_{t+1}^P + \frac{1}{2\psi^2} \sigma_{g,t}^2 \\ & + \frac{1}{2} \left(\frac{1}{\psi} - \gamma \right)^2 \sigma_{\tilde{v},t}^2 + \frac{1}{2} \sigma_{r,t}^2 + \frac{1}{\psi} \left(\frac{1}{\psi} - \gamma \right) \sigma_{g\tilde{v},t} + \left(\frac{1}{\psi} - \gamma \right) \sigma_{\tilde{v}r,t} + \frac{1}{\psi} \sigma_{gr,t}, \end{aligned} \quad (14)$$

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

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

$$\begin{aligned} \mathbf{E}_t g_{t+1} = & \psi(\mathbf{E}_t r_{t+1}^P - \log(\tilde{\delta}_{t+1})) + (1 - \gamma\psi) \mathbf{E}_t \tilde{v}_{t+1} \\ & + \frac{1}{2\psi} \sigma_{g,t}^2 + \frac{\psi}{2} \left[\left(\frac{1}{\psi} - \gamma \right)^2 \sigma_{\tilde{v},t}^2 + \sigma_{r,t}^2 + \left(\frac{1}{\psi} - \gamma \right) \sigma_{\tilde{v}r,t} \right] \\ & + \left(\frac{1}{\psi} - \gamma \right) \sigma_{g\tilde{v},t} + \sigma_{gr,t}. \end{aligned} \quad (15)$$

Equation (15) highlights the identification problem. If the expected portfolio return, the time discount factor, and the conditional variances are constant over time, then the expected consumption growth rate $\mathbf{E}_t g_{t+1}$ is constant and for any value of ψ there is a corresponding time discount factor δ that delivers the same level of $\mathbf{E}_t g_{t+1}$. Without additional restrictions on δ or ψ these two parameters cannot be separately identified, as shown by Kocherlakota (1990) and Svensson (1989).

Equation (15) also shows three ways in which the problem can be resolved. First, the variance terms may change over time, and generally do so in a model with risky labor income and wealth accumulation. Changes in variances tend to be quantitatively more important when wealth is rapidly being accumulated, but less so when retirement savings are already adequate (Gomes and Michaelides 2005). Second, there may be time variation in the expected portfolio return. Even though our model has no exogenous variation in expected returns for any individual asset, we have endogenous variation driven by changes

¹⁵Our model also has short-sales constraints on risky asset holdings, but these do not bind for the middle-aged households we are considering.

in the agent’s optimal portfolio as financial wealth accumulates and future labor income declines. This effect also tends to be important when wealth is being accumulated and is low relative to income. Third, there may be time variation in the effective time discount factor $\tilde{\delta}_{t+1} = \delta p_{t,t+1}$. Even though the underlying time preference rate is constant in our model, $\tilde{\delta}_{t+1}$ is adjusted for the survival probability $p_{t,t+1}$ which is a function of age. This source of variation becomes increasingly important at older ages.

All three sources of variation imply that the profile of the wealth-income ratio as a function of age is affected in different ways by the TPR and the EIS. Our identification strategy builds on this intuition.

4.2 Identification Strategy

We illustrate the promise of our identification strategy by running a series of regressions based on simulated data from the model. More specifically we regress the underlying preference parameters that were used to generate those simulations against a series of moments from the simulated data. The values for the preference parameters are the same grid points that we consider in our estimation: 1,848 combinations of 12 values of RRA ranging from 2 to 10, 11 values of the TPR ranging from -0.05 to 0.22, and 14 values of the EIS ranging from 0.1 to 2.5. For each of these 1,848 preference parameter combinations we consider all 4,151 combinations of the initial wealth-income ratio and other group characteristics that we observe in the data.¹⁶

4.2.1 Moments

To build intuition we consider four moments in our regressions. The first moment is the initial wealth-income ratio which determines the initial conditions in our simulations $((W/Y)_{i0})$. The second moment is the average risky share for household i over the 8 years in our sample:

$$\bar{\alpha}_i = \left(\frac{1}{8} \sum_{t=1}^8 \alpha_{it} \right), \tag{16}$$

¹⁶Simulated moments are obtained by averaging 10,000 simulations. In this exercise, unlike our empirical analysis, we use the observed wealth-income ratio only in the first year, and take wealth-income ratios in subsequent years from the simulated data rather than from the observed data.

which should provide strong identification of the risk aversion parameter. The third moment that we consider is the 8-year cumulative growth of the wealth-income ratio,

$$grWY_i = \left[\left(\frac{W}{Y} \right)_{i8} / \left(\frac{W}{Y} \right)_{i0} \right]. \quad (17)$$

Finally, to capture age variation in the rate of wealth accumulation, the fourth moment we consider is

$$convexWY_i = \frac{\frac{1}{2} \left[\frac{W}{Y}_{i1} + \frac{W}{Y}_{i9} \right]}{\frac{W}{Y}_{i5}} - 1. \quad (18)$$

This measures the convexity (curvature) of the wealth-income ratio as a function of age.¹⁷

As an alternative to these four moments, we also consider the average risky share and the average wealth-income ratio in every year, thus giving us a total of sixteen moments ($\{\alpha_{it}\}_{t=1,8}$ and $\{(W/Y)_{it}\}_{t=1,8}$). These are the moments we actually use in our estimation.

4.2.2 Regressions on simulated data

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

$$\gamma_i = k_\gamma^0 + k_\gamma^1 \bar{\alpha}_i + e_i. \quad (19)$$

Panel A of Table 4 reports the estimation results. Confirming that the average risky share is a very good moment for identifying the risk aversion parameter, the adjusted R^2 from this regression is 72.4%. 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 also depends on the initial wealth-income ratio.

In our second specification we add the initial wealth-income ratio as an explanatory variable, while in the third specification we add the other two moments ((17) and (18)); the adjusted R^2 statistics increase to 74.0% and 75.0%, respectively. Finally, in the last column we report results using all 16 moments from our estimation ($\{\alpha_{it}\}_{t=2,9}$ and $\{(W/Y)_{it}\}_{t=2,9}$).

¹⁷An 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.

The adjusted R^2 is now 80.0%, which we again re-emphasize as an impressive number given that these are simple linear regressions.

Having confirmed that risk aversion is well identified, we now turn our attention to the time preference rate. Estimation results are reported in Panel B of Table 4. In the first specification we regress the TPR on the cumulative growth rate of the wealth-income ratio and obtain an adjusted R^2 of 60.6%. Although lower than the values obtained in the risk aversion regression, this is again an impressive number considering that this is a simple univariate linear specification. When we add the initial wealth-income ratio the adjusted R^2 increases to 64.5% while the other two moments, convexity and the average risky share, add very little explanatory power.

Finally we consider the EIS. Motivated by the discussion in the previous section, we first consider a specification that includes only the cumulative growth rate and the convexity of the wealth-income ratio. The adjusted R^2 , reported in Panel C of table 4 is only 3.2%. When we add the initial wealth-income ratio and the initial risky share, the explanatory power is still only 5.6%. These low values are however obtained in the context of a simpler linear regression, and where we only have cumulative growth and convexity to capture the full path of the wealth-income profile. When we consider all sixteen moments ($\{\alpha_{it}\}_{t=1,8}$ and $\{(W/Y)_{it}\}_{t=1,8}$), so that we fully capture the evolution of the wealth-income ratio over time, the adjusted R^2 increases to 14.1%.

Panel C still imposes a linear relationship between the observed moments and the EIS. In the final panel of Table 4, Panel D, we relax the linearity assumption along one single dimension, by estimating regressions within seven different ranges of values of the initial wealth-income ratio. The values for the adjusted R^2 statistics are between 19.0% and 32.7%.

Finally we consider the fact that since the impact of the EIS on savings depends on the difference $E_t r_{t+1}^P - \log(\tilde{\delta}_{t+1})$, it will depend on the values of the other two preference parameters. Therefore, in the final set of regressions we also condition on a specific value of the time discount factor, choosing 0.97 as a specific illustration. The adjusted R^2 now lies between 30.5% and 82.9%, and it is monotonically decreasing in the initial wealth-income ratio. This relation with the wealth-income ratio is intuitive since two of the three mechanisms for identification of the EIS weaken as the level of wealth increases, so in the limit of extremely high wealth only the age-dependent mortality rate allows us to identify the EIS. Overall, our simulation results confirm that, conditional on having identified the other preference parameters, the EIS is also well identified in our framework.

4.3 Indirect Inference Estimator

4.3.1 Definition

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

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

$$\hat{\mu}_{1,t}^g = \frac{\sum_{h=1}^{N^g} W_{h,t}}{\sum_{h=1}^{N^g} Y_{h,t}}, \quad (20)$$

and (ii) the group's risky share:

$$\hat{\mu}_{2,t}^g = \frac{\sum_{h=1}^{N^g} \alpha_{h,t} W_{h,t}}{\sum_{h=1}^{N^g} W_{h,t}}. \quad (21)$$

We stack these auxiliary statistics into the *empirical auxiliary estimator*

$$\hat{\mu}^g = (\hat{\mu}_{1,1}^g, \dots, \hat{\mu}_{1,T}^g, \hat{\mu}_{2,1}^g, \dots, \hat{\mu}_{2,T}^g)'$$

By construction, $\hat{\mu}^g$ has $q = 16$ components.¹⁸ The auxiliary statistics $\hat{\mu}_{1,t}^g$ and $\hat{\mu}_{2,t}^g$ provide reliable measures of risk-taking and wealth accumulation based on group aggregates. We note that $\hat{\mu}_{1,t}^g$ and $\hat{\mu}_{2,t}^g$ 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.

