Household Portfolio Allocation in the EU

Andrew Davis — Acadia University

Jiseob Kim — Yonsei University

June 18, 2017*

Abstract

We exploit the Eurosystem Household Finance and Consumption Survey to analyze household portfolio choices across the EU. Largely, the basic patterns of household portfolio allocation amongst the participating countries mirror those found in the US, for example in the U-shaped life-cycle profile of the share of the household portfolio comprised of 'riskier' assets.

We find that there is significant across-country hetereogeneity in risky asset shares, and show that country-specific factors including the provision of social insurance and the fiscal strength of government as measured by the price of default insurance predict household asset allocation choices, something that cannot be observed in the US-centric household portfolio literature.

Finally, we examine whether several US-data results on household portfolio allocation hold in this European dataset.

*Andrew Davis: andrew.davis@acadiau.ca. Jiseob Kim: jiseob83@gmail.com. We thank participants at the 2016 Canadian Economics Association Meetings for helpful comments and advice, particularly Sally Shen and Louis Raes. Any and all mistakes are our own.

1 Introduction

Significant attention in the literature has been paid to understanding the decisions households make in terms of their personal portfolios. Is wealth held in real or financial assets? In safer or riskier assets? Increasingly, there is a growing suite of options and vehicles for a household's savings, and the difficulty of accessing them continues to decline with the proliferation of online discount brokerages, financial education and literacy,

Abstractly, what might influence household portfolio asset allocation? Certainly, some households will be more risk tolerant than others. But the personal circumstances of households might also play a significant role. For example, consider two households of the same age, same income, and same assets. However, suppose they have different expectations of the volatility of their future income - very naturally, one might imagine the household with less certainty over future income might choose to save using less risky vehicles.

Some such components of the household decision to hold a certain mixture of assets are well identified and understood. For example, there is a robust consensus that risky asset portfolio shares - the fraction of the household portfolio held in financial assets that are considered 'risky', typically dominated by equities[1] - are increasing in wealth, e.g. Wachter and Yogo (2010); Calvet and Sodini (2014). Health status (Hugonnier, Pelgrin, St-Amour (2012)), marital and family status (Love (2010)), debt (Becker and Shabani (2010)), mental health (Bogan and Fertig (2013), Lindeboom and Melynchuk (2015)), and labor market uncertainty (Chang, Hong, Karabarbounis (2015)) have also been shown to contribute to a household's asset allocation decision-making.

The extent of this literature demonstrates the importance and interest in the topic, but the discussion is dominated by the Federal Reserve's Survey of Consumer Finances (SCF). While the SCF contains rich detail on household portfolios and is available over several

[1]Definitions may differ somewhat by dataset. See section 2 for details.

iterations (although with only a very limited panel element), being limited to one country prevents it from shedding insight on questions which might plausibly be a function of country-specific effects. Take Maurer, Mitchell, and Rogalla (2010), a brief theoretical discussion on the impact of social security on household portfolio allocation. Given that social security in the US is a federal program, the SCF provides relatively little ability to examine household portfolios under different social insurance systems, since we observe no variation in social security treatment - all members of the SCF are members of the same social security system.

There have been other datasets used in this literature, including the Health and Retirement Survey, but they typically again face the same problem - no cross-country variation, which means no ability to say anything about national-level policies or economic conditions.

This paper investigates household portfolio allocation decisions using a newly available European data set - the Eurosystem Household Finance and Consumption survey, hereafter the EU HFCS, which captures similar data to the SCF across fifteen European countries[2]. We find that, broadly, the basic descriptive patterns of how households carry out asset allocation are similar in the EU and the USA, for example in how households generally choose to increase risky asset holdings over the lifecycle and begin to wind them down sometime after retirement, generating a vaguely U-shaped profile.

We also show that country-specific effects, such as how much social insurance is provided by federal government and expectations over the government's ability to meet future promises as measured by the price of default insurance on sovereign debt, also have tangible influences on how a household chooses to save. When the government provides more social insurance, and when savers have more confidence in the government's fiscal situation, households shift towards investing in more risky assets.

