

Who supports liberal policies? A tale of two referendums in Italy

Leonardo Madio^a, Francesco Principe^b

^a*University of Padova, Department of Economics and Management, Via del Santo,
33, Padova, Italy*

^b*University of Bergamo, Department of Economics, Via dei Caniana, 2, Bergamo, Italy*

Abstract

We leverage a unique dataset at the municipality level in Italy to examine the factors that drive support for two separate referendum campaigns - one on the decriminalization of cannabis cultivation and the other on physician-assisted suicide. Using machine learning techniques, we identify key predictors of support for both referendums, including income, population density, and political leaning of the municipality. Our analysis also highlights that local economic conditions, such as the number of firms, and educational attainment, along with exposure to organized crime, are critical factors driving mobilization in favor of the cannabis referendum. In contrast, support for legalizing assisted suicide is more likely to be explained by religiosity.

Keywords: Referendum, cannabis, liberal policy, direct voting, euthanasia.

JEL codes: D72, I10, K4

1. Introduction

The utilization of referendums as a form of direct democracy is commonly acknowledged as a method of eliciting voters' preferences regarding policies that focus on a single issue. Nonetheless, collecting data at an individual level to understand the factors influencing support for such policies, including abortion, drug legalization, and gun ownership, has proven to be difficult due to privacy concerns.

This paper tackles this challenge by leveraging a unique municipal-level dataset that records the universe of signatures and individual-level data from supporters of

two polarizing referendum campaigns in Italy. Italy holds the second position in terms of the number of referendums conducted, following Switzerland.¹ The first referendum concerns the decriminalization of physician-assisted suicide as, by law, anyone helping a person to commit suicide, even if affected by an incurable disease, can be sentenced to prison. The associated referendum campaign lasted from July to September 2021, with approximately 1.2 million signatures being collected both offline and online. The second referendum concerns the decriminalization of cannabis cultivation for personal use, aiming at easing sanctions on other cannabis-related crimes, with offenders no longer risking prison sentences for selling small amounts of the drug. Its campaign was conducted online in September 2021 and approximately 600.000 signatures were collected.² Despite their vast support, the Italian Constitutional Court deemed both referendum questions inadmissible in 2022. Yet, the data on the universe of signatures in support for these two policies provides a unique opportunity to understand the determinants of political mobilization in support of polarizing topics.³

We employ state-of-art machine-learning techniques to identify the primary area-level factors that predict support for the two referendum campaigns. Machine learning methods perform significantly better than traditional methods in establishing non-linear relationships between variables (Varian, 2014; Mullainathan and Spiess, 2017), though lacking causal inference. After considering different models, we rely on Random Forests (RF) to determine the most important municipal- and province-level features that help explain participation in the referendum.

To identify areas with high support for the referendum campaigns, we calculate a locally-adjusted threshold that measures the number of signatures needed for a suc-

¹Note that this is not unexpected given that the Italian Republic was founded after the approval of a referendum in 1946.

²In order to support the referendum, a voter could use her/his Public Digital Identity System (SPID), which is a simple, fast and secure access key to digital services of local and central administrations (e.g., hospitals, tax agency, as well as to benefit from subsidies).

³The two referendums were not directly backed by the major parties. In our analysis, we show the surprising role of a small liberal, pro-cannabis and pro-euthanasia, party (+Europa) in predicting support for the two referendums.

successful referendum campaign, assuming that each municipality contributes equally based on the number of signatures per inhabitant towards the required 500,000 signatures. We then predict the success of a given campaign based on whether the number of signatures collected in a particular municipality is above (or below) the locally-adjusted threshold. Additionally, we conduct an analysis of the number of signatures per capita collected in each municipality as an alternative specification, showing that our findings hold for both the intensive and extensive margins of referendum support.

Our results indicate that certain common predictors, such as the average disposable income, the number of taxpayers, population density, and the support for left-leaning parties in the previous European election are important predictors of the success of referendum campaigns. However, there are issue-specific dimensions that also play a crucial role in mobilizing voters. For instance, in Southern regions, the referendum on cannabis received greater support than the one on euthanasia. We show that this variation can be attributed to differences in the degree of religiosity between Northern and Southern regions. For instance, we find that religiosity is a critical factor influencing support for the euthanasia referendum, whereas educational attainment and economic factors such as the number of firms drive mobilization in favor of the decriminalization of cannabis cultivation. In addition, the municipality's vulnerability to organized crime also plays a significant role together with the local access to shops selling cannabidiol-based "light" cannabis.

Our study highlights significant polarization within the Italian electorate, as observed across age groups, education levels, and regions of residence, for the two referendum campaigns examined. Our findings offer important insights into the complex dynamics of grassroots mobilization for contentious social issues and their support for more *liberal* policies.

