Value Fractionalization: A New Measure of Diversity

Author:
Dean Weesner

Advisor:
Quamrul Ashraf

Williams College
Williamstown, Massachusetts

May 11, 2011

Many thanks to my advisor, Quamrul Ashraf for his continual guidance throughout this project.

1 Introduction

In *The Protestant Ethic and the Spirit of Capitalism*, Max Weber argued that the Protestant world view lent a moral significance to work and the creation of wealth which then spurred the remarkable economic success of capitalism in northern Europe. The idea that the wealth and economic fate of nations is determined not only by the accumulation of capital, technology, access to trade, etc., but also significantly by the character and values of individuals, adds a new dimension of study to development economics. Embracing this idea, however, poses some challenges to those who wish to construct a rigorous model of economies – while the volume of a nation's trade or the worth of its combined capital is easily enough measured and quantified (on a conceptual if not necessarily practical level), the zeitgeist of a nation cannot easily be defined. Thus, many social scientists since Weber have tackled the questions of what specific aspects of a nation's cultural composition are most economically significant and how those dimensions can be measured.

There exists a sizable literature discussing what types of values or world views are most conducive to economic success – is it better to be individually oriented or community oriented? Adhere to tradition or encourage self-expression? Another dimension of culture that has been explored is diversity. Do highly diverse areas see economic dynamism from the exchange of ideas or stagnation due to inter-group conflict? Often, diversity is defined in terms of fractionalization between discrete groups in a population, defined by ethnicity, primary language, religion, or some other trait. This paper attempts to combine these two directions of research – the effects of types of values and the effects of diversity – by creating a novel measure of value fractionalization. This is accomplished by analyzing survey responses regarding world views to quantify values at an individual level and then measuring the dispersion of individuals in “value-space” to create a measure of diversity. The working definition of “values” that I will be using is a broad one and is meant to denote any internal traits relating to opinions, beliefs, and preferences; I use “values” and “world view” synonymously.

Previous discussions of the economic significance of diversity have been based on “external” traits such as ethnicity, religion, or language. A diversity measure based on values, an “internal” trait, has the potential to tell us something very distinct from ethnic fractionalization, for example. Ethnic or religious affiliation may certainly be correlated with one’s world view as external measures of diversity may serve as a proxy for diversity of thought and behavior, but it is certainly not the same thing. External traits are readily observable and are the primary drivers of group affiliation within a society and hence lie behind economic effects due to inter-group interactions and group identification. In contrast internal traits, i.e. world views, can be seen as a closer measure to diversity in thought and action, which may have economic significance distinct from that of mere group identification. Having a measure of value diversity in addition to traditional fractionalization statistics thus will allow for more specific inferences into the role of diversity on economic development. In particular, measures such as ethnic fractionalization are often applied to growth regressions, allowing a researcher to conclude whether ethnic diversity is a positive or negative contributor to economic success, as indicated by the sign of the OLS coefficient on that variable [3, 4, 6, 16]. I use the the same strategy of including the value diversity measure as an explanatory variable in cross-country growth regressions and find tentative evidence of a negative net effect of value diversity that is distinct from the effects due to ethnic diversity.

2 Background

There are many intuitive reasons why diversity (broadly construed) could effect economic development in both a positive and negative way. New York City's economic dynamism, for example, is often claimed to be the result of its extraordinarily diverse population [15]. Diversity in observable dimensions such as ethnicity and language may imply diversity in ways of thinking and problem solving, leading to skill complementarities between members of different groups and increased productivity. Additionally, diversity may directly improve welfare by increasing the range of consumption choices available – consider the privilege of choosing between take-out menus for cuisine of any continent. On the other hand, diversity may engender inter-group conflict which would diminish growth. Easterly and Levine cite this as a major cause of Africa's lack of growth during the 20th century [6]. These conflicts could be explicit such as civil war and ethnic violence, or less obvious in the case of groups promoting rent-seeking public policies that enrich one group at the expense of overall economic progress. On a more micro level, people tend to trust members of other groups less which effectively increases transaction costs between individuals of different groups. Alesina and La Ferrara (2005) model the trade-offs of diversity as positively influencing "private" consumption through enhanced productivity but negatively influencing "public" consumption as diversity in preferences makes providing a public good less efficient [3]. Among these hypothesized mechanisms that relate diversity and economic activity, the benefits of diversity tend to be related to differences in internal traits, to the effect of representing different behavior, while the negative effects of diversity tend to be related to group identification in and of itself which would be associated with more apparent external traits. This distinction is supported by Ratna and Grafton (2009) who find that in U.S. states racial diversity, a highly-salient group identifier, negatively impacts growth while linguistic diversity, more indicative of cultural diversity and also easily obscured by use of English as a common language, positively impacts growth [17].

Past research on the net effect of diversity on growth has found that diversity is generally negatively associated with GDP growth, with some caveats. In Sub-Saharan Africa, the negative relationship between ethnolinguistic fractionalization and economic development is well documented, even after controlling for other factors [3, 6]. However, comparisons within the United States have yielded different results: Ottavino and Peri (2005) find a positive relationship between country of birth fractionalization and wage and rent growth in American cities [16]. Alesina and La Ferrara (2005) find a negative effect of ethnic fractionalization on income growth in a broad cross-country study, but a positive interaction of diversity and income level, implying that wealthier countries may experience diversity differently than poorer ones [3].

