

# Moral Universalism and the Structure of Ideology

BENJAMIN ENKE

*Harvard University and NBER*

RICARDO RODRÍGUEZ-PADILLA

*Harvard University*

and

FLORIAN ZIMMERMANN

*University of Bonn and Institute on Behavior and Inequality*

*First version received September 2022; Editorial decision June 2022; Accepted September 2022 (Eds.)*

Throughout the Western world, people's policy views are correlated across domains in a strikingly similar fashion. This article proposes that what partly explains the structure of ideology is *moral universalism*: the extent to which people exhibit the same level of altruism and trust towards strangers as towards in-group members. In new large-scale multinational surveys, heterogeneity in universalism descriptively explains why some people support redistribution, health care, environmental protection, affirmative action, and foreign aid, while others advocate for spending on the military, law enforcement, and border protection. Universalism is a substantially stronger predictor of policy views and ideological constraints than variables such as income, wealth, education, religiosity, or beliefs about government efficiency. Consistent with the idea that universalism shapes policy views, we further document that the left–right divide on redistribution, environmental protection, or foreign aid strongly attenuates or even reverses when people evaluate less universalist implementations of these policies.

*Key words:* Moral universalism, Ideological constraint, Behavioural political economy.

*JEL Codes:* D72, D01

## 1. INTRODUCTION

A key object of interest in the study of political economy is *ideological constraint*: whether people's policy views are systematically correlated across economic and social domains, or rather relatively independently distributed across seemingly unrelated policy topics (e.g. Ansolabehere, Rodden and Snyder, 2008; Lenz, 2013). We implement rich new large-scale surveys in Australia, Germany, France, the US, and Sweden to document that people in one cluster generally desire government expenditure on foreign aid, affirmative action, environmental protection, welfare, and universal health care, while people in another cluster support government spending on the military, police, and law enforcement, and border control.

---

*The editor in charge of this paper was Nicola Gennaioli.*

While these ideological clusters may seem obvious to some because they have almost come to define how the public thinks about political disagreements, their structure is also somewhat puzzling. First, the prevailing ideological clusters are very similar across many rich Western countries, even though these often exhibit considerable differences in electoral systems, party structures, and ethnic composition. Second, despite the widespread occurrence of these clusters, there is no obvious “law of nature” that explains why these *particular* bundles of policy views should prevail in the first place. For example, it is not obvious why support for redistribution is always correlated with support for environmental protection rather than support for a strong military. Still, the striking similarity of the correlations in issue positions across countries suggests that these bundles reflect more than coincidence.

The central proposition of this article is that what imposes this particular structure on the space of policy views is heterogeneity in *moral universalism*: the extent to which people exhibit the same level of altruism and trust towards strangers as towards in-group members. Universalism is not about a person’s overall *level* of altruism or trust, but instead about its *slope* as a function of social distance, such that group- or place-based identities are less important to universalists than to communitarians.

Our analysis of the role of universalism for policy views is motivated by a growing body of work that argues that people’s political ideology and voting behaviour to a large extent reflect heterogeneity in moral values related to universalism, rather than most material concerns (e.g. Graham, Haidt and Nosek, 2009; Haidt, 2012; Waytz, Iyer, Young, Haidt and Graham, 2019; Enke, 2020). Building on this body of work, our investigation is based on the simple intuition that one of the two aforementioned ideological clusters consists of policies that require high altruism towards, and trust in, socially distant individuals, while the other cluster comprises policies that are primarily aimed at protecting, and caring for, certain in-groups. For instance, foreign aid, affirmative action, environmental protection, and federal redistribution are all highly universalist policies in that their beneficiaries will often be socially or geographically distant strangers. On the other hand, public spending on the military or border control is often designed to erect boundaries between “us” and “them,” which may at least partly reflect the communitarian moral stance that compatriots (in-group members) deserve higher priority, as well as the belief that outsiders cannot be trusted.

Based on these hypotheses, this article presents a new set of descriptive stylized facts to make four contributions. First, we collect rich new survey data to show that people’s desired expenditure shares are strongly partitioned into two distinct clusters, and that these clusters are extremely similar across Australia, France, Germany, Sweden, and the US. Second, by supplementing our data-collection exercise with rich and experimentally validated measures of universalism in altruism and universalism in trust, we show that universalism is strongly predictive of policy views in a way that rationalizes the observed clusters, in an almost identical fashion across countries. Third, we provide an extensive set of benchmarking analyses that suggest that universalism is quantitatively substantially more important for explaining policy views and ideological constraints than traditional political economy variables such as income, wealth, equity-efficiency preferences, or beliefs about government efficiency. Fourth, in an attempt to disentangle supply- and demand-side explanations, we document that people’s views on broad policy domains such as redistribution or foreign aid strongly depend on how universalist vs. communitarian the specific implementation of a policy is. For example, while left- and right-wingers differ in their views on national redistribution in familiar ways, this relationship strongly attenuates or even reverses when people are polled about local redistribution.

Our internet surveys are pre-registered and comprise nationally representative samples in each of five countries, for a total of about 11,000 respondents. To measure universalism in altruism, we implement structured decision tasks. In each task, a respondent is endowed with the hypothetical

sum of \$100 and is asked to split the money between two equally rich individuals: (i) a randomly selected member of a specific social (in-) group who lives in their own country of residence and (ii) a randomly selected person who lives in their own country of residence. For example, in one question, a US participant is asked to split hypothetical \$100 between a member of their extended family and a randomly selected person from the US. In addition to these questions that measure “domestic universalism”, we also measure “foreign universalism” and “global universalism” through money allocation tasks that involve different types of foreigners. While all of our survey questions are hypothetical in nature, they underwent an extensive selection and experimental validation procedure and have been shown to be correlated with real donation decisions (see Enke, Rodríguez-Padilla and Zimmermann, 2022).

Using an analogous procedure, we estimate respondents’ universalism in trust by asking them to indicate who of the individuals described above they trust more. In our data, universalism in altruism and universalism in trust are highly correlated, which suggests that they capture the same underlying psychology, which we refer to as “moral universalism”. Respondents exhibit large variation in universalism: some participants always split their money or trust equally, while others consistently share more money with, and trust more, members of their in-groups.

We supplement these measures of universalism with detailed questions on respondents’ policy views. To balance the richness and quantitative interpretability against the cognitive burden placed on participants, we elicit both quantitative information on desired government expenditure levels for specific policies, and simple Likert scale measures.

Looking at the link between policy views and universalism, we find that desired expenditures are strongly correlated with universalism in the ways we hypothesized and pre-registered. Universalism is *positively* correlated with desired expenditure on welfare payments, environment, affirmative action, foreign aid, and—to a lesser extent—universal health care. Moreover, universalism is *negatively* correlated with desired expenditure on border control, military, and law enforcement and police. In this sense, universalism reproduces the structure of policy views that we attempt to explain in this article. In terms of quantitative magnitudes, moving from the universalism of zero (allocating all money or trust to the in-group) to equal splits is associated with changes in policy views of between 10% and 80% of a standard deviation.

These correlations are robust and general in the following three ways. (1) The results are qualitatively very similar when we consider either universalism in altruism or universalism in trust, though quantitatively universalism in altruism is typically about twice as important. (2) The relationship between universalism and policy preferences is robust against controlling for rich measures of income, wealth, religiosity, education, urbanicity, and beliefs about government efficiency, among others. (iii) The results are strikingly similar across the US, Australia, France, Germany, and Sweden.

To put our results in perspective, we implement a series of exercises that benchmark universalism against a rich set of individual characteristics. In the first step, we assess the degree to which universalism and other characteristics contribute to the magnitude of ideological constraint, defined as the average intracorrelation of people’s policy views. In conservative estimations, we find that ideological constraint would decrease by 25% if there was no heterogeneity in universalism. To put this magnitude into perspective, our estimations suggest that ideological constraint would not change at all if there was no heterogeneity in, for example, income, wealth, education, equity-efficiency preferences, or beliefs about government efficiency.

In a second step of benchmarking our results on universalism, we use simple machine learning techniques to identify a candidate set of “predictors” of policy views, which partly consists of complex interactions of individual characteristics. We then study whether any of these candidate predictors meaningfully predict the structure of policy views. We find that the candidate

predictors are often correlated with desired expenditures in important and known ways. At the same time, universalism is the only variable in our data that meaningfully organizes the key pattern we are trying to explain: simultaneous support for government spending in the domains of welfare, universal health care, environmental protection, affirmative action, and foreign aid, but opposition to large government spending in the domains of military, police, and border control.

A potential alternative view of the formation of ideological clusters that contrasts with our emphasis on a demand-side mechanism are supply-side accounts, according to which people identify with a party and then “learn” from elites which bundles of policy positions they are supposed to hold. While such an account does not explain why we observe the *specific* ideological clusters that we do, we acknowledge and discuss in detail the potential role of supply-side mechanisms.

To provide direct evidence that a demand-side mechanism is at least a part of the story, we leverage the observation that many broad policies can be implemented in a more or less universalist vs. communitarian fashion. For instance, redistribution can be organized at an impersonal, federal level, or it can be implemented locally. Similarly, environmental protection policies can aim at preventing global climate change or saving local forests and rivers. We hypothesize that the traditional left–right divide on these and other topics attenuates or even reverses when canonical left-wing policies are implemented in a more communitarian fashion and when canonical right-wing policies are implemented in more universalist ways.

To test these hypotheses, we elicit respondents’ desired spending levels for specific policy implementations, where some policies are more universalist than others. For example, within the broad domain of welfare payments, we separately elicit the desired spending on “Redistributing local tax revenues as welfare payments across all communities nationwide” and “Redistributing local tax revenues as welfare payments only within the local communities they were raised.” We design these survey questions holding fixed the efficiency of the redistributive system. In these exercises, the relationship between respondents’ political orientation and policy views predictably attenuates or even reverses, depending on whether the specific policy proposal is more or less universalist. For example, respondents who identify as right-wing are equally likely to support redistribution or environmental protection as left-wingers once it takes place locally.

We view these results as suggesting that the link between universalism and policy views at least partly reflects a demand-side mechanism. To further corroborate this, we also show that all of our results on the link between universalism and policy views hold controlling for respondents’ political ideology on a left–right spectrum.

Linking our work to the literature, much research in economics, political science, and moral and political psychology has highlighted the role of morality, identity, and social preferences in political attitudes. Most closely related to our approach is recent work on moral universalism, in both psychology (e.g. Graham *et al.*, 2009; Haidt, 2012) and economics (Enke, 2020). We contribute to this line of research (1) by providing a quantitative analysis examining the internal structure of policy views rather than voting; (2) not just in the US but in the Western world more generally; and (iii) by working with direct economic measures of preferences and beliefs rather than psychological questionnaires. Much of our approach is inspired by the model in Tabellini (2008).