This study contributes to the growing body of research on the determinants of support for single-issue policies. Previous studies, such as [Becker et al. \(2017\)](#) and [Alabrese et al. \(2019\)](#), have investigated the factors that influence support for Brexit using individual and regional data. Our study extends this line of research by utilizing both municipal- and regional-level data to demonstrate the significance of local conditions as a source of heterogeneity in fully mobilizing citizens. Furthermore,

our analysis sheds new light on the determinants of support for cannabis legalization, an issue that has previously been studied using individual-level data (Williams et al., 2016; Palali and van Ours, 2017). In contrast, our approach employs a wide range of area-level indicators to elicit support for legalization.

The remainder of the paper is structured as follows. In Section 2, we describe data and methodology. Section 3 presents the results. In Section 4 we provide concluding remarks.

2. Data and Methodology

We use individual-level data on the universe of the signatures in support of two referendums on cannabis and euthanasia, which were made available by the *Associazione Luca Coscioni*⁴, the main proponent of the two referendum campaigns. We combine them with municipal- and province-level information on demographics, infrastructure, and aggregate socio-economic conditions, provided by the Italian National Institute of Statistics (ISTAT).

2.1. Methods

We use Random Forest (RF) to estimate the importance of area-level indicators in predicting the local support for the two referendums.⁵ As for local predictors, we rely on the following municipal and province-level data, which have been widely used in the most recent literature (e.g., Becker et al. 2017; Carrieri et al. 2021; Micevska 2021):

- *Institutional characteristics*, which capture the political and administrative characteristics of the municipality. These variables include information on the local administrators such as their age, education level, and gender. We also include data on the age and gender of the mayor, as well as the average age of council members. Additionally, we consider the representation of women in

⁴Additional information are available at: <https://www.associazionelucacoscioni.it/> (last accessed on July 18th, 2023).

⁵More details on the methods are provided in the Online Appendix.

the council, which is an important indicator of gender equality in local politics. These variables are sourced from the Italian Ministry of the Interior. Moreover, we also now include additional information on the political-leaning of the municipality, providing information on the election results of the major political parties in the 2019's European elections, which we aggregate as follow: Centre-Left (i.e., Democratic Party and +Europa, a centrist party with a strong pro-euthanasia and pro-cannabis platform), Five Star Movement, and Centre-Right (i.e., Forza Italia, Fratelli d'Italia, and Northern League).

- *Demography and socio-economic characteristics*, include data on internal and external net migration in 2019, population density, composition of the population by age (i.e., 18-34, 35-65, and over-65), number of museums and museum visitors in 2018, university enrollment in 2017, the share of religious marriage ceremonies over the total and the number of all marriages in 2017, the proportion of inhabitants with a secondary school diploma, the share of taxpayers per capita, number of established firms, and access to the Internet (i.e., via ADSL and FTTH locations).
- *Geography*, which includes the altimetric zone, whether a highly urbanized location, whether a coastal city, and whether the municipality is located in a island.

We also include additional variables that are referendum-specific. For example, we include a province-level measure of the number of hospital beds, which serves as a proxy for local access to healthcare, and the number of light cannabis shops per province.⁶ Moreover, for the referendum on cannabis, we also include the and the Index of Permeability of Territories to Organized Crime (IPCO), sourced from Eurispes. Descriptive statistics of all variables are provided in the Online Appendix - Table [A.1](#).

⁶The presence of light cannabis shops may serve as a measure of the social acceptability of the cannabis market and, in turn, influence support for the referendum.

To ensure the success of a referendum campaign in Italy, a minimum of 500,000 valid signatures must be obtained nationally. However, to account for differences in local populations of eligible voters, we construct an equivalent threshold at the municipality level, which measures the number of signatures per inhabitant that ensures that each municipality contributed equally towards the 500,000 valid signatures. Formally, for a given municipality $i \in [1, N]$, where N represents the total number of municipalities in Italy, we have calculated the locally-adjusted threshold (LAT_i) using the following formula:

$$LAT_i = 500,000 \times \frac{POP18_i}{\sum_j^N POP18_j} \quad \forall i \in [0, N],$$

where $POP18$ refers to the population over 18 years of age in each municipality.

To predict the success of a referendum campaign at the municipal level, we then construct a target binary variable D , which takes a value of 0 if the number of locally-collected signatures n_i is lower than the locally-adjusted threshold LAT_i , and 1 if it is greater than or equal to LAT_i . Formally, the algorithm predicts whether the following

$$D_i = \begin{cases} 0, & \text{if } n_i < LAT_i \\ 1, & \text{if } n_i \geq LAT_i \end{cases} .$$

3. Analysis and Results

We first present some descriptive evidence. Figure 1 illustrates the local support for cannabis and euthanasia across all Italian municipalities, weighted by the legally resident population. Notably, there is significant geographic heterogeneity in the support for these two policies. Furthermore, political participation tends to be higher in urban areas. Moreover, Southern Italy saw much greater mobilization in support of the referendum on cannabis relative to the referendum on euthanasia.⁷

Moving on to Figure 2, we observe the number of signatures in support of the

⁷In the Online Appendix (see Figure A.1), we also illustrate the differences in the local support for the two referendums, confirming the North-South divide.

two referendums, divided by gender and age group. The data highlights a generational gap, with the peak number of signatures for cannabis falling within the 18-25 age band and for euthanasia within the 25-34 age band. Support for cannabis also appears to be more divisive in terms of age and gender, with older individuals exhibiting a sharper drop in signatures compared to euthanasia. Women’s participation is only slightly higher than men’s for the euthanasia referendum, whereas there are no substantial differences between men and women for the other age bands. However, the gender gap for cannabis is more pronounced among younger individuals, and it tends to even out with age.

