

ORIGINAL ARTICLE

A nonparametric entropy-based measure of mass political polarization

Le Bao¹  and Jeff Gill²

¹Massive Data Institute, McCourt School of Public Policy, Georgetown University, Washington, D.C., USA and ²Department of Government, Department of Mathematics & Statistics, American University, Washington, D.C., USA

Corresponding author: Le Bao; Email: le.bao@georgetown.edu

(Received 21 March 2023; revised 8 December 2023; accepted 5 February 2024)

Abstract

Political polarization has become an increasingly salient issue worldwide, but a systematic examination of the variation and sources of mass polarization across countries is limited by current measurement methods. This work proposes a nonparametric, entropy-based measure of mass political polarization. It exploits the specific structure of ordinal distributions in public opinion data, makes no prior assumptions about the form and spacing of the data, and can still draw reliable measures of issue-based polarization. We demonstrate the theoretical and practical superiority of the measure with analytical comparisons and simulations. We then apply the proposed measure to questions about mass polarization in the USA, the relationship between radical parties and polarization in Europe, and cross-country trends in affective and ideological polarization.

Keywords: comparative polarization; entropy; ideology; issue opinion; mass polarization; measurement; nonparametric

1. Introduction

The past decade witnessed growing scholarly interest in mass polarization and concerns that divided electorates are pulling political systems apart. While early debates were centered more on *whether* electorates have become more polarized (e.g., Abramowitz and Saunders, 2008; Fiorina *et al.*, 2010), other works have shifted to emphasizing the group-based explanations of polarization (e.g., Iyengar *et al.*, 2012; Mason, 2015; Abramowitz and Webster, 2016; Wagner, 2021). Yet the question still remains about whether and to what extent ideology and issue preferences are polarized, especially in the sense of how they connect to the growing identity-based antagonism. This requires researchers to take a zoom-out approach from time to time and compare polarization across countries. However, with a few exceptions (e.g., Pontusson and Rueda, 2008; Lupu, 2015; Bosancianu, 2017; Bischof and Wagner, 2019), research on issue-based polarization tends to focus on single-country contexts (mostly the USA) with predefined context and groups, where issue opinions and group identity are mixed together (Dias and Lelkes, 2022). Limitations in our knowledge of the variation and sources of mass political polarization stem mainly from this lack of appropriate measurement of issue-based polarization across countries.

In this work, we propose a new nonparametric, entropy-based measure of *mass political polarization*. It follows the fundamental conceptualization of polarization in the literature but relaxes the unnecessary assumptions about the specific distributional characteristics. This proposed measure exploits the specific structure of ordinal variables in public opinion survey data, makes no prior assumptions about the distribution and spacing of the data, and is able to capture important features of polarization. We demonstrate that the proposed measure is theoretically and conceptually more relevant to the standard intuition of issue-based polarization, and thus is always able to draw reliable and clear measures of aggregated polarization.

We focus on issue and ideology polarization that is generally identified using ordinal survey items. Many of the existing statistical measures of polarization (variances, kurtosis) are not suitable for ordinal variables common in survey data because they assume continuous (interval) data. In addition, nearly all published metrics of polarization are borrowed from statistics because there is some overlap of the statistical properties and the concept of polarization. These metrics tend to only capture or even partially capture one of the aspects of polarization. Even worse, they often make strong assumptions about the specific forms of the distribution assigned and can be biased for different choices of distributions (Downey and Huffman, 2001). As a result, there is no guarantee that empirically these metrics can capture the complete dynamics of polarization.

The entropy-based measure of polarization developed here provides a more intuitive and direct representation of polarization (rather than relying on artificial distinctions to use existing statistical metrics). It is able to emphasize both the concentration and ordering of ordinal data simultaneously and eventually represent polarization. We utilize multiple approaches including hypothetical distributions, simulated data, and crowd-sourcing validation to demonstrate the properties and advantages of our proposed measure. Then, we apply this entropy-based measure of polarization to questions about mass polarization in the USA, the relationship between radical party and polarization in Europe, and cross-country trends in affective and ideological polarization.

2. Existing measurements of polarization

The basic definition of polarization is not very controversial as it emphasizes simply to what extent preferences or opinions are opposed (DiMaggio *et al.*, 1996; Fiorina and Abrams, 2008). The measurement issue of polarization is also not particularly pronounced when the group context is defined, where polarization can simply be measured by the Euclidean distance between group means (e.g., the partisan difference between Democratic and Republican voters in the USA), while the literature has pointed out that this group mean approach is limited in its ability to reflect the full measure of polarization (Levendusky and Pope, 2011).¹ To measure ideological or issue polarization more broadly without predefined group context, scholars tend to rely on different characterizations or conceptualizations of polarization in order to utilize existing metrics. Scholars often distinguish different principles of polarization such as dispersion, bimodality, divergence, spread, regionalization, fragmentation, distinctness, and so on (Bramson *et al.*, 2016; Lelkes, 2016). While some principles can have particular theoretical significance, in reality, many measurements significantly overlap with each other and yet each single aspect cannot reflect real polarization.

For example, studies commonly use the rubrics of *dispersion* and *bimodality* to represent polarization. The former concept refers to the breadth of preferences: to what extent preferences are diverse and “far apart.” The latter concept captures the fact that, when being polarized, people with different positions cluster into separate camps (DiMaggio *et al.*, 1996). However, either of these aspects in fact only partially reflects polarization and it’s even more difficult to disentangle them empirically. As the distribution becomes increasingly bimodal it will naturally in some way go through the process of dispersion as few values land in between the poles² and existing measures often show confused pictures of polarization motivating some scholars to use multiple limited views in an effort to compensate (DiMaggio *et al.*, 1996; Bramson *et al.*, 2016; Lelkes, 2016). To some extent, these distinctions are a by-product of scholarly attempts to use existing statistical metrics to capture polarization since there is no direct metric that provides a fully coherent

¹The idea of comparing groups also applies to other more principled approaches, such as overlapping coefficient and ideal point estimation.

²Fiorina and Abrams (2008) show a bimodal distribution can be a necessary condition for polarization but not a sufficient one.

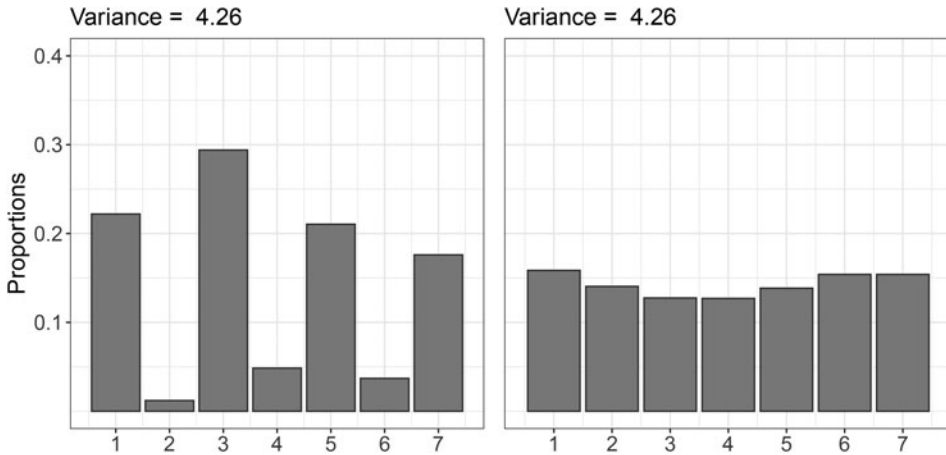


Figure 1. Bimodal distribution versus equal dispersion.

summary measure of polarization. Moreover, these existing metrics also impose additional assumptions about the data and distribution, which is often violated in the data used for studying polarization.

2.1 The limitations of variance and bimodality measures of polarization

Variance is one of the most commonly used measures for issue-based polarization because of its feature of measuring data variability in the common sense. DiMaggio *et al.* (1996) originally used the variance to represent the extent to which respondents are likely to differ in their opinions (e.g., Mouw and Sobel, 2001; Levendusky and Pope, 2010; Hill and Tausanovitch, 2015; Bosancianu, 2017; Bischof and Wagner, 2019). The variance (and of course the mean that its calculation contains) assumes the presence of continuous data or equal spacing of discrete values, and is therefore not mathematically correct as a way to measure opinions and preferences from typical Likert scale questions ubiquitous in survey data measuring ideology or issue opinion (Blair and Lacy, 2000; Downey and Huffman, 2001; Homola *et al.*, 2016).

