Culture, Ethnicity and Diversity*

Klaus Desmet
Universidad Carlos III and CEPR

Ignacio Ortuño-Ortín
Universidad Carlos III

Romain Wacziarg
UCLA, NBER and CEPR

First Draft - April 2014

Abstract

We investigate the relationship between ethnicity and culture, defined as a vector of traits reflecting norms, attitudes and preferences. Using surveys of cultural attitudes and values, we find that ethnic identity is a significant predictor of cultural attitudes, yet that cultural diversity is uncorrelated with ethnic diversity. The reason is that the degree of cultural heterogeneity across ethnic groups is small relative to total cultural heterogeneity. We propose new measures of the degree of overlap between culture and ethnicity, χ^2 and F_{ST}, that stem directly from a simple model of social antagonism. We study the cross-country correlates of cultural diversity and of the overlap between culture and ethnicity. Finally, we find that our proposed overlap measures are strong and robust predictors of civil conflict, that cultural diversity tends to reduce the incidence of conflict, and that ethnic fractionalization is not significantly related to conflict. Thus, civil conflict is more likely when culture and ethnicity reinforce each other.

*Desmet: Department of Economics, Universidad Carlos III de Madrid, 28903 Getafe, Madrid, Spain, klaus.desmet@uc3m.es; Ortuño Ortín: Department of Economics, Universidad Carlos III de Madrid, 28903 Getafe, Madrid, Spain, iortuno@eco.uc3m.es; Wacziarg: UCLA Anderson School of Management, 110 Westwood Plaza, Los Angeles CA 90095, wacziarg@ucla.edu.

1 Introduction

1.1 Motivation

Ethnic diversity affects a wide range of political economy outcomes, but the mechanisms through which it operates are not well understood. Ethnic diversity could work through transaction costs: groups in ethnically heterogeneous societies may share common goals, but experience difficulties in coordinating due to racial animosity, linguistic barriers, or lack of trust. In this case, ethnic heterogeneity will matter even in societies where groups share the same norms, values and preferences - the broad set of preferences and attitudes that we will refer to as "culture". Alternatively, ethnically distinct groups may have different values or preferences over social goals, encompassing public goods and broader policy objectives. In this case, it is cultural diversity, not ethnic diversity *per se* that should matter to explain political economy outcomes. Another hypothesis, which we propose and explore here, is that dysfunction occurs only when there is some overlap between culture and ethnic identity, neither form of diversity being sufficient in its own right to affect political economy outcomes.

In this paper, we conduct a systematic exploration of the links between culture and ethnicity. We propose new measures capturing cultural diversity as well as the overlap between culture and ethnicity, and examine how these measures vary across countries. In doing so, we shed light on the channels linking human heterogeneity to political economy outcomes. As an application, we highlight the important role played by the degree of overlap between culture and ethnicity as a determinant of civil conflict.

The idea that ethnic and linguistic groups differ in terms of norms, values and preferences is an often unstated and usually untested pillar of the growing economics literature on the effects of ethnolinguististic diversity. For instance, economists usually assume that ethnically heterogeneous communities tend to spend less on public goods because they have different preferences over such public goods (Alesina, Baqir and Easterly, 1999, Alesina, Baqir and Hoxby, 2004).[1] Similarly, economists usually take it as given that ethnic groups may engage in violent conflict because they disagree strongly over shared policies and public goods (Esteban and Ray, 2011, Esteban, Mayoral

[1] The opening sentences in Alesina, Baqir and Easterly (1999) are: "When individuals have different preferences, they want to pull fewer resources together for public projects. This paper argues that certain public goods (...) supplied by US cities are inversely related to ethnic fragmentation in those cities."

and Ray, 2012).[2] The common assumption is that heterogeneous ethnolinguistic groups do not share the same preferences, broadly construed to include social goals, norms and values. Ethnic differences are merely a proxy for diversity in preferences, norms and values. Under this view, we should expect cultural diversity to matter for conflict, not ethnic diversity *per se*.

Ethnic heterogeneity could also be associated more directly with conflict between groups. Groups may have identical preferences and attitudes, but differ racially or linguistically. These differing attributes, which are purely ethnic markers, may foster conflict by allowing one group to exclude the other from the spoils of war (Caselli and Coleman, 2013). Ethnic identity can also introduce barriers between groups, increasing animosity and reducing trust and coordination independently from any fundamental differences in preferences (Glaeser, 2005, Guiso, Sapienza and Zingales, 2009).[3] In this case, we should expect more conflict even in societies where cultural diversity is low, as long as ethnic diversity is high.

A final possibility, which we propose and investigate in this paper, is that ethnic diversity matters more for conflict when norms, values and preferences overlap with ethnic identity. Heterogeneity in preferences, norms and values may have benign effects if society is ethnically homogeneous, as coordination among groups may be facilitated. Conversely, ethnic heterogeneity may not *per se* impede coordination when there is social consensus. Under this view, what matters for conflict and economic outcomes is when ethnic and cultural heterogeneity reinforce each other: animosity between ethnic groups only arises when these groups also differ in terms of values, norms and preferences. A related idea exists in the political science literature on cross-cutting cleavages, but with important differences. This literature has explored whether the extent to which multiple dimensions of heterogeneity cross-cut or reinforce each other matters for political economy outcomes (Rae and Taylor, 1970, Selway, 2010). However, it has not examined the specific dimensions of heterogeneity we are interested in, namely heterogeneity in cultural attitudes and in ethnic composition. We also depart from much of this literature by developing theoretically grounded measures of the overlap between these multiple dimensions of heterogeneity.

[2]For instance, Esteban, Mayoral and Ray (2012) state: "[we] employ the linguistic distance between two groups as an appropriate indicator for their difference in preferences over public goods."

[3]For an early survey of the effect of culture on economic outcomes see Guiso, Sapienza and Zingales (2006).

1.2 Contribution

This paper seeks to make several contributions to the literature on ethnic divisions. We first explore systematically the relationship between ethnolinguistic identity and culture, using individual-level data from various surveys such as the World Values Survey. We seek to explain answers on norms, values and preferences using a respondent's economic and demographic characteristics, among which are ethnic and linguistic indicators, and to evaluate the joint statistical significance of the latter. We find that ethnicity dummy variables are jointly significant predictors of responses in about half of the cases, although this average masks significant heterogeneity across countries. Thus, on average ethnic identity appears to be an important determinant of cultural norms, values and preferences.

Second, we compute measures of cultural fractionalization, using the answers to survey questions as the basis for identifying heterogeneity in culture. That is, we now seek to understand the relationship, not between identity and culture, but between cultural heterogeneity and ethnic diversity. We find that heterogeneity in norms, values and preferences is uncorrelated with ethnolinguistic fractionalization across countries. Taken together, these results show that even though culture does differ across ethnolinguistic groups, cultural fractionalization and ethnolinguistic fractionalization are not related. In contrast to what the literature has usually posited, then, ethnic fractionalization cannot so readily be taken as a proxy for overall cultural and preference heterogeneity.

How can these seemingly contradictory results be reconciled? Within-group heterogeneity in culture may account for the low correlation between cultural heterogeneity and ethnolinguistic diversity, without precluding the possibility that ethnic identity has predictive power for cultural attitudes: the degree of between-ethnic group cultural heterogeneity could be small relative to total heterogeneity, yet have significant predictive power for various political economy outcomes. Thus, in a third step we propose new indices of the degree of overlap between ethnicity and culture, derived from a simple model of social antagonism. The specific measure we derive from the model is a χ^2 index, capturing the average distance between the answers of each ethnic group and the answers in the overall population. A low value of the index indicates that groups reflect the countrywide distribution of answers, while a high value indicates a lot of group-specificity with respect to the distribution of answers. The second index, developed in the context of population genetics, is known as a *fixation index*, or F_{ST}. It captures the share of between-group variance in answers to survey questions to overall variance. A value of zero indicates that there is no informational content to knowing an individual's ethnic identity, while a value of one indicates that answers can

be perfectly predicted based on knowing an individual's ethnic identity. These indices will be more formally defined below. For now, it suffices to say that, heuristically, these measures capture the degree of overlap between cultural attitudes and ethnolinguistic identity. We find that the degree to which cultural and ethnic cleavages overlap is small, yet there is substantial variation across countries in the F_{ST} and χ^2 measures, and this variation is related in interesting ways to some salient cross-sectional characteristics of countries.

In a fourth and final step, we use our new measures in a specific application. We explore the effect of ethnic heterogeneity, cultural heterogeneity and the degree to which cultural attitudes and ethnic identity reinforce each other on the onset and incidence of civil conflict. In principle, civil wars could arise when preferences are very diverse (high cultural heterogeneity), when ethnic groups experience a high degree of antagonism (high ethnic diversity), or when culture and ethnicity reinforce each other. Empirically, we find that both cultural and ethnic diversity have benign effects on civil conflict. If anything, cultural diversity has a pacifying effect. However, the degree of overlap between cultural attitudes and ethnic identity has a strong and robust effect on civil wars: when culture and ethnicity reinforce each other (i.e. for high values of F_{ST} or χ^2) violent conflict becomes more likely.

1.3 Related Literature

This paper is related to various strands of the literature on ethnolinguistic diversity. The first strand studies the relationship between ethnolinguistic diversity and socioeconomic outcomes, using conventional measures of diversity such as fractionalization (for instance, Easterly and Levine, 1997, Alesina, Baqir and Easterly, 1999, Alesina et al., 2003, Alesina and La Ferrara, 2005, among many others). Our paper is related to this literature as we examine the effect of ethnic and cultural fractionalization on one particular outcome, civil conflict. By explicitly considering cultural diversity and its relation with ethnic heterogeneity, we cast light on the mechanisms that led to the empirical regularities uncovered in the earlier literature.

The second strand seeks to advance the measurement of diversity by considering alternative measures that improve on simple fractionalization. These measures take different forms, accounting for distance between groups (Esteban and Ray, 1994, 2004, Bossert, d'Ambrosio and La Ferrara, 2011), looking at income inequality between ethnic groups (Alesina, Michalopoulos and Papaioannou, 2012) or the historical depth of ethnic cleavages (Desmet, Ortuño-Ortín and Wacziarg, 2012).

Our paper is related to this measurement literature because we propose a new measure of heterogeneity in cultural attitudes and propose measures of the degree of overlap between culture and ethnicity. These measures shed new light on the complex empirical relationship between culture and ethnicity.

A third strand of the literature examines the relationship between culture and economic outcomes. This literature usually examines the effect of a particular historically-determined trait on current outcomes, rather than the effect of cultural *diversity* as we do. This is, again, a vast literature, but salient examples include Alesina, Giuliano and Nunn (2011, 2013) on the historical legacy of the heavy plough on values affecting fertility and female labor force participation; Giuliano (2007) on the effect of culture on living arrangements; Fernandez and Fogli (2009) on culture, fertility and female labor force participation; Luttmer and Singhal (2011) on culture and the taste for redistribution; Tabellini (2010) on cultural traits and economic performance across the regions of Europe; and Guiso, Sapienza and Zingales (2009) on trust and bilateral trade. In contrast to this literature, we study the effect of cultural heterogeneity rather than the effect of a specific cultural trait.

Finally, a recent literature seeks to relate genetic differences - a measure associated with cultural differences - with political and economic outcomes, including conflict. For instance, Spolaore and Wacziarg (2009) look at the effect of genetic distance between countries on the diffusion of the Industrial Revolution and Spolaore and Wacziarg (2013) study the effect of genetic distance between countries on interstate conflict and war. While these two studies also use F_{ST} as a measure of distance between groups, this F_{ST} is based on genetic rather than cultural data, and it is used to study interactions between countries rather than between groups within countries. Ashraf and Galor (2013) investigate the effect of genetic diversity, used as a broader measure of diversity in both cultural and biological traits within countries, on historical and contemporary economic performance, finding a non-monotonic effect. In Arbati, Ashraf and Galor (2013), the same measure of genetic diversity is found to have a positive effect on the probability of civil conflict. In contrast to these papers, we measure cultural diversity directly using responses to surveys on norms, attitudes and preferences, rather than using genetic data.

1.4 Outline

This paper is organized as follows. In Section 2, we use individual level data from a variety of cultural attitudes surveys to examine systematically the relationship between ethnic identity and cultural attitudes. In Section 3, we introduce a simple model of social antagonism leading to three classes of measures of heterogeneity hypothesized to affect socioeconomic outcomes. We show how to operationalize these theoretically derived measures using data on ethnicity and cultural traits. In Section 4, we introduce our new measures of heterogeneity, compute them using the World Value Survey, and describe their interrelationships and determinants. In Section 5, as an application for our new measures, we explore empirically the effect of cultural and ethnic heterogeneity on civil conflict. Section 6 concludes.

2 Identity and Culture

2.1 Methodology

In this section we use a variety of surveys to examine the relationship between ethnic identity and cultural attitudes. The exercise requires individual level data on answers to questions on norms, values and preferences, and corresponding data on the respondent's ethnic or linguistic identity. We seek to examine the joint significance of indicators of ethnolinguistic identity on survey responses, proceeding question by question and country by country and controlling for observable individual characteristics. In principle, 5% of the questions should feature a significant joint effect of ethnic identity if the statistical criterion is 95% confidence and there was in fact no association between cultural attitudes and ethnicity. We ask whether the share of questions for which there is a significant effect of ethnicity is actually higher, and find that it is in fact much higher, on the order of half in general. We also examine whether the importance of identity for culture varies in systematic ways across question types, countries, continents, etc.

For each survey, each question and each country, we estimate the following specification:

$$Q_m = \alpha + \sum_{s=1}^{S} \beta_s D_m^s + \boldsymbol{\gamma}' \mathbf{X}_m + \varepsilon_m \tag{1}$$

where m denotes a respondent, $s = 1, ..., S$ indexes ethnolinguistic groups, Q_m is individual m's answer to the question under consideration, D_m^s is equal to one if respondent m is part of group s, zero otherwise, and $\mathbf{X}_m$ is a vector of controls. Estimation is by least squares.

We test for the joint significance of the β_s parameters using conventional F-tests. We do so for all questions and countries, and then examine the share of questions for which ethnolinguistic identity is a significant predictor of cultural attitudes at the 5% level. We compute these shares over different categories of questions, for each country separately, and for different regions.

2.2 World Values Survey: Data and Results

Our main survey is the Integrated World Values Survey-European Values Survey (WVS-EVS) dataset covering 1981 to 2008 and five survey waves. In order to examine the relationship between ethnicity and culture in a systematic way, we choose to focus on the broadest set of available questions without casting judgment on which ones are more representative of attitudes and preferences: we let the dataset largely guide our choice of questions, as opposed to making *ad hoc* choices ourselves. In the WVS-EVS integrated dataset, there is a total of $1,031$ fields, or questions. Some of these fields are not survey questions but instead refer to socio-demographic characteristics of the respondent or the interviewer, and some have zero observations. We confine attention to survey questions identified by the survey itself as pertaining to norms, values and attitudes (these come grouped by the survey organization into question categories labelled from A to G), and with a nonzero number of respondents. Among those, in very rare cases some questions were asked in a slightly different manner in some countries (Colombia, Hong Kong, Mexico, Iraq), and those were dropped (19 questions). We also dropped questions that asked about circumstances specific to a given country, i.e. questions that could not conceivably be asked in more than one country (74 questions). In the end we were left with 808 questions.

Among these remaining questions, there were three types: those with a binary response (yes/no, agree/disagree: 252 questions), those with an ordered response (where answers are on a scale of, say, 1 to 10: 496 questions), and those with strictly more than two possible responses that are not naturally ordered (60 questions). The first two categories can be used readily as dependent variables. For the third category, we cannot directly estimate the joint effect of ethnicity on unordered responses, so we transformed each possible response into a series of binary response questions.[4] Thus, the 60 questions with unordered responses resulted in 193 new binary questions,

[4]For instance, question C009 asks "Regardless of whether you're actually looking for a job, which one would you, personally, place first if you were looking for a job?" and offers the following choices: "a good income", "a safe job with no risk", "working with people you like", "doing an important job", "do something for community". We define 5 binary response questions, where, for instance, for "a good income", the response value is 1 if the respondent

leading to a total of 941 questions. Of course, not every one of these questions was asked in every country, or in every wave. We keep all questions irrespective of where or when they were asked. In the end, out of 941 questions, on average 294 were asked in each country (the number of questions per country varied between 81 and 447 - Appendix Table A1 provides the exact count, country by country). When combined across all waves, the average number of respondents across the countries in the sample, and across all questions, was $1,497$. There is some heterogeneity around this number as some questions were asked in more waves than others, and the number of surveyed individuals varies across countries and waves.

An important aspect of our exercise is to correctly code ethnolinguistic identity in order to estimate the joint effect of ethnicity dummies on responses. To do so, we have to define ethnicity. The WVS/EVS asks respondents to report both their ethnicity and language. In some cases, the reported ethnic categories do not appropriately capture ethnic identity. For instance, in many cases for African countries the WVS/EVS integrated survey reports ethnicities as White / Black. For instance in Zambia, 99.47% of respondents are Black, while there are 0.27% Asians and 0.27% Whites. Most ethnographers agree that for Africa, language is a better measure of ethnic identity than race. For Zambia, WVS/EVS respondents speak 18 separate languages, the largest of which (Bemba) represents 36.6% of the respondents. The opposite problem exists in Latin America, where language is not usually used as a measure of ethnic affiliation, and race defines ethnic identity instead. For instance, in Venezuela 100% of respondents report speaking Castilian. However the largest racial group is coded as "Colored (light)", representing 42.7% of respondents.