The data-generating process is based on the policy functions of households with preference parameter vector θ , the return process, and the labor income process defined in earlier sections. As the number of households in the group goes to infinity, the empirical auxiliary

¹⁸We could potentially also include the risky share in the initial year (α_{i0}), since that is also an endogenous moment from the simulations. We exclude it in order to have an equal number of auxiliary statistics related to the wealth-income ratio and to the risky share.

estimator $\hat{\mu}^g$ converges to the *binding function* $\mu^g(\theta) \in \mathbb{R}^q$ with components

$$\mu_{1,t}^g(\theta) = \frac{\mathbf{E}_\theta^g(W_t)}{\mathbf{E}_\theta^g(Y_t)}, \quad (22)$$

$$\mu_{2,t}^g(\theta) = \frac{\mathbf{E}_\theta^g(\alpha_t W_t)}{\mathbf{E}_\theta(W_t)}, \quad (23)$$

where $\mathbf{E}_\theta^g(\cdot)$ denotes the cross-sectional mean of households in the group. The population expectations in (22) and (23) are computed under the assumption that all households in the group earn the riskfree rate R_f and the synthetic excess risky return $R_{N,t}^e$ on their risky asset holdings.

We generate the simulated auxiliary estimator as follows. Using the parameters from Table 2 as inputs, we solve the life-cycle model for the 4151 different household groups. For each group g and for each preference parameter θ , we compute the wealth-income ratio and the risky share predicted by the model for every year between 1999 and 2007.

For every group g and each year t , we simulate 10,000 individual paths/households taking as given past realizations of labor income and the group's *empirical* wealth-income ratio at the end of period $t-1$. We feed in year t 's empirical risky asset return and group-level income shock and, using the data-generating process described in earlier sections, we simulate the idiosyncratic permanent and transitory income shocks of each household in the group. We apply the policy function to obtain each household-level consumption and risky share during year t , as well as the wealth and risky share at the end of year t . Finally, we aggregate up households to obtain the group's simulated wealth-income ratio and wealth-weighted average risky share at the end of year t . The simulated moments are estimates of the wealth-income ratio and risky share predicted by the model conditional on the information available at the end of year $t-1$. The indirect inference estimator compares this one-year ahead prediction to its realization.

For every year t , we implement this procedure in five steps.

i. The starting point is an information set at the end of year $t-1$ containing: (a) the empirical wealth-income ratio of the group at the end of year $t-1$, and (b) the empirical group-level income shock in year t . Consistent with the life-cycle model, we assume that households have this much advance information about wages and hours.

ii. We simulate $S = 10,000$ households/paths in the group over year t . For each simulated unit $i \in \{1, \dots, S\}$, we sample wealth $\tilde{W}_{i,t-1}$ from the empirical distribution of household wealth at the end of period $t-1$, and we simulate labor income and permanent income $(\tilde{Y}_{i,t}, \tilde{Y}_{i,t}^P)$ in period t . Using the lifecycle model's policy functions $\alpha_t^*(\cdot)$ and $C_t^*(\cdot)$, we compute the

risky share and consumption

$$\begin{aligned}\tilde{\alpha}_{i,t-1} &= \alpha_t^*(\tilde{Y}_{i,t}, \tilde{W}_{i,t-1}, \tilde{Y}_{i,t}^P; \theta), \\ \tilde{C}_{i,t} &= C_t^*(\tilde{Y}_{i,t}, \tilde{W}_{i,t-1}, \tilde{Y}_{i,t}^P; \theta),\end{aligned}$$

of each simulated unit during year t . Consistent with the model, the simulated unit sets both quantities at the end of year $t - 1$ and keeps them constant throughout year t .

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

$$\hat{W}_{i,t} = (R_f + \tilde{\alpha}_{i,t-1} R_{p,t}^e)(\tilde{W}_{i,t-1} + \tilde{Y}_{i,t} - \tilde{C}_{i,t}), \quad (24)$$

as equation (12) implies. The prediction incorporates empirical financial returns in year t , assumed equal for all groups.

iv. This method produces the group's predicted wealth-income ratio:

$$\tilde{\mu}_{1,t}^g(\theta) = \frac{\sum_{i=1}^S \hat{W}_{i,t}}{\sum_{i=1}^S \tilde{Y}_{i,t}} \quad (25)$$

at the end of year t .

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

$$\tilde{\mu}_{2,t}^g(\theta) = \frac{\sum_{i=1}^S \tilde{\alpha}_{i,t} \tilde{W}_{i,t}}{\sum_{i=1}^S \tilde{W}_{i,t}} \quad (26)$$

at the end of year t .

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

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

$$\hat{\theta}^g = \arg \min_{\theta} [\tilde{\mu}_S^g(\theta) - \hat{\mu}^g]' \Omega [\tilde{\mu}_S^g(\theta) - \hat{\mu}^g], \quad (27)$$

where Ω 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.

In practice, we use a diagonal weighting matrix Ω that is common to all groups. Each diagonal element of Ω is a scale factor that converts the wealth-income ratios and risky shares into comparable units. Specifically, we let:

$$\Omega = \text{diag}(\omega_1, \dots, \omega_1, \omega_2, \dots, \omega_2), \quad (28)$$

where ω_i is the inverse of the square of the average corresponding auxiliary statistic across groups and time—that is, $\omega_1 = \left(\frac{1}{GT} \sum_{g=1}^G \sum_{t=1}^T \hat{\mu}_{1,t}^g\right)^{-2}$ and $\omega_2 = \left(\frac{1}{GT} \sum_{g=1}^G \sum_{t=1}^T \hat{\mu}_{2,t}^g\right)^{-2}$.

The matrix Ω normalizes all observations into comparable units. In practice, we obtain weights such that $(\omega_2/\omega_1)^{1/2} = 7.57$. This result is intuitive since the average risky share is around 0.5 and the average wealth to income ratio is around 3.5 across groups. The use of a common weighting matrix Ω implies that the objective function defining the indirect inference estimator in equation (27) is comparable across groups.

The numerical implementation of the estimation methodology proceeds as follows. For each group, we simulate the lifecycle model on a grid of preference parameters. The grid contains 12 values of the RRA ranging from 2 to 10, 11 values of the TPR ranging from -0.05 to 0.22, and 14 values of the EIS ranging from 0.1 to 2.5. For each of the 1,848 ($= 12 \times 11 \times 14$) parameter vectors θ on the grid, we calculate the value of the objective function $[\tilde{\mu}_S^g(\theta) - \hat{\mu}^g]' \Omega [\tilde{\mu}_S^g(\theta) - \hat{\mu}^g]$.

We next evaluate the objective function on a finer grid by interpolating the values obtained on the original grid. The finer grid contains 81 equally-spaced values of the RRA, 241 values of the EIS, and 251 values of the TPR with the same endpoints as the original grid, or a total of 4,899,771 ($= 81 \times 241 \times 251$) preference parameter vectors. To mitigate oscillation problems, we adopt the modified Akima interpolation method. We obtain the indirect inference estimator by determining the parameter vector on the finer grid that minimizes the objective function. This value may occasionally be slightly negative due to interpolation error. For inference purposes, such as calculations of the root mean squared error or the Jacobian matrix, we compute each of the 16 auxiliary statistics by a separate cubic spline interpolation. We refer the reader to the online appendix for further details.

4.3.2 Asymptotic Properties

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, as we now show.

As the group size N^g goes to infinity, the empirical auxiliary estimator $\hat{\mu}^g$ is asymptoti-

cally normal:

$$\sqrt{N^g} [\hat{\mu}^g - \mu^g(\theta)] \rightarrow^d \mathcal{N}(0, W_g). \quad (29)$$

This result follows from the delta method and the fact that the auxiliary statistics (20) and (21) can be interpreted as ratios of sample moments. We estimate the asymptotic variance covariance matrix of $\hat{\mu}^g$ by the jackknife estimator

$$\frac{\hat{W}_g}{N^g} = \frac{N^g - 1}{N^g} \sum_{i=1}^{N^g} (\hat{\mu}_{[i]}^J - \bar{\mu}^J)(\hat{\mu}_{[i]}^J - \bar{\mu}^J)', \quad (30)$$

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

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:

$$\sqrt{N^g} (\hat{\theta}^g - \theta^g) \rightarrow^d \mathcal{N}(0, V^g). \quad (31)$$

The asymptotic variance-covariance matrix is given by

$$V^g = (1 + s_g^{-1}) (D_g \Omega D_g')^{-1} D_g \Omega W_g \Omega D_g' (D_g \Omega D_g')^{-1}, \quad (32)$$

where the ratio $s_g = S/N_g$ accounts for simulation noise and $(D_g)' = \partial \mu^g(\theta^g) / \partial \theta'$ is the Jacobian matrix of the binding function $\mu^g(\cdot)$ evaluated at the true parameter θ^g . In practice, we estimate the asymptotic variance-covariance matrix of V_g by

$$\hat{V}^g = (1 + s_g^{-1}) (\hat{D}_g \Omega \hat{D}_g')^{-1} \hat{D}_g \Omega \hat{W}_g \Omega \hat{D}_g' (\hat{D}_g \Omega \hat{D}_g')^{-1}, \quad (33)$$

where \hat{D}_g is a finite-difference approximation of D_g that we discuss in the appendix.

When the size of each group g is large, we could achieve efficient estimation by a two-stage estimation of the parameter vector. That is, we could set the second-stage weighting matrix equal to the inverse of the jackknife estimator: $\Omega^{(2)} = \hat{W}_g^{-1}$, and then solve the optimization problem (27) to obtain the second-stage efficient estimator of the preference parameter vector θ_g . This procedure generates an asymptotically efficient estimator as N^g goes to infinity.

