ONLINE APPENDIX

Appendix 1. Additional Notes on Cumulative Entropy

Appendix 1.1 Intuitive Example

For an intuitive example of cumulative entropy, we can look at extreme data cases with a 7-point ordinal variable ($k = 7$) as a way to understand the properties of E_c. In the most polarized case we have equal proportions at the poles, $[0.5, 0, 0, 0, 0, 0, 0.5]$, so that: $E_c = 1$, and a slightly less polarized case is given by $[0, 0.5, 0, 0, 0, 0.5, 0]$, providing: $E_c = 0.833$. Conversely, if all of the responses are concentrated at *any* single value, such as the second place, then: $E_c = 0$. Importantly, the relative space between modes is critical, rather than absolute positions, meaning that $[0.5, 0, 0, 0.5, 0, 0, 0]$ and $[0, 0, 0, 0.5, 0, 0, 0.5]$ reflect the same level of polarization giving the same value of E_c.

Appendix 1.2 Spacing, Cutpoints, and Latent Scale

The proposed cumulative entropy is invariant to the spacing between these categories or the latent scale underlying the ordinal responses because it does not exist to include, meaning that the same E_c value is returned for any arrangement of the unknown but true underlying spacings between categories, because it is strictly focused on aggregate ordinal counts. This is seen in the two examples in Figure A.1 where both panels show counted survey choices from a five-point ordinal question where the red lines indicate delineations between categories, and the grey bars indicate the identical counts of cases within the categories. Even though there is very different unseen spacing between ordinal categories, the E_c values will be identical: E_c treats the two situations in the Figure A.1 above identically because the only known information from the survey is the height of the grey bars and their ordering. This highlights the difference between the entropy measure and models whose goal is estimating the location of such cut-points in other contexts, such as ordered logit/probit as well as item response specifications plus related methods.

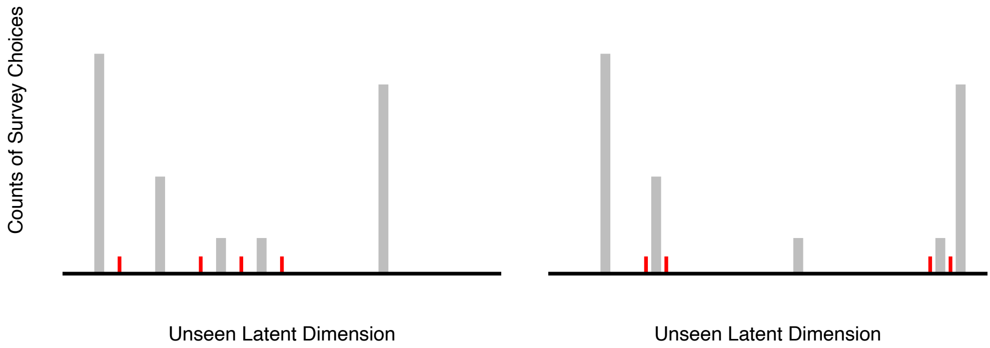

Figure A.1. Spacing Example: Identical E_c Value

Appendix 2. Additional Details about Simulation

Appendix 2.1 Simulation Setup

The simulation setup involves two major steps. We first use Gaussian mixture distributions to simulate continuous data with predefined number of clusters. This way we can define the "truth" of polarization as the distance between the means of two Gaussian distributions. The process can be described as follows:

$$x_i \sim F(x)$$
$$F(x) = \sum_{j=1}^{K} \lambda_j \mathcal{N}(x|\mu_j, \sigma_j) \tag{1}$$

where $x_1, x_2, ..., x_n$ is an array of n values drawing from a Gaussian mixture distribution. The Gaussian mixture has a total cluster of K (K = 2 for most of the time except for the trimodal distribution) and different means and variances for subcomponents. The polarization is thus computed as the difference between the means of the two Gaussians. We start with a basic configuration with two clusters with equal variance followed by other four different settings (see Figure A.2) to investigate how different shapes of distributions, such as trimodal and unequal standard deviation and size, affect the performance of the measures.

The second step is to cluster the continuous data into multiple categories. It is more challenging to cluster one-dimensional data as it generally conveys less information. This is also well-known that this leads to an *NP-hard* problem in a Euclidean space, even when the number of clusters k is 2 (Aloise et al. 2009). Therefore, the conventional heuristic k-means clustering tends to heavily rely on the initial cluster centers, producing neither optimal nor repeatable results. In this case, we use a modified optimal k-means algorithm with dynamic programming to classifying the continuous data into ordinal data (Wang and Song 2011). In other words, the goal here is to assign elements of $x_1, x_2, ..., x_n$ into k clusters (k here corresponds to the eventual clusters of ordinal responses, and thus, $k \neq K$). This is done by first sorting the data vector, $x_1, x_2, ..., x_n$, by non-descending order. Then, the optimized algorithm will calculate the minimum within-cluster sum of square iteratively for the recurrent substructure $\mathbf{D}[i, m]$ from the original problem $\mathbf{D}[n, k]$:

$$\mathbf{D}[i, m] = \min_{m \leq c \leq i} \{\mathbf{D}[c-1, m-1] + d(x_c, ..., x_i)\} \quad (1 \leq i \leq n, 1 \leq m \leq k) \tag{2}$$

