Education and the geography of Brexit

Robert Calvert Jump[*] Jo Michell[†]

November 18, 2020

Abstract

While it is well established that educational attainment is highly correlated with Brexit voting patterns, the predictive capacity of education has attracted less attention. Using full-sample and split-sample exercises, this paper demonstrates that educational attainment alone can correctly classify over 90% of local authorities by voting outcome in the 2016 referendum to leave the EU, depending on the prediction model and classification method used. This illustrates the importance of education as a key factor in the geography of Brexit.

Keywords: Brexit, Education, Geography, Polarisation.

Acknowledgements: We thank Thiemo Fetzer, four anonymous referees and the editor for comments. Any remaining errors are the responsibility of the authors.

[*]Institute of Political Economy, Governance, Finance and Accountability, University of Greenwich. Email: r.g.calvertjump@greenwich.ac.uk.

[†]Department of Accounting, Economics and Finance, University of the West of England. Email: jo.michell@uwe.ac.uk.

1 Introduction

In October 2017, the Conservative MP for Daventry wrote to the vice chancellors of British universities to ask for the names of any professors involved in the teaching of European affairs, "with particular reference to Brexit" (Fazackerley, 2017). Discussing this letter on the BBC's Sunday Politics programme, Barry Sheerman, the Labour MP for Huddersfield, argued that, "The truth is that when you look at who voted to Remain, most of them were the better-educated people in our country". This observation provoked widespread debate, with at least one Conservative MP accusing Sheerman of snobbery (Baynes, 2017).

Since this short-lived media storm, and others like it, careful empirical work has established educational attainment as one of the strongest correlates of Brexit vote shares across British voting districts (e.g. Becker et al., 2017; Beecham et al., 2018; Zhang, 2018), using regression models with and without control variables.[1] In particular, Becker et al. (2017) demonstrate that R^2 values between measures of educational attainment and vote shares across British voting districts are greater than 60%, with demographic variables dominating other predictors by this metric. In comparison, variables proxying the exposure of voting districts to the European Union, including measures of immigration and trade exposure, collectively explain less than 50% of the variance in vote shares (Becker et al., 2017).

While the effect sizes and measures of fit reported in the existing literature on the electoral geography of Brexit offer important insights into the relationship between educational attainment and Brexit vote shares, they offer little insight into the ability of education levels to predict which areas were more likely to vote Leave and which areas were more likely to vote Remain. Consider figure 1, for example, which plots the relationship between educational attainment and Leave voting in English and Welsh local authorities. The left panel plots Leave vote shares against adult population shares educated to degree-level or above, overlaid with the estimated conditional expectation function from a simple linear regression model. The right panel plots Leave vote shares against the predicted Leave vote shares from the same linear regression model.

The R^2 metric – the coefficient of determination – which the existing literature focuses on, is calculated using squared prediction errors. This means that a small number of large prediction errors can substantially reduce the coefficient of determination, even if the majority of prediction errors are small. In the left panel of figure 1, for example, the large prediction errors associated with Remain-voting authorities in the lower right quadrant will significantly reduce the R^2, even though the prediction errors are much lower in the rest of the distribution. Moreover, despite the fact that the prediction errors are relatively large for authorities with high Remain shares, the linear regression model still gives the correct

[1] Other papers exploring the determinants of Brexit include Goodwin & Heath (2016), Harris & Charlton (2016), Hobolt (2016), Manley et al. (2017), Colantone & Stanig (2018), Fetzer (2019), Johnston et al. (2018), and Lee et al. (2018).

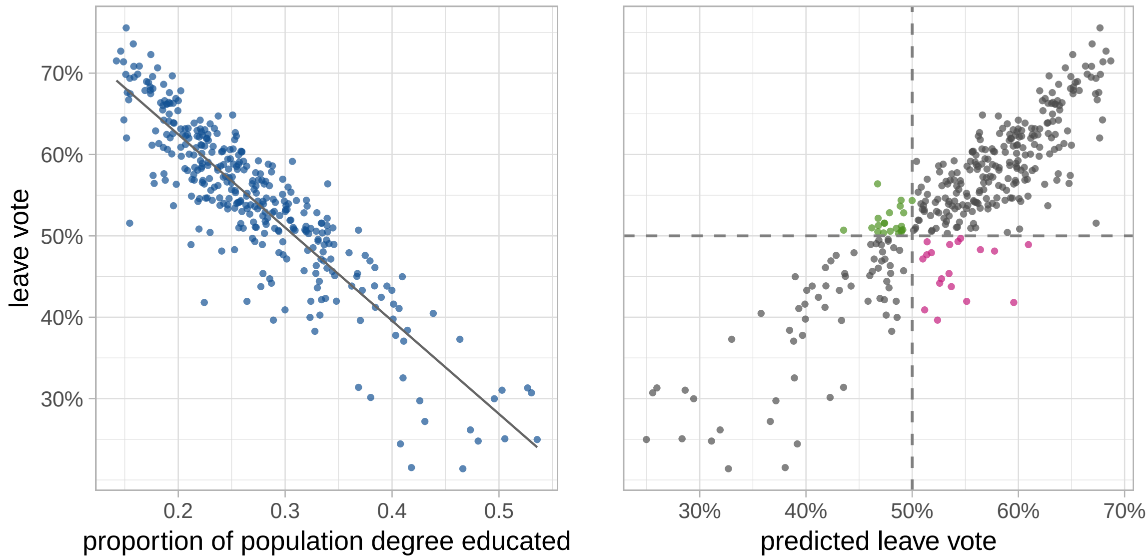

Figure 1: Predicting local authority Brexit referendum voting in England and Wales using a simple linear model. City of London is not shown in the scatter because it is a substantial outlier with 68% of the population degree educated.