Intuitively, as motivated by the mentioned literature that discusses the connection of

[2]The structure of the survey will be further discussed in the next section. A full backgrounding on the EU HFCS can be found here `https://www.ecb.europa.eu/pub/economic-research/research-networks/html/researcher_hfcn.en.html`

health, debt, wealth, family, and labor force risk to a household's asset allocation, we think of the true household portfolio comprising real assets, financial assets, and intangible assets. What we observe in terms of financial asset allocation choices is the result of an unobserved maximization over risk and return along multiple dimensions. If some countries have more social insurance than others, then as we show, the individual treats that as a substitute for safe assets and shifts more of the financial component of their portfolio towards risky assets.

We also connect our results to the literature previously mentioned by examining whether the US-centric results of the existing household portfolio allocation literature are robust to our European dataset.

Beyond these papers, Arrondel et al (2014) is by far the most closely related work to our own, a wide-ranging and mostly descriptive survey of what the HFCS says about household portfolio allocation. We have a much narrower focus, looking primarily at the share of financial assets held in risky instruments, doing so in more detail, and making the connection to the literature on household portfolios derived from the SCF.

Other literature that deals with household portfolio allocation questions includes Campanale (2011), who rationalizes observed patterns of risky asset holding over the life cycle through ambiguity aversion. However, presumably ambiguity aversion would not vary systematically across developed countries, and thus

Going forward, we hope to construct a formal, structural model of the interaction between government policies and household portfolio decisions as a test of whether our proposed channels for the empirical results are plausible. We also hope to motivate more use of the EU HFCS data in general.

2 Household Portfolio Data in the HFCS

The main data source in this paper is the Eurosystem Household Finance Consumption Survey (HFCS). The HFCS covers 15 different eurozone countries, comprising about 62,000 households. Currently, the survey consists of a single cross-section, which was collected largely between 2009 and 2011. The survey is administered at the country level by a variety of different statistical agencies and then collated and released by the Household Finance and Consumption Network as an agency of the European Central Bank.

While the sample is dispersed across 15 countries, the actual weight of the sampling is concentrated in a few selected countries. France comprises about 24% of the sample, Finland about 17.5%, Italy 12.8%, and Spain 10% - collectively almost two-thirds of respondents. Conversely, Luxembourg, Malta, and Slovenia all recorded less than a thousand households.

One immediate concern with the data is that while the questions asked where identical across countries, there were some differences in how surveying was done, as each country participating in the survey collected data somewhat differently. Countries differed in terms of how computerized the interviews were, how much training the interviewers were given, and response rates varied dramatically, from a low of 18.7% in Germany to 82.2% in Finland. In practice, however, there is nothing we can do to mitigate these problems or bound their magnitude.

The HFCS is multiply-imputed five times, though the number of missing elements is quite small for most variables. All standard errors throughout the paper are thus calculated following Rubin (1987).

Our goal is to analyze household portfolio choices. Unfortunately, we cannot identify household asset holdings as precisely as we might want to, given the constraints of the data - we cannot see every security holding, but only the amount reported in various categories. This is similarly a limitation in the SCF.