To correctly characterize ethnic identity in a systematic way, we again rely on existing classifications rather than our own judgement. We examine the ethnic and linguistic classifications in the integrated WVS/EVS file and see which one is closest to either the Alesina et al. (2003) or the Fearon (2003) classifications, which are widely used in the literature.[5] We choose either ethnic identity or linguistic identity depending on which one gives us a classification and a distribution of individuals across groups that most resembles the Alesina et al. and Fearon classifications. In the above example, ethnic identity in Zambia is coded using the language spoken at home variable, while ethnic identity in Venezuela is coded as the ethnic group to which a respondent belongs. The

answered "a good income" to question C009, and zero otherwise, and so on for the other answer categories.

[5]The WVS/EVS question on ethnic group is question x051 while the language spoken at home question is g016. These are the two questions we use to code a respondent's ethnic identity.

idea is that a measure of ethnolinguistic fractionalization computed from the resulting group shares in the WVS/EVS dataset should be highly correlated with fractionalization measures derived in Alesina et al., and Fearon. Indeed our ethnic classification results in fractionalization measures that are 74% correlated with fractionalization in Alesina et al., and 73% correlated with fractionalization according to Fearon - this despite the data coming from very different sources (a survey for WVS/EVS, mostly census for the other two sources).

Finally, control variables in the WVS/EVS dataset consist of the respondent's age (question x003), sex (x001), highest educational level attained (x025) and household income (x047).

The results for the WVS/EVS dataset are presented in Tables 1 and 2. Table 1 presents the overall share of regressions where ethnicity dummies are jointly significant at the 5% level, as well as a regional breakdown. Table 2 presents a breakdown by question category (using the classification of questions provided by the WVS/EVS) and by question type (binary, scale, and binary constructed from multiple response questions). Additionally, Appendix Table A1 presents the share of regressions with significant ethnic dummies country by country.

Interesting findings emerge. Firstly, the average number of questions for which ethnicity dummies are jointly significant, across all countries, is 43%. Thus, ethnic identity is an important determinant of responses to many questions. Secondly, this average masks interesting variation across regions. In South Asia, East Asia and Sub-Saharan Africa, the shares are much higher, respectively 67%, 63% and 62%. In Latin America and Western Europe, the shares are much lower, at 17% and 31% respectively. The small share in Latin America could be due to the fact that, despite racial heterogeneity, linguistic and religious identity in Latin America is much more homogeneous than in places where ethnic identity is a stronger predictor of culture, for instance Africa. The Latin American exception is not specific to the New World, as North America (defined here as Canada and the US) displays a relatively high share (51%). Thirdly, the breakdown by question category shows little variation. We find that ethnic identity matters a bit more for questions pertaining to religion and morals, as well as (predictably) for those pertaining to national identity, and a bit less for questions related to work. Otherwise, there is substantial homogeneity across categories. We conducted the same breakdown by question category continent by continent, finding again little variation in the share of regressions with significant ethnic dummies. These findings suggest that the choice of questions is not very material to the issue of whether ethnic identity affects norms,

values and preferences, as regional patterns are stable across question categories.[6]

2.3 Other Surveys

Afrobarometer. We also conduct a similar analysis using the 4^{th} wave of Afrobarometer (2008), the latest round available as of this writing. The wave covers 20 Sub-Saharan African countries. There are 182 questions on values, norms and preferences in this survey, 7 admitting binary response, 155 where the response is on an ordered scale and 20 which admit more than two unordered responses. These 20 questions were converted into a set of 117 binary questions in the same way as was done for the WVS/EVS, resulting in a total of 279 available questions for us to analyze. A specificity of Afrobarometer is that most questions were asked in every country, so there is little heterogeneity in the set of questions used for the analysis for this dataset (the mean number of questions asked in each of the 20 surveyed countries is 271, with a standard deviation of 5).

For ethnic identity, we relied entirely on the classification provided by Afrobarometer (variable Q79: "What is your tribe or ethnic group?"). The number of ethnic groups varies by country, ranging from 5 to 38. Finally, the regression specification includes as controls the respondent's age (Q1), gender (Q101) and present living conditions (Q4B) as a proxy for income (no measures of human capital or direct measures of income are available).

Results obtained using the Afrobarometer survey are presented in Appendix Table A2. Confirming results for Sub-Saharan Africa from the WVS/EVS, the share of questions for which ethnicity significantly predicts responses is high, on average 57% (versus 62% in the WVS/EVS). This average masks interesting variation across countries. Some countries like South Africa and Nigeria display very high shares (88% and 91%, respectively), while some small countries such as Lesotho and Cape Verde display much smaller shares (respectively 10% and 20%).[7] Interestingly, the results for

[6] Similarly, we find little variation across types of questions - binary, scale or binary constructed from unordered response questions. Ethnicity predicts answers to scale questions slightly more frequently than for binary questions, but the difference is not large. This again suggests that the specific choice of questions is not very material to our results.

[7] For Lesotho, the ethnic nomenclature provided by Afrobarometer is at the level of clans. Lesotho is considered by ethnographers to be a very ethnically homogeneous country (99.7% of the population belongs to the Sotho ethnic group, according to the CIA World Factbook, 2009). This homogeneity is reflected in the fact that different clans do not seem to differ much in terms of values, norms and preferences. Despite the questionable ethnic classification proposed for Lesotho by the Afrobarometer, we chose to leave this country in our sample. Excluding Lesotho from the

South Africa and Nigeria mirror those from the WVS/EVS, where the corresponding shares were 88% and 84%, despite different survey questions, ethnic nomenclature, methods and samples.

Latinobarómetro. Finally, we used the Latinobarómetro, which covers Latin American countries. The survey asks a question on ethnic identity only since 2007, so we are constrained to using waves for 2007, 2008, 2009 and 2010, covering 18 countries. The survey classifies questions into categories, and we focus on the categories that refer to values, opinions, and preferences. Among those, however, a few questions about the factual situation of the respondent crept in, and we removed them from consideration. We ended up with 231 binary response questions, 358 questions for which the response is on an ordered scale, and 19 questions with more than two unordered response categories. The latter were transformed into a series of 96 binary response questions, as before. We ended up with a total of 685 usable questions. As was the case for the Afrobarometer, the availability of questions across countries did not vary much. The average country had 678 questions (with a standard deviation of less than 2).

Ethnic identity is as defined by Latinobarómetro, and represents a classification very similar to the one we used for Latin American countries in the WVS/EVS. There are seven ethnic categories, corresponding to the respondent's race (the variable coding ethnicity is named A505206). The categories are Asian, Black, Indigenous, Mestizo, Mulatto, White and other race. The regression specification includes ethnic dummies, sex (S01), age (S02), education (S51) and socioeconomic level (S62) as a proxy for income.

Results obtained using the Latinobarómetro survey are presented in Appendix Table A3. The share of questions for which ethnicity dummies are jointly significant correlates of answers is 32.5%. As in the WVS/EVS survey, this average does not seem to vary much across question categories or question type (although, as before, the share is slightly smaller for binary questions created from underlying unordered multiple response questions). There appears to be only limited variation across countries. Argentina displays the lowest share in this sample (17.8%) while the Dominican Republic displays the largest (60%), but most shares are comprised in a tight band between 20% and 40%. The relatively low share obtained with Latinobarómetro is in line with results obtained previously using the WVS/EVS.

Afrobarometer sample raises the average share of regressions where ethnicity significantly predicts culture to 59.4%.

2.4 Summary

Several interesting lessons emerge from this exercise. First, on average, there is a relatively large share of questions for which ethnic identity is a significant predictor of cultural attitudes - often more than 50%. Second, there is substantial variation across regions, with particularly high overlap between ethnic identity and culture in Asia and Sub-Saharan Africa, and relatively low overlap in Latin America. Europe and North Africa are somewhere in between. Third, the results are largely unchanged across different surveys: with Afrobarometer and Latinobarómetro we largely confirm results found using the WVS/EVS. Fourth, this exercise allowed us to characterize whether ethnicity was a statistically significant predictor of culture, but we did not characterize the magnitude of the overlap between these two dimensions of heterogeneity. In what follows we turn to a novel measurement framework in which we will explicitly quantify the extent to which culture and ethnicity overlap with each other.

3 Measuring Heterogeneity

This section is about measurement. We previously argued that the extent to which ethnic identity can predict answers to questions on attitudes, values and preferences varies across countries. Here we seek to quantify the extent of this overlap between culture and ethnicity. We present a simple model of social antagonism to guide the choice of functional forms for the heterogeneity measures used in the empirical investigation that follows. From this framework, under various assumptions about the source of heterogeneity giving rise to antagonism, we derive indices of ethnic diversity, cultural diversity and the overlap between culture and ethnicity: F_{ST} and χ^2. We then show how to operationalize these theoretical indices using data.

In a nutshell, we assume that individuals feel antagonism towards people who are different from them. *Social antagonism* is the sum of all the individual levels of antagonism in society, as in the alienation framework of Esteban and Ray (1994). We adopt a broad interpretation of what antagonism captures. It could represent feelings of alienation felt toward groups with different cultures or different ethnicities. Antagonism could also stem from barriers that prevent fruitful interactions between groups, for instance due to an inability to communicate or trust each other. We consider three distinct types of societies depending on how various dimensions of heterogeneity give rise to antagonism. For each type of society we derive an index measuring the level of social antagonism. Later, we will calculate these indices and relate them with the probability of civil

conflict.

It is useful to start with some notation. A country is composed of n individuals characterized by the ethnic group to which they belong and by their cultural values or preferences. There are S ethnic groups, indexed by $s = 1, ..., S$. The share of each ethnic group in the population is w^s. Cultural values (or traits) are the answers to the q questions in the WVS (or any other survey of cultural attitudes), each indexed by $i = 1, ..., q$. Each question i has $r(i)$ different possible answers, indexed by j. Focusing on a given country, w_{ij} is the share of the population that gives answer j to question i. Finally, w_{ij}^s is the share of individuals from ethnic group s that gives answer j to question i.

The type of an individual, k, is given by his ethnic group s and his answers to the q questions. We define a vector ω_k of dimension $1 + q$ where the first component is a number from $\{1, 2, ..., S\}$ and denotes his ethnic group, and the remaining q components represent the answers to each of the corresponding q questions. For example, if there are two ethnic groups, $S = 2$, and three questions, $q = 3$, and each question has two answers, $r(i) = 2$, the vector $\omega_1 = \{1, 2, 1, 1\}$ characterizes the type of an individual (i.e. type 1) who belongs to the first ethnic group and who gives answers $2, 1, 1$ to the first, second and third question, respectively. Since we have a finite number of individuals, n, as well as a finite number of questions and answers, the total number of different types of individuals is finite. We denote by K the number of different types and by n_k the number of individuals of type k, so $\sum_{k=1}^{K} n_k = n$. The population share of individuals of type k is denoted by $w_k = n_k/n$, where of course $\sum_{k=1}^{K} w_k = 1$. We denote as $\xi(k, i)$ the answer given by an agent of type k to question i, and as $s(k)$ the ethnic group of a type k agent: $\omega_k = (s(k), \xi(k, 1), \xi(k, 2), ..., \xi(k, q))$.

3.1 The Cultural Heterogeneity Channel

We first assume that only cultural values matter for antagonism. Belonging to a different ethnic group s does not generate any antagonism *per se*. Thus, an individual's antagonism is given by the share of individuals in society who have preferences different from his. More formally, the antagonism of an agent of type k depends, for each of the q questions, on how many people respond in the same way as him. The population share of individuals that give the same answer to question i as agent of type k is $w_{i,\xi(k,i)}$.

We give the same weight to all the q questions. Thus, for an agent of type k the average share, over the q traits, of individuals giving the same answer as him is $\frac{1}{q}\sum_{i=1}^{q} w_{i,\xi(k,i)}$. Hence, his level

of antagonism, v_k, is given by:

$$v_k = 1 - \frac{1}{q}\sum_{i=1}^{q} w_{i,\xi(k,i)} \tag{2}$$

Here, individuals feel antagonism if they live in the same society as other individuals who have different cultural characteristics. Ethnicity does not matter. In this case v_k measures the average probability, over all questions, that a randomly chosen citizen disagrees with the answers given by agent k. Social antagonism is the summation of all the individual levels of antagonism, normalized by the population size n, i.e. average individual antagonism:

$$v = \sum_{k=1}^{K}\left(1 - \frac{1}{q}\sum_{i=1}^{q} w_{i,\xi(k,i)}\right)\frac{n_k}{n} \tag{3}$$

To operationalize this measure of cultural antagonism, notice that (3) can be written as:

$$v = \sum_{k=1}^{K}\left(1 - \frac{1}{q}\sum_{i=1}^{q} w_{i,\xi(k,i)}\right)w_k = 1 - \frac{1}{q}\sum_{i=1}^{q}\sum_{k=1}^{K} w_{i,\xi(k,i)}w_k$$

Next, note that:

$$\sum_{k=1}^{K} w_{i,\xi(k,i)}w_k = \sum_{j=1}^{r(i)}\sum_{k:\xi(k,i)=j} w_{i,\xi(k,i)}w_k = \sum_{j=1}^{r(i)} w_{ij}\sum_{k:\xi(k,i)=j} w_k = \sum_{j=1}^{r(i)} w_{ij}^2$$

Thus, v becomes the following easy-to-calculate index of cultural fractionalization (CF):

$$CF = \frac{1}{q}\sum_{i=1}^{q}\left(1 - \sum_{j=1}^{r(i)} w_{ij}^2\right) \tag{4}$$

The cultural fractionalization (CF) index measures the average probability that two randomly drawn individuals from a population give different answers, on average, to questions from the WVS/EVS. Thus, if we believe that antagonism is driven exclusively by differences in preferences and cultural values, the index of cultural fractionalization, CF, should matter for socioeconomic outcomes. This is what is assumed (more or less explicitly) in various contributions to the study of ethnicity in economics, such as Alesina, Baqir and Easterly (1999), Alesina, Baqir and Hoxby (2004), Esteban and Ray (2011) and Esteban, Mayoral and Ray (2012).

3.2 The Ethnic Heterogeneity Channel

Alternatively, we assume that antagonism stems only from ethnic differences, not from cultural differences. Again, this antagonism could stem from animosity *vis-à-vis* other ethnic groups (racial hatred and prejudice) or from barriers that impede interactions between groups because of lack of

communication or trust. The probability that a randomly chosen individual belongs to the ethnic group $s(k)$ is $w^{s(k)}$. We postulate that in this society the level of antagonism of an individual of type k is:

$$v_k = 1 - w^{s(k)} \tag{5}$$

Thus, under this assumption individual antagonism is just the probability that a person meets or is matched with another person from a different ethnic group. Social antagonism, v, is the average of this probability over all individuals:

$$v = \sum_{k=1}^{K} (1 - w^{s(k)}) w_k \tag{6}$$

It is easy to see that this is just the traditional index of ethnolinguistic fractionalization. Indeed, we have:

$$v = \sum_{k=1}^{K} (1 - w^{s(k)}) w_k = 1 - \sum_{s=1}^{S} \sum_{k:s(k)=s} w^{s(k)} w_k = 1 - \sum_{s=1}^{S} w^s \sum_{k:s(k)=s} w_k = 1 - \sum_{s=1}^{S} (w^s)^2 \tag{7}$$

Hence, v in this case becomes the common ELF index of ethnic fractionalization:

$$ELF = 1 - \sum_{s=1}^{S} (w^s)^2 \tag{8}$$

Thus, if we believe that antagonism is driven purely by ethnic animosity or barriers between ethnic groups, without any role for cultural differences, the conventional index of ethnolinguistic fractionalization, ELF, is what should matter for socioeconomic outcomes. Such is the case in research where antagonism can stem from ethnic differences *per se* rather than any underlying cultural differences. Studies in this spirit include Caselli and Coleman (2013), Glaeser (2005) and Guiso, Sapienza and Zingales (2009).

3.3 The Overlap Channel

3.3.1 Deriving a Measure of Overlap Between Ethnicity and Culture

As a third alternative, we assume that an individual's antagonism depends on how culturally different her group is from other ethnic groups. An individual does not experience any antagonism if people from other ethnic groups answer the questions in the WVS in the same way as people in her own ethnic group. Ethnicity only matters if ethnic groups differ in their cultural values. We also assume that in this society cultural differences between the members of the same ethnic group do not increase the level of antagonism.

Take agent of type k and question i. Suppose first that type k only interacts with agents of her own ethnic group $s(k)$. Remember that we denote by $\xi(k,i)$ the answer that agent of type k gives to question i. In this case, by definition the share of people within group $s(k)$ with the same answer to question i as agent of type k is $w^{s(k)}_{i,\xi(k,i)}$. In other words, this is the probability that a randomly chosen agent from the ethnic group $s(k)$ agrees with an agent of type k on question i.

Now assume that an agent of type k is equally likely to interact with anybody in society. In this case the probability an agent of type k agrees on question i with a randomly chosen individual in society is $w_{i,\xi(k,i)}$. If the probability $w_{i,\xi(k,i)}$ is equal to $w^{s(k)}_{i,\xi(k,i)}$, an agent of type k does not see any difference between her own ethnic group and society overall.[8] However, if the proportion of people in society overall answering $\xi(k,i)$ is lower than the corresponding proportion within her own ethnic group, the agent experiences antagonism. In particular we assume that antagonism for question i and an agent of type k, v_{ik} depends on the (relative) difference between these two shares:

$$v_{ik} = \frac{w^{s(k)}_{i,\xi(k,i)} - w_{i,\xi(k,i)}}{w_{i,\xi(k,i)}} \tag{9}$$

Notice that if $w^{s(k)}_{i,\xi(k,i)} < w_{i,\xi(k,i)}$ the individual experiences *negative antagonism*, i.e. she is happy to interact with people in society who give the same answer as she does in greater proportion than people in her own group. Suppose that I trust people, and that 50% of those in my ethnic group trust people. I feel antagonism toward the rest of society if the share of people in the rest of society that trust people is 20%, but I am quite happy if the share of people in the rest of society that trust people is 60%.

