

Dutt, Satyajit; Radermacher, Jan W.

Working Paper

Age, wealth, and the MPC in Europe: A supervised machine learning approach

SAFE Working Paper, No. 383

Provided in Cooperation with:

Leibniz Institute for Financial Research SAFE

Suggested Citation: Dutt, Satyajit; Radermacher, Jan W. (2023) : Age, wealth, and the MPC in Europe: A supervised machine learning approach, SAFE Working Paper, No. 383, Leibniz Institute for Financial Research SAFE, Frankfurt a. M.

This Version is available at:

<https://hdl.handle.net/10419/268744>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Satyajit Dutt | Jan W. Radermacher

Age, Wealth, and the MPC in Europe

A Supervised Machine Learning Approach

SAFE Working Paper No. 383 | February 2023

Leibniz Institute for Financial Research SAFE
Sustainable Architecture for Finance in Europe

Age, Wealth, and the MPC in Europe

A Supervised Machine Learning Approach^{*}

Satyajit Dutt[†] Jan W. Radermacher[‡]

February 13, 2023

Abstract

We investigate consumption patterns in Europe with supervised machine learning methods and reveal differences in age and wealth impact across countries. Using data from the third wave (2017) of the Eurosystem’s Household Finance and Consumption Survey (HFCS), we assess how age and (liquid) wealth affect the marginal propensity to consume (MPC) in the Netherlands, Germany, France, and Italy. Our regression analysis takes the specification by Christelis et al. (2019) as a starting point. Decision trees are used to suggest alternative variable splits to create categorical variables for customized regression specifications. The results suggest an impact of differing wealth distributions and retirement systems across the studied Eurozone members and are relevant to European policy makers due to joint Eurozone monetary policy and increasing supranational fiscal authority of the EU. The analysis is further substantiated by a supervised machine learning analysis using a random forest and XGBoost algorithm.

^{*}We thank Andreas Hackethal, Nathanael Vellekoop, Gyozo Gyöngyösi, Fabian Nemecek, and the participants of the EFL Jour Fixe at Goethe University Frankfurt on 18 January, 2021 for helpful comments and suggestions. Financial support from the Leibniz Institute for Financial Research SAFE is gratefully acknowledged. This paper uses data from the Eurosystem’s Household Finance and Consumption Survey (HFCS).

[†]sunnydutt@gmail.com

[‡]SAFE, GSEFM, and Goethe University, radermacher@safe-frankfurt.de

1 Introduction

In the EU, pension systems differ between the member states but monetary policy is conducted jointly. It is therefore important to understand differences in consumption patterns - particularly, as the group of retirees is relatively growing in most societies.

Special emphasis lies on the marginal propensity to consume (MPC), i.e. the consumption rate/fraction out of a change in income. Because pension amounts are often at the discretion of governments, policy makers are particularly concerned about their ramifications on macroeconomic variables such as aggregate demand. But also additional/removed taxes, introduced/suspended transfers, rising/falling inflation, or appreciating/depreciating foreign exchange rates can pose income shocks that affect the purchasing power of individuals and are therefore at the concern of policy makers.

This study analyzes consumption patterns in four of the five largest Eurozone economies: Germany, France, Italy, and the Netherlands.¹ Because traditional consumption theory (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978; Deaton, 1991; Carroll, 1997) does not account for country-specific differences, we ask: how does the MPC differ per Eurozone country? What are the impacts of important demographic variables such as age and wealth? Do country differences in pension schemes, life expectancy, age or wealth distributions matter for consumption behavior?

We follow Christelis et al. (2019) in their basic econometric set-up and use data from the third wave (2017) of the Eurosystem's Household Finance and Consumption Survey (HFCS) (HFCN, 2020b). Utilizing decision trees, a corner stone of supervised machine learning, we derive age and wealth dummies for regression specifications that outperform the specifications suggested by Christelis et al. (2019). Decision trees are discussed as simple, effective, and interpretable tools to deal with non-linear data and we suggest their use in combination with the traditional and more rigid method of linear regression. Because decision trees are also the building blocks of many advanced supervised machine learning models like random forests or XGBoost, this paper also has an educational purpose: to show how these simple data-driven methods can be used in combination with traditional econometrics and to build the intuition for the more advanced machine learning methods. Both random forest and XGBoost are also introduced and applied to complement the analysis.

In our initial regression analysis, we could not reproduce the results

¹Spain had to be excluded due to missing data.

found by Christelis et al. (2019) fully: age-dummies, as suggested by the authors, were not significant in our data for the Netherlands. Using decision trees to derive data-driven dummies at alternative points in the distribution rendered significant coefficients and higher or equal model fit for all countries. Suggested upper age splits were in the Netherlands and Germany close to the expected future retirement age at 73 and 63, respectively, and in France and Italy below normal pension age at 58 and 56. In France and Italy, aggregate replacement ratios upon retirement are much higher than in the Netherlands or Germany, suggesting that retirement should have a weaker impact on income and consumption. Similarly, for (liquid) wealth we derived dummies based on percentiles different from those used by Christelis et al. (2019). We interpret our results as a data-driven way of showing how differences between countries in age and wealth impact the MPC and how to generate customized models. Furthermore, the results from machine learning provide additional support for country-specific differences and drill down on individual model choices by means of a SHAP (SHapley Additive exPlanations) analysis.

Since we are working with survey data, we want to point out that there are several flaws coming with this data type and that interpretation should be conducted with the usual care. Regarding parametrization of machine learning models, we use cross validation and a parameter grid to make optimal use of the available data and to try several parameter combinations.

The remainder of this paper is structured as follows. Chapter 2 introduces the literature and sets the theoretical background. Chapter 3 presents the data and methodology. Chapter 4 features the analysis, which contains (i) a simple linear model, (ii) a decision tree analysis, (iii) resulting customized regression set-ups, (iv) an analysis using random forests, and (v) an analysis using XGBoost. Chapter 5 discusses limitations and suggestions for further research. Chapter 6 concludes.

2 Literature Review and Theoretical Background

2.1 Consumption Theory

Consumption theory for the last five decades has been centered around the concept of the Life Cycle/Permanent Income Hypothesis (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978). In (the imaginary world assumed by) these models, rational agents are equipped with a utility func-

tion² which they use for inter-temporal consumption optimization to derive patterns that yield highest utility for them. If one accepts all of these assumptions (which implicitly many economists do as the concept is, in some form, included in most contemporary macroeconomic models (Smets and Wouters, 2003; Kaplan et al., 2018)), one can solve the maximization problem. The implied optimal solution is that the highest possible utility level is achieved when consumption is the same in every period, which yields perfect consumption smoothing.

Common sense suggests (to most people) that this is not exactly how people behave in reality and that empirical evidence in its support will be hard to find. To economics, however, the lack of empirical support was a problem, so the theory has been tinkered with frequently in the past decades to accommodate the mismatch when being confronted with real data. Most notable augmentations to the theory were (i) the acknowledgement of liquidity constraints (Dolde and Tobin, 1971; Hubbard et al., 1986), (ii) the introduction of the buffer-stock model (Deaton, 1991; Carroll, 1997), (iii) the suggestion of the rule-of-thumb approach (Campbell and Mankiw, 1989), as well as the notion of the hand-to-mouth (iv) both poor and wealthy (Kaplan et al., 2014), as well as illiquid and liquid (Olafsson and Pagel, 2018a). One can argue that, with all these adjustments, the theory is now back at a relation between consumption and disposable income (instead of permanent income), which was originally suggested by Keynes (1936) almost one hundred years ago.

