You've Got a (Coarsened Exact) Match!

Non-Parametric Imputation of European Abstainers' Vote

Damien Bol*

Marco Giani†

October 16, 2018

Abstract

There is a long tradition of imputation studies looking at how abstainers would vote if they had to. This is crucial for democracies because when abstainers and voters have different preferences, the electoral outcome ceases to reflect the will of the people. In this paper, we apply a new causal inference method to revisit existing evidence. We impute the vote of abstainers in 15 European countries using Coarsened Exact Matching (CEM). While traditional imputation methods rely on the choice of voters that are, on average, like abstainers and automatically simulate full turnout, CEM only imputes the vote of the abstainers that are exactly like voters, and allows to simulate the electoral outcome under varying levels of turnout, including levels that more credibly simulate compulsory voting. We find that higher turnout would benefit social democratic parties while imposing substantial losses to extreme left and green parties.

Keywords: Turnout, Representation, Elections, Imputation, Coarsened Exact Matching

*Department of Political Economy, King's College London, damien.bol@kcl.ac.uk.
†Department of Government, London School of Economics, m.giani@lse.ac.uk.

Abstention can lead to a major problem of political representation. If abstainers and voters do not have the same preferences, the electoral outcome ceases to represent the will of the people, and only represents the will of voters. This is a well-known problem that has pushed prominent political scientists to make big claims in favor of compulsory voting (Lijphart, 1997). Some studies evaluate the extent to which abstention constitutes a normative issue for political representation. Scholars in the field typically rely on survey analysis in order to infer the likely voting choice of abstainers. In particular, they estimate regression models predicting the party choice of voters based on a set of covariates available in the surveys, and reconstruct the likely choice of abstainers on the basis of the model's prediction (e.g., Bernhagen and Marsh 2007). Hereafter, we call this strategy *standard parametric imputation*.

In this paper, we apply a new causal inference method in order to revisit old evidence. We use Coarsened Exact Matching (CEM) and survey data from 15 European democracies and 30 elections between 1998 and 2014 in order to match survey abstainers with survey voters. The voting choice of abstainers is imputed based on the voting choice of the voters with whom they are matched.[1] In line with the conventional wisdom regarding abstainers' political preference, we focus on the score of left-wing parties, in distinguishing between social democratic, extreme left and green parties.

Using CEM as an imputation strategy instead of standard parametric imputation has two key advantages. Firstly, CEM uses an exact matching algorithm. Therefore, it imputes the voting choice of abstainers based on voters that are *exactly* equal to them on a set of (possibly coarsened) covariates, rather than imputing the voting choice based on voters that are on only equal to them *on average*. This is an important advantage in terms of internal validity, as traditional determinants predict voting choices not only directly, but also in interactions with each other. Secondly, CEM permits to simulate compulsory voting under varying levels of turnout, whereas traditional imputation studies equate compulsory voting with full turnout. This is another important advantage, since, in reality, turnout is never full. In Belgium, where compulsory has

[1]The identification strategy is logically equivalent to comparing observations randomly assigned to different treatments in an experiment. For other studies using similar identification strategies, though in other contexts and using different matching methods (Dehejia and Wahba, 2002).

been used for more than a century, turnout varies between 88 and 95%. Moreover, comparative studies show that citizens with a lower-level of education attainment tend to vote less, even in countries where voting is compulsory (Katz and Levin, 2018). In other words, several covariates continue to predict turnout even in these situations. Simulating what would the electoral outcome under compulsory is thus particularly relevant. Along this line, some recent studies exploit the historical abolishment of compulsory voting in a few European countries to estimate what would have been the voting choice of abstainers (Fewerda, 2014; Bechtel, Handgartner and Schmid, 2015; Miller and Dassonneville, 2016).[2]

Most post-2000 standard parametric imputation studies find that, in Europe and the United States, abstainers and voters have similar political preferences, and thus conclude that abstention does not threaten political representation (Bernhagen and Marsh, 2007; Brunell and DiNardo, 2004; Citrin, Schickler and Sides, 2003).[3] Instead, we find that the imputed voting choice of abstainers substantially differs from the one of voters. Social democratic parties would gain from higher turnout, whereas extreme left and green parties would be worse off. Hence, our result implies that left-wing partisan preferences are differently misrepresented in final electoral outcomes. Counter intuitively, it is the moderate left, rather than the extreme one, that pays the price of low turnout.

