NBER WORKING PAPER SERIES

MORAL UNIVERSALISM AND THE STRUCTURE OF IDEOLOGY

Benjamin Enke
Ricardo Rodríguez-Padilla
Florian Zimmermann

Working Paper 27511
http://www.nber.org/papers/w27511

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2020

We acknowledge funding from the William F. Milton Fund, the Harvard Asia Center, and a Harvard Inequality in America Grant. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Moral Universalism and the Structure of Ideology
Benjamin Enke, Ricardo Rodríguez-Padilla, and Florian Zimmermann
NBER Working Paper No. 27511
July 2020
JEL No. D01,D72

ABSTRACT

Throughout the Western world, people's policy views are correlated across domains in a strikingly similar fashion. This paper proposes that what partly explains the structure of ideology is moral universalism: the extent to which people's altruism and trust remain constant as social distance increases. In new large-scale multinational surveys, heterogeneity in universalism descriptively explains why the left and right both simultaneously support and oppose different types of government spending. Moreover, the left-right divide on topics such as redistribution strongly depends on whether people evaluate more or less universalist policies. Large-scale donation data provide additional evidence for the political left's universalism.

Benjamin Enke
Department of Economics
Harvard University
Littauer M008
Cambridge, MA 02138
and NBER
enke@fas.harvard.edu

Ricardo Rodríguez-Padilla
Harvard University
rodriguezpadilla@g.harvard.edu

Florian Zimmermann
Institute on Behavior
& Inequality (briq)
Schaumburg-Lippe-Str. 5-9
53113 Bonn
Germany
and University of Bonn
florian.zimmermann@briq-institute.org

An online appendix is available at http://www.nber.org/data-appendix/w27511

1 Introduction

A key stylized fact in the study of political economy is the existence of *ideological constraint*: people's policy views are correlated across domains, so that an individual's self-identification as "left" or "right" carries information about an entire vector of policy views. Strikingly, the internal structure of these ideological clusters is very similar across Western countries. As we confirm using new large-scale survey data from multiple Western democracies, people in a left cluster generally desire government expenditure on foreign aid, affirmative action, environmental protection, welfare, and universal health care, while people in a right cluster always support government spending on the military, police and law enforcement, and border control. While these clusters appear to have become more pronounced over the last 40 years, the basic qualitative structure of ideology has been remarkably constant in recent history, across both time and space.

Yet, it is not immediately obvious why these *particular* bundles of policy views would prevail in the first place. A prominent view – which we confirm in our data – is that people differ in their overall preferences for "big government." However, views about the size of government as a whole do not rationalize why, in terms of expenditure shares, demand for redistribution is always correlated with demand for environmental protection rather than support for a strong military. A fortiori, the fact that the left desires a larger government overall does not explain why in some policy domains (such as law enforcement) the left actually demands a lower level of spending than the right. Still, the striking similarity of ideological clusters across countries with vastly different electoral systems and party structures suggests that these bundles reflect a systematic core rather than coincidence. This paper attempts to identify this core and to partly explain what it ultimately means to be "left" or "right," beyond the mechanical description of policy views associated with these labels.

Our central proposition is that what imposes the particular structure on the space of policy views is heterogeneity in *moral universalism*, by which we mean the extent to which people's altruism and their trust in others remain constant as social distance increases. Universalism is not about a person's overall *level* of altruism or trust, but instead about its *slope* as a function of social distance. Universalists are not more or less moral people, they just allocate a given altruism or trust budget more uniformly. Based on this definition, we first conceptualize the link between universalism and policy views in a simple model. We then test the resulting predictions in large-scale surveys to provide evidence that heterogeneity in universalism descriptively explains the structure of ideology observed in the Western world, in an almost identical fashion across countries. Using those same surveys, we also document that the canonical left-right divide on policy views substantially attenuates or even reverses once traditionally conservative policy

domains are recast as universalist policies, or once traditionally left-wing policies are implemented in non-universalist ways. Finally, we leverage large-scale donation data to provide complementary field evidence for the link between heterogeneity in universalism and political behavior. The entire paper is descriptive in nature and offers a new set of stylized facts.

To formalize how we think about the link between universalism and a vector of policy views, we first introduce a simple model that builds on Tabellini (2008). In the model, the key primitives are two parameters that govern an agent's universalism in altruism and in trust. Universalism in altruism determines the welfare weights that an agent assigns to other agents, as a function of social distance. Universalism in trust determines the extent to which trust in others declines as a function of social distance. Both universalist and non-universalist agents have rational expectations about the overall rate of cheating in society, yet they differ in their beliefs about who the cheaters are, where a full universalist believes that cheating is uncorrelated with social distance from her.

Agents evaluate two potential policies, where Policy B is "risky" in that it introduces a scope for free-riding or cheating by individual members of society. Policy A is "safe" in that it reduces the scope for cheating but is associated with other societal costs. For example, in the domain of welfare, Policy B corresponds to a system with more extensive welfare payments, which introduces scope for claiming benefits one is not entitled to. Policy A, on the other hand, corresponds to a smaller redistributive system with less scope for free-riding, yet this introduces the social cost that random income shocks cannot be equalized ex post. In this setup, less universalist agents oppose welfare because they are more likely to believe that the socially distant are likely to free-ride on the agent's in-groups. Thus, in the model, universalism in altruism and trust both lead to a stronger demand for welfare.

To further illustrate the logic of the model, consider the domain of police and law enforcement. Here, the risky Policy B corresponds to a system with less police presence, which introduces scope for stealing. The safe Policy A, meanwhile, corresponds to more police presence, which eliminates the scope for stealing but introduces the societal cost of paying for a law enforcement system. Here, less universalist agents again support the safe Policy A because they worry that their in-group members get exploited by socially distant agents. The takeaway is that, in our framework, less universalist agents (defined in both altruism and trust space) sometimes support and sometimes oppose government spending, purely depending on whether it introduces or prevents free-riding opportunities. While our formal model emphasizes the role of cheating, we also discuss informally how universalism in altruism could explain variation in policy preferences absent opportunities for cheating. The common thread that runs through our formal and informal applications is that a person's universalism should be predictive of their support for con-

temporary "left" policies.

We test our predictions about the link between universalism and a vector of policy views in pre-registered representative large-scale surveys in five Western countries: United States, Australia, France, Germany, and Sweden. We further include Brazil and South Korea as two non-Western countries in our sample. Non-Western countries typically do not exhibit the particular ideological clusters observed in the West, so that the link between universalism and policy should be weaker or absent in these countries. In total, we survey $N \approx 15,000$ individuals. We measure respondents' universalism in altruism and trust, along with their policy views.

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. Each respondent makes ten allocation decisions across which the social group (i) varies. The list of groups is based on an ex-ante crowdsourcing exercise and includes the respondent's extended family; neighbors; friends of the family; colleagues; members of the same organization; or people who share the respondent's hobbies; religious beliefs; age; political views; and race. For example, in one question, a U.S. participant is asked to split hypothetical $100 between a member of their extended family and a randomly-selected person from the United States. In addition to these 10 questions that measure "domestic universalism", we also measure "foreign universalism" and "global universalism" through money allocation tasks that involve different types of foreigners. From all of these questions, we construct an individual-level summary statistic of universalism in altruism. 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 et al., 2020).

Using an analogous procedure, we estimate respondents' universalism in trust by asking them to allocate 100 trust points between the individuals outlined above, to indicate whom the respondent trusts more. These questions again deliver measures of domestic, foreign, and global universalism in trust. 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." In our data, respondents exhibit large variation in universalism: some participants always split the money or trust points equally, while others consistently share more money with, and trust more, members of their own in-groups. Consistent with prior findings, universalism does not just reflect favorable economic conditions: if anything, individuals with higher income and wealth are less universalist.

We supplement these measurements of universalism with detailed questions on respondents' policy views. To this effect, we solicit quantitative responses about how much money the government should collect on average from each citizen to fund *specific expenditure categories*. Here, a respondent states a per capita dollar amount that they would like to see collected and spent on each of welfare payments; universal health care; affirmative action; military; law enforcement and police; border control; foreign aid; and environmental protection.

Using these data, the empirical analysis begins with a principal component analysis that analyzes the structure of ideology. In line with prior findings, we find that people's desired expenditure levels can be summarized by two intuitively appealing main components that are strikingly similar across Western countries: (i) a big-vs.-small government component that captures how much money respondents would like to spend overall and (ii) desired expenditure shares conditional on overall spending, which exhibit the familiar structure described in the opening paragraph.

Looking at the link between policy views and universalism, we find that the structure of this second component is strongly correlated with universalism in the ways predicted by the model and our pre-registration. Universalism is *positively* correlated with desired expenditure levels on welfare payments, environment, affirmative action, foreign aid, and – to a lesser extent – universal health care. Moreover, as we pre-registered, universalism is *negatively* correlated with desired expenditure levels 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 paper.

These correlations are robust and general in the following three ways. (i) The results are almost identical when we consider either universalism in altruism or universalism in trust, as predicted by our model. (iii) 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 United States, Australia, France, Germany, and Sweden. In the two non-Western countries in our sample, Brazil and Korea, where policy views generally cannot be grouped according to the Western left-vs.-right divide, heterogeneity in universalism does not explain much of the variation in policy views. At the same time, among the rich and well-educated elites in these countries, universalism is correlated with policy views in a very similar fashion to the patterns in Western countries. This may suggest that the role of morality for policy preferences is a "luxury good".

As has long been known, various sociodemographics, beliefs and preferences are correlated with the left-right divide. To put our results on universalism in perspective, we implement a series of benchmarking exercises against variables such as age, religiosity, education, income and wealth, equity-efficiency preferences and beliefs about the effi-

ciency of government. In our data, these variables are all reasonably strongly correlated with respondents' self-positioning on a left-vs.-right scale, which suggests that we measure them in meaningful ways. We also find that these variables are often correlated with desired expenditure levels in important and known ways. At the same time, universalism is the only variable in our data that 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. While many other variables plausibly affect policy views on single or multipe issues, none of them gets close to (correlationally) producing the characteristic structure of ideology that is our focus here.

In a next step, we make the link between universalism and policy views more direct by manipulating people's support for broad policy domains by proposing specific (non-) universalist implementations of different broad policies. To take an example, universalists may well be in favor of specific universalist policies within the general domain of the military, and non-universalist conservatives may be supportive of redistribution once it is implemented in a local, "communal" fashion. We elicit respondents' desired spending levels for specific policy proposals within each broad policy domain, where some proposals are more universalist than others. For example, within the domain of the military, we separately elicit desired spending levels on "Peacekeeping and humanitarian missions by the military abroad" and "Ensuring American defense and security." Likewise, within the broad domain of welfare payments, we separately elicit desired spending levels 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."

In these exercises, the relationship between universalism and policy views can be predictably attenuated or even reversed, depending on whether the specific policy proposal is more or less universalist. To take a few examples, universalists (left-wingers) are more supportive of military expenditure than non-universalists (right-wingers) once the military is said to focus on humanitarian missions. Similarly, conservatives are equally likely to support redistribution or environmental protection as left-wingers once it takes place locally. These results further strengthen the empirical case for the idea that what matters for the support of a policy is at least partly whether it is universalist in nature.

In the final part of the paper, we complement the survey analysis with field evidence. We estimate the universalism of U.S. Congressional Districts (CDs) using large-scale donation data from DonorsChoose, an American non-profit organization providing an online "crowdfunding" platform for public school teachers. On this website, individual donors give money to specific funding requests that are posted by teachers. As a proxy

for aggregate universalism, we estimate the extent to which a CD's donations decline as a function of the geographic distance to the recipient CD. As in our surveys, we only leverage variation in *towards whom* a given donor CD donates, not how much they donate (or receive) overall.

We find that a CD's universalism is strongly correlated with Democratic vote shares: Republican CD's donate relatively more money locally and less money to faraway places. That is, as in our surveys, left-wingers tend to treat their local community relatively poorly also in terms of actual donations. This raw correlation is robust against leveraging only within-state variation, and against controlling for variables such as local education expenditure or income.

Linking our work to the literature, much research in political science 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. Popular accounts often distinguish between an "economic" and a "cultural" or "social" axis, yet these descriptive classifications do not explain (i) why economic and social views are correlated in systematic ways and (ii) why certain types of social views tend to go together.