Averaging v_{ik} over all the q questions, giving the same weight to all of them, individual antagonism is:

$$v_k = \frac{1}{q}\sum_{i=1}^{q} \frac{w^{s(k)}_{i,\xi(k,i)} - w_{i,\xi(k,i)}}{w_{i,\xi(k,i)}} \tag{10}$$

Notice that if all ethnic groups are identical, i.e., if for each question the distribution of answers is independent of the distribution of ethnic groups, we have that $v_k = 0$. We add up the individual levels of antagonism across k, weighing by the population shares of each type k, to obtain social antagonism:

$$v = \frac{\sum_{k=1}^{K} v_k n_k}{n} = \frac{\sum_{k=1}^{K} \frac{1}{q}\sum_{i=1}^{q} \frac{w^{s(k)}_{i,\xi(k,i)} - w_{i,\xi(k,i)}}{w_{i,\xi(k,i)}} n_k}{n} \tag{11}$$

[8] Here we assume that the agent takes all other ethnic groups as being a unique group.

Again, if the distribution of answers within each group is the same as the distribution of answers in society overall, $v = 0$. If, on the contrary, culture and ethnicity overlap strongly, then v will be large.

To operationalize the v measure as a measure that can be calculated from data, it can be rewritten as:

$$\chi^2 = \frac{1}{q} \sum_{i=1}^{q} \sum_{s=1}^{S} \sum_{j=1}^{r(i)} \frac{w^s (w_{ij} - w_{ij}^s)^2}{w_{ij}} \tag{12}$$

The proof, which is in the Appendix, is conceptually easy but cumbersome in terms of notation.

Thus, if we believe that antagonism is driven purely by differences in culture across ethnic groups, we should observe a positive correlation between the χ^2 index of overlap and socioeconomic outcomes. Defining and using this index to measure the overlap between culture and ethnicity is the main contribution of this paper.

3.3.2 Heuristic Discussion of the χ^2 Index

To complement the discussion above, it is useful to give a heuristic sense of the meaning of the χ^2 index. χ^2 is based on comparing the distribution of average answers for a given group to the distribution of answers in the overall population. If the distribution of answers in a given ethnic group is exactly the same as in the entire population, then knowing a person's ethnic identity conveys no information about his cultural attributes. If instead the distributions are distinct, then there is overlap between ethnic identity and cultural attributes.

To measure the overlap between ethnolinguistic diversity and preference diversity we can compare the distribution of answers across groups. This is what the χ^2 accomplishes.[9] Let n_{ij}^s be the number of individuals who belong to ethnic group s and give answer j to question i. We write $n_i^s = \{n_{i1}^s, n_{i2}^s, ...n_{ir(i)}^s\}$. Under independence, the expected number of individuals that belong to ethnic group s and give answer j to question i should be $w_{ij}n^s$, while the observed frequency is n_{ij}^s. The χ^2 index is based on the difference between the observed number of individuals of an ethnic group s that give answer j and the corresponding expected number of individuals under the assumption of independence between ethnicity and answers. So for each question i:

$$\chi_i^2 = \sum_{s=1}^{S} \sum_{j=1}^{r(i)} \frac{(w_{ij}n^s - n_{ij}^s)^2}{w_{ij}n^s} \tag{13}$$

[9] See Selway (2010) for a previous use of this index to examine the overlap between religion and ethnicity.

The value of χ_i^2 depends on the group sample sizes n^s. Since different countries have different sample sizes and we want to compare different values of χ_i^2 across countries, it is better to work from group shares than from the number of individuals in each group. Thus, we can divide the χ_i^2 index by n to obtain a normalized index:[10]

$$\phi_i^2 = \sum_{s=1}^{S}\sum_{j=1}^{r(i)} \frac{n^s(w_{ij} - w_{ij}^s)^2}{nw_{ij}} = \sum_{s=1}^{S}\sum_{j=1}^{r(i)} \frac{w^s(w_{ij} - w_{ij}^s)^2}{w_{ij}} \tag{14}$$

where $w^s = n^s/n$.

If we combine the ϕ_i^2 from the different questions, we obtain the Chi-squared index derived above, χ^2:[11]

$$\chi^2 = \frac{1}{q}\sum_{i=1}^{q} \phi_i^2 = \frac{1}{q}\sum_{i=1}^{q}\sum_{s=1}^{S}\sum_{j=1}^{r(i)} \frac{w^s(w_{ij} - w_{ij}^s)^2}{w_{ij}} \tag{15}$$

Thus, the χ^2 index depends on the average difference between the observed shares w_{ij}^s and the expected shares w_{ij} that we would observe if the distribution of ethnicity and the distribution of culture were independent. This index has a minimum value of zero when there is no overlap. The maximum value depends on the number of ethnic groups, S, and the number of answers of each question, $r(i)$.[12]

3.3.3 An Alternative: The Fixation Index or F_{ST}

An alternative to the χ^2 index is F_{ST}, an index commonly used in population genetics to measure genetic differentiation or distance between groups (see Wright, 1949, and Nei, 1973). In genetics, F_{ST} is a measure of relative heterogeneity: it is the ratio of between-group heterogeneity in genetic characteristics to total heterogeneity. Analogously, here we compute a cultural F_{ST} - the ratio of

[10] In statistical parlance, this ϕ index is known as a weighted mean squared contingency index.

[11] Strictly speaking, the index we use is ϕ^2 not χ^2. However, for simplicity we will use the term χ^2 in referring to the index based on population shares rather than number of individuals.

[12] An alternative to this index is Cramér's V, which is defined as $V = \sqrt{\frac{\chi^2}{n\ t}}$, where t is the smaller of the number of ethnic groups minus one or the number of answers minus one (Cramér, 1946). The index V is always between zero and one, and corrects for the different number of ethnic groups or answers. However, this index is hard to interpret as it does not derive from our model. Another alternative to χ^2 is the Mutual Information Index, originating from the concept of entropy in information theory. In our case this is a measure of the amount of information that ethnicity contains about values, i.e., the reduction in the uncertainty about how an individual answers the questions due to the knowledge of his ethnicity. It can be shown that χ^2 is up to an order of approximation equal to the Mutual Information Index (Cover and Thomas, 2006, p. 400).

between-group cultural heterogeneity to total heterogeneity: when F_{ST} is 0, ethnic identity conveys no information about cultural attitudes, norms and values. In contrast, if F_{ST} is equal to 1, knowing someone's ethnolinguistic identity allows a perfect prediction of their cultural attributes. F_{ST} is therefore a measure of overlap between cultural values and ethnolinguistic identity. An added advantage of F_{ST} is that it relates neatly to the already described measures of cultural diversity in terms of functional form - namely, it isolates the part of the variation in overall cultural diversity that occurs between groups.

To define F_{ST}, we start from the probability that two randomly drawn individuals from ethnic group s give a different answer to question i (the within-group cultural diversity of group s):

$$CF_i^s = 1 - \sum_{j=1}^{r(i)} \left(w_{ij}^s\right)^2 \tag{16}$$

The population-weighted average of the within-group cultural fractionalization for question i can be written as:

$$CF_i^W = \sum_{s=1}^{S} w^s CF_i^s = \sum_{s=1}^{S} w^s \left(1 - \sum_{j=1}^{r(i)} \left(w_{ij}^s\right)^2\right) \tag{17}$$

The share of the total population's cultural fractionalization that is not due to within-group fractionalization for each question i is then:

$$F_{ST_i} = \frac{CF_i - CF_i^W}{CF_i} \tag{18}$$

This is, for each question, the ratio of between-group cultural fractionalization divided by total fractionalization. Averaging over all questions gives us Wright's fixation index, F_{ST} (Wright, 1949, Nei, 1973):[13]

$$F_{ST} = \frac{1}{q} \sum_{i=1}^{q} F_{ST_i} \tag{19}$$

The F_{ST} index measures the share of between group heterogeneity in total cultural heterogeneity. If all ethnic groups were as heterogeneous as the total population, F_{ST} would be equal to 0, and there would be no between-group heterogeneity. In that case, cultural cleavages and ethnolinguistic cleavages cross-cut. Knowing someone's ethnolinguistic identity would give no information about his preferences or culture. Instead if all ethnic groups were to be homogeneous ($CF_i^S = 0$), F_{ST}

[13] There are of course many ways to average across questions. For instance, Cavalli-Sforza et al. (1994) separately average the numerator and the denominator of equation (18), and then take the ratio. We adopt the simpler method of averaging the question by question F_{ST}^i.

would be equal to 1, and all heterogeneity would be between groups. In that case, cultural cleavages and ethnolinguistic cleavages would be reinforcing.

The advantage of F_{ST} is that it is well-known and captures intuitively a simple concept, as it represents how much one can predict answers to questions on norms, attitudes and preferences simply by knowing a respondent's ethnolinguistic identity. In the case of two ethnic groups and one question with only two possible answers this index ranges from 0 to 1. With two groups and more than two possible answers, or more generally when the number of groups exceeds the number of answers, there is always some within-group fractionalization and the index cannot reach 1.

While the F_{ST} index is very commonly used in population genetics, it does have some drawbacks, as explained in, for example, Jost (2008), Meirmans and Hedrick (2011) and Jakobsson et al. (2013). The most important drawback, outlined in Jost (2008), relates to the properties of the decomposition of within and between fractionalization in a context where these measures are bounded above by 1. To illustrate this potential problem, let us denote between-group fractionalization by D_i for question i. Such between-group fractionalization is defined by subtracting within-group fractionalization CF_i^W from total fractionalization CF_i, i.e., $D_i \equiv CF_i - CF_i^W$. Thus, this approach relies upon the additive decomposition of total fractionalization, but CF_i^W and D_i are not independent because we always have that $D_i + CF_i^W \leq 1$.[14] This constraint implies that D_i declines with within-group fractionalization CF_i^W regardless of the degree of cultural differentiation of ethnic groups.

A numerical example is useful to illustrate this drawback. Suppose that there is just one question and two ethnic groups of the same size. The question has four possible answers, a, b, c and d. The vector of answers for individuals from the first ethnic group is $\{0.1, 0.9, 0, 0\}$, i.e., 10% of them answer a, and 90% answer b. For the second ethnic group the vector of answers is $\{0.9, 0.1, 0, 0\}$. It is easy to check that in this society, $F_{ST} = 0.64$. Suppose a second society where those two vectors of answers are $\{0.5, 0.5, 0, 0\}$ and $\{0, 0, 0.5, 0.5\}$. It is clear that in this society culture and ethnicity overlap more strongly than in the first society. However, in this case we have $F_{ST} = 0.33$.[15] The reason is that the second society displays a much higher degree of within-group heterogeneity than

[14] Jost (2008, pp 4018) provides a complete explanation of this constraint and its implication: "Additive partitioning of heterozygosity does not produce pure within-and between-subpopulation components; it is an incomplete partitioning". In our case cutural fractionalization is a parallel concept to heterozygosity in population genetics.

[15] A similar numerical example appears in Wright (1978).

the first (a high CF^W), which drives down F_{ST} in spite of the higher degree of overlap between culture and ethnicity.

Our first overlap measure, χ^2, is not subject to this drawback, but as we will see empirically it does not matter which index we use: while the χ^2 index comes out directly from our model of antagonism, empirically χ^2 and F_{ST} are almost perfectly correlated (in our sample the correlation is 97%, a result that also holds in simple simulations).[16]

3.3.4 Conceptual Discussion of the Overlap Measures

Relationship with existing measures of cross-cuttingness. Our proposed indices of overlap capture how much someone's ethnolinguistic identity reveals about his culture or preferences. High values imply that ethnolinguistic and cultural cleavages are reinforcing, whereas low values imply that they are cross-cutting. Our indices are thus related to an existing literature in political science concerned with the measurement of cross-cutting cleavages, starting with Rae and Taylor (1970, chapter 4). In what follows we start by discussing the Rae and Taylor measure of cross-cuttingness, and we then discuss the similarities and differences with our indices of overlap.

Consider two cleavages. In our terminology the first cleavage could refer to ethnicity and the second to culture (defined on a single dimension for now - say on a generic question i from the WVS). Assume there are S ethnic groups and $r(i)$ cultural groups. Fractionalization on cleavage 1 is ELF and fractionalization on cleavage 2 is CF_i, as previously defined. In Rae and Taylor's definition, if all those from a given ethnic group are also in a given cultural group, cleavages are perfectly reinforcing. They define cross-cutting XC as the "proportion of all pairs of individuals whose two members are in the same group of one cleavage but in a different group of the other cleavage" (p. 92), and show that, for a large enough population, we can write:

$$XC_i = \sum_{j=1}^{r(i)} w_{ij}^2 + \sum_{s=1}^{S} (w^s)^2 - 2\sum_{j=1}^{r(i)}\sum_{s=1}^{S} \left(w^s w_{ij}^s\right)^2 = 2F_i^c - CF_i - ELF, \tag{20}$$

where

$$F_i^c = 1 - \sum_{j=1}^{r(i)}\sum_{s=1}^{S} \left(w^s w_{ij}^s\right)^2. \tag{21}$$

Heuristically, F_i^c is fractionalization computed over all groups defined by both ethnicity and culture (so, for example, if $S = 3$ and $r(i) = 4$, there are 12 distinct groups defined by heterogeneity in

[16] In the case of a question with two possible answers F_{ST} and χ^2 coincide exactly (see Workman and Niswander, 1970). Details of our simulations are available upon request.

both ethnicity and the answer to question i): F_i^c measures the probability that two randomly chosen individuals answer question i in a different way *or* belong to a different ethnic group. It is a measure of fractionalization where belonging to a different ethnic group or to a different cultural group defines different groups of individuals symmetrically, with the same weight.[17] With the definition above, it can be easily seen that XC_i is indeed the probability that two randomly chosen individuals in the population belong to the same group on one cleavage but to a different group on the other cleavage. XC_i can be averaged over questions i to obtain an overall index XC.

Intuitively, being a measure of cross-cutting cleavages, XC should be negatively correlated with the χ^2 index (as well as F_{ST}) which are measures of reinforcing cleavages. Both types of indices are also quite different conceptually. This was already noted by Rae and Taylor (1970) and further discussed in Selway (2010, 2011). To put the distinction in stark focus with a simple example, consider the following distribution of individuals over two ethnicities and three possible answers to a cultural question:

(Entries are # of people)	Answer 1	Answer 2	Answer 3
Ethnic group 1	2	2	2
Ethnic group 2	1	1	1

Here both χ^2 and F_{ST} are obviously zero: both groups have the same distribution of answers as the population overall, and the share of between-group variance in total variance is zero (i.e. there is no between-group variance). However, F_i^C is 0.815, ELF is 0.444 and CF_i is 0.667, so that XC_i equals 0.518. As expected, a low value of χ^2 (or F_{ST}) corresponds to a high value of XC. When we double the size of ethnic group 1 proportionally for each possible answer, the values of χ^2 and F_{ST} are unchanged, as expected: they continue to be zero, since doubling the size of ethnic group 1 does not affect how informative ethnicity is about culture. In contrast, XC_i increases from 0.518 to 0.560, because the probability of two individuals sharing one cleavage but not the other increases. As this example illustrates, XC is sensitive to changes in group sizes that are not associated with changes in the degree to which ethnicity is informative about a person's cultural attitudes, the concept we have sought to capture in this paper so far. In fact, the example shows that XC has different properties from χ^2, an index that comes out directly from a simple model of social antagonism. Nevertheless, for completeness, in the empirical section we also compute the

[17] In contrast ELF gives no weight to cultural heterogeneity while CF_i gives no weight to ethnic heterogeneity.

XC index, and while it will not be the focus of our analysis, we will briefly discuss its empirical properties.

Uses of F_{ST} in the past literature. We conclude this conceptual section by discussing a few papers that have used F_{ST} to measure between-group cultural heterogeneity, noting that their goals and methods are very different from ours. Bell et al. (2009) study inter-group competition and analyze whether there is more scope for selection based on cultural traits rather than on genetic traits. They use the WVS to compute a cultural F_{ST} measure between 150 pairs of neighboring countries. They show that this measure is an order of magnitude larger than an analogous measure of F_{ST} based on genetic data, suggesting a greater scope for cultural rather than genetic selection. In contrast to our approach, they measure cultural heterogeneity between countries rather than between groups within countries. In another paper, Ross et al. (2013) compute a measure of cultural F_{ST} based on between-group variation in folktales across different European ethnic groups. Again they are interested in comparing patterns of cultural F_{ST} to those of genetic F_{ST}, finding some similarities between the two. In contrast to our work, neither of these papers is interested in using cultural F_{ST} to measure the degree of overlap between ethnicity and culture, or in understanding how cultural F_{ST} relates to overall cultural heterogeneity and overall ethnolinguistic diversity. Instead, their focus is on the importance of cultural F_{ST}, relative to genetic F_{ST}. Finally, Spolaore and Wacziarg (2009, 2013) use a genetic F_{ST} as a measure of intergenerational divergence in a wide range of human traits transmitted culturally or biologically, in order to estimate the effects of barriers between populations on political and economic outcomes. In contrast to our approach, their F_{ST} is based on neutral genes, not cultural attitudes, and it measure distance between countries rather than between ethnic groups, within countries.

4 Ethnic Heterogeneity and Cultural Diversity

In this section we calculate and describe the measures of heterogeneity derived in Section 3. We show that, contrary to the assumption of much of the past literature, measures of ethnic diversity and cultural diversity are uncorrelated with each other. At the same time, we know from Section 2 that ethnic identity does help to predict cultural attitudes. To reconcile these seemingly contradictory results, we show that between-group heterogeneity in cultural attitudes is small compared to total heterogeneity. Measures of the overlap between culture and ethnicity show considerable variation across countries, variation that mirrors that found in Section 2. We explore the correlates of these

new measures, uncovering interesting patterns concerning characteristics of countries with a high degree of cultural diversity as well as those with a high degree of overlap between culture and ethnicity.