where c is the index of the smallest value in cluster m that guarantees the optimal within-cluster sum of square of $\mathbf{D}[i, m]$ and $d(x_c, ..., x_i)$ is the sum of squared distances from $x_c, ..., x_i$ to their mean. By this recurrence and the final minimum within-cluster sum of square $\mathbf{D}[n, k]$, we can obtain the new clustered data, $x'_1, x'_2, ..., x'_n$ ($x' \in 1, 2, ..., k$). Then, we can compute the proportions of each category and the eventual polarization metrics. It is important to note that k-means cluster estimation comes with strong assumptions that are often ignored but fit here because we control the structure of the simulation (the number of clusters is known in advance, the clusters are spatially grouped, the clusters are of similar size, and the within-cluster variance is the same). Figure A.2 visualizes two-step procedures and five configurations. Each configuration is repeated 5000 times with a sample of 2000 data points each time and ranked by both the true and calculated polarization, with higher ranks indicating higher level of polarization.

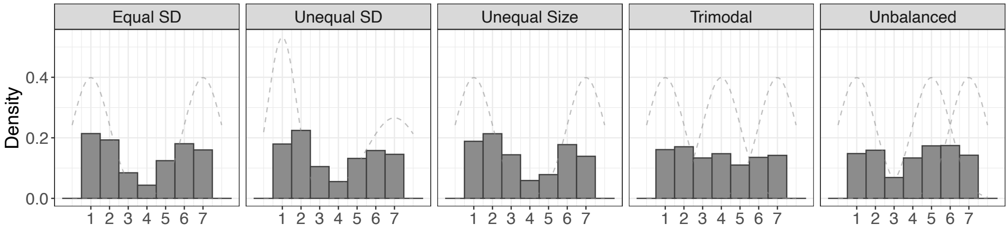

Figure A.2. Simulation Setups. The dashed curves represent the Gaussian mixture distributions. The bar plots are the ordinal data after clustering.

Appendix 2.2 Additional Simulation Results with Different Numbers of Categories

Additionally, we also test whether different numbers of categories would affect the results. Figure A.3 compares 4, 5, 6, 7, and 10 categories, which are all very common in surveys. As we can see, the entropy-based measure of polarization clearly outperforms both of the other measures in all five settings. Generally speaking, we would suspect that, as the number of categories increases, the distribution of the ordinal data should be closer to the pattern of "continuous" data (i.e. more equally spaced, having a meaningful average point, etc.) However, even with 10 categories, we still see that the entropy-based measure performs better than the other two statistics that are based on continuous data[1].

Appendix 3. Crowd-Sourcing Validation

We use a "wisdom-of-the-crowd" validation approach to further compare the metrics of polarization with more intuitive human judgments. The goal of this crowd-sourcing approach is to solicit human judgments of different scenarios of polarization, aggregate this information based on collective opinions from multiple individuals and multiple rounds of evaluations, and finally, benchmark different metrics of polarization against the crowd assessment. Whenever MTurk samples are used

1. We can see variance performs better when number of category equals four. This is largely due to the setup of simulations and the classification algorithm we used: it tends to generate more evenly distributed four categories, and thus, favors variance in this case.

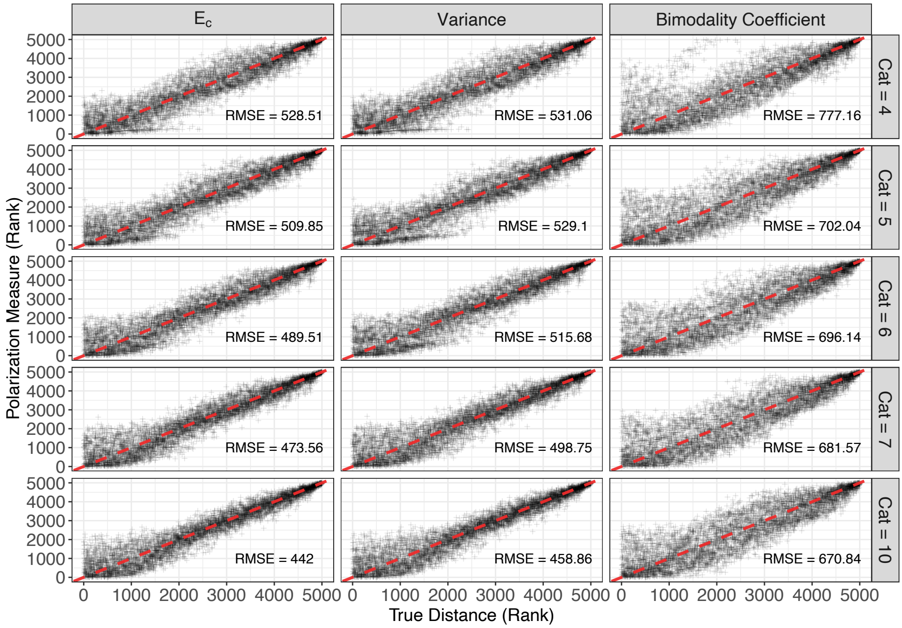

Figure A.3. Simulations of Ordinal Data and Polarization Measures.
Note: Each row of plots represent one configuration. Each configuration is repeated 5000 times with a sample of 2000 data points each time and ranked by both the true and calculated polarization, with higher ranks indicating higher level of polarization. The blue dashed lines represent when measured and defined polarization is perfectly aligned ($y = x$). RMSEs are calculated based on errors of the rank, which scale both the polarization measures and true distances to the same unit.

for analysis caveats are required, even in our case which is solely for a validation check. In addition to paying for the "high quality" respondents offered by MTurk and preparing a detailed informational statement for them, we carefully checked their responses for inconsistencies and inattention as well as basic understanding of polarized distribution.

