Machine Learning and Financial Literacy: An Exploration of Factors Influencing Financial Knowledge in Italy

Susanna Levantesi *[ID] and Giulia Zacchia [ID]

Department of Statistics, Sapienza University of Rome, 00161 Rome, Italy; giulia.zacchia@uniroma1.it
* Correspondence: susanna.levantesi@uniroma1.it

Abstract: In recent years, machine learning techniques have assumed an increasingly central role in many areas of research, from computer science to medicine, including finance. In the current study, we applied it to financial literacy to test its accuracy, compared to a standard parametric model, in the estimation of the main determinants of financial knowledge. Using recent data on financial literacy and inclusion among Italian adults, we empirically tested how tree-based machine learning methods, such as decision trees, random, forest and gradient boosting techniques, can be a valuable complement to standard models (generalized linear models) for the identification of the groups in the population in most need of improving their financial knowledge.

Keywords: financial literacy; machine learning; decision trees; random forest; gradient boosting

1. Introduction

In the wake of the global financial crisis of 2007–2008 and of the recent events concerning the COVID-19 global pandemic crisis, the debate on the importance of financial literacy (FL) has gained further momentum, because more vulnerable and less informed investors are the most exposed to crises, and not only financial ones. As of May 2020, more than 70 countries were designing or implementing national strategies for financial literacy. Thus, the OECD developed a set of recommendations to assist governments or other public authorities to design, implement, and evaluate policies to support financial resilience and well-being, in addition to addressing the needs of vulnerable groups (OECD 2020). In fact, those who are financially illiterate have been proven to have a lower ability to cope with emergency expenses and income shocks (Hasler et al. 2018), and a lower propensity to withdraw deposits from distressed banks (Brown et al. 2016) and to leave the stock market before crashes (Guiso and Viviano 2015), than the more financially conscious population. In this sense, financial knowledge can help individuals with the process of financial decision-making and savings because it enables them to plan for wealth accumulation (Ameriks et al. 2003), to be more financially included (Grohmann et al. 2018), and to choose investments that are the most suitable for their needs, considering all the possible risks (Bianchi 2018).

At the micro level, several papers document a positive correlation between measures of FL and wiser financial decisions in various domains about both assets and debts. For example, individuals with higher levels of financial knowledge are more likely to participate in financial markets and to invest in stocks (Christelis et al. 2010; Yoong 2011; Van Rooij et al. 2011), to have better diversified portfolios (Guiso and Jappelli 2008; Von Gaudecker 2015), and to earn higher yields on deposit accounts (Deuflhard et al. 2019). Hsiao and Tsai (2018) provide evidence of a positive impact of FL on trading in leveraged derivative products, an important means of hedging financial risks in portfolios. Financial literates are less prone to over-indebtedness (Lusardi and Tufano 2009; Lusardi and Scheresberg 2013; Lusardi et al. 2016) and to choose adjustable rate mortgages instead of less risky mortgages (Gathergood and Weber 2017). They usually better perform in peer-to-peer lending markets (Chen et al. 2018) and choose mutual funds with lower fees (Hastings and

Tejeda-Ashton 2008). They are more likely to plan for retirement (Goda et al. 2020) and, as a result, to better allocate resources over their lifetimes in a world that, especially in recent years, is increasingly complex and uncertain (Clark et al. 2015; Behrman et al. 2012). Recently, Feng et al. (2019), following a Bayesian two-part latent variable modelling approach, identified the simultaneous impact of FL on household debt and assets. They found that households with insufficient financial knowledge are more financially vulnerable because they are more likely both to have fewer assets and to choose high-cost unsecured debts that expose them more to potential financial constraints.

Another stream of literature investigates the impact of FL at the macro level (Lusardi and Mitchell 2014). For example, Gerardi et al. (2010) stresses how limited financial knowledge can be considered a cause of the 2007 U.S. financial crisis. Fornero and Lo Prete (2019) found a clear association between the FL of the electorate (or the ability to understand essential concepts of economic reforms, mainly regarding the pension systems) and electoral outcomes. They conclude that "financial illiteracy may also harm reformist efforts and has clear policy implications" (Fornero and Lo Prete 2019, p. 24) in terms of successful implementation of economic reforms. Moreover, using a life-cycle approach, Lusardi et al. (2017) show that gaps in FL amplify differences in wealth accumulation patterns and the consequent perpetuation in wealth inequality. In this direction, Lo Prete (2013, 2018) empirically tested how the ability to take advantage of different financial opportunities, measured by financial knowledge, may help to reduce inequality across countries and over time. She found that the level of economic literacy is associated with income inequality across countries, using a sample of advanced and developing countries observed over the 1980–2007 period.

All of these studies have evident policy implications because inequality appears to decrease not when more complex and sophisticated financial instruments are available but only when the ability to understand and use these instruments increases among all of the population. In fact, the debate on the relationship between finance and inequality poses FL as relevant to the policy agenda of many countries, as defined above. Consequently, as the OECD (2020) underlined in the Recommendation on Financial Literacy, it is crucial to collect high-quality, comparable data on levels of financial knowledge and to analyze these "data to identify aspects of financial literacy that cause particularly significant issues as well as the groups in the population in most need of improving" (OECD 2020, p. 7).

This paper aims to contribute to the analysis of FL, extending the common methodology to machine learning (ML) techniques. Although ML has been widely used in finance (e.g., see, Dixon et al. 2020; Bracke et al. 2019; Bazarbash 2019), to the best of our knowledge there are still no analyses of ML techniques applied to financial knowledge. Nonetheless, we state that ML techniques can be valuable as a complement to standard parametric models in the study of financial literacy. Demonstrating that analytical steps of the econometric processes, like the logistic regression model that we apply to real data, has a homologous step in ML analyses, we clearly find a correspondence between parametric and ML techniques, with the goal to also facilitate and reconcile the adoption of ML techniques in the context of financial literacy. ML can meet the need of in-depth investigation, which is of paramount importance in financial literacy analyses. Due to its flexibility, the ML framework can provide more information about the heterogeneity and commonality across different subpopulations and can help researchers and policy-makers to understand the characteristics of individuals with lower levels of financial literacy and therefore at higher risk of financial fragilities. Our analysis provides preliminary evidence that ML techniques can produce reliable information for financial literacy that is consistent with the literature, although it can also identify different patterns of correlations than traditional parametric models (i.e., high variable importance of financial behavior and attitude as determinants of financial knowledge).

In detail, we propose a comparison among a parametric model (a logistic regression model) and ML models to identify the precision accuracy of the different models. We also use tree models to assess model selection and the approximate direction and functional

form of the relationships between the inputs and the output, discussing the measures of variable importance in three tree models: decision trees (Breiman et al. 1984), random forest (Breiman 2001), and gradient boosting machine (Friedman 2001). These models are classified on the basis of the outcome variable type: classification models in the case of categorical variables and regression models otherwise. Because we refer to categorical variables, the algorithms used for this study are classification models. We test empirically these models using data available for Italy collected by the Bank of Italy.

We concentrate on Italy because is a negative outlier among the most advanced economies considering the level of financial competencies of adults (Klapper et al. 2015; Di Salvatore et al. 2018). According to the Standard & Poor's Ratings Services Global Financial Literacy Survey (S&P Global FinLit Survey), only 37 per cent of adult Italians correctly understand basic financial concepts, compared with 52 per cent on average in the EU. In addition, the G20/OECD International Network for Financial Education (INFE) report on adult financial literacy shows a very low level of financial literacy in Italy compared with the G20 average (Figure 1). Thus, the financial knowledge score in Italy is 3.5 out of a maximum of 7 points on average, compared with a G20 average of 4.3. According to Di Salvatore et al. (2018), the lower level of financial knowledge in Italy can be explained by the higher share of individuals with low levels of education, in fact "about 47 per cent of the adult Italian population has a primary level of education, while the same group accounts for only 14 per cent of the population in Germany and does not exceed 10 per cent in Canada and the UK" (p. 9). We can add that Italy has also higher unemployment rates than most of the countries compared in Figure 1. Therefore, our analysis aims to test parametric and ML techniques to define the main determinants of financial literacy gaps among Italians, who on average are less financially educated than G20 citizens.
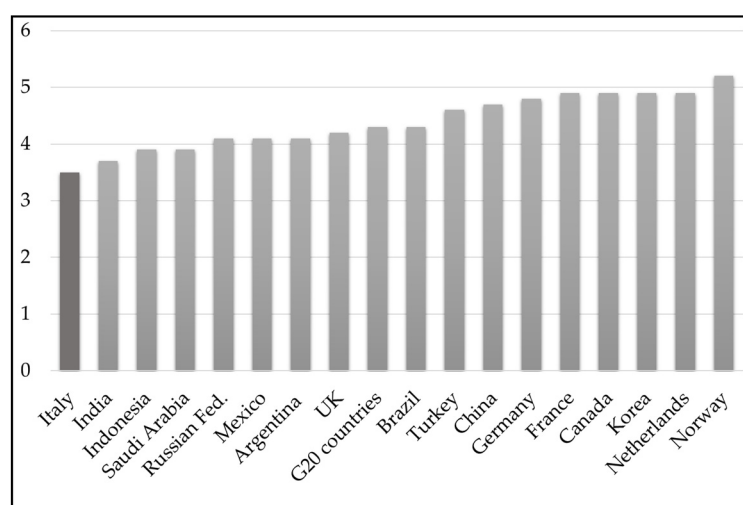
Figure 1. Financial knowledge in Italy compared to G20 countries (averages; weighted data). *Source:* OECD (2017).