Efficient estimation, however, is not feasible in our sample because most groups are too small to obtain a reliable estimator of W_g^{-1} . Indeed, the median group size is 63, while the symmetric matrix W_g contains 136 ($= 16 \times 17/2$) distinct elements. In a large fraction of groups, the jackknife estimator \hat{W}_g is close to singularity and its inverse contains extreme

elements. A related problem is that in many groups, the weighting matrix $\Omega^{(2)} = \hat{W}_g^{-1}$ assigns almost all the weight to the risky share, while the wealth-income ratio plays essentially no role in estimation. Efficient estimation is therefore unsatisfactory in our sample on statistical and economic grounds. These difficulties are consistent with the finite-sample inaccuracy of two-step generalized method of moments studied in Hwang and Sun (2018). For these reasons, we henceforth focus on the one-step estimation of θ^g based on the diagonal weighting matrix Ω defined in equation (28). This approach does not allow us to conduct global specification tests based on the value of the objective function (27). For this reason, we will focus instead on measures of goodness of fit based on root mean squared error or economic significance.

5 Empirical Results

This section reports empirical results. We first summarize the cross-sectional distribution of our preference parameter estimates and relate the estimates to observable household characteristics. We then discuss hypothesis tests based on asymptotic standard errors, and summarize the fit of our life-cycle model in relation to alternatives. Finally we conduct a simple Monte Carlo analysis of our estimation method and present some robustness analyses.

5.1 The Cross-Sectional Distribution of Preference Estimates

Tables 5 through 8 and Figures 3 and 4 summarize the cross-sectional distributions of our estimated preference parameters. Table 5 reports the cross-sectional means, medians, and standard deviations of the estimated parameters along with four summary statistics of the data we use to estimate these parameters: the average risky share, the initial wealth-income ratio, the cumulative growth rate of the wealth-income ratio, and the convexity of the wealth-income ratio. Table 6 reports the cross-sectional correlations of the estimated parameters and summary statistics. Table 7 reports multiple regressions of the estimated parameters onto the four summary statistics and control variables. Table 8 repeats this regression analysis using income risk and education as explanatory variables. Figure 3 plots the univariate cross-sectional distributions of each estimated preference parameter, and Figure 4 plots selected bivariate distributions. A number of interesting patterns are visible in these tables and figures.

Table 5 reports a moderate mean for relative risk aversion (RRA) of 5.24, close to the median estimate of 5.30 and in the middle of the range from 1 to 10 defined as reasonable by Mehra and Prescott (1985). The cross-sectional standard deviation of estimated RRA is

modest at 0.47, less than a tenth of the mean and median estimates.

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

The other two preference parameters have much greater cross-sectional variation, and both are strongly right-skewed. The median time preference rate (TPR) is 4.08%, considerably lower than the mean of 6.18%, and the cross-sectional standard deviation of TPR is 6.03%. Similarly, the median EIS is 0.42, considerably lower than the mean of 0.99, and the cross-sectional standard deviation of the EIS is 0.96. This cross-sectional standard deviation is twice as large for the EIS as for RRA in absolute terms, and over 20 times as large in proportional terms; this contrasts with the prediction of a power utility model, which would imply equal proportional standard deviations for RRA and the EIS since one parameter is the reciprocal of the other.

Figure 3 illustrates the cross-sectional distributions of the three preference parameters. The distribution of RRA is approximately normal although it has multiple modes that probably reflect the binning procedure by which we sorted households into groups by their initial risky shares. The distribution of the TPR has fat tails, particularly at the right but also to some extent at the left: 6.5% of our TPR estimates are actually negative. In interpreting this fact, it is important to keep in mind that our model omits any explicit bequest motives and therefore our TPR estimates will be lower to the extent that such motives are important. The distribution of the EIS is U-shaped, with probability mass concentrated below 1 and near the upper edge of our parameter space which we set to 2.5. We discuss these upper-edge estimates in greater detail in the next subsection. The EIS is less than one for 59.7% of households, and even less than the reciprocal of risk aversion for 34.7% of households, while it is greater than one for 40.3% of households.

Table 6 shows that our estimates of preference parameters are only weakly cross-sectionally correlated. Relative risk aversion and the EIS have a low positive correlation of 0.11, a finding that contrasts with the perfect negative correlation between the logs of RRA and the

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

Figure 4 illustrates bivariate relationships between estimated preference parameters using heat maps which indicate data density with a logarithmic color scale. The top panel of the figure reports the EIS on the horizontal axis and RRA on the vertical axis. The far left and far right bins on the horizontal axis contain households where the EIS is estimated at exactly its lower limit of 0.1 or its upper limit of 2.5 respectively, while all remaining bins are marked with the upper edge value. The heat map shows that the dispersion of RRA estimates is similar across all levels of the EIS, and it illustrates the concentration of EIS estimates at lower values and upper-edge values.

The bottom panel of Figure 4 reports the EIS on the horizontal axis and the TPR on the vertical axis. Here we see a very different pattern, with much greater dispersion of the TPR among households with a low EIS. This reflects the fact, shown in equation (15), that the EIS scales the impact of the TPR on planned consumption growth, so that much more extreme TPR values are consistent with reasonable consumption and savings behavior when the EIS is low than when it is high.

5.2 Preference Estimates and Household Characteristics

Our parameter estimates have interesting correlations with observable variables that help us understand what drives the estimates. The lower portion of Table 6 explores these correlations.

Looking first at the correlations among observables, the initial wealth-income ratio has a correlation of -0.52 with the average risky share and a correlation of -0.69 with the average growth rate of the wealth-income ratio. These correlations are consistent with the predictions of our life-cycle model that the risky share declines with the level of financial wealth in relation to human capital, and that households that enter the sample with low financial wealth have a strong motive to accumulate wealth to finance retirement. Unsurprisingly given these patterns, the average risky share and the average growth rate of the wealth-income ratio have a strong positive correlation of 0.62. The convexity of the wealth-income ratio is only weakly correlated with the other three observable variables.

Our estimate of RRA is negatively correlated (-0.41) with the average risky share, an

intuitive result that is also consistent with our identification analysis. RRA is also weakly negatively correlated (-0.09) with the initial wealth-income ratio. Mechanically, this reflects the fact that households who enter the sample with high wealth have risky shares that are insufficiently lower than the risky shares of other households to be consistent with the same level of RRA. Finally, RRA is negatively correlated (-0.21) with the convexity of the wealth-income ratio.

Our estimate of the TPR is negatively correlated (-0.30) with the initial wealth-income ratio and positively correlated (0.35) with the average growth rate of the wealth-income ratio in our sample period. Mechanically, this is due to the fact that households that enter our sample with low initial wealth accumulate wealth more rapidly than average households, but not as rapidly as they would do if they had an average 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 (0.28) with the average risky share, reflecting the lower average wealth-income ratio of impatient households that justifies a riskier investment strategy.

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

Table 7 relates our parameter estimates to observable variables in a different way, reporting multiple regressions of the estimates onto the four observables discussed above. In the even-numbered columns we also include control variables: dummies for the 9 categories of income risk and education, and cohort fixed effects. The coefficients on the observable variables are robust to the inclusion of these controls, although the controls do increase the explanatory power of the regressions, particularly those predicting RRA.

The first two columns of Table 7 show that RRA is predicted negatively by both the average risky share and the initial wealth-income ratio, consistent with the correlations reported in Table 6. Controlling for these effects, the average growth rate of the wealth-income ratio predicts RRA positively, as implied by the theory of precautionary savings. The R^2 of the regressions are 0.33 without control variables and 0.85 with them.

The next two columns of Table 7 show that TPR is predicted positively by the average risky share and the average growth rate of wealth. Both variables are positively correlated with the initial wealth-income ratio and they drive out any predictive effect of that variable. The R^2 coefficients are 0.23 without control variables and 0.27 with them.

The last two columns of Table 7 show that EIS is predicted positively by the initial wealth-income ratio, consistent with Table 6, and negatively by the convexity of the wealth-income ratio. There is only a weak positive effect of the growth rate of the wealth-income ratio, and the R^2 of the regressions are modest at 0.11 with or without control variables. The limited ability of linear regressions to explain the EIS estimates is consistent with our analysis of identification in section 4.2.

Another interesting question is how our estimates are related to households' income risk and education. In Table 8 we explore this by regressing preference estimates on labor income volatility and the level of education. All regressions also include cohort fixed effects, but the cohort coefficients are not reported.

RRA is the preference parameter most strongly related to these observables. Households with riskier labor income tend to have lower risk aversion. Mechanically, this results from the fact documented in Table 3 that income volatility has little effect on the risky share: if risk aversion were the same for safe and for risky occupations, then the risky share should fall with income risk. 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. Beyond this, more educated people also tend to have lower risk aversion. The R^2 of the regression for risk aversion is almost 38%.

Households with high income risk also tend to have a lower TPR, but the explanatory power of this regression is less than 7%. Finally, the EIS has almost no relation to income risk, education, or cohort fixed effects; the explanatory power of the regression for EIS is less than 1%.

5.3 Parameter Uncertainty

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

As a first exercise, we plot parameter standard errors against the corresponding parameter estimates in three heat maps in Figure 5. The top panel of the figure shows that most RRA estimates have modest standard errors, and there is no obvious cross-sectional correlation between the levels of the estimates and their standard errors. The middle panel shows that TPR estimates in the middle of the range tend to be precisely estimated with low standard

errors, while unusually low and high TPR estimates tend to have high standard errors with the exception of a few households whose TPR is estimated at the lower bound. The bottom panel shows that low EIS estimates tend to be precisely estimated, while high EIS estimates tend to have large standard errors.

