NBER WORKING PAPER SERIES

WILL BREXIT AGE WELL? COHORTS, SEASONING AND THE AGE-LEAVE GRADIENT, PAST, PRESENT AND FUTURE

Barry Eichengreen
Rebecca Mari
Gregory Thwaites

Working Paper 25219
http://www.nber.org/papers/w25219

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2018

We thank John Curtice, Maria Grasso, Ashoka Mody and James Tilley for helpful advice. All errors remain our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Will Brexit Age Well? Cohorts, Seasoning and the Age-Leave Gradient, Past, Present and Future
Barry Eichengreen, Rebecca Mari, and Gregory Thwaites
NBER Working Paper No. 25219
November 2018
JEL No. D72,I38,J11

ABSTRACT

In the UK's 2016 referendum on EU membership, young voters were more likely than their elders to vote Remain. Applying new methods to a half century of data, we show that this pattern reflects both ageing and cohort effects. Although voters become more Eurosceptical as they age, recent cohorts are also more pro-European than their predecessors. Much of the pro-Europeanism of these recent cohorts is accounted for by their greater years of education. Going forward, the ageing of the electorate will thus be offset at least in part by the replacement of older cohorts with younger, better-educated and more pro-European ones. But we also document large nationwide swings in sentiment that have little to do with either seasoning or cohort effects. Hence these demographic trends are unlikely to be the decisive determinants of future changes in European sentiment. Rather, nationwide changes in sentiment, reflecting macroeconomic or other conditions, and the age-turnout gradient will be key.

Barry Eichengreen
Department of Economics
University of California, Berkeley
549 Evans Hall 3880
Berkeley, CA 94720-3880
and NBER
eichengr@econ.Berkeley.edu

Gregory Thwaites
LSE Centre for Macroeconomics
London School of Economics
Houghton Street
London WC2A 2AE
United Kingdom
g.d.thwaites@lse.ac.uk

Rebecca Mari
BAFFI-Carefin Centre
Bocconi University
Milan, 20136
Italy
rebeccamariamari@gmail.com

1 Introduction

In its 23 June 2016 referendum, the UK voted to leave the European Union. One clear pattern in the returns is a positive correlation between age and the probability of voting Leave. 25% of those aged 18-24 voted Leave, for example, compared to 56% of those 50 and older.[2]

How should we interpret this pattern? Does it mean that voters become more Eurosceptic as they age (that the observed pattern is a seasoning effect)? Or does it indicate that recent cohorts are more pro-European (that the pattern is a cohort effect)? And if cohort effects dominate, can they be explained by changes in the characteristics of different generations, such as their education, in ways that enable us to reliably anticipate the attitudes of future cohorts?

These different interpretations have very different implications for the development of public opinion in the UK. If the observed pattern reflects seasoning, then the electorate is likely to become more Eurosceptic as it ages. On the other hand, if the pattern reflects a cohort effect and the attitudes of future cohorts resemble those of recent cohorts, then the UK will become more Europhile as recent pro-European cohorts supplant their Eurosceptic forebears.

Figure 1: Referendum breakdown by age[1]

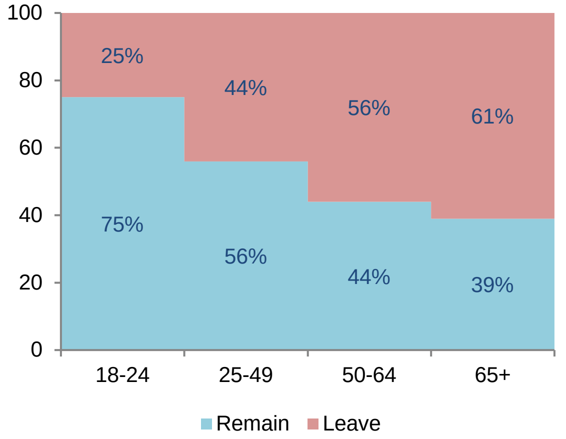

It is not straightforward to disentangle seasoning and cohort effects when there also exist electorate-wide swings in sentiment. Given the available data, one cannot observe the same cohorts at two different ages at the same time, or at the same age at two different times. To make progress, researchers must therefore make assumptions about the form of age, period and cohort (APC) effects and to what observables they are related. For example, authors analysing general election voting in the UK and other countries have attempted to proxy for period effects with observable variables such as GDP growth, allowing them to back out the associated age and cohort effects.[3]

[2] Liberini, Oswald, Proto and Redoano analyze not the referendum vote itself but the Understanding Society dataset, a UK household survey conducted between January 2016 and December 2017, which asked participants questions about Brexit and found little evidence of an age-opinion gradient. According to their findings, only the very youngest voters were significantly more pro-EU; after the age of 24 they find little evidence of opinions changing between voters of different ages. Here we analyze a much larger dataset spanning a considerable longer period of time in an effort to establish robust patterns.

In this paper we assemble panel and repeated cross-section datasets of UK sentiment towards the EU. In all we have more than 100,000 observations spanning the period 1963-2017. To analyse them, we develop a new method for separately identifying APC effects, adapting the Kalman filter to a panel context and relying on the assumption that age and cohort effects evolve smoothly.

Using this approach and also more conventional methods, we establish five findings.

- UK voters grow more Eurosceptic as they age.
- The attitudes of past cohorts trace out a U-shaped pattern. Baby boomers are more Eurosceptic than both earlier and later generations. Thus, both seasoning and cohort effects contribute to the observed age-voting gradient.
- Education, as captured by the number of years of schooling, is positively associated with pro-Europeanism.
- Relatedly, the greater pro-Europeanism of relatively recent cohorts is largely accounted for by their greater educational attainment.
- Finally, there have been large nationwide swings in sentiment that have little do with either seasoning or cohort effects. Any analysis that focuses on demographic trends to the exclusion of these nationwide shifts in attitudes and conditions is likely to mislead.

In conjunction with ONS population projections, our estimates can be used to generate forecasts of EU sentiment. Specifically, if past behaviour is a guide then the ageing of the electorate will partly cancel out the replacement of older Eurosceptic cohorts with newer, better-educated, more pro-European cohorts. But the balance of these two effects is small relative to the age-turnout gradient and the volatility and uncertainty surrounding period effects. It follows that demographic trends are unlikely to be the decisive determinants of future changes in European sentiment. Rather, changes in nationwide attitudes and in the age-turnout gradient will be key.

2 Literature

Since the 2016 referendum, a number of studies have analyzed the determinants of the vote. Fetzer (2018) finds that local and individual exposure to welfare reforms was an important determinant of voting patterns.[4] Becker et al. (2017) show that education, historical dependence on manufacturing employment and economic deprivation were key drivers of the vote at a local level. Antonucci et al.

[3] When doing so they have found evidence of both ageing and cohort effects (see e.g. Tilley and Evans 2014 and Dohmen et al. 2017).

[4] Fetzer in fact looks at the UKIP share of the popular vote in various elections, but this is highly correlated with the Leave vote in the EU referendum.

(2017) argue that a 'squeezed middle' class formed an important component of the Leave vote. In addition to these and related economic factors and characteristics, Swales (2017) finds that questions of national identity were associated with voting patterns. Lee et al. (2018) find that 'locally rooted' individuals living in their county of birth were more likely to vote Leave, but that this tendency was concentrated among respondents in areas experiencing relative economic decline or increases in migration. Garretsen et al. (2018) show that local prevalence of the 'Openness' personality characteristic (in the sense of the Big 5 personality traits or the 'five factor' personality model) was positively associated with the Remain vote share. Finally, Janmaat et al. (2018) find that support for Leave was most volatile among 'engaged and satisfied' voters.

We are not aware of studies that look specifically at APC effects in the context of the EU Referendum or EU sentiment. In passing, McGill (2016) and Zhang (2018) note that the observed age-Remain gradient weakens or disappears once one controls for factors such as education; what might appear to be evidence of a seasoning or cohort effect could actually reflect differences in other voter characteristics insofar as educational attainment and related variables differ across generations.[5] Whether this conclusion is robust to the addition of a separate cohort effect is unclear.

There is of course a large literature on APC effects in other contexts. Studies of ageing (seasoning) effects have related them to psychological changes experienced by individuals on getting older. It is sometimes argued that people grow more conservative, cognitively inflexible and close-minded as they age (see e.g. Feather 1977). Donellan and Lucas (2008) find that individuals become less "open" as they age, irrespective of gender or education. Alternatively, attitudinal shifts could reflect changes in personal circumstances associated with age (see Binstock and Quadagno 2001).

Cohort effects, meanwhile, could reflect the tendency for durable political attitudes to form early in life, influenced perhaps by circumstances prevailing when cohorts enter the electorate (see e.g. Green et al. 2002). Cohorts reaching voting age in recent decades, whose members flew around Europe on Ryanair and had their educations shaped by the Bologna Process, could conceivably view EU membership more favourably than members of earlier cohorts.[6]

Thus, the related literature suggests that both cohort and seasoning effects could lie behind the observed age-Leave gradient but provides no evidence on their relative importance. It is to this task that we now turn.

[5] And since young voters may not have completed their schooling (or had the opportunity to obtain advanced education) at the time they are surveyed.
[6] The same may be true of members of cohorts reaching voting age in the 1950s, when the EU was seen as a solution to the age-old problem of war in Europe. We return to this possibility below.

3 Data

Our principal dataset is assembled from the British Election Study (BES). The BES is a series of partially overlapping panel and repeated cross-section surveys conducted around every UK General Election since 1964 up to an including 2017, as well as at the time of the 1975 and 2016 EU referenda. The long time span of the data permits us to observe cohorts ageing over 54 years. We corroborate our main findings using the British Social Attitudes/European Social Survey (BSA/ESS), as detailed in an appendix.

Survey responses include information on the religion, class, age and educational background of respondents, together with answers to a varying list of attitudinal questions. Usefully for our purposes, in each round except 1979 BES respondents were asked whether the UK should become or remain a member of the EU and its antecedents.[7] We use these responses to code a variable representing the respondent's preference for being part of the EU, giving us a total of 115,134 observations. Our dependent variable is bounded by zero and one.

The precise question posed varies from year to year, although it was always something along the lines of 'Do you think the UK should join/remain in the Common Market/European Community/European Union?' Some change is unavoidable, since the nature of the EU and its predecessors itself changed over the 54-year period covered by the data. Furthermore, not wanting to join (the meaning of a zero until the 1974 survey) is not the exactly same as wanting to leave (from 1974 onwards). These changes in the question are likely to affect the direct comparability over time of our estimates of the time effect, as at different points in time being pro-European might mean being in favor of a higher or lower degree of European integration.[8]

We truncate our data below age 18 and above 80, where we have relatively few observations. Education is proxied by the age at which the respondent left full time school minus 5 or the number of years of full time education, depending on whether or not the individual has completed his studies.[9] We use ONS annual population composition data and forecasts to construct contributions to past and future average sentiment.

---

[7] In the BES, individuals are queried twice, and then replaced; it is an overlapping panel. If they do not respond the second time, then they are replaced prematurely. This may bias our standard errors, but only to a limited extent insofar as the sample is very large. Panel data estimates with fixed and random effects are shown in the appendix.

[8] Furthermore, they may introduce some bias in the estimates of the age and cohort effects. For example, the observed decrease in pro-European sentiment over time could be picking up the greater degree of integration in the European Union – i.e. a more demanding measure of pro-Europeanism.