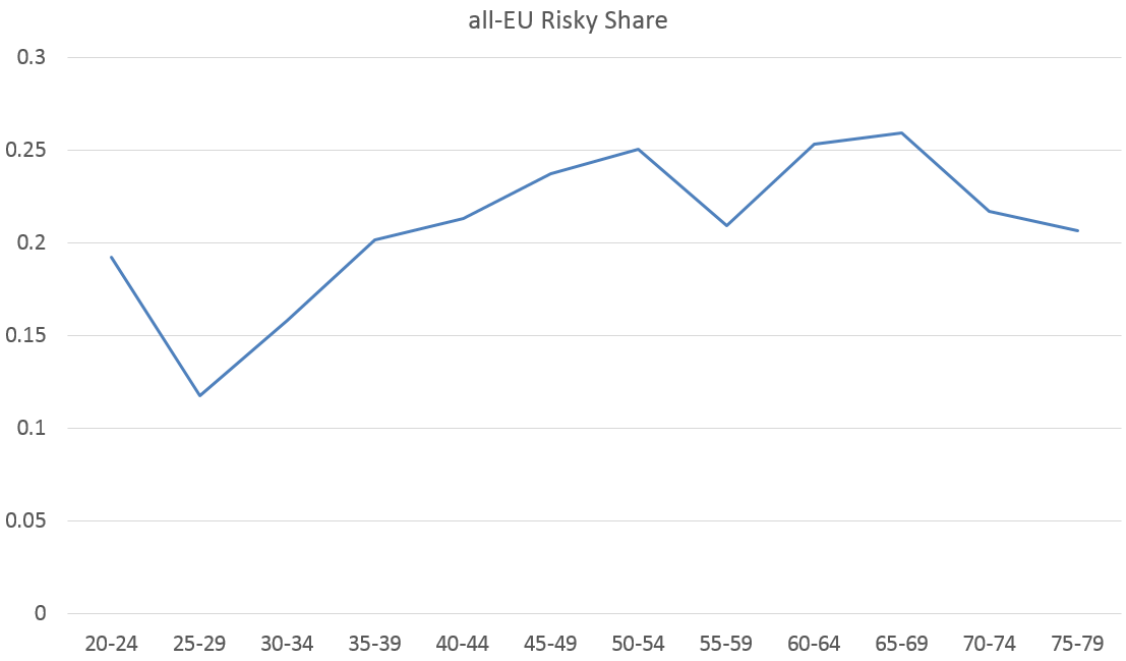

Figure 1: Fraction of all household financial assets held in 'risky' vehicles, by age group, EU HFCS and author's calculations.

To that end, mimicking commonly used definitions in the SCF, we define *risky financial assets* as the sum of mutual funds, publicly traded equity, investment accounts managed by a third party, and private non-ownership stakes in firms. That leaves the sum of cash, deposits, bonds, owed debts, private pension or life insurance assets, and miscellaneous assets as *safe financial asset* holdings. Clearly, this may not be perfect - managed accounts could contain cash, for example - but is the best classification we can make conditional on the information we have and in terms of maintaining comparability to calculations made using the SCF.

One well-known stylized fact from the SCF literature (and other American sources) is that life-cycle profiles of risky asset shares - the fraction of the household's financial portfolio held in these risky assets - follows an increasing and U-shaped pattern, slowly growing throughout the life cycle and beginning to decline sometime after retirement. This is immediately seen in the EU HFCS data as well, as presented in Figure 1.

If we plotted the conditional asset share profile - when we would only include households

