

Targeting policy-compliers with Machine Learning: An application to a tax rebate program in Italy*

Monica Andini[†]

Bank of Italy

Emanuele Ciani[†]

Bank of Italy

Guido de Blasio[†]

Bank of Italy

Alessio D'Ignazio[†]

Bank of Italy

Viola Salvestrini[¶]

London School of Economics

First and preliminary version: June 2017.

Please do not quote or circulate without permission.

Abstract

Machine Learning (ML) can be a powerful tool to inform policy decisions. Those who are treated under a program might have different propensities to put in practice the behavior that the policy maker wants to incentivize. ML algorithms can be used to predict the policy-compliers; that is, those who most likely behave in the way desired by the policy maker. When the design of the program is tailored to target the policy-compliers, the policy overall effectiveness is increased. This paper proposes an application of ML targeting that uses a massive tax rebate scheme (“80 euros”) introduced in Italy in 2014.

Keywords: machine learning, prediction, program evaluation, fiscal stimulus

JEL Codes: C5, H3

* This paper was partly written while Viola Salvestrini was an intern at the Structural Economic Analysis Directorate of the Bank of Italy. We would like to thank Audinga Baltrunaite, Andrea Neri, Stefano Gagliarducci, Christian Hansen, Matthew Harding, Enrico Rettore, Paolo Sestito and the participants at the *IX ESPAnet Conference* (Macerata, September 2016), the *First Dondena Workshop on Public Policy* (Milan, November 2016), the *Bank of Italy Seminar Series* (Rome, November 2016), the *Royal Economic Society Annual Conference* (Bristol, April 2017), the *Bank of Italy Internal Workshop on “Big Data”* (Rome, June 2017) for suggestions and comments. The views expressed in this paper are those of the authors and do not necessarily reflect those of the institutions they are affiliated with.

[†] Bank of Italy, Directorate General for Economics, Statistics and Research, Structural Economic Analysis Department, Via Nazionale 91, 00184 Rome, Italy.

[¶] London School of Economics and Political Science, Department of Economics, Houghton Street, London WC2A2AE, United Kingdom.

1. Introduction

Machine Learning (ML) algorithms have been developed in the computer science and statistical literature. Differently from the econometrics literature, which mainly points towards reducing the estimator bias, the focus of ML algorithm is to minimize the out-of-sample prediction error (Athey and Imbens, 2016b; Mullainathan and Spiess, 2017). In this respect, they do not propose new solutions to the fundamental identification issues when dealing with causal effects (see, however, Varian, 2014). Nevertheless, the success of any public intervention depends on the actual implementation, which often requires a decision about who to target. This decision is a “prediction policy problem”, a term coined by Kleinberg et al. (2015). The basic idea is that those who are treated under a program might have different propensities to put in practice the behavior that the policy maker wants to incentivize, or have different payoffs from a given treatment. In this framework, ML algorithms can help targeting the program towards the policy-compliers, that is those who are most likely to behave in the desired way, or have the highest payoff, so to increase the overall effectiveness. The potential of ML predictions for policy decisions cannot go unnoticed. Evaluations might deliver a rigorous assessment of whether a program worked or not, but are often less useful when it comes to advising the policy makers on how to ensure that a given program will work. In particular, even if evaluation studies often propose rich heterogeneity analyses of the policy effect in different sub-groups, they rarely advise on how to build a better targeting rule. In this respect, ML algorithms provide us with some powerful techniques to predict which individuals are more likely to benefit the policy using all the available information.

Early applications of ML targeting include: (i) predicting the riskiest patients for which a joint replacement would be futile (Kleinberg et al., 2015); (ii) improving over judges’ decision on whether to detain or release arrestees as they await adjudication of their case (Kleinberg et al., 2017); (iii) targeting restaurant hygiene inspections (Kang et al., 2013); (iv) predicting highest risk youth for anti-violence interventions (Chandler et al., 2011); (v) predicting the effectiveness of teachers in terms of value added (Rockoff et al., 2011); (vi) hiring police officers who will not behave violently, as well as promoting the best teachers only (Chalfin et al., 2016). We extend this literature by focusing on a tax rebate scheme introduced in Italy in 2014 with the (main) purpose of boosting household consumption. The program, named the “80 euro” bonus,¹ has been a centerpiece of the Italy’s government’s policy efforts to counterbalance the negative consequences of the Great Recession. The scheme was financially considerable: the tax credit was provided to all employees whose annual income ranged between 8,145 and 26,000 euro. According to government’s estimates (Ministry of Economics and Finance, 2015), it entailed a transfer of almost 7 billion euro in 2014, equivalent to 0.4% of Italian GDP.

The effect of the scheme on consumption depends on the allocation rule. If the bonus already targets the recipients that would have benefited the most in terms of increased consumption, then the policy effect is maximized, otherwise there is room for improvement. Our exercise is conducted on data from the Bank of Italy’s Survey on Household Income and Wealth (SHIW).

¹ Hereafter, the expressions *tax credit*, *tax rebate*, and *80 euro bonus* are used interchangeably.

Descriptive evidence suggests that those households who received the bonus in 2014 belong to different points of the household income distribution; not all of them can be classified as needy. In line with previous results about the effectiveness of the policy, we show that having difficulties in making ends meet, as reported by the household head, is associated with a larger impact of the bonus on consumption. We therefore use such a variable to proxy the consumption constrained status. Then, we turn to our core “prediction policy problem”. We assume to be in the position of the Italian Government at the outset of the 2014, when the scheme was designed. We use the SHIW waves (2010 and 2012) available at that time to build a forecasting model to identify the households most likely to be consumption constrained on the basis of observable variables. We also limit the information set to variables that are not sensitive to privacy issues or ethical concerns. Consistently with the current policy we maintain the focus on employees, hence we restrict our analysis and our proposed alternative allocation rule to households with at least one employee.

Our analysis is mostly based on decision tree algorithms (see Hastie et al., 2009). We find them to be ideal for our targeting purposes because they provide an assignment mechanism that can be transparently communicated to the general public by an accountable policy maker. However, we also present, for the sake of comparisons, evidence resulting from non ML methods, linear probability models, and black-box ML routines, such as clustering algorithms. We show that the gains obtained by using the targeting rule given by the decision tree algorithm can be substantial: our results suggest that 29% of the actual expenditure (about 2 billion euro, yearly) has been allocated to recipients that are not identified as the best targets. Very importantly, the decision tree provides us with an allocation rule that involves few variables related to income and financial wealth and it is easily interpretable. We also discuss how to use our insights to design *ex-ante* a possibly more (consumption) efficient allocation for the case of pensioners.

The impact of the “80 euro” bonus on consumption is currently being studied by Neri et al. (2015), Gagliarducci and Guiso (2015) and Pinotti (2015). Neri et al. (2015) use a Difference-in-Differences approach with the longitudinal component of the SHIW survey, and suggest that households who received the bonus spent around 50-60% of it. Gagliarducci and Guiso (2015) use data from the Survey on Household Consumption of the National Statistical Office, matched with administrative data on labor income. They apply a Regression Discontinuity Design exploiting the fact that the bonus depends on labor income thresholds, and find a positive impact of the bonus on food and on mortgage payments, suggesting that the entire bonus goes to consumption. However, using similar data, Pinotti (2015) argues that the effect on total consumption cannot be estimated with precision, and it may even be zero. Differently from these papers, we do not directly focus on the ex-post impact of the bonus on consumption, but we rather study how the policy could be improved by changing the allocation rule. Our evidence on futile expenditures, however, provides boundaries for the magnitude of the ex-post average impact on the treated. For instance, if 29% of the recipients are unlikely to increase consumption out of the bonus, as our ML exercise suggests, then estimated average treatment effects close to unity are not consistent with our study and should be judged as impractical.