4.1 Cultural Diversity and Ethnolinguistic Fractionalization

Before describing the indices, some comments on the data are in order. First, we use the same baseline set of questions from the integrated WVS/EVS dataset as in Section 2. Second, not all WVS/EVS questions are asked in all countries. In our benchmark analysis we drop questions that are not asked in at least 50 countries, to ensure cross-country comparability of the indices.[18] Third, since we are interested in relating cultural fractionalization to ethnolinguistic fractionalization, we focus exclusively on countries in the WVS/EVS for which we have ethnolinguistic information.[19] Taken together, this gives us information on 76 countries. Finally, there is no need here to convert questions that admit multiple unordered answers into series of binary questions. For these questions we compute diversity metrics directly as per equation (4).

Figure 1 shows a world map with the values of cultural heterogeneity in the 76 countries in our sample, and Table 3 displays some simple summary statistics.[20] Darker-colored countries are more culturally diverse than lighter-colored ones. The most culturally diverse country is Zambia ($CF = 0.602$), and the least culturally diverse country is Jordan ($CF = 0.427$). Other interesting data points are France and India, with relatively high cultural heterogeneity, and Egypt, Indonesia and China, with relatively low numbers (a high degree of cultural conformism). The average value of CF across countries is 0.529.

The standard assumption in the literature is that cultural heterogeneity (CF) should be highly correlated with ethnolinguistic heterogeneity (ELF). Comparing the map of ELF in Figure 2 with the one of CF in Figure 1, it becomes immediately obvious that there are important differences.

[18] In further robustness analysis, we use different thresholds for when to include a question or not, and we also recompute our indices for different types of questions (e.g., focusing exclusively on binary questions).

[19] In principle we could of course use data on ethnolinguistic fractionalization from sources other than the WVS. We refrain from doing here so for two reasons. First, using the same data source makes the two indices, cultural and ethnolinguistic fractionalization, more easily comparable. Second, we later analyze the overlap between culture and ethnicity. For that we need to use ethnolinguistic identity and cultural values at the individual level, from the WVS/EVS.

[20] See Table A4 in the Appendix for the underlying values country by country.

Countries such as Pakistan and Egypt have high levels of ethnolinguistic heterogeneity but low levels of cultural heterogeneity. At the other extreme are countries such as Germany and South Korea, which are ethnolinguistically fairly homogeneous but culturally diverse. The lack of a relationship between both types of heterogeneity is not limited to these few examples. The correlation between CF and ELF is essentially zero, -0.030 to be exact. Ethnolinguistic diversity is therefore not associated with culturally diversity. This result is unexpected in light of our previous finding in Section 2 that an individual's ethnolinguistic identity helps to predict answers in about half of the questions in the WVS/EVS. We return to this important point below.

If cultural heterogeneity is not correlated with ethnolinguistic heterogeneity, what might it be related to? Figure 3 shows the correlation between cultural fractionalization and several variables. Four correlates stand out: countries with a higher proportion of Muslims exhibit a lower CF (the correlation is -0.597); partly reflecting the previous correlation, countries located in North Africa and the Middle East show a similar negative correlation (-0.529); more democratic countries (measured by the Polity 4 index) have higher cultural fractionalization (the correlation is 0.586); and richer countries exhibit a greater CF (the correlation is 0.382). The other correlations are mostly small and statistically insignificant. Countries in North America, Europe and Central Asia, and countries at higher latitudes, have, on average, greater cultural fractionalization. As for other measures of diversity, although the correlation with ELF computed from WVS/EVS data is zero, there is a small, positive and statistically significant correlation between CF and religious fractionalization computed from the data in Alesina et al. (2003).

4.2 The Overlap Measures

Figure 4 shows a map of the χ^2 index for all countries in our database, and Table 3 reports summary statistics. Several observations are in order. First, χ^2 takes on low average values, indicating that the relative difference between the within-group heterogeneity and the overall heterogeneity is small (the mean value of χ^2 in our sample of 76 countries is 0.029). However, there is substantial variation in χ^2, with Asia (especially South Asia and Southeast Asia) and Sub-Saharan Africa displaying high values, while Europe, Russia and Latin America display relatively low values. Notable data points with high χ^2 values include India, Thailand and Zambia. Countries with low values include Japan, Russia, Poland and Italy. These patterns closely mirror those uncovered in Section 2. The regions where ethnicity could significantly predict responses to survey questions about values,

norms and preference are the same regions where χ^2 takes on higher values.[21] These patterns help to explain why cultural diversity and ethnic diversity are uncorrelated, even though ethnic identity helps predict a large share of answers to questions on cultural attitudes: most of the heterogeneity is within groups.

Second, the ranking of countries is very similar across both the χ^2 and F_{ST} measures (Table 3, Panel B). The correlation between the two is 0.981 (and so is the Spearman rank correlation).[22] This gives us confidence, despite very different functional forms, that these two measures capture common features of the data regarding the overlap between ethnicity and culture. In particular, the drawback of F_{ST} identified above does not seem very relevant empirically, since χ^2 and F_{ST} are very highly correlated, and χ^2 is not subject to the drawback. As with χ^2, the mean value of F_{ST} is low: the share of between-group variance in cultural attitudes relative to the overall variation is 0.012. A similar result is well-known in population genetics, where within-group variation in genetic characteristics swamps between-group variation (Cavalli-Sforza et al., 1994).

Third, reinforcing cleavages have a weak, positive correlation with cultural fractionalization. This can be seen by comparing Figure 4 with Figure 1, or by noting, from Table 3 Panel B, that the correlation between χ^2 and cultural fractionalization is 0.219 (this correlation is statistically significant at the 10% level). The corresponding number for F_{ST} is 0.179 (this correlation is not statistically significant at the 10% level). An example of this positive correlation is Malaysia, a country that is culturally heterogeneous ($CF = 0.563$) and where knowing someone's identity is relatively informative about that individual's preferences ($\chi^2 = 0.092$). But other examples show the lack of a strong correlation. Germany is culturally heterogeneous ($CF = 0.576$), but has a low χ^2 ($\chi^2 = 0.009$).

[21] In fact, if you consider, for each country in our sample, the share of WVS/EVS questions for which ethnicity dummies are jointly significant predictors of individual responses, and correlate this share across countries with our measures of χ^2 and F_{ST}, you obtain correlations of 0.73 and 0.77, respectively. These high correlations exist despite the vastly different methodologies used in Section 2 and Section 4 to capture the degree of overlap between culture and ethnicity.

[22] In addition to empirically analyzing χ^2 and F_{ST}, we also calculated the XC index of Rae and Taylor (1970). Since χ^2 and F_{ST} are measures of reinforcing cleavages, whereas XC is a measure of cross-cuttingness, the correlations between both types of indices is expected to be negative. This is indeed the case: the correlations of XC with χ^2 and F_{ST} are, respectively, -0.43 and -0.45. Since we have only micro-founded the χ^2 index in our theory, in the rest of our paper we focus mainly on χ^2, although we also discuss F_{ST} since it is highly correlated with χ^2 and it is a well-known index from population genetics.

Fourth, as expected, reinforcing cleavages are stronger in countries that are more ethnically diverse. The correlation between χ^2 and ethnolinguistic fractionalization is 0.620 (statistically significant at 1%); and the corresponding figure for F_{ST} is the same. This positive correlation can also be perceived by comparing Figure 4 with Figure 2. Ethnolinguistically diverse countries such as India, Philippines or Ethiopia also tend to have high levels of χ^2 or F_{ST}.

Finally, we investigate the correlates of our overlap measures. Figure 5 displays the quantitative magnitudes of the simple correlations between χ^2 and a set of country characteristics (the correlations with F_{ST} are very similar, and available upon request). We find interesting descriptive patterns. Consistent with our previous findings, χ^2 is higher in South Asia, East Asia and Sub-Saharan Africa. It is lower in Latin America and Europe. χ^2 is also higher in countries with English legal origins, partly reflecting these spatial patterns. Per capita income is negatively associated with χ^2, as is latitude and to a lesser extent democracy, indicating that poorer countries tend to display more overlap between culture and ethnicity. An interesting hypothesis, which we do not pursue in depth but which is consistent with these correlations, is that modernity severs the link between ethnicity and cultural values. In contrast, recall that modernity was positively correlated with overall cultural diversity (CF). Among variables describing the prevalence of various religions, only the percentage of Catholics is significantly (and negatively) correlated with χ^2, although it is hard to disentangle this correlation from the fact that χ^2 is lower in Latin America than elsewhere.

5 Culture, Ethnicity and Civil Conflict

In this subsection we examine the relationship between our various measures of diversity and civil conflict. There remains a debate on whether ethnolinguistic diversity in fact affects civil conflict. Results seem to depend on which measures of diversity are used: using fractionalization, Fearon and Laitin (2003) find little evidence of an effect on conflict onset. Collier and Hoeffler (2004), using a measure of social fractionalization that combines ethnic and religious dimensions, find evidence that greater diversity reduces the probability of a civil war. In contrast, using a measure of ethnic polarization, Montalvo and Reynal-Querol (2005) find that it has a significantly positive effect on civil war incidence. Esteban, Mayoral and Ray (2012), using different measures of diversity jointly, found that they were significantly associated with civil conflict. Of particular note, a recent paper by Huber and Mayoral (2013) examines the role of income inequality between and within ethnic groups as a determinant of civil conflict, finding that within-group inequality affects

conflict positively. Our paper shares with theirs a decomposition of overall differences between and within ethnic groups, and also uses survey data. However, the application is very different as their paper is interested in within- and between-ethnic group income inequality, as opposed to cultural differences.[23]

How might ethnolinguistic diversity affect civil conflict in our framework? First, ethnolinguistically heterogeneous societies may have more diverse preferences and values, leading to increased overall antagonism in society. Second, ethnolinguistic fractionalization may matter *per se* because of direct animosity, hatred or barriers between different ethnolinguistic groups. Third, civil conflict may arise more frequently when ethnic divisions and cultural differences reinforce each other. To evaluate which of these channels operates most strongly, we introduce measures of these three dimensions of heterogeneity simultaneously in regressions explaining civil conflict. If the first explanation is valid, then cultural fractionalization (CF) should affect civil conflict. If the second explanation is valid, then ethnolinguistic fractionalization (ELF) should matter. Finally, if the third explanation is valid, then the χ^2 (or F_{ST}) should matter.

5.1 Data and Specification

In our application, the aforementioned contributions to the study of civil conflict constitute the methodological starting point. Following the literature, we define a dummy variable C_{ct} equal to 1 if country c experiences a civil war in year t (to explore conflict incidence). We also define a separate dummy C^o_{ct} equal to 1 if in a given year a country experiences the onset of a new civil war (to study onset only). We relate these outcomes to our three sets of measures of diversity, CF, ELF and χ^2 (or F_{ST}):

$$C_{ct} = \beta_0 + \beta_1 CF_c + \beta_2 ELF_c + \beta_3 \chi^2_c + \boldsymbol{\beta}_4' \mathbf{Z}_{ct} + \varepsilon_{ct} \tag{22}$$

where $\mathbf{Z}_{ct}$ is a vector of control variables commonly used in the literature. In particular, we use an expansive set of controls very close to the ones used in Fearon and Laitin (2003), Esteban, Mayoral and Ray (2012) and Desmet, Ortuño-Ortín and Wacziarg (2012). These include a variety of geographic variables, lagged per capita GDP, as well as lagged conflict, legal origins, and dummy variables for major geographic regions.

The data on civil conflict and the control variables come from Fearon and Laitin (2003). In this database, a country is coded as being in a civil conflict when the conflict overall killed over $1,000$

[23] The functional forms are also different as they use a decomposition of the Gini index, not fractionalization.

people, with an average of at least 100 deaths a year and at least 100 deaths on both sides of the conflict. As an alternative source of data we use a database from the Peace Research Institute Oslo (PRIO), which has been used in recent contributions (for instance Esteban, Mayoral and Ray, 2012, and Huber and Mayoral, 2013). There, a civil war is defined as "a contested incompatibility that concerns government and/or territory where the use of armed force between two parties, of which at least one is the government of a state, results in at least 25 battle-related deaths." We also use a more stringent threshold of $1,000$ battle related deaths. When looking at conflict incidence, in our sample of countries, the correlation between PRIO25 and the Fearon-Laitin variable is 0.69, while the correlation between PRIO1000 and the Fearon-Laitin variable is 0.47 (the correlation between PRIO25 and PRIO1000 is 0.51).

5.2 Results

The results are reported in Tables 4 through 7. In Table 4 we examine the determinants of civil conflict incidence, introducing cultural fractionalization, ethnolinguistic fractionalization and χ^2, first individually and then jointly (columns 1 through 4). We find that, when introduced individually, χ^2 is a significant predictor of conflict incidence (at the 1% level), and this continues to be the case when all three measures are introduced jointly. In fact in that case the logit marginal effect of χ^2 becomes twice as large (column 4): it is equal to 0.870. That implies that a one standard deviation change in χ^2 (equal to 0.025) raises the probability of conflict by 2.175 percentage points. The baseline probability of being in a civil conflict is 14.173% in the sample for this regression, so the standardized effect of χ^2 amounts to about 15% of the probability of conflict. Our interpretation of this result is that ethnic divisions matter for civil conflict, but only when they overlap with cultural cleavages.

Another notable result from Table 4 is that cultural fractionalization (CF) tends to reduce the incidence of civil wars (even when controlling for several variables previously found to be correlates of CF, such as democracy, per capita income and region dummies). In column 4, the marginal effect of CF is -0.391. The standard deviation of CF is 0.037, which means that a one standard deviation increase in CF is associated with an 10% reduction in the probability of conflict. One interpretation of this finding is that cultural diversity is the sign of a society that is tolerant of a multiplicity of values and preferences, and this tolerance reduces the incidence of civil conflict. A related interpretation is that cultural diversity is the sign of a society that embraces modernity more generally, and modernity is not fully captured by the included controls that correlate with

CF. Ethnolinguistic fractionalization is insignificant, and has an unstable sign across specifications, although it tends to bear a negative sign when all measures of heterogeneity are entered together (columns 4-6). In columns (5) and (6) of Table 4 we add legal origins and GDP growth and its lag to the baseline specification, with little effect on the estimated coefficient on χ^2 and CF. None of these results change much when using probit rather than logit.

Table 5 looks at the determinants of civil war onset. The results are very much in line with those for incidence. Of course, conflict onset is a much rarer event than conflict incidence, with the percentage of country-year observations featuring the onset of a civil war being equal to 1.780%. Given this fact, the marginal effect of χ^2 on conflict onset is found to be 0.175 (column 4), implying that a one standard deviation increase in χ^2 reduces the probability of civil war onset by 25%, a sizeable effect. On the other hand, while it is still negative, the effect of cultural diversity on conflict onset is generally not significant at the 5% level. Finally, the effect of ethnolinguistic fractionalization continues to be insignificant, in line with results on civil war onset in Fearon and Laitin (2003).

5.3 Endogeneity of CF and χ^2

As in most of the literature on civil conflict, we have so far treated our heterogeneity measures as exogenous to conflict. As long as we limit attention to ethnolinguistic fractionalization, and include a suitably expansive set of controls, this can be justified as ethnolinguistic fractionalization is very time-persistent and is unlikely to change very much as a result of conflict. The same cannot necessarily be said of cultural heterogeneity and the overlap of culture and ethnicity. The experience of civil wars can lead people to change their values and preferences, and respond differently to questions from the World Values Survey. This in turn can lead measures of cultural diversity such as CF to change as a result of conflict (though it is not clear *a priori* in what direction, and hence the sign of the endogeneity bias is not clear a priori). Similarly, civil war, especially when there is an ethnic component, can change the salience of ethnic identity, leading ethnic groups to adopt values and attitudes that differ more than they did prior to the war, or on the contrary can lead the victor to impose their values and preferences on the vanquished, therefore affecting the degree of overlap between ethnicity and culture (again in an ambiguous direction).[24] This could be a problem particularly in our setting as the questions from the World Values Survey used to

[24] This is a prevalent theme in constructivist approaches to ethnicity.

characterize preferences and values are from survey waves from 1981 to 2008, while our main civil war dataset (from Fearon and Laitin, 2003) covers 1945 to 1999. We already partly address this problem in the regressions shown so far as we adopt a dynamic specification for the incidence of civil wars, i.e., we include a term for lagged civil war on the right-hand side of the specification, in keeping with the usual practice in the literature (see in particular Fearon and Laitin, 2003, p. 84 and Esteban, Mayoral and Ray, 2012, p. 1318). Since civil war incidence is highly autocorrelated, this purges CF and χ^2 of much of their variation attributable to past wars.

To deal with any remaining endogeneity, we adopt a three-pronged approach. First, we focus on questions that display a high degree of persistence in cultural fractionalization across waves, and are thus less likely than other questions to respond endogenously to external events such as civil wars. For each question i, country c and WVS wave w, consider cultural fractionalization CF_{icw}. For each question i and each country c we compute the coefficient of variation of CF_{icw} across waves w, and average this coefficient of variation across countries for each question i. This gives a measure of persistence for each question i. We then remove from consideration every question with a coefficient of variation in excess of 0.1, which leaves us with about 60% of the questions used previously - the ones with the most time persistent value of CF_{icw}. We reran our baseline specification (the one in column 4 of Table 4) with CF and χ^2 computed from this restricted set of questions. The results are presented in column (1) of Table 6. The signs of our main effects remain the same, namely CF affects conflict negatively (albeit the effect is no longer statistically significant at the 5% level) and χ^2 affects conflict incidence positively and remains significant at the 5% level. The magnitude of the effect falls slightly in standardized terms, with a one standard deviation increase in χ^2 associated with a 10% increase in the probability of conflict.