Appendix 3.1 Validation Tasks and Online Experiment

We designed an online validation experiment in which respondents compare polarization of different scenarios. The design can be described as follows:

First, we generate ordinal data based on Dirichlet distributions:

$$(p_1, p_2, ...p_k) \sim Dir(\alpha) \; (k = 7, \; \alpha > 0) \tag{3}$$

where, $p_1, p_2, ...p_k$ are proportions for a ordinal variable with seven categories ($k = 7$). Then, using these proportions, we can formulate validation tasks as an evaluation of which scenario is more polarized. Study participants were recruited through Amazon Mechanical Turk ($N = 250$). For each task, each respondent was presented with a pair of barplots as in Figure A.4 with a contrived context of ideology or issue opinion and asked to choose a more polarized scenario according to the barplots of ordinal distribution. Each respondent was asked to evaluate eight tasks in addition to one baseline task (Figure A.5) that simply tests whether the respondents understand the most basic meaning of polarization.

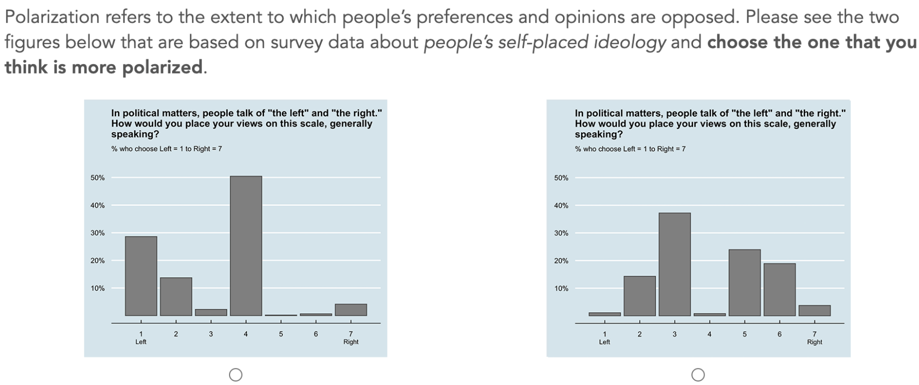

(a) Ideology Context

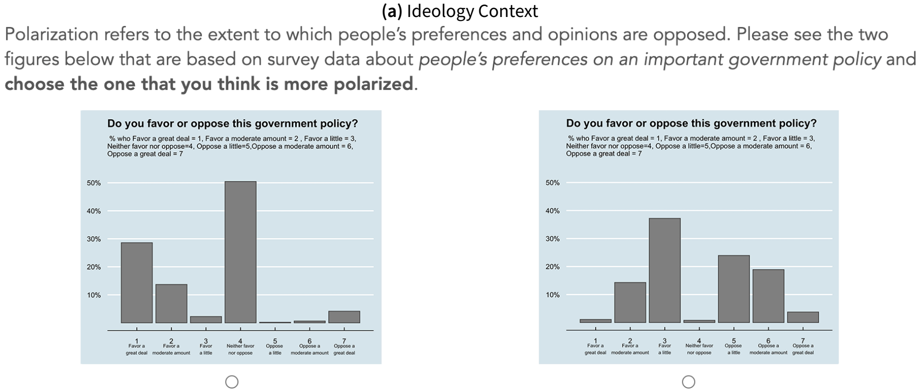

(b) Issue Opinion Context

Figure A.4. Sample Crowd-Sourcing Validation Tasks.
Note: The plots are the screen shots taken from the online experiment. They are intentionally created to look like charts from some popular magazine.