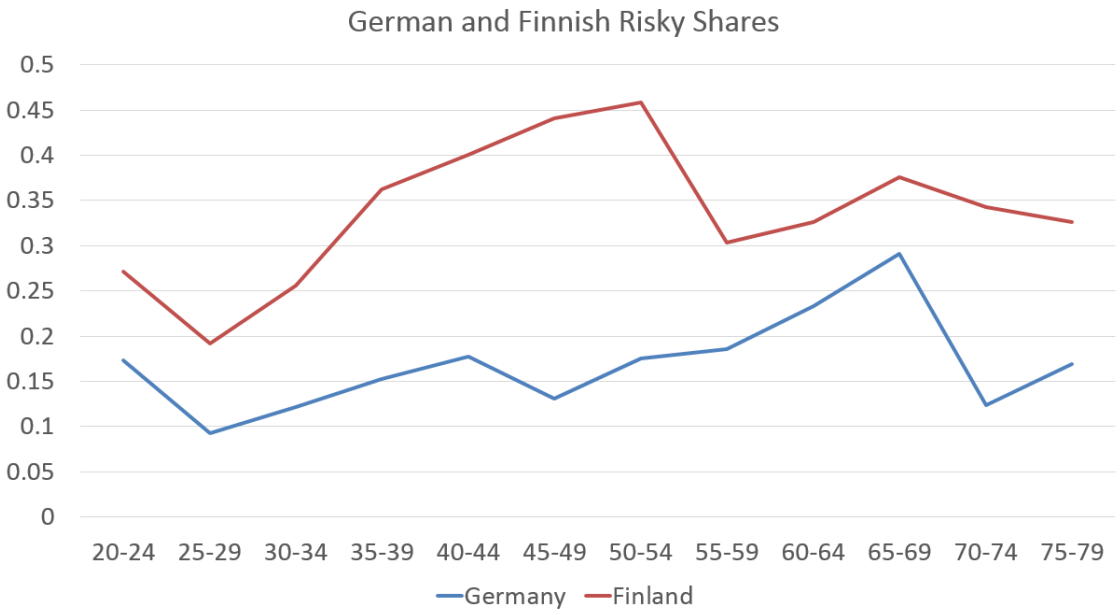

Figure 2: Fraction of all household financial assets held in 'risky' vehicles, by age group, Finland and Germany subsamples, EU HFCS and author's calculations.

that hold some assets - the shape of the profile looks effectively identical, except shifted up about 0.1 as simply a bunch of zeros are being dropped. The shape here is entirely consistent with the literature coming out of the SCF, e.g. as per Chang, Hong, and Karabarbounis (2015) and elsewhere.

There is, however, a nontrivial degree of disparity across countries. For a motivating example, consider the life-cycle risky shares of German and Finland, presented in Figure 2. We can see immediately that the general profiles are of the same shape - increases over the lifespan, some winding down after retirement - much as we see in the US data.

There are, of course, many, many ways that we could present descriptive statistics out of the EU HFCS. However, as mentioned, Arrondel et al (2014) have done that at considerable length, and thus we only present key motivating pieces from the survey.

As motivation for our regression exercises, consider the following figure, where we plot basic OECD data on government spending against country average

We now move on to a series of basic regressions based around exploiting cross-country

Figure 3: Fraction of all household financial assets held in 'risky' vehicles, by age group, Finland and Germany subsamples, EU HFCS and author's calculations.

variation in household portfolio holdings.

3 Estimation

We now move to carrying out some estimation. We use standard multiple-imputation methods for all reported regressions. Since we have significant data across numerous countries, unlike in the SCF, we can examine the impacts of government level policy on asset allocation decisions. Abstractly, a household's true portfolio includes many types of assets. Even if two households earn the same income, expect the same future income, and hold the same amount of savings, many factors could influence their savings choices, as mentioned in the introduction. If the government provides a large safe asset on the household balance sheet through a very generous social safety net, the household may feel comfortable taking on more risk in its personal savings.

In order to tackle the questions we are particularly interested in - whether or not government policy influences household portfolio allocation decisions - we need to supplement the HFCS. We collate country-level data on various government expenditures from the OECD and on credit default swap (CDS) prices on sovereign debt. In the process, we drop Cyprus, Malta, and Luxembourg on the basis of data issues - either the OECD does not report a breakdown of government spending, or CDS data is not available. None of these drops significantly affect our sample size.

We estimate a series of specifications that try to illustrate the impact of federal-level decisions and policies on household choices. Our estimation takes the form of:

$$rshare_{ij} = \beta X_{ij} + \gamma Z_j + \epsilon_{ij}$$

where $rshare$ denotes the fraction of a country-j household-i financial portfolio held in risky assets, X is a set of household-specific characters, and Z is a set of country-wide characteristics. Table 1 below reports coefficients and standard errors for a set of different specifications.

Wealth and income refer to the household as a whole. Age refers to the age of the survey respondent, who is presumably more likely to be involved with the financial matters of the household. Degree also refers to the respondent, and takes value 1 if they have a university degree, and 0 otherwise. We cannot use the more traditional years of education, as education in the survey is given as a 4-step categorical variable based upon UN ISCED coding. Attitude refers to self-reported willingness to take 'above average' risks, which is coded as 1, and 0 otherwise. % social GDP measures the total fraction of GDP spent by government (old age, survivor and death, disability, health, family, labor market, and housing benefits). CDS price measures the average basis point cost of buying insurance on 5-year sovereign debt, taken over the last 18 months to the end of 2009, trying to capture attitudes at the time of

