User Guide for our Estimates of Policy Ideology in Europe

March 15, 2019

This guide helps users to access and interpret our measures. We also offer guidance on how to modify or extend the measures, including estimating ideology for different numbers (or types) of dimensions and potential extensions to other demographic subgroups, countries and future surveys.

1 Accessing our Measures of Ideology by Country and Period

The scales by country and two-year period are available for download from `https://tomogradypolitics.wordpress.com/data-on-european-ideology/`.

2 Interpretation of the Scales

Here we offer a few notes on how to interpret our scales for use in substantive applications

2.1 The role of the original survey data

It is worth emphasizing that, as with any scaling method, our eventual measures are mainly determined by the underlying data. Figure 1 in the main paper shows that the raw survey data exhibit the same temporal and cross-country patterns as our scales. For instance, our finding that Northern European countries – including Scandinavian countries – are more economically

conservative than Southern European countries is a clear feature of the original survey data and is not 'imposed by the model.' Of course, our choice of grouping the data into four scales affected the results that we found. But once the questions were grouped, our scales accurately reflected the data that comprises each one.

Hence our findings, or those of anyone choosing to modify or extend our analysis, can only ever be as 'good' or as 'accurate' as the survey data itself. There are a number of potential limitations that could affect the cross-country or over-time comparability of all cross-national survey data or measures derived from it, including our scales and previous measures of European ideology such as those derived from the Eurobarometer's question asking respondents to place themselves on a left-right scale. These limitations include: differences in sampling procedures or survey response patterns that lead to measured cross-country or over-time differences in opinion in the absence of genuine differences; differential item functioning, such that different people in different countries or periods interpret the same questions differently;[1] or differential influence from the 'status quo' across countries (see below).

2.2 Interpreting Economic Mood vs. Conservatism

As discussed in the main paper, it is important to be cautious when interpreting cross-country differences in our economic mood measure. Whilst our economic conservatism measure should be comparable in an absolute sense, the mood measure is only interpretable relative to the status quo within a country. Because relative preferences depend on the policy status quo, two individuals from countries with different policies may well differ in their relative preferences even if they share the same absolute preferences. This means that, for instance, the fact that Swedes are amongst the most conservative in Europe in terms of mood does not necessarily mean that the average Swede would opt for less redistribution or government spending in an absolute sense than citizens of other countries. It merely means that the average Swede currently wants cuts to Swedish government spending more than citizens of other countries want cuts in their own governments' spending. This is very likely, of course, to reflect the fact that Swedish government spending is relatively

[1] Although, as explained in the paper, we think that the questions making up our surveys are likely to be much less subject to differential item functioning than questions asking for self-described left-right positions

very high. For substantive applications, therefore, it would make sense to focus on within-country variation in mood, perhaps using two-way fixed effects as we do in Figure 10 of the paper.

3 Modifying or Extending our Measures: Considerations

Because our estimates were produced with publicly-available software, it is logistically quite easy for future users to modify or extend our measures for their own applications. However, the availability of data – both survey data and demographic data for post-stratification – does impose certain constraints. Our replication archive contains the code needed to replicate our scales. Here, we offer some more general guidance on modifications or extensions of our work, particularly the role played by data availability.

3.1 Producing estimates for different ideological domains, or modifying our existing domains

As we explain in the paper, our choice of a three-dimensional structure (with economics further divided into 'absolute' and 'relative') is based on an ex-ante classification that is consistent with the literature on European politics, rather than data-driven criteria for selecting the 'correct' number of latent dimensions. While our analysis of inter-item correlations in the Supplementary Information provides an additional statistical justification for our choices, we certainly do not view our paper as the final word on the dimensionality of European politics. Other analysts may wish to make different choices, and here we explain the issues that arise in doing so.

Logistically, it is very easy to add or delete variables from our existing scales. Our code takes all of the individual cross-national surveys and amalgamates them into a single large dataset. From there, this single dataset is broken down into dimension-specific datasets (e.g. economic issues). One need only change 2-3 lines of our code in order to reassign a variable from one dimension to another, or exclude a variable altogether. Adding in a new variable from one of the surveys requires only a little more work: the additional step of extracting it from the original datasets and potentially re-coding it such that higher values indicate more conservative opinions.