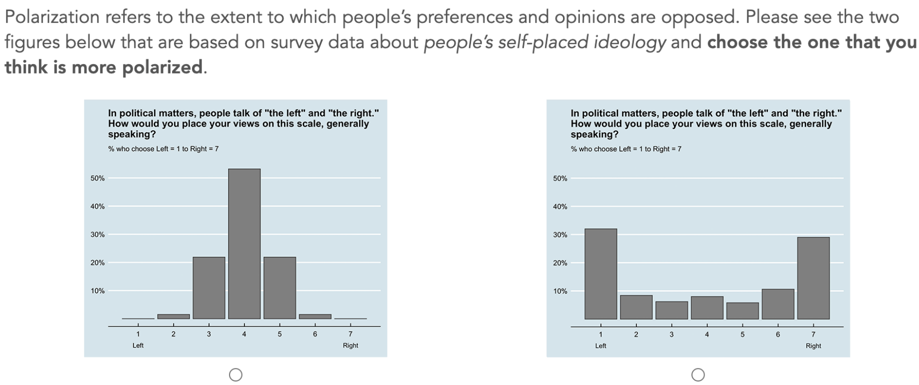

(a) Ideology Context

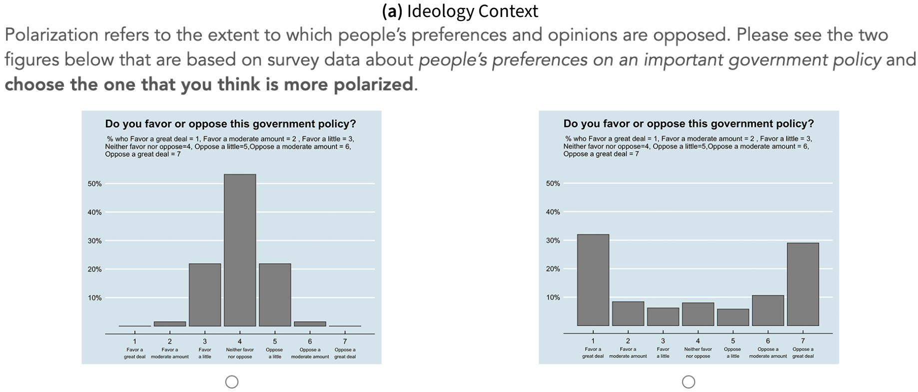

(b) Issue Opinion Context

Figure A.5. Baseline Crowd-Sourcing Tasks

Appendix 3.2 The Calculation of Agreement Rates

To analyze the crowd-sourcing data, we calculate the agreement rates between the metrics of polarization (variance, bimodality coefficient, entropy-based measure) and crowd-sourcing evaluations. Specifically, for respondents $i = 1, \dots, n$, they evaluated $j = 1, \dots, m$ tasks, we can determine the agreement rates as:

$$
\begin{aligned}
X_{ij}^{E_c} &= \mathbb{1}(Y_{ij} = Y_{ij}^{E_c}) \\
X_{ij}^{Var} &= \mathbb{1}(Y_{ij} = Y_{ij}^{Var}) \\
X_{ij}^{Bimod} &= \mathbb{1}(Y_{ij} = Y_{ij}^{Bimod}) \qquad (4)
\end{aligned}
$$

where X is an indicator of agreement for respondent i and task j. When the metric and the response agrees on which one is more polarized, $X_{ij} = 1$, otherwise, 0. Then, we compute and compare the overall agreement rates by the average agreement rate for each metric:

$$\begin{aligned}\widehat{E_c} &= \frac{1}{nm}\sum_{i=1}^{n}\sum_{j=1}^{m} X_{ij}^{E_c} \\ \widehat{\mathrm{Var}} &= \frac{1}{nm}\sum_{i=1}^{n}\sum_{j=1}^{m} X_{ij}^{\mathrm{Var}} \\ \widehat{\mathrm{Bimod}} &= \frac{1}{nm}\sum_{i=1}^{n}\sum_{j=1}^{m} X_{ij}^{\mathrm{Bimod}}\end{aligned} \tag{5}$$

Appendix 4. Mass Polarization in the U.S.

Appendix 4.1 Data Source and Survey Instruments

Data for the example of mass polarization in the U.S. comes from the American National Election Studies (ANES). We combined the 1948-2016 cumulative data with the newly released 2020 Time Series Study. The data is based on work supported by the National Science Foundation under grant numbers SES 1444721, 2014-2017, the University of Michigan, and Stanford University.

We use six survey instruments in this study: liberal-conservative scale (7-point), government services-spending (7-point), opinions on abortion (4-point), aid to blacks (7-point), protect homosexuals against discrimination (4-point), opinions on immigration (6-point).

Appendix 4.2 Detailed Description of Each Item

To show the dynamics of responses behind the polarization measures, in Figure A.6 we show the detailed barplots for all the ANES items that we used in Section 5.1. In a more nuanced way, they show how our entropy-based measure of polarization can pick up the subtle dynamics that are neglected and wrongly captured by other measures.

For the partisanship and ideology items, the barplots of distribution show that the entropy-based measure of polarization is able to capture both the long-term trend and short-term nuances of