The idea that social groups and identity play an important role in understanding contemporary policy views runs through various recent contributions and reviews (Shayo, 2009; Grossman and Helpman, 2018; Bonomi *et al.*, 2021; Besley and Persson, 2019; Guriev and Papaioannou, 2022). Relatedly, large literatures explain variation in demand for redistribution through ethnic divisions, religion and citizenship (e.g. Alesina, Baqir and Easterly, 1999; Luttmer, 2001; Alesina and Glaeser, 2004; Scheve and Stasavage, 2006; Gilens, 2009;

Chen and Lind, 2019).<sup>1</sup> Our central contribution to these literatures is to highlight the importance of studying *heterogeneity in how much people care about group- and place-based identities*, as this shapes an entire vector of policy views in a strikingly similar fashion across Western democracies.

In political science, a large amount of work has been devoted to studying the internal structure of elite opinion (Poole and Rosenthal, 2000), but there is no extant theory that convincingly explains the internal structure of mass opinion. A popular view in the sub-field of political behaviour was (and is) that citizens usually do not hold internally coherent policy views, and that, by and large, the genuine ideological constraint does not exist. Rather, people are believed to follow the cues of party leaders (e.g. Lenz, 2013; Erikson and Tedin, 2015; Broockman and Butler, 2017). However, Ansolabehere *et al.* (2008) documented that such inferences about lack of preference stability and ideological constraint are largely driven by measurement error in surveys. We build on this insight by illuminating how universalism structures an entire vector of (richly measured) policy views, and by showing that demand-side explanations contribute to this phenomenon.

The remainder of the article proceeds as follows. Section 2 summarizes the internal structure of ideology. Section 3 states our hypotheses. Sections 4–6 describe the design and results of our surveys, and Section 7 concludes.

## 2. THE STRUCTURE OF WESTERN POLITICAL IDEOLOGY

We illustrate the structure of political ideology in rich Western societies using our own survey data, as they are substantially richer than standard cross-national political science datasets such as the CSES. The data cover the US, Australia, France, Germany, and Sweden, for a total of approximately 11,000 respondents. We elicited respondents' desired per capita expenditure for eight domains: welfare payments; universal health care; affirmative action; environmental protection; foreign aid; military; police and law enforcement; and border control.

To probe the correlation structure of policy views, we implement principal component analyses (PCA), separately in each country. We find that, in each country, the first principal component of  $\text{arsinh}(\text{desired expenditures})^2$  across domains exhibits an unsurprising and almost identical structure: it loads positively and with essentially equal weights on desired expenditure levels in the eight categories. This first component captures “big vs. small government” views. The second principal component closely corresponds to our object of interest: in each country, it loads negatively on desired expenditure levels for military, police and law enforcement, and border control, and almost always positively on welfare, universal health care, affirmative action, environmental protection, and foreign aid. This second component, by virtue of being orthogonal to the first one, intuitively captures desired expenditure *shares*.

To make this point more explicit, we perform a principal component analysis directly on desired shares of overall spending, see Figure 1. Border control, military, and police and law enforcement all receive negative weights in each country, while foreign aid, affirmative action,

1. The broader concept of social capital has received substantial attention in the political economy literature (Putnam, 2000). For example, Dal Bó, Finan, Folke, Persson and Rickne (2018) and Algan, Beasley, Cohen and Foucault (2018) document that far-right voters exhibit lower trust. Somewhat relatedly, a number of social theorists have argued that what fundamentally distinguishes the left from the right is that people on the left believe that human nature is fundamentally “good”, while people on the right believe that people are “flawed” and need control (e.g. Sowell, 2007; Lakoff, 2010). We differ from these contributions in that we emphasize the relevance of the *slope* of social capital, rather than its level.

2. Throughout the article, we transform desired expenditures as  $\text{arsinh}(x) = \ln(x + \sqrt{x^2 + 1})$ .




FIGURE 1

Factor loadings of the first principal component of desired expenditure shares. Sign convention: the loading on “Border” is always non-positive, and the other signs are determined accordingly.

environmental protection, welfare payments, and universal health care almost always receive positive weights.

This structure is reminiscent of intuitive notions of “left” and “right”. To confirm this intuition, we elicited from our respondents how they would position themselves on a left-vs.-right Likert scale. Figure 2 summarizes the relationship between respondents’ ideology and their desired expenditure levels. In all countries, more pronounced left-wing identification is correlated with *higher* desired expenditure levels for canonical left-wing policies and *lower* desired expenditure levels for canonical conservative policies.

Indeed, Figure 2 informally suggests that when respondents tell us that they are “left” or “right”, they appear to refer more to *how* they would like to use a given government budget rather than the *overall size* of government. To make this argument more formal, we compute the pairwise correlations between people’s left-vs.-right self-positioning, the first principal component of desired expenditure levels (the “big-vs.-small-government” component), and the first principal component of desired expenditure shares. We find that the correlation between the left-right-scale and the big-vs.-small-government component ranges between  $\rho = 0.14$  in the US and  $\rho = 0.02$  in France. In contrast, the correlation between the left-right-scale and the expenditure-shares-component ranges between  $\rho = 0.49$  in the US and  $\rho = 0.30$  in Australia. This suggests that a large part of people’s self-identification as “left” and “right” relates to *how* a given budget is spent, rather than how big the budget is in the first place.

While we make use of our own survey data to exposit the structure of ideological clusters, the existence of ideological constraint is well documented. While early scholars believed that ideological constraint is relatively weak, Ansolabehere *et al.* (2008) showed that such inferences are largely driven by measurement error. Once political views are elicited using rich and multiple measures, they tend to be highly intracorrelated. For example, in the 2016 American National Election Survey, standard composite measures of respondents’ economic conservatism and their social conservatism (Ansolabehere, Rodden and Snyder Jr, 2006) exhibit a correlation of  $\rho = 0.48$ , suggesting strong ideological constraint. Draca and Schwarz (2020) and Wu (2020) also provide evidence for the existence of ideological constraint across multiple Western democracies. While some recent research suggests that the magnitude of the intra-correlations between people’s




FIGURE 2

The figure plots the OLS regression coefficients of univariate regressions of desired arsinh expenditure levels for each policy domain on self-positioning on a left-right scale (0–10). The dependent variables are standardized into z-scores. Error bars indicate 95% confidence intervals using robust standard errors. The “All countries” specification includes country-fixed effects.

policy views has increased over the last 40 years, the basic qualitative structure of ideology has been relatively stable over time (e.g. Kozłowski and Murphy, 2021; Wu, 2020).

### 3. HYPOTHESES

*Conceptualizing universalism.* Figure 3 illustrates how we think about heterogeneity in universalism, which is a slight modification of the setup in Tabellini (2008). A person’s degree of universalism concerns the slope of their altruism (or trust) as social distance increases, holding fixed the overall level of altruism (or trust). According to this conceptualization, which will be reflected in our empirical measurement, universalists are not “more or less moral,” they just allocate a given level of altruism more uniformly, and trust people at varying social distances more uniformly. Intuitively, a universalist might argue that it is appealing to treat everyone equally, while communitarians might point out that the universalist does not treat her friends very well. Indeed, in Enke *et al.* (2022) we show that universalists have fewer friends and spend less time with them.

In this conceptualization, “social distance” is a stand-in for different types of social identities and group memberships, including family, friendship, ethnicity, religious beliefs, values, hobbies, nationality, etc. Since our interest is in both domestic and foreign policies, it will be useful to broadly distinguish between domestic in-groups, domestic strangers as well as global in-groups and global strangers.




FIGURE 3

Illustration of heterogeneity in universalism. In the empirical analysis, we do not impose that the domestic stranger is socially closer than the global in-group.

*Universalism and policy views.* To articulate why universalism should affect the formation of an entire vector of policy views, we rely on two complementary strategies. First, we here briefly discuss simple intuitions. Second, in Supplementary Appendix A, we provide a formal model based on Tabellini (2008) that also generates our hypotheses.

Our starting point is the observation that many policy domains involve the welfare and anticipated behaviour of people who differ in how socially close they are from the decision-maker. This suggests that the decision-maker's universalism in altruism (who he cares about) and his universalism in trust (who he believes will misbehave) will matter for how he thinks about policy proposals.

Consider first domestic policies that have a redistributive flavour, such as *welfare payments* or *universal health care*. In Western democracies, these are implemented in a highly impersonal fashion, whereby people's tax money is used to fund payments to others that they have potentially no connection. This is in stark contrast to the more community-based redistributive systems that have dominated human history for centuries. We argue that the decision-maker's support for such highly impersonal redistributive systems will partly depend on (1) whether he actually cares about people that are socially distant from him and (2) whether he believes that these socially distant people will "cheat on the system" (e.g. by claiming benefits they are not entitled to, or reducing their supply of labour as a form of moral hazard).

Very similar arguments apply to *affirmative action* and *foreign aid*: these are—almost by definition—policies that increase the welfare of people who are socially distant from most decision-makers. Thus, universalist individuals should be more likely to support them, both because they internalize the welfare of socially distant people to a larger degree, and because they are less inclined to believe that these strangers will "cheat" by reducing effort after receiving favourable treatment. Again relatedly, from the perspective of Western citizens, *environmental protection* arguably also largely benefits socially distant strangers, including future generations



and people in developing countries whose geographic location or lack of funds makes them especially vulnerable to climate change.

All of the preceding hypotheses concern policy domains in which universalists are more supportive of government spending than communitarians. Yet, we hypothesize that the opposite holds true for expenditure categories that carry signatures of “us vs. them.” First, because universalists internalize the welfare of foreigners more, and believe that they are as trustworthy as compatriots, universalists will be less supportive of strong *border protection* than decision-makers who place little weight on the welfare of foreigners (and potentially believe that they are more likely to commit crimes or to misbehave in other ways). Similarly, universalists who place high weight on the welfare of foreign nations, and who believe that these nations are trustworthy partners rather than potential aggressors, will be less inclined to spend money on a strong *military*.

Finally, regarding *police and law enforcement*, communitarian decision-makers believe that socially distant groups are more likely to cheat on society and commit crimes, which comes at the expense of the communitarian’s in-groups. Yet, the communitarian fundamentally dislikes this idea because he primarily cares about his in-group members and hence supports a strong police force.