We are conscious that we focus on a limited case study, but we think that it can be seen as a first step to encourage the adoption of ML techniques in applied economics and among researchers in the context of financial literacy. ML is, in fact, a transparent research tool with an important role to play because it has the advantage of: (i) focusing on out-of-sample predictability over variance adjudication; (ii) using computational methods to avoid relying on (potentially unrealistic) assumptions; (iii) having the ability to "learn" complex specifications, including non-linear, hierarchical, and non-continuous interaction effects in a high-dimensional space; and (iv) featuring importance analyses robust to multicollinearity. For all of the above reasons, ML can be useful to researchers, and policy makers or financial analysts, to analyze complex data and a large volume of information

simultaneously, thus providing a more nuanced and detailed picture of the phenomenon of financial literacy.

This paper is organized into three main sections. The first section summarizes research findings of recent literature about the main determinates of FL, providing an accurate mapping of methodologies and the main variables used to explain the phenomenon. The second section provides readers with foundational knowledge of the ML algorithms used. The third section introduces the data used and the main results of the empirical analysis. The final section summarizes research findings and identifies future research needs. The proposed ML methodology can be used above and beyond our empirical analysis, because ML offers the opportunity to gain insight from: (a) new datasets that cannot be modelled with econometric methods; and (b) old datasets that incorporate complex relationships that are still unexplored.

2. Factors Influencing Financial Knowledge: A Literature Review

Financial literacy, as described in the introduction, is increasingly attracting the attention of international organizations, financial regulators, policymakers, and academics (for a review of the most cited papers on the issue see Goyal and Kumar 2020). Findings around the world are sobering. FL is low even in advanced economies with well-developed financial markets. On average, only about one-third of the global population has a familiarity with the basic concepts that underlie everyday financial decisions (Lusardi and Mitchell 2011; Lusardi 2019).

Despite the importance of the issue, there is still no consensus on its best definition and the most suitable tools to measure the level of financial knowledge (Rieger 2020). Different data and definitions have been used, from mathematical skills at school age in Programme for International Student Assessment (PISA) test scores (Jappelli and Padula 2013) to numerical ability and other dimensions of cognitive function in older adults (Banks and Oldfield 2007). However, in the literature, an increasing number of papers (Bianchi 2018; Fornero and Monticone 2011; Kadoya and Khan 2019; Klapper and Panos 2011) uses the same measure for assessing the level of financial knowledge of adults, based on three basic concepts, commonly called the "Big Three" (Lusardi and Mitchell 2008). These three concepts, that can be easily applied to every context and economic environment, are: (1) numeracy, the capacity to do interest rate calculations, and to understand how to calculate interest compounding; (2) the knowledge of inflation and how it interacts with purchasing power; and (3) the comprehension of the importance of portfolio diversification to reduce risks. At the international level, OECD International Network for Financial Education (INFE) integrates the understanding of the three basic concepts described above with measures of financial attitude and behavior necessary to make sound financial decisions and ultimately achieve individual financial wellbeing.

Considering the main determinants of different levels of financial knowledge, the variables used in the literature are heterogeneous, depending on the countries analyzed and the different perspectives of researchers. Highlighting some common trends in determinants and main results, we can summarize an extensive literature on financial literacy by dividing the variables that correlate to financial literacy into seven categories. These categories are as follows:

1. Gender: One of the main common results in the literature is that women have lower FL than men. In fact, in 2011 the OECD found a gender gap in FL in 13 countries, with Hungary the only exception (Atkinson and Messy 2012). Bucher-Koenen et al. (2017), extending the evidence for other countries, found that only ex-Soviet countries (Russia, Romania, and East Germany) have an equal distribution of financial knowledge between sexes. However, recent literature stresses how, when asked to answer questions that measure knowledge of basic financial concepts, women are less likely than men to indicate that they do not know the answer (Bucher-Koenen et al. 2017; Kim and Mountain 2019; Ooi 2020). Therefore, the lower scores of women compared to men in financial literacy surveys reflects more the differences in the genders' self-

reported confidence than the gender differences in their actual level of financial knowledge. Al-Bahrani et al. (2020) found the origins of the gender-based financial literacy gap early in life (early college age), before individuals have the opportunity to develop financial skills through experience or specialization in household roles. Jappelli and Padula (2013) explain the gender gap in the fact that women generally have less wealth than men and therefore fewer incentives to invest in FL.

2. Education: Higher education is usually reported as one of the most important factors in ensuring an adequate understanding of financial concepts. Many studies have shown that individuals with higher levels of education, i.e., who completed a university or college degree, are the most likely to be financially literate (Lusardi and Mitchell 2008; Cole et al. 2011). In addition, Mandell (2008) and Al-Bahrani et al. (2020) have shown that the correlation between financial literacy and education is present at the early stages of lifecycle, and is highly correlated with mathematics ability. Morgan and Trinh (2019), using the OECD/INFE data for Cambodia and Viet Nam, found that both financial literacy and general education levels are found to be positively and significantly related to savings behavior and financial inclusion, also controlling for possible endogeneity of financial literacy.

3. Financial fragility: Financial knowledge is usually associated with household's income levels and financial fragility. The concept of financial fragility is of paramount importance in the period of crises (such as the COVID-19 pandemic) to understand whether households lack capacity to face shocks. The concept, as defined in Demertzis et al. (2020), encompasses the state of household balance sheets, including indebtedness, and also relies on individual perceptions of the ability to rely on family and friends and other methods to deal with shocks. Previati et al. (2020) examined financial fragility in Italy using pre-COVID-19 data, and documented the strong link between financial fragility and financial literacy: almost 45% of low financial educated Italian households do not have sufficient financial resources to cover a lack of income even for short periods (2 months or less). Therefore, households with a low level of financial education are also less resilient.

4. Age: The impact of age is controversial, even if the age effect is widely described as an inverse U-shaped pattern (Kadoya and Khan 2019; Klapper and Panos 2011; Fornero and Monticone 2011; Boisclair et al. 2017). In fact, younger and older respondents usually have a lower share of correct answers about financial issues in contrast to the working age class. Jappelli and Padula (2013) stressed how financial knowledge changes over people's life cycle and that early-life cognition and schooling are strongly correlated with late-life FL.

5. Employment status: This is also an important determinant of financial knowledge, with the lowest level of FL usually recorded among those who are not in the formal paid labor markets (Kadoya and Khan 2019). However, retired people have higher levels of financial knowledge, perhaps due to the increasing privatization of national pension systems, which implies a personal choice among different pension investment plans and solutions for retirement.

6. Family status: Mixed effects are reported in the literature with reference to marital status and family size. According to Jappelli and Padula (2013) and Klapper and Panos (2011), singles have a significant propensity for lower financial literacy levels compared those who are married. In contrast, Bianchi (2018), for France, finds that financial knowledge is negatively correlated with marital status. Moreover, Jappelli and Padula (2013) report a significant negative correlation between financial literacy and family size, whereas Klapper and Panos (2011), for Russia, find a positive but not significant relation.

7. Geography: In addition to personal characteristics, recent literature demonstrates how different cultural backgrounds and embedded social norms can impact on financial knowledge and skills, and thus the importance of analyzing data disaggregated by different geographical contexts (Brown et al. 2018; De Beckker et al. 2020). For Italy,

Fornero and Monticone (2011), exploiting data from the Bank of Italy's Survey on Household Income and Wealth, found evidence of main differences within the same national territory: they identified a significant difference in FL among residents of different regions, with North-Central Italian residents having higher literacy levels than those of the South of Italy. They also reported a positive correlation between individual FL and the household level of digital alphabetization (measured by the presence of at least one member of the household using a computer).

Different studies have also analyzed the relationship among financial knowledge, attitude, and behavior with mixed results. Xiao et al. (2011) found that financial knowledge predicts financial attitude and the latter contributes to the financial behavior of a person. Chaulagain (2017), instead, argues that behaviours are influenced by literacy but not by attitude, and vice versa. Finally, Kadoya and Khan (2019), for Japan, emphasized the importance of psychological variables, in addition to demographic and socio-economic variables, as determinants of FL.

With the exception of psychological variables, for which data are not readily available, we apply all of the dimensions described above, including other controls (see Table A1 in the Appendix A) to define how ML techniques can be used to describe the financially literate population and how accurate they are. Machine learning has been used in the financial services industry for over 40 years, however, it is only in recent years that it has become more pervasive across investment management and trading. Several recent articles have been published that provide evidence of superior performance of non-linear regression techniques for fundamental factor models, such as regression trees (López de Prado 2019; Jain and Jain 2019). Many contributions apply machine learning for predicting portfolio returns. Among others, Moritz and Zimmermann (2016) predict portfolio returns considering tree-based models, Gu et al. (2020) address the prediction of individual stock returns and compare the forecasting performance of different machine learning methods for aggregate portfolio returns to ordinary least squares (OLS) regression, obtaining better accuracy. We apply ML techniques to finance literacy data to show that they can be a useful tool for integrating traditional econometric analysis.

3. Estimation Techniques: Machine Learning

Before empirically applying non-parametric models to study FL and test their predictive performance compared with parametric models, we briefly explain the main characteristics and differences of ML approaches, in particular decision trees (Section 3.1), random forest (Section 3.2), and gradient boosting machine (Section 3.3) techniques.

3.1. Decision Tree

Following a hierarchical structure, a decision tree (DT) partitions the predictor space $\mathbb{R}$ by a sequence of binary splits, giving rise to a tree (Hastie et al. 2016). In this manner, the predictor space is recursively split into simple regions, and the response for a given observation can be predicted using the mean of the training observations in the region to which that observation belongs (James et al. 2017).