These studies have used a similar research framework. First, a fractionalization index is created to measure diversity along some dimension; the index used is the opposite of the Herfindahl-Hirschman concentration index (HHI). The index represents the probability that with M different groups, two randomly selected individuals will belong to different groups. It is defined in Equation 1, where s_i is the share of the total population that group i has.

$$Div = 1 - \mathrm{HHI} = 1 - \sum_{i=1}^{M} s_i^2 \tag{1}$$

This measure ranges from 0 (when all members belong to the same group), to 1 (when each group's share of the population is infinitesimally small). To construct this index, we must assume that everyone in a population can be sorted into one of many discrete groups. The dimension along which we measure diversity could be ethnic group, linguistic group, nation of birth, religion, or any other observable and distinguishing feature. This measure is then used as an explanatory variable (along with other controlling variables) in an OLS regression on an economic measure such as income per capita, wage growth, average education level, etc. The sign on the coefficient of the diversity measure in such a regression estimate would then be an indication of whether diversity helps or harms economic growth.

This methodology has several shortcomings. First, group affiliation (ethnic or otherwise), can sometimes be subjective, and distinctions between groups can be the result of an arbitrary decision made during data collection. This problem can arise from the level of disaggregation: should a person be measured as protestant or Lutheran? Black or Nigerian? The most salient level of distinction can also change over time: if you asked someone in Somalia before 1990 what ethnic group they affiliated with, he or she would most likely say "Somali." Now, that answer is much more likely to be a local clan [6]. There is no objectively correct answer to the question of group definition, and datasets from different countries make different distinctions between groups, further complicating cross-country comparisons. Additionally, what dimension of diversity (ethnic, linguistic, religious, etc.) is most relevant is unclear. Alesina et al. (2003) compare several different types of fractionalization and note that the effect of diversity on economic measures is sensitive to which dimension diversity is defined on [4].

The second shortcoming is this measure does not capture "distance" between groups. Consider: a community of equal numbers of Korean and Japanese would have the same measured diversity as a community made of equal numbers of Irish and Cambodian. There is a very strong intuition that some group distinctions are more salient than others, but this is not captured by the fractionalization index typically used. Fearon (2003) attempts to address this issue by using structural differences in language as a proxy for cultural distance, and weighting the diversity measure accordingly [7]. This remains a crude solution, however.

Third and last, does ethnicity (or a similar variable) capture the essential features of diversity we're interested in, or is it a stand-in for something deeper and not readily observable? The hypothesized effects of diversity generally fall into two categories: those resulting from group identification itself, or complementarities in production due to diversity of thought. Insofar as we are concerned about the former effect, ethnicity may be an appropriate explanatory variable because individuals often recognize member and non-members of their group based on such readily observable characteristics. However, the second hypothesized effect of diversity posits that people don't just appear different but in fact act different as a result of diversity. If we are concerned about deeper, more fundamental differences in world views and attitudes, ethnicity may be only a crude proxy for what we wish to measure. Certainly it's reasonable to expect values to be correlated with group membership, but knowing ethnicity, language, or religion only gives a hazy picture of a given individual's behavior. By measuring values directly instead of through the lens of ethnicity, we would have a more accurate measure of these differences and furthermore be able to distinguish between effects from external and internal measures of diversity.

3 Measuring Value Diversity

3.1 Quantifying Values

While thoughts and values may seem rather intangible, many researchers have been able to incorporate measures of world views into empirical work. Usually, this is accomplished by using survey data to quantify some aspect of (selfreported) behavior or opinion. For example, Guiso et al. (2003) identify intensity of religious belief in populations by using survey data on frequency of religious service attendance [9]. Likewise, to construct a measure of trust for a population, one could ask survey respondents a question such as "Can most strangers be trusted?" and let the percentage replying in the affirmative be the statistic. Such measures can then be compared to economic outcomes to see the effect of a specific dimension of values [10].

While a single survey question gives a narrow view of a particular dimension of one's world view, appropriately combining responses from many questions allows one to take a more complete picture of a person's beliefs. Inglehart and Baker (2000) describe a method for measuring and quantifying values holistically based on data from the World Values Survey (WVS) – though their methods would also be applicable to other similar datasets such as the General Social Survey [11]. The WVS samples households in many different countries, asking them questions about their backgrounds and personal beliefs which affords a wide array of dimensions to capture a respondent's world view. The WVS is an survey of 87 nations repeated over 5 "waves": 1981-1984, 1989-1993, 1994-1999, 1999-2004, and 2005-2008. The first wave of surveys from 1981-1984 included few countries, however, and is excluded from my analysis. Over a thousand different questions are asked across the five waves on topics such as personal history, family structure, demographics and, most importantly, beliefs and values. These questions were designed to be broadly relevant, allowing for cross-country comparisons of world views.

The method for measuring and quantifying world views works as follows. First, select a subset of survey questions which encompass important dimensions of a world view. I base my analysis on the same questions as Inglehart and Baker (2000), taking ten questions that are consistent across waves and encompass many aspects of belief and opinion. The questions are [1]:

- How important is God in your life?
- 4-item autonomy index [2]
- Is abortion justifiable?
- How proud of you of your nationality?
- Would greater respect for authority be a good thing?
- 4-item post-materialist index [3]
- How often do you feel happy?
- Have you ever signed a petition?
- Is homosexuality justifiable?
- Can most people be trusted?

Since the responses to the above questions are coded numerically, we can observe systematic correlations between the response variables. Presumably,

[1] The WVS codes of the questions are: f063, y003, f120, g006, e018, y002, a008, e025, f118, a165

[2] Composite score of questions related to autonomy. See WVS codebook for full description.

