Wealth and Risk-Taking:
A Cross-Country Study of European Households

Master's Thesis

Department of Finance

Stockholm School of Economics

Gustaf Östblom*

December 2022

Abstract

This study investigates the relationship between wealth and risk-taking among European households. Using data from 22 European countries, I show that households with more investable wealth are more likely to invest in risky assets and have a higher risky share conditional on participation. I also show that the elasticity of investable wealth with respect to the risky share is positive and that it increases slightly with investable wealth, which suggests that the share of investable wealth invested in risky assets is an increasing and convex function of investable wealth. Furthermore, I document the importance of considering commercial real estate as an investment vehicle when drawing conclusions on the risky share, as opposed to focusing strictly on financial assets. The results are indicative of decreasing relative risk aversion among European households.

Tutor: Paolo Sodini

Acknowledgements: I would like to thank my tutor Paolo Sodini, Professor of Finance at the Stockholm School of Economics, for providing valuable guidance, assistance, and constructive feedback. I would also like to thank August Hansson, Senior Research Analyst at the Swedish House of Finance, for his guidance in managing the code used in this project. Lastly, I want to thank Frida for her endless encouragement and support.

*24002@student.hhs.se

Contents

Abbreviations

AT	Austria
BE	Belgium
CY	Cyprus
DE	Germany
EE	Estonia
ES	Spain
FI	Finland
FR	France
GR	Greece
HR	Croatia
HU	Hungary
IE	Ireland
IT	Italy
LT	Lithuania
LU	Luxembourg
LV	Latvia
MT	Malta
NL	Netherlands
PL	Poland
PT	Portugal
SI	Slovenia
SK	Slovakia

Acronyms

CRRA	Constant relative risk aversion
DRRA	Decreasing relative risk aversion
HFCN	Household Finance and Consumption Network
HFCS	Household Finance and Consumption Survey

1 Introduction

Households hold a significant portion of the financial assets on the market in aggregate. Yet, households have traditionally not had a substantial presence in the area of financial economics when compared to corporations (Tufano, 2009). Although many areas within corporate finance are applicable to households as well, there are several characteristics that set them apart. Households hold a large portion of their wealth in human capital, which is a non-traded asset that can be defined as the present value of labor income that the household expects to earn. Households are also subject to restrictive borrowing constraints and information barriers, and they are often challenged by a limited understanding of financial instruments. In addition, the financial decisions of households are shaped by institutions that are not the focus of corporate finance. In recent years, however, the area of household finance has experienced an increase in academic interest. First, households are becoming more directly engaged in financial decisions, and financial innovation has resulted in a wider array of products available. Second, there has been a substantial growth in the availability of data on household finances (Guiso and Sodini, 2013).

An important financial decision that households face is how much to invest in risky assets. It is well-documented that household portfolios are often under-diversified and that many households invest less in risky assets than suggested by normative models. This puzzle has led to numerous studies trying to identify the determinants of risk-taking among households. One characteristic that has garnered particular attention is wealth, since the relationship between wealth and risk-taking carries information regarding relative risk aversion. Constant relative risk aversion (CRRA), a key feature of the power utility function, implies that wealth does not affect the optimal allocation to risky assets. However, empirical findings such as counter-cyclical risk premiums are difficult to reconcile with CRRA. Therefore, decreasing relative risk aversion (DRRA) has become a popular alternative as is predicts a lower aversion to risk when wealth increases. Although some studies have provided empirical evidence of DRRA, fewer studies have looked at cross-country data in order to control for potential heterogeneity between countries.

This study uses a rich set of cross-sectional data from 22 European countries to study the relationship between investable wealth, defined as the sum of financial wealth and commercial real estate wealth, and household risk-taking. Specifically, I look at participation, i.e. the decision to own risky assets, as well as the share of investable wealth allocated to risky assets. The source of data is the Household Finance and Consumption Survey (HFCS), which is administered by the European Central Bank. The survey is unique in that is provides harmonized data from a large number of countries. Furthermore, the HFCS employs a complex survey design in order to minimize bias.

I find that investable wealth has a positive and highly significant effect on risk-taking in household portfolios, both in terms of participation and in terms of the risky share conditional on participation. The relationship holds when controlling for financial and demographic characteristics as well as risk preferences. Since the study controls for risk preferences, the impact of investable wealth on risk-taking is difficult to reconcile with CRRA utility. Hence, the study provides additional empirical support households generally express decreasing relative risk aversion. I also find that the elasticity of investable wealth with respect to the risky share is positive and that it increases slightly with investable wealth. In other words, the effect of a proportional shock to investable wealth on the risky share is greater for wealthier households. This suggests that the risky share could be is an increasing and convex function of investable wealth.

2 Literature Review

2.1 Risk-Taking and Utility

An important mechanism affecting the demand for risky assets is the functional form of utility. Many models rely on the assumption of constant relative risk aversion (CRRA). The so-called power utility function displays CRRA:

$$U(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma} \tag{1}$$

where C_t is consumption in period t and γ is a risk aversion parameter. Furthermore, relative risk aversion is defined as

$$\text{RRA} = -\frac{C_t U''(C_t)}{U'(C_t)} \tag{2}$$

and aims to capture how averse an investor is to a proportional change in consumption. As implied by the name, CRRA utility functions such as (1) are reduced to constants when substituted into (2). In other words, CRRA utility implies that economic agents enjoying different levels of consumption are equally averse to a proportional shock to consumption. A popular alternative to CRRA utility is decreasing relative risk aversion (DRRA) utility. The defining characteristic of DRRA utility is that RRA decreases with consumption (Guiso and Sodini, 2013).

Whether investors display CRRA utility or DRRA utility has significant impacts on the demand for risky assets. As shown by Samuelson (1969), investors with CRRA utility will choose the same allocation of assets regardless of wealth. Based on this insight, Merton (1969) derived household i's optimal risky share ω_i based on (1):

$$\omega_i = \frac{E\left[r_i^e\right]}{\gamma_i \sigma_i^2} \tag{3}$$

where $E\left[r_i^e\right]$ is the expected excess return of risky assets and σ_i^2 is the expected return variance. Furthermore, it is often assumed in literature that investors have homogeneous expectations (Guiso and Sodini, 2013). Thus, the i subscripts in (3) can be dismissed when it comes to expected return and variance, and the expression reduces to:

$$\omega_i = \frac{E\left[r^e\right]}{\gamma_i \sigma^2} \tag{4}$$

An important implication of the framework is that all heterogeneity in the demand for risky assets among CRRA investors is explained by differences in risk aversion. In other words, financial characteristics such as wealth and income has no effect on the optimal risky share. On the other hand, a consequence of assuming DRRA utility is that wealthy investors are less sensitive to proportional shocks compared to poorer DRRA investors. DRRA therefore predicts a higher allocation to risky assets among wealthy investors (Guiso and Sodini, 2013). There are several utility frameworks that incorporate DRRA. One example is habit formation (Campbell and Cochrane, 1999). This framework incorporates the following utility function:

$$U(C_t) = \frac{(C_t - X_t)^{1-\gamma}}{1-\gamma}$$

where X_t represents habit, which can be seen as a baseline level of consumption that an economic agent wants to exceed at all costs. In other words, economic agents do not only assign utility to the level of consumption, but also based on its distance from this baseline level.

The empirical relationship between wealth and the demand for risky assets is meaningful because it carries information about whether investors display CRRA or DRRA utility. Some empirical studies have shown that there is a positive relationship between risk-taking and wealth (e.g. Calvet et al., 2007; Guiso et al., 2003a). However, as noted by Calvet and Sodini (2014), the mere presence of such a positive relationship does not refute CRRA utility in the cross-section as it could be that wealth and risk tolerance are cross-sectionally correlated. In order to distinguish between CRRA and DRRA, one needs to control for risk preferences. There are several methods by which this can be achieved. One method is to use panel data, which allows for the possibility to study the effects of changes in wealth on the risky share for a single investor over time. Calvet and Sodini (2014) uses a different approach by employing a twin study. When observing identical twins, it is possible to control for much of the variation by using a twin pair fixed effect. Since identical twins have the same genetic heritage and presumably similar upbringings, it is much less likely that differences in the demand for risky assets within the twin pair is explained by differences in personal traits and characteristics, such as tolerance for risk. The study finds a positive relationship between financial wealth and the demand for risky assets, which implies DRRA utility. They also find that the elasticity of the risky share with respect to financial wealth is positive but that it decreases with financial wealth. Thus, they conclude that the risky share is a concave and increasing function of financial wealth. If households are more averse to risk when times are bad, they demand a higher risk premium for holding risky assets when times are bad. In other words, DRRA can be reconciled with counter-cyclical risk premiums (Calvet and Sodini, 2014). Yet another method to control for risk aversion is by elicitation. This can be done qualitatively, which involves asking investors to assess different statements regarding their propensity to assume risk. It can also be done quantitatively, which entails the construction of experimental gambles from which quantified measures of risk aversion can be extracted (Guiso and Sodini, 2013).

2.2 Housing

Housing is an important asset class for households, and it also has been shown to impact financial decisions with regards to risk-taking. Home ownership provides more stable and predictable residential costs compared to renting. House prices and rents also tend to co-vary, meaning that home ownership provides a hedge against rent fluctuations if one were to decide to rent instead of owning in the future. As such, home ownership can allow for more risk-taking in financial portfolios compared to non-ownership (Sinai and Souleles, 2005). On the other hand, home ownership can also be seen as a speculative investment. Cocco (2005) shows that house price risk leads to a crowding-out effect whereby the limitations that home ownership imposes on financial wealth, especially for poor households, discourages stock market participation. These findings are further evidences by Yao and Zhang (2005), who show that investors who are indifferent between owning and renting tend to invest less in the stock market when owning. Nonetheless, Calvet and Sodini (2014) distinguish between residential real estate, which provide hedging benefits and housing services, from commercial real estate and find that only commercial real estate is negatively related to investment in risky financial assets.