Method

In this paper, we revisit existing evidence with the help of a new method. We use the well-known dataset from the European Social Survey (ESS). For a description and justification of the use of the dataset, see appendix A1. Our goal is to impute the voting choice of abstainers, and compare it with the one of voters. We do so by relying on CEM, and follow the step-by-step procedure of Iacus, King and Porro (2012). CEM can be used within an observational study to mimic experimental methods. The basic idea is simple and powerful: matching untreated

[2]Note that these studies find effects pointing at different directions, compulsory voting giving an advantage to left-wing parties or not.

[3]Note that Brunell and DiNardo (2004) use a propensity-score matching method to impute the voting choice of abstainer. Although based on a matching algorithm like CEM, this method is parametric and functions like regression imputation.

observations (here: abstainers) that are exactly similar to treated observations (here: voters) on relevant covariates, thus forming strata of observations.

The first step is to select a set of covariates. We select socioeconomic variables and political interest that are known to be strong predictors of voting behavior (Blais, 2000). We organize them into three matching specifications. In the basic specification, we match voters and abstainers according to their age, gender, ethnic group, household status, highest education attainment, and the subjective feeling of income insecurity.[4] In an augmented specification, we add the main source of income of the household and unemployment status. We acknowledge that the socioeconomic status is not the only determinant of voting behavior. Consistently, in the full specification, we add self-reported political interest. In the full specification, the combination of covariates score gives rise to 32,372 strata, 4,141 of which include observations. Details about the covariates can be found in the appendix A2. The table shows that abstainers are systematically different from voters on most of the selected covariates. There are less educated, younger, more likely to be unemployment and from a minority background, feel less secured economically and less interested in politics.

The second step is to match, for each election and within each country, voters and abstainers that have equal scores on the selected dummy covariates and similar values on the selected continuous covariates, in the spirit of CEM. Concretely, consider the basic matching specification presented above. Abstainers are matched with voters with equal gender, and ethnic minority and household status. Matching requires abstainers and voters to also share a close value on continuous covariates. Age is coarsened by working age categories (15-24; 25-34...). Education attainment and subjective feeling of income insecurity, which range respectively from 1-7 and 1-4, are coarsened using the Scott-Break method that maximizes the trade-off between homogeneous and sufficiently populated strata (Blackwell et al., 2009).

In applying CEM, we identify three groups of individuals: (1) the *certain voters* whose co-

[4]The ESS also provides information about objective income decile, but this variable has many more missing data.

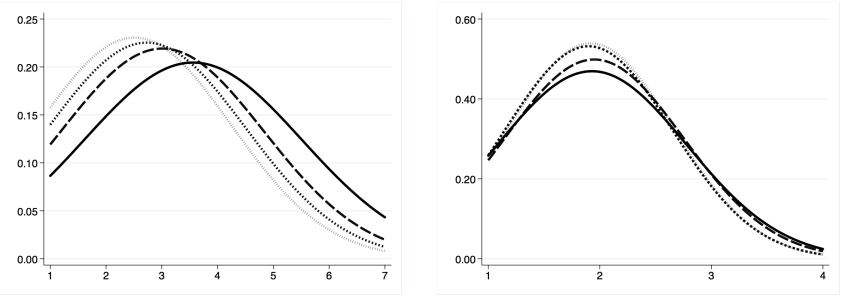

(a) Education (b) Income

Figure 1: Imbalance in education and income between matched individuals and the entire sample. Entire sample (solid line), matched individuals according to the basic (long-dashed), augmented (dashed) and full (dotted) matching specification.

variates are such that they are not matched with any of the abstainers, (2) the *certain abstainers* whose covariates are such that they are not matched with any of the voters, and (3) the *marginal voters*.[5] This last group is of key importance for our analysis, as it includes voters and abstainers who share equal, or close to equal in the case of coarsened continuous variables, scores on the relevant covariates. We think of marginal voters as those who sometimes vote, sometimes not. Their decision cannot be predicted by the selected covariates. As such, we are treating their turnout decision as conditionally random: in a world where voting behavior is determined by these covariates, marginal voters have equal *ex-ante* turnout probability.