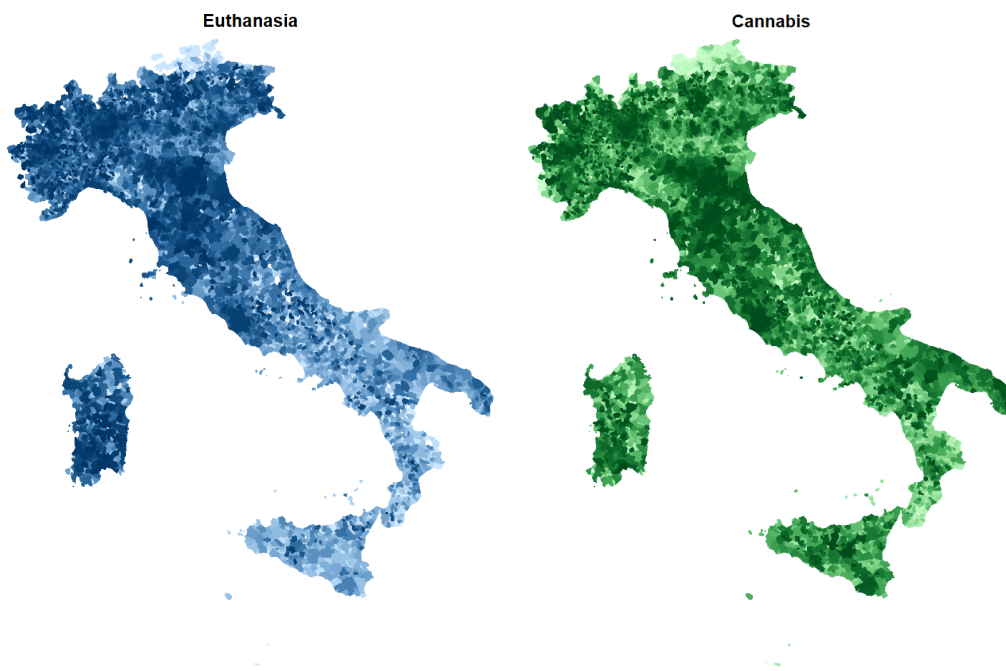


Figure 1: Support for the two referendums.

Note: the figure displays the number of signatures in support for the two referendums, weighted by the local population, across Italian municipalities. Darker areas imply more support for the referendums.

We can now delve into potential explanations for this heterogeneity. We use RF and present the main findings in Figure 3. It displays the relative importance of the 15 most relevant predictors of $D_i = 1$ for any $i \in (0, N]$ in the two referendum




Figure 2: Number of signatures per gender and age band.

Note: the figure displays the number of signatures in support for the two referendums by gender and age band.

campaigns. In other words, we identify the main features that predict whether each municipality meets the locally-adjusted threshold.

We observe some common predictors, such as income (i.e, the proportion of taxpayers over the municipal population and disposable income per capita) and demography (i.e., population density). These represent long-standing fundamental determinants that predict political mobilization in support of any of the two referendums. Furthermore, although the referendums did not receive backing from traditional political parties, we have observed that municipalities with a center-left leaning tend to explain the support for both platforms. It is noteworthy that support for the Five Star Movement, a populist movement, only emerges as a predictor for the euthanasia referendum. However, this finding may also be correlated with the party’s strong presence in Southern Italy and the under-performance of the referendum campaign in these regions. In Figure 4 we present the coefficients of an OLS regression, which highlights that municipalities with a center-right leaning exhibit lower support for both referendums. Notably, local support for +Europa, a centrist party advocating for a liberal, pro-cannabis, and pro-euthanasia political platform, plays a pivotal role in influencing the results.