Let $(R_j)_{j \in J}$ be the partition of $\mathbb{R}$, where J is the number of distinct and non-overlapping regions. The DT estimator, given a set of variables $x = x_1, \ldots, x_p$, is defined as follows:

$$\hat{f}^{DT}(x) = \sum_{j \in J} \hat{Y}Rj \mathbb{1}_{\{x \in R_j\}} \tag{1}$$

where $\mathbb{1}\{.\}$ is the indicator function. The regions $(R_j)_{j \in J}$ are found by minimizing the residual sum of squares (RSS): $\sum_{j \in J} \sum_{j \in Rj} (y_i - \hat{y}_{Rj})^2$. The estimation of the target variable $\hat{y}_{Rj}$ is identified by the average values of the variable belonging to the same region R_j.

The size of the tree is controlled by a stopping criterion that sets a limit to its growth, to prevent the splitting process continuing until the terminal nodes of the tree become pure (a node is pure when all of the data belong to the same class). The number of terminal nodes is represented by the complexity parameter cp. Small values of cp produce large

trees, increasing the risk of overfitting, whereas large values can underfit the response variable. DTs have the main advantage of being easily interpreted and able to capture any kind of correlation in data. However, they lack robustness in predicting data and small input modification can lead to very different trees. This drawback is due to the use of locally optimal solutions that could be unable to guarantee globally optimal trees. The DT predictive performance can be improved by aggregating many decision trees, thus reducing the variance with respect to a single tree. This technique is behind the ensemble methods, which also include random forest and gradient boosting machine.

3.2. Random Forest

Random forest (RF) is an ML technique consisting of the aggregation of many DTs, obtained by generating bootstrap training samples from the original dataset (Breiman 2001). The idea behind this algorithm is to insert a random perturbation in the learning system to differentiate the trees and combine their predictions through an aggregation technique. The RF technique is based on a bootstrap aggregation (bagging), but its peculiarity is the way it considers the predictors: at each split the algorithm selects a random subset of predictors as split candidates from the final set of predictors, thus preventing the predominance of strong predictors in the splits of each tree (James et al. 2017). Specifically, the random subset consists of two-thirds of the data that are sampled with replacement for training, while the remaining third of the data (called "out-of-bag" observations) are excluded for validation. Therefore, in each bootstrap sample, the data of the training set that are not in the sample can be used as a test set. This technique is called out-of-bag (OOB) and allows for easy estimation of the prediction errors.

The RF is defined by:

$$\hat{f}^{RF}(x) = \frac{1}{B} \sum_{b=1}^{B} \hat{f}^{DT}(x|b) \tag{2}$$

where B is the number of bootstrap samples and $\hat{f}^{DT}(x|b)$ is the decision tree estimator developed on the sample b. The number of trees in the forest must be chosen with the goal of explaining the largest percentage of variance and the lowest mean of squared residuals (MSR). It should be quite large so that each predictor has enough possibilities to be selected, although a relatively smaller number of trees (a few hundred) could be sufficient to achieve high accuracy (Probst and Boulesteix 2018). To understand the relevance of the variables for prediction, we refer to the Mean Decrease Gini (MDG), which is a variable importance measure based on the Gini impurity index, i.e., the average (over the forest) of the decrease in the Gini impurity index for a predictor. Let $i(t)$ be the Gini impurity in node t, we denote $\Delta i(s_t, t)$ as the decrease in impurity of a binary split s_t dividing node t into a left node t_l and a right node t_r. We define $\Delta i(s_t, t)$ as follows:

$$\Delta i\ (s_t,\ t)\ =\ i(t)\ -\ p\ (t_l) \cdot i\ (t_l)\ -\ p\ (t_r) \cdot i\ (t_r) \tag{3}$$

where $p(t_l) = \frac{N_{tl}}{N}$ is the proportion of samples reaching the left node t_l and $p(t_r) = \frac{N_{tr}}{N}$ the proportion of samples reaching the right node t_r, with N the sample size, and N_{tl} and N_{tr} the number of samples reaching the left and right node, respectively. Hence, MDG evaluates the importance of a given variable, x_m, in predicting the response variable and is defined as follows:

$$MDG(x_m) = \frac{1}{N_T} \sum_T \sum_{t \epsilon T: v(s_t) = x_m} p(t) \Delta i(s_t,\ t) \tag{4}$$

where N_T is the number of trees in the forest, $v(s_t)$ is the variable used to split node t and $p(t) = \frac{N_t}{N}$ is the proportion of samples reaching the node t.

3.3. Gradient Boosting Machine

The gradient boosting machine (GBM) is a tree-based algorithm proposed by Friedman (2001) that essentially uses decision trees of a fixed size as weak learners. The prediction is

obtained by a sequential approach and not by parallelizing the tree-building process as in RF. More precisely, in GBM, each decision tree uses the information from the previous decision tree to improve the current fit, i.e., "boosting (improving) the error (gradient)" (Ayyadevara 2018, p. 117). In the following, we briefly describe the algorithm's functioning. Given a current model fit, F_{m-1}, GBM provides a new estimate, F_m, as follows:

$$F_m(x) = F_{m-1}(x) + \lambda \cdot \gamma_m \cdot h_m(x) \tag{5}$$

where λ is the learning rate scaling the contribution of each weak learner and $h_m(x)$ is the weak learners, defined as:

$$h_m = -\sum_{i=1}^{p} \nabla_F L(y_i, F_{m-1}(x_i)) \tag{6}$$

representing the negative gradient of the loss function, $L(y_i, F_{m-1}(x_i))$, evaluated at the current model F_{m-1}. In summary, the new weak predictor h_m tries to minimize the loss function L, given the previous ensemble F_{m-1}.

The accuracy of GBM depends on three fundamental parameters: the number of trees, their depth (i.e., the maximum nodes for each tree), and the learning rate, usually called shrinkage. It is important to choose the right number of trees to obtain a high reduction of the error on the training set. A high number of trees (at least 500) is generally preferable, as a low number might induce overfitting. However, to achieve the minimum predictive error, an appropriate combination of number of trees, tree complexity, and learning rate is necessary.

4. Data and Methods

We use data from the Bank of Italy's 2017 survey that investigates FL and inclusion among Italian adults, with a questionnaire developed by the INFE. The Italian sample consists of about 2500 persons interviewed using two different methods: 40 per cent were interviewed face-to-face whereas the remainder used a tablet to record their responses. The survey questionnaire, designed according to the INFE framework, measures financial knowledge, behavior, and attitudes. We focus our analysis on the knowledge component that assesses the understanding of basic concepts that are a pre-requisite for making sound and conscious financial decisions (Lusardi and Mitchell 2011): understanding simple and compound interest, inflation, and the benefits of portfolio diversification. There were 7 questions about financial knowledge; we calculated from this dimension a composite FL index that ranges from 0 to 7. Since the average score for Italy is 3.5 out of a maximum of 7 points, lower than the G20 average of 4.3 (see Figure 1), we split the responders into two groups: those with higher financial literacy than the average of Italians—namely those who correctly answered at least 4 questions—and those who are less financial educated. To define the main determinants of higher financial education of Italian adults, we consider a set of personal observable characteristics commonly used in the literature and described in Section 1, such as gender, age, education, household composition, and employment status. We also controlled for migrant status, because migrants are usually more exposed to financial exclusion, and for geographic macro areas of Italy, given the evident macroeconomic gaps among different areas of Italy, specifically the Northern and Southern/"Mezzogiorno" areas (the so-called Italian socio-economic dualism). We also used two different variables to assess the financial fragility of responders: the household economic stress, or whether, in the 12 months before the interview, the household income was insufficient to cover monthly expenses, and risk capacity, or the ability to sustain unexpected expenses without asking for formal or informal loans. We enriched the analysis by also considering financial variables such as financial behavior and attitudes, the propensity for pension planning (to have private pension plans, any pension product, or savings for retirement), and respondents' high self-assessment of financial knowledge (on a scale ranging from 1 to 5). We used the International Network for Financial Education (INFE) framework to measure the three areas of financial literacy:

knowledge, behavior, and attitudes (OECD INFE 2011). Therefore, the behavior index was based on questions assessing whether people manage household financial resources by formulating a budget, are able to pay their debts and utilities with no concerns, and acquire information before making investments. Following the OECD/INFE framework, the Bank of Italy measures financial behavior by incorporating a variety of questions to identify three potentially prudent financial behaviours, namely:

- Saving, financial assets, and long-term planning: a set of questions is used to understand if individuals purchased financial assets in the two years before the survey, therefore, if they are actively saving or borrowing, and whether they set themselves long-term financial goals.
- Making considered purchases: there are questions that explore if individuals make informed decisions before making a purchase of financial products and services.
- Keeping track of cash flow: some questions are asked to understand if individuals keep a watch of financial affairs, and if they pay their bills on time. The ability to manage financial resources properly was measured, as for the OECD INFE (2011), on a scale of 0 to 9. Financial attitude instead evaluated personal traits such as preferences, beliefs, and non-cognitive skills, which are likely to affect personal well-being, on a scale from 0 to 5; the main driver of the index is a positive saving orientation, mainly for the long term. Because Di Salvatore et al. (2018) found that "the response behavior of Italian respondents appears to be influenced by the survey mode" (p. 8), we also included in our estimates a dummy variable to identify if the responder had a face-to-face interview or used a tablet to record their responses (in Appendix A, Table A1 provides a full description of the variables considered). It is clear from the first descriptive statistics in Table 1 that the level of FL is not uniform throughout the population in Italy. Although small, there are gender gaps in financial knowledge, with men slightly more financially literate than women. In addition, we find the above-cited reverse U-shaped curve for age, because financial knowledge increases with age but decreases for older adults, with a peak for the working age group 40–49 years old. FL is higher for those employed in paid work but lower among those in unpaid domestic work and those unemployed or seeking their first employment. FL is higher in the North Western regions of Italy. However, on average a low share of Italians (8%) rates their financial knowledge as being high. Finally, among financial literates, the average levels of good financial behaviours and attitudes are still low (4.5 on a scale of 0–9 and 2.1 on a scale of 0–5, respectively), but their ability to cope with unexpected expenses without asking for formal or informal loans or to cover monthly expenses is quite high, a peculiar characteristic of Italians, who achieve, on average, a high level of savings.