As a second exercise, in Table 9 we report summary statistics describing the results of hypothesis tests based on our asymptotic standard errors and using 5% significance levels. The top panel of the table focuses on hypotheses about the time preference rate. The first row of the table reports the fact we have already mentioned that 6.5% of households are in groups estimated to have a negative TPR. The second row of the table reports that we can reject the null of a positive TPR at the conventional 5% significance level for only 2.1% of households; thus, a significantly negative TPR is a rare occurrence in our sample. The third row of the table shows that we can reject the null of a negative TPR at the 5% level for 61.1% of households, and the fourth row shows that we can reject the null of a zero TPR using a two-sided test for 58.7% of households. Thus we have the reassuring result that the TPR is significantly positive for the majority of the Swedish population.

The next panel of the table considers hypotheses about the EIS. We have already reported that 59.7% of households are in groups with estimated EIS less than one. The following rows show that we can reject the null of an EIS greater than one for 49.0% of households, and can reject the null of an EIS less than one for only 17.2% of households. Using a two-sided test, we can reject the null that the EIS equals one for 63.3% of households. There is considerable heterogeneity in the EIS, but it is far more common for Swedish households to have an EIS significantly below one than an EIS significantly above one.

Turning to a different null, that of power utility, we have already reported that 34.7% of households have an estimated EIS that is lower than the reciprocal of risk aversion. The table shows that we can reject the null that EIS exceeds the reciprocal of risk aversion for 5.4% of households, and can reject the null that the EIS is lower than the reciprocal of risk aversion for 33.1% of households. We can reject the power utility null using a two-sided test for 34.0% of households. Power utility, despite the tight restrictions it imposes on preference parameters, is a reasonable description of preferences for the majority of the Swedish population.

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

equal the cross-sectional median estimates, treating the medians as known for simplicity. Overall, the table presents strong statistical evidence against homogeneity of preferences within the framework of our life-cycle model.

A third use of our asymptotic standard errors is to adjust our estimates of the heterogeneity in true preference parameters. Table 5 and Figure 3 describe the cross-sectional distribution of our parameter estimates, but this is increased by noise in the estimation procedure. Since our asymptotic standard errors estimate the noise for each group, in principle we can correct for the effect of noise on the estimated cross-sectional variance of parameters by subtracting the cross-sectional average squared standard error from the cross-sectional variance of our estimates.

A practical difficulty in doing this is that some groups have extremely high standard errors. Although these high standard errors are not pervasive enough to undermine our ability to reject homogeneous preferences for most households in the group-specific tests reported in Table 9, they do have a strong influence on the cross-sectional average of squared standard errors. In fact, if we do not limit the influence of outliers the average squared standard error is higher than the cross-sectional variance of estimates for TPR and EIS, implying a negative cross-sectional variance for true TPR and EIS. We obtain more reasonable results if we winsorize the group-specific standard errors at the 90th percentile of the cross-sectional distribution, trimming the 10% largest standard errors. This procedure implies a cross-sectional standard deviation of 0.46 for RRA, 3.69% for the TPR, and 0.54 for the EIS, as compared with the cross-sectional standard deviations of estimates reported in Table 5 which are 0.47, 6.03%, and 0.96 respectively. We present further details of this analysis in the online appendix.

5.4 Model Fit

In this subsection we consider measures of model fit more directly. We begin by describing the cross-sectional distribution of the errors our model makes in fitting the 16 auxiliary statistics that are the target of our estimation procedure. We take the 8 wealth-income ratios and the 8 risky shares, and for each of these variables we calculate the root mean squared error (RMSE), the square root of the average squared deviation of the model-fitted variable from the observed variable. The results are reported in percentage points in the first two rows of Table 10. The mean RMSE across all groups is 30.8% for the wealth-income ratio and 4.7% for the risky share. In other words, the average error in fitted wealth is just under 4 months of income and the average error in the risky share is just below 5% of wealth. The RMSE distribution is somewhat right-skewed as indicated by the fact that the median RMSEs are below the mean RMSEs at 19.9% and 4.1% respectively.

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

Our estimation procedure takes into account that the wealth-income ratio and the risky share have different units, and scales them in proportion to their grand cross-sectional means. The next two rows of Table 9 similarly divide the RMSEs for the wealth-income ratio and risky share by their grand means of 4.93 and 0.65, respectively, to express them in proportional units. The mean proportional RMSE is 6.3% for the wealth-income ratio and 7.2% for the risky share.

Finally, we report a transformation of the objective function that is rescaled to express it in RMSE-equivalent units. The objective function is the sum of squared proportional errors, so we divide by the number of auxiliary statistics (16) and take the square root, then multiply by 100 to express the RMSE-scaled objective function in percentage points. Group by group, the result is not exactly the average of the proportional errors for the wealth-income ratio and the risky share because of the interpolation method we use in estimation; and the quantiles of the cross-sectional distribution also may refer to different groups. Nonetheless, the bottom row of Table 10 is similar to an average of the previous two rows. The mean RMSE-scaled objective function is 6.7%, with a moderately right-skewed distribution.

Table 11 shows how the cross-sectional distribution of model fit, as measured by the RMSE-scaled objective function, deteriorates if we suppress cross-sectional heterogeneity in preference parameters. The top row of Table 11 repeats the bottom row of Table 10, and subsequent rows set particular preference parameters to their estimated cross-sectional means (treated as known quantities). The mean RMSE-scaled objective function more than doubles to 15.8% if we fix RRA at its cross-sectional mean. Fixing TPR at its cross-sectional mean is less deleterious but still produces a mean RMSE-scaled objective function of 7.9%. The smallest impact comes from restricting the EIS to its cross-sectional mean, which delivers a mean RMSE-scaled objective function of 7.6%.

Both the TPR and the EIS have a strong impact on savings decisions, so a model that fixes one of these can compensate to some degree by varying the other. Unsurprisingly, if we restrict both these parameters simultaneously to their cross-sectional means, the model fit deteriorates further and we get a mean RMSE-scaled objective function of 10.2%. Finally, fixing all parameters at their cross-sectional means is disastrous in the sense that it increases

the mean RMSE-scaled objective function to 26.6%. A life-cycle model with homogeneous preferences, under our maintained assumption of homogeneous rational beliefs, delivers an extremely poor fit to the cross-section of household financial behavior.

Figures 6 and 7 illustrate these results in greater detail using a series of heat maps. In Figure 6 the three columns correspond to restrictions on RRA (left column), the TPR (middle column), and the EIS (right column). The top row shows the proportional RMSE for the wealth-income ratio, the middle row shows the proportional RMSE for the risky share, and the bottom row shows the RMSE-scaled objective function. Within each heat map, the horizontal axis indicates the RMSE of our unrestricted model, and the vertical axis indicates the RMSE of the particular restricted model under consideration. The logarithmic color scale indicates data density as in Figures 4 and 5.

The heat maps in Figure 6 show that suppressing heterogeneity in risk aversion has a particularly harmful effect on the overall fit of the risky share, as one might expect; but it also causes the fit of the wealth-income ratio to deteriorate on average. Suppressing heterogeneity in the TPR or the EIS also tends to worsen the fit of the risky share, as the model adjusts risk aversion to generate precautionary saving, fitting the wealth-income ratio at the expense of the fit to the risky share. There are a few cases where one or the other auxiliary statistic is better fit by a restricted model than by the unrestricted model, but such cases are always more than offset by the deterioration in the other RMSE so that restrictions always increase the RMSE-scaled objective function. Parameter restrictions tend to have the worst impact on those groups which are best fit by the unrestricted model, at the left of each heat map.

Figure 7 has the same structure as Figure 6, but restricts both the TPR and the EIS simultaneously in the left column, and all three preference parameters in the right column. As previously noted, these simultaneous parameter restrictions are even more destructive of model fit than restricting parameters one at a time.

5.5 Monte Carlo Analysis

Our estimation procedure is based on asymptotic analysis as the number of households in each group increases. In order to evaluate the finite-sample performance of our procedure for groups of the sizes we encounter empirically, we undertake a simple Monte Carlo exercise. For each group in our sample, we simulate our model under the group's initial conditions and the preference parameters we estimated for the group. We combine simulated households into hypothetical groups each containing N_g^* households, where N_g^* is a measure of the effective empirical group size. We repeat this procedure to obtain 1,000 hypothetical groups

and calculate the mean parameter estimate. A comparison of this mean with the preference parameters under which the model was simulated allows us to assess finite-sample bias in our estimation method.

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

Table 12 reports regression coefficients of Monte Carlo mean parameter estimates on the parameter estimates that were used to generate the simulated data (“true” parameters for the purpose of this exercise). The results are very good for RRA, which has a slope coefficient of 1.006, insignificantly different from one, and an R^2 statistic of 93%. The regression for TPR has a slope coefficient of 0.931 and an R^2 statistic of 87%. Results are not quite as good for EIS, which has a slope coefficient of 0.712 and an R^2 statistic of 0.608. This regression places most of its weight on the high EIS estimates, which are noisy; but results are similar if we run the regression using the log of the EIS.

An important lesson of these results is that small-sample bias cannot explain the substantial cross-sectional heterogeneity in our preference parameter estimates. There is almost no small-sample bias for RRA, and minimal bias for the TPR; and while there is some bias in our EIS estimates, a bias correction would have little effect on the cross-sectional dispersion of the EIS.