In Figure 1, we show the distribution of subjective feeling of income insecurity of marginal voters (that is, matched individuals) for each matching specification (basic, augmented and full). For the sake of comparison, we also show the distribution of these variables for the entire sample. It reveals that marginal voters always have a lower income security and education level than the entire sample. Further, it also shows that this difference increases when we match on more covariates (from basic to full specifications). In other words, Figure 1 shows that the profile of marginal voters become closer to the one of abstainers when we add enough covariates to distinguish them from voters.

[5]We borrow the expression 'marginal voter' to Fowler (2015). In this paper, marginal voters are those who sometimes vote, sometimes not, depending on the circumstance (e.g., if it rains).

The third step is to impute the voting of each abstainer based on the voting choice of their matches.[6] Concretely, the imputation looks like this: suppose that we match one abstainer and three voters, two voting for Party A, and one voting for Party B; we impute the vote choice of the abstainer as being Party A with a probability of 2/3, and Party B with a probability of 1/3. In turn, this means that the vote share of each party remains the same within each stratum of marginal voters. We thus have a vote distribution for abstainers that we can compare to the one of voters. In an attempt to keep results synthetic, we show the difference of vote distribution between voters and abstainers for three groups of parties: (1) social democratic, extreme left, and green parties. The classification of the parties in the countries covered in this paper can be found in the appendix A3.

In total, we are able to match between 42.7% and 82.7% of the abstainers, depending on the specification. It is important to note that the voting choice of certain abstainers (that is, unmatched abstainers) is not imputed. Certain abstainers can be seen as individuals who would not vote even if voting was compulsory. Related to this, the proportion of imputed observations is almost nil in countries where voting is already compulsory (Belgium and Luxembourg) under the full matching specification (3% and 7% respectively). However, the size of the group of certain abstainers depends on the number of covariates upon which we require matching. CEM is a flexible method that permits to set different matching specifications that, in turn, yield different levels of simulated turnout. The more covariates we use for matching voters with abstainers, the lower the number of matched abstainers. However, it is important to keep in mind that when the number of covariates is small, the strata of matched individuals are less homogeneous, which threatens the validity of the imputation (see below).

Results

Table 1 summarizes our results. We compare the party scores among all voters in the survey with the party scores among matched abstainers, across the three matching specifications. In addition, we contrast our results with those obtained when we impute the vote of abstainers from

[6]Note that we do not restrict matching to be one-to-one: abstainers can be matched with more than one voter, and vice-versa.

	Basic	Augmented	Full
% Social democratic parties			
Voters	25.4		
Abstainers (matching)	26.3	26.9	27.7
Abstainers (parametric)	25.8	26.4	25.9
% Extreme left parties			
Voters	5.5		
Abstainers (matching)	4.8	4.6	3.9
Abstainers (parametric)	6.1	6.5	6.0
% Green parties			
Voters	5.5		
Abstainers (matching)	4.9	4.6	4.0
Abstainers (parametric)	5.2	5.2	4.8
% Turnout			
Sample Turnout	80.4		
Compulsory (matching)	96.6	92.7	88.7
Compulsory (parametric)	100	100	100

Covariates in Basic specification: education, age, gender, household status, minority status, feeling of income insecurity. Augmented: add source of income and unemployment. Full: add political interest. For CEM, we match units within each election. For parametric imputation, we use a binary logit regression with election fixed effects.

Table 1: Party choice of voters and abstainers

standard parametric imputation. To do so, we estimate a logit regression predicting the choice of voters for social democratic, extreme left, and green parties with the same covariates than in the three matching specifications. Then, we simulate the vote of abstainers using the predicted probabilities as given by the coefficient estimates.