Perhaps worse, the variance does not directly capture polarization in the way that political scientists want—to distinguish between polarization and voters simply holding diverse views since the latter is what we would expect for citizens in a democratic system. For instance, a uniform distribution of responses, across say five or seven ordered response choices, will lead to a high variance but such a pattern does not comport with the idea of polarization that most scholars of public opinion have. Figure 1 demonstrates that variance cannot in fact distinguish between a polarized bimodal distribution (left) and equal dispersion (right) since they produce the exact same variance value.³

For measuring bimodality, kurtosis and related statistics are choices that measure the “tailedness” (not peakedness) of distribution by comparing the tail density of a given frequency–distribution curve to any normal probability density function based on the scaled fourth moment: $\kappa = \mu_4/\sigma^4$, where μ_4 is the fourth central moment⁴ and σ^2 is the variance of the observed data (Pearson, 1905; Chissom, 1970). Since the kurtosis of any normal distribution is 3, higher values than this indicate heavier tails on either side of the mean and the distribution is then called

³Note that the equal spacing between categories in this and later figures are done merely by convenience and tradition; we are *not* implying or assuming equal spacing on the true latent scale.

⁴The p th central moment of a given distribution is given by $\mu_p = (1/n) \sum_{i=1}^n [(X_i - \bar{X})/\sigma]^p$.

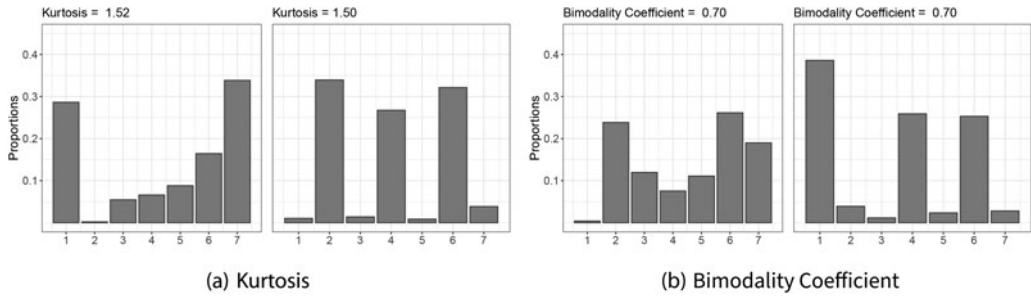


Figure 2. Bimodality versus trimodality: (a) kurtosis and (b) bimodality coefficient.

leptokurtic. Notice that no part of this technical definition addresses actual bimodality. To address this problem scholars as early as 1945 (Mosier *et al.*, 1945) proposed the *bimodality coefficient* to measure bimodality, which combines the kurtosis and skewness of distribution:

$$B_c = \frac{g^2 + 1}{\kappa^* + (3(n-1)^2)/((n-2)(n-3))}, \quad (1)$$

where $g = \mu^3/\sigma^3$ is the skewness (the scaled third moment), κ^* is the excess kurtosis subtracting 3 from κ , and n is the size of the data. A uniform distribution gives a B_c of approximately 0.555, where lower values are evidence of unimodality, and greater values are evidence of bimodality. The statistic was eventually adopted by studies to assess the degree of bimodality for mass political polarization (see Lelkes, 2016).

The logic behind the bimodality coefficient is that a bimodal distribution will have high skewness, low kurtosis, or both. Kurtosis and the bimodality coefficient indeed can capture some characteristics of the distribution regarding polarization, but *they are still based on continuous data measurement and are tied to the normal distribution as a reference point*. Moreover, bimodality measures pose specific assumptions about the distribution of the data, and this distribution is not necessarily bimodal. It can be trimodal or multimodal. Downey and Huffman (2001) note that kurtosis-based measures of polarization provide misleading results in the presence of more than two modes. Figure 2 shows that bimodal and trimodal structures return nearly identical numerical values for both measures.

In sum, since both data variability and bimodality metrics are not originally designed for measuring polarization, in terms of construct validity, they are each only capable of capturing one of the many aspects of polarization. The variance and related measures are designed to describe dispersion in the classic sense regardless of modal distributional features, and the kurtosis-based measures including bimodality coefficient indicate the strengths of modes regardless of dispersion. Consequently, these measures are not sensitive to the distributional differences that are associated with the dynamics of polarization. Figure 3 shows again that very different patterns of distributions can essentially result in the same variance and kurtosis values. Therefore, although conventional measures may occasionally reflect partial trends of polarization, there is no guarantee they can provide a comprehensive and accurate account of polarization. There have also been other efforts to develop measures for ordinal dispersion–concentration similar to the intuition of kurtosis (Leik, 1966; Blair and Lacy, 2000). The key objective here is summarizing the distributional information of ordinal data by its *cumulative relative frequency*. This is a compromise between imposing an assumption about the nature of the continuum underlying the categories and totally neglecting the ordering of the categories, which is the task that we take on here.

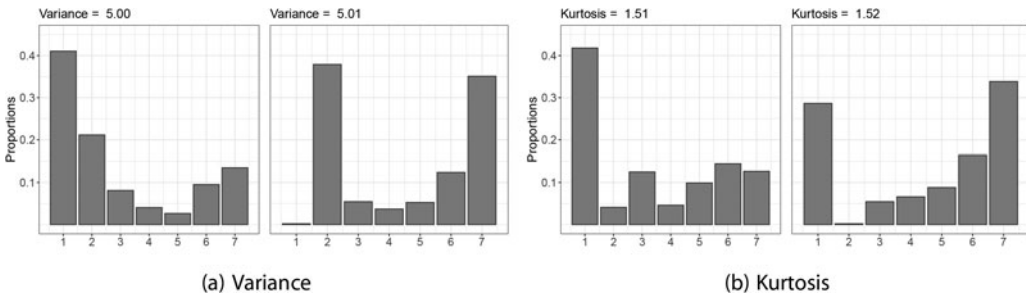


Figure 3. Different patterns of distribution but similar variance and kurtosis: (a) variance and (b) kurtosis.

2.2 What should an effective measure of polarization look like?

To measure polarization in its original sense of identifying “the extent to which preferences or opinions are opposed in relation to some theoretical maximum” (DiMaggio *et al.*, 1996), first, while aspects like dispersion or bimodality can be indicative of polarization, the measurement should primarily focus on the original idea of polarization, rather than being narrowly confined to specific attributes. This is not to understate the significance of specific attributes or subcomponents. Indeed, there are instances where a researcher’s attention might be primarily directed toward transitions from unimodal to bimodal distributions. Nevertheless, for a polarization measure, the primary emphasis should be on capturing the overall pattern and dynamics of polarization. *Second*, an effective measurement of polarization should reflect the inherent nature of ordinal data from standard survey instruments. For a 7-point item, 1 and 7 are obviously more extreme than 2 and 6. Therefore, the measure of polarization should capture the aggregate distribution of ordinal data via such differential spacing and allow for the presence of nonmonotonic patterns as described by Petrocik (1974). *Third*, an ideal measure should describe the distribution of opinions or preferences as a macro-level concept—different from individual-oriented measures of polarization such as ideological constraint/affective polarization (Abramowitz and Saunders, 2008; Baldassarri and Gelman, 2008), or perceptions and attitudes toward out-/in-group member (Iyengar *et al.*, 2012; Druckman and Levendusky, 2019). *Finally*, the measure of polarization also needs to be comparable across time, measurement level, and circumstances.

3. A new entropy-based measure of mass polarization

Here, we propose an entropy-based measure of mass polarization. Contrary to the previous measurement that relies on existing variability- and kurtosis-based metrics, this measure instead focuses on using the features of entropy to describe the aggregate, ordinal distribution *as a whole*. The proposed measure achieves this by emphasizing both the concentration and ordering simultaneously. Note that we employ concentration and ordering not as separate concepts or elements of polarization, but as intrinsic features of the ordinal distribution that researchers can use to identify patterns that might be polarization. In this sense, this is fundamentally different from the previous approach where studies need to conceptualize and measure different subcomponents of polarization separately such as dispersion versus bimodality. Instead, this entropy-based measure provides a more natural and comprehensive since it emphasizes the entirety of the ordinal distribution.

3.1 Entropy background

It is often the case that social and political data are not continuous, especially data generated from survey research. Shannon (general) entropy (Shannon, 1948) is the classic means of describing

information in discrete streams of possible outcomes (x_1, \dots, x_n) occurring with probability $p(x_1), \dots, p(x_n)$, giving $E_S = -\sum_{i=1}^n p(x_i) \log p(x_i)$. This simple formula belies its power to explain natural and human-generated data. General entropy and its modified forms, such as Simpson’s index (Simpson, 1949), have been one of the most popular metrics to describe categorical data structure since it is a direct measure of uncertainty for discrete random variables (Gill, 2005). It increases as every category of the responses becomes more equally likely and decreases as values concentrate in fewer categories (Homola *et al.*, 2016). Entropy also has a naturally intuitive range: the measure is minimized when all values fall into a single category and it is maximized when the values are uniformly distributed across all categories. However, there is one obvious limitation of directly using entropy and related indices for measuring polarization: it cannot detect the *ordering* (or *direction*) of the dispersion hindering its use as a direct measure of polarization. We show here that this is corrected by using an ordinal modification of entropy starting with a cumulative statement.

