Age Isn't Just A Number: A Comparative Age-Period-Cohort Analysis of Political Ideology in Western Europe

Phil Swatton
University of Essex
philswtton@gmail.com

Abstract: Age, period, and cohort effects all represent possible long-term drivers of both change and stability in political ideology in the electorate. However, the question as to the extent that the explanations these effects offer is consistent across countries has not been addressed. I therefore perform a comparative APC analysis of left-right political ideology. I run two side-by-side APC analyses of left-right positions, constraining effects to be common across countries in one and nesting them within country in the other. I pay special attention to the issue of how measures of ideology can be meaningfully compared, and develop a measure of relative ideology. I find evidence for ageing effects and life-cycle effects. Moreover, I find that while the constraint of common cohort effects is not a strong one, the constrain of common period effects is over-strong. Future research should focus first on better understanding this contrasting result, and second on developing absolute measures to better understand patterns of change and continuity in left-right ideological positions in the public.

Working paper, subject to change. If citing, please ensure that you include the version date. Version 28-01-2024

1 Introduction

A week is both a long and a short time in politics. It is a long time in the sense of events: unexpected occurrences can dramatically shift electoral realities in a matter of hours. However, politics is also slow: outside of day-to-day drama, patterns are slow to change in politics. Recent examples might include the rise of the radical right or the growing importance of educated voters in the electorate (see Ford and Jennings, 2020). These trends developed in the course of years, even if they are more visible in the drama of particular days.

It is not for nothing therefore that political scientists should be interested in the dynamics of long-term stability and change in political ideology and behaviour. In this paper, I am focussed on the former. One question that has continually perplexed political scientists is whether long-term trends and changes in political ideology are better explained by ageing effects, period effects, or cohort effects. The first describes processes of psychological ageing and predictable changes during the individual life-cycle. The second describes the effect of the 'mood of the moment': a transient effect which changes through time. The last describes the long-lasting effect of early influences that remain with a generation.

One reason APC analysis has presented greater difficulty than other statistical analyses is the difficulty involved in disentangling the three effect types. In their linear, continuous forms the three effects are perfectly multicollinear with one another. As a consequence, APC analysis has developed not merely as a form of substantive research but also has an entire sub-methodology associated with it. The primary feature of this sub-methodology has been a strong emphasis on functional form and modelling assumptions has emerged.

Extant APC analyses have thus far broadly found in favour of cohort effects on left-right positions. However, the question remains as to the extent that

similar countries enjoy similar cohort and period trends in political ideology. Many of the longer-term events experienced in politics are experienced cross-nationally, from the social democratic moment in the post world war 2 era to today's present political trends. It is therefore reasonable to wonder as to the extent that these trends are common across nations. Thus far, prior analyses have typically been confined to single country case studies and have made modelling and theoretical assumptions that mean their results do not generalise. In the one case where a comparative analysis has been performed, aggregate cohort effects were not reported, and two cohort effects were included in a single model. I therefore seek to assess in a comparative context the extent to which common cohort and period trends occur, and the extent to which APC analyses help us to understand long-term patterns of change and continuity in left-right ideology.

I begin with a general outline of APC analysis. I discuss the substantive interpretations of the three effect types, the identification problem that arises between them, and the results in past APC analyses of political ideology. I then turn to the issues of defining and measuring ideology in comparative longitudinal research. When making comparisons across contexts, the issue of differences in meaning arises. I discuss two approaches that arise from this: relative and absolute ideology. I further discuss survey measurement issues such as differential item functioning that commonly arise in survey contexts. From here, I proceed to outlining my methodological approach. I use Aldrich-McKelvey scaling to produce a DIF-corrected measure of left-right positions. From here, I use Hierarchical APC models to perform a comparative APC analysis. As part of this, a decision needs to be made regarding the treatment of country contexts. Should cohort and period effects be constrained to be similar across countries, or should they be nested within countries? Lacking a good a priori justification

either way, I run both models and compare the results.

I find that there is good evidence for ageing, life-cycle, and cohort effects in terms of relative ideological positions. There are few differences in the inferences made between countries in terms of cohort effects regardless of the model specification. This implies that Western Europe does experience common socialising effects. This goes against the interpretation of past country case studies, which have interpreted cohort effects in terms of country-specific political influences. By contrast, period effects appear to differ much more substantially from country to country once freed to do so, implying that here a country-specific interpretation is more likely correct. I conclude by arguing that future APC and comparative research should focus on the development of absolute measures of ideology, so that the results presented here can be better understood.

2 Age-Period-Cohort Analysis

The purpose of age-period-cohort (APC) analysis is to distinguish between age, period, and cohort effects on a dependent variable of interest. In this section I discuss the theoretical distinctions between these effects and give an overview of research results regarding APC effects on ideology.

2.1 Theorising APC

'Age effects' broadly captures two theoretical perspectives on the role of ageing. In the first, the physical process of ageing is the causal variable of interest (Glenn, 1974). In political science, this will in typically be understood as psychological ageing. In the second, the individual's progression through stages of the life-cycle represents the causal variable of interest. Here, predictable changes over an individual's life-cycle such as marriage, increases in income, promotions at work, children, and home ownership have an effect on their political ori-

entations (Glenn, 1974; Tilley, 2005; Tilley and Evans, 2014). In both cases, chronological ageing is an imperfect correlate of these processes and never the actual quantity of interest (Glenn, 1974).

'Cohort effects'[1] instead emphasise persistent generational differences. Within the social sciences, these effects are typically considered in terms of 'socialisation', where individuals are socialised into holding certain views as a consequence of influences from their formative years which are retained over time (see Mannheim, 1970; Dawson and Prewitt, 1968). Various potential sources of socialisation have been identified, including historical events (Mannheim, 1970), parental influence (Campbell et al., 1960; Butler and Stoke, 1974), education (Stubager, 2008; Surridge, 2016), and peer groups (Hooghe, 2004). Of primary interest here however is the role of the first: historical events. This is because this type of socialisation represents the closest correspondence to the notion of cohort effects - lasting formative influences unique to a given generation. Past research supports the notion that the long-term influence of historical events is strongest during the formative years between adolescence and young adulthood (Jennings, 2007, p. 35; Rekker, 2016, p. 121).

'Period effects' represent the effect of a given time period on the dependent variable of interest. Similarly to age and cohort effects, interest is not in the chronological time period itself but rather in the predominant features of that time period (Glenn, 2005). Unlike cohort effects however, its influence is taken to be temporary, rather than lasting. While interest is sometimes in age alone, while in other cases interest is in all three effects as potential explanations for long-term patterns of stability and change, ts inclusion is necessary in all cases. This is because of the fact that the three effects are potential mutual confounders of one another. This mutual confounding gives rise to the identification problem

[1]For clarity, 'cohort' and 'generation' can be considered interchangeable for the purpose of this paper

which characterises APC analysis.

In all three cases, the numbers associated with these effects - age, birth year, time period - are simply ways of capturing these effect types. But these numbers are not the source of our interest in themselves, and hence giving these effects careful intereptation during is an important task of APC analysis.

2.2 Identification

The APC identification problem emerges from the fact that in their continuous, linear forms

$$C = P - A \tag{1}$$

where A is age, P is the time period, and C is cohort membership (i.e. birth year). None of these effects are necessarily mutually exclusive, but in this format they possess perfect multicollinearity with each other. Due to this, a unique solution does not exist and thus the model cannot be estimated. All APC analyses must therefore tackle this identification problem with some set of assumptions regarding some or all of the three effects that allows for statistical identification of the model.