[9] We drop observations claiming more than 35 years of full-time education.

3 Estimating APC effects

We model the pro-European sentiment of individual *i* born in year t (y_{it}) as the linear sum of an age or seasoning effect a_{t-i}, a period effect τ_t and cohort effects c_i:

$$y_{it} = c_i + a_{t-i} + \tau_t + \epsilon_{it} \quad (1.1)$$

where y_{it} is the response variable of interest (our indicator of pro-Europeanism) of individual or cohort i at time t, c_i is a cohort effect specific to individuals in cohort with date of birth i, a_{t-i} is a seasoning effect associated with cohort i being $t - i$ years old in year t, τ_t is a time effect common to all individuals in the sample at time t, and ϵ_{it} is an idiosyncratic error.

No restriction is placed on the shape of $\{c, a, \tau\}$. However, the framework embeds the restriction that the effects are additive and that seasoning effects are constant across cohorts.[10] Furthermore, we are assuming a linear probability model rather than a nonlinear logit or probit specification for y_{it} to enable us to implement a filtering algorithm. Reassuringly, when we compare linear and probit estimates using our alternative methods, as set out in the appendix, the results are broadly similar.

Figure 2 shows simple averages of y_{it} by time, age and date of birth. Evidently, attitudes toward membership in the EU and its antecedents have fluctuated over the last 55 years, with no clear trend. Sentiment was unusually negative in the 1970s, positive in the 1990s, and negative again around 2010. The 2016 value for the BES was 7 points below the simple average.[11]

[10] Both assumptions are unlikely to hold exactly – it is likely that seasoning effects will operate differently over time and across different households, for example.
[11] Here, as below, there is broad agreement between the BES and the BSA/ESS.

Figure 2 – simple average of y_{it} by age, date of birth and time

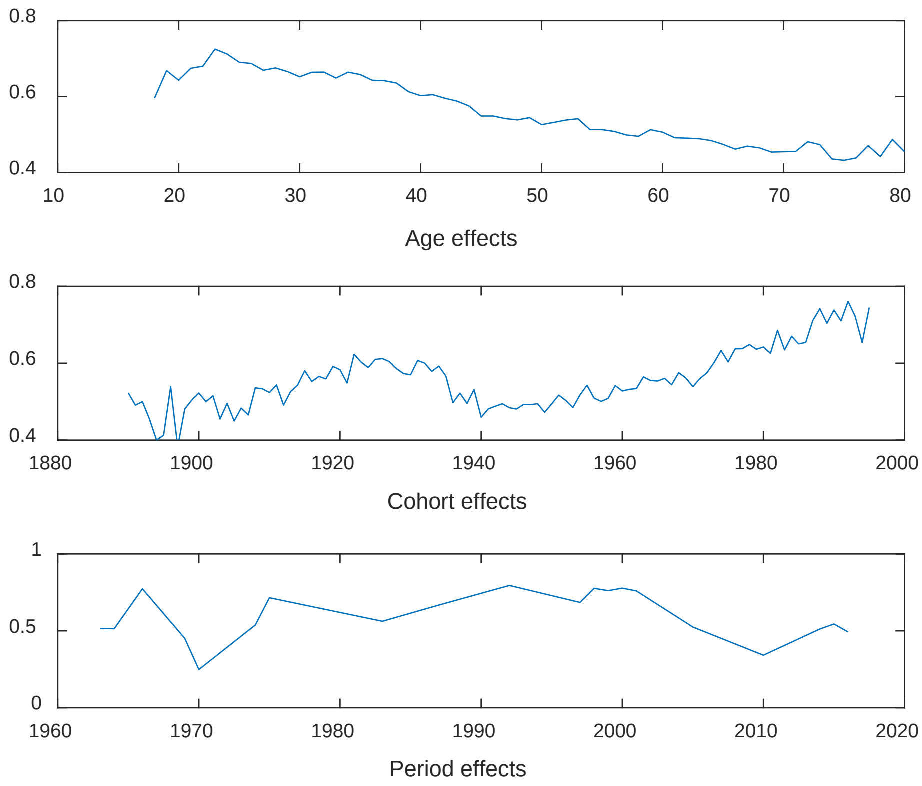

Figure 2 also shows that pro-Europeanism rises appreciably between age 18, below which we truncate our data, and the mid-20s. Given that we document below a strong relationship between schooling and pro-Europeanism, this could be explained by individuals acquiring additional schooling while in their 20s. Beyond their mid-20s, however, older voters have typically been more Eurosceptic, with a fall of 20-30 points between ages 18 and 80. The figure shows a more complex relationship between date of birth and EU sentiment. Away from the extrema of our sample, where the cohorts are too small to give a clear signal, we see a U-shaped cohort effect, with the immediate post-war generations reporting more Euroscepticism than their predecessors or successors.

These simple averages will not in general correspond to the underlying APC effects. We cannot infer from them that voters become more Eurosceptic as they age: older respondents are also necessarily born earlier (they are members of an earlier cohort). Thus, the pattern in the figure could be driven by age effects, in that more recently-born cohorts are only observed when young.

Lastly, it is at least conceivable that the time-series variation in Figure 2 is driven by changes in the age and cohort composition of the sample. Given, however, that the cohort and age composition of the UK electorate has changed only slowly, these effects would have to be very large and volatile to generate the pattern observed in the time series. In practice, there is no reason to think that the attitudes of 19 year olds are radically different from those of 20 year olds, and that the attitudes of individuals born in 1985 are radically different from those born in 1986. Indeed, the idea that these age and cohort effects are likely to be smooth is key to the identification method we now propose.

It is well known that a model in which age, period and cohort effects are all present and unrestricted is not identified (Neundorf and Niemi 2014). Researchers must restrict some or all of the APC effects, either by grouping, omitting or proxying them or else by restricting their functional form.

We take the second of these two approaches. We assume that age and cohort effects, by their nature, are likely to be smooth, in the sense that they are unlikely to change violently between adjacent ages or cohorts. Concretely, we estimate a trend that is close to fitting the data while penalizing its second difference (the change in the change). Note that assuming that age and cohort effects evolve smoothly and gradually does not prevent them from being non-linear or non-monotone. Reversals in the direction of change can still occur, but not too frequently and sharply over a short time span. In contrast, time effects are unrestricted. They may exhibit a trend and vary sharply between successive years. Intuitively, they could vary sharply because an expanding economy enters a recession, because a stable financial system becomes unstable, or because a native population is suddenly be confronted by a large refugee influx.

The problem of dividing an observable variable into a smooth trend and residual "cycle" is isomorphic to one common to time-series econometrics, which involves decomposing a macroeconomic time series y_t into a smooth trend x_t and a cycle $\hat{y}_t$:

$$y_t = x_t + \hat{y}_t \qquad (1.2)$$

As in our problem, with no other information this problem is not identified - infinitely many sequences $\{x, \hat{y}\}$ satisfy this equation. Similarly, in equation (1.1) there are infinitely many sequences $\{c, a, \tau\}$ which yield well-behaved error terms.

We thus adapt an algorithm from time-series econometrics (the Hodrick-Prescott or HP filter) commonly used to identify a smooth trend x_t in equation (1.2) in order to obtain smooth age and cohort effects in equation (1.1). The standard Kalman filter x_t is the solution to the following constrained minimisation problem, minimising the variance of the cycle subject to keeping the trend smooth (i.e. minimising its second differences).

$$x_t = argmin \sum_t (y_t - x_t)^2 + \lambda(\Delta y_t - \Delta y_{t-1})^2$$

where Δ is the time difference operator, and λ determines the size of the penalty on the trend's second differences and therefore the smoothness of the trend. In the limit ($\lambda \to \infty$), the trend becomes linear.

Our age and cohort effects are accordingly identified as the solution to the following minimisation problem

$$c_i, a_{t-i} = argmin \sum_i \sum_t (y_{it} - c_i - a_{t-i})^2 \ + \lambda_c(\Delta_i c_i - \Delta_i c_{i-1})^2 + \lambda_a(\Delta_{t-i} a_{t-i} - \Delta_{t-i} a_{t-i-1})^2$$

where λ_c and λ_a are parameters governing the desired smoothness of the age and cohort effects and Δ_i and Δ_{t-i} denote differences with respect to age and cohort respectively. The solution to this problem is derived in Appendix 1.

The resulting annual age and cohort effects are generalisations of the linear trends in e.g. Tilley and Evans (2014). Equally, our approach is related to the grouping assumption used by Grasso (2014), where cohorts are binned into groups (a 5-year birth cohort, for example) within which there is no variation in cohort effects, but between which the variation is unrestricted.[12] Our method allows for variation between adjacent ages or dates of birth but limits the change in the change of these effects.

Figure 3 shows estimated age, period and cohort effects for different values of λ_c, λ_a. We begin by setting $\lambda_c, \lambda_a = 6.25$, the standard value for the HP filter using annual data, and experiment with larger and smaller values. As shown, larger values of λ yield more nearly linear age and cohort effects. One reassuring insight from Chart D is that the estimated time effects are almost completely insensitive to the choice of λ. This is important, since these time effects will be prominent in our estimates and important for the associated forecasting exercises.

When the smoothness parameter is relatively low, we see that cohort effects are nonmonotone, peaking for the generation born in the interwar period, falling for the baby boomers, and rising again in the most recent cohort.[13] The min-max range of these cohort effects is about 0.1 – other things equal, the probability of supporting EU membership varies by about 10 percentage points according to date of birth. Age effects are larger – they fall within a range of 0.3 - and monotonically negative for a wider range of the smoothing parameters. This echoes the finding in Tilley and Evans (2014) that voters are more likely to vote Conservative as they age. There is some evidence that ageing

[12] See Appendix 3 for a complementary approach using 5-year cohorts and established identification methods

[13] We relate this cohort effect to observable characteristcs – namely education - below.

effects fall off (or flatten) for the very young and the very old. In the latter case, this may be driven in part by differential mortality among different demographic groups.

A value of $\lambda_c, \lambda_a = 0.625$ yields time and age effects with a reasonable degree of smoothness and responsiveness to the data. We therefore take this as our baseline value and conduct sensitivity analyses when making projections.