Table 1: Regressions on household portfolio risky shares

	(1)	(2)	(3)	(4)	(5)
Wealth (000s)	0.00002	0.00002	0.00002	0.00002	0.00002
	(5.8E-6)***	(5.6E-6)***	(5.6E-6)***	(5.3E-6)***	(5.3E-6)***
Income (000s)	0.0006	0.0006	0.0005	0.0004	0.0004
	(7.4E-5)***	(7.6E-5)***	(7.6E-5)***	(6.9E-5)***	(6.8E-6)
Age	-	0.0003	0.002	0.0005	0.0005
	-	(9.7E-5)***	(0.006)***	(9.5E-5)***	(9.5E-5)***
Ageˆ2	-	-	-0.00001	-	-
	-	-	(5.7E-6)**	-	-
Degree	-	-	-	0.050	0.049
	-	-	-	(0.0047)***	(0.0047)***
Attitude	-	-	-	-	0.047
	-	-	-	-	(0.0082)***
% social GDP	0.233	0.234	0.236	0.192	0.203
	(0.034)***	(0.034)***	(0.034)***	(0.032)***	(0.034)***
CDS price (100bps)	-0.019	-0.020	-0.021	-0.013	-0.020
	(0.0049)***	(0.0049)***	(0.0049)***	(0.0047)***	(0.0050)***

Stars on standard errors (bracketed) denote statistical significance as follows: *: $p < 0.9$, **: $p < 0.05$, ***: $p < 0.01$.

the survey.

These results are generally robust to the inclusion of a much wider set of controls. Note that the coefficients on our key variables of interest are quite common. Unsurprisingly, we immediately mirror the results out of the prior literature on wealth and income. The more affluent a household is, the larger fraction of their household financial portfolio they dedicate to riskier asset classes.

It may be hard to parse what some of these coefficients imply. Remember that our risky share variable is measured as a percentage of the household's financial portfolio. Consider moving a household from Italy to Finland. Italian debt costs 80bps more to insure against in our sample, and Finland spends an additional 14.6% of GDP on social programs. Using specification (5) from the table, the conditional expectation is that such a household, holding wealth, income, age, education, and risk preference constant, would invest an extra 2.96%

into risky assets due to the extra social insurance provided by government, plus 1.6% more due to the inherent riskiness of government promises: 4.56% in total.

However, remember that as a baseline, only approximately 20% of financial assets in our survey are defined to be in risky asset classes. Therefore, an increase of 4.56 percentage points in a risky share represents an increase in risky asset holdings of approximately 23% relative to average, which we take as a very significant magnitude.

For comparison, consider the coefficients on wealth. As mentioned, the connection between wealth and riskier asset holdings has been well documented. However, an increase in wealth of a million euros - a rather significant increase in assets - only corresponds to approximately the same effect on savings behavior as a 100 basis points change in the cost of insuring 5-year sovereign debt.

4 Extending Prior Results

From our first results, we already see that household risky asset portfolio shares are increasing in wealth, as already found in US data by Wachter and Yogo (2010) and Calvet and Sodini (2014), among others. This is not a surprising result, and is rationalized in these papers and elsewhere.

We now examine whether several other results in the household portfolio hold, given the constraints of our data in the EU HFCS. Unfortunately, we do not have the ability to attempt to replicate and extend certain studies previously mentioned, notably those dealing with health issues - the information simply isn't collected in the HFCS.

4.1 Family Status (Love 2010)

The potential connection between family status and portfolio choices is quite intuitive - families are both large financial assets and large financial liabilities on the household balance

sheet, potentially necessitating adjustments on other parts of the balance sheet to achieve the desired asset mix. For example, wanting to save for a kid's future education may shift towards holding more in financial assets than a childless couple who save more in home equity.

Love (2010), unlike most of the papers we discuss, actually relies on stockholding information from the PSID (Panel Survey of Income Dynamics) and the HRS (Health and Retirement Survey), not the SCF. This makes extending his results somewhat imprecise - we cannot precisely map HFCS variables into PSID/HRS variables, nor do we have access to a panel element that we can exploit.