**Hypothesis 1.** *Universalism is positively linked to support for welfare, universal health care, affirmative action, environmental protection, and foreign aid, but negatively linked to support for border protection, a strong military, and police and law enforcement. This holds for universalism in both altruism and trust.*

*Discussion: Default implementation of policies.* In practice, many of the policies discussed above contain both universalist and communitarian elements, or could be implemented in more or less universalist ways. For example, the US military not only defends American security (a communitarian cause) but also frequently engages in humanitarian or peacekeeping missions (a universalist cause). Redistribution could be implemented not just at the national level (a universalist system) but also at the level of small, local communities (a communitarian system). Environmental protection programs can be targeted at preventing global climate change (a universalist objective) but also at conserving local forests and rivers (a more communitarian objective). Foreign aid could be given to those in most need (a universalist implementation) but also to those who are international allies (a more communitarian implementation).

As a result, the development of our hypotheses above implicitly relies on the idea that policies are usually implemented—or framed—in a particular fashion, which we refer to as a policy’s *default implementation*. For instance, when people think of welfare, they arguably have in mind the type of redistributive system that is currently in place (a national one). Indeed, our analysis not only rests on the assumption of there being default implementations—we also assume that these are similar across the five countries that we study. We deem this assumption plausible, also given the cultural and economic similarities between the countries in our sample. Indeed, for some domains, this assumption is demonstrably true: in all countries in our sample, redistribution and health care do take place at the national level rather than locally.

*Manipulating policy implementations.* The fact that policies often contain both universalist and communal elements is useful because it allows us to elicit support for specific implementations of policies that are more or less universalist in nature.

**Hypothesis 2.** *The positive correlation between left- versus right-wing political ideology and support for canonical “left-wing” policies (such as support for welfare, environmental protection, and foreign aid) attenuates or even reverses once these policies are implemented in less*

*universalist ways. Similarly, the negative correlation between left-wing ideology and support for canonical “right-wing” policies (such as a strong military) attenuates or reverses once these policies are implemented in more universalist ways.*

To sum up, our empirical analyses will leverage variation in both (1) moral preferences and beliefs and (2) the policies that respondents are asked to assess. This is useful because approach (1) is necessarily correlational in nature (we cannot change people’s universalism), while approach (2) allows for experimental manipulations.

#### 4. SURVEY DESIGN

##### 4.1. Logistics

In the summer of 2019, we implemented internet surveys in Australia, France, Germany, Sweden, and the US through the infrastructure of the market research panel of *Dynata*. The original survey was developed in English, translated into other languages by *Dynata*, and then checked by us using native speakers.

The survey consisted of four components: (1) an introductory screen that elicited demographics and routed respondents into or out of the survey; (2) decision screens to measure universalism and other social preferences; (3) screens to measure policy views; and (4) a questionnaire to elicit additional information and covariates. The order of parts (2) and (3) was randomized across respondents.<sup>3</sup> We also randomized the order in which universalism in altruism and universalism in trust were elicited.<sup>4</sup>

We took two measures to ensure quality control. First, every respondent who completed the survey in less than 400 s was dropped and replaced by *Dynata*. Second, the survey contained two attention check questions, interspersed throughout the survey. Whenever a respondent answered an attention check incorrectly, they were immediately routed out of the survey and replaced by *Dynata*.

We contracted with *Dynata* for nationally representative samples of  $N = 1,700$  citizens aged at least 18 in each country (see details on the pre-registration below). However, because constructing a sample that is nationally representative along the lines of age, gender, ethnicity, income, employment status, and education is logistically difficult, *Dynata* eventually supplied a larger sample to us (total  $N = 11,063$ ), a subset of which ( $N = 8,500$ ) makes up the more representative samples that we pre-registered. Since we view throwing away data as scientifically questionable, all analyses reported in the main text make use of the full sample. In the Supplementary Appendix, we replicate all analyses using the pre-registered (smaller) representative samples. The results are always very similar. Sample characteristics are summarized in Supplementary Appendix B.1.

Finally, we implemented identical surveys also in Brazil and South Korea, see Section 6 and Supplementary Appendix B.7 for an analysis of these “non-Western” countries.

##### 4.2. Measurement of universalism

###### 4.2.1. Decision tasks.

*Universalism in altruism.* We rely on a set of structured experimentally validated decision tasks to measure universalism (Enke *et al.*, 2022). Respondents completed a total of 16

3. In Supplementary Appendix B.7, we provide histograms that show that the distribution of policy views is unaffected by whether policy views are elicited before or after universalism.

4. A permanent link for the US version of our survey is: [https://harvard.az1.qualtrics.com/jfe/form/SV\\_aftuqgHsyIASHkp](https://harvard.az1.qualtrics.com/jfe/form/SV_aftuqgHsyIASHkp).

hypothetical money allocation tasks that allow us to construct a summary statistic of universalism in altruism. Analogous to Figure 3 and our formal framework in Supplementary Appendix A, the construction of the decision tasks is organized along four different types of groups: domestic in-groups, domestic strangers, global in-groups, and global strangers.

First, to estimate *domestic* universalism, respondents made 10 decisions. In each of them, they were asked to split hypothetical \$100 between (1) a randomly selected person from their country of residence and (2) a randomly selected member of one of their social groups, who also resides in the respondent's country of residence. Across the ten questions, the social groups included extended family, friends of family, neighbours, colleagues at work or school, same organization (e.g. club), same age, same ethnic background or race, same political views, same hobbies, and same religious beliefs. For example, in one question, respondents in the US were asked to split \$100 between a randomly selected person who lives in the US and a member of their extended family, such as a cousin. The average allocation to the randomly selected person across the 10 questions then makes up the domestic universalism measure.

Second, to estimate *foreign* universalism, respondents were asked to split \$100 between (1) a randomly selected person from their country of residence and (2) a randomly selected person who lives anywhere in the world. Foreign universalism then corresponds to the monetary amount sent to the global stranger.

Third, to estimate *global* universalism, respondents made five decisions, in each of which they were asked to split hypothetical \$100 between (1) a randomly selected person who lives anywhere in the world and (2) a randomly selected person who lives anywhere in the world and is a member of the respondent's social groups. Across the five questions, the social groups included the same language, same religious beliefs, same ethnic background, same values, and same occupation. The average amount of money sent to the randomly selected world citizen makes up the global universalism measure.

For the purpose of these tasks, respondents were always asked to assume (1) that both individuals are equally rich (addressing income effects) and (2) that neither of these individuals would find out who sent them the money (ruling out reciprocity considerations). The order of questions was randomized across respondents. Figures 1 and 2 in Supplementary Appendix B.2 show example decision screens.

The money allocation decisions, and in particular the domestic, foreign, and global universalism summary components are all highly positively correlated with each other. To reduce the dimensionality of the data and minimize measurement error, we average the three components into a pre-registered summary statistic of universalism in altruism.

*Universalism in trust.* Respondents again completed a total of 16 tasks. The procedure was identical to the one described for altruism above, except that in a given task respondents were asked to allocate 100 points (rather than \$100) between two individuals, to express whom of the two they trust more. This was explained as indicating “how much you trust that different people will not cheat on you or take advantage of you.” This procedure again yields domestic, foreign, and global universalism components, which we average into a summary statistic of universalism in trust. Figure 3 in Supplementary Appendix B.2 shows an example decision screen.

*Composite measure of universalism.* Universalism in altruism and trust exhibit a correlation of  $\rho = 0.70$  after accounting for measurement error using the obviously related instrumental variables technique of Gillen, Snowberg and Yariv (2019). To reduce the dimensionality of the analysis, we often work with a composite measure, which consists of the unweighted average of universalism in trust and universalism in altruism. We always reference robustness checks that use the altruism and trust measures separately, see Section 5.4.

*Relationship to concepts in moral psychology.* Prior work on moral universalism in politics usually relies on psychological questionnaires such as the Moral Foundations Questionnaire (e.g. Graham *et al.*, 2009; Haidt, 2012; Enke, 2020). While these questionnaires have the advantage of capturing the full richness of the human moral mind, the underlying concepts are only vaguely related to economic concepts and utilitarian models of preferences and beliefs. Indeed, many concepts in psychological questionnaires such as the MFQ capture deontological motivations. At the same time, arguably the main idea behind, for example, the moral foundations of Haidt (2012) is the distinction between a relationship- or group-specific morality and a universalist morality.<sup>5</sup> Thus, our main objectives in designing our measures were (1) to separately capture preferences and beliefs; and (2) to capture the broad main distinction between a universalist and a relationship-specific morality.<sup>6</sup>

*Construct validity.* We validate the universalism measures along two dimensions. See Enke *et al.* (2022) for details. (1) We implemented an *ex ante* experimental validation procedure. First, we show that, over a 1-week horizon, our hypothetical measure of universalism in altruism is highly correlated with a financially incentivized measure of universalism, which consists of the same questions with real incentives. Second, we document that behaviour in our trust task is highly correlated with trust beliefs in a structured cheating task that is standard in the experimental economics literature. (2) We also show that our survey measure of moral universalism predicts real donation decisions: while universalists donate less to local community organizations, they donate more to nationwide and international charities.

**4.2.2. Descriptives.** Figure 4 shows a histogram of the composite universalism measure, pooled across countries. Numbers around 50 imply on average equal allocations of money and trust points to in-groups and strangers. Numbers below 50 indicate a tendency to allocate more money and trust points towards in-groups. Numbers above 50 correspond to the (largely counterfactual) case that someone allocates more money and trust points to socially more distant individuals. Supplementary Appendix B.3 shows histograms for each country separately.

Table 1 reports correlations with demographics. To avoid focusing on variables that we select based solely on our intuitions, we instead select demographics (or their interactions) through a LASSO regression, see Supplementary Appendix Table 9. The strongest correlations are with age and wealth, both of which correlate negatively with moral universalism. Similarly, men, higher-income individuals, and the religious exhibit lower universalism. Of potential interest is also the significant conditional correlation between education and universalism: fixing people's income and wealth, higher education is associated with higher universalism. This is potentially relevant for understanding recent changes in political cleavages (Gethin, Martínez-Toledano and Piketty, 2022). Supplementary Appendix B.5 reports correlations between universalism and an expanded set of survey variables.

Finally, while our study was designed to study within- rather than across-country variation, we note that the average composite universalism in each country is as follows: Australia 0.38, France 0.41, Germany 0.40, Sweden 0.38, and US 0.39.

5. For example, Enke (2020) collapsed Haidt's first four moral foundations into a one-dimensional summary measure because much psychological research suggests that such a "principal component" captures a large fraction of the variation in human morality (as far as ideas related to universalism are concerned).

6. Universalism is potentially related to the concept of individualism, which captures whether people define themselves as "I" or "we". We view the individualism–collectivism cleavage as capturing two distinct elements: whether people care more about their in-group (which captures universalism) and whether people are expected to primarily look after themselves rather than care about others. Our notion of universalism holds the latter constant.