It has also been shown that leverage discourages financial risk-taking. Grossman and Vila (1992) shows that leverage alters the optimal risky share in the presence of borrowing constraints. This holds even if the constraint is not yet binding; even if an investor has more room to borrow, the optimal strategy changes because of the possitiblity that the borrowing ceiling will be reached in the future. The idea that leverage discourages risk-taking was shown empirically by Calvet and Sodini (2014).

2.3 Human Capital

Human capital refers to individual attributes such as skills, knowledge, and education. These attributes are encapsulated in the ability of the individual to earn income. Therefore, human capital is often quantitatively represented by the present value of future labor income. Human capital is a unique asset for a few different reasons. First, it accumulates very slowly. Second, it is difficult to evaluate given the uncertainty in forecasting every income stream over an individual's entire lifespan. Third, it can neither be traded nor easily liquidated. Finally, it carries risk that cannot be easily hedged. Thus, individuals are often not protected against income shocks other than the possibility of receiving unemployment benefits (Guiso and Sodini, 2013). Nonetheless, it is an important asset and is often substantial for young households, since they have more future income streams to receive. Campbell and Viceira (2002) derive an optimal risky share that assumes the presence of human capital for investor i:

$$\omega_{i,HC} = \left(1 + \frac{HC_i}{W_i}\right)\omega_i - \beta_{i,HC}\frac{HC_i}{W_i} \tag{5}$$

where HC_i is the value of human capital, W_i is total wealth, ω_i is the optimal risky share without the presense of human capital as defined in (3), and β_{HC} is the sensitivity of human capital with respect to market returns. In other words, human capital increases the optimal risky share unless it has enough exposure to market risk. The intuition is that a high $\beta_{i,HC}$ means that the investor is already exposed to systematic risk through her income and need not take on as much additional exposure through financial investments. Furthermore, Cocco (2005) shows that $\beta_{i,HC}$ is close to zero for the average household, implying that human capital has an incontestably positive relationship with the optimal risky share as the second component in (5) vanishes. A resulting implication is that young households should have a higher risky share, since they have more wealth in human capital. The empirical evidence on the effect of human capital on risk-taking is not conclusive. On the one hand, Guiso et al. (2003b) show that middle-aged investors participate to a greater extent than young households, which opposes the prediction that young households should invest more in risky assets. On the other hand, Calvet and Sodini (2014) find that human capital has a positive impact on the risky share. Furthermore, Heaton and Lucas (2000) show that households with a high income volatility invest less in stocks than equally wealthy households with a lower income volatility, although they participate to a large extent. This result is expected in theory, since a higher income volatility discounts the present value of future income streams and therefore reduces the value of the human capital stock.

2.4 Participation Costs

An implication of (4) is that all households should own at least some risky assets. Non-participation can only be explained by a complete absence of a risk premium, an infinite volatility, or an infinite risk aversion, all of which are unrealistic assumptions. Yet, empirical studies have shown that many households do not own any risky assets (see e.g. Canner et al., 1997). This misalignment has been labeled the participation puzzle. Several theories have tried to explain this puzzle. Vissing-Jorgensen (2003) looks specifically at stock market participation and suggests that participation is associated with fixed costs. Monetary costs can be incurred in the shape of e.g. administrative costs when opening an investment account. In addition, informational costs are incurred as households need to spend time learning about financial markets. A prediction when introducing participation costs is that wealthy and financially literate households are more likely to participate. However, as noted by Guiso and Sodini (2013), participation costs cannot explain the magnitude of observed cross-country differences in participation.

2.5 Additional Factors

Several studies that have tried to explain the low participation in risky asset markets draw inspiration from behavioral concepts. Dimmock et al. (2016) show that aversion to ambiguity is associated with low participation. In this context, ambiguity refers to unknown probabilities of future events. Ambiguity is different from risk in the sense that a future payoff with a known probability less than one is risky but not ambiguous. The study also suggests that the negative effect of ambiguity aversion on risky asset market participation is reduced when financial literacy is high. Another proposed explanation is counter-cyclical risk aversion, which refers to the notion that risk aversion is higher when times are bad, and vice versa. Cohn et al. (2015) show that this can lead to feedback loops whereby market downturns increase risk aversion, which leads to a further reduction in the willingness to participate in risky asset markets. Furthermore, Campbell et al. (2011) suggest that individual preferences are subject to present-bias, meaning that current consumption is preferred even if it damages future welfare, which implies that households would rather consume today than invest in financial assets for future consumption. Other behavioral attributes that have been found to affect risk-taking are sensation-seeking and overconfidence Grinblatt and Keloharju (2009).

2.6 Cross-Country Comparisons

Although many studies have tried to explain the determinants of household portfolios, fewer have made cross-country comparisons. Some, however, have looked specifically at stock market participation. Using data from 1998 and 1999, Guiso et al. (2003a) looked at stock market participation in six European countries: France, Germany, Italy, the Netherlands, Sweden, and the UK. The results indicate a large variation regarding participation due to differences in perceived benefits of stock ownership, differences in perceived participation costs, and differences in actual participation costs. Guiso et al. (2003b) further investigate stockholding in France, Germany, Italy, the Netherlands, the UK, and the US. The study pools households from all countries and employs a country fixed effect in order to observe household investment decisions within countries. They find that participation rates differ significantly between the countries; however, all countries show similar patterns when relating participation to certain demographic variables. For example, wealth and education are both correlated with participation rates in all countries.

Christelis et al. (2013) expand the scope by including other household balance sheet items, mainly privately owned businesses, real estate, and mortgages. Their data includes 13 countries, 12 of which are European countries and one being the United States. Their main focus is on differences between Europe and the US, as opposed to differences within European countries. Similarly to Guiso et al. (2003a) and Guiso et al. (2003b), they find significant differences in participation probabilities, not only when it comes to stock market participation but other asset classes as well. They find that households in the U.S. invest more in stocks and less in real estate compared to demographically similar European households. U.S. households also tend to have larger mortgages. Their data, however, consists of multiple surveys for different countries, which raises questions regarding the extent to which cross-country data is comparable. Furthermore, their data only includes people aged 50 or more, thereby limiting the ability to draw conclusions regarding households of other ages. Lastly, Guiso et al. (2002) provide an overview of household portfolios in several European countries. Although not a comparative work, it includes insights regarding cross-country differences. For example, the chapter by Guiso and Jappelli (2002) suggests that the historic thinness and volatility of the Italian stock market has discouraged households from owning stocks. Therefore, direct investment in bonds has been a more common form of saving, although trends are showing that stocks are becoming more and more common. This highlights the importance of the institutional environment, as well as cultural heritage, in explaining cross-country heterogeneity.

3 Data and Methodology

3.1 Household Finance and Consumption Survey

3.1.1 Background

The Household Finance and Consumption Network (HFCN) was founded in 2006 with the purpose of creating and maintaining the Eurosystem Household Finance and Consumption Survey (HFCS), which is the source of data used in this study. The survey data includes demographic information as well as household assets, liabilities, labor income, pension and insurance policies, consumption, and more. The HFCN consists of researchers from national statistical institutions, national central banks, as well as the European Central Bank (ECB). All countries in the eurozone as well as some non-eurozone European countries participate in the survey. While the survey is coordinated centrally by the ECB and HFCN, it is decentrally conducted on a country bases by national statistical institutions and central banks. A primary target of the HFCS is to produce output variables that are harmonized across countries. To achieve this goal, the conducting institution in each country uses the same blueprint questionnaire, which is then modified in order to cater to cross-country differences. For example, questions may be adapted according to cross-country differences in availability of certain financial products. The survey has been published in three waves. Most of the fieldwork took place in 2010 to 2011 for the first wave, 2013 to 2015 for the second wave, and 2017 for the third wave.

3.1.2 Survey Design

The survey is mainly characterized by its two sections: the personal section and the household section. Information regarding demographics, employment, pensions, and life insurance policies is recorded at the personal level. All members of a sampled household aged 16 or older are asked to participate in the personal section. The member of the household who is deemed most knowledgeable with respect to the financial matters of the household is then invited to participate in the household section of the survey. This person will hereafter be referred to as the *household head*. The household section is the primary focus of this study and contains information regarding assets, liabilities, and consumption for the household as a whole. The household level sample size and the personal level sample size are reported in Table 1.

Table 1: Sample Size*

	First Wave		Second Wave		Third Wave	
	Household	Personal	Household	Personal	Household	Personal
AT	2380	5014	2997	6189	3072	6414
BE	2327	5506	2238	5200	2329	5370
CY	1237	4169	1289	4223	1303	4188
DE	3565	8134	4461	10201	4942	11251
EE	NA	NA	2220	5709	2679	6724
ES	6106	15852	6120	15536	6413	16335
FI	10989	27009	11030	27142	10210	24818
FR	15006	35729	12035	28845	13685	32799
GR	2971	7740	3003	7744	3007	7463
HR	NA	NA	NA	NA	1357	3699
HU	NA	NA	6207	14623	5968	13937
IE	NA	NA	5419	14546	4793	12778
IT	7951	19836	8156	19366	7420	16462
LT	NA	NA	NA	NA	1664	3729
LU	950	2540	1601	4444	1616	4384
LV	NA	NA	1202	2814	1249	2824
MT	843	2307	999	2703	1004	2632
NL	1301	2962	1284	2835	2556	5250
PL	NA	NA	3455	9035	5858	15017
PT	4404	11126	6207	16513	5924	15079
SI	343	964	2553	7245	2014	5405
SK	2057	5351	2135	5433	2179	5307

* Missing values represent non-participation in the survey.