[3] Composite score of questions related to post-materialist values. See WVS codebook for full description.

these correlations are the result of the observed responses being driven by some underlying traits. Since the same underlying traits may contribute to several responses, there is some redundancy in the observed response variables. Principal-component factor analysis is a technique that allows us to extract the unobserved and underlying traits, or factors, from the observed variables. This analysis takes the correlations between survey question responses and generates hypothetical unobserved variables (factors) which when linearly combined explain the variation in the observed data. For example, with k observed variables (survey responses) and two factors, the factor analysis generates a model like the one below.

$$
\begin{aligned}
Q_1 &= \alpha_1 + \beta_{1,1} F_1 + \beta_{2,1} F_2 + \epsilon \\
Q_2 &= \alpha_2 + \beta_{1,2} F_1 + \beta_{2,2} F_2 + \epsilon \\
&\vdots \\
Q_k &= \alpha_k + \beta_{1,k} F_1 + \beta_{2,k} F_2 + \epsilon
\end{aligned}
$$

As there are fewer factors than questions, by using these unobserved factors instead of the responses directly we can summarize many questions into a smaller number of variables. At the same time the significance of the questions is weighted by how well they correspond to the factors, called the factor loading. For example, how often a respondent goes to church and whether the respondent believes in an afterlife could both be explained by an unobserved variable “religiousness”. By using factors instead of individual survey responses, we can describe people in terms of more fundamental traits as opposed to just a particular behavior. Additionally, Varimax rotation is used to improve the significance of the results and ensure the generated variables are orthogonal. Knowing the relationship between the survey responses and hypothetical factors, for each individual in the WVS dataset we can solve for the specific values of the factors using the known survey responses. This generates new variables for the dataset, called factor scores or in this context value scores, which describe the traits underlying the above survey questions. Factor analysis on the above WVS questions yielded two factors which together explain 39% of the variance in survey responses.

Several decisions must be made while generating the factor scores. First, how many factors should be used? Fewer factors affords heuristic simplicity at the expense of explanatory power; more factors diminish the significance of a given factor. To resolve this question I use the Kaiser Rule which says to take only the factors which have eigenvalues (a measure of statistical significance) greater than 1. Though unimportant to eventual discussions of diversity, one may also want to know what dimension a factor corresponds to. Factor analysis is blind to the real-world significance of the data, but by examining the factor loadings one can come up with a heuristic definition. For example, Inglehart and Baker find two statistically significant factors using the same questions (though without the 2005-2008 WVS wave data) and name them the “traditional - rational/secular” dimension and the “survival - self-expression” dimension [11].

The second, and more difficult decision, is the selection of questions on which the analysis will be based. What questions should we include? How do we know if those we've chosen cover all the important dimensions of world views? There is no objectively correct way to select the questions for analysis; which observed variables get chosen is a judgment call left to the researcher. This is because, fundamentally, “values” is a somewhat amorphous concept; it should be expected that, say, norms of conformity be included in a discussion of values but persuasive arguments could be made for or against including something like love of the outdoors as part of a world view. The selection of survey questions indeed creates some room to fudge the results. Yet, there are reasons to believe that this particular question selection is an appropriate one. With 2 factors, the factor scores for each household in the survey place the observations into a 2-dimensional value-space. When the households of a country are plotted together, they form a (hopefully representative) distribution of a population's values. Summary statistics can be computed for this distribution which tell us about the overall values for the population. Mean scores can be computed allowing for cross-country comparisons of “average” values; when these national scores are plotted together, countries cluster together in an intuitive way. Figure 1 is such a plot based on Inglehart and Baker's analysis of this data: nations which would be expected to have similar cultures - Latin American, Protestant European, ex-communist Europe, and Confucian counties - sit close to each other in “value-space” [13, 11, 12]. We thus see that the factor scores seem to

Figure 1: Distribution of Mean Value Scores [13]

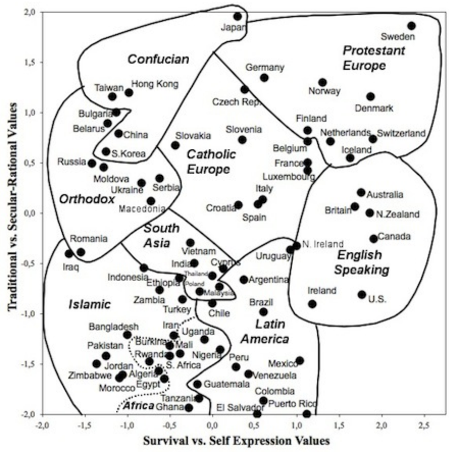

have the desired real-world significance and are not merely a statistical artifacts.

If different questions were chosen, what difference would it make? I check the robustness of question selection by trying another factor analysis using a larger set of questions[4]. This larger question set produces more significant factors, but the diversity statistic that is ultimately calculated from the factor scores ends up being decently correlated ($r = .7$). So while the factor scores themselves are sensitive to the selection of survey questions, the value fractionalization statistic that is calculated from them is robust. While it may seem that adding additional questions is weekly improving (more potential dimensions of world view to describe), I nevertheless use the 10-question set because the significance of its factors has already been studied and vindicated [11, 12, 13].

[4]WVS codes: f063 y003 f120 g006 e018 y002 a008 e025 f118 a165 a035 a001-a006 a062 a008 a009 a165 a173 a170 c060 c002 f022 f063 d017 d018 d019 a025 a030 a032 a034 a035 a038 a039 a040 a042 e023 e037 y001 e014 e016 e018 e022 e104 e108 f118 f120 f121 g006

3.2 Measuring Heterogeneity In Values

A measure of value diversity in a population can be generated by calculating the fractionalization of households' factor scores. In contrast to the fractionalization measures based on ethnicity, there are no pre-determined groupings with which to construct a Herfindahl index. Thankfully, cluster analysis provides a means to construct such groupings while preserving information about the distance between the constructed groups. There are several variations on the techniques for cluster analysis, but all essentially use the following steps. First, a measure of dissimilarity is chosen. I use the Euclidean distance between the value coordinates of two given observations. That is, for agents i and j, with value scores x and y (in the case of there being two value factors), the measure of dissimilarity, d, is:

$$d = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} \tag{2}$$

Next, for a given d, groups can be constructed according to a chosen linkage method. I use the average linkage method, which creates the least number of groups possible such that the mean dissimilarity of observations within a group is at most d. So, for $d = 0$ each entity is its own group, and as d increases the groups agglomerate into larger groups until at some critical value, d_c, all entities belong to the same group. This process can be illustrated with a dendrogram diagram. Figure 2 is an example of a cluster analysis dendrogram using the United States data (only the top 20 groupings are shown for clarity).