Considering the data described above, we formulated our model to estimate the main determinants of FL in Italy as follows:

Higher financial literacy~gender + education + risk.capacity + HE.stress + age + employment.status + household.composition + geographic.area + native + financial.behavior + financial.attitude + FL.self.assessment + pension.savings + pension products in the last 2 years (PP.in.the.last.year)+ pension.fund +interview.type

We split the data set into a training set and a test set, according to the common splitting rule of 80–20%. Therefore, the training and the test sets consisted of 1901 and 475 observations, respectively[1]. We are conscious that the size of the sample considered is relatively small but, due to cross-validation, machine learning can be used to validate the predictive accuracy without problems for small datasets.

---

[1] Our results are obtained from special R packages suitably implemented for tree-based ML algorithms: the rpart package developed by Therneau et al. (2017) for DT, the randomForest package developed by Liaw (2018) for RF, and the gbm package developed by Ridgeway (2007) for GBM. A brief description of the parameters and their optimal tuning for the different ML algorithms is provided in Table A2 in the Appendix A.

Table 1. Main characteristics of adult Italians with higher financial knowledge than average.

Personal Characteristics	%	Personal Characteristics	%
1. Gender		*5. Employment status*	
Men	51%	Employee	39%
2. Education		Self-employed	12%
University degree/some university studies	26.7%	Unemployed	9%
Secondary school (completed)	42.4%	Unpaid domestic work	9%
Some secondary school	25.3%	Retired	23%
Primary school (completed)	5.2%	*6. Family status*	
Some primary school	0.4%	Single	11%
3. Financial fragility		*7. Geography*	
Household economic stress (HE.stress)	26%	Centre	19%
Risk capacity	57%	South	21%
4. Age		North-West	27%
<30	13%	North-East	21%
30–39	15%	Islands	11%
40–49	22%	*Other*	
50–59	20%	Native	98%
60–69	17%	Financial behavior (mean)	4.5
70–79	10%	Financial attitude (mean)	2,1
>80	4%	FL self-assessment	8%

5. Results

The results obtained for the tree-based ML algorithm (see Figure A1 in Appendix A) depict the best decision tree for FL data used. We see that the best tree has 10 terminal nodes (nine splits) and the root node splits on risk capacity = 0 in yes and no. Each node shows the predicted class (1 or 0) and the percentage of observations in the node. To assess the performance of the model, we refer to the OOB technique.

5.1. Predictive Quality: Models' Validation, Accuracy and Performance Evaluation

The tree-based algorithms are usually validated using the OOB score, that is, the average prediction error calculated on each training sample x_j, using only the trees that did not have x_j in their bootstrap sample. Sub-sampling allows one to define an OOB estimate of the prediction performance improvement by evaluating predictions on those observations that were not used in the building of the next base learner.

The variation of the OOB error with respect to the number of trees used in the RF algorithm shows that the OOB error rate stabilizes around 0.4 when 100 trees are used for building the forest, suggesting a good capacity of the RF algorithm to predict the FL (Figure 2, panel a). Panel b in Figure 2 shows the GBM performance evolution, based on the Bernoulli deviance, when the algorithm combines a progressively larger number of weak learners. Smaller deviance values indicate better performance. The black line represents the training Bernoulli deviance, whereas the green line shows the testing Bernoulli deviance, which is the result of the cross validation. The blue dashed line shows the optimal number of iterations. The plot highlights that, beyond a certain point (in our case, 58 trees), the model generalization power starts decreasing, explaining only the training data. This point represents the optimal number of iteration.
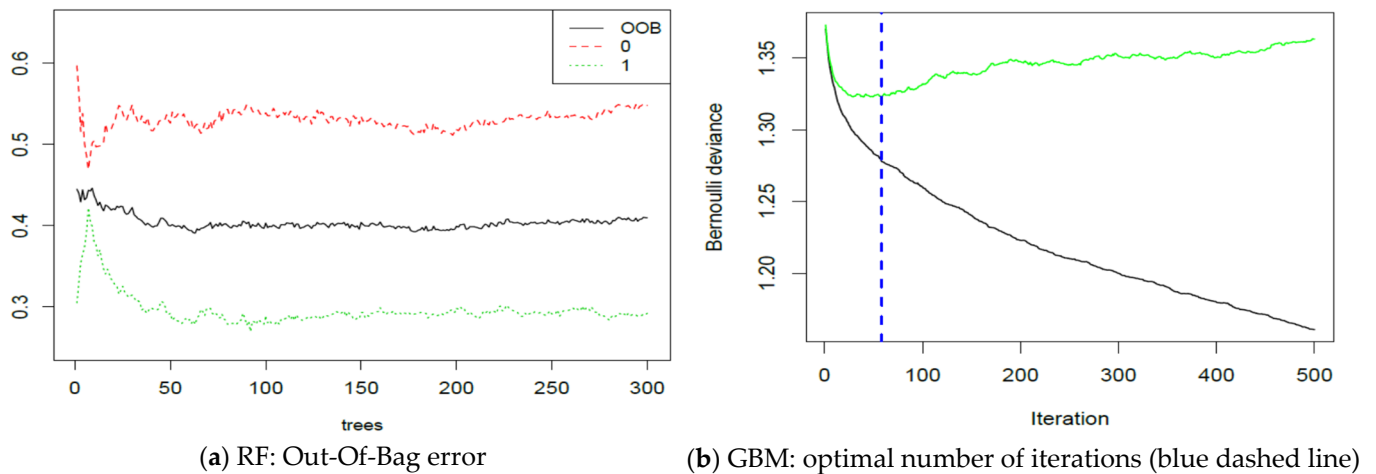
(**a**) RF: Out-Of-Bag error

(**b**) GBM: optimal number of iterations (blue dashed line)

Figure 2. Random forest (RF) and gradient boosting machine (GBM) performance.

The model's accuracy is measured according to a set of indicators that can be easily determined by the confusion matrix, reporting the number of observations correctly or incorrectly classified (Table 2). The diagonal elements of the confusion matrix indicate correct predictions, whereas the other elements indicate incorrect predictions.

Table 2. Confusion matrix.

	Predicted Negative	**Predicted Positive**
Actual negative	True negatives (TN)	False positives (FP)
Actual positive	False negatives (FN)	True positives (TP)

The metrics used in this paper are listed below, expressed according to the elements of the confusion matrix:

- Accuracy (acc): $\frac{TP+TN}{(TP+TN+FP+FN)}$
- True Positive Rate (TPR), also called sensitivity: $\frac{TP}{(TP+FN)}$
- False Positive Rate (FPR): $\frac{FP}{(FP+TN)}$
- True Negative Rate (TNR), also called specificity (or 1-FPR): $\frac{TN}{(FP+TN)}$
- Precision: $\frac{TP}{(TP+FP)}$

Note that the accuracy, i.e., the proportion of correct predictions, can be written as: acc $= \frac{1}{N}\sum_{i=1}^{N}\mathbb{I}(y_i = \hat{y}_i)$, where ($\cdot$) is the indicator function. The overall performance of the ML algorithms, summarized over all possible thresholds, can be represented by the Receiver Operating Characteristics (ROC) curve and in particular by the Area Under (this) Curve (AUC), that is, the integral area of plotting the sensitivity (TPR) on the *y*-axis vs. 1-specificity (FPR) on the *x*-axis. Specifically, ROC shows how TPR and FPR vary with different threshold values and can be used to compare different classification algorithms.

The values of the accuracy measures applied to the FL data in Italy are reported in Table 3 for the ML algorithms and the logistic regression model. The ML algorithm's performance is compared to the results of a logistic regression model (LR), that is, a generalized linear model (GLM) with a logit link function g(.) = logit and a binomial distribution for the response (binary) variable Y. Letting μ denote the expectation of the response variable Y, the structure of a logistic model is:

$$\text{logit}\,(\mu)\ =\ \beta_i + \sum_{i}^{p}\beta_i \cdot X_i \tag{7}$$

where $\beta_1, \ldots, \beta_{p-1}$ are the regression parameters that need to be estimated and β_0 is the intercept. The covariates enter a logistic regression model through the linear predictor logit (μ), leading to interpretable effects of the explanatory variables on the response.

Table 3. Accuracy measures.

Model	Sensitivity	Specificity	Precision	Accuracy	AUC
DT	0.6250	0.6462	0.7172	0.6337	0.6313
RF	0.6478	0.7184	0.7992	0.6737	0.6702
GBM	0.7225	0.6268	0.5656	0.6653	0.7231
Logit	0.6092	0.6283	0.7090	0.6168	0.6908

The RF algorithm accurately identifies the individuals who are financially literate in the test set with an accuracy level equal to 67.37%, which is the highest among the set of models taken into account.

To measure the precision's accuracy, we show the ROC curve for all of the models considered in Figure 3. According to the ROC curve, the GBM model provides the highest AUC, hence resulting in the best model. We can conclude that machine learning can improve the accuracy of some standard parametric models in the estimation of the main determinants of financial literacy.