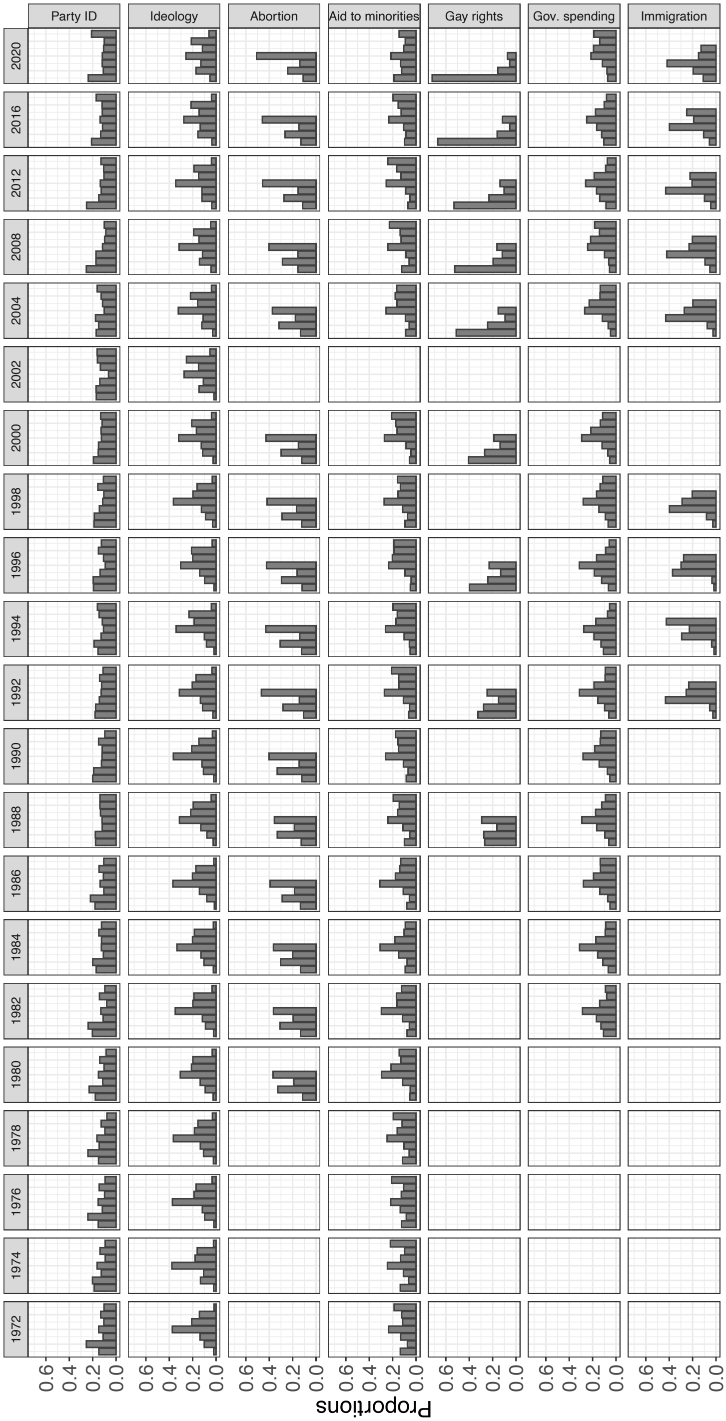

Figure A.6. The barplots of the responses across years (1972-2020). The years in which the questions weren't asked are omitted. The codings all start with one but different questions have different numbers of categories.

polarization. The barplots of abortion opinions also demonstrate that E_c is able to reflect the relatively stable trend of the distributions of responses since 1980s and the recent slight decreasing polarization as more people move to the pro-choice side. For the aid to minority item, the barplots show the decreasing trend of polarization from 1970s to early 2000s, and the increasing polarization in the most recent decade, which are also reflected in the metrics of polarization. The distributions of gay rights in the barplots clearly show the once relatively polarized issue has been forming unprecedented agreement as the majority of the people becomes on the same page of supporting for protecting gay rights. As demonstrated in the barplots, all the metrics of polarization are able to capture the dynamics of government spending. For the immigration issue, the barplots show an increasing trend of polarization from early 2000s to 2016 as more people became neutral or pro-immigration and a decrease in 2020 as it forms a uniform distribution, which is only picked up by the entropy-based measure.

References

American National Election Studies, University of Michigan, & Stanford University. 2021. Time Series Cumulative Data File (1948-2016) [data file and codebook]. https://electionstudies.org/data-center/anes-time-series-cumulative-data-file/

American National Election Studies, University of Michigan, & Stanford University. 2021. 2020 Time Series Study [data file and codebook]. https://electionstudies.org/data-center/2020-time-series-study/

Appendix 5. Radical Parties and Mass Polarization in Europe

Appendix 5.1 Data Source

The data for the analysis of radical parties and mass polarization in Europe comes from the replication materials of the original study by Bischof and Wagner (2019). All variables and coding are the same as the original study except for the additional measure on the ideology item.

Appendix 5.2 Additional Analysis

In Bischof and Wagner (2019), the authors also identified that electoral thresholds may make entrances of new parties more difficult, and thus, give the rise of radical parties additional impact on polarization.

They report the results for both the entire sample and countries with electoral threshold, and identify that the entrance of radial-right parties have even larger effects on mass polarization for countries with threshol. In the main text, due to limited space, we only report the results for the whole sample. Table A.1 shows the comparison of E_c and variance estimates for the country sample with electoral thresholds. In line with the original study, we also find the slightly larger estimates for the models with E_c as the outcome. We again confirm the entropy-based measure results in smaller standard errors, suggesting an improvement in the efficiency of the estimation.

Table A.1. OLS estimates for the effects of radical party entrance on polarization. The table reports the results for only countries with electoral threshold as defined in the original study. Standard errors are clustered by country/election.