Firstly, we find evidence that abstainers would be more supportive of the social democratic parties in the 30 elections covered in the data using CEM. The difference between voters and abstainers goes from .8 to 2.3%-points depending on the matching specification. Interestingly, the exact the opposite happens for extreme left and green parties. Abstainers would vote substantially less for these parties (from .8 to 1.6%-points). Higher turnout would thus only increase the score of social democratic parties, at the expenses of other left-wing parties. In appendix A4, we show that these differences are statistically significant at a level of $p < .01$. In contrast, when we impute the voting choice of abstainers with the standard parametric imputation, we only find, consistently with the literature, a small difference between voters and abstainers.

Secondly, we observe differences between matching specifications. The more covariates we use

to match voters and abstainers, the larger the difference in party scores. This result holds true for every party. Intuitively, this means that increasing the number of covariates makes the group of marginal voters smaller, and more homogeneous, which decreases the simulated turnout rate, from 88.7% with the full specification to 96.6% with the basic one. This is an important advantage as the researcher can construct a simulation with varying turnout levels. For example, the turnout rate with the full specification can be seen as credible simulation of compulsory voting. In contrast, the simulated turnout is always the same when we use the parametric imputation method. It is automatically 100% (full turnout).

We also conduct several validation tests of our novel imputation method that discuss in the appendix. Firstly, we show that the validity of the method increases with the number of covariates (A5). The strata are increasingly homogeneous in vote shares when we go from a basic to a full specification. Secondly, we also perform an out-of-sample prediction validation test (A6). We find that CEM is able to correctly recollect the known voting choice of actual voters

Conclusion

Our contribution is threefold. First, from a normative point of view, we find that representation is problematic in European democracies: abstainers would vote differently than voters. This means that, without compulsory voting, the electoral outcome does not adequately represent the will of the people. Second, from a positive point of view, we document that compulsory voting does not change the overall score of left-wing parties, but would affect its composition, favoring social democratic parties at the expenses of extreme left and green parties. Third, from a methodological point of view, we demonstrate how CEM can be used to impute missing values. To our knowledge, our study is the first one to use CEM for that purpose. This method is particularly appealing when missing data are, themselves, at the center of the research question, like in this paper.

References

Bechtel, Michael M., Dominik Handgartner and Lukas Schmid. 2015. "Does Compulsory Voting Increase Support for Leftist Policy?" *American Journal of Political Science* 60(3):752–767.

Bernhagen, Patrick and Michael Marsh. 2007. "The Partisan Effects of Low Turnout: Analyzing Vote Abstention as a Missing Data Problem." *Electoral Studies* 26(3):548–560.

Blackwell, Matthew, Stefano Iacus, Gaery King and Guiseppe Porro. 2009. "CEM: Coarsened Exact Matching in Stata." *The Stata Journal* 9(4):524–546.

Blais, André. 2000. *To Vote or Not to Vote: The Merits and Limits of Rational Choice Theory.* Pittsburgh: University of Pittsburgh Press.

Brunell, Thomas L. and John DiNardo. 2004. "A Propensity Score Reweighting Approach to Estimating the Partisan Effects of Full Turnout in American Presidential Elections." *Political Analysis* 12(1):28–45.

Citrin, Jack, Eric Schickler and John Sides. 2003. "What if Everyone Voted? Simulating the Impact of Increased Turnout in Senate Elections." *American Journal of Political Science* 47(1):75–90.

Dehejia, Rajeev H. and Sadek Wahba. 2002. "Propensity Score Matching Methods for Nonexperimental Causal Studies." *Review of Economics and Statistics* 84(1):151–161.

Fewerda, Jeremy. 2014. "Electoral Consequences of Declining Participation: A Natural Experiment in Austria." *Electoral Studies* 35(3):242–252.

Fowler, Anthony. 2015. "Regular Voters, Marginal Voters and the Electoral Effects of Turnout." *Political Science Research and Methods* 3(2):205–219.

Iacus, Stefano M., Gary King and Giusepppe Porro. 2012. "Causal Inference Without Balance Checking: Coarsened Exact Matching." *Political Analysis* 15(1):1–24.