Love (2010) tackles personal portfolio allocation by dividing stocks (which comprises directly held stocks, mutual funds, and investment trusts), by net financial wealth, defined as 'stockholding share'. The mapping from the HRS to the EU HFCS cannot be perfect, but we feel this roughly approximates our 'risky asset share'. As evidence that the two measures are close cousins of each other, Love plots his measure of asset allocation by age in Table 4 and finds a very similar U-shaped pattern to the EU life-cycle profile of risky asset shares, as is consistent throughout the asset allocation literature.

While most of Love (2010) is a formal modeling exercise, we are mostly interested in the empirical results. Comparison should be drawn to his Table 9 and Table 10. We first tackle asset allocation by looking at the effects of marriage. In our Table 2, Married refers to marriages and legally recognized civil unions. Given that Love focuses on interaction terms between marital status and education, we expand the educational variables to include a 'no high school (HS)' variable. Given the UN-standardized coding in the dataset, our 'no HS' variable refers generally to students who dropped out of school before the age 16. Other variables are defined as in the previous section.

Our results, compared to Love (2010), are generally consistent with respect to marriage. He finds mixed, largely insignificant effects of marriage on stockholding, and here we find

Table 2: Regressions on household portfolio risky shares: Marriage and divorce

	(1: OLS)	(2:OLS)	(3:Tobit)
Wealth (000s)	0.00002	0.0002	0.00005
	(5.2E-6)***	(5.1E-6)***	(8.7E-6)***
Income (000s)	0.0005	0.0004	0.0019
	(7.3E-5)***	(7.0E-5)***	(0.0003)***
Age	0.0005	0.0007	0.0027
	(9.6E-5)***	(0.0001)***	(0.0005)***
Married	-0.0054	-	-
	(0.0038)	-	-
Divorced	-0.0052	-	-
	(0.0064)	-	-
Married*No HS	-	0.0092	0.0059
	-	(0.0038)**	(0.017)
Married*Degree	-	-0.0050	-0.054
	-	(0.0094)	(0.031)*
no HS	-	0.032	-0.233
	-	(0.0038)***	(0.020)***
Degree	0.0052	0.050	0.241
	(0.0047)***	(0.0076)***	(0.017)***

Stars on standard errors (bracketed) denote statistical significance as follows: *: $p < 0.9$, **: $p < 0.05$, ***: $p < 0.01$.

Table 3: Regressions on household portfolio risky shares: Widowed

	(1: OLS)	(2:OLS)	(3:Tobit)	(4:Tobit)
Wealth (000s)	0.00002	0.00002	0.0005	0.00005
	(5.4E-6)***	(5.4E-6)	(6.6E-6)***	(6.6E-6)***
Income (000s)	0.0004	0.0004	0.0015	
	(6.9E-5)***	(6.9E-5)	(0.0002)***	
Age	0.0007	0.0007	0.002	0.002
	(0.0001)***	(0.0001)***	(0.0006)***	(0.0006)***
Widow	-0.001	-	-0.049	-
	(0.0061)	-	(0.029)***	-
Widow*Degree	-	-0.014	-	-0.132
	-	(0.005)***	-	(0.039)***
no HS	-0.025	-0.021	-0.210	-0.172
	(0.0039)***	(0.0039)***	(0.022)***	(0.023)***
Degree	0.052	0.052	0.240	0.242
	(0.0058)***	(0.0058)***	(0.243)***	(0.019)***

Sample includes only the married and the widowed. Stars on standard errors (bracketed) denote statistical significance as follows: *: $p < 0.9$, **: $p < 0.05$, ***: $p < 0.01$.

similarly with respect to risky asset holding. Note that the coefficients for wealth, income, age, and degree are largely unchanged.

However, Love's least ambiguous empirical results are for the widowed. He argues that widowers significantly reduce their risky asset holdings. Restricting our sample to the married and the widowed, we see whether such a story appears in the HFCS. Given that the HFCS reports widowed as a specific response to their marital status question, we can identify widowers precisely.