		E_c			**Variance**	
	(1)	**(2)**	**(3)**	**(1)**	**(2)**	**(3)**
Radical-right enter	0.036	0.054	0.057	0.126	0.161	0.174
	(0.013)	(0.013)	(0.015)	(0.053)	(0.045)	(0.048)
GDP growth			−0.001			−0.005
			(0.001)			(0.005)
Unemployment (t-1)			−0.001			−0.003
			(0.001)			(0.005)
Party system polarization (t-1)			−0.0002			−0.001
			(0.001)			(0.004)
Party system fragmentationt (t-1)			−0.006			−0.022
			(0.005)			(0.017)
Constant	0.429	0.522	0.552	2.088	2.505	2.610
	(0.009)	(0.028)	(0.038)	(0.036)	(0.095)	(0.139)
N	253	253	243	253	253	243
R-squared	0.078	0.582	0.608	0.061	0.646	0.670

References

Bischof, Daniel, and Markus Wagner. 2019. “Do Voters Polarize When Radical Parties Enter Parliament?” *American Journal of Political Science* 63, no. 4: 888–904.
Bischof, Daniel, and Markus Wagner. 2019. "Replication Data for: Do Voters Polarize When Radical Parties Enter Parliament?". *Harvard Dataverse*, V3. https://doi.org/10.7910/DVN/DZ1NFG.

Appendix 6. Cross-Country Trends in Ideological and Affective Polarization

Appendix 6.1 Data Sources and Survey Instruments

The data for cross-country trends of affective polarization shared by Boxell, Gentzkow and Shapiro (forthcoming)[1]. The data for cross-country trends of ideological polarization is complied by the authors. The question is generally worded as "in politics people sometimes talk of left and right. Where would you place yourself on this scale?" with some small variations between surveys. The data sources and scales is described in Table A.2:

Table A.2. Data Sources for Cross-Country Ideological Polarization

Country	Sources
Australia	CSES (1996, 2004, 2007, 2013); Australian Election Study (1993, 1998, 2001, 2010, 2016, 2019)
Britain	CSES (1997, 2005, 2015); The British Election Study (1992, 2001, 2010, 2017)
Canada	CSES (1997, 2004, 2008, 2011, 2015); Canadian Election Study (1993, 2000, 2019)
Denmark	CSES (1998, 2001, 2007); Denmark Election Project (1994, 2005, 2011, 2015)
France	CSES (2002, 2007, 2012, 2017)
Germany	CSES (1998, 2002, 2005, 2009, 2013, 2017); Polibarometers (1990, 1991, 1992, 1994, 1995, 1996, 1997, 1999, 2000, 2001, 2003, 2004, 2006, 2007, 2008, 2010, 2011, 2012, 2014, 2015, 2018)
Japan	CSES (2007, 2013); World Value Survey (1996 (1995), 2004 (2005));
New Zealand	CSES (1996, 2002, 2008, 2011, 2014, 2017); World Value Survey (1999 (1998), 2005 (2004))
Norway	CSES (1997, 2001, 2005, 2009, 2013, 2017); European Value Survey (1993 (1996))
Sweden	CSES (1998, 2002, 2006); European Value Survey (1991 (1990)), 1994 (1996), 2010 (2010)
Switzerland	CSES (1999, 2003, 2007, 2011); European Value Survey (1995 (1996))
United States	CSES (1996, 2004, 2008, 2012, 2016, 2020); ANES (1990, 1992, 1994, 1998, 2000)

Appendix 6.2 Trend Estimation

We follow a similar strategy as in Boxell, Gentzkow and Shapiro (forthcoming) to estimate the linear trend of both affective and ideological polarization. They fitted a bivariate linaer regression with affective polarization as the outcome variable and year as the explanatory variable. Because the sample size is in fact relatively small (N = 149) here[2], we take a further step and fit a Bayesian hierarchical

1. Available at: https://scholar.harvard.edu/files/shapiro/files/data-for-cross-polar.zip
2. There are 12 countries in total. Germany has the most data available for a total of 27 years from 1990 to 2020. Japan and France both only have four years of data.

partial-pooling model for both affective and ideological polarization, respectively as follows:

$$y_i = a_{j[i]} + \beta_j t_i + \epsilon_i$$

$$\alpha_j \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha)$$

$$\beta_j \sim \mathcal{N}(\mu_\beta, \sigma_\beta)$$

where i refers to the individual country-year case, j refers to the country as grouping unit, and the outcome variables (y_i) are either affective or ideological polarization and predicted by the survey years (t_i). Both intercepts (α_j) and slopes (α_j) are allow to vary in the models, which substantively means countries can have different baselines of polarization as well as different rates of change. Because of the features of the partial-pooling model, this allows us to draw strength from the overall trends for those cases with scarce data—or in other words, for those cases, their estimates of trend will shrink towards the means. Overall, the hierarchical model allows us to draw inferences using both between- and within-country variations while still accounting for country-specific patterns and trends of polarization.