We have also explored using our Monte Carlo procedure to calculate standard errors for preference parameters. We find the Monte Carlo standard errors to be much smaller than our asymptotic standard errors, implying even stronger rejections of the hypothesis of homogeneous preferences across groups. These small Monte Carlo standard errors reflect the low degree of uncertainty that we include in our simulations. The simulations assume that each household starts each year with the group average wealth-income ratio, that our estimated labor income process holds exactly for each household, that each household in the group has an identical financial and real estate portfolio so that it earns an identical

portfolio return that we feed into the simulation, and that households receive no other wealth shocks such as inheritances or other transfers from relatives. We make these assumptions for tractability in our Monte Carlo analysis, but clearly they abstract from important sources of household-specific variation in the data. The low degree of uncertainty becomes even more significant in our context because our preference parameters are estimated on a grid, hence small perturbations can lead to identical parameter estimates. Given these concerns we do not use Monte Carlo standard errors in our analysis, and likewise we do not attempt to use the Monte Carlo procedure to conduct formal specification tests of our model.

5.6 Robustness

In this subsection we briefly report some robustness checks, giving more complete details in the online appendix.

As a first check, we exclude households whose initial wealth-income ratios are in the highest decile. These households are challenging for our model because our identification strategy is likely to be less effective when wealth is very high relative to income, and because any mismeasurement of risky returns has the largest impact on WY when wealth is already high. Excluding these households, we find that the distribution of risk aversion estimates is fairly stable, but the TPR distribution shifts slightly upwards and the EIS distribution shifts slightly downwards. Despite this, the fraction of households for which we can reject an EIS less than one remains stable at around 17%, reflecting the fact that high-WY households often have high standard errors as well as high levels of their EIS estimates. Finally, the fit of our life-cycle model improves considerably when we exclude these households. The mean RMSE for WY, reported in the first row and first column of Table 10, declines from 30.8% in the base case to 23.3%.

As a second check, we increase the assumed Sharpe ratio of the composite risky asset from 0.26 to 0.40. Unsurprisingly, this change increases the risk aversion needed to rationalize the portfolio choice of Swedish households. Mean risk aversion increases from 5.24 to 7.94, and median risk aversion from 5.30 to 7.90. The improved investment opportunity set also increases the time preference rate needed to fit the accumulation of wealth (given that we estimate most households to have an EIS less than one). The mean TPR increases from 6.2% to 8.7%, and the median TPR from 4.1% to 6.2%. Our EIS estimates decline slightly, with the mean falling from 0.99 to 0.67 and the median declining from 0.42 to 0.27.

The change in the Sharpe ratio does not, however, do much to the cross-sectional dispersion of parameter estimates that is the main focus of this paper. The cross-sectional standard deviation of our risk aversion estimates rises from 0.47 to 0.74, but this increase is

roughly proportional to the increase in the mean. The cross-sectional standard deviation of the TPR estimates is slightly higher and that of the EIS estimates is slightly lower than in the base case, but the distributions of our estimates are qualitatively similar. It remains true that the TPR is negatively correlated with both risk aversion and the EIS, and in fact the negative correlation is slightly stronger than in the base case. The parameter estimates also have similar correlations with observables. We continue to find very strong evidence of heterogeneity in preference parameters using tests based on our asymptotic standard errors of the sort reported in Table 9.

As a third check, we extend the above analysis to consider a simple form of heterogeneity in beliefs about the available Sharpe ratio. We simulate the model under three alternative assumptions about the Sharpe ratio: the base value of 0.26, the high value of 0.40, and a low value of 0.15. Then, for each group we pick the assumption and set of preference parameters that minimizes the objective function. The base case Sharpe ratio is selected for groups representing 49% of households, the low Sharpe ratio for 28% of households, and the high Sharpe ratio for 23% of households.

Allowing for heterogeneity in household beliefs about the Sharpe ratio improves the fit of our model modestly. The RMSE-scaled objective function that was reported in the bottom row of Table 10 falls from a mean of 6.67% and a median of 5.92% in the base case, to a mean of 5.43% and a median of 4.79% in the richer model with heterogeneous beliefs.¹⁹ However, heterogeneous beliefs do not reduce the heterogeneity we estimate in preference parameters. In fact, the cross-sectional standard deviation of RRA is almost four times larger, at 1.91, in the heterogeneous-beliefs model, and the cross-sectional standard deviations of TPR and EIS are very slightly larger. The explanation for this pattern is that the model uses heterogeneous beliefs to improve its ability to fit wealth accumulation, and offsets belief heterogeneity with RRA heterogeneity to avoid counterfactual heterogeneity in the risky share.

6 Conclusion

In this paper we have 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 house-

¹⁹The improvement in fit is concentrated among households whose behavior is already well fit by the base model; the right tail of the distribution of the RMSE-scaled objective function does not change much. Related to this, the improvement in fit is sensitive to the way in which we interpolate the objective function on a fine grid in the last step of our estimation procedure.

holds 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 assumption we maintain for most of this 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 relative risk aversion (RRA) of 0.47 around a mean of 5.24, a standard deviation in the time preference rate (TPR) of 6.03% around a mean of 6.2%, and a standard deviation in the elasticity of intertemporal substitution (EIS) of 0.96 around a mean of 0.99. The cross-sectional distributions of both the TPR and the EIS are right-skewed, so their medians are lower than their means at 4.08% and 0.42 respectively.

Our estimates of risk aversion and the EIS are almost uncorrelated across households, 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 negatively correlated with both risk aversion and the EIS.

We estimate a negative correlation between income volatility 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. Income volatility is also negatively related to the TPR and weakly positively related to the EIS. More educated households tend to have lower risk aversion, controlling for their income volatility, but we do not find that educated households are more patient; if anything, the relationship is the opposite when we control for income volatility.

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

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

Even if individual households have constant preference parameters, cross-sectional heterogeneity in these parameters can break the relation between household preferences and the implied preferences of a representative agent. In a representative-agent economy, preferences with habit formation are needed to generate countercyclical variation in the price of risk (Constantinides 1990, Campbell and Cochrane 1999), but in heterogeneous-agent economies, countercyclical risk premia can arise from time-variation in the distribution of wealth across agents with different but constant risk preferences (Dumas 1989, 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), Meeuwis et al (2018), and Giglio et al (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).

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

Our model omits some other features of the household decision problem that may potentially be important and deserve further research. We assume that financial market returns are iid rather than time-varying, ignoring intertemporal effects on asset allocation discussed by Campbell and Viceira (1999, 2002). In our base case we assume homogeneous beliefs about financial market returns. As a robustness check we have allowed for a simple form of belief heterogeneity, but certainly it would be possible to go further in this direction, for example along the lines suggested by Calvet et al (2020). We fix preference parameters for each household and do not allow them to vary with wealth at the household level, contrary to evidence that relative risk aversion declines with wealth (Carroll 2000, 2002, Wachter and Yogo 2010, Calvet and Sodini 2014). We model labor income risk using normally distributed shocks rather than the skewed distributions estimated by Guvenen, Ozkan, and Song (2014) and others. We treat labor income as exogenous and do not consider the possibility that the household can endogenously vary its labor supply (Bodie, Merton, and Samuelson 1992, Gomes, Kotlikoff and Viceira 2008). We ignore the possibility that some components of consumption involve precommitments or generate habits that make them costly to adjust (Gomes and Michaelides 2003, Chetty and Szeidl 2007, 2010). We do not model homeownership jointly with other financial decisions as in Cocco (2005). Household-level data on asset allocation and wealth accumulation and our structural approach to estimation of a life-cycle model provide a laboratory in which these aspects of household financial decision-making can be explored.

References

- Ameriks, John, and Stephen Zeldes, 2004, How do household portfolio shares vary with age?, unpublished paper, Columbia University.
- 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, 2020, Rich pickings? Risk, return, and skill in household wealth, *American Economic Review* 110, 2703–2747.
- Bansal, Ravi and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- Barsky, Robert, Thomas Juster, Miles Kimball, and Matthew Shapiro, 1997, Preference parameters and behavioral heterogeneity: An experimental approach in the Health and Retirement Study, *Quarterly Journal of Economics* 112, 537–579.
- 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 C  lerier, Paolo Sodini, and Boris Vall  e, 2020, Can security design foster household risk-taking?, 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. and Luis Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors* (Oxford University Press).
- 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.
- Curcuro, 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.
- Dahlquist, Magnus, Ofer Setty, and Roine Vestman, 2018, “On the Asset Allocation of a Default Pension Fund”, *Journal of Finance* 73, 1893–1936
- 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.
- Ghosh, Anisha, Arthur Korteweg, and Qing Xu, 2020, Recovering heterogeneous beliefs and preferences from asset prices, unpublished paper, McGill University and University of Southern California.
- 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).
- Guiso, Luigi and Paolo Sodini, 2013, Household finance: an emerging field, in George Constantinides, Milton Harris, and Rene Stulz eds., *Handbook of the Economics of Finance, Volume 2B*, North-Holland, 1397–1531.
- Guvenen, Fatih, 2009, A parsimonious macroeconomic model for asset pricing, *Econometrica* 77, 1711–1750.
- Guvenen, Fatih, Serdar Ozkan, and Jae Song, 2014, The nature of countercyclical income risk, *Journal of Political Economy* 122, 621–660.
- Haliassos, Michael, and Carol Bertaut, 1995, Why do so few hold stocks?, *Economic Journal* 105, 1110–1129.
- Hall, Robert E., 1988, Intertemporal substitution in consumption, *Journal of Political Economy* 96, 221–273.
- Heaton, John, and Deborah Lucas, 2000, Portfolio choice and asset prices: The importance of entrepreneurial risk, *Journal of Finance* 55, 1163–1198.
- 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.
- Hwang, Jungbin, and Yixiao Sun, 2018, Should we go one step further? An accurate comparison of one-step and two-step procedures in a generalized method of moments framework, *Journal of Econometrics* 207, 381–405.
- Kocherlakota, Narayana R., 1990, Disentangling the coefficient of relative risk aversion from the elasticity of intertemporal substitution: An irrelevance result, *Journal of Finance* 45, 175–190.
- 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.

