Nelson, Kennedy, Krueger (in press 2016) paper presented at SBP-BRiMS 2016, 28 June-1 July, Washington, DC.

Exploratory Models of Trust with Empirically-Inferred Decision Trees

John B. Nelson, William G. Kennedy, and Frank Krueger

George Mason University
{jnelso11, wkennedy, fkrueger}@gmu.edu

Abstract. What is the relationship between an individual's values and their propensity to trust other people? To explore this question, we built decision trees on the microdata provided by the World Value's Survey. Our findings confirm the extant literature while also hinting at cultural heterogeneity. We propose that studying nationally-specific decision trees based on survey data allows for easy-to-intuit representations of complex social problems. Moreover, for the sake of pragmatism, decision trees developed in this manner offer researchers a good tool in terms of cost-to-benefits.

Introduction

Social theorists have long documented the role of trust in social stability. Trust, a form of social capital, affords the opportunity for cooperation and collective action (Coleman 1990). By contrast, with distrust comes disunity, and all the maladies entailed (Putnam, Leonardi, and Nanetti 1994). This motivates a simple research question: *why are some people trusting, and others not?* The answer – accounting for the drivers of individual-level trust – allows the researcher to understand cultural-level problems.

At the psychological level, there is evidence of universal cognitive structures which relate identifiable personal values to one another (Schwartz 1992). But, the very recognition of cultural variation suggests group-level, sociological effects. To a computational social scientist, methodological individualism demands accounting for agents, with local experiences and information processing (Cioffi-Revilla 2013). The assumption is that the cultural variation emerges from the structural variations which pattern individual interactions. Recognizing this, we first explore the relationship between individual values – and, the effect of different values on trust – within a set of cultures. The larger goal is to develop an agent-based model of trust. Instead of reporting the larger (in-progress) findings, this paper documents how inferred decision trees allow for pragmatic preliminary analysis. That is, the tools of machine learning rapidly identify important features of the data and hint as to possible behavior, which informs subsequent theorizing and development.

Background

In human psychology, some needs are more important than others (Maslow 1943). According to Schwartz (1992), expressed values similarly form a hierarchy, as they are directly related to needs. Expounding, values are beliefs which motivate behavior. Contrast this to a belief that something is or is not true, but is otherwise disconnected from motive. Values drive and justify behavior whereas beliefs express general expectations.

Trust appears to be a mixture of basic values, in particular, honesty, fairness, and benevolence. Yet, there are broadly two kinds of social trust: trust in institutions, and trust in individuals. Critically, they may be antagonistic. A person who values conformity and tradition – that is, assigns this concept high importance in their value system – tends to trust institutions. Whereas, a person who values autonomy and responsibility tends to trust in individuals (Devos, Spini, and Schwartz 2002).

Scholars often rely upon surveys to characterize the relationships between expressed values. Particularly tailored to this task is the World Values Survey (WVS), which surveys individual values across time and space (World Values Survey Association 2016). At the time of writing, the six WVS waves (i.e. survey periods) span 192 countries and 34 years. These are microdata, with each record reporting the responses of a single individual. On a per-country basis, there are typically thousands of records (i.e. individuals surveyed).

Using the WVS, Tausch (2015) finds two dimensions which explain the variation in human values well. The first dimension is that of "traditional vs secular-rational" values; the second, "survival vs self-expression." Further demonstrating the prior art in modeling trust, both Jen et al. (2010) and Morselli, Spini, and Devos (2012) use the same data to find support for a link between trust and individual health, while also testing Schartz's theory.

Methodology

During preliminary investigation, there are benefits to rapid development and exploration. For reasons detailed in the later discussion, this motivated the use of decision tree learning. A variety of tree inference algorithms exist (Mitchell 1997, Ch 3). We chose the well-tested and popular scikit-learn Python package (Pedregosa et al. 2011), which implements a decision tree classifier that uses an optimized version of the CART algorithm.[1]

Succinctly, CART performs a search over attributes when assembling trees. At each stage, the attribute selected as the next node is the one that best partitions

[1] The open-source tool built and used to generate the trees is available at https://bitbucket.org/johnnybjorn/wvstreeview. It provides an IPython notebook GUI interface to the underlying WVS data.

the space of examples with respect to the target variable. Here, best means that, by performing the candidate split, the resulting sets are more "pure." As the purity criteria