2.2 Topics around Financial Wealth and Age

Liquidity constraints (Keynes, 1936; Dolde and Tobin, 1971; Hubbard et al., 1986) describe situations where individuals cannot seamlessly borrow money as suggested by the Life Cycle/Permanent Income Hypothesis (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978). Recognizing the existence of liquidity constraints has led to the hand-to-mouth model for poor people (Johnson et al., 2006; Kaplan et al., 2014; Kaplan and Violante, 2014).

More recently, awareness has grown about the impact of wealth inequality on the MPC. Carroll et al. (2017) build a consumption model that captures the wealth distribution and explains resulting heterogeneity in the MPC.

Regarding the effect of age, empirical evidence has shown that consumption decreases upon retirement and individuals increase savings (Hamer-

²The utility function usually satisfies two conditions: (i) more consumption leads always to higher utility, and (ii) the utility increase per unit of consumption decreases. Mathematically speaking, the first derivative has to be positive and the second derivative negative.

mesh, 1982; Mariger, 1987; Robb and Burbidge, 1989; Banks et al., 1998; Bernheim et al., 2001; Schwerdt, 2005; Haider and Stephens Jr, 2007; Olafsson and Pagel, 2018b). This is the opposite of what Life Cycle/Permanent Income Hypothesis models (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978; Deaton, 1991; Carroll, 1997; Kaplan et al., 2014) would suggest because retirement is generally a predictable event.

Normal pension age varies across Europe. In 2016, around the time when the data for the third (2017) wave of the HFCS (which is used in this study), was collected, it was at 65.5 in the Netherlands, 65 years in Germany, 61.6 years in France, and in Italy at 66.6 years for men and 65.6 years for women (OECD, 2017). Future normal pension ages for job starters aged 20 years in 2016 varied even more and were at 71 years in the Netherlands, 65 years in Germany, 63 years in France, and 71.2 years in Italy. Likewise, there is large variation in pension adequacy between these four countries. The European Commission (2018) reports aggregate replacement ratios for 2016 at 50% in the Netherlands, 46% in Germany, 68% in France, and 69% in Italy.

2.3 Machine Learning Methods in Economics and MPC Research

In economics, machine learning methods have long been frowned upon as being opaque (black-box) and not concerning causality. The introduction of the causal random forest (Athey and Imbens, 2015, 2016; Wager and Athey, 2018; Athey and Wager, 2019; Athey et al., 2019) as an extension of the random forest (Ho, 1995; Breiman, 2001) that focuses on causality paved the way for its adoption in economics and finance. Since then, causal forests and similar techniques related to heterogeneous treatment effects have been frequently applied in economics and finance (Mullainathan and Spiess, 2017; Davis and Heller, 2017; Bertrand et al., 2017; Strittmatter, 2018; Medina and Pagel, 2021; Bernard, 2022).

In consumption research, Andini et al. (2018) use decision trees to predict consumption responses of individuals to a tax rebate. Their goal is to identify a group to which the tax rebate is most useful. Lewis et al. (2019) use an unsupervised clustering method to identify latent heterogeneity in the MPC distribution. Dutt and Radermacher (2022) analyze MPC heterogeneity in varying income shock sizes by means of regularized regression and a random forest.

3 Data and Methodology

3.1 Data & Summary Statistics

We use data from the third wave (HFCN, 2020b) of the Eurosystem’s Household Finance and Consumption Survey (HFCS) for the Netherlands, Germany, France, and Italy. The data samples are relatively representative for the overall populations in these countries.

Figure 1 on page 6 shows histograms and bilateral scatter plots for the Dutch data set. The histograms and bilateral scatter plots for Germany, France, and Italy are depicted in the appendix in Figure A.1, Figure A.2, and Figure A.3, respectively.




Figure 1: Descriptive plots for the Netherlands. CashOnHand is winsorized at 1% and 98% for visualization purposes. All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

The descriptive plots for the four Eurozone countries are broadly similar. The MPC is spiking at 0, 50, and 100 for Germany, France and Italy. For the Netherlands, there are almost no responses at 100 and spikes are at 0, 25, and 50, instead. Age is similarly distributed across countries as a continuous curve with most observations around age 40 to 70. The relative share above 70 is larger for Italy than for the other three countries, which generally reflects life expectancy and the demographic situation. Cash-on-hand has been winsorized at the 1% and 98% level for visualization purposes. The distributions look similar, but are easily distorted by extremes such that we are cautious regarding over-interpretation. The gender split is at roughly two-thirds vs. one-third for all countries, indicating that a majority of respondents has been male. A clear majority of households have one or two members, about a quarter have a family size of three or more.

3.2 Decision Trees

Decision trees were developed by Fisher (1936) and popularized in a machine learning context by Breiman et al. (1984). Decision trees are graph networks with parent-nodes, bifurcating edges, and child-nodes, which inductively become new parent-nodes. The objective at each bifurcation/split is to minimize impurity in the subsequent child-nodes. Splitting (and tree growing) generally stops when all nodes are pure, i.e. all elements in the node have the same class. Decision trees can fit the training data perfectly, but are prone to overfitting. They are therefore regarded as weak learners. By parameter choices, decision trees can be contained (pruned) to limit an overfitting of the data. Likewise, several decision trees can be combined into ensemble methods such as a random forest to derive more robust and generalizable results.

3.3 Random Forest

A random forest is an ensemble method combining several decision trees into a joint model. The method has been first used by Ho (1995) and had been substantially extended by Breiman (2001). The idea is to combine several weak learners (decision trees - because they are prone to overfitting) into a combined tool that uses the strengths of the individual learners (very good mapping of the underlying data) in a more robust manner (all decision trees vote on outcomes and majority wins) that produces more generalizable results.

The random forest is the workhorse of data scientists. Unlike economists (who use regression because they care about generalization and causality), data scientist are concerned about prediction and model performance on unseen data.

3.4 XGBoost

Whereas random forests use bootstrap aggregation (bagging) to grow several decision trees in parallel and independent of each other, it is also possible to grow trees sequentially taking into account prior information and performance. This method is referred to as boosting. XGBoost is a popular algorithm which makes use of boosting in a gradient-descent manner.³ It was developed by Chen et al. (2015) at Microsoft and performed very successfully at several online machine learning competitions, where it gained its popularity.

4 Analysis

4.1 Standard Linear Model

We start with an attempt to reproduce the results by Christelis et al. (2019) for the MPC in response to a one-time windfall. Christelis et al. (2019) use data from an online-survey of a representative sample of Dutch households (Center Internet panel) conducted in 2015. They apply the following regression:

$$\begin{aligned}
 MPC_i = & \beta_0 + \beta_1 Age_BTW_35_50_i \\
 & + \beta_2 Age_BTW_50_65_i \\
 & + \beta_3 Age_GEQ_65_i \\
 & + \beta_4 Male_i \\
 & + \beta_5 FamilySize_i \\
 & + \beta_6 CashOnHand_Q2_i \\
 & + \beta_7 CashOnHand_Q3_i \\
 & + \beta_8 CashOnHand_Q4_i + \epsilon_i
 \end{aligned} \tag{1}$$

³See Hastie et al. (2009) for a general introduction, Schapire (1990) for the initial idea on boosting, and Friedman (2001) for the initial idea on gradient boosting.

The dependent variable *MPC* is here the marginal propensity to consume out of a windfall equal to the size of the household's monthly net income (Christelis et al., 2019).⁴ The variables *Age_BTW_35_50*, *Age_BTW_50_65*, and *Age_GEQ_65* are dummies indicating respondents being aged 35-49, 50-64, and 65+, respectively. The variables *CashOnHand_Q2*, *CashOnHand_Q3*, and *CashOnHand_Q4* indicate that the household's cash-on-hand (i.e., liquid household wealth) lies in the second, third, and fourth quartile of the sample population, respectively. The categorization of the continuous variables, age and cash-on-hand, into the dummies, as presented in Equation 1, seems reasonable but there is no explicit motivation for it by the authors. One might ask why the variables were split specifically into quartiles as opposed to, e.g., tertiles, quintiles, or deciles. And why age splits were made at 35, 50, and 65; and not at, e.g., 25, 40, 55, or 70?