- Meeuwis, Maarten, Jonathan A. Parker, Antoinette Schoar, and Duncan I. Simester, 2018, Belief disagreement and portfolio choice, NBER working paper 25108.
- Mehra, Rajnish and Edward Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.
- Parker, Jonathan, Antoinette Schoar, and Yang Sun, 2020, Retail financial innovation and stock market dynamics: the case of target date funds, NBER working paper 28028.
- 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: Size-Weighted Wealth-Income Ratios and Risky Shares by Levels of Education and Income Volatility

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

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

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

Table 2: Percentage Volatilities of Income Shocks

	Total			Systematic		
	No High School	High School	Post-High School	No High School	High School	Post-High School
Low	13.59	13.42	15.85	2.76	2.88	3.49
Medium	17.20	16.31	16.69	2.86	3.37	3.44
High	20.15	19.83	21.48	3.08	3.27	3.91

	Idiosyncratic permanent			Idiosyncratic transitory		
	No High School	High School	Post-High School	No High School	High School	Post-High School
Low	7.48	6.88	5.27	11.01	11.15	14.53
Medium	8.84	7.8	5.64	14.47	13.93	15.33
High	6.95	7.35	5.91	18.66	18.13	20.28

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

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

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

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

Table 4: Regressions of Preference Parameters on Simulated Moments

Panel A. RRA Regressions.

Average RS	Yes	Yes	Yes	No
Initial WY	No	Yes	Yes	No
Growth of WY	No	No	Yes	No
Convexity of WY	No	No	No	No
16 moments	No	No	No	Yes
R^2	0.724	0.740	0.750	0.800

Panel B. TPR Regressions.

Average RS	No	No	Yes	No
Initial WY	No	Yes	Yes	No
Growth of WY	Yes	Yes	Yes	No
Convexity of WY	No	No	Yes	No
16 moments	No	No	No	Yes
R^2	0.606	0.645	0.647	0.643

Panel C. EIS Regressions.

Average RS	No	No	Yes	No
Initial WY	No	Yes	Yes	No
Growth of WY	Yes	Yes	Yes	No
Convexity of WY	Yes	Yes	Yes	No
16 moments	No	No	No	Yes
R^2	0.032	0.033	0.056	0.141

Panel D. EIS Regressions Part II.

WY range	≤ 1	(1, 2]	(2, 3]	(3, 5]	(5, 7]	(7, 10]	> 10
All RRA and TPR	0.243	0.247	0.327	0.278	0.289	0.319	0.190
TPR = 0.03	0.829	0.758	0.685	0.617	0.584	0.538	0.305

This table reports the R^2 statistics of regressions in simulated data using all preference parameters on a grid containing 12 values of relative risk aversion (RRA) ranging from 2 to 10, 11 values of the time preference rate (TPR) ranging from -0.05 to 0.22, and 14 values of the elasticity of intertemporal substitution (EIS) ranging from 0.1 to 2.5. For each of the 1,848 combinations of preference parameters we consider all initial levels of the wealth-income ratio (WY) observed among Swedish household groups. (*Continued on the next page.*)

(Continued.) We regress the preference parameters on simulated moments including the average risky share (RS), the initial WY, the 8-year cumulative growth of WY defined in equation (17) in the text, the convexity of WY defined in equation (18) in the text, and all 16 moments (8 values of RS and 8 values of WY) used in our empirical analysis. The dependent variable in the regressions is RRA in Panel A, the TPR in Panel B, and the EIS in Panel C. The four columns in Panels A, B, and C include different combinations of explanatory variables. Panel D runs EIS regressions separately, using all 16 moments, for simulated groups with initial WY in different bins indicated in the columns. The first row of panel D reports the R^2 statistic when the regressions include simulated groups with all possible values of RRA and the TPR, while the second row reports the R^2 statistic when the regressions include only simulated groups whose TPR equals 0.03.

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

	Mean	Median	Std. Dev.	10%	25%	75%	90%
RRA	5.24	5.30	0.47	4.60	4.90	5.50	5.90
TPR (%)	6.18	4.08	6.03	1.01	3.15	7.42	18.39
EIS	0.99	0.42	0.96	0.10	0.10	1.97	2.50
Average RS	0.65	0.62	0.17	0.45	0.52	0.75	0.89
Initial WY	4.13	2.88	3.78	0.96	1.69	4.71	7.67
Growth of WY	1.08	1.07	0.05	1.03	1.05	1.10	1.14
Convexity of WY	0.24	0.23	0.10	0.15	0.19	0.28	0.34

This table reports the mean, median, standard deviation, and 10th, 25th, 75th, and 90th percentiles of estimated preference parameters and group financial characteristics. All statistics weight groups by their size to recover the underlying cross-sectional distributions at the household level. Growth of WY is defined in equation (17) and convexity of WY is defined in equation (18) in the text. There are 4,151 groups containing 291,488 households.

Table 6: Size-Weighted Cross-Sectional Correlations of Estimated Preference Parameters and Group Financial Characteristics

	RRA	TPR	EIS	Average RS	Initial WY	Growth of WY	Convexity of WY
RRA	1.000						
TPR	-0.275***	1.000					
EIS	0.113***	-0.289***	1.000				
Average RS	-0.412***	0.472***	-0.155***	1.000			
Initial WY	-0.088***	-0.296***	0.320***	-0.524***	1.000		
Growth of WY	0.003	0.352***	-0.179***	0.621***	-0.685***	1.000	
Convexity of WY	-0.205***	0.095***	0.034***	0.163***	0.182***	0.029***	1.000

This table reports the cross-sectional correlations across estimated preference parameters and group financial characteristics. Correlations weight groups by their size to recover the underlying cross-sectional correlations at the household level. Growth of WY is defined in equation (17) and convexity of WY is defined in equation (18) in the text. Statistical significance levels of correlation coefficients are indicated with stars: * denotes 1-5%, ** 0.1-1%, *** less than 0.1% significance. There are 4,151 groups containing 291,488 households.

Table 7: Size-Weighted Cross-Sectional Regressions of Estimated Preference Parameters on Group Financial Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
	RRA	RRA	TPR	TPR	EIS	EIS
Average RS	-1.900*** (0.061)	-1.985*** (0.053)	0.137*** (0.009)	0.124*** (0.009)	-0.004 (0.111)	-0.054 (0.115)
Initial WY	-0.034*** (0.003)	-0.030*** (0.004)	-0.001* (0.000)	0.000 (0.000)	0.098*** (0.006)	0.098*** (0.006)
Growth of WY	2.609*** (0.383)	3.212*** (0.478)	0.098** (0.037)	0.201*** (0.036)	1.820*** (0.510)	1.960*** (0.525)
Convexity of WY	-0.223** (0.070)	-0.475*** (0.051)	0.022 (0.012)	0.019 (0.012)	-0.368** (0.140)	-0.350** (0.135)
Constant	3.852*** (0.404)	3.286*** (0.495)	-0.135*** (0.040)	-0.230*** (0.039)	-1.293* (0.548)	-1.474** (0.562)
Control variables	No	Yes	No	Yes	No	Yes
R^2	0.326	0.849	0.230	0.271	0.107	0.111

This table reports the cross-sectional regression coefficients across estimated preference parameters and group financial characteristics. All regressions weight groups by their size, to recover the underlying cross-sectional relationships at the household level. Growth of WY is defined in equation (17) and convexity of WY is defined in equation (18) in the text. Standard errors are reported in parentheses and statistical significance levels are indicated with stars: * denotes 1-5%, ** 0.1-1%, *** less than 0.1% significance. There are 4,151 groups containing 291,488 households. Control variables are 9 income risk/education categories and cohort fixed effects.

Table 8: Size-Weighted Cross-Sectional Regressions of Preference Parameters on Income Volatility and Education

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

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

Table 9: Size-Weighted Statistical Test Results for Estimated Preference Parameters

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

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

Table 10: Size-Weighted Cross-Sectional Distributions of Model Fit Indicators

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

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

Table 11: Size-Weighted Cross-Sectional Distributions of RMSE-Scaled Objective Functions for Alternative Model Specifications

	Mean	Median	Std. Dev.	10%	25%	75%	90%
Unrestricted	6.67	5.92	5.42	2.42	4.31	7.99	11.59
Fixed RRA	15.81	13.89	9.70	6.02	9.11	20.20	27.66
Fixed TPR	7.91	6.79	5.51	3.81	5.11	9.27	13.36
Fixed EIS	7.61	6.66	5.28	3.72	5.01	9.06	12.49
Fixed TPR and EIS	10.19	8.86	5.68	5.87	7.35	11.43	16.82
All Parameters Fixed	26.56	20.25	20.10	8.59	12.80	33.08	54.26