classification for these authorities, in the sense that the predicted referendum result is still Remain.

The only authorities for which the regression model gives the wrong classification – a Leave authority is predicted to vote Remain or vice versa – are highlighted in green and pink in the right hand panel of figure 1. A natural measure of predictive accuracy in this context is the percentage of local authorities for which the model gives the correct classification. The coefficient of determination is not very informative in this respect, since it penalises large prediction errors in correctly classified authorities while ignoring small prediction errors in incorrectly classified authorities. For the linear regression model shown in figure 1, 89% of areas are correctly classified, but the coefficient of determination is only 79%.

A classification exercise of this kind has not yet been applied to the electoral geography of Brexit. To fill this gap in the literature, we use probit models and linear discriminant analysis to classify British local authorities into Leave and Remain areas using demographic predictors. We demonstrate that for England and Wales in particular, average educational attainment is a powerful predictor of the voting direction of local authorities. Our results complement the existing results in Becker et al. (2017), Beecham et al. (2018), and Zhang (2018), and the classification exercises on individual voting behaviour reported in Alabrese et al. (2018). Our results emphasise the importance of education as a key predictor of the geography of Brexit, and support the wider literature linking education and Euroscepticism (e.g. Lubbers & Scheepers, 2010; Hakhverdian et al., 2013).

2 Data

The data are comprised of Leave vote shares and a set of demographic characteristics, observed at local authority level in Great Britain. The latter include the population shares born in the UK, identifying as ethnically white, identifying as male, median age, and the working-age population share with a 'lower supervisory and technical', 'semi-routine', or 'routine' occupation. These are all standard predictors in the literature. The measure of educational attainment is the adult population share educated to degree-level or above. Voting data for the 2016 referendum are taken from the Electoral Commission and the demographic variables are taken from the 2011 census.

We confine our analysis to Britain because data for Northern Ireland are unavailable at local authority level. We also present separate results for England and Wales alongside results for the whole of Britain, because every local authority in Scotland voted Remain. Scotland as a whole is therefore an outlier compared with Britain as a whole, so it is worth examining whether or not its inclusion affects the results. We control for regional effects at the NUTS 1 level, which comprises large non-administrative regions in England (e.g. 'South East', 'West Midlands', 'Yorkshire and The Humber') plus Wales and Scotland. Finally, while it would be preferable to use demographic data from 2016, rather than census data from 2011, the Annual Population Survey from which these data could be drawn does not include occupational data using the National Statistics Socioeconomic Classification (NS-SEC). We discuss this further in section 4.

3 Classification Approach

We use probit models and linear discriminant analysis to classify local authorities into 'Leave' and 'Remain' areas. Section 3.1 outlines the probit model for binary and fractional data, section 3.2 outlines predictive linear discriminant analysis, and section 3.3 outlines our full-sample and split-sample classification strategies.

3.1 Probit models

We use probit models for a binary dependent variable of the form,

$$\Pr[Y = 1|\mathbf{x}] = \Phi(\mathbf{x}\beta + u), \tag{1}$$

and probit models for a fractional dependent variable of the form,

$$\mathrm{E}[S|\mathbf{x}] = \Phi(\mathbf{x}\beta + u), \tag{2}$$

where Φ is the standard Normal cumulative distribution function, $\mathbf{x}$ is a vector of demographic regressors and u is a fixed or random regional effect. In (1), the dependent variable $Y = 1$ if a local authority voted Leave, and 0 otherwise. This model yields a set of predicted probabilities $\hat{\Pr}$. In (2), the dependent variable S is the share of Leave votes in valid votes. This model yields a set of predicted shares $\hat{s}$. The standard probit model is estimated using maximum likelihood, while the fractional probit model is estimated using quasi-maximum likelihood.[2]

3.2 Linear discriminant analysis

The probit model in (1) estimates the probability of class membership, where the set of classes {Leave, Remain} is used to define the binary dependent variable Y, by fitting a sigmoid function of the regressors in $\mathbf{x}$. In comparison, discriminant analysis uses Bayes' theorem to estimate the the probability of class membership,

$$\Pr[Y = 1|\mathbf{x}] = \frac{f(\mathbf{x}|Y = 1)\Pr[Y = 1]}{f(\mathbf{x}|Y = 0)\Pr[Y = 0] + f(\mathbf{x}|Y = 1]\Pr[Y = 1]}, \tag{3}$$

where $f(\mathbf{x}|Y = y)$ are the probability densities of the regressors in $\mathbf{x}$ conditional on the referendum result, and $\Pr[Y = y]$ are the unconditional probabilities of the referendum result. We use linear discriminant analysis, which specifies $f(\mathbf{x}|Y = y)$ as multivariate normal with identical covariance matrices for the two classes.[3] As with the probit model in (1), the model in (3) yields a set of predicted probabilities $\hat{\Pr}$.