Katz, Gabriel and Ines Levin. 2018. "A General Model of Abstention Under Compulsory Voting." *Political Science Research and Methods* 6(3):489–508.

Lijphart, Arend. 1997. "Unequal Participation: Democracy's Unresolved Dilemma." *American Political Science Review* 91(1):1–14.

Miller, Peter and Ruth Dassonneville. 2016. "High Turnout in the Low Countries: Partisan Effects of the Abolition of Compulsory Voting in the Netherlands." *Electoral Studies* 44(2):132–143.

Appendix A1: Data

This paper relies on data from the European Election Study (ESS). The ESS regularly conducts face-to-face surveys on representative national samples in European countries. For the sake of comparability, we restrict our analysis to the 15 countries that were members of the European Union before the enlargement of 2004. In each of them, we analyze the two latest national elections available in the data until the release of the 7th round of the ESS. In total, we analyze 30 elections, two per country. It is important to mention that we only include national elections because: (1) the ESS lacks of systematic data for other elections, and (2) voting behavior at the regional and European level tends to follow different logics due to the second-order nature of the elections. Also, as to minimize memory issues, we only analyze the surveys collected right after to each national election. In total, our analysis includes $57,670$ respondents, among which $11,009$ report an abstention.[7] Table A1 provides detailed information about each of the election that we use in the analysis. We report, for each country belonging to EU15, the year of the two elections that we analyze. For each of this election, we use the closest available round of the ESS. We also report, for each election, the real turnout and the sample turnout.

	1st Election				2nd Election			
	Year	N Obs.	Real	Sample	Year	N Obs.	Real	Sample
Austria (**AT**)	2008	1703	78.8	72.0	2013	1483	74.9	73.8
Belgium (**BE**)	2007	1416	89.3	91.1	2014	1514	88.5	89.8
Germany (**DE**)	2008	2395	70.8	78.3	2013	2570	71.5	81.8
Denmark (**DK**)	2007	1432	86.5	93.9	2011	140	87.7	93.5
Spain (**ES**)	2008	1933	73.9	80.3	2011	1609	68.9	72.9
Finland (**FI**)	2007	1719	67.9	80.9	2011	1815	70.4	82.5
France (**FR**)	2007	1690	60.2	74.2	2012	1608	57.2	65.9
United Kingdom (**GB**)	2005	2170	61.4	70.6	2010	2212	65.1	70.5
Greece (**GR**)	2007	1558	74.1	83.6	2009	2528	70.9	78.4
Ireland (**IE**)	2007	1575	67.0	78.7	2011	2297	70.0	71.9
Italy (**IT**)	2001	740	81.4	82.6	2013	725	75.2	73.9
Luxembourg(**LU**)	1999	972	86.5	58.4	2004	1121	91.4	70.2
Netherlands (**NE**)	2010	1722	74.7	82.9	2012	1707	74.3	82.7
Portugal (**PT**)	2009	1541	59.7	65.7	2011	1556	58.1	58.9
Sweden (**SE**)	2010	1556	84.6	90.4	2014	1579	85.8	91.0

One advantage of the ESS is that interviewers make strong efforts not to have a sample

[7] We also exclude non eligible voters.

skewed towards politically-interested individuals. Politically-interested individuals and voters are often over-represented in surveys that are specifically about political issues like the American National Election Study. Table A1 shows that self-reported sample turnout follows closely actual turnout. Although the reported turnout rates are often higher than the actual turnout rates, the differences between the two are small. The average difference is 4.45 points for the first election, and 3.2 points for the second one. Even in Belgium that uses compulsory voting, the proportion of abstainers in the survey is substantial, and very much in line with the proportion of abstainers in reality. All in all, the issue of turnout over-reporting is not severe in our sample.

Looking at each country separately, we observe that the sample turnout often exceeds the real one. This excess sample turnout, however, differs from country to country. In particular, the latter is negligible in Belgium, Italy and Portugal while severe in Finland, Greece, and the Netherlands. In Austria, instead, sample turnout is lower than real turnout, though the two are very close. Note that Luxembourg is a strong outlier, as the actual turnout rate is higher than the one reported in the survey. This probably due to the high proportion of non-citizens living in the country. Note that Luxembourg is a strong outlier, as the actual turnout rate is higher than the one reported in the survey. This is probably due to the high proportion of non-citizens living in the country.