Various literatures in economics, political science, and moral and political psychology have highlighted the role of morality, identity and social preferences for political attitudes, though none of them attempts to explain the internal structure of ideology. Enke (forthcoming) studies the supply of and demand for universalist vs. communal moral values in U.S. presidential elections using a psychological (non-utilitarian) framework of moral values (Haidt, 2012). We innovate on this work (i) by examining not just voting behavior but the internal structure of specific policy views; (ii) not just in the U.S. but in the Western world more generally; and (iii) by operating with a utilitarian framework of universalism and corresponding measurements.[1] Much of our approach is inspired by the model in Tabellini (2008).

The idea that social groups and identity play an important role in politics runs through various recent contributions and reviews (Shayo, 2009; Grossman and Helpman, 2018; Gennaioli and Tabellini, 2019; Kranton and Sanders, 2017; Besley and Persson, 2019; Guriev and Papaioannou, 2020). For instance, large literatures explain variation in demand for redistribution through ethnic divisions and citizenship (Alesina et al., 1999, 2018; Luttmer, 2001; Alesina and Glaeser, 2004; Gilens, 2009; Fehr et al., 2019), or social preferences (Kerschbamer and Müller, 2020; Epper et al., 2020; Fisman et al.,

[1]A popular view in political psychology is that political affiliation correlates with "negativity bias" (Hibbing et al., 2014) or "threat sensitivity" (Jost et al., 2009). Waytz et al. (2019) show that U.S. liberals express greater moral concern toward friends relative to family, and the world relative to the nation. We believe that our notions of universalism in altruism and trust can be usefully understood as capturing many of the more fine-grained psychological concepts that have received attention in this literature.

2017). Similarly, the broader concept of social capital has received substantial attention in the political economy literature (Putnam, 2000). For example, Dal Bó et al. (2018) and Algan et al. (2018) document that far-right voters exhibit lower trust. We differ from all these contributions in that we emphasize the relevance of universalism (the gradient of social capital, rather than its level) for an entire vector of policy views.[2]

Finally, we interpret our results as linking to a recent broader social science literature that emphasizes the importance of affective polarization (Iyengar et al., 2019; Boxell et al., 2020) and the role of emotions and morality in political disagreement (Haidt, 2012). Some of the hostility in political conflict may stem from people having a hard time understanding that those on the other side of the aisle are not selfish but instead act on different moral priorities, emphasizing either those that are close to them or impartial treatment.

The remainder of the paper proceeds as follows. Section 2 summarizes the internal structure of ideology. Section 3 offers a formal framework. Sections 4 and 5 describe the design and results of our surveys. Section 6 offers field evidence and Section 7 concludes.

2 The Structure of Western Political Ideology

To illustrate our motivating observation on the structure of political ideology in rich Western societies, we work with our own survey data, described in detail in Section 4. The data cover the United States, Australia, France, Germany, and Sweden, along with the non-Western countries Brazil and South Korea, for a total of approximately 15,000 respondents. We elicited respondents' desired per capita expenditure levels for eight domains: welfare payments; universal health care; affirmative action; environmental protection; foreign aid; military; police and law enforcement; and border control. That is, respondents provided a per capita amount that they would like their national government to collect and spend on each of these domains.

To probe the correlation structure of policy views, we implement principal component analyses (PCA) separately in each country. The first principal component (first eigenvector) is that convex combination of the underlying variables that accounts for as much variation in the data as possible. It hence assigns similar weights to highly correlated variables. The second principal component is that convex combination of the underlying variables that explains as much of the residual variation as possible, conditional on being orthogonal to the first eigenvector.

We find that, in each Western country, the first principal component of (log) desired expenditures across domains exhibits an unsurprising and almost identical structure: it

[2]More generally, our paper links to a recent literature on the economics of morality (e.g., Bénabou et al., 2020; Dal Bó and Dal Bó, 2014; Bénabou and Tirole, 2011).

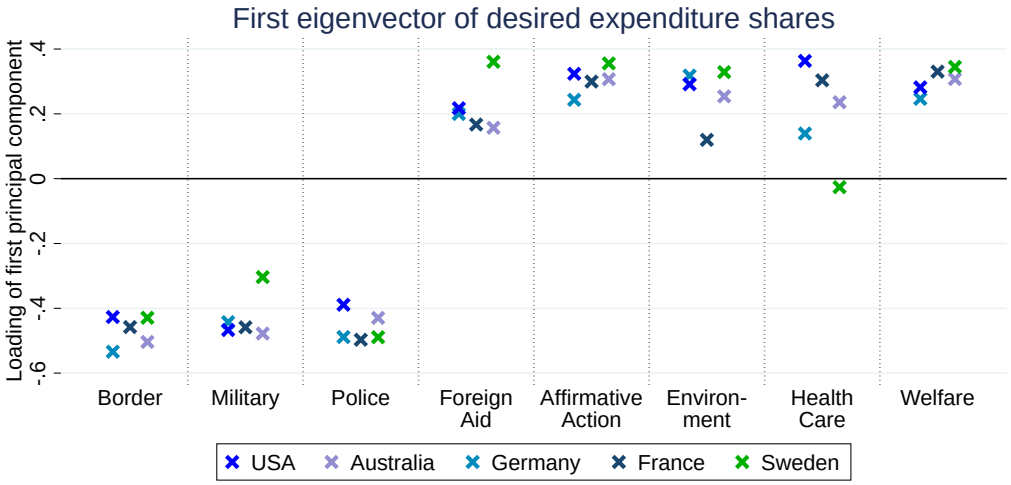

Figure 1: Factor loadings of the first principal component of desired expenditure shares. Sign convention: the loading on "Border" is always non-positive, and the other signs are determined accordingly.

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, on the other hand, 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 desired expenditure levels for 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, computed as desired expenditure level in a given domain divided by total desired expenditure on all eight domains. Figure 1 presents the loadings of the first principal component for the Western countries. Border control, military, and police and law enforcement all receive negative weights in each country, while foreign aid, affirmative action, environmental protection, welfare payments, and universal health care almost always receive positive weights.

The structure of this eigenvector is reminiscent of intuitive notions of "left" and "right." To confirm this intuition, we elicited from our respondents how they would position themselves on an 11-point left-vs.-right Likert scale. Figure 2 summarizes the relationship between respondents' self-positioning and their desired expenditure levels. In all Western 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

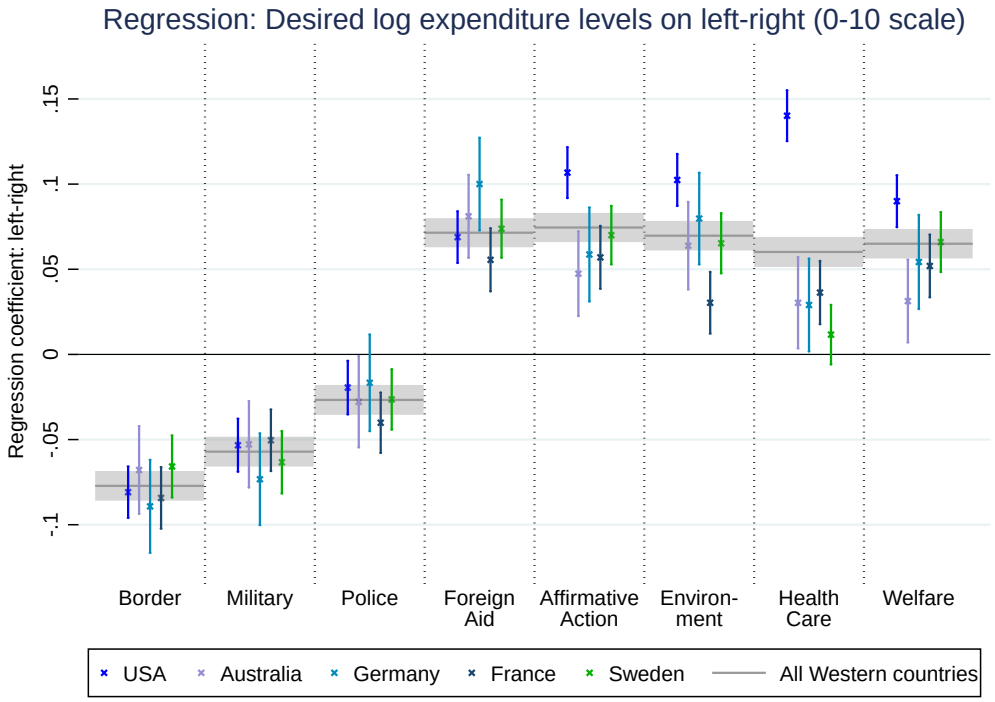

Figure 2: The figure plots the OLS regression coefficients of univariate regressions of desired log expenditure levels for each policy domain on self-positioning on a left-right scale (0–10). The dependent variables are standardized into z-scores. Error bars indicate 95% confidence intervals using robust standard errors. The "All western countries" specification includes country fixed effects.

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 U.S. 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 U.S. and $\rho = 0.30$ in Australia. This suggests that at least a considerable 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. The objective of this paper is to understand why policy views exhibit this particular correlation structure, in a strikingly similar fashion across Western countries.

Two comments on the scope of our analysis are in order. First, we attempt to understand the structure of contemporary, rather than historical, ideology. This being said, while recent research suggests that the magnitude of the intra-correlations between people's policy views has increased over the last 40 years, the qualitative structure of ideology has remained remarkably constant over this period (e.g., Rehm and Reilly, 2010; Kozlowski and Murphy, 2019; Wu, 2020; Draca and Schwarz, 2020).

Second, we only attempt to understand the structure of Western, rather than global, ideology. Figures 19 and 20 in Appendix C.4 replicate the analyses above for the two non-Western countries in our sample. Similarly to the results found in other survey datasets (Malka et al., 2019), we see that the structure of policy views outside the West is considerably less pronounced, and there is no clear relationship with people's left-vs.-right self-assessment.[3]

3 Theoretical Framework

This section develops a simple framework to clarify how we think about the relationship between policy preferences and universalism in both altruism and trust. Our setup builds on Tabellini (2008). In the model, agents choose between two policy options, yet we will argue that the structure of these two policies captures an essential feature of all eight policy domains discussed in the previous section.

The main ingredients of the model are: (i) agents live on a rectangle and hence at different (social) distance to different members of humanity, where social distance could capture distance along the lines of family, ethnicity, religion, language, values, geography etc.; (ii) the two policy options differ in the extent to which they enable or rule out free-riding or cheating; (iii) in terms of timeline, agents first vote on a policy and then decide whether to cheat on society; and (iv) agents differ in the extent to which their altruism and trust are universalist. Our object of interest is how an agent's universalism affects their choice between the two policies. We relegate derivations to Appendix A.

3.1 Social Distance and Preferences

Let I be a set consisting of N agents from two separate countries, where for analytical convenience we assume that $N \to \infty$ and that N is a multiple of four. We formalize countries and social distances by allocating agents in equal proportion to the vertices of a rectangle of length d_l and width d_w where $d_w < d_l$ and $d_w + d_l = 1$. The social distance between agents i and j is $d_{i,j}$, where distance is measured along the edges of the rectangle. Agents in the same country are connected by the short end of the rectangle.

We assume that each vertex of the rectangle corresponds to a social group. Agents who populate the same vertex are said to belong to the same domestic in-group (say, the same neighborhood or the same set of religious beliefs). Agents at distance d_w can be thought of as domestic out-group. Likewise, we think of agents at distance d_l as global

[3]We confirm that very similar results on the difference between Western and non-Western countries hold in a much larger sample of countries in the Comparative Study of Electoral Systems (CSES) dataset, and in the World Values Survey (WVS) longitudinal dataset, see Appendix B.

in-group (say, people who live in a different country but adhere to the same values) and at distance $d_l + d_w = 1$ as global out-group.

Agents care about their own consumption and the consumption of others, though to potentially heterogeneous degrees. Our formalization of universalism is similar to Tabellini (2008); also see Enke (2019) for a recent cultural economics application. Define $J_i = I \setminus \{i\}$ to be the set of $N-1$ people in the population other than i and by D_i the set of $N/2-1$ domestic people other than i. Let x_i denote the consumption of agent i. The utility function of agent i is given by

$$u_i(x_i, x_{-i}) = x_i + \beta_i \sum_{j \in J_i} x_j a_{i,j}(d_{i,j}, \theta_i) \tag{1}$$

$$a_{i,j}(d_{i,j}, \theta_i) = \frac{1+\theta_i}{2} - \theta_i d_{i,j} \tag{2}$$

The parameter $\beta_i \in (0,1]$ scales agent i's *level* of altruism, while $\theta_i \in (0,1]$ governs the *slope* of altruism as a function of social distance. Figure 3 illustrates. We construct $a_{i,j}$ such that (i) altruism declines linearly as a function of distance; and (ii) the function integrates to a constant (1/2). This clarifies that the universalism parameter θ_i does not scale who is "more or less moral," but only how uniformly an agent distributes a given altruism budget.[4] Intuitively, a full universalist might argue that it is appealing to treat everyone equally, while others might point out that the universalist's moral compass is distorted in that she treats her friends not very well. Indeed, in Enke et al. (2020) we show that universalists have fewer friends and spend less time with them, compared to less universalist people.