3.3 Classification strategies

We are interested in the ability of the demographic variables in $\mathbf{x}$ to classify British local authorities into Leave and Remain voting areas, using observations on the 380 local authorities that existed in Britain as of June 2016. Given these data and the models described in sections 3.1 and 3.2, there are two basic approaches to measuring classification success: full-sample and split-sample. Full-sample classification estimates the models on the entire sample, then defines a correctly classified Leave area as one for which $y = 1$ and $\hat{\Pr} > 0.5$ (or $s > 0.5$ and $\hat{s} > 0.5$), and a correctly classified Remain area as one for which $y = 0$ and $\hat{\Pr} \leq 0.5$ (or $s \leq 0.5$ and $\hat{s} \leq 0.5$). The percentage of correctly classified Leave areas is then the number of local authorities for which $y = 1$ and $\hat{\Pr} > 0.5$ (or $s > 0.5$ and $\hat{s} > 0.5$) divided by the number of local authorities for which $y = 1$. The percentage of correctly classified Remain areas, and the percentage of correctly classified areas overall, are calculated in a similar way.

[2] The models were estimated in Stata and R; code and data are available on request.

[3] Quadratic discriminant analysis, in which the covariance matrices are allowed to differ between classes, yielded similar results. See e.g. Härdle & Simar (2015) for a discussion of discriminant analysis.

Full-sample classification measures the ability of a model estimated on a sample A to classify observations in the same sample A.[4] This approach risks over-fitting, which occurs when the model is successful at classifying observations within the sample A but performs very poorly when used to classify observations in a second sample B. In this sense, the full-sample percentages of correctly classified areas might over-estimate the accuracy of the models in sections 3.1 and 3.2. An alternative measure of classification accuracy is split-sample validation, in which the sample A is split into two sub-samples A_1 and A_2. The models are estimated using observations in A_1 and used to classify observations in A_2. As this approximates a situation in which the model is estimated using the full sample A and then used to classify observations in a second sample B, it reduces the risk of over-fitting. The split-sample percentages of correctly classified areas should therefore provide a better measure of model accuracy.

We use two methods to compute split-sample classification success: method (A) estimates the model on all local authorities from NUTS 1 regions outside the South East, and uses this model to predict voting outcomes for authorities inside the South East. We use the South East because its Leave share of 51.78% was very similar to the overall Leave share in England and Wales of 51.89%, and it has the largest electorate of the NUTS 1 regions. Method (B) uses a repeated split-sample approach in which the model is estimated on two thirds of the sample, and this model is used to predict voting outcomes on the remaining observations. This is similar to the approach used in Alabrese et al. (2018), and we report classification results averaged over 10,000 randomly chosen estimation samples.

4 Results

Tables 1 and 2 present full-sample and split-sample classification results using standard, fixed effects and random effects probit models, a fixed effects fractional probit model, and linear discriminant analysis. In the full-sample results, all of the models that use education as their sole predictor correctly classify over 90% of English and Welsh local authorities that voted Leave in the 2016 referendum, and – with the exception of dicriminant analysis – over 75% of English and Welsh local authorities that voted Remain. Moreover, these models are highly competitive with the models that use all of the demographic variables to predict voting outcomes. The fixed effects probit model with education, for example, correctly classifies 96.2% of English and Welsh Leave-voting authorities and 80% of English and Welsh Remain-voting authorities, or 253 of the 263 Leave-voting areas and 68 of the 85 Remain-voting areas. In comparison, the fixed effects probit model using all of the demographic variables correctly classifies 97% of Leave-voting authorities and 82.4% of Remain-voting authorities, or 255 and 70 correctly classified Leave and Remain areas respectively. As such, the marginal gain from the full set of demographic variables is two correctly classified Leave

[4] Full-sample classification is known as 'resubstitution classification' in the discriminant analysis literature.

Table 1: Full-sample results: per cent of local authorities correctly classified

	England and Wales			Great Britain		
	All	Leave	Remain	All	Leave	Remain
Standard probit:						
Education only	90.8	95.4	76.5	82.9	92.4	61.5
All demography	92.2	97.0	77.6	88.2	94.3	74.4
All demog. minus educ.	85.9	93.5	62.4	77.6	92.4	44.4
Fixed effects probit:						
Education only	92.2	96.2	80.0	84.2	90.1	70.9
All demography	93.4	97.0	82.4	91.1	93.9	84.6
All demog. minus educ.	89.9	95.4	72.9	81.3	89.4	63.2
Fixed effects fractional probit:						
Education only	91.4	93.2	85.9	84.7	87.5	78.6
All demography	93.1	95.4	85.9	88.4	91.3	82.1
All demog. minus educ.	89.4	92.8	78.8	81.8	85.2	74.4
Random effects probit:						
Education only	92.5	97.0	78.8	92.9	96.6	84.6
All demography	93.1	97.3	80.0	94.2	97.0	88.0
All demog. minus educ.	89.1	95.4	69.4	90.5	95.8	78.6
Linear discriminant analysis:						
Education only	89.1	96.6	65.9	84.2	95.1	59.8
All demography	91.7	98.9	69.4	86.3	98.1	59.8
All demog. minus educ.	85.3	94.3	57.6	78.9	94.7	43.6