The resulting Table 3 raises some questions. Simply throwing in a dummy for those reporting widowed has mixed results depending on whether we use a simple OLS or a tobit. We adopt the tobit here because Love relies on a random effects tobit. Since we have no panel element, we rely on simple tobit. However, Love (2010) reports only coefficients on widowed interacting with education, and we see those are indeed significantly negative, as he finds. Specification (4) here is closest to his results. Note that the interpretation of the

coefficients changes in the tobit model.

Overall, we consider this somewhat mixed but overall positive evidence in support of Love (2010). There is nothing that we find that's starkly in contrast to his results. However, while approximate replication of his specification here generates similar effects, the magnitude of the coefficients suggests that marital status is not that important a determinant of household portfolio allocation, especially compared to education or our cross-country variables from the prior section.

4.2 Labor Market Uncertainty (Chang, Hong, Karabarbounis 2015)

Much like Love (2010), Chang, Hong, and Karabarbounis (2015), hereafter CHK, is again mostly a model-theoretic exercise. We simply attempt to see if our data holds any support for whether or not labor market uncertainty can actually explain any variation in household portfolios.

One extremely simple way to proxy for future labor market uncertainty would be to use age. As households age, much of the uncertainty with respect to their labor market opportunities and outcomes is resolved. This suggests that older households, with less labor market risk in their effective portfolio, should be open to holding riskier assets. As seen earlier in the paper, this is exactly what we find.

However, the casuality is not clear. Age could be capturing other things, such as increased knowledge of financial markets. A slightly more sophisticated approach - given the limitations of the data we have available in the EU HFCS - would be to actually look at household income volatility. Unfortunately, since the survey has no panel element, this is not something we can measure directly.

To try and examine the impact of labor market uncertainty, we make use of several other variables that may represent household stability and expectations over the future. In particular, we use whether a household considered their income to be 'normal' in the survey

Table 4: More regressions on household portfolio risky shares

	(1: OLS)	(2:Tobit)
Wealth (000s)	9.0E-6	0.00002
	(3.8E-6)**	(9.5E-6)*
Income (000s)	0.0005	0.0016
	(8.4E-5)***	(0.0003)***
Age	0.0004	0.0012
	(0.0002)*	(0.001)
Income Shock	0.011	0.068
	(0.0042)***	(0.018)***
Temporary Job	-0.0015	-0.090
	(0.008)	(0.047)*
Job Tenure	0.0009	0.005
	(0.0002)***	(0.0009)***
no HS	-0.019	0.172
	(0.005)***	(0.030)***
Degree	0.046	0.226
	(0.006)***	(0.020)***

Sample includes only the married and the widowed. Stars on standard errors (bracketed) denote statistical significance as follows: *: $p < 0.9$, **: $p < 0.05$, ***: $p < 0.01$.

year, whether their employment is permanent or temporary, and how long they've been working at their current job.

Results are seen in Table 4. In particular, 'Income Shock' takes a value of 1 for households that reported income to be either above or below expectations, suggesting general volatility in income, and 0 if income was in line with their expectations. 'Temporary Job' refers to whether or not the respondent is on a temporary (instead of a permanent) job contract.

Obviously, none of these proxies are (even close to) perfect. However, the signs raise some concerns. Households on the receiving end of an income shock - suggesting more labor market uncertainty - actually invest more into risky assets, opposite of our priors on the coefficient. Whether or not the job is temporary is insignificant, whereas we would have ex-ante suspected that those on permanent contracts would be more comfortable investing

more in risky financial vehicles, given that they have a large, safe asset in their job already on their balance sheet.

We do find that job tenure is positively correlated with increased allocation towards riskier assets in the household's financial portfolio, in line with our priors. With job tenure presumably comes learning about the career path, and less uncertainty about future financial rewards in the labor market. With that risk increasingly resolved, the household feels comfortable shifting into a more aggressive financial portfolio.

In sum, we find mixed, weak evidence in favor of a labor market uncertainty hypothesis. Given the limitations of our data, however, we cannot speak strongly on the issue, as none of our proxies are really a perfect measure for how much uncertainty is in a household's expectations over their future life-cycle income.

4.3 Debt (Becker and Shabani 2010)

To be completed.