Over time, a variety of 'solutions' have been proposed. Broadly speaking, these typically require that some kind of assumption regarding the APC effects are required (Bell, 2020). The weakest assumptions typically focus on functional form, often assuming that period and/or cohort effects are non-linear in nature. The strongest typically require assuming that one of the three effects is 0 and thus can be ignored (Bell, 2020). A recent controversial methodology is hierarchical APC (HAPC) models, which estimates cohort and period effects as random effects. Although hotly debated (see full discussion below), this requires assuming non-linear effects in cohorts and period.

2.3 APC Analysis of Political Ideology

In political science, APC analysis has typically been used to assess long-term patterns of continuity and change in political ideology and behaviour. In this paper, I focus on the former. Broadly speaking, past research has tended to find in favour of the presence of cohort effects. How this this should generalise across countries and contexts - if at all - is however not always consistent across studies.

In terms of single-country case studies, a large number of APC analyses on ideology have been performed in the United Kingdom. The earliest of these is an analysis by Tilley (2005) on 'Libertarian-Authoritarian' attitudes. Tilley assumes away a psychological process for aging and thus includes only life-cycle indicators rather than age in itself. Tilley finds that cohort effects and not life-cycle effects or cohort composition drive age differences in these attitudes. Newer generations are increasingly libertarian over time - in line with general political changes over time. Here, it is clear that broadly we should expect similar findings in other nations.

More recently, Grasso et al. (2019) similarly find in favour of cohort effects and against ageing effects. Grasso et al. use a generalised additive model (GAM) methodology, first developed by Grasso (2014). Unlike Tilley however, Grasso et al. find increasing right-wing and authoritarian attitudes among those coming of age during the Thatcher and Blair years. Similarly unlike Tilley who uses simple five-year groupings, Grasso et al. group survey respondents' cohort memberships according to political periods distinct to the United Kingdom. In a single-country case study however, it is impossible to distinguish between general versus country-specific cohort trends.

In a comparative context, Down and Wilson (2013, 2017) perform an APC analysis on support for the European Union. Across their two papers, they

distinguish between utilitarian and affective attitudes to the European Union and find especially strong cohort effects for the latter, though they are present for both. Generally speaking, more recent generations are ceteris paribus more in favour of European Integration as they have been raised in a context where it was more established.

Shorrocks (2018) examines cohort differences in gender gaps in left-right ideology in Europe and Canada. Shorrocks finds cohort effects in this gap not captured by aggregate-level analysis: older women tend to be more right-wing than older men, while younger women tend to be more left-wing than younger men. However, Shorrocks's model combines two cohort trends: a linear cohort trend interacted with gender, and the cohort random effect more typical of HAPC methodology. The inclusion of a linear cohort trend is justified only on the grounds of creating the interaction, rather than through an assumption of linearity. Nor is it clear how the inclusion of two separate cohort trends should be theoretically understood. Moreover, Shorrocks only reports the changing gap between the genders: not the actual cohort or period effects.

In all of these papers, either comparative analysis is not performed, or where it is performed cohort effects are either constrained to be the same across countries or not but the effect of such a constrain (or its absence is never explored). In many of the single-country case studies, cohort specifications rest on theoretical justifications specific to that country, and so the results cannot be understood to generalise. I therefore aim to fill this gap, by performing an APC analysis of left-right ideology in a comparative setting. I aim not only to learn which the APC effects act as drivers of left-right ideology, but also the extent to which trends are common across countries within Western Europe.

3 Political Ideology in a Comparative Setting

Political ideology is a difficult concept to utilise in a comparative setting. Ideology is an abstract concept, and the precise meaning of left and right varies between contexts. Since APC analysis always requires several different time periods (by definition), and the extension to comparative research requires the addition of multiple country contexts, this is a clear problem for APC analysis of political ideology. In this section, I therefore address the multiple issues that arise from attempting to quantitatively measure ideology in a comparative setting.

3.1 Defining Ideology

One approach to defining political ideology is to emphasise the role of *constraints*. The seminal work on this is Converse (1964), where ideology is defined as a '*configuration of ideas and attitudes in which the elements are bound together by some form of constraint or functional interdependence*'. If we know one opinion an individual holds, we are better placed to guess another opinion they hold. Peffley and Hurwitz (1985) expands this model of constraints to a hierarchical model of public opinion, wherein an overall left-right dimension drives domain attitudes which in turn drive specific issue attitudes. This is visualised in figure 1 below:

Analogous to the hierarchical model of public opinion is the basic space theory of political ideology. Here, also following on from Converse's model, where individuals have structured belief systems they can be represented in a low-dimension space (Poole, 1998). This low-dimensional space was referred to as the basic space by Ordeshook (1976) and as a predictive dimension by Hinich and colleagues (Hinich and Pollard, 1981). The primary difference between the hierarchical model and the basic space theory is that where the hierarchical

Figure 1: Hierarchical Model of Public Opinion

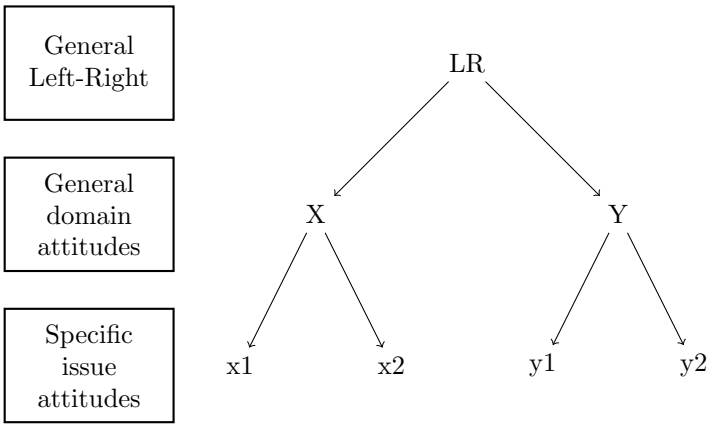

model assumes causal pathways from more abstract dimensions to more concrete dimensions, the basic space theory is causally agnostic and treats the abstract dimension as a representation of several smaller ones.

There does exist some debate as to the exact extent this holds true. Using an ordinal probit model with random effects, Lauderdale, Hanretty and Vivyan (2018) find that just 1/7th of variation in survey responses corresponds to a summary dimension as described above. Another 3/7ths corresponds to idiosyncratic variation, while the last 3/7ths corresponds to response instability (Lauderdale, Hanretty and Vivyan, 2018). By contrast, using a novel mixture model methodology (Fowler et al., 2022) find that approximately 70% of US citizens are in fact best described by a single ideological dimension, while another 30% are either idiosyncratic[2] or simply random in their responses.

Both studies were conducted in the US, and focus on slightly different quantities - the percentage of variation in responses versus how respondents are best described. Nonetheless, these different methodologies provide vastly different answers as to how well a single dimension does describe variation in ideology.

[2]i.e. consistent in their views, but not in the traditional left-right manner

The important point however, is that to some extent or other we can usefully speak of a left-right summary dimension. I do not require that the causal pathways of the hierarchical model hold: only that ideology can be usefully described in in the higher dimensionality of left and right in the manner described above. There are merits in conducting (comparative) APC analyses on other dimensions as have indeed been done - but my focus here is on the left-right dimension.

