How Polarised are Citizens? Measuring Ideology from the Ground-Up *

Mirko Draca,† Carlo Schwarz‡

April 21, 2023

(First Version: April 2018)

Abstract

We investigate whether the ideological polarisation of citizens has increased in Western democracies. We propose a novel methodology to identify individual ideologies by applying Latent Dirichlet Allocation to political survey data. This approach indicates that questions related to confidence in institutions play a leading role in defining citizen ideologies, in addition to the questions associated with the traditional left-right scale. We decompose the shift in ideological positions across the population over time and measure polarisation. This reveals evidence of a 'disappearing centre' in a sub-group of countries with citizens shifting away from centrist ideologies into anti-establishment 'anarchist' ideologies. This trend is especially pronounced for the US.

JEL *classification*: D72, C81.
Keywords: Polarisation, Ideology, Populism, Unsupervised Learning.

*We are grateful to Arun Advani, Felix Forster, Lucie Gadenne, Nicola Gennaioli, Andrew Oswald, Debraj Ray, Ken Shotts, Guido Tabellini and seminar participants at the Berkeley Haas, Chicago Harris, Queen Mary University, LSE Management, Lund University, Barcelona Summer Forum, De Montfort University, Geneva University, Warwick University and the Galatina Summer Meetings 2018 for their helpful suggestions. Schwarz was supported by a Doctoral Scholarship from the Leverhulme Trust as part of the *Bridges* program.

†Department of Economics, University of Warwick, Centre for Competitive Advantage in the Global Economy (CAGE)
‡Università Bocconi, Department of Economics, IGIER, PERICLES, CEPR, CAGE, carlo.schwarz@unibocconi.it.

1 Introduction

In political terms, we seem to be living in the midst of the proverbial 'interesting times'. Across established democracies, there appear to be strong trends of political populism and ideological polarisation. In the US, a large body of evidence indicates that the political positions taken by elected representatives in legislatures have sharply polarised. For example, this is apparent in recent work examining partisanship in the use of political language (Jensen et al., 2012; Gentzkow et al., 2019b). In particular, Gentzkow et al. (2019b) isolate this increase as occurring from the mid-1990s onwards, a period when the nature of political communication changed as parties became more acutely strategic with their use of language. Further evidence of 'elite polarisation' is also found in the extensive literature (following Poole and Rosenthal, 1985) that has measured the evolving ideological positions of elected representatives using data on Congressional roll call voting.

By comparison, the evidence about political polarisation amongst the general public (or 'citizens') is more contested than the findings that have emerged for political elites. In the US, contributions such as Fiorina and Abrams (2008) make the point that both the underlying distribution of views across issues and the level of self-identification with 'strong' political positions have been stable over time.[1] In Europe, recent contributions by Algan et al. (2017) and Guiso et al. (2017) have documented a strong pattern of populist politics across the continent that appears to have roots in changing economic conditions. However, these populist trends are not necessarily symptomatic of ideological polarisation. For example, Algan et al. (2017) detect no significant shift in political positioning along the left-right scale in their cross-country sample and pick up a decline in close party identification.

A major issue for the analysis of polarisation is measurement. Most commonly, research has focused on party affiliation or positions on specific issues as the target of analysis. In this paper, we focus instead on *ideology* as a summary measure of political positions that can be used to evaluate polarisation. Following Benabou (2008), ideologies can be defined as 'collectively sustained sets of beliefs' and have most often been defined in terms of a left-right scale. However, some researchers have argued that the concept of citizen ideology itself might be misleading as citizens can mix issue positions from both sides of the traditionally defined ideological spectrum. (Converse, 2006; Kinder and Kalmoe, 2017). Further to this, it is not clear how to position developments such as the recent surge in populist politics in relation to the left-right scale.[2]

[1] Similar scepticism about citizen polarisation in the US is also evident in the studies of Glaeser and Ward (2006), and Ansolabehere et al. (2006), while a recent analysis by Kaplan et al. (2022) emphasises an important trend of rising within-state polarisation.

[2] For example, some recent work has suggested that the left-right ideological model needs to be extended by either incorporating voters' identity (Bonomi et al., 2021), a "globalists" vs "nativists" dimension (Gethin et al., 2022) or moral

In this paper, we therefore propose a new and more flexible approach to measuring citizen ideology and political polarisation, which allows for mixtures of ideologies. In short, the core of our approach is based on applying Latent Dirichlet Allocation (LDA) topic models (Blei et al., 2003) to individual-level survey responses ('issue positions) across a typical range of social and economic questions as a measure of an individual's political views. Topic models are mainly known in the social sciences for their use in the analysis of text data, in particular for their capacity in identifying the latent topic structure that underpins the generation of documents across various corpora.[3]

However, rather than analysing text, we instead make individual-level responses from survey data the main objects of analysis, interpreting the latent topics as *political ideologies* that underpin the generation of individual political beliefs amongst the general public.[4] The advantage of this particular approach is that it is based on a probabilistic generative model of ideology, analogous to a simple structural model. Specifically, the generative model we outline follows a two-stage sequential choice structure, as per Munro and Ng (2020)'s recent topic model of categorical survey responses. The first choice relates to 'type' or ideology assignment, while the second involves the choice of a specific response conditional on the first stage assignment.

This generative model provides a framework for understanding the statistical pattern of political views across the population. Individual beliefs are explained as mixtures of latent ideologies, which in turn can be defined as an importance ranking over issue positions. That is, our ideology or 'type vectors' give the probability that an individual will take a particular issue position, given that they are drawing from a certain type. The overall framework therefore provides useful information on the individual-level mixture of ideologies (the type shares) and the nature of the ideologies (the type vector). In the traditional left-right model, there are policies or positions that are left or right. Hence, we can represent individuals on a 1-simplex (i.e., a line segment). Similarly, we can think about our ideological types as points in a higher dimensional simplex rather than a 'line'.

We use this methodology to explore two main questions. Firstly, we ask: to what extent do the general public hold beliefs that can be summarised as statistically coherent 'ideologies'? Further to this point, to what extent do the latent ideologies found in the data conform to the traditional left-right ideological line that dominates both popular discourse and classic formal models in the spirit of Downs (1957)? The second main question we address is then: how do the empirically-based citizen ideologies

foundations Enke (2020).

[3]Applications of topic modelling have thus proliferated recently with empirical studies of text data across a range of social science questions (Gentzkow et al., 2019a).

[4]Within economics, the general approach we take here for analysing discrete, non-text data is closest to Bandiera et al. (2020)'s empirical model of behavioural manager 'types' in CEO time-use data.

we identify vary across countries and over time? In practical terms, this involves studying the factors that determine the ideological mixture of views held by citizens at the individual level and assessing the extent to which aggregate shifts can be explained in terms of changing demographics or other observables.

The main data source we use in our analysis is the cross-country World Values Survey (WVS) which provides a wide-ranging set of consistently asked questions from the late 1980s onwards. In answer to our first major question, a series of coherent citizen ideologies do indeed emerge from our modelling. These ideologies are alternatively defined by a set of issue positions that are consistent with traditional 'left-right' perspectives on social issues and then by issue positions that relate to citizen confidence in institutions. We generically label the ideological types that are characterised by low confidence in institutions as 'anarchist', but note that the broad position that these types represent is consistent with the anti-establishment or populist positions that have been the focus of recent research (Acemoglu et al., 2013; Gethin et al., 2022; Rodrik, 2018).[5]

This central result regarding the defining role of confidence in institutions for citizen ideologies is robust to a comprehensive set of exercises where questions are left-out individually or according to sub-groups. A similar pattern of 'low-trust' types also emerges when we apply our model to a different dataset (the European Social Survey (ESS) with a comparable set of questions.

The unsupervised machine learning model allows us to identify a clear hierarchy of empirical ideologies that emerges as we fit models that are premised on different numbers of *ex-ante* types. Following Chang et al. (2009), we implement a 'topic cohesion' measure that allows us to select an empirical model of ideology that is based on 4 types. We label these types as Liberal Centrist, Conservative Centrist, Left Anarchist, and Right Anarchist.

Next, we use our findings regarding the structure of the ideologies to analyse the variation of ideologies across countries and time periods. Firstly, at the level of the latent ideologies, we find that our 4 main ideological types are stable over time with limited 'within ideology' changes, as measured by the weighting of different issue positions. The most notable finding here is an increase in the intensity of socially liberal attitudes across most types. For example, the Conservative Centrist type shifts in their attitudes on issues such as homosexuality and abortion.

Secondly, we use the information on individual type shares (the mixture parameter in our LDA model) to measure how prevalent different ideologies are across countries and how this changes over

[5]The anarchist label that we use is meant to avoid pejorative interpretations of terms such as 'populist' and emphasise opposition to current institutional structures as the defining feature of this ideological type. For example, see media critiques such as 'Populism: It's the BBC's new buzzword, being used to sneer at the 'uneducated' 17 million who voted for Brexit' from the UK's *Daily Mail* (Murray (2016)).

time. The general pattern is consistent with the existing literature - for example, northern European countries are more liberal, while countries with stronger religious traditions are more conservative. In turn, this is reinforced by a sensible pattern of correlations between individual-level characteristics and type shares (e.g., women are more liberal and conservatism increases with age). Our main finding here is that the composition of the aggregate type shares is stable across time for most countries. However, a notable exception is the US, where the total type share for the two Anarchist types increases from around 30% in the 1989-1993 wave to 50% by the fifth WVS wave in 2005-2009. The majority of this increase is accounted for by the Right Anarchist type.

The ideological type shares also have interesting relationships with variables representing self-positioning on the Left-Right scale and the probability of voting for 'populist' political parties. We find a strong relationship between type shares and Left-Right self-positioning, and the anarchist ideological type shares also prove to be strong predictors of populist voting. For example, we estimate that an individual with a 50% type share in either the Left or Right Anarchist ideologies has a 38% higher probability of voting populist relative to the mean, even after controlling for Left-Right self-positioning and other covariates.

The final part of our analysis then uses the outputs of the empirical LDA model to devise two further measures of ideological structure. The first 'citizen slant' measure provides a within-person measure of ideological concentration and is constructed following a basic Gini index logic. It directly exploits the mixed membership format of our unsupervised learning framework to capture how partisan individuals are in their ideological views. We find that the mean citizen slant across types, countries and years is relatively high at around 0.75 on a 0-1 scale. The degree of slant or within-person concentration has also increased over the time window we consider. There is a slight increase in the case of Europe, but much stronger shifts are apparent in the US. The rise in the US is also focused heavily on the Anarchist types (which increased their slant by around 15%) as well as the Centrist Conservative type (a 5% increase).

The second societal polarisation measure that we put forward builds on the framework of Esteban and Ray (1994), and Duclos et al. (2004). This framework allows us to develop a novel, multi-polar analysis of ideology in terms of own-group *identification* and between-group *alienation*. Practically, this is achieved by leveraging the information on relative group size within countries (where group membership is defined according to the dominant type share), alongside the other information from the LDA model outlined above. In this way, our approach allows us to study cross-country trends in ideological polarisation, which complement existing international measures of affective polarisation that are principally based on party identification (see for example Boxell et al., 2020). We find that changes in the level of polarisation over time are muted. Again, the US stands out as experiencing the sharpest increase, chiefly driven by the compositional change in type shares noted above. Interestingly, the nature of the US polarisation experience is more characteristic of a 'disappearing centre' driven by the growth of anarchist types than it is by a traditional left-right division.

Related Literature. This paper is related to several strands of literature. Three areas to highlight are the following. Firstly, there is the literature on democratic politics and populism, with recent examples that include: Acemoglu et al. (2013), Algan et al. (2017), Buisseret and Van Weelden (2020), Bursztyn et al. (2020), Dal Bó et al. (2017), Dal Bó et al. (2018), Guiso et al. (2017), and Rodrik (2018)[6]. Earlier work by Benabou (2008) discusses citizen ideology in the context of positions on the state versus the market. As discussed, our work sheds light on the potential long-run ideological underpinnings of these various political trends in the population.

Secondly, there is a fast-growing literature that studies aspects of ideology, policy-making and political communication using tools from machine learning and natural language processing. This includes the already noted Gentzkow et al. (2019b) and Jensen et al. (2012), as well as other text-based studies such as: Ash (2015), Grimmer (2009), Hansen et al. (2014), Cagé et al. (2020) and Jelveh et al. (2015). Another branch of this overall literature (Blaydes and Grimmer (2013), Gross and Manrique-Vallier (2012), Hill and Tausanovitch (2015), Wang et al. (2017), Munro and Ng (2020), Desmet et al. (2022)) has also begun to explore the application of unsupervised learning tools to survey response data. To the best of our knowledge, we are the first to use LDA for the study of polarisation in survey data.

Finally, there is a large literature that explicitly addresses polarisation and fractionalisation along political, ethnic and cultural lines. This literature often focuses on measuring group structure in societies and relating this to patterns of conflict. An indicative list includes: Alesina et al. (2003), Bossert et al. (2011), Caselli and Coleman (2013), Duclos et al. (2004), Esteban and Ray (1994) and Montalvo and Reynal-Querol (2005), Canen et al. (2020). Some recent work of interest here includes Bertrand and Kamenica (2018), who measure 'cultural distance' between population sub-groups in the US and find a constant relationship amongst most outcomes and group splits. They do however note divergences in social attitudes based on political ideology and income. Desmet and Wacziarg (2021) also examine cultural distance, again finding stability across most dimensions.

The approach to identifying social sub-groups in a purely data-driven way has the potential to inform the emerging literature on identity politics (Atkin et al., 2021; Grossman and Helpman, 2021; Shayo, 2009). Currently, this literature has focused on identity groups whose definition hinges on

[6]A range of studies that have looked at the recent determinants of voting patterns are also relevant here: Becker et al. (2017), Dippel et al. (2015), Dorn et al. (2020) and Che et al. (2016).

ex-ante characteristics (e.g., race, gender, income class). Our methodology shows that there is scope to define latent social sub-groups based on observable positions. We also note that the 'identification and alienation' framework used in our polarisation measures is directly analogous to key concepts in the identity literature (Akerlof and Kranton, 2000) and therefore provides metrics to study the potential frictions between social groups over time.

Structure. The paper is organised in the typical way. In section 2, we outline the main data used, namely the World Values Survey (WVS), as well as our approach to defining answers to survey questions as 'issue positions'. Section 3 describes our unsupervised learning methodology for studying this issue-position data. This includes details on how we develop a hierarchy of ideological types and select the optimal number of topics in our LDA models. Section 4 outlines the results, and section 5 concludes.

2 Data

World Values Survey

For our main analysis, we use data from the World Values Survey (WVS) and the European Values Study (EVS). These surveys are an output of a global research project conducted by a large network of social scientists and run via a non-profit association based in Stockholm.[7] The WVS consists of 7 waves from 101 countries, while the EVS consists of 5 Waves from 48 countries. We construct what is formally known as the Integrated Value Survey (IVS) by combining the two datasets. The resulting dataset contains the 4 EVS waves and the corresponding waves 1, 2, 4, 5 and 7 from the WVS. For the sake of simplicity, we refer to this combination of the data as the 'World Values Survey (WVS)'.

The set of questions asked and countries covered differs across successive waves of the WVS. We therefore develop a sample of WVS observations based on the principle of capturing the widest range of consistently asked questions over waves and across countries[8]. Since the first and seventh waves have more limited country and question coverage[9], we construct our sample from the second to the fifth wave and develop a set of 17 countries in Europe and North America (Austria, Belgium, Canada, Denmark, Finland, France, Germany, Iceland, Ireland, Italy, Malta, Netherlands, Portugal, Spain, Great Britain, United States, North Ireland) and 29 questions. The selected questions cover issues such as abortion,

[7]They have been widely deployed in social science research, and some prominent studies using the data include: Alesina et al. (2013, 2001); Blanchflower and Oswald (2008); Inglehart (1997); and Norris (2016).

[8]We provide additional details on the selection of questions in Appendix A. We also later find (in Appendix E) that the basic structure of the ideological clusters we identify is robust to the inclusion or exclusion of questions.

[9]In the first wave, the countries Austria and Portugal, as well as 7 complete questions, are not available. The 7th wave misses Belgium, Canada, Ireland, Malta, and Northern Ireland.

immigration, sexuality, the role of government, and confidence in institutions. The resulting dataset contains a total of 82,338 observations over 3 waves spanning the years from 1989 to 2010.

A 7th wave of the WVS was conducted to cover the period 2017-2020. However, at the time of writing, the available 7th wave data still lacked information on 4 countries. Hence, in the online Appendix H we provide separate results for the 7th wave of the WVS. These results are directly in line with those in our main analysis below.

Construction of Features

As part of the data preparation, we unify the coding of the questions and convert them to the same scale. The intention here is to represent the answers to the survey questions as discrete 'features' for the subsequent topic modelling. Specifically, we recode the responses for each of the 29 questions into two indicator variables expressing either support for or opposition to each issue, for example, an indicator variable if the person believes that abortion is justifiable and a second indicator variable if the person opposes abortion. In cases where a person expressed neither support nor opposition to an issue, both binary variables are coded as zero.

Summary statistics for the 58 recoded issue positions can be found in Appendix Table A.1. Importantly, the features cover a broad range of salient political issues. Several questions deal with what would be typically classified as 'social issues' such as abortion, prostitution and attitudes towards minority groups, while three questions deal with classic economic questions relating to the role of government, private sector competition and support for the welfare state. Finally, there is a set of questions dealing with confidence in a comprehensive set of social and political institutions.

The information in Appendix Table A.1 indicates a rich mix of positions across political issues. There is a current of anti-foreigner sentiment, with 12.3% of respondents preferring not to have immigrants as neighbours, and this is backed up by an overwhelming 60% endorsing a priority for native workers in the allocation of scarce jobs. However, most respondents either hold liberal or neutral views on leading social issues such as abortion and prostitution. There also is a widespread lack of confidence in key institutions, with only around 35-45% expressing a favourable view of the press, parliaments, the civil service and major companies.

3 Discovering Latent Ideology

The empirical model of ideology that we present below posits citizens as selecting particular issue positions based on a two-step process. The first step involves choosing a type or ideology assignment.

This assignment which provides a probabilistically structured information source for the second step of choosing the actual issue position. The types or ideologies therefore represent belief systems that are used as resources for individuals to choose a vector of issue positions.

3.1 Discovering Citizen Ideology via Latent Dirichlet Allocation (LDA)

Underlying our approach is a general model of individual choices over policy positions. Each individual $i \in I$ in our data can be thought of as maximising her utility by choosing a set of policy positions $r_{i,n}$. A policy position could, for example, be opposition to abortion. In principle, the individual is free to choose any number of policy position N_i among the $Q = 58$ recoded issue positions in our data.[10] As per Bénabou and Tirole (2016), individuals can have either instrumental or affective motives for holding beliefs or policy positions. That is, they can subscribe to a policy position because of its direct tangible benefits (instrumental motive) or because of its consumption value (affective).

Following Munro and Ng (2020)'s topic model of categorical survey responses, an individual follows a sequential choice approach in determining their set of policy positions. The first step involved here is type or ideology assignment. When formulating a response $r_{i,n}$ for policy position Q, an individual uses information from one of $t \in T$ ideological types to make their decision. Specifically, individual i chooses an ideological type assignment $z_{i,n} \in T$ for each potential policy position $n \in N_i$:

$$z_{i,n} = \arg\underbrace{\max}_{t \in T} U_i(t) = \sum_{t=1}^{T} I(t = j)(e_{i,t,n}) \tag{1}$$

where $e_{i,t,n}$ captures the idiosyncratic utility individual i receives if he chooses type t for response n. An overall ideological type share for individual i can then be calculated as the share of responses that individual i draws from type t. Denoting this as $\theta_{i,t}$ it is written simply as:

$$\theta_{i,t} = \frac{1}{N_i} \sum_{n \in N_i} I(z_{i,n} = t) \tag{2}$$

$$\theta_{i,t} = \frac{1}{N_i} \sum_{n \in N_i} P((e_{i,t,n} = \underbrace{\max}_{t \in T}((e_{i,t,n}))) \tag{3}$$

Note that this structure means that an individual can freely mix policy positions from the T ideologies with a vector θ_i of length T characterising the share of policy positions individual i is drawing from each of the ideologies. For example, in a simple 'left-right' case where $T = 2$, the θ_i vector could describe an individual's responses as 20% 'conservative' and 80% 'liberal'. The concept of ideology we use here is closely related to that of Downs (1957), who defines ideologies as the

[10] Given that the responses are derived from a survey, the individual number of responses N_i varies between 20 and 29 for nearly all i. See notes to Table A.1 for more details on the coding of the questions.

party platforms or belief systems that help an individual reduce the information costs of weighing all competing policies against each other. The T latent ideologies in our model work in the same way to help the individual to structure her responses.[11]

The second step in this sequential choice approach involves the actual choice of an optimal response $r_{i.n}$ from the set of possible responses Q such that:

$$r_{i,n} = \arg \underbrace{\max}_{v \in Q} B_i(v) = \sum_{u=1}^{Q} I(v = u)(q_{z_{i,n},u} + s_{i,u}) \tag{4}$$

where $s_{i,u}$ captures the idiosyncratic utility from giving response $v \in Q$. Similarly, $q_{z_{i,n},u}$ is the utility the individual receives for giving a response based on the type assignment $z_{i,n}$ of this response.
The generative process underlying the data is given as:

1. For each ideological type t draw a distribution over policy position Q such that $\beta_t \sim Dir(\gamma)$, where γ is a hyperparameter

2. For each individual i in the data draw ideological type shares $\theta_i \sim Dir(\alpha)$, where α is a hyperparameter

3. For each of the $n \in N_i$ responses of individual i which we refer to as $r_{i,n}$:

    - Draw a type assignment $z_{i,n} \sim Mult(\theta_i)$
    - Draw a response $r_{i,n}$ from $P(r_{i,n}|z_{i,n}, \beta)$

Given this generative process, the probability of the observed survey responses is:

$$P(\theta, z, r, \alpha, \beta) = \prod_{t=1}^{T} P(\beta_t|\gamma) \prod_{i=1}^{I} P(\theta_i|\alpha) \left(\prod_{n=1}^{N_i} P(z_{i,n}|\theta_i) P(r_{i,n}|z_{i,n}, \beta) \right) \tag{5}$$

The first term is the probability of observing the ideological type vector β_t. The second term describes how likely it is to observe an individual's ideological type shares θ_i. The third term in brackets is the probability of observing the responses of individual i. LDA identifies ideological types by finding parameter values for β_t and θ_i such that this probability is maximised. Due to dimensionality, simply maximising this likelihood for the relevant parameters is computationally unfeasible. LDA therefore

[11]The advantage of this concept of ideology is that it is completely general, and in addition to concepts like party platforms, it can also encompass ideas like the moral foundations (Haidt, 2012), class and cultural identity (Bonomi et al., 2021), narratives (Akerlof and Rayo, 2020) or moral universalism (Enke et al., 2022), as all of these presume common forces that shape the ideologies of individuals.

makes use of an approximate inference algorithm. We use the inference algorithm developed by Hoffman et al. (2010, 2013) and implemented by Pedregosa et al. (2011).[12]

As in nearly all clustering algorithms, LDA itself does not provide any topic labels, and the standard machine learning topic labelling approaches (e.g. Lau et al., 2011; Aletras et al., 2014) are not applicable in our setting. Therefore, it is up to the user to interpret and judge the focus of each topic.

In our application, as in most applications of LDA, the assumption of the independence of responses does not strictly hold. If a question has been answered, the same question cannot be answered again by the same person. We discuss this in detail in Appendix B, with specific reference to the survey data application of Gross and Manrique-Vallier (2012). In short, the inference of LDA is nonetheless still valid, and $z_{i,n}$ still represents the correct probability of a person belonging to one ideological type. Only the interpretation of the β vectors changes. We provide an exposition of this specific point about the β vectors in Appendix C.

Comparison with Other Clustering Methods. The generative approach of LDA has the key advantage of offering a potential microfoundation with parameters that have a direct empirical interpretation. In contrast, while both Principal Component Analysis (PCA) or Factor Analysis (FA) have been widely used to either identify the big 5 personality traits (e.g. Tupes and Christal, 1961; Norman, 1963) or the general intelligence factor g (e.g. Spearman, 1904), neither models the latent types of each individual directly. Moreover, both PCA and FA use linear transformations of the data, while LDA allows for non-linear relationships. Overall, LDA is hence better suited for categorical data than either PCA or FA. Another advantage of LDA is that it is a mixed membership model which describes every observation as a mixture of types rather than in terms of some attachment to a single type or category, as in Latent Class Analysis, k-means, or spectral clustering. The method of Relational Class Analysis (RCA), as applied by Baldassarri and Goldberg (2014), has elements of a mixed membership structure but is not formally probablistic.

Relation to Earlier Definitions of Ideology Our approach to defining and measuring ideology aligns most closely with Converse (2006) within political science who focuses on 'belief systems' as held by elites and the mass public. Converse defines a belief system as configuration of ideas in which elements are bound together by a functional interdependence or constraint. Practically, this 'constraint' implies predictability: if ones knows an individual's position on a given issue then their position on other issues

[12] α, γ are the Dirichlet priors, are we set to $\alpha = 0.25$ and $\gamma = 0.1$. In Appendix E, we provide robustness for alternative priors. Overall, the resulting types are highly robust.

can be successfully forecast. This work also suggests a difference between elite and popular ideologies. These themes were followed up by Kinder and Kalmoe (2017) who argue that there are significant ‘non-ideological’ groups amongst the US public.

3.2 Determining the Optimal Number of Types

LDA makes it possible to estimate any number of ideological types. Therefore, the question of model selection is crucial for understanding which level of topic model best describes the data. In recent years, several methods for the understanding of topic cohesion in text data have been developed (e.g. Chang et al., 2009; David Newman et al., 2010; Aletras and Stevenson, 2013; Lau et al., 2014). We modify these methods for the application to our "issue position" data. One advantage of our topic cohesion approach is that it is also applicable to any LDA model, especially non-text data. As such, our approach could find application in any setting where previously the number of topics was simply chosen by the authors.

Our approach follows standard k-fold cross-validation principles. K-fold cross-validation works by fitting models to different parts or ‘folds’ of data. These models are then evaluated against each other based on an appropriate measure of model fit. As is standard in machine learning, the model with the best fit is used for analysis.

In our case, we first randomly split the data from each wave in our sample into 10 folds (each 10% of the data). Nine folds are then grouped into a training sample, and the remaining becomes the test sample. Afterwards, we fit 10 LDA models with different numbers of types (1 type up to 10 types) to the training sample. In each run of LDA, a different test sample is chosen, and we evaluate the fit of each model relative to this hold-out data.

The optimal number of ideological types is then automatically chosen based on the cohesion of the generated types. A type is more cohesive if the issue positions with the largest weight for that type also frequently appear together in the held-out survey responses of WVS participants. The intuition behind this is that more cohesive ideological types should put more weight on issue positions that people frequently hold together, e.g., the co-occurrence of the views that abortion and suicide are not justifiable. This approach is preferable to evaluating the likelihood, or the perplexity of the model in the hold-out data since the hold-out likelihood is not necessarily a good predictor for human judgment of topic cohesion (see for example Chang et al., 2009).

As a measure of the co-occurrence of issue positions, we use Normalised Pointwise Mutual Information (NPMI). NPMI is defined as:

$$NPMI_{k,l} = \frac{PMI_{k,l}}{-\ln(p(k,l))} = \frac{\ln\left(\frac{p(k,l)}{p(k)\cdot p(l)}\right)}{-\ln(p(k,l))} \tag{6}$$

Pointwise Mutual Information (PMI) is simply defined as the log-ratio of the joint and marginal probabilities. Hence, PMI measures how probable it is that two features k and l appear together in comparison to how often we would expect them to appear together if the features were independent of each other. NPMI additionally normalises PMI between $[-1, 1]$. If two features always appear together, their NPMI will be 1. In the case where two features never appear together, their NPMI will be -1.[13]

The average NPMI for all pairwise combinations of the B most important issue positions of an ideological type t is then given by:

$$\overline{NPMI_t} = \frac{\sum_k^B \sum_{l \neq k}(NPMI_{k,l})}{B \cdot (B-1)} \tag{7}$$

Similarly, the overall cohesion for a model with M ideological types can be calculated from the hold-out sample as:

$$Cohesion_m = \frac{\sum_t^M \overline{NPMI_t}}{M} \tag{8}$$

Follow the findings of Lau and Baldwin (2016), we average our measure of cohesion over different numbers of features $B \in (5, 10, 15, 20)$. As we discuss later, based on these scores, we choose the 4-type LDA specification as our benchmark model since it seems to describe the pattern of responses across citizens best.

3.2.1 Dynamic Type Models - Ideological Change Over Time.

The three waves of the WVS that we use stretch over 20 years. For our analysis, we want to allow for the ideological types to change over time. We do this by fitting LDA models separately to the 3 waves in our sample and only linking the ideological types together afterwards based on the similarity of their issue positions. Our approach is more generic than a dynamic topic model (Blei and Lafferty, 2006) or continuous topic model (Wang et al., 2008) since we neither impose any assumptions on the dynamics of the ideological types nor on the shares of the types over time. The general structure of our approach is most closely related to the topic chains suggested in Kim and Oh (2011) and has the advantage of allowing for completely different ideological patterns to emerge in each wave. But, as we will see, the ideological types in our WVS data display a high degree of stability over time.

[13] More details on the topic cohesion literature and an example for the calculation of NPMI can be found in Appendix D.