Appendix 6.3 Descriptive Trends of Ideological Polarization

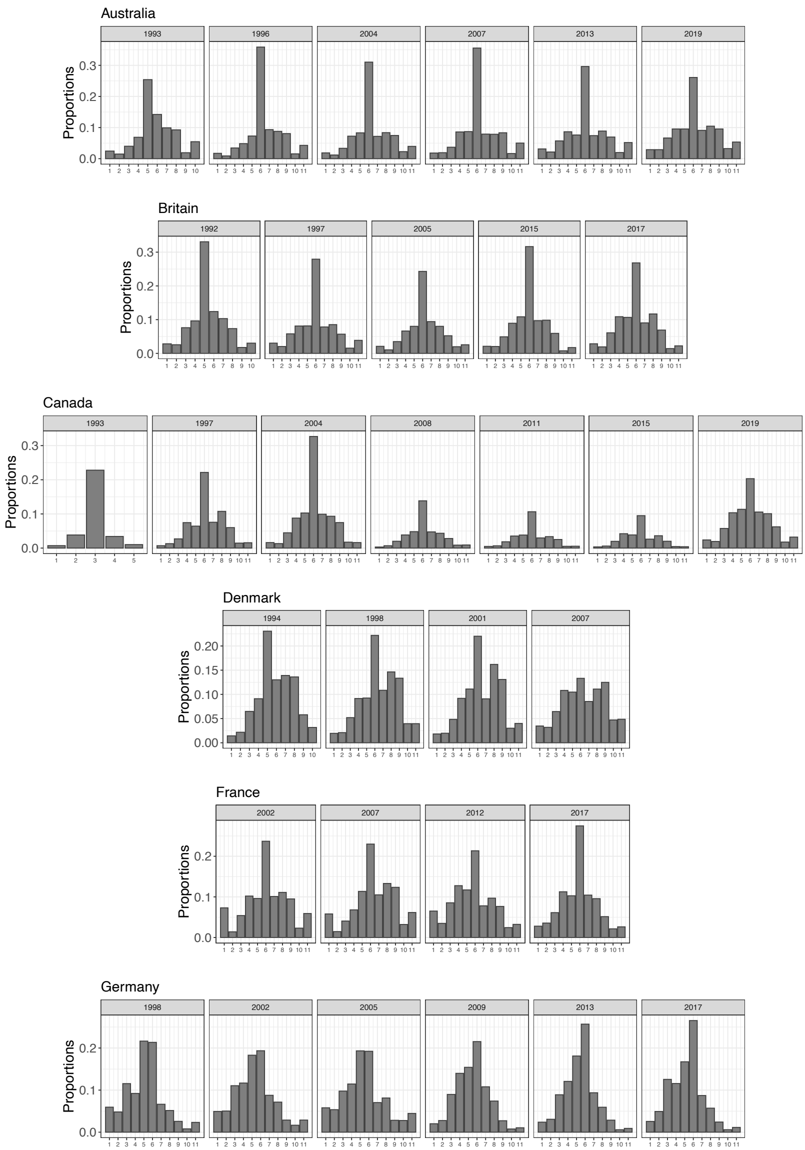
Australia
1993
1996
2004
2007
2013
2019
Proportions
Britain
1992
1997
2005
2015
2017
Proportions
Canada
1993
1997
2004
2008
2011
2015
2019
Proportions
Denmark
1994
1998
2001
2007
Proportions
France
2002
2007
2012
2017
Proportions
Germany
1998
2002
2005
2009
2013
2017
Proportions

Figure A.7

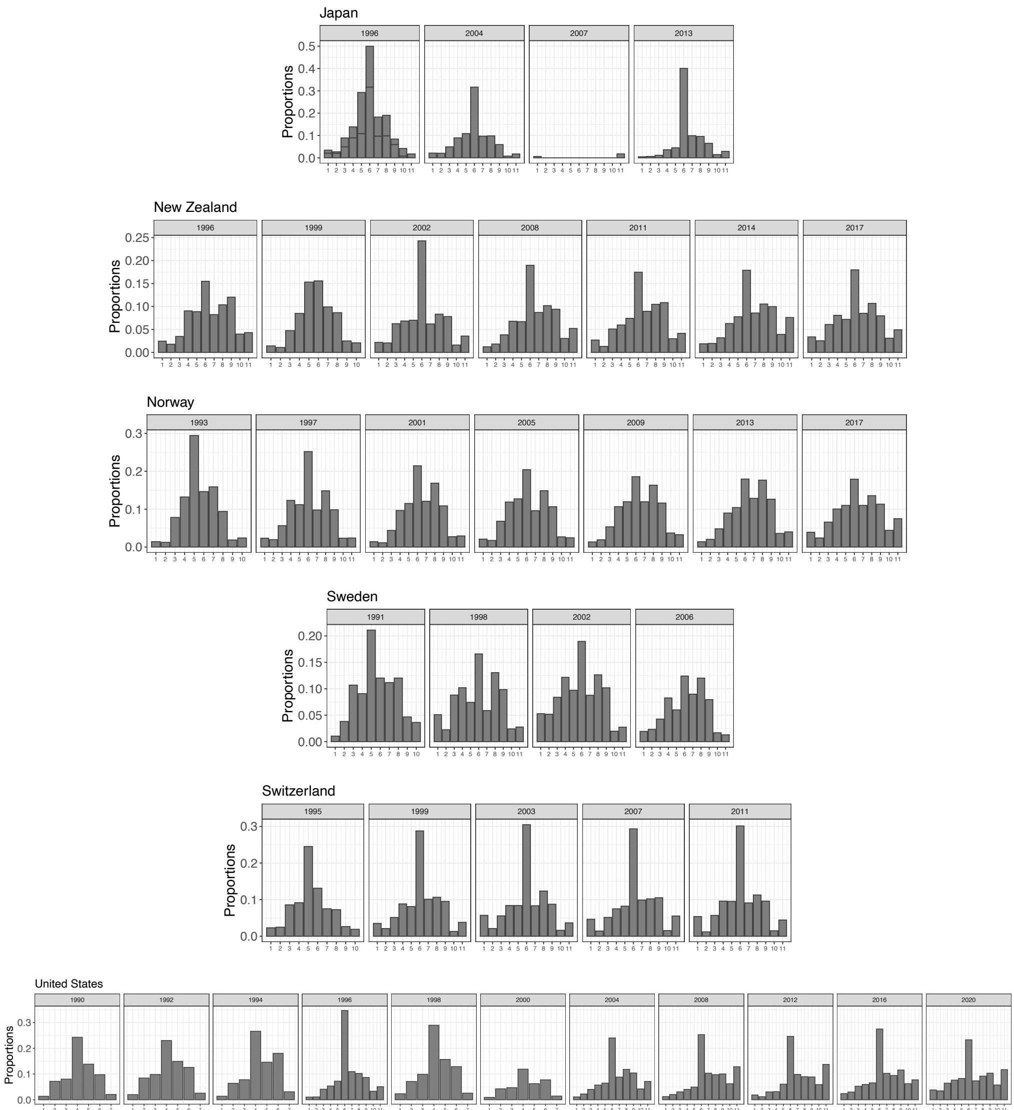
Japan
1996
2004
2007
2013
Proportions
New Zealand
1996
1999
2002
2008
2011
2014
2017
Proportions
Norway
1993
1997
2001
2005
2009
2013
2017
Proportions
Sweden
1991
1998
2002
2006
Proportions
Switzerland
1995
1999
2003
2007
2011
Proportions
United States
1990
1992
1994
1996
1998
2000
2004
2008
2012
2016
2020
Proportions