Notes: The standard probit models do not include fixed or random regional effects. The random effects models are estimated using `meprobit` in Stata. The models for the full-sample results are estimated using all local authorities in the sample.

areas and two correctly classified Remain areas.

While the full-sample predictive power of the set of demographic variables excluding education is competitive, there is no model in table 1 for which the set of demographic variables excluding education is a superior predictor to education by itself. In fact, in some cases the set of demographic variables excluding education is quite a poor predictor. This is most obviously the case in linear discriminant analysis, in which the set of demographic variables excluding education correctly classifies less than 60% of Remain authorities in England and Wales. It is also the case in the random effects probit model, which correctly classifies 69.4% of Remain authorities in England and Wales when the the set of demographic

Table 2: Split-sample results: per cent of local authorities correctly classified

	England and Wales			Great Britain		
	All	Leave	Remain	All	Leave	Remain
Standard probit (A):						
Education only	91.0	88.4	95.8	88.1	83.7	95.8
All demography	88.1	90.7	83.3	83.6	90.7	70.8
All demog. minus educ.	80.6	81.4	79.2	67.2	69.8	62.5
Standard probit (B):						
Education only	90.0	94.6	75.9	82.9	92.5	61.4
All demography	91.4	96.1	77.1	86.5	93.0	72.2
All demog. minus educ.	86.1	93.6	63.0	76.9	91.5	44.3
Linear discriminant analysis (A):						
Education only	89.6	93.0	83.3	89.6	88.4	91.7
All demography	80.6	95.3	54.2	82.1	95.3	58.3
All demog. minus educ.	74.6	86.0	54.2	70.1	81.4	50.0
Linear discriminant analysis (B):						
Education only	88.9	96.7	65.1	83.6	94.5	59.4
All demography	90.7	98.5	66.5	85.8	97.5	59.6
All demog. minus educ.	85.3	94.3	57.7	77.4	93.3	42.1

Notes: The models do not include fixed or random regional effects. The models for the split-sample (A) results are estimated on all NUTS 1 regions other than the South East, which is used for the classification exercise. The models for the split-sample (B) results are estimated on randomly chosen subsets with the remaining observations used for the classification exercise. The reported results are averages over 10,000 estimations.

variables excluding education is used, compared to 78.8% when education is used by itself. This is quite a remarkable result: using our measures of classification success, local authority population shares educated to degree-level or above contain more information on 2016 vote shares than the population shares born in the UK, identifying as ethnically white, identifying as male, median age, and those with a blue-collar occupation.

Importantly, the measures of split-sample classification success in table 2 are competitive with the full-sample results in table 1. In most cases the percentages of correctly classified areas are lower in table 2 compared to table 1, as expected, but in some cases the percentage of correctly classified Remain areas is actually higher when using the split-sample measures of success than when using the full-sample measures. The conclusions regarding the relative superiority of education as a predictor still hold, as do the conclusions regarding the relative inferiority of linear discriminant analysis. Finally, we note that classification success is