The HFCS uses stratified sampling, which is a commonly used sampling technique whereby the population of interest is divided into exhaustive and mutually exclusive subgroups, also known as strata, prior to sampling.[1] The HFCS applies multi-stage stratified sampling, which means that a stratified sampling process iterates multiple times to create subgroups within each strata. For example, one can first divide a population based on geographic region, and then divide each geographic region into groups based on wealth. The criteria by which the stratification occurs is different between countries; see Household Finance and Consumption Network (2020) for a more detailed description.

Furthermore, sample weights are reported for each household. The weight assigned to an observation in a sample represents the number of units in the population that is represented by that observation (Brick and Kalton, 1996). The HFCS data is weighted

[1]See Fuller (2009) for a theoretical overview of the concept.

Figure 1: Multiple Imputation Illustration

The node to the left represents the oroginal, non-imputed dataset. The second node illustrates how several datasets are created as the multiple imputation technique is applied. Estimations are then performed on each dataset before being pooled into a final point estimate.

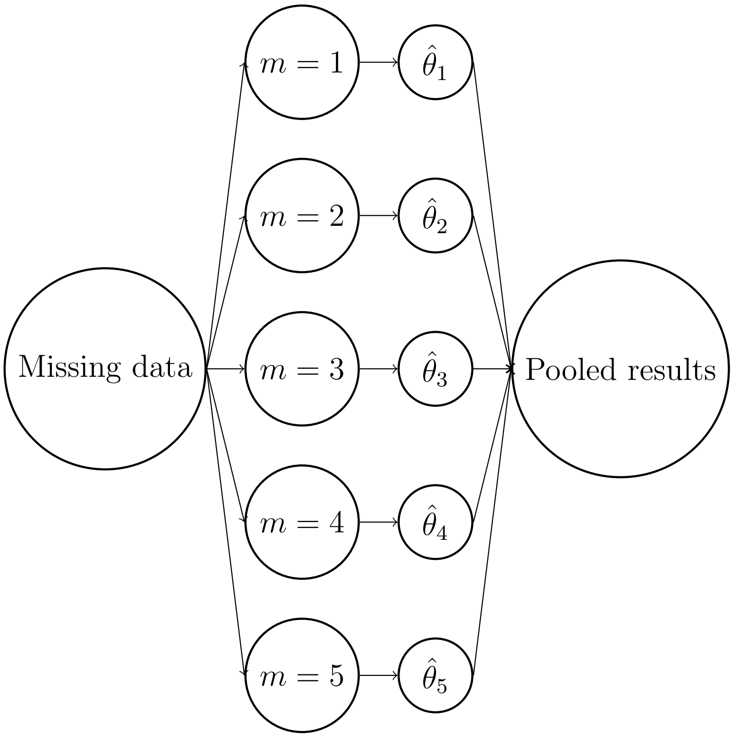

through a multi-step procedure that starts with the inverse of the selection probability. In other words, if the probability that household h of stratum s will be selected to participate in the survey is $\pi_{h,s}$, the household will be given the weight $\omega_{h,s} = \pi_{h,s}^{-1}$. These weights have been modified to account for e.g. coverage and discrepancies between characteristics of respondents and non-respondents. It also uses a calibration approach, which means that weights are calibrated using auxiliary information about the population (Särndal, 2007). For example, if it is known that a stratum accounts for a certain proportion of the population, the weights of observations sampled from that stratum are adjusted to reflect that proportion.

Another important feature of the HFCS is the use of multiply imputed data. Imputation refers to the act of replacing a missing or otherwise deficient value with an estimation. Single imputation, the simplest imputation method, means the value is replaced only once. For example, one can replace the missing value with the mean of observed values, or use a more advanced model such as a linear regression in order to predict the missing value.

Imputing missing data can reduce bias and allow for the use of complete-data methods of analysis. However, there are a few drawbacks associated with the use of single imputation. In short, a missing value is, by definition, not truly known. Thus, the method by which the value is imputed has an inherent uncertainty which needs to be taken into account when performing statistical inference. Imputing the value only once can fail to capture this uncertainty. Multiple imputation methods mediate this issue by imputing missing data multiple times using models that include an uncertainty component. Thus, the process yields multiple estimations for each missing value (Rubin, 1987). In practice, if missing values have been imputed M times, one ends up with M data sets, which are often called implicates. The non-imputed values are identical, but the imputed values differ based on the choice of imputation model. The data generated in the imputation process is not meant to represent individual observations observations; rather, the technique enables the researcher to perform analyses with less bias and more appropriate standard errors (Rubin, 1996). Figure 1 illustrates an example using five implicates, which is the case with the HFCS data. If one is interested in obtaining a point estimate of a parameter θ, such as the mean age of a population, one can estimate the parameter in each data set separately as if it were a complete data set. Assuming there are five data sets, the process therefore yields five point estimates denoted $\hat{\theta}_m$ for implicate m. These point estimates can then be pooled to a final point estimate $\hat{\theta}$.

As defined by Rubin (1987), the point estimate of multiply imputed data is the arithmetic average of the point estimate of each implicate:

$$\hat{\theta} = \frac{1}{M} \sum_{m=1}^{M} \hat{\theta}_m \tag{6}$$

It should be highlighted that Equation 6 holds for other point estimates than the mean, such as the median.

When it comes to standard errors of multiply imputed estimates, it is not suitable to take the arithmetic average of the standard errors of each implicate as there is additional uncertainty that comes from the fact that the imputed values have been estimated with some degree of uncertainty (Enders, 2022). As further described by Rubin (1987), one needs to separate the variance that occurs *within* each implicate from the variance that occurs *between* implicates. The within-imputation variance is defined as the arithmetic average of the squared standard error of each implicate:

$$\bar{V}_W = \frac{1}{M} \sum_{m=1}^{M} SE_m^2$$

Furthermore, the between-imputation variance is reminiscent of the sample variance formula:

$$V_B = \frac{1}{M-1} \sum_{m=1}^{M} \left(\hat{\theta}_m - \hat{\theta}\right)^2 \tag{7}$$

where $\hat{\theta}$ is the multiply imputed point estimate as defined in Equation 6. The within-imputation variance and between-imputation variance are combined to form the total standard error of an estimate:

$$SE = \sqrt{\bar{V}_W + V_B + \frac{V_B}{M}}. \tag{8}$$

All point estimates reported in this study will be calculated based on (6), and the reported standard errors of these estimates will be calculated based on (8). The HFCS dataset has been multiply imputed five times on a country-level using stochastic imputation. This means that each imputation has been performed using a regression technique that includes a normally distributed random noise term. Imputing missing data five times is generally accepted as being adequate (Rubin, 1996).

The HFCS also incorporates resampling. Resampling refers to the act of simulating many samples using only one observed sample. When performing statistical inference, one could obtain many samples from that population and reach conclusions based on the sampling distribution. However, obtaining many samples is a costly and time-consuming activity. An alternative method is to simulate the process using resampling. The HFCS uses the Rao-Wu bootstrap method to rescale the weights of each observation. When estimating some parameter θ, this process does not affect the estimated value of θ; rather, it mainly adjusts the standard errors.

Resampling starts with viewing the existing sample as a population itself. Then, a so called pseudosample, i.e. a sample of a sample, is drawn from each stratum. In the HFCS, each pseudosample is of size $n_h - 1$, where n_h is the number of observations in the stratum from which the pseudosample is drawn. The pseudosample is drawn with replacement, which means that one observation from the parent sample can be drawn more than once. The sample weight of the observation in the parent sample is then rescaled based on an expression that relies on the number of times that the observation occurs in the pseudosample. The result is that randomness is added to the weight of each sample. This process is then repeated K times, which means that K weights are generated for each observation. This is analogous to multiple imputation, where M values are imputed for each missing value. Statistical analyses such as point estimations are then performed based on each sampling weight. The results are then pooled to a final point estimate along with standard errors that now take the uncertainty of the sampling method into account. With M implicates and K rescaling iterations, one essentially works with $M \times K$ different versions of the data. Although computationally intense, the process generates standard errors that reflect both the fact that the imputed values are uncertain and that the fact a different sample drawn from the population might yield different estimations. Thus, resampling is comparable to simulating a sampling variance with only one sample.

The expression by which the replicate weight of observation i of stratum h is calculated is as follows:

$$w_{hi}^{*} = \left[1 - \sqrt{\frac{m_h \left(1 - \frac{n_h}{N_h}\right)}{n_h - 1}} + \left(\sqrt{\frac{m_h \left(1 - \frac{n_h}{N_h}\right)}{n_h - 1}} \right) \left(\frac{n_h}{m_h} \right) m_{hi}^{*} \right] w_{hi}$$

where m_h is the size of the pseudosample, n_h is the number of sampled units from stratum h, N_h is the number of units in the population stratum, m_{hi}^{*} is the number of times that unit i of sample h occurs in the pseudosample, and w_{hi} is the original, non-rescaled sample weight (Rao et al., 1992; Household Finance and Consumption Network, 2020).