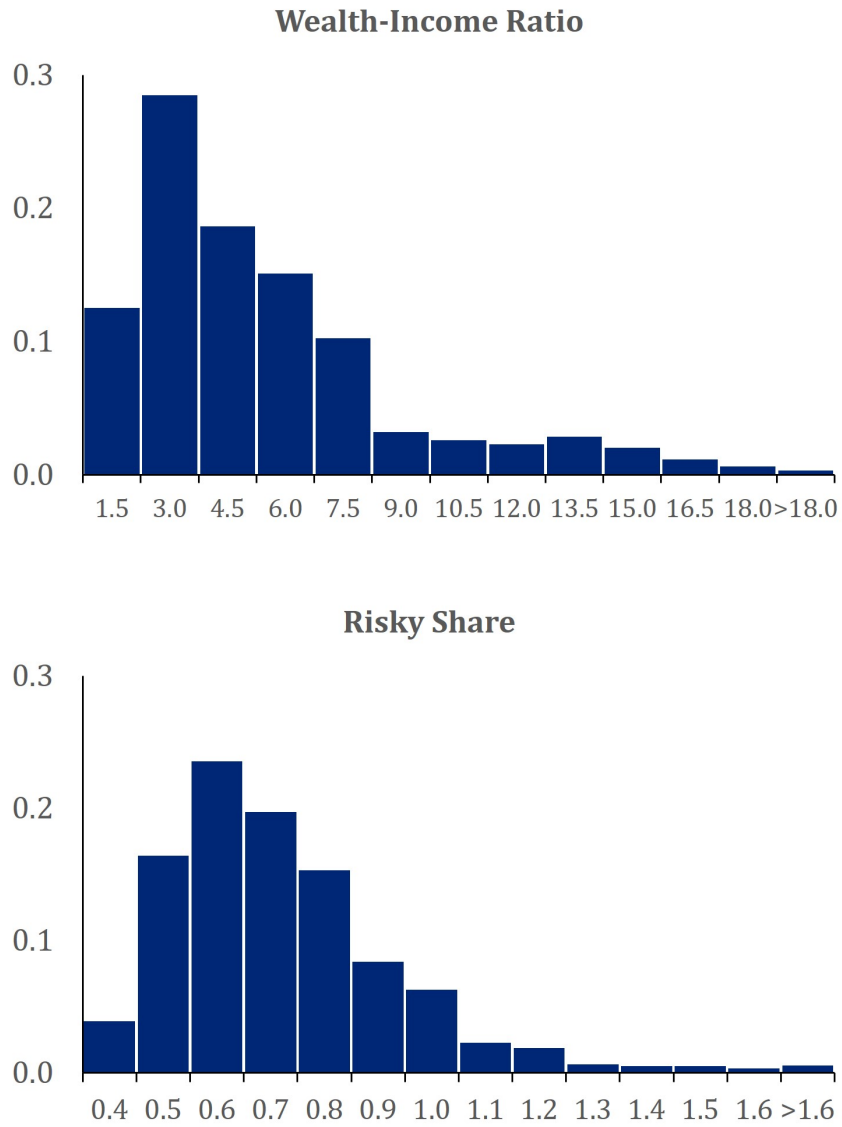
This table reports the mean, median, standard deviation, and 10th, 25th, 75th, and 90th percentiles of the RMSE-scaled objective function for several alternative model specifications. All statistics weight groups by their size to recover the underlying cross-sectional distributions at the household level. The RMSE-scaled objective function is the square root of the objective function divided by 4 and multiplied by 100 to express it in RMSE-equivalent percentage units. The results in the first row are for the unrestricted model estimated in Table 9. The results in subsequent rows are for models that fix selected parameters at their size-weighted cross-sectional means estimated in the unrestricted model. There are 4,151 groups containing 291,488 households.

Table 12: Size-Weighted Cross-Sectional Regressions of Mean Monte Carlo Preference Parameter Estimates on Indirect Inference Estimates

	(1)	(2)	(3)	(4)
	RRA	TPR	EIS	Log EIS
Slope coefficient	1.006*** (0.005)	0.930*** (0.007)	0.704*** (0.012)	0.732*** (0.011)
Constant	-0.030 (0.026)	0.006*** (0.001)	0.314*** (0.017)	-0.014 (0.012)
R^2	0.935	0.868	0.606	0.673


This table reports cross-sectional regressions of average Monte Carlo estimates of preference parameters on the preference parameters used to generate simulated data, which are set equal to indirect inference parameter estimates for each group. Monte Carlo simulations use the effective group size, the reciprocal of the sum of squared wealth shares of individual households in each group. Average estimates are calculated from 1,000 simulations of each group. All regressions weight groups by their size, to recover the underlying cross-sectional relationships at the household level. Standard errors are reported in parentheses and statistical significance levels are indicated with stars: * denotes 1-5%, ** 0.1-1%, *** less than 0.1% significance. There are 4,151 groups containing 291,488 households.

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




This figure presents histograms for the wealth-income ratio (WY) and risky share (RS) across 4,151 groups of Swedish households, size-weighted to recover the underlying distribution across households under the assumption that WY and RS are homogeneous within groups. Each bin is labeled on the horizontal axis with the upper cutoff value of WY or RS at the right edge of the bin, except the extreme right bin which captures all groups above the previous bin's cutoff. The vertical axis shows the size-weighted fraction of the sample in each bin.

Figure 2: Estimated Age-Income Profiles



This figure presents estimated age-income profiles, including replacement ratios in retirement, for Swedish households with three levels of education: no high school (HS), high school, and post-high-school. The estimates are based on a labor income process specified in equations (6)-(10) in the text.


Figure 3: Distribution of Estimated Preference Parameters



This figure presents histograms for estimates of relative risk aversion (RRA), the time preference rate (TPR), and the elasticity of intertemporal substitution (EIS) across 4,151 groups of Swedish households, size-weighted to recover the underlying distribution across households under the assumption that preferences are homogeneous within groups. *(Continued on the next page.)*

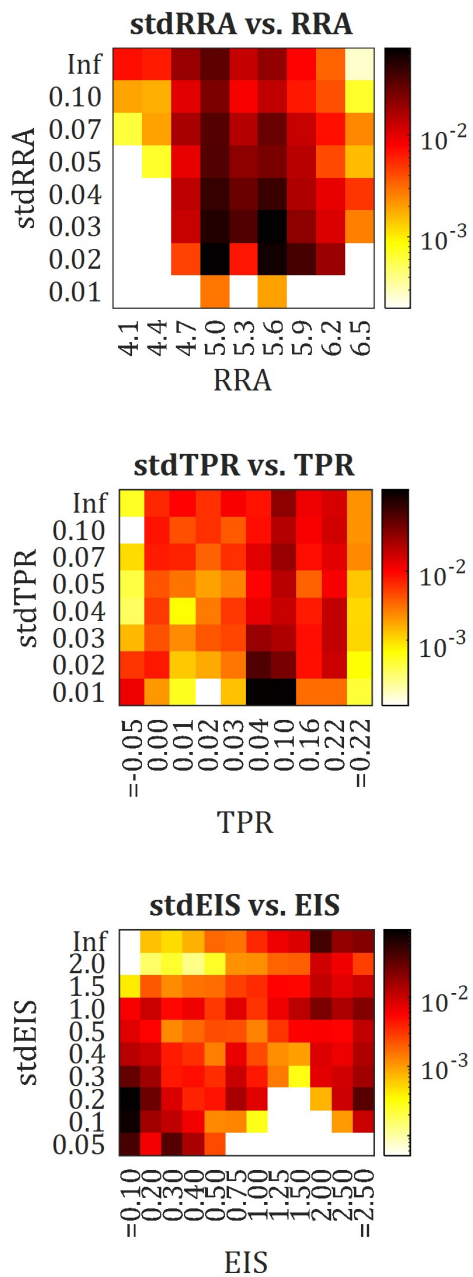
(Continued.) Each horizontal axis label shows the upper cutoff value at the right edge of the bin above the label. The vertical axis shows the size-weighted fraction of the sample in each bin. The estimation procedure allows RRA to range from 2 to 10, the TPR from -0.05 to 0.22, and the EIS from 0.1 to 2.5.

Figure 4: Joint Distribution of Estimated Preference Parameters



This figure presents bivariate heat maps for estimates of RRA and EIS (top panel) and TPR and EIS (bottom panel) across 4,151 groups of Swedish households, size-weighted to recover the underlying distribution across households under the assumption that preferences are homogeneous within groups. Each axis label shows the upper cutoff value of the corresponding bin, except for labels beginning with = which indicate that the bin contains only estimates of the exact value indicated by the label. The logarithmic color scheme indicates the fraction of the sample in each bin. This fraction is equal to 8.3% and 9.6% for the darkest color in the top and bottom panels respectively and 0.0% for the brightest color in both panels.