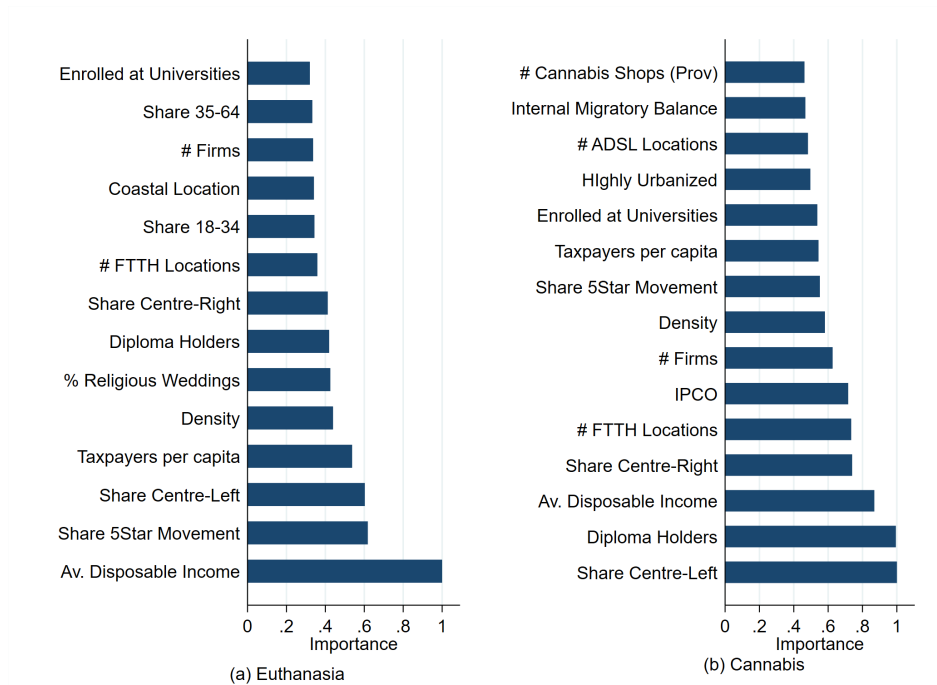


Figure 3: Random Forest: 15 most important features

Note: the figure displays the feature importance for the 15 most important features to predict $D_i = 1$ in the referendum campaigns for Euthanasia (panel a) and Cannabis (panel b). Random forest trained on 70% of observations and tested on the remaining 30%.

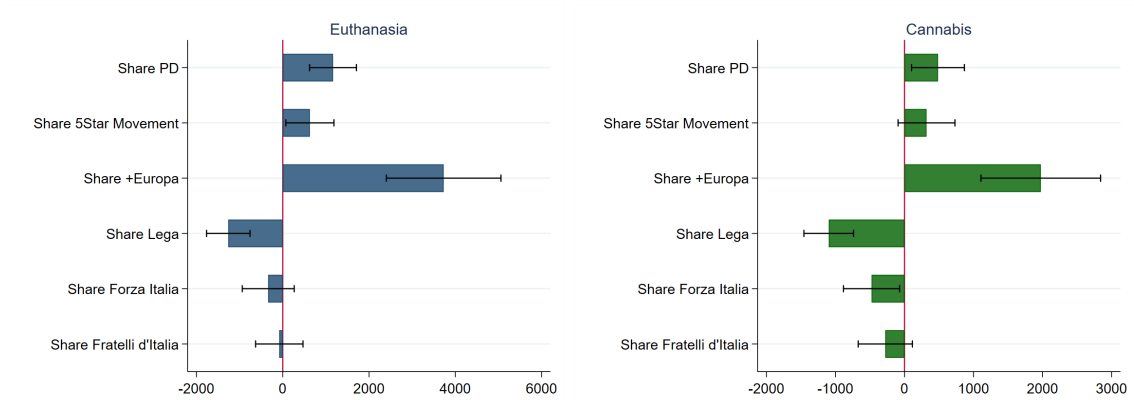


Figure 4: OLS Coefficients

Note: The figure plots the coefficients of an OLS regression where the dependent variable is the number of signature per 100,000 inhabitants at municipality level. Covariates only include the single parties' shares of votes and province fixed effects. Bars display 95% confidence intervals.

The heterogeneity observed in the two maps in Figure 2 can be partially explained by major differences between the two campaigns. Culturally-related features are important predictors for decriminalizing physician-assisted suicide, as the petition was opposed by the Catholic Church. Therefore, the share of religious ceremonies appears as one of the most significant predictors, which is not the case for the cannabis referendum. This difference may partially explain the North-South divide in support of the assisted suicide referendum.

Predictors that are more related to the state of the economy and the education level of the resident population seem to play a major role instead in the support for legalizing cannabis. For instance, the number of companies and the level of education (both in terms of diploma holders and university enrollments) have a stronger impact on support for this referendum than on the euthanasia referendum. We also find that the permeability of the municipality to organized crime plays a role. Results from an OLS regression (Table A.2) in the Online Appendix show that the direction of this effect is negative.

The RF also shows that the familiarity of the local population with the cannabis market, which has undergone an unintended liberalization for what concerns light

cannabis flowers, is key in predicting support for legalization. These shops, whose presence has partially displaced the supply of illegal marijuana (Carrieri et al., 2019), and found support also among non-cannabis consumers for the pain-relief effects (Carrieri et al., 2020) might have set the ground for the social acceptability of the legalization of cannabis.