Christelis et al. (2019) find *Age_GEQ_65* and the constant to be significant at the 1% significance level, *Age_BTW_50_65* and *Male* to be significant at the 5% significance level, and *Age_BTW_35_50* to be significant at the 10% significance level. The signs of the respective coefficients of the age dummies are all positive, which they interpret as being in line with Life Cycle/Permanent Income Hypothesis models (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978). Because life expectancy decreases with age, consumption shares out of transitory income shocks should increase due to the shorter remaining spending period.

We repeat this analysis with a similar dataset: the Dutch part of the third (2017) wave of the HFCS (HFCN, 2020b). This wave also contains an MPC question on a similar shock as in the one-month-up scenario for the CentER Internet panel, which is used by Christelis et al. (2019). Yet, the response categories are more limited. The CentER Internet panel distinguishes between (i) saving, (ii) debt repay, (iii) consumption of durable goods, and (iv) consumption of non-durable goods (Christelis et al., 2019). The HFCS question only distinguishes between (i) spending on goods and services and (ii) saving, investing, or repaying debt out of a one-time windfall equal to monthly household net income (HFCN, 2020a).

To replicate Christelis et al. (2019), we make the simplifying assumption that their variable of non-durable consumption is similar to our variable of spending on goods and services. This is equivalent to interpreting durable consumption more as investing rather than as spending. We are aware that this is not entirely correct, but also not completely wrong. We thus regard it as a justifiable step.

⁴This regression corresponds to the second column in Table 3 of Christelis et al. (2019).

We attain some further reassurance by comparing the histogram of our MPC variable from the HFCS (top-left in Figure 1 on page 6) with the histogram of their MPC variable (top left in Figure 3 of Christelis et al. (2019)). The right-hand side ($MPC > 50$) is very similar for both distributions. The left-hand side ($MPC < 50$), however, looks quite different. Almost 40% of respondents in the variable used by Christelis et al. (2019) report a value of 0, about 10% of 10 and 20, and about 5% for 30, 40, and 50, each. In our variable, about 25%, 30%, and 35% report MPC scores of 0, 25, and 50, respectively. Even if the variable for durable consumption in the dataset used by Christelis et al. (2019) was part of the spending category in our dataset, this would not square the two different datasets. MPC scores just seem to differ somewhat between the two samples.⁵

Using the HFCS dataset (HFCN, 2020b), we repeat the analysis performed by Christelis et al. (2019). The first column in Table 1 on page 12 corresponds to Equation 1 on page 8. Like Christelis et al. (2019) we do not find significance for the cash-on-hand dummies in the Dutch data, but we are also not able to reproduce the significant results for the age dummies that were found by the authors. We also repeat our analysis for the three largest Eurozone countries, Germany, France, and Italy. The results are presented in the appendix in Table B.1, Table B.2, and Table B.3, respectively. For these three countries, we do find significant coefficients for the age dummies.⁶ In case of Germany and France, also with the same positive sign as Christelis et al. (2019) found for the Netherlands. According to them, the interpretation would be in line with Life Cycle/Permanent Income Hypothesis models (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978). For Italy, the signs are negative. Hence, at odds with these models.

4.2 Decision Trees

Since we did not find the same results for the Netherlands as Christelis et al. (2019), we considered that this might be due to the way the age and cash-on-hand dummies were generated. In order to detect natural splitting points in the continuous variables, we make use of simple decision trees on two hypothetical classification problems: (i) predicting low MPC scores

⁵Both data sources are a hypothetical question in a survey. We should be therefore be careful when interpreting the data. This is discussed further in the discussion section.

⁶Note that sample sizes are 4940 for Germany, 13685 for France, and 7420 for Italy. They are thus significantly larger than our sample size for the Netherlands at 1735 and that of Christelis et al. (2019) at 1208. Note also that with larger samples, p-values decrease and variables are therefore more likely to be interpreted as significant (Demidenko, 2016).

(<35) and (ii) predicting high MPC scores (>65). The trees are depicted in Figure 2 on page 11.

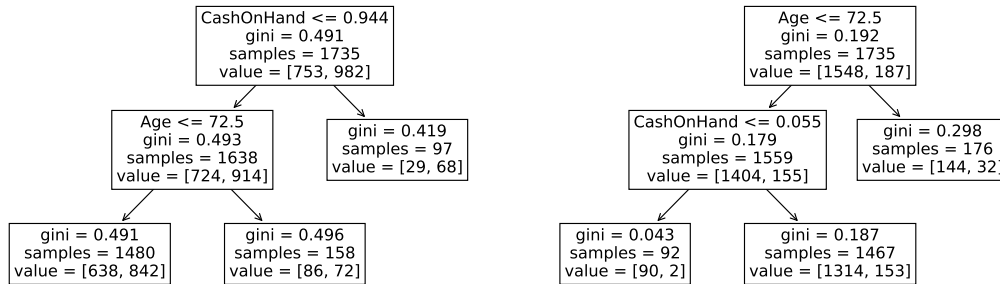


Figure 2: Decision trees for the Netherlands. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCN (HFCN, 2020b).

4.3 Customized Regressions

The splitting rules of the trees are used to determine the age and cash-on-hand dummies. The resulting regressions are shown in Table 1 on page 12. The second column uses the splits suggested by the decision tree of the classification on low-MPC and the third column uses the splits suggested by the decision tree of the classification on high-MPC. The fourth column uses all the variables that were significant in column two and three.

Our final specification includes the following significant dummies for age and cash-on-hand: *Age_GEQ_73* (age greater than or equal to 73 years), *CashOnHand_LEQ_p055* (cash-on-hand percentile less than or equal to 0.055), and *CashOnHand_GEQ_p944* (cash-on-hand percentile greater than or equal to 0.944). The signs of the coefficients for *Age_GEQ_73* and *CashOnHand_GEQ_p944* are in line with Life Cycle/Permanent Income Hypothesis models (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978). The sign of the coefficient of *CashOnHand_LEQ_p055* is not.

We repeat this procedure for our samples from Germany, France, and Italy. Decision trees and regression results are in the appendix in Figure A.4 and Table B.1, Figure A.5 and Table B.2, and Figure A.6 and Table B.3, respectively.

For Germany, *Age_LEQ_37* (age less than or equal to 37 years), *Age_GEQ_63* (age greater than or equal to 63 years), and *CashOnHand_GEQ_p856* (cash-

	(1)	(2)	(3)	(4)
	MPC	MPC	MPC	MPC
Age_BTW_35_50	0.581 (2.273)			
Age_BTW_50_65	0.538 (2.101)			
Age_GEQ_65	2.608 (2.045)			
Male	-2.994** (1.428)	-2.606* (1.405)	-2.941** (1.405)	-2.672** (1.401)
FamilySize	0.506 (0.620)	0.533 (0.571)	0.525 (0.571)	0.568 (0.569)
CashOnHand_Q2	1.441 (1.815)			
CashOnHand_Q3	1.259 (1.831)			
CashOnHand_Q4	0.2962 (1.866)			
Age_GEQ_73		6.826*** (2.137)	5.999*** (2.138)	6.517*** (2.135)
CashOnHand_LEQ_p055			-7.843*** (2.804)	-8.408*** (2.798)
CashOnHand_GEQ_p944		-9.691*** (2.766)		-10.148*** (2.764)
Intercept	31.960*** (2.233)	33.436*** (1.496)	33.643*** (1.505)	33.925*** (1.501)
Obs.	1735	1735	1735	1735
R-squared	0.004	0.014	0.012	0.019

Table 1: Regression results for the Netherlands. Regression (1) takes the specification by Christelis et al. (2019), regressions (2) and (3) take splits suggested by the decision trees, regression (4) is a combination of (2) and (3). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

on-hand percentile greater than or equal to 0.856) are significant. All signs of the coefficients are in line with Life Cycle/Permanent Income Hypothesis models (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978).