The remainder of the paper proceeds as follows. Section 2 provides the relevant details on the design and the implementation of the policy. Section 3 sketches the SHIW features and gathers

information on how the tax credit is detected in the 2014 wave. Section 4 contains descriptive evidence on the effect of the income tax credit on consumption. Section 5 describes the empirical framework used for our targeting analysis and presents the results. Section 6 focuses on the case of retirees, the consequences of the ML targeting for labor taxation, and information requirements. Section 7 concludes.

2. The “80 euro” bonus: institutional details

During the second dip of the economic crisis, total consumption of Italian households dropped from nearly 970 billion Euros in 2010 to 909 in 2013. While foreign demand kept supporting exports, the recession was prolonged by the stagnation of internal demand. A government crisis led to the appointment of a new Prime Minister in February 2014. One of the first announcements was the proposal of a new monetary transfer to households, aimed at boosting consumption and reducing the effective tax rate on labor income. This proposal was reiterated on several occasions during the early weeks of the new government, and formally announced to the press as a “80 euros” transfer on the 12th of March. Figure 1, which describes Google trends for search and news about “80 euros”, shows clearly how the press and general public became fast aware of the bonus. The transfer was finally implemented as a tax credit with the Decree Law 66/2014, which was approved by the Government on the 24th of April, and later ratified by the Parliament. There was no debate about this bonus earlier than the change of government in 2014, and therefore our analysis should not be affected by possible anticipation effects, as it focuses on the entire year 2014 without distinguishing between specific months.

[Figure 1]

The benefit was designed as a tax-credit. It targets employees and holders of similar income² with gross annual income between €8,145³ and €26,000. In particular, the tax credit amounts at €640 per year if the gross annual income ranges between €8,145 and €24,000; for earnings between €24,000 and €26,000, the amount of the benefit is calculated as follows: $80 \times (26,000 - \text{income}) / 2,000$. The tax credit is acknowledged automatically by the employer on the monthly salary paid without a specific request by the beneficiary, on the basis of the predicted annual income according to the current contract and as long as the gross (predicted) tax is greater than the tax deduction for employee income. An individual may receive the bonus on the basis of this predicted income, but will have to give it back if the actual income at the end of the year is outside the eligible range. This will happen at the moment of the annual tax return, which has to be submitted between April and the beginning of July 2015 for the 2014 annual income. Several factors contributed to determine the allocation rule. First of all, the bonus was thought as a tax rebate: employees were therefore chosen as recipients since such tax credit could

² For instance, freelancers, priests, cooperative workers, recipients of unemployment or disability benefits, recipients of scholarships and unemployment insurance (*Cassa Integrazione*).

³ This threshold is €8,145 for individuals who worked the entire year, but it might be lower for workers who have been employed for less than 365 days. Indeed, the bonus was granted under the condition of a net tax greater than zero, which could happen for incomes below the tax area if the employment spell was less than 365 days.

be automatically acknowledged by the withholding agent. This was clearly aimed at making the implementation of the program much easier and faster. Following the same rationale, the “*incapienti*” (namely, those who earn so little that they do not pay tax) were left out, so to avoid the withholding agent to pay the transfer out of pocket. Both mechanisms were introduced to speed up the allocation of the bonus, but they potentially reduced the ability to target households that are consumption constrained.

3. The bonus in the SHIW

The Survey on Household Income and Wealth (SHIW) is a statistical survey conducted on a biennial basis by Bank of Italy, to gather information on the economic behavior of the Italian families at the microeconomic level. The SHIW (www.bancaditalia.it/statistiche) collects the following information: characteristics of the household and of its members (number of income earners, gender, age, education, job status, and characteristics of the dwelling); income (wage and salaries, income from self-employment, pensions and other financial transfers, income from financial assets and real estates); consumption and saving (food consumption, expenses for housing, health, insurance, spending on durable goods, and household saving); wealth in terms of real estate, financial assets, liabilities.

The 2014 wave contains some specific questions on the income tax credit: households were asked if they received the bonus, how many beneficiaries were present within the family and the overall amount received. The overall size of the sample for the 2014 wave is 8,156 households. Since a necessary condition to be eligible for the bonus is to be employed,⁴ we consider only households with at least one employee, for a total of 3,646 observations. Since the survey data we rely on were collected between January and July 2015, that is when the incomes referring to the previous year were secured, we can reasonably assume that individuals knew at the moment of the interview if they were entitled to the bonus or they would have to give it back.

We also make use of previous waves to create the forecasting model aimed at improving the targeting of the bonus. In order to keep information as homogeneous as possible, we only use the other two waves collected after the beginning of the recession (that is, 2010 and 2012). Although the dataset provides us with a large set of covariates, the actual income variable that determines the eligibility for the bonus is not observed. The reasons for this are two-fold. Firstly, the survey collects only income net of taxes and social contributions, while the bonus was determined on the basis of gross income. It is not possible to simply invert the tax formula, because it depends on a full set of deductions that are household-specific (and not reported in the survey). Secondly, the bonus was assigned according to the predicted income, which may differ from the actual annual income.

The SHIW collects both annual expenditure on durable goods and the average monthly non-durable consumption during the year. Households are also asked to report average monthly expenditure on food consumed at home and, separately, on food consumed outside home. The main limitation of these questions is that they are retrospective. The Survey on Household

⁴ We do not consider the other categories of bonus recipients, which however represent a small fraction of the beneficiaries.

Consumption, carried out by the National Statistical Institute and used by Gagliarducci and Guiso (2015) and Pinotti (2015), is based instead on diaries filled in by a sample of households, who are asked to report detailed expenditures during a single month. Although this reduces the risk of misreporting, the fact that consumption refers to a single month of the year increases the volatility of the measure, thus reducing the ability to detect the effect of the 80 euro bonus. In this respect, the nature of the consumption variables in SHIW, which are referred to the whole year (although reported as totals or monthly averages), may be more adequate for our purpose.

4. Descriptive evidence

Using the 2014 SHIW wave, we first show that the households who received the bonus are not always the needy ones, and this signals serious problems of targeting. Figure 2 presents the distribution of the sample by bonus recipients: nearly 40% of the sample received the bonus. Figure 3 reports the difficulty that households face making ends meet by treatment status: there does not seem to be a relevant difference in the distribution of the sample between beneficiaries and non-beneficiaries. Figure 4 presents the distribution of beneficiaries and non-beneficiaries who reported to be liquidity-constrained. Among the recipients nearly 6% reported to be liquidity-constrained, against 2.6% among the non-recipients. Finally, Figure 5 reports the distribution of bonus recipients and non-recipients by income quartiles. Among the recipients, nearly 16% are in the top income quartile. This is due to the fact that we are considering household average income rather than individual income.⁵

[Figure 2 to 5]

We can recover an estimate of the impact of the bonus on consumption by relying on selection on observables. Let $ammbonus_i$ be the amount in euro of the bonus received by household i . The main outcome of interest is consumption c_i (measured in euro) while x_i is a vector of household characteristics, including average income within the household, household size (and its square) and several other characteristics, plus a constant. We then estimate

$$c_i = \delta ammbonus_i + \beta_x x_i + \varepsilon_i \quad (1)$$

$$E[\varepsilon_{it} | ammbonus_i, x_i] = 0 \quad (2)$$

Given that both c_i and $ammbonus_i$ are in euro, δ can be interpreted as the fraction of the bonus spent on consumption, conditional on x_i . In principle, the assumption of selection on observables is appropriate for the policy under scrutiny because the bonus was automatically

⁵ In order to check that our findings are not driven by the wealthier households, we replicate the above statistics excluding the top 5% of the income distribution. Statistics (not shown for space reasons) provide a picture summarized as follows. The percentage of bonus recipients slightly increases while the percentage of households making ends meet easily and very easily is reduced among both recipients and non-recipients. Interestingly, excluding the top 5% of the income distribution does not alter the number of households that report to be liquidity-constrained. Finally, the distribution of non-recipients by income quartiles appears now more uniform while that of recipients is nearly unchanged.