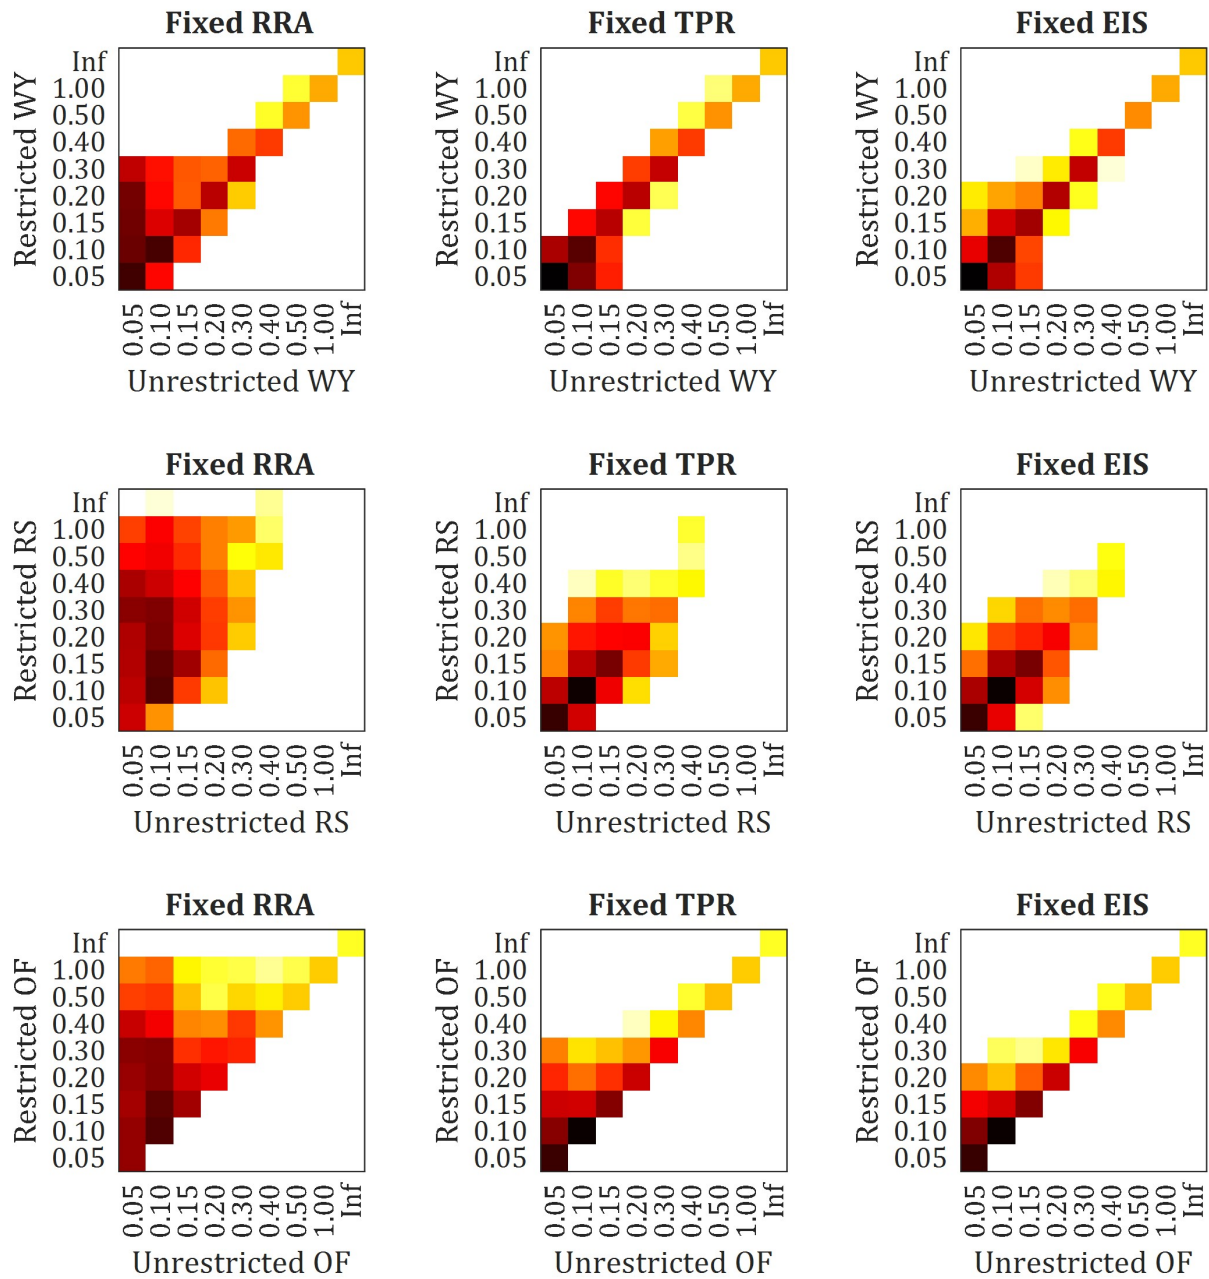
Figure 5: Joint Distribution of Estimated Preference Parameters and Respective Standard Errors



(Figure note on the next page.)

This figure presents bivariate heat maps for estimates of RRA and its standard error (top panel), the TPR and its standard error (middle panel), and the EIS and its standard error (bottom panel) across 4,151 groups of Swedish households, size-weighted to recover the underlying distribution across households under the assumption that preferences are homogeneous within groups. Each axis label shows the upper cutoff value of the corresponding bin, except for labels beginning with = which indicate that the bin contains only estimates of the exact value indicated by the label, and the label Inf which indicates that the bin has no upper cutoff but contains all values above the previous bin's cutoff. The logarithmic color scheme indicates the fraction of the sample in each bin. This fraction is equal to 7.2%, 15.1% and 9.5% for the darkest color in the top, middle and bottom panels respectively and 0.0% for the brightest color in all three panels.

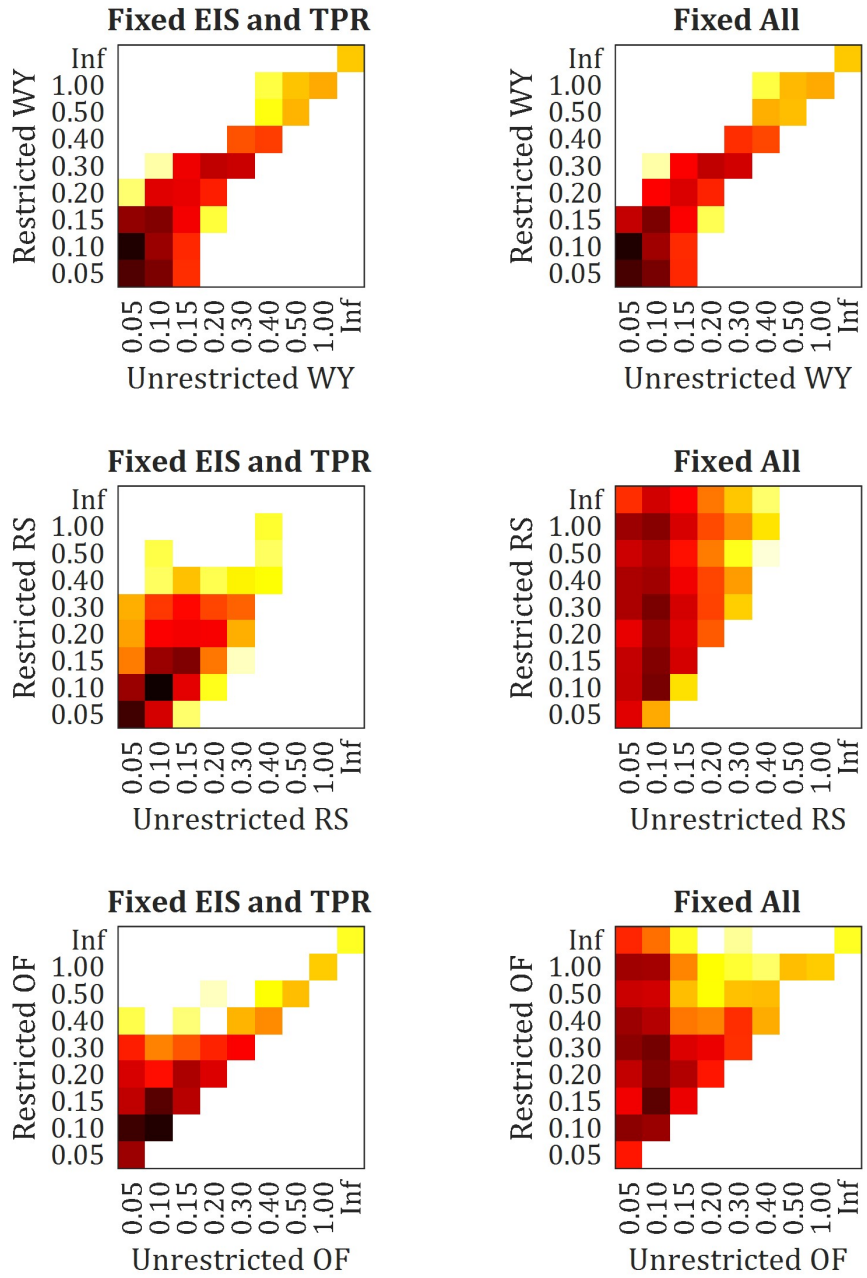
Figure 6: Impact on Model Fit of Imposing Homogeneity of a Single Preference Parameter



(Figure note on the next page.)

This figure presents bivariate heat maps for measures of model fit comparing the unrestricted model with three free preference parameters (horizontal axis) and a restricted model with one preference parameter restricted to equal the cross-sectional mean from the unrestricted model (vertical axis). The heat maps refer to 4,151 groups of Swedish households, size-weighted to recover the underlying distribution across households under the assumption that preferences are homogeneous within groups. The left column of heat maps restricts RRA, the middle column restricts the TPR, and the right column restricts the EIS. The measures of model fit are the scaled root mean squared error (RMSE) for WY in the top row, the scaled RMSE for RS in the middle row, and the RMSE-scaled objective function (OF) in the bottom row. Each axis label shows the upper cutoff value of the corresponding bin, except for the label Inf which indicates that the bin has no upper cutoff but contains all values above the previous bin's cutoff. The logarithmic color scale is the same as in Figures 4 and 5 and indicates the fraction of the sample in each square of the heat map.

Figure 7: Impact on Model Fit of Imposing Homogeneity of Multiple Preference Parameters



(Figure note on the next page.)

This figure presents bivariate heat maps for measures of model fit comparing the unrestricted model with three free preference parameters (horizontal axis) and a restricted model with two or three preference parameters restricted to equal their cross-sectional means from the unrestricted model (vertical axis). The heat maps refer to 4,151 groups of Swedish households, size-weighted to recover the underlying distribution across households under the assumption that preferences are homogeneous within groups. The left column of heat maps restricts both the EIS and the TPR but leaves RRA free, while the right column restricts all three preference parameters. The measures of model fit are the scaled root mean squared error (RMSE) for WY in the top row, the scaled RMSE for RS in the middle row, and the RMSE-scaled objective function (OF) in the bottom row. Each axis label shows the upper cutoff value of the corresponding bin, except for the label Inf which indicates that the bin has no upper cutoff but contains all values above the previous bin's cutoff. The logarithmic color scale is the same as in Figures 4 and 5 and indicates the fraction of the sample in each square of the heat map.