For France, *Age_LEQ_47* (age less than or equal to 47 years), *Age_GEQ_58* (age greater than or equal to 58 years), *CashOnHand_LEQ_p144* (cash-on-hand percentile less than or equal to 0.144), and *CashOnHand_GEQ_p525* (cash-on-hand percentile greater than or equal to 0.525) are significant. The signs of the coefficients for *Age_LEQ_47* and *Age_GEQ_58* are in line with Life Cycle/Permanent Income Hypothesis models (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978). The signs of the coefficients of *CashOnHand_LEQ_p144* and *CashOnHand_GEQ_p525* are not.

For Italy, *Age_GEQ_68* (age greater than or equal to 68 years), *CashOnHand_LEQ_p371* (cash-on-hand percentile less than or equal to 0.371), and *CashOnHand_GEQ_p560* (cash-on-hand percentile greater than or equal to 0.560) are significant. The signs of the coefficients for *CashOnHand_LEQ_p371*

and *CashOnHand_GEQ_p560* are in line with Life Cycle/Permanent Income Hypothesis models (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978). The sign of the coefficient of *Age_GEQ_68* is not.

Apart from these implications for consumption theory, a comparison of the four regression specifications in Table 1 on page 12 is useful. The model fit increases from an R-squared of 0.004 for the original specification by Christelis et al. (2019) to 0.019 for the specification that combines the insights from the two simple classification trees. Critics might call this p-value hacking⁷, but we (as proponents) would like to advocate this as a data-driven way to generate customized regression specifications. Understanding this concept will also help in understanding how more complex machine learning algorithms work.

Additionally, the data-splits suggested by the decision trees can be interpreted as information revealed about country/sample differences in age and liquid wealth distributions. Legal age of pension entry, life expectancy, living standard, wealth level, etc. differ across Germany, France, Italy, and the Netherlands. So why should we apply the same splits on continuous variables to create dummies? The results suggest that for the Netherlands, one age split at 73 years and two liquid wealth splits at the 5.5-percentile and 94.4-percentile fit the data better and provide significant results in a traditional econometric manner than the splits suggested by Christelis et al. (2019) at an age of 35 years, 50 years, and 65 years; and a cash-on-hand at the 25.0-percentile, 50.0-percentile, and 75.0-percentile. For Germany, this holds true if two age splits are made at 37 years and 63 years and one cash-on-hand split at the 85.6-percentile. For France and Italy, the specifications perform equally well as the one suggested by Christelis et al. (2019) if age is split at 47 and 58 years and cash-on-hand at the 14.4-percentile and 52.5-percentile for France; and age at 68 years and cash-on-hand at the 37.1-percentile and 56.0-percentile for Italy.

4.4 Random Forest

We further substantiate our analysis of how the MPC is affected by age and (liquid) wealth, i.e. cash-on-hand, by the use of more robust (and more complex) machine learning methods. Random forests, consisting of several trees, are grown for the two classification problems (low and high MPC).

⁷To these critics we kindly suggest the perusal of Demidenko (2016) and Hahn and Ang (2017) on problems with p-values altogether and new approaches to statistical reporting. Problems around publication bias due to p-values are also present in top economic journals (Brodeur et al., 2020).

Cross validation is used to optimally utilize the available data and to select a best-performing model. Model quality scores are presented in Table B.4 in the appendix.

The random forest model for Dutch data is decent considering the relatively small sample size, the few explanatory variables, and the social science context.⁸ Because of these deficiencies, however, results should not be over-interpreted.




Figure 3: Feature importance of random forest analysis for the Netherlands. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

Figure 3 on page 14 depicts the feature importance for the random forest analysis of the two classification problems (low and high MPC). Cash-on-hand is in both cases the variable with the highest feature importance, followed by age, and family-size. Being male does not seem to matter. But this is almost by design in this random forest set-up. Because being male is a binary variable, it can only be used at one (great-)parental node for each child in a decision tree. The variables cash-on-hand and age can be split much more often and can therefore be frequently used in the trees and resulting forests.

SHAP (SHapley Additive Explanations) (Lundberg and Lee, 2017) enables us to look at the effect that individual observations have on decision outcomes. These values are depicted in Figure 4 on page 15. For the low-MPC classification (left), we observe that high levels for cash-on-hand have large positive effects, i.e. affecting the model output towards low-MPC levels. Likewise, high age and large family-size had negative effects, i.e. affecting the model output towards non-low-MPC levels (i.e. towards medium or high MPC levels). The results for the high-MPC classification (right) can be interpreted in a similar way. All suggested results are in line with Life

⁸Model quality increases with sample size as can be seen in Table B.4 in the appendix for Germany, France, and Italy.




Figure 4: SHAP values of random forest analysis for the Netherlands. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

Cycle/Permanent Income Hypothesis models (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978). The results for Germany, France and Italy are in Figures A.7 to A.12 in the appendix and can be interpreted in a similar way.

4.5 XGBoost

Ordinary random forests draw trees in parallel and independent of each other. XGBoost takes accrued information into account and builds trees sequentially to optimize results overall.

Like the random forest model, also the XGBoost model for Dutch data is decent considering the relatively small sample size, the few explanatory variables, and the social science context.⁹ Because of these deficiencies, however, results should not be over-interpreted.

The results show that with XGBoost also the other variables gain importance. This makes sense given that XGBoost grows trees sequentially, "specializing" on new information that has not been included yet. As model performance of the XGBoost models are comparable to those of the random forests, this makes a case for the importance of information in the variables *FamilySize* and *Male*.

The interpretation of the SHAP values of the XGBoost models in Figure 6 on page 16 is similar to the interpretation of the SHAP values of the random forests in Figure 4 on page 15. Likewise, the results are in line with Life Cycle/Permanent Income Hypothesis models (Modigliani and Brumberg, 1954; Friedman, 1957; Hall, 1978). The results for Germany, France and

⁹Model quality increases with sample size as can be seen in Table B.4 in the appendix for Germany, France, and Italy.




Figure 5: Feature importance of XGBoost analysis for the Netherlands. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

Italy are in Figures A.13 to A.18 in the appendix and can be interpreted in a similar way.




Figure 6: SHAP values of XGBoost analysis for the Netherlands. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

5 Limitations and Discussion

5.1 Survey Data

Survey data is frequently and popularly used in research on the MPC (Carroll, 1997; Parker, 1999; Shapiro and Slemrod, 2003; Johnson et al., 2006; Parker et al., 2013; Kaplan et al., 2014; Christelis et al., 2019; Lewis et al., 2019; Christelis et al., 2019; Jappelli and Pistaferri, 2020). But survey data can also come with severe measurement errors as discussed by Karlan and Zinman (2008); Olafsson and Pagel (2018a); Radermacher (2022). Not all of these can be regarded as random noise, some might lead to structural biases. Common issues on the respondent-side (which might also apply to our data) are: (i) concerns about data privacy, (ii) social stigma about responses, (iii) lack of incentive for providing correct answers, (iv) misunderstanding the question, and (v) a gap between reported and actual behavior. Since we did not elicitate the data ourselves, we can further make no absolute claims on the quality of the survey conduct, i.e. errors on the surveyor-side. All these aspects should be kept in mind when interpreting our data and results.