3.2 Relative and Absolute Ideology

In APC analysis and many other comparative analyses, survey respondents from several time periods are pooled together. In a comparative APC analysis, they are pooled from both several time periods and several countries. Given the above conceptualisation of left-right ideology as a summary dimension defined by constraints above, the question emerges how to approach this comparatively. In a given context, the relevant issues, and therefore issue domains will differ. Similarly, the weights on those issues and domains up to the left-right dimension will also differ. The EU is much more important to left-right ideology in the UK in 2019 than in say, 1997. Likewise, French voters who see Marine Le Pen as being to the right of Emmanuel Macron are not doing so on the grounds of their respective economic positions.

I therefore introduce the concepts of relative and absolute ideology. With 'absolute' ideology, the meaning and interpretation of measures of a left-right dimension should be constant over time and space. Following the above discussion, raw survey data does not typically meet this criteria. It will therefore be necessary to rescale the data such that all data points share an interpretation. One method of doing this is via anchoring vignettes (King et al., 2004; King and Wand, 2007; Hopkins and King, 2010). Here, vignettes are provided and scaled

by the respondents. Since the information in the vignettes are constant over time, this information becomes the 'anchor' against which all other data can be rescaled. Other options for rescaling the data may be discovered in the future: the important point is that there is some piece of information that is constant over time that raw data can be rescaled against. Likewise, the method of choice for producing this rescaling is immaterial, and many such methods exist (see e.g. Wand, King and Lau, 2011; Bakker, Jolly, Polk and Poole, 2014; Bakker, Edwards, Jolly, Polk, Rovny and Steenbergen, 2014).

By contrast, when measuring ideology in the 'relative' sense, each context is allowed to retain its own meaning and interpretation of the left-right dimension. This should ideally however be relative to a reference point. We may for example seek to use data measured relative to the political center of that context. One of my contentions in this paper is that raw survey data is in fact relative in nature, but implicitly rather than explicitly so. In this paper, I use this type of comparative ideology - largely due to data constraints. I develop my measure of relative ideology by rescaling data with respect to standardised party distributions (not too distant from the idea of scaling against the political center of a given context). I discuss this further in the methodology section.

The conceptual discussion here raises a question: will absolute versus relative ideology produce different results? It seems obvious that it should be so: if the interpretation of a variable changes over time, then surely the nature of the relationship between other variables and itself should also change over time. In APC analysis this seems especially pronounced: are people becoming more right-wing as they age, or is the political system shifting 'left' around them? Without access to measures of both, we cannot answer this question. We can however still gain some interesting results from a single measure, but we must be careful in how we interpret those results. I now turn to a final measurement

concern for this paper: differential item functioning.

3.3 Differential Item Functioning

Closely related to the problem of item meaning is the issue of scale perception. Differential item functioning (DIF) is a measurement issue that arises when for the same underlying perception, different survey respondents give different answers (King et al., 2004). Even where two respondents come from the same context, and share the same underlying meaning of left and right, and share the same underlying ideological position, DIF means that they will place themselves on the survey scale in different locations to one another. Insofar DIF is purely random, this will result in attenuation bias. Insofar as there are systematic patterns in DIF, this will result in biased results. And indeed, more politically informed respondents may for instance use a survey scale in a different manner to less informed respondents.

Along with survey item meaning, researchers seeking to study ideology must also contend with the problem of DIF. Solutions typically focus on finding some objective external *anchor* on which to rescale responses (King et al., 2004). Several methods for solving DIF exist, but the broad concept is the same. If respondents are asked to locate one or more external stimuli on the same scale, these placements can be used to reveal both the 'true' location of the external stimuli and through this a corrected measure of the respondent's location can be produced. One popular approach to this is the use of anchoring vignettes (King et al., 2004; King and Wand, 2007; Hopkins and King, 2010). Another approach is to use real-world stimuli, such as political parties or elites (Aldrich and McKelvey, 1977; Poole, 1998; Hare et al., 2015). It is this latter approach that I take in this paper.

4 Methodology

4.1 Case Selection and Data

To perform a comparative APC analysis of left-right ideology, I utilise the Comparative Study of Electoral Systems (CSES) integrated module dataset. This is a dataset collected alongside the election studies of participant countries, meaning that survey responses are always from the context of the heat of an electoral campaign. I chose this dataset for three reasons. First, it spans a good number of years: from 1996 to 2016. This is essential for APC analysis, which requires a large time span in order for analysis to be effective. Secondly, it contains a self-reported measure of left-right positions which can be corrected for differential item functioning. Third and finally, as I discuss below, this correction process can be used to produce a clear measure of relative ideology.

I select Western Europe as a case study because it is a solid testing ground for the notion that cohort and period effects may be similar between countries sharing common political trends. Not only does Western Europe broadly enjoy this reality, it also represents a set of country cases that should be reasonably similar in terms of the relationship between age, generation, and political ideology. Introducing post-soviet countries could potentially introduce very different dynamics in terms of the relationship between age and political ideology. After filtering for required survey questions, the 15 countries included for analysis are Austria, Belgium, Denmark, Finland, France, Germany, Great Britain, Iceland, Ireland, Netherlands, Norway, Portugal, Spain, Sweden, and Switzerland. Any West European countries not included in the analysis are either not included in the CSES, or did not contain the requisite survey questions.

4.2 Measuring Left-Right Ideology

As discussed above, when measuring ideology in a comparative context the issue of survey item meaning emerges. I approach both this problem and the problem of DIF simultaneously by utilising Aldrich-McKelvey (AM) scaling (Aldrich and McKelvey, 1977). Aldrich-McKelvey scaling is a methodology developed to correct differential item functioning in respondent placements of political parties. As part of the scaling process, respondent-specific parameters are recovered which can in turn be used to generated DIF-corrected measures for survey respondents on the same scale. In other words, the corrected political party placements are used as external anchors, respondent-specific parameters are generated by regressing these on the respondent's own placements, these parameters are then applied to respondents' self placements.

DIF can thus be corrected by running Aldrich-McKelvey scaling within each country-year subsample. The question then remains how this might become a true measure of relative ideology. When measuring political party positions, it is common practice to standardise them due to the absence of a natural 0 point (see e.g. Hanretty, 2022). Since Aldrich-McKelvey scaling returns political party placements with a mean 0 distribution by construction (Aldrich and McKelvey, 1977), all that remains is to divide party positions by their standard deviation. Since the respondents' recovered placements are on the same scale, respondents' positions can similarly be divided by the party standard deviation. The interpretation of respondents' positions is then their placements relative to their country's standardised party system in that year.

In practice, I do not expect strong differences between models based on the raw survey data and the rescaled relative data. This is in part because DIF is largely treated as a noisy process in Aldrich-McKelvey scaling, and so should not particularly alter the measure other than to remove some noise. It is also

because the raw survey data is in itself a form of relative data. Both methods therefore capture a contextual ideology and should not particularly differ in the results they provide. To test the notion that both the raw data and the scaling data provide results relative to that given context, I therefore present models with both the raw data and the rescaled data from the Aldrich-McKelvey procedure outlined above. In the case where the raw data is used, I have standardised it across the entire dataset to be mean 0 and standard deviation 1. This transformation does not affect the relationships between variables, but will make the recovered parameters more similar in size to the rescaled data and thus easier to compare.

For the purposes of Aldrich-McKelvey scaling, the respondents' left-right placements of political parties in CSES were used. Only parties where at least 40% of respondents had placed the party were utilised for this purpose. This was a fairly arbitrary choice. The threshold was chosen to be large enough to remove parties that very few respondents placed, while remaining small enough to avoid removing too many political parties from the scaling procedure. This did however necessitate further filtering of respondents for missing data, as Aldrich-McKelvey scaling requires that respondents place all political parties used.