Appendix A2: Covariates

	Voters			Abstainers		
Variables	N. obs	Mean	Std. Dev	N. obs	Mean	Std. Dev
Education (1 low - 7 high)	34,358	3.69	1.97	8,606	3.02	1.78
Gender (0 Male - 1 Female)	46,377	.53	.50	11,285	.55	.50
Age (15 - 114)	46,287	51.60	17.35	11,247	43.85	18.83
Household status (0 no - 1 yes, children)	46,358	.63	.48	11,272	.64	.48
Minority status (0 No - 1 yes)	46,086	.03	.16	11,122	.06	.23
Income feeling (1 positive - 4 negative)	46,149	1.93	.85	11,133	2.19	.88
Unemployment (0 yes - 1 no)	46,219	.74	.44	11,212	.66	.47
Political interest (1 high - 4 low)	46,284	2.40	.90	11,231	3.01	.89

Table A2 summarizes the descriptive statistics of covariates. For the purpose of evaluating imbalance, we provide separate descriptive statistics for voters and abstainers. From a demographic perspective, abstainers are on average younger, more likely to be women and to have ethnic minority background. From a socioeconomic perspective, abstainers are poorer, less educated, more likely to be unemployed, and feel less secure economically speaking. Finally, abstainers' level of political interest is lower.

Appendix A3: Classification of parties

In the analysis, we only include parties that have a parliamentary representation. The *social democratic* parties are: SPÖ (Austria), PS and SPa (Belgium), SD (Denmark), SDP (Finland), PS (France), SPD (Germany), Pasok (Greece), Labour (Ireland), PD (Italy), LSAP (Luxembourg), PvdA (the Netherlands), PS (Portugal), PSOE (Spain), SSDP (Sweden), and Labour (Great Britain).

The *extreme left* parties included in the analysis are: PTB (Belgium), SF (Denmark), VAS (Finland), PCF and FG (France), Linke (Germany), KKE and Syriza (Greece), Rifondazione Comunista (Italy), Déi Lénk (Luxembourg), SP (the Netherlands), BE (Portugal), IU (Spain), V (Sweden).

The *green* parties are: Die Grunen (Austria), Ecolo and Groen (Belgium), Enhlø(Denmark), VIHR (Finland), EELV (France), Die Grunen (Germany), Green Party (Ireland), Girasole (Italy), Déi Gréng (Luxembourg), GL (the Netherlands), MP (Sweden), Green Party (Great Britain).

Appendix A4: Statistical significance

	Basic	Augmented	Full
Δ Social democratic parties			
Matched - Unmatched	.027***	.027***	.030***
Standard error	(.004)	(.004)	(.005)
Δ Extreme left parties			
Matched - Unmatched	-.021***	-.017***	-.022***
Standard error	(.002)	(.002)	(.002)
Δ Green parties			
Matched - Unmatched	-.020***	-.016***	-.021***
Standard error	(.002)	(.002)	(.002)

N = 57,676. Covariates in Basic specification: education, age, gender, household status, minority status, feeling of income insecurity. Augmented: add source of income and unemployment. Full: add political interest. For CEM, we match units within each election. For parametric imputation, we use a binary logit regression with election fixed effects. *** $p < .01$

In Table A4, we test whether the voting behavior of marginal voters (matched voters/abstainers) differs significantly from those of certain voters (unmatched voters), using a standard two sided $z-$test. The reported coefficients refer to the difference between the score of each party in the two groups. For instance, the number .027 in the first entry means that the score of the social democratic party among marginal voters is higher than the one among certain voters by 2.7%-points.

All in all, the outcomes presented here confirm those in Table 1 in the paper: social democratic parties are stronger among marginal voters than among certain voters, while the opposite holds true for the extreme left and green parties. Standard errors are relatively low, and each outcome is statistically significant at a level of $p < .01$. Alike what we observed in Table 1, matching on political interest makes the distance between marginal voters and certain voters greater.