The method for point estimations and calculation of standard errors is reminiscent of the multiple imputation methods described in Section 3.1.2. Each sample weight yields a unique point estimation, and the pooled point estimation is the arithmetic average as in (6):

$$\bar{\hat{\theta}} = \frac{1}{K} \sum_{k=1}^{K} \hat{\theta}_k \tag{9}$$

where K is the number of resampling iterations. The so-called bootstrap variance relies on the squared differences between the point estimate of each resampling iteration and the pooled point estimate as described in (9). It is calculated similarly to (7):

$$V_B = \frac{1}{K-1} \sum_{k=1}^{K} \left(\hat{\theta}_k - \bar{\hat{\theta}} \right)^2$$

Table 2: Investment Assets

Asset	Assumed Risky
Deposits	
Sight accounts	No
Saving accounts	No
Mutual Funds	
Equity funds	Yes
Bond funds	Yes
Money market funds	No
Real estate funds	Yes
Hedge funds	Yes
Other funds	Yes
Bonds	Yes
Stocks	Yes
Non-self-employment private business	Yes
Managed accounts	Yes
Money owed to household	Yes
Voluntary pension and whole life insurance	Yes
Other financial assets	Yes
Commercial real estate	Yes

Lastly, the survey does have a panel component in the sense that a fraction of households have participated in multiple waves of the survey. However, this study will not place emphasis on the panel component as the share of observations with the panel structure is restricted to only a few countries and a fraction of observations within those countries.

3.2 Definition of Variables

It is important to determine which assets to consider when drawing conclusions about risk-taking. As previuosly mentioend, many studies have only looked at direct and indirect ownership of stocks. Although they provide valuable insights regarding investor behavior, the extent to which conclusions can be generalized to other risky assets is limited. More recent studies (e.g. Calvet and Sodini, 2014) only look at liquid financial assets. I, on the other hand, will consider all reported assets in the data that are primarily used as investment vehicles. This definition includes liquid financial assets, such as stocks and mutual funds, as well as less liquid financial assets, such as privately owned businesses in which the sample unit is not employed.[2] Furthermore, real estate used for commercial purposes will

[2]It is assumed that private businesses in which household members are employed provide functions in addition to serving as investment vehicles.

be considered as part of a household's investable portfolio. The sum of financial assets and commercial real estate will be referred to as *investable wealth.* Furthermore, *investment assets* refers to the asset allocation of investable wealth. The *risky share* is defined as the proportion of investable wealth held in assets that carry risk. A list of assets comprising investable wealth is presented in Table 2. In addition, the *conditional risky share* refers to the risky share of the subset of households that own risky assets, i.e. the risky share conditional on participation. Studying heterogeneity in the conditional risky share is more meaningful as there is no variation in the risky share among non-participants by definition. Therefore, the risky share will refer to the conditional risky share unless otherwise stated. Also, money owed to the household will be considered as part of the risky share. Even if lending money privately may not primarily be an investment decision, the fact that money owed carries risk takes precedence.

As previously mentioned, habit and human capital have been shown to have an effect on risk-taking. However, habit is not observable in the HFCS data. Calvet and Sodini (2014) use the household's average income over the past three years as a proxy for internal habit and the average income of the municipality in which the household resides over the past three years as a proxy for external habit. Furthermore, the value of human capital is derived by estimating the present value of future labor income. Making such estimations can be done using methods such as autoregressive modelling; however, this requires panel data. A subsection of surveyed households have participated in more than one wave, meaning that a panel structure exists in some cases. Nevertheless, given that the panel structure is only present for a few countries and a small subset of households within those countries, the data on consumption on income as they are defined in Table A1 will be used in place of habit and human capital.

The survey includes a risk assessment exercise that will be used as a proxy for relative risk aversion. The variable is an ordinal scale whereby the respondent is asked which of the following alternatives comes closest to describing the amount of financial risk that the respondent willing to take when making investments: 1) take substantial financial risks expecting to earn substantial returns; 2) take above average financial risks expecting to earn above average returns; 3) take average financial risks expecting to earn average returns 4) not willing to take any financial risk. The variable is coded such that the value n represents the n-th option, meaning that 1 represents the lowest level of risk aversion and 4 the highest level of risk aversion This elicitation method is not unique to the HFCS and have been shown to be predicative of risk-related behavior (Guiso and Sodini, 2013).

In order to control for heterogeneity in price levels and the distribution of wealth across countries, all control variables that refer to euro amounts are demeaned by country and wave according to the below function:

$$f(\theta_{h,c,t}) = \frac{\theta_{h,c,t}}{\mu_{\theta,c,t}} \tag{10}$$

where $\theta_{h,c,t}$ represents the value of a variable for household h in country c at time t and $\mu_{\theta,c,t}$ is the average value of θ across all households in country c at time t. This transformation will be applied to all independent variables that represent euro amounts, i.e. investable wealth, residential real estate wealth, consumption, and income.

The definition of the investable wealth elasticity to the risky share takes inspiration from Calvet and Sodini (2014) but is modified to incorporate (10):

$$\eta_{h,c,t} = \frac{d \log(\omega_{h,c,t})}{d\xi_{h,c,t}} \tag{11}$$

where $\omega_{h,c,t}$ is the risky share and $\xi_{h,c,t}$ is the log of the function defined in Equation 10 applied to investable wealth, i.e. $\xi_{h,c,t} = \log(f(W_{h,c,t}))$.

3.3 Model Specifications

I make the following specification in order to model variations in the risky share:

$$\log(\omega_{h,c,t}) = \alpha_w + \delta_c + \eta\, \xi_{h,c,t} + \beta' x_{h,c,t} + \varepsilon_{h,c,t} \tag{12}$$

where α_w is a fixed effect specific to the w-th wave, δ_c is a fixed effect specific to country c, $\xi_{h,c,t}$ is the log transformation of investable wealth as seen in Equation 11, and $x_{h,c,t}$ is a vector of additional control variables. Table A1 in the appendix provides definitions of each control variable. Lastly, η and β are coefficients quantifying the sensitivity of $\log(\omega_{h,c,t})$ to changes in $\xi_{h,c,t}$ and $x_{h,c,t}$, respectively.

Note that Equation 11 can simply be rearranged to

$$d \log(\omega_{h,c,t}) = \eta_{h,c,t} d \log(f(W_{h,c,t})). \tag{13}$$

In other words, by using the log of the risky share in the model specification, the coefficient

η in the specification is a direct representation of the investable wealth elasticity to the risky share. Assuming that the average investable wealth $\mu_{W,c.t}$ is held constant, the fact that wealth has been demeaned does not affect the elasticity.

Furthermore, a probit model is specified in order to model the participation choice. The participation of household h in period t is a binary variable denoted $Y_{h,c,t}$ with the following definition:

$$Y_{h,c,t} = \begin{cases} 1, & \text{if } \omega_{h,c,t} > 0 \\ 0, & \text{if } \omega_{h,c,t} = 0 \end{cases} \tag{14}$$

where $\omega_{h,c,t}$ is the risky share. The aim of the model is to describe the probability that a randomly selected household participates in risky financial markets given a set of regressors denoted $X_{h,c,t}$. Thus, if $X_{h,c,t}$ represents the complete set of regressors and β_* represents the coefficients on each regressor, the specification becomes:

$$P\left(Y_{h,c,t} = 1 \mid X_{h,c,t}\right) = \Phi\left(\beta_*' X_{h,c,t}\right) \tag{15}$$

where Φ is the cumulative distribution function of the standard normal distribution.

Some variations of the above models will be tested for comparative purposes. Specifically, the models will be tested without including the control variables $x_{h,c,t}$ and the country fixed effect δ_c. In addition, a model will be tested using $\xi^2_{h,c,t}$ as an additional regressor:

$$\log\left(\omega_{h,c,t}\right) = \alpha_w + \delta_c + \eta_1\, \xi_{h,c,t} + \eta_2\, \xi^2_{h,c,t} + \beta' x_{h,c,t} + \varepsilon_{h,c,t} \tag{16}$$

A positive coefficient on a regressor in combination with a positive (negative) coefficient on its square implies that there is a positive relationship between the regressor and the independent variable and that the marginal effect of a unit increase in the regressor increases (decreases) as the regressor increases. Since both $\omega_{h,c,t}$ and $W_{h,c,t}$ have been transformed to logs in the regression, this technique allows me to investigate whether elasticity is constant across different levels of investable wealth. As previously mentioned, Calvet and Sodini (2014) finds that the elasticity of financial wealth of the risky share is positive but decreases with wealth. If this holds, η_1 will be positive and η_2 will be negative.

The above models will also be tested using a different definition of the risky share. In this alternative definition, commercial real estate is excluded from the risky share as well as from investable wealth, meaning that investable wealth will only consist of nontangible assets similarly to the setup in Calvet and Sodini (2014). Instead, commercial real estate wealth will be used as a control variable similarly to residential real estate. This alternative specification is used for comparative purposes in order to investigate what impact the definition of the risky share has on conclusions regarding elasticity.

Lastly, given that the HFCS covers 22 countries, investigating the drivers behind unobserved cross-country differences is an extensive task as it requires an analysis of the investment climate of each country. Conducting such an analysis with an adequate level of detail is therefore beyond the scope of this project.

4 Results

4.1 Cross-Country Characteristics

There is substantial heterogeneity in participation and the risky share across European countries. The difference in participation can be read from Panel A in Figure 2. The panel also shows the proportion of households that own any investment assets, regardless of risk. In several countries, almost all households have investment assets. In some countries, such as Greece, Croatia, and Hungary, a substantial proportion (although still a minority) does not have any investment assets, meaning that these households do not even have a non-empty bank account. It should be noted that the HFCS does not cover physical currency, so these households may still hold cash and real assets that are not primarily used as investment vehicles. Nonetheless, there are no clear indications that households in countries where more (fewer) households own investment assets have a higher (lower) proportion of participating households. If one were to view the proportion of households owning investment assets as an indication of the level of financial market development in a country, the cross-country differences in participation are puzzling. For example, the difference in participating households between Austria and Belgium are very significant.