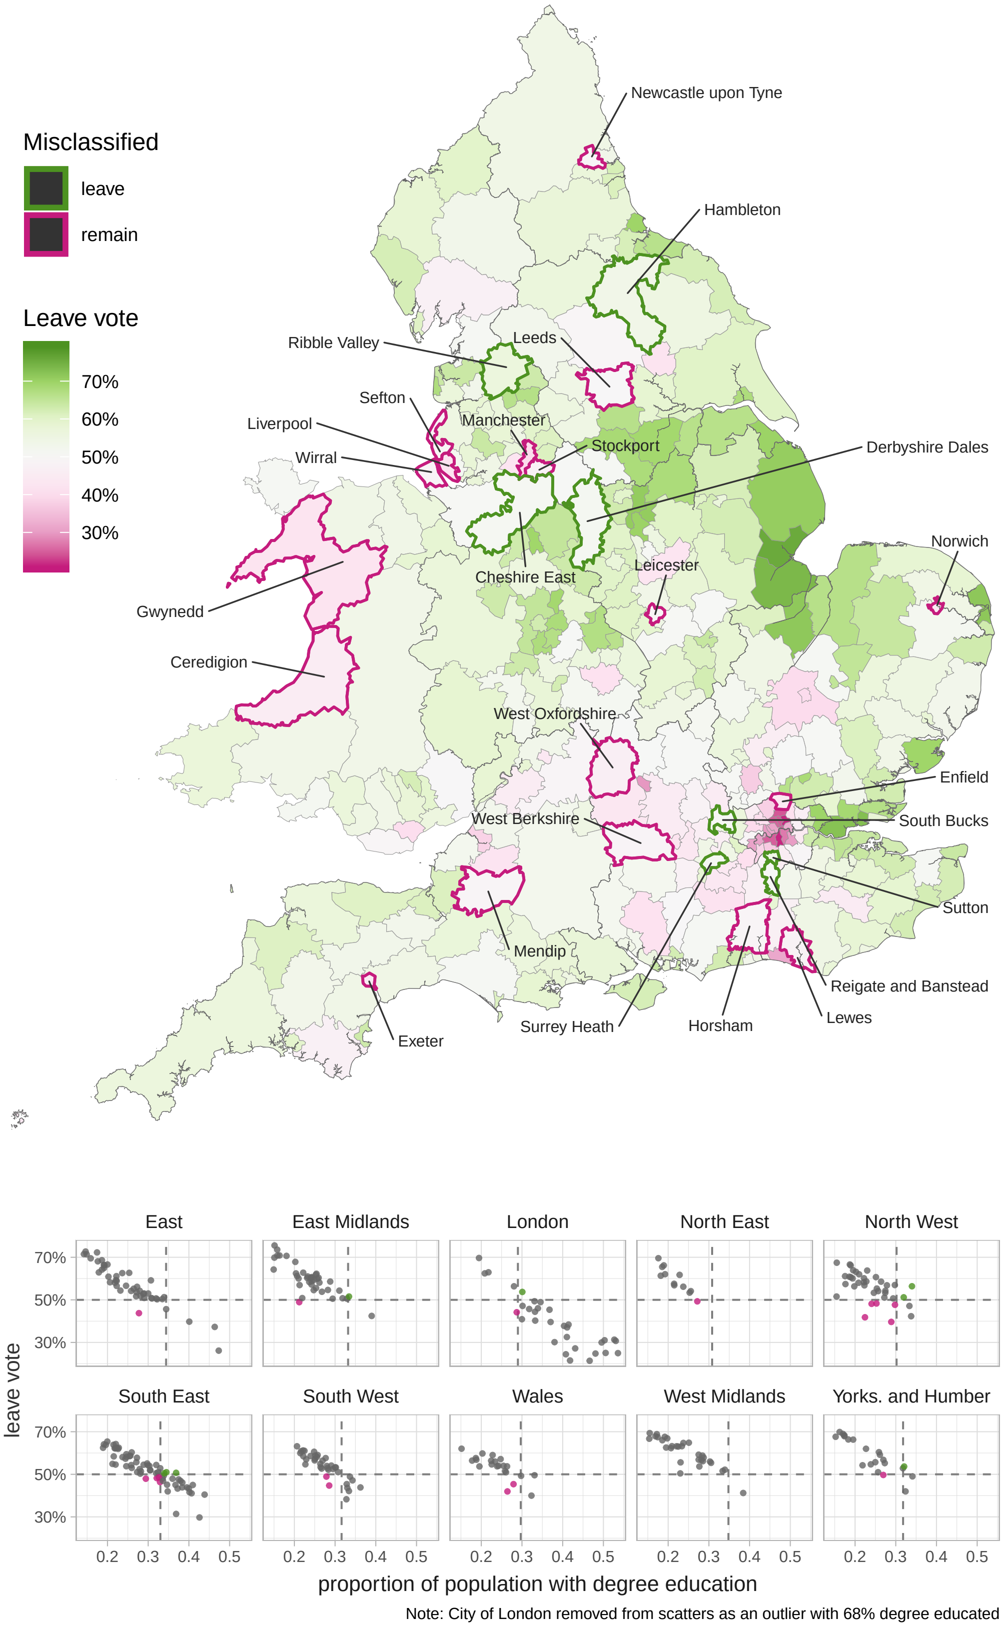

Figure 2: Map of Leave vote shares highlighting incorrectly classified local authorities using random effects probit model. The scatter plots show how regional education thresholds identified by the model are used to classify local authorities. Incorrectly classified local authorities are highlighted.

Table 3: Correlation matrix of demographic variables, England and Wales. The correlation matrix for the whole of Britain is almost identical.

	median age	birthplace	education	ethnicity	gender	occup.
median age	1.00					
birthplace	0.74	1.00				
education	−0.22	−0.54	1.00			
ethnicity	0.77	0.92	−0.34	1.00		
gender	−0.44	−0.35	0.21	−0.32	1.00	
occupation	0.22	0.52	−0.90	0.38	−0.09	1.00

generally lower when we include Scotland in the analysis. In the standard probit and linear discriminant analysis using the set of demographic variables excluding education, for example, the full-sample percentages of correctly classified Remain areas are less than 45%. This suggests that the relationship between voting and education in the 2016 referendum follows a different distribution in Scotland – in which every local authority voted Remain – compared with England and Wales.