Figure 3 – HP filtered estimates of APC effects, BES survey, for different values of λ

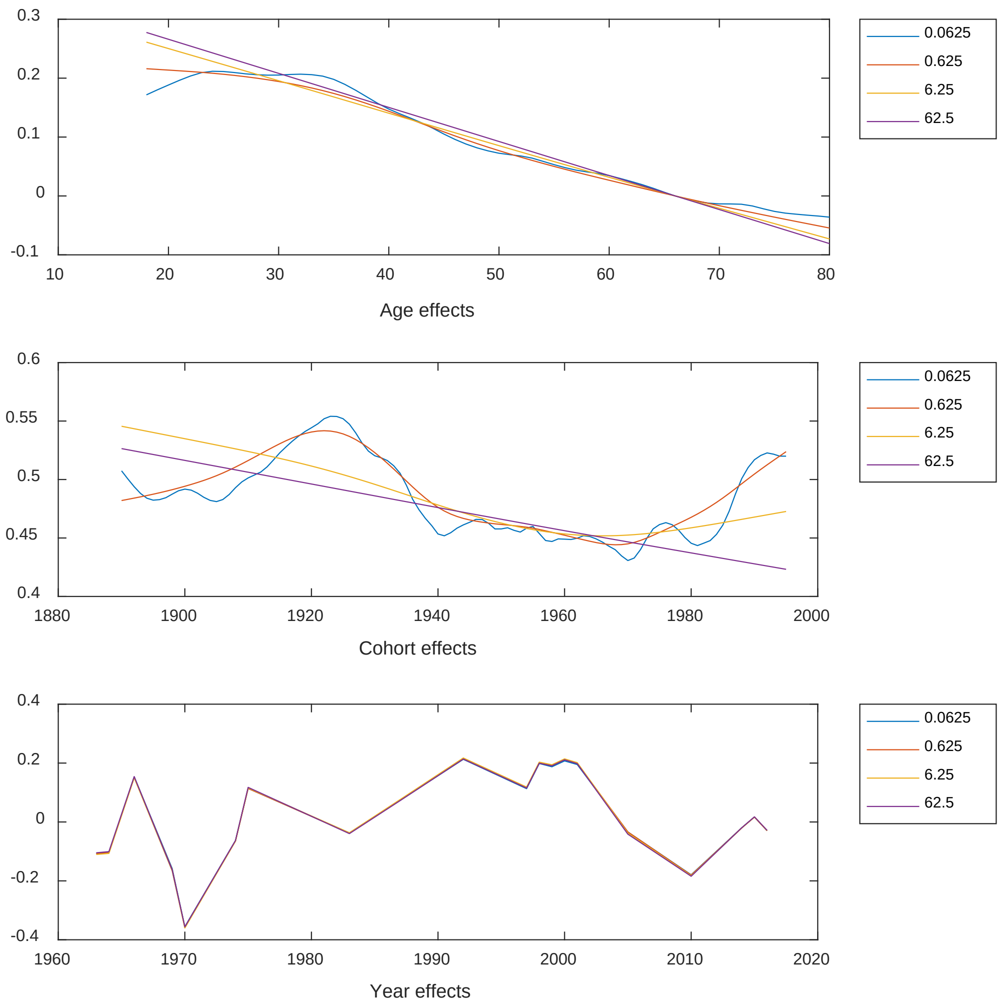

In addition to uncertainty arising from the smoothing parameter, our estimated APC effects are uncertain due to sampling error. Figure 4 displays bootstrapped 90% confidence intervals where we normalize our age effects to be zero at age 65 (the most common age of respondents in our sample). Reassuringly, these intervals are relatively narrow. So while our estimates of age and cohort effects are sensitive to assumptions about their smoothness, uncertainty due to sampling variability is minor.[14]

Figure 4 – bootstrapped confidence intervals for age and cohort effects

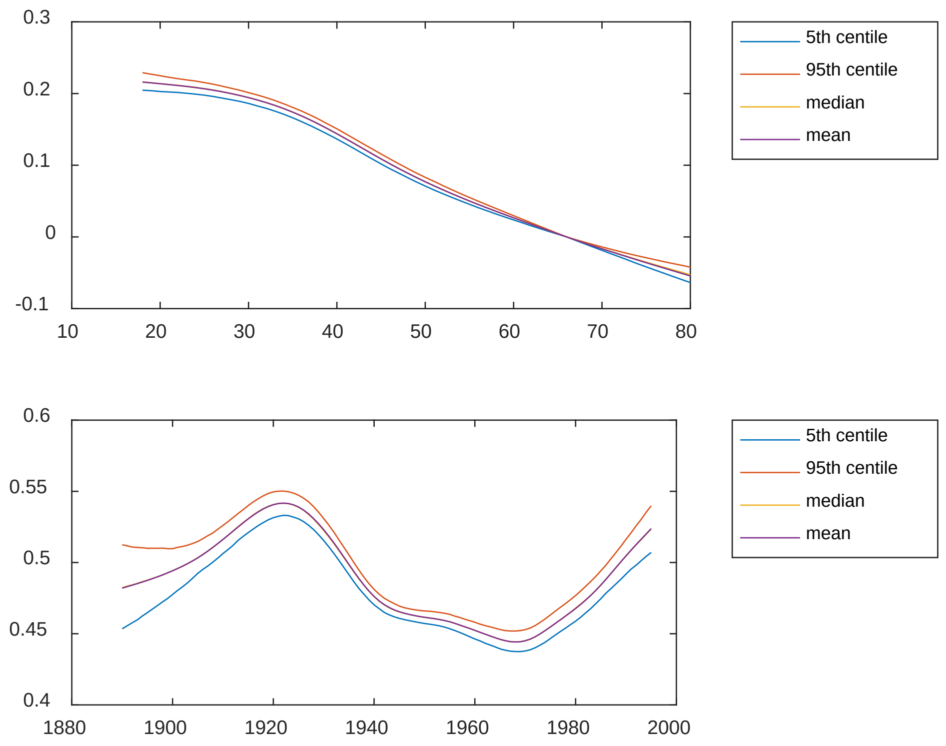

Figure 5 compares our filtered estimates to those produced by simple univariate and bivariate models of each effect, where we regress our indicator y_{it} on one or two sets of dummy variables for age, period and/or cohort effects. Two points emerge from the chart. First, the bottom panel shows, as above, that our estimates of period effects are relatively insensitive to what we assume about cohort and age effects. This is consistent with the fact that the slowly-moving age and cohort effects cannot possibly account for the sharp swings we have seen in national surveys. Second, the middle panel shows that a model without age effects estimates much steeper cohort effects. This reflects the fact that recent cohorts are only observed when young, whereas the obverse is true for the earliest cohorts. This underlines the importance of accounting for age and cohort effects simultaneously.

[14] To obtain our confidence intervals we constructed 1000 synthetic samples of data by resampling with replacement from each year of data, keeping each year's sample size the same. We ran our Kalman filter algorithm for each of these samples and then calculated centiles of the resulting distribution

Figure 5 – univariate, bivariate and filtered estimates of APC effects

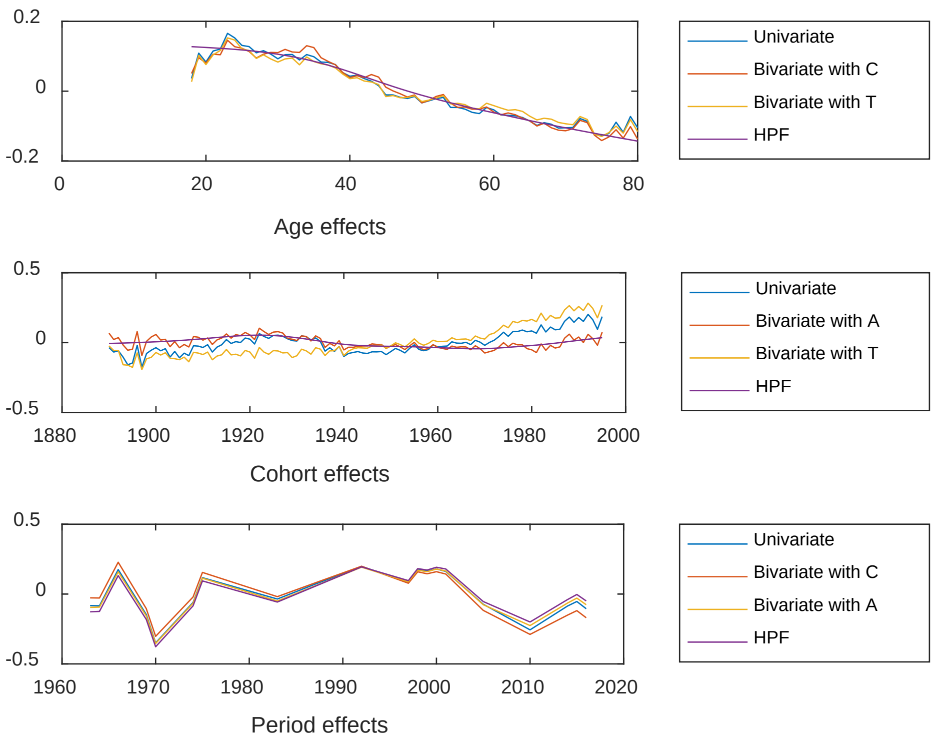

Consistent with this, a bivariate model of y_{it} that omits period effects and depends only on age and cohort comes close to our preferred filtered estimates. One way of examining the robustness of our results is therefore to calculate the bias that would be obtained in a bivariate model when there are omitted time effects. Appendix 2 shows that the bias that would result from an omitted time effect depends on when the time effects happen. A negative bias to the age gradient would require a negative time trend, with large negative time effects at the end of the sample. In fact, all the measures we have of time effects suggest that they are small in 2016 and reasonably stationary, in contrast to the trend needed to generate our results. This again serves to reassure.

As a further robustness checks, Figure 6 plots the time series of average pro-European sentiment for 10-year birth cohorts, normalized for the average sentiment in the population at the time it was measured. It shows how the pro-Europeanism of successive cohorts has evolved over their lifetimes controlling for time effects. The picture indicates that the average trend in pro-EU sentiment over a typical cohort's life is negative, albeit highly variable.

Finally, in Appendix 3 we report estimates from a model that proxies for time effects with observable variables, such as unemployment and GDP growth, and finds strong negative age effects.

Figure 6 - evolution of pro-EU sentiment across the lifetime, by cohort - *BES*

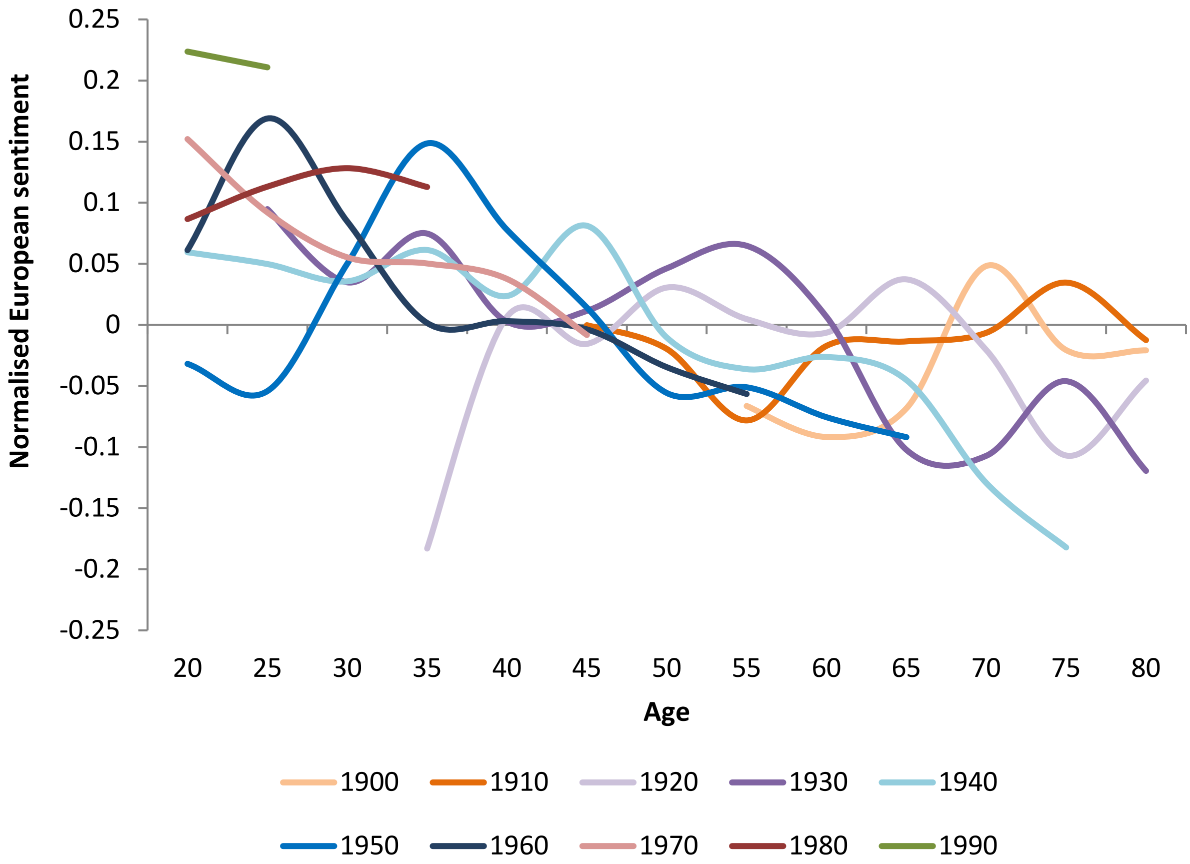

4 Cohort effects and schooling

While we have identified differences in pro-European sentiment across cohorts at a given age, we have not explained their presence. But it may be possible to relate them either to specific events that happen to those cohorts at formative points in their lives (see e.g. Malmandier and Nagel 2009, Malmandier et al. 2017) or to changes in the composition of observable characteristics between generations. Explaining cohort effects in this way may help to forecast their future evolution.