Our second approach is to limit our sample to the post-1970 period. The idea is that if endogeneity were a strong concern, we should find different results in this subsample compared to the full sample. The argument could take various forms. On the one hand, if we limit attention to wars that occur closer to the date when we observe values, there is perhaps greater potential for recent wars to affect values, and then in turn cultural diversity and χ^2. On the other hand, if the lag with which war may affect values is substantial, by focusing on a recent sample, values may not yet have had time to change, and therefore CF and χ^2 may not yet have changed in response to civil war. In either case, if reverse causality were a concern we would observe *different* estimates of the effect of CF and χ^2 in the post-1970 sample and in the whole sample. Column (2) of Table

6 presents the results, which are almost unchanged compared to the baseline regression of Column 4 of Table 4. Indeed, the standardized effect of χ^2 stands equal to the one previously calculated, at 15%. The standardized effect of CF on the probability of civil war also remains equal to -10%. These results have the added advantage of showing the stability and robustness of our estimated effects to the sample period under consideration.

Out third approach is a variation on the previous one. Here, we limit attention to respondents born before 1950 *and* to the post-1970 sample. The effects could once again go in a variety of directions, but the argument is again that the estimates would be *different* if endogeneity were a big concern. On the one hand, if respondents' values are formed in early adulthood and change little after that, since every respondent in the sample would be at least 20 years old in the event of a civil war, their cultural values may respond less than younger individuals to the event of a war. On the other hand, if one's view was that these individuals were the most likely to be affected by a civil war because they were most likely to be combatants or to be affected by the war in adult age, their values may be most likely to be affected by the war. Either way, war would affect values, and therefore potentially CF and χ^2 also (although, again, in unknown directions). Column (3) of Table 6 presents the results, but once again we find very little evidence of a different effect of cultural diversity and χ^2 on the probability of war: the standardized effects of these two variables are, respectively, -9% and 13%, very close to those found in column (4) of Table 4.

While we do not want to place too much weight on any one of these tests in isolation, taken together they do suggest that our main results are remarkably stable when looking only at the post-1970 sample, when looking only at respondents born before 1950, and when including only questions for which question-by-question cultural diversity CF_i is most stable across survey waves. Under reverse causality, if war had a strong causal effect on CF or χ^2, we would have expected estimates under each of these modifications to differ from the baseline. They do not differ materially.

5.4 Robustness Checks

In Table 7, we carry out a series of additional robustness tests. First, we estimate our baseline regression using F_{ST} rather than χ^2 as the measure of overlap between culture and ethnicity (column 1). The results do not change in any substantive way, as expected because F_{ST} is so highly correlated

with χ^2.[25] The standardized effect of F_{ST} is 16% while the effect of CF is -10%. Second, we change our data source for civil wars to data from (PRIO). With the 25 battle deaths threshold (column 2), we find a standardized effects of χ^2 and CF that are much larger than in our baseline (respectively 27% and -20%). With the $1,000$ battle deaths threshold the corresponding effects stand at 14% and -10%, respectively, very similar to the baseline (column 3). We conclude that the signs and significance of the estimated effects are not sensitive to changes in the definition of a civil war.

In a final robustness check, we expand the number of questions included to calculate our measures of cultural diversity and overlap to those that were asked in at least 30 countries, rather than the more stringent criterion of 50 countries used previously. This expands the set of questions used in our calculations, at the cost of greater heterogeneity across countries in the set of questions. The results appear in column (4) of Table 7. Reassuringly, nothing changes much: the standardized effects of χ^2 and CF are largely unchanged at 17% and -9%, respectively.[26] This greatly increases our confidence that our results do not depend on the set of questions used to calculate cultural diversity and overlap.

6 Conclusion

In this paper we studied the complex relationship between ethnicity and culture, defined as a vector of answers to a broad set of questions about norms, values and preferences. We uncovered novel results. First, ethnicity does serve to significantly predict cultural attitudes, to an extent that varies across geographic regions. Second, the share of variation in culture that is explained by ethnicity is relatively small. As a result, cultural diversity, defined as the average probability that two randomly chosen individuals respond differently to a question from the World Values Survey, is not correlated with ethnic diversity. Thus, ethnic fractionalization cannot readily be taken as a proxy for diversity in values, attitudes and preferences. Third, we derived and calculated several new indices measuring the extent of overlap between culture and ethnicity, stemming from a simple

[25] When using the Rae and Taylor (1970) measure of cross-cuttingness, XC, instead of χ^2 or F_{ST}, we find that XC has a negative effect on the probability of civil war, as expected, but is not robustly significant across specifications corresponding to those in Table 4. However, as already noted, XC has very different properties from our proposed indices and does not come out of our model of antagonism.

[26] Reassuringly also, the correlation in our sample of 76 countries between CF using the 50-countries threshold and CF using the 30-countries threshold is very high - at 0.95. Similarly, the two versions of χ^2 bear a 0.99 correlation with each other.

model of social antagonism. These measures display interesting geographic variation, with the degree of overlap being greatest in Sub-Saharan Africa and Asia, and smallest in Latin America. Fourth, as an application we used our new measures of cultural diversity and overlap to study the determinants of civil conflict, finding that ethnic fractionalization has no predictive power for civil conflict, but that cultural diversity has, if anything, a pacifying effect. Our new measures of overlap between culture and ethnicity, χ^2 and F_{ST}, have a positive effect on the probability of civil conflict onset and incidence, indicating that ethnic divisions matter for conflict when they are associated with cultural differences across ethnic groups. Hence, we have identified the degree of overlap between culture and ethnicity as an new and robust determinant of civil conflict.

Our results parallel a famous debate in population genetics on within group versus between group genetic differentiation, going back to Lewontin (1972). Lewontin pointed out that between-race genetic variation was a very small part of overall variation, and that within-group diversity accounted for a much larger share of overall genetic variation. This led Lewontin to question the validity of the very concept of race. In a series of rejoinders, Edwards (2003), Dawkins (2005) and others argued that while between-group variation was small, it could still be a relevant part of the variation: humans share up to 99% of their DNA with animals, yet the 1% that differs matters a lot to set the two groups apart. Lewontin's point on genetics mirrors our finding that between-ethnic group cultural variation is a small part of overall cultural variation, and that most of this variation occurs within-groups. Edward's and Dawkins' argument also finds an echo in our work, since we argue that between-group variation, while a small share of the overall variation, matters greatly for civil conflict.

The question we posed here is also related to a continuing debate in the social sciences as to whether ethnic, linguistic and religious identities are "constructed" or reflect "primordial" differences between different groups of humans. Each of these traditions reflect a variety of viewpoints on the persistence of ethnic and cultural identities and a wide range of theories on the factors that gave rise to both ethnic and cultural differentiation. However, drawing a stark distinction between these two broad categories of views helps bring into focus a fundamental difference separating them: the primordialist view holds that ethnolinguistic divisions reflect deep differences between humans, the result of historical separation which allowed for cultural drift over centuries and millennia, so that the resulting ethnic divisions are associated with stark and persistent differences in culture, norms,

values and preferences.[27] In contrast, constructivists view ethnic identities as the endogenous result of shifting patterns of power, some very recent, so that the association between ethnic identity and cultural differences, if there is one at all, would be context-dependent, malleable, and fleeting.

Our paper provides evidence consistent with a synthesis of both views: ethnicity is indeed associated with fundamental differences in values, attitudes and preferences, in line with a primordialist viewpoint. Moreover, to the extent that ethnic divisions matter for conflict, they only do so when they overlap with cultural cleavages, once again a result with primordialist connotations. However, there are many other sources of variation in culture, not associated with ethnic identity: the magnitude of our χ^2 and F_{ST} indices tends to be small, indicating that the extent to which ethnicity is informative for culture is limited, a result that is more in line with the constructivist view. Moreover, some regions like Latin America feature a weak degree of association between culture and ethnicity, while others like Sub-Saharan Africa and Asia feature stronger degrees of overlap. The degree to which ethnic classifications reflect deep differences in cultural attitudes varies across regions, so the extent to which ethnic identities are primordially given or constructed varies across locations. Future work should continue to study the complex relationship between ethnicity, culture and diversity, a subject that had so far remained missing from the economics literature on ethnic heterogeneity.

[27] See the voluminous and growing literature on ethnic heterogeneity in economics, among which the recent paper by Michalopoulos (2012) constitutes an excellent illustration of the primordialist orientation of the economics approach to ethnicity.

References

Alesina, A. and E. La Ferrara (2005), "Ethnic Diversity and Economic Performance," *Journal of Economic Literature*, vol. 43, pp. 762-800.

Alesina, A., R. Baqir and W. Easterly (1999), "Public Goods and Ethnic Divisions," *Quarterly Journal of Economics*, vol. 114, no. 4, pp. 1243-1284.

Alesina, A., R. Baqir and C. Hoxby (2004), "Political Jurisdictions in Heterogeneous Communities," *Journal of Political Economy*, vol. 112, no. 2, pp. 348-396.

Alesina, A., A. Devleeschauwer, W. Easterly, S. Kurlat and R. Wacziarg (2003), "Fractionalization," *Journal of Economic Growth*, vol. 8, no. 2, pp. 155-194.

Alesina, A., P. Giuliano and N. Nunn (2013), "On the Origins of Gender Roles: Women and the Plough," *Quarterly Journal of Economics*, vol. 128, no. 2, pp. 469-530.

Alesina, A., S. Michalopoulos and E. Papaioannou (2012), "Ethnic Inequality," *NBER Working Paper* #18512.

Arbatli, C. E., Q. Ashraf and O. Galor (2013), "The Nature of Civil Conflict," *working paper*, Brown University.

Ashraf, Q. and O. Galor (2013), "The Out of Africa Hypothesis, Human Genetic Diversity, and Comparative Economic Development," *American Economic Review*, vol. 103, no. 1, pp. 1-46.

Bell, A., P. J. Richerson and R. McElreath, (2009) "Culture Rather than Genes Provides Greater Scope for the Evolution of Large-Scale Human Prosociality," *Proceedings of the National Academy of Sciences*, vol. 106, no. 2, pp. 17671-17674.

Bossert, W., C. D'Ambrosio and E. La Ferrara (2011), "A Generalized Index of Fractionalization," *Economica*, vol. 78, no. 312, pp. 723-750.

Caselli, F. and W. J. Coleman (2013), "On the Theory of Ethnic Conflict," *Journal of the European Economic Association*, vol. 11, no. S1, pp. 161-192.

Cavalli-Sforza, L. L., P. Menozzi and A. Piazza (1994), *The History and Geography of Human Genes*, Princeton: Princeton University Press.

Central Intelligence Agency (2009), *The World Factbook 2009*, Washington, DC: Central Intelligence Agency, 2009, https://www.cia.gov/library/publications/the-world-factbook/index.html

Collier, P. and A. Hoeffler (2004), "Greed and Grievance in Civil War," *Oxford Economic Papers*, vol. 56, no. 4, pp. 563-595.

Cover, T. M. and J. A. Thomas (2006), *Elements of Information Theory*, Hoboken, New Jersey: John Wiley & Sons

Cramér, H. (1946), *Mathematical Methods of Statistics*. Princeton: Princeton University Press.

Dawkins, R. (2005), *The Ancestor's Tale: A Pilgrimage to the Dawn of Evolution*, New York: Houghton Mifflin Harcourt.

Desmet, K., I. Ortuño-Ortín and R. Wacziarg (2012), "The Political Economy of Linguistic Diversity," *Journal of Development Economics*, vol. 97, no. 2, pp. 322-338.

Easterly, W. and R. Levine (1997), "Africa's Growth Tragedy: Policies and Ethnic Divisions," *Quarterly Journal of Economics*, vol. 112, no. 4, November, pp. 1203-1250.

Edwards, W. A. F. (2003), "Human Genetic Diversity: Lewontin's Fallacy," *BioEssays*, vol. 25, no. 8, pp. 798-801.

Esteban, J. M., and D. Ray (1994), "On the Measurement of Polarization," *Econometrica*, vol. 62, no. 4, pp. 819-851.

Esteban, J. M., and D. Ray (2004), "Polarization: Concepts, Measurement, Estimation," *Econometrica*, vol. 72, no. 6, pp. 1737-1772.

Esteban, J. and D. Ray (2011), "Linking Conflict to Inequality and Polarization," *American Economic Review*, vol. 101, no. 4, pp. 1345-74.

Esteban, J., L. Mayoral and D. Ray (2012), "Ethnicity and Conflict: An Empirical Study," *American Economic Review*, vol. 102, no. 4, pp. 1310-1342.

Fearon, J. and D. Laitin (2003), "Ethnicity, Insurgency, and Civil War," *American Political Science Review*, vol. 97, no. 1, February, pp. 75-90.

Fearon, J. (2003), "Ethnic and Cultural Diversity by Country", *Journal of Economic Growth*, vol. 8, no. 2, June, pp. 195-222.

Fernandez, R. and Alessandra F. (2009), "Culture: An Empirical Investigation of Beliefs, Work, and Fertility," *American Economic Journal: Macroeconomics, vol.* 1, no. 1, pp. 146–177.

Glaeser, E. L. (2005), "The Political Economy of Hatred," *Quarterly Journal of Economics*, vol. 120, no. 1, pp. 45-86.

Giuliano, P. (2007), "Living Arrangements in Western Europe: Does Cultural Origin Matter?" *Journal of the European Economic Association*, vol. 5, no. 5, pp. 927-952.

Guiso, L., P. Sapienza and L. Zingales (2006), "Does Culture Affect Economic Outcomes," *Journal of Economic Perspectives*, vol. 20, no. 2, pp. 23-48.

Guiso, L., P. Sapienza and L. Zingales (2009), "Cultural Biases in Economic Exchange?" *Quarterly Journal of Economics*, vol. 124, no. 3, pp. 1095-1131.

Huber, J. D. and L. Mayoral (2013), "Inequality, Ethnicity and Civil Conflict," *Barcelona GSE Working Paper*, no. 744.

Jakobsson, M., Edge, M. D. and N. A. Rosenberg, (2013), "The Relationship Between F_{ST} and the Frequency of the Most Frequent Allele," *Genetics*, vol. 193, pp. 512-528.

Jost, L. (2008), "GST and Its Relatives Do Not Measure Differentiation," *Molecular Ecology*, vol. 17, no. 18, pp. 4015-4026, September.

Lewontin, R. (1972), "The Apportionment of Human Diversity," *Evolutionary Biology*, vol. 5, pp. 381–398

Luttmer, E. F. P., and M. Singhal (2011), "Culture, Context, and the Taste for Redistribution," *American Economic Journal: Economic Policy*, vol. 3, no. 1 pp. 157-79.

Meirmans P. G. and Hedrick P. W. (2011) "Assessing Population Structure: FST and Related Measures," *Molecular Ecology Resources*, vol. 11, pp. 5-18.

Michalopoulos, S. (2012), "The Origins of Ethnolinguistic Diversity," *American Economic Review*, vol. 102, no. 4, pp. 1508–1539.

Montalvo, J. G. and M. Reynal-Querol (2005), "Ethnic Polarization, Potential Conflict and Civil War," *American Economic Review*, vol. 95, no. 3, June, pp. 796-816.

Nei, M (1973) "Analysis of Gene Diversity in Subdivided Populations," *Proceedings of the National Academic of Sciences, USA*, vol. 70, pp. 3321-3323.

Rae, D. W. and M. Taylor (1970), *The Analysis of Political Cleavages*, New Haven: Yale University Press,

Ross R. M., S. J. Greenhill and Q. D. Atkinson (2013), "Population Structure and Cultural Geography of a Folktale in Europe," *Proceedings of the Royal Society B,* vol. 280.

Selway, J. (2010) "Cross-Cuttingness, Cleavage Structures and Civil War Onset," *British Journal of Political Science,* vol. 41, pp. 111–138

Selway, J. (2011) "The Measurement of Cross-Cutting Cleavages and Other Multidimensional Cleavage Structures," *Political Analysis*, vol. 19, pp. 48-65.

Spolaore E. and R. Wacziarg (2009), "The Diffusion of Development," *Quarterly Journal of Economics*, vol. 124, no. 2, May, pp. 469-529.

Spolaore E. and R. Wacziarg (2013), "War and Relatedness," *working paper, UCLA and Tufts University.*

Tabellini, G. (2010), "Culture and Institutions: Economic Development in the Regions of Europe," *Journal of the European Economic Association*, vol. 8, no. 4, pp. 677-716.

Workman, P. L. and J. D. Niswander (1970), "Population Studies on Southwestern Indian Tribes. II. Local Genetic Differentiation in the Papago," *American Journal of Human Genetics*, vol. 22, no. 1, pp. 24–49.

Wright, S. (1949). "The Genetical Structure of Populations," *Annals of Human Genetics*, vol. 15, pp. 323–354.

Wright, S. (1978), *Evolution and the Genetics of Populations, Volume 4: Variability Within and Among Natural Populations.* Chicago: University of Chicago Press.

Appendix 1

Proof that $v = \chi^2$.