Heterogeneity in Moral Universalism




FIGURE 4

Distribution of the composite measure of moral universalism, pooled across all Western countries. The amounts reflect allocations to random strangers, so that the measure is decreasing in in-group favouritism. 50 corresponds to an equal split of money or trust points.

TABLE 1  
*Individual-level correlates of universalism*

	Correlation between composite measure of universalism and:						
	Age	Religiosity (z-score)	Wealth Index (z-score)	Income Index (z-score)	Male (0–1)	Age × Male × Income (z-score)	College (0–1)
<i>Raw correlation</i>	−0.16***	−0.10***	−0.12***	−0.07***	−0.07***	−0.01	0.01
<i>OLS coeff. (w/Country FEs)</i>	−0.12***	−1.16***	−1.51***	−0.84***	−1.78***	−0.06	0.21
<i>OLS coeff. (multivariate) (w/Country FEs)</i>	−0.09***	−1.04***	−0.82***	−0.50***	−1.11***	0.33***	1.23***

*Notes:* The first row reports the Pearson’s raw correlation between individual characteristics and the composite measure of universalism ( $N = 11,063$ ). The second row reports OLS coefficients from individual regressions of the composite measure of universalism on the given characteristic, including country-fixed effects; this row thus presents by how many dollars/trust points universalism increases for a one unit change in the demographic variable. The third row reports OLS coefficients from a multivariate regression of the composite measure of universalism on all characteristics at once, including country fixed effects. See Supplementary Appendix E for details on the construction of the demographic variables. All z-scores are computed separately within each country. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

4.3. *Measurement of political attitudes*

We measure policy views in two complementary ways: (1) by eliciting measures of desired government expenditure amounts and (2) through standard Likert scale questions in which people express their support for certain policies. Both of these techniques have well-understood advantages and disadvantages, which is why we leverage both. The main upside of (1) is that the measures are quantitative in nature and do not rely on the interpretational ambiguity associated with responses on Likert scales. The main downside of these measures, however, is that they are cognitively more complex.

*Desired expenditures.* Respondents were instructed to imagine they could decide the average amount of money that their federal or national government collects per year from each citizen to spend on each of the eight policy categories. We asked respondents to assume that all dollar amounts collected for a category would be spent only on this particular category, without any waste. In addition, we provided respondents with a reference value: annual per capita spending on education in their country of residence.

Respondents were asked to enter eight monetary amounts to indicate their desired per capita spending levels for each of welfare, universal health care, foreign aid, environmental protection, affirmative action,<sup>7</sup> military and counterintelligence, police and law enforcement, and border control.<sup>8</sup> The order of these categories on the computer screen was randomized. Figure 4 in Supplementary Appendix B.2 shows an example decision screen, and histograms of desired expenditure amounts for each policy category and country are presented in Supplementary Appendix B.6.<sup>9</sup>

Frequently, we will be interested in linking universalism to desired expenditure *shares* rather than simple levels. For this purpose, we simply compute the fraction of money a respondent desires to spend on a given policy domain.

*Likert scale questions.* Respondents were asked to indicate their support for each of the eight policies described above on a 0–10 scale. We again frequently work with normalized measures, referred to as *relative policy support*, that consist of support for a given policy divided by the sum of support across all policies.

#### 4.4. *Covariates*

Even though this article is descriptive in nature, we seek to assess the extent to which a potential relationship between universalism and policy views is driven by omitted variables. Our survey hence elicits rich measures of covariates, including age, gender, ethnicity/race, educational attainment, income (two measures), wealth and asset ownership (three measures), religiosity (three measures), urbanicity, employment status, marital status, migration background, belief about whether the government is efficient or wasteful (on a scale 0–10), beliefs about whether the respondent is likely to personally benefit from government expenditure in a given category, and measures of altruism, generalized trust, and equity–efficiency preferences. See Supplementary Appendix E.

#### 4.5. *Pre-registration*

The survey was pre-registered on EGAP, see <http://egap.org/registration/5792>. The pre-registration contained (i) the desired sample size; (ii) the precise construction of the summary statistics of universalism in altruism and trust; (iii) predictions about how we expected

7. The term “affirmative action” was never used in the survey. Instead, this was described as “Measures to ensure no individual is disadvantaged in access to education, the labour force, and marriage.”

8. This selection of policies was motivated by two considerations. First, to include some of the most important policies in terms of government spending. Second, to restrict attention to those that are politically contentious in all countries that we study (which for example excludes gun control and abortion).

9. We implement the following procedure to all desired expenditure amounts referenced in the article. First, we perform a PPP conversion to USD. Because of the free-entry format, responses to these questions are subject to large outliers. As such, we then winsorize the desired PPP spending levels at  $\pm 3$  SD of the within-country mean, as specified in our pre-registration. This affects 2.2% of all responses. Third, we compute the inverse hyperbolic sine ( $\text{arsinh}$ ) of these winsorized, PPP desired expenditure amounts, and finally standardize them into z-scores within each country.

universalism to be correlated with support for each of the eight policy domains, as stated in Section 3; (iv) the construction of a summary statistic of policy views discussed below; and (v) the analysis of the specific policy proposals in Section 6.1.

## 5. RESULTS

### 5.1. *Universalism and self-reported ideology*

Our ultimate interest in this article is to understand how universalism potentially shapes an entire vector of policy views. However, because today clusters of policy views are commonly associated with the labels “left” and “right”, we begin by documenting the link between universalism and people’s self-assessment on a 0–10 left-vs.-right scale. Figure 5 summarizes the results separately for each country by showing binned scatter plots. We see that, in each country, the composite universalism measure is strongly correlated with left-wing identification. The correlation in the full sample is  $\rho = 0.25$ , conditional on country fixed effects. This correlation is statistically highly significant in each country and ranges between  $\rho = 0.19$  in France and  $\rho = 0.29$  in the US. While we provide more sophisticated benchmarking analyses later, it is perhaps informative that the corresponding correlations between left-vs.-right-wing ideology and other commonly analysed variables are substantially smaller in magnitude:  $\rho = -0.01$  for the belief that government is efficient vs. wasteful,  $\rho = 0.02$  for college degree,  $\rho = -0.06$  for age,  $\rho = -0.09$  for an income index, and  $\rho = -0.13$  for a wealth index (all correlations computed conditional on country fixed effects).

An obvious question is whether particular components of the universalism measure drive this result. Instead, as Supplementary Appendix C.1 shows, *every single one* of our money and trust tasks is significantly correlated with people’s self-reported ideology in the same direction. This suggests that a universalist mindset is a general characteristic of people who view themselves as “left”.

### 5.2. *Universalism and policy views*

Figure 6 shows the link between composite universalism and policy views. Figure 6A and B reports the results of regressions of desired expenditure shares on composite universalism, the only difference being that the regressions in Figure 6B control for age, gender, income, wealth, college, urbanicity, religiosity, equity-efficiency preferences, altruism, trust, beliefs about the

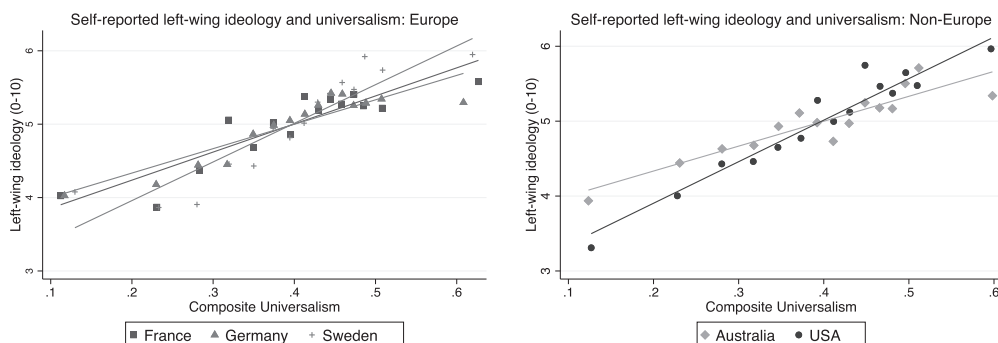


FIGURE 5

Binscatter plot of individuals’ left-right self-assessment against composite universalism. Both universalism and left-wing ideology are residualized of country fixed-effects.

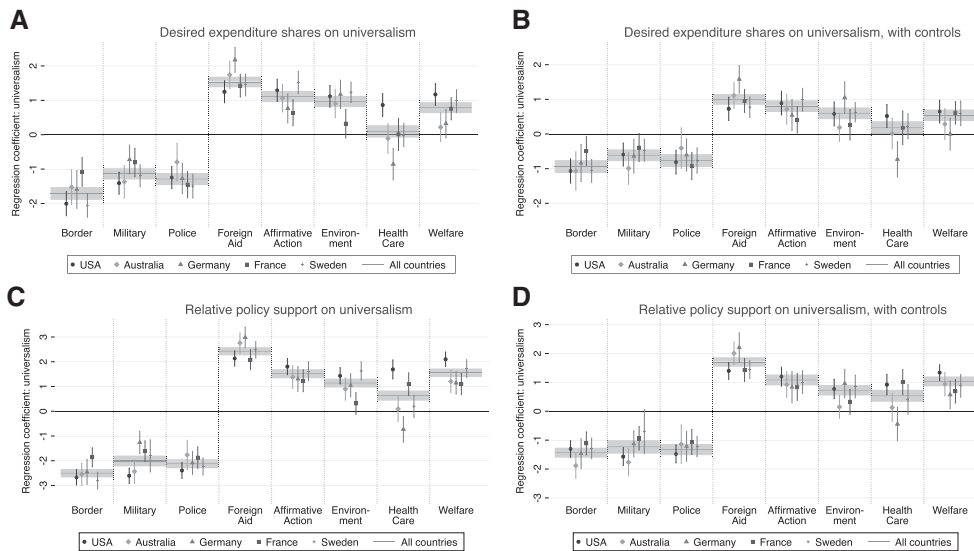


FIGURE 6

OLS regression coefficients of desired expenditure shares (A and B) or relative policy support (C and D) on composite universalism. The left panels only control for country FE, while the right panels also control for age, gender, income, wealth, college, urbanicity, religiosity, equity-efficiency preferences, altruism, trust, beliefs about the efficiency of government, and beliefs about whether one will personally benefit from government expenditure in each domain. In the top panels, the dependent variables are desired expenditure shares for each policy domain, while in the bottom panels, the dependent variables are relative policy support as computed from the Likert scale questions. Universalism is in  $[0,1]$  and the dependent variables are standardized into z-scores within each country. Error bars indicate 95% confidence intervals using robust standard errors.

efficiency of government, and beliefs about whether one will personally benefit from government expenditure in each domain. Figure 6C and D follows an analogous logic, except that the dependent variables are given by the relative policy support measures as computed from the Likert scale questions. All dependent variables are standardized into z-scores. Each panel shows the results of 48 regressions: eight policy domains in five countries each, plus a full sample specification.