4 Results

We report our results across four linked sub-sections. In the first sub-section, we show the results of our LDA models in terms of different variants of types - from 2-types to 5-types. The results here indicate a coherent hierarchy of types across the models such that types can be seen to 'split off' into related families as we move to higher-order models. The second sub-section then applies the NPMI model selection criteria outlined above to the different orders of type models with the conclusion that the 4-type model is the most preferred.

We then use the 4-type model as our main vehicle of analysis in the third sub-section, focusing on within-type and between-type differences over time. To guide the reader, this boils down to a close study of the β type vectors in the LDA model, that is, the ranking of issue positions of each estimated type. In the final subsection, we focus on how the distribution of type shares - essentially the θ_i values outputted by the LDA model - play out over countries and time. In turn, this leads to our analysis of within-person slant and country-level polarisation.

4.1 Hierarchy of Ideological Types.

In Table 1, we summarise the results of various orders of LDA models, reporting the 'top ten' features for each type. These top ten features represent the issue positions with the highest probability values in the β type vectors and are effectively the defining features of each ideological type. We present the results as separate panels in the table per order of type model.[14]

Panel (a) shows the results for the basic 2-type model in the first column. These two types are distinguished by stances on social issues - for example, a liberal attitude towards minority groups (e.g., reporting 'no problems' with neighbours who are homosexuals or immigrants) by one type and conservative positions on social issues such as abortion and prostitution by the other type. We therefore label these types in panel (a) generically as 'Left' and 'Right'. Across the 58 features, the β topic vectors for these types have a correlation of 0.39, indicating that they have some common positions.[15]

The second column of panel (a) then reports the top features for the 3-type model. Two 'Left' and 'Right' types distinguished mainly by their positions on social issues such as sexuality, race and abortion are still apparent. However, the most striking result from this model is the nature of the third

[14]The type hierarchy we show in Table 1 relates to wave 5 only. We focus on this wave as it represents the latest version or 'most recent evolution' across time of our basic types. However, as we show across multiple exercises (cohesion-based model selection, pooled wave model), the type structure is very stable over time. For completeness, we report the top ten features for the pooled wave model in Appendix Table F.2.

[15]Note that while economic issues like beliefs about competition do not appear among the 10 most important features for the centrist types, they are among the features that separate these types (see Appendix Table F.5).

Table 1: Hierarchy of Types (Top Ten Features)

(a) 2-4 Type Model

2 Type Model	3 Type Model	4 Type Model
Left	**Liberal Centrist**	**Liberal Centrist**
No Problem Neighbours: Homosexuals	Confidence: Police	Confidence: Police
No Problem Neighbours: People different race	Confidence: Justice system/courts	No Problem Neighbours: Homosexuals
No Problem Neighbours: People AIDS	No Problem Neighbours: Homosexuals	No Problem Neighbours: People different race
No Problem Neighbours: Immigrants/foreign workers	No Problem Neighbours: People different race	Justifiable: Divorce
Justifiable: Divorce	No Problem Neighbours: Immigrants/foreign workers	Proud of nationality
Not Justifiable: Someone accepting a bribe	No Problem Neighbours: People AIDS	No Problem Neighbours: People AIDS
Proud of nationality	Proud of nationality	Not Justifiable: Someone accepting a bribe
Justifiable: Homosexuality	Not Justifiable: Someone accepting a bribe	No Problem Neighbours: Immigrants/foreign workers
Justifiable: Euthanasia	Confidence: The civil services	Not Justifiable: Claiming government benefits
Not Justifiable: Claiming government benefits	Not Justifiable: Cheating on taxes	Confidence: Justice system/courts
Right	**Conservative Centrist**	**Conservative Centrist**
Not Justifiable: Someone accepting a bribe	Not Justifiable: Abortion	Confidence: Police
Proud of nationality	Not Justifiable: Prostitution	Confidence: Churches
Not Justifiable: Suicide	Not Justifiable: Suicide	Confidence: Armed forces
Not Justifiable: Cheating on taxes	Proud of nationality	Not Justifiable: Suicide
Not Justifiable: Prostitution	Not Justifiable: Someone accepting a bribe	Not Justifiable: Prostitution
Not Justifiable: Avoiding a fare on public transport	Not Justifiable: Avoiding a fare on public transport	Not Justifiable: Abortion
Not Justifiable: Claiming government benefits	Not Justifiable: Cheating on taxes	Proud of nationality
Not Justifiable: Abortion	Not Justifiable: Claiming government benefits	Confidence: Justice system/courts
No Problem Neighbours: People different race	Not Justifiable: Homosexuality	Not Justifiable: Someone accepting a bribe
Confidence: Police	Not Justifiable: Euthanasia	Confidence: The civil services
	Anarchist	**Left Anarchist**
	No Confidence: Civil services	No Confidence: Churches
	No Confidence: Parliament	Justifiable: Divorce
	No Confidence: Churches	No Problem Neighbours: Homosexuals
	No Confidence: Major companies	No Problem Neighbours: People AIDS
	No Confidence: Justice system/courts	No Problem Neighbours: People different race
	No Confidence: The press	No Problem Neighbours: Immigrants/foreign workers
	No Problem Neighbours: Homosexuals	No Confidence: Parliament
	No Problem Neighbours: People different race	Justifiable: Homosexuality
	No Problem Neighbours: People AIDS	No Confidence: Armed forces
	No Confidence: Labour unions	No Confidence: Major companies
		Right Anarchist
		No Confidence: Parliament
		No Confidence: Civil services
		No Confidence: Justice system/courts
		No Confidence: The press
		No Confidence: Labour unions
		No Confidence: Major companies
		Not Justifiable: Someone accepting a bribe
		Not Justifiable: Claiming government benefits
		Not Justifiable: Avoiding a fare on public transport
		Not Justifiable: Cheating on taxes

type. Rather than being a simple mixture of the basic Left-Right types of the earlier model the third type draws on a qualitatively different set of issue positions for its top features. Specifically, the third type draws heavily on features that represent low confidence in major institutions such as parliament, the civil service, the press and major companies. We provide a more detailed discussion of the rationale for our type labels in the next sub-section, but here we flag type 3 as an 'Anarchist' type to reflect this type's opposition to the current workings of major social institutions. In contrast, the main left and right types in the 3-type model report confidence in institutions across the majority of features in this category. We label these types as 'Liberal Centrist' and 'Conservative Centrist' to reflect their contrasting positions on social issues but common pattern of support for established political institutions.[16]

The top features for the 4-type model are reported in the third column of panel (a), Table 1. The type structure continues to evolve here. Most notably, two anarchist types now become apparent, again distinguished by contrasting views on social issues but similar positions in terms of confidence (or lack thereof) in institutions. Note also that the anarchist types distrust different institutions. For example, the left anarchist expresses distrust in churches and the military while the right anarchist lacks trust in civil services, labour unions and the press. Both express distrust in the parliament and major companies. These are labelled 'Left Anarchist' and 'Right Anarchist' to reflect this.[17] Intuitively, the top ten features reported in panel (c) suggest a splitting of the Anarchist type from the 3-type model has occurred. Further evidence on which features separate the types for this 4-type model can be found in Table F.5 in the Online Appendix, which reports the most important type differences.[18]

We can validate this by examining the cross-model correlations in the weights on issue positions in the β type vectors. These correlations are useful for indicating how close the individual types in the 4-type model are to those in the lower-order 3-type order. We report these in Figure 1. In line with the intuitive 'eyeballing' of the top features, the Left Anarchist and Right Anarchist types are most strongly correlated with the Anarchist type from the 3-type model, with correlation measures of 0.86 and 0.68, respectively. This splitting of the Anarchist type is reinforced by the continuity in the Liberal Centrist and Conservative Centrist types as we go from the 3-type to 4-type model. These two types can be tracked across the different hierarchies of type model, with correlations of 0.97 (Liberal Centrist) and 0.88 (Conservative Centrist) across the models.

[16]Note here that the Conservative Centrist type in the 3-type model reports confidence in the churches, armed forces and police as its 11th-13th ranked features.

[17]In terms of social issues, the Right Anarchist type also reports pride in his nationality, opposition to prostitution and drug addicts in the 11th-15th ranked features.

[18]This role for trust in shaping personal ideology has some explicit precedents in political science research. See the body of work by Marc Hetherington, such as Hetherington (2005,2008)

Figure 1: Hierarchy of Types as created by LDA

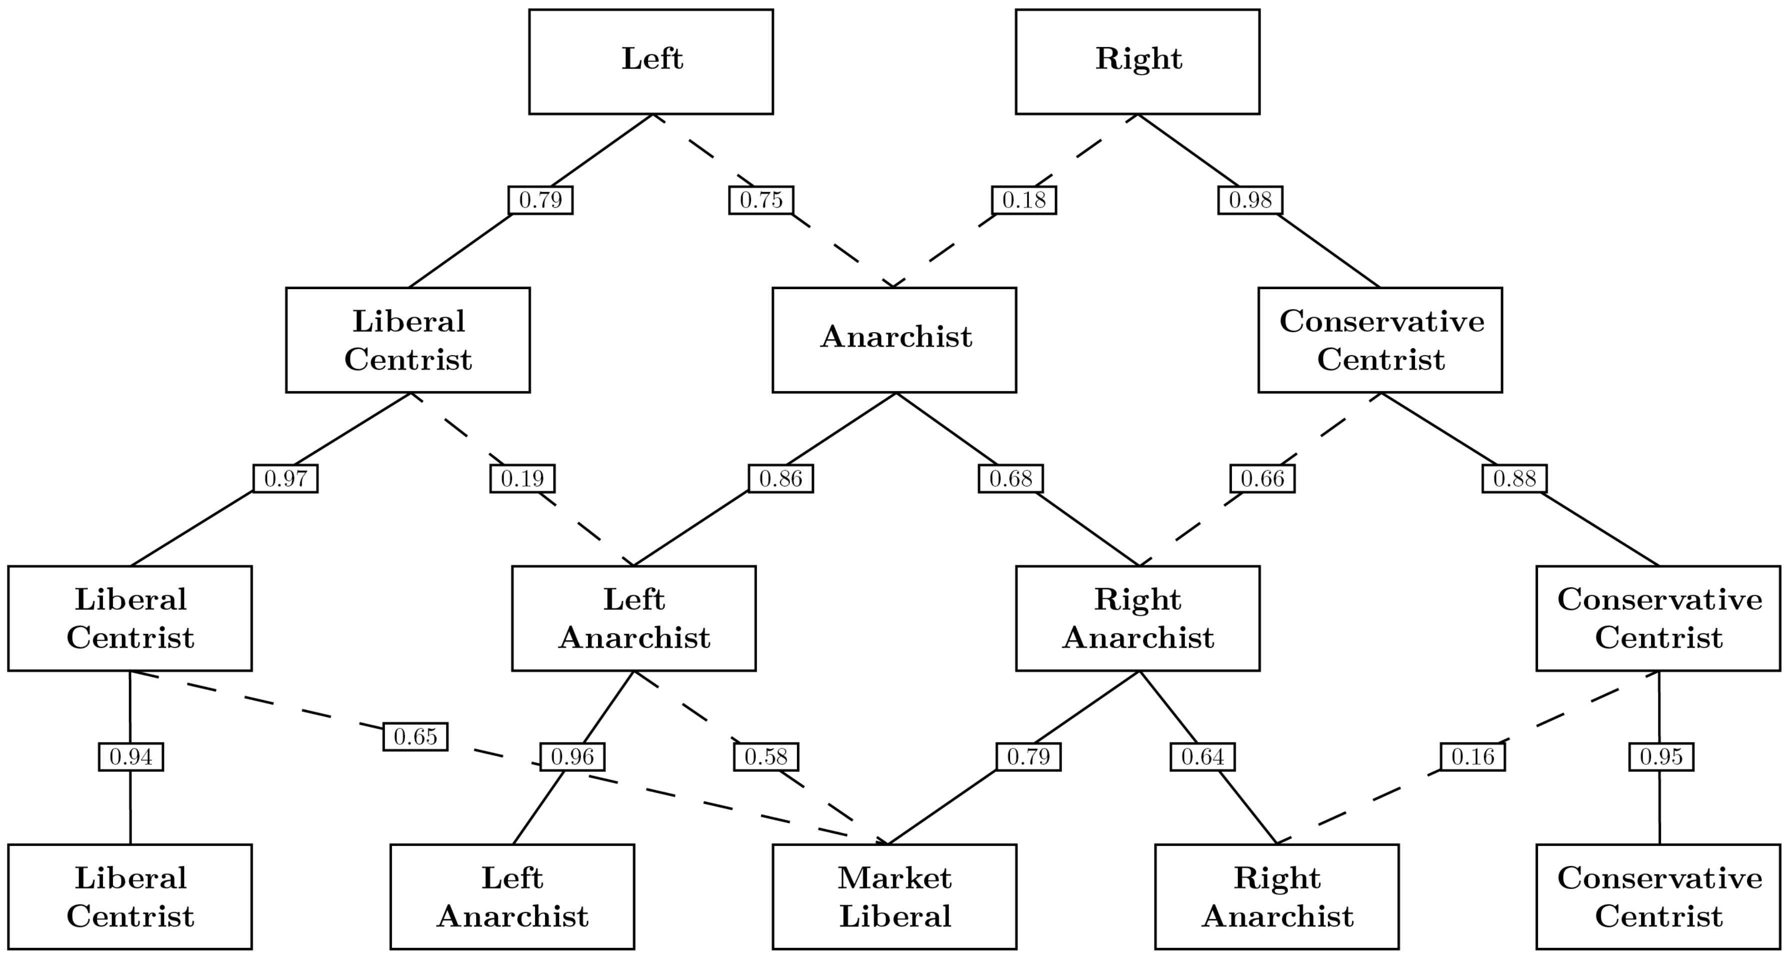

Notes: This graphic shows the hierarchy of types as created by Latent Dirichlet Allocation (LDA) for the different orders of topic model reported in Table 1. The values reported amongst the lines connecting the boxes record the correlation in the β issue-position probability vectors across types as a measure of type similarity.

The top features for a further 5-type model are reported in Appendix Table F.1. Additional nuances in the types become evident here. The set of Liberal Centrist, Conservative Centrist and Left Anarchist types remain intact relative to the 4-type model, but there is a split of the Right Anarchist type. Two variants of the Right Anarchist emerge. One variant still expresses a lack of confidence in institutions but appears to be liberal on social issues and is economically liberal in terms of attitudes towards unions and the claiming of government benefits. We label this type as 'Market Liberal'[19]. The other variant of the Right Anarchist is not socially liberal, with a string of conservative positions on minorities and social issues amongst its top ten features. The correlations indicate that this type is strongly correlated (0.64) with the original Right Anarchist from the 4-type model but negatively correlated with the Liberal Centrist (-0.195) and Left Anarchist types (-0.295) types.

Further results on potential 6-type and 7-type models are reported in Appendix Table F.1. The basic set of types is preserved such that we can directly label the types in line with those identified in the 4-type and 5-type models. The evolution of the hierarchy is apparent in the further splitting of the Right Anarchist type (6-type model) and the emergence of a nihilistic 'Super Anarchist' type (7-type model).

Overall, the results presented above indicate that both the low-order (2 or 3 types) and higher-order (4 plus types) models offer plausible sets of types and, considered together, can be interpreted in terms of a coherent hierarchy. We next turn to the question of formal model selection using the NPMI framework outlined previously.

4.2 Model Selection and Type Labelling

4.2.1 Automatic Model Selection

Our NPMI framework for assessing model cohesion is based on comparing predictions of feature co-occurrence in hold-out data. Simply put, the approach asks: to what extent do the (say) top 10 features suggested by our type models occur together in data held out from the original estimation of the given model? Figure 2 reports the results of this exercise for all waves of the WVS. The x-axis denotes the order of the model we are estimating (going from a 1-type model up to a 10-type model), while the y-axis denotes the resulting cohesion score.

The cohesion scores show an inverse U-shape pattern. At first, the cohesion score increases with the number of topics. After the number of topics increases beyond 4, the cohesion scores begin to fall. The inverse u-shape pattern arises since an increasing number of topics at first allows the model to better describe the data, but at some point, the model will over-fit the idiosyncrasies of the training sample, which will make the resulting types less cohesive.

[19]Our nickname for this anarchistic, pro-market, and socially liberal type is 'George Mason University (GMU) Blogger'.

Figure 2: Average Cohesion of Ideological Types for Different LDA Models

(a) Wave 2

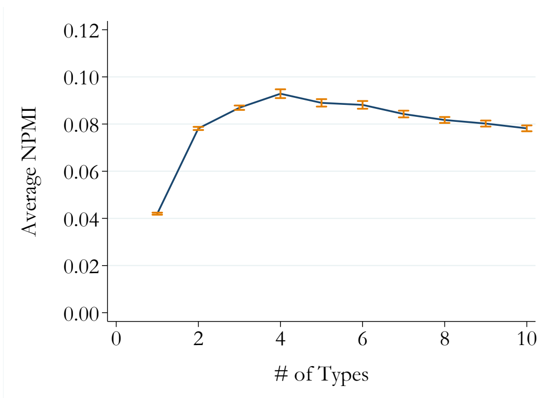

(b) Wave 4

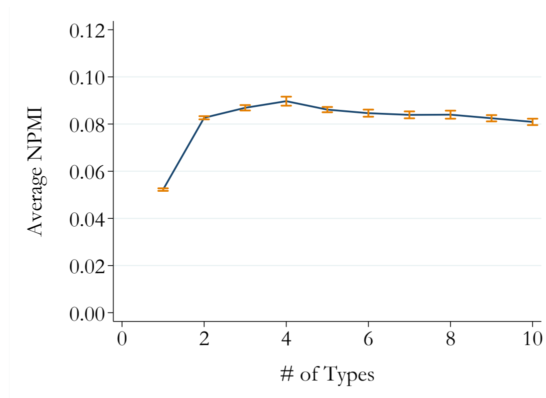

(c) Wave 5

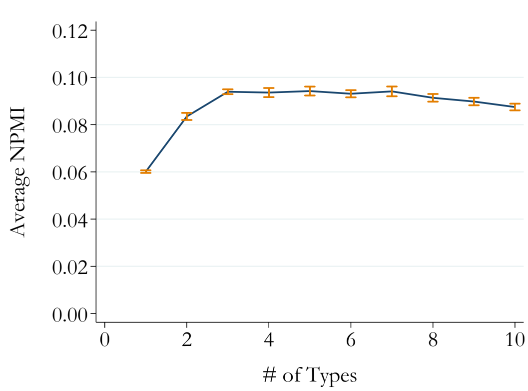

Notes: This figure show the topic cohesion scores calculated for models with $M \in \{1-10\}$ types for the 2nd/4th/5th wave. The topic cohesion is calculated based on different numbers of features $B \in (5, 10, 15, 20)$. Afterwards, the average of the cohesion scores for different values of B is taken. See Section 4.2.1 for a more detailed description. The confidence intervals are based on 100 iterations of the model.

Overall, the most cohesive models (M) appear to be either the 4 or 3-type models. We decided to use the 4-type model since it delivers a higher cohesion score in two of the three waves. Given this evidence, our analysis from this point therefore employs the 4-type model comprised of the Liberal Centrist, Conservative Centrist, Left Anarchist, and Right Anarchist types.

That said, note that our results do not hinge on this aspect of model specification and are qualitatively very similar if we use models of higher or lower numbers of types. The reason for this is that the types develop as part of a coherent hierarchy (see Figure 1). Moving to a 3 or 5-type model hence does not fundamentally change (for example) the prevalence of the Anarchist types in the data. We will return to this point about model specification again as we discuss various results (e.g., on 'slant', polarisation, and links with populist voting).

4.2.2 Type Labelling

The labelling of our LDA-derived types is a point for discussion. An important advantage of our approach is that it is based on ideologies that emerge from the 'bottom-up' collection of views amongst the general public. The topics that we identify are empirical ideologies and may not necessarily have a tight mapping to traditional taxonomies of ideology[20].

Our labelling attempts to capture the main empirical differences in issue positions between types. Note that we primarily use labels to simplify the exposition since we otherwise would have to refer to the types by numbers. Furthermore, the type labelling does not drive any of our results. Arguably, the main issue here is the labelling of types 3 and 4, which we have dubbed Left and Right Anarchist, respectively. These two types are strongly distinguished by issue positions that hinge on (low) confidence in institutions[21]. We use the term 'anarchist' to denote a pattern of opposition to current structures of political authority and hierarchical organization. That is, our use of the term is meant to be distinct from historical uses of the label, as per early socialist or syndicalist thinkers such as *inter alia*, Proudhon, Bakunin, or Kropotkin.

Other plausible labels for these types are 'Populist', 'anti-Establishment', or 'anti-system' (see Hopkin, 2020). In particular, the fact that recent studies of populism (such as Algan et al. (2017)) have directly leveraged data on institutional trust provides some foundation for such branding. However, we choose Anarchist as our label for this type because (i) it is a more generic and potentially neutral term

[20]These taxonomies, covered in texts such as Vincent (2009) and Geoghegan (2003), are centred on 'classical' ideologies (e.g., liberalism, conservatism, socialism) that are often explicitly articulated as bodies of thought by key writers (Locke, Burke, Stuart Mill, Marx), as well as modern 'post-materialist' ideologies (e.g., environmentalism, feminism).

[21]See Appendix Table F.5 for a breakdown of the largest differences in β issue-position weights across pairs of types. This shows the points of separation between the Centrist and Anarchist types.

for the concept of opposition to existing institutional structures[22], and (ii) the types that we identify are apparent from the early 1990s, thereby pre-dating the latest wave of populist politics. As a result, our empirical findings indicate that there may be some clear historical roots to the current populist trend, extending at least as far back as the late 1980s.

4.2.3 Alternative Models

In the appendix, we provide extensive robustness checks for the stability of our LDA findings. Firstly, in Appendix E, we test whether the resulting types are similar when we use different Dirichlet priors (α and γ) and random seeds. The results presented in Figure E.1 indicate that these choices do not make any qualitative difference for the resulting ideological types.

Secondly, in Appendix Table E.2 we look at the sensitivity of our baseline 4-type model to the removal of features (question). The baseline model is very robust with types from iterative 'leave one out' models showing a high correlation with the types in the original model. To further rule out that our Anarchist types are the artifact of the number of trust questions in the WVS, we run further robustness checks in which we consecutively remove more and more of the trust questions (see Figure E.3). As it turns out, the basic structures of ideologies are preserved even when we remove several trust questions at a time.

Lastly, we investigate if our ideological types are robust to the addition of features. We again find that the relative ordering of β weights is preserved when we substantially widen the feature set (i.e., add lots more questions - see Appendix Table E.1).[23] Both of these exercises provide reassurance that our overall baseline feature set is comprehensive enough to identify stable types in the data.

A second modelling issue that we examine (in Appendix L) is a comparison of our LDA approach with other unsupervised learning methods. Specifically, we apply principal components analysis (PCA), factor analysis (FA) and k-means clustering to the same discretised feature data as our LDA models. As we discuss in Appendix L, these alternative approaches are distinct from LDA in that they are linear methods and capture mixed membership relationships in a less explicit way. For example, a method such as PCA will pick out linear combinations of features with the highest degrees of variance in the data and therefore may not parse more complex data-generating processes.

[22]As mentioned, the term 'populist' can be considered pejorative - see the blunt critique of UK's *Daily Mail* (Murray (2016)). The term 'anti-establishment' is subject to similar concerns, with competing claims of who the elite or establishment are.

[23]Among the questions in the widened feature set are many that no do not directly relate to political ideologies and which were therefore excluded from our baseline model. Further, some of the added questions are missing for close to 50% of the data, and nearly a third of all questions are missing for more than 10% of the sample. Hence the model with the widened feature set requires extensive imputation and does not lend itself to be used as a baseline model.

An additional advantage of LDA relative to PCA and FA comes from the fact that LDA identifies co-occurrence instead of linear combinations. Co-occurrences arguably better describe political ideologies as there can be many issues even opposing ideologies agree upon. For example, even though there are considerable differences between the positions of Republicans and Democrats in the United States, most supporters of both parties would like to profess pride in their nationality.

This is borne out in the types derived from these models which are reported in Appendix Tables L.1, L.2 and L.3. The PCA models tend to identify conservative and anti-establishment types as part of the main model components, with no clear centrist or socially liberal types emerging. The FA and k-means models produce similar results. Further to this, no plausible hierarchy or 'family' of types emerges from these alternative models. Again, this provides reassurance that our LDA models - which are, after all, specifically intended for the analysis of discrete multinomial data - identify stable and interpretable types that are difficult to pin down using other methods.

4.2.4 Cross-Check Exercise Using the European Social Survey

As further validation of our findings, we cross-check our results using the European Social Survey (ESS). This exercise serves two purposes. First, we want to demonstrate that our LDA methodology extends to other survey data. Second, we aim to understand if it is possible to recover similar ideological types using a different data set and a comparable set of questions. If our main LDA exercise is picking up valid latent types from the WVS data then this 'signal' should be apparent in other datasets.

Appendix G provides additional details on the ESS data and reports the full list of questions we selected for the exercise. Overall, the results from this exercise (see Appendix G for details) are striking. Although we use a completely different data set and varied the set of questions, the types that emerge from the ESS are broadly similar to those we identify in the WVS. In particular, we again find a split of the ideological types along the left-right spectrum and see a set of types characterised by their distrust of institutions.

4.3 Changing Ideologies?

Given the baseline 4-type model established above, our next exercise examines within and between-type shifts across the different waves of the WVS. Our approach here is to estimate the 4-type model separately for each wave and compare the β type vectors over time.

The first point to note is that our main types are stable and repeat themselves across waves. This is evident in Table 2a, which reports the correlations between the separately estimated types across waves. It is straightforward to pin down similar types across waves since the correlations are high, for example, the Liberal Centrist type showing a correlation of 0.96 between waves 2 and 4 or the Right Anarchist type reporting a correlation of 0.94 between waves 2 and 5.

Table 2: Type Correlations

(a) Between-Wave Type Correlations

	Centrist Liberal Wave 2	Centrist Conservative Wave 2	Left Anarchist Wave 2	Right Anarchist Wave 2
Wave 4	0.963	0.986	0.961	0.980
Wave 5	0.929	0.978	0.961	0.943

Notes: This table reports the correlation of the β issue-position probability weights across types estimated in separate waves. That is, we identify 4 types in the initial Wave 2 (1989-1993) and correlate their β weights with the most similar types estimated separately on Waves 4 (1999-2004) and 5 (2004-2009).

(b) Within-Wave Type Correlations

	Liberal Centrist	Conservative Centrist	Left Anarchist	Right Anarchist
Wave 2				
Liberal Centrist	1.000			
Conservative Centrist	0.418	1.000		
Left Anarchist	0.225	-0.526	1.000	
Right Anarchist	0.191	0.267	0.130	1.000
Wave 4				
Liberal Centrist	1.000			
Conservative Centrist	0.468	1.000		
Left Anarchist	0.322	-0.504	1.000	
Right Anarchist	0.251	0.289	0.178	1.000
Wave 5				
Liberal Centrist	1.000			
Conservative Centrist	0.523	1.000		
Left Anarchist	0.257	-0.408	1.000	
Right Anarchist	0.224	0.287	0.265	1.000

Notes: This table shows the correlation of the β issue-position probability weights amongst types in the same wave. That is, we estimate our 4 types using data on a single wave and then correlate the β weights across pairs of types in the same wave.

These high correlations also imply that there are fairly limited 'within-type' shifts over time, as measured by the probability weights in the β type vectors. Since we are using the same issue-position features across waves we can directly report the shifts in probabilities per feature. To facilitate the interpretation we have re-scaled the β vectors as described in Appendix C[24].

[24] Since the rescaling of the β vectors is a non-linear transformation, the changes between the re-scaled β vectors and unscaled β vectors can differ. See Appendix C for further technical details.

These probabilities can be interpreted as statistics for the approximate 'likelihood of expression' for a given issue-position conditional on drawing on a latent type. For example, a (re-scaled) β weight of 0.46 for 'Confidence: Labor Unions' within the Liberal Centrist type indicates that an individual drawing on this type to form their issue position will express confidence in this institution 46% of the time.

The ten largest shifts in probabilities for the Wave 2-5 difference are shown in Figure 3 for each type. The baseline numbers are also reported in Appendix Table F.4 along with the changes. Note that given the coding of each question into two features, the issues of decreasing importance will be approximately the opposite of the increasing features. Again, this relates to our adaptation of the LDA model for studying survey questions, which we outline in Appendix B and Appendix C. The most noticeable trend is an increase in socially liberal attitudes across types with, for example, the Conservative Centrist increasing their weights on issue positions such as 'No Problem Neighbours: Homosexual' and 'No Problem Neighbours: People with AIDS'. Also notable is the Right Anarchist type, which shows higher confidence in the police and armed forces over time, along with more intense hostility towards immigration. Some of these changes are nominally large, with 10-15% increases in liberal attitudes on gay rights for the Conservative Centrist and 20-30% increases in confidence for the armed forces and police for the Right Anarchist. However, the overall changes in the β weights have not been pervasive enough to drastically shift the between-wave correlations evident in Table 2a[25].