3.2 Cumulative entropy

Following the literature on ordinal dispersion and concentration (Leik, 1966; Blair and Lacy, 2000), we develop the cumulative entropy to measure polarization. A generalized version is called Tsallis entropy (Tsallis, 2011). Cumulative entropy measures are also used in some natural science fields in this context and appear to have been originally developed in chemistry (e.g., Pace *et al.*, 1955; Wynblatt, 1969). There are three required methodological steps to produce a *cumulative entropy measure* to account for polarization in the way discussed in the last section.

First, define the *binary entropy measure*, which shows to what degree the two categories of a dichotomous outcome are similar in magnitude:

$$H(p, 1 - p) = 2^{-[p \cdot \log_2(p) + (1-p) \cdot \log_2(1-p)]} - 1 \tag{2}$$

where p and $1 - p$ correspond to the proportions of the two categories from some survey or other data source. The exponential component, $-[p \cdot \log_2(p) + (1 - p) \cdot \log_2(1 - p)]$, is a basic binary entropy function, which is a special case of Shannon entropy for a Bernoulli process with p and $1 - p$ as the probabilities of an event landing in either category. Mathematically, $H(p, 1 - p)$ has a maximum value when $p = 1 - p = 0.5$, and a minimum value if either p or $1 - p$ is zero. Here, we require the added assumption that $0 \times \log_2(0) = 0$. This is a common assumption in the literature, and it is not material in real survey data settings since there are no zero response categories except for trivially small number of subjects in a study. Using 2 as a base for the exponent means that the first term of $H(p, 1 - p)$ is scaled between 2^0 and 2^1 (a *typical set* in information theory, see Cover and Thomas, 1991). The form of $H(p, 1 - p)$ makes no parametric assumptions about the underlying scale of uncertainty (Shannon, 1948; Jaynes, 1968, 1982). Also, both base 2 in the logarithms and base 2 in the exponent reflect the diverging notion of pushing mass to the two extremes. The choice of logarithm base here is arbitrary, but using 2 leads to a more intuitive measure for our purposes.

Second, generalize the binary case to an ordinal measured variable with k discrete categories. The corresponding observed proportions for each category are denoted p_1, p_2, \dots, p_k . For the j th category, $j \in \{1, \dots, k\}$, define two complementary cumulative response proportions from an ordinal variable in the data as:

$$S_j = \sum_{i=1}^j p_i \quad S_{\neg j} = \sum_{i=j+1}^k p_i = 1 - S_j, \tag{3}$$

where j and $\neg j$ indicate the two summed lower and upper mass regions ($1 \leq j < k$; $2 \leq \neg j \leq k$) for which we will apply the binary entropy in Equation 2 $k - 1$ times. Note that any selection of $j = 1$,

..., k produces a pair of cumulative values. Also observe that the full set of these S_j, S_{-j} pairs is $k - 1$ in length since S_{-j} does not exist for the k th category, and are analogous to the thresholds on the latent dimension in ordered logit/probit regression models. Note that for any j the sum of these two terms is always equal to 1: $S_{-j} = 1 - S_j$.⁵

Third, using the $k - 1$ pairs defined by Equation 3 the *cumulative entropy* is defined by feeding each of the (S_j, S_{-j}) pairs into Equation 2 and summing the results:

$$E_c = \frac{\sum_{j=1}^{k-1} H(S_j, S_{-j})}{k - 1} = \frac{H(S_1, S_{-1}) + H(S_2, S_{-2}) + \dots + H(S_{k-1}, S_{-k})}{k - 1}, \tag{4}$$

where dividing by $k - 1$ is just a scaling factor. Hence the (S_j, S_{-j}) contrast is compared for each of the $k - 1$ pairs providing a nonparametric ordinal description of the distribution moving from left to right in the sum. So E_c is a summary measure of both concentration and ordering (and in some sense both modal features and dispersion) across the range of the item.

Mathematically, at each of the three steps the calculated values have convenient limits by design:

$$H(p, 1 - p) \in [0 : 1] \quad \text{by log rules and subtracting 1} \tag{5}$$

$$S_j, S_{-j} \in [0 : 1] \quad \text{since } S_j + S_{-j} = 1 \quad \forall (k - 1) \text{ pairs} \tag{6}$$

$$E_c \in (0 : 1] \quad \text{from standardization of the sum by } k - 1. \tag{7}$$

In addition, the magnitude of j does not alter these properties. Thus, E_c is easily interpretable for any value of j , and can be used for cases from low-category Likert scales ($k = 3, 5, 7, \dots$) to higher numbers of categories (although the differing index on the sums means that comparisons between variables/cases should be confined to the same j).⁶ For questions with relatively few categories, the measure can still yield insightful results about polarization, though the granularity might differ from items with more categories. In the latter type of instruments it is likely that the sum in E_c will have a lengthy sum of small probabilities. Online Appendix 1.1 provides an intuitive example, which also shows that the absolute positions of modes do not affect the eventual outcome. This feature is important for comparing polarization across different issues and contexts.

Substantively, the comparison of cumulative proportions describes a distributional difference for a given category along the ordinal scale. Then the cumulative entropy measures the concentration of this distribution: if each side has half of the distribution, then it will result in the maximum of cumulative entropy. Also, the sum of entropy for cumulative proportions will reflect the ordering of the categories. When the distribution is highly dispersed and concentrated on two poles, it will start with a large value in the cumulated process and carry the large value until the other end. When it is concentrated in one category, the entropies for all the cumulative proportions will be very small and it will result in a small value of the measure. As a result, the cumulative entropy measure can capture the complete dynamics of polarization considering both the concentration of particular categories (modes) and to what degree the modes stick together or fall

⁵As a simple illustration, consider category 3 from a total of 7 categories ($j = 3, k = 7$), so that $S_3 = p_1 + p_2 + p_3$ and $S_{-3} = p_4 + p_5 + p_6 + p_7$.

⁶ E_c can theoretically handle survey items with two or more categories. For $k = 2$, it is reduced to the binary entropy, which can capture polarization (or uncertainty) in a manner consistent with the cumulative entropy approach, albeit without the cumulation process.

apart. At the same time, the calculations are based only on the proportion of cases in each category of an ordinal measure, which by definition does not include spacing between these categories. In this sense, it is not based on the unknown cut-points between categories or more generally respondents' personal assessments of the distances between categories (individually or on average) on the latent scale underlying the ordinal responses. In fact, the measure proposed is invariant to such information because it does not exist to include (see online Appendix 1.2). Therefore, this entropy-based measure of polarization imposes no assumption about the central tendency, spacing between categories, and modality of the distribution and is therefore fully nonparametric.

4. Illustration, simulation, and validation

One challenge in developing such a measure is how to verify its performance since there is no such phenomenon as *absolute* polarization with real data as the objective benchmark. In this section, we use multiple approaches toward benchmarking the competing metrics of polarization including hypothetical distributions, systematic simulations, and additional crowd-sourcing validation to comprehensively demonstrate this entropy-based measure's consistency and validity in measuring polarization from different aspects.

4.1 Hypothetical distributions applying E_c

Here, we set up some hypothetical distributions to illustrate the properties of the cumulative entropy and compare it with conventional metrics that were used to measure polarization including the variance and the bimodality coefficient. Figure 4 shows eight sets of contrived barplots with 7-choice ordinal data. Figure 4(a) illustrates the definition of polarization—movement toward the poles of distribution. Figure 4(b) demonstrates the (nearly) maximum value, (nearly) minimum value, equally spread distribution, and trimodal distribution, respectively. Notice here that both the variance and the bimodality coefficient do a poor job relative to the E_c cumulative