The local authorities incorrectly classified using the full-sample method by the random effects probit model are highlighted in figure 2 using a choropleth and regional scatters. The latter show the regional thresholds on educational attainment identified by the probit model: for any given proportion of the population educated to degree-level or above, the Midlands, the East and the South East are more likely to vote Leave, while London, Wales and the North East are less likely to. Misclassified Remain-voting areas include those with large universities, suggesting a possible influence of education that is not captured by the proportion of degree-educated adults. On the other hand, most of the other incorrectly classified Remain areas are geographically contiguous with areas with particularly strong Remain votes. There is, therefore, clearly information being ignored by our models. Nonetheless, the results presented here illustrate the remarkable capacity of educational attainment to predict the geographic distribution of Britain's vote to leave the EU. This complements, and is consistent with, the existing literature on the unconditional and conditional association between education and the geography of Brexit discussed in Becker et al. (2017) Alabrese et al. (2018), Beecham et al. (2018), and Zhang (2018).

Further results and robustness checks are reported in an online appendix. These include estimated partial effects, plus discriminant function coefficients from the linear discriminant analysis and full-sample classification results for each model using observations on average educational attainment from the 2016 Annual Population Survey, instead of the 2011 census. In addition, we report test results for residual spatial correlation in the standard and fixed effects probit models, and demonstrate that this is not a material issue once regional effects are accounted for.

Table 4: Results of fixed effects fractional probit model of demographic variables on Leave voting, England and Wales.

	education only	all demog. vars.	all demog. vars. except education
education	-1.07^{**}	-1.31^{**}	
median age		0.01^{**}	0.00^{**}
birthplace		-0.15	0.52^{**}
ethnicity		0.15^{**}	-0.28^{**}
gender		1.44^{**}	0.03
occupation		-0.21^{*}	0.98^{**}

Notes: * $p < 0.05$, ** $p < 0.01$. The reported results are average partial effects, i.e. $\mathbb{E}[dy/dx]$, for the fractional probit model with regional fixed effects.

5 Discussion

How should these results be interpreted? The narratives that have been proposed to explain the vote can be divided into those emphasising cultural divergence between socially liberal Remain voters and socially conservative Leave voters (e.g. Goodwin & Heath, 2016), and those emphasising economic drivers such as inequality, austerity, and the effects of globalisation (e.g. Fetzer, 2019). Educational attainment can be invoked in favour of either set of narratives, as it is closely correlated with both social attitudes and economic success. A third possibility is that voters with less formal education are more likely to change their voting intentions in response to political campaigns or other signals; this argument has been invoked by Vasso Ioannidou in her discussion of Becker et al. (2017).

Isolating the effects of different channels of influence from education to Brexit voting patterns is made difficult by the presence of confounding variables. Table 3 illustrates this problem. Perhaps surprisingly, population shares educated to degree-level or higher are not strongly correlated with median age. However, average education levels are strongly correlated with working-age population shares with a 'lower supervisory and technical', 'semi-routine', or 'routine' occupation. Although this association makes it difficult to provide a straightforward interpretation for the predictive power of education, it is worth noting that there is enough information in the data to separate the partial effects of educational attainment and occupational class. This is illustrated in table 4, which tabulates estimated partial effects from the fixed effects fractional probit model using educational attainment, every demographic variable, and every demographic variable except educational attainment in the set of predictors.[5] Interestingly, when educational attainment is excluded from the

[5]We report further measures of multicollinearity in the online appendix. As a robustness check, partial effects are also estimated from the fixed effects fractional probit model where the data are weighted using valid votes cast. These are also reported in the appendix, and are not materially affected by the weighting.

model – in column 3 of table 4 – the partial effect of occupational class on the Leave vote share is large, positive, and significant. When educational attainment is included – in column 2 of table 4 – the partial effect of occupational class on the Leave vote share is negative, while the effect of educational attainment is strongly negative. As educational attainment and occupational class are negatively correlated, this suggests that the positive partial effect of occupational class on Leave voting when educational attainment is excluded is simply picking up the partial effect of the excluded variable. In other words, the positive correlation between the working-age population share with a blue-collar occupation and Brexit voting is potentially due to bias resulting from the exclusion of educational attainment. When educational attainment is included in the model, local authorities with more blue-collar workers are predicted to have lower Leave shares.

Despite difficulties in isolating causal channels, the predictive capacity of education is important in itself. This is because it indicates the presence of persistent social and political divides, despite evidence that polarisation over issues such as gender roles, homosexuality, and race is in long-run decline (Duffy et al., 2019). The Equalities and Human Rights Commission noted in 2011 that the educational attainment gap between individuals of differing socio-economic backgrounds remains wide in Britain, even after a long period of improvement (EHRC, 2011). Since 2010, state education has been significantly affected by austerity (Granoulhac, 2017), with adult participation in further and higher education faring particularly badly (Lupton et al., 2015). Education is therefore likely to remain a key dividing line in British politics and society for the foreseeable future. Although the analysis presented in this paper cannot disentangle the precise causal channels connecting these phenomena, it does illustrate the importance of education as a key predictor of the geography of Brexit.