The between-type differences can also be summarised using correlations across the β type vectors within the WVS waves and we show these in Table 2b. The increase in the intensity of socially liberal issue positions is now most clearly seen via the increasing closeness between the Conservative Centrist type and the two left-wing types. Between waves 2 and 5 the negative correlation with the Left Anarchist type moderates (going from -0.526 to -0.408) while the correlation with the Liberal Centrist type strengthens (from 0.418 to 0.523). Hence, at the between-type level defined by the β-vectors, we can say that there has been some convergence in the overall ideologies driven in part by attitudes on social issues. Despite this convergence, note that the types remain clearly distinct and opposite to each other on many issues. As an illustration, in Appendix Table F.5 we report the most important differences between the 4 types for the 5th wave.

[25]In the case of Right Anarchist attitudes towards the police and armed forces it should be noted that this shift brings this type closer to the mean β for these issue positions displayed by the two Centrist types.

Figure 3: Within-Type Changes in Issue-Position Weights (Wave 2 to 5)

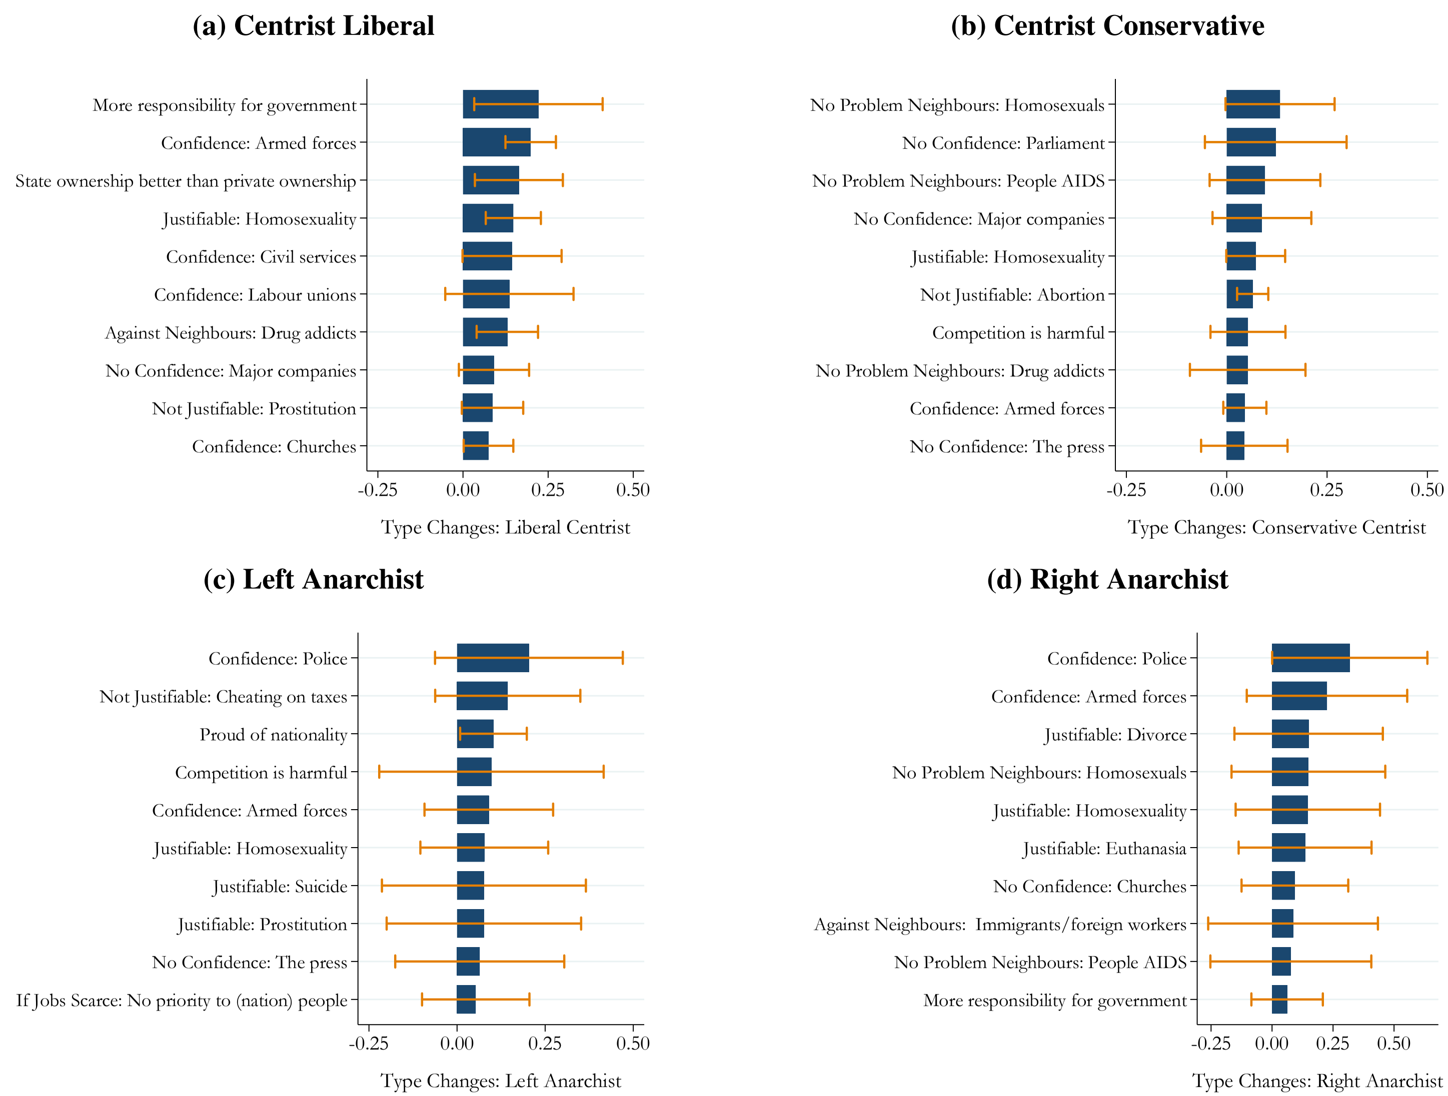

Notes: This figure reports the largest changes in the β issue position weights per ideological type from Waves 2 (1989-1993) to Wave 5 (2005-2009). We report the top ten changes per type amongst the 58 features. The scale is set to 0-0.4 to facilitate direct comparisons across types. Confidence intervals are based on 100 bootstrap iterations.

4.4 Analysis of Type Shares

4.4.1 Correlates of Type Shares and Country Differences

We start the analysis of the θ_i individual type shares by studying the micro-level correlates. In particular, we estimate regressions of the individual type shares of the following form:

$$y^t_{icw} = X'_{icw}\delta + \tau_w + \mu_c + \epsilon_{icw} \tag{9}$$

where X_{icw} is a vector of covariates including the gender, age and the employment status.[26] τ_w and μ_c are wave and country fixed effects. The dependent variable y^t_{icw} is the share of type t of individual i in country c and wave w. Since the dependent variable is a continuous share the regression tells us how the intensity of ideological positions changes with different covariates. The results are reported in Table 3.

Table 3: Correlates of Individual-level Type Shares

	(1) Liberal Centrist	(2) Conservative Centrist	(3) Left Anarchist	(4) Right Anarchist
Female	0.015***	0.010***	-0.011***	-0.013***
	(0.002)	(0.002)	(0.002)	(0.002)
Age	-0.003***	0.004***	-0.003***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Unemployed	-0.073***	-0.003	0.049***	0.027***
	(0.005)	(0.005)	(0.004)	(0.005)
Wave 4	0.086***	-0.055***	0.007***	-0.039***
	(0.003)	(0.003)	(0.002)	(0.003)
Wave 5	0.069***	-0.076***	0.032***	-0.024***
	(0.003)	(0.003)	(0.002)	(0.003)
Country FE	Yes	Yes	Yes	Yes
Observations	81,141	81,141	81,141	81,141
Mean of DV	0.298	0.269	0.170	0.263
R^2	0.14	0.11	0.11	0.06

Notes: Each column reports the regression results for individual level regression of Equation (9). The dependent variable are the type shares for one of the 4 types created by LDA. Robust standard errors are used. Significance levels: *** p<0.01, ** p<0.05, and * p<0.1. The data come from the World Value Survey and the European Value Survey.

We run four regressions corresponding to each type. These indicate some intuitively plausible correlations - women are more liberal and centrist, with a magnitude of 1.5% points, and the unemployed have higher shares in the two anarchist ideologies. Furthermore, there are clear shifts in the distribution of type shares over time. After conditioning on covariates, it is evident that the Liberal Centrist share

[26] We only use a limited number of covariates in this exercise because these are the most complete ones available in the WVS in terms of missing values. When we run similar exercises with additional variables on reduced samples (circa N = 50,000) we get similar results (e.g., low incomes are positively correlated with Anarchist shares, high education is positive with Liberal Centrism). These results are available on request.

increases by around 6.9% points after Wave 2 with the Conservative Centrist share falling by a similar amount. Following a similar pattern, the Left Anarchist share rises in Waves 4 and 5 while the Right Anarchist share falls.

Hence, across the sample of countries, the net result is a growth in the share of the two left-wing types (i.e., Liberal Centrist and Left Anarchist). However, there are also significant country-level factors evident from the individual-level analysis. The country fixed effects in Table 3 account for 50-75% of the explained variation and we plot the country-level means by type in Figure 4. The country-level type shares are constructed by calculating the average of the individual-level type shares within each country.[27] This again shows some expected relationships - northern European countries (e.g., Denmark, Finland, Netherlands) are more liberal while countries with strong religious traditions (Malta, Ireland, Portugal) are more conservative.

To summarise the changes across countries over time we implement some splits along different ideological types. Firstly, in Figure 5a, we examine the left-right distinction and pool the type shares for the left-wing Liberal Centrist and Left Anarchist types.[28] The plot of changes in these pooled type shares between Waves 2 (1989-1993) and 5 (2005-2009) shows that most countries have moved left ideologically, with a mean shift of 8% points. In Figure 5b we then plot the changes for the pooled Left and Right Anarchist types. This provides an indicator of the overall strength of anti-establishment ideological sentiment across countries. The results show a large increase in the Anarchist type shares for the US (around 16% points), with significant increases also evident for the Anglo-Celtic domains (Great Britain, Northern Ireland) and the Netherlands. In net terms, however, the anarchist trend is more muted across countries.

In Figure 6, we further probe the sharp increase in the Anarchist ideologies for the US. The clearest development is the growth in the US Right Anarchist share, which increases from a 25% share in Wave 2 (1989-1993) to 37% in Wave 5 (2005-2009). Note here however that this increase took place as a single-step change between Waves 4 (1999-2004) and 5 (2005-2009). By comparison, the rise of the US Left Anarchist share from 8% to 13% was more gradual across the waves.

As discussed in Section 2, the set of countries available in wave 7 (post-2017) of the WVS is limited. However, in Appendix H we provide evidence that for the available countries the patterns in the 7th wave are consistent with the trends we document for wave 5. For example, the Anarchist type

[27] As such a country in which each person exhibits a 25% share in each type will have the same type shares as a country with 4 equal-sized population groups each of which has a 100% in one type. We analyse these differences in Section 4.4.4 in which we construct a measure of within-person concentration.

[28] The type differences are based on the sum of the average type shares for the Liberal Centrist and Left Anarchist types for each country and wave. The figure then plots the difference between a country's average type share in wave 5 and wave 2.

Figure 4: Country-Level Type Shares

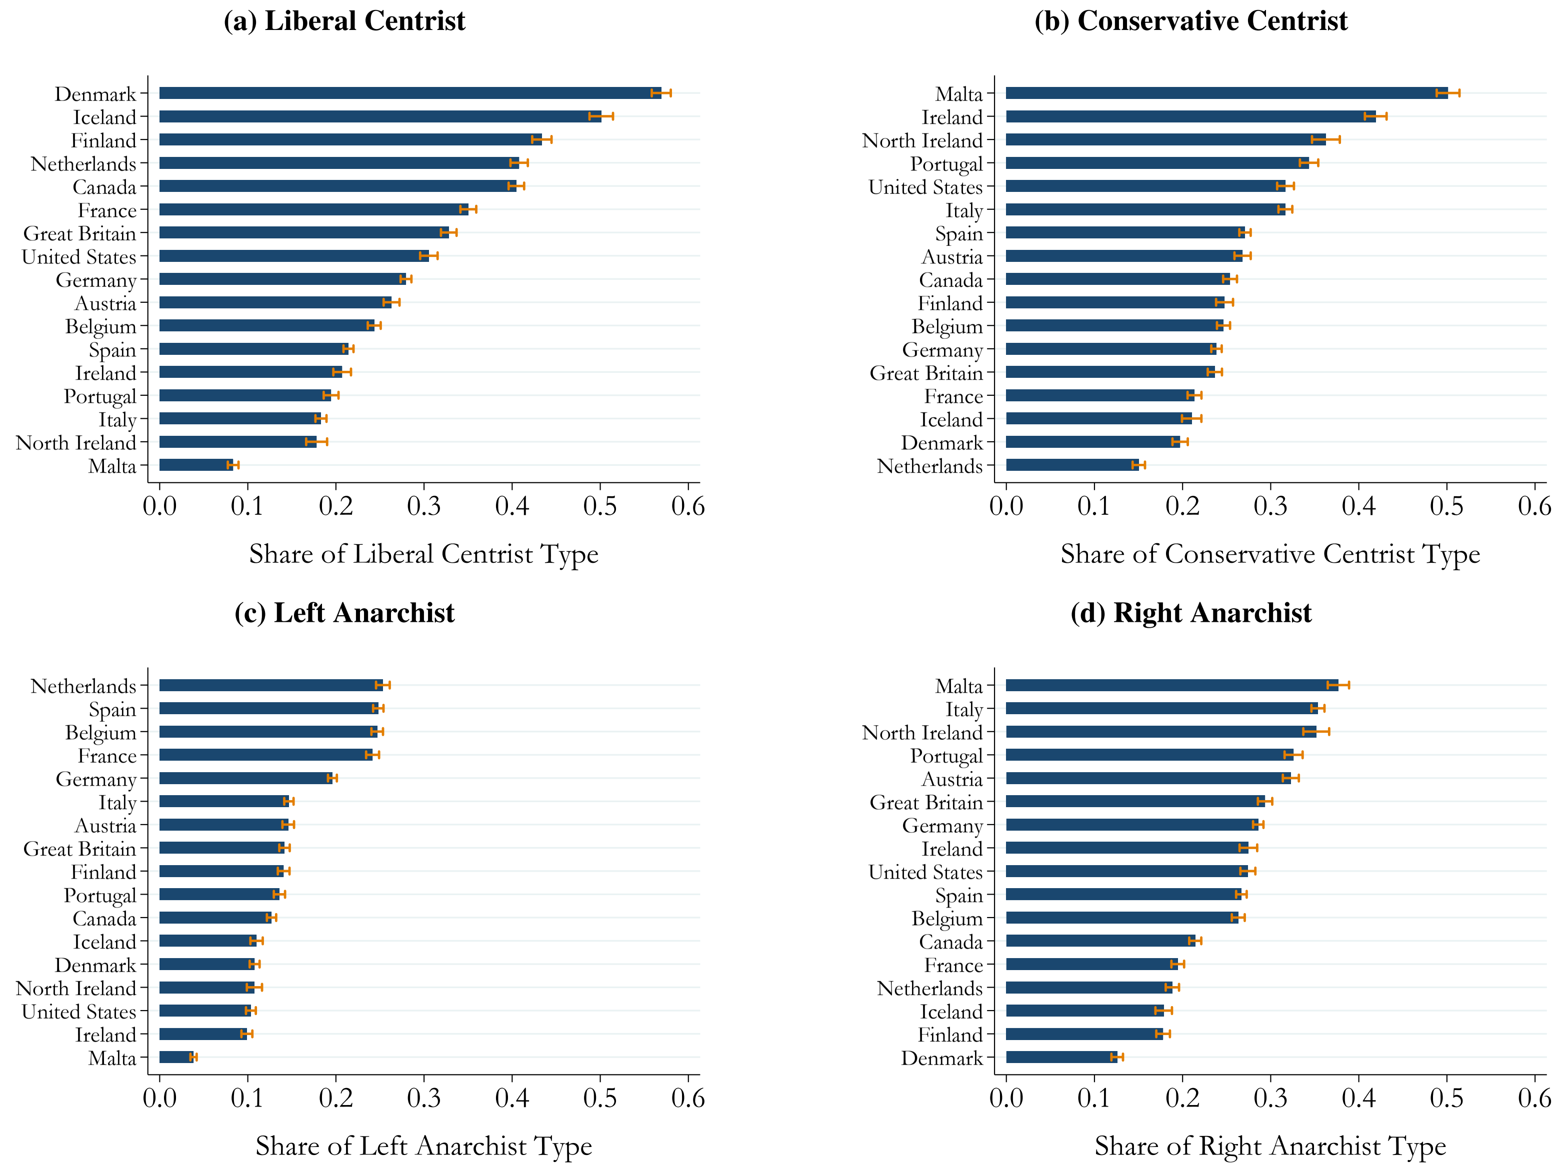

Notes: This figure shows the average country-level θ type shares aggregated over individuals. 95% confidence intervals are reported in orange. Country means and confidence intervals are calculated using WVS sample weights. 95% confidence intervals are reported in orange.

Figure 5: Changes of Types over Time

(a) Change in Left-Wing Share **(b) Change in Anarchist Share**

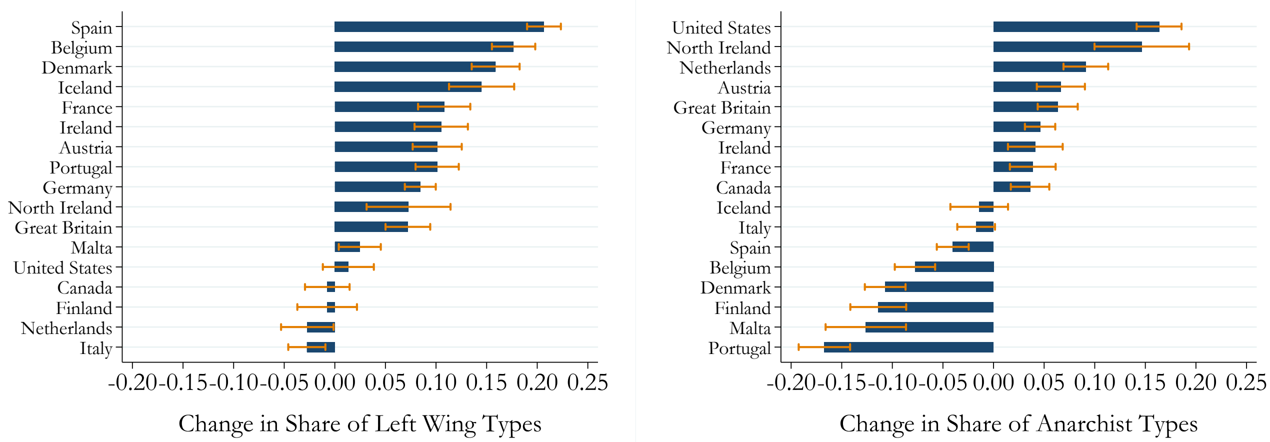

Notes: This figure shows the change in θ type shares by country between Waves 2 (1989-1993) and 5 (2005-2009) in the WVS. In 4(a) we pool the type shares for the Liberal Centrist and Left Anarchist types. In figure 4(b) we show the pooled change in the Left Anarchist and Right Anarchist types. 95% confidence intervals are reported in orange.

Figure 6: Type Shares - US vs non-US

(a) Left Anarchist Types **(b) Right Anarchist Types**

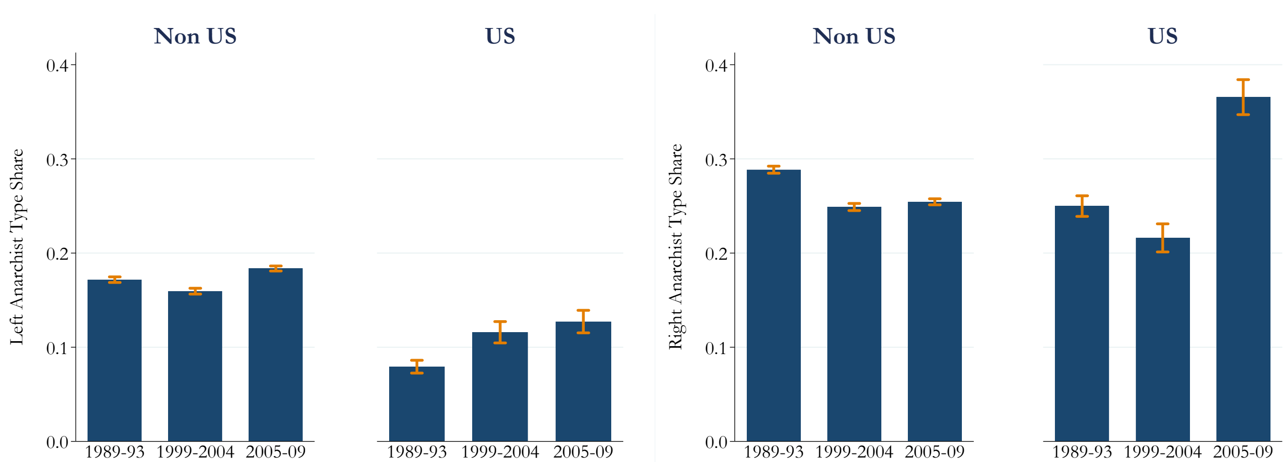

Notes: This figure compares the levels of θ type shares across waves for the Left Anarchist and Right Anarchist types. We pool all 16 non-US countries (effectively all Western European countries apart from Iceland and Canada) and contrast them to the US. The pooling for the non-US sample is based on WVS sample weights. The timing of the waves is Wave 2 (1989-1993), Wave 4 (1999-2004) and Wave 5 (2005-2009). 95% confidence intervals are reported in orange.

shares in the US increase further.

Overall, these country-level findings are generally consistent with other international studies of shifts in political attitudes (Inglehart (1997); Inglehart et al. (2010)). Taken together with the within-type analysis, the basic message on the ideological change that follows from our methodology so far is one of a stable structure of ideologies over 20 years and some increase in social liberalism. This increase in social liberalism has occurred both on the intensive margin (i.e., the weights on liberal issue positions in the β vectors) as well as the extensive margin (the growing individual-level type shares for the Liberal Centrist and Left Anarchist types). The other major development in the data on type shares is the strong tilt towards anti-establishment Anarchist ideologies in some countries - particularly the US.

4.4.2 Ideological Types and the Left-Right Scale.

We next analyse the relationship between our ideological types and the self-positioning of individuals on a left-right scale. For this analysis, we make use of question E033 in the WVS, which asks people to position themselves on a scale of 1 (left) to 10 (right). Recall that this measure of self-positioning is not one of the ingredients in the feature set for our LDA analysis. It is held out from the estimation of the ideological clusters and therefore provides a useful test of validity.

The mean left-right scores according to the dominant type are telling. Individuals with a dominant Left Anarchist type position themselves furthest to the left (mean: 4.33), followed by the Liberal Centrist (5.24), Right Anarchist (5.49) and the Conservative Centrist (mean: 5.74).[29] In line with our previous findings, average political attitudes are moving leftwards with a shift of -0.17 on the left-right scale between waves 2 and 5.

In Figure 7, we visualise the relationship between the right-wing (Conservative Centrist and Right Anarchist) type shares and the left-right scale. We find a strong relationship between the type shares and the political orientation of individuals. As expected, the larger the share of the right-wing types in an individual, the further right they place themselves on the political spectrum. The inverse mechanically holds for the left-wing type shares (not shown). This provides further validation for our type labels as they appear to align with the classic left-right ideological spectrum.

4.4.3 Ideological Types and Populism.

Lastly, we investigate whether our anarchist types are associated with stronger support for populist parties. To do so we consider the question "which political party would you vote for" (question code

[29] Note that, while the mean difference in positioning appears nominally small in these comparisons this is because answers are clustered on middle values: more than 55% of the people position themselves between 4 and 6 on the Left-Right line.

E179/E179WVS). We recode the responses of individuals as populist parties based on the classification by Rooduijn et al. (2019) (See Appendix I for additional details). In Figure 7 panel (b) we then plot the support for populist parties conditional on the anarchist type shares of the individual. Again a clear positive relationship between voting for populist parties and the anarchist ideologies emerges. This suggests that the anarchist types might indeed be a base for populist political mobilisation.

As further evidence, we estimate a simple linear probability model (LPM) of voting for populist parties in Appendix Table I.2, contrasting the explanatory power of our ideological type share variables with that of the left-right self-positioning question. Appendix Table I.2 shows a strong positive relationship between the anarchist type shares and populist voting across all specifications. Interestingly, the left-right self-positioning question only has a significant association when we specify the variable either as a set of dummies for far left or far right positions (column 4) or as a step function for each value (Figure 7c).

In particular, this Figure 7c step function shows a U-shape in the probability of populist voting with respect to the left-right scale. That is, people located near the centre of the scale are the least likely to vote populist. One possibility here is that the anarchist type shares are proxying for extreme left or right positioning. However, as noted above, in column (4) of Appendix Table I.2 we control for indicator variables of extreme positioning and this has minimal effects on the previous association. The anarchist type share variables appear to pick up tendencies for populist support from across the left-right spectrum. This association is non-negligible: based on the estimates in Appendix Table I.2, we calculate that an individual with a 50% type share in one of the anarchist type has a 38% higher probability of voting for a populist party relative to the sample baseline.[30]

In Appendix Table I.3, we report a more extensive analysis of the associations between the ideological types and voting behaviour. For this analysis, we have obtained information on party families from the Manifesto Project (Lehmann et al., 2022) and the Chapel Hill Expert Survey (Jolly et al., 2022). We then hand-coded the parties in the WVS into the respective party families and report the correlations of the party families with the ideological types (relative to the liberal centrists).[31] Overall, the correlations between party families and ideological types are highly intuitive, independent which source for party families we are using.

[30]As an additional exercise, we tested whether these findings also held across the entire left-right spectrum by running regressions that split for each value of the left-right scale. For nearly all values of the scale, a higher anarchist type share is associated with more support for populist parties. These results are available by request, though note that the level shift in the U-shape plotted in Figure 7c directly corroborates this point.

[31]Note that there is no 1:1 match between the parties in the Manifesto Project and the World Value Survey. We are hence not able to code all responses in the WVS to a party family.

Figure 7: Self-positioning on Left-Right Scale and Support for Populist Parties

(a) Self-positioning on Left-Right Scale

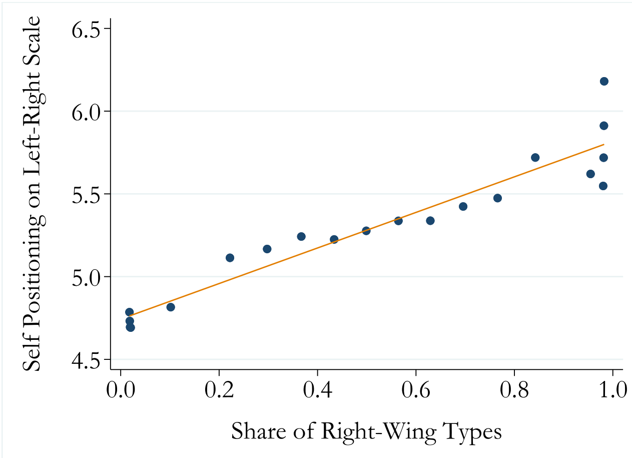

(b) Support for Populist Parties

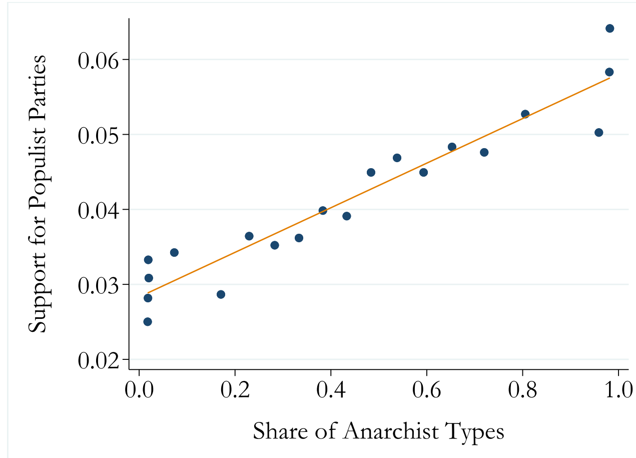

(c) Left-Right Scale and Support for Populist Parties

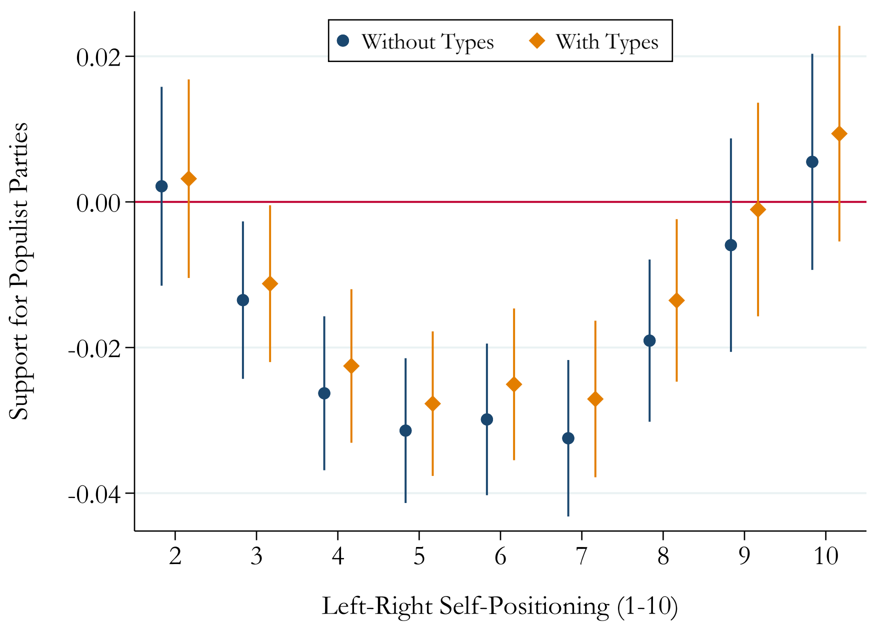

Notes: The binscatter in panel (a) visualises the relationship between the individual-level type share of the right-wing types (Conservative Centrist and Right Anarchist) and the self-positioning of individuals on a 1 (left) to 10 (right) scale based on question E033 from the World Value Survey. Panel (b) shows the relationship between the individual-level share of anarchist types and the voting for populist parties coded according to Rooduijn et al. (2019). See Table I.1 in the online appendix for the full list of parties. Panel (c) plots the coefficients and 95% confidence intervals for individual-level regression, where the dependent variable is an indicator variable for the support of a populist party. The independent variable is a full set of indicator variables for an individual's positioning along the left-right spectrum, the excluded category being 1 (far left). The reported coefficients in orange additionally controls for the individual level type shares θ_i.