In this section we ask how much of the cohort effect can be explained by changes in years of schooling. Figure 7 shows a non-monotone relationship in 2016 between age (and hence date of birth) and average years of schooling, with average years of schooling rising through the mid-20s as people typically complete their education, and then falling until around age 80, reflecting rises in the school leaving age and recent increases in participation in tertiary education. [15]In addition there is likely to be a stronger correlation between mortality and education among the very old. Note that a similar relationship exists between average pro-European sentiment and age.

Figure 7 - age left school and share pro-EU by age, 2016

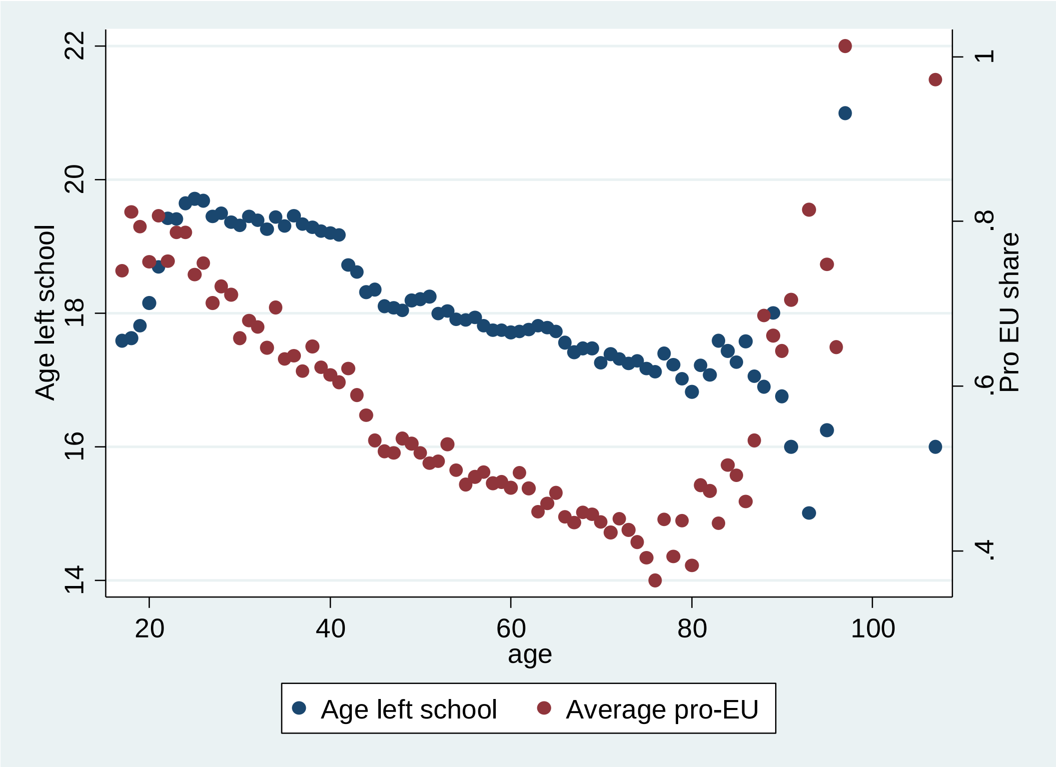

Figure 7 is consistent with education being an important determinant of pro-European sentiment, and therefore with the hypothesis that changes in education explain part of the behaviour of different cohorts. But it is also consistent with age/cohort and education exhibiting common trends.

To go further we must compare individuals with given ages or dates of birth but different levels of education. We therefore estimate a version of our bivariate age-cohort effect model (i.e. omitting period effects), breaking down cohort effects into an effect common to all members of the cohort and a term that depends on years of schooling but varies across cohorts[16]

$$y_{it} = c_i + a_{t-i} + c_i \times ageleft_i$$

[15] Beyond age 80 there are few households in the sample, increasing sampling error.
[16] We omit period effects here in order to be able to identify and decompose age and cohort effects. Figure 5 above shows that this provides a reasonable approximation to our filtered model.

The last term is an interaction between cohort and the age at which the individual left full-time education. We can interpret this as the effect of one extra year of schooling on a cohort born in a given year.

Figure 8 plots the central estimates and confidence intervals for this set of coefficients. Education is positively associated with pro-EU sentiment throughout the sample, with the effect rising from 0.02 (additional decimal points in favour of Remain per extra year of schooling) for pre-war cohorts to over 0.05 for the baby boomers, and then falling back to around 0.03 for millennials.

Figure 8 - effect of an extra year of schooling on probability of supporting EU membership, by birth cohort

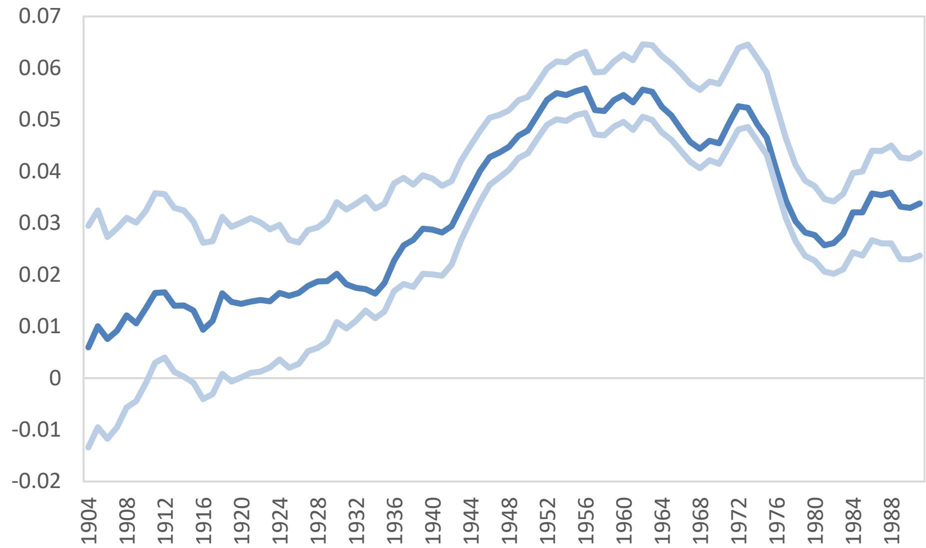

We can use these coefficients to account for the effect of changing education levels on the pro-Europeanism of successive cohorts. Figure 9 shows the cohort effect for a hypothetical individual with no education (the blue line), the simple cohort effect from a model without education (the purple line), the cohort effect for an individual with an average level of education for someone with their date of birth (the green line), and the cohort effect averaging education across the sample (the red line). The change in the pure cohort effect denoted by the blue line is more pronounced once we allow for the rising amount and, for much of the sample, effect of education within cohorts. This is in line with the results shown also in Appendix 3 using an alternative model.

The fact that education has been a reliable predictor of pro-European sentiment throughout the sample suggests that future cohorts are likely to remain relatively pro-European, assuming that they will be at least as well-educated as existing cohorts. And insofar as levels of education are unlikely to change radically, it provides a basis for projecting cohort effects over time.

Figure 9 – impact of education on cohort effects

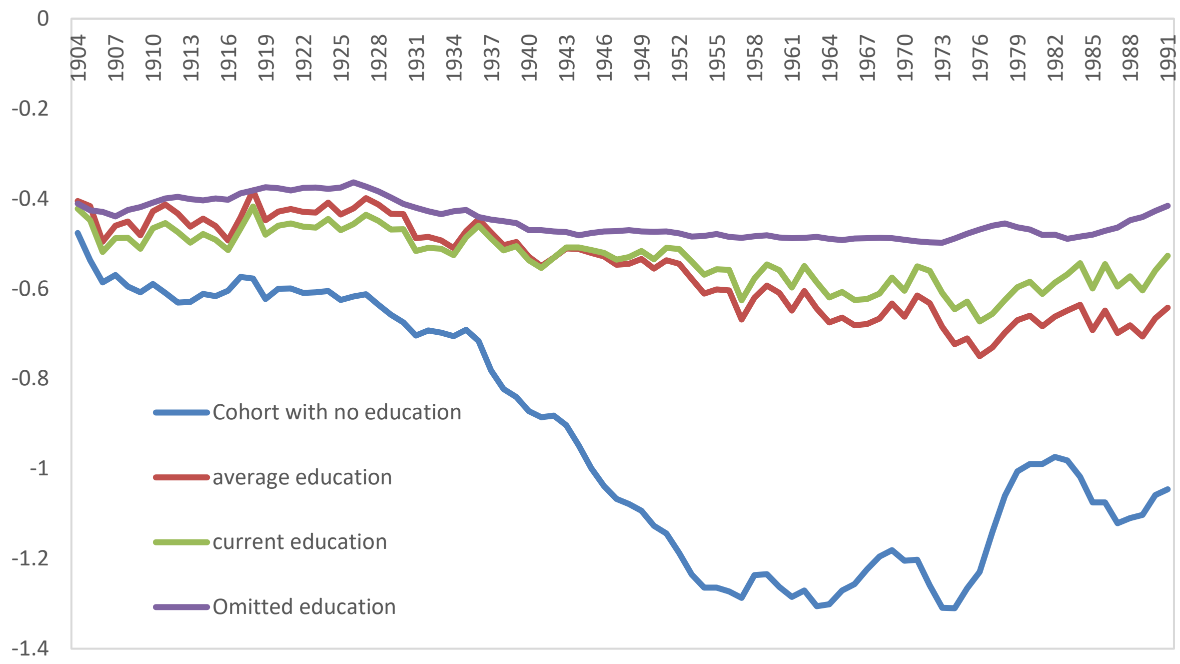

5 Forecasting Pro-EU sentiment

We now use our estimates of age and cohort effects to produce illustrative forecasts of pro-European sentiment. We obtain population data and projections from the Office of National Statistics for 1971-2068. Figure 10 shows how the age distribution of the electorate has and is forecast to change. Falling birth rates and rising longevity have both contributed to ageing of the population, a trend that is forecast to continue over coming generations. Given the negative ageing effects found in the previous section, this is likely to have a negative effect on aggregate sentiment toward the EU.