Proof. Recall that the χ^2 index is given by:

$$\chi^2 = \frac{1}{q}\sum_{i=1}^{q}\sum_{s=1}^{S}\sum_{j=1}^{r(i)}\frac{w^s(w_{ij}-w_{ij}^s)^2}{w_{ij}}$$

We also have:

$$v = \sum_{k=1}^{K}\frac{1}{q}\sum_{i=1}^{q}\frac{w_{i,\xi(k,i)}^{s(k)}-w_{i,\xi(k,i)}}{w_{i,\xi(k,i)}}w_k = \frac{1}{q}\sum_{i=1}^{q}\sum_{k=1}^{K}\frac{w_{i,\xi(k,i)}^{s(k)}-w_{i,\xi(k,i)}}{w_{i,\xi(k,i)}}w_k$$

Thus, it is enough to show that:

$$\sum_{k=1}^{K}\frac{w_{i,\xi(k,i)}^{s(k)}-w_{i,\xi(k,i)}}{w_{i,\xi(k,i)}}w_k = \sum_{s=1}^{S}\sum_{j=1}^{r(i)}\frac{w^s(w_{ij}-w_{ij}^s)^2}{w_{ij}}, \text{ for all } i=1,2,...,q \tag{23}$$

For each question $i = 1, 2, ..., q$ the right-hand term in (23) can be written as:

$$\sum_{s=1}^{S}\sum_{j=1}^{r(i)} \frac{w^s(w_{ij} - w_{ij}^s)^2}{w_{i,j}} = \sum_{s=1}^{S}\sum_{j=1}^{r(i)} \frac{w^s(w_{ij} - w_{ij}^s)(w_{ij} - w_{ij}^s)}{w_{ij}} \tag{24}$$

$$= \sum_{s=1}^{S}\sum_{j=1}^{r(i)} w^s(w_{ij} - w_{ij}^s) - \sum_{s=1}^{S}\sum_{j=1}^{r(i)} \frac{w^s w_{ij}^s(w_{ij} - w_{ij}^s)}{w_{ij}} \tag{25}$$

$$= \sum_{s=1}^{S}\sum_{j=1}^{r(i)} w^s w_{ij} - \sum_{s=1}^{S}\sum_{j=1}^{r(i)} w^s w_{ij}^s - \sum_{s=1}^{S}\sum_{j=1}^{r(i)} \frac{w^s w_{ij}^s(w_{ij} - w_{ij}^s)}{w_{ij}} \tag{26}$$

$$= -\sum_{s=1}^{S}\sum_{j=1}^{r(i)} \frac{w^s w_{ij}^s(w_{ij} - w_{ij}^s)}{w_{ij}}$$

and the left-hand term in (23) can be written as:

$$\sum_{k=1}^{K} \frac{w_{i,\xi(k,i)}^{s(k)} - w_{i,\xi(k,i)}}{w_{i,\xi(k,i)}} w_k = \sum_{s=1}^{S} \sum_{k:s(k)=s} \frac{w_{i,\xi(k,i)}^{s(k)} - w_{i,\xi(k,i)}}{w_{i,\xi(k,i)}} w_k \tag{27}$$

$$= \sum_{s=1}^{S}\sum_{j=1}^{r(i)} \sum_{k:s(k)=s,\xi(k,i)=j} \frac{w_{i,\xi(k,i)}^{s(k)} - w_{i,\xi(k,i)}}{w_{i,\xi(k,i)}} w_k$$

$$= \sum_{s=1}^{S}\sum_{j=1}^{r(i)} \frac{w_{ij}^s - w_{ij}}{w_{ij}} \sum_{k:s(k)=s,\xi(k,i)=j} w_k$$

Notice that $\sum_{k:s(k)=s,\xi(k,i)=j} w_k$ is the population share of individuals who belong to ethnic group s and give answer j, thus:

$$\sum_{k:s(k)=s,\xi(k,i)=j} w_k = w^s w_{ij}^s \tag{28}$$

and from (27) and (28) we have:

$$\sum_{k=1}^{K} \frac{w_{i,\xi(k,i)}^{s(k)} - w_{i,\xi(k,i)}}{w_{i,\xi(k,i)}} w_k = \sum_{s=1}^{S}\sum_{j=1}^{r(i)} (\frac{w_{ij}^s - w_{ij}}{w_{ij}}) w^s w_{ij}^s = -\sum_{s=1}^{S}\sum_{j=1}^{r(i)} \frac{w^s w_{ij}^s(w_{ij} - w_{ij}^s)}{w_{ij}}$$

Thus, equality (23) holds for all questions i. ■

Table 1 – Joint Significance of ethnolinguistic dummies in questions from the World Values / European Values integrated surveys, by region

	# of regressions	Share of regressions with jointly significant ethnic dummies
Whole sample	**21,469**	**0.430**
Africa	**3,623**	**0.548**
Of which: Sub-Saharan Africa	2,724	0.616
Of which: North Africa	899	0.344
Europe	**7,769**	**0.373**
Of which: Western and Southern Europe	2,369	0.313
Of which: Eastern and Central Europe	5,400	0.399
Asia	**5,656**	**0.571**
Of which: East and Southeast Asia	2,090	0.626
Of which: South Asia	852	0.667
Of which: Southwestern and Central Asia	1,511	0.479
Of which: Middle East	1,203	0.525
America	**3,749**	**0.235**
Of which: North America	741	0.513
Of which: Latin America and Caribbean	3,008	0.166
Oceania	**672**	**0.342**

Note: North America is defined here as Canada and the US. Mexico is included with Latin America and the Caribbean.

Table 2 - Joint Significance of ethnolinguistic dummies in questions from the World Values / European Values integrated surveys, by question category and question type

	Number of regressions	Share of regressions with jointly significant ethnic dummies
Breakdown by Question Category		
A: Perceptions of Life	4,382	0.425
B: Environment	971	0.427
C: Work	2,409	0.398
D: Family	1,319	0.445
E: Politics and Society	9,046	0.409
F: Religion and Morals	2,316	0.516
G: National Identity	1,026	0.495
Breakdown by Question Type		
Binary questions	4,551	0.426
Binary from unordered response questions	7,029	0.362
Scale questions	9,889	0.479

Note: This result does not change if you break it down by continent: there is little difference in shares of questions with significant ethnolinguistic dummies when the breakdown by category is done continent by continent.

Table 3 – Summary Statistics for the Main Indices of Ethnic Heterogeneity, Cultural Diversity, FST and χ^2.

Panel A: Summary Statistics

	Cultural Fractionali-zation	ELF	FST	χ^2
Mean	0.5291	0.3896	0.0116	0.0291
Standard Deviation	0.0366	0.2584	0.0110	0.0254
Minimum	0.4273	0.0000	0.0000	0.0000
Maximum	0.6024	0.8517	0.0588	0.1281

(Summary statistics based on 76 observations)

Panel B: Correlations

	Cultural Fractionali-zation	ELF	FST	χ^2
Cultural Fractionalization	1			
ELF	-0.0303	1		
FST	0.1787	0.6198**	1	
χ^2	0.2191	0.6203**	0.9813**	1

(** Significant at the 1% level; correlations based on 76 observations)

Table 4 - Incidence of Civil Conflict and Diversity
(Dependent variable: incidence of civil conflict)

	(1) Incidence	(2) Incidence	(3) Incidence	(4) Incidence	(5) Incidence	(6) Incidence
Cultural Fractionalization	**-0.187**			**-0.391****	**-0.472****	**-0.417****
	[-1.60]			**[-3.33]**	**[-3.49]**	**[-3.65]**
Ethnolinguistic Fractionalization		**0.019**		**-0.037**	**-0.041**	**-0.035**
		[1.09]		**[-1.77]**	**[-1.88]**	**[-1.78]**
Chi Square			**0.472****	**0.870****	**0.992****	**0.840****
			[2.74]	**[3.82]**	**[4.00]**	**[3.96]**
Lagged War	0.860**	0.868**	0.864**	0.840**	0.840**	0.833**
	[30.66]	[33.74]	[33.33]	[27.88]	[26.28]	[25.78]
Log Lagged GDP per capita	-0.005	-0.008	-0.007	0.003	0.005	0.003
	[-0.90]	[-1.74]	[-1.47]	[0.62]	[0.98]	[0.78]
Log Lagged Population	0.011**	0.012**	0.013**	0.012**	0.015**	0.011**
	[3.69]	[3.98]	[4.33]	[4.18]	[4.32]	[4.32]
% Mountainous Terrain	0.000*	0.000*	0.000*	0.000**	0.000	0.000*
	[2.53]	[2.34]	[2.34]	[2.60]	[1.84]	[2.02]
Non Contiguous	0.025	0.017	0.015	0.026	0.025	0.025
	[1.61]	[1.27]	[1.16]	[1.60]	[1.61]	[1.54]
Oil	0.020	0.021	0.025	0.022	0.014	0.015
	[1.56]	[1.50]	[1.80]	[1.62]	[1.34]	[1.33]
New State	0.200*	0.222*	0.240*	0.210*	0.255*	0.143*
	[2.14]	[2.29]	[2.36]	[2.14]	[2.16]	[2.11]
Instability	-0.009	-0.009	-0.009	-0.009	-0.010	-0.010
	[-1.00]	[-0.94]	[-1.00]	[-1.08]	[-1.03]	[-1.34]
Democracy Lagged (Polity 2)	0.001	0.000	0.000	0.000	0.001	0.000
	[0.90]	[0.56]	[0.20]	[0.49]	[0.75]	[0.47]
Latin America and Caribbean	0.017	0.007	0.018	0.058	0.042	0.066*
	[0.98]	[0.48]	[0.99]	[1.92]	[1.54]	[1.99]
Sub-Saharan Africa	0.022	0.001	0.000	0.051	0.064	0.060
	[0.90]	[0.05]	[0.02]	[1.53]	[1.53]	[1.65]
East and Southeast Asia	-0.008	-0.010	-0.017*	-0.016	-0.017	-0.015
	[-0.68]	[-1.12]	[-2.25]	[-1.86]	[-1.84]	[-1.86]
UK Legal Origin					0.009	
					[0.26]	
French Legal Origin					0.032	
					[0.97]	
Socialist Legal Origin					0.010	
					[0.30]	
GDP Growth						-0.119**
						[-3.50]
GDP Growth Lagged						-0.003
						[-0.08]
Observations	2,921	2,921	2,921	2,921	2,705	2,850
Pseudo R-squared	0.752	0.752	0.754	0.758	0.754	0.771

Logit estimation, based on at most 69 countries from 1945 to 1999, standard errors clustered at country level. The columns report marginal effects.
Robust z statistics in brackets. * significant at 5%; ** significant at 1%

Table 5 - Onset of Civil Conflict and Diversity
(Dependent variable: onset of civil conflict)

	(1) Onset	(2) Onset	(3) Onset	(4) Onset	(5) Onset	(6) Onset
Cultural Fractionalization	**-0.024**			**-0.075**	**-0.120***	**-0.090**
	[-0.62]			**[-1.56]**	**[-2.25]**	**[-1.91]**
Ethnolinguistic Fractionalization		**0.005**		**-0.009**	**-0.008**	**-0.008**
		[0.94]		**[-1.11]**	**[-0.92]**	**[-1.11]**
Chi Square			**0.088**	**0.175***	**0.254****	**0.179***
			[1.82]	**[2.05]**	**[2.95]**	**[2.12]**
Lagged War	-0.006**	-0.006**	-0.007**	-0.007**	-0.008**	-0.007**
	[-3.13]	[-3.22]	[-3.27]	[-3.32]	[-3.56]	[-3.05]
Log Lagged GDP per capita	-0.005*	-0.005**	-0.004**	-0.003	-0.001	-0.002
	[-2.49]	[-3.17]	[-2.85]	[-1.50]	[-0.58]	[-1.14]
Log Lagged Population	0.003**	0.003**	0.003**	0.003**	0.004**	0.003**
	[3.55]	[3.79]	[4.06]	[3.69]	[3.70]	[3.49]
% Mountainous Terrain	0.000**	0.000**	0.000**	0.000**	0.000	0.000**
	[2.94]	[2.74]	[2.92]	[2.97]	[1.73]	[2.77]
Non Contiguous	0.005	0.004	0.004	0.006	0.006	0.005
	[1.08]	[0.94]	[0.92]	[1.06]	[1.15]	[1.01]
Oil	0.021**	0.021**	0.025**	0.024**	0.018**	0.019*
	[2.80]	[2.73]	[2.96]	[2.79]	[2.65]	[2.09]
New State	0.075*	0.077*	0.080*	0.074	0.081	0.060
	[1.98]	[1.99]	[2.01]	[1.91]	[1.88]	[1.82]
Instability	0.008*	0.008*	0.008*	0.008	0.009*	0.006
	[2.00]	[2.04]	[2.05]	[1.93]	[2.10]	[1.70]
Democracy Lagged (Polity 2)	0.000	0.000	-0.000	0.000	0.000	0.000
	[0.42]	[0.26]	[-0.04]	[0.02]	[0.41]	[0.08]
Latin America and Caribbean	-0.001	-0.002	-0.001	0.004	0.003	0.006
	[-0.20]	[-0.47]	[-0.17]	[0.44]	[0.36]	[0.62]
Sub-Saharan Africa	-0.002	-0.004	-0.003	0.002	0.007	0.003
	[-0.40]	[-1.19]	[-1.00]	[0.33]	[0.77]	[0.49]
East and Southeast Asia	-0.001	-0.001	-0.001	-0.002	-0.003	-0.001
	[-0.23]	[-0.17]	[-0.53]	[-0.55]	[-1.00]	[-0.40]
UK Legal Origin					-0.006	
					[-0.75]	
French Legal Origin					0.003	
					[0.26]	
Socialist Legal Origin					0.002	
					[0.17]	
GDP Growth						-0.027*
						[-2.30]
GDP Growth Lagged						-0.008
						[-0.70]
Observations	2,921	2,921	2,921	2,921	2,705	2,850
Pseudo R-squared	0.147	0.148	0.153	0.158	0.157	0.165

Logit estimation, based on at most 69 countries from 1945 to 1999, standard errors clustered at country level. The columns report marginal effects.
Robust z statistics in brackets. * significant at 5%; ** significant at 1%

Table 6 – Endogeneity of Chi-Square and Cultural Fractionalization
(Dependent variable: incidence of civil conflict)

	(1) Persistent Questions	(2) Post-1970	(3) Post-1970 Birth Cohort < 1950
Cultural Fractionalization	**-0.222**	**-0.504****	**-0.358****
	[-1.49]	**[-2.64]**	**[-2.74]**
Ethnolinguistic Fractionalization	**-0.014**	**-0.022**	**-0.017**
	[-0.76]	**[-0.64]**	**[-0.66]**
Chi Square	**0.458***	**1.071****	**0.370****
	[2.21]	**[2.60]**	**[2.65]**
Lagged War	0.826**	0.870**	0.858**
	[27.71]	[30.48]	[26.32]
Log Lagged GDP per capita	-0.004	-0.013	0.003
	[-1.03]	[-0.92]	[0.25]
Log Lagged Population	0.012**	0.017**	0.015**
	[3.75]	[3.67]	[4.12]
% Mountainous Terrain	0.000**	0.001**	0.000*
	[2.94]	[3.27]	[2.02]
Non Contiguous	0.002	0.072*	0.038
	[0.24]	[2.12]	[1.73]
Oil	0.018	0.006	0.006
	[1.55]	[0.25]	[0.29]
New State	0.219*	0.326	0.358*
	[2.15]	[1.85]	[2.03]
Instability	-0.005	-0.017	-0.002
	[-0.58]	[-1.10]	[-0.15]
Democracy Lagged (Polity 2)	0.000	0.001	0.001
	[0.05]	[0.46]	[0.45]
Latin America and Caribbean	0.030	0.057	0.003
	[1.49]	[1.67]	[0.20]
Sub-Saharan Africa	0.011	0.028	0.001
	[0.73]	[0.67]	[0.06]
East and Southeast Asia	-0.007	-0.033*	-0.013
	[-0.76]	[-2.53]	[-1.18]
Observations	2,677	1,785	1,636
# of countries	64	69	64
Pseudo R-squared	0.758	0.793	0.796

Logit estimation, standard errors clustered at country level. The columns report marginal effects.
Robust z statistics in brackets. * significant at 5%; ** significant at 1%

Table 7 – Robustness Tests
(Dependent variable: incidence of civil conflict)

	(1) FST	(2) PRIO25	(3) PRIO1000	(4) 30 Questions
Cultural Fractionalization	**-0.361** **[-3.20]**	**-0.894** **[-3.27]**	**-0.127** **[-1.92]**	**-0.323** **[-2.80]**
Ethnolinguistic Fractionalization	**-0.038** **[-1.85]**	**-0.058** **[-1.34]**	**-0.004** **[-0.39]**	**-0.035** **[-1.66]**
Overlap measure[a]	**2.021** **[4.34]**	**1.683** **[4.41]**	**0.238*** **[2.27]**	**0.855** **[3.71]**
Lagged War	0.842** [26.74]	0.683** [15.80]	0.398** [5.78]	0.848** [30.34]
Log Lagged GDP per capita	0.002 [0.54]	0.019* [1.98]	0.003 [1.28]	0.000 [0.11]
Log Lagged Population	0.012** [4.17]	0.013** [2.76]	0.003** [2.79]	0.012** [4.01]
% Mountainous Terrain	0.000** [2.84]	0.001** [3.20]	0.000** [3.88]	0.000** [2.92]
Non Contiguous	0.028 [1.63]	0.055* [2.18]	0.008 [1.67]	0.022 [1.46]
Oil	0.019 [1.47]	0.055 [1.86]	0.022* [2.03]	0.028 [1.85]
New State	0.213* [2.15]	0.133 [1.73]	0.029 [1.45]	0.217* [2.22]
Instability	-0.009 [-0.99]	0.017 [1.23]	0.005 [1.14]	-0.01 [-1.12]
Democracy Lagged (Polity 2)	0.000 [0.46]	0.000 [0.06]	0.000 [0.24]	0.000 [0.35]
Latin America and Caribbean	0.056 [1.89]	0.112* [2.16]	0.005 [0.50]	0.049 [1.79]
Sub-Saharan Africa	0.057 [1.62]	0.177* [2.06]	0.046 [1.56]	0.038 [1.39]
East and Southeast Asia	-0.014 [-1.67]	-0.020 [-1.30]	0.001 [0.30]	-0.017* [-2.09]
Observations	2,921	2,834	2,834	2,921
Pseudo R-squared	0.758	0.591	0.514	0.757

Logit estimation, based on at most 69 countries from 1945 to 1999, standard errors clustered at country level. The columns report marginal effects.
Robust z statistics in brackets. * significant at 5%; ** significant at 1%
a: FST in column (1), Chi-square in all other columns

Appendix Table A1 - Joint Significance of ethnolinguistic dummies in questions from the World Values / European Values integrated surveys, country by country

Country	# of regressions / questions	Share of regressions with jointly significant ethnic dummies	# of ethnic groups
Albania	350	0.109	3
Algeria	287	0.296	4
Andorra	285	0.379	5
Armenia	274	0.255	4
Australia	405	0.437	7
Azerbaijan	272	0.665	10
Bangladesh	256	0.141	4
Belarus	274	0.336	5
Bosnia and Herzegovina	346	0.676	4
Brazil	282	0.106	5
Bulgaria	267	0.352	4
Burkina Faso	282	0.585	9
Canada	396	0.639	3
Chile	349	0.186	6
Cyprus	277	0.650	5
Czech Republic	271	0.059	4
Dominican Republic	271	0.089	6
Egypt	226	0.624	5
Estonia	274	0.701	3
Ethiopia	281	0.712	8
Finland	403	0.231	5
France	193	0.233	6
Georgia	398	0.445	12
Germany	287	0.206	5
Ghana	254	0.748	6
Great Britain	198	0.369	7
Guatemala	291	0.210	2
India	398	0.990	16
Indonesia	365	0.770	9
Iran	242	0.603	9
Iraq	160	0.825	5
Israel	81	0.642	2
Jordan	235	0.149	7
Kyrgyzstan	287	0.551	4
Latvia	273	0.407	11
Lithuania	273	0.282	2
Macedonia	345	0.739	6
Malaysia	276	0.699	15

Country	# of regressions / questions	Share of regressions with jointly significant ethnic dummies	# of ethnic groups
Mali	281	0.306	9
Mexico	406	0.259	6
Moldova	401	0.546	6
Morocco	386	0.215	6
New Zealand	267	0.199	6
Nigeria	320	0.838	5
Norway	285	0.225	2
Pakistan	198	0.697	8
Peru	245	0.269	7
Philippines	290	0.617	20
Poland	277	0.061	3
Puerto Rico	271	0.089	6
Romania	266	0.147	5
Russian Federation	322	0.522	4
Saudi Arabia	208	0.418	11
Serbia	276	0.370	7
Singapore	217	0.705	6
Slovakia	271	0.421	5
Slovenia	279	0.090	4
South Africa	447	0.884	12
Spain	431	0.548	5
Sweden	287	0.220	7
Switzerland	240	0.575	5
Taiwan	371	0.337	5
Tanzania	290	0.190	3
Thailand	287	0.976	7
Trinidad and Tobago	278	0.237	6
Turkey	280	0.493	5
Uganda	289	0.346	9
Ukraine	395	0.565	3
United States	345	0.368	6
Uruguay	398	0.106	6
Venezuela	217	0.078	6
Viet Nam	284	0.342	2
Zambia	280	0.782	18

Note: Using an alternative, more restricted set of questions, we find these shares to be remarkable stable. The correlation between the two series is 98.84%, despite using only half the questions, and leaving out the transformed multinomial questions.