5.2 Parameter Choices

Outcomes of machine learning models such as random forests or XGBoost vary with the parametrization. To make optimal use of the available data, cross validation is used. For the random forest, five-fold cross-validation is used on a grid of 18 parametrization candidates, which consists of 10 or 100 trees in the forest; 4, 8, or 16 levels as maximum tree depth; and 10, 50, and 100 samples as minimum per node. For XGBoost, the parameter grid is the same except for the minimum nodes per tree, which is replaced by a learning rate of 0.1, 0.2, and 0.3.

For the MPC classification problem, MPC scores are split into low (below 35), medium (35 to 65), and high (above 65) in accordance with the observation that for most countries, MPC scores spike at 0, 50, and 100. Given that the MPC is not an exact measure, but only a subjective response to a hypothetical question, the choice seems reasonable. Since MPC scores for the Netherlands do not spike at 100, but are mostly at or below 50, another categorization could be tried in subsequent research. Robustness checks for the other countries with differing splits (e.g. at 30 and 70; or at 25, 50, 75; or derived from a decision tree) could also be appropriate.

5.3 Why Do Age and Wealth Splits Differ per Country?

Table 2 presents key statistics (aggregate replacement rate, normal retirement age, future normal retirement age) for the Netherlands, Germany, France and Italy in combination with the upper age splits used by Christelis et al. (2019) and those suggested by the decision trees which were used in our final model.

	NL	DE	FR	IT
Aggregate replacement ratio for pensions in 2016	0.50	0.46	0.68	0.69
Normal retirement age in 2016	65.5	65	61.6	66.6 (65.6)
Future normal retirement age for those born in 1996	71	65	63	71.2
Highest age split suggested by Christelis et al. (2019)	65 [0.004]	65 [0.009]	65 [0.074]	65 [0.037]
Highest age split suggested by decision tree in final model	73 [0.019]	63 [0.016]	58 [0.074]	56 [0.039]

Table 2: Pension data and model outputs. Round brackets indicate differing retirement age for women. Square brackets indicate R-squared score of the respective regression. Data on the replacement ratio is taken from EuropeanCommission (2018). Data on retirement age is taken from OECD (2017). Data for the regression models is taken from HFCN (2020b).

The age split applied by Christelis et al. (2019) is close to the normal retirement ages in 2016 for the Netherlands, Germany and Italy. Our age splits for the Netherlands and Germany are very close to their future normal retirement ages. These two countries also have an aggregate replacement rate at or below 50%, significantly lower than France and Italy. In these countries, the replacement ratios are around 70% - pension entry should therefore not have such severe financial impacts as in the Netherlands and Germany. Our upper age splits are significantly lower than those applied by Christelis et al. (2019) and we still produce the same or a slightly better regression fit although we use less variables. For the Netherlands and Germany the improvement in regression fit is even stronger from 0.004 to 0.019 and from 0.009 to 0.016, respectively.

We interpret the results in a way that differing wealth distributions and retirement systems across the studied Eurozone members have an impact on consumption behavior. This should be relevant to policy makers to take inter-country heterogeneity in consumption patterns into account when

conducting joint Eurozone monetary policy and introducing increasingly supranational fiscal tools at the authority of the EU.

5.4 Integrating Decision Trees into Econometric Analysis

I showed how simple decision trees can be used for econometric analysis, when one desires to create data-driven dummies out of continuous variables. Although decision trees are not very robust, they are extremely intuitive, transparent, and easy to communicate to regulators. Applying them in combination with traditional econometric analysis is a good mix between (black-box) machine learning and traditional econometrics without the need for large data sources or much computing power.

Andini et al. (2018) of the Italian central bank, also use decision trees in their analysis of whom best to target for a tax rebate in Italy. They suggest that policies should be better targeted to an audience that benefits from the policy. Identifying who could best benefit from a policy is seen as a classical prediction task, where machine learning excels at. Lewis et al. (2019) also use machine learning algorithms on consumption response data and discover considerable levels of heterogeneity in the response to transitory income shocks.

Finally, we would like to emphasize that also traditional econometric approaches apart from OLS might be beneficial for future MPC research. Censored regression could be appropriate since the MPC scores are on a censored scale from 0 to 100. Logistic regression can be used for binary classification instead of decision trees. And quantile regression can be a good alternative to deal with non-linear data. Comparing the outcomes of these methods to the OLS-models that were augmented by machine learning would be a good exercise for future research.

6 Conclusion

We analyze consumption patterns across Germany, France, Italy, and the Netherlands assessing the age and wealth impact on the marginal propensity to consume (MPC). We take Christelis et al. (2019) as a starting point for our regression analysis and make use of decision trees to derive data-driven distribution splits for age and wealth dummies for subsequent analysis. Our customized regression specifications feature higher or equal model fit than the initial specifications and also significant coefficients.

We interpret the results as indication for structural differences in age and wealth across the investigated Eurozone countries, which affect the MPC. These differences can arise from distinct regulatory systems, e.g. concerning retirement age or pension type. Our analysis is further substantiated by the use of more advanced supervised machine learning methods which are built upon decision trees. Also the results of this analysis point out country-differences and heterogeneity regarding the MPC across Europe.

Our findings are of particular importance to European policymakers who conduct joint monetary policy for the Eurozone and see increasing fiscal influence on the supranational EU level. Comparing consumption patterns between countries enables to better assess the impact of changes in taxes, transfers, inflation, interest, and exchange rates on individual purchasing power and key macroeconomic variables like aggregate demand.

Using decision trees, this paper has additional methodological and educational purposes. Methodologically, it is shown how decision trees as simple, flexible, and intuitive tools can be used in combination with the more rigid traditional linear regression. This is particularly useful when dealing with non-linear data. Educationally, decision trees are interesting as they are the building blocks of more complex supervised machine learning tools. Its use in this paper provides intuition on what algorithms like random forests or XGBoost actually do.

References

- Andini, M., Ciani, E., de Blasio, G., D'Ignazio, A., and Salvestrini, V. (2018). Targeting with machine learning: An application to a tax rebate program in Italy. *Journal of Economic Behavior & Organization*, 156:86–102.
- Athey, S. and Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360.
- Athey, S. and Imbens, G. W. (2015). Machine learning methods for estimating heterogeneous causal effects. *stat*, 1050(5):1–26.
- Athey, S., Tibshirani, J., and Wager, S. (2019). Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178.
- Athey, S. and Wager, S. (2019). Estimating treatment effects with causal forests: An application. *Observational Studies*, 5(2):37–51.
- Banks, J., Blundell, R., and Tanner, S. (1998). Is there a retirement-savings puzzle? *American Economic Review*, pages 769–788.
- Bernard, R. (2022). Mental accounting and the marginal propensity to consume. Technical report, Working Paper.
- Bernheim, B. D., Skinner, J., and Weinberg, S. (2001). What accounts for the variation in retirement wealth among US households? *American Economic Review*, 91(4):832–857.
- Bertrand, M., Crépon, B., Marguerie, A., and Premand, P. (2017). Contemporaneous and post-program impacts of a public works program.
- Breiman, L. (2001). Random forests. *Machine learning*, 45(1):5–32.
- Breiman, L., Friedman, J., Olshen, R., and Stone, C. (1984). Classification and regression trees.
- Brodeur, A., Cook, N., and Heyes, A. (2020). Methods matter: P-hacking and publication bias in causal analysis in economics. *American Economic Review*, 110(11):3634–60.
- Campbell, J. Y. and Mankiw, N. G. (1989). Consumption, income, and interest rates: Reinterpreting the time series evidence. *NBER macroeconomics annual*, 4:185–216.