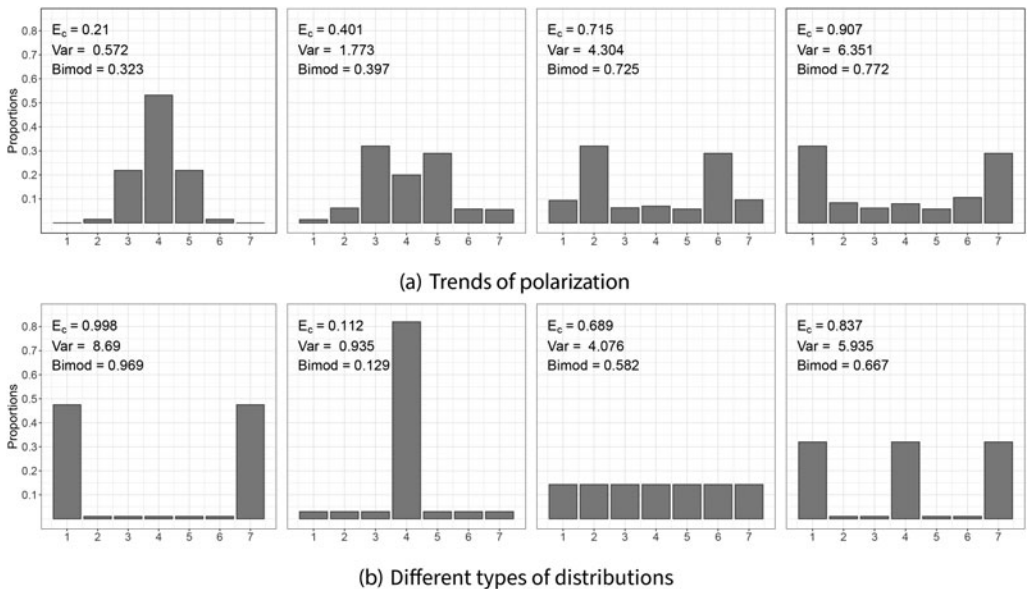


Figure 4. Illustration of different types of distributions: (a) trends of polarization and (b) different types of distributions.

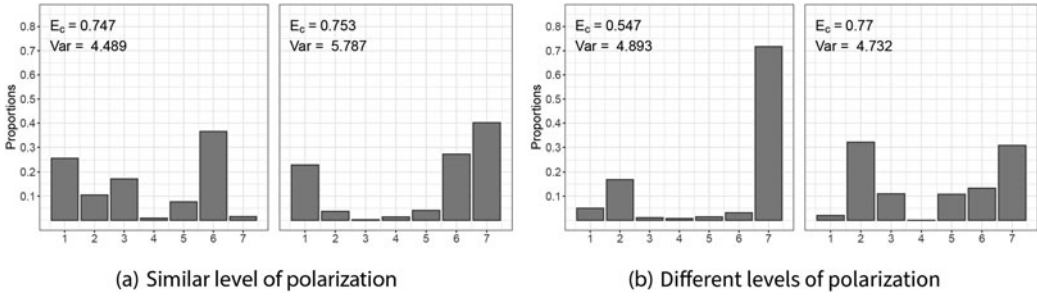


Figure 5. Comparing E_c and variance: (a) similar level of polarization and (b) different levels of polarization.

entropy measure in that their relative magnitudes do not monotonically reflect the intuition of the polarization changes.

Next, in Figures 5 and 6 we show how E_c is able to reveal both more complicated and nuanced dynamics of polarization comparing to traditional variance and the bimodality coefficient. Figure 5 focus on the comparison between entropy and variance. While E_c was able to reflect the similarity in Figure 5(a) and the differences in Figure 5(b), the variance cannot distinguish the dispersion and polarization. For example, the left panel in Figure 5(b) shows the majority of the responses concentrates on one category, which does not indicate high-level polarization. On the contrary, the right panel shows that two nearly equally sized camps are relatively far apart. Similarly, Figure 6 demonstrates that the bimodality coefficient fails to capture the nuances in the changes of distributions when it is not perfectly bimodal. It is important to note that we are making these comparisons for 7-point Likert scales in the graph and in other examples for comparability, and because it is the most common ordinal data type in survey research. Comparisons made with 5-point and 9-point scales reveal the same differences but are scaled differently numerically due to the summation function employed by both measures, and this also means that comparisons between differently sized scales are not appropriate. Variance, skewness, and kurtosis (with only theoretical minimum values) are also not comparable if the original scales are different.

4.2 Simulations of E_c

Since there is no absolute polarization as “ground truth” that can be used as a benchmark from real data, we employ a two-step procedure creating quasi-true values and compare different metrics of polarization. First, we use normal mixture distributions to simulate continuous data with predefined clusters and the “true polarization” is defined as the distance between the

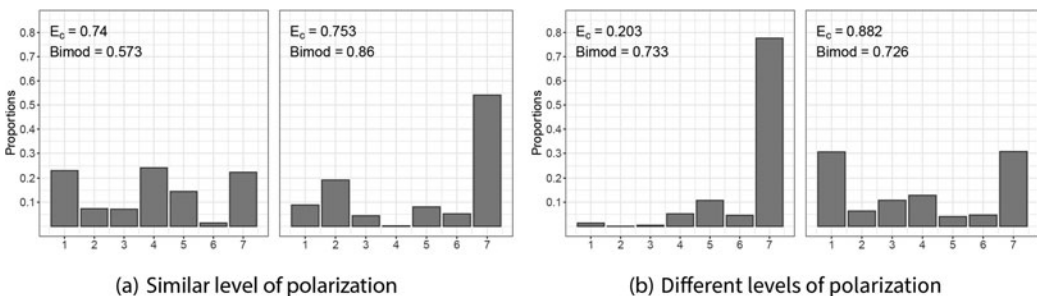


Figure 6. Comparing E_c and bimodality coefficient: (a) similar level of polarization and (b) different levels of polarization.

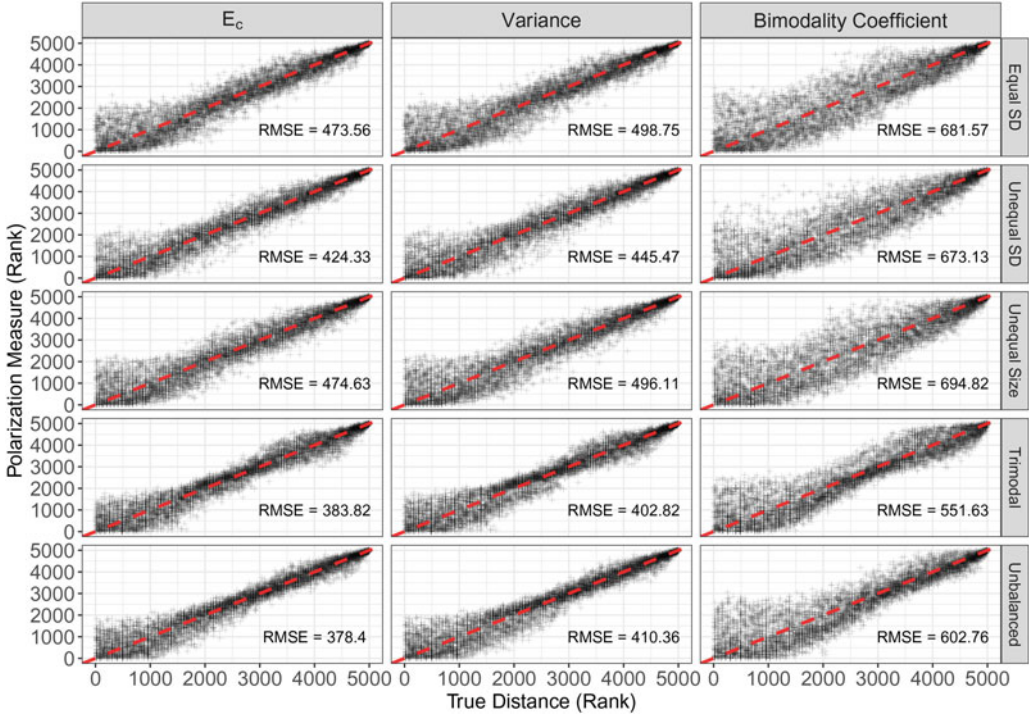


Figure 7. Simulations of ordinal data and polarization measures. The blue dashed lines represent when measured and defined polarization is perfectly aligned ($y=x$). RMSEs are calculated based on errors of the rank, which scale both the polarization measures and true distances to the same unit.

means (modes) of two normal distributions. The next task is to cluster the continuous data into multiple categories, focusing on the 7-category case. We also use a modified optimal k -means algorithm with dynamic programming (Wang and Song, 2011) for the clustering to overcome the challenge of clustering one-dimensional data as it generally conveys less information, and leads to an NP -hard problem in a Euclidean space (Aloise *et al.*, 2009) as discussed further in online Appendix 2.1. It is also important to note that k -means cluster estimation comes with strong assumptions that are often ignored in common practice but fits here because we control the structure of the simulation. We use five basic configurations to simulate normal mixture distributions: (1) equal standard deviation, (2) unequal standard deviation, (3) unequal size, (4) trimodality, and (5) unbalanced (middle point) as illustrated in Figure A.2. In online Appendix 2.2, we also test whether different numbers of categories would affect the results.