3.2 Domestic Policy

3.2.1 Domestic Policy Options

Agents first vote for one of two policies in a simple majority system, where voting is assumed to be sincere. After each agent casts a vote $v_i \in \{A, B\}$, depending on which policy was selected, agents potentially take an action $q_i \in \{0,1\}$ that we will think of as free-riding or cheating. The "safe" policy option A enforces that nobody can cheat on society. However, the enforcement of this policy is costly, and that cost is shared equally among all domestic agents for a per capita cost c that is deducted from the baseline consumption level that is normalized to zero. The "risky" option B does not impose a per capita cost on each agent, yet each agent can free-ride or cheat on society. Cheating

[4] A potential micro-foundation for such type-dependent altruism is that agents exhibit greater altruism towards those agents that they believe to be "good" types, as in the model of Levine (1998). Then, our utility function corresponds to a reduced-form version of a model in which beliefs about the types of others vary as a function of social distance, as in Section 3.2 below.

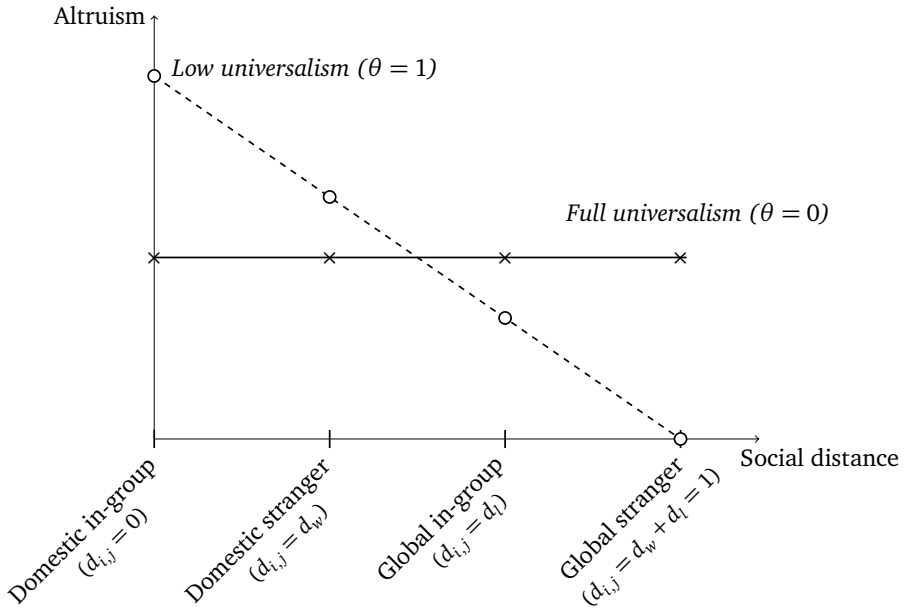

Figure 3: Illustration of heterogeneity in universalism. In the empirical analysis, we do not impose that the domestic stranger is socially closer than the global in-group.

by agent i delivers an extra rent $s > 0$ for agent i but imposes an overall externality of $e > 2s$, which is shared equally among all domestic agents for a per capita externality of $2e/N$. As will become clear, we only use the terminology "safe" and "risky" policy to point out the scope for cheating that is implied by the policies – it will sometimes be the case that what we call the "safe policy" is riskier in respects other than cheating, but this is immaterial for our purposes.

As explained in Table 1, we argue that these abstract features of the two policies map into some of the structural features of each of the eight policy domains discussed above. For example, in the case of welfare payments, the safe option A corresponds to a system with few welfare payments, so that agents cannot cheat on society by claiming benefits they are not entitled to. On the other hand, this causes a societal loss because random income shocks cannot be equalized. The risky option B, on the other hand, corresponds to a more expansive welfare state, which opens up the possibility of cheating.

In this model, option A and option B are not defined by the implied level of government spending. Indeed, as can be seen in Table 1, in the domain of welfare, the safe option A corresponds to lower spending, while in the domain of police and law enforcement, the safe policy A corresponds to higher spending.

3.2.2 Beliefs and Equilibrium

In order to calculate valuations of each policy, a decision maker must form beliefs about who would cheat under policy option B. We model decision makers that have rational ex-

Table 1: Mapping of policy domains to abstract model policies

Policy domain	**Risky Option B**	**Safe Option A**
Abstract framework	No fixed cost, but agents can cheat and hence earn rent s by imposing per capita externality of e	Cheating impossible, but agents pay fixed per capita cost of c
Domestic policies		
Welfare	Expansive welfare state: Agents can cheat on society by claiming benefits they are not entitled to (and hence reap rent s); this causes per capita externality e	No welfare state: Agents cannot cheat by claiming benefits they are not entitled to, yet this imposes a per capita cost c because in the absence of welfare payments, random income shocks cannot be equalized ex post
Universal health care	Same logic as for welfare	
Affirmative Action	Extensive AA: Agents who benefit from AA can cheat by reducing effort because they know that they will get promoted either way; this imposes a cost on other agents	No AA: Agents cannot reduce effort while still getting promoted; yet absence of AA also entails a social cost because disadvantaged groups in society cannot live up to their potential
Police and law enforcement	Weak police: Stealing and fraud possible	Strong police: Stealing is impossible; but entails a cost because police needs to be paid for
Foreign policies		
Effective border control	Weak border control: Increase in number of people who could come into country and free ride on others' efforts	Strong border control: less immigration, but this entails a per capita cost because border control is expensive, and because some immigrants are truly in need
Military	Weak military: Other countries can cheat or exploit	Strong military: Foreigners cannot exploit domestic people; entails per-capita cost because military needs to get paid for
Environmental protection	Strong regulation: Other countries can cheat by de-regulating and hence growing their economy at expense of domestic agents	Weak regulation: Foreign countries cannot exploit domestic regulation; yet this entails cost because environmental degradation might have economic or health impacts on domestic agents
Foreign aid	Extensive aid: Foreigners can cheat by claiming aid they are not entitled to or by misusing funds	No aid: Foreigners cannot cheat; yet this entails a cost because lack of aid could cause increased migration or wars

pectations about the *overall* fraction of agents who will cheat under option B. However, to allow for an analysis of the role of universalism in trust beliefs, we assume that decision makers may not form correct beliefs about *which* agents will cheat. For example, casual introspection suggests that people differ dramatically in whether they believe that immigrants are more likely to be criminals than natives. Formally, the subjective probability that agent i assigns to agent j not cheating under policy B is

$$b_{i,j}(d_{i,j}, \delta_i) = \gamma^* + \frac{d_w}{2}\delta_i - \delta_i d_{i,j} \tag{3}$$

where $\delta_i \in (0, 1]$ controls the rate at which the belief that an agent will not cheat falls as a function of social distance. We think of δ_i as the inverse of universalism in trust. Note that beliefs are defined analogously to altruism above, and can be graphically represented analogously to Figure 3. As in the case of altruism, universalists and non-universalists do not differ in their overall level of trust: the belief function in equation (3) integrates to the constant γ^*, which is endogenous and corresponds to the fraction of agents who do not cheat in equilibrium. Thus, in equilibrium, agents have rational expectations about the overall rate of cheating in society but not necessarily about how cheating is correlated with social distance from them. We assume that $(\beta_i, \theta_i, \delta_i)$ are positive independent joint uniform.

3.2.3 Domestic Policy Views

We solve the game by backward induction. Denote by $E_i[\cdot]$ the subjective "expectations operator" that applies the belief function in equation (3). Further denote by L_{-i} the hypothetical losses that agents incur due to the cheating of agents other than i. In the second stage of the game, if the risky policy is implemented, agent i cheats iff

$$E_i[u_i(q_i = 1)] = \left(s - \frac{2e}{N} - L_{-i}\right) + \beta_i \sum_{j \in D_i} \left\{[1 - b_{i,j}(\delta_i)] \cdot s - \frac{2e}{N} - L_{-i}\right\} \cdot a_{i,j}(\theta_i) \tag{4}$$

$$> E_i[u_i(q_i = 0)] = -L_{-i} + \beta_i \sum_{j \in D_i} \{[1 - b_{i,j}(\delta_i)] \cdot s - L_{-i}\} \cdot a_{i,j}(\theta_i) \tag{5}$$

which delivers the vector q^* of individual cheating decisions $q_i^*(\theta_i)$. The resulting losses (externalities) that each agent incurs are denoted by $L^*(q^*(\theta)) \equiv (1 - \gamma^*(\theta))e$.

In the first stage of the game, an agent votes for the safe policy A iff

$$u_i(A) = -c + \beta_i \sum_{j \in D_i} \{-c\} \cdot a_{i,j}(\theta_i) \tag{6}$$

$$> E_i[u_i(B)]$$

$$= [s \cdot q_i^*(\theta_i) - L^*(q^*(\theta))] + \beta_i \sum_{j \in D_i} \{[1 - b_{i,j}(\delta_i)] \cdot s - L^*(q^*(\theta))\} \cdot a_{i,j}(\theta_i) \tag{7}$$

which delivers the vote $v_i^*(\theta_i, \delta_i)$ as a function of universalism and other parameters. Under the parameter assumptions discussed in Appendix A, we get:

Prediction. *Individuals with higher universalism exhibit a stronger preference for the risky*

domestic policy B: welfare, universal health care, affirmative action, and weak police and law enforcement. These predictions hold for universalism in both altruism and trust.

See Appendix A.2 for a proof. The intuition behind this prediction is straightforward. All else equal, a decision maker who is less universalist will believe that out-group agents are more likely to cheat. This, in turn, implies a redistribution of resources away from the agent's in-group members to the out-group, which all agents who are not fully universalist dislike. As a consequence, agents who are less universalist in altruism or trust (or both) prefer option A more.

3.3 Foreign Policy

Decision makers are again presented with a choice between two policy options. Under the safe policy option A, domestic and foreign agents receive their baseline consumption x. Domestic agents additionally pay a per capita cost c.[5] Under the risky policy B, domestic agents do not have to pay c. However, in this regime foreign agents can cheat and get s by imposing an overall cost of e on all domestic people, which is again equally shared. Table 1 explains how this abstract structure maps into the domains of military, border control, foreign aid, and environmental protection. As with the domestic policies above, note that the risky policy B sometimes corresponds to big and sometimes to small government. Again, the key defining characteristic that matters for our analysis is whether a policy introduces or prevents cheating opportunities.

The mechanics of the foreign policy analysis are very similar to the domestic case. We again assume that agents' beliefs about the overall fraction of cheaters are correct, but that they have heterogeneous beliefs about how cheating is correlated with distance from them. We exposit the details in Appendix A, and state the main prediction here.

Prediction. *Individuals with higher universalism exhibit a stronger preference for the risky policy B: weak border control, weak military, stringent environmental protection, and expansive foreign aid. These predictions hold for universalism in both altruism and trust.*

3.4 Informal Discussion

Our formal model highlights the role of cheating opportunities because this allows us to tie together different types of policy preferences in a simple formal framework, in a way that accommodates both altruism and trust. At the same time, universalism in altruism as such can also plausibly explain heterogeneity in some policy preferences

[5] In some of the foreign policy domains we consider, c is likely to be paid by both domestic and foreign agents. Our main predictions remain unchanged if we assume that c is paid by both domestic and foreign agents.

absent cheating opportunities. For instance, non-universalists may intrinsically dislike the idea that their own tax money goes to people that they are socially distant from them. Similarly, non-universalists may support tight border control even in the absence of cheating opportunities for immigrants, if they view immigrants as crowding out domestic in-groups' consumption of the country's resources, like jobs. To take a final example, non-universalists might dislike the concept of foreign aid simply because they value the welfare of the domestic poor higher than that of the foreign poor. These examples illustrate that our general idea – that "left" policies are appealing to universalists – does not appear to hinge on our specific cheating framework but comfortably accommodates other intuitions as well.