As hypothesized, in all countries, we observe a strong negative relationship between universalism and desired expenditure shares or policy support for the three “right-wing” policy domains, while the relationship is generally positive and statistically significant for the five “left-wing” domains.<sup>10</sup> In Figure 6A, out of the 40 regression coefficients for the individual countries reported, 37 have the expected (pre-registered) sign. Of these, 33 are statistically significant at least at the 10% level. Once our battery of controls is added in Figure 6B, 38 of these coefficients have the expected sign, out of which 31 are statistically significant at least at the 10% level.

An immediate question is whether these results are only driven by universalism in altruism or universalism in trust alone. Indeed, while we construct Figure 6 based on the composite

10. A notable exception occurs in the domain of universal health care as elicited using expenditure shares, where the relationship is strongly positive in the US but either not statistically significant or even negative in the other countries. This pattern might arise because, in contrast to the US, all of these countries have had versions of universal health care for decades, which may generate less heterogeneity in views on universal health care across the political spectrum. Indeed, Figure 2 shows that, in all countries except the US, the link between self-reported left-wing ideology and desired health care expenditure levels is quantitatively very small, and about 1/4 the magnitude as in the US.



TABLE 2  
*Universalism and policy views*

	Panel A							
	Dependent variable: desired expenditure shares (Z-scores)							
	Border control	Military	Police	Foreign aid	Aff. action	Environment	Health care	Welfare
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Universalism in altruism	−0.84*** (0.08)	−0.55*** (0.08)	−0.57*** (0.08)	0.69*** (0.08)	0.51*** (0.08)	0.46*** (0.08)	0.16* (0.08)	0.25*** (0.09)
Universalism in trust	−0.59*** (0.10)	−0.28*** (0.10)	−0.36*** (0.10)	0.39*** (0.09)	0.30*** (0.09)	0.25*** (0.09)	0.030 (0.10)	0.35*** (0.11)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,881	10,881	10,881	10,881	10,881	10,881	10,881	10,881
R <sup>2</sup>	0.07	0.05	0.05	0.08	0.03	0.03	0.01	0.04

	Panel B							
	Dependent variable: relative policy support (Z-scores)							
	Border control	Military	Police	Foreign aid	Aff. action	Environment	Health care	Welfare
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Universalism in altruism	−1.33*** (0.10)	−1.09*** (0.09)	−0.97*** (0.09)	1.40*** (0.08)	0.84*** (0.09)	0.60*** (0.12)	0.32*** (0.11)	0.54*** (0.11)
Universalism in trust	−0.62*** (0.14)	−0.41*** (0.12)	−0.56*** (0.11)	0.36*** (0.09)	0.24*** (0.11)	0.30* (0.16)	0.27** (0.12)	0.59*** (0.15)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,971	10,971	10,971	10,971	10,971	10,971	10,971	10,971
R <sup>2</sup>	0.13	0.10	0.11	0.18	0.05	0.04	0.03	0.09

*Notes:* OLS estimates, robust standard errors in parentheses. In Panel A, the dependent variables are desired expenditure shares for each of the eight policy categories, normalized into z-scores within each country. In Panel B, the dependent variables are given by responses to the policy Likert scale questions, normalized by the sum of responses across Likert scale questions. These measures are also normalized into z-scores within each country. Universalism in altruism and universalism in trust are both in [0,1]. Controls include age, gender, income, wealth, college, urbanicity, religiosity, equity-efficiency preferences, altruism, trust, beliefs about the efficiency of government, and beliefs about whether one will personally benefit from government expenditure in each domain. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

universalism measure, we highlight that our data allow us to separately consider preferences and trust beliefs. Table 2 summarizes the results of OLS regressions, in which we link a respondent's policy views (normalized into z-scores) to their universalism in altruism and trust, controlling for our full set of covariates. In panel A, we investigate desired expenditure shares, which directly corresponds to Figure 6. In panel B, the dependent variables are the analogous Likert scale questions, again normalized into measures of relative support across policy domains. All regressions include the full set of controls discussed above.

We find that the coefficients of universalism in altruism and universalism in trust are almost always statistically significant and sizable in magnitude. This suggests that even though universalism in altruism and universalism in trust are positively correlated, they each capture distinct variation that is relevant for understanding policy views.

The coefficient magnitudes are consistently larger for universalism in altruism than for universalism in trust, which suggests that group-specific altruism weights are more important for policy views than group-specific trust levels. We view it as one strength of our approach that we can separately quantify the importance of beliefs and preferences, while for example the psychological approach of Haidt (2012) does not allow such conceptual distinctions.

Comparing magnitudes across policy domains, for both the desired expenditure share data and for the Likert scale questions, we find that universalism appears to be most important for




FIGURE 7

The figure plots the OLS regression coefficients of univariate regressions of desired arsinh expenditure levels for each policy domain on composite universalism. Universalism is in  $[0,1]$  and the dependent variables are standardized into z-scores within each country. Error bars indicate 95% confidence intervals using robust standard errors. The “All countries” specification includes country-fixed effects.

foreign aid, affirmative action, border control, military and police, and law enforcement. We did not pre-register predictions about differential magnitudes across policy domains, and we can only speculate about what drives them. One potential reason is that some policies (like foreign aid) are primarily about in-group vs. stranger tradeoffs, while others (such as welfare) also involve many other considerations and preferences, such as one’s own relative income or views on self-sufficiency.

We consistently find that (1) the magnitude of the OLS coefficient of universalism is 50–100% larger for the Likert scale questions and (2) the R-squared is also usually twice as large. A plausible interpretation of this is that measurement error in the Likert scale questions is lower than in the more complex quantitative desired expenditure amounts questions. Another contributing factor is that the Likert scale questions have fewer outliers. While we winsorized the desired expenditure data at  $\pm 3$  SD of the mean, this still generates large outliers that reduce the variance explained.

*Desired expenditure levels.* Figure 7 reproduces the left panel of Figure 6, except that now the dependent variables are desired (arsinh) expenditure levels rather than implied shares. Again, we standardize desired spending into z-scores within each country to keep the magnitudes comparable across countries. The results show that universalists desire higher government spending in the canonical left-wing policy domains, yet lower government spending in the canonical conservative domains. Thus, universalists do not *always* desire higher government spending, but rather only in those policy domains that we argue have a universalist “default implementation” (see Section 3).

### 5.3. *Benchmarking exercises*

In the following, we benchmark the link between universalism and policy views against the link between universalism and other individual characteristics, showing that universalism is quantitatively substantially more important for explaining policy views and ideological constraint.

*Benchmarking of “effect sizes”.* We first compare our estimated magnitudes of the link between universalism and policy views with the link between policy views and people’s left-vs.-right ideology (0–10). We do not do this because we believe that all variation in left-vs.-right ideology should itself be considered primitive—as discussed above, we suspect that when people say they are “left” they partly mean that they are universalists. Instead, we view people’s self-reported ideology as close to an upper bound for how much universalism can possibly explain, given how strongly people’s policy views are clustered around their political ideology. We additionally consider income as a benchmarking variable, given its salience in political economy analyses.

Figure 8 summarizes the results. Each coefficient stems from a separate regression of desired expenditure shares on either composite universalism, people’s left–right self-assessment, or their income. To aid comparability, all variables are expressed as z-scores, and the signs on the coefficients on income are all negated (i.e. they show the expected change in desired expenditure shares given a one standard deviation *decrease* in our income index). We see that, on average across policy domains, the point estimate of universalism is about two-thirds as large as that of people’s self-reported political ideology, while that of our income index is only about one-fourth as large. We view these results as supportive of the idea that universalism quantitatively matters for policy views.<sup>11</sup>

*Benchmarking of contribution to ideological constraint.* In a second step, rather than asking how much universalism matters for policy views considered in isolation from each other, we now study how much it matters for the *intracorrelations* among policy views observed in the data. To study this, we ask by how much the correlations among policy views (ideological constraint) would decrease if there was no heterogeneity in universalism. That is, we aim to construct a counterfactual measure of ideological constraint for a world in which everyone has the same degree of universalism.

To this effect, we proceed in two steps. First, we construct a measure of predicted ideological constraint.<sup>12</sup> We compute the fitted values of regressions of desired expenditure shares on a set of predictors. The average intracorrelation of these predicted policy views tells us how strongly the predictable component of policy views are correlated in our data, on average across all policy domains.<sup>13</sup> Second, to derive a quantitative indication of the importance of universalism, we repeat the entire procedure, except that the predicted policy views are computed after collapsing variation in universalism to the sample mean. Thus, this measure tells us how strongly the predictable component of policy views would be intracorrelated if everyone had the same degree

11. Supplementary Appendix Figure 21 shows the variance explained of each of the three predictors across policy domains.

12. For the purpose of this analysis, we focus on predicted rather than actual ideological constraint because our main interest is in understanding how much universalism matters for patterns that are, at least in principle, explainable based on individuals’ characteristics.

13. To compute this measure of average predicted intracorrelation, we recode the correlations such that they are positive when the prevailing structure of ideology predicts it, and as negative otherwise. For example, if predicted views on welfare and military were negatively correlated, we would recode the correlation as positive. If predicted views on welfare and foreign aid were negatively correlated, we code it as negative in computing average correlations.




FIGURE 8

This figure plots coefficients from separate regressions of desired expenditure shares on respondents' left-right self-assessment, composite universalism, and income, controlling for country fixed effects. All variables are standardized into z-scores within each country. Error bars indicate 95% confidence intervals using robust standard errors. Sign convention: the coefficients on income are all negated for easier comparability of magnitudes across predictors. As such, they should be interpreted as the expected change in desired expenditure share for each policy domain following a one standard deviation *decrease* in income.

of universalism. We perform the same procedure by collapsing other demographic variables to the sample mean, as natural benchmarks. An obvious limitation of such an analysis is that most variables are measured with error—as a result, even if universalism was the only determinant of ideological constraint (which we do not claim it is), we would still see intracorrelations among policy views when all heterogeneity in (partly mismeasured) universalism is collapsed.