Interestingly, features related to the quality of the local administration, proxied by information of the city council and the mayor, do not appear among the most important features. Figure A.2 in the Online Appendix summarizes the relative importance of all features included in our analysis.

As an alternative specification, we use RF to predict a continuous variable capturing the number of signatures per capita in each Italian municipality. Results are shown in Figure 5 and confirm qualitatively the main results we have already discussed. Interestingly, access to healthcare, proxied by the number of beds available in the respective province, is also an important determinant in support of euthanasia.

To address concerns about the selection of the best-performing algorithm and to gain a better understanding of the predictive ability of the Random Forest, we analyze a Receiver Operating Characteristic (ROC) curve. The ROC curve is generated by plotting the true positive rate against the false positive rate for various alternative models, including a linear probability model, Probit, LASSO, RIDGE, and ENET.⁸ This graphical representation provides a visual indication of the models' performance. Specifically, if the curve is positioned significantly higher than the diagonal line, it indicates a greater predictive power. The area under the curve (AUC) summarizes the strength of the prediction. Figure 6 shows the ROC curves for the models discussed. For both referendum campaigns, RF is the best-performing model, with a corresponding AUC of 0.97. Concerning other performance measures, the F1-score, which combines precision and recall in a single metric, reaches a value of 0.92 and 0.89 in the euthanasia and cannabis model specifications, respectively.

⁸Details on how these methods perform are provided in the Online Appendix.

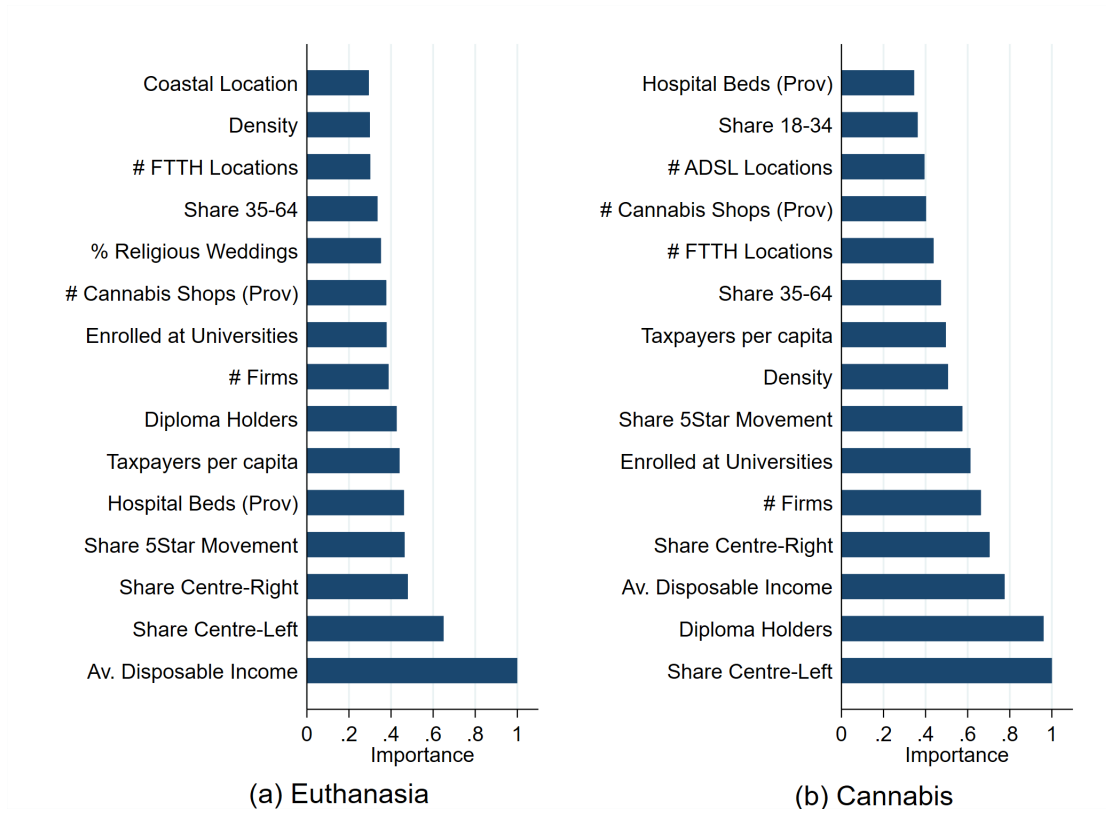


Figure 5: Alternative specification on the number of signatures per capita in each municipality.

Note: the figure displays the feature importance for the 15 most important features to predict the number of signatures per 100,000 inhabitants in each municipality using the RF for the referendum campaigns for Euthanasia (panel a) and Cannabis (panel b). Random forest trained on 70% of observations and tested on the remaining 30%.