distributed to all eligible individuals on the basis of their tax-relevant information and therefore no self-selection occurred. For the reasons outlined in Section 2, our dataset does not provide information on the precise individual income variables involved in the allocation rule. Nevertheless, we observe a complete set of variables related to it and that, at the same time, have an impact on consumption (usually modeled as well as measured at the household level). It should be noted that in terms of observable characteristics there is a large overlapping between the group of households who received the bonus and the others, and therefore we can compare households who are quite similar but differ by bonus receipt. Notice also that having an higher external validity, as it is the case for selection on observables, is desirable for our targeting analysis, which elaborates on the heterogeneity of the effect. More internally valid estimates - in particular, those obtained with a Regression Discontinuity Design - would provide us with 'local' estimates, making it more difficult to estimate heterogeneity and to make predictions on the overall sample. Having said that, our estimates, as it will be showed in this section, are pretty much in line with those provided by Neri et al. (2015), which adopt a more rigorous identification framework.

We start by estimating the effect of the bonus on average monthly total non-durable consumption (Table 1; Table A.1 provides a description of the variables used in the baseline specifications). We focus only on non-durable consumption because, as discussed by Neri et al. (2015), durable consumption tends to be more volatile and therefore it is difficult to detect the impact of the 80 euro bonus. The total consumption on non-durables excludes rents (imputed or actual), mortgages, and in-kind benefits from the employer. To begin with, the average monthly amount of the bonus perceived and household annual disposable income (net of the bonus) are considered (column (1)). Since current consumption depends - possibly not linearly - on the size of the family, we also control for the number of components and its square (column (2)). Then, we subsequently add a rich set of demographic characteristics such as age, education, gender, marital status (columns (3)-(6)). In column (7), regional fixed effects are included to capture specific factors that might affect all the people residing in the same area. Then, in line with the current debate on the destination of the bonus (Gagliarducci and Guiso, 2015; Pinotti, 2015), we repeat the above estimation on monthly spending for food eaten at home (Table 2). While the effect of the bonus on total consumption is generally not statistically significant and oscillates in magnitude, when considering food expenditure the effect of the bonus is quite consistent throughout the specifications: in particular, for every additional euro received as bonus, roughly 31.5 cents are spent in food consumed at home. This is in line, although slightly larger, with the results by Neri et al. (2015). In a nutshell, we find evidence that the bonus has an effect on food consumption, it is statistically significant even when we introduce additional controls and its economic magnitude is stable across specifications. The effect of the tax rebate on total consumption is not easily detectable in our data because total consumption is likely to be more volatile than food expenditure. In the remaining of the paper, we will mainly focus on food consumption in estimating the inefficiency in the current allocation of the bonus.

[Table 1]

[Table 2]

Using the specification in column (7) of the above tables, we then investigate whether the bonus has a heterogeneous effect on consumption according to different definitions of “needy” households (Table 3). For the estimation in Panel A and B, we construct two indicators taking value 1 for households who face some difficulty making ends meet given the household’s income. In particular, in Panel A $\text{difficulty} = 1$ if the household makes ends meet with great difficulty, with difficulty or with some difficulty. Noting that a considerable number of households reported to face some difficulty getting through the month, in Panel B we slightly modify the definition of difficulty so to include only households making ends meet with great difficulty or with difficulty. In both specifications, total consumption does not seem to significantly react to the tax rebate. Nevertheless, the bonus seems to increase food consumption for all households and the effect is larger and statistically significant for families facing difficulties. In Panel C, we estimate the effect of the bonus among households that report to be liquidity-constrained and households that do not. Differently from the difficulty in making ends meet, this indicator provides an assessment of the overall households’ wealth, rather than income only. The bonus has a significant and positive effect only on food consumption for both constrained and unconstrained households, and such increase is larger for constrained families. In our data, being consumption-constrained is a relevant phenomenon, characterizing about 60 percent of households; on the other hand, only 4 percent of households report themselves to be liquidity-constrained.

[Table 3]

In the remaining of the paper, our preferred indicator of needy households is the difficulty to make ends meet. First of all, this indicator comes from a questionnaire variable available for each wave in the same manner. SHIW 2010 and 2012 contain hypothetical questions about the propensity to consume out of an income shock that are possibly closer to our scope and that have been used also by Jappelli and Pistaferri (2010). However, the question changes significantly between the two waves and it is not available in 2014, making impossible to estimate the heterogeneity in the effect of the bonus according to the different answers. Secondly, we prefer to work with a reasonable number of observations of needy households. This excludes the liquidity constrained indicator that identifies few households as needy, as shown in Panel C of Table 3. Nevertheless, about 90% of the liquidity constrained households also report to have difficulties in making ends meet. Thirdly, we prefer to use a variable that allows us to estimate an effect of the bonus on consumption which is in line with the other papers focused on estimating the effect of the bonus. As one can see from Table 3, the results in Panel A obtained using our preferred indicator are the closest to the ones obtained by Neri et al. (2015). Fourthly, we prefer to use the unrestricted version of the difficulty of make ends meet indicator as the restricted one likely brings us to identify the poorest households, which is not our purpose ex ante.

5. Targeting analysis

5.1 Variable selection

For the targeting analysis, we want to be able to identify in the 2014 sample those households that are more likely to be needy. In this way, we can evaluate the efficiency of the current allocation and suggest some possible alternatives. Although we observe needy households also in the 2014 sample, it is not useful for targeting: we need to rely on information available prior to the start of the policy, which could have been used by the policy maker. We therefore focus on a pooled dataset of 2010 and 2012 waves and estimate models that allow us to predict the needy status on the basis of a set of observable covariates.⁶ We only maintain one rule of the current policy, that is the focus on employees. Hence we select only households with at least one employee.

In order to choose the covariates set, we consider variables that are recorded in both the pooled 2010-12 dataset and the 2014 one, so as to predict the needy status in 2014 using the prediction model estimated on the 2010-12 dataset. Among those variables, we select only those that are observable by the policy maker and not sensitive to privacy issues or ethical concerns so as to end up with a feasible targeting rule. That is, an assignment mechanism that can be transparently communicated to the general public by an accountable policy maker. Finally, we dismiss all the variables that are excluded for collinearity reasons by running a simple regression of our needy status proxy on the covariates set.⁷ The complete list of variables used for prediction can be found in the Table A.2. They essentially refer to household income, wealth, and demographic characteristics.

5.2 Prediction

ML techniques rely on highly flexible functional forms. The degree of flexibility is the result of the following trade-off: as we allow for more flexibility in the model, we improve the in-sample fit at the cost of reducing the out-of-sample fit (over-fitting). In particular, each ML algorithm comes with a regularization of the complexity level: the less we regularize the greater is the ability of the model to approximate the variable to target within the sample; at the same time, an increased level of complexity comes at the risk of lowering the out-of-sample performance. In order to choose the level of complexity, ML algorithms rely on empirical tuning, where the model performance is evaluated over a small portion (randomly chosen) of the dataset: this procedure is repeated several times and the regularization parameter chosen is the one characterized by the best performance on average (cross validation). Given that the main purpose is out-of-sample prediction, we estimate and tune the models on a training subsample, composed of a randomly selected 2/3 of the 2010-12 pooled sample. The remaining 1/3 of such dataset constitutes the testing subsample.

Our main ML algorithm is the *decision tree* (Hastie et al., 2009), which allows in principle to reach a perfect in-sample fit by adding more and more leaves, while the regularization is made

⁶ The sample includes a longitudinal component, but we ignore it because the sample size would be too small for a reliable targeting analysis. Therefore, i univocally identifies a household-year pair.