4 Survey Design

4.1 Logistics

We implemented internet surveys in Australia, France, Germany, Sweden, the United States, Brazil, and South Korea through the infrastructure of the market research panel of *Dynata*. The surveys were implemented between June and August 2019. The original survey was developed in English, translated into other languages by *Dynata*, and then checked by us using native speakers. The median completion time was 20 minutes.

The survey consisted of four components: (i) an introductory screen that elicited demographics and routed respondents into or out of the survey; (ii) decision screens to measure universalism and other social preferences; (iii) screens to measure policy views; and (iv) a questionnaire to elicit additional information and covariates. The order of parts (ii) and (iii) was randomized across respondents. We also randomized the order in which universalism in altruism and universalism in trust were elicited.[6]

We took two measures to ensure quality control. First, every respondent who completed the survey in less than 400 seconds 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 $N = 14,731$), a subset of which

[6]A permanent link for the U.S. version of our survey is: `https://harvard.az1.qualtrics.com/jfe/form/SV_aftuqgHsyIAShkp`.

makes up the more representative samples that we pre-registered. The physical process was that *Dynata* kept sampling respondents until our pre-specified quotas were satisfied. "Surplus" respondents came free of charge for us. Since we view throwing away data as scientifically questionable, all analyses reported in the main text make use of the full sample. In the Appendix we replicate all analyses using the pre-registered (smaller) representative samples. The results are always extremely similar.

As a final remark on the sample, *Dynata* had considerably more difficulty in constructing representative samples in Brazil and South Korea than in the other countries, which we did not anticipate when we initially contracted with them. Thus, the final samples sent to us skew young, rich, and employed in Brazil and Korea. The sample characteristics are summarized in Appendix C.1.

4.2 Measurement of Universalism

Our objective is to measure the empirical analogue of the universalism parameters θ and δ in the theoretical framework in Section 3. This requires measuring how altruism and trust vary as social distance increases, holding fixed the overall level of altruism and trust. We rely on a new set of structured experimentally-validated survey games to measure an individual's universalism. Our main goals when designing these games were to use survey games that (i) are conceptually closely linked to the model; (ii) capture a broad set of in-groups; and (iii) can be deployed at scale in online surveys relatively easily. To conserve space and focus, we relegated the development, experimental validation, and testing of these survey measures to a separate paper (Enke et al., 2020). We summarize the key aspects below.

4.2.1 Survey Games

Universalism in altruism. Respondents completed a total of 16 hypothetical money allocation tasks that allow us to construct a summary statistic of universalism in altruism (θ in the model). The construction of the survey games is closely tied to the theoretical framework in Section 3 in that it makes use of four different types of groups: domestic in-groups, domestic strangers, global in-groups, and global strangers. From these four types of groups, we construct three universalism components: domestic universalism, foreign universalism, and global universalism.

First, to estimate *domestic* universalism, respondents made ten decisions. In each of them, they were asked to split hypothetical \$100 between (i) a randomly-selected person from their country of residence and (ii) a randomly-selected member of one of their social groups, who also resides in the respondent's country of residence. We based the selection of in-groups on an ex-ante crowd-sourcing exercise (see Enke et al., 2020,

for details). Across the ten questions, the social groups included extended family, friends of family, neighbors, 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 U.S. were asked to split $100 between a randomly-selected person who lives in the U.S. and a member of their extended family, such as a cousin. The average allocation to the randomly-selected person across the ten questions then makes up the domestic universalism measure.

Second, to estimate *foreign* universalism, respondents were asked to split $100 between (i) a randomly-selected person from their country of residence and (ii) 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 (i) a randomly-selected person who lives anywhere in the world and (ii) 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 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 (i) that both individuals are equally rich (addressing income effects) and (ii) 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. Figure 13 in Appendix C.2 shows an example decision screen.

As discussed in detail in Enke et al. (2020), the separate money allocation decisions, and in particular the domestic, foreign, and global universalism summary components are all highly positively correlated with each other in a representative sample of the U.S. population. This is also true in our multinational dataset. To reduce the dimensionality of the data and minimize measurement error, we hence average the three components into a summary statistic of universalism in altruism. The construction of this summary statistic was pre-registered, see below. To document the validity of this procedure, some of the analyses below will also work with the separate universalism components.

Universalism in trust. Respondents completed a total of 16 tasks from which we estimate an individual's universalism in trust as empirical analogue of δ in the model. The procedure was identical to the one described for altruism above, except that in a given game respondents were asked to allocate 100 "trust points" (rather than $100) between two individuals, to express whom they trust more. This again yields domestic, foreign, and global universalism components, which we average into a summary statistic of uni-

versalism in trust. Again, the construction of this summary statistic was pre-registered.

Composite measure of universalism. Universalism in altruism and trust exhibit a correlation of $\rho = 0.62$ after accounting for measurement error using the obviously-related instrumental variables technique of Gillen et al. (2015). To reduce the dimensionality of the analysis, in most analyses below we work with a composite measure of universalism, which consists of the unweighted average of universalism in trust and universalism in altruism. At the same time, we reference robustness checks that use the altruism and trust measures separately, see Section 5.6.

4.2.2 Construct Validity

We validate the universalism measures along three dimensions. See Enke et al. (2020) for details. (i) *Experimental validation.* We implemented an ex-ante experimental validation procedure. Specifically, we show that, over a one-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 behavior in our trust point allocation game is highly correlated with trust beliefs in a structured cheating task that is standard in the experimental economics literature. (ii) *Correlation with real donation decisions.* We also show that our survey measure of moral universalism predicts real donation decisions: while universalists donate less than non-universalists to local community organizations, they donate more to nationwide and international charities. (iii) *Choice of social groups.* We document that an individual's degree of universalism with respect to the set of fifteen domestic and foreign groups that we implement is highly correlated with their universalism with respect to a more comprehensive set of forty social groups.

4.2.3 Descriptives

Figure 4 shows a histogram of the composite universalism measure, pooled across all Western 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. Appendix C.3 shows the corresponding histograms in each country separately. Table 2 reports correlations with demographics. 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. These results are consistent with those documented in Enke et al. (2020)

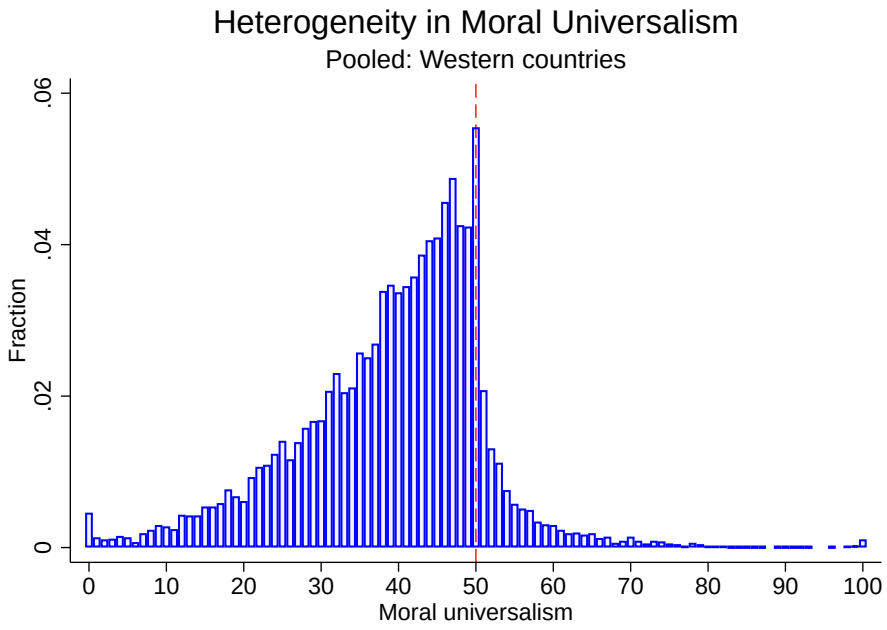

Figure 4: Distribution of the composite measure of moral universalism, pooled across all Western countries. The amounts reflect allocations to random strangers, so that the measure is decreasing in in-group favoritism. 50 corresponds to an equal split of money or trust points.

for a U.S. sample. Importantly, these correlations highlight that heterogeneity in universalism does not simply pick up variation in income or education (as in economic stories of the "left behind") – if anything, individuals with higher income and wealth are *less* universalist.

4.3 Measurement of Political Attitudes

Measures of Support for Expenditure Categories. 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 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, military, police and law enforcement, and border control. The order of these categories on the computer screen was randomized.

Figure 16 in Appendix C.2 provides a screenshot. Naturally, because of the free-entry format, responses to these questions are subject to large outliers. To account for these

Table 2: Individual-level correlates of universalism: Western countries

	Correlation between composite measure of universalism and:						
	Age	Female (0-1)	Income Index (z-score)	Wealth Index (z-score)	College (0-1)	Religiosity (z-score)	Urbanicity (z-score)
Raw correlation	-0.16^{***}	0.07^{***}	-0.08^{***}	-0.12^{***}	0.01	-0.10^{***}	0.03^{***}
OLS coeff. (w/ Country FEs)	-0.12^{***}	1.78^{***}	-0.93^{***}	-1.48^{***}	0.21	-1.16^{***}	0.38^{***}
OLS coeff. (multivariate) (w/ Country FEs)	-0.09^{***}	1.13^{***}	-0.54^{***}	-0.79^{***}	1.14^{***}	-1.04^{***}	0.19

Notes. The first row reports the Pearson 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 Appendix E for details on the construction of the demographic variables. All z-scores are computed separately within each country. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

outliers, we winsorize the desired spending levels at $+/-$ 3 standard deviations of the within-country mean, as specified in our pre-registration (discussed below). That is, we replace each dollar amount above (below) the amount that corresponds to 3 SD above (below) the mean with this value. This affects 1.6% of all responses.

Summary statistic of policy views. As specified in our pre-registration, we compute a simple summary statistic of policy views across all policy domains, which is computed from the desired expenditure shares:

$$\text{Summary statistic of policy views} = \frac{\text{Foreign aid} + \text{Environment} + \text{Aff. action} + \text{Welfare} + \text{Health care}}{5} - \frac{\text{Military} + \text{Police} + \text{Border control}}{3} \tag{8}$$

where each policy denotes share of desired expenditure that goes to a domain. Pooling data across all countries, this summary statistic exhibits a correlation of $\rho = 0.40$ with respondents' self-positioning on a left-right scale (0–10). We pre-specified the summary statistic in this particular way because it corresponds very closely to the structure of policy views in the Western countries discussed in Section 2. We standardize all political attitudes variables into z-scores, separately within each country.[7]

[7]As a second, and complementary measure of policy views, we elicit respondents' level of support for the eight policy domains above using Likert scale questions. These directly ask participants to indicate whether they strongly support or strongly oppose a given policy, on a scale from zero to ten. As specified in a pre-registration (see below), we use these measures as instruments to be able to account for measurement error using "Obviously-Related Instrumental Variables" analyses (Gillen et al., 2019).

Measures of Support for Specific Policy Proposals. The measures reported in the previous section aim at capturing a respondent's support for broad policy domains. In addition, we measured respondents' preferences over more specific policy proposals. After respondents had indicated their desired spending levels for the eight broad policy domains, we asked them how much money they would like to see collected and spent on two specific projects or policy proposals within each broader policy domain. We constructed these proposals such that one was more universalist than the other, yet both focused on the same policy domain. We present the policy proposals in Table 3. To illustrate, take the example of welfare payments. We elicited desired spending levels for (i) "Redistributing local tax revenues as welfare payments across all communities nationwide" and (ii) "Redistributing local tax revenues as welfare payments only within the local communities they were raised." Figure 17 in Appendix C.2 provides a screenshot.

As pre-registered, we again winsorize the data at +/− 3 sd. of the within-country mean, which affects 0.1% of all responses. We also standardize these variables into z-scores, separately within each country.

4.4 Covariates

Even though this paper is descriptive in nature, we seek to assess to which extent 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. All of these covariates and their construction are described in detail in Appendix E.

To highlight just a few, we compute income, wealth and religiosity indices using principal component analyses. An income index is computed as first principal component of two questions that ask respondents (i) for a continuous estimate of their household income and (ii) to place themselves into income buckets. The wealth index is the first principal component of the z-scores of (i) respondents' estimates of net worth, (ii) whether they owned a home and (iii) whether they own stocks. The religiosity index is constructed as first principal component of the z-scores of (i) a self-assessment of religiosity (scale 0–10), (ii) frequency of church attendance, and (iii) a binary indicator for whether the respondent considers themselves to be an Atheist.