Figure 9 summarizes the results. The left panel shows the results from predicting policy views based on universalism and the set of individual characteristics discussed above. The leftmost bar shows how strongly predicted policy views are correlated in our data when none of the variables is collapsed to its sample mean. The second bar shows the average intracorrelation when age is collapsed to its sample mean, and equivalently for the other variables. We see that heterogeneity in universalism is the only variable that visibly reduces predicted ideological constraint, from around 0.57 to 0.43 (which corresponds to 25%). All other variables, including measures of income, wealth, religiosity, and beliefs in government efficiency, do not appreciably contribute to predicted ideological constraint, and sometimes collapsing variation in these variables even *increases* the observed intracorrelations.<sup>14</sup>

14. This happens, for example, when a variable is positively correlated with one policy view and negatively with another policy view, even though the two policy views “should” be positively correlated from the perspective of today’s political clusters.

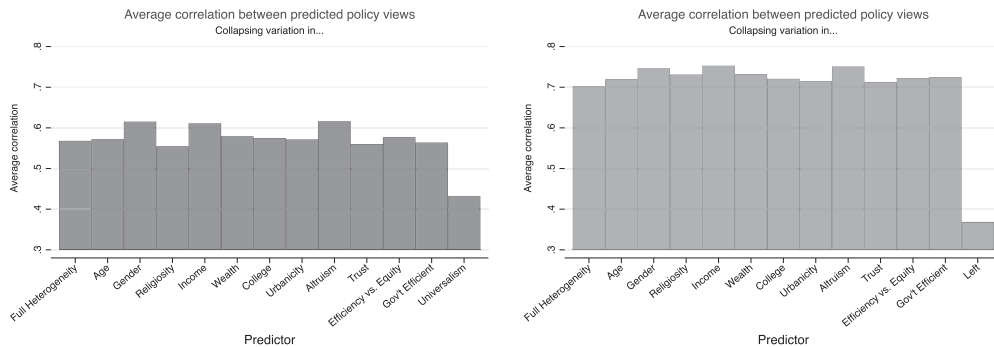


FIGURE 9

This figure plots counterfactual, average correlations between desired expenditure shares across policy domains when we collapse variation in demographics, beliefs, and preferences one at a time. The figure is constructed as follows. First, we regress the z-scores of desired expenditure shares for each of our eight policy domains on all of the variables shown in the left panel (or right panel), controlling for country FE. We then obtain the fitted values from these regressions when we collapse one of the variables in the regression at a time. Finally, we compute the correlation between these predicted desired expenditure shares for each policy domain (including country-fixed effects). We average these correlations across policy domains (recoding correlations so they are positive if in the hypothesized direction and negative otherwise), so that we obtain *counterfactual* clusters of ideology when the effect of heterogeneity in a given demographic, preference, or belief is removed. In the left panel, we document how predicted ideological constraint changes when we collapse variation in these variables in a regression that includes composite universalism. In the right panel, the regression instead includes the left–right self-assessment.

The right panel provides another comparison, by repeating the same analysis except that now predicted policy views are computed based on people’s self-reported left–right assessment, in combination with the same individual characteristics as in the left panel. Again, we here view people’s left–right self-assessment as a useful comparison because it is a plausible upper bound for how much can be explained with (noisy) survey measures of individual characteristics. There are two main takeaways. First, predicted ideological constraint increases by about 25% when self-reported ideology rather than universalism is taken into account. Second, predicted ideological constraint decreases by about 48% when variation in ideology is collapsed to its sample mean.

Our takeaway from this analysis is that universalism contributes substantially more to the phenomenon of ideological constraint than all other “primitive” individual characteristics in our data. Universalism does not explain as much as self-reported ideology, which is plausibly an upper bound for how much could in principle be explained based on individual characteristics.

*Do other individual characteristics produce the observed policy clusters?* In a final step of our benchmarking analyses, we study whether other individual-level variables also produce the characteristic pattern shown in Figure 6: positive correlations with the five “left-wing” policies and negative correlations with the three “right-wing” policies.

Given that our dataset includes a large set of individual characteristics, all of which could matter either in isolation or in the form of interaction effects, we proceed by selecting a set of candidate variables using simple machine learning techniques. Specifically, we implement a LASSO regression of respondents’ left–right self-assessment on all individual characteristics except universalism, in a fully interacted model. In other words, we do not only allow individual characteristics to matter but also their (potentially higher-order) interactions with each other. We then select those 10 regressors that exhibit the largest regression coefficients (all variables




FIGURE 10

Benchmarking analyses. We report the standardized beta coefficients and confidence intervals for regressions of arsinh desired expenditure level for a policy domain on different individual-level characteristics, conditional on country fixed effects. These characteristics were selected through a LASSO regression of people's left-vs.-right self-assessment on all possible interactions of standardized demographics, beliefs, and preferences. Out of those variables selected by the regression, we plot the top 10. Each estimate corresponds to a separate regression. To obtain residual altruism and trust, we respectively computed the residuals of dictator game allocations and generalized trust with respect to universalism.

were standardized into z-scores). The selected regressors are religiosity, wealth, residual trust, college degree, efficiency–equity preferences times beliefs about government efficiency, age times beliefs about government efficiency, income times beliefs about government efficiency, income, male times beliefs about government efficiency, and male. Note that while the majority of these regressors turn out to be simple individual characteristics, we did not exogenously impose this—it just turns out that in the data these individual characteristics are more important than their interactions.

In Figure 10, the leftmost panel serves as a reminder about the patterns for universalism. We find that none of the other ten individual characteristics that got picked by the LASSO produces the characteristic pattern that universalism successfully reproduces, at least in a statistically or even economically significant manner. In other words, other variables are often significantly correlated with policy views in meaningful and known ways—we are not trying to argue that they are unimportant for understanding policy views. However, our results show that they do not generate the characteristic internal structure of ideology that we are interested in here.

#### 5.4. Robustness checks

Supplementary Appendix contains five sets of robustness checks. First, some of our main analyses employed the composite measure of universalism. As specified in our pre-registration, Supplementary Appendix C.4 shows that very similar results hold if we work with universalism in altruism or universalism in trust separately.

Second, to tentatively probe the time stability of our results, Supplementary Appendix C.5 documents that the link between universalism and policy views is very similar among young and old survey participants.

Third, as we pre-registered, we employ instrumentation strategies from Gillen *et al.* (2019) to address the effects of measurement error in our elicitations of policy views and universalism. Results using multiple elicitations for both outcome and explanatory variables are very similar, see Supplementary Appendix D.

Fourth, we contracted with *Dynata* for  $N=1700$  respondents in each country, stratified to match the population on a number of dimensions. In Supplementary Appendix C.6, we replicate the analysis using these more representative samples, with very similar results.

Fifth, we address potential concerns over multiple hypothesis testing that may arise from our procedure of linking universalism to a vector of eight policy views (though we note that each of these analyses was pre-registered). To reduce such concerns, we first collapse the eight policy variables into a summary statistic, such that only one set of regressions needs to be run. This summary statistic was also pre-registered and is given by the average desired expenditure share for the five left-wing domains minus the average desired expenditure share for the three right-wing domains. Supplementary Appendix Table 11 shows that this summary statistic is strongly correlated with both universalism in altruism and universalism in trust.

Finally, as an additional remedy against multiple testing concerns, Supplementary Appendix C.8 implements the false discovery rate (FDR) procedure proposed by Benjamini, Krieger and Yekutieli (2006) and Anderson (2012). The adjusted  $p$ -values that result from this procedure are very similar to the unadjusted ones reported above.

## 6. DEMAND- AND SUPPLY-SIDE EXPLANATIONS

In trying to better understand the structure of policy views in the Western world, this article focuses on a demand-side explanation: that universalism systematically shapes (dis)approval of policies. We recognize, however, that supply-side mechanisms could also be driving some of the observed patterns. As posited in the political behaviour literature, voters might identify with a political party and then simply follow the party line on different policy domains. While it is not entirely clear why such a supply-side story should generate the particular clusters that we observe (in an identical fashion across countries), supply and demand-side mechanisms could also interact in important ways. For instance, it is conceivable that heterogeneity in universalism determines people's policy views on "core" issues and therefore drives the sorting of voters into parties, but that voters simply follow the party line in other domains. Or, political views may exert a feedback effect on basic moral values (Hatemi, Crabtree and Smith, 2019), which suggests complex interactions between the supply- and demand-side of politics. We seek to provide additional evidence that underscores the relevance of the demand-side mechanism we propose, rather than to rule out that the supply side also matters.<sup>15</sup>

15. One piece of suggestive evidence that supply-side mechanism may play a role stems from our own surveys. In addition to the Western countries for which we reported the results above, we also implemented the same survey in Brazil and South Korea. As is well-known in comparative political science, intracorrelations among policy views outside the cultural West do not follow the same patterns as in the West and cannot be neatly organized according to a left-vs.-right axis. We therefore pre-registered that the link between universalism and policy views would be different in Brazil and South Korea. Supplementary Appendix C.9 replicates our main analysis for these two countries. The relationships between universalism and policy preferences are all weaker in magnitude and sometimes opposite in sign relative to those observed in Western countries. While these results do not allow us to disentangle between supply- and demand-side mechanisms, they are at least consistent with the view that a difference between Western and non-Western political

### 6.1. *Specific policy implementations*

As we discussed in Section 3, many of the policy domains we consider entail both universalist and communitarian elements, or could be implemented in more or less universalist ways. Hence, our hypotheses implicitly rely on the assumption that there exist *default implementations* for these policies. In the following, we leverage that policies often contain both universalist and communal elements and propose to respondents specific implementations of policies that are more or less universalist in nature. One way of thinking about these empirical exercises is that they are meant to “overwrite” people’s mental default implementations of the different policies.

To do so, our survey additionally elicited respondents’ desired government spending level for specific policy proposals (two for each broad policy domain). The general objective of this exercise was to manipulate the implementation of a policy, such that one was unambiguously more universalist than the other. We then investigate whether self-reported left-wingers become more supportive of canonical conservative policies once they are implemented in a more universalist fashion, and whether right-wingers become more supportive of canonical left-wing policies once they are implemented in a communitarian fashion.

To systematically manipulate the implementation of a policy, we manipulated the social distance between the respondent and the beneficiaries of a policy. For welfare, environmental protection and health care, we implemented this by manipulating *geographic scope*. To illustrate, consider the example of welfare payments. We elicited desired spending levels separately for (1) “Redistributing local tax revenues as welfare payments across all communities nationwide” and (2) “Redistributing local tax revenues as welfare payments only within the local communities they were raised.”<sup>16</sup> We emphasized to respondents that every other aspect of the policy, e.g., collection of money and efficiency of redistribution, was identical between the two implementations. Thus, the two policies plausibly only differ in universalism-related aspects: who the money goes to (which concerns universalism in altruism) and whose tax cheating matters (which concerns universalism in trust). We used similar language to elicit the desired spending levels for health care that is funded through local vs. national tax revenues, and for local environmental protection vs. preventing global climate change, see Supplementary Appendix Table 17.