4.4.4 'Citizen Slant' - Within-Person Concentration

Our analysis so far has focused on changes at the level of the ideological types as well as the total shares in the types across the sample. However, for the analysis of the polarisation, the loadings of individuals on the four types is of key importance. In particular, the 'mixed membership' structure of our approach means that two countries with the same overall type distribution can have completely different individual type compositions.

For example, imagine a country which has an overall 50% share in Type 1 and 50% share in a second Type 2. This country can either consist of completely identical individuals with 50% shares in the two types or it could consist of half the population holding a 100% share in Type 1 and another half with a 100% share in Type 2. These two possible type compositions have very different implications for the understanding of societal polarisation. A country with two separate sets of 'pure' homogeneous types is plausibly more vulnerable to political conflict than a country where there is more ideological heterogeneity at the individual level.

We therefore develop a measure of how strongly an individual is loading on one of the four ideological types by constructing a Gini-style measure of within-person concentration or 'slant'. We define the within-person concentration G_i of individual i as:

$$G_i = \frac{\sum_{t=1}^{4} \sum_{s \neq t} |\theta_i^t - \theta_i^s|}{2(n-1) \sum_{t=1}^{4} \theta_i^t} \tag{10}$$

where θ_i^t and θ_i^s are the types shares of individual i. In short, this approach is aggregating the absolute pairwise difference in ideological shares that exist at the individual level. The measure of within-person concentration G_i is monotonically increasing the more an individual loads on one of the ideological types. If a person has a 100% share in one type then G_i will be equal to 1, while $G_i = 0$ implies shares of 25% in all types.

In this way, our measure captures how segregated type shares are on a within-individual basis. Furthermore, it allows us to analyse which groups exhibit a particularly high tendency towards within-person concentration. We plot the (Wave 5) country means of the G_i measure in Figure 8 along with the changes between Wave 2 and 5. This shows that G_i is relatively high across the sample with a mean of around 0.75. However, between-country differences are not dramatic. There is only a 7% gap between the most and least concentrated countries and the ordering does not suggest that any particular ideological types are driving concentration. That is, amongst the most concentrated or 'slanted' countries we see cases of both relatively Conservative and Liberal countries defined in terms of the mean type shares seen earlier. The major, positive country-level shifts in slant over the waves occurred in Denmark, Finland and the US (Figure 8a) but the changes were muted for most countries.

Figure 8: Citizen Slant by Country

(a) Change in Slant by Country **(b) Ranking of Countries by Slant (Wave 5)**

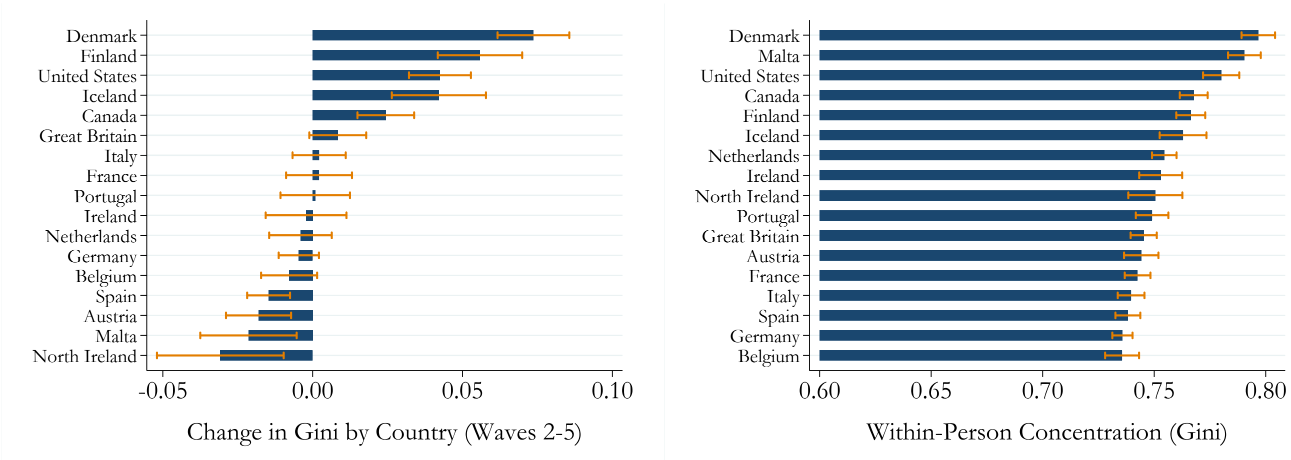

Notes: Panel (a) shows the change in our G_i Gini within-person ideological concentration measure ('citizen slant') from Waves 2 (1989-1993) to 5 (2005-2009) by country. Panel (b) shows the level of the within-person Gini measure by country in Wave 5. 95% confidence intervals are reported in orange. Country means and confidence intervals are calculated using WVS sample weights.

To study the importance of individual characteristics on within-person concentration (as well as the development of G_i over time) we estimate the following regression equation:

$$G_{icw} = X'_{icw}\delta + \tau_w + \mu_c + \epsilon_{icw} \tag{11}$$

where X_{icw} is a vector of covariates[32], τ_w are wave dummies, μ_c are the country dummies and G_{icw} is the of Gini coefficient of individual i in country c and wave w. The results are reported in Table 4 with controls for the type shares and with the Liberal Centrist set as the baseline type. The purpose of the type share controls is to allow us to study whether G_i concentration is increasing with shares of certain types. The results indicate that the Left Anarchist is the least concentrated type followed by the Right Anarchist. In turn, this means that the individuals with larger shares in either of the two anarchist ideologies are more likely to mix different ideological types than the centrist types.

After controlling for the available individual characteristics we find a 1.6% increase in G in wave 4 and an 0.6% increase in wave 5[33]. The results for the analysis of the US are similar overall except that the increases of G concentration in Waves 4 and 5 are far larger, standing at 2.8% and 5.0% (Column 3). To further probe the increases in G over time we estimate Equation (11) separately for individuals

[32]We use the same individual covariates as in Table 3.

[33]We suppress the reporting of the individual attribute coefficients in Equation (11) to avoid clutter. The basic result for these covariates is that only gender (female) and unemployment contribute significantly to within-person concentration but with small magnitudes. They enter with positive and negative signs respectively.

conditional on their main type and also broken down according to the US and non-US samples. The results in Appendix Table J.1 show that the increases in G within the US are predominantly driven by the two Anarchist types, both of which exhibit increases in concentration around 13% relative to the overall sample mean (0.753).

To clarify, note that the message from the earlier table was that the Anarchist types are less concentrated in the cross-section (hence the positive coefficients on these type variables in the associated regressions). In contrast, the regressions in Table J.1 track how concentration developed over time on a per-type basis. The simple story then is that, when they do manifest, Anarchist views are becoming more concentrated or 'purer' at the individual level rather than being spread out amongst more people.

Table 4: Correlates of 'Citizen Slant' (Gini Concentration)

	All Countries		USA	
	(1) Gini	(2) Gini	(3) Gini	(4) Gini
Conservative Centrist	0.000		-0.013**	
	(0.002)		(0.006)	
Left Anarchist	-0.035***		-0.060***	
	(0.002)		(0.010)	
Right Anarchist	-0.024***		-0.045***	
	(0.002)		(0.007)	
Female	0.003**	0.003***	0.001	0.003
	(0.001)	(0.001)	(0.005)	(0.005)
Age	0.000***	0.000***	0.000	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)
Unemployed	-0.006**	-0.009***	-0.006	-0.010
	(0.002)	(0.002)	(0.012)	(0.012)
Wave 4	0.016***	0.016***	0.028***	0.028***
	(0.002)	(0.002)	(0.006)	(0.006)
Wave 5	0.006***	0.006***	0.050***	0.042***
	(0.001)	(0.001)	(0.006)	(0.006)
Country FE	Yes	Yes		
Observations	81,141	81,141	4,197	4,197
Mean of DV	0.753	0.753	0.759	0.759
R^2	0.02	0.01	0.04	0.02

Notes: Each column reports the regression results for individual-level regression. The dependent variable is the Gini Coefficient of the individual type shares as a measure of polarisation. Columns (1) and (2) use all data, and columns (3) and (4) restrict the sample to the USA. Robust standard errors are used. Significance levels: *** p<0.01, ** p<0.05, and * p<0.1. The data come from the World Value Survey and the European Value Survey.

In effect, this evidence implies that the Anarchist types have become an even more dominant ideology for people who had already shown a lack of trust in social and political institutions. While in earlier waves, this section of the population might still have shown large type shares in Centrist ideologies this potentially moderating centrist influence became less evident in more recent years. The findings for the US also contrast fairly strongly with the results for the non-US sample, where the increase in concentration for the Anarchist types is more muted and, in any case, is matched by increases for the Liberal Centrist type as well (see Appendix Table J.1, panel (B), first column).

4.5 Societal Polarisation

While the above measure of within-person concentration describes the strength of the individual loadings on the four ideological types, it does not fully summarise the extent of the divisions between citizens within a society. A society in which there are sub-groups of individuals that heavily load on the same ideological type may not necessarily be dramatically polarised. The extent of polarisation would hinge on how big these 'purist' sub-groups are relative to the full set of ideological sub-groups across the population. As an example, the country-level type share plots in Figure 4 indicate that some countries are characterised by widely represented types with aggregate type shares around the 50% mark, such as Liberal Centrists in Denmark and Conservative Centrists in Malta. At face value, these countries could be plausibly classified as 'unipolar' and less vulnerable to group conflict, no matter how concentrated the different types are in terms of citizen slant.

We therefore study polarisation by adapting the measures proposed by Esteban and Ray (1994) and Duclos et al. (2004) to our setting with 4 ideological types. These measures have the feature of accommodating two forces that define polarisation as a general concept. Firstly, there is *identification* which occurs amongst individuals with a common characteristic and is an increasing function of the total number of common individuals (that is, group size). Secondly, there is *alienation* which accounts for the social detachment that individuals feel towards others who do not share some common characteristic. Again, the strength of the alienation effect will depend on (relative) group size as well as the 'distance' between groups in the key characteristic of concern.

Using the example of income inequality, Esteban and Ray (1994) prove that any measure of polarisation P that accurately accounts for own-group identification as well as alienation in relation to an out-group and fulfils 3 'reasonable' axioms has to be of the form[34]:

$$P(\pi, y) = \kappa \sum_{i=1}^{n} \sum_{j=1}^{n} \pi_i^{1+\nu} \pi_j |y_i - y_j| \tag{12}$$

where π are the number of people in the groups, y is the amount of income of each group, κ is a normalizing constant, and ν is the polarisation sensitivity, which parameterises how the polarisation measure shifts with group sizes. Many measures of polarisation fall into this category. For example, the affective polarisation measure in Boxell et al. (2020) is a special case of this measure for $\nu = 0$ and $\kappa = (\sum_{t=1}^{4} \pi_t)^{-2}$.

[34]The Axioms put forward in Esteban and Ray (1994) are explained in more detail in Appendix K.

The Esteban and Ray (1994) polarisation measure P was constructed for a one-dimensional variable y, for example, income. Polarisation, in our case, has to be measured over all four ideological types. To do this, we divide people into meaningful ideological groups based on their dominant type share. That is, we allocate individuals to one of our four groups based on their highest type share at the individual i-level. We then add up the θ_i type shares amongst the defined group members to get the mean type share, defined as $\tilde{\theta}_t$. This differs from the mean type shares $\bar{\theta}_t$ we have presented earlier on the basis that we are only taking the mean over individuals with the *same dominant type* rather than over the whole population.

Given this modification, our polarisation measure is defined as:

$$P(\pi, \theta) = \kappa \sum_{t=1}^{4} \sum_{j=1}^{4} \pi_t^{1+\nu} \pi_j (|\tilde{\theta}_{t1} - \tilde{\theta}_{j_1}|\rho_{t1} + |\tilde{\theta}_{t2} - \tilde{\theta}_{j_2}|\rho_{t2} + |\tilde{\theta}_{t3} - \tilde{\theta}_{j_3}|\rho_{t3} + |\tilde{\theta}_{t4} - \tilde{\theta}_{j_4}|\rho_{t4})$$

where π_t and π_j are the number of people who have the dominant type share t and j. The means of the type shares in each of the four dominant type groups are $\tilde{\theta}_t$ for its own type and $\tilde{\theta}_j$ for a generic other type. The second subscript on $\tilde{\theta}_t$ and $\tilde{\theta}_j$ represents the dominant type group we are conditioning on when calculating the absolute distance between groups. Finally, $\rho_{tj} = \frac{3 - corr(\beta_t, \beta_j)}{2}$ uses information from the β type vectors. As such, ρ_{tj} is a measure of the similarity of types based on the correlation of types rescaled to be contained in the $[1, 2]$ interval. Individuals of dominant type t weight differences in type t by 1 while all other type differences have weight strictly larger than one. As an example, consider setting type t as the Liberal Centrist and j is the Conservative Centrist. We index the Liberal Centrist as the type 1 in the second conditioning subscript. The calculation $|\tilde{\theta}_{t1} - \tilde{\theta}_{j_1}|$ then represents the (absolute) difference between the mean Liberal Centrist type share for dominant Liberal Centrist individuals and the mean Liberal Centrist type share for dominant Conservative Centrist individuals. This can be interpreted as a measure of how close different ideological groups are despite their contrasting dominant type shares. That is, a Liberal Centrist and a Conservative Centrist are more likely to 'get along' if they have high minority-type shares in each other's ideology.

The other components of $P(\pi, \theta)$ are the polarisation sensitivity parameter ν, which we fix at $\nu = 0.5$, and the constant $\kappa = (\sum_{t=1}^{4} \pi_t)^{-(2+\nu)}$ that serves to normalise the polarisation measure by population size. We provide more detail and show how P varies with different values of ν in Appendix K.

Intuitively, the polarisation measure will be largest for the case where there are two major type share groups of identical size who exhibit completely different type shares. An example of this would be a bipolar Liberal Centrist and Right Anarchist society where each type group had very small minority shares in the other type. This provides a natural link back to our earlier measure of citizen slant. Since an increase in citizen slant implies an increase in the means for $\tilde{\theta}_t$ and $\tilde{\theta}_j$, absolute differences in type shares between the groups increase, and polarisation rises due to stronger alienation effects.

It is also useful to note how polarisation also depends on the relative sizes of the groups within a population, as measured by π_t and π_j. For example, given the same between-group differences in types, a country in which two groups each make up 50% of the population will be more polarised than a country with 4 groups each making up 25% of the population.

We calculate the polarisation measure separately for each country and wave in our sample. The ranking of the countries based on their polarisation in each wave is shown in Figure 9. The ranking of countries according to Wave 5 polarisation levels is distinct from the earlier ranking for citizen slant. Denmark, which has the lowest level of polarisation, provides an instructive example of how the P polarisation measure combines information. The inputs into the result for Denmark are its high Liberal Centrist type share (above 0.5 - see Figure 4) and high level of within-person concentration or slant (which intensifies over time - see Figure 8). Hence the low Danish P measure reflects a case of ideological consensus where there is a major plural type (Liberal Centrist) that is strongly held by individuals (as manifested in a high slant).

Figure 9: Polarisation by Country

(a) Change in Polarisation by Country **(b) Ranking of Countries by Polarisation (Wave 5)**

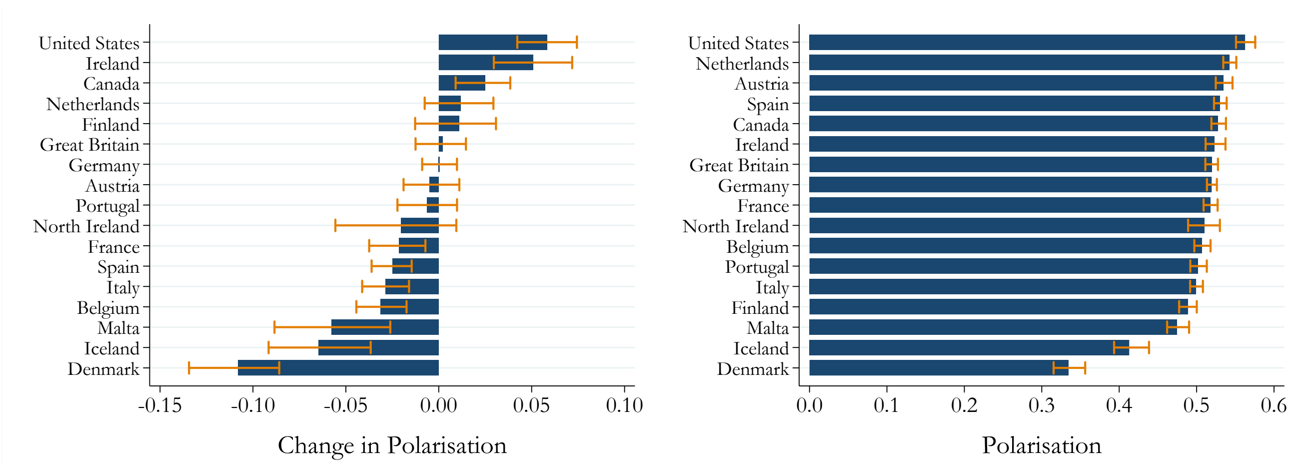

Notes: Panel (a) shows the change in country-level polarisation measures from Waves 2 (1989-1993) to 5 (2005-2009) calculated following Esteban and Ray (1994). Panel (b) shows the level of the country-level polarisation measure in Wave 5. Confidence intervals are based on 500 bootstrap iterations.

The US, which sits at the top of the polarisation ranking in Wave 5, provides a sharp comparison that again illustrates the mechanics of the P measure. It has a relatively even spread of type shares, with shares of around 30% for the Liberal Centrist, Conservative Centrist and Right Anarchist types. Hence, the group size effect picked up by the π_t and π_j terms is stronger in the US compared to unipolar cases such as Denmark. Overall, the increase in polarisation in the US is mainly driven by the rise in slant over time (Figure 8), which contributes by intensifying the alienation effect. The changes in dominant type composition in the US only have a smaller influence on the polarisation measure.

However, it should be noted that, across countries, the changes in polarisation over time are not dramatic, with most of the shifts occurring in the 2-5% range relative to baseline values in Wave 2. A key point to note is that the defining feature of some of the polarisation episodes seen in the data is the *qualitative* content of developments. The US is the banner example here since the increase in polarisation was driven by an increase in the presence of Anarchist types. Hence the US experience with polarisation has the extra dimension of also reflecting the trend of a 'disappearing centre', which is likely to have additional consequences for social cohesion over and above the increase in P that we measure statistically.

In unreported results, we also calculated our polarisation measure using the 3, 5, 6 and 7-type models. The US is consistently either the top country or within the top 3 of the most polarised countries across these models.

5 Conclusion

In this paper, we have proposed a new way to identify the latent ideologies of individuals from survey data. Our approach does not presuppose any ideological structure for individuals. Nonetheless, we are able to identify interpretable, consistent and stable ideological types in the data. The findings generally align with the left-right framing frequently used in the social sciences, but we also identify anti-establishment 'anarchist' ideologies that are characterised by their distrust in important societal institutions.

The approach taken in this paper can also be extended in a number of directions. Firstly, the basic approach outlined here can be applied to other survey datasets, both for the countries studied here and for those outside North America and Western Europe. Indeed, our measure of topic cohesion might be used for any topic modelling application with non-text data. Secondly, the approach is general enough to be used to study questions beyond political ideologies, such as clusters of cognitive and non-cognitive skills, behavioural patterns or personality types.

This latter point about extensions that cover subjects apart from political views is potentially very rich. Contributions such as Ortoleva and Snowberg (2015); Chapman et al. (2018) and Enke (2020) have identified behavioural foundations of political views that systematically map into voting and other outcomes. However, we see it as plausible that our low-trust 'anarchist' types may have some underpinnings in a further layer of personality or behavioural characteristics. Given sufficiently rich data, these layers could be modelled and validated using the hierarchical, out-of-sample approach we have outlined in this paper. The concept of a *hierarchic system* – suggested by Simon (2019) – is a potential model for such future work, and we think that our paper shows the potential of unsupervised learning methods to model such latent, unobserved characteristics at a new level of depth.

References

D. Acemoglu, G. Egorov, and K. Sonin. A political theory of populism. *The Quarterly Journal of Economics*, 128(2):771–805, 2013.

G. A. Akerlof and R. E. Kranton. Economics and identity. *The Quarterly Journal of Economics*, 115(3): 715–753, 2000.

R. Akerlof and L. Rayo. Narratives and the economics of the family. *CEPR Discussion Paper No. DP15152*, 2020.

A. Alesina, E. L. Glaeser, and B. Sacerdote. Why doesn't the united states have a european-style welfare state? *Brookings Papers on Economic Activity*, 2001(2):187–277, 2001.

A. Alesina, A. Devleeschauwer, W. Easterly, S. Kurlat, and R. Wacziarg. Fractionalization. *Journal of Economic growth*, 8(2):155–194, 2003.

A. Alesina, P. Giuliano, and N. Nunn. On the origins of gender roles: Women and the plough. *Quarterly Journal of Economics*, 128(2):469–530, 2013.

N. Aletras and M. Stevenson. Evaluating topic coherence using distributional semantics. *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013)*, 2013.

N. Aletras, T. Baldwin, J. H. Lau, and M. Stevenson. Representing topics labels for exploring digital libraries. *Proceedings of the 14th ACM/IEEE-CS Joint Conference on Digital Libraries,IEEE Press*, pages 239–248, 2014.

Y. Algan, S. Guriev, E. Papaioannou, and E. Passari. The european trust crisis and the rise of populism. *Brookings Papers on Economic Activity*, 2017(2):309–400, 2017.

S. Ansolabehere, J. Rodden, and J. M. Snyder. Purple america. *Journal of Economic Perspectives*, 20 (2):97–118, 2006.

E. Ash. The political economy of tax laws in the us states. Technical report, Working Paper, Columbia University, 2015.

D. Atkin, E. Colson-Sihra, and M. Shayo. How do we choose our identity? a revealed preference approach using food consumption. *Journal of Political Economy*, 129(4):1193–1251, 2021.

O. Bandiera, A. Prat, S. Hansen, and R. Sadun. Ceo behavior and firm performance. *Journal of Political Economy*, 128(4):1325–1369, 2020.

S. O. Becker, T. Fetzer, and D. Novy. Who voted for brexit? a comprehensive district-level analysis. *Economic Policy*, 32(92):601–650, 2017.

R. Benabou. Joseph schumpeter lecture: ideology. *Journal of the European Economic Association*, 6 (2-3):321–352, 2008.

R. Bénabou and J. Tirole. Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–64, 2016.

M. Bertrand and E. Kamenica. Coming apart? cultural distances in the united states over time. Technical report, National Bureau of Economic Research, 2018.

D. G. Blanchflower and A. J. Oswald. Is well-being u-shaped over the life cycle? *Social science & medicine (1982)*, 66(8):1733–1749, 2008.

L. Blaydes and J. Grimmer. Political cultures: Exploring the long-run determinants of values transmission. In *Annual Meeting of the International Political Economy Society, Claremont Graduate School, Claremont, CA*, 2013.

D. M. Blei and J. D. Lafferty. Dynamic topic models. *Proceedings of the 23rd international conference on Machine learning*, pages 113–120, 2006.

D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.

G. Bonomi, N. Gennaioli, and G. Tabellini. Identity, beliefs, and political conflict. *The Quarterly Journal of Economics*, 136(4):2371–2411, 2021.

W. Bossert, C. D'Ambrosio, and E. La Ferrara. A generalized index of fractionalization. *Economica*, 78(312):723–750, 2011.

L. Boxell, M. Gentzkow, and J. M. Shapiro. Cross-country trends in affective polarization. *The Review of Economics and Statistics*, pages 1–60, 2020.

P. Buisseret and R. Van Weelden. Crashing the party? elites, outsiders, and elections. *American Journal of Political Science*, 64(2):356–370, 2020.

L. Bursztyn, G. Egorov, and S. Fiorin. From extreme to mainstream: The erosion of social norms. *American economic review*, 110(11):3522–48, 2020.

J. Cagé, N. Hervé, and M.-L. Viaud. The production of information in an online world. *The Review of Economic Studies*, 87(5):2126–2164, 2020.

N. Canen, C. Kendall, and F. Trebbi. Unbundling polarization. *Econometrica*, 88(3):1197–1233, 2020.

F. Caselli and W. J. Coleman. On the theory of ethnic conflict. *Journal of the European Economic Association*, 11(suppl. 1):161–192, 2013.

J. Chang, J. Boyd-Graber, D. M. Blei, C. Wang, and S. M. Gerrish. Reading tea leaves: How humans interpret topic models. *Advances in neural information processing systems*, pages 288–296, 2009.

J. Chapman, M. Dean, P. Ortoleva, E. Snowberg, and C. Camerer. Econographics. *National Bureau of Economic Research*, 2018.

Y. Che, Y. Lu, J. R. Pierce, P. K. Schott, and Z. Tao. Does trade liberalization with china influence us elections? Technical report, National Bureau of Economic Research, 2016.

P. E. Converse. The nature of belief systems in mass publics (1964). *Critical review*, 18(1-3):1–74, 2006.

E. Dal Bó, F. Finan, O. Folke, J. Rickne, and T. Persson. Economic losers and political winners: The rise of the radical right in sweden. Technical report, mimeo Uppsala University, 2017.

E. Dal Bó, P. Dal Bó, and E. Eyster. The demand for bad policy when voters underappreciate equilibrium effects. *The Review of Economic Studies*, 85(2):964–998, 2018.

David Newman, Jey Han Lau, Karl Grieser, and Timothy Baldwin. Automatic evaluation of topic coherence. *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the ACL,*, pages 100–108, 2010.

K. Desmet and R. Wacziarg. The cultural divide. *The Economic Journal*, 131(637):2058–2088, 2021.

K. Desmet, I. Ortuno-Ortin, and R. Wacziarg. Endogenous partitions. *Working Paper*, 2022.

C. Ding and X. He. K-means clustering via principal component analysis. In *Proceedings of the twenty-first international conference on Machine learning*, page 29. ACM, 2004.

C. Dippel, R. Gold, and S. Heblich. Globalization and its (dis-) content: Trade shocks and voting behavior. Technical report, National Bureau of Economic Research, 2015.

D. Dorn, G. Hanson, K. Majlesi, et al. Importing political polarization? the electoral consequences of rising trade exposure. *American Economic Review*, 110(10):3139–83, 2020.

A. Downs. An economic theory of political action in a democracy. *Journal of Political Economy*, 65(2): 135–150, 1957.

J.-Y. Duclos, J. Esteban, and D. Ray. Polarization: Concepts, measurement, estimation. *Econometrica*, 72(6):1737–1772, 2004.

B. Enke. Moral values and voting. *Journal of Political Economy*, 128(10):3679–3729, 2020.

B. Enke, R. Rodríguez-Padilla, and F. Zimmermann. Moral Universalism and the Structure of Ideology. *The Review of Economic Studies*, 09 2022.

J.-M. Esteban and D. Ray. On the measurement of polarization. *Econometrica*, pages 819–851, 1994.

M. P. Fiorina and S. J. Abrams. Political polarization in the american public. *Annual Review of Political Science*, 11:563–588, 2008.

M. Gentzkow, B. Kelly, and M. Taddy. Text as data. *Journal of Economic Literature*, 57(3):535–74, 2019a.

M. Gentzkow, J. M. Shapiro, and M. Taddy. Measuring group differences in high-dimensional choices: Method and application to congressional speech. *Econometrica*, 87(4):1307–1340, 2019b.

V. Geoghegan. *Political ideologies: An introduction*. Routledge, 2003.

A. Gethin, C. Martínez-Toledano, and T. Piketty. Brahmin left versus merchant right: Changing political cleavages in 21 western democracies, 1948–2020. *The Quarterly Journal of Economics*, 137(1):1–48, 2022.

E. L. Glaeser and B. A. Ward. Myths and realities of american political geography. *Journal of Economic Perspectives*, 20(2):119–144, 2006.