Table 3: Specific policy proposals

Policy domain	**More universalist**	**Less universalist**
Military and counterintelligence	Peacekeeping and humanitarian missions by the military abroad	Ensuring [American, French, etc.] defense and security
Welfare payments	Redistributing local tax revenues as welfare payments across all communities nationwide	Redistributing local tax revenues as welfare payments only within the local communities they were raised
Effective border control	Identifying and admitting into the country only those immigrants with the highest need for help	Identifying and admitting into the country only those immigrants who would be good citizens (e.g., be likely to pay taxes and refrain from engaging in criminal activities)
Environmental protection	Preventing global climate change	Cleaning and conserving forests and rivers in local communities in [the U.S., France, etc.]
Universal healthcare	Using local tax revenues to fund health insurance across all communities nationwide	Using local tax revenues to fund health insurance only within the local communities they were raised
Police and law enforcement	Sensitivity training for the police to ensure justice and equal treatment of all	Increasing the capabilities of the police to prevent and prosecute criminal or suspicious behavior
Foreign aid	Sending foreign aid to countries that are in most need of help	Sending foreign aid to countries that are our international allies
Measures to ensure no individual is disadvantaged in access to education, the labor force, and marriage	Measures to ensure no individual is disadvantaged in access to education, the labor force, and marriage	Measures to ensure no one of your same background (e.g., gender, ethnic background or ancestry) is disadvantaged in access to education, the labor force, and marriage

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 universalism to be correlated with support for each of the eight policy domains, based on the model in Section 3; (iv) the construction of the summary statistic of policy views discussed above; (v) the prediction that universalism would be more positively correlated with the more universalist implementations of policy domains than their less universalist counterparts; and (vi) an analysis of whether the patterns in Brazil and South Korea are different from those in the Western countries.

Two remarks regarding the relationship between the pre-registration and the analyses in this paper are in order. First, as discussed above, our sample turned out to be larger than anticipated, for reasons beyond our control. We report robustness checks using the

smaller more representative samples in the Appendix.

Second, we pre-specified that we expect all of our hypotheses to be true for both universalism in altruism and universalism in trust. To conserve space and reduce the dimensionality of the analysis, we mostly work with a composite measure of universalism that averages universalism in altruism and trust. We replicate these analyses with the separate universalism measures in the Appendix. The results are always very similar.

5 Survey Results

5.1 Summary Statistic of Policy Views

We begin by considering the summary statistic of policy views, where higher values indicate higher desired expenditure shares for the canonical left-wing policies. We first pool the data across Western countries and then disaggregate the results in a second step.

Table 4 presents the results of a set of OLS regressions of the summary statistic of policy views on each of the separate universalism measures detailed in Section 4.2. The composite universalism measure is constructed as average of universalism in altruism and trust. The universalism measures are all in [0,1], where zero means that all money and trust points are allocated to the in-group member in a given game, 0.5 means that the money and the trust points are split equally, on average, and one corresponds to the (counterfactual) case that someone always allocates all money and trust points to the socially more distant individual.

We find a strong positive relationship between universalism and the summary statistic of policy views. This is true for each individual component of universalism, regardless of whether it is measured in the altruism or trust space. In fact, as we document in Figure 21 in Appendix C.4.1, this pattern is even more general than what is suggested by the results in Table 4: out of the 32 different allocation decisions in our survey from which we estimate universalism in altruism and trust, *all* are significantly correlated with the summary statistic of policy views, such that a higher allocation towards the socially more distant individual is correlated with a "higher" score on the summary statistic of policy views. This provides evidence that our results are not driven by a just a few in-groups but reflect a general psychological tendency. Given the similarity of results across different universalism components, to average out measurement error, and to reduce the dimensionality of the analysis, we focus on the composite measure of universalism in what follows.

As we document in column (10), the relationship between universalism and our summary statistic of policy views is robust against controlling for age, gender, income,

Table 4: Summary statistic of policy views and different universalism measures, pooled across countries

	Dependent variable: Summary statistic of policy views									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Domestic universalism in altruism	0.90*** (0.06)									
Foreign universalism in altruism		0.94*** (0.04)								
Global universalism in altruism			1.19*** (0.06)							
Composite universalism in altruism				1.57*** (0.07)						
Domestic universalism in trust					1.01*** (0.08)					
Foreign universalism in trust						0.94*** (0.06)				
Global universalism in trust							1.19*** (0.08)			
Composite universalism in trust								1.50*** (0.09)		
Composite universalism									2.09*** (0.09)	1.64*** (0.09)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographic Controls	No	No	No	No	No	No	No	No	No	Yes
Observations	10881	10881	10881	10881	10881	10881	10881	10881	10881	10881
R^2	0.02	0.05	0.04	0.06	0.02	0.03	0.03	0.04	0.06	0.11

Notes. OLS estimates, robust standard errors in parentheses. Data are pooled across all five Western countries. The dependent variable is the summary statistic of policy views, constructed as described in Section 4.3 and standardized into a z-score within each country. The construction of each universalism measure is outlined in Section 4.2. Demographic controls include age, gender, income, wealth, college, urbanicity, religiosity, equity-efficiency preferences, altruism, trust, and beliefs about the efficiency of government. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

wealth, college education, urbanicity, religiosity, equity-efficiency preferences, altruism, trust, and beliefs about the efficiency of government. Conditional on country fixed effects, the composite universalism measure exhibits a partial correlation with the summary statistic of policy views of $\rho = 0.25$. While we provide more sophisticated benchmarking analyses later, it is perhaps informative that the corresponding correlation for the belief that government is efficient vs. wasteful is $\rho = 0.15$, the one for college degree $\rho = 0.05$, for age $\rho = -0.12$, for the religiosity index $\rho = -0.10$, for the income index $\rho = -0.07$, and for the wealth index $\rho = -0.12$.

The analysis reported here correlates universalism with the pre-registered summary statistic of policy views. An alternative approach is to link universalism to the first principal component of desired expenditure shares derived in Section 2. These correlations are also always positive and statistically highly significant. The correlations range between $\rho = 0.19$ in France and $\rho = 0.33$ in Sweden. In other words, universalism is significantly correlated with an unsupervised summary statistic of the data. In a next step, we study the different policy issues one-by-one.

5.2 Separate Policy Views

Figure 5 summarizes the results for the separate policy categories. The underlying OLS regressions relate the desired *share* of overall desired expenditure for each policy (standardized into z-scores) to universalism, separately for each country and all Western countries combined, for a total of 48 regressions. The left panel shows the results of univariate regressions, while the point estimates in the right panel stem from multivariate regressions that 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.

As hypothesized, in all Western countries, we observe a strong negative relationship between universalism and desired expenditure shares for the three "right-wing" policy domains, while the relationship is generally positive and statistically significant for the five "left-wing" domains. That is, viewed through the lens of the theoretical framework in Section 3, lower universalism is associated with decreased support for "risky" policies that introduce cheating opportunities. In terms of quantitative magnitude, the estimated regression coefficients suggest that increasing universalism from zero to 1/2 (and hence moving from 100:0 to 50:50 allocation decisions) is associated with a 0.25–1.0 standard deviation change in each of the policy views.[8]

Overall, universalism is consistently correlated with policy views in the ways we hypothesized. Out of the 40 regression coefficients for the individual countries reported in the left panel of Figure 5, 37 have the expected sign. Of these, 33 are statistically significant at least at the 10% level. Once our battery of controls is added in the right panel, 38 of these coefficients have the expected sign, out of which 31 are statistically significant at least at the 10% level.

We proceed by looking at desired expenditure *levels*. To this effect, Figure 6 reproduces the left panel of Figure 5, except that now the dependent variables are desired (log) expenditure levels rather than implied shares. 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 than non-universalists – just in policy domains that have a universalist orientation. In this sense, universalism directly reproduces the pattern reported in Figure 2 in Section 2 that motivates our paper.

[8]A notable exception occurs in the domain of universal health care, where the relationship is strongly positive in the U.S. but either not statistically significant or even negative in the other countries. This pattern might arise because, in contrast to the United States, 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. It probably also implies that respondents outside the U.S. interpret survey questions about "universal health care" in a different fashion than Americans.

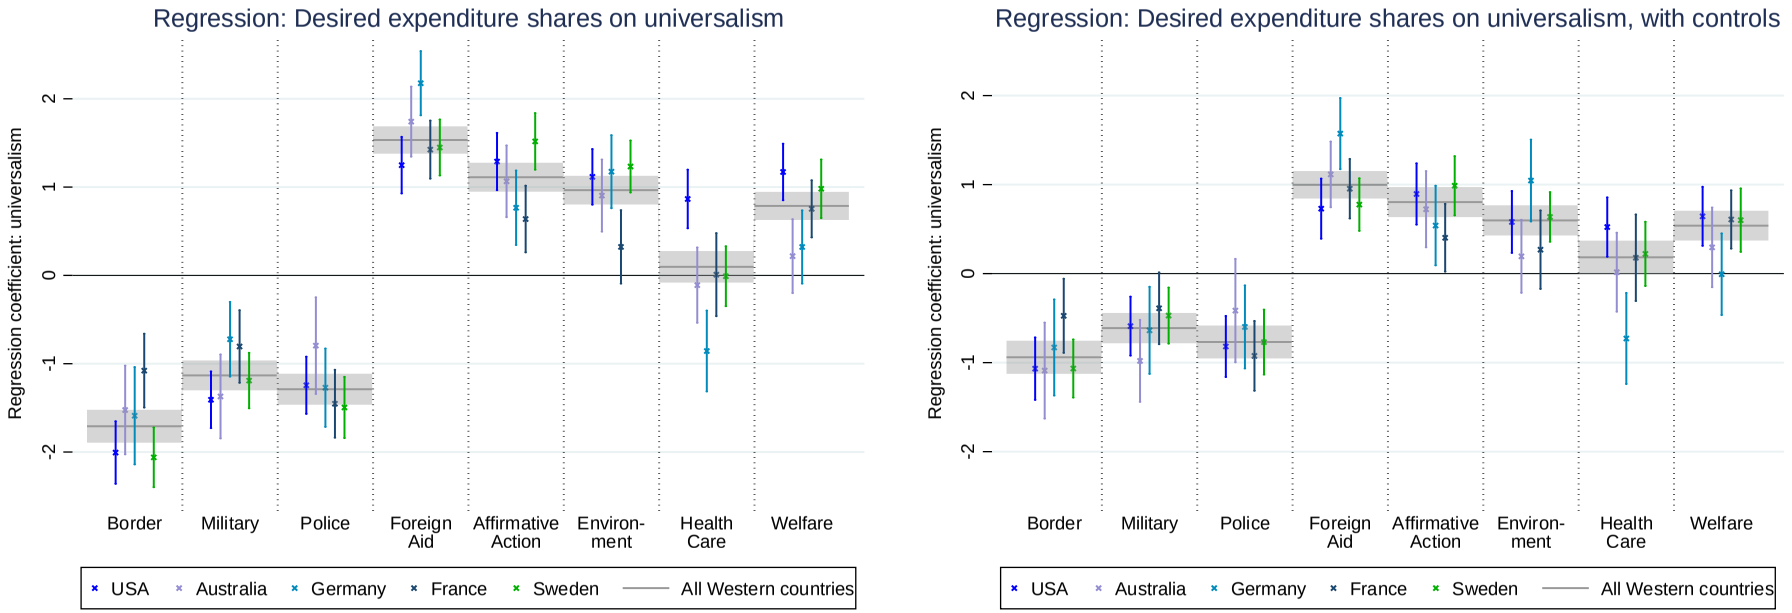

Figure 5: The left panel plots the OLS regression coefficients of univariate regressions of desired expenditure shares for each policy domain (as a fraction of overall desired government spending for the eight policy domains) on composite universalism. The right panel plots the analogous coefficients of multivariate regressions, in which we 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. See Appendix E for details on the construction of these variables. 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 western countries" specifications include country fixed effects.