⁷ This step allows us to select only one variable amongst a set of variables that represent the same thing: for instance, age and year of birth; education represented with different levels of accuracy, and so on.

by pruning the tree. Decision trees are particularly appropriate for applications in which the assignment mechanism needs to be transparent; for instance, when the results need to be shared in order to facilitate decision making (Lantz, 2013). As it will be clear in the Section 5.3, the output of a decision-tree algorithm can be easily described in a graph.

The algorithm divides data into progressively smaller subsets to identify patterns that can be used for predicting a specific binary output. In our case, the algorithm creates a decision rule which partitions the observations according to their needy status on the basis of the values of the observable covariates (z_i). Non-linearities and interactions are captured by the sequence of splits. Following a top-down approach, at each step the algorithm selects a variable z_{gi} from z_i and splits the observations into two groups according to a threshold z_g (or according to a subset of values in case of a multinomial discrete variable). The variable used to split and the threshold are chosen to obtain the largest possible reduction in heterogeneity (impurity) of the variable to be predicted (Siroky, 2009). In the decision tree algorithm that we use,⁸ the degree of impurity at each node (leaves) is measured using a heterogeneity index. The algorithm then proceeds to the next step by further splitting the sub-samples at each terminal node. It stops when the degree of impurity of a terminal node is as low as possible. A high number of levels in a tree is likely to overfit the data. This could lead to a model which performs very well in the training sample, but gives highly imprecise predictions out-of-sample (Athey and Imbens, 2016a; Lantz, 2013; Breiman et al., 1984). A solution to this problem is to reduce the complexity of the tree by setting a complexity parameter (cp) and use it to prune the tree. We choose the optimal cp by using a rule of thumb suggested in the literature (Hastie et al., 2009).⁹

We compare the findings obtained with the decision tree with those deriving from another ML algorithm, the k-Nearest Neighbors (k-NN), and a standard Linear Probability Model (LPM). In the k-NN algorithm (Lantz, 2013), the trade-off between overfitting and out-of sample prediction is solved by choosing the optimal number of neighbors (i.e. the level of k). For each observation in the testing sample, the algorithm identifies the k closest observations from the training sample (the so-called nearest neighbors) and assigns a prediction on the basis of a majority rule, i.e. takes as prediction the most frequent outcome among those of the nearest neighbors. We chose the optimal number k of neighbors by using 10-fold cross-validation. Following the work of Chandler et al. (2011), we also make use of a LPM prediction.¹⁰ To make it more comparable with ML predictions, we includes all the variables depicted in Table A.2, the squares and cubes of the continuous variables, plus all interactions between themselves and all interactions between them and the discrete covariates. In the case of LPM the prediction is continuous, so we consider the dummy needy having the value one if the predicted probability is larger than 0.5. k-NN and LPM are used essentially to probe the prediction quality of the decision tree. In our case, they

⁸ We use the R package “rpart” [<https://cran.r-project.org/web/packages/rpart/rpart.pdf>].

⁹ First, the complexity parameter associated to the smallest cross-validation error (say $errmin$) is found. Then, the optimal cp is the one that has a cross-validated error which is the closest to $errmin + standard_error(errmin)$. The rule of thumb leads to a simpler tree, because the cross-validation error curve tends to be flat around its minimum, hence there is a small gain in picking exactly the minimum, while there is a higher risk of over-fitting.

¹⁰ The results are unchanged if we use a Probit model. We then decide to rely on a LPM since it is easier to interpret the coefficients.

cannot be considered as real alternatives, as we are looking for a transparent assignment mechanism.

5.3 Empirical findings

The decision tree leads to the assignment mechanism shown in Figure 6. It depends on few variables, essentially referring to household income and wealth. The targeted households would be: (a) those that have financial assets lower than 13,255 euro; among these ones, the needy are those that either perceive income lower than 36,040 euro yearly or those that earn more than 36,040 euro but the maximum income perceived within the households is lower than 34,500 euro; (b) those that have financial assets higher than 13,255 euro; among these ones, the needy are those that earn less than 52,591 euro yearly and have an income from financial assets lower than 432.9 euro together with a minimum income perceived within the household lower than 13,895 euro.¹¹ As for a comparison, using either LPM or k-NN to target households would be a much more challenging task. Both methods do not select a subset of the variables, and therefore the actual allocation of the bonus would require acquiring a larger amount of information on each household. Furthermore, (1) both methods require cumbersome computations to obtain the actual index that is used for the allocation and (2) they do not provide clear insights (or not at all, in the case of k-NN) on which characteristics of the household are pivotal in the selection rule.

[Figure 6]

Table 4 compares the performance of the three models in terms of correctly predicting the “needy” status. Notwithstanding its simplicity, the decision tree correctly identifies 74.1% of the observations, a share very close to that of its alternatives (respectively, 73% and 75% for the k-NN and the LPM). Since we are using 2010-2012 information to predict 2014 needy households, we also investigate whether the association between the actual needy status and tree-selected predictors is stable. We run two separate LPM regressions for the 2010 and 2012 subsamples, using as dependent variable the dummy for difficulty in making ends meet and as covariates the variables selected by the tree. The relationship between the observables and the needy status appears to be quite stable, as coefficients change only marginally. Results are presented in Table A.3.

[Table 4]

We proceed to estimate with the 2014 data the effect of the bonus between the households that according to our decision-tree assignment, should have received the bonus, as they would be predicted to be needy, and those that should have not. Table 5 reports the results of the estimation. The effect of the bonus for food consumption is positive and significant for the households that would have been targeted with our assignment rule. The effect is instead neither

¹¹ In principle, targeted households may also include the *incapienti* (see Section 2). Nevertheless, we cannot argue on the actual presence of *incapienti* among members of targeted households as the decision rule we suggest is based on household rather than individual income.

statistically nor economically different from zero for households that received the bonus without being consumption constrained according to the decision tree rule. In particular, households predicted to be needy spend on average 36.9% of the bonus in food consumption. This share is very close to the one estimated by using 2014 data (see Table 2).¹²

[Table 5]

Table 6 provides the percentage of overlap between predicted status (i.e. being needy or not) and the receipt of the bonus. The overlap includes households that: (i) both receive the bonus and are predicted to be needy, and (ii) both do not receive the bonus and are predicted to be non-needy. This fraction is quite low, around 49%. This implies that several households received the bonus but would have had not if the allocation rule was the one that we propose. Given that we find evidence of an impact on consumption only among those predicted to be needy, this implies that there were margins to improve the total effect.

[Table 6]

In order to capture this misallocation, we focus on a measure of spending inefficiency due to the actual allocation rule. As shown in Table 6, 70.9% of the households that receive the bonus are predicted to be needy by the decision tree algorithm. Our spending inefficiency measure refers to the remaining 29.1%. We look at the amount that was spent for the bonus recipients that the decision tree does not identify as needy households. The way we compute such a measure is as follows. Let A be the number of bonus recipients in our dataset, and B the subset of A made up of predicted needy households. The total expenditure for the tax rebate is given by

$$E_{total} = \sum_{i=1}^A ammbonus_i \quad (3)$$

while the “efficient” expenditure (namely, the amount spent for the predicted consumption constrained households) is given by

$$E_{correct} = \sum_{i=1}^B ammbonus_i \quad (4)$$

Therefore, the percentage of expenditure that has been allocated inefficiently can be computed as

¹² One issue is that predictors and the needy status are both measured at the same time. In principle, one would predict the needy status with variables that have been already observed at the time the policy is implemented. Our data do not allow us to follow such a strategy. However, note that the selected predictors such as income and wealth are characterized by a high degree of persistence. In particular, we use the panel component of the dataset and regress each predictor measured in 2014 on its 2010-12 average value. Such an estimate is roughly 0.9 for the two main predictors (income and financial assets).

$$\frac{E_{total} - E_{correct}}{E_{total}} \quad (5)$$

This share turns out to be equal to 29% of the total expenditure. In order to maximize the coverage of the program, this amount could be reallocated to those households that are predicted to be needy but did not receive the bonus. One possibility is to endow this group with a transfer which is set to be equal to the per capita transfer received by households belonging to *B* (i.e. roughly 57 euro). In such a case, keeping fixed the total public expenditure for this transfer, we could reach 30% of predicted needy households that did not receive the bonus. In this way, 60% of the households we predict as needy would be endowed with a bonus.