Note that some countries have very small numbers because of a lopsided distribution of respondents across ethnic groups. Examples include Czech Republic, Poland, Slovenia, where there are few groups and a very small number of respondents in some groups.

Appendix Table A2 - Afrobarometer: Overall Results and Breakdown by Country

	# of regressions	Share of jointly significant ethnic dummies	# of ethnic groups
Overall sample	**5,427**	**0.569**	**-**
Binary response questions	138	0.522	-
Binary from unordered multiple response questions	2,200	0.441	-
Scale response questions	3,089	0.662	-
Benin	273	0.733	13
Botswana	273	0.505	24
Burkina Faso	273	0.509	22
Cape Verde	252	0.198	11
Ghana	271	0.609	25
Kenya	273	0.700	21
Lesotho	273	0.095	31
Liberia	268	0.463	16
Madagascar	273	0.692	21
Malawi	273	0.429	16
Mali	273	0.645	20
Mozambique	266	0.590	20
Namibia	273	0.546	20
Nigeria	277	0.906	31
Senegal	273	0.374	10
South Africa	272	0.879	14
Tanzania	272	0.647	38
Uganda	273	0.908	26
Zambia	273	0.516	31
Zimbabwe	273	0.403	13

Notes: 20 countries. Based on data from the 4th Afrobarometer wave only (2008).
Breakdown question category not available for this dataset; Afrobarometer does not break down questions into categories.
Regression specification includes ethnic dummies, age (Q1), present living conditions (Q4B) as a proxy for income, and gender of respondent (Q101).

Appendix Table A3 - Latinobarometer: Overall Results and Breakdown by Question Category, Question Type and Country

	# of regressions	Share of jointly significant ethnic dummies
Across All Questions and Countries:	**12,210**	**0.325**
Breakdown by Question Category:		
A: Democracy, Participation, Social Values, Trust	6,546	0.327
B: Public Policies, Corruption, Labor	3,497	0.326
C: Economics, Development, Entrepreneurship	1,099	0.354
D: Means of communication	852	0.285
E: Political Developments	216	0.292
Breakdown by Question Type:		
Binary	4,132	0.309
Scale	6,350	0.367
Binary from unordered multiple response questions	1,728	0.213
Breakdown by Country		
Argentina	679	0.178
Bolivia	682	0.453
Brazil	677	0.236
Colombia	680	0.228
Costa Rica	676	0.249
Chile	677	0.297
Ecuador	680	0.415
El Salvador	679	0.432
Guatemala	680	0.296
Honduras	678	0.355
Mexico	679	0.199
Nicaragua	679	0.247
Panama	678	0.289
Paraguay	676	0.408
Peru	680	0.334
Uruguay	675	0.215
Venezuela	677	0.427
Dominican Republic	678	0.600

Notes: 18 countries. This is based on 2007, 2008, 2009, 2010 waves of Latinobarometro, the only available waves where the ethnicity question was asked.
Regression specification includes ethnic dummies, sex (S01), age (S02), respondent education (S51) and socioeconomic level (S62) as a proxy for income.

Table A4 – Indices of Cultural Diversity, Ethnic Diversity and Overlap

Iso Code	Country	Cultural fractionalization (CF)	Ethnolinguistic fractionalization (ELF)	FST	Chi-Square
ALB	Albania	0.5239	0.0269	0.0017	0.0055
DZA	Algeria	0.4856	0.2639	0.0055	0.0175
AND	Andorra	0.5409	0.6166	0.0103	0.0267
ARG	Argentina	0.5381	0.1313	0.0032	0.009
ARM	Armenia	0.5244	0.0643	0.0038	0.0112
AUS	Australia	0.5524	0.2441	0.0058	0.015
AZE	Azerbaijan	0.493	0.2495	0.0152	0.0363
BGD	Bangladesh	0.4356	0.1507	0.0038	0.01
BLR	Belarus	0.5257	0.3833	0.0064	0.0155
BIH	Bosnia Herzogovina	0.5542	0.6261	0.0151	0.0307
BRA	Brazil	0.5632	0.582	0.0052	0.0132
BGR	Bulgaria	0.5377	0.2955	0.0104	0.023
BFA	Burkina Faso	0.5427	0.6674	0.0208	0.0588
CAN	Canada	0.5583	0.6256	0.0119	0.0254
CHL	Chile	0.565	0.1621	0.0042	0.0118
CHN	China	0.4939	0.1356	0.0126	0.0357
CYP	Cyprus	0.583	0.5038	0.0306	0.0688
CZE	Czech Republic	0.5374	0.0071	0.0025	0.0071
EGY	Egypt	0.4299	0.635	0.007	0.0156
EST	Estonia	0.5174	0.4862	0.0228	0.0462
ETH	Ethiopia	0.5453	0.694	0.0195	0.0576
FIN	Finland	0.5519	0.0531	0.004	0.0122
FRA	France	0.5825	0.1352	0.0086	0.0294
GEO	Georgia	0.4947	0.2502	0.0078	0.0191
GER	Germany	0.576	0.0729	0.0028	0.0093
GHA	Ghana	0.5293	0.582	0.0162	0.0446
GBR	Great Britain	0.5728	0.1274	0.009	0.0278
GTM	Guatemala	0.5129	0.4093	0.0023	0.0064
IND	India	0.5655	0.8517	0.0588	0.1281
IDN	Indonesia	0.4616	0.5956	0.0167	0.0379
IRN	Iran	0.4998	0.6136	0.008	0.0211
IRQ	Iraq	0.4892	0.296	0.0227	0.0415
ITA	Italy	0.5574	0	0	0
JPN	Japan	0.5273	0.0043	0.0009	0.0028
JOR	Jordan	0.4273	0.4976	0.0066	0.0165
KGZ	Kyrgyzstan	0.5489	0.586	0.0193	0.042
LVA	Latvia	0.5302	0.5869	0.0173	0.0469
LTU	Lithuania	0.5031	0.1697	0.0042	0.0098

Iso Code	Country	Cultural fractionalization (CF)	Ethnolinguistic fractionalization (ELF)	FST	Chi-Square
MKD	Macedonia	0.5428	0.4398	0.0344	0.0674
MYS	Malaysia	0.5632	0.6528	0.0336	0.0919
MLI	Mali	0.5657	0.4331	0.0128	0.036
MEX	Mexico	0.5476	0.6426	0.0057	0.0148
MDA	Moldova	0.545	0.3695	0.0108	0.0282
MAR	Morocco	0.4452	0.2235	0.0031	0.0084
NZL	New Zealand	0.5379	0.0408	0.0065	0.0227
NGA	Nigeria	0.4987	0.7675	0.0103	0.0231
NOR	Norway	0.5394	0.0661	0.0028	0.0092
PAK	Pakistan	0.4451	0.7562	0.0198	0.0482
PER	Peru	0.5234	0.5767	0.0081	0.0219
PHL	Philippines	0.5168	0.7655	0.0295	0.0765
POL	Poland	0.5443	0.0199	0.0021	0.006
PRI	Puerto Rico	0.5059	0.6382	0.0054	0.0169
ROM	Romania	0.5229	0.1387	0.0049	0.0145
RUS	Russian Federation	0.5388	0.4837	0.0035	0.0098
SAU	Saudi Arabia	0.5082	0.5148	0.0137	0.0345
SER	Serbia	0.5869	0.1896	0.0102	0.0299
SGP	Singapore	0.5099	0.6778	0.0321	0.0629
SVK	Slovak Republic	0.5418	0.1339	0.01	0.0305
SVN	Slovenia	0.5469	0.0559	0.002	0.0062
ZAF	South Africa	0.5552	0.8389	0.0191	0.0445
KOR	South Korea	0.5518	0	0	0
ESP	Spain	0.5522	0.2944	0.0052	0.0125
SWE	Sweden	0.5438	0.2277	0.0086	0.0277
CHE	Switzerland	0.5566	0.638	0.0139	0.0312
TWN	Taiwan	0.516	0.5394	0.0089	0.0203
TZA	Tanzania	0.4651	0.2744	0.0044	0.0149
THA	Thailand	0.5611	0.6979	0.0495	0.1218
TTO	Trinidad and Tobago	0.5325	0.6209	0.0094	0.027
TUR	Turkey	0.5186	0.2041	0.0092	0.0253
UGA	Uganda	0.5166	0.8114	0.0185	0.0469
UKR	Ukraine	0.5429	0.5018	0.0062	0.0135
URY	Uruguay	0.5437	0.1126	0.0041	0.0111
USA	USA	0.5512	0.3359	0.0067	0.0169
VEN	Venezuela	0.5454	0.6602	0.005	0.0156
VNM	Vietnam	0.5039	0.0814	0.0016	0.0043
ZMB	Zambia	0.6024	0.8045	0.0293	0.0834

Figure 1: Cultural Fractionalization (CF)

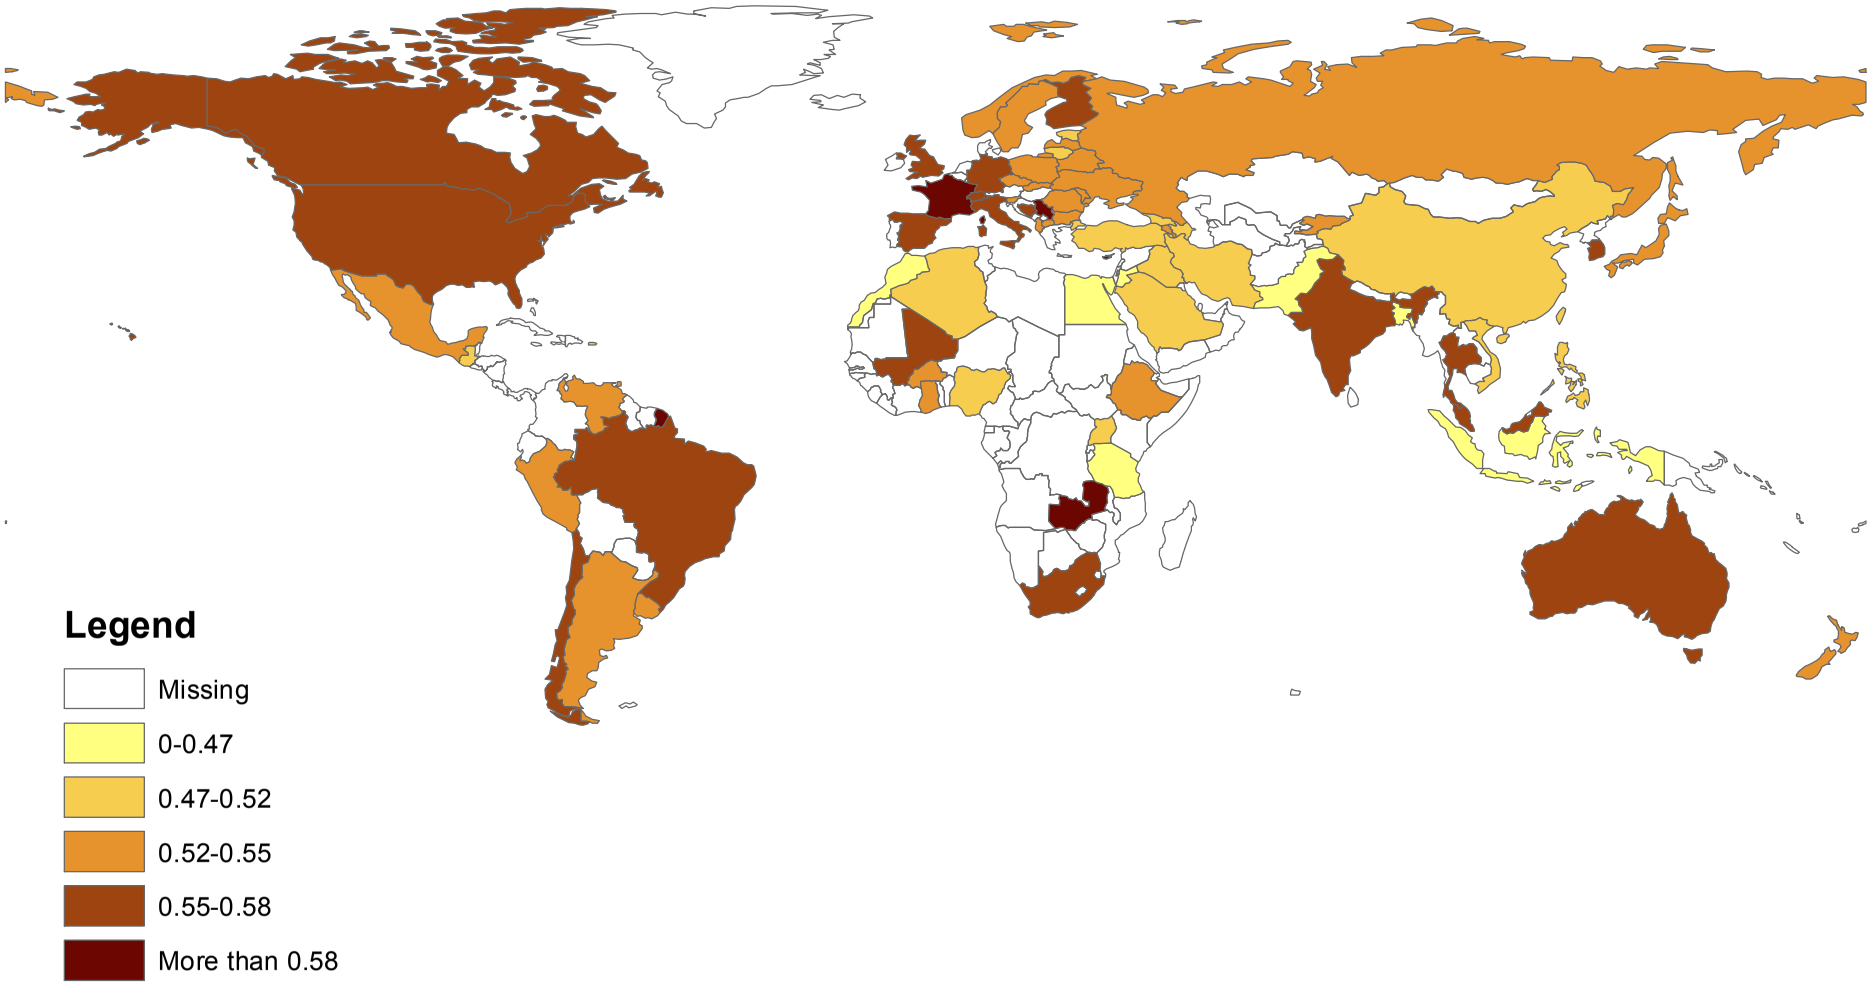

Figure 2: Ethnolinguistic Fractionalization (ELF)