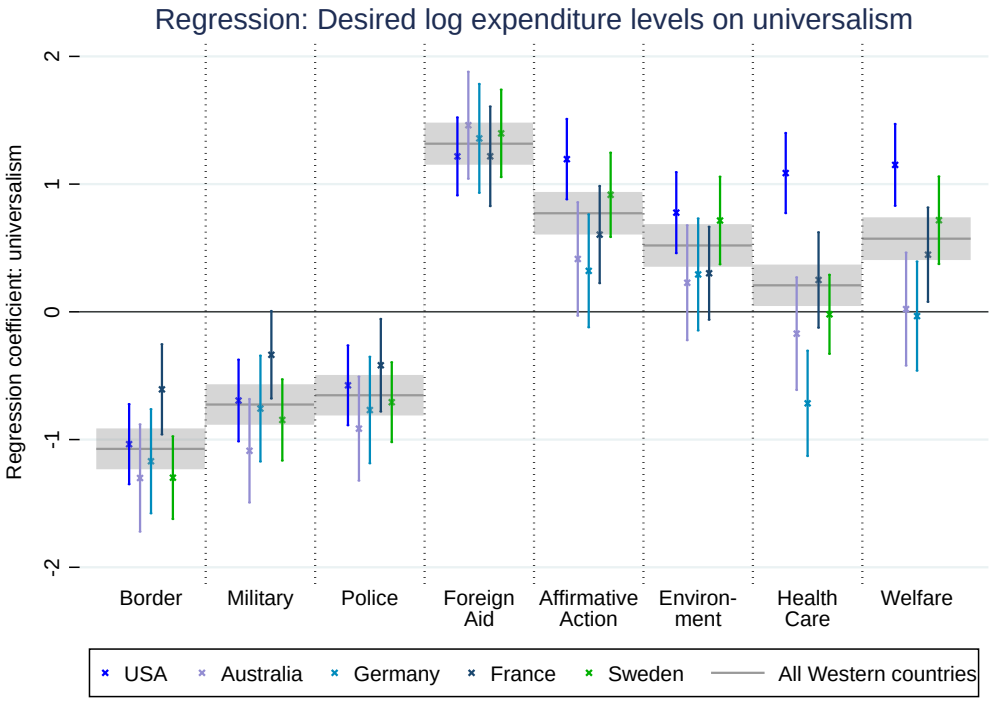

Figure 6: The figure plots the OLS regression coefficients of univariate regressions of desired log expenditure levels for each policy domain on composite universalism. Universalism is in [0,1] and the dependent variables are standardized into z-scores within each country. Error bars indicate 95% confidence intervals using robust standard errors. The "All western countries" specification includes country fixed effects.

5.3 Benchmarking Exercises

An immediate question is whether other individual characteristics could also produce the patterns we are trying to explain. To address this question, Figure 7 summarizes the relationship between desired (log) expenditure levels and eleven individual characteristics. For simplicity, we pool the data across Western countries for this analysis. In terms of demographics, we focus on age, religiosity, income, wealth, completion of a college degree, and urbanicity. In terms of beliefs and preferences, we consider residual measures of altruism and of generalized trust, the respondent's preferences over equity vs. efficiency, strength of belief that the government works efficiently, and strength of the belief that one might personally benefit from government spending on each policy domain.[9] We selected this set of variables for the benchmarking exercise because they are commonly associated with an individual's position on the political spectrum. Indeed, in our data, conditional on country fixed effects, a respondent's self-assessment on an 11-point left-vs.-right scale exhibits correlations of: $\rho = 0.07$ with income, $\rho = 0.13$ with wealth and $\rho = 0.22$ with religiosity. This suggests that we measure these variables

[9]We employ *residual* measures of altruism and trust because both our dictator game and our elicitation of generalized trust are framed vis-à-vis a randomly-selected stranger. Thus, by construction, these raw measures partly include universalism.

in meaningful ways and that they are generally predictive of broad ideological views.

In Figure 7, the leftmost panels serve as reminder and show the pattern we are trying to explain: we are looking for a variable that is negatively correlated with desired spending levels for military, police and law enforcement and border control, but positively correlated with desired spending on welfare, health care, environmental protection, affirmative action and foreign aid. We find that none of the other eleven variables produces the characteristic pattern that universalism successfully reproduces. 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.[10]

5.4 Specific Policy Proposals

The claim of our paper, in particular viewed through the lens of the formal framework in Section 3, is not that universalists approve or disapprove of certain policy domains per sé, but rather that this is the case because each domain is implemented in a predominantly (non-) universalist way. If this was true, then it should be possible to manipulate people's support for broad policy domains such as the military or welfare by having them consider particularly universalist or non-universalist counterfactual implementations of these policies. For this purpose, as described in Section 4.3, we asked respondents to indicate their desired government spending level for 16 specific policy proposals (two for each of the eight broad policy domains), where one proposal was more universalist than the other.

To analyze whether this affects people's stated policy preferences, Figure 8 plots the OLS regression coefficients of universalism for each of the specific policy proposals. Here, the left panel reports the results for the more universalist policies and the right panel those for the less universalist policies.

Focusing first on the left panel, we find strong and positive relationships between desired expenditure levels for each of the eight policy domains and universalism. For example, in contrast to the baseline analysis above, universalists are now *more* likely to endorse a strong military than non-universalists once the military is said to focus on humanitarian missions and peacekeeping abroad. Looking at the right panel, we find that the relationship between universalism and policy views is substantially shifted downwards, relative to the more universalist proposals. That is, the correlations are substantially attenuated and in many cases even reverse. For example, while non-universalists

[10]While it may appear puzzling that income and wealth are not correlated with support for welfare payments, this is merely a result of looking at desired expenditure *levels* rather than shares; once we look at shares, support for welfare payments decreases significantly with wealth and income.

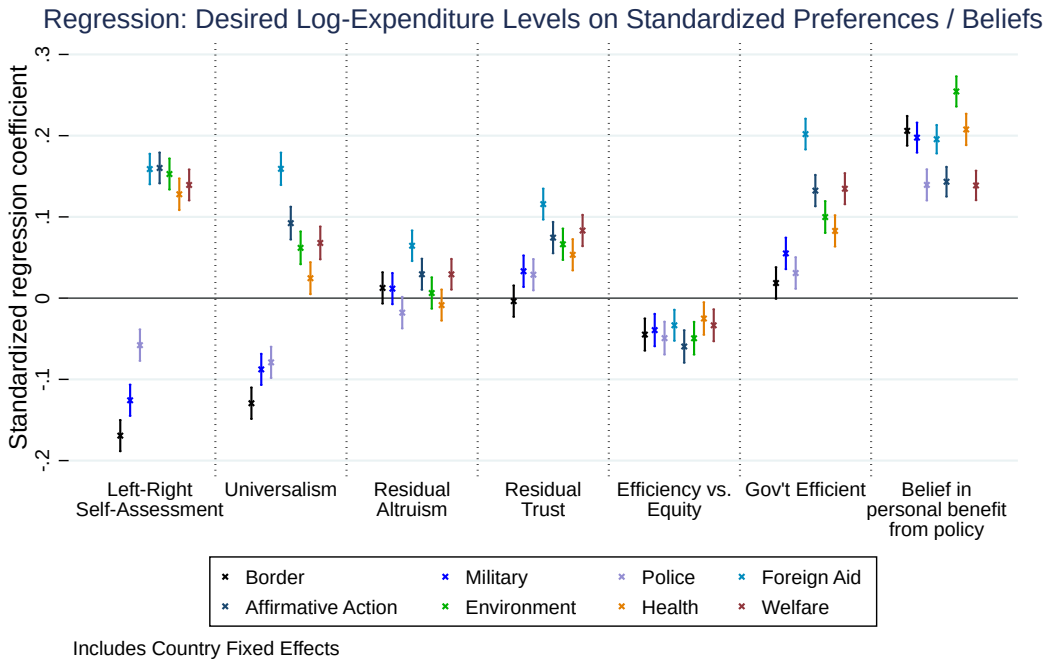

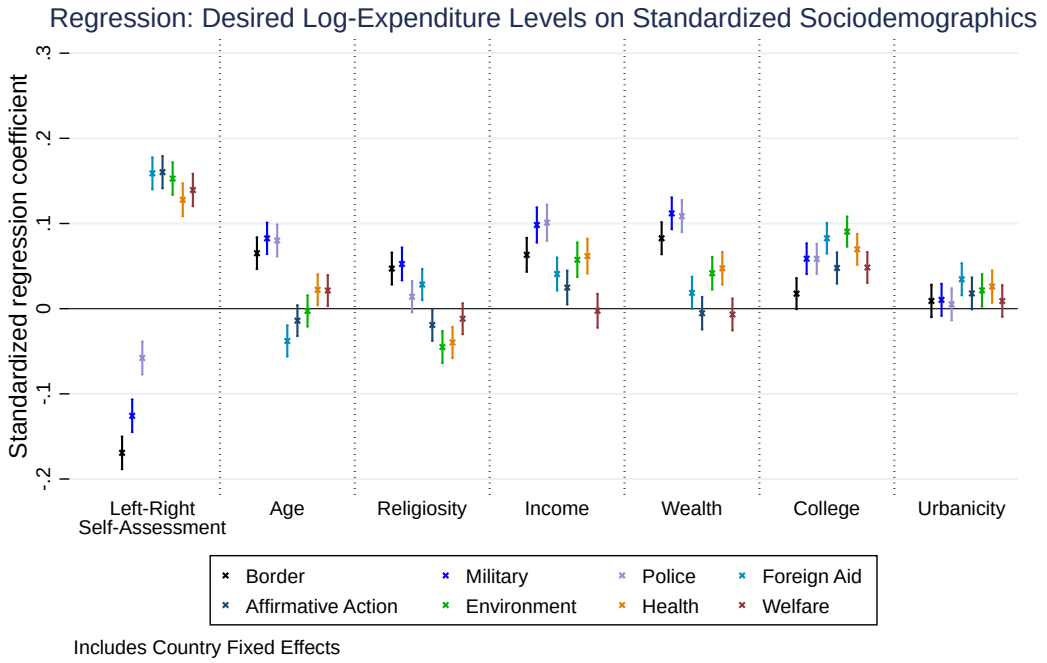

Figure 7: Benchmarking analyses. We report the standardized beta coefficients and confidence intervals for regressions of log desired expenditure level in a category on different individual-level chatracteristics, conditional on country fixed effects. All variables are standardized into z-scores within countries. The top panel considers the preferences and beliefs of respondents; the bottom panel considers demographics. 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.

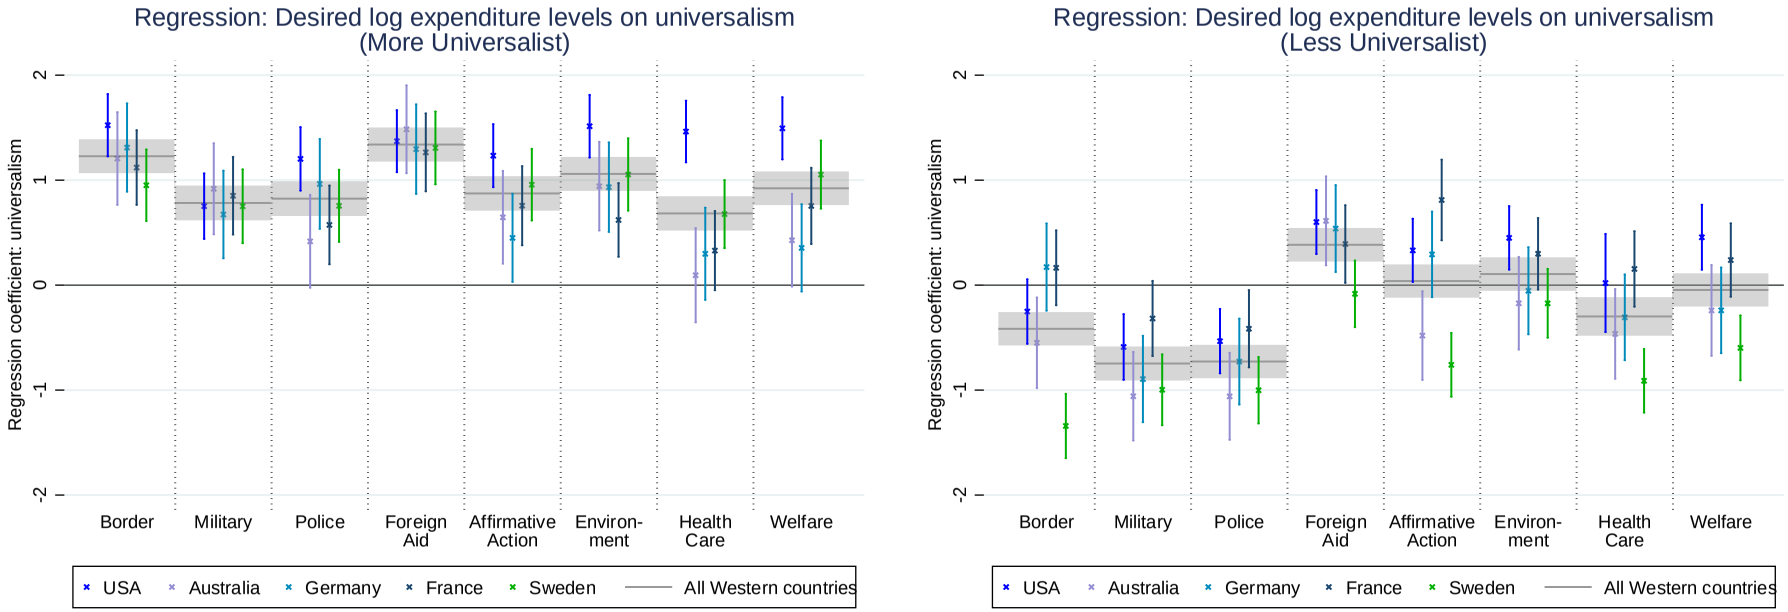

Figure 8: This figure plots the OLS coefficients of regressions of log desired expenditure levels for specific policy proposals on universalism, separately for each country. The left panel shows the results for the more universalist policies and the right panel those for the less universalist ones. See Table 3 for the wording of each of the policy proposals. 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 western countries" specification includes country fixed effects.

are generally opposed to welfare and environmental protection (compare Figure 5), they are as supportive of local redistribution and protecting the local environment as universalists.[11]

There is only one instance in which the coefficient on universalism is lower in the left panel than in the right panel (affirmative action in France). Otherwise, the OLS coefficient of universalism is between 0.16 and 2.29 units of a standard deviation larger in the left panel than in the right panel. Table 14 and Figure 26 in Appendix C.5 show that this difference in coefficient magnitudes is statistically significant in almost all cases.