4.3 The Hierarchical Age-Period-Cohort (HAPC) Model

The most recent - and controversial - development in APC research is the hierarchical age-period-cohort (HAPC) model (Yang and Land, 2006, 2008, 2013). In HAPC models, the cohort and period effects are assumed to be non-linear

and are modelled as random effects. This is given by (2):

$$
\begin{aligned}
Y_{ijk} = \boldsymbol{X}_{ijk}\boldsymbol{\beta} + cohort_j + period_k + u_{ijk} \qquad (2)\\
cohort_j \sim N(0, \sigma_{cohort})\\
period_k \sim N(0, \sigma_{period})\\
u_{ijk} \sim N(0, \sigma_u)
\end{aligned}
$$

where Y_{ijk} is the outcome of interest, $\boldsymbol{X}_{ijk}$ is a vector of covariates for individual i, $\boldsymbol{\beta}$ is the vector of fixed effects, $cohort_j$ is the cohort random effects, $period_k$ is the period random effects, and u_{ijk} is the error term in the model. Note that variation occurs across three levels: individuals, cohort, and time period.

An implication of this is that cohort and period membership are treated as contexts within which individual survey respondents are nested (Yang and Land, 2006, p. 85). Strictly speaking, the model is always identified as the three effects are not linear and additive at the same level of analysis (Yang and Land, 2013). However, there has has been substantial controversy as to whether the HAPC model correctly identifies APC effects. In the first critique, HAPC models allow researchers to estimate a model without properly stating their assumptions around the non-linearity of period and cohort effects (Bell, 2020). Where the model's assumption of non-linear effects are correct, it will be correct. However, a model with linear period and cohort effects will still be estimated.

This leads to the second, older, and more serious critique of the model. Here, the concern is that HAPC models misallocate APC effects despite the apparent breaking of linear dependency (Bell and Jones, 2013, 2014*a*,*b*, 2015). The debate that followed generated more heat than light, but broadly some points of consensus do emerge. First, agreement is found on the treatment of cohort

and period effects as random effects. Second, in the presence of exact algebraic linear effects the HAPC model will fail (Reither, Masters, Yang, Powers, Zheng and Land, 2015; Bell and Jones, 2015; Reither, Land, Jeon, Powers, Masters, Zheng, Hardy, Keyes, Fu, Hanson et al., 2015). The point on which consensus does not exist is the exact conditions under which the model would enter difficulty beyond exact linear dependency. The critics of HAPC models argue that it is enough that period and cohort effects are monotonic (i.e. always increasing or decreasing, rather than merely linear) for collinearity to occur (Bell and Jones, 2014*a*,*b*, 2015). Moreover, it is argued that HAPC models typically find in favour of period effects because there are typically many more years (and thus groupings) covered by the cohort effects (Bell and Jones, 2018).

I utilise the HAPC model because there are good a priori reasons to expect that cohort and period effects are unlikely to be linear (or even monotonic) in practice. The tide of history does sweep forever in one direction, but ebbs and flows in unpredictable ways. If the post world war 2 era was characterised by greater social democracy, it was also characterised by a level of social authoritarianism that today's center right by contrast would not accept. Likewise, even as during the 1970s the neoliberal turn begun, this was concurrent with a shift in more socially liberal directions. It would be surprising to witness linear cohort and period trends both: it is for this reason I favour the HAPC specification. Some questions however remain. First, how should cohort membership and periods be created? Second, how should the inclusion of multiple countries be incorporated into the HACP model?

There are broadly two competing views on the specification of cohort memberships and time periods. The first suggests we should theoretically specify the cohorts and time periods based on a priori knowledge. The alternative perspective, and the one I endorse, takes the view that however cohorts and

periods are defined there will necessarily be arbitrariness at the boundaries at the boundaries between cohorts and periods (Spitzer, 1973, 1355). Hence, many researchers specify cohorts and periods in arbitary five-year groups. Given the well-acknowledged arbitrariness, I therefore adopt this approach.

4.4 Comparative APC

The extension of APC analysis to a comparative setting requires the country contexts survey respondents are nested in to be considered. An advantage of utilising the HAPC methodology is that countries can be treated as another random effect in the model. This is demonstrated in (3) below:

$$
\begin{aligned}
Y_{ijkl} = \boldsymbol{X}_{ijkl}\boldsymbol{\beta} + cohort_j + period_k + country_l + u_{ijkl} \qquad (3)\\
cohort_j \sim N(0, \sigma_{cohort})\\
period_k \sim N(0, \sigma_{period})\\
country_l \sim N(0, \sigma_{country})\\
u_{ijk} \sim N(0, \sigma_u)
\end{aligned}
$$

where the primary addition is the country random effect $country_l$ and the index denoting that level of variation, l.

However, the question of the relationship between country random effects and the cohort and period random effects arises. In the model above, cohort and period effects do not vary by country and thus are constrained to be the same across countries. Given that cohort membership and time periods are treated as contexts within which individuals are nested, it is not immediately clear whether these should be nested within countries or not. In other words, are cohort and period effects unique to the countries in question? I therefore

run both the constrained model in (3) and the nested model below:

$$
\begin{aligned}
Y_{ijkl} = \boldsymbol{X}_{ijkl}\boldsymbol{\beta} + cohort_{jl} + period_{kl} + country_{l} + u_{ijkl} \qquad (4)\\
cohort_{jl} \sim N(0, \sigma_{cohort})\\
period_{kl} \sim N(0, \sigma_{period})\\
country_{l} \sim N(0, \sigma_{country})\\
u_{ijk} \sim N(0, \sigma_{u})
\end{aligned}
$$

Note that the only difference in (4) relative to (3) is that the cohort and period effects are now free to vary by country.

4.5 Gender Generation

Although not the focal point of this paper, Shorrocks (2018) establishes the presence of a gender-generation gap in political ideology that varies on cohort lines. However, as discussed above Shorrocks includes an additional linear cohort term in the fixed portion of the HAPC model and interacts it with gender. Instead, I take an approach more in line with the assumptions and theoretical motivations for specifying a HAPC model and utilise cohort random slopes on gender. This allows the effect of gender to vary from cohort to cohort, without requiring that it vary linearly.

4.6 Life-Cycles

The final set of modelling decisions regards additional variables to include alongside age. Over an individual's life-cycle, many important changes can occur: university, increases in income, children, marriage, home ownership. I include marital status, income, and university education in the models below as important life-event variables that correlate with age. I do not include children or

home ownership as these are not measured in the CSES.

Of these, 'university education' is least straightforwardly interpretable as a life-cycle effect. Higher education is not pursued by everyone, but has been disproportionately pursued by the present younger generation (Ford and Jennings, 2020). Its effect is therefore more a correlate of cohort membership than of age.

5 Results

Table 1 presents the results of the HAPC regression models. The first two models on the left-hand side are the set of models with shared cohort and period effects. The next two models on the right-hand side are those with nested cohort and period effects. Within these groups, the first left-hand model is that with the raw response data as the dependent variable, while the second on the right-hand side is that with the rescaled response data following the Aldrich-McKelvey procedure outlined above. The reference categories are single for marital status, no education for education level, and the 1st (i.e. lowest) income quintile for income level. Results are reported at 3 decimal points to avoid rounding some effects to 0. 95% confidence intervals are reported alongside the parameter estimates, and a star is used to denote when the null hypothesis of 0 falls outside this interval.