T. L. Griffiths and M. Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101(suppl 1):5228–5235, 2004.

J. Grimmer. A bayesian hierarchical topic model for political texts: Measuring expressed agendas in senate press releases. *Political Analysis*, 18(1):1–35, 2009.

J. H. Gross and D. Manrique-Vallier. A mixed-membership approach to the assessment of political ideology from survey responses. In *Individual Presentation, Society for Political Methodology, 29th Annual Summer Meeting, Chapel Hill, NC*. Citeseer, 2012.

G. M. Grossman and E. Helpman. Identity politics and trade policy. *The Review of Economic Studies*, 88(3):1101–1126, 2021.

L. Guiso, H. Herrera, M. Morelli, and T. Sonno. Demand and supply of populism. *IGIER (Innocenzo Gasparini Institute for Economic Research), Bocconi University*, 2017.

J. Haidt. *The righteous mind: Why good people are divided by politics and religion*. Vintage, 2012.

S. Hansen, M. McMahon, and A. Prat. Transparency and deliberation within the fomc: a computational linguistics approach. *The Quarterly Journal of Economics*, 2014.

S. J. Hill and C. Tausanovitch. A disconnect in representation? comparison of trends in congressional and public polarization. *The Journal of Politics*, 77(4):1058–1075, 2015.

M. Hoffman, F. R. Bach, and D. M. Blei. Online learning for latent dirichlet allocation. In *advances in neural information processing systems*, pages 856–864, 2010.

M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *the Journal of machine Learning research*, 14(1):1303–1347, 2013.

J. Hopkin. Anti-system politics: the crisis of market liberalism in rich democracies. *Oxford University Press*, 2020.

R. Inglehart. *Modernization and postmodernization: Cultural, economic, and political change in 43 societies*. Princeton University Press, 1997.

R. F. Inglehart, M. Basanez, and A. Moreno. *Human Values and Beliefs: A Cross-Cultural Sourcebook*. University of Michigan Press, Ann Arbor, 2010. ISBN 978-0-472-10833-6.

Z. Jelveh, B. Kogut, and S. Naidu. Political language in economics. *Working Paper*, 2015.

J. Jensen, S. Naidu, E. Kaplan, L. Wilse-Samson, D. Gergen, M. Zuckerman, and A. Spirling. Political polarization and the dynamics of political language: Evidence from 130 years of partisan speech [with comments and discussion]. *Brookings Papers on Economic Activity,*, Fall:1–81, 2012.

S. Jolly, R. B. L. H. G. M. J. P. J. R. M. Steenbergenand, and M. A. Vachudova. Chapel hill expert survey trend file, 1999-2019. *Electoral Studies 75*, 2022. doi: https://doi.org/10.1016/j.electstud.2021.102420.

E. Kaplan, J. L. Spenkuch, and R. Sullivan. Partisan spatial sorting in the united states: A theoretical and empirical overview. *Journal of Public Economics*, 211:104668, 2022.

D. Kim and A. Oh. Topic chains for understanding a news corpus. *International Conference on Intelligent Text Processing and Computational Linguistics, Springer*, pages 163–176, 2011.

D. R. Kinder and N. P. Kalmoe. Neither liberal nor conservative: Ideological innocence in the american public. *University of Chicago Press*, 2017.

J. H. Lau and T. Baldwin. The sensitivity of topic coherence evaluation to topic cardinality. *Proceedings of NAACL-HLT*, pages 483–487, 2016.

J. H. Lau, K. Grieser, D. Newman, and T. Baldwin. Automatic labelling of topic models. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1, Association for Computational Linguistics*, pages 1536–1545, 2011.

J. H. Lau, D. Newman, and T. Baldwin. Machine reading tea leaves: Automatically evaluating topic coherence and topic model quality. *Proceedings ofthe 14th Conference ofthe European Chapter ofthe Association for Computational Linguistic*, pages 530–539, 2014.

P. Lehmann, T. Burst, T. Matthieß, S. Regel, A. Volkens, B. Weßels, and L. Zehnter. Manifesto project: Version 2022a. 2022. doi: https://doi.org/10.25522/manifesto.mpds.2022a.

J. G. Montalvo and M. Reynal-Querol. Ethnic polarization, potential conflict, and civil wars. *American economic review*, 95(3):796–816, 2005.

C. Mudde. The populist zeitgeist. *Government and opposition*, 39(4):541–563, 2004.

E. Munro and S. Ng. Latent dirichlet analysis of categorical survey responses. *Journal of Business & Economic Statistics*, 2020.

D. Murray. Populism: It's the bbc's new buzzword, being used to sneer at the 'uneducated' 17 million who voted for brexit. *Daily Mail*, 16:December, 2016.

W. T. Norman. Toward an adequate taxonomy of personality attributes: Replicated factor structure in peer nomination personality ratings. *The Journal of Abnormal and Social Psychology*, 66(6):574, 1963.

P. Norris. *Democratic deficit: Critical citizens revisited*. Cambridge University Press, Cambridge [etc.], 2016. ISBN 978-0521127448.

P. Ortoleva and E. Snowberg. Overconfidence in political behavior. *American Economic Review*, 105 (2):504–35, 2015.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, and V. Dubourg. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(Oct):2825–2830, 2011.

K. T. Poole and H. Rosenthal. A spatial model for legislative roll call analysis. *American Journal of Political Science*, pages 357–384, 1985.

D. Rodrik. Populism and the economics of globalization. *Journal of international business policy*, 1(1): 12–33, 2018.

M. Rooduijn, S. Van Kessel, C. Froio, A. Pirro, S. De Lange, D. Halikiopoulou, P. Lewis, C. Mudde, and P. Taggart. The populist: An overview of populist, far right, far left and eurosceptic parties in europe. *The PopuList. https://popu-list. org*, 2019.

C. Schwarz. ldagibbs: A command for topic modeling in stata using latent dirichlet allocation. *Stata Journal*, 18(1):101–117, 2018.

M. Shayo. A model of social identity with an application to political economy: Nation, class, and redistribution. *American Political science review*, 103(2):147–174, 2009.

H. A. Simon. The sciences of the artificial. *MIT press*, 2019.

C. Spearman. " general intelligence," objectively determined and measured. *The American Journal of Psychology*, 15(2):201–292, 1904.

J. Su and W.-P. Liao. Latent dirichlet allocation for text and image topic modeling, 2013.

E. Tupes and R. Christal. Stability of personality factors based on trait ratings. *USAF ASD Tech. Rep*, 1961.

A. Vincent. *Modern political ideologies*. John Wiley & Sons, 2009.

C. Wang, D. M. Blei, and D. Heckerman. Continuous time dynamic topic models. *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*, pages 579–586, 2008.

Y. S. Wang, R. L. Matsueda, E. A. Erosheva, et al. A variational em method for mixed membership models with multivariate rank data: An analysis of public policy preferences. *The Annals of Applied Statistics*, 11(3):1452–1480, 2017.

F. Zhao, X. Ren, S. Yang, and X. Yang. On privacy protection of latent dirichlet allocation model training. *arXiv preprint arXiv:1906.01178*, 2019.

Online Appendix

This appendix presents further details on data, the LDA model, additional robustness exercises, and additional results:

- Appendix A provides details on the selection of features from the WVS.
- Appendix B provides additional detail on the LDA model.
- Appendix C discusses the interpretation of the β vectors in our setting.
- Appendix D discusses the construction of our type cohesion measure.
- Appendix E shows sensitivity checks for our LDA model.
- Appendix F shows additional information on the type hierarchies.
- Appendix G provides a replication of results using data from the ESS.
- Appendix H shows the results for the 7th Wave of the WVS.
- Appendix I provides additional details on our definition of populist parties.
- Appendix J provides additional citizen slant results.
- Appendix K provides additional details on the polarisation measure.
- Appendix L shows clustering results using other clustering algorithms.

A Appendix: Additional Details on the Selection of Question from the WVS

This section describes in more detail, the selection process that lead to the 29 questions that are used in the paper. There are 7 waves of the World Value Survey (WVS) and 5 waves of the European Value Survey (EVS). The 5 Waves of the EVS correspond to the 1st, 2nd, 4th, 5th and 7th wave of the WVS. When constructing the Integrated Value Survey by combining the WVS and EVS we excluded the 1st wave since it contained a smaller set of countries and questions. The Integrated Value Survey (WVS) in total contains 971 different items grouped in 13 different categories (number of questions in brackets): Environment (25), Family (64), National Identity (105), Perceptions of life (210), Politics and Society

(267), Religion and Morale (122), Science (2), Security (22), Socio-demographics (38), Special Indexes (3), Structure of the file (25), Sylatech module (42) and Work (46).[35]

We limited the set of questions to those questions which were consistently asked in the 2nd, 4th and 5th wave of WVS. This already reduced the set of possible questions down to 92. From these 92 questions, we chose our 29 based on which questions are most salient for the evaluation of a person's ideological type. The excluded question are listed in Table A.2. For example, we exclude questions about family structure (e.g., single parenting, beliefs in marriage), questions about non-political moral values (e.g., 'important child qualities'), life satisfaction, and generic trust in others. This is because our aim is to model the latent structure of political opinions that are most analogous to the concept of ideology. In the conclusion of the paper, we describe possible extensions of our general approach that would accommodate interactions between (say) behavioural characteristics and political beliefs.

In Appendix E we show that the selection of these 29 questions is not crucial for our findings and that the ideological types are very similar if we use all 92 questions. We further show that also removing any of the 29 questions from our data has no bearing on our results.

A further point is that LDA does not allow for missing responses in the data. If we simply excluded all observations with any missing responses and restricted ourselves to observations with complete sets of answers, we would need to drop sizable fractions of the WVS data. We instead impute a small set of missing responses with the sample mean of the non-missing data in the same wave. This treatment of missing data allows LDA to use the information from this observation across other questions that have non-missing values. Moreover, the imputation has only a minimal effect on the LDA classification since the sample mean does not influence the classification of each individual. Imputation with the mean is also preferable to an alternative approach where we would simply replace all missing responses with 0s because the 0s would bias the classification.

In Table F.3, we report the resulting type hierarchy for the approximate 50% of observations in the sample that do not require any imputation. As it turns out, the resulting type hierarchy is nearly identical to the one achieved with imputation. Also, the resulting individual-level type shares are very similar. This suggests that the imputation of missings is not having a major influence on our results.

B Appendix: Additional Details on the LDA Model Inference

One difference between our application and the standard use of LDA is that in our case features can only appear once for each observation, that is, individuals can only answer each question once in our

[35]The categories socio-demographics, special indexes, structure of the study, Sylatech module and work do not contain any questions concerning the values of people.

Table A.1: Summary Statistics, WVS Questions

Code	Question	Scale	Share For	Share Against
	On this list are various groups of people. Could you please mention any that you would not like to have as neighbours?	0,1		
A124_02	People of a different race		0.097	0.903
A124_06	Immigrants/foreign workers		0.123	0.877
A124_07	People who have AIDS		0.208	0.792
A124_08	Drug addicts		0.638	0.362
A124_09	Homosexuals		0.219	0.781
C002	Do you agree, disagree or neither agree nor disagree with the following statements?: "When jobs are scarce, employers should give priority to people of this country over immigrants."	1-3	0.600	0.305
E036	Rate your view on a 1 to 10 scale between the positions: "Private ownership of business and industry should be increased" vs. "Government ownership of business and industry should be increased"	1-10	0.506	0.255
E037	Rate your view on a 1 to 10 scale between the positions: "Government should take more responsibility to ensure that everyone is provided for" vs. "People should take more responsibility to provide for themselves"	1-10	0.376	0.469
E039	Rate your view on a 1 to 10 scale between the positions: "Competition is good. It stimulates people to work hard and develop new ideas" vs. "Competition is harmful. It brings out the worst in people"	1-10	0.613	0.215
	Could you tell me how much confidence you have in these organizations:	1-4		
E069_01	Church		0.519	0.481
E069_02	Armed forces		0.567	0.433
E069_04	The press		0.356	0.644
E069_05	Labour unions		0.385	0.615
E069_06	The police		0.704	0.296
E069_07	Parliament		0.413	0.587
E069_08	The civil services		0.451	0.549
E069_13	Major Companies		0.432	0.568
E069_17	Justice system/courts		0.533	0.466
	Please tell me for each of the following actions whether you think it can always be justified, never be justified, or something in between:	1-10		
F114	Claiming government benefits		0.076	0.869
F115	Avoiding a fare on public transport		0.086	0.842
F116	Cheating on taxes		0.106	0.828
F117	Someone accepting a bribe		0.035	0.931
F118	Homosexuality		0.407	0.432
F119	Prostitution		0.196	0.663
F120	Abortion		0.348	0.458
F121	Divorce		0.496	0.280
F122	Euthanasia		0.418	0.430
F123	Suicide		0.149	0.730
G006	How proud are you of your nationality?	1-4	0.885	0.115

Notes: This table reports summary statistics for the recoded questions from the WVS. The third column reports the original coding of the question in the WVS. Questions with a binary or 1–4 coding are recoded into two indicator variables expressing either support or opposition to each issue. Questions with 1–3 or 1–10 allow for a neutral coding if the answer is coded as 3 or 5 in which case both indicator variables are coded as zero. The fourth (fifth) column contains the share of people that are coded as a positive (negative) response to the question.

Table A.2: List of Excluded Questions

Code	Questions	Code	Questions
A001	Important in Life: Family	D057	Being a housewife just as fulfilling
A002	Important in Life: Friends	E001	Aims of Country: First choice
A003	Important in Life: Leisure time	E002	Aims of Country: Second choice
A004	Important in Life: Politics	E003	Aims of Respondent: First choice
A005	Important in Life: Work	E004	Aims of Respondent: Second choice
A006	Important in Life: Religion	E005	Most Important: First choice
A008	Feeling of happiness	E006	Most Important: Second choice
A009	State of health (subjective)	E012	Willingness to fight for country
A029	Important Child Qualities: Independence	E015	Future Changes: Less importance placed on work
A030	Important Child Qualities: Hard work	E016	Future Changes: More emphasis on technology
A032	Important Child Qualities: Feeling of responsibility	E018	Future Changes: Greater respect for authority
A034	Important Child Qualities: Imagination	E019	Future Changes: More emphasis on family life
A035	Important Child Qualities: Tolerance and respect for other people	E022	Opinion about scientific advances
A038	Important Child Qualities: Thrift saving money and things	E023	Interest in politics
A039	Important Child Qualities: Determination perseverance	E025	Political Action: signing a petition
A040	Important Child Qualities: Religious faith	E026	Political Action: joining in boycotts
A041	Important Child Qualities: Unselfishness	E027	Political Action: attending lawful/peaceful demonstrations
A042	Important Child Qualities: Obedience	E033	Self positioning in political scale
A124_03	Neighbours: Heavy drinkers	E035	Income equality
A124_05	Neighbours: Muslims	E069_11	Confidence: The government
A165	Most people can be trusted	E069_12	Confidence: The political parties
A170	Satisfaction with your life	E069_18	Confidence: The European Union
A173	How much freedom of choice and control	F001	Thinking about meaning and purpose of life
B001	Would give part of my income for the environment	F028	How often do you attend religious services
B002	Increase in taxes if used to prevent environmental pollution	F034	Religious person
B003	Government should reduce environmental pollution	F035	Churches give Answers: moral problems
C001	Jobs Scarce: Men should have more right to a job than women	F036	Churches give Answers: the problems of family life
C006	Satisfaction with financial situation of household	F037	Churches give Answers: people´s spiritual needs
C059	Fairness: One secretary is paid more	F038	Churches give Answers: the social problems
D018	Child needs a home with father and mother	F063	How important is God in your life
D022	Marriage is an out-dated institution	F065	Moments of prayer, meditation...
D023	Woman as a single parent		

Notes: This table contains the questions that were excluded from the baseline LDA model.

setting, while words can appear more than once in a text document. This section describes how the LDA model inference approach we use nonetheless remains valid and compares our implementation of LDA to the model suggested by Gross and Manrique-Vallier (2012).

Initially, it is instructive to note the fundamental source of similarity between the two settings (text data versus survey responses). It is also the case in textual data that many words appear only once in each document. As a result, LDA mostly infers the importance of a word for a particular topic based on how many documents use a given word. Words that are used in nearly all documents that relate to a topic will get a higher weight than words that are rarely used. For example, while any paper on taxation will use the word "tax" at some point in their overall narrative, only very few will use the word "persuasion". Reflecting this, some applications of LDA use a binary matrix of word occurrences (e.g., Su and Liao, 2013; Zhao et al., 2019). This is equivalent to words only being allowed to appear once per document. In our setting, an issue position will get a higher weight within a topic if more people holding a particular ideology express that issue position. In contrast, issue positions that are only expressed by few people of a particular ideology will receive a lower weight.

To formally understand why the approximation of the likelihood works even when features can only appear once, it is helpful to analyse the updating steps of the approximation algorithms. Amongst the existing LDA approximation algorithms, the collapsed Gibbs sampling algorithm developed by Griffiths and Steyvers (2004) (and for example used in Schwarz (2018)) provides the clearest insight into the workings of LDA. The collapsed Gibbs sampler works by consecutively sampling a new topic (type) assignment for each feature – words in text or individual's question responses – based on the current topic assignment of all other features.[36]

In our application, the Gibbs sampler would calculate the probabilities for $z_{i,n}$ – the type assignment of response n and individual i – conditional on $z_{-(i,n)}$ the current type of all other features and the given response r based on the following equation:

$$P(z_i = t | z_{-r}, \beta) \propto \frac{\eta_t^{(r)} + \gamma}{\eta_t^{(.)} + Q\gamma} \cdot \frac{n_i^{(t)} + \alpha}{n_i^{(.)} + T\alpha}$$

where $n_t^{(r)}$ is the number of times response r is currently assigned to type t and $\eta_t^{(.)}$ is the number of times any response is assigned to type t. Similarly, $\eta_t^{(i)}$ is the number of responses of individual i assigned to type t and $\eta_i^{(.)}$ is the total number of responses given by individual i. α, γ are the Dirichlet priors, Q is the number of Questions (58 in our case) and T is the number of types. The first term in the expression above therefore captures the probability of observing response r conditional on type t, while the second term captures the probability of type assignment t for individual i.

[36]In the beginning, the Gibbs sampler is initialised by randomly assigning features to topics (types)

After calculating $P(z_i = t|z_{-r}, w)$ for all $t \in T$, the Gibbs sampler randomly draws a new type assignment based on the calculated probabilities. In other words, the Gibbs sampler exploits how likely $P(z_i = 1|z_{-r}, w)$ is relative to $P(z_i = 2|z_{-r}, w), \cdots, P(z_i = T|z_{-r}, w)$. Hence, the assignment of $z_{i,q}$ to types captures the relative frequency of response r conditional on types. Since all individuals can give response r at most once the type assignments are valid. The type assignments would only be biased if individuals differed in how often they could give response r.

Gross and Manrique-Vallier (2012) developed an alternative model to model survey responses, which leads to an alternative updating equation:

$$P(z_i = t|z_{-r}, \beta) \propto \frac{\eta_t^{(r)} + \gamma}{\eta_t^{(\neg r)} + R\gamma} \cdot \frac{\eta_i^{(t)} + \alpha}{\eta_i^{(.)} + T\alpha}$$

, where $\eta_t^{(\neg r)}$ is the number of times another response is given to question q, and R is the number of possible responses to question q. Therefore, the difference between the two updating equations is that LDA uses the probability of response r relative to all other responses given by type t, while the Gross and Manrique-Vallier (2012) model uses the probability relative to other responses given to the same question. Hence, in LDA the probabilities of responses sum to one across all questions ($\sum_{q=1}^{Q} \beta_q = 1$). Whereas in the (Gross and Manrique-Vallier, 2012) model, the probabilities of responses sum to one for each individual question ($\sum_{r=1}^{R} \beta_{q,r} = 1$).

Put more simply, in the Gross and Manrique-Vallier (2012) model, every question is treated separately. In the LDA model, putting weight on one issue reduces the weight on other issues. In this way, the LDA model creates a ranking of issue positions and their importance for the ideological types. Our recoding of questions into 2 features also naturally incorporates this feature of LDA, as an individual who, for example, states trust in the government cannot also state distrust in the government.

C Appendix: Interpretation of the β Vectors

LDA allows for repeated draws of a feature, while in our application, people can only answer a question once. As already discussed in the main part of the paper and Appendix B, this does not influence the validity of LDA, since LDA exploits how often features appear relative to each other.

However, this difference influences the interpretation of the β vectors. The β vectors capture the probability that a response is drawn in each of the 29 draws (questions) asked to an individual, e.g. how likely it is that an individual will answer that he is opposed to abortion in each of the 29 draws. Therefore, the β vectors do not take into account that once a person has answered a question, the same person cannot answer the same question again.

As a result, the β still capture which groups are more likely to exhibit an ideological position, but the values do not have a natural interpretation within our setting. If necessary, one can scale up the β probabilities to give them a more natural interpretation within our setting. To do this, one has to calculate the probability that a feature shows up in any of the 29 draws of the LDA taking into account that a question can only be answered once. Given this intuition $P_{f,t}$, the overall probability that a feature f appears if the chosen type is t, can be expressed as $P_{f,t} = \sum_{d=1}^{29}(1-\beta_{f,t})^{d-1}\beta_{f,t}$, where d is the number of the draw (question) and $\beta_{f,t}$ is the value of the β vector for feature f and type t. In this expression, $(1-\beta_{f,t})^{d-1}$ is the probability that the response has not been given in any previous draw, and $\beta_{f,t}$ is the probability that the response will be given in the current draw.

As an example to illustrate this calculation, consider the question of ”Confidence in the Police”. In the 5th wave, the liberal centrist has a value of $\beta_{14,1} = 0.0408$ and the value for the left anarchist is $\beta_{14,3} = 0.0089$. This difference in the β values translate into the following overall differences in probability. While a liberal centrist will express confidence in the police with a probability of 70.1%, the probability that a left anarchist will express similar views is only 22.8%.

This scaling-up does not take into account that some features are mutually exclusive. Hence, the scaled-up probability of the features ”Confidence in the Police” and “No Confidence in the Police” will not necessarily add up to 1.

D Additional Details on Topic Cohesion

D.1. Automatic Evaluation of Topic Model Cohesion

The main theme of the literature on the cohesion of topic models is that humans judge topics to be more consistent based on word co-occurrences (Chang et al., 2009; David Newman et al., 2010; Lau et al., 2014; Lau and Baldwin, 2016). Consider, for example, a topic containing words like ‘labour’, ‘wage’ and ‘firm’, which often appear together in a text, will be judged as highly coherent by humans. An alternative topic that contains words like ‘inflation’, ‘agriculture’ and ‘hospital’ appears incoherent since these words are not used together as frequently.

Given this approach, it is possible to automatically calculate measures of topic cohesion that are highly correlated with human judgment. These measures are usually based on the most frequently occurring words in each topic. One standard approach is to calculate how often words appear together using the Wikipedia corpus (David Newman et al., 2010). The title and sub-sections of the Wikipedia article are used as ‘tags’ for discrete, human-curated topics. The more frequently that words within an

LDA-derived topic appear together in a Wikipedia article (or within a sub-section of an article) then the more coherent the automatically defined topic is judged to be.

In our specific case of using survey response data, there is no equivalent, human-curated outside corpus available to guide analysis. We therefore take the approach of using hold-out samples from within our data to calculate the cohesion scores. Our method thereby exploits the same intuition normally used in the literature on topic model cohesion. The key here is the β issue-position weights can be used as predictions of feature co-occurrence in the hold-out data. A political ideology is judged to be more coherent if people frequently hold issue positions together. We use Normalised Pointwise Mutual Information (NPMI) as our score of topic cohesion since NPMI has been shown to outperform other information metrics such as PMI or Pairwise Log Conditional Probability (LCP) and shows similar performance to pairwise distributional similarity (Aletras and Stevenson, 2013; Lau et al., 2014).

D.2. Making Sense of the NPMI Values

The calculation of the NPMI is based on the independent and joint probabilities of given features k and l. The probability $p(k)$, for example, could capture the share of the population that believes abortion is not justifiable, while $p(l)$ captures the probability that a person has confidence in the church. The joint probability $p(k, l)$ then captures how many people believe that abortion is unjustifiable and have confidence in the church at the same time. The larger the joint $p(k, l)$ is in relation to $p(k)$ and $p(l)$, the higher is NPMI score of the two features.

Re-capping the basic equation from the main paper NPMI is defined as:

$$NPMI_{k,l} = \frac{PMI_{k,l}}{-\ln(p(k,l))} = \frac{\ln\left(\frac{p(k,l)}{p(k)\cdot p(l)}\right)}{-\ln(p(k,l))} \tag{D.1}$$

As an illustration, Table D.1 shows two examples of NPMI scores for different values of $p(k)$ and $p(l)$, as well as different joint probabilities $p(k, l)$. In the first example, both features appear with a probability of 0.2. In the situation where all people who are against abortion also have confidence in the church, the joint probability of the features is 0.2, and the NPMI value will be 1. If the two features were independent of each other one would expect them to appear together in the data with a frequency of $(0.2 \cdot 0.2) = 0.04$. In this situation, the calculated NPMI will be 0. If the joint probability is larger than the probability in the case of independence, then NPMI will be positive. The final two rows of Example 1 in Table D.1 illustrate this relationship.

A technical point to note here is that the exact value of the NPMI depends on the individual as well as the joint probabilities. This is illustrated via the second example reported in Table D.1. Note

Table D.1: Example Calculation NPMI

Case	$p(k)$	$p(l)$	$p(k,l)$	PMI	$NPMI$
Example 1					
Perfect Co-Occurrence	0.2	0.2	0.2	1.609	1
Independence	0.2	0.2	0.04	0	0
$p(k,l) > Independence$	0.2	0.2	0.06	0.405	0.144
$p(k,l) < Independence$	0.2	0.2	0.02	-0.693	-0.177
Example 2					
Perfect Co-Occurrence	0.6	0.6	0.6	0.511	1
Independence	0.6	0.6	0.36	0	0
$p(k,l) > Independence$	0.6	0.6	0.54	0.405	0.658
$p(k,l) < Independence$	0.6	0.6	0.18	-0.693	-0.404

that in both Example 1 and Example 2, the third-row cases are characterised by a joint probability that is 50% larger than in the case of independence. The PMI is identical across the two different 'third-row' cases, but the NPMI is different. Two pairs of features will only have the same NPMI if $log_{p(k,l)}(p(k), p(l)) = log_{p(k,l)}(p(k), p(l))$. In other words, the NPMI is identical if you have to raise the joint probability to the same power to get the product of the individual probabilities.

E Sensitivity Checks LDA Model

In this section, we analyse how sensitive our baseline 4-type model is to changes in the priors and random seeds of LDA as well as the removal and addition of features. The feature addition and removal exercises we run here can be interpreted as a leverage or influence analysis on the statistical definition of our ideological clusters. We are unaware of formal model robustness statistics of this nature in the literature on LDA. Hence, while we think that the exercises below are very promising in terms of the robustness of the basic clusters that they reveal, they should be considered indicative.

E.1. Priors and Seeds

First, to address the concern that our findings could be driven by the choice of priors for the Dirichlet distribution or could be the artefact of the specific random seed, we rerun the LDA for all values of α and γ in the 0-1-interval in steps sizes of 0.1.[37] In total, we fit 100 LDA models, all of which use a different random seed. As such, this analysis will inform us if the same ideological types arise independent of the priors of the LDA model.

[37] Technically speaking α and γ are vectors. As is standard in the LDA literature, we set symmetric prior for each topic and response. Hence, the priors are scalars.