5 Extension: A Formal Model

The empirical exercises undertaken in the prior two sections are not complicated, with nothing being done in the way of formal modelling. Several papers have written down formal models of how households might make portfolio decisions, including Love (2010) discussed above. One reason for us to write down a formal model would be as a way to try and achieve a sort of experimental identification - we don't have an extremely compelling casual story coming through our data exercise, but if we wrote down a model might provide some circumstantial evidence if the model results were in the ballpark

6 Conclusion

We exploit the Eurosystem Household Finance and Consumption Survey to analyze household portfolio choices across the European Union. We find that there is significant across-country hetereogeneity in the fraction of financial assets a household chooses to hold in risky instruments, and that the hump-shaped life-cycle profile of these risky asset shares observed in the US data is largely mirrored in Europe, thus extending the general reach of several results that deal with the American profile of risky asset holdings over the life cycle.

In particular, we show that the degree of social insurance provided through government influences asset allocation, as does the fiscal strength of government as measured by the price of insuring sovereign debt. This suggests households are actively making sophisticated financial decisions incorporating expectations over future government promises. We believe our results here are entirely new to the household portfolio literature. That said, the magnitude of the coefficient here - from either specification - is quite small, despite its statistical significance, and it does not appear to be a significant channel of action.

We also test whether several other American-centric results, typically based on the Survey of Consumer Finances or the Health and Retirement Survey, hold across the EU: that risky asset shares are a function of wealth (e.g. Calvet and Sodini 2014), debt (Becker and Shabani 2010), and labor market uncertainty (Chang, Hong, and Karabarbounis 2015). Collectively, our evidence on whether these results hold more generally beyond the US is mixed.

References

[1] Ameriks, J., and Zeldes, S. P. How do household portfolio shares vary with age. Tech. rep., working paper, Columbia University, 2004.

[2] Arrondel, L., Bartiloro, L., Fessler, P., Linder, P., Mathä, T. Y., Rampazzi, C., Savignac, F., Schmidt, T., Schürz, M., and Vermeulen, P. How do households allocate their assets? stylised facts from the eurosystem household finance and consumption survey.

[3] Becker, T. A., and Shabani, R. Outstanding debt and the household portfolio. *Review of Financial Studies* (2010), hhq023.

[4] Bogan, V. L., and Fertig, A. R. Portfolio choice and mental health. *Review of Finance 17*, 3 (2013), 955–992.

[5] Calvet, L. E., and Sodini, P. Twin picks: Disentangling the determinants of risk-taking in household portfolios. *The Journal of Finance 69*, 2 (2014), 867–906.

[6] Campanale, C. Learning, ambiguity and life-cycle portfolio allocation. *Review of Economic Dynamics 14*, 2 (2011), 339–367.

[7] Chang, Y., Hong, J. H., and Karabarbounis, M. Labor-market uncertainty and portfolio choice puzzles.

[8] Hugonnier, J., Pelgrin, F., and St-Amour, P. Health and (other) asset holdings. *The Review of Economic Studies* (2012), rds033.

[9] Lindeboom, M., and Melnychuk, M. Mental health and asset choices. *Annals of Economics and Statistics*, 119-120 (2015), 65–94.

[10] Love, D. A. The effects of marital status and children on savings and portfolio choice. *Review of Financial Studies 23*, 1 (2010), 385–432.

[11] Maurer, R., Mitchell, O. S., and Rogalla, R. The effect of uncertain labor income and social security on life-cycle portfolios. Tech. rep., National Bureau of Economic Research, 2010.

[12] Rubin, D. B. *Multiple imputation for nonresponse in surveys*. John Wiley & Sons, 1987.

[13] Wachter, J. A., and Yogo, M. Why do household portfolio shares rise in wealth? *Review of Financial Studies 23*, 11 (2010), 3929–3965.

A Additional Information on the Eurosystem Household Finance and Consumption Survey

The European Central Bank, which collates and releases the EU HFCS from individual member nations, does not freely distribute the data. Data may be made available based on internal judgment and privacy conditions following submission of a project proposal to the ECB.