In principle one could easily estimate ideology across fewer dimensions than we do. The single large dataset that we create could be used to produce a uni-dimensional measure of ideology, although we think that our results show that such an enterprise would be inappropriate given the very different cross-national and over-time patterns across dimensions. A more reasonable enterprise might be to combine the immigration and social domains into a single second dimension, given that the cross-country patterns are similar across both domains (even though over-time patterns are very different). Again, this can be achieved with a single line of our code.

It is more difficult, but by no means impossible, to produce estimates for a wider set of dimensions. The constraint that arises is the potential sparse availability of data. In our scales, survey data are available for virtually every two-year period across a reasonably large set of countries. One reason for this is that the scales amalgamate a relatively diverse set of questions. For instance, the 'social issues' scale includes questions on gender equality, gay rights, abortion, euthanasia, marijuana legalization, environmental issues and civil liberties. This is in line with standard definitions of the 'socio-cultural axis'. It could be more challenging to estimate a scale for a single issue like abortion or gay rights simply because questions on those issues have been asked less often. We emphasize that such an enterprise is by no means impossible. We merely mean that the results would be based on less data and more imputation than is the case with our scales.

Indeed, one more general issue with an enterprise of this type is that there is a lag between a set of issues becoming politically salient and their inclusion in cross-national opinion surveys over time and countries. The fact that we are only able to estimate our immigration scale from 1989 is telling in its own right. It simply wasn't asked about in cross-national surveys before then. A contemporary example would be authoritarian values. Questions on authoritarianism have not been asked with much frequency in the past, making it impossible to estimate an 'authoritarian' scale with our method at present. Assuming that such questions do begin to be asked more often from now on, it might soon be possible to do so.

3.2 Producing estimates with different demographic subgroups

With our approach, in principle there are two reasons to estimate ideology at the sub-group level. One is that an appropriate choice of subgroups can help to more accurately estimate aggregate opinion. When the groups themselves hold different opinions (e.g., opinions on social issues clearly differ by age group), estimating ideological positions for each subgroup first and then combining them with post-stratification should yield better estimates of national ideology. The second reason could be that, rather than being interested in aggregate ideology, the ideological positions of sub-groups are of interest directly. For instance, one might wish to investigate gender differences in ideology. In such a case, there is no need to carry out the post-stratification step.

The distinction between these two cases matters because the latter case requires less data. Extending our measures to other sub-groups would only require choosing groups that are measured consistently across all of the surveys in our data, or at least across a sufficiently large subset of them. Besides age and gender (which we use) education would be an obvious candidate for further sub-group analysis. Virtually every survey measures whether or not respondents have a degree, or have completed secondary education. Other possibilities include employees of the public and private sectors, trade union and non-union members, or urban and rural dwellers. These demographic variables are recorded in many of the surveys.

When using further demographic subgroups as a building block in estimating overall national ideology, one must also post-stratify the subgroup estimates, which imposes constraints on feasibility. It requires demographic data on the proportion of the population of each country in each cell defined by the subgroups. In our case, we needed data on the population shares of women aged 16-34, men aged 16-34, women aged 35-59, men aged 35-59, women aged 60+ and men aged 60+. This data is readily available from censuses and population surveys and is collated across Europe back to the 1960s by Eurostat. However, consistent data on the population shares of men and women of different ages with certain educational qualifications is *not* readily available, which would make it challenging to use education as a grouping variable. That is why we did not use education in our analysis, because our ultimate interest was in national ideologies rather than those of sub-groups.

A final consideration when it comes to estimation for further subgroups is computing time. In our experience, increasing the number of subgroups being estimated leads to substantial increases in computing time. While we managed to estimate our models with six subgroups within a reasonable time-frame, even adding one more grouping variable such as whether respondents have a degree would double the number of subgroups from six to twelve. This is likely to lead the models to require substantially more computational time.