Figure 2: Cluster analysis dendrogram

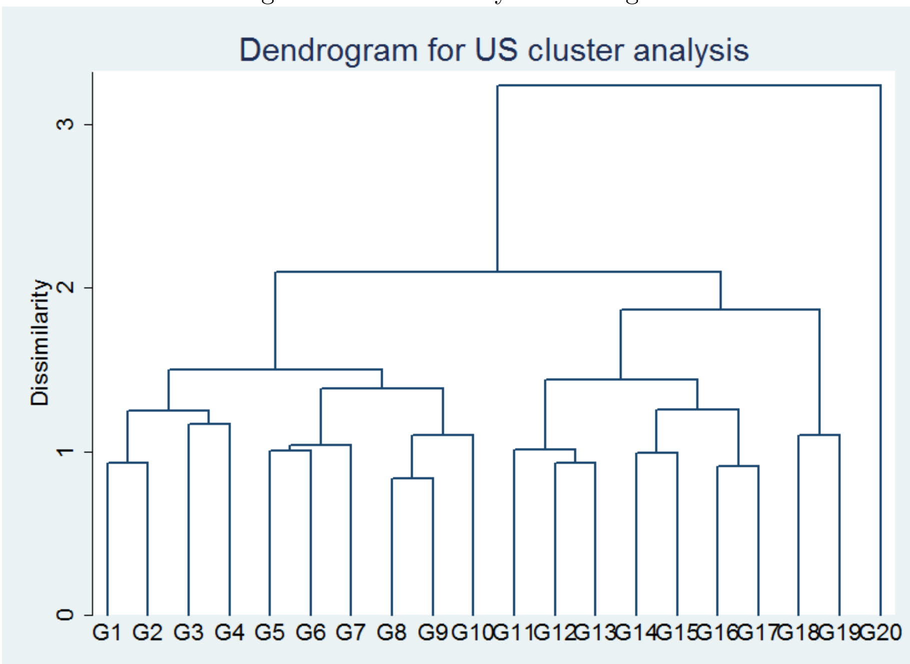

The dendrogram is incredibly useful for understanding the results of a cluster analysis: it visualizes how the cluster analysis describes the entire hierarchical structure of a population's possible groupings. Each vertical line represents a group or cluster of observations, and as the dissimilarity tolerance increases (moving up in the diagram) groups combine to make larger clusters. We can imagine making a horizontal "cut" on the dendrogram (i.e. choosing a value for d) and thus defining a possible partitioning of the population; choosing a different value of d would define a different level of aggregation and a different partitioning. With the population partitioned into groups based on similar world views, it is possible to compute a fractionalization index as in Equation 1.

However, as the dendrogram shows, there are many possible ways to partition the population, and thus many possible values for the measured fractionalization. It would thus be poor methodology to base an analysis upon an arbitrarily chosen dissimilarity level. Because the groupings change at discrete values of d - not continuously - the group assignment and thus computed fractionalization may be very sensitive to the choice of d in the vicinity of where groups get combined. From examination of a dendrogram one may be able to hazard an opinion as to a "representative" value of d, but this sort of ad hoc decision would be impossible to defend when a large number of populations are being compared against each other.

The following methodology resolves this problem by constructing a statistic which includes all values of d – essentially, by integrating fractionalization over dissimilarity. For the following discussion, assume that the populations are large [5] so that when each entity is its own group i.e. $d = 0$, the HHI is,

$$HHI = \lim_{N\to\infty} \sum_{i=1}^{N} s_i^2 = \lim_{N\to\infty} N \left(\frac{1}{N}\right)^2 \approx 0 \tag{3}$$

Now, recall that each d defines a partition of the population (a "cut" anywhere on the dendrogram yields some set of groups). In turn, each partition defines an HHI value. As a result, we can imagine a function for each country j, $H^j(d)$ (superscript emphasizes that the function will be different for each population), taking d as its argument and returning the HHI for the groups defined

[5] Relaxing this assumption slightly biases smaller populations towards being less fractionalized

at that d [6]. For each population, $H^j(d)$ increases monotonically from 0 at $d = 0$ to 1 at $d = d_c$. (Because groups combine at discrete events, this will technically be a step-wise function.) Since we are concerned with measures of diversity (as opposed to concentration or similarity), we use $1 - H^j(d)$ which of course ranges from 1 to 0. A hypothetical graph of one such function is sketched in Figure 3.

This function provides us with a lot of information about the structure of the population. The value of the function describes how fractionalized the population is at a given dissimilarity tolerance; how fast it decreases with d tells us how far apart those groups are. Finally, we define the following statistic for country j:

$$V = \int_0^\infty (1 - H^j(d))\mathrm{d}d \tag{4}$$

Since H^j is not explicitly defined, the statistic V can be evaluated with numeric integration by calculating HHI for the groups generated at small increments of d. The formula defines an infinite upper bound for the integral, but one can conclude numeric integration once the integrand equals zero because it will remain zero for all larger values of d (this occurs at $d = d_c$ when all individuals

[6] More formally: Let I be the set of all individuals in a given population. There exists a function (determined by the cluster analysis) $U : \mathbb{R}_+ \to P$, $P \in C$, with C the collection of all partitions of I, such that $|U(0)| = |I|$ and $|U(d_c)| = 1$ for some finite d_c. Define,

$$HHI(P) = \sum_{i=1}^{N} \left(\frac{|S_i|}{|I|}\right)^2 \;;\; P = \{S_1, \ldots, S_N\}$$

$$H(d) = HHI(U(d))$$

Note,

$$\bigcup_{i=1}^{N} S_i = I$$

Therefore $U(d_c) = \{I\}$ which implies $U(d_c) = HHI(U(d_c)) = 1$. Furthermore, this implies that $|S_i| = 1$ when $S_i \in U(0)$. So as I becomes large, $H(0) = HHI(U(0)) \approx 0$.

