Targeting policy-compliers with machine learning: an application to a tax rebate programme in Italy

M. Andini†, E. Ciani†, G. de Blasio†,
A. D'Ignazio†, V. Salvestrini‡

†Bank of Italy, ‡LSE

Bank of Italy Workshop on Big Data
Rome, 26-27 March 2018

Prediction policy problems

- The impact of a policy depends on **who benefits** from it
  - Some individuals may have a higher payoff
  - The effect may be zero on some groups, also because some may be less likely to put into practice the incentivized behavior
- Predicting who is more likely to belong to these groups is a **prediction policy problem** (Kleinberg et al., 2015)
  - Predicting effectiveness of teachers in terms of value added (Rockoff et al., 2011)
  - Targeting of households at risk of poverty in developing countries (McBride and Nichols, 2015)
- How can Machine Learning (ML) help us with this task?

Machine Learning for Policy Prediction

- ML methods can help us because they are suited for prediction:
    - A targeting rule has to work well **out-of-sample**, and ML models are trained and validated to this purpose
    - ML allows us to use **data driven specifications**
- But via prediction we may not obtain what we want
    - Predicting individuals at risk is not the same as predicting individuals with stronger effect (Ascarza, 2017; Athey, 2018)
    - We need to assess whether **targeting-on-prediction increases effectiveness** by exploiting the causal inference toolbox
    - Hence out-of-sample accuracy is not sufficient
    - We try to reconcile the two by combining ex-ante prediction and ex-post evaluation

Our paper

- Our paper provides an application of ML targeting to a massive tax-rebate introduced in Italy in 2014
- We consider the hypothetical situation in which its only purpose was to increase consumption
- We set ourselves in an ex-ante situation in which theory and previous evidence suggests which group should be targeted to reach this purpose, but this group cannot be directly observed and needs to be predicted
- We use supervised ML methods to carry out this prediction and we check whether it would have improved the impact on consumption

The tax rebate

- In Spring 2014 the Italian Government introduced a large scale tax credit
- The tax rebate aimed at reshaping labor taxation and also aimed at providing a short-run boost to consumption
- It was channeled to employees with annual income between 8,145 and 26,000 euro
- 640 euro/year if income $\leq 24,000$ and declined to zero until the 26,000 threshold
- The total transfer was almost 7 billion euro in 2014 (0.4% of Italian GDP)

The target group

- The impact of a fiscal policy on consumption also depends on the heterogeneity of marginal propensity to consume (Jappelli and Pistaferri, 2014)
- Neri et al (2018) find that the tax rebate had an impact on consumption, but it was stronger for **consumption constrained** households
- In the Bank of Italy's Survey on Household Income and Wealth (SHIW) we have a (reasonably good) proxy of consumption constraints
- We look at **Needy** which is a dummy for the household reporting to make ends meet with at least some difficulty (needy = consumption constrained from now on)

Targeting with ML

- The policy-maker cannot obviously know who are the households with difficulties in making ends meet
- Our targeting exercise aims at identifying them, through the following steps:
  1. We use the 2010 and 2012 SHIW to train and test a model for the needy (constrained) status on the basis of observed covariates
  2. We predict this status on the 2014 wave and we assess whether the impact of the tax rebate was larger among predicted needy households
  3. We show how much the predicted needy households overlap with actual recipients of the rebate

Step 1

- We use the 2010 and 2012 SHIW waves to train a model to predict “needy” households on the basis of observable characteristics (income, demographic, etc) in a 2/3 randomly selected training sample
- We use a decision tree, because it allows more transparency (more on this later)
- We correctly predict 74.1% of observations in the 1/3 testing sample (somewhat similars to other methods - but slightly worse)

Step 1 (cont.) Classification tree for needy households

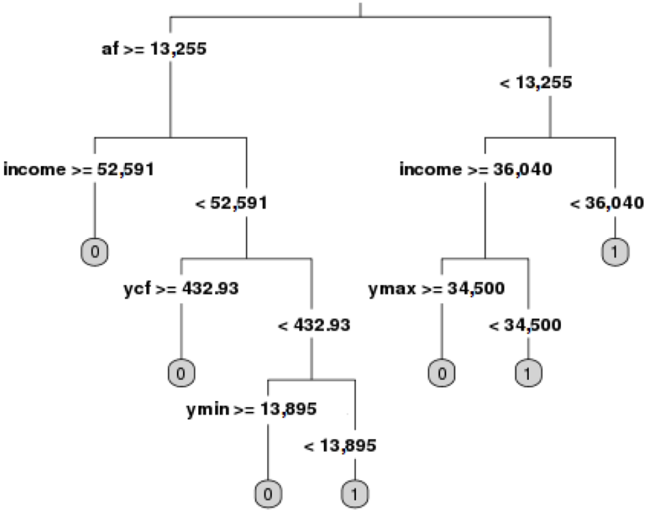

Step 2

- We switch to 2014 and we predict the needy status
- We provide evidence about the effect of the tax rebate on consumption according to the predicted status, estimating:

$$c_i = \delta \text{Rebate amount}_i + \boldsymbol{\beta}_{\mathbf{x}} \mathbf{x_i} + \varepsilon_i \quad (1)$$

$$E[\varepsilon_i | \text{Rebate amount}_i, \mathbf{x_i}] = 0, \quad (2)$$

- Assuming selection on observables, δ recovers the impact of the rebate on consumption.
  - the tax rebate was automatically distributed to all eligible individuals, so no self-selection occurred
  - we observe a set of relevant variables related to the rebate that impact on consumption

Step 2 (cont.)

	Total consumption		Food consumption	
	Not_needy	*Needy*	*Not_needy*	*Needy*
	(1)	(2)	(3)	(4)
Rebate amount	-0.527	0.710**	0.009	0.369***
	(0.563)	(0.315)	(0.184)	(0.111)
Controls	Y	Y	Y	Y
N	1146	2500	1146	2500
R^2	0.459	0.415	0.356	0.442

- We find evidence that the effect is indeed heterogeneous
- This ensures that our prediction satisfies our purposes

Step 3

We show how much the predicted needy households overlap with the actual tax rebate recipients

		Predicted status (decision tree)			
		Not_needy	*Needy*	Total	Overlap (%)
	Not Recipient	715	1446	2161	33.0%
Tax rebate	Recipient	431	1054	1485	70.9%
	Total	1146	2500	3646	
	Overlapping (%)	62.4%	42.1%		48.5%

- The overlap (i.e. households that: (i) receive the rebate and are predicted to be constrained, and (ii) did not receive the rebate and are predicted to be non-constrained) is 48.5%
- 29.1% of the households that received the rebate are predicted not to be constrained

Issues

- Omitted payoffs (Kleinberg et al, 2017)
  - A full evaluation of the proposed targeting also needs to consider the other targets of the policy
  - In terms of poverty, our targeting would reach poorer households; in terms of labor supply a full assessment requires a structural model to assess both the extensive and intensive margins (and the second earner's reaction)
- Prediction stability (Athey, 2018)
  - We assessed whether the association between the variable used by the decision tree and the needy status is stable between 2010 and 2012
  - The effectiveness check on 2014 also confirms stability

Issues (cont.)

- Manipulability (Athey, 2018)
    - Households may manipulate their variables to (ex-post) be part of the predicted constrained
    - Even the current rule is subject to manipulation, which is made more difficult in our targeting rule because more variables would have to be altered
- Transparency
    - We used decision tree as they provide a more transparent decision rule
    - This implies a loss in prediction accuracy with respect to more black-box models
    - Formal vs actual transparency: targeting via ML forces us to clearly indicate the purpose and assess whether the rule is fit for it

Conclusions

- We discuss how to use ML algorithms on survey data to improve the targeting of a cash benefit
- Our approach shows that they can increase effectiveness, but we need to be careful in how we validate the model: prediction accuracy is not enough!
- One alternative is to work ex-post in a causal-inference framework and assess the heterogeneity of the effect (using ML), but this might be costly and have low external validity
- Our application is not the best to show the gains of ML, as we have few observed covariates (and a small dataset)
- We are currently working on the targeting of a public credit guarantee scheme

Out of sample performance (2010-12 testing dataset)

		Real Status		
		Not Needy	Needy	Total
		Panel A: LPM		
Predicted status	Not Needy	608	208	816
	Needy	447	1358	1805
	Total	1055	1566	2621
	% Correctly Predicted	57.6%	86.7%	75.0%
		Panel B: k-NN		
Predicted status	Not Needy	593	244	837
	Needy	462	1322	1784
	Total	1055	1566	2621
	% Correctly Predicted	56.2%	84.4%	73.0%
		Panel C: Decision tree		
Predicted status	Not Needy	608	232	840
	Needy	447	1334	1781
	Total	1055	1566	2621
	% Correctly Predicted	57.2%	85.1%	74.1%