In Figure 7, the entropy-based measure of polarization E_c clearly performs the best in all five settings where the E_c estimates are closer to the $y=x$ lines, indicated by the smallest root mean squared errors (RMSEs) of ranks. Variance is second, in quality even though we know it violates the underlying assumption of ordinal data. It performs nearly as well in the equal standard deviation setting but becomes worse when the distributions are less perfect.⁷ The bimodality coefficient performs much worse in all five configurations. Note that the simulation procedure would generally favor the measures of variance and the bimodality coefficient because we generate two spread-out clusters and equally spaced categories in order to use the difference in means as the benchmark. Yet even in this setting, E_c performs consistently better, and we would also expect so

⁷It is expected variance will perform well in some scenarios as it measures the data variability and should be able to pick up some of the dynamics of polarization while still violating the assumption of ordinal data.

Table 1. Agreement rates between measures of polarization and crowd-sourcing evaluations

Metrics	Agreement rates	
	Full sample	Baseline-satisfied
E_c	0.650 (0.0106)	0.674 (0.0110)
Variance	0.421 (0.0110)	0.410 (0.0115)
Bimod. Coef.	0.403 (0.0109)	0.384 (0.0114)

in more complex real-world data with more nuances. It is also worth mentioning that one generally cares more about the right halves of each panel in terms of measuring polarization. The right halves are more about the dynamics of polarization while the left halves contain more uniform distributions. It is generally less informative to compare two near uniform distributions in the context of polarization. In addition, In Figure A.3 with different numbers of categories, the entropy-based measure of polarization still outperforms both of the other measures regardless of the number of categories.

4.3 Crowd-sourcing validation of E_c

As it is impossible to define a perfect “ground truth” of polarization, even in just a simulated setting, we turn to a crowd-sourcing approach to benchmark the metrics of polarization with more intuitive judgments by 250 humans. We designed an online MTurk validation task in which respondents compare graphs reflecting different levels of polarization. The intention here is not about expert knowledge about polarization or political issues in general. Rather, we can seek the “wisdom of the crowd” and ask the respondents to perform a basic cognitive task: to identify which of two graphical distributions appeared more polarized, in line with the idea that most people can distinguish when presented with two contrasting distributions, most should instinctively determine which exhibits greater polarization (Fiorina and Abrams, 2008). For the validation task, each respondent was presented with a pair of barplots given in online Appendix 3.1, Figure A.4 with a contrived context of either ideology or issue opinions and asked to choose a more polarized scenario according to the distributions. We provide the respondents with the most fundamental definition of polarization and also an extremely obvious baseline task as in Figure A.5.

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 using the standard methodology shown in online Appendix 3.2. The results, as reported in Table 1, show that E_c performs significantly better than the other two measures, and in addition it has a smaller standard error than the variance and the bimodality coefficient. So about two-thirds of the time, E_c and the testers agree on which scenario is more polarized, but for the other two measures it is *notably worse than flipping a coin*.

5. Empirical applications

In this section, we apply the proposed measure to three contexts of mass polarization and compare it with the conventional measurements of polarization. Our objective is to show both similarities and differences with the previously used measures with real data as users of our new E_c would encounter.

5.1 Mass polarization in the USA

Mass polarization has been an extremely salient and important topic in American politics. To depict polarization in the different dimensions, we apply both the entropy-based measure (and

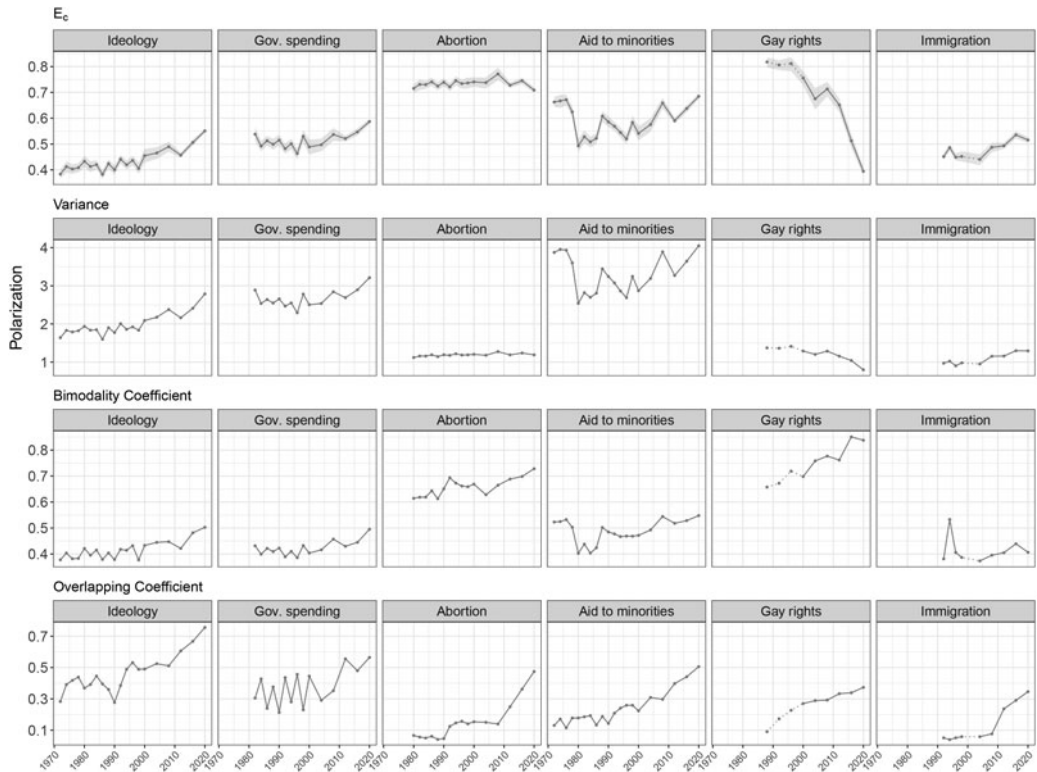


Figure 8. Mass polarization in the American public. Data: ANES (1972–2020). The 95 percent nonparametric bootstrapping confidence intervals are shown for E_c . The dotted lines indicate that there are disconnections between the years that ask the questions and are fitted using linear interpolation. The overlapping coefficient is reversed from the original formula to reflect polarization (nonoverlapping region) for easy comparison.

the conventional metrics of polarization) to the ideology (liberal-conservative) self-placement and issue opinions in longitudinal ANES surveys from 1972 to 2020 in [Figure 8](#). In addition to the aforementioned variance and bimodality coefficient for measuring general polarization, we also include *overlapping coefficient* that is previously used for measuring partisan polarization for comparison (Levendusky and Pope, 2011; Lelkes, 2016). It is important to note that overlapping coefficient requires probability density of continuous data and a predefined two fractions (i.e., two distributions separately from Democrats and Republicans) to calculate the overlapping region, which is mathematically incorrect for ordinal data and theoretically different from what the proposed metric measures.

For ideology, E_c , variance, and bimodality coefficient present an overall similar trend of general ideological polarization where there is moderately increasing trend starting in the late 1990s after remaining at a relatively low level for a long period, while the overlapping coefficient shows a more dramatic increase in the divergence of ideology between Democrats and Republicans. This gap between general and partisan polarization may suggest that, while there is an increasing trend of partisan sorting, the societal polarization remains relatively moderate. It is not surprising that the mathematically incorrect variance and the bimodality coefficient share some features with E_c since they are designed to capture two different effects that are incorporated simultaneously in E_c (dispersion and modality). Also, data in these two dimensions are relatively balanced and has a middle point, which makes it resemble some features of continuous data. This is in comparison to other contexts where the majority of the country can lean left or right and the

middle category is not perfectly aligned with the mean of the distribution due to the categorical nature of the data. Nevertheless, considering the dynamics of other domains and overall scale,⁸ E_c seems to reflect a clearer trend of the recent moderate increase and overall relatively low-level ideological polarization, which is consistent with what other conventional studies suggest (Fiorina and Abrams, 2008; Hill and Tausanovitch, 2015; Lelkes, 2016).

The issue polarization is where metrics start to differ substantively. These salient issues can become rather complex in terms of measuring polarization. They are measured in different scales and some issues do not have a theoretical middle or some do not have a balanced distribution between two sides. The spacing between categories in such settings are also generally more complex than ideological and partisanship items. Comparing the metrics, we see both similarities and differences in terms of both the levels and trends of polarization. However, for some issues, there are both big and small discrepancies. For example, for the abortion issue, E_c and the variance suggest a relatively stable trend since the 1980s, but E_c is able to capture more subtle dynamics for the uneven decline since 2008. While the E_c measure finds a gradual polarization decrease in recent waves of the survey, the bimodality coefficient shows an increasing trend since 1998. The E_c measure is clearly more in line with what the substantive research suggests: the abortion issue, albeit divisive, certainly has not become more polarized (Mouw and Sobel, 2001; Carsey and Layman, 2006; Fiorina *et al.*, 2010), and there is some evidence that the proportion of citizens who favor more abortion rights has increased during this period (as shown in the detailed barplots in Figure A.6), which should result in less polarization. We also know from a vast literature that support for abortion rights is not as stable over time as the variance measure implies here. The overlapping coefficient again depicts a different trend than all other metrics suggesting a tension between partisan and societal polarization on the issue, but it should be noted that overlapping coefficient may exacerbate the polarization as it is *invariant* to the size of the two distributions.