Figure 3: Fractionalization vs. Dissimilarity

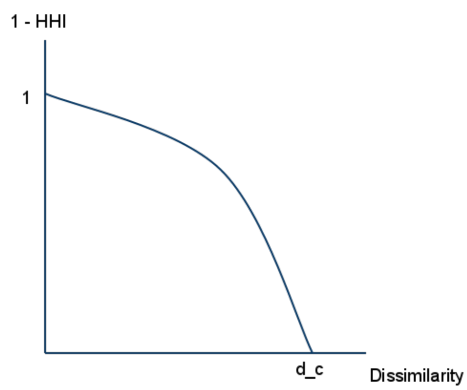

belong to the same group). Thus V in Equation 4 is the value fractionalization statistic that will measure value diversity.

This measure addresses the previously discussed shortcomings of ethnic (or similar) fractionalization measures. First and most importantly, this new statistic measures a different kind of diversity, that related to internal beliefs as opposed to external appearance. Sort of extensive field observation of individuals, this is the best indicator of how people behave and think differently, as opposed to superficial characteristics which could only serve as a proxy for these differences. Second, the approach is immune to critiques of groups being defined arbitrarily or subjectively. Using cluster analysis ensures that group distinctions have mathematically objective foundations, and the integrating over dissimilarity in the diversity statistic allows us to examine all possible levels of group aggregation at once. Third and last, assigning group membership based on quantified individual characteristics means that it is possible to capture a measure of distance between groups. Since populations whose members are further apart in value-space will have a larger diversity statistic, distance between groups is appropriately accounted for.

This methodology may seem excessively complex, but it captures important information that a more naive measure of diversity might leave out. The following comparison will demonstrate how it differs from other measures and illustrate how the value fractionalization statistic is calculated. Consider an alternative diversity measure that takes the mean pairwise distance between every individual in a population: d from Equation 2 is the distance between any two individuals so $E(d)$, the expected value of d is the alternative measure[7]. Recall that the average linkage clustering method creates groups such that the mean pairwise distance is at most d; thus, $E(d)$ is the same as the dissimilarity where the cluster analysis puts all individuals in the same group or the height of the dendrogram. So using the average distance diversity measure is like looking at only the height of the dendrogram while losing all the information contained in its body. Suppose there are three individuals in a population and their value scores position them as shown in Population 1 and Population 2:

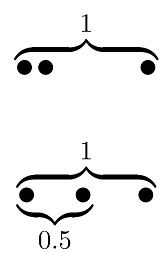

For both populations $E(d) = \frac{2}{3}$ so they would have the same measured diversity. But clearly, the spatial structure is different between the two: the second population is more spread out in value-space and more hence more diverse. For the first population there are two groups for $d < \frac{2}{3}$ and one thereafter. So the value fractionalization is,

[7] Since the measure should be extendable to arbitrarily many factors or dimensions, standard deviation is a desireable alternative measure here.

$$
\begin{aligned}
V_1 &= \frac{2}{3}\left(1 - \left[\left(\frac{2}{3}\right)^2 + \left(\frac{2}{3}\right)^2\right]\right) \\
&= \frac{10}{27} \approx 0.370
\end{aligned}
$$

For the second population, there are three groups for $d < \frac{1}{2}$, two for $\frac{1}{2} \leq d < \frac{2}{3}$, and one thereafter. So the value fractionalization is,

$$
\begin{aligned}
V_2 &= \frac{1}{2}\left(1 - 3\left(\frac{1}{3}\right)^2\right) + \frac{1}{6}\left(1 - \left[\left(\frac{2}{3}\right)^2 + \left(\frac{1}{3}\right)^2\right]\right) \\
&= \frac{23}{54} \approx 0.426
\end{aligned}
$$

We see that the second population is more diverse according to the value fractionalization measure; value fractionalization is a more complete description of the group structure of a population than a simpler dispersion measure.

4 Global Value Fractionalization and the World Values Survey

So what can value diversity tell us about the world? In this section and the next I will discuss the interesting cross-country comparisons that value diversity as computed from the World Values Survey (WVS) provides. I calculate value fractionalization two ways based on the results of the factor analysis. First, value fractionalization is calculated for each country by survey wave. Since only four waves are used, this does not provide very much time variation so the second analysis is based on a synthetic panel: the WVS has data on date of birth for respondents, so I generate cohort groups for people who would be between 20 and 60 years old at five year intervals from 1960 to 2005. Previous research suggests that world views are largely fixed after entering adulthood, so even though the data is based on contemporary cohorts it may still be an accurate read on the composition of that cohort years ago [1, 2, 11, 19]. I choose the age range 20-60 because those are the ages where individuals are working and have the largest economic impact. (Note that individuals are included in multiple cohorts in the synthetic panel.) This by-country, by-cohort panel allows more time variation, having data from 1960 to 2005 and nine time periods.

Table 1 gives basic summary statistics of the by-cohort and by-wave value fractionalization. As the table shows, aggregate measures are very similar for both data sets. The by-cohort data has a few low-diversity outliers in early years of the panel. Figure 4 shows the distribution of diversity scores in the by-cohort panel in more detail.