For foreign aid and affirmative action, we likewise manipulated the social distance between respondent and beneficiaries of a policy, except that this was done through other means than geography. For foreign aid, we separately elicit desired spending for (1) “Sending foreign aid to countries that are in most need of help” and (2) “Sending foreign aid to foreign countries that are our international allies”, which manipulates social distance. Similarly, for affirmative action, we separately elicit support for policies that ensure that (1) “no individual is disadvantaged” vs. that “no one of your same background (e.g. gender, ethnicity, or ancestry) is disadvantaged.” Again, this manipulates the universalist appeal of the policy domain because the latter proposal is directed at improving the welfare of in-group members.<sup>17</sup> See Supplementary Appendix Table 17. We acknowledge that not all of these specific policy implementations are as tight as the ones for welfare, health care, and environmental protection, in the sense that they sometimes change more than one aspect of a policy at the same time. However, we nonetheless view these exercises as

competition is that parties in the West more aggressively frame political issues around themes related to universalism. This may of course itself be partly driven by parties’ recognition of demand-side forces related to universalism.

16. Figure 5 in Supplementary Appendix B.2 provides a screenshot.

17. We process the data as follows. First, we again first perform a PPP conversion to USD. Second, because there are some huge outliers in the free text entry format, we delete all responses that are larger than 500 times the within-country median response. This affects 0.2% of responses. Third, as pre-registered, we winsorize the data at  $\pm 3SD$  of the within-country mean, which affects 1.3% of all responses. Fourth, we take the inverse hyperbolic sine of these winsorized PPP amounts, and finally standardize these variables into z-scores, separately within each country.



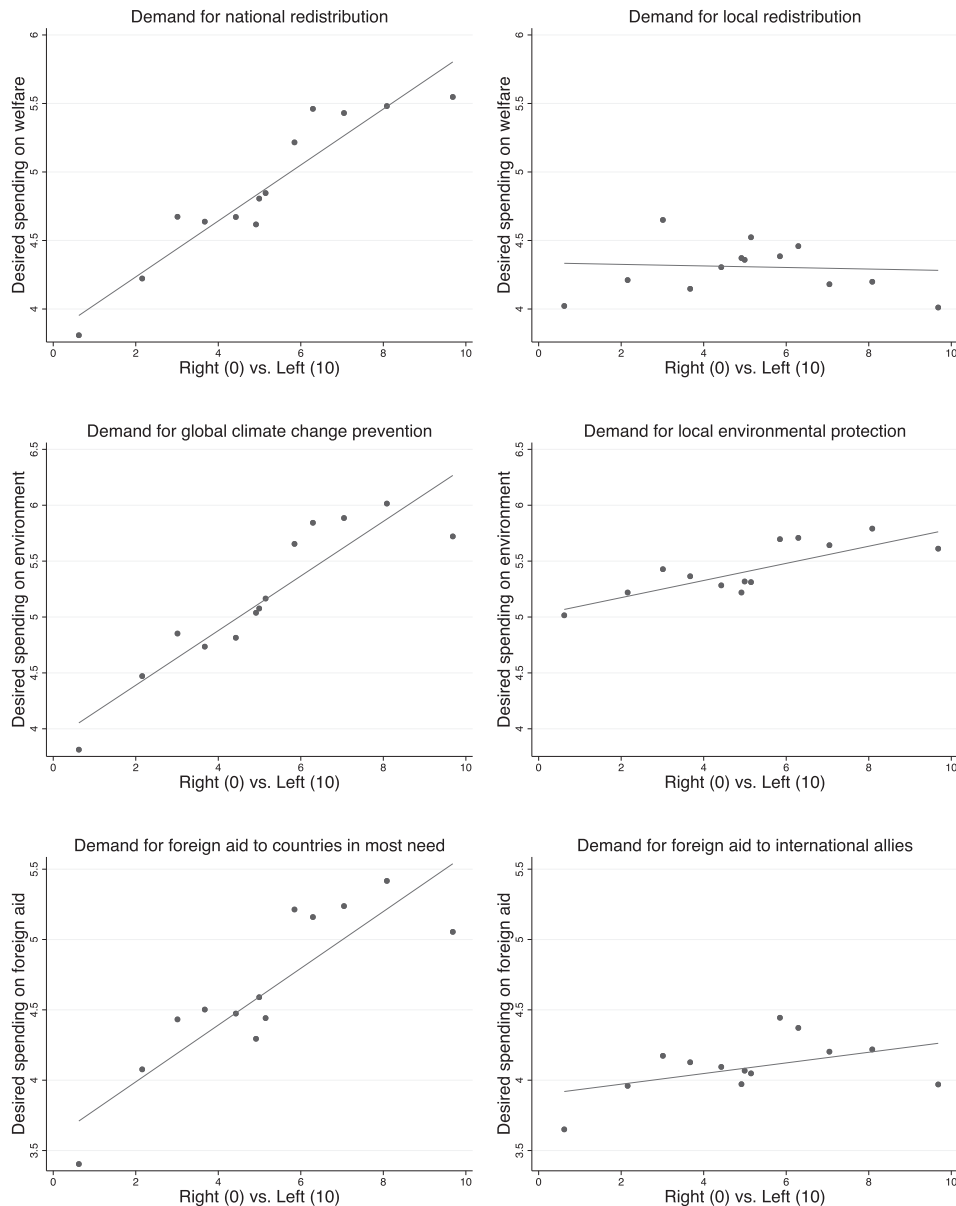


FIGURE 11

This figure shows binned scatter plots between arsinh desired spending levels for specific policy proposals and respondents' political ideology. In each row, the left panel shows the more universalist and the right panel the more communitarian implementation. The data are pooled across countries, and all panels are constructed controlling for country-fixed effects.

helpful because they suggest that making policies more universalist (in one or multiple ways) changes *who* supports a policy.

Figure 11 illustrates the results by showing binned scatter plots that link arsinh desired spending levels for specific policy proposals to respondents' self-reported left-vs.-right ideology (0–10), where higher values mean that the respondent considers himself more left-wing.

To conserve space, we show the results for welfare, environmental protection, and foreign aid; the results for healthcare and affirmative action look very similar, see Supplementary Appendix C.10. In each row, the left panel shows the link between left-vs.-right ideology and the more universalist policy implementation, while the right panel shows results for the more communitarian implementations.

Unsurprisingly, we see that a left-wing political orientation is strongly correlated with desired spending for national redistribution, global climate change prevention and foreign aid to the most needy. These are essentially the standard correlations that we see in the political discourse, as the more universalist policies arguably closely correspond to what we referred to as “default manipulation” (or default framing) in Section 3. However, much more interestingly, and as we hypothesized, we also see that self-identified left-wingers are no more supportive of redistribution, environmental protection and foreign aid once these policies are implemented in more communitarian ways. The partial correlations (conditional on country fixed effects) between the self-reported left-vs.-right scale and desired spending levels drops from  $\rho = 0.16$  for national redistribution to  $\rho = 0.00$  for local redistribution, from  $\rho = 0.20$  for global climate change prevention to  $\rho = 0.07$  for local environmental protection, and from  $\rho = 0.17$  for foreign aid to the most needy to  $\rho = 0.03$  for foreign aid to international allies.

In Supplementary Appendix C.10, we implement similar exercises for the traditional conservative spending categories of border control, military and police and law enforcement. In line with the results summarized above, we find that self-identified left-wingers become substantially more likely to support these expenditures once they focus on universalist objectives, such as peacekeeping by the military.

Overall, these patterns are suggestive that people’s policy views systematically depends on whether the policies are implemented in a universalist or communitarian fashion. This seems at odds with a pure supply-side account according to which people support certain government expenditure categories because they “learned” from their party that this is the right thing to do.

## 6.2. *Universalist minorities and support for affirmative action*

As a second strategy to document at least a partial role for demand-side explanations, we leverage that the role of moral universalism in some policy domains should be mediated by demographics. This is most evident for affirmative action, where, arguably, the link between universalism and support for affirmative action should be weaker among the beneficiaries of affirmative action (e.g. racial minorities, women).

To formally investigate this, in Columns (1) and (3) of Table 3, we regress support for affirmative action on moral universalism, interacted with a gender dummy. Columns (2) and (4) implement a similar analysis for racial minorities. Here, we restrict attention to the US, for which we arguably have a clear hypothesis about who the racial minorities are that stereotypically benefit from affirmative action (Hispanics, African Americans, and Native Americans).<sup>18</sup> According to our demand-side account, the coefficient on universalism should be smaller for the minority group, whereas a supply-side story would predict similar effects for minorities and non-minorities.

Overall, we see that the “effect” of universalism on affirmative action is weaker among racial minorities and women. We interpret these results as providing additional evidence for the partial relevance of a demand-side story.

18. In contrast, in the other countries in our sample, the race/ethnicity variable does not allow us to cleanly identify minorities/beneficiaries of affirmative action. For example, when a respondent in Germany indicates that s/he is not German but “Other European”, then this person may be a minority (because s/he may be a foreigner) but also a majority (because s/he is likely white).

TABLE 3  
*Universalism and support for affirmative action*

	Dependent variable: Z-scores of support for affirmative action:			
	Desired arsinh expenditure		Likert scale	
	(1)	(2)	(3)	(4)
Composite universalism	0.54*** (0.12)	0.90*** (0.18)	0.74*** (0.12)	1.06*** (0.17)
Female	0.044 (0.07)		0.21*** (0.07)	
Female × Universalism	−0.054 (0.17)		−0.37** (0.16)	
Racial minority		0.68*** (0.19)		0.59*** (0.16)
Racial minority × Universalism		−1.36*** (0.46)		−1.09*** (0.37)
Country FE	Yes	No	Yes	No
Controls	Yes	Yes	Yes	Yes
Observations	11063	2934	11063	2934
R <sup>2</sup>	0.05	0.09	0.19	0.28

*Notes:* OLS estimates, robust standard errors in parentheses. In Columns (2) and (4), the sample is restricted to US respondents. Racial minorities include African-Americans, Hispanics, and Native Americans. Universalism is in [0,1]. Controls include age, gender, income, wealth, college education, urbanicity, religiosity, equity-efficiency preferences, altruism, trust, beliefs about the efficiency of government, and beliefs about whether one will personally benefit from government expenditure in affirmative action. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

### 6.3. Replicating results controlling for political ideology

Finally, we attempt to shut down the supply-side channel by looking at the relationship between universalism and policy preferences *conditional* on political affiliation. To this effect, Supplementary Appendix C.11 replicates all results reported above, controlling for fixed effects for respondents' self-positioning on a 0–10 left-right scale. There are three results. First, universalism remains a statistically significant predictor of policy views in the same manner as identified in Section 5, see Supplementary Appendix Figure 32.