- Carroll, C., Slacalek, J., Tokuoka, K., and White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3):977–1020.
- Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly journal of economics*, 112(1):1–55.
- Chen, T., He, T., Benesty, M., Khotilovich, V., Tang, Y., Cho, H., Chen, K., et al. (2015). Xgboost: extreme gradient boosting. *R package version 0.4-2*, 1(4):1–4.
- Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., and Van Rooij, M. (2019). Asymmetric consumption effects of transitory income shocks. *The Economic Journal*, 129(622):2322–2341.
- Davis, J. and Heller, S. B. (2017). Using causal forests to predict treatment heterogeneity: An application to summer jobs. *American Economic Review*, 107(5):546–50.
- Deaton, A. (1991). Saving and liquidity constraints. *Econometrica: Journal of the Econometric Society*, pages 1221–1248.
- Demidenko, E. (2016). The p-value you can't buy. *The American Statistician*, 70(1):33–38.
- Dolde, W. and Tobin, J. (1971). Wealth, liquidity, and consumption. Technical report, Cowles Foundation for Research in Economics, Yale University.
- Dutt, S. and Radermacher, J. W. (2022). Heterogeneity and the mpc: Observations from machine learning. Technical report.
- EuropeanCommission (2018). Social protection committee annual report 2018. 2018 spc annual review of the social protection performance monitor (sppm) and developments in social protection policies. report on key social challenges and main messages.
- Fisher, R. (1936). Linear discriminant analysis. *Ann. Eugenics*, 7:179.
- Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232.
- Friedman, M. (1957). *Theory of the consumption function*. Princeton University Press.

- Hahn, E. D. and Ang, S. H. (2017). From the editors: New directions in the reporting of statistical results in the journal of world business. *Journal of World Business*, 52(2):125–126.
- Haider, S. J. and Stephens Jr, M. (2007). Is there a retirement-consumption puzzle? evidence using subjective retirement expectations. *The review of economics and statistics*, 89(2):247–264.
- Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence. *Journal of political economy*, 86(6):971–987.
- Hamermesh, D. S. (1982). Consumption during retirement: The missing link in the life cycle. Technical report, National Bureau of Economic Research.
- Hastie, T., Tibshirani, R., and Friedman, J. H. (2009). *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer.
- HFCN (2020a). Euro area questionnaire on household finance and consumption plus paradata. Technical report, ECB Statistics Paper.
- HFCN (2020b). The household finance and consumption survey: methodological report for the 2017 wave. Technical report, ECB Statistics Paper.
- Ho, T. K. (1995). Random decision forests. In *Proceedings of 3rd international conference on document analysis and recognition*, volume 1, pages 278–282. IEEE.
- Hubbard, R. G., Judd, K. L., Hall, R. E., and Summers, L. (1986). Liquidity constraints, fiscal policy, and consumption. *Brookings Papers on Economic Activity*, 1986(1):1–59.
- Jappelli, T. and Pistaferri, L. (2020). Reported mpc and unobserved heterogeneity. *American Economic Journal: Economic Policy*, 12(4):275–97.
- Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5):1589–1610.
- Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.
- Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

- Kaplan, G., Violante, G. L., and Weidner, J. (2014). The wealthy hand-to-mouth. Technical report, National Bureau of Economic Research.
- Karlan, D. and Zinman, J. (2008). Lying about borrowing. *Journal of the European Economic Association*, 6(2-3):510–521.
- Keynes, J. M. (1936). *The general theory of employment, interest, and money*. MacMillan.
- Lewis, D. J., Melcangi, D., and Pilosoph, L. (2019). Latent heterogeneity in the marginal propensity to consume. *FRB of New York Staff Report*, (902).
- Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- Mariger, R. P. (1987). A life-cycle consumption model with liquidity constraints: theory and empirical results. *Econometrica: Journal of the Econometric Society*, pages 533–557.
- Medina, P. C. and Pagel, M. (2021). Does saving cause borrowing? Technical report, National Bureau of Economic Research.
- Modigliani, F. and Brumberg, R. (1954). Utility analysis and the consumption function: An interpretation of cross-section data. *Franco Modigliani*, 1(1):388–436.
- Mullainathan, S. and Spiess, J. (2017). Machine learning: an applied econometric approach. *Journal of Economic Perspectives*, 31(2):87–106.
- OECD (2017). *Pensions at a Glance 2017: OECD and G20 Indicators*. OECD Publishing Paris.
- Olafsson, A. and Pagel, M. (2018a). The liquid hand-to-mouth: Evidence from personal finance management software. *The Review of Financial Studies*, 31(11):4398–4446.
- Olafsson, A. and Pagel, M. (2018b). The retirement-consumption puzzle: New evidence from personal finances. Technical report, National Bureau of Economic Research.
- Parker, J. A. (1999). The reaction of household consumption to predictable changes in social security taxes. *American Economic Review*, 89(4):959–973.

- Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, 103(6):2530–53.
- Radermacher, J. W. (2022). Mamma mia! pca-biplot reveals yet another consumption puzzle: Low mpc levels among italy's elderly poor. Technical report.
- Robb, A. L. and Burbidge, J. B. (1989). Consumption, income, and retirement. *Canadian Journal of Economics*, pages 522–542.
- Schapire, R. E. (1990). The strength of weak learnability. *Machine learning*, 5(2):197–227.
- Schwerdt, G. (2005). Why does consumption fall at retirement? evidence from germany. *Economics Letters*, 89(3):300–305.
- Shapiro, M. D. and Slemrod, J. (2003). Consumer response to tax rebates. *American Economic Review*, 93(1):381–396.
- Smets, F. and Wouters, R. (2003). An estimated dynamic stochastic general equilibrium model of the euro area. *Journal of the European economic association*, 1(5):1123–1175.
- Strittmatter, A. (2018). What is the value added by using causal machine learning methods in a welfare experiment evaluation? *Working Paper*.
- Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

Appendices

A Descriptive Plots

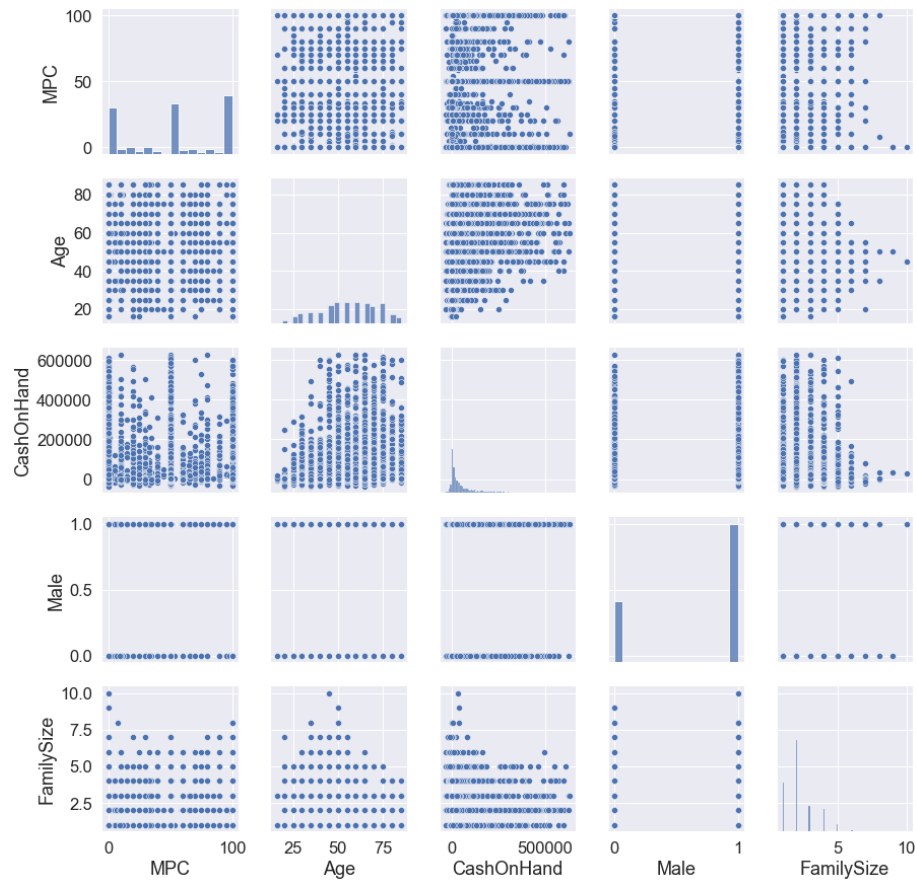


Figure A.1: Descriptive plots for Germany. CashOnHand is winsorized at 1% and 98% for visualization purposes. All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