The variances and covariances of the random effects are reported in the tables. Plots of predicted random effects for cohort and period random intercepts are presented throughout the main analysis. Since the gender-generation gap is not a focus of my analysis, I present the plots of the cohort random slopes for gender in the appendix of this paper. Likewise, since I am not substantively focussing on the random intercepts of the various countries in the analysis, these are also presented in the appendix of this paper.

Starting with age effects, in all four models the coefficient for age is positive

Table 1: HAPC Results

	Constrained		Nested	
	Raw	Scaled	Raw	Scaled
Age	0.005*	0.005*	0.004*	0.003*
	[0.004; 0.007]	[0.003; 0.007]	[0.004; 0.005]	[0.002; 0.004]
Married	0.064*	0.092*	0.062*	0.091*
	[0.040; 0.089]	[0.066; 0.119]	[0.038; 0.087]	[0.064; 0.117]
Divorced/Separated	0.012	0.042*	0.001	0.029
	[−0.026; 0.049]	[0.002; 0.083]	[−0.036; 0.038]	[−0.011; 0.069]
Widowed	0.092*	0.115*	0.102*	0.120*
	[0.047; 0.136]	[0.067; 0.164]	[0.057; 0.147]	[0.072; 0.169]
Primary Education	−0.034	−0.075*	−0.021	−0.056
	[−0.099; 0.031]	[−0.145; −0.004]	[−0.088; 0.045]	[−0.127; 0.015]
Secondary Education	−0.065	−0.164*	−0.047	−0.124*
	[−0.131; 0.000]	[−0.235; −0.093]	[−0.114; 0.020]	[−0.197; −0.051]
Post-Secondary Education	−0.008	−0.079*	−0.003	−0.050
	[−0.075; 0.059]	[−0.152; −0.006]	[−0.072; 0.065]	[−0.124; 0.024]
University Education	−0.178*	−0.328*	−0.172*	−0.291*
	[−0.245; −0.112]	[−0.400; −0.255]	[−0.240; −0.104]	[−0.365; −0.218]
Other Education	−0.382	−0.247	−0.345	−0.047
	[−0.916; 0.151]	[−0.825; 0.331]	[−0.877; 0.188]	[−0.623; 0.528]
2nd Income Quintile	0.022	0.051*	0.025	0.049*
	[−0.006; 0.049]	[0.021; 0.080]	[−0.003; 0.052]	[0.019; 0.078]
3rd Income Quintile	0.056*	0.071*	0.056*	0.061*
	[0.028; 0.084]	[0.041; 0.101]	[0.028; 0.085]	[0.031; 0.092]
4th Income Quintile	0.096*	0.107*	0.097*	0.095*
	[0.066; 0.125]	[0.075; 0.139]	[0.068; 0.127]	[0.064; 0.127]
5th Income Quintile	0.260*	0.283*	0.260*	0.274*
	[0.229; 0.290]	[0.250; 0.316]	[0.229; 0.290]	[0.241; 0.307]
Var: Cohort	0.007	0.009		
Var: Country:Cohort			0.015	0.016
Var: Cohort (Gender)	0.010	0.008		
Var: Country:Cohort (Gender)			0.023	0.027
Cov: Cohort	−0.003	0.002		
Cov: Country:Cohort			−0.008	−0.006
Var: Period	0.000	0.002		
Var: Country:Period			0.005	0.029
Var: Country	0.036	0.060	0.033	0.035
Var: Residual	0.945	1.112	0.934	1.087
N	55833	55833	55833	55833
AIC	155569.601	164664.656	155335.450	163881.254
BIC	155757.133	164852.189	155522.982	164068.787
Log Likelihood	−77763.800	−82311.328	−77646.725	−81919.627

* Null hypothesis value outside the confidence interval.

and significant at the 95% confidence level. Also positive and significant at the same level across all models, though relatively small, are the coefficients for those who are married and for those who are widowed relative to those who are single. There is therefore some initial evidence here the notions that there are both ageing and life-cycle effects in terms of relative ideology. In other words -

despite potential changes in the political system around them, individuals still move to the (relative) right as they age.

Similarly, across all four models there are positive and significant effects for education and income level. These can less straightforwardly be considered as life-cycle effects, but there are also at present large generational differences - especially in terms of education level. Notably, the effects for university education and the top two quintiles are reasonably large relative to the other effects. Also worth acknowledgement is the fact that unlike other life-cycle effects, education moves individuals to the relative political left rather than right. Overall then, there are strong life-cycle effects in terms of relative ideology. We should not however immediately neglect the coefficient for age on the grounds that it is small: humans enjoy long lifespans, and the shift to the right predicted here will happen over a lifetime. Figure 2 plots the predicted values on the relative left-right scale as someone ages, with other variables set to their mean values.

Figure 2: Predicted Left-Right Position by Age

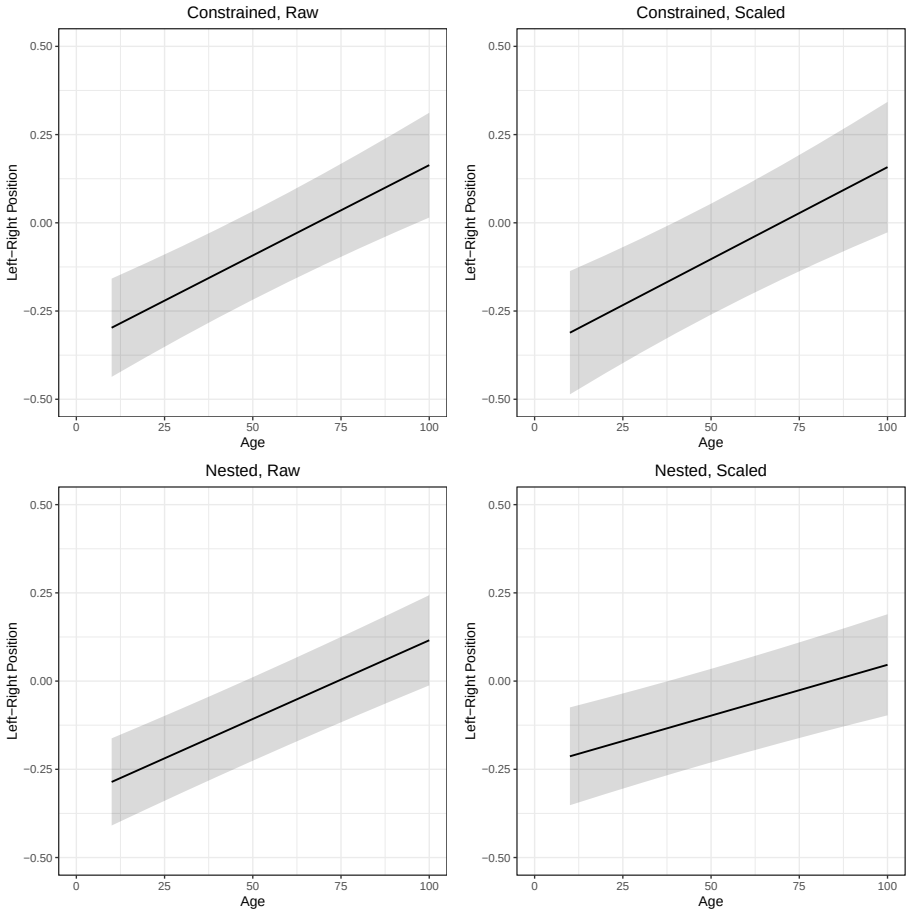

This plot shows the predicted relative left-right position an indivudal will hold as they age. All other variables are set to their means in the sample. The left plot displays this for the constrained model, while the right plot displays this for the nested model.