Panel B shows the conditional risky share. Interestingly, there seems to be less variation in the conditional risky share than participation. As noted by Guiso and Sodini (2013), cross-country differences in participation costs are unlikely to explain the extent to which participation varies across countries. Nonetheless, cross-country differences in fixed participation costs would predict variation in participation but not variation in the conditional risky share. A supplementary illustration of the country-level relationship between participation and the risky share is provided in Appendix Figure A1.

The risky part of investment assets is decomposed in Figure 3 for each country and wave. First, there is substantial heterogeneity between countries in terms of the composition of risky portfolios, which highlights the importance of using data from multiple countries when researching risk-taking among households. As previously mentioned, studies have often focused on stock ownership, even though it is clear that stocks do not necessarily constitute a significant portion of risky portfolios for the average household in several countries. Secondly, the figure highlights the need of considering commercial real estate when researching risky portfolios, as commercial real estate is an important investment vehicle for the average household. In Greece, the average weight of commercial

Figure 2: Participation and the Risky Share

The first panel shows the proportion of households that 1) own financial assets of any kind and 2) own risky financial assets. The second panel shows the average share of financial portfolios invested in risky financial assets conditional on participation in risky financial markets. The error bars represent 95% confidence intervals.

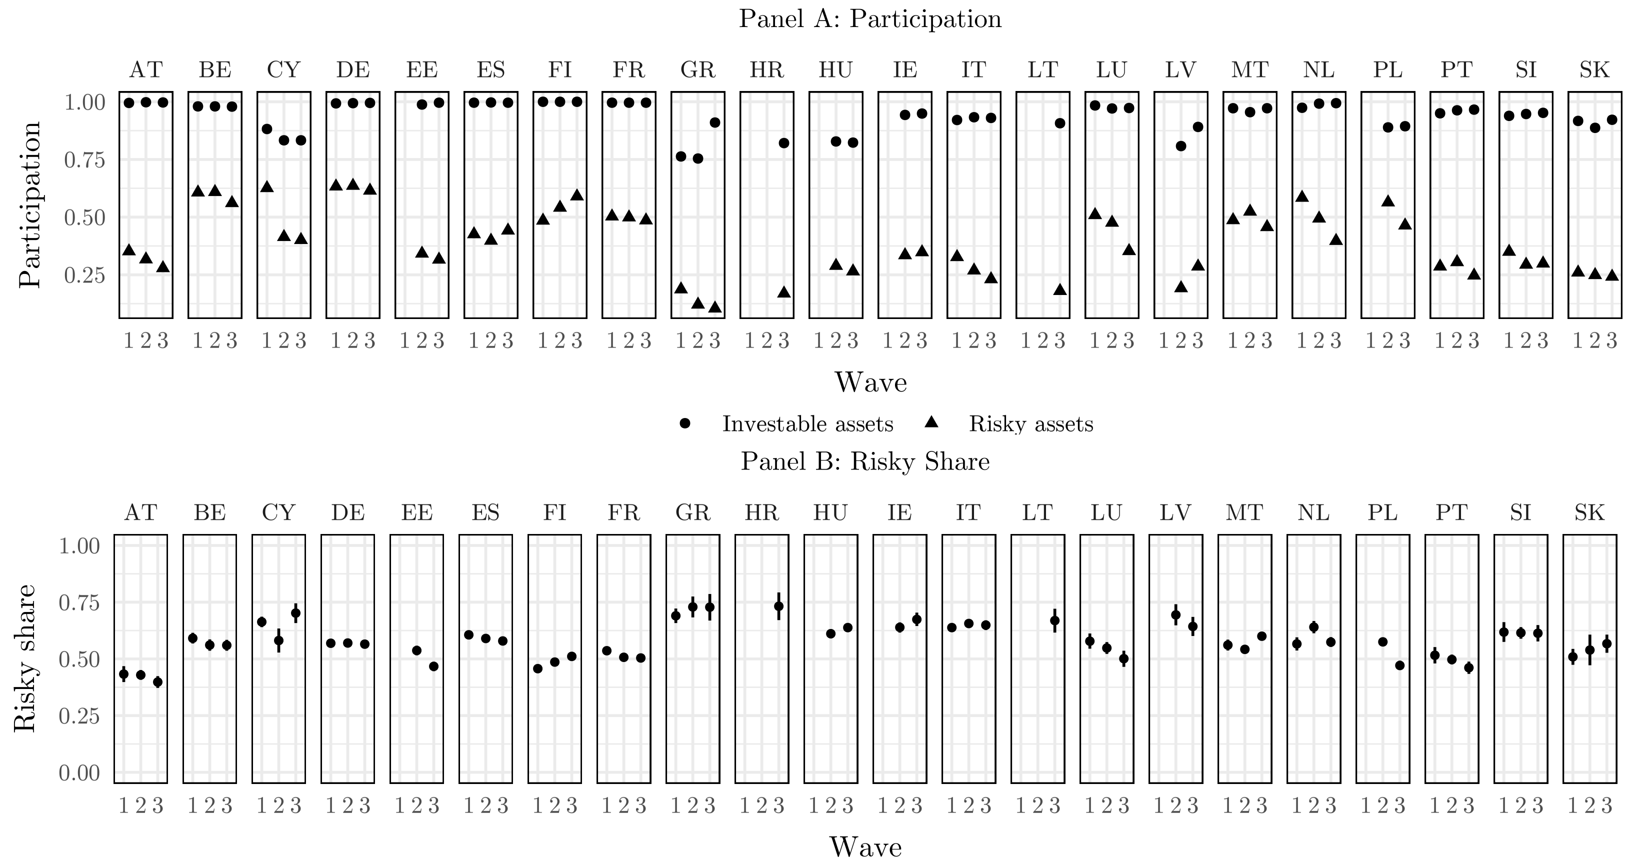

Figure 3: Decomposition of Risky Portfolios

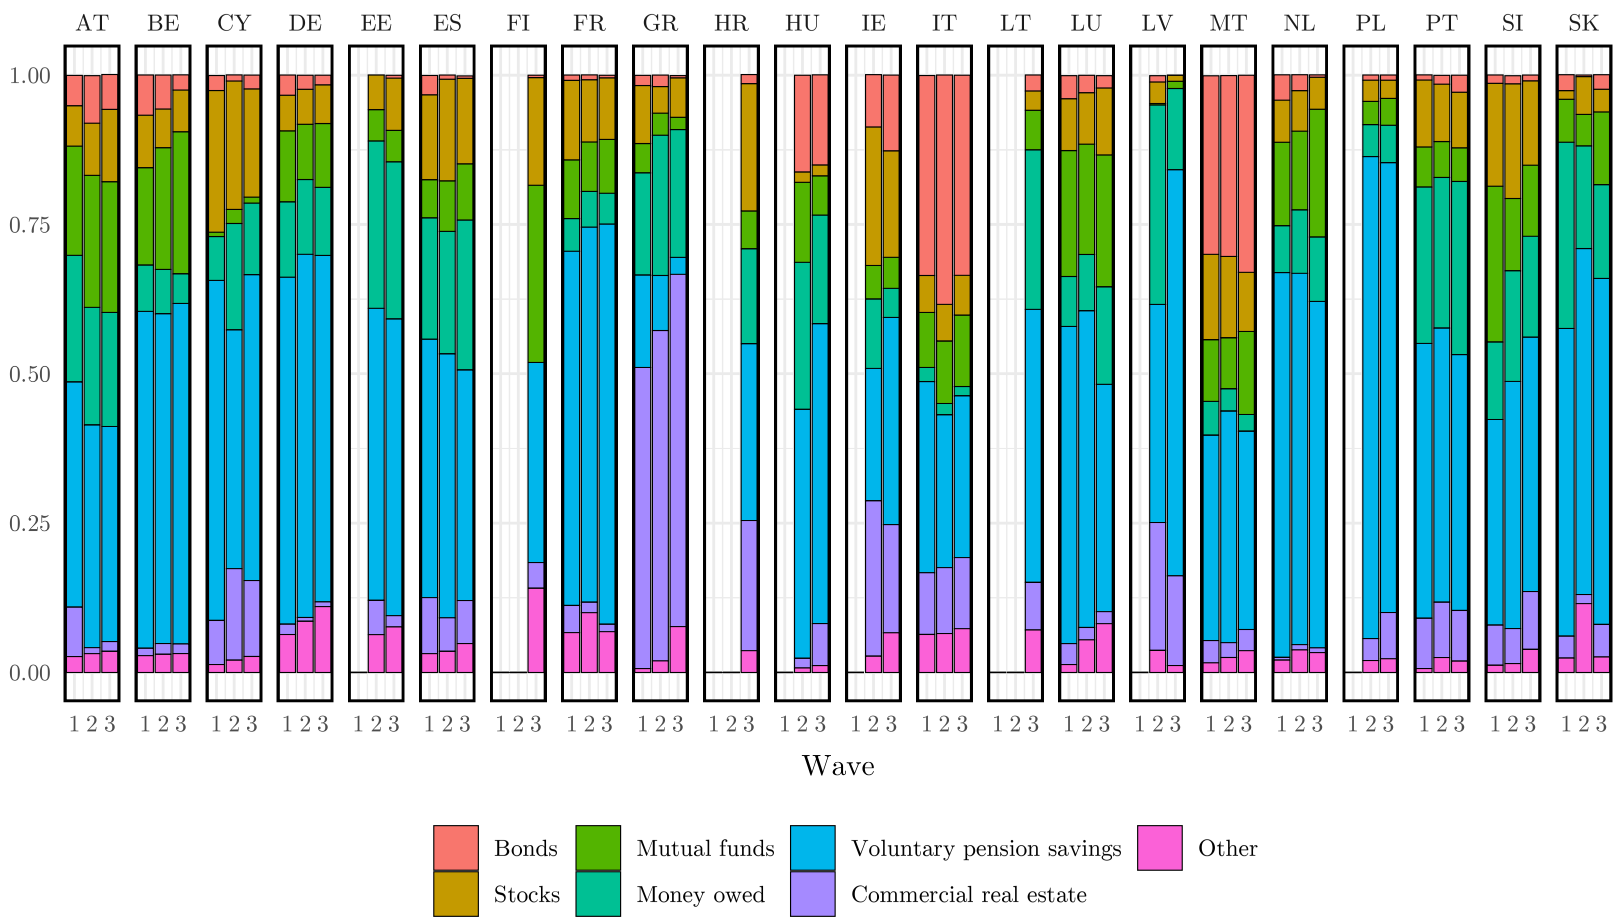

real estate within risky portfolios is substantially higher than any other asset, which is yet another clear indication of the existence of cross-country heterogeneity. Greek households that decide to invest in risky assets are evidently motivated by the possibility of investing in real estate, which helps explain both the low participation and the high risky share in Greece. Buying real estate for commercial purposes entails a large investment, which discourages households from participating while significantly increasing the risky share for those who do participate. The underlying reasons for this phenomenon is not observed in the household-level data; however, they will be absorbed by the country fixed effect in the model specification detailed in Section 3.3. Furthermore, these unobserved variables are likely to be numerous, and identifying them does not lie within the scope of this thesis but could be areas of interest for future research.

Secondly, Figure 3 highlights a drawback of the HFCS data, which is that voluntary pension savings are not decomposed into different asset classes. In most countries, voluntary pension savings is the most important risky asset class for the average household. It should be highlighted that this asset class, by the definition of the HFCS, does not include savings managed by the household itself. Rather, this asset class refers to pension investments managed by pension funds and other financial institutions. It could be that households have limited awareness regarding the underlying assets of these funds, which would impede their ability to give reliable information regarding their decomposition in a survey setting. On the other hand, it could be possible to impute this data using auxiliary information.

4.2 Investable Wealth and Risk-Taking

In Table 3, I estimate the probit model (15) of participation in columns 1-4 and the linear model (12) of the investment wealth elasticity of the risky share in columns 5-8. The coefficient η is significant in all specifications. As previously mentioned, the coefficient can be directly interpreted as the elasticity in the risky share regressions in columns 5-7. Thus, the elasticities are estimated to be 0.031, 0.018, and 0.083, respectively. The positive relationship holds even when controlling for demographic and economic characteristics such as age, income, family size, and leverage. The elasticity becomes larger when adding the control variables. This implies that there is some correlation between the regressor and the control variables whose relationship with the risky share is opposite of the relationship between investment wealth and the risky share. Furthermore, the coefficient on the squared regressor, η_2, is significantly positive, meaning that the investment wealth elasticity of the risky share increases with wealth. Thus, a change in investment wealth has a larger impact on the risky share for wealthy households. The results suggest an increasing and slightly convex relationship between investment wealth and the risky share, which is contrary to the concave relationship proposed by Calvet and Sodini (2014). It should be noted that the positive coefficient on the squared regressor is very close to zero, and that it is only significant on the 0.05 level. Furthermore, the significance of the coefficients are robust to the exclusion of money owed in the risky share and investable wealth definitions.