(**a**) Decision tree

(**b**) Random Forest

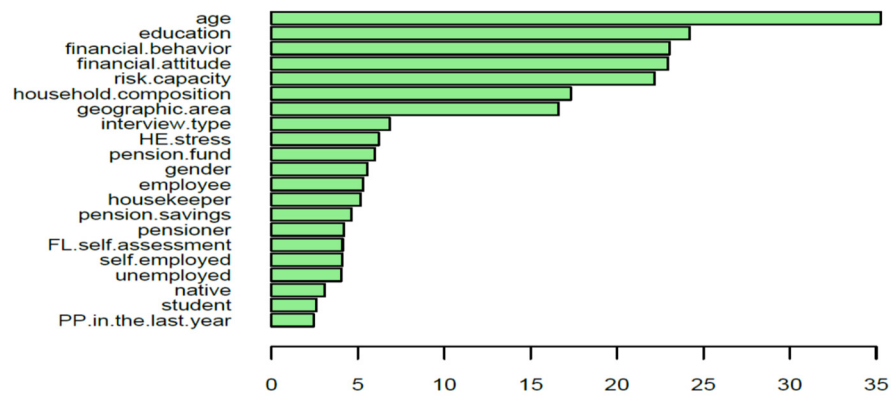(**c**) Gradient boosting machine
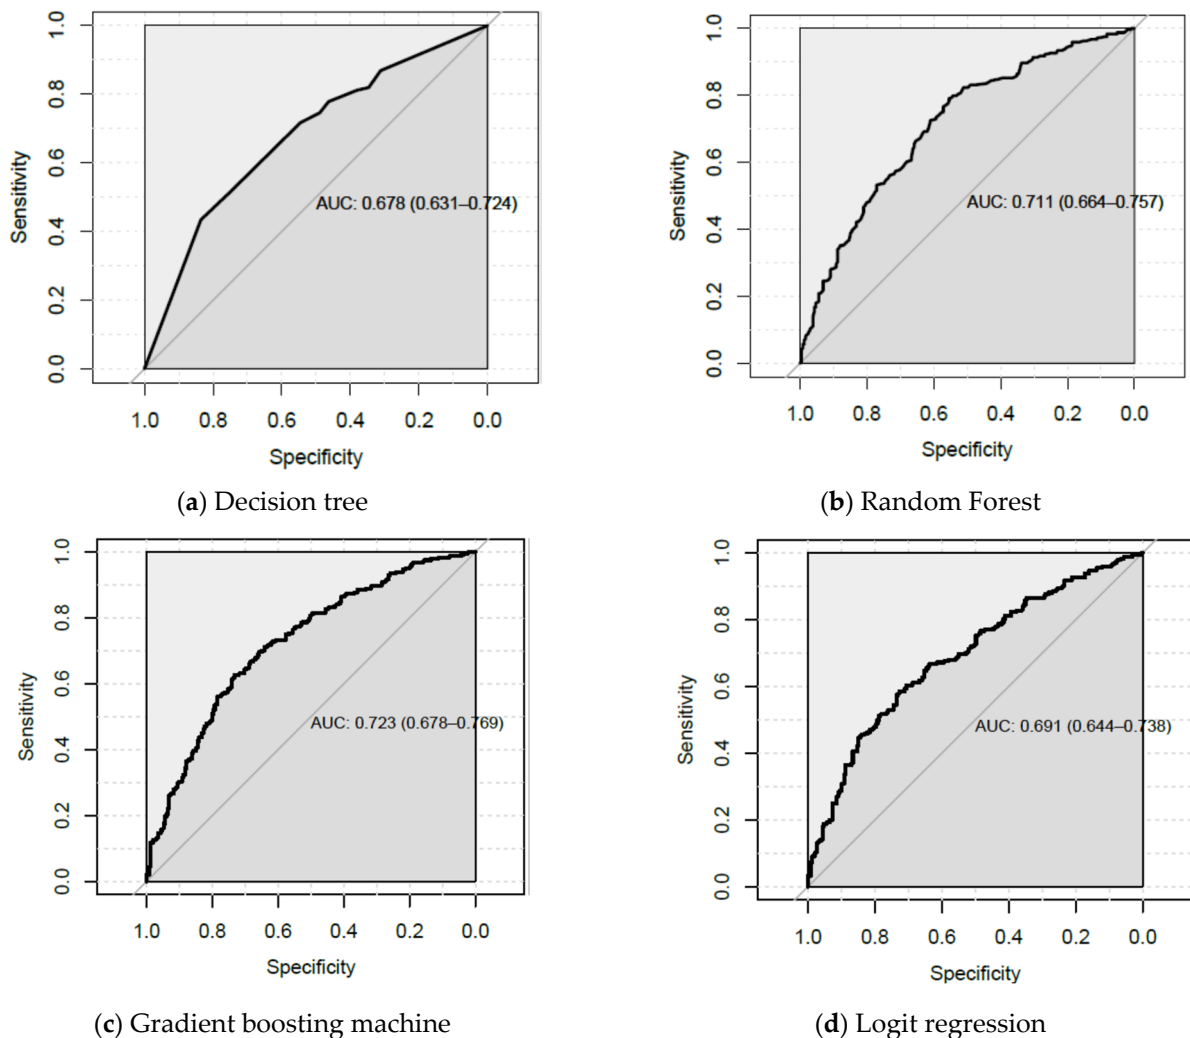
(**d**) Logit regression

Figure 3. Receiver Operating Characteristics (ROC) curves for machine learning (ML) estimators (decision tree (DT), RF, GBM) and logistic regression model (LR).

5.2. Variable Importance and Partial Dependence

ML algorithms are usually viewed as a black box because gaining insight into a RF prediction rule is hard due to the large number of trees. One of the most common approaches to extract interpretable information on the contribution of different variables from the random forest consists of the computation of the so-called variable importance measures. Variable importance is determined according to the relative influence of each predictor, by measuring the number of times a predictor is selected for splitting during the tree-building process, weighted by the squared error improvement in the model each split, and averaged over all trees. We plot in Figure 4 the relative importance of the predictors for different ML techniques. The most important variables are at the top of each plot, and the less important are at the bottom. From the results we observe the predominance of age in determining the FL, especially for RF and GBM algorithms. Education and financial behavior follow. The gradient boosting machine (GBM) model, which we identified as the best model in the previous session, highlights the importance of financial attitude and financial behavior to explain different levels of financial knowledge. It is also interesting to note that gender is not among the most relevant dimensions to explain Italian adults' financial literacy and the geographical distribution in the national territory is more relevant.
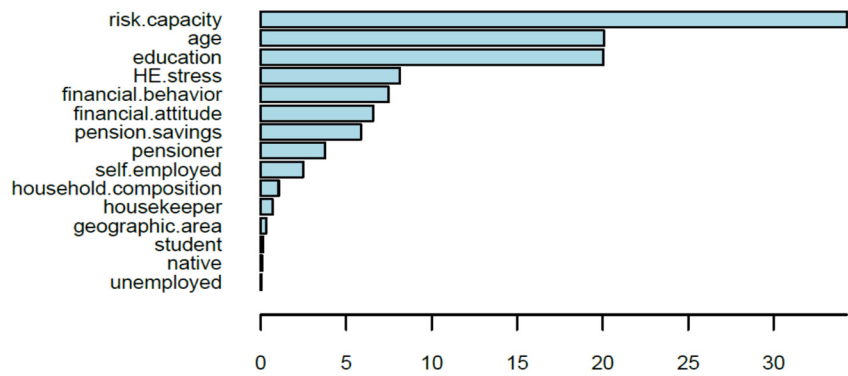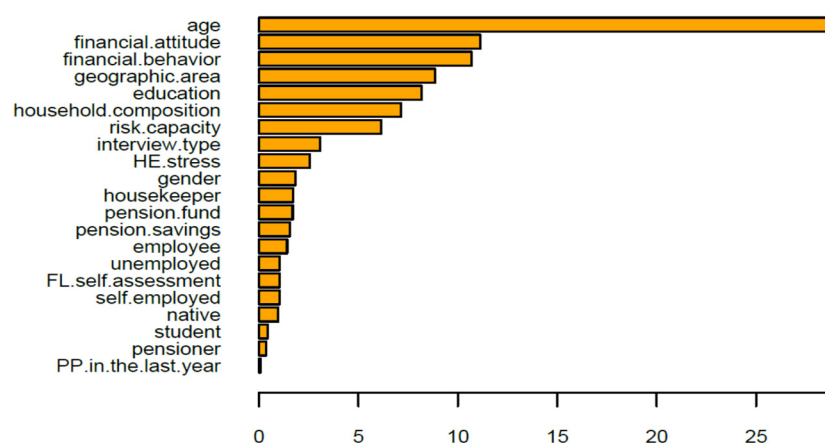

(**a**) Decision tree (DT)



(**b**) Random forest (RF)

Figure 4. *Cont.*

(**c**) Gradient boosting machine (GBM)

Figure 4. Variable importance with ML models.