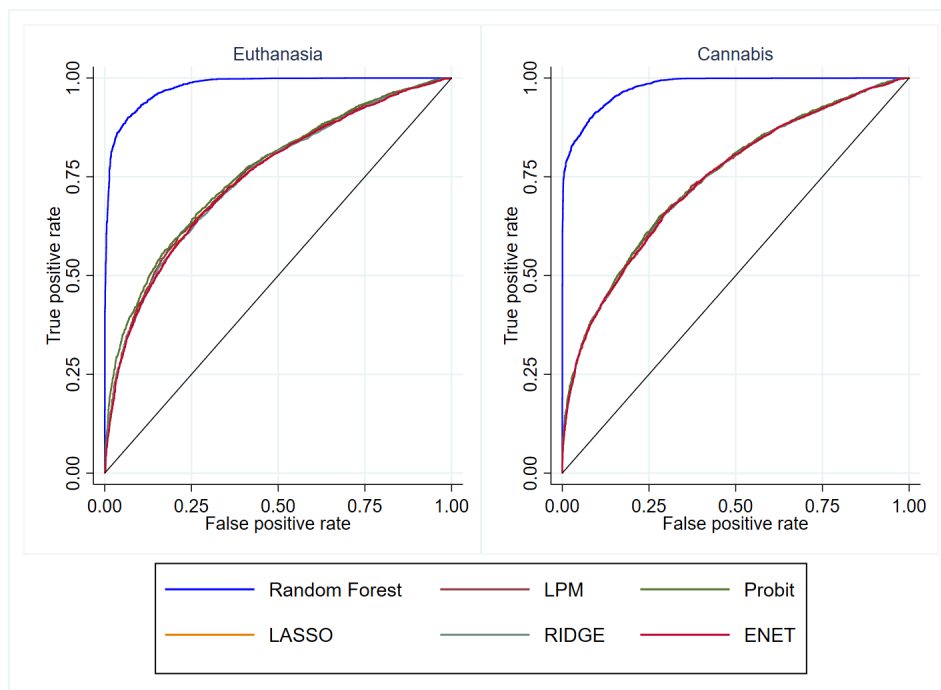


Figure 6: ROC curves for different models

Note: Receiver Operating Characteristics for the different model specifications considered in our analysis. Random Forest (blue curve) outperform any other model.

4. Discussion

Our study provides insight into the primary determinants of support for liberal policies, using a unique dataset of signatures in support of referendums on the decriminalization of cannabis cultivation and physician-assisted suicide. Our results indicate that socio-economic characteristics are significant predictors of political participation and support for more liberal policies. However, while education and the state of the economy are key determinants of support for the cannabis legalization referendum, religiosity plays a more prominent role in support for the euthanasia referendum. Additionally, we found that issue-specific local features, such as the presence of organized crime and the local availability of light cannabis, play a more important role in predicting support for the cannabis referendum, which also receives

more support in left-leaning municipalities. Overall, our findings highlight the existence of an interplay between long-standing primary determinants, such as income and education, and local factors in shaping support for liberal policies.

Acknowledgments. We thank the Editor and an anonymous referee for useful comments that have improved the quality of the manuscript. We also thank Giacomo Battiston, Edoardo Grillo, Giuliano Resce, and Orestis Trompounis for insightful comments on previous versions of the paper. The usual disclaimer applies.

References

- Alabrese, E., Becker, S. O., Fetzer, T., and Novy, D. (2019). Who voted for Brexit? individual and regional data combined. *European Journal of Political Economy*, 56:132–150.
- Becker, S. O., Fetzer, T., and Novy, D. (2017). Who voted for Brexit? a comprehensive district-level analysis. *Economic Policy*, 32(92):601–650.
- Carrieri, V., Lagravinese, R., and Resce, G. (2021). Predicting vaccine hesitancy from area-level indicators: A machine learning approach. *Health Economics*, 30(12):3248–3256.
- Carrieri, V., Madio, L., and Principe, F. (2019). Light cannabis and organized crime: Evidence from (unintended) liberalization in italy. *European Economic Review*, 113:63–76.
- Carrieri, V., Madio, L., and Principe, F. (2020). Do-it-yourself medicine? the impact of light cannabis liberalization on prescription drugs. *Journal of Health Economics*, 74:102371.
- Micevska, M. (2021). Revisiting forced migration: A machine learning perspective. *European Journal of Political Economy*, 70:102044.
- Mullainathan, S. and Spiess, J. (2017). Machine learning: an applied econometric approach. *Journal of Economic Perspectives*, 31(2):87–106.
- Palali, A. and van Ours, J. C. (2017). Cannabis use and support for cannabis legalization. *Empirical Economics*, 53(4):1747–1770.
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2):3–28.
- Williams, J., van Ours, J. C., and Grossman, M. (2016). Attitudes to legalizing cannabis use. *Health Economics*, 25(9):1201–1216.