Since the survey includes a risk assessment exercise, I can control for risk preference and conclude that investment wealth increases the risky share, even when risk preference is held fixed. This is highly indicative of DRRA utility. The conclusion does, however, rely on the validity of the risk assessment exercise. It could be that the exercise is a proxy for some other variable, or that one's self-perceptions and beliefs regarding risk are not reflected in real-world scenarios. Using twin-pair fixed effects as dony by Calvet and Sodini (2014) is another way to control for risk tolerance. In any case, future research should continue to apply different risk elicitation methods.

Next, I document the results using an alternative set of regressors. Instead of investable wealth, i.e. financial assets and commercial real estate wealth, the main regressor will only include financial wealth. Thus, the risky share will only refer to the risky share of financial assets. Instead, commercial real estate wealth is used as a control variable. This specification is similar to specification used by Calvet and Sodini (2014). The results are reported in Table 4. The effect of financial wealth on participation is positive and highly significant, as seen in columns 1-4. The result is robust to the inclusion of control variables, including commercial real estate wealth. However, the coefficients are insignificant in the risky share regressions, as seen in columns 5-8. Thus, the results fail to replicate the significant financial wealth elasticity of the risky share shown in the twin-pair fixed regressions by Calvet and Sodini (2014). There are a few reasons that could explain this discrepancy. First, the authors only look at Swedish data. Sweden is not a part of the HFCN, so there is no data on Swedish households in the HFCS. The coefficients on the dummy variables from the country fixed effects are in general highly significant. Therefore, much of the variation in the risky share is attributed to unobserved cross-country differences. This study does not capture such differences that are applicable to Swedish households. Second, the significant financial wealth elasticity of the risky share found by

Figure 4: Increase in Elasticity

The figure compares the investment wealth elasticity of the risky share between two model specifications (17) and (18). Each dot represents a country. The commercial real estate wealth ratio on horizontal axis represents the aggregate share of commercial real estate wealth relative to investment wealth within that country. The vertical axis shows the increase in investment wealth elasticity of the risky share that occurs when including commercial real estate wealth in the risky share and investment wealth.

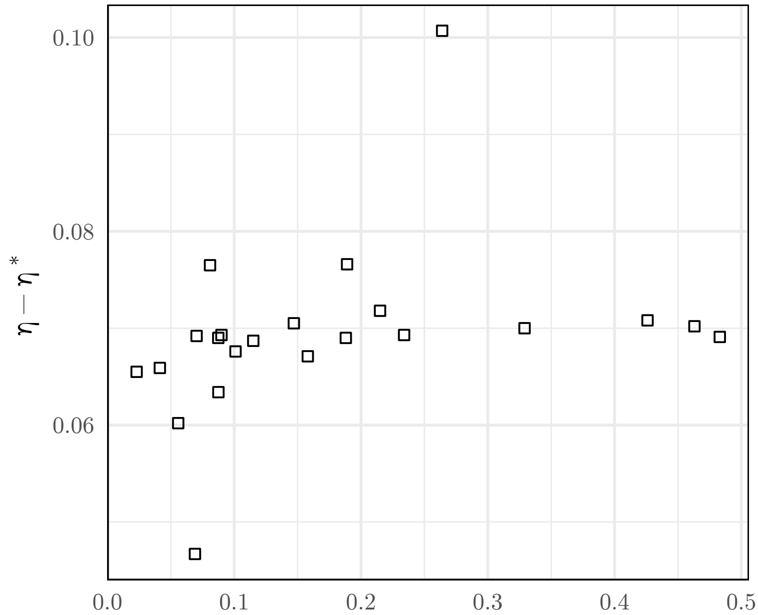

Calvet and Sodini (2014) is more similar to the results shown in Table 3. This indicates that the exclusion of commercial real estate wealth in the risky share definition has more significant consequences in the countries surveyed in the HFCS than it does in Sweden, perhaps because commercial real estate is a less common investment vehicle in Sweden. In order to further investigate the impact of including commercial real estate wealth in the risky share, I test the below two specifications for each country:

$$\log\left(\omega_{h,t}\right) = \alpha_t + \eta\ \xi_{h,t} + \beta' x_{h,t} + \varepsilon_{h,c,t} \tag{17}$$

$$\log\left(\omega^*_{h,t}\right) = \alpha_t + \eta^*\ \xi^*_{h,t} + \beta' x^*_{h,t} + \varepsilon_{h,t} \tag{18}$$

where the asterisks indicate that commercial real estate wealth is a control variable and not a part of investable wealth or the risky share. I then compare the estimated elasticities for each country by subtracting η^* from η. Given the difference in elasticities shown in Table 3 and 4, it is expected that the elasticity increases more in countries where households in

Table 3: Investable Wealth Regressions

The table shows the probit participation regressions (columns 1-4) and the log risky share regressions conditional on participation (columns 5-8). All standard errors account for resampling and between-imputation variance. The R^2 values reported for the OLS regressions are obtained by running the regression on each implicate separately and calculating the average R^2 from the five implicates.