In summary, this analysis documents that we can manipulate the relationship between universalism and policy views in predictable ways by asking people to consider more or less universalist implementations of each expenditure category. These results suggest that the (moral) conflict between the left and the right is not over abstract notions of the military or redistribution as such, but at least partly about which specific (non-) universalist form they take.

5.5 Non-Western Countries

Up to this point, our analyses have focused on the five Western countries in our sample. In this section, we comment briefly on the relationship between policy preferences and universalism in Brazil and South Korea. Figure 24 in Appendix C.4.2 plots the coefficients of regressions of desired expenditure shares on universalism in all countries, including Brazil and Korea. Here, we observe that the relationships between universalism and policy preferences are all weaker in magnitude and sometimes opposite in sign relative to those observed in Western countries. Furthermore, Figures 26 and 28 in Appendix C.5 show that policy preferences in Brazil and Korea are not observably affected by whether these policy domains are implemented in a more or less universalist way through specific policy proposals, as we did in Section 5.4.

These patterns might be unsurprising because (as discussed in Section 2 and Appendix C.4) the very clusters of policy views that we attempt to rationalize in this paper are absent in these countries. Put simply, if a baseline pattern is not observed, then it cannot be explained by universalism.

While we pre-registered that the relationship between universalism and policy views might look different in Korea and Brazil than in the Western countries, it seems worth discussing why these differences exist. A possible conjecture is related to a large and

[11] A related question is whether we can predictably manipulate the link between policy views and *self-reported political orientation*. Figure 29 in Appendix C.5 reproduces Figure 8, except that it plots the relationship between policy preferences and people's self-positioning on an 11-point left-right scale. Here, very similar patterns hold: left-wingers become much more likely to endorse a policy domain once the specific policy is universalist, and self-reported right-wingers become much more supportive of redistribution, health care, foreign aid, and environmental protection once it is implemented in less universalist ways.

influential body of cross-cultural work. Using datasets such as the World Values Survey, researchers have documented that over the past 50 years Western societies have increasingly moved towards emphasizing "post-material" values rather than purely material considerations, but that such a transition has not taken place outside the West in a comparable fashion (Inglehart, 1997). It is therefore conceivable that the relevance of moral considerations for political decision making only emerges when a certain level of income or other broader cultural factors have materialized.

To investigate this hypothesis, we analyze whether rich and well-educated elites in Brazil and Korea exhibit patterns that look similar to those established for the Western countries. As we discuss in detail in Appendix C.4.2, this is indeed the case: the correlations between universalism and policy views are significantly larger for richer and more well-educated individuals. Moreover, in both Brazil and Korea universalism directionally correlates with exactly the same cluster of policy views as in the West. In other words, non-Western elites appear similar to Western populations in terms of how their moral views relate to policy views. A potential tentative interpretation of these patterns is that morality may be a "luxury good" (Friedman, 2006) and only matters for voting once a certain level of economic security is attained.

5.6 Robustness Checks

The Appendix contains three further sets of robustness checks. First, our main analysis employed the composite measure of universalism. As specified in our pre-registration, Appendices C.4.1 and Appendix C.5.3 show that very similar results hold if we work with universalism in altruism or universalism in trust separately.

Second, 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 Appendix D.

Third, we contracted with *Dynata* for $N = 1,700$ respondents in each country, stratified to match the population on a number of dimensions. In Appendix C.6, we replicate the analysis using these more representative samples, with very similar results.

6 Field Evidence

We complement the survey analysis with field evidence. Here, we estimate the aggregate universalism of entire Congressional Districts (CDs) using large-scale donation data and link these to administrative data on local vote shares. The objective is to use financially

incentivized field choices to study whether – in line with the analysis above – more universalist regions vote left in higher proportions.

6.1 Data

To estimate a CD's universalism in altruism, we leverage data from DonorsChoose, an American non-profit organization providing an online "crowdfunding" platform for public school teachers.[12] On this platform, teachers can post funding requests for a wide variety of classroom "projects," such as field trips, classroom furniture, and purchases of basic school supplies or technology. Potential donors visit the website and donate to individual projects. Appendix F.4 provides screenshots of the layout and functionality of the platform. Notably, potential donors' ability to search through and filter projects based on location is a salient (usually, the highest) option available on the website.

The geographic scope of the data is broad and comprehensive: DonorsChoose reported in June 2019 that since the platform's inception in 2000, teachers in 82% of public schools in the United States had posted 1.4 million projects, reaching 34 million students and involving nearly 3.8 million donors, who had contributed $838 million. We use publicly available data to match all individual donations made on DonorsChoose between March 2000 and October of 2016 to their recipient projects. These data report the school's location (latitude and longitude) and the first three digits of each donor's ZIP code. We drop all observations for which the donor ZIP code is missing. Appendix F.1 reports summary statistics.

The geographic measures enable us to investigate how a CD's altruism towards another CD changes as a function of distance to the recipient. To perform this analysis, we aggregate individual donation data at the CD level to construct a dyadic dataset, where each observation represents every possible unique donor-recipient CD pair.

We work with two different measures of distance. First, the simple geographic distance between the CD's centroids. Second, a measure of friendship distance that was recently constructed from Facebook data by Bailey et al. (2018). This measure gives the probability that two randomly drawn individuals from two CDs are friends on Facebook. We view this measure of friendship distance as a summary statistic of social distance that aggregates a wide variety of demographic and social dimensions, such as ethnic distance, age distance, ideological distance, income distance, educational distance, etc.

[12]We are indebted to Ray Fisman for suggesting this analysis to us.

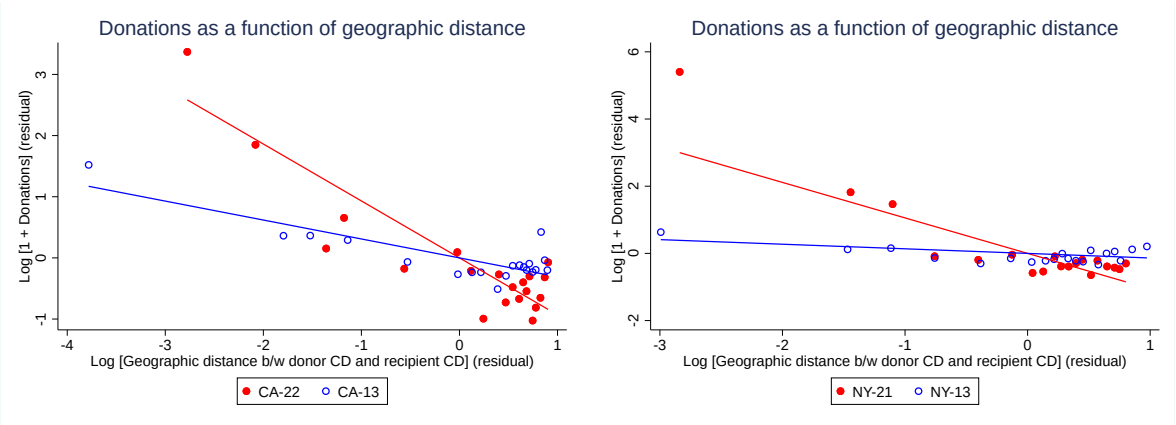

Figure 9: This figure illustrates regression equation (9) for four CDs. The left panel presents a binned scatter plot of all donations from both a Democratic and a Republican CD (based on 2016 presidential vote shares) in California against geographic distance to the respective recipient CDs. The right panel presents an analogue for New York state. All data are residualized of donor and recipient CD fixed effects.

6.2 Empirical Approach: Identifying Universalism in Altruism

To begin, we estimate a CD's universalism in altruism as (the negative of) the extent to which donations from a given donor CD decline as a function of geographic distance. Figure 9 illustrates this approach for four donor CDs from California and New York. For each donor CD, we provide a binned scatter plot of the log donation amount as a function of geographic distance to the recipient. Our interest is then in the *slope* of this function, where – as in the model in Section 3, we define a CD as being less universalist if it exhibits a steeper slope. In these scatter plots, the donation and distance data are residualized from donor and recipient fixed effects. That is, as explained below, we hold fixed the level of donations from and to a given CD, and only exploit variation in the slope.

Formally, for each donor CD i and recipient CD j, denote the log distance measure by $d_{i,j}$ and the log total dollar amount of donations by $p_{i,j}$. Further denote by $S_i \in \{0,1\}$ an indicator variable for each donor CD i and by $R_j \in \{0,1\}$ an indicator variable for each recipient CD j. Our estimating equation is then given by:

$$p_{i,j} = \sum_i \boldsymbol{\theta_i} \left[d_{i,j} \times S_i \right] + \sum_i \alpha_i S_i + \sum_j \varphi_j R_j + \varepsilon_{i,j} \tag{9}$$

The primary measure of interest is the vector of θ_i, which captures the extent to which donations from i to j decline as distance increases.

The estimating equation includes donor and recipient fixed effects to control for spatial variation in donation rates due to causes unrelated to universalism. For instance, a given donor CD may have disproportionately many users of DonorsChoose or be rich on average, hence leading to higher overall donation amounts. Similarly, a given recipient CD may post many projects on the DonorsChoose website or be very poor and hence

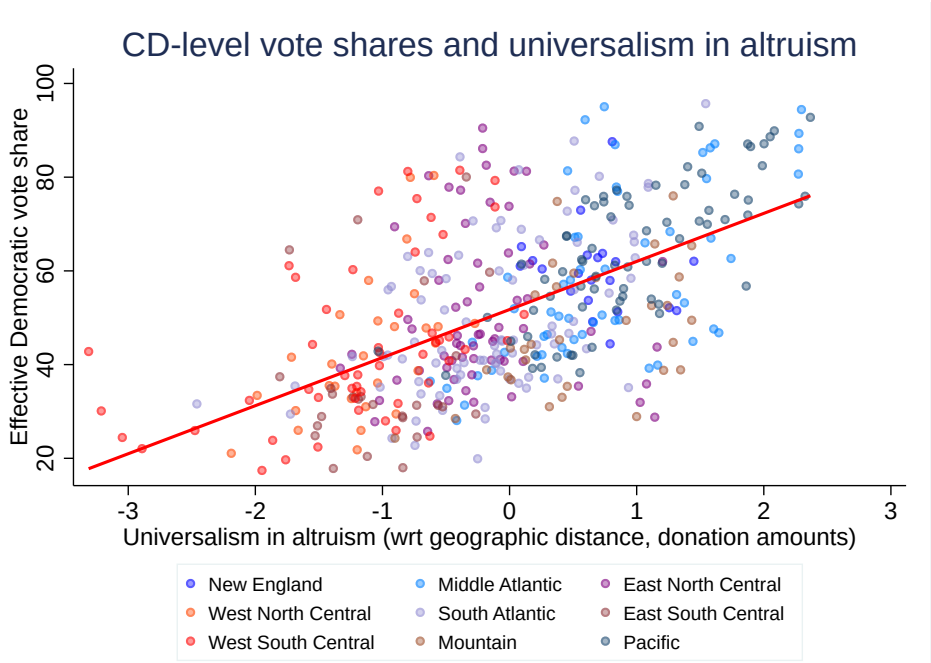

Figure 10: Relationship between universalism in altruism and CD-level vote shares. Universalism is the negative of $\hat{\theta}_i$ in equation (9).

receive many donations. Our specification nets out these level effects and only identifies the responsiveness of donations to distance, holding fixed both the level of donations from the donor and the amount of money a given recipient receives.[13]

6.3 Results: Universalism and Vote Shares

Figure 10 shows the raw correlation between universalism (standardized into a z-score) and 2016 Democratic vote shares ($\rho = 0.57$). Table 5 provides a regression analysis. Using the baseline measure of universalism developed above, columns (1)–(4) document that a one-standard-deviation increase in a CD's universalism is associated with a 10 to 13 percent higher Democratic vote share in that CD. Columns (2)–(4) show that the result is robust to including state fixed effects. The regressions also control for the CD's level of donations on DonorsChoose, median household income, the fraction of the population with at least a college degree, geographic controls, and racial fractionalization.