Once visualised across the span of a human life-time, it becomes clear that there are indeed substantial ageing effects present in terms of relative ideology. Although the year-on-year difference is small, the gradual accumulation over a lifetime results in a shift on average from left to right. Relative to a standardised party distribution, this will be approximate to a notion of moving from center-

left to center-right as the average positions here are within the -1 and 1 standard deviations. Of course, it is possible that what is occurring is that the political contexts individuals find themselves in shift left as they age. It will require the creation of an asbolute measure of ideology to better understand the relationship between age an ideology shown here.

Pointed out but not explicitly analysed thus far is the fact that the results from both the raw data and the scaled data from the Aldrich-McKelvey procedure produce the same inferences. This offers evidence that the raw data can indeed by considered as a form of relative data. This does not mean that the exercise in rescaling the data was pointless: instead, it has demonstrated the contention of this paper that we do not know how APC results would look with absolute measures of the data. Moreover, it is *transparently* relative. Instead of providing such results without discussion, the rescaling ensures that it is clear to the readers how they should interpret the results. Since the results continue to be the same between the raw and rescaled data throughout the rest of the analysis, for the goal of concise presentation plots of the random effects from the raw model are presented in the appendix.

What then about cohort and period effects? I begin by examining cohort effects for the constrained model with the scaled data. The predicted cohort effects are plotted in figure 3:

Figure 3: Predicted Constrained Scaled Cohort Effects

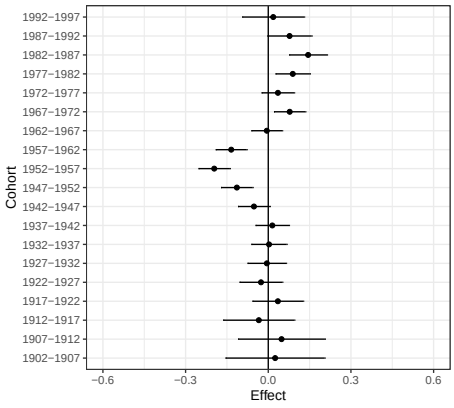

Across the generations, a reasonably clear pattern emerges in figure 3. For the earliest generations, there is a clear null effect. However, for the generations born in the 1940s to the early 1960s, there is a clear left-wing cohort effect. Relative to the political system of their day, this model suggests that individuals belonging to these generations are more supportive of left-wing politics net of other relevant factors. By contrast, the generations born in the late 1960s to the early 1990s are more supportive of right-wing politics net of other relevant factors. This could be interpreted as the respective effects of the post-war period of social democracy, followed by the neoliberal turn from the late 1960s onwards. What is fascinating about these effects is that given the relative interpretation of ideology being used, these effects hold even as the nature of left-right ideology changes around them. The question then emerges: how, if it all, does this model change in the nested model where cohort effects are free to vary country by country? Figure 4 plots the predicted cohort effects from the nested model with the scaled data:

The trends in figure 4 are somewhat mixed. Where individual countries

Figure 4: Predicted Nested Scaled Cohort Effects

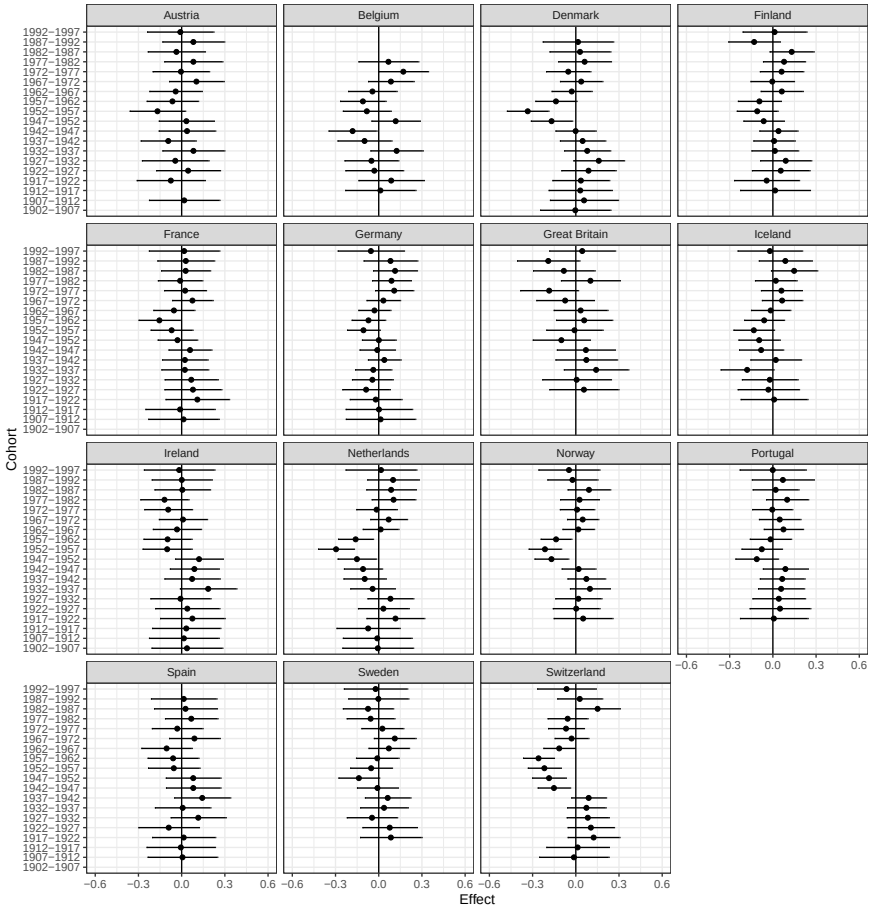

contain statistically significant cohort effects, they follow the same pattern as in figure 3. In other countries the observed cohort effects are not statistically distinguishable from 0. The constraint of similar cohorts across countries is therefore not a particularly restrictive one - at least insofar as relative ideology is concerned.. Although the magnitude of cohort effects differ from country to country, the pattern here is clear enough that it brings into question country-specific theories of cohort effects, such as that in Grasso et al. (2019). The final set of effects left to analyse is that of period effects. Figure 5 plots the period

effects from the constrained model for the scaled data:

Figure 5: Predicted Constrained Scaled Period Effects

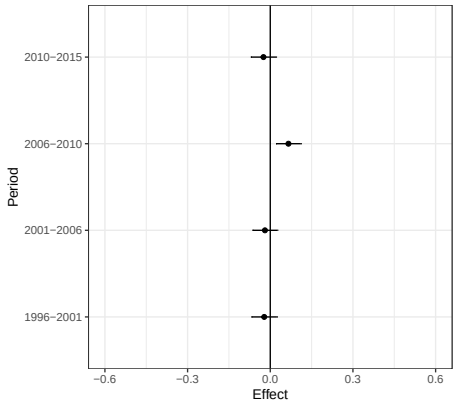

With the exception of the 2006-2010 period, the effects here are all null effects. Similarly, the effect for 2006-2010 is small, though perhaps should not be entirely ignored. Overall however, the result here is a fairly simple one: there is little difference from time period to time period in relative left-right ideology - at least for the period in which these surveys were conducted. The evidence here therefore suggests that there is little effect from 'the mood of the moment' on ideological positions. This may of course be driven by the fact that the survey period relative to the wider range of cohorts in the CSES is fairly short. It may be that period effects do exert an effect, but this would only be observable with a longer survey time period. The obvious question that follows is does this result hold if we nest the period effects in countries? Figure 6 therefore plots the predicted period effects from the nested model:

As compared to figure 5, the results presented in figure 6 offer a more mixed set of results. As in figure 5 Belgium, Finland, Britain, Ireland, Iceland, Portugal, Sweden, and Switzerland show null results and thus no difference from time

Figure 6: Predicted Nested Period Effects

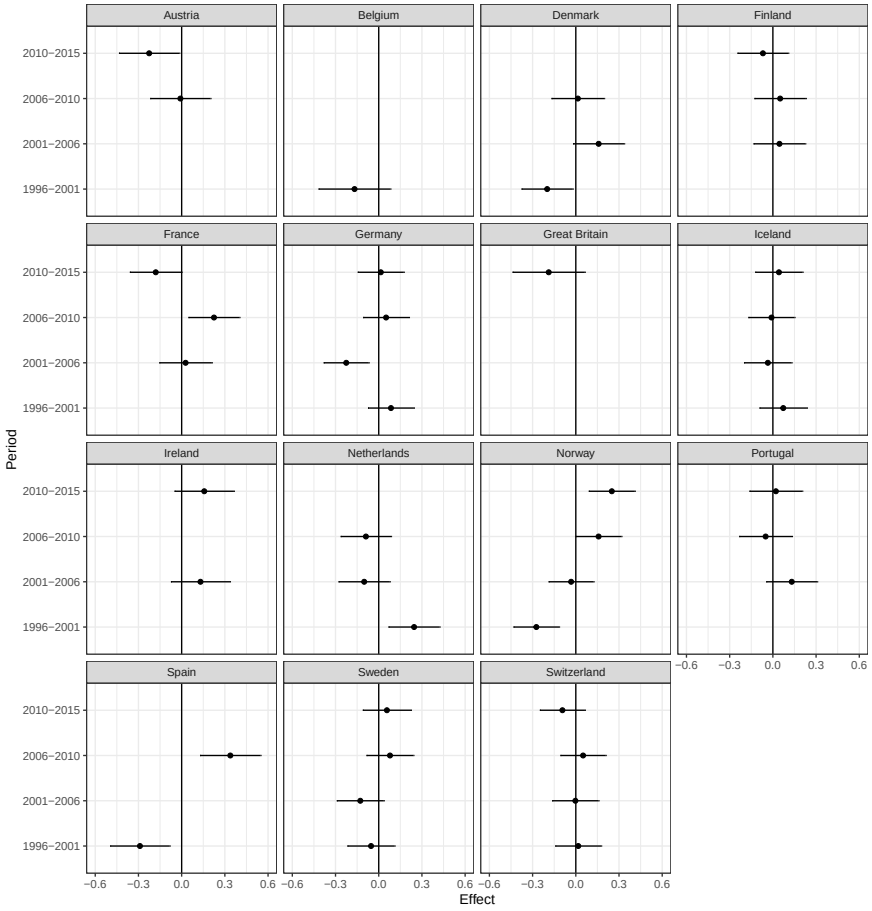

period to time period in the average positioning of the electorate relative to the party system. However, many countries do show interesting trends. Spain and Norway both show a shift from left to right over the years in which they are surveyed. Many countries show their own period trends, such as Germany, France, Austria, and Denmark. These results are therefore in line with a theory of period effects which emphasises the role of political conditions specific to that country. Interpreting an individual set of country results will therefore require country-specific expertise, which is beyond the scope of this paper. Attempting

to constrain period effects to remain the same across countries is therefore a much stronger assumption than it is for cohort effects.

5.1 Robustness

A key argument in my model specification is that the boundaries between cohorts and periods are arbitrary and thus cohort and periods can also be arbitrarily specified into 5-year periods. There is nothing special per se about 5 year time periods: it is simply a common convention in the APC literature. It is perhaps advantageous because it places enough respondents within a group to be useful, but not so many as to begin losing too much detail. This assumption is easily testable however. I therefore re-ran the analysis with 7-year cohorts and time periods. The plots showing the predicted random intercepts for these are presented in the appendix, and they broadly corroborate the patterns found in the analysis here. The results are therefore robust to the way in which cohorts and periods are specified, as we should theoretically expect.

6 Conclusion

In summary, the results presented here show clear evidence in favour of ageing effects, life-cycle effects, and cohort effects in terms of relative ideology. In other words, even as the political system changes around them in terms of the constellation of issue saliences and the 'center' position of the time, people do on average still shift in the direction the political right as they age. They also however retain an initial socialising influence, which is notable given the relative interpretation of the measure here. This would suggest a persistent political influence that remains robust to changes in the precise make-up of political ideology. This in itself is fascinating: it suggests that the socialising effect remains robust to later changes in the nature of political ideology in that

country.

There are few differences between cohort effects from the constrained to the nested models: either they are null results or they show the same pattern. This in itself is an important insight: it implies that past theories conceptualising cohort effects in terms of country-specific political socialisation are flawed. Instead, countries in West Europe appear to follow common political trends in terms of socialising influences and cohort effects - at least insofar as relative ideology is concerned. By contrast, the same dynamic is not observed in terms of period effects: where the constrained results are null, in the nested model the countries with significant effects exhibit fairly different patterns to one another. It is probable that were we to compare more different parts of the world: for example Western Europe to Eastern Europe, we would expect the notion of common cohorts to also break down.

This analysis therefore shows that at least within Western Europe, constraining cohorts effects to be the same across countries is not a strong constraint. This implies that, within Western Europe, the early political influences that remain with a generation are not specific from country to country. This may in part explain why we witness common waves of events such as the summer of 1968, or today's rise of the radical right. Identifying exactly what these common influences are should be a field of future investigation.

By contrast, period effects show a more mixed pattern of results once allowed to be free from country to country. This is in line with a transient interpretation of these effects: they are likely driven by local political events, politicians, and elections, insofar as they are present. The fact that there are reasonably common cohort effects but not period effects is in itself interesting, and worthy of further investigation in future research.

Due to the unavailability of such a measure, it is a shame that the results for

relative ideology here could not be compared against results using a measure of absolute ideology. Indeed, insofar relative ideology captures how people change or stay in their ideological position as the nature of political ideology changes over time, it would be interesting to compare this against results showing how they change or stay given a single, constant, context-independent measure. Indeed, this would inform us to the extent that the results observed here in terms of ageing effects are driven by changes in the political system rather than actual movement in a right-wing direction per se. Future APC - and indeed comparative research more broadly - should focus on this as a matter of urgency.

References

Aldrich, John H and Richard D McKelvey. 1977. "A Method of Scaling with Applications to the 1968 and 1972 Presidential Elections." *American Political Science Review* 71(1):111–130.

Bakker, Ryan, Erica Edwards, Seth Jolly, Jonathan Polk, Jan Rovny and Marco Steenbergen. 2014. "Anchoring the experts: Using vignettes to compare party ideology across countries." *Research & Politics* 1(3):2053168014553502.

Bakker, Ryan, Seth Jolly, Jonathan Polk and Keith Poole. 2014. "The European common space: Extending the use of anchoring vignettes." *The Journal of Politics* 76(4):1089–1101.