Online Appendix

Details on machine learning algorithms

Random Forest (RF). Random Forest (RF) combines the power of decision trees and the concept of ensemble learning. It works by constructing multiple decision trees, where each tree is trained on a different subset of the training data. The final prediction is made by aggregating the predictions of individual trees. RF provides a measure of feature importance, allowing therefore to assess the relevance of different variables in the prediction process. The features included in our analysis are detailed in Table A.1 and, for expositional reasons, we provided only an illustration of the 15 most important features in Figure 3 and 5.

It is important to note that in our main analysis (Figure 5), we dealt with a discrete dependent variable. In this case, RF is used for classification tasks. The algorithm constructs an ensemble of decision trees, and the final prediction is determined by majority voting among the individual tree predictions. This voting mechanism allows Random Forests to handle classification problems efficiently and accurately.

In our robustness analysis, whose results are shown in Figure 5, we also use a continuous dependent variables. In this case, RF builds a collection of decision trees, and the final prediction is obtained by averaging the predictions of all the trees. This ensemble approach ensures a smooth and continuous output.

LASSO (Least Absolute Shrinkage and Selection Operator). LASSO is a regularization technique in machine learning that effectively combines feature selection and regularization by introducing a penalty term into the ordinary least squares objective function. By shrinking the coefficients of less important predictors towards zero, LASSO promotes sparsity and facilitates the identification of essential features for enhanced interpretability and predictive performance.

RIDGE regression. RIDGE is a regularization method in machine learning and statistics. It addresses multicollinearity and overfitting by adding a penalty term to the ordinary least squares objective function, which helps to reduce the magnitude of the coefficients. This regularization technique aims to strike a balance between model complexity and prediction accuracy by controlling the amount of shrinkage applied to the coefficients.

ENET (Elastic Net). ENET is a powerful regularization technique that combines the strengths of both LASSO and RIDGE regression. It addresses multicollinearity and performs feature selection by adding a penalty term that encourages sparsity in

the model. ENET allows for simultaneous variable selection and coefficient shrinkage, providing a flexible approach for controlling model complexity and improving prediction accuracy.

The performance for alternative models (RF, LPM, LASSO, RIDGE, and ENET) is evaluated with standardized features.

Predictors	Mean	S.D.
% Women in the City Council	0.323	0.130
% Degree Holders in the City Council	0.289	0.191
Average Age in the City Council	47.81	4.643
Mayor with Degree	0.439	0.496
Mayor's Age	52.97	10.95
Hospital Beds (Prov.)	2735.7	3040.5
# Cannabis Shops (Prov.)	10.26	14.66
Share Age 18-34	0.165	0.0239
Share Age 35-64	0.427	0.0254
Share Age 65+	0.252	0.0490
External Migratory Balance	20.93	222.3
Internal Migratory Balance	0.183	118.7
Population Density	1.942	4.020
Visits Museums	17139.85	362764.6
% Religious Weddings	0.471	0.305
# Weddings	388.4	496.8
Enrolled at Universities	218.1	1396.1
# Museums	0.634	2.551
Diploma Holders	34.09	7.203
# FTTH Locations per 100.000 inhabitants	2075.446	32988.51
# ADSL Locations per 100.000 inhabitants	5119.25	102884.
# Firms	632.0	4601.5
Taxpayers per capita	0.724	0.0729
Average Disposable Income	12951.18	3282.815
IPCO	99.61	3.247
High Urbanization	0.0322	0.177
Coastal Location	0.149	0.356
Island	0.00397	0.0629
Altimetric Zone)	3.058	1.537
Share Movimento 5 Stelle	0.155	0.0905
Share PD	0.186	0.0714
Share Forza Italia	0.0952	0.0608
Share +Europa	0.0239	0.0209
Share Lega	0.387	0.135
Major is a Woman	0.146	0.353

Table A.1: Descriptive Statistics




Figure A.1: Differences in the support for two referendums.

Note: the figure plots differences in the number of signatures per capita between the two referendums. Support is larger for the referendum on euthanasia (respectively, cannabis) in the blue (respectively, green) areas.