References

Alabrese, E., Becker, S., Fetzer, T., & Novy, D. (2018). Who voted for Brexit? Individual and regional data combined. *European Journal of Political Economy*, *56*, 132–150.

Baynes, C. (2017). Brexit: Labour MP who said Remain voters were better educated 'has facts on his side', says former YouGov boss. *The Independent, Monday 30th October*.
URL https://www.independent.co.uk/news/uk/politics/brexit-remain-voters-better-educated-labour-mp-barry-sheerman-huddersfield-yougov-facts-leave-a8027121.html

Becker, S., Fetzer, R., & Novy, D. (2017). Who voted for Brexit? A comprehensive district-level analysis. *Economic Policy*, *32*(92), 601–650.

Beecham, R., Slingsby, A., & Brunsdon, C. (2018). Locally-varying explanations behind the United Kingdom's vote to the Leave the European Union. *Journal of Spatial Information Science*, *16*, 117–136.

Colantone, I., & Stanig, P. (2018). Global competition and Brexit. *American Political Science Review*, *112*(2), 201–218.

Duffy, B., Hewlett, K., McCrae, J., & Hall, J. (2019). Divided Britain? Polarisation and fragmentation trends in the UK. *King's College London Policy Institute Report*.

EHRC (2011). How fair is Britain? Equality, Human Rights and Good Relations in 2010. Tech. rep., Equality and Human Rights Commission.
URL https://www.equalityhumanrights.com/sites/default/files/how-fair-is-britain.pdf

Fazackerley, A. (2017). Universities deplore 'McCarthyism' as MP demands list of tutors lecturing on Brexit. *The Guardian, Tuesday 24th October*.
URL https://www.theguardian.com/education/2017/oct/24/universities-mccarthyism-mp-demands-list-brexit-chris-heaton-harris

Fetzer, T. (2019). Did austerity cause Brexit? *American Economic Review*, *109*(11), 3849–86.

Goodwin, M., & Heath, O. (2016). The 2016 referendum, Brexit, and the left behind: An aggregate-level analysis of the result. *The Political Quarterly*, *87*(3), 323–332.

Granoulhac, F. (2017). Excellence at What Cost? Austerity and the Reform of the School System in England (2010–2016). *The Political Quarterly*, *88*(3), 434–442.

Hakhverdian, A., Van Elsas, E., Van der Brug, W., & Kuhn, T. (2013). Euroscepticism and education: A longitudinal study of 12 EU member states, 1973–2010. *European Union Politics*, *14*(4), 522–541.

Härdle, W., & Simar, L. (2015). *Applied Multivariate Statistical Analysis*. Springer.

Harris, R., & Charlton, M. (2016). Voting out of the European Union: Exploring the geography of Leave. *Environment and Planning A*, *48*(11), 2116–2128.

Hobolt, S. (2016). The Brexit vote: A divided nation, a divided continent. *Journal of European Public Policy*, *23*(9), 1259–1277.

Johnston, R., Manley, D., Pattie, C., & Jones, K. (2018). Geographies of Brexit and its aftermath: voting in England at the 2016 referendum and the 2017 general election. *Space and Polity*, *22*(2), 162–187.

Lee, N., Morris, K., & Kemeny, T. (2018). Immobility and the Brexit vote. *Cambridge Journal of Regions, Economy and Society*, *11*(1), 143–163.

Lubbers, M., & Scheepers, P. (2010). Divergent trends of euroscepticism in countries and regions of the European Union. *European Journal of Political Research*, *49*(6), 787–817.

Lupton, R., Unwin, L., & Thomson, S. (2015). The Coalition's record on Further and Higher Education and Skills: Policy, Spending and Outcomes 2010-2015. Tech. rep., London School of Economics and Political Science, Centre for Analysis of Social Exclusion.
URL http://sticerd.lse.ac.uk/dps/case/spcc/WP14.pdf

Manley, D., Jones, K., & Johnston, R. (2017). The geography of Brexit - What geography? Modelling and predicting the outcome across 380 local authorities. *Local Economy*, *32*(3), 183–203.

Zhang, A. (2018). New findings on key factors influencing the UK's referendum on leaving the EU. *World Development*, *102*, 304–314.

Online appendix

Table A1: Full results for models using all demographic variables, England and Wales.

	Standard probit	Fixed effects probit	Random effects probit	Fractional probit	L.D.A.
education	-4.74^{**}	-4.09^{**}	-4.30^{**}	-1.31^{**}	1.25
median age	0.02^{**}	0.03^{**}	0.03^{**}	0.01^{**}	-0.44
birthplace	-0.32	0.25	-0.03	-0.15	0.17
ethnicity	0.24	-0.11	0.03	0.15^{**}	-0.30
gender	7.73^{**}	7.32^{**}	7.09^{**}	1.44^{**}	-0.29
occupation	-1.79^{**}	-1.04	-1.29^{*}	-0.21^{*}	0.34

Notes: * $p < 0.05$, ** $p < 0.01$. Neither the standard probit model nor the linear discriminant analysis include fixed or random regional effects. The four columns of results for the probit models report average partial effects, i.e. $\mathbb{E}[dy/dx]$. The final column reports the standardized canonical discriminant function coefficients returned from the linear discriminant analysis.