Figure E.1: Robustness Dirichlet Priors

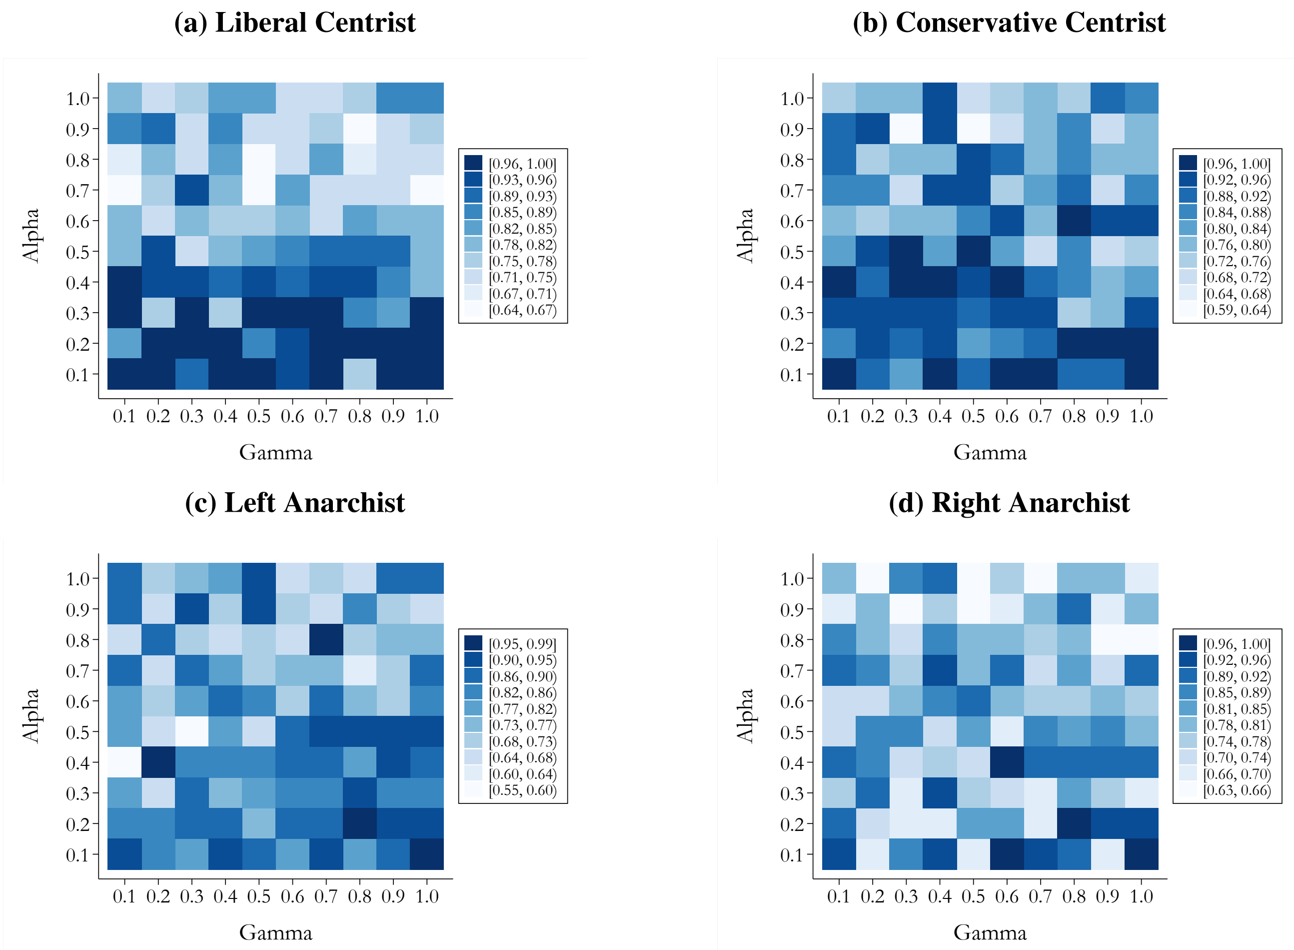

Notes: These heatplots show the correlation of the β vectors of our baseline model with topic models in which we varied the priors for α and γ. Each of the 100 topic models was fit with a different random seed. The darker colours indicate a higher correlation.

The results from this exercise are visualised in Figure E.1. Each square in the heatmaps represents the similarity (as measured by the correlation) of an LDA model using the priors indicated on the x and y-axis to the corresponding baseline type.[38] It is immediately apparent that the types are, on average, very stable independent of the prior and the seed. For example, the lowest similarity for the liberal centrist type is 0.64.

E.2. 'Leave One Out' Clusters

As the next exercise, we re-estimate the 4-type model removing 1 of the 29 questions (2 of the 58 features) at a time. Afterwards, we compare the original model to the new 'leave one out' model based

[38]We always report the correlation to the most similar type among the newly created types. This does not always lead to a 1-to-1 correspondence between the new and old types.

on the similarity of the β vectors, as measured by their correlation. Figure E.2 reports the results of this exercise.

Overall, independent of the particular removed question, we find high correlations between the different β vectors. This is strongest for the Liberal Centrist type, which has an average correlation of 0.979 between the original and leave-one-out models across all dropped questions. This indicates that the types generated by LDA are very closely comparable across the different models. The highest degrees of sensitivity relate to the confidence in institutions questions (where the β correlations are between 0.70-0.80 for three of the types). Another point of sensitivity is questions relating to foreigners/immigration in the case of the Right Anarchist. Given the centrality of the confidence and immigration questions to the character of different types, these sensitivities are within expectations. To provide further robustness, in particular for the anarchist types, we next turn to an exercise in which remove more than one trust question at a time.

E.3. 'Leave Trust Out' Clusters

As the trust questions are of particular importance to the anarchist types, we additionally want to ensure that these types are not an artefact of the number of trust questions we include in the LDA algorithm. We therefore repeat our 'leave question out' exercise but instead remove all possible combinations of the 9 trust questions. This exercise involves estimating a total of 502 LDA models. We then again calculate the similarity of the resulting types to our baseline model.

The results are represented in Figure E.3. The histograms represent the distribution of the similarities over the 502 LDA models. It is immediately apparent that the centrist types are highly stable, and most models exhibit a similarity above 0.9. As expected, the Anarchist types show a larger variance, in particular, the right anarchist type exhibits a distribution centred around 0.8, but in most cases, highly similar types arise.

Note that some of these models are estimated without any trust questions at all. This implies that even without the trust question, very similar question response profiles arise. The fact that within these profiles, the trust questions then play a central role gives us further confidence that institutional trust is an important determinant of political ideology. This leads us to the next issue of how the types might change when we add more information to the feature set.

E.4. Widening the Feature Set

In the next exercise, we investigate how the structure of our clusters changes when we include additional features in the topic model. As described in Appendix A, there are a total of 92 questions that are

Figure E.2: Leave One Out

(a) Liberal Centrist

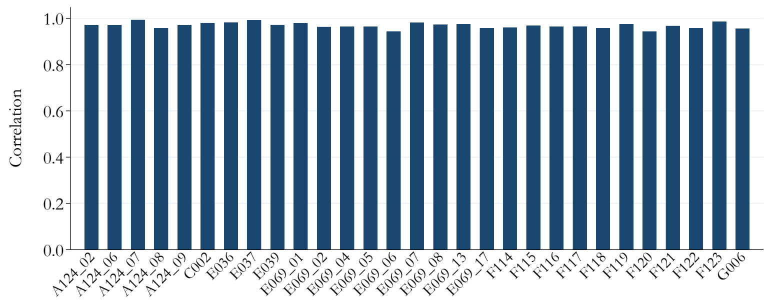

(b) Conservative Centrist

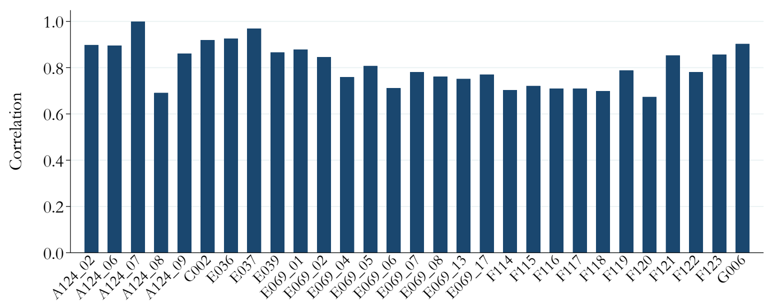

(c) Left Anarchist

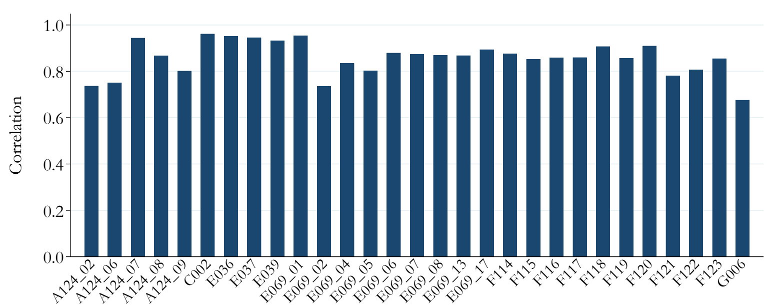

(d) Right Anarchist

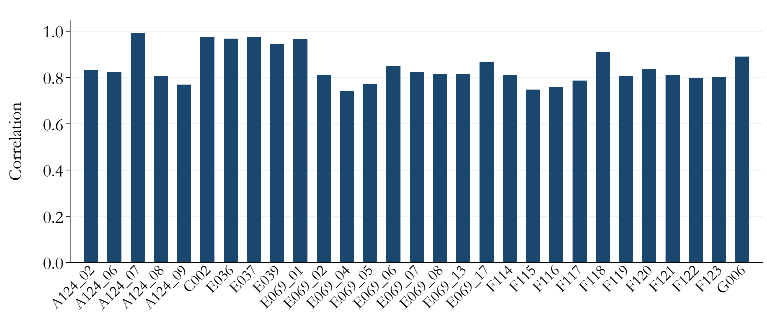

Notes: This figure reports the correlation of the β vectors of our baseline model with topic models in which 1 of the 29 questions from the baseline model was removed (indicated on the x-axis).

Figure E.3: Leaving Trust Questions Out

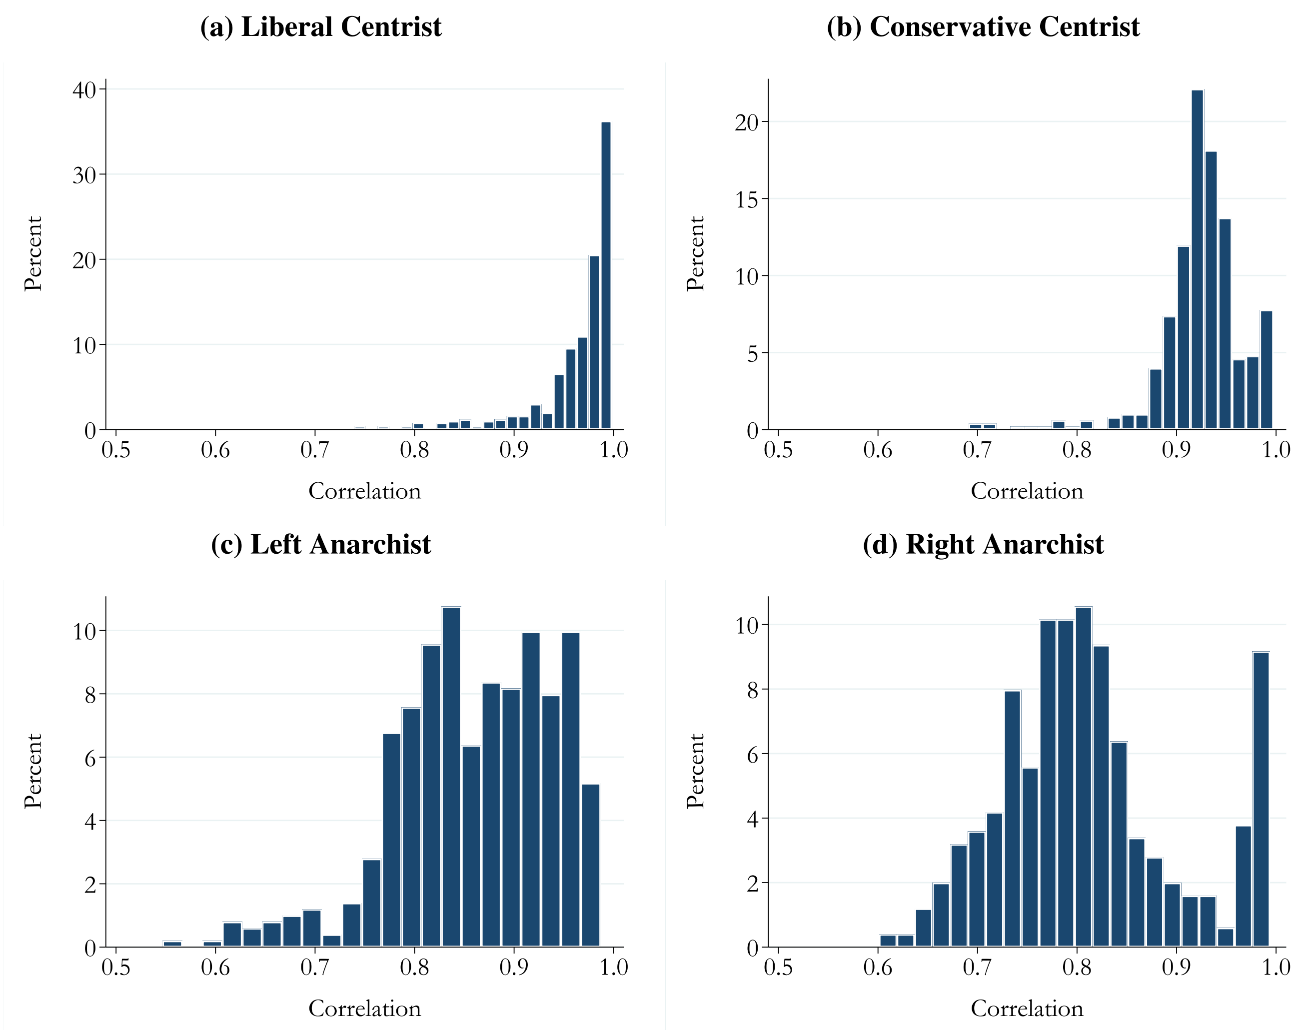

Notes: The histograms plot the correlation of the β vectors of our baseline model with topic models in which we remove all possible combination of the trust questions. In total, we fit 502 separate LDA models.

available across all 3 waves of the WVS used in this paper. As an additional robustness check, we include all these 92 questions in our topic model and create an extended type hierarchy. We then correlate the weights on the β positions between the original and extended models where they overlap.

Practically, this exercise allows us to ask whether the relative ordering of the original β issue-position weights changes as we add more features to the model. Note that this is more of an 'add them all in' rather than an iterative 'add one in' exercise. We adopt this approach both for the sake of brevity as well as to see how our original 4-type model is affected by a large, lateral addition of information. The concern would be that the addition of many extra features would fundamentally change the structure of the clusters and shift the ordering of the initial set of features.

Table E.1 reports the correlations between the β-vectors from the baseline type hierarchy and those from the extended-feature type hierarchy. Obviously, the correlation coefficients can only be calculated on the basis of the 29 original questions used in the baseline hierarchy. The correlations are very high across all the hierarchy models. Overall, we find these results to be encouraging. The same basic type structure is intact even when adding in a large amount of information. This is compatible with the idea that the extra questions/features fit in as new responses that tap into a stable set of underlying latent types.

We stress though that the exercises we present here are indicative with limited formal precedents in the LDA literature. One interesting pattern here is that the Centrist types are less sensitive to changes in features relative to the Anarchist types. This fits with the intuition that the Centrist types are well-established and better defined with the Anarchist types still being more fluid. The tendency of the Anarchist types to split as we consider higher-order models (e.g., 5, 6, and 7-type models) is also consistent with this assessment.

Table E.1: Sensitivity to Additional Features

	Type 1’	Type 2’	Type 3’	Type 4’	Type 5’
2 Type Model					
Left		0.985			
Right	0.983				
3 Type Model					
Liberal Centrist		0.947			
Conservative Centrist			0.951		
Anarchist	0.923				
4 Type Model					
Liberal Centrist		0.944			
Conservative Centrist			0.937		
Left Anarchist	0.829				
Right Anarchist				0.631	
5 Type Model					
Liberal Centrist	0.877				
Conservative Centrist			0.941		
Left Anarchist					0.800
Right Anarchist				0.970	
Market Liberal		0.987			

Notes: This table reports the correlation of the β vectors of the type hierarchy from the main paper and the type hierarchy of a topic model including all 92 consistent questions from the WVS. The prime’ notation indicates the types estimated using the 92-feature topic model. We report the highest cross-model correlations for the overlapping β weights, except for the 4-type Left Anarchist case where (in the interests of exposition) we report the three highest correlations.

F Additional Type Hierarchy Information

Table F.1: Extended Hierarchy of Types (Top Ten Features)

5 Type Model		
Liberal Centrist	**Conservative Centrist**	**Left Anarchist**
Confidence: Police	Not Justifiable: Abortion	No Confidence: Armed forces
Confidence: Justice system/courts	Not Justifiable: Euthanasia	No Confidence: Churches
Confidence: The civil services	Not Justifiable: Prostitution	No Confidence: Parliament
Justifiable: Divorce	Not Justifiable: Suicide	No Confidence: Police
Confidence: Parliament	No Problem Neighbours: People different race	No Confidence: Major companies
Proud of nationality	Confidence: Churches	No Confidence: Justice system/courts
No Problem Neighbours: People different race	No Problem Neighbours: Immigrants/foreign workers	No Confidence: Civil services
Confidence: Armed forces	Not Justifiable: Someone accepting a bribe	No Problem Neighbours: Homosexuals
Not Justifiable: Someone accepting a bribe	Confidence: Police	No Problem Neighbours: People AIDS
No Problem Neighbours: Homosexuals	Not Justifiable: Cheating on taxes	No Problem Neighbours: People different race
Right Anarchist	**Market Liberal**	
Against Neighbours: People AIDS	No Confidence: Parliament	
Against Neighbours: Homosexuals	No Confidence: Civil services	
Against Neighbours: Immigrants/foreign workers	No Confidence: The press	
Against Neighbours: Drug addicts	No Problem Neighbours: Homosexuals	
If Jobs Scarce: priority to (nation) people	No Problem Neighbours: People different race	
Not Justifiable: Homosexuality	No Problem Neighbours: People AIDS	
No Confidence: Parliament	Not Justifiable: Claiming government benefits	
Not Justifiable: Suicide	Not Justifiable: Someone accepting a bribe	
Against Neighbours: People different race	No Confidence: Labour unions	
No Confidence: Labour unions	No Problem Neighbours: Immigrants/foreign workers	
6 Type Model		
Liberal Centrist	**Conservative Centrist**	**Left Anarchist**
Confidence: The civil services	Not Justifiable: Abortion	No Confidence: Armed forces
Confidence: Parliament	Confidence: Police	Justifiable: Divorce
Confidence: Justice system/courts	Not Justifiable: Prostitution	No Confidence: Churches
Confidence: Police	Confidence: Churches	No Confidence: Major companies
No Problem Neighbours: Homosexuals	Confidence: Armed forces	Justifiable: Homosexuality
Proud of nationality	Not Justifiable: Suicide	No Problem Neighbours: Homosexuals
No Problem Neighbours: People different race	Not Justifiable: Someone accepting a bribe	Justifiable: Euthanasia
No Problem Neighbours: People AIDS	Not Justifiable: Cheating on taxes	No Problem Neighbours: People different race
No Problem Neighbours: Immigrants/foreign workers	Not Justifiable: Avoiding a fare on public transport	Justifiable: Abortion
Not Justifiable: Someone accepting a bribe	Not Justifiable: Claiming government benefits	No Problem Neighbours: Immigrants/foreign workers
Market Liberal	**Right Anarchist ('Soft')**	**Right Anarchist ('Hard')**
No Confidence: The press	No Confidence: Justice system/courts	Against Neighbours: Immigrants/foreign workers
Proud of nationality	No Confidence: Armed forces	Against Neighbours: People AIDS
No Confidence: Parliament	No Confidence: Parliament	Justifiable: Avoiding a fare on public transport
Confidence: Armed forces	No Confidence: Civil services	Against Neighbours: People different race
Confidence: Police	No Confidence: Police	Justifiable: Cheating on taxes
Not Justifiable: Someone accepting a bribe	No Confidence: Labour unions	Against Neighbours: Homosexuals
Not Justifiable: Claiming government benefits	Not Justifiable: Suicide	If Jobs Scarce: priority to (nation) people
No Confidence: Labour unions	No Confidence: The press	Justifiable: Claiming government benefits
No Problem Neighbours: Homosexuals	No Confidence: Major companies	Against Neighbours: Drug addicts
No Problem Neighbours: People different race	Not Justifiable: Someone accepting a bribe	Justifiable: Euthanasia

continues on next page

Table F.1: Continued from Previous Page

7 Type Model		
Liberal Centrist	**Conservative Centrist**	**Left Anarchist**
Confidence: The civil services	Confidence: Police	No Confidence: Armed forces
Confidence: Parliament	Not Justifiable: Abortion	No Confidence: Churches
Confidence: Justice system/courts	Confidence: Churches	Justifiable: Divorce
Confidence: Police	Not Justifiable: Euthanasia	No Confidence: Major companies
Proud of nationality	Confidence: Armed forces	Justifiable: Homosexuality
No Problem Neighbours: People different race	Not Justifiable: Suicide	No Confidence: Parliament
Not Justifiable: Someone accepting a bribe	Not Justifiable: Prostitution	No Problem Neighbours: Homosexuals
No Problem Neighbours: Homosexuals	No Problem Neighbours: People different race	Justifiable: Euthanasia
No Problem Neighbours: Immigrants/foreign workers	Not Justifiable: Cheating on taxes	No Problem Neighbours: People AIDS
Justifiable: Divorce	Not Justifiable: Someone accepting a bribe	No Problem Neighbours: People different race
Market Liberal	**Right Anarchist (‘Soft’)**	**Right Anarchist (‘Hard’)**
No Confidence: Parliament	No Confidence: Parliament	Against Neighbours: People AIDS
No Confidence: The press	No Confidence: Civil services	Against Neighbours: Homosexuals
No Problem Neighbours: People different race	No Confidence: Justice system/courts	Against Neighbours: Immigrants/foreign workers
No Problem Neighbours: Homosexuals	No Confidence: Armed forces	Against Neighbours: Drug addicts
Confidence: Police	Not Justifiable: Suicide	If Jobs Scarce: priority to (nation) people
Proud of nationality	No Confidence: Major companies	Against Neighbours: People different race
Confidence: Armed forces	No Confidence: Labour unions	Not Justifiable: Homosexuality
No Problem Neighbours: People AIDS	No Confidence: The press	Proud of nationality
Not Justifiable: Someone accepting a bribe	No Problem Neighbours: People different race	Confidence: Armed forces
No Confidence: Labour unions	Not Justifiable: Prostitution	Not Justifiable: Someone accepting a bribe
Super Anarchist (’Rage Against the Machine’)		
Justifiable: Avoiding a fare on public transport		
Justifiable: Cheating on taxes		
Justifiable: Claiming government benefits		
Justifiable: Accepting a bribe		
Justifiable: Euthanasia		
If Jobs Scarce: priority to (nation) people		
Proud of nationality		
Justifiable: Prostitution		
Justifiable: Divorce		
No Problem Neighbours: People different race		

Notes: This table reports the 10 most important features for a n-type LDA model, where $n \in \{5, 6, 7\}$. The types are labelled on the basis of their β-weight correlation with types in the previous level. For example, the 6-type model Liberal Centrist has a 0.96 correlation with the 5-type model Liberal Centrist.

Table F.2: Type Hierarchy – All WVS Waves Pooled

2 Type Model	3 Type Model	4 Type Model
Left	**Liberal Centrist**	**Liberal Centrist**
No Problem Neighbours: Homosexuals	Confidence: Police	Confidence: Police
No Problem Neighbours: People different race	Confidence: Justice system/courts	No Problem Neighbours: Homosexuals
No Problem Neighbours: People AIDS	No Problem Neighbours: Homosexuals	No Problem Neighbours: People AIDS
No Problem Neighbours: Immigrants/foreign workers	No Problem Neighbours: People different race	No Problem Neighbours: People different race
Justifiable: Divorce	No Problem Neighbours: People AIDS	No Problem Neighbours: Immigrants/foreign workers
Not Justifiable: Someone accepting a bribe	No Problem Neighbours: Immigrants/foreign workers	Justifiable: Divorce
Justifiable: Euthanasia	Proud of nationality	Not Justifiable: Someone accepting a bribe
Justifiable: Homosexuality	Not Justifiable: Someone accepting a bribe	Proud of nationality
Not Justifiable: Claiming government benefits	Confidence: The civil services	Not Justifiable: Claiming government benefits
Proud of nationality	Not Justifiable: Cheating on taxes	Not Justifiable: Cheating on taxes
Right	**Conservative Centrist**	**Conservative Centrist**
Not Justifiable: Someone accepting a bribe	Not Justifiable: Homosexuality	Confidence: Police
Not Justifiable: Suicide	Not Justifiable: Abortion	Confidence: Churches
Proud of nationality	Not Justifiable: Suicide	Not Justifiable: Suicide
Not Justifiable: Prostitution	Not Justifiable: Prostitution	Proud of nationality
Not Justifiable: Avoiding a fare on public transport	Proud of nationality	Confidence: Justice system/courts
Not Justifiable: Claiming government benefits	Not Justifiable: Someone accepting a bribe	Not Justifiable: Prostitution
Not Justifiable: Cheating on taxes	Not Justifiable: Avoiding a fare on public transport	Confidence: The civil services
Not Justifiable: Abortion	Not Justifiable: Claiming government benefits	Not Justifiable: Someone accepting a bribe
Not Justifiable: Homosexuality	Not Justifiable: Cheating on taxes	Not Justifiable: Abortion
No Problem Neighbours: People different race	Not Justifiable: Euthanasia	Confidence: Armed forces
	Anarchist	**Left Anarchist**
	No Confidence: Civil services	Justifiable: Divorce
	No Confidence: Parliament	No Confidence: Churches
	No Confidence: Churches	No Confidence: Armed forces
	No Confidence: Justice system/courts	No Problem Neighbours: Homosexuals
	No Problem Neighbours: Homosexuals	No Confidence: Parliament
	No Confidence: Armed forces	No Confidence: Civil services
	No Confidence: Major companies	No Problem Neighbours: People different race
	No Problem Neighbours: People different race	No Problem Neighbours: People AIDS
	No Problem Neighbours: People AIDS	No Problem Neighbours: Immigrants/foreign workers
	No Confidence: The press	Justifiable: Euthanasia
		Right Anarchist
		No Confidence: Parliament
		No Confidence: Civil services
		No Confidence: Labour unions
		No Confidence: The press
		No Confidence: Justice system/courts
		Not Justifiable: Someone accepting a bribe
		Not Justifiable: Suicide
		Not Justifiable: Avoiding a fare on public transport
		Not Justifiable: Claiming government benefits
		Proud of nationality

Notes: This table reports the 10 most important features based on the β vectors for a n-type LDA model fitted to all waves of the WVS pooled together, where $n \in \{2, 3, 4\}$.

Table F.3: Type Hierarchy – No Imputation

2 Type Model	3 Type Model	4 Type Model
Left	**Liberal Centrist**	**Liberal Centrist**
No Problem Neighbours: Homosexuals	Confidence: Police	Confidence: Police
No Problem Neighbours: People different race	Confidence: Justice system/courts	No Problem Neighbours: Homosexuals
No Problem Neighbours: People AIDS	No Problem Neighbours: Homosexuals	No Problem Neighbours: People AIDS
No Problem Neighbours: Immigrants/foreign workers	No Problem Neighbours: People different race	No Problem Neighbours: People different race
Justifiable: Divorce	No Problem Neighbours: Immigrants/foreign workers	Proud of nationality
Not Justifiable: Someone accepting a bribe	No Problem Neighbours: People AIDS	No Problem Neighbours: Immigrants/foreign workers
Justifiable: Homosexuality	Proud of nationality	Not Justifiable: Someone accepting a bribe
Proud of nationality	Not Justifiable: Someone accepting a bribe	Not Justifiable: Claiming government benefits
Justifiable: Euthanasia	Not Justifiable: Cheating on taxes	Not Justifiable: Cheating on taxes
Not Justifiable: Claiming government benefits	Confidence: The civil services	Justifiable: Divorce
Right	**Conservative Centrist**	**Conservative Centrist**
Not Justifiable: Someone accepting a bribe	Not Justifiable: Abortion	Confidence: Police
Proud of nationality	Not Justifiable: Prostitution	Confidence: Churches
Not Justifiable: Suicide	Not Justifiable: Suicide	Confidence: Armed forces
Not Justifiable: Cheating on taxes	Proud of nationality	Not Justifiable: Prostitution
Not Justifiable: Prostitution	Not Justifiable: Someone accepting a bribe	Not Justifiable: Abortion
Not Justifiable: Avoiding a fare on public transport	Not Justifiable: Cheating on taxes	Not Justifiable: Suicide
Not Justifiable: Claiming government benefits	Not Justifiable: Avoiding a fare on public transport	Proud of nationality
Not Justifiable: Abortion	Not Justifiable: Homosexuality	Not Justifiable: Cheating on taxes
No Problem Neighbours: People different race	Not Justifiable: Claiming government benefits	Not Justifiable: Someone accepting a bribe
Confidence: Police	Not Justifiable: Euthanasia	Not Justifiable: Euthanasia
	Anarchist	**Left Anarchist**
	No Confidence: Civil services	Justifiable: Divorce
	No Confidence: Parliament	Justifiable: Homosexuality
	No Confidence: Churches	Justifiable: Abortion
	No Confidence: Major companies	Justifiable: Euthanasia
	No Confidence: Justice system/courts	State ownership better than private ownership
	No Problem Neighbours: Homosexuals	No Problem Neighbours: Homosexuals
	No Confidence: The press	No Problem Neighbours: People different race
	No Problem Neighbours: People different race	No Problem Neighbours: Immigrants/foreign workers
	No Problem Neighbours: People AIDS	No Confidence: Churches
	No Problem Neighbours: Immigrants/foreign workers	No Problem Neighbours: People AIDS
		Right Anarchist
		No Confidence: Parliament
		No Confidence: Civil services
		No Confidence: Justice system/courts
		No Confidence: Labour unions
		No Confidence: The press
		No Confidence: Major companies
		Not Justifiable: Someone accepting a bribe
		Not Justifiable: Claiming government benefits
		No Confidence: Police
		No Problem Neighbours: People different race

Notes: This table reports the 10 most important features based on the β vectors for a n-type LDA model fitted to the subset of observation in the WVS that do not require imputation of missing values, where $n \in \{2, 3, 4\}$.