Some interesting trends appear when looking at how value diversity varies geographically. Table 2 shows the most and least diverse countries based on average national diversity scores across different cohorts. We see that the most diverse countries tend to be more developed, and Westernized counties while the least diverse are less developed and more traditional/non-Western. A similar trend is seen in measures of religious fractionalization, and it has been hypothe-

Figure 4: Distribution of Diversity Scores

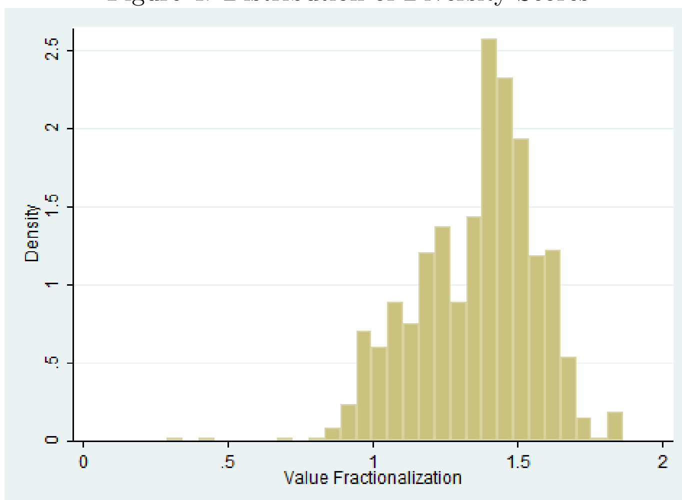

sized that this is due to more developed countries being more accommodating of religious diversity [4]; likewise, developed and liberalized countries may be more tolerant and encouraging of varied world views. Looking at average regional value diversity, we see a trend with European and Western countries being the most diverse (Table 3)[8].

Other interesting trends can be seen in the correlations of value diversity with other measures (Table 4). One WVS question[9] asks whether "Tolerance and respect for other people" is an important quality for a child to learn; the percentage of respondents answering in the affirmative can be used as a measure of tolerance. Table 4 shows a positive correlation between tolerance and value diversity which supports the idea that more pluralistic societies can support more variety in world views. Interestingly, value diversity has a weekly negative

[8]Europe and West includes European countries as well as Canada, the United States, Australia, and New Zealand; Mideast region includes North Africa.

[9]WVS survey code a035.

correlation with ethnic fractionalization and almost no correlation with religious fractionalization – this is reassuring, because it indicates that the metric of value diversity is measuring something distinct from other measures of diversity[10]. Lastly, there is a positive correlation between value diversity and real GDP per capita which is the subject of the next section.

5 Economic Development and Value Diversity

As an application of the value diversity measure, I include value diversity as an explanatory variable in growth regressions of the countries included in the WVS. I base my regression models on similar ones used by past research on the effect of ethnic fractionalization on growth rates [4, 6, 7]. The general model is given by Equation 5. Annual growth rate of real GDP per capita (in US dollars) is the dependent variable and measure of economic development. $V_{i,t}$ denotes the value fractionalization for country i in time period t. The time period may correspond to a five-year period for the by-cohort dataset or a four- to six-year period for the by-wave dataset. The estimated coefficient β is then the marginal effect of an increase in value diversity on the growth rate. I also include a vector of controls, X, which includes initial log of GDP per capita, telephones per 100 persons (a measure of infrastructure), percent of population living in urban areas, and education (as measured by percentage of primary school completion). The data for these controls come from the World Bank Development Indicators (WDI) except for educational attainment data which

[10] Data for ethnic and religious fractionalization from Alesina et al. (2003).

is from Barro and Lee [5, 20]. Since the time periods in the regression model are longer than one year, I averaged data that was collected annually to obtain multi-year data where necessary. The WDI and WVS contain data on different sets of countries, so regression estimates use only those 68 countries that are in both datasets.

$$\text{Growth}_{i,t} = \alpha + \beta V_{i,t} + \Pi X_{i,t} + \epsilon_{i,t} \tag{5}$$

Table 5 shows the results of several regression specifications for the by-cohort dataset. Regression (1), which is without any sort of fixed effects included, shows a coefficient of -2.14 for the effect of value diversity. To give some perspective on the magnitude of the effect, the coefficient implies that increasing value diversity by one standard deviation would lead to an expected decrease of 0.45% in GDP per capita growth. Adding decade fixed effects doesn't change the diversity effect much, but adding regional fixed effects destroys the significance of value diversity. This could be because within-region variation in value diversity is less than overall variation (Europe has particularly low variation and also has the most observations) so regional fixed effects absorb much of any detectable effect of value diversity. Since Africa is an outlier in terms of growth, Regression (5) includes a regional dummy only for Africa to explicity control for any results that are driven by an "Africa effect". In this case, the magnitude and significance of the value diversity diminishes somewhat, but remains negative.

Table 6 adds additional variables of interest to the model. Since past research into ethnic fractionalization and growth has shown differing effects for advanced and developing economies, Regression (6) includes an interaction between value diversity and GDP per capita (in 1000's of 2000 U.S. dollars). Interestingly, there is a negative interaction between value diversity and wealth, meaning that homogeneity is better in advanced economy in terms of growth. As mentioned, the WVS contains data on tolerance of other cultures. Regression (7) includes an interaction between value diversity and tolerance to test whether more tolerant societies benefit more from diversity. Contrary to intuition, including the tolerance interaction shows a net benefit of value diversity, but with more tolerant societies experiencing a more negative effect from diversity. Regression (8) controls for ethnic fractionalization and the results of this specification are consistent with the those of the basic specifications and past results related to ethnic fractionalization – both have negative effects. Table 7 shows results of some of the basic specifications applied to the by-wave dataset but there are almost no significant results, possibly due to the smallish sample size.

Lastly, I replicate the analysis of Easterly and Levine (1997) with the addition of the value diversity metric using their original dataset which contains countries different than the ones in the WDI as well as additional control variables [11]. Easterly and Levine used data from African and Caribbean countries from the 1960's through 1970's and found a robust negative association between ethnic fractionalization and growth. I run regressions that include: (13) ethnic fractionalization, (14) value diversity, and (15) both ethnic fractionalization and value diversity. We again see a negative effect of ethnic fractionalization but no

[11] Some variables are omitted from the regression table for brevity. See Easterly and Levine (1997) for the full list.

effect of value diversity in either specification.