For aid to minorities the E_c measure and the variance are very similar and pick up the same fluctuations, whereas the bimodality coefficient suggests a much more stable picture of polarization as a policy issue, at least since 1992. This implies that the polarization story here is more about dispersion than multimodality. For the issue of gay rights, the four measures show very different patterns. The entropy-based measure depicts a long-term, sharp decrease in the level of polarization, which is consistent with academic (Bishin *et al.*, 2021) and journalistic accounts (shown in Figure A.6 with detailed distributions). The variance shows a slow decline which is not consistent with such accounts. More incorrectly the bimodality coefficient gives an overall increase in polarization around gay rights since 1988 because it is overly sensitive to an immobile but shrinking mode of opposition. Similarly, the overlapping coefficient also presents an increase in partisan polarization of the issue of gay rights as it neglects the changes in the sizes of two parties. Interestingly, the patterns for polarization over this period are essentially the same across time for views on government spending except that overlapping coefficient shows some dramatic fluctuations.

Online Appendix 4 provides barplots to describe the detailed distributions for all the issues as well as ideology, which can further demonstrate that E_c is able to provide more reasonable accounts for the dynamics and overall trends of the ordinal distributions. Finally, it is important to remember that similarity between measures does not imply the same quality of underlying theories and assumptions: the variance and the bimodality coefficient routinely violate features of the data as noted in Section 2; the overlapping coefficient may highlight the partisan differences yet overlook the overall societal trends.⁹

⁸Variance does not have a theoretical maximum and fixed scale so it is also not possible to compare variance across issue domains if the original scales are different.

⁹This gap between partisan and general societal polarization, signified by the differences between overlapping coefficient and E_c , may invite future studies to revisit studies on partisan polarization in the USA. The partisan polarization, measured

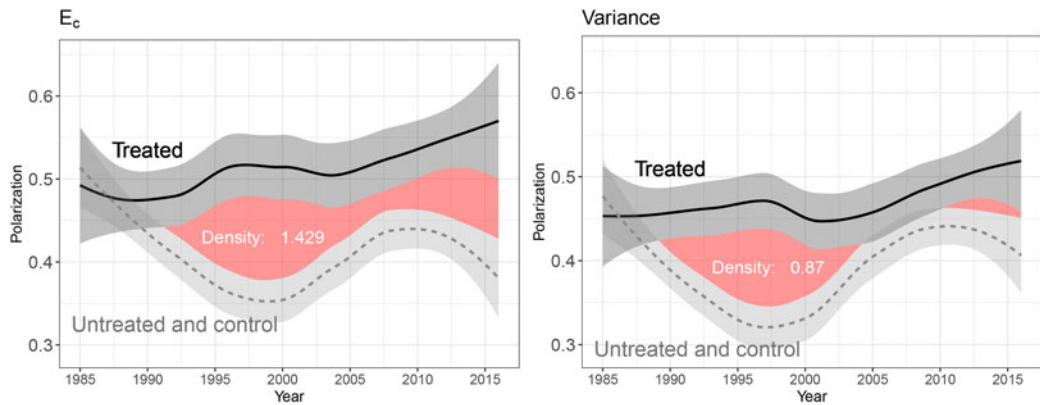


Figure 9. Descriptives of mass polarization in Europe. The figures show lines fitted with local polynomials and corresponding 95 percent confidence intervals. Solid curves represent the treated cases and dashed curves represent the control and untreated cases. Red shades represent the differences between treated and untreated trends after separation. Since variance does not have a bounded scale, it is normalized to 0 and 1 for comparison.

5.2 Radical parties and mass polarization in Europe

In this section, we focus on a more generalized and analytic example that compares the polarization across countries. There has indisputably been rise of extreme parties and radical political elites as well as increasing polarization across continents. Building on the theories that party and elite polarization is conducive to fueling the polarization in the mass public, Bischof and Wagner (2019) employ a time-series cross-sectional analysis on European countries and demonstrate that the mass ideological polarization will increase after a radical-right party gains power in the legislature. For the key outcome, they use the standard deviation of left–right self-placements in each country-year unit to measure public polarization. We replicate their descriptive, inferential, and causal results using the proposed measure and compare them with the original variance measure in their paper.

Figure 9 describes the polarization time trends between treated cases (those encountered entries of radical parties, solid curve) and controls (those never experienced or had not yet experienced, dashed curve). The pattern is similar in general trends between E_c and the original authors' use of the variance, indicating that dispersion not multimodality is dominant in these data. The entropy measure finds a greater divergence between the two groups from 1985 onward (as the density of the differences between two trends indicates), meaning that radical parties have had an even greater impact on polarization than Bischof and Wagner found. They also estimate linear regression models for the relationship between the entrance of radical-right party and mass ideological polarization. Table 2 reproduces the original results (right half) and replicates the models using the entropy-based measure for polarization as the outcome variable (left half).¹⁰ The overall results using E_c are consistent with the original findings, suggesting the entrance of radical-right parties has a reliable effect on the increase of polarization. However, the E_c measure finds noticeably greater separation in later years where we know that the effect is greater. Also, the entropy-based measure results in relatively smaller standard errors (comparing to the scale of coefficient estimates), which can mean an improvement in the efficiency of the estimation. Additionally, the different measurement of the outcome also changes the relationship with other explanatory variables such as *unemployment*.

by overlapping coefficient, is invariant to the dynamics of macropartisanship and thus may overestimate the societal polarization.

¹⁰Online Appendix 5 reproduces the additional results for countries with electoral threshold as in the original study.

Table 2. Ordinary least squares estimates: increasing polarization after entrance of radical right party

	E_c			Variance		
	(1)	(2)	(3)	(1)	(2)	(3)
Radical-right enter	0.031 (0.009)	0.043 (0.009)	0.046 (0.009)	0.090 (0.034)	0.116 (0.032)	0.131 (0.033)
GDP growth			-0.002 (0.001)			-0.006 (0.003)
Unemployment ($t - 1$)			0.0004 (0.001)			0.003 (0.003)
Party system polarization ($t - 1$)			0.0001 (0.001)			-0.001 (0.002)
Party system fragmentation ($t - 1$)			-0.004 (0.003)			-0.018 (0.013)
Constant	0.412 (0.006)	0.423 (0.014)	0.444 (0.029)	2.055 (0.021)	2.103 (0.051)	2.203 (0.111)
R^2	0.056	0.661	0.683	0.035	0.674	0.690
N (elections)	164	164	145	164	164	145
N	534	534	503	534	534	503
Country-fixed effects		✓	✓		✓	✓
Decade-fixed effects		✓	✓		✓	✓

Note: Standard errors are clustered by country/election.

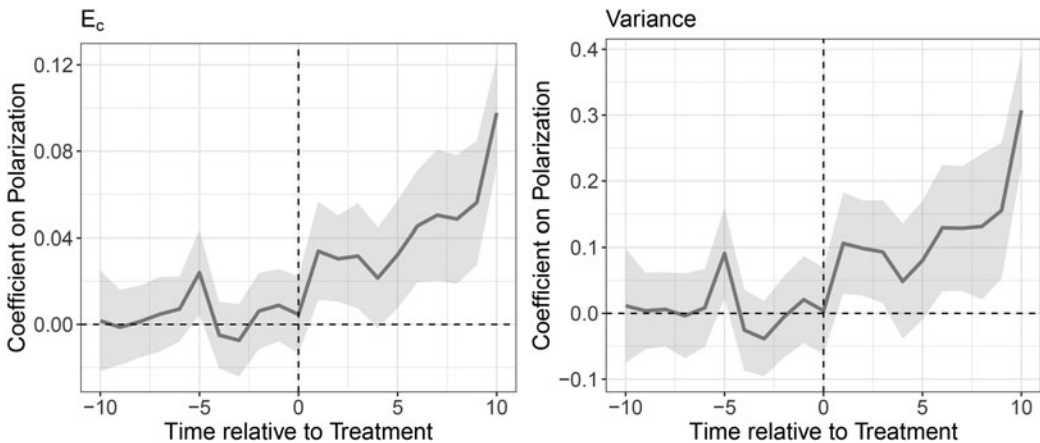


Figure 10. GSCM estimates: effects of radical-right parties on polarization.