Figure 10 – age distribution of the UK population age 18-80

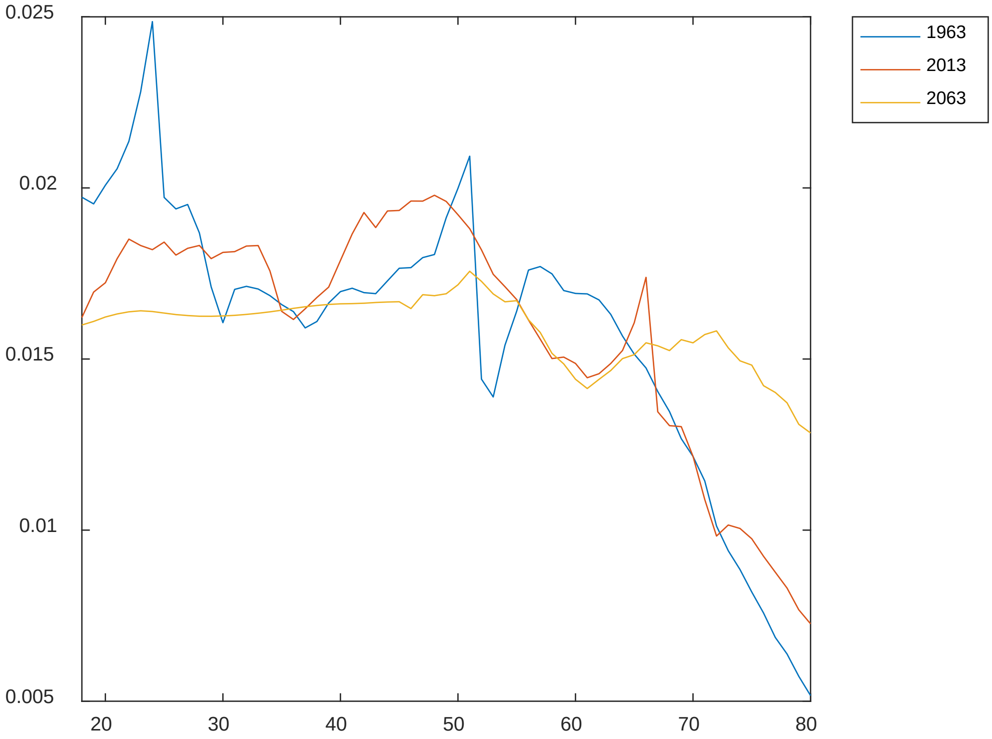

Figure 11 shows the mean of y_{it} by age in the 2016 BES, alongside our estimates of the age effect and the cohort effects that apply to those ages in 2016, both of which were estimated over the entire 1963-2016 sample. The sum of these effects does a good job at capturing the 2016 age-Leave gradient. Our coefficient estimates suggest that most of the age-leave gradient observed in 2016 is attributable to age rather than cohort effects.

Figure 11 – estimated and actual age profile of leave sentiment in 2016

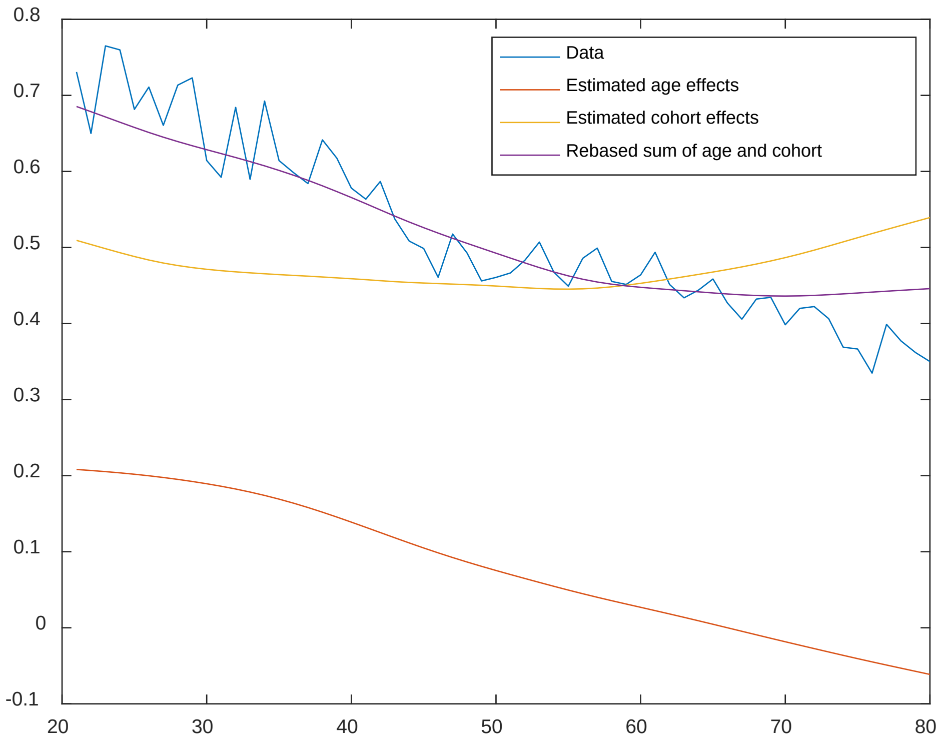

As a neutral assumption we hold future cohort effects fixed at the level of current young cohorts. That cohort effects are largely accounted for by differences in education levels, which are unlikely to change radically, lends weight to this assumption.

Figure 12 shows the contributions of age and cohort effects to average sentiment relative to the population distribution in 1971, when our population data begin. The solid lines show population-weighted averages, while the dotted lines show voting-population-weighted averages, where we obtain the age-turnout gradient from the 2016 British Social Attitudes Survey. Evidently, cohort effects have affected sentiment toward the EU over the past half-century in increasingly negative ways, as the relatively pro-European pre-war generation was progressively replaced by more Eurosceptic baby boomers. The contribution of this cohort effect is now about to turn positive, but it will be partly offset by an increasingly negative contribution from ageing, which to date has not had a large impact on the aggregate. The net impact of these two effects has been to reduce pro-European sentiment by about 3 percentage points over the past half-century, an effect that will have reversed by the middle of this century.

Figure 12 – time series of estimated age and cohort effect contributions[17]

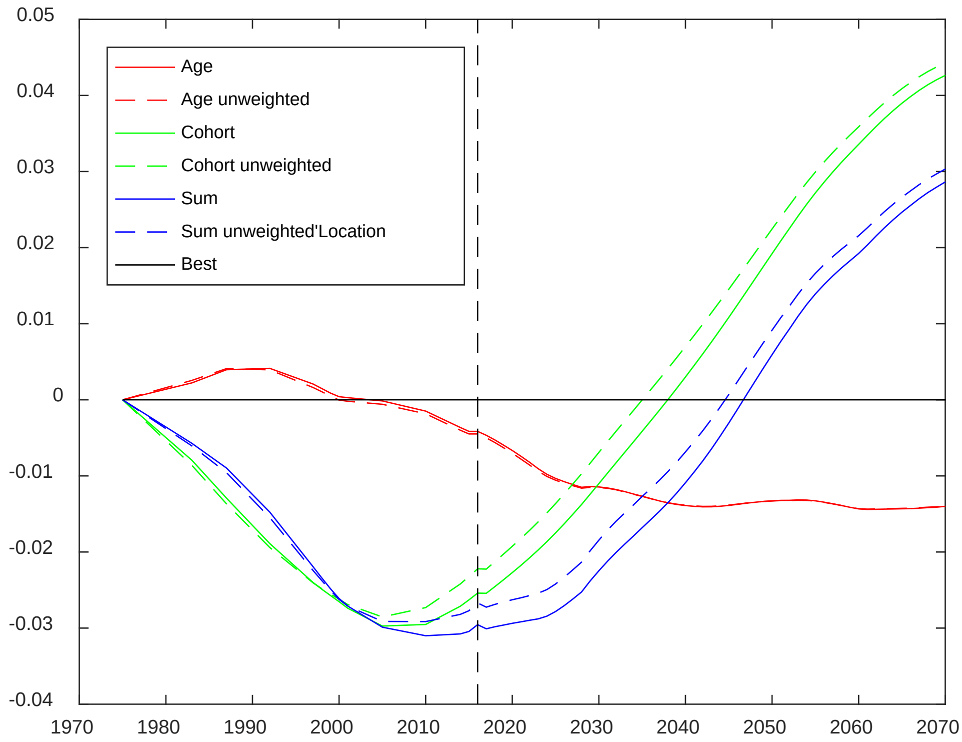

Figure 13 displays several alternative decompositions and forecasts of the change in pro-EU sentiment since our population data began in 1971. We estimate sampling error using our bootstrapped estimates of APC effects to calculate alternative contributions of the age and period effects. The 5th and 95th centile are shown as the "low bootstrap" and "high bootstrap" in the figure. Consistent with the narrow confidence intervals displayed in the previous section, sampling error is relatively small.

We can further explore the impact of model uncertainty by calculating the sum of age and cohort contributions using a larger value of the smoothing parameter and a bivariate age-cohort model that omits period effects. The resulting range is somewhat wider – the high λ_c, λ_a model projects no rise in pro-European sentiment, while the bivariate model predicts a sharper rise than the baseline.

Especially important, our analysis reveals, is the contribution of time effects. To show this, we linearly interpolate the estimates for survey years to obtain an annual dataset, fit a second-order autoregressive model to the interpolated data, simulate draws from this series, and add the 5th and 95th percentiles to the baseline forecast (denoted "low time effect" and "high time effect" in the

[17] The kink in the chart is due to a discontinuity in the population backdata vis-a-vis the population projections.

figure). The difference between the “low time effect” and “high time effect” is very large. If past volatility is a guide, then these effects and the uncertainty surrounding them swamp the impact of changes in the age and cohort effects over the entire plausible range for the latter. This suggests focusing in future work on the underlying drivers of these time or period effects, since the latter are likely to be the dominant determinants of sentiment toward the UK going forward.