We also define the partial dependence plots to show the marginal effect of the selected predictor on the target variable averaged over the joint values of the other predictors provided by a tree structure (see Friedman (2001) for further details). The function explaining the partial dependence is: $\hat{f}_s(x_s) = \frac{1}{n} \sum_{i=1}^{n} f(x_s, x_{i,C})$, where x_s is the variable of interest and $x_{i,C}$ is the complementary variable in the dataset. Figure 5 illustrates three one-way partial dependence plots for our dataset, with the GBM regressor (the best model for the precision's accuracy, as shown in the next section). The plots show that the most important predictors are age, financial attitude, and financial behavior. The results are in line with the main results of the literature in the field. We confirm the correlation of age with a higher level of FL among working age adults and a lower level for younger and older adults, as described in Section 1. Moreover, there is a clear positive correlation between financial knowledge and both financial attitude and financial behavior. In the latter case, because one of the elements contributing to the good behavior score we used is the purchase of financial assets in the two years before the survey, we can speculate that experience has a positive effect in the acquisition of financial knowledge. In that sense, many studies suggest that experience plays an important role in a person's motivation to become financially literate. For example, Mandell (2008) found that financial education programs that include experiential components have a higher impact; for example, participation in a stock market game results in a 6–8% improvement in FL among respondents. Frijns et al. (2014) suggest that "people with more financial experience acquire more financial knowledge either through self-education or by becoming more receptive to financial education programmes" (p. 125). Our results do not identify the causal relationship between financial knowledge and behaviours; however, we can speculate that causality runs in both directions, either when more financially literate people engage in more financial activity and therefore become more experienced, or when people may learn from their financial experiences and therefore become more literate. The main implication of these results is that policy makers should consider ways to increase the financial experience of people, through experiences in real-world situations, as a way of improving FL.

(**a**) Age (normalised)

(**b**) Financial attitude

(**c**) Financial behavior
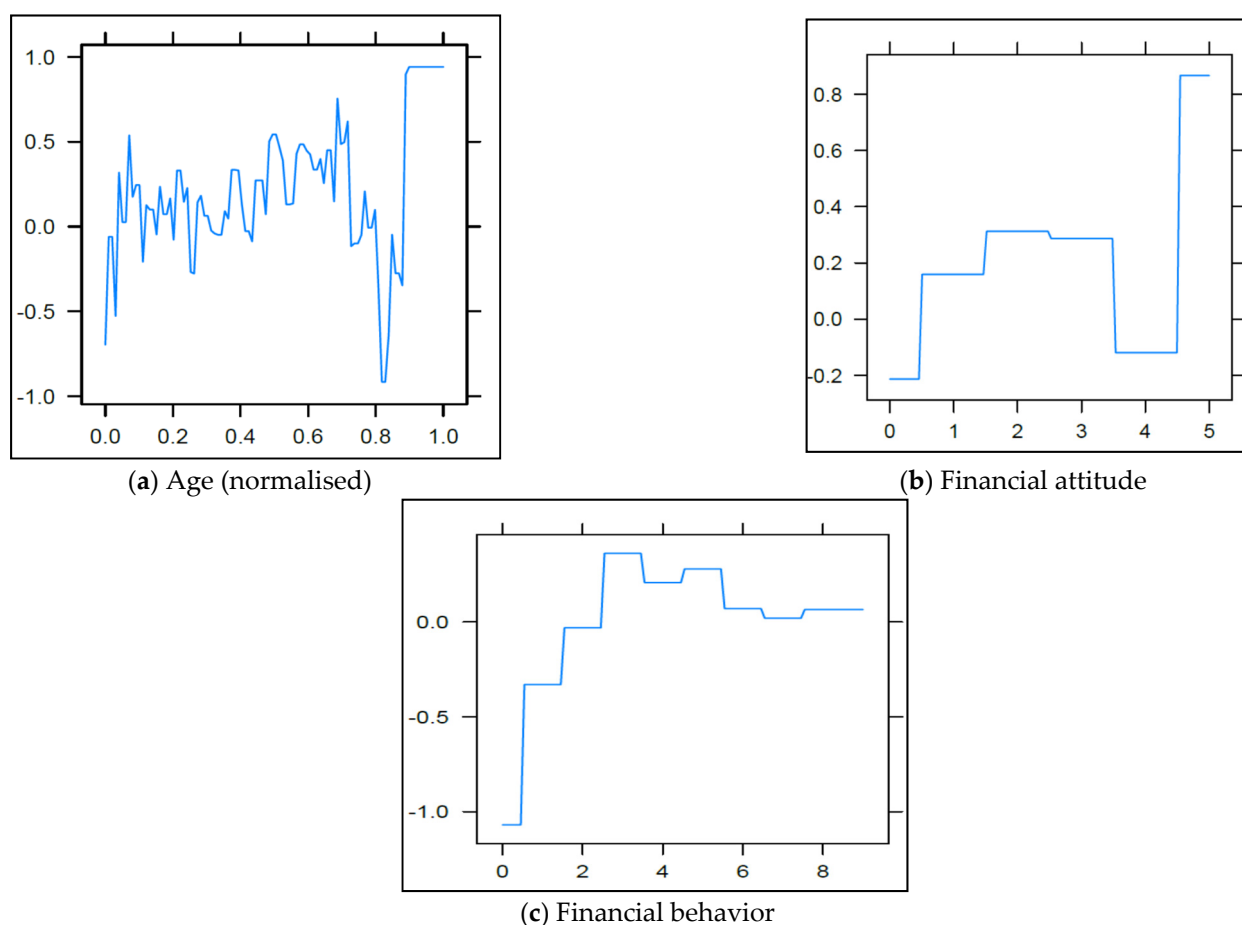
Figure 5. Single variable partial dependence plots for the three most important predictors.

6. Conclusions

One of the main recent developments in financial research is the availability of new administrative, unstructured, micro-level data that are difficult to analyze with traditional econometric models. In this scenario, machine learning techniques can offer the capabilities and functional flexibility needed to identify complex patterns in a high-dimensional spaces and datasets. There are clearly advantages and disadvantages for both parametric and ML models. The latter are nonparametric and do not postulate a functional form linking the target variable to the explanatory variables, so their main strength is their high flexibility in learning from data and their high predictive performance. Their main drawbacks are the risk of overfitting and the interpretability of the results generated by the algorithms. In contrast, parametric models, such as the generalized linear model (GLM), have the advantages of being parsimonious and easy to interpret and estimate; their drawbacks are that they have limited complexity and generally poor predictive power.

We do not wish to define which of the two methods, the parametric or ML model, is the best approach. However, this study examined how they can be used and integrated with each other to gain a better understanding of the phenomenon of financial literacy. We demonstrated that analytical steps of the econometric processes, such as the logit analysis that we applied to the Italian data on financial literacy, has a homologous step in ML analyses. By clearly stating this correspondence, we hope that the adoption of ML techniques in the context of financial literacy will be facilitated.

In detail, we tested the improvement in the accuracy in explaining the determinants of FL using not only the decision tree, but also two more powerful ML algorithms: random forest and gradient boosting. Our results demonstrate that the gradient boosting machine methodology outperforms conventional methods. Moreover, ML analyses produce reliable information consistent with the literature, because FL is highly correlated with demo-

graphic variables such as educational attainment, age, and household financial fragility. The results of ML models also highlight, in contrast to the traditional parametric model, the importance of financial behaviours in defining the level of financial knowledge. Because we used the INFE-OECD's definition of financial behavior, which accounts for the purchase of financial assets in the two years before the survey, we can speculate that experience has a positive effect in the acquisition of financial knowledge. Therefore, these results have policy implications because they suggest that effective strategies to tackle financial illiteracy should involve experiences in real-world situations. In that sense, banks and financial institutions could play an essential role in the field of education and training in FL (Trunk et al. 2017).

We are conscious that we tested ML models based on a limited case study, using the few available microdata distributed by the Bank of Italy on the levels of adults' financial literacy in Italy. However, we hope that this could be a first step to encourage the adoption of ML techniques in applied economics and among finance researchers and policy makers in the context of financial literacy. We can conclude that machine learning techniques can be valuable as a complement to standard models, which can be further extended in several directions. The ML approaches can be useful for analyzing complex data structures and a large amount of information simultaneously. Thus, they can provide a more nuanced picture of the phenomenon to give policy makers, national bodies, and financial institutions a clearer framework for effectively targeting the problem of financial illiteracy in accordance with the OECD's Recommendation on Financial Literacy. In the era of Big Data, where massive amounts of very high-dimensional or unstructured data are continuously produced and stored, ML techniques provide new opportunities in data analysis, both for exploring the hidden structures and correlation of each variable considered, which traditionally has not been feasible, and for extracting important common features across many subpopulations, even when there are large individual variations. However, substantial efforts are also required for advancement in data collection and the availability of individual information on financial behaviours, attitude, and knowledge at the national and international levels.

Based on our analysis, further applications of machine learning methods to high-dimensional data (big data) on financial literacy, once available, would help understanding the heterogeneity and commonality in levels of financial literacy across different subpopulations. This is particularly relevant in the wake of the COVID-19 pandemic, which is exacerbating social and economic inequalities globally.

Appendix A

Table A1. List of variables.