6 Conclusions

I have outlined here a methodology for quantifying diversity of values and world views, a potentially important aspect of a population's cultural dynamics which could otherwise be easily dismissed as subjective and unmeasurable. When value diversity is used as an explanatory variable in growth regressions, we see that there appears to be a net negative association of this type of diversity – the costs of mitigating inter-group conflict are larger than the gains to productivity. While the hypothesized mechanisms relating economic activity and diversity predicted a more positive effect for an internal trait such as values, this appears not to be the case. The significance of this result is not very robust to the specification, however, so the result is somewhat tentative. An equally important result is that we observe distinct effects from value and ethnic fractionalization. This means that the previously documented effects of ethnic diversity were not due to ethnicity acting as a proxy for internal traits, but rather ethnic identification itself. There of course remain some unresolved issues with the value diversity measure which future researchers hoping to use it could explore. First, how does culture and thus value diversity react to economic changes? To what extent is value diversity endogenous to the growth models used here? Lacking a natural experiment or instrument, the direction of causality between economic growth and diversity cannot be stated with certainty. Second, how reliable is the survey data? A core assumption is that the survey responses comprise a representative depiction of a population's world views. If people do not respond honestly to the surveys or the households sampled are not representative of the larger population, this would skew any inferences made from the data. The general idea, however, of using the dispersion of factor scores as a measure of diversity for difficult to quantify "soft" variables has broad applications to development economics and any other area of social science where a metric of opinion diversity would be useful.

Table 1: Value Diversity Summary Statistics

	Mean	Std. Deviation	Min	Max
By Wave	1.417	0.221	0.894	1.904
By Cohort	1.365	0.212	0.400	1.864

Table 2: Mean National Diversity

Most Diverse		Least Diverse	
Croatia	1.861	Jordan	0.931
Germany	1.774	Tanzania	0.931
Spain	1.737	Egypt	0.979
Argentina	1.702	Ghana	0.988
Slovenia	1.701	Bangladesh	1.011
Uruguay	1.697	Indonesia	1.015
Finland	1.696	Rwanda	1.016
Canada	1.690	Pakistan	1.031
Switzerland	1.659	Zimbabwe	1.044
Netherlands	1.634	Thailand	1.075

References

[1] Adorno, T. W., Frenkel-Brunswik, E., Levinson, D. J., & Sanford, R. N. (1950). The authoritarian personality. New York: Harper & Brothers.

[2] Altemeyer, B. (1998). The other authoritarian personality. In M. P. Zanna (Ed.). Advances in experimental social psychology (Vol. 30, pp. 4791). New

Table 3: Value Diversity By Region

Region	Mean	Std. Dev.	Freq.
Africa	1.216	0.186	80
Americas	1.353	0.181	110
Asia	1.176	0.167	120
Europe and West	1.498	0.098	320
Mideast	1.237	0.327	50

Table 4: Value Correlations

	Value Frac.	Tolerance	Ethnic. Frac.	Religious Frac.	GDP/ Capita
Value Frac.	1.000				
Tolerance	0.405	1.000			
Ethnic Frac.	-0.322	-0.241	1.000		
Religious Frac.	0.187	0.196	0.111	1.000	
GDP/Capita	0.503	0.599	-0.359	0.198	1.000

Table 5: Basic Regression Specifications

GDP Growth	(1)	(2)	(3)	(4)	(5)
Value Diversity	-2.14***	-2.07***	-0.40	-0.24	-1.26
	(0.948)	(0.905)	(1.006)	(0.957)	(0.907)
lnGDPcap	0.17	0.10	0.16	0.14	-0.09
	(0.219)	(0.225)	(0.227)	(0.233)	(0.225)
Tele	0.00	0.00	-0.01	-0.01	0.01
	(0.013)	(0.015)	(0.015)	(0.016)	(0.015)
Urban	0.00	0.00	-0.01	-0.01	-0.01
	(0.10)	(0.010)	(0.011)	(0.010)	(0.010)
Education	0.03***	0.03***	0.02*	0.02**	0.03***
	(0.012)	(0.012)	(0.013)	(0.012)	(0.012)
Decade Fixed Effects	No	Yes	No	Yes	Yes
Region Fixed Effects	No	No	Yes	Yes	No
Africa					-2.19***
					(0.486)
Constant	3.32***	5.07***	0.31	1.78	6.45***
	(1.512)	(1.497)	(1.572)	(1.558)	(1.503)
Observations	522	552	552	552	552
R-squared	0.02	0.13	0.08	0.19	0.16
Adj. R-squared	0.011	0.11	0.07	0.17	0.14

Standard errors in parentheses
*** p<0.05, ** p<0.10, * p<0.15

Table 6: Additional Regression Specifications

GDP Growth	(6)	(7)	(8)
Value Diversity	-1.93***	1.67	-2.06***
	(0.901)	(1.211)	(0.900)
lnGDPcap	0.44**	0.38**	0.08
	(0.255)	(0.230)	(0.224)
Tele	0.03**	0.02	-0.00
	(0.018)	(0.015)	(0.015)
Urban	-0.01	-0.00	0.00
	(0.011)	(0.010)	(0.010)
Education	0.03***	0.04***	0.02***
	(0.012)	(0.012)	(0.012)
Decade Fixed Effects	Yes	Yes	Yes
ValueDiv.*GDP	-0.069***		
	(0.025)		
ValueDiv.*Tolerance		-5.32***	
		(1.169)	
Ethnic Frac.			-1.78***
			(0.674)
Constant	2.83**	2.65**	6.25***
	(1.688)	(1.564)	(1.554)
Observations	552	552	552
R-squared	0.14	0.16	0.14
Adj. R-squared	0.12	0.14	0.12