Figure A.8

References

Bean, C., Gow, D., McAllister, I. 1999. Australian Election Study 1998 [computer file], January 1999.

Bean, C., Gow, D., McAllister, I. 2002. Australian Election Study 2001 [computer file], April 2002.

Blais, A, Brady, H, Gidengil, E, Johnston, R and Nevitte, N. 1994. The 1993 Canadian Election Study [dataset]. Toronto, Ontario, Canada: Institute for Social Research [producer and distributor].

Blais, A, Gidengil, E, Nadeau, R and Nevitte, N. 2001. The 2000 Canadian Election Study [dataset]. Toronto, Ontario, Canada: Institute for Social Research [producer and distributor].

Boxell, Levi, Matthew Gentzkow, and Jesse M. Shapiro. Forthcoming. "Cross-Country Trends in Affective Polarization." *Review of Economics and Statistics.*

Clarke, H. et al. British General Election Study, 2001; Cross-Section Survey [computer file]. Colchester, Essex: UK Data Archive [distributor], March 2003.

Fieldhouse, E., Green, J., Evans, G., Schmitt, H., van der Eijk, C., Mellon, J., Prosser, C. 2018. British Election Study, 2017: Face-to-Face Post-Election Survey [data collection]. Version 1.0.

Forschungsgruppe Wahlen, Mannheim. 2020. Partial Cumulation of Politbarometers 1977-2018. GESIS Data Archive, Cologne. ZA2391 Data file Version 11.0.0.

Heath, A. et al. British General Election Study, 1992; Cross-Section Survey [computer file]. Colchester, Essex: UK Data Archive [distributor], April 1993.

Inglehart, R., Haerpfer, C., Moreno, A., Welzel, C., Kizilova, K., Diez-Medrano J., M. Lagos, P. Norris, E. Ponarin & B. Puranen et al. (eds.). 2020. World Values Survey: All Rounds – Country-Pooled Datafile. Madrid, Spain & Vienna, Austria: JD Systems Institute & WVSA Secretariat [Version: http://www.worldvaluessurvey.org/WVSDocumentationWVL.jsp]

Jones, R., McAllister, I., Denemark, D., Gow, D. 1993. Australian Election Study 1993 [computer file], August 1993.

McAllister, I., Bean, C., Gibson, R., Pietsch, J. 2011. Australian Election Study 2010 [computer file], May 2011.

McAllister, I., Pietsch, J., Bean, C., Gibson, R., Makkai, T. 2017. Australian Election Study 2016 [computer file], February 2017.

McAllister, I., Sheppard, J., Bean, C., Gibson, R., Makkai, T. 2019. Australian Election Study 2019 [computer file], December 2019.

Stephenson, Laura B., Allison Harell, Daniel Rubenson and Peter John Loewen. The 2019 Canadian Election Study – Phone Collection. [dataset]

The Comparative Study of Electoral Systems (www.cses.org). CSES Integrated Module Dataset [dataset and documentation].

The Comparative Study of Electoral Systems (www.cses.org). CSES Module 5 [dataset and documentation].

Whiteley, P.F. and Sanders, D. British Election Study, 2010: Face-to-Face Survey [computer file]. Colchester, Essex: UK Data Archive [distributor], August 2014.