Finally, we reproduce the authors’ analysis using generalized synthetic control methods (GSCM) (Xu, 2017), which can provide causal inference with interactive fixed-effect models and exploit synthesized counterfactuals for treated units based on information from untreated groups. Figure 10 reports the GSCM estimates using both entropy-based measure as the outcome (left panel) and the original standard deviations (right panel). It again shows very similar patterns between two measures with the entropy-based measure providing some more nuanced dynamics. The increased efficiency from the measure results in a more steady trend in both the pre- and post-treatment periods with E_c farther away from zero in the post-treatment period. This efficiency difference is not just an artifact of this one example. Since the variance measure enforces (assumes) equal spacing it will always provide greater dispersion over a truly ordinal measure as the distribution deviates from uniformity, which is what we see in Figure 7.

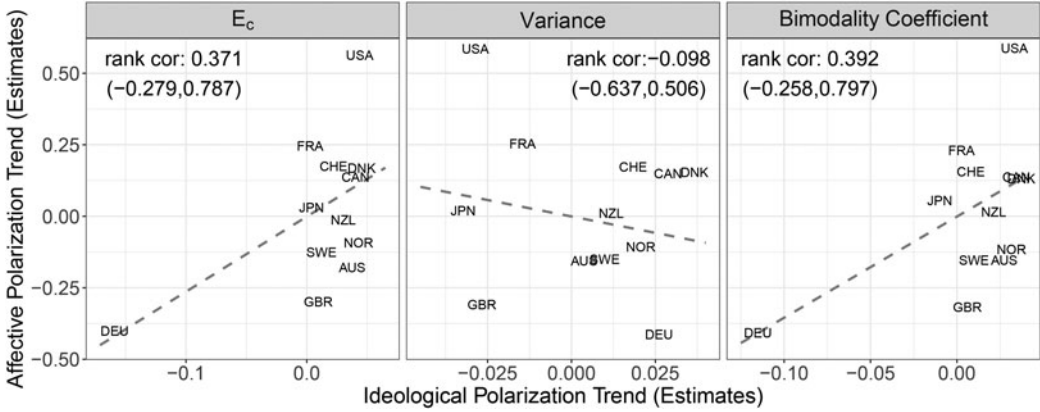


Figure 11. Trends in affective and ideological polarization. Both axis variables are the estimates of linear trend based on Bayesian hierarchical partial-pooling models. The dashed line displays a fitted bivariate linear regression. The correlation estimates are from the Spearman rank correlation and 95 percent confidence intervals are reported in parenthesis.

5.3 Cross-country trends in ideological and affective polarization

Recent work in the polarization literature has observed that mass polarization is not only about where people stand on the issues, but also about how people emotionally dislike those from rival parties (Iyengar *et al.*, 2012). The relationship between ideological and affective polarization has generated ongoing debates in American political studies (see Rogowski and Sutherland, 2016; Mason, 2018). However, cross-country comparisons are elusive, possibly due to the lack of comparable measures of mass ideological polarization. We apply the metrics of polarization to data assembled from multiple survey projects and compare its trend with the recent finding of affective polarization in 12 Organization for Economic Co-operation and Development (OECD) countries for the past three decades.

The affective polarization is defined as the weighted average of respondents’ partisan affect and individual partisan affect is measured by the extent to which an individual expresses a more favorable attitude toward their own party than toward other parties (see also Iyengar *et al.*, 2019; Gidron *et al.*, 2020; Boxell *et al.*, 2022). For ideological polarization, we assemble data from multiple surveys and match them with the affective polarization data set (see online Appendix 6.1). We again apply three polarization metrics to survey items about left–right position and measure the polarization trends of ideology across years for each country. The survey items have similar wordings that ask respondents’ about self-identified left–right positions but different numbers of categories across surveys.

Figure 11 shows a scatter plot and the Spearman rank correlation estimates to compare the trends between ideological polarization (X-axis) and affective polarization (Y-axis) for each measure of polarization. The correlation coefficients between the variables are not statistically reliable in all three cases as the confidence intervals all contain zero. The E_c measure and the bimodality coefficient suggest a positive relationship between ideological polarization and affective polarization while the variance suggests a negative relationship. This again shows that variance as a statistic assuming interval measured data is not responsive to distributional differences in ordinal data across countries. While there are only small differences between E_c and the bimodality coefficient the latter also assumes interval measured data, which obviously is not appropriate here.

Figure 12 shows the ideological polarization for each of the 12 OECD countries from 1990 to 2020. For the time patterns of ideological polarization, the three measures depict both similarity and differences. Note that E_c produces a more consistent linear pattern of polarization for each country, reflecting in that dots in the figure closely center around the line. This further

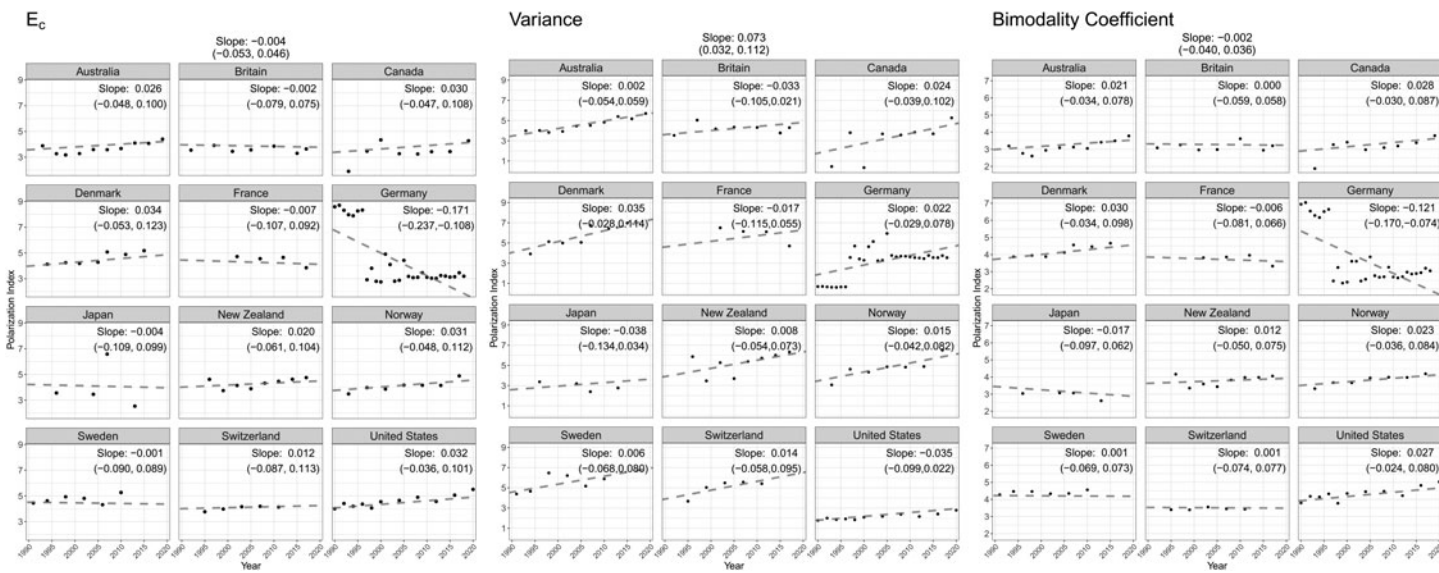


Figure 12. Trends in ideological polarization by country. The slope estimates are based on Bayesian hierarchical partial-pooling models and also depicted by the dashed lines. The 95 percent credible intervals are reported in parenthesis.

demonstrates the construct validity of E_c as a measure of aggregate polarization since a country's polarization at the aggregate level should show a certain level of consistency and pattern over time rather than random fluctuation. Each plot also includes an estimated linear time trend and reports the associated average slope coefficient and individual slope for each country, where the variance indicates a substantially greater increasing trend than the other two (see online Appendix Figures A.7 and A.8). The other two measures are aligned with detailed single in-depth single-country studies using alternative country-specific measures (see Lelkes, 2016; Merkle, 2021), which suggests there is little evidence that indicates the mass public in these types of countries is significantly more ideologically polarized than 30 years ago.

This example further demonstrates the validity and consistency of the entropy-based measure: the E_c provides a more reasonable depiction here of ideological polarization and its connection to affective polarization than the variance, according to both the original data distribution and previous findings. The reason why the bimodality coefficient is similar to the E_c in this example is that cross-country comparisons provide more distributional variation with regards to modality. The alternative proxy statistics of polarization only capture some components of the dispersion and distribution in the ordinal data, which partially describes polarization. Thus, the use of E_c is even more important for comparative studies as countries can present a more heterogeneous set of distributional patterns of preferences and opinions.