Figure 13 – past uncertainty and future volatility in change in weighted average y_{it} due to sampling variability, model uncertainty and volatility of time effects

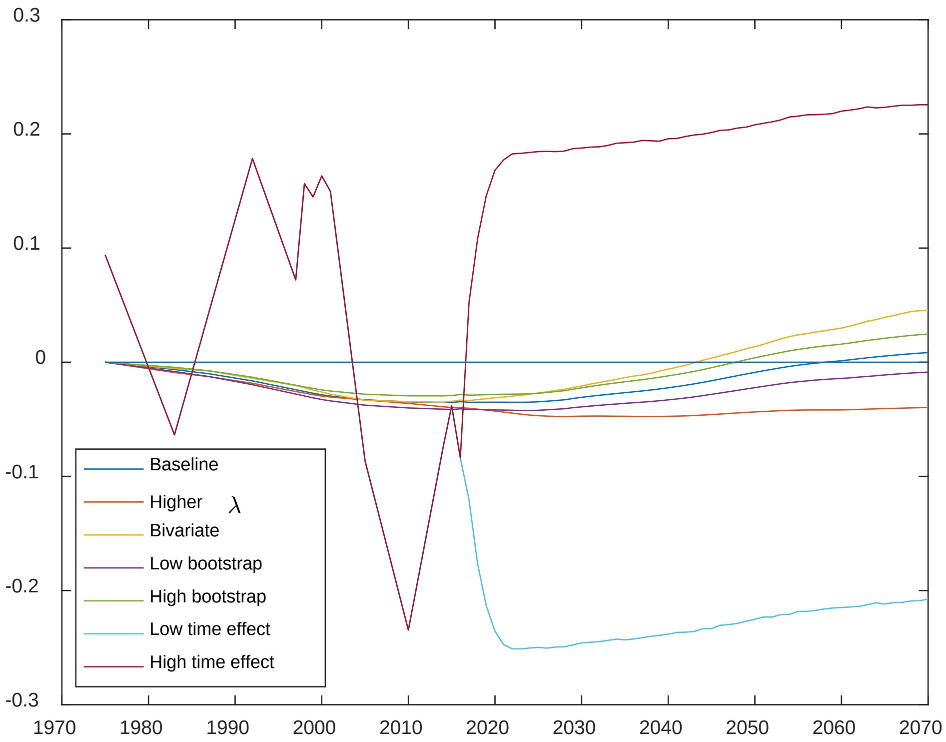

6 Conclusion

In the UK’s 2016 Brexit referendum, young voters were more likely than their elders to support remaining in the European Union. Using a half-century of data and new techniques, we find that recent cohorts tend to be more pro-European than their predecessors, but that voters also become more Eurosceptic as they age. Furthermore, we show that much of the pro-Europeanism of recent cohorts is associated with greater years of education. Finally, we document large nationwide swings in sentiment that have little to do with seasoning or cohort effects. These time effects are plausibly associated with, inter alia, macroeconomic fluctuations, financial conditions and geopolitical circumstances, but they also could have other sources. They are volatile, and they dominate the impact of the estimated seasoning and cohort effects. These results are robust to the use of alternative data sets, methods and estimators.

We use these estimates to provide a new accounting of past trends and a forecast of future sentiment. Our baseline estimates of age and cohort effects suggest that together they have reduced pro-EU sentiment by about 3 percentage points over the past 40 years, as the electorate has aged and Eurosceptic baby-boomers have replaced their more pro-European parents. If the attitudes of future cohorts resemble those of the currently young, this effect will have been reversed by the middle of the present century, as further ageing of the electorate is more than offset by the passing of the relatively Eurosceptic baby-boomer generation.

However, these effects are small compared to the volatility and uncertainty surrounding time effects. It follows that demographic trends are unlikely to be the decisive determinants of future changes in European sentiment. Rather, nationwide changes in attitudes and in the age-turnout gradient will be critical.

As well as the future paths of time and cohort effects, there are further sources of uncertainty that we do not model here. There is uncertainty about turnout, fertility, mortality and migration. But, notwithstanding such uncertainty, our results point to the time effects as key. These can and have varied over a wide range. That variation is thus likely to swamp the combined impact of age and cohort effects for all reasonable values of the latter. Understanding what lies behind these time or period effects in order to forecast them more accurately must be at the top of the agenda for those concerned to understand the evolution of UK public opinion towards the European Union.

References

Alwin, D.F. and J.A. Krosnick (1991), “Ageing, Cohorts and the Stability of Socio-Political Orientations Over the Life Span,” *American Journal of Sociology* 97, pp.169-195.

Antonucci, L., Horvath, L., Kutiyski, Y. and A. Krouwel (2017), “The malaise of the squeezed middle: Challenging the narrative of the ‘left behind’ Brexiter,” *Competition and Change* 21, pp. 211-229.

Becker, S., Fetzer, T. And D. Novy, (2017). "Who Voted for Brexit? A Comprehensive District-Level Analysis," *Economic Policy* 32, pp. 601-650.

Binstock, R.H. and J. Quadagno (2001), “Aging and Politics,” in R.H. Binstock and L.K. George (eds), *Handbook of Aging and the Social Sciences*, San Diego: Academic Press, pp.333-351.

Donellan, M. B. And R. E. Lucas, (2008), “Age Differences in the Big Five Across the Life Span: Evidence from Two National Samples,, *Psychological Aging* 23, pp. 558-566.

Dohmen, Thomas, Armin Falk, Bart H. H. Golsteyn, David Huffman, and Uwe Sunde (2017), “Risk Attitudes Across the Life Course,” *Economic Journal* 127, pp.F95-F116.

Feather, N.T. (1977), “Value Importance, Conservatism and Age,” *European Journal of Social Psychology* 7, pp.241-245.

Fetzer, T. (2018), “Did Austerity Cause Brexit?” CESifo Working Paper no. 7159 (July).

Garretsen H, J I Stoker, D Soudis, R L Martin and P J Rentfrow (2018), “Brexit and the Relevance of Regional Personality Traits: More Psychological Openness could have Swung the Regional Vote,” *Cambridge Journal of the Regions, Economy and Society* 11, pp.165-175.

Grasso, M.T. (2014) “Age, Period and Cohort Analysis in a Comparative Context: Political Generations and Political Participation Repertoires in Western Europe," *Electoral Studies* 33, pp.63-76.

Green, D., B. Palmquist and E. Schickler, *Partisan Hearts and Minds*, New Haven: Yale University Press.

Hobolt, S (2016), “The Brexit Vote: A Divided Nation, a Divided Continent,” *Journal of European Public Policy* 23, pp.1259-1277.

Janmaat, J. G., Melis, G., Green, A. and N. Pensiero (2018), “Changing Preferences for Brexit: Identifying the Groups with Volatile Support for ‘Leave’”, LLAKES Research Paper 65, University College London: Centre for Learning and Life Chances in Knowledge Economies and Societies.

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

Liberini, F., A. Oswald, E. Proto and M. Redoano (2017), "Was Brexit Caused by the Unhappy and the Old?" IZA Discussion Paper no.11059 (September).

Malmandier, Ulrike and Stefan Nagel (2009), "Depression Babies: Do Macroeconomics Experiences Affect Risk-Taking?" *Quarterly Journal of Economics* 126, pp.373-416.

Malmandier, Ulrike, Stefan Nagel and Zhen Yan (2017), "The Making of Hawks and Doves: Inflation Experience on the FOMC," NBER Working Paper no.23228 (March).

McGill, Andrew (2016), "Who Voted for Brexit?" *The Atlantic* (25 June), https://www.theatlantic.com/international/archive/2016/06/brexit-vote-statistics-united-kingdom-european-union/488780/.

Neundorf, A. and R. Niemi (2014), "Beyond Political Socialization: New Approaches to Age, Period, Cohort Analysis," *Electoral Studies* 33, pp.1-6.

Swales, K. (2017), "Understanding the Leave Vote," London: National Centre for Social Research

Tılley, James, and Geoffrey Evans (2014), "Ageing and Generational Effects on Vote Choice: Combining Cross-Sectional and Panel Data to Estimate APC effects" E*lectoral Studies* 33, pp.19-27.

Van der Brug, Wouter, and Sylvia Kritzinger (2012), "Special Symposium: Generational Differences in Electoral Behaviour," *Electoral Studies* 31, pp.245-468.

Zhang, Aihua (2018), "New Findings on Key Factors Influencing the UK's Referendum on Leaving the EU," *World Development* 102, pp.304-31.

Appendix 1 – Derivation of multivariate panel Kalman filter

Formally, if we write a time series as a vector $y = (y_1, \dots, y_T)'$ the HP filter is the solution to the following problem

$$\min_x V = (y - \tilde{y})'(y - \tilde{y}) + \alpha \tilde{y}' P' P \tilde{y} \quad (1.3)$$

with the smoothing parameter governing the trade-off between the variance of the cyclical component $\hat{y}_t = y - \tilde{y}_t$ and the smoothness of the cyclical component , calculated in turn by the second difference operator

$$P = \begin{bmatrix} 1 & -2 & 1 & 0 & \dots & 0 \\ 0 & 1 & -2 & 1 & \dots & 0 \\ 0 & 0 & 1 & -2 & \dots & 0 \\ & & \dots & & & \\ 0 & 0 & 0 & 0 & \dots & 1 \end{bmatrix} \quad (1.4)$$

a matrix of order $(T - 2) \times T$

With this is mind, we can vectorise our unbalanced panel response variable $y = (y_{11}, \dots, y_{1T_1}, y_{21}, \dots, y_{2T_2}, \dots, y_{NT_N})'$ and the cohort and age effects $c = (c_1, \dots, c_C)'$, $a = (a_1, \dots, a_A)'$ and write down our problem as follows

$$\min_{a,c,\beta} V = (y - S_a a - S_c c - z\beta)'(y - S_a a - S_c c - z\beta) + \alpha_a a' P_a{}' P_a a + \alpha_c a' P_c{}' P_c a \quad (1.5)$$

where the second-difference matrices P_a, P_c are defined as above and selection matrices S_a, S_c map the vectors of age and cohort effects into corresponding panel response variable observations. We add the option of including observable explanatory variables z with a fixed linear association β to the response variables.

The solution to this problem is as follows

$$\begin{bmatrix} c \\ a \end{bmatrix} = x = [I - (S'S + P'P)^{-1} S'Z(Z'Z)^{-1}Z'S]^{-1} \times (S'S + P'P)^{-1} S'[I - Z(Z'Z)^{-1}Z']\, y \quad (1.6)$$

$$\beta = (Z'Z)^{-1}Z'(y - Sx) \quad (1.7)$$

$$S = [S_c \quad S_a] \quad (1.8)$$

$$P = \begin{bmatrix} \alpha_c P_c & 0 \\ 0 & \alpha_a P_a \end{bmatrix} \quad (1.9)$$

Appendix 2 – Potential bias in bivariate age-cohort model

Suppose the true APC model is $y = X_a a + X_p p + X_c c + \varepsilon$ where $\{a, p, c\}$ are vectors of APC effects and $\{X_a, X_p, X_c\}$ matrices of dummy variables. The bias in a bivariate model containing only age and cohort effects is given by

$$E \begin{bmatrix} a \\ c \end{bmatrix} = \begin{bmatrix} a \\ c \end{bmatrix} + ([X_a\, X_c]'[X_a\, X_c])^{-1}([X_a\, X_c]'X_p)\, p$$

and accordingly depends on the covariance of the period effects with age and cohort effects, and the size of the period effects.

Figure A3 below calculates the bias in the age and cohort effects if there is a 1 unit time effect in the first and last years of our sample.[18] The figure shows that a negative bias to the age gradient would require a negative time trend, with large negative time effects at the end of the sample. Intuitively, a positive time effect late in the sample (e.g. a generalised rise in pro-Europeanism in 2016) would be picked up as rise in the pro-Europeanism of the recent cohorts, because only they appear only in the recent data. But given that the 2016 age-pro gradient hasn't changed (i.e. the time effect hits everyone by definition), this must push down on the estimated age effect for the young, which has knock-on effects for the rest of the age and cohort effects in order to minimise deviations from the other observations. Conversely, one observes the opposite pattern for a time effect at the beginning of the sample (it's smaller because the sample was smaller that year), then the bias from time effects in the middle of the sample is small.

What does this mean in practice? First, all of our measures of time effects suggest these are small in 2016 and fairly stationary, in contrast to the trend needed to generate our results. Second, this exercise shows what a sceptic would have to appeal to - a pattern of rising time effects - for our result on negative ageing effects to be wrong.