Variables Used	Description	Minimum Value	Maximum Value	Mean (%)	Standard Deviation
Dependent Variable: **Higher financial literacy**	Responder who correctly answered at least 4 questions out of 7 questions about financial knowledge (three topics: understanding simple and compound interest, inflation and the benefits of portfolio diversification)	0	1	53.49	0.5
1. Gender		0 (woman)	1 (man)	49	0.5
2. Education	7-levels	1 university degree or more	7 no complited education	3.03	1.26
3. Financial fragility					
(a) Household economic stress	In the past 12 months household income was not sufficient to cover monthly expenses?	0 (no)	1 (yes)	28	0.45
(b) Risk capacity	Ability to sustain unexpected expenses without asking for formal or informal loans	0 (no)	1 (yes)	48	0.5
4. Age (normalized)		18	80 more	50.34	17.09
5. Employment status					
(a) selfemployed	dummy variable	0 (no)	1 (yes)	11	0.31
(b) employee	dummy variable	0 (no)	1 (yes)	36	0.48
(c) housekeeper	dummy variable	0 (no)	1 (yes)	11	0.31
(d) unemployed	dummy variable	0 (no)	1 (yes)	10	0.29
(e) pensioner	dummy variable	0 (no)	1 (yes)	24	0.43
6. Family status					
Household composition	6 classes	1 (1 member)	6 (6 or more)	2.95	1.22
7. Geography					
area 1	North-West	0 (no)	1 (yes)	27	0.44
area 2	North-East	0 (no)	1 (yes)	20	0.4
area 3	Centre	0 (no)	1 (yes)	19	0.39
area 4	South	0 (no)	1 (yes)	22	0.42
area 5	Islands	0 (no)	1 (yes)	11	0.32
8. Native	dummy variableBorn:	0 Not in Italy	1 In Italy	97	0.16
9. Financial variables					
FL self-assessment	Respondent's high level of self-assigned financial knowledge (well above average and above the average).	0 (no)	1 (yes)	7	0.25
Pension savings	Responder is confident that she/he has done a good job of making financial plans for her/his retirement.	0 (no)	1 (yes)	11	0.32
Pension products in the last 2 years	In the last two years the responder has bought a pension or retirement product.	0 (no)	1 (yes)	25	0.16
Pension fund	Responder is funding her/his retirement with a private pension plan.	0 (no)	1 (yes)	10	0.3
Financial behavior	9 questions assessing whether people are able to formulate a budget, to pay their debts and utilities with no concerns, and acquire information before making investments.	0	9	4.48	1.71
Financial attitude	5 questions about personal attitude towards precautionary saving and long run savings.	0	5	1.92	1.36
9. Interview type	Survey mode	0 (tablet)	1 (ftf)	40	0.49

Table A2. Optimal tuning parameters for DT, RF and GBM algorithms.

	Parameter Description	Optimal Tuning Parameter
DT	Complexity parameter	cp = 0.006
RF	Number of trees	ntree = 300
	Minimum number of observations in a terminal node	nodesize = 11
	Number of input variables in each node	mtry = 2
GBM	Number of trees	ntree = 500
	Maximum nodes for each tree	interaction.depth = 3
	Learning rate	shrinkage = 0.1
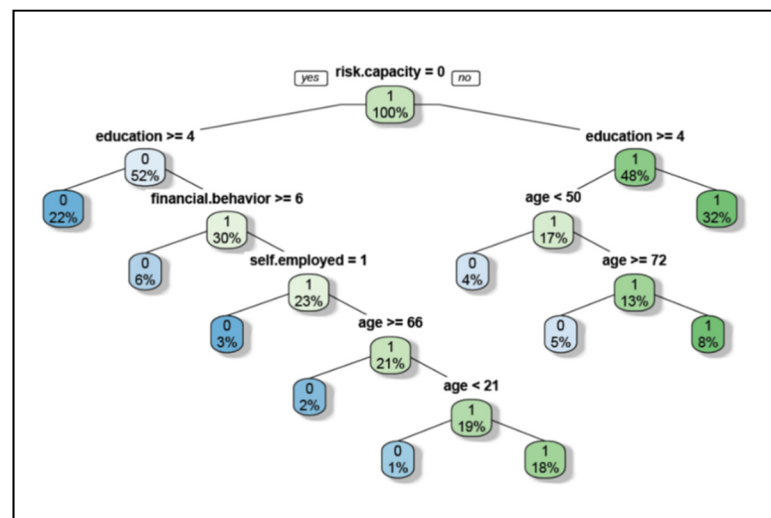
Figure A1. Decision tree for FL.

References

Al-Bahrani, Abdullah, Buser Whitney, and Darshak Patel. 2020. Early Causes of Financial Disquiet and the Gender Gap in Financial Literacy: Evidence from College Students in the Southeastern United States. *Journal of Family and Economic Issues* 41: 558–71. [CrossRef]

Ameriks, John, Caplin Andrew, and John Leahy. 2003. Wealth accumulation and the propensity to plan. *Quarterly Journal of Economics* 118: 1007–47. [CrossRef]

Atkinson, Adele, and Flore-Anne Messy. 2012. Measuring Financial Literacy: Results of the OECD/International Network on Financial Education (INFE) Pilot Study. In *OECD Working Papers on Finance, Insurance and Private Pensions*. No. 15. Paris: OECD Publishing.

Ayyadevara, V. Kishore. 2018. *Pro Machine Learning Algorithms*. Berkeley: Apress.

Banks, James, and Zoe Oldfield. 2007. Understanding pensions: Cognitive function, numerical ability and retirement saving. *Fiscal Studies* 28: 143–70. [CrossRef]

Bazarbash, Majid. 2019. FinTech in Financial Inclusion: Machine Learning Applications in Assessing Credit Risk. In *IMF Working Paper 19/109*. Washington, DC: IMF.

Behrman, Jere R. Mitchell Olivia S., Soo Cindy K, and David Bravo. 2012. How financial literacy affects household wealth accumulation. *American Economic Review* 102: 300–4. [CrossRef]

Bianchi, Milo. 2018. Financial literacy and portfolio dynamics. *The Journal of Finance* 73: 831–59. [CrossRef]

Boisclair, David, Lusardi AnnaMaria, and Pierre Carl Michaud. 2017. Financial literacy and retirement planning in Canada. *Journal of Pension Economics & Finance* 16: 277–96.

Bracke, Philippe, Datta Anupam, Jung Carsten, and Shayak Sen. 2019. Machine learning explainability in finance: An application to default risk analysis. In *Bank of England Staff Working Paper*; No. 816; London: Bank of England.

Breiman, Leo, Friedman Jerome, R. Olshen, and Charles J. Stone. 1984. *Classification and Regression Trees*. Boca Raton: Chapman & Hall/CRC.

Breiman, Leo. 2001. Random forests. *Machine Learning* 45: 5–32. [CrossRef]

Brown, Martin, Guin Benjamin, and Stefan Morkoetter. 2016. Deposit withdrawals from distressed commercial banks: The importance of switching costs. In *University of St. Gallen, School of Finance Research Paper*. St. Gallen: University of St. Gallen, pp. 2013–19.

Brown, Martin, Henchoz Caroline, and Thomas Spycher. 2018. Culture and financial literacy: Evidence from a within-country language border. *Journal of Economic Behavior and Organization* 150: 62–85. [CrossRef]

Bucher-Koenen, Tabea, Lusardi Annamaria, Alessie Rob, and Maarten Van Rooij. 2017. How financially literate are women? An overview and new insights. *Journal of Consumer Affairs* 51: 255–83. [CrossRef]

Chaulagain, Ramesh Prasad. 2017. Relationship between Financial Literacy and Behavior of Small Borrowers. *NRB Economic Review* 29: 33–53.

Chen, Jia, Jiang Jiajun, and Yu-jane Liu. 2018. Financial literacy and gender difference in loan performance. *Journal of Empirical Finance* 48: 307–20. [CrossRef]

Christelis, Dimitrios, Jappelli Tullio, and Mario Padula. 2010. Cognitive abilities and portfolio choice. *European Economic Review* 54: 18–38. [CrossRef]

Clark, Robert L., Lusardi Annamaria, and Olivia S. Mitchell. 2015. Financial knowledge and 401 (k) investment performance: A case study. *Journal of Pension Economics and Finance* 16: 1–24. [CrossRef]

Cole, Shawn, Sampson Thomas, and Bilal Zia. 2011. Prices or Knowledge? What Drives Demand for Financial Services in Emerging Markets? *The Journal of Finance* 66: 1933–67. [CrossRef]

De Beckker, Kenneth, De Witte Kristof, and Geert Van Campenhout. 2020. The role of national culture in financial literacy: Cross-country evidence. *Journal of Consumer Affairs* 54: 912–30. [CrossRef]

Demertzis, Maria, Domínguez-Jiménez Marta, and Anna Maria Lusardi. 2020. The financial fragility of European households in the time of COVID-19. In *Policy Contribution 2020/15*. Brussels: Bruegel.

Deuflhard, Florian, Georgarakos Dimitris, and Roman Inderst. 2019. Financial Literacy and Savings Account Returns. *Journal of the European Economic Association* 17: 131–64. [CrossRef]

Di Salvatore, Antonietta, Franceschi Francesco, Neri Andrea, and Francesca Zanichelli. 2018. Measuring the financial literacy of the adult population: The experience of the Bank of Italy. *IFC Bulletins* 47: 1–35. [CrossRef]

Dixon, Matthew F., Halperin Igor, and Bilokon Paul. 2020. *Machine Learning in Finance*. Berlin: Springer International.

Feng, Xiangnan, Bin Lu, Xinyuan Song, and Shuang Mad. 2019. Financial Literacy and Household Finances: A Bayesian Two-Part Latent Variable Modeling Approach. *Journal of Empirical Finance* 51: 119–37. [CrossRef]

Fornero, Elsa, and Chiara Monticone. 2011. Financial literacy and pension plan participation in Italy. *Journal of Pension Economics and Finance* 10: 547–64. [CrossRef]

Fornero, Elsa, and Anna Lo Prete. 2019. Voting in the aftermath of a pension reform: The role of financial literacy. *Journal of Pension Economics and Finance* 18: 1–30. [CrossRef]

Friedman, Jerome H. 2001. Greedy function approximation: A Gradient Boosting Machine. *Annals of Statistics* 29: 1189–232. [CrossRef]

Frijns, Bart, Gilbert Aaron, and Tourani-Rad Alireza. 2014. Learning by doing: The role of financial experience in financial literacy. *Journal of Public Policy* 34: 123–54. [CrossRef]

Gathergood, John, and Jorg Weber. 2017. Financial literacy, present bias and alternative mortgage products. *Journal of Banking and Finance* 78: 58–83. [CrossRef]

Gerardi, Kristopher, Goette Lorenz, and Stephan Meier. 2010. Financial Literacy and Subprime Mortgage Delinquency: Evidence From a Survey Matched to Administrative Data. In *Federal Reserve of Atlanta WP 2010-10*. Darby: DIANE Publishing.