Second, in a difference-in-difference analysis, universalism is strongly correlated with the degree to which respondents prefer the more universalist implementation of a policy relative to the more communitarian one, conditional on political ideology fixed effects, see Supplementary Appendix Figure 33.

Third, and relatedly, Supplementary Appendix Figures 34 through 36 show the link between universalism and support for the specific universalist or communitarian policy implementations discussed in Section 6.1. We find that the relationship between policy views and universalism attenuates or reverses as we switch from a universalist to a more communitarian implementation of a policy, even when we include left–right fixed effects.

## 7. CONCLUSION

For decades, philosophers have argued about the scope of our moral obligations. Should we always be impartial? Do we have some special connection to those that are socially close to us, in a way that should make us trust them more and feel more morally responsible for them? These deep normative questions have direct empirical counterparts because many policies and government expenditure categories that are of central concern to economists implicitly or explicitly concern tradeoffs between strangers and those that are close to us. Yet, empirically, people exhibit large disagreement over what the morally appropriate way to treat different social groups is. This article

has proposed that this individual-level heterogeneity in universalism accounts for the particular structure of policy views that Western democracies have seen over the last few decades.

We believe that this empirical insight is valuable for at least three reasons. First, it helps us to understand where political differences come from. Second, it sheds light on how policy-makers can appeal to voters by proposing policies with a particularly universalist or communitarian flavour. Third, our article arguably does some heavy lifting in terms of a dimensionality reduction of political views: rather than merely explain attitudes on one topic (e.g. redistribution) with people's underlying moral preferences and beliefs, we endogenize a vector of eight (and perhaps more) policy views as a function of morality.

As discussed in Section 2, our analysis is conditional on two restrictions. First, we only analyse the structure of ideology as it has prevailed over the last 40–50 years. We do not have much to say about whether or how universalism mattered for policy in more distant history. This being said, there is some evidence that suggests that the relevance of universalism for politics has increased over time. The Democratic “loss of the South” and subsequent polarization were largely tied to ideas related to (non-) universalism (Kuziemko and Washington, 2018). Furthermore, Enke (2020) documents using text analyses that Republicans and Democrats used universalist vs. communal moral language in roughly equal frequencies until the mid-60's but steadily diverged thereafter, which could be understood as suggesting that heterogeneity in universalism is more relevant politically today than in the past.

Second, our analysis deliberately focused on the Western world. The connection between the structure of ideology and morality might be different outside the West for various reasons. Future research is needed to understand why the political systems of non-“WEIRD” societies (Henrich, 2020) do not give rise to the ideological clusters that dominate the West, and which role morality plays in these contexts.

*Acknowledgments.* For valuable comments and discussions we thank Nicola Gennaioli, four very constructive referees, Alberto Alesina, Roland Bénabou, Maxim Boycko, Matt Gentzkow, Ed Glaeser, Vincent Pons, Jesse Shapiro, Andrei Shleifer, Jim Snyder, Guido Tabellini, Marco Tabellini, and many seminar audiences. Phoebe Cai, Aditi Chitkara, Robert Fluegge, and Jan Knuf provided outstanding research assistance. We also thank Spencer Kwon, Joakim Semb, and Julien Senn for help with the surveys. We acknowledge funding from the William F. Milton Fund, the Harvard Asia Center, and a Harvard Inequality in America Grant. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1– 390838866. Funding by the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (Project A01) is gratefully acknowledged.

### Supplementary Data

Supplementary data are available at *Review of Economic Studies* online. And the replication packages are available at <https://dx.doi.org/10.5281/zenodo.6804787>.

### Data Availability Statement

All data and code that are used for this article are available as online replication package at <https://doi.org/10.5281/zenodo.6804787>.

### REFERENCES

- ALESINA, A. and GLAESER, E. (2004), *Fighting Poverty in the US and Europe: A World of Difference* (Oxford, UK: Oxford University Press).
- , BAQIR, R. and EASTERLY, W. (1999), “Public Goods and Ethnic Divisions”, *The Quarterly Journal of Economics*, **114**, 1243–1284.
- ALGAN, Y., BEASLEY, E., COHEN, D. and FOUCAULT, M. (2018), “The Rise of Populism and the Collapse of the Left-Right Paradigm: Lessons from the 2017 French Presidential Election” (CEPR Discussion Paper, DP13103).
- ANDERSON, M. L. (2012), “Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects”, *Journal of the American Statistical Association*, **103**, 1481–1495.

- ANSOLABEHERE, S., RODDEN, J. and SNYDER, J. M. (2008), "The Strength of Issues: Using Multiple Measures to Gauge Preference Stability, Ideological Constraint, and Issue Voting", *American Political Science Review*, **102**, 215–232.
- , ———, AND SNYDER, J. M. JR. (2006), "Purple America", *Journal of Economic Perspectives*, **20**, 97–118.
- BENJAMINI, Y., KRIEGER, A. M. and YEKUTIELI, D. (2006), "Adaptive Linear Step-Up Procedures that Control the False Discovery Rate", *Biometrika*, **93**, 491–507.
- BESLEY, T. and PERSSON, T. (2019), "The Rise of Identity Politics" (Working Paper, IIES, Stockholm University). <http://perseus.iies.su.se/~tpers/courses/politec2019/BesleyPersson190330.pdf>.
- BONOMI, G., GENNAIOLI, N. and TABELLINI, G. (2021), "Identity, Beliefs, and Political Conflict", *The Quarterly Journal of Economics*, **136**, 2371–2411.
- BROCKMAN, D. E. and BUTLER, D. M. (2017), "The Causal Effects of Elite Position-Taking on Voter Attitudes: Field Experiments with Elite Communication", *American Journal of Political Science*, **61**, 208–221.
- CHEN, D. L. and LIND, J. T. (2019), "The Political Economy of Beliefs: Why Fiscal and Social Conservatives/Liberals (Sometimes) Come Hand-in-Hand" (Working Paper).
- DAL BÓ, E., FINAN, F., FOLKE, O., PERSSON, T. and RICKNE, J. (2018), "Economic Losers and Political Winners: Sweden's Radical Right" (Working Paper).
- DRACA, M. and SCHWARZ, C. (2020), "How Polarised Are Citizens? Measuring Ideology from the Ground Up" (Working Paper).
- ENKE, B. (2020), "Moral Values and Voting", *Journal of Political Economy*, **128**, 3679–3729.
- , RODRÍGUEZ-PADILLA, R. and ZIMMERMANN, F. (2022), "Moral Universalism: Measurement and Economic Relevance", *Management Science*, **68**, 3590–3603.
- ERIKSON, R. S. and TEDIN, K. L. (2015), *American Public Opinion: Its Origins, Content and Impact* (New York, NY, USA: Routledge).
- GETHIN, A., MARTÍNEZ-TOLEDANO, C. and PIKETTY, T. (2022), "Brahmin Left versus Merchant Right: Changing Political Cleavages in 21 Western Democracies, 1948–2020", *The Quarterly Journal of Economics*, 2022, **137**, 1–48.
- GILENS, M. (2009), *Why Americans Hate Welfare: Race, Media, and the Politics of Antipoverty Policy* (Chicago, IL, USA: University of Chicago Press).
- GILLEN, B., SNOWBERG, E. and YARIV, L. (2019), "Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study", *Journal of Political Economy*, **127**, 1826–1863.
- GRAHAM, J., HAIDT, J. and NOSEK, B. A. (2009), "Liberals and Conservatives Rely on Different Sets of Moral Foundations", *Journal of Personality and Social Psychology*, **96**, 1029.
- GROSSMAN, G. and HELPMAN, E. (2018), "Identity Politics and Trade Policy" (Working Paper).
- GURIEV, S. and PAPAIOANNOU, E. (2022), "The Political Economy of Populism", *Journal of Economic Literature*, **60**, 753–832.
- HAIDT, J. (2012), *The Righteous Mind: Why Good People are Divided by Politics and Religion* (New York City, USA: Vintage).
- HATEMI, P. K., CRABTREE, C. and SMITH, K. B. (2019), "Ideology Justifies Morality: Political Beliefs Predict Moral Foundations", *American Journal of Political Science*, **63**, 788–806.
- HENRICH, J. (2020), *The WEIRDest People in the World: How the West Became Psychologically Peculiar and Particularly Prosperous*, (New York City, USA: Macmillan).
- KOZLOWSKI, A. and MURPHY, J. P. (2021), "Issue Alignment and Partisanship in the American Public: Revisiting the 'Partisans without Constraint' Thesis", *Social Science Research*, **94**. <https://doi.org/10.1016/j.ssresearch.2020.102498>.
- KUZIEMKO, I. and WASHINGTON, E. (2018), "Why Did the Democrats Lose the South? Bringing New Data to an Old Debate", *American Economic Review*, **108**, 2830–2867.
- LAKOFF, G. (2010), *Moral Politics: How Liberals and Conservatives Think* (Chicago, IL, USA: University of Chicago Press).
- LENZ, G. S. (2013), *Follow the Leader?: How Voters Respond to Politicians' Policies and Performance*, University of Chicago Press.
- LUTTMER, E. F. P. (2001), "Group Loyalty and the Taste for Redistribution", *Journal of Political Economy*, **109**, 500–528.
- POOLE, K. T. and ROSENTHAL, H. (2000), *Congress: A Political-Economic History of Roll Call Voting* (Oxford, UK: Oxford University Press on Demand).
- PUTNAM, R. D. (2000), *Bowling Alone: The Collapse and Revival of American Community* (New York, USA: Simon and Schuster).
- SCHEVE, K. and STASAVAGE, D. (2006), "Religion and Preferences for Social Insurance", *Quarterly Journal of Political Science*, **1**, 255–286.
- SHAYO, M. (2009), "A Model of Social Identity with an Application to Political Economy: Nation, Class, and Redistribution", *American Political Science Review*, **103**, 147–174.
- SOWELL, T. (2007), *A Conflict of Visions: Ideological Origins of Political Struggles* (New York, USA: Basic Books).
- TABELLINI, G. (2008), "The Scope of Cooperation: Values and Incentives", *Quarterly Journal of Economics*, **123**, 905–950.
- WAYTZ, A., IYER, R., YOUNG, L., HAIDT, J. and GRAHAM, J. (2019), "Ideological Differences in the Expanse of the Moral Circle", *Nature Communications*, **10**, 1–12.
- WU, A. (2020), "Political Polarization" (Cambridge, MA, USA: Undergraduate Senior Thesis, Harvard University).