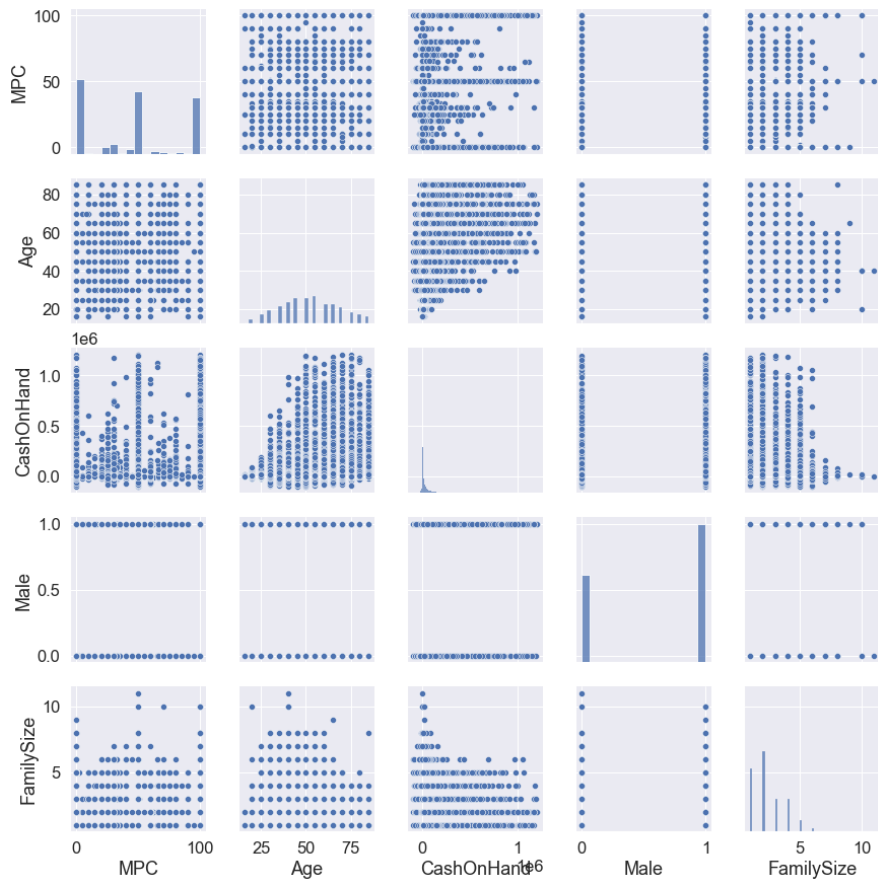


Figure A.2: Descriptive plots for France. CashOnHand is winsorized at 1% and 98% for visualization purposes. All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

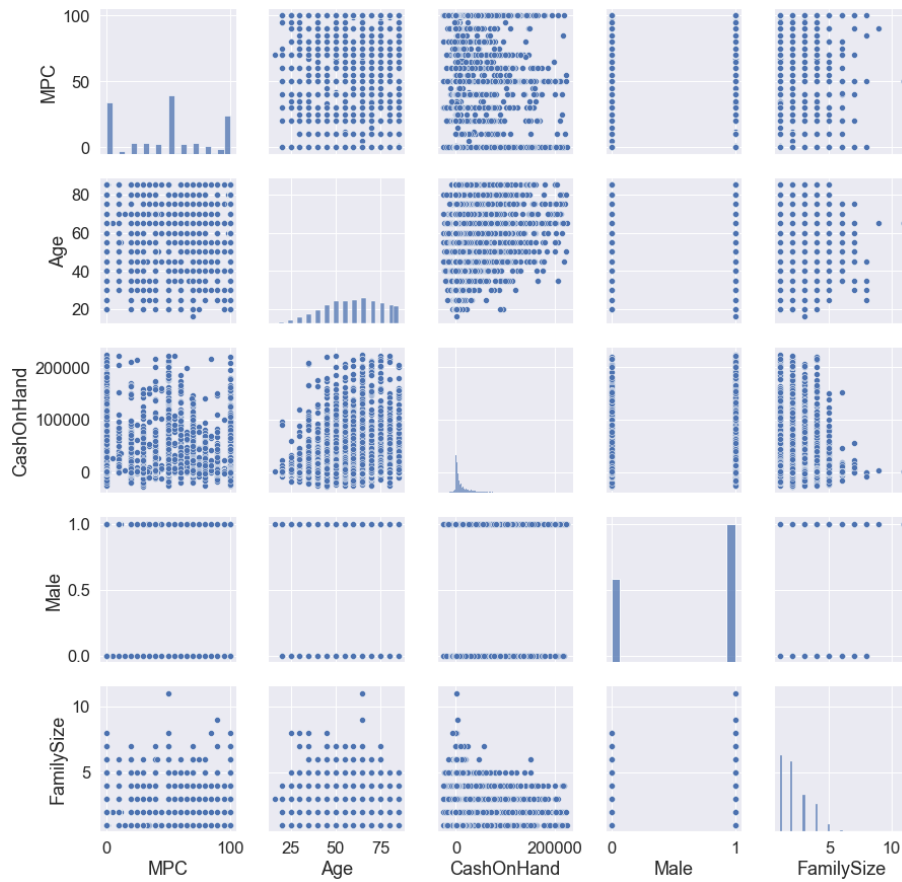


Figure A.3: Descriptive plots for Italy. CashOnHand is winsorized at 1% and 98% for visualization purposes. All data is taken from the third wave (2017) of the HFCN (HFCN, 2020b).

B Decision Trees




Figure A.4: Decision trees for Germany. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCN (HFCN, 2020b).




Figure A.5: Decision trees for France. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCN (HFCN, 2020b).