Table F.4: Issues of Increasing Importance between Wave 2 and Wave 5

Question	Baseline	Change
Liberal Centrist		
More responsibility for government	0.004	0.222
Confidence: Armed forces	0.430	0.199
State ownership better than private ownership	0.000	0.164
Justifiable: Homosexuality	0.477	0.148
Confidence: The civil services	0.425	0.144
Confidence: Labour unions	0.324	0.136
Against Neighbours: Drug addicts	0.456	0.130
No Confidence: Major companies	0.285	0.091
Not Justifiable: Prostitution	0.323	0.087
Confidence: Churches	0.328	0.075
Conservative Centrist		
No Problem Neighbours: Homosexuals	0.402	0.133
No Confidence: Parliament	0.009	0.122
No Problem Neighbours: People AIDS	0.445	0.095
No Confidence: Major companies	0.150	0.088
Justifiable: Homosexuality	0.000	0.072
Not Justifiable: Abortion	0.595	0.065
Competition is harmful	0.206	0.053
No Problem Neighbours: Drug addicts	0.278	0.052
Confidence: Armed forces	0.618	0.045
No Confidence: The press	0.301	0.044
Left Anarchist		
Confidence: Police	0.024	0.204
Not Justifiable: Cheating on taxes	0.277	0.144
Proud of nationality	0.365	0.103
Competition is harmful	0.444	0.098
Confidence: Armed forces	0.002	0.090
Justifiable: Homosexuality	0.570	0.077
Justifiable: suicide	0.372	0.076
Justifiable: Prostitution	0.463	0.076
No Confidence: The press	0.487	0.064
If Jobs Scarce: no priority to (nation) people	0.391	0.053
Right Anarchist		
Confidence: Police	0.000	0.319
Confidence: Armed forces	0.169	0.225
Justifiable: Divorce	0.045	0.150
No Problem Neighbours: Homosexuals	0.372	0.149
Justifiable: Homosexuality	0.000	0.147
Justifiable: Euthanasia	0.058	0.135
No Confidence: Churches	0.378	0.094
Against Neighbours: Immigrants/foreign workers	0.225	0.086
No Problem Neighbours: People AIDS	0.431	0.077
More responsibility for government	0.258	0.062

Notes: This table reports the 10 features of each type which show the biggest increase in weight from wave 2 to wave 5. Column 2 reports the baseline value in wave 2 and column 3 reports the change from wave 2 to wave 5.

Table F.5: Issue Position Differences between Types (4 Type Model)

	Difference Lib. Centrist	Difference Cons. Centrist	Difference Left Anarchist	Difference Right Anarchist
L. Centrist		Justifiable: Divorce Justifiable: Euthanasia Justifiable: Homosexuality Justifiable: Abortion No Confidence: Churches No Problem Neighbours: Homosexuals If Jobs Scarce: no priority to natives No Problem Neighbours: People AIDS Justifiable: Prostitution Competition is good	Confidence: Justice system/courts Confidence: Police Confidence: Armed forces Confidence: The civil services Confidence: Parliament Competition is good Against Neighbours: Drug addicts Confidence: Major Companies More responsibility for people Not Justifiable: avoiding fare on pub. trans.	Confidence: Justice system/courts Justifiable: Divorce Confidence: The civil services Justifiable: Abortion Justifiable: Homosexuality Confidence: Police Justifiable: Euthanasia Confidence: Parliament Confidence: Major Companies Confidence: Labour unions
C. Centrist	Not Justifiable: Abortion Not Justifiable: Euthanasia Not Justifiable: Homosexuality Not Justifiable: Divorce Confidence: Churches Not Justifiable: Prostitution Not Justifiable: Suicide If Jobs Scarce: priority to natives Against Neighbours: People AIDS Against Neighbours: Homosexuals		Confidence: Churches Not Justifiable: Abortion Not Justifiable: Prostitution Not Justifiable: Euthanasia Confidence: Armed forces Confidence: The civil services Confidence: Justice system/courts Confidence: Police Not Justifiable: Suicide Confidence: Parliament	Confidence: Justice system/courts Confidence: The civil services Confidence: Parliament Confidence: Major Companies Confidence: Police Confidence: Churches Confidence: Labour unions Confidence: Press Confidence: Armed forces Not Justifiable: Euthanasia
L. Anarchist	No Confidence: Armed forces No Confidence: Justice system/courts No Confidence: Civil services Competition is harmful No Confidence: Parliament No Confidence: Police No Confidence: Churches No Problem Neighbours: Drug addicts Justifiable: avoiding fare on pub. trans. More responsibility for government	No Confidence: Churches Justifiable: Divorce No Confidence: Armed forces Justifiable: Euthanasia No Confidence: Civil services Justifiable: Homosexuality Justifiable: Abortion No Confidence: Parliament No Confidence: Justice system/courts Justifiable: Prostitution		Justifiable: Abortion Justifiable: Divorce Justifiable: Homosexuality Justifiable: Euthanasia Justifiable: Prostitution Competition is harmful No Problem Neighbours: Drug addicts Justifiable: suicide Justifiable: avoiding fare on pub. trans. No Confidence: Churches
R. Anarchist	No Confidence: Justice system/courts No Confidence: Civil services No Confidence: Parliament Not Justifiable: Abortion No Confidence: Labour unions No Confidence: The press No Confidence: Major companies Not Justifiable: Homosexuality No Confidence: Police Not Justifiable: Euthanasia	No Confidence: Civil services No Confidence: Justice system/courts No Confidence: Parliament No Confidence: Major companies No Confidence: Labour unions No Confidence: The press No Confidence: Churches No Confidence: Police No Confidence: Armed forces Against Neighbours: Drug addicts	Not Justifiable: Prostitution Not Justifiable: Suicide Not Justifiable: Abortion Against Neighbours: Drug addicts If Jobs Scarce: priority to natives Not Justifiable: Homosexuality Not Justifiable: avoiding fare on pub. trans. Competition is good Not Justifiable: Euthanasia No Confidence: Labour unions	

Notes: This table reports the 10 features for which there exist the largest differences between the 4 ideological types created by LDA. The model is fitted to the 5th wave of the sample. The type labels are chosen by the author.

G Cross-Check of Results with European Social Survey

The European Social Survey (ESS) is a biannual survey of 37 European countries covering the years from 2002 until 2016. For our replication exercise, we use all countries in the ESS that also appear in the WVS and were used in our main analysis. Overall, 13 of our original 17 countries also appear in the ESS (Germany, Great Britain, France, Denmark, Spain, Finland, Portugal, Austria, Belgium, Italy, Ireland, Netherlands, and Iceland). Similarly, we create a subsample of the ESS waves that aligns with the waves of the WVS. We use ESS rounds 1 to 5 (2002 - 2010) which are comparable to the 4th and 5th wave of the WVS (Wave 4: 1999-2004 and Wave 5: 2005-2009). We further select a set of questions from the ESS that cover similar issues to those we used from the WVS. Table G.1 provides an overview of the ESS questions as well as their scale. Identical to the main results of the paper, we recode questions into 2 binary features indicating support and opposition to issues.

As the next step, we fit LDA models with an increasing number of types to the ESS data to produce a type hierarchy. Identical to the results in Table 1, we report the 'top ten' features for each ESS type in Table G.2. It should be apparent that since we use a different set of questions the ESS types can never be identical to the WVS types. What is important for our purpose is that the resulting types recover a similar ideological spectrum.

For the basic 2-type model in the first column, the two types are distinguished by their trust in institutions. While the first type, which we label as 'Centrist' trusts the police, the legal system and is satisfied with the democracy in the country, the second type (labelled as 'Anarchist') does not trust politicians and political parties and is unsatisfied with the national government. Interestingly, this shows that in the ESS data, the 'Anarchist' type already arises in the 2-type model.

The second column reports the top features for the 3-type model. The 'Centrist' type remains more or less unchanged, but we observe a split of the 'Anarchist' type along immigration issues. On the one hand, the 'Left Anarchist' supports immigration and gay rights. Moreover, this type considers it important to take care of people and treat them equally. The 'Right Anarchist' on the other hand opposes immigration and puts a larger weight on security and safety.

In the third column, we show the top features for the 4-type model. In the 4-type model, a new split between two 'Centrist' types emerges. One important difference between the two 'Centrist' types is the importance of religions, traditions and customs. Further, the two types differ based on the importance they attribute to safety, but both profess trust in the legal system.

Overall, the type structure that emerges from the ESS is reasonably similar to the types that emerge in the WVS. We again find that types split apart based on their trust in institutions. This allows us to label types as ‘Centrist’ and ‘Anarchist’. Additionally, we observe type characteristics that are broadly in line with the left-right spectrum. For example, one of the important dividing issues is immigration. The social issues which define the types differ across the two datasets but this is mainly a result of the differences in the question set. The ESS simply does not contain questions concerning support and opposition to abortion and suicide, and neither does the WVS contain a detailed set of questions concerning immigration. We therefore view this exercise as useful corroboration that our core finding of ideological types that are differentiated by trust in institutions holds across independent datasets.

Table G.1: Selected Question from the ESS

ESS Variable Code	Question	Scale
ppltrst	Most people can be trusted or you can't be too careful	
trstprl	Trust in country's parliament	
trstlgl	Trust in the legal system	
trstplc	Trust in the police	
trstplt	Trust in politicians	
trstprt	Trust in political parties	
trstep	Trust in the European Parliament	
trstun	Trust in the United Nations	
stflife	How satisfied with life as a whole	
stfeco	How satisfied with present state of economy in country	
stfgov	How satisfied with the national government	
stfdem	How satisfied with the way democracy works in country	
stfedu	State of education in country nowadays	
stfhlth	State of health services in country nowadays	
gincdif	Government should reduce differences in income levels	
freehms	Gays and lesbians free to live life as they wish	
imsmetn	Allow many/few immigrants of same race/ethnic group as majority	
imdfetn	Allow many/few immigrants of different race/ethnic group from majority	
impcntr	Allow many/few immigrants from poorer countries outside Europe	
imbgeco	Immigration bad or good for country's economy	
imueclt	Country's cultural life undermined or enriched by immigrants	
imwbcnt	Immigrants make country worse or better place to live	
rlgdgr	How religious are you	
rlgatnd	How often attend religious services apart from special occasions	
pray	How often pray apart from at religious services	
ipeqopt	Important that people are treated equally and have equal opportunities	
impsafe	Important to live in secure and safe surroundings	
ipfrule	Important to do what is told and follow rules	
ipudrst	Important to understand different people	
ipgdtim	Important to have a good time	
impfree	Important to make own decisions and be free	
iphlppl	Important to help people and care for others well-being	
ipstrgv	Important that government is strong and ensures safety	
ipbhprp	Important to behave properly	
iprspot	Important to get respect from others	
iplylfr	Important to be loyal to friends and devote to people close	
impenv	Important to care for nature and environment	
imptrad	Important to follow traditions and customs	

Notes: This table reports the questions selected from the European Social Survey.

Table G.2: Type Hierarchy as Created with ESS Data

2 Type Model	3 Type Model	4 Type Model
Centrist	**Centrist**	**Conservative Centrist**
Satisfied with life as a whole	Satisfied: Democracy in country	Satisfied: Democracy in country
Trust in the police	Trust in the police	Trust in the police
Important: To be loyal to friends	Satisfied with life as a whole	Important to behave properly
Satisfied: Democracy in country	Trust in the legal system	Satisfied with life as a whole
Important: That people are treated equally	Trust in country's parliament	Important to follow traditions and customs
Important: To help people and care for others	Satisfied with state of education	Trust in the legal system
Important: To care for environment	Important: To be loyal to friends	Important: To help people and care for others
Important to understand different people	Trust in the United Nations	Important: To care for environment
Trust in the legal system	Satisfied with state of health services	Important: To be loyal to friends
Important to make own decisions and be free	Important: To care for environment	Important: That government ensures safety
Anarchist	**Left Anarchist**	**Left Anarchist**
No Trust in politicians	Allow immigrants of same race/ethnic group	Allow immigrants of same race/ethnic group
No Trust in political parties	Allow immigrants of different race/ethnic group	No Trust in politicians
Important: To be loyal to friends	Allow immigrants from poorer countries	Allow immigrants of different race/ethnic group
Important: That people are treated equally	No Trust in politicians	Allow immigrants from poorer countries
Not satisfied with the national government	Important: That people are treated equally	No Trust in political parties
Important: To help people and care for others	Important: To be loyal to friends	Important: That people are treated equally
Important: To care for environment	Gays and lesbians free to live life as they wish	Important: To help people and care for others
Important to make own decisions and be free	Important: To help people and care for others	Important: To be loyal to friends
Important to live in secure and safe surroundings	Important to understand different people	Important: To care for environment
Important: That government ensures safety	Important to make own decisions and be free	Important to understand different people
	Right Anarchist	**Right Anarchist**
	Not Allow immigrants of different race/ethnic group	Not Allow immigrants of different race/ethnic group
	Not Allow immigrants from poorer countries	Not Allow immigrants from poorer countries
	Immigrants make country worse place to live	Not Allow immigrants of same race/ethnic group
	Not Allow immigrants of same race/ethnic group	Immigrants make country worse place to live
	Important: To be loyal to friends	No Trust in politicians
	No Trust in politicians	Immigration bad for country's economy
	Important to live in secure and safe surroundings	No Trust in political parties
	Important: That government ensures safety	Important: To be loyal to friends
	Immigration bad for country's economy	Important to live in secure and safe surroundings
	No Trust in political parties	Important: That government ensures safety
		Liberal Centrist
		Not Important to follow traditions and customs
		Not Important to do what is told and follow rules
		Not often pray apart from at religious services
		Not often attend religious services
		Not Important to behave properly
		Not religious
		Gays and lesbians free to live life as they wish
		Not Important to live in secure and safe surroundings
		Not Important: That government ensures safety
		Trust in the legal system

Notes: This table reports the 10 most important features for a n-type LDA model fit to the European Social Survey, where $n \in \{2, 3, 4, \}$.

H Robustness 7th Wave of the WVS

Our main analysis is based on the 2nd, 3rd and 5th wave of the World Value Survey (WVS) and European Value Study (EVS). The combination of these two surveys significantly increases the number of European countries that are covered in our data. For this reason, we excluded the 4th and 6th wave of the WVS as there are no corresponding waves of the EVS. For the 7th wave of the WVS there is a corresponding wave of the EVS. We can extend our analysis past 2010 by making use of the 7th wave of the WVS, but in this process, the sample of available countries shrinks by Belgium, Canada, Ireland, Malta, Northern Ireland. Due to the changing sample of countries we decided to relegate the analysis of the 7th wave to the appendix.

Fitting a 4-type LDA topic model to the 7th wave of the WVS leads to very similar types. We again observe a liberal centrist and conservative centrist type in the data (correlations of 0.98 and 0.97 respectively with their wave 5 equivalent). Also, the right anarchist and left anarchist type emerge in the LDA model (correlation of 0.92 and 0.93 with wave 5 equivalent).

We fit a separate LDA model to generate new individual-level type shares for wave 7. Based on the resulting type shares, we analyse if any further changes in the type composition occurred in wave 7. In particular, we reproduce Figure 6 based on the countries that are available in all 4 waves. The results are presented in Figure H.1.

Overall, our findings are similar when we include the 7th wave. The right anarchist types stabilised at a high level in the US, while we observe a further increase in the share of the left anarchist type. For the other countries in the sample, we do not observe any major shifts in the prevalence of the anarchist types. If anything the right anarchist type shares appear to decrease slightly. This decrease is offset by a slight increase in the left anarchist type share.

We also repeat our polarisation analysis for the countries that remain in our sample. Again the US appears as the most polarised country in our sample, while the Nordic countries (Denmark, Iceland, Finland) exhibit the lowest level of polarisation.

Figure H.1: Type Shares - US vs non-US (Wave 7)

(a) Left Anarchist Types

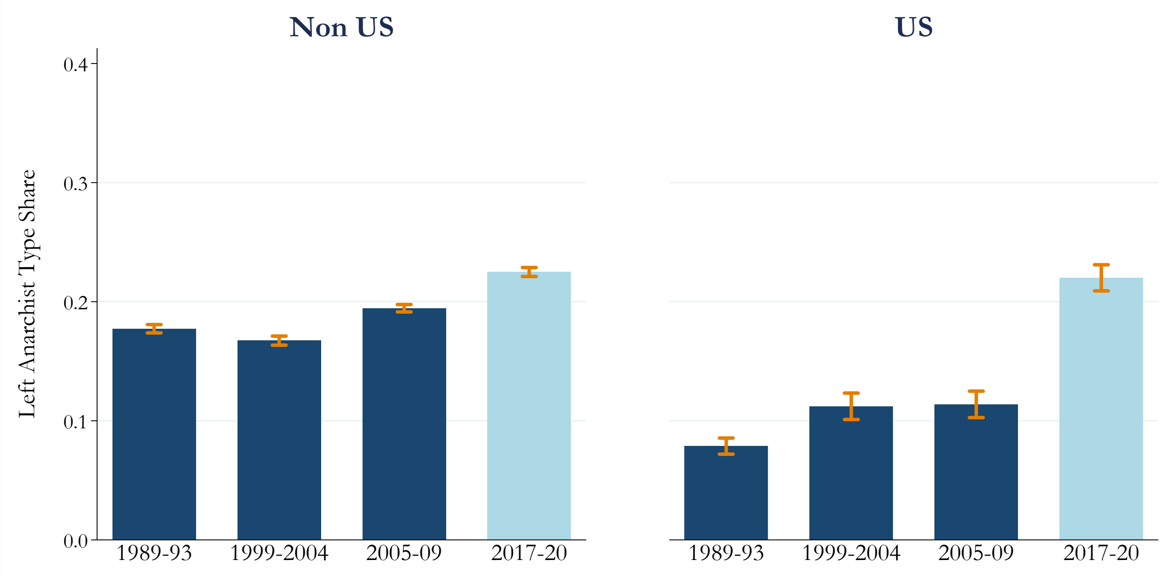

(b) Right Anarchist Types

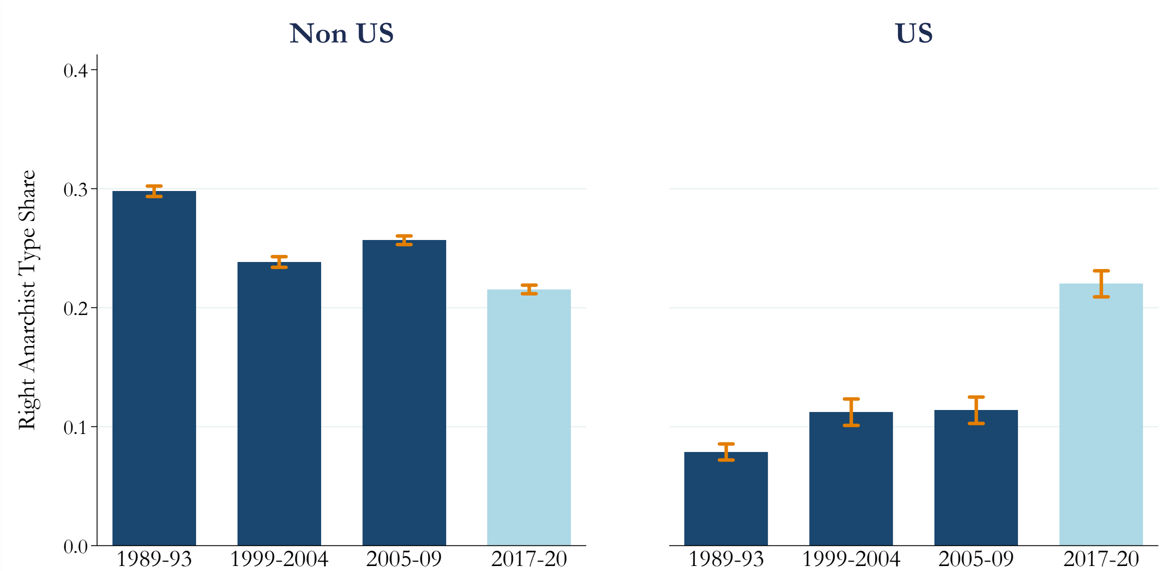

Notes: This figure compares the levels of θ type shares across waves for the Left Anarchist and Right Anarchist types. We pool all non-US countries and contrast them to the US. The pooling for the non-US sample is based on WVS sample weights. The timing of the waves is Wave 2 (1989-1993), Wave 4 (1999-2004), Wave 5 (2005-2009) and Wave 7 (2017-2020). 95% confidence intervals are reported in orange.

Figure H.2: Polarisation 7th Wave

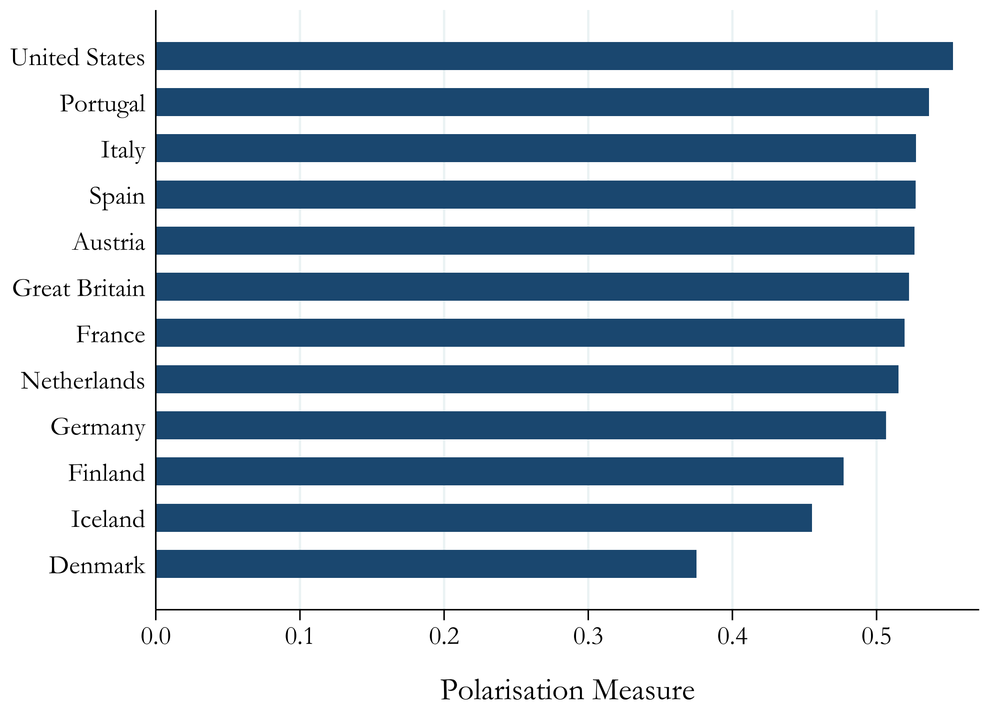

Notes: The figure shows the country-level polarisation measures from Waves 7 (2017-2020) calculated following Esteban and Ray (1994).

I Additional Details on Populist Parties

A list of European parties that can be classified as populists in 2019 was prepared by Rooduijn et al. (2019). Their classification is based on the following definition:

> "Populist parties: parties that endorse the set of ideas that society is ultimately separated into two homogeneous and antagonistic groups, "the pure people" versus "the corrupt elite," and which argues that politics should be an expression of the volonté générale (general will) of the people (Mudde, 2004)."

As the list does not contain any information for parties outside of Europe, we further code the Reform Party in the US as a populist party based on the (see http://www.reformparty.org/). Lastly also the NDP in Canada is classified as populist as it exhibited populist tendencies during our observation period (see https://www.thecanadianencyclopedia.ca/en/article/populism).

To achieve a consistent coding of parties across waves, we also classify predecessor parties as populist. For example, the German party "Die Linke" is listed in Rooduijn et al. (2019). Hence, we also code the party "Partei des demokratischen Sozialismus" as populist.

Table I.1: List of Populist Parties

Country	Party
Austria	FPÖ
Austria	Alliance for the Future of Austria
Austria	Dr. Martin's List - For Democracy
Belgium	Front National
Belgium	Vlaams Blok
Belgium	Vlaams Belang
Canada	NDP
Denmark	Danish People Party
Denmark	Progress Party
Finland	True Finns
France	Front National
France	Le Front National de Jean-Marie le Pen
France	Le Front National de Bruno Megret
Germany	Partei des demokratischen Sozialismus
Iceland	Citizen Movement
Ireland	Sinn Fein
Italy	Forza Italia
Italy	Northern League
Netherland	Party for Freedom
Netherland	Socialistische Partij
United Kingdom	UK Independence Part
United Kingdom	Sinn Fein
United States	Reform Party

Notes: This table reports the parties that were coded as populist based on the information from Rooduijn et al. (2019)

Table I.2: Support for Populist Parties

Panel A: All Countries

	(1)	(2)	(3)	(4)
Conservative Centrist	-0.007**		-0.007**	-0.008***
	(0.003)		(0.003)	(0.003)
Left Anarchist	0.037***		0.031***	0.033***
	(0.004)		(0.004)	(0.004)
Right Anarchist	0.033***		0.033***	0.032***
	(0.004)		(0.004)	(0.004)
Distance from Centre		0.008***	0.008***	
		(0.001)	(0.001)	
I[Far Left]				0.024***
				(0.004)
I[Far Right]				0.028***
				(0.004)
Country FE	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Observations	67,666	67,757	67,666	67,666
Mean of DV	0.042	0.042	0.042	0.042
R^2	0.05	0.05	0.05	0.05

Panel B: All Countries except USA

	(1)	(2)	(3)	(4)
Conservative Centrist	-0.007**		-0.007**	-0.008**
	(0.003)		(0.003)	(0.003)
Left Anarchist	0.037***		0.032***	0.033***
	(0.005)		(0.005)	(0.005)
Right Anarchist	0.036***		0.036***	0.035***
	(0.004)		(0.004)	(0.004)
Distance from Centre		0.009***	0.009***	
		(0.001)	(0.001)	
I[Far Left]				0.025***
				(0.004)
I[Far Right]				0.031***
				(0.004)
Country FE	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Observations	63,825	63,915	63,825	63,825
Mean of DV	0.044	0.044	0.044	0.044
R^2	0.05	0.05	0.05	0.05

Notes: Each column reports the regression results for individual-level regression. The dependent variable is an indicator variable for the support of a populist party. Robust standard errors are used. Significance levels: *** p<0.01, ** p<0.05, and * p<0.1. The data come from the World Value Survey and the European Value Survey. Note that the 67,666 sample presented here is smaller than our main 81,141 sample due to missing values and non-responses for the populist voting and left-right scale questions.

Table I.3: Ideological Types and Support for Party Families

(a) Manifesto Project

	(1) Ecological	(2) Socialist	(3) Social Democratic	(4) Liberal	(5) Christian Democrat	(6) Conservative	(7) Nationalist	(8) Agrarian	(9) Ethnic	(10) Special Issue
Conservative Centrist	-0.054***	-0.016***	-0.044***	-0.059***	0.135***	-0.001	0.025***	0.022***	-0.008***	0.000
	(0.006)	(0.005)	(0.011)	(0.006)	(0.009)	(0.007)	(0.005)	(0.003)	(0.002)	(0.002)
Left Anarchist	0.155***	0.143***	0.092***	-0.080***	-0.192***	-0.146***	-0.000	-0.005*	0.017***	0.016***
	(0.010)	(0.009)	(0.014)	(0.008)	(0.010)	(0.008)	(0.006)	(0.003)	(0.004)	(0.003)
Right Anarchist	-0.038***	0.004	-0.066***	-0.045***	0.077***	-0.031***	0.073***	0.013***	0.003	0.009***
	(0.006)	(0.005)	(0.012)	(0.007)	(0.010)	(0.007)	(0.006)	(0.003)	(0.003)	(0.003)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,258	32,258	32,258	32,258	32,258	32,258	32,258	32,258	32,258	32,258
Mean of DV	0.076	0.065	0.318	0.101	0.180	0.154	0.066	0.019	0.011	0.010
R^2	0.09	0.13	0.08	0.19	0.24	0.33	0.23	0.19	0.05	0.05

(b) Chapel Hill Expert Survey

	(1) Radical Right	(2) Conservative	(3) Liberal	(4) Christian Democratic	(5) Socialist	(6) Radical Left	(7) Green	(8) Regionalist	(9) No Family	(10) Confessional	(11) Agrarian
Conservative Centrist	0.010*	0.011	-0.078***	0.142***	-0.047***	-0.014***	-0.052***	-0.005*	-0.002***	0.018***	0.018***
	(0.005)	(0.007)	(0.008)	(0.009)	(0.013)	(0.005)	(0.006)	(0.003)	(0.001)	(0.003)	(0.003)
Left Anarchist	0.010*	-0.109***	-0.096***	-0.190***	0.078***	0.144***	0.155***	0.016***	0.001	0.000	-0.011***
	(0.006)	(0.007)	(0.010)	(0.010)	(0.015)	(0.009)	(0.011)	(0.004)	(0.001)	(0.001)	(0.003)
Right Anarchist	0.084***	-0.006	-0.058***	0.052***	-0.079***	0.007	-0.037***	0.011***	0.001	0.012***	0.011***
	(0.006)	(0.008)	(0.008)	(0.010)	(0.014)	(0.006)	(0.007)	(0.004)	(0.001)	(0.002)	(0.003)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	26,724	26,724	26,724	26,724	26,724	26,724	26,724	26,724	26,724	26,724	26,724
Mean of DV	0.044	0.131	0.115	0.174	0.364	0.060	0.078	0.017	0.001	0.003	0.013
R^2	0.10	0.33	0.15	0.28	0.05	0.11	0.12	0.05	0.02	0.03	0.28

Notes: Each column reports the regression results for individual-level regression. The dependent variable is an indicator variable for the support of a party of the party family indicated at the top. In panel (a), party families are defined based on the classification of the Manifesto Project. Panel (b) uses the party classification from the Chapel Hill Expert Survey. Robust standard errors are used. Significance levels: *** p<0.01, ** p<0.05, and * p<0.1.