Standard errors in parentheses
*** p<0.05, ** p<0.10, * p<0.15

Table 7: Basic Specifications: By-Wave

GDP Growth	(1)	(2)	(3)	(4)
Value Diversity	-1.44	-0.23	0.25	-1.44
	(1.699)	(1.768)	(1.954)	(1.696)
lnGDPcap	-0.09	-0.58	-0.19	-0.03
	(0.439)	(0.467)	(0.444)	(0.441)
Tele	0.01	0.02	0.02	0.00
	(0.030)	(0.030)	(0.033)	(0.030)
Urban	-0.03**	-0.02	-0.05***	-0.04**
	(0.019)	(0.019)	(0.021)	(0.020)
Education	0.04	0.04*	0.03	0.03
	(0.029)	(0.029)	(0.030)	(0.030)
Region Fixed Effects	No	No	Yest	No
Wave Fixed Effects	No	Yes	No	No
Africa				-1.29
				(1.097)
Constant	7.86***	7.73***	6.47**	8.17***
	(3.136)	(3.495)	(3.462)	(3.142)
Observations	113	113	113	113
R-squared	0.10	0.18	0.15	0.11
Adj. R-squared	0.05	0.11	0.08	0.06

Standard errors in parentheses
*** p<0.05, ** p<0.10, * p<0.15

Table 8: Easterly & Levine (1997)

GDP Growth	(13)	(14)	(15)
Ethnic Frac.	-0.01**		-0.01**
	(0.007)		(0.008)
Value Diversity		0.01	0.00
		(0.011)	(0.011)
Africa	-0.01	-0.02**	-0.01
	(0.010)	(0.010)	(0.010)
Latin America	-0.02***	-0.01***	-0.02***
	(0.005)	(0.005)	(0.005)
lnGDP	0.09***	0.11***	0.09***
	(0.031)	(0.031)	(0.033)
$lnGDP^2$	-0.01***	-0.01***	-0.01***
	(0.002)	(0.002)	(0.002)
School	0.00	0.00	0.00
	(0.006)	(0.007)	(0.007)
Constant	-0.22**	-0.31***	-0.23**
	(0.124)	(0.125)	(0.131)
Observations	102	99	99
R-squared	0.71	0.69	0.70
Adj. R-squared	0.66	0.64	0.65

Standard errors in parentheses
*** p<0.05, ** p<0.10, * p<0.15

York: Academic Press.

[3] Alesina, Alberto, and Eliana La Ferrara. "Ethinic Diversity and Economic Performance." *Journal of Economic Literature* 43 (2005): 762-800.

[4] Alesina, Alberto, Arnaud Devleeschauwer, William Easterly, Sergio Kurlat, and Romain Wacziarg. "Fractionalization." *Journal of Economic Growth* 8 (2003): 155-194.

[5] Barro, Robert and Jong-Wha Lee, April 2010, "A New Data Set of Educational Attainment in the World, 1950-2010." NBER Working Paper No. 15902.

[6] Easterly, William, and Ross Levine. "Africa's Growth Tragedy: Policies and Ethic Division." *The Quarterly Journal of Economics* (Nov. 1997): 1203-1250.

[7] Fearon, James. "Ethnic and Cultural Diversity by Country." *Journal of Economic Growth* 8 (2003): 195-222.

[8] Granato, Jim, Ronald Inglehart, and David Lebland. "The Effect of Cultural Values on Economic Development: Theory, Hypotheses, and Some Emperical Tests." *American Journal of Political Science* 40 (1996): 607-631.

[9] Guiso, Luigi, Paola Sapienza, and Luigi Zingales. "People's Opium? Religion and Economic Attitudes." *Journal of Monetary Economics* 50 (2003): 225-282.

[10] Guiso, Luigi, Paola Sapienza, and Luigi Zingales. "Does Culture Affect Economic Outcomes?" *Journal of Economic Perspectives* 20 (2006): 23-48.

[11] Inglehart, Ronald, and Wayne Baker. "Modernization, Cultural Change, and the Persistence of Traditional Values." *American Sociological Review* 65 (2000): 19-51.

[12] Inglehart, Ronald. *Modernization and Postmodernization.* Princeton, NJ: Princeton University Press, 1997.

[13] Inglehart, Ronald and Marita Carballo. "Does Latin America Exist? (And Is There a Confucian Culture?): A Global Analysis of Cross-Cultural Differences." *Political Science and Politics* 30 (1997): 34-47.

[14] Inglehart, Ronald and Christian Welzel. "Political Culture and Democracy." *Comparative Politics* 36 (2003): 61-79.

[15] Jacobs, J. *The Economy of Cities.* New York: Random House, 1996.

[16] Ottaviano, Gianmarco, and Giovanni Peri. "The Economic Value of Cultural Diversity: Evidence From U.S. Cities." *Journal of Economic Geography* 6 (2005): 9-44.

[17] Ratna, Nazmun, Quentin Grafton, and Tom Kompas. "Is Diversity Bad For Economic Growth? Evidence From State-Level Data in the US." *The Journal of Socio-Economics* 38 (2009): 859-870.

[18] Sidanius, J., and Pratto, F. (1999). Social dominance: An intergroup theory of social hierarchy and oppression. New York: Cambridge University Press.

[19] Weber, Max. *The Protestant Ethic and the Spirit of Capitalism*. New York: Scribner, 1930.

[20] World Bank, World Development Indicators (WDI) December 2010 , ESDS International, (Mimas) University of Manchester.

[21] World Values Survey 1981-2008 Official Aggregate v.20090901, 2009. World Values Survey Association (www.worldvaluessurvey.org). Aggregate File Producer: ASEP/JDS, Madrid.