6. Extensions and discussions

6.1 Retirees

We sketch how a prediction exercise may be implemented for the retirees, which have not been considered under the scheme so far but may be included in the near future. We consider the households in which there is at least one retiree and create an allocation rule which mostly resembles the one currently applied to employees, given that also for the pensioners we do not observe gross income. To be sure, we consider the distribution of the average income from pensions within the household and we divide it into quartiles. In line with the 80 euro allocation rule, we then exclude the poorest and richest households. In other words, we assume that the bonus recipients are those households whose average income from pension is between the second and the third quartile of the distribution. The income thresholds thus obtained are consistent with those defined for the 80 euro bonus.

In order to check whether there is any baseline difference in food consumption behavior between potential recipients and non-recipients, we regress food expenditure on income and the same set of controls we used in our main exercise. Table 7 (Panel A) shows that such difference seems not to exist.¹³ We then assume that potential recipients are those households who are “needy” according to the definition we used in our main exercise (i.e., households that make ends meet with some difficulty, with difficulty or with great difficulty). Panel B provides evidence of a different food consumption behavior between needy and not needy households. Namely, the effect of an income perturbation on food expenditure is higher for needy households than for the non-needy ones.

[Table 7]

We then apply the decision tree algorithm to the retirees sample. Figure 7 shows the profile of the predicted “needy” retirees. The targeted households in this case are those perceiving income

¹³ We also consider the possibility to allocate the bonus to the poorest retirees and thus re-run the above exercise assuming that recipients are those households whose average income from pension is in the first quartile of the distribution. Even in this case, we find no evidence of a statistically significant difference in food consumption behavior between potential recipients and non-recipients. This reinforces our argument that allocating the bonus on the basis of income only may be misleading.

lower than 25,509 euro yearly or those that earn more than 25,509 euro but have income from financial assets lower than 126.6 euro together with an household yearly income from retirement lower than 29,700 euro.

[Figure 7]

Table 8 reports the results of the estimation using the predicted “needy” status. Again, the effect of an income increase on food consumption is higher for those households that are predicted to be needy. It is worth noting that targeting the needy households according to the decision tree algorithm delivers a positive and statistically significant effect of an income increase on total consumption as well.

[Table 8]

6.2 *Taxing labor*

So far, we have neglected the fact that the bonus was also intended to reduce the fiscal burden on wages, although we have maintained the basic rule of the current policy by focusing only on households with at least one employee. Within this group, our decision-tree rule is designed to maximize the consumption impact of the bonus. With respect to the current allocation, our proposal is also more tailored towards the bottom of it, hence it entails a redistributive effect which is likely to be stronger than the current one.

It is more complicated to understand how the two policies differ with respect to their impact on labor supply. The actual effect of the current policy on employment is, from a theoretical standpoint, ambiguous. On the one hand, the 80 euro reduce the overall taxation on wages for eligible individuals, hence they increase the probability that an individual accepts an employment offer with a given salary.¹⁴ On the other hand, the bonus may have negative effects on the intensive margin, because for individuals whose gross wage is between 24,000 and 26,000 the effective tax rate, which accounts for the decrease in the bonus amount, is extremely high.¹⁵ The overall effect on labor supply of the current scheme depends on the combination of the two mechanisms. Also the rule that we suggest influences labor supply through these two channels. Nevertheless, it is hard to say ex-ante which rule would be preferable from this point of view. Given that our rule favors low income and low wage individuals, the increase in the probability of accepting an employment offer is likely to be larger than with the actual policy. However, it may also cause a stronger decrease in the intensive margin. Furthermore, looking at the household level response, our proposed rule has a drawback with respect to the current one. Being designed on family characteristics, eligibility depends not only on the individuals’ income, but on the household’s. As a result, it may have a negative effect on the second earner, usually the female, as her decision to work may lead the household not to be eligible anymore. A full discussion requires a behavioral micro-simulation model that simulates the agents’ behavioral

¹⁴ Obviously, nothing changes if the offered wage would bring the individual above the income threshold.

¹⁵ As intensive margin we refer to overtime work, because the switch from part-time to full-time is already accounted for in the actual rule.

response accounting for the heterogeneity in their budget constraints (see, for instance, Pacifico, 2009).

In any case, one should consider that the adoption of the ML rule delivers substantial fiscal savings. 29% of the allocated expenditures did not serve to increase consumption. That is, the public budget that can be devoted to other targets is now increased. This means that more resources to cut tax (both on households and firms) are available.

6.3 Data requirement

The decision tree rule is based on information at household level that, at least in principle, is observable by a policy maker. As an matter of example, the equivalent economic situation indicator (i.e., the so called 'ISEE') enables the policy makers to collect information on income and wealth at household level. We are aware that implementing the targeting rule we suggest may increase the costs of the policy in the short term because it would require, using the same example, to know the ISEE of all Italian households. However, the use of household-level information is also in line with other recent proposals to review some assistance benefits policies aiming at the use of eligibility criteria that approximate the ISEE or, more generally, the household economic condition. In short: data defined at the household level are going to be collected anyway to comply with a more efficient welfare system.

Note also that having only a subset of the (few) variables included in the decision tree assignment rule will deliver lower but still sizable benefits. Indeed, a useful feature of the decision tree algorithm is the possibility to compute the fraction of households that would be incorrectly identified as needy by observing only a subset of characteristics among those involved in the tree. For instance, let us assume that the policy maker can observe household financial assets and disposable income only. In this case, her decision rule to identify needy households could be based only on the financial assets and income thresholds given by the tree. In terms of Figure 2, needy households would be those that have financial assets lower than 13,255 and disposable income lower than 36,040 euro and non-needy households would be those that have financial assets at least equal to 13,255 and disposable income at least equal to 36,040 euro yearly. Clearly, these groups do not overlap with the groups of predicted needy and predicted non-needy households identified through the use of all the variables involved in the tree. Using a decision rule based on financial assets and disposable income only, 22% of the households would be allocated to a status that does not correspond to the one predicted by the use of all the variables (i.e., the entire tree). If the policy maker observes the maximum income perceived within the household too, and constructs a decision rule also based on this variable using the thresholds given by the tree, then the fraction of incorrectly allocated households decreases to 5%. Finally, the fraction of incorrectly identified needy households is 0 in case both income from financial assets and minimum income perceived within the household are observable because the decision rule now coincides with the entire tree.

7. Conclusions

During economic downturns, well-designed programs may contribute to the recovery. A key ingredient for an effective policy is an accurate targeting of beneficiaries, who should behave in the way the policy-maker wants to incentivize. This ideal framework unlikely corresponds to the actual one because of a trade-off between ease and accuracy of the targeting rule. Machine Learning algorithms help addressing such a trade-off as they allow to target units that most likely behave in the desired way or to gain more from the policy.

In this paper, we focus on a massive tax rebate program recently implemented in Italy. We make use of ML techniques to identify the households that would have benefited the most from the program in terms of increased consumption. To do so we use a decision tree and find that 29% of the actual expenditure has been allocated to recipients that are not the best target for policy effectiveness.