	Participation				Risky Share			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log investable wealth	0.297***	0.427***	0.488***	0.484***	0.031***	0.018*	0.083***	0.095***
	(0.006)	(0.010)	(0.014)	(0.035)	(0.006)	(0.007)	(0.010)	(0.012)
Squared log investable wealth				-0.001				0.006*
				(0.006)				(0.003)
Number of observations	222,443	222,443	222,443	222,443	169,503	169,503	169,503	169,503
R^2					0.0398	0.057	0.0610	0.0640
Wave fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country fixed effects	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Controls:								
Log residential real estate wealth	No	No	Yes	Yes	No	No	Yes	Yes
Log consumption	No	No	Yes	Yes	No	No	Yes	Yes
Log disposable income	No	No	Yes	Yes	No	No	Yes	Yes
Log leverage	No	No	Yes	Yes	No	No	Yes	Yes
Number of adults	No	No	Yes	Yes	No	No	Yes	Yes
Number of children	No	No	Yes	Yes	No	No	Yes	Yes
Age	No	No	Yes	Yes	No	No	Yes	Yes
Risk aversion	No	No	Yes	Yes	No	No	Yes	Yes

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4: Financial Wealth Regressions

The table shows the probit participation regressions (columns 1-4) and the risky share OLS regressions conditional on participation (columns 5-8). All regressions uses wave fixed effects. Columns 2-4 and 6-8 also use country fixed effects. All standard errors account for resampling and between-imputation variance. The R^2 values reported for the OLS regressions are obtained by running the regression on each implicate separately and calculating the average R^2 from the five implicates.

	Participation				Risky Share			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log financial wealth	0.276***	0.416***	0.471***	0.498***	0.011	-0.002	0.008	0.025
	(0.006)	(0.010)	(0.037)	(0.049)	(0.006)	(0.007)	(0.034)	(0.035)
Squared log financial wealth				0.008				0.009*
				(0.009)				(0.004)
Number of observations	222,133	222,133	222,133	222,133	135,008	135,008	135,008	135,008
R^2					0.0432	0.0612	0.0745	0.0746
Wave fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country fixed effects	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Controls:								
Log commercial real estate wealth	No	No	Yes	Yes	No	No	Yes	Yes
Log residential real estate wealth	No	No	Yes	Yes	No	No	Yes	Yes
Log consumption	No	No	Yes	Yes	No	No	Yes	Yes
Log disposable income	No	No	Yes	Yes	No	No	Yes	Yes
Log leverage	No	No	Yes	Yes	No	No	Yes	Yes
Number of adults	No	No	Yes	Yes	No	No	Yes	Yes
Number of children	No	No	Yes	Yes	No	No	Yes	Yes
Age	No	No	Yes	Yes	No	No	Yes	Yes
Risk aversion	No	No	Yes	Yes	No	No	Yes	Yes

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5: Risk Aversion

The table shows the coefficient on the relative risk aversion proxy used in the HFCS and the log risky share. The column Columns 1-4 are probit regressions on participation as specified in Equation 15. Columns 5-8 are OLS regressions on the log risky share conditional on participation as specified in Equation 12. The regressions are equivalent to those presented in Table 3, but the relative risk aversion parameter is used as the main regressor and log investable wealth is used as a control variable. All standard errors account for resampling and between-imputation variance. The R^2 values reported for the OLS regressions are obtained by running the regression on each implicate separately and calculating the average R^2 from the five implicates.

	Participation				Risky Share			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Risk aversion	-0.333***	-0.435***	-0.238***	-0.237***	-0.144***	-0.118***	-0.114***	-0.112***
	(0.012)	(0.016)	(0.023)	(0.022)	(0.013)	(0.013)	(0.017)	(0.017)
Number of observations	222,443	222,443	222,443	222,443	169,503	169,503	169,503	169,503
R^2					0.0089	0.0239	0.0530	0.0536
Wave fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country fixed effects	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Controls:								
Log investable wealth	No	No	Yes	Yes	No	No	Yes	Yes
Squared log investable wealth	No	No	No	Yes	No	No	No	Yes
Log residential real estate wealth	No	No	Yes	Yes	No	No	Yes	Yes
Log consumption	No	No	Yes	Yes	No	No	Yes	Yes
Log disposable income	No	No	Yes	Yes	No	No	Yes	Yes
Log leverage	No	No	Yes	Yes	No	No	Yes	Yes
Number of adults	No	No	Yes	Yes	No	No	Yes	Yes
Number of children	No	No	Yes	Yes	No	No	Yes	Yes
Age	No	No	Yes	Yes	No	No	Yes	Yes

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

aggregate own more commercial real estate relative to other investment assets. The results are visualized in Table 4. First of all, $\eta > \eta^*$ in every country. Secondly, $\eta - \eta^*$ generally increases as the commercial real estate wealth ratio increases. The result supports the idea that including commercial real estate wealth in the definitions of the risky share and investable wealth has a larger impact on the estimated elasticity in countries where households own more commercial real estate. Another possible explanation of the non-significance of the financial wealth elasticity to the risky share in Table 4 is that the standard errors around the coefficients take both resampling and between-imputation variance into account. Since both resampling and multiple imputation considers sources of uncertainty that is ignored in single-imputation methods, it is possible that the standard errors reported in this study are larger than they would have been if missing data had been imputed only once.

In order to establish the validity of the relative risk aversion proxy used in the HFCS, regressions equivalent to those presented in Table 3 have been performed using the risk aversion parameter as the main regressor instead of log investable wealth. The results are presented in Table 5 and suggest that the relative risk aversion elicited in the HFCS has validity. If it is true that the variable is a proxy for relative risk aversion, it must follow that the variable has a significant negative impact on risk-taking, both in terms of participation and in terms of the conditional risky share. Regardless of whether an investor displays CRRA or DRRA utility, a higher RRA means that the investor is more sensitive to a proportional change in wealth and will therefore choose to invest less in risky assets. The results show a highly significant relationship between the risk aversion parameter and participation, as well as a highly significant relationship between the risk aversion parameter and the conditional risky share. Significance holds when adding control variables. Combining these results with those presented in Table 3 provides a strong indication that the average household displays decreasing relative risk aversion. These findings are in line with the empirical findings by Calvet and Sodini (2014).

The last section of the result will briefly discuss the country fixed effects. The coefficients of each country dummy variable from the regressions in Table 3 are presented in Table 6. The coefficients are highly significant in most cases, which implies that there is much remaining variation in both participation and the conditional risky share that is explained by unobserved cross-country differences. In other words, the financial and demographic characteristics of households as well as their level of risk tolerance are not

Table 6: Country Fixed Effects

The table shows the coefficients on each dummy variable resulting from the use of country fixed effects. Austra is the base-level case and is therefore omitted. The first column refers to the participation regression in column 4 in Table 3, and the second column refers to the conditional risky share regression in column 8 in Table 3.

			Continued		
Country	(1)	(2)	Country	(1)	(2)
BE	1.158***	0.462***	IT	-0.032	0.550***
	(0.070)	(0.060)		(0.042)	(0.059)
CY	0.902***	0.314***	LT	-0.222	0.680***
	(0.115)	(0.077)		(0.191)	(0.076)
DE	1.166***	0.422***	LU	0.537***	0.238***
	(0.039)	(0.050)		(0.044)	(0.063)
EE	0.184*	0.056	LV	0.414***	0.676***
	(0.084)	(0.103)		(0.095)	(0.066)
ES	1.556***	0.534***	MT	0.637***	0.366***
	(0.055)	(0.066)		(0.086)	(0.095)
FI	1.448***	0.342***	NL	0.686***	0.380***
	(0.045)	(0.063)		(0.142)	(0.072)
FR	1.584***	0.383***	PL	0.394***	0.231***
	(0.058)	(0.055)		(0.061)	(0.060)
GR	-0.586***	0.594***	PT	0.008	0.097
	(0.045)	(0.057)		(0.053)	(0.073)
HR	-0.142	0.570***	SI	0.278**	0.296**
	(0.167)	(0.148)		(0.086)	(0.085)
HU	0.118*	0.512***	SK	-0.449***	0.123*
	(0.050)	(0.071)		(0.093)	(0.051)

enough to explain cross-country differences with respect to risk-taking. These findings are in line with Guiso et al. (2003a), who also employed a country fixed effect regression on a smaller set of countries, albeit with fewer control variables. Given the significance of the country dummy variables, future studies should investigate the impact of different institutional environments across countries in order to provide further insights regarding the determinants of household risk-taking.

5 Conclusions

Using data from 22 European countries, I show that households with more investable wealth are more likely to hold risky assets. In addition, participating households with more investable wealth have a higher conditional risky share. The positive relationship holds when controlling for risk aversion, which is highly indicative of decreasing relative risk aversion. In addition, the results are robust to the inclusion of country-specific fixed effects. Furthermore, I find that the inclusion of commercial real estate in the risky share is a prerequisite for observing the positive investable wealth elasticity of the risky share. Viewing commercial real estate as an investment vehicle similarly to financial assets is motivated by the observation that the average household allocates a significant portion of their investable wealth to commercial real estate.

I further observe significant cross-country differences in both participation and the risky share, in spite of controlling for financial and demographic characteristics. Thus, many underlying drivers of differences across countries with respect to participation and the conditional risky share are still left unexplored. A suggestion for future research investigating cross-country differences is therefore to complement the micro-level data observed in surveys such as the HFCS with country-level data in order to disentangle the unobserved heterogeneity. Lastly, the HFCN is expected to publish a fourth wave of the HFCS in 2023. As more waves of the HFCS are being published, it is likely that the panel component of the data will be strengthened in the future. Therefore, it would be interesting to employ fixed effects specific to individual households in order to investigate how investment decisions within households evolve over time.

References

Brick, J. and Kalton, G. (1996). Handling Missing Data in Survey Research. *Statistical Methods in Medical Research*, 5(3):215–238.

Calvet, L., Campbell, J., and Sodini, P. (2007). Down or Out: Assessing the Welfare Costs of Household Investment Mistakes. *Journal of Political Economy*, 115(5):707–747.

Calvet, L. E. and Sodini, P. (2014). Twin Picks: Disentangling the Determinants of Risk-Taking in Household Portfolios. *The Journal of Finance*, 69(2):867–906.