Bell, Andrew. 2020. "Age period cohort analysis: a review of what we should and shouldn't do." *Annals of Human Biology* 47(2):208–217.

Bell, Andrew and Kelvyn Jones. 2013. "The impossibility of separating age, period and cohort effects." *Social science & medicine* 93:163–165.

Bell, Andrew and Kelvyn Jones. 2014*a*. "Another 'futile quest'? A simulation

study of Yang and Land's Hierarchical Age-Period-Cohort model." *Demographic Research* 30:333–360.

Bell, Andrew and Kelvyn Jones. 2014*b*. "Don't birth cohorts matter? A commentary and simulation exercise on Reither, Hauser, and Yang's (2009) age–period–cohort study of obesity." *Social Science & Medicine* 101:176–180.

Bell, Andrew and Kelvyn Jones. 2015. "Should age-period-cohort analysts accept innovation without scrutiny? A response to Reither, Masters, Yang, Powers, Zheng and Land." *Social Science & Medicine* 128:331–333.

Bell, Andrew and Kelvyn Jones. 2018. "The hierarchical age–period–cohort model: Why does it find the results that it finds?" *Quality & quantity* 52(2):783–799.

Butler, David and Donald Stoke. 1974. *Political change in Britain: The Evolution of Electoral Change*. St. Martin's Press.

Campbell, Angus, Philip E Converse, Warren E Miller and Donald E Stokes. 1960. *The American Voter*. University of Chicago Press.

Converse, Philip E. 1964. "The nature of belief systems in mass publics." *Critical review* 18(1-3):1–74.

Dawson, Richard E and Kenneth Prewitt. 1968. *Political socialization: an analytic study*. Little, Brown.

Down, Ian and Carole J Wilson. 2013. "A rising generation of Europeans? Life-cycle and cohort effects on support for 'Europe'." *European Journal of Political Research* 52(4):431–456.

Down, Ian and Carole J Wilson. 2017. "A rising generation of Europeans? Revisited." *European Journal of Political Research* 56(1):199–214.

Ford, Robert and Will Jennings. 2020. "The changing cleavage politics of Western Europe." *Annual Review of Political Science* 23:295–314.

Fowler, Anthony, Seth J Hill, Jeffrey B Lewis, Chris Tausanovitch, Lynn Vavreck and Christopher Warshaw. 2022. "Moderates." *American Political Science Review* pp. 1–18.

Glenn, Norval D. 1974. "Aging and conservatism." *The ANNALS of the American Academy of Political and Social Science* 415(1):176–186.

Glenn, Norval D. 2005. *Cohort analysis.* Vol. 5 Sage.

Grasso, Maria T. 2014. "Age, period and cohort analysis in a comparative context: Political generations and political participation repertoires in Western Europe." *Electoral Studies* 33:63–76.

Grasso, Maria Teresa, Stephen Farrall, Emily Gray, Colin Hay and Will Jennings. 2019. "Thatcher's children, Blair's babies, political socialization and trickle-down value change: An age, period and cohort analysis." *British Journal of Political Science* 49(1):17–36.

Hanretty, Chris. 2022. "Party system polarization and the effective number of parties." *Electoral Studies* 76:102459.

Hare, Christopher, David A Armstrong, Ryan Bakker, Royce Carroll and Keith T Poole. 2015. "Using Bayesian Aldrich-McKelvey Scaling to Study Citizens' Ideological Preferences and Perceptions." *American Journal of Political Science* 59(3):759–774.

Hinich, Melvin J and Walker Pollard. 1981. "A new approach to the spatial theory of electoral competition." *American Journal of Political Science* pp. 323–341.

Hooghe, Marc. 2004. "Political socialization and the future of politics." *Acta Politica* 39(4):331–341.

Hopkins, Daniel J and Gary King. 2010. "Improving anchoring vignettes: Designing surveys to correct interpersonal incomparability." *Public opinion quarterly* 74(2):201–222.

Jennings, M Kent. 2007. Political socialization. In *The Oxford handbook of political behavior*.

King, Gary, Christopher JL Murray, Joshua A Salomon and Ajay Tandon. 2004. "Enhancing the validity and cross-cultural comparability of measurement in survey research." *American political science review* 98(1):191–207.

King, Gary and Jonathan Wand. 2007. "Comparing incomparable survey responses: Evaluating and selecting anchoring vignettes." *Political Analysis* 15(1):46–66.

Lauderdale, Benjamin E, Chris Hanretty and Nick Vivyan. 2018. "Decomposing public opinion variation into ideology, idiosyncrasy, and instability." *The Journal of Politics* 80(2):707–712.

Mannheim, Karl. 1970. "The problem of generations." *Psychoanalytic review* 57(3):378–404.

Ordeshook, Peter C. 1976. "The spatial theory of elections: A review and a critique." *Party Identification and Beyond," John Wiley & Sons, New York* .

Peffley, Mark A and Jon Hurwitz. 1985. "A hierarchical model of attitude constraint." *American Journal of Political Science* pp. 871–890.

Poole, Keith T. 1998. "Recovering a basic space from a set of issue scales." *American Journal of Political Science* pp. 954–993.

Reither, Eric N, Kenneth C Land, Sun Y Jeon, Daniel A Powers, Ryan K Masters, Hui Zheng, Melissa A Hardy, Katherine M Keyes, Qiang Fu, Heidi A Hanson et al. 2015. "Clarifying hierarchical age–period–cohort models: A rejoinder to Bell and Jones." *Social science & medicine* 145:125–128.

Reither, Eric N, Ryan K Masters, Yang Claire Yang, Daniel A Powers, Hui Zheng and Kenneth C Land. 2015. "Should age-period-cohort studies return to the methodologies of the 1970s?" *Social Science & Medicine* 128:356–365.

Rekker, Roderik. 2016. "The lasting impact of adolescence on left-right identification: Cohort replacement and intracohort change in associations with issue attitudes." *Electoral Studies* 44:120–131.

Shorrocks, Rosalind. 2018. "Cohort change in political gender gaps in Europe and Canada: The role of modernization." *Politics & Society* 46(2):135–175.

Spitzer, Alan B. 1973. "The historical problem of generations." *The American Historical Review* 78(5):1353–1385.

Stubager, Rune. 2008. "Education effects on authoritarian–libertarian values: a question of socialization1." *The British journal of sociology* 59(2):327–350.

Surridge, Paula. 2016. "Education and liberalism: pursuing the link." *Oxford Review of Education* 42(2):146–164.

Tilley, James and Geoffrey Evans. 2014. "Ageing and generational effects on vote choice: Combining cross-sectional and panel data to estimate APC effects." *Electoral studies* 33:19–27.

Tilley, James R. 2005. "Research note: Libertarian-authoritarian value change in Britain, 1974–2001." *Political Studies* 53(2):442–453.

Wand, Jonathan, Gary King and Olivia Lau. 2011. "Anchors: Software for anchoring vignette data." *Journal of Statistical Software* 42:1–25.

Yang, Yang and Kenneth C Land. 2006. "A mixed models approach to the age-period-cohort analysis of repeated cross-section surveys, with an application to data on trends in verbal test scores." *Sociological methodology* 36(1):75–97.

Yang, Yang and Kenneth C Land. 2008. "Age–period–cohort analysis of repeated cross-section surveys: fixed or random effects?" *Sociological methods & research* 36(3):297–326.

Yang, Yang and Kenneth C Land. 2013. *Age-period-cohort analysis: New models, methods, and empirical applications.* CRC press.