Figure A1 – hypothetical bias from omitted time effects on age and cohort effects in bivariate model

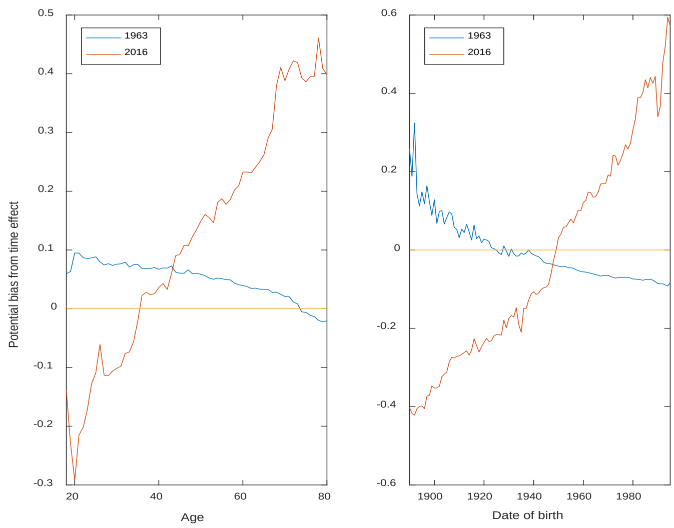

[18] Suppose the true APC model is $y = X_a a + X_p p + X_c c + \varepsilon$ where $\{a, p, c\}$ are vectors of APC effects and $\{X_a, X_p, X_c\}$ matrices of dummy variables. The bias in a bivariate model containing only age and cohort effects is given by $E\begin{bmatrix} a \\ c \end{bmatrix} = \begin{bmatrix} a \\ c \end{bmatrix} + ([X_a\ X_c]'[X_a\ X_c])^{-1}\left([X_a\ X_c]'X_p\right) p$, and accordingly depends on the covariance of the period effects with age and cohort effects, and the size of the period effects

Appendix 3 – Alternative data and estimators

Previous researchers have attempted to solve the problem of perfect collinearity between age, cohort and year effects by proxying time effects using observable variables such as GDP growth. For example, Dohmen et al. (2017) estimate cohort and seasoning effects in risk attitudes by using GDP growth as a proxy for time effects. Since GDP growth is a nonlinear function of time, the problem of perfect collinearity in the APC terms dissolves.

The performance of this approach depends on how good a proxy for time effects is available. In our context it can be plausibly argued that time-varying economic conditions, as reflected in GDP growth or the unemployment rate, shape variations over time in attitudes toward the EU, perhaps because unemployment or slow growth is blamed on the EU, perhaps because unemployment or slow growth affect support for domestic political parties that are pro- and anti-EU, and perhaps for other reasons.

We therefore present an alternative set of estimates using this approach. Unemployment rate and GDP growth are entered both separately and interactively in each year to capture time fixed effects. This allows for direct comparison with the results discussed in the main body of the paper. We find that our central results for seasoning and cohort effects are unchanged.

In addition, we present results when applying our analysis to other datasets: the European Social Survey (ESS) and the British Social Attitudes Survey (BSA). Both the ESS and BSA are repeated cross-sections. While the former has been administered biannually since 2002, the latter has been administered annually since 1983 except for 1988 and 1992. 2016 is the latest wave of the ESS-BSA.[19]

Since they are similar in question-and-answer structure, we merge the subsets of the ESS and BSA containing our variables of interest. The key variable is a binary dummy coding answers to the question "do you want to remain member of the EU/EEC?" This is observed in 1983-87, 1989-91, 1997, and 2013-16. Despite the longer and more frequent time series, the ESS-BSA pooled cross-section provides valid responses from 24,875 individuals, a much smaller sample than the BES.

In both the ESS-BSA and the BES the dependent variable is bounded by zero and one. We therefore adopt a probit model for our estimates.[20] We also estimate logit regressions to test the sensitivity of our results to the choice of link function, finding no significant differences.

---

[19] The 2016 wave was administered in July-October 2016 (small number of interviews in November) for the BSAS, in September 2016-March 2017 for the ESS, and in 24th June-4th July 2016 for the BES.

[20] In Table A2 we show that employing panel data techniques for BES analysis doesn't add much value, it limits rather the comparability to the ESS-BSAS results.

We define cohorts as 5 year intervals. This is a reasonable compromise between accuracy (given by the homogeneity in unobservable characteristics affecting EU sentiment linked to the birth year) and significance (or statistical precision, coming from the size of the cohort).

The model is specified as:

$$(1)\quad ProEU_{it} = \alpha + \beta Age_{it} + \sum_{j=1980}^{2000} \theta Cohort_{j+5,it} + \gamma GDPg_{it} + \delta UR_{it} + \varphi(GDPg \times UR)_{it} + \varepsilon_{it}$$

In addition, we control or a third order polynomial of education:

$$(2)\quad ProEU_{it} = \alpha + \beta Age_{it} + \sum_{j=1980}^{2000} \theta Cohort_{j+5,it} + \mu EduY_{it} + \pi EduY_{it}^2 + \rho EduY_{it}^3 + \gamma GDPg_{it}$$
$$+\delta UR_{it} + \varphi(GDPg \times UR)_{it} + \varepsilon_{it}$$

$proEU_{it}$ is a dummy for being favourable to the European Union

Age_{it} is a variable for the age of the respondent

$GDPg_{it}$ is a variable for real GDP growth rate in the observation's year

UR_{it} is a variable for unemployment rate in the observation's year

$Cohort_{j+5,it}$ is a dummy variable for the respondent's 5-year cohort

$EduY_{it}$ is a variable for the respondent's number of years in full time education

The earliest BES datapoint is 1963, 20 years earlier than first available ESS-BSAS observation. This introduces a different relationship observed between GDP growth and unemployment rate in the years and after the early 1980s. In the 1970s the UK experienced an atypical increase in unemployment coincident with an acceleration in GDP growth. Including pre-1980 BES data therefore produces different point estimates on these variables in the two full data sets. In Table A3 below, we also provide estimates using only post-1980 data. That table shows that none of the other BES estimates is sensitive to truncating the sample this way.

Table A1 shows the results for the full time span including data for the 1960s and 1970s, substituting GDP growth, unemployment rate and their interaction for the time effects, while Figure A4 shows the corresponding marginal effects.

Table A1 - Proxy time effects regressions including 1960s and 1970s data

	Probit			
	(1)	(2)	(1)	(2)
VARIABLES	Pro-EU	Pro-EU	Pro-EU	Pro-EU
Age	-0.0473***	-0.0481***	-0.0562***	-0.0563***
	(0.000808)	(0.000778)	(0.000841)	(0.000805)
Years of full time education			-0.0947***	-0.196***
			(0.0325)	(0.00906)
Years of full time education^2			0.0179***	0.0234***
			(0.00189)	(0.000590)
Years of full time education^3			-0.000437***	-0.000529***
			(3.35e-05)	(1.25e-05)
GDP growth	0.128***	0.118***	0.118***	0.107***
	(0.0139)	(0.0137)	(0.0143)	(0.0141)
Unemployment rate	0.0186***	0.00810	0.00713	-0.00322
	(0.00555)	(0.00531)	(0.00571)	(0.00546)
UR x GDPg	-0.0232***	-0.0219***	-0.0232***	-0.0217***
	(0.00165)	(0.00162)	(0.00169)	(0.00166)
Pre1983		-1.876***		-2.105***
		(0.0471)		(0.0486)
Pre1983 x GDP growth		0.0617**		0.0126
		(0.0243)		(0.0247)
Pre1983 x UR x GDPg		-0.0871***		-0.0704***
		(0.00752)		(0.00764)
Constant	1.612***	6.599***	0.850***	7.854***
	(0.0497)	(0.339)	(0.181)	(0.340)
Observations	165,581	178,166	159,459	171,984
Cohort effects	YES	YES	YES	YES
Time span	1983-2016	1963-2016	1983-2016	1963-2016
Survey	BES	BES	BES	BES

Observations weighted by the size of the pooled sample
Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

In line with the results in Section 3, age is negatively associated with pro-EU sentiment. The age effects are stable over time.

Although we trim the data at 80 years of age in our regressions, as in the main body of the paper, we investigated whether an increasing skepticism towards the EU persists also in later age. The negative age effect tends to reverse after 80 years. This is consistent with what is considered in economic literature as the threshold above which wealth tends to dramatically increase survival rates. [21] This probably occurs because while wealth and education are correlated, our measure of education may not fully capture the status of an individual in society and therefore true economic wealth. There are different types of university bachelor courses, for example, with different prospects of future

[21] This last pattern is particularly evident in the BES dataset.

employment and income. Unfortunately, this variation is not captured in our data. Neither the BES nor ESS-BSA provides the income or wealth level of respondents. This prevents us from assessing the size of the underlying wealth effect.[22]

Cohort effects are also significant as a group and individually, although there are too many to report here.

Finally, pro-EU sentiment is positively associated with education in both datasets, with the magnitude of the effect diminishing with years of schooling. When controlling for education with a third degree polynomial, however, the individual coefficients become more sensitive to the underlying shape of the educational distribution (Figure A2). This explains the differences in the estimated coefficients across the two datasets.

Figure A2 - years of education by age

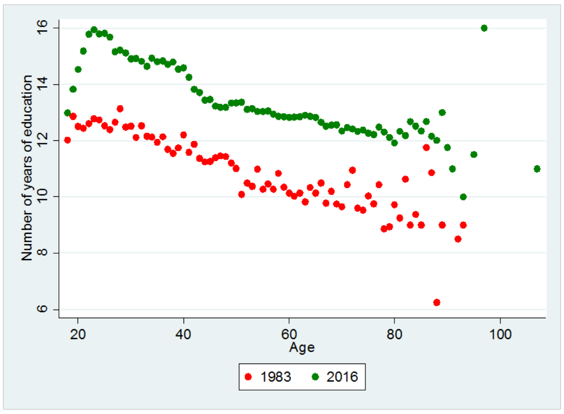

BES data

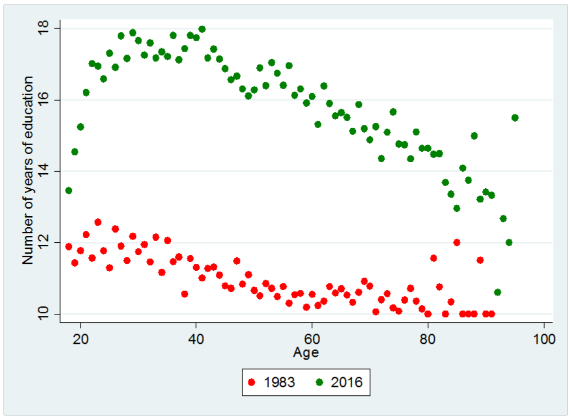

ESS-BSAS data

As noted in Section 4, years of schooling vary with age. Figure A2 shows how access to education has been increasing in these alternative data sets, with older residents showing average levels of education lower than younger generations, except for young individuals not yet having yet completed school.