Campbell, J. and Cochrane, J. (1999). By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy*, 107(2):205–251.

Campbell, J. Y., Jackson, H. E., Madrian, B. C., and Tufano, P. (2011). Consumer Financial Protection. *Journal of Economic Perspectives*, 25(1):91–114.

Campbell, J. Y. and Viceira, L. M. (2002). *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press, New York.

Canner, N., Mankiw, N. G., and Weil, D. N. (1997). An Asset Allocation Puzzle. *The American Economic Review*, 87(1):181–191.

Christelis, D., Georgarakos, D., and Haliassos, M. (2013). Differences in Portfolios across Countries: Economic Environment versus Household Characteristics. *Review of Economics and Statistics*, 95(1):220–236.

Cocco, J. F. (2005). Portfolio Choice in the Presence of Housing. *Review of Financial Studies*, 18(2):535–567.

Cohn, A., Engelmann, J., Fehr, E., and Maréchal, M. A. (2015). Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals. *American Economic Review*, 105(2):860–885.

Dimmock, S. G., Kouwenberg, R., Mitchell, O. S., and Peijnenburg, K. (2016). Ambiguity Aversion and Household Portfolio Choice Puzzles: Empirical Evidence. *Journal of Financial Economics*, 119(3):559–577.

Enders, C. K. (2022). *Applied Missing Data Analysis*. Methodology in the social sciences. The Guilford Press, New York, second edition.

Fuller, W. A. (2009). *Sampling Statistics.* Wiley series in survey methodology. Wiley, Hoboken, N.J. OCLC: ocn299695775.

Grinblatt, M. and Keloharju, M. (2009). Sensation Seeking, Overconfidence, and Trading Activity. *The Journal of Finance*, 64(2):549–578.

Grossman, S. J. and Vila, J.-L. (1992). Optimal Dynamic Trading with Leverage Constraints. *The Journal of Financial and Quantitative Analysis*, 27(2):151.

Guiso, L., Haliassos, M., and Jappelli, T., editors (2002). *Household Portfolios.* MIT Press, Cambridge, Mass. Meeting Name: Conference on Household Portfolios.

Guiso, L., Haliassos, M., and Jappelli, T. (2003a). Household Stockholding in Europe: Where Do We Stand and Where Do We Go? *Economic Policy*, 18(36):125–170.

Guiso, L., Haliassos, M., and Jappelli, T. (2003b). Stockholding: A European Comparison. In Guiso, L., Haliassos, M., and Jappelli, T., editors, *Stockholding in Europe*, pages 3–29. Palgrave Macmillan UK, London.

Guiso, L. and Jappelli, T. (2002). Household Portfolios in Italy. In Guiso, L., Haliassos, M., and Jappelli, T., editors, *Household portfolios.* MIT Press, Cambridge, Mass. Meeting Name: Conference on Household Portfolios.

Guiso, L. and Sodini, P. (2013). Household Finance: An Emerging Field. In *Handbook of the Economics of Finance*, volume 2, pages 1397–1532. Elsevier.

Heaton, J. and Lucas, D. (2000). Portfolio Choice in the Presence of Background Risk. *The Economic Journal*, 110(460):1–26.

Household Finance and Consumption Network (2020). ECB Statistics Paper Series. Technical Report 35, Household Finance and Consumption Network.

Merton, R. C. (1969). Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case. *The Review of Economics and Statistics*, 51(3):247.

Rao, J. N. K., Wu, C. F. J., and Yue, K. (1992). Some Recent Work on Resampling Methods for Complex Surveys. *Survey Methodology*, 18(2):209–217.

Rubin, D. B. (1987). *Multiple Imputation for Nonresponse in Surveys.* Wiley series in probability and mathematical statistics. Wiley, New York.

Rubin, D. B. (1996). Multiple Imputation after 18+ Years. *Journal of the American Statistical Association*, 91(434):473–489.

Samuelson, P. A. (1969). Lifetime Portfolio Selection By Dynamic Stochastic Programming. *The Review of Economics and Statistics*, 51(3):239.

Sinai, T. and Souleles, N. S. (2005). Owner-Occupied Housing as a Hedge Against Rent Risk. *The Quarterly Journal of Economics*, 120(2):763–789.

Särndal, C.-E. (2007). The Calibration Approach in Survey Theory and Practice. *Survey methodology*, 33(2):99–119.

Tufano, P. (2009). Consumer Finance. *Annual Review of Financial Economics*, 1(1):227–247.

Vissing-Jorgensen, A. (2003). Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions. *NBER Macroeconomics Annual*, 18:139–194.

Yao, R. and Zhang, H. H. (2005). Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints. *Review of Financial Studies*, 18(1):197–239.

Figure A1: Participation and the Risky Share

Each panel represents one wave of data. The horizontal axis represents the proportion of participating households. The vertical axis shows the conditional risky share.

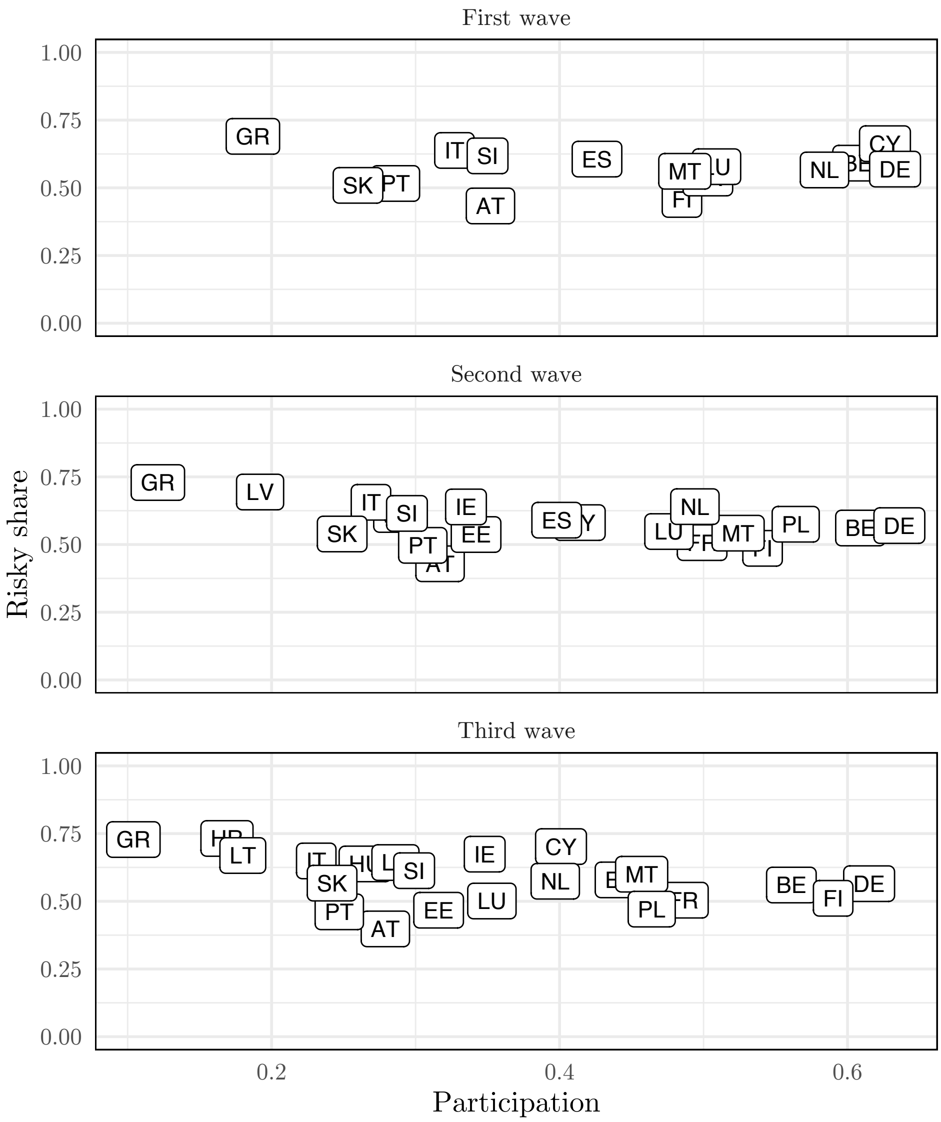

Table A1: Definition of Variables

The table shows a definition of control variables used in regressions. Note that not all control variables are used in all regressions. The selection of control variables used in each regression is specified in each regression table. The table also indicates whether the variable has been transformed. *Demeaned* refers to the transformation defined in Equation 10. This transformation is performed before the logarithmic transformation.

Asset	Description	Demeaned	Log
Investable assets	Sum of financial assets and commercial real estate	Yes	Yes
Residential real estate	Value of household's main residence and other real estate not used for business activities	Yes	Yes
Commercial real estate*	Value of real estate used for business activities	Yes	Yes
Consumption	Amount spent on food, utilities, trips, holidays, goods, services, charities, and alimonies	Yes	Yes
Income	Total household gross income	Yes	Yes
Leverage	Ratio between total liabilities and total gross assets	No	Yes
Number of adults	Number of household members 16+	No	No
Number of children	Number of children in household (0-13)	No	No
Age	Age of the household head	No	No
Risk aversion	See section 3.2	No	No
Gender	Dummy variable equal to 1 if the household head is female	No	No
Entrepreneur dummy	Dummy variable equal to 1 if the household head is self-employed	No	No
Unemployed dummy	Dummy variable equal to 1 if the household head is unemployed	No	No
College dummy	Dummy variable equal to 1 if the household has completed post-secundary education	No	No