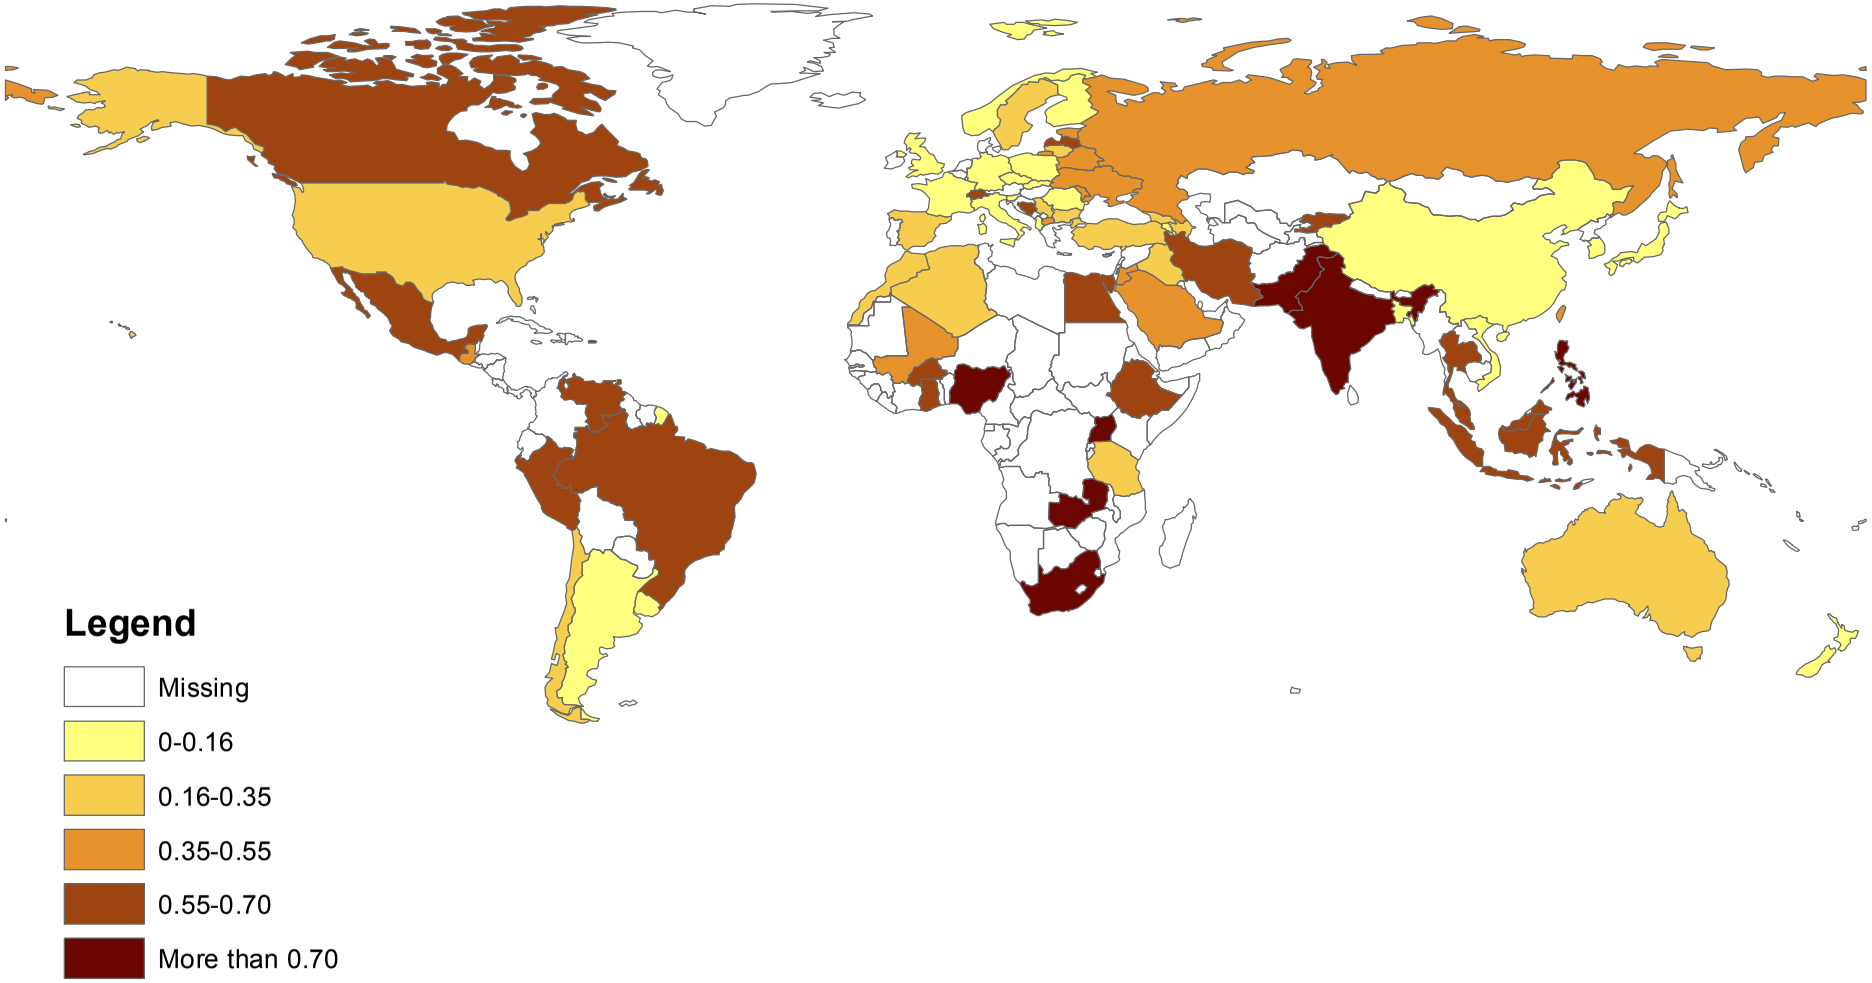

Figure 3: Correlates of Cultural Fractionalization

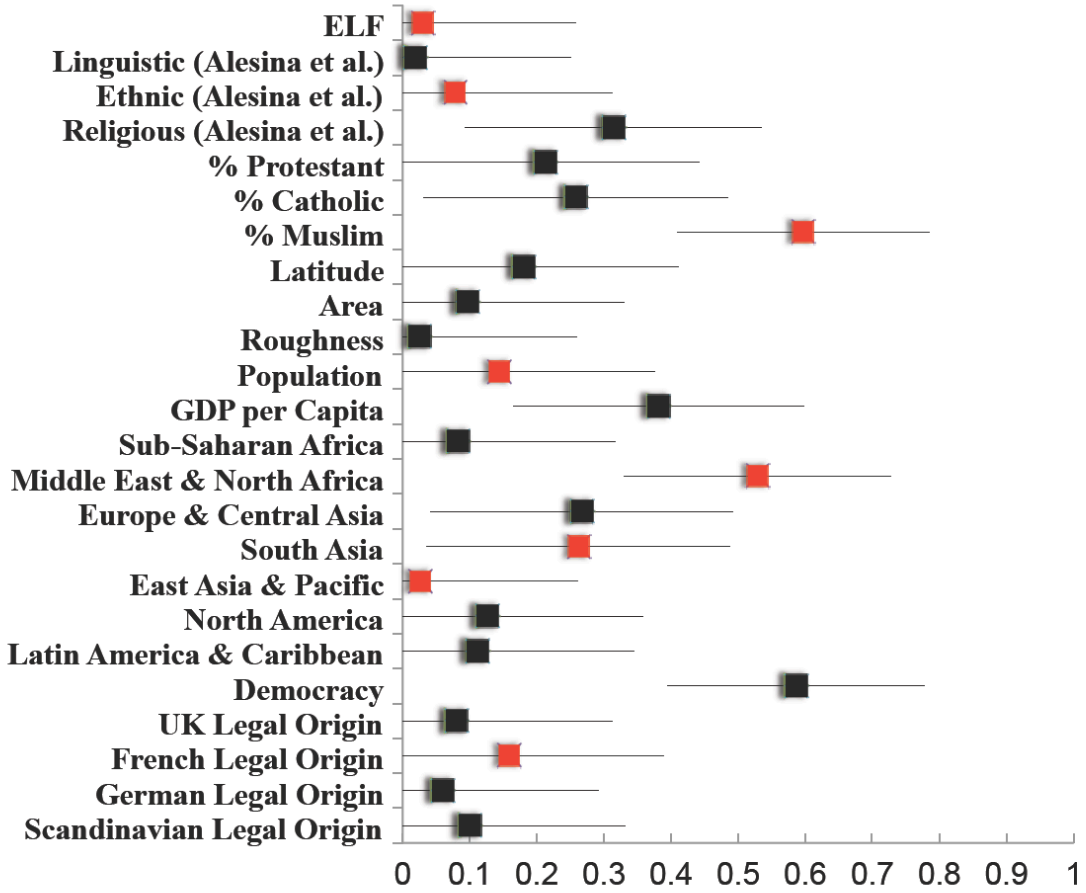

Note: Figure shows absolute values of correlations, where red denotes a negative correlation. Lines are 95% confidence intervals.

Figure 4: Chi-Square

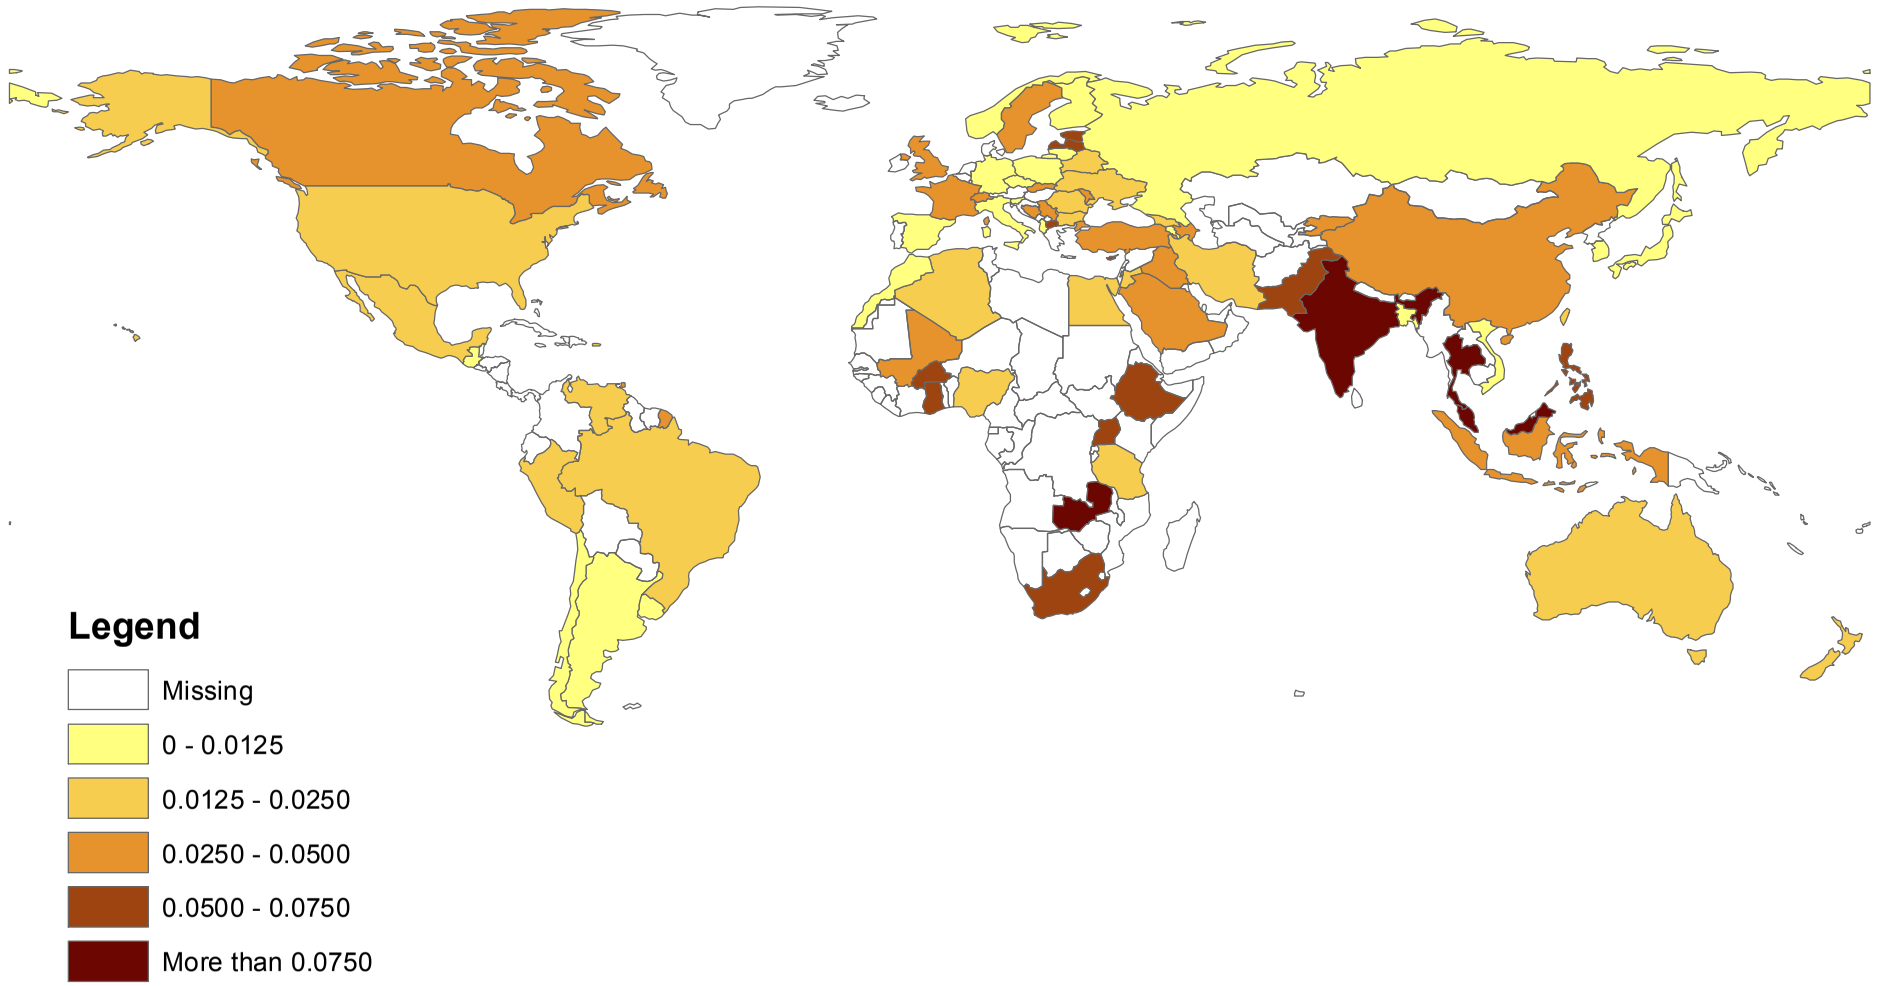

Figure 5: Correlates of χ^2

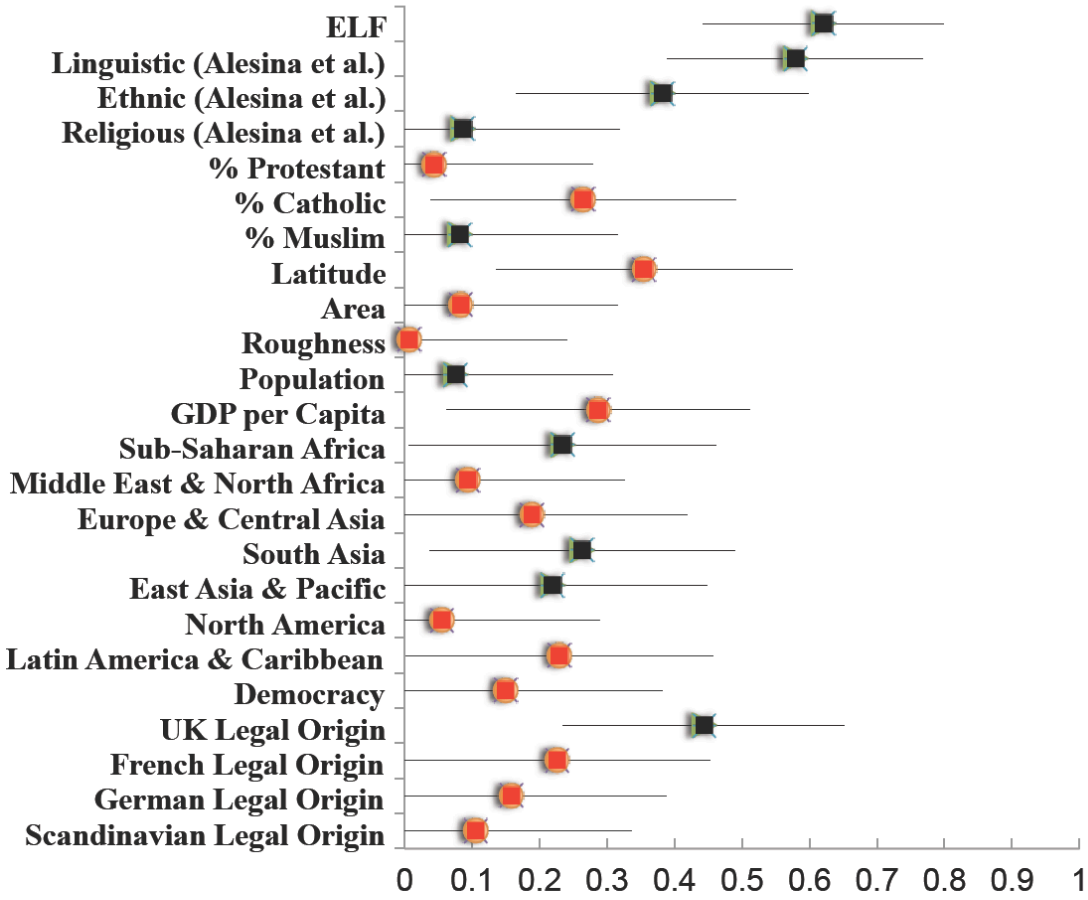

Note: Figure shows absolute values of correlations, where red denotes a negative correlation. Lines are 95% confidence intervals.