Appendix A5: Number of covariates and validity

We calculate, in each stratum of matched abstainers, the proportion of votes for social democratic parties (see Figure A6).[8] The proportions are either very close to 0 or 1. This means that matching on these covariates strongly discriminates supporters of these parties. Of all strata with observations, 42.9% are composed of either exactly 0 or 100% of social-democratic voters when we use the basic matching specification. Also, we show that the strata are increasingly homogeneous in vote shares depending on the number of covariates (from basic to full specification). In other words, the more covariates included, the more valid the imputation is. This time, 67.7% of all strata are composed of either 0 or 100% of social democratic voters.

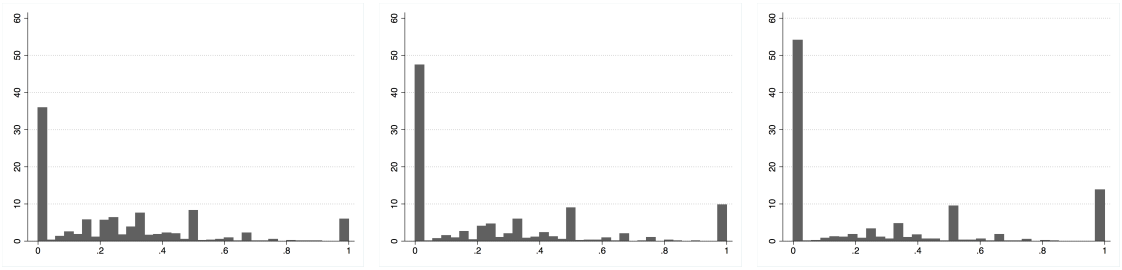

(a) Basic specification (b) Augmented specification (c) Full specification

[8]The results are similar for extreme left and green parties. However, since there are many less voters for these two groups of parties the figure is not as easy to read.

Appendix A6: Out-of-sample prediction validation

We perform an out-of sample prediction validation test. We focus on the respondents for which we have information regarding voting choice (i.e., we exclude abstainers) and evaluate how accurately CEM imputes a voting choice for them. To do so, we randomly split voters into two groups: a test group (N = 100 in each election) and a training group (N=sample size−100 in each election). We remove the voting choice of the test group and treat them as abstainers. We then perform the CEM imputation as described above, and compare the imputed voting choice of individuals of the test group to their actual voting choice. In addition, we perform the same test using the standard parametric imputation described above (prediction from logit regressions). We repeat this training exercise 100 times with 100 different random samples.

	Real	Matching	Parametric
% Social democratic parties			
	25.4	25.9	25.0
% Extreme left parties			
	5.5	5.1	6.1
% Green parties			
	5.5	5.0	5.9

N = 57,676. Entries are real voting choices, or CEM and parametric predictions, in the test group. Based on augmented specification. Coefficients average out 1,000 trials.

From Table A6, we observe that the average difference for social democratic, extreme left and green parties is relatively small when we use CEM as an imputation method. The difference between imputed and actual voting distribution is never larger than 1%-point. However, Table A6 also reveals that a standard parametric imputation performs, on average, as good as CEM in this training exercise. This is not surprising given that a key advantage of CEM is that that it only imputes a value for the missing untreated observations that are matched to at least one treated non-missing observation. In this training exercise, the treated (training group) and untreated observations (test group) are necessarily the same because they are random subsamples of the same larger sample of voters. Hence, CEM matches a large fraction of untreated observations (86.8% on average), and imputes a value for each of them. However, even in this situation that is not in its advantage, the method performs at least as good as standard parametric imputation. This suggests that an alternative to our method would be to use CEM as

a preprocessing technique to reduce a sample to observations that are exactly the same on key covariates (as shown by Ho et al 2007). The imputation resulting from logit regression prediction ran on this reduced sample would be similar to those given by CEM. Hence, this alternative is more labor-intensive, as it necessitates the mobilization of two different methods instead of one.

Ho, Daniel, Kosuke Imai, Gary King, and Elizabeth Stuart. 2007. "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference." *Political Analysis* 15(3): 199–236.