J Additional Results Citizen Slant

Table J.1: 'Citizen Slant' - US vs non-US Comparison

Panel A: United States				
	(1) Liberal Centrist	(2) Conservative Centrist	(3) Left Anarchist	(4) Right Anarchist
Wave 4	0.036***	0.008	0.037	0.046***
	(0.009)	(0.010)	(0.024)	(0.012)
Wave 5	0.006	0.040***	0.098***	0.095***
	(0.012)	(0.012)	(0.023)	(0.011)
Controls	Yes	Yes	Yes	Yes
Observations	1,412	1,408	267	1,110
Mean of DV	0.778	0.762	0.724	0.739
R^2	0.02	0.02	0.09	0.09

Panel B: Non United States				
	(1) Liberal Centrist	(2) Conservative Centrist	(3) Left Anarchist	(4) Right Anarchist
Wave 4	0.045***	-0.007**	0.014***	0.006**
	(0.003)	(0.003)	(0.004)	(0.003)
Wave 5	0.029***	-0.013***	0.004	0.001
	(0.003)	(0.003)	(0.004)	(0.003)
Country FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Observations	24,339	21,326	10,766	20,513
Mean of DV	0.764	0.767	0.726	0.740
R^2	0.06	0.04	0.03	0.02

Panel C: All Countries				
	(1) Liberal Centrist	(2) Conservative Centrist	(3) Left Anarchist	(4) Right Anarchist
Wave 4	0.044***	-0.005*	0.015***	0.009***
	(0.003)	(0.003)	(0.004)	(0.003)
Wave 5	0.028***	-0.010***	0.006*	0.007***
	(0.003)	(0.003)	(0.004)	(0.003)
Country FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Observations	25,751	22,734	11,033	21,623
Mean of DV	0.765	0.767	0.726	0.740
R^2	0.05	0.04	0.03	0.02

Notes: Each column reports the regression results for individual-level regression. The dependent variable is the Gini Coefficient of the individual type shares as a measure of polarisation. Column (1) use all US data and column (2), (3) and (4) restrict the sample to the individuals based on their dominant type. Robust standard errors are used. Significance levels: *** $p<0.01$, ** $p<0.05$, and * $p<0.1$. The data come from the World Value Survey and the European Value Survey.

K Additional Details on the Polarisation Measure

The Esteban and Ray (1994) measure of polarisation is based on three axioms. These three axioms aim to capture sensible assumptions about how own-group identification and out-group alienation contribute to an overall index of polarisation.

Figure K.1 illustrates the three axioms of Esteban and Ray (1994) graphically. The first axiom states that polarisation increases if two small masses *b* and *c* that are close to each other are joined at their midpoint (see panel (a) of Figure K.1). The intuition behind this axiom is that the joining of the masses increases the own-group identification of the now joined smaller masses, while the average distance and out-group alienation with respect to other major societal groups *a* stay unchanged.

The second axiom states that polarisation increases if a small mass of people *b* moves closer to the side of the spectrum where fewer people are concentrated (see panel (b) of Figure K.1). Put simply, this change increases polarisation because while the mass *b* has moved closer to group *c* it has also moved further away from another group *a*. Since mass *a* is larger than mass *c*, the overall alienation effect increases.

The third axiom states that polarisation increases if mass is shifted equally from a central mass *b* to two lateral masses *a* and *c* that are each equally far away from the central mass (see panel (c) of Figure K.1). This axiom captures the effect of the disappearing centre. If mass shifts equally from the centre to the fringes of the spectrum the own-group identification at the fringes increases while the overall out-group alienation increases as well.

Esteban and Ray (1994) prove that any measure of polarisation that fulfils these three axioms must be of the form:

$$P(\pi, y) = \kappa \sum_{i=1}^{n} \sum_{j=1}^{n} \pi_i^{1+\nu} \pi_j |y_i - y_j| \tag{K.1}$$

The axioms hold for values of $\nu \in [0, 1.6]$. The sensitivity parameter ν also influences the maximal possible value of the polarisation measure. Note that the measure will not be bounded between [0,1]. Esteban and Ray (1994) suggest a potential fourth axiom that would make it possible to narrow the possible interval of $\nu \in [1, 1.6]$.

This fourth axiom is illustrated in Figure K.2. The axiom states that moving mass from a small mass *a* to a larger mass *c* will increase polarisation. Hence, the axiom makes an assumption on the importance of small groups within a society. On the one hand, moving mass from *a* to *c* reduced the distance between the groups and therefore lowered polarisation. On the other hand the mass *a* is small

Figure K.1: Axioms of Esteban & Ray 1994

(a) Axiom 1

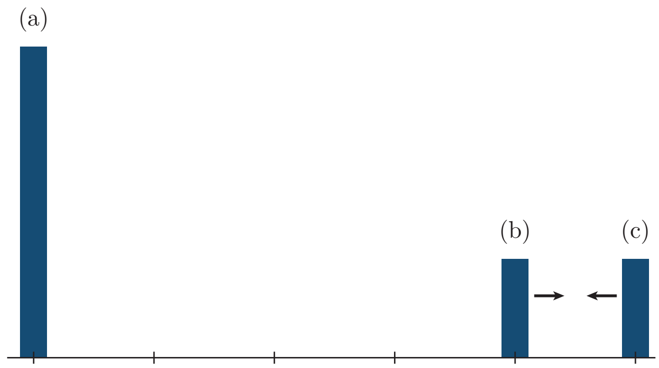

(b) Axiom 2

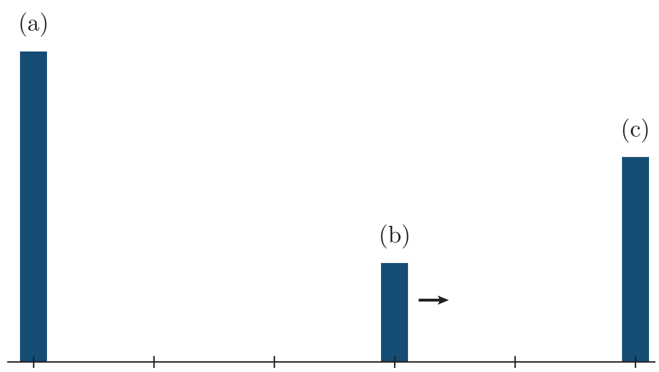

(c) Axiom 3

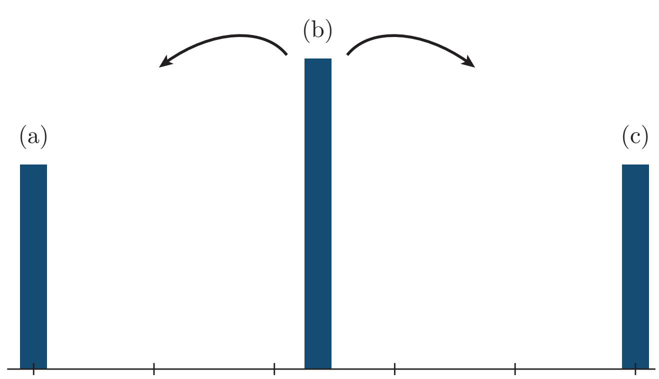

Notes: This figure illustrates the 3 main axioms use in Esteban and Ray (1994) to derive the polarisation measure.

in comparison to b and c and hence the effect of group a for overall polarisation might be negligible while increasing the mass of c can increase societal tension.

The polarisation sensitivity parameter ν here captures the relative sizes of a and c for which polarisation will increase. The larger is ν the smaller is the importance of a for overall polarisation. It is a priori not clear whether this axiom is sensible in our context. Hence, we do not restrict the range of polarisation sensitivity to $\nu > 1$.

Figure K.2: Additional Axiom of Esteban & Ray 1994

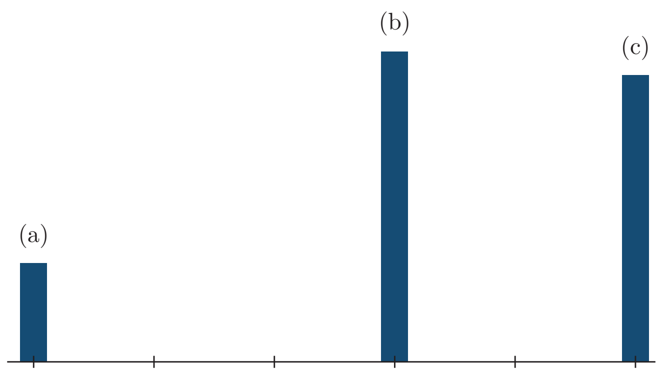

Notes: This figure illustrates the 4th axiom suggested in Esteban and Ray (1994). This axioms is not necessary to derive the form of the polarisation measure but it allows for restrictions to the possible range of ν.

K.1. Extending the Esteban and Ray (1994) Measure to Higher Dimension

The Esteban and Ray (1994) measure was originally constructed for one-dimensional indicators (e.g. the income distribution). Our measure extends the measure to the four dimensions of our ideological type space. We assume that an individual identifies with groups based on his or her dominant type share since in our model the four ideological types are the most natural line for group delineations.

Theoretically, it would also be possible to define groups based on discrete intervals of the type share distribution, such that a type would be defined by a specific interval in the four-dimensional ideological type space (e.g. [0,0.1] Liberal Centrist, [0.2,0.3] Conservative Centrist, [0.4,0.5] Left and Right Anarchist). This would lead to a far greater number of ideological groups. The problem with this approach is that it is not obvious to decide on an interval length such that we can plausibly assume sufficient degrees of separation between these groups.

If the groups are defined by the dominant type share of each individual, intuitively, the alienation between these groups will be based on differences in type shares. The only alteration to the original measure then is the fact that in our case the groups can differ along four dimensions rather than a single variable y. We hence define the overall out-group alienation as the sum of the type share differences between different groups.

K.2. Robustness Esteban-Ray Measure

So far we have not addressed the question of the choice of ν. As explained above any $\nu \in [0, 1.6]$ leads to a measure of polarisation that fulfils the axioms of Esteban and Ray (1994). As a robustness exercise, we calculate the Esteban-Ray measure for several values of ν. Table K.1 reports the ranking of countries by their polarisation over the three waves conditional on the choice of ν. It is important to note that the values of the polarisation measure are not comparable across different ν, since dependent on ν the maximal possible polarisation level varies.

Our main finding for the rising level of polarisation in the US holds for all except the largest values of ν. As long as $\nu < 1$ the US emerges as the most polarised country in our sample. The results for $\nu = 1.6$ differ, since for high values of ν the importance of small groups in society is diminished. Hence, in this case, the polarisation P measure for the US - where we observe four comparably sized ideological groups - is lower than for other values of ν. In contrast, measured polarisation is higher in countries with one large ideological group, e.g. the Conservative Centrist in Malta or Liberal Centrist in Denmark.

Overall, the results seem to point towards the fact that values of $\nu < 1$ lead to a more balanced polarisation ranking across countries. The fact that for $\nu = 1.6$ countries such as Denmark, Iceland, Finland and Canada - all of which are usually considered harmonious societies - end up at top of the ranking seems counterintuitive. Based on these findings we set $\nu = 0.5$ as the baseline value for polarisation sensitivity in our main P measure.

Table K.1: Esteban-Ray Polarisation Measure for different ν

Panel A: Wave 2

$\nu = 0$		$\nu = 0.5$		$\nu = 1$		$\nu = 1.6$	
Country	Pol. Measure	Country	Pol. Measure	Country	Pol. Measure	Country	Pol. Measure
Spain	1.077	Spain	0.555	Malta	0.356	Malta	0.226
France	1.059	Austria	0.539	North Ireland	0.315	North Ireland	0.178
Belgium	1.058	France	0.539	Portugal	0.308	Portugal	0.177
Italy	1.024	Belgium	0.538	Austria	0.301	Ireland	0.167
Netherlands	1.024	Malta	0.532	Netherlands	0.296	Netherlands	0.161
Germany	1.017	Netherlands	0.531	Spain	0.293	Austria	0.154
Austria	1.006	North Ireland	0.530	United States	0.287	United States	0.150
Great Britain	0.990	Italy	0.528	Ireland	0.285	Canada	0.143
North Ireland	0.958	Germany	0.519	Italy	0.282	Denmark	0.142
Canada	0.954	Great Britain	0.518	France	0.281	Spain	0.139
Finland	0.929	Portugal	0.508	Belgium	0.279	Iceland	0.138
United States	0.921	United States	0.504	Great Britain	0.278	Italy	0.138
Iceland	0.902	Canada	0.503	Canada	0.278	Great Britain	0.135
Portugal	0.898	Finland	0.478	Germany	0.271	France	0.132
Ireland	0.853	Iceland	0.478	Iceland	0.266	Finland	0.131
Malta	0.849	Ireland	0.472	Finland	0.258	Belgium	0.130
Denmark	0.827	Denmark	0.442	Denmark	0.255	Germany	0.128

Panel B: Wave 4

$\nu = 0$		$\nu = 0.5$		$\nu = 1$		$\nu = 1.6$	
Country	Pol. Measure	Country	Ray Measure	Country	Pol. Measure	Country	Pol. Measure
Spain	1.151	Spain	0.576	Malta	0.349	Malta	0.230
Austria	1.070	Austria	0.553	North Ireland	0.300	Iceland	0.188
France	1.066	Great Britain	0.546	Ireland	0.297	Denmark	0.182
Belgium	1.057	France	0.540	United States	0.295	Netherlands	0.163
Germany	1.053	Germany	0.537	Austria	0.294	Ireland	0.162
Great Britain	1.052	Italy	0.537	Canada	0.291	Canada	0.161
Italy	1.038	United States	0.533	Great Britain	0.291	North Ireland	0.157
United States	1.005	North Ireland	0.530	Spain	0.289	Finland	0.157
North Ireland	0.987	Belgium	0.529	Italy	0.287	United States	0.151
Canada	0.952	Ireland	0.514	Finland	0.286	Portugal	0.150
Portugal	0.946	Malta	0.509	Netherlands	0.284	Austria	0.141
Ireland	0.945	Canada	0.507	France	0.281	Great Britain	0.140
Finland	0.935	Finland	0.499	Germany	0.280	Italy	0.139
Netherlands	0.918	Portugal	0.498	Portugal	0.279	France	0.133
Malta	0.794	Netherlands	0.487	Iceland	0.279	Germany	0.131
Iceland	0.755	Iceland	0.428	Belgium	0.266	Spain	0.126
Denmark	0.669	Denmark	0.376	Denmark	0.254	Belgium	0.117

Panel C: Wave 5

$\nu = 0$		$\nu = 0.5$		$\nu = 1$		$\nu = 1.6$	
Country	Pol. Measure	Country	Pol. Measure	Country	Pol. Measure	Country	Pol. Measure
United States	1.068	United States	0.563	Malta	0.320	Malta	0.208
Netherlands	1.063	Netherlands	0.543	United States	0.306	Denmark	0.175
Austria	1.057	Austria	0.534	Canada	0.291	Iceland	0.168
Spain	1.054	Spain	0.530	North Ireland	0.291	Finland	0.162
Germany	1.032	Canada	0.528	Ireland	0.291	North Ireland	0.157
France	1.020	Ireland	0.523	Netherlands	0.285	United States	0.152
Belgium	1.000	Great Britain	0.520	Finland	0.284	Ireland	0.150
Canada	0.999	Germany	0.519	Portugal	0.283	Canada	0.149
Great Britain	0.999	France	0.518	Great Britain	0.280	Portugal	0.147
Ireland	0.985	North Ireland	0.510	Austria	0.275	Great Britain	0.138
Italy	0.954	Belgium	0.507	Italy	0.273	Italy	0.138
North Ireland	0.949	Portugal	0.502	France	0.270	Netherlands	0.137
Portugal	0.930	Italy	0.499	Spain	0.269	France	0.128
Finland	0.916	Finland	0.489	Germany	0.264	Austria	0.126
Iceland	0.767	Malta	0.475	Belgium	0.263	Belgium	0.124
Malta	0.756	Iceland	0.413	Iceland	0.258	Spain	0.120
Denmark	0.595	Denmark	0.334	Denmark	0.233	Germany	0.118

Notes: This table reports the polarisation measure for different ν. For more details, see the text.

L Comparison of LDA to PCA, Factor Analysis and K-means

This section provides a comparison between Latent Dirichlet Allocation (LDA) and the other alternative machine learning dimensionality reduction techniques, specifically Principal Component Analysis (PCA), Factor Analysis (FA) and k-means clustering. At their core, all of these techniques aim to reduce high dimensional data to a set of more easily interpretable topics, components, factors or clusters. Differences arise in the way these lower-dimensional representations of the data are constructed.

As we have outlined in detail in the main part of the paper, LDA relies on a generative model that makes assumptions about the data-generating process and allows for a direct interpretation of the latent objects as topics. Furthermore, the LDA model was specifically designed for the analysis of sparse multinomial data.

PCA, on the other hand, relies on a truncated singular value decomposition to derive components that explain the maximum possible amount of variance in the data while keeping all components orthogonal to each other. The truncated singular value decomposition is based on decomposing the original $O \times F$ data matrix D of rank R with O observation and F features into three matrices such that $D = U\Sigma W^T$, where U is a $O \times R$ orthogonal matrix, W^T is a $R \times F$ orthogonal matrix, and Σ is a $R \times R$ diagonal matrix. Afterwards, PCA truncates the resulting matrices by removing the rows and columns associated with the smallest eigenvalues in the matrix Σ. This truncation process reduces the dimensions of the matrices to a user-chosen number of components C, such that U becomes U_C of dimension $O \times C$, Σ becomes Σ_C of dimension $C \times C$, and W^T becomes W_C^T of dimension $C \times F$.

Each of the resulting components are orthogonal to each other and represent a linear combination of the original data weighted by eigenvectors. This highlights two important limitations of PCA for our application. Neither is it obvious that the ideological types (components) we want to find in the data should be orthogonal to each other nor are they necessarily a linear combination of the data. As a result, the ideological type hierarchy created by PCA (see Table L.1) is less coherent than the types created by LDA.

Similar problems arise when using FA. FA represents the original data as a linear combination of factors such that $D = C + \beta \cdot F + \epsilon$, where D is the original data matrix, C is a vector of constants F is the factor matrix, β are the factor loadings and ϵ a vector of Gaussian noise. The advantage of FA in comparison to PCA is that it accounts for random measurement error through the ϵ vector and hence allows for heteroscedastic noise. Nevertheless, FA still uses a linear model to decompose the data. Due to the linear model, the ideological type generated by FA (see Table L.2) are less coherent than the LDA results. Note that the change in types 1 and 2 from the 2-type to the 3-type model is driven by the change in the signs of the factor loadings. The factors still load on the same features, they just point in the opposite direction.

Last, k-means is a clustering algorithm that minimises the distance of the original data to a user-chosen number of centroids. As with any other clustering algorithm, k-means assigns each observation to a unique cluster. This seems counterintuitive in our case since people do not necessarily subscribe to a single political ideology. For example, people might be liberal when it comes to social issues but conservative with regard to economic questions. While LDA captures this its mixture of ideological types, k-means cannot account for this.[39] Moreover, as discussed by Ding and He (2004) k-means clustering represents a discrete cluster solution to the components derived by PCA. As such k-means suffers from similar shortcomings as PCA, and the derived ideological types (see Table L.3) also are less coherent in comparison to LDA.

[39] PCA and FA also allow for 'mixed membership' through different component and factor loadings.

Table L.1: Hierarchy of Types (Top Ten Features) as created by PCA

2 Type Model	3 Type Model	4 Type Model
Type 1	**Type 1**	**Type 1**
No Confidence: Churches	No Confidence: Churches	No Confidence: Churches
No Confidence: Civil services	No Confidence: Civil services	No Confidence: Civil services
No Confidence: Parliament	No Confidence: Parliament	No Confidence: Parliament
No Confidence: Armed forces	No Confidence: Armed forces	No Confidence: Armed forces
No Confidence: Justice system/courts	No Confidence: Justice system/courts	No Confidence: Justice system/courts
No Confidence: Police	No Confidence: Police	No Confidence: Police
No Confidence: Major companies	No Confidence: Major companies	No Confidence: Major companies
Justifiable: Euthanasia	Justifiable: Euthanasia	Justifiable: Euthanasia
Justifiable: Abortion	Justifiable: Abortion	Justifiable: Abortion
Justifiable: Divorce	Justifiable: Divorce	Justifiable: Divorce
Type 2	**Type 2**	**Type 2**
Not Justifiable: Abortion	Not Justifiable: Abortion	Not Justifiable: Abortion
Not Justifiable: Homosexuality	Not Justifiable: Homosexuality	Not Justifiable: Homosexuality
Not Justifiable: Euthanasia	Not Justifiable: Euthanasia	Not Justifiable: Euthanasia
No Confidence: Justice system/courts	No Confidence: Justice system/courts	No Confidence: Justice system/courts
No Confidence: Parliament	No Confidence: Parliament	No Confidence: Parliament
No Confidence: Civil services	No Confidence: Civil services	No Confidence: Civil services
Not Justifiable: Divorce	Not Justifiable: Divorce	Not Justifiable: Divorce
No Confidence: Labour unions	No Confidence: Labour unions	No Confidence: Labour unions
Not Justifiable: Prostitution	Not Justifiable: Prostitution	Not Justifiable: Prostitution
No Confidence: The press	No Confidence: The press	No Confidence: The press
	Type 3	**Type 3**
	More responsibility for people	More responsibility for people
	Against Neighbours: Drug addicts	Against Neighbours: Drug addicts
	Competition is good	Competition is good
	Private better than state ownership	Private better than state ownership
	If Jobs Scarce: priority to (nation) people	If Jobs Scarce: priority to (nation) people
	Justifiable: Euthanasia	Justifiable: Euthanasia
	Against Neighbours: People AIDS	Against Neighbours: People AIDS
	Confidence: Armed forces	Confidence: Armed forces
	No Confidence: Labour unions	No Confidence: Labour unions
	Against Neighbours: Immigrants/foreign workers	Against Neighbours: Immigrants/foreign workers
		Type 4
		More responsibility for government
		Against Neighbours: People AIDS
		Confidence: Labour unions
		If Jobs Scarce: priority to (nation) people
		Against Neighbours: Homosexuals
		Against Neighbours: Immigrants/foreign workers
		Confidence: Press
		Competition is harmful
		State ownership better than private ownership
		Against Neighbours: Drug addicts

Notes: This table reports the 10 most important features for a n-type Principal Component Analysis model, where $n \in \{2, 3, 4\}$.

Table L.2: Hierarchy of Types (Top Ten Features) as created by Factor Analysis

2 Type Model	**3 Type Model**	**4 Type Model**
Type 1	**Type 1**	**Type 1**
Confidence: Churches	No Confidence: Civil services	No Confidence: Civil services
Confidence: The civil services	No Confidence: Churches	No Confidence: Churches
Confidence: Parliament	No Confidence: Parliament	No Confidence: Parliament
Not Justifiable: Abortion	No Confidence: Justice system/courts	No Confidence: Justice system/courts
Confidence: Armed forces	No Confidence: Armed forces	Justifiable: Abortion
Confidence: Justice system/courts	No Confidence: Police	No Confidence: Police
Confidence: Police	Justifiable: Abortion	No Confidence: Armed forces
Not Justifiable: Prostitution	Justifiable: Euthanasia	Justifiable: Euthanasia
Not Justifiable: Euthanasia	Justifiable: Divorce	Justifiable: Divorce
Confidence: Major Companies	No Confidence: Major companies	Justifiable: Homosexuality
Type 2	**Type 2**	**Type 2**
Justifiable: Abortion	Not Justifiable: Abortion	Not Justifiable: Abortion
Justifiable: Homosexuality	No Confidence: Parliament	No Confidence: Parliament
Confidence: Parliament	No Confidence: Civil services	No Confidence: Civil services
Confidence: The civil services	Not Justifiable: Homosexuality	Not Justifiable: Homosexuality
Justifiable: Divorce	Not Justifiable: Euthanasia	No Confidence: Justice system/courts
Confidence: Justice system/courts	No Confidence: Justice system/courts	Not Justifiable: Euthanasia
Justifiable: Euthanasia	Not Justifiable: Divorce	Not Justifiable: Divorce
Confidence: Labour unions	Not Justifiable: Prostitution	No Confidence: Labour unions
Confidence: Police	No Confidence: Labour unions	Not Justifiable: Prostitution
Confidence: Press	Not Justifiable: Suicide	No Confidence: The press
	Type 3	**Type 3**
	Against Neighbours: Immigrants/foreign workers	Against Neighbours: Immigrants/foreign workers
	Against Neighbours: People different race	Against Neighbours: People different race
	Against Neighbours: People AIDS	Against Neighbours: People AIDS
	Against Neighbours: Homosexuals	Against Neighbours: Homosexuals
	If Jobs Scarce: priority to (nation) people	If Jobs Scarce: priority to (nation) people
	Not Justifiable: Homosexuality	Not Justifiable: Homosexuality
	Against Neighbours: Drug addicts	Against Neighbours: Drug addicts
	Not Justifiable: Abortion	Not Justifiable: Abortion
	Not Justifiable: Divorce	Not Justifiable: Divorce
	No Confidence: Justice system/courts	No Confidence: Justice system/courts
		Type 4
		Not Justifiable: Cheating on taxes
		Not Justifiable: Claiming government benefits
		Not Justifiable: Avoiding a fare on public transport
		Not Justifiable: Someone accepting a bribe
		Justifiable: Homosexuality
		No Problem Neighbours: Homosexuals
		No Problem Neighbours: People AIDS
		Justifiable: Divorce
		No Confidence: Major companies
		Competition is good

Notes: This table reports the 10 most important features for a n-type Factor Analysis model, where $n \in \{2, 3, 4\}$.

Table L.3: Hierarchy of Types (Top Ten Features) as created by K-means

2 Type Model	**3 Type Model**	**4 Type Model**
Type 1	**Type 1**	**Type 1**
Not Justifiable: Someone accepting a bribe	No Problem Neighbours: People different race	Proud of nationality
No Problem Neighbours: People different race	Confidence: Police	Not Justifiable: Someone accepting a bribe
No Confidence: Parliament	No Problem Neighbours: Homosexuals	Confidence: Police
No Problem Neighbours: Homosexuals	No Problem Neighbours: Immigrants/foreign workers	Not Justifiable: Cheating on taxes
No Problem Neighbours: Immigrants/foreign workers	Proud of nationality	Not Justifiable: Claiming government benefits
Not Justifiable: Claiming government benefits	Not Justifiable: Someone accepting a bribe	Not Justifiable: Avoiding a fare on public transport
No Problem Neighbours: People AIDS	No Problem Neighbours: People AIDS	Not Justifiable: Suicide
Proud of nationality	Not Justifiable: Claiming government benefits	No Problem Neighbours: People different race
No Confidence: Civil services	Not Justifiable: Cheating on taxes	Confidence: Armed forces
Not Justifiable: Cheating on taxes	Confidence: Justice system/courts	Confidence: The civil services
Type 2	**Type 2**	**Type 2**
Proud of nationality	No Problem Neighbours: People different race	No Problem Neighbours: People different race
Not Justifiable: Someone accepting a bribe	No Confidence: Parliament	No Problem Neighbours: Homosexuals
Confidence: Police	Not Justifiable: Someone accepting a bribe	Confidence: Police
No Problem Neighbours: People different race	No Problem Neighbours: Homosexuals	No Problem Neighbours: Immigrants/foreign workers
Not Justifiable: Cheating on taxes	No Problem Neighbours: People AIDS	Not Justifiable: Someone accepting a bribe
Not Justifiable: Claiming government benefits	No Problem Neighbours: Immigrants/foreign workers	No Problem Neighbours: People AIDS
No Problem Neighbours: Immigrants/foreign workers	No Confidence: Civil services	Proud of nationality
Not Justifiable: Avoiding a fare on public transport	No Confidence: Major companies	Not Justifiable: Claiming government benefits
Confidence: Armed forces	Not Justifiable: Claiming government benefits	Not Justifiable: Cheating on taxes
No Problem Neighbours: Homosexuals	No Confidence: The press	Justifiable: Divorce
	Type 3	**Type 3**
	Not Justifiable: Someone accepting a bribe	No Problem Neighbours: Homosexuals
	Proud of nationality	No Problem Neighbours: People different race
	Not Justifiable: Cheating on taxes	No Confidence: Parliament
	Not Justifiable: Suicide	No Problem Neighbours: People AIDS
	Not Justifiable: Claiming government benefits	No Problem Neighbours: Immigrants/foreign workers
	Not Justifiable: Avoiding a fare on public transport	Not Justifiable: Someone accepting a bribe
	Not Justifiable: Prostitution	No Confidence: Civil services
	No Problem Neighbours: People different race	No Confidence: Churches
	Confidence: Police	No Confidence: Major companies
	Not Justifiable: Abortion	Not Justifiable: Claiming government benefits
		Type 4
		Not Justifiable: Someone accepting a bribe
		Proud of nationality
		Not Justifiable: Claiming government benefits
		Not Justifiable: Cheating on taxes
		Not Justifiable: Avoiding a fare on public transport
		No Confidence: Parliament
		Not Justifiable: Suicide
		No Problem Neighbours: People different race
		Not Justifiable: Prostitution
		No Confidence: The press

Notes: This table reports the 10 most important features for a n-type k-means model, where $n \in \{2, 3, 4\}$.