Figure A.6: Decision trees for Italy. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

C Regression Tables

	(1)	(2)	(3)	(4)
	MPC	MPC	MPC	MPC
Age_BTW_35_50	3.614*			
	(2.172)			
Age_BTW_50_65	3.805*			
	(1.983)			
Age_GEQ_65	9.520***			
	(1.977)			
Male	-2.215*	-2.253*	-2.390*	-2.126*
	(1.228)	(1.220)	(1.228)	(1.223)
FamilySize	0.493	0.5120	0.618	0.538
	(0.541)	(0.515)	(0.516)	(0.515)
CashOnHand_Q2	-1.129			
	(1.554)			
CashOnHand_Q3	-1.604			
	(1.571)			
CashOnHand_Q4	-6.220***			
	(1.597)			
Age_LEQ_37		-4.248***		-4.519***
		(1.587)		(1.598)
Age_GEQ_63		5.614***	6.186***	5.794***
		(1.279)	(1.220)	(1.285)
CashOnHand_LEQ_p272			3.946*	1.859
			(2.399)	(1.289)
CashOnHand_GEQ_p337			0.778	
			(2.257)	
CashOnHand_GEQ_p856		-10.716***		-10.205***
		(1.571)		(1.611)
Intercept	48.886***	52.262***	48.096***	51.512***
	(2.164)	(1.627)	(2.582)	(1.708)
Obs.	4940	4940	4940	4940
R-squared	0.009	0.015	0.007	0.016

Table B.1: Regression results for Germany. Regression (1) takes the specification by Christelis et al. (2019), regressions (2) and (3) take splits suggested by the decision trees, regression (4) is a combination of (2) and (3). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

	(1)	(2)	(3)	(4)
	MPC	MPC	MPC	MPC
Age_BTW_35_50	4.055*** (1.141)			
Age_BTW_50_65	10.743*** (1.102)			
Age_GEQ_65	23.175*** (1.155)			
Male	0.043 (0.674)	0.736 (0.674)	0.241 (0.674)	0.444 (0.673)
FamilySize	0.019 (0.285)	0.162 (0.276)	0.379 (0.280)	0.549 (0.281)
CashOnHand_Q2	3.478*** (0.907)			
CashOnHand_Q3	7.226*** (0.910)			
CashOnHand_Q4	12.779*** (0.935)			
Age_LEQ_47			-6.603*** (0.853)	-6.681*** (0.851)
Age_GEQ_58		16.473*** (0.738)	13.112*** (0.871)	12.673*** (0.871)
CashOnHand_LEQ_p144		-7.418*** (1.030)		-7.579*** (1.028)
CashOnHand_GEQ_p455		6.754*** (0.729)		2.142 (1.333)
CashOnHand_GEQ_p525			8.547*** (0.657)	4.642*** (1.294)
Intercept	27.585*** (1.262)	35.142*** (0.967)	37.334*** (1.010)	38.773*** (1.060)
Obs.	13685	13685	13685	13685
R-squared	0.074	0.068	0.069	0.074

Table B.2: Regression results for France. Regression (1) takes the specification by Christelis et al. (2019), regressions (2) and (3) take splits suggested by the decision trees, regression (4) is a combination of (2) and (3). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

	(1)	(2)	(3)	(4)
	MPC	MPC	MPC	MPC
Age_BTW_35_50	-4.712** (2.160)			
Age_BTW_50_65	-7.478*** (2.077)			
Age_GEQ_65	-9.968*** (2.022)			
Male	0.357 (0.863)	0.070 (0.863)	0.349 (0.862)	0.347 (0.862)
FamilySize	1.160*** (0.371)	0.898** (0.367)	0.934*** (0.361)	0.933*** (0.361)
CashOnHand_Q2	-4.931*** (1.121)			
CashOnHand_Q3	-11.476*** (1.127)			
CashOnHand_Q4	-13.768*** (1.134)			
Age_LEQ_62		0.700 (1.301)		
Age_GEQ_68		-4.892*** (1.345)	-5.404*** (0.885)	-5.400*** (0.885)
CashOnHand_LEQ_p364		11.156*** (0.832)		1.500 (4.765)
CashOnHand_LEQ_p371			7.352*** (1.120)	5.881*** (4.806)
CashOnHand_GEQ_p560			-5.626*** (1.088)	-5.626*** (1.088)
Intercept	59.362*** (2.155)	42.050*** (1.432)	46.157*** (1.323)	46.157*** (1.323)
Obs.	7420	7420	7420	7420
R-squared	0.037	0.035	0.039	0.039

Table B.3: Regression results for Italy. Regression (1) takes the specification by Christelis et al. (2019), regressions (2) and (3) take splits suggested by the decision trees, regression (4) is a combination of (2) and (3). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

D Random Forests




Figure A.7: Feature importance of random forest analysis for Germany. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).




Figure A.8: SHAP values of random forest analysis for Germany. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).




Figure A.9: Feature importance of random forest analysis for France. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).




Figure A.10: SHAP values of random forest analysis for France. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).




Figure A.11: Feature importance of random forest analysis for Italy. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).




Figure A.12: SHAP values of random forest analysis for Italy. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

E XGBoost Models




Figure A.13: Feature importance of XGBoost analysis for Germany. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).




Figure A.14: SHAP values of XGBoost analysis for Germany. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).




Figure A.15: Feature importance of XGBoost analysis for France. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).




Figure A.16: SHAP values of XGBoost analysis for France. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).




Figure A.17: Feature importance of XGBoost analysis for Italy. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).




Figure A.18: SHAP values of XGBoost analysis for Italy. The left graph shows a classification on low MPC values (< 35), the right graph shows a classification on high MPC values (> 65). All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

		NL	DE	FR	IT
Random Forest: Low MPC	AUC on training data	0.623	0.602	0.646	0.624
	AUC on test data	0.492	0.565	0.637	0.586
Random Forest: High MPC	AUC on training data	0.754	0.599	0.683	0.654
	AUC on test data	0.503	0.548	0.670	0.614
XG Boost: Low MPC	AUC on training data	0.648	0.642	0.661	0.646
	AUC on test data	0.502	0.557	0.635	0.578
XG Boost: High MPC	AUC on training data	0.785	0.663	0.701	0.665
	AUC on test data	0.476	0.533	0.672	0.615
Number of Observations		1735	4940	13685	7420

Table B.4: Model fit for supervised machine learning models. All data is taken from the third wave (2017) of the HFCS (HFCN, 2020b).

Recent Issues

No. 382	Jan Wedigo Radermacher	Mamma Mia! Revealing Hidden Heterogeneity by PCA-Biplot - MPC Puzzle for Italy's Elderly Poor
No. 381	Olga Goldfayn, Nathanael Vellekoop	The Role of Personality Traits in Household Loan Expectations and Borrowing Constraints
No. 380	Alexander Hillert, Alexandra Niessen-Ruenzi, Stefan Ruenzi	Mutual Fund Shareholder Letters: Flows, Performance, and Managerial Behavior
No. 379	Irene Mecatti, Tobias Tröger	Who Should Hold Bail-Inable Debt and How can Regulators Police Holding Restrictions Effectively?
No. 378	Sabine Bernard, Benjamin Loos, Martin Weber	How Speculative Asset Characteristics Shape Retail Investors' Selling Behavior
No. 377	Konstantin Egorov	Optimal Policy under Dollar Pricing
No. 376	Raphael Abiry, Marien Ferdinandusse, Alexander Ludwig, Carolin Nerlich	Climate Change Mitigation: How Effective is Green Quantitative Easing?
No. 375	Olivia S. Mitchell, Orly Sade, Abigail Hurwitz	Testing Methods to Enhance Longevity Awareness
No. 374	Franco Fiordelisi, Giulia Fusi, Angela Maddaloni, David Marqués-Ibáñez	Pandemic Lending: Micro and Macro Effects of Model-Based Regulation
No. 373	Emily Kormanyos, Tobin Hanspal, Andreas Hackethal	Do Gamblers Invest in Lottery Stocks?
No. 372	Alessandro Pollastri, Paulo Rodrigues, Christian Schlag, Norman Seeger	A Jumping Index of Jumping Stocks? An MCMC Analysis of Continuous-Time Models for Individual Stocks
No. 371	Mariano Massimiliano Croce, Tatyana Marchuk, Christian Schlag	The Leading Premium
No. 370	Benjamin M. Abdel-Karim, Kevin Bauer, Oliver Hinz, Michael Nofer	The Effects of Discontinuing Machine Learning Decision Support
No. 369	Katja Langenbucher	Consumer Credit in The Age of AI – Beyond Anti-Discrimination Law