Table A2: Full-sample results: per cent of local authorities correctly classified using 2016 Annual Population Survey data on qualifications.

	England and Wales			Great Britain		
	All	Leave	Remain	All	Leave	Remain
Standard probit:						
Education only	89.3	94.7	72.6	81.0	90.9	58.6
Fixed effects probit:						
Education only	89.9	95.1	73.8	82.3	90.1	64.7
Fixed effects fractional probit:						
Education only	89.4	92.0	81.2	83.9	87.5	76.1
Random effects probit:						
Education only	89.1	94.7	71.8	90.5	95.1	80.3
Linear discriminant analysis:						
Education only	89.6	96.6	67.9	82.1	93.5	56.0

Notes: The standard probit model does not include fixed or random regional effects. The random effects model is estimated using `meprobit` in Stata. The data are the percentage of local authority populations aged 16-64 with a degree or equivalent qualification, and correspond to the 12 months up to December 2016 (so the 2016 referendum is roughly in the middle of the data collection period). The Isles of Scilly are not included, as there is no observation for this local authority. The data were sourced from NOMIS.

Table A3: Results for the fixed effects fractional probit model, England and Wales, with observations weighted by valid votes cast.

	education only	all demog. vars.	all demog. vars. except education
education	-1.15^{**}	-1.31^{**}	
median age		0.01^{**}	0.00^{**}
birthplace		-0.10^{**}	0.62^{**}
ethnicity		0.15^{**}	-0.30^{**}
gender		1.70^{**}	0.52^{**}
occupation		-0.18^{**}	1.03^{**}

Notes: * $p < 0.05$, ** $p < 0.01$. The reported results are average partial effects, i.e. $\mathbb{E}[dy/dx]$, for the fractional probit model with regional fixed effects and frequency weights defined by valid votes cast per local authority. So Birmingham, for example, in which 450,702 valid votes were cast, receives a weighting approximately four times that of Hull, in which 113,355 valid votes were cast.

Table A4: Residual spatial correlation tests for the standard and fixed effects probit models.

	England and Wales	
	I^2	p-value
Standard probit:		
Education only	4.304	0.038
All demography	10.051	0.002
All demog. minus educ.	22.879	0.000
Fixed effects probit:		
Education only	0.185	0.668
All demography	2.243	0.134
All demog. minus educ.	2.563	0.109

Notes: Following Amaral, Anselin & Arribas-Bel (2013) [Testing for spatial error dependence in probit models, *Letters in Spatial and Resource Sciences*, 6(2), 91-101], we report the I^2 test statistic from the Moran's I test for models with binary dependent variables proposed by Kelejian & Prucha (2001) [On the asymptotic distribution of the Moran I test statistic with applications, *Journal of Econometrics*, 104(2), 219-257]. The test statistic is,

$$I^2 = \frac{(e_1' W e_1)^2}{\mathrm{tr}\,(W\Sigma W\Sigma + W'\Sigma W\Sigma)},$$

where the elements of the column vector e_1 are residuals defined by,

$$e_{1i} = y_i - \Phi(\mathbf{x}_i\hat{\beta}),$$

for observation i, and Σ is a diagonal matrix containing the variances of the residuals on the main diagonal, $\hat{\sigma}_i^2 = \hat{\Phi}_i(1 - \hat{\Phi}_i)$. The matrix W is a spatial distance matrix with zeros on the main diagonal. We use an inverse distance matrix, where the distances are measured from the population-weighted centroids of each local authority district, and we normalise W such that its largest eigenvalue is 1.

The test statistic is asymptotically distributed as $\chi^2(1)$ under the null of no spatial autocorrelation. As reported in Amaral et al. (2013), the test is unbiased across a wide range of sample sizes, and achieves its asymptotic distribution under the null for considerably smaller sample sizes than we use in this paper. In addition, it has good power against the alternative, and the test does not appear to be affected by spatial correlation in the regressors.

Table A5: Multicollinearity statistics for all demographic variables, England and Wales.

	VIF	SQRT VIF	Tolerance	R^2
education	7.46	2.73	0.1341	0.8659
median age	2.89	1.70	0.3465	0.6535
birthplace	11.28	3.36	0.0886	0.9114
ethnicity	9.29	3.05	0.1077	0.8923
gender	1.37	1.17	0.7308	0.2692
occupation	6.17	2.48	0.1622	0.8378

Notes: R^2 is calculated by regressing the variable in question on all remaining variables; tolerance is given by $1 - R^2$; VIF (variance inflation factor) is given by $(1 - R^2)^{-1}$. A tolerance of less than 0.1, or equivalently a VIF of greater than 10, is usually seen as evidence of multicollinearity. VIF can be interpreted as the factor by which OLS estimator variances are inflated due to the presence of multicollinearity, so the square root of the VIF is the factor by which OLS estimator standard deviations are inflated. Note that this does not directly translate to the limited dependent variable models used in the paper.