6. Conclusion

We introduce a nonparametric, entropy-based method for measuring issue-based mass political polarization that is completely new to the literature. We demonstrate here that the proposed measure is theoretically and conceptually more appropriate for the intuition and structure of polarization, and further, it measures this phenomenon in a way that does not rely on the confusing distinction between dispersion and bimodality typically used in this literature. Unlike these previous methods, our measure exploits the structure of ordinal variables in public opinion surveys such that polarization is revealed in a novel way where it captures the concentration and ordering of the data at the same time. The new measure makes no *a priori* assumptions about the central tendency, spacing between categories, specific forms of distributions, and is therefore fully nonparametric. The hypothetical illustrations, the simulation analysis, and the crowdsourcing validation exercise all demonstrate that our measure is able to reliably reveal the nuanced and complicated dynamics of polarization with different types of empirical distributions. We also apply the measure to three different examples to demonstrate the utility of the entropy approach with real data.

Current studies of polarization mostly focus on single cases, which rely on predefined political and social contexts usually by nation. This does not answer the big questions: *why are some countries increasingly polarized, and how are political systems being stressed by polarization*. At the same time, empirically, there is another layer of its connections to affective and group-based polarization. To investigate such topics requires a reliable measure of mass polarization that can be applied to cross spatial contexts, as we have provided here.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/psrm.2024.24>. To obtain replication material for this article, <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/ATBJNO>.

References

- Abramowitz AI and Saunders KL (2008) Is polarization a myth?. *The Journal of Politics* **70**, 542–555.
- Abramowitz AI and Webster S (2016) The rise of negative partisanship and the nationalization of us elections in the 21st century. *Electoral Studies* **41**, 12–22.

- Aloise D, Deshpande A, Hansen P and Popat P** (2009) NP-hardness of Euclidean sum-of-squares clustering. *Machine Learning* 75, 245–248.
- Baldassarri D and Gelman A** (2008) Partisans without constraint: political polarization and trends in American public opinion. *American Journal of Sociology* 114, 408–446.
- Bischof D and Wagner M** (2019) Do voters polarize when radical parties enter parliament?. *American Journal of Political Science* 63, 888–904.
- Bishin BG, Freebourn J and Teten P** (2021) The power of equality? Polarization and collective mis-representation on gay rights in congress, 1989–2019. *Political Research Quarterly* 74, 1009–1023.
- Blair J and Lacy MG** (2000) Statistics of ordinal variation. *Sociological Methods & Research* 28, 251–280.
- Bosancianu CM** (2017) A growing rift in values? Income and educational inequality and their impact on mass attitude polarization. *Social Science Quarterly* 98, 1587–1602.
- Boxell L, Gentzkow M and Shapiro JM** (2022) Cross-country trends in affective polarization. *Review of Economics and Statistics* 106, 557–565.
- Bramson A, Grim P, Singer DJ, Fisher S, Berger W, Sack G and Flocken C** (2016) Disambiguation of social polarization concepts and measures. *Journal of Mathematical Sociology* 40, 80–111.
- Carsey TM and Layman GC** (2006) Changing sides or changing minds? Party identification and policy preferences in the American electorate. *American Journal of Political Science* 50, 464–477.
- Chissom BS** (1970) Interpretation of the kurtosis statistic. *The American Statistician* 24, 19–22.
- Cover TM and Thomas JA** (1991) Information theory and statistics. *Elements of Information Theory* 1, 279–335.
- Dias N and Lelkes Y** (2022) The nature of affective polarization: disentangling policy disagreement from partisan identity. *American Journal of Political Science* 66, 775–790.
- DiMaggio P, Evans J and Bryson B** (1996) Have Americans' social attitudes become more polarized?. *American Journal of Sociology* 102, 690–755.
- Downey DJ and Huffman ML** (2001) Attitudinal polarization and trimodal distributions: measurement problems and theoretical implications. *Social Science Quarterly* 82, 494–505.
- Druckman JN and Levendusky MS** (2019) What do we measure when we measure affective polarization?. *Public Opinion Quarterly* 83, 114–122.
- Fiorina MP and Abrams SJ** (2008) Political polarization in the American public. *Annual Review of Political Science* 11, 563–588.
- Fiorina MP, Abrams SJ and Pope JC** (2010) *Culture War?: The Myth of a Polarized America*. New York, NY: Pearson Longman.
- Gidron N, Adams J and Horne W** (2020) *American Affective Polarization in Comparative Perspective*. New York, NY: Cambridge University Press.
- Gill J** (2005) An entropy measure of uncertainty in vote choice. *Electoral Studies* 24, 371–392.
- Hill SJ and Tausanovitch C** (2015) A disconnect in representation? Comparison of trends in congressional and public polarization. *The Journal of Politics* 77, 1058–1075.
- Homola J, Jackson N and Gill J** (2016) A measure of survey mode differences. *Electoral Studies* 44, 255–274.
- Iyengar S, Lelkes Y, Levendusky M, Malhotra N and Westwood SJ** (2019) The origins and consequences of affective polarization in the United States. *Annual Review of Political Science* 22, 129–146.
- Iyengar S, Sood G and Lelkes Y** (2012) Affect, not ideology: a social identity perspective on polarization. *Public Opinion Quarterly* 76, 405–431.
- Jaynes ET** (1968) Prior probabilities. *IEEE Transactions on Systems Science and Cybernetics* 4, 227–241.
- Jaynes ET** (1982) On the rationale of maximum-entropy methods. *Proceedings of the IEEE* 70, 939–952.
- Leik RK** (1966) A measure of ordinal consensus. *Pacific Sociological Review* 9, 85–90.
- Lelkes Y** (2016) Mass polarization: manifestations and measurements. *Public Opinion Quarterly* 80, 392–410.
- Levendusky MS and Pope JC** (2010) Measuring aggregate-level ideological heterogeneity. *Legislative Studies Quarterly* 35, 259–282.
- Levendusky MS and Pope JC** (2011) Red states vs. blue states: go beyond the mean. *Public Opinion Quarterly* 75, 227–248.
- Lupu N** (2015) Party polarization and mass partisanship: a comparative perspective. *Political Behavior* 37, 331–356.
- Mason L** (2015) 'I disrespectfully agree': the differential effects of partisan sorting on social and issue polarization. *American Journal of Political Science* 59, 128–145.
- Mason L** (2018) Ideologues without issues: the polarizing consequences of ideological identities. *Public Opinion Quarterly* 82, 866–887.
- Merkley E** (2021) Ideological and partisan bias in the Canadian public. *Canadian Journal of Political Science* 54, 267–291.
- Mosier CI, Myers MC and Price HG** (1945) Suggestions for the construction of multiple-choice test items. *Educational and Psychological Measurement* 5, 261–271.
- Mouw T and Sobel ME** (2001) Culture wars and opinion polarization: the case of abortion. *American Journal of Sociology* 106, 913–943.

- Pace EL, Dennis KS and Berg WT** (1955) Thermodynamic properties of krypton adsorbed on titanium dioxide (rutile). *The Journal of Chemical Physics* **23**, 2166–2168.
- Pearson K** (1905) 'Das fehlergesetz und seine verallgemeinerungen durch fechner und pearson'. a rejoinder. *Biometrika* **4**, 169–212.
- Petrocik JR** (1974) An analysis of intransitivities in the index of party identification. *Political Methodology* **1**, 31–47.
- Pontusson J and Rueda D** (2008) Inequality as a source of political polarization: a comparative analysis of twelve OECD countries. In Beramendi P and Anderson CJ (eds), *Democracy, Inequality, and Representation: A Comparative Perspective*, New York, NY: Russell Sage Foundation, pp. 312–353.
- Rogowski JC and Sutherland JL** (2016) How ideology fuels affective polarization. *Political Behavior* **38**, 485–508.
- Shannon CE** (1948) A mathematical theory of communication. *The Bell System Technical Journal* **27**, 379–423.
- Simpson EH** (1949) Measurement of diversity. *Nature* **163**, 688–688.
- Tsallis C** (2011) The nonadditive entropy S_q and its applications in physics and elsewhere: some remarks. *Entropy* **13**, 1765–1804.
- Wagner M** (2021) Affective polarization in multiparty systems. *Electoral Studies* **69**, 102–199.
- Wang H and Song M** (2011) Ckmeans.1d.dp: optimal k -means clustering in one dimension by dynamic programming. *The R Journal* **3**, 29.
- Wynblatt P** (1969) On the formation and migration entropies of vacancies in metals. *Journal of Physics and Chemistry of Solids* **30**, 2201–2211.
- Xu Y** (2017) Generalized synthetic control method: causal inference with interactive fixed effects models. *Political Analysis* **25**, 57–76.