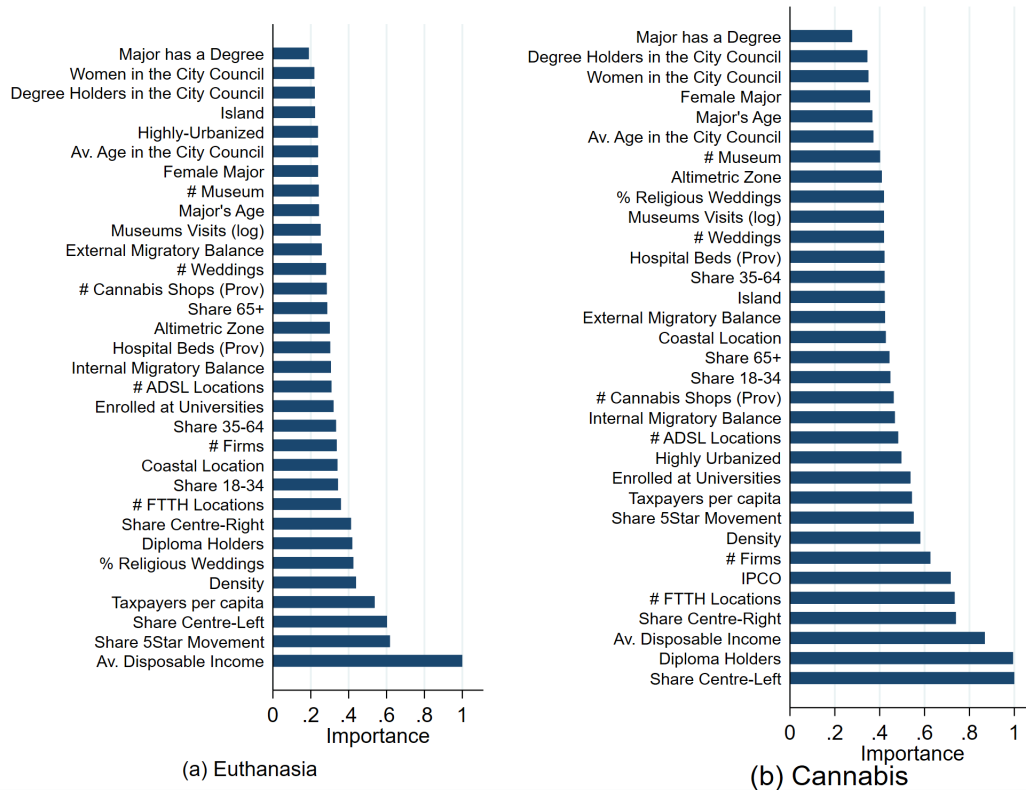


Figure A.2: Random Forest: all features

Note: the figure displays the feature importance for all the features to predict $D_i = 1$ in the referendum campaigns for Euthanasia (panel a) and Cannabis (panel b). Random forest trained on 70% of observations and tested on the remaining 30%.

	Euthanasia	Cannabis
% Women in the City Council	0.008 (1.44)	0.006 (0.97)
% Degree Holders in the City Council	0.004 (0.76)	0.008 (1.24)
Average Age in the City Council	0.000 (0.07)	0.010 (1.67)
Mayor with Degree	-0.006 (-0.58)	-0.007 (-0.63)
Mayor's Age	-0.018*** (-3.45)	-0.016** (-2.84)
Hospital Beds (Prov.)	2.445*** (6.86)	-0.082 (-0.20)
# Cannabis Shops (Prov.)	-2.784*** (-6.20)	0.122 (0.24)
Share Age 18-34	0.004 (0.42)	0.034*** (3.59)
Share Age 35-64	0.018* (2.27)	0.024** (2.91)
Share Age 65+	-0.010 (-0.75)	0.016 (1.13)
External Migratory Balance	-0.052** (-2.59)	-0.107*** (-3.54)
Internal Migratory Balance	0.032*** (4.44)	0.023* (2.30)
Population Density	0.009 (1.62)	0.021*** (3.76)
Visits Museums	0.011 (1.36)	-0.021 (-1.85)
% Religious Weddings	-0.007 (-1.11)	0.001 (0.20)
# Weddings	-0.005 (-1.00)	-0.008 (-1.29)
Enrolled at Universities	0.052 (1.80)	0.066 (1.67)
# Museums	-0.012 (-1.20)	0.026* (2.23)
Diploma Holders	0.080*** (9.36)	0.107*** (12.14)
# FTTH Locations	0.001 (0.18)	-0.014 (-1.55)
# ADSL Locations	0.001 (0.10)	0.003 (0.14)
# Firms	-0.003 (-0.13)	0.051 (1.28)
Taxpayers Per Capita	-0.008 (-0.79)	0.007 (0.65)
Average Disposable Income	0.081*** (5.71)	0.049*** (3.44)
Highly Urbanized	-0.055* (-2.28)	0.009 (0.31)
Coastal Location	-0.057** (-2.97)	-0.059** (-3.16)
Island	-0.120 (-1.47)	-0.048 (-0.58)
Altimetric Zone	0.024*** (4.91)	0.006 (1.08)
Share Movimento5Stelle	0.001 (0.05)	-0.012 (-0.66)
Share Centre-Right	-0.107*** (-4.76)	-0.135*** (-6.20)
Share Centre-Left	-0.005 (-0.36)	0.003 (0.25)
Mayor is Woman	0.016 (1.23)	0.008 (0.52)
IPCO		-0.097* (-2.49)
Constant	0.843*** (8.62)	0.503*** (5.32)
Province FE	YES	YES
Obs	7091	7091
R2	0.375	0.293

Table A.2: Regression Analysis.

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable: $D_i = 1$. All variables are standardized to have mean zero and standard deviation of 1 in order to directly compare the point estimates and their respective magnitude. The model includes province fixed effects. Robust standard errors in parenthesis.