A potential concern is that our results are merely a mechanical result of the differing geographic distributions of Democratic and Republican CDs—Democratic CDs could lie farther from projects available for donations. Column (4) shows that the results are robust to controlling for the average distance from a given CD to all projects.

[13]To mitigate measurement error in the estimation of CD-level coefficients θ_i, we shrink these coefficients to the sample mean by their signal-to-noise ratio, see Appendix F.2.1. Universalism is measured very precisely at the CD level due to the large underlying sample of donations, so the shrinkage does not

Table 5: Vote shares and universalism in altruism across Congressional Districts

	Dependent variable: Effective Democratic vote share 2016 (in %)					
	(1)	(2)	(3)	(4)	(5)	(6)
Universalism in altruism (wrt geographic distance)	10.3***	13.5***	11.1***	9.54***		
	(0.66)	(1.18)	(1.47)	(1.40)		
Universalism in altruism (wrt friendship distance)					8.83***	3.55***
					(0.72)	(1.10)
Log [1 + Total donations]			2.81**	1.97**		2.87**
			(1.09)	(0.99)		(1.19)
Log [Median household income]				-45.1***		-45.1***
				(5.32)		(5.53)
Fraction of population with college degree				79.8***		89.4***
				(12.84)		(13.13)
Latitude				1.08**		0.48
				(0.55)		(0.60)
Log [Distance to coast]				-2.04***		-2.06***
				(0.57)		(0.71)
Racial fractionalization				18.9***		20.2***
				(5.84)		(6.03)
Log [Average distance to all projects]				72.2***		64.4***
				(15.11)		(17.96)
State FE	No	Yes	Yes	Yes	No	Yes
Observations	436	436	436	436	436	436
R^2	0.33	0.48	0.49	0.64	0.25	0.61

Notes. OLS estimates, robust standard errors in parentheses. Effective Democratic vote shares are given by Demoractic vote share as a fraction of Democratic and Republican vote share. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Finally, we present an extension in which universalism in altruism is computed based on *social* rather than geographic distance. When estimating equation (9), we use as $d_{i,j}$ the probability that two individuals from different CDs are friends on Facebook (Bailey et al., 2018); Appendix F.2.2 describes this measure in greater detail. Columns (5)–(6) of Table 5 show that very similar results hold with this alternative distance measure. This shows that our results do not merely reflect the fact that Democrats' friends are located further away than Republicans' friends. Instead, *holding fixed a given level of friendship distance*, Democrats give relatively less if friendship distance is small and relatively more if friendship distance is large. That is, Republicans treat close friends "better" than Democrats, but Democrats treat distant strangers "better" than Republicans.

meaningfully impact our results – the correlation between the raw and shrunk measures is 0.99.

6.4 Robustness Checks

Controlling for local sources of education funding. A limitation of our analysis is that we estimate universalism only from DonorsChoose data, and do not observe giving outside of this platform. This would be problematic if, for example, variation in universalism across CDs was generated only as an artefact of variation in amounts given locally through other means in each CD. A prime candidate in this respect is the public school funding system, e.g., payments through local property taxes. Table 16 in Appendix F.3 shows that controlling for the per capita amount of primary and secondary education spending derived from local revenue sources does not affect the results.

Geographic distributions of CDs by party. Another potential concern pertains to differences in the geographic distribution of red and blue CDs. To address this, we implement two robustness checks. First, we re-estimate universalism after re-coding geographic distance into a binary variable, based on a distance threshold of 50 miles. Thus, this measure of universalism only leverages variation in whether donations are "local" or "distant." Long-distance coast-to-coast donations are hence treated just like other non-local donations. As a second robustness check, we add state-pair fixed-effects to the baseline analysis. That is, our analysis fixes a donor state and a recipient state and only leverages variation in distance within these states, say from Massachusetts to Vermont. The results in these two robustness checks are very similar. See Appendix F.3 and Table 16 for details.

7 Conclusion

This paper has proposed that individual-level heterogeneity in universalism in both altruism and trust accounts for the particular structure of policy views observed in the West. As discussed in Section 2, our analysis is conditional on two restrictions. First, we only analyze 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 (forthcoming) 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. As discussed in Section 5.5, the connection between the structure of ideology and morality might be different outside the West for various reasons. While we find that for the majority of the population in Brazil and Korea universalism is not very predictive of policy views, rich and educated elites appear quite similar to Western populations in how universalism is related to policy preferences. A possible conjecture is that as these countries get richer and / or also undergo a transformation towards "post-material" values (Inglehart, 1997), correlations between universalism and an internally consistent cluster of policy views may emerge also for the broader population.

References

Alesina, Alberto and Edward Glaeser, *Fighting poverty in the US and Europe: A world of difference*, Oxford University Press, 2004.

_ , Armando Miano, and Stefanie Stantcheva, "Immigration and redistribution," Technical Report, National Bureau of Economic Research 2018.

_ , Reza Baqir, and William Easterly, "Public goods and ethnic divisions," *The Quarterly Journal of Economics*, 1999, *114* (4), 1243–1284.

Algan, Yann, Elizabeth Beasley, Daniel Cohen, and Martial Foucault, "The rise of populism and the collapse of the left-right paradigm: lessons from the 2017 French presidential election," 2018.

Bailey, Michael, Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong, "Social Connectedness: Measurement, Determinants, and Effects," *Journal of Economic Perspectives*, 2018, *32* (3), 259–280.

Bénabou, Roland and Jean Tirole, "Identity, morals, and taboos: Beliefs as assets," *The Quarterly Journal of Economics*, 2011, *126* (2), 805–855.

_ , Armin Falk, and Jean Tirole, "Narratives, Imperatives and Moral Persuasion," Technical Report, Working Paper 2020.

Besley, Timothy and Torsten Persson, "The Rise of Identity Politics," 2019.

Bó, Ernesto Dal and Pedro Dal Bó, "Do the right thing: The effects of moral suasion on cooperation," *Journal of Public Economics*, 2014, *117*, 28–38.

_ , Frederico Finan, Olle Folke, Torsten Persson, and Johanna Rickne, "Economic Losers and political winners: Sweden's radical right," *Manuscript in preparation*, 2018.

Boxell, Levi, Matthew Gentzkow, and Jesse M Shapiro, "Cross-country trends in affective polarization," Technical Report, National Bureau of Economic Research 2020.

Chetty, Raj and Nathaniel Hendren, "The impacts of neighborhoods on intergenerational mobility II: County-level estimates," *The Quarterly Journal of Economics*, 2018, *133* (3), 1163–1228.

Draca, Mirko and Carlo Schwarz, "How polarised are citizens? Measuring ideology from the ground up," *Working Paper*, 2020.

Enke, Benjamin, "Kinship, cooperation, and the evolution of moral systems," *The Quarterly Journal of Economics*, 2019, *134* (2), 953–1019.

_ , "Moral Values and Voting," *Journal of Political Economy*, forthcoming.

_ , **Ricardo Rodríguez-Padilla, and Florian Zimmermann**, "Moral Universalism: Measurement and Heterogeneity," *Working Paper*, 2020.

Epper, Thomas, Fehr Ernst, and Senn Julien, "Other-regarding preferences and redistributive politics," Technical Report, Department of Economics-University of Zurich 2020.

Fehr, Dietmar, Johanna Mollerstrom, and Ricardo Perez-Truglia, "Your Place in the World: The Demand for National and Global Redistribution," Technical Report, National Bureau of Economic Research 2019.

Fisman, Raymond, Pamela Jakiela, and Shachar Kariv, "Distributional preferences and political behavior," *Journal of Public Economics*, 2017, *155*, 1–10.

Friedman, Benjamin M, *The moral consequences of economic growth*, Vintage, 2006.

Gennaioli, Nicola and Guido Tabellini, "Identity, Beliefs, and Political Conflict," *Working Paper*, 2019.

Gilens, Martin, *Why Americans hate welfare: Race, media, and the politics of antipoverty policy*, University of Chicago Press, 2009.

Gillen, Ben, Erik Snowberg, and Leeat Yariv, "Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study," *Working Paper*, 2015.

_ , _ , **and** _ , "Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study," *Journal of Political Economy*, 2019, *127* (4), 1826–1863.

Grossman, Gene and Elhanan Helpman, "Identity Politics and Trade Policy," *Working Paper*, 2018.

Guriev, Sergei and Elias Papaioannou, "The political economy of populism," 2020.

Haidt, Jonathan, *The Righteous Mind: Why Good People are Divided by Politics and Religion*, Vintage, 2012.

Hibbing, John R, Kevin B Smith, and John R Alford, "Differences in negativity bias underlie variations in political ideology," *Behavioral and Brain Sciences*, 2014, *37*, 297–350.

Inglehart, Ronald, *Modernization and postmodernization: Cultural, economic, and political change in 43 societies*, Princeton university press, 1997.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J Westwood, "The origins and consequences of affective polarization in the United States," *Annual Review of Political Science*, 2019, *22*, 129–146.

Jost, John T, Christopher M Federico, and Jaime L Napier, "Political ideology: Its structure, functions, and elective affinities," *Annual review of psychology*, 2009, *60*, 307–337.

Kerschbamer, Rudolf and Daniel Müller, "Social preferences and political attitudes: An online experiment on a large heterogeneous sample," *Journal of Public Economics*, 2020, *182*, 104076.

Kozlowski, Austin and James P Murphy, "Issue Alignment and Partisanship in the American Public: Revisiting the'Partisans without Constraint'Thesis," 2019.

Kranton, Rachel E and Seth G Sanders, "Groupy versus non-groupy social preferences: Personality, region, and political party," *American Economic Review: Papers & Proceedings*, 2017, *107* (5), 65–69.

Kuziemko, Ilyana and Ebonya Washington, "Why Did the Democrats Lose the South? Bringing New Data to an Old Debate," *American Economic Review*, 2018.

Levine, David K, "Modeling altruism and spitefulness in experiments," *Review of economic dynamics*, 1998, *1* (3), 593–622.

Luttmer, Erzo FP, "Group loyalty and the taste for redistribution," *Journal of political Economy*, 2001, *109* (3), 500–528.

Malka, Ariel, Yphtach Lelkes, and Christopher J Soto, "Are cultural and economic conservatism positively correlated? A large-scale cross-national test," *British Journal of Political Science*, 2019, *49* (3), 1045–1069.

Poole, Keith T and Howard Rosenthal, *Congress: A political-economic history of roll call voting*, Oxford University Press on Demand, 2000.

Putnam, Robert D, *Bowling alone: The collapse and revival of American community*, Simon and schuster, 2000.

Rehm, Philipp and Timothy Reilly, "United we stand: Constituency homogeneity and comparative party polarization," *Electoral Studies*, 2010, *29* (1), 40–53.

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

Tabellini, Guido, "The Scope of Cooperation: Values and Incentives," *Quarterly Journal of Economics*, 2008, *123* (3), 905–950.

Waytz, Adam, Ravi Iyer, Liane Young, Jonathan Haidt, and Jesse Graham, "Ideological differences in the expanse of the moral circle," *Nature Communications*, 2019, *10* (1), 1–12.

Wu, Alex, "Political polarization," *Undergraduate senior thesis, Harvard University*, 2020.