The bias associated with the exclusion of education from regression (1) varies from variable to variable. For obvious reasons, a positive relationship between age and number of years of education holds most strongly for young voters. So long as the average respondent completes his or her

[22] There are of course questions related to the individual's wealth, e.g. homeowner or tenant, employed or unemployed, size of house etc. Conceivably, these could be used to construct a personal wealth index. But insofar as our goal is to predict the evolution of sentiment toward the EU, forecasting the evolution of wealth would be difficult – even more difficult than forecasting demographic structure or educational attainment.

education by 25 years of age, in every new survey wave from when an individual is 18 years or older we will detect a tendency for years of education to rise with age.[23]

Controlling for education also influences the marginal cohort effects, given how access to education has improved over time (so that educational attainment and birth cohort are correlated). As shown in Figure A6, including education reduces the marginal cohort effects by as much as 4.4 percentage points. The 1990 cohort is the most negatively affected, while the 1890 cohort (recall that we are analyzing surveys taken as early as 1980) is the only one to benefit. The reduction in negative impact from 1990 onward is due to the fact that most recent cohorts are still in the process of completing their education and therefore show lower levels of educational attainment than pre-1990s cohorts.

Figure A3 - hypotetical omitted variable bias in cohort effects arising from exclusion of education

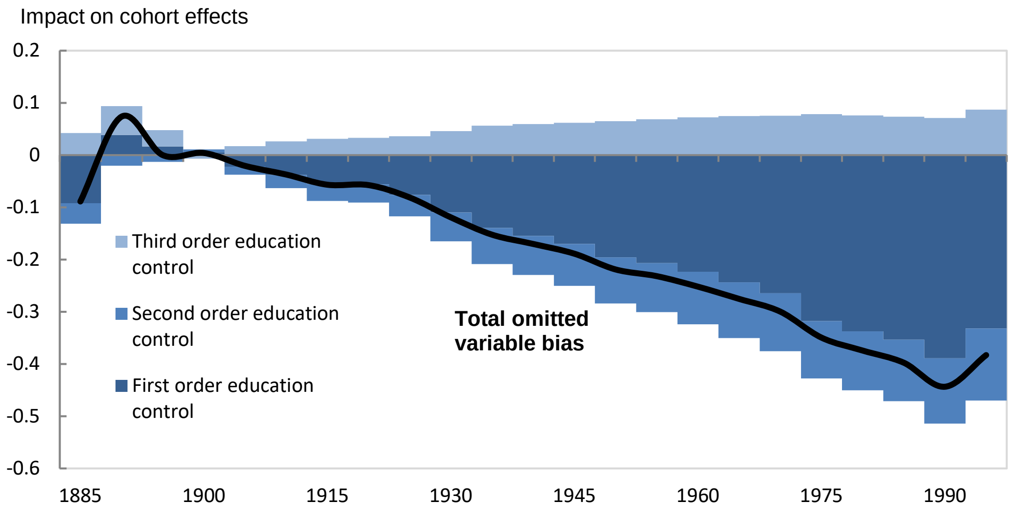

BES dataset. The size of OVB is obtained from linear regressions specified in the same way as those shown in Table 1

[23] This is why the estimated age effect is larger (in absolute value) in regression (2) than regressions (1). .

Figure A4 - Visual display of marginal effects

ESS-BSAS - Probit prediction | *BES* - Probit prediction

...by age

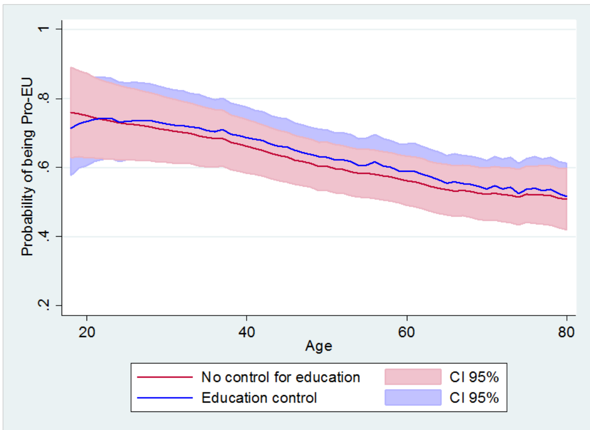

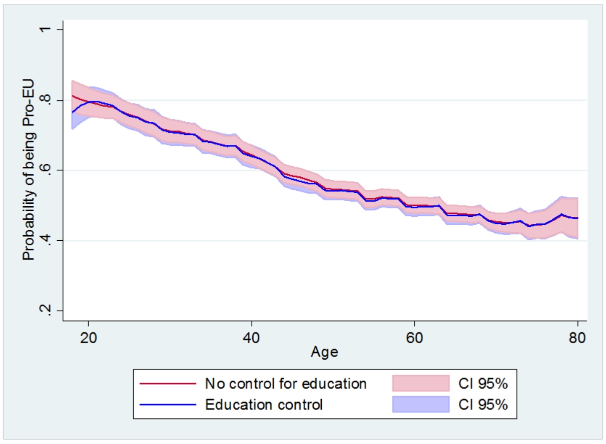

...by cohort

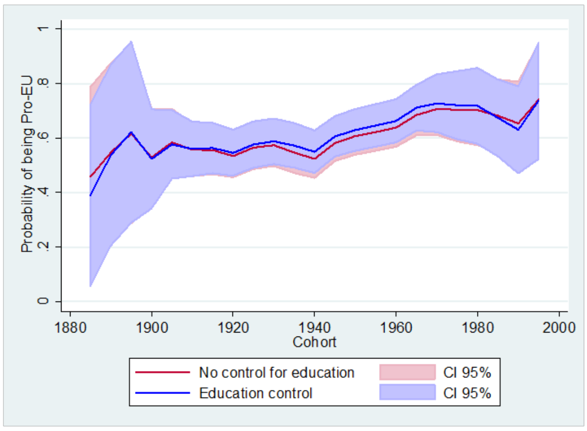

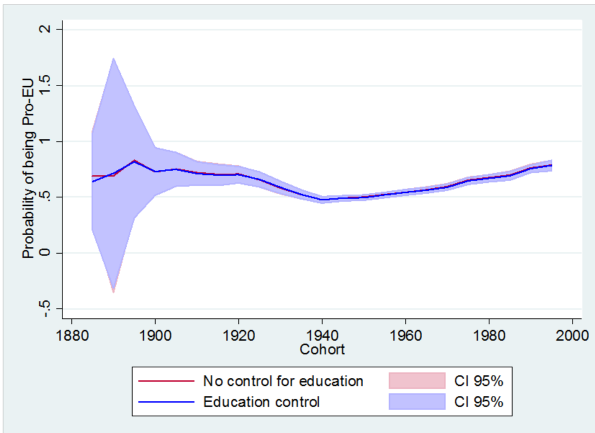

...by year

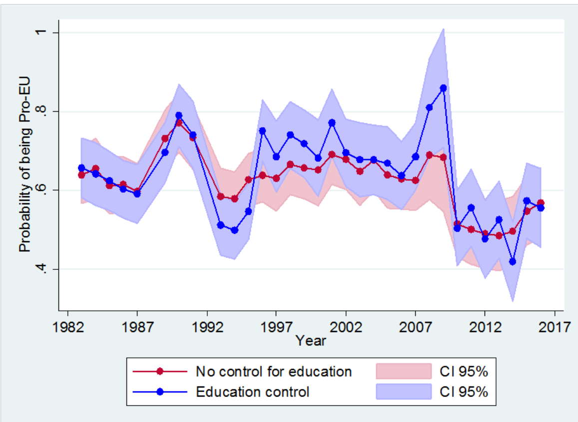

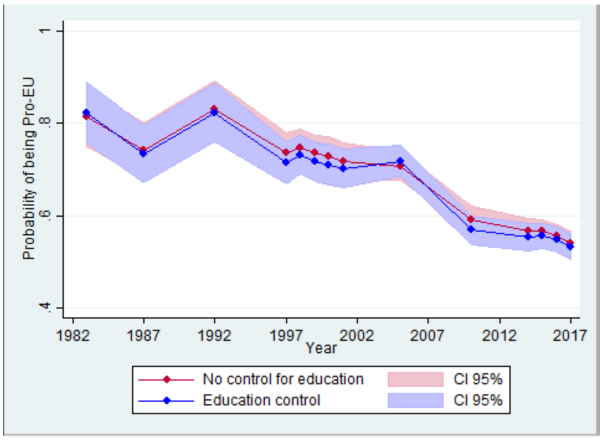

Table A2 - BES panel data analysis with random and fixed effects based on ID code

	(1)	(2)	(3)	(4)
VARIABLES	OLS Fixed effects	OLS Random effects	Probit Population-averaged	Probit Random effects
	Pro-EU	Pro-EU	Pro-EU	Pro-EU
Age	-0.0142***	-0.0157***	-0.0422***	-0.185***
	(0.000863)	(0.000273)	(0.000782)	(0.00290)
Years of full time education	-0.0186	0.0265**	0.0128	0.265**
	(0.0260)	(0.0118)	(0.0351)	(0.116)
Years of full time education^2	0.00102	0.00202***	0.00858***	0.0211***
	(0.00144)	(0.000685)	(0.00203)	(0.00669)
Years of full time education^3	-1.61e-05	-7.15e-05***	-0.000242***	-0.000727***
	(2.44e-05)	(1.21e-05)	(3.56e-05)	(0.000117)
GDP growth	0.0637***	0.0469***	0.127***	0.120***
	(0.00535)	(0.00368)	(0.0102)	(0.0350)
Unemployment rate	-0.0111***	-0.00865***	-0.0198***	-0.0899***
	(0.00164)	(0.00131)	(0.00362)	(0.0124)
UR x GDPg	-0.0103***	-0.00846***	-0.0235***	-0.0463***
	(0.000621)	(0.000442)	(0.00123)	(0.00421)
Constant	0.386	0.551***	4.290***	1.169*
	(0.293)	(0.0655)	(0.553)	(0.641)
Observations	159,459	159,459	159,459	159,459
Cohort effects	YES	YES	YES	YES
Time span	1983-2017	1983-2017	1983-2017	1983-2017
Survey	BES	BES	BES	BES

Observations weighted by the size of the pooled sample
Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table A3 – Results with Data Limited to Post-1980 Period

	Probit			
	(1)	(2)	(1)	(2)
VARIABLES	Pro-EU	Pro-EU	Pro-EU	Pro-EU
Age	-0.0359***	-0.0550***	-0.0473***	-0.0562***
	(0.00133)	(0.00150)	(0.000808)	(0.000841)
Years of full time education		0.144**		-0.0947***
		(0.0645)		(0.0325)
Years of full time education^2		0.000275		0.0179***
		(0.00395)		(0.00189)
Years of full time education^3		-5.76e-05		-0.000437***
		(7.71e-05)		(3.35e-05)
GDP growth	-0.125***	-0.465***	0.128***	0.118***
	(0.0399)	(0.0423)	(0.0139)	(0.0143)
Unemployment rate	-0.117***	-0.206***	0.0186***	0.00713
	(0.0138)	(0.0144)	(0.00555)	(0.00571)
UR x GDPg	0.00770*	0.0422***	-0.0232***	-0.0232***
	(0.00439)	(0.00463)	(0.00165)	(0.00169)
Constant	4.813***	5.892***	1.612***	0.850***
	(0.265)	(0.434)	(0.0497)	(0.181)
Observations	24,874	24,796	165,581	159,459
Cohort effects	YES	YES	YES	YES
Time span	1983-2016	1983-2016	1983-2017	1983-2017
Survey	ESS-BSAS	ESS-BSAS	BES	BES

Observations weighted by the size of the pooled sample
Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1