Goda, Shah Gopi, Levy Matthew R, Manchester Colleen Flaherty, Sojourner Aaron, and Joshua Tasoff. 2020. Who is a passive saver under opt-in and auto-enrollment? *Journal of Economic Behavior and Organization* 173: 301–21. [CrossRef]

Goyal, Kirty, and Satish Kumar. 2020. Financial literacy: A systematic review and bibliometric analysis. *International Journal of Consumer Studies* 45: 80–105. [CrossRef]

Grohmann, Antonio, Klühs Theres, and Lukas Menkhoff. 2018. Does financial literacy improve financial inclusion? Cross country evidence. *World Development* 111: 84–96. [CrossRef]

Gu, Shihao, Bryan Kelly, and Dacheng Xiu. 2020. Empirical Asset Pricing via Machine Learning. *The Review of Financial Studies* 33: 2223–73. [CrossRef]

Guiso, Luigi, and Eliana Viviano. 2015. How much can financial literacy help? *Review of Finance* 19: 1347–82. [CrossRef]

Guiso, Luigi, and Tullio Jappelli. 2008. Financial literacy and portfolio diversification. In *EUI Working Paper (ECO 2008/31)*. Florence: European University Institute.

Hasler, Andrea, Lusardi Annamaria, and Noemi Oggero. 2018. Financial fragility in the US: Evidence and implications. In *GFLEC working Paper n. 2018-1*. Washington, DC: Global Financial Literacy Excellence Center, The George Washington University School of Business.

Hastie, Trevor, Tibshirani Robert, and Jerome Friedman. 2016. *The Elements of Statistical Learning. Data Mining, Inference, and Prediction*. New York: Springer, ISBN 10: 0387848576.

Hastings, Justine S., and Lydia Tejeda-Ashton. 2008. *Financial Literacy, Information, and Demand Elasticity: Survey and Experimental Evidence from Mexico*; NBER Working Papers 14538; Cambridge: National Bureau of Economic Research.

Hsiao, Yu-Jen, and Wei-Che Tsai. 2018. Financial literacy and participation in the derivatives markets. *Journal of Banking & Finance* 88: 15–29.

Jain, Prayut, and Shashi Jain. 2019. Can machine learning-based portfolios outperform traditional risk-based portfolios? The need to account for covariance misspecification. *Risks* 7: 74. [CrossRef]

James, Gareth, Witten Daniela, Hastie Trevor, and Robert Tibshirani. 2017. *An Introduction to Statistical Learning: With Applications in R*. Springer Texts in Statistics. Berlin: Springer, ISBN 10: 1461471370.

Jappelli, Tullio, and Mario Padula. 2013. Investment in financial literacy and saving decisions. *Journal of Banking & Finance* 37: 2779–92.

Kadoya, Yoshihiko, and Mostafa Khan. 2019. What determines financial literacy in Japan? *Journal of Pension Economics and Finance*, 1–19.

Kim, Namhoon, and Travis P. Mountain. 2019. Financial Knowledge and "Don't Know" Response. *Journal of Consumer Affairs* 53: 1948–69. [CrossRef]

Klapper, Leora, and Georgios A. Panos. 2011. Financial literacy and retirement planning: The Russian case. *Journal of Pension Economics & Finance* 10: 599–618.

Klapper, Leora, Lusardi Annamaria, and Peter Van Oudheusden. 2015. Financial Literacy around the World: Insights from The Standard & Poor's Ratings Services. In *Global Financial Literacy Survey*. Washington, DC: Global Financial Literacy Excellence Center, the George Washington University.

Liaw, Andy. 2018. Package Randomforest. Available online: https://cran.r-project.org/web/packages/randomForest/randomForest.pdf (accessed on 15 February 2021).

Lo Prete, Anna. 2013. Economic literacy, inequality, and financial development. *Economics Letters* 118: 74–76. [CrossRef]

Lo Prete, Anna. 2018. Inequality and the finance you know: Does economic literacy matter? *Economia Politica* 35: 183–205. [CrossRef]

López de Prado, Marcos. 2019. Beyond Econometrics: A Roadmap Towards Financial Machine Learning. *SSRN*. Available online: https://ssrn.com/abstract=3365282 (accessed on 10 October 2020). [CrossRef]

Lusardi, Annamaria, and Carlo de Bassa Scheresberg. 2013. Financial literacy and high-cost borrowing in the United States. In *NBER Working Paper*; Cambridge: National Bureau of Economic Research.

Lusardi, Annamaria, and Olivia S. Mitchell. 2008. Planning and financial literacy: How do women fare? *American Economic Review* 98: 413–17. [CrossRef]

Lusardi, Annamaria, and Olivia S. Mitchell. 2011. Financial literacy around the world: An overview. *Journal of Pension Economics & Finance* 10: 497–508.

Lusardi, Annamaria, and Olivia S. Mitchell. 2014. The Economic Importance of Financial Literacy: Theory and Evidence. *Journal of Economic Literature* 52: 5–44. [CrossRef]

Lusardi, Annamaria, and Peter Tufano. 2009. Debt Literacy, Financial Experience and Over-Indebtedness. In *NBER Working Paper, 14808*; Cambridge: National Bureau of Economic Research.

Lusardi, Annamaria, de Bassa Scheresberg Carlo, and Oggero Noemi. 2016. Student loan debt in the US: An analysis of the 2015 NFCS Data. *GFLEC Policy Brief*, November 14.

Lusardi, Annamaria, Michaud Pierre-Carl, and Olivia S. Mitchell. 2017. Optimal financial knowledge and wealth inequality. *Journal of Political Economy* 125: 431–77. [CrossRef] [PubMed]

Lusardi, Annamaria. 2019. Financial literacy and the need for financial education: Evidence and implications. *Swiss Journal of Economics and Statistics* 155: 1–8. [CrossRef]

Mandell, Lewis. 2008. *The Financial Literacy of Young American Adults: Results of the 2008 National Jumpstart Coalition Survey of High School Seniors and College Students*. Seattle: University of Washington and the Aspen Institute.

Morgan, Peter J., and Long Q. Trinh. 2019. Determinants and Impacts of Financial Literacy in Cambodia and Viet Nam. *Journal of Risk and Financial Management* 12: 19. [CrossRef]

Moritz, Benjamin, and Tom Zimmermann. 2016. Tree-Based Conditional Portfolio Sorts: The Relation between Past and Future Stock Returns. *SSRN*. Available online: https://ssrn.com/abstract=2740751 (accessed on 15 January 2021).

OECD INFE. 2011. *Measuring Financial Literacy: Core Questionnaire in Measuring Financial Literacy: Questionnaire and Guidance Notes for conducting an Internationally Comparable Survey of Financial literacy*. Paris: OECD.

OECD. 2017. *G20/OECD INFE Report on Adult Financial Literacy in G20 Countries*. Paris: OECD.

OECD. 2020. *Recommendation of the Council on Financial Literacy*. OECD/LEGAL/0461. Paris: OECD.

Ooi, Elizabeth. 2020. Give mind to the gap: Measuring gender differences in financial knowledge. *Journal of Consumer Affairs* 54: 931–50. [CrossRef]

Previati, Daniele A., Ricci Ornella, and Lopes Stentella Lopes. 2020. La capacità delle famiglie italiane di assorbire lo shock pandemico: Il ruolo dell'alfabetizzazione finanziaria. In *L'Italia ai Tempi del Covid-19*. Edited by Paoloni Mauro and Marco Tudino. Toronto: Wolters Kluwer Italia srl, vol. 2.

Probst, Phillipp, and Anne-Laure Boulesteix. 2018. To Tune or Not to Tune the Number of Trees in Random Forest. *Journal of Machine Learning Research* 18: 1–18.

Ridgeway, Greg. 2007. Generalized Boosted Models: A Guide to the gbm Package. Available online: https://cran.r-project.org/web/packages/gbm/gbm.pdf (accessed on 21 May 2018).

Rieger, Marc Oliver. 2020. How to Measure Financial Literacy? *Journal of Risk and Financial Management* 13: 324. [CrossRef]

Therneau, Terry M., Elizabeth J. Atkinson, and Mayo Foundation. 2017. An Introduction to Recursive Partitioning Using the RPART Routines. Available online: https://cran.r-project.org/web/packages/rpart/vignettes/longintro.pdf (accessed on 16 September 2020).

Trunk, Ales, Kocar Sergeja, and Nada Trunk. 2017. Education and training for financial literacy: The role of banks case study Slovenia. International Journal of Innovation and Learning. *Inderscience Enterprises* 22: 385–406.

Van Rooij, Maarten, Lusardi Annamaria, and Rob Alessie. 2011. Financial literacy and stock market participation. *Journal of Financial Economics* 101: 449–72. [CrossRef]

Von Gaudecker, Hans-Martin. 2015. How does household portfolio diversification vary with financial literacy and financial advice? *Journal of Finance* 70: 489–507. [CrossRef]

Xiao, Jing Jian, Serido Joyce, and Soyeon Shim. 2011. Financial education, financial knowledge, and risky credit behaviour of college students. In *Financial Decisions Across the Lifespan: Problems, Programs, and Prospects*. Edited by Douglas J. Lamdin. New York: Springer, pp. 113–28.

Yoong, Joanne. 2011. Financial illiteracy and stock market participation: Evidence from the RAND American Life Panel. In *Financial Literacy: Implications for Retirement Security and the Financial Marketplace*. Edited by Olivia S. Mitchell and Lusardi Annamaria. Oxford: Oxford University Press, pp. 76–100.