References

- Athey, S. and Imbens, G. (2016a). Recursive partitioning for heterogeneous causal effects. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 113, pages 7355–7360.
- Athey, S. and Imbens, G. (2016b). The state of applied econometrics. Unpublished manuscript.
- Breiman, L., Friedman, J., Stone, C. J., and Olshen, R. A. (1984). *Classification and regression trees*. Belmont: Wadsworth.
- Chalfin, A., Danieli, O., Hillis, A., Jelveh, Z., Luca, M., Ludwig, J. and Mullainathan, S. (2016). Productivity and selection of human capital with machine learning. *American Economic Review: Papers & Proceedings*, 106(5): 124-127.
- Chandler, D., Levitt, S. D., and List, J. A. (2011). Predicting and preventing shootings among at-risk youth. *The American Economic Review*, 101(3):288–292.
- Coady, D., Grosh, M. E., and Hoddinott, J. (2004). *Targeting of transfers in developing countries: Review of lessons and experience*. Washington D.C.: World Bank.
- Gagliarducci, S. and Guiso, L. (2015). Gli 80 euro? Spesi al supermercato. *Lavoce.info*.
- Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning*, 2nd edition. Springer.
- Jappelli, T. and Pistaferri, L. (2010). The consumption response to income changes. *Annual Review of Economics*, 2(1):479–506.
- Jappelli, T. and Pistaferri, L. (2014). Fiscal policy and MPC heterogeneity. *American Economic Journal: Macroeconomics*, 6(4):107–136.
- Kang, J. S., Kuznetsova, P., Luca, M., and Choi, Y. (2013). Where not to eat? Improving public policy by predicting hygiene inspections using online reviews. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics. Stroudsburg, pages 1443–1448.
- Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., and Mullainathan, S. (2017). Human decisions and machine predictions. NBER Working Paper No. 23180, February 2017, National Bureau of Economic Research .
- Kleinberg, J., Ludwig, J., Mullainathan, S., and Obermeyer, Z. (2015). Prediction Policy Problems. *American Economic Review: Papers & Proceedings*, 105(5):491–495.

- Lantz, B. (2013). *Machine learning with R*. Birmingham: Packt Publishing Ltd.
- McBride, L. and Nichols, A. (2015). Improved poverty targeting through machine learning: An application to the U.S. aid poverty assessment tools. Unpublished manuscript. Available at: http://www.econthatmatters.com/wp-content/uploads/2015/01/improved_targeting_21jan2015.pdf.
- Ministry of Economics and Finance (2015). Statistiche sulle dichiarazioni fiscali. Analisi dei dati Irpef - Anno d'imposta 2014. Technical report. Available at: http://www1.finanze.gov.it/finanze2/analisi_stat/v_4_0_0/contenuti/analisi_dati_2014_irp-ef.pdf?d.
- Mullainathan, S. and Spiess, J. (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2):87–106.
- Neri, A., Rondinelli, C., and Scoccianti, F. (2015). The marginal propensity to consume out of a tax rebate: the case of Italy. Unpublished manuscript. Available at: https://www.bancaditalia.it/pubblicazioni/altri-atti-convegni/2015-analysis-household-finances/papers/2.Neri_Rondinelli_Scoccianti.pdf.
- Pacifico, D. (2009). A behavioral microsimulation model with discrete labour supply for Italian couples, CAPP Working Paper No. 65, March 2009, Department of Economics, University of Modena and Reggio Emilia.
- Pinotti, P. (2015). 80 euro rimasti nel portafoglio. *Lavoce.info*.
- Rockoff, J. E., Jacob, B. A., Kane, T. J., and Staiger, D. O. (2011). Can you recognize an effective teacher when you recruit one? *Education*, 6(1):43–74.
- Siroky, D. S. (2009). Navigating random forests and related advances in algorithmic modeling. *Statistics Surveys*, 3(0):147–163.
- Varian, H. R. (2014). Big data: new tricks for econometrics. *Journal of Economic Perspectives*, 28(2): 3–28.

Figure 1: Google trends for “80 euro”, Italy

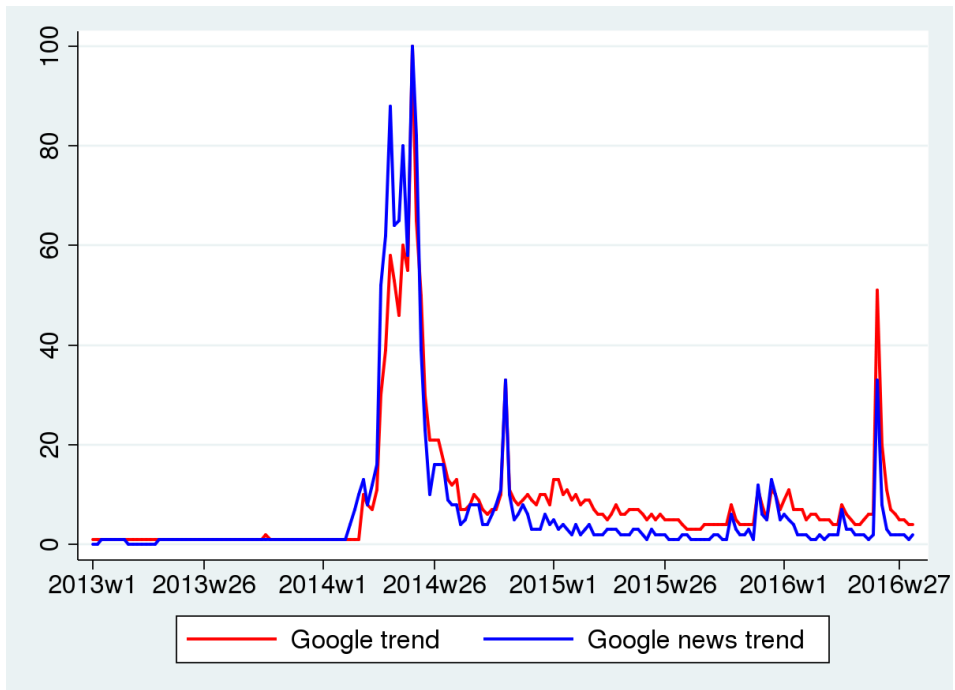


Figure 2: Bonus Recipients

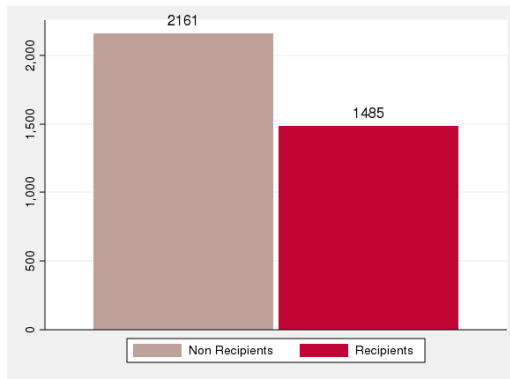


Figure 3: Difficulty making ends meet by treatment status

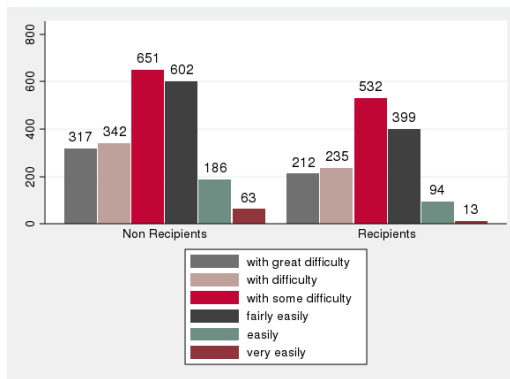


Figure 4: Liquidity constraints by treatment status

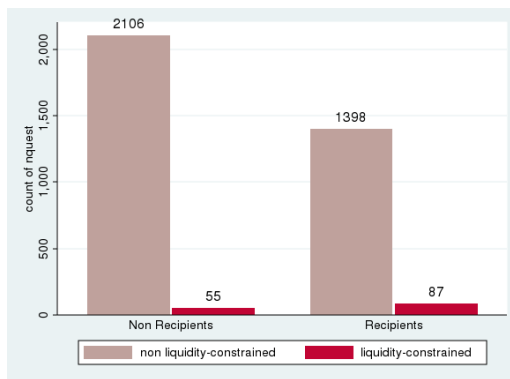


Figure 5: Income quartiles by treatment status

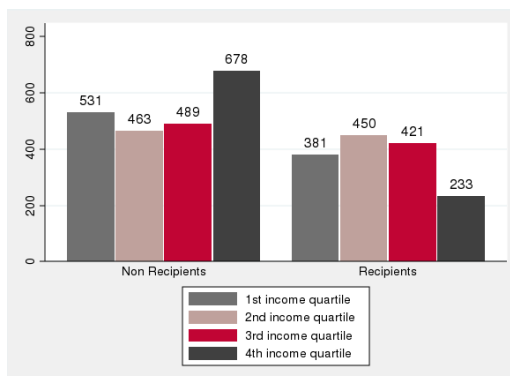
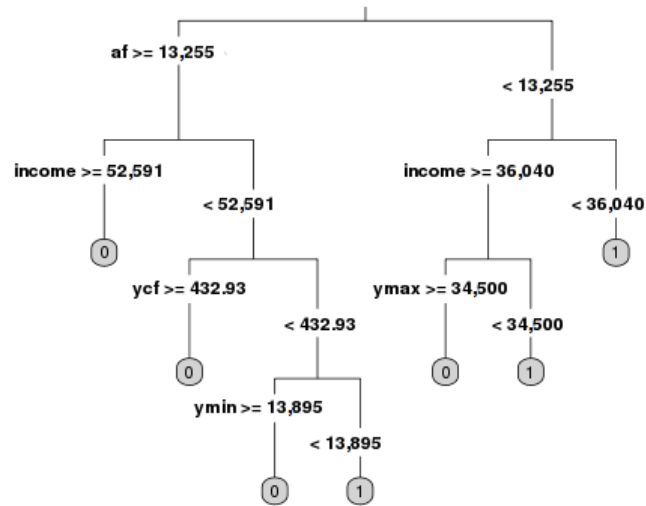


Figure 6: Decision tree output


Classification Tree for Needy Households



Legend	
af	household yearly financial assets
income	income household yearly disposable
ycf	household income from financial assets
ymin	minimum individual labor income within the household
ymax	maximum individual labor income within the household

Figure 7: Retirees: Decision tree output

Retirees: Classification Tree for Needy Households



Legend	
income	income household yearly disposable
ycf	household income from financial assets
ytp	household yearly income from retirement

Table 1: Effect of the bonus on total non-durable consumption

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ammbonus	0.537*	0.114	0.134	0.217	0.214	0.208	0.373
	(0.281)	(0.273)	(0.275)	(0.275)	(0.276)	(0.276)	(0.276)
income	0.0231***	0.0217***	0.0216***	0.0205***	0.0204***	0.0204***	0.0202***
	(0.000859)	(0.000885)	(0.000922)	(0.000992)	(0.00100)	(0.00100)	(0.00106)
ncomp		159.4***	157.5***	176.8***	177.1***	159.3***	153.5***
		(27.04)	(26.74)	(27.13)	(27.15)	(28.92)	(28.45)
ncomp2		-10.20**	-9.796**	-11.73***	-11.78***	-9.947**	-8.189*
		-4.379	-4.335	-4.303	-4.309	-4.381	-4.200
age			0.797	2.042**	2.024**	1.972**	2.600***
			(0.764)	(0.866)	(0.866)	(0.862)	(0.894)
diploma				69.53***	69.59***	69.73***	68.24***
				(21.14)	(21.14)	(21.13)	(20.96)
degree				150.7***	151.5***	152.0***	148.7***
				(35.21)	(35.27)	(35.30)	(35.35)
male					6.974	2.076	9.219
					(18.58)	(18.93)	(19.06)
married						34.39	42.03*
						(22.45)	(22.35)
Constant	551.4***	247.1***	211.5***	94.08*	91.90*	109.4**	145.8**
	(27.72)	(34.22)	(44.67)	(55.16)	(55.26)	(55.47)	(74.23)
Regional FE	NO	NO	NO	NO	NO	NO	YES
N	3646	3646	3646	3646	3646	3646	3646
R2	0.470	0.493	0.493	0.497	0.497	0.497	0.510

Notes: Estimation on the 2014 dataset. *-**-*** denotes statistical significance at 10%-5%-1%. Robust standard errors in parentheses.

Table 2: Effect of the bonus on food consumption

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ammbonus	0.497*** (0.101)	0.207** (0.0933)	0.255*** (0.0944)	0.263*** (0.0941)	0.265*** (0.0940)	0.262*** (0.0939)	0.315*** (0.0960)
income	0.00527*** (0.000271)	0.00426*** (0.000253)	0.00400*** (0.000253)	0.00387*** (0.000276)	0.00389*** (0.000278)	0.00386*** (0.000278)	0.00375*** (0.000286)
ncomp		114.2*** (11.19)	109.5*** (10.62)	111.5*** (10.72)	111.3*** (10.74)	102.3*** (11.38)	102.8*** (11.07)
ncomp2		-7.850*** (1.890)	-6.868*** (1.796)	-7.040*** (1.793)	-7.007*** (1.795)	-6.072*** (1.816)	-5.951*** (1.754)
age			1.936*** (0.259)	2.097*** (0.286)	2.109*** (0.285)	2.082*** (0.284)	2.237*** (0.285)
diploma				11.98* (6.892)	11.94* (6.893)	12.01* (6.888)	9.926 (6.803)
degree				16.45 (11.46)	16.00 (11.48)	16.25 (11.50)	17.90 (11.43)
male					-4.279 (5.911)	-6.773 (6.000)	-1.840 (5.981)
married						17.51** (7.633)	21.95*** (7.635)
Constant	292.0*** (9.319)	78.00*** (13.03)	-8.416 (15.83)	-24.11 (18.50)	-22.77 (18.73)	-13.87 (19.15)	-30.81 (25.05)
Regional FE	NO	NO	NO	NO	NO	NO	YES
N	3646	3646	3646	3646	3646	3646	3646
R2	0.269	0.388	0.397	0.398	0.398	0.398	0.416

Notes: Estimation on the 2014 dataset. *-**-*** denotes statistical significance at 10%-5%-1%. Robust standard errors in parentheses.

Table 3: Effect of the bonus on consumption: heterogeneity analysis

	Total Consumption		Food Consumption	
	(1)	(2)	(3)	(4)
<i>Panel A</i>	No Difficulty	Difficulty	No Difficulty	Difficulty
ammbonus	-0.222 (0.492)	0.503 (0.333)	0.0518 (0.161)	0.357*** (0.117)
N	1357	2289	1357	2289
R2	0.485	0.438	0.385	0.437
<i>Panel B</i>	No Difficulty (restricted)	Difficulty (restr.)	No Difficulty (restr.)	Difficulty (restr.)
ammbonus	0.0815 (0.340)	0.635 (0.444)	0.151 (0.116)	0.486*** (0.166)
N	2540	1106	2540	1106
R2	0.478	0.474	0.398	0.450
<i>Panel C</i>	Liquidity unconstrained	Liquidity constrained	Liquidity unconstrained	Liquidity constrained
ammbonus	0.292 (0.273)	-0.721 -2.040	0.294*** (0.0987)	1.057*** (0.393)
N	3504	142	3504	142
R2	0.522	0.365	0.413	0.613

Notes: Estimation on the 2014 dataset. All controls of specification (7) in Tables 1-2 included. In Panel A, *difficulty* =1 if the household makes ends meet with great difficulty, with difficulty or with some difficulty. In Panel B, *difficulty* =1 if the household makes ends meet with great difficulty or with difficulty. In Panel C, *liquidity constrained* =1 if the household was at least partially rejected a request for a mortgage, or would have liked to apply for it but had not because they thought they would have been rejected. *-*-*** denotes statistical significance at 10%-5%-1%. Robust standard errors in parentheses.

Table 4: Decision Tree, k-NN and LPM models performance compared

		Real status		
		Not Needy	Needy	Total
<i>Panel A: decision tree</i>				
Predicted Status	Not Needy	608	232	840
	Needy	447	1334	1781
	Total	1055	1566	2621
	% Correctly Predicted	57.6%	85.1%	74.1%
<i>Panel B: k-NN</i>				
Predicted Status	Not Needy	593	244	837
	Needy	462	1322	1784
	Total	1055	1566	2621
	% Correctly Predicted	56.2%	84.4%	73.0%
<i>Panel C: LPM</i>				
Predicted Status	Not Needy	608	208	816
	Needy	447	1358	1805
	Total	1055	1566	2621
	% Correctly Predicted	57.6%	86.7%	75.0%

Notes: Out-of-sample estimation on the testing subsample of the 2010-2012 pooled dataset. *Needy* =1 if the household makes ends meet with great difficulty, with difficulty or with some difficulty.

Table 5: Effect of the bonus on consumption by predicted needy households (decision tree)

	Total Consumption		Food Consumption	
	(1)	(2)	(3)	(4)
	Not Needy	Needy	Not Needy	Needy
ammbonus	-0.527 (0.563)	0.710** (0.315)	0.00907 (0.184)	0.369*** (0.111)
N	1146	2500	1146	2500
R2	0.459	0.415	0.356	0.442

Notes: Estimation on the 2014 dataset. All controls of specification (7) in Tables 1-2 included. *Needy* =1 if according to the decision tree algorithm the household is predicted to make ends meet with great difficulty, with difficulty or with some difficulty. *_**-* denotes statistical significance at 10%-5%-1%. Robust standard errors in parentheses.

Table 6: Decision Tree rule: predicted status and real status

		Predicted status			% Overlapping
		Not Needy	Needy	Total	
Real Status	Not Recipient	715	1446	2161	33.0%
	Recipient	431	1054	1485	70.9%
	Total	1146	2500	3646	
	% Overlapping	62.4%	42.1%		48.5%

Notes: Estimation on the 2014 dataset. *Needy* =1 if the household is predicted to make ends meet with great difficulty, with difficulty or with some difficulty.

Table 7: Retirees: effect of an income shock on food consumption by potential recipients

	Food Consumption	
	(1)	(2)
<i>Panel A: allocation rule based on income quartiles</i>	Not Recipient	Recipient
income	0.00302*** (0.000474)	0.00325*** (0.000740)
N	2380	2373
R2	0.427	0.427
<i>Panel B: allocation rule based on the needy status</i>	Not Recipient	Recipient
income	0.00197*** (0.000401)	0.00619*** (0.000445)
N	1778	2975
R2	0.323	0.474

Notes: Estimation on the retirees 2014 dataset, controlling for the presence of bonus recipients. All controls included. In Panel A, *Recipient* =1 if average income from pension within the household is between the 2nd and 3rd quartile of the distribution. In Panel B, *Recipient* =1 if the household makes ends meet with great difficulty, with difficulty or with some difficulty. *_**_*** denotes statistical significance at 10%-5%-1%. Robust standard errors in parentheses.

Table 8: Retirees: effect of an income shock on consumption by predicted needy households

	Total Consumption		Food Consumption	
	(1)	(2)	(3)	(4)
	Not Needy	Needy	Not Needy	Needy
income	0.0115*** (0.00180)	0.0185*** (0.00121)	0.00155*** (0.000373)	0.00544*** (0.000478)
N	1475	3278	1475	3278
R2	0.353	0.463	0.235	0.453

Notes: Estimation on the retirees 2014 dataset, controlling for the presence of bonus recipients. All controls included. *Needy* =1 if the household is predicted to make ends meet with great difficulty, with difficulty or with some difficulty. *-*-*-*** denotes statistical significance at 10%-5%-1%. Robust standard errors in parentheses.

Table A.1: Variables description: baseline regressions

total consumption	average monthly spending on all items
food consumption	average monthly spending on food eaten at home
ammbonus	overall amount of the bonus received monthly by the household
income	household annual disposable income (net of the bonus)
ncomp	number of household members
age	age of the head of the household
male	=1 if the head of the household is male
diploma	=1 if the head of the household has a upper secondary school diploma
degree	=1 if the head of the household has a university degree or more
married	=1 if the head of the household is married
liquidity constrained	=1 if the household is liquidity-constrained
difficulty	=1 if the household makes ends meet with great difficulty, difficulty or some difficulty
difficulty (restricted)	=1 if the household makes ends meet with great difficulty or with difficulty

Table A.2: Variables description: Machine Learning dataset

godabit	=1 if the household flat is property of household members
nfigli	number of household sons and daughters [used as discrete variable]
carta	=1 if some member of the household holds a credit card
bancomat	=1 if some member of the household holds a debit card
cartapre	=1 if some member of the household holds a prepaid card
altrab	=1 if some member of the household holds properties different than residence house
debita1	=1 if the household has house-related debts (acquisition or restructuring)
ncomp	number of household components [used as discrete variable]
income	household annual disposable income (net of the bonus)
yl	household income from employment
ytp	household income from retirement
ym	household income from self-employment
yca	household income from real estate
ycf	household income from financial assets
af	household financial assets
ymin	minimum individual labor income within the household
ymax	maximum individual labor income within the household
native	=1 if the head of the household is Italian
staciv	civil status of the head of the household
age	age of the head of the household
q	employment condition of the head household [used as discrete variable]
nperc	number of income perceivers within the household [used as discrete variable]
acom4c	dimensional class of the household municipality of residence [used as discrete variable]
degree	=1 if the head of the household has a university degree or more
diploma	=1 if the head of the household has a upper secondary school diploma
compulsory	=1 if the head of the household has a compulsory education
africa	=1 if the head of the household is African
asia	=1 if the head of the household is Asian
east europe	=1 if the head of the household is East-European
south america	=1 if the head of the household is South-American
south	=1 if the head of the household lives in the South of Italy

Table A.3: Probability of being a needy household

	(1)	(2)	(3)
af	0.0000149 (0.000117)	-0.00000999 (0.000174)	0.000107 (0.000186)
income	-0.00873*** (0.000426)	-0.00873*** (0.000653)	-0.00874*** (0.000535)
ycf	-0.00159 (0.00513)	-0.0116** (0.00483)	-0.00131 (0.00519)
y _{max}	0.000227 (0.00152)	-0.000242 (0.00126)	0.000424 (0.00252)
y _{min}	-0.00653*** (0.00140)	-0.00812*** (0.00112)	-0.00518** (0.00227)
Constant	1.035*** (0.0203)	1.040*** (0.0175)	1.037*** (0.0296)
N	7802	3939	3863
R ²	0.223	0.229	0.222

Notes: Columns (1), (2) and (3) are estimated on the 2010-2012 pooled dataset, the 2010 dataset and the 2012 dataset, respectively. We focus on a Linear Probability Model to ease the interpretation of the coefficients. All coefficients (excluding the constant term) have been multiplied by 1000. *Needy* =1 if the household makes ends meet with great difficulty, with difficulty or with some difficulty. *-*-*-*** denotes statistical significance at 10%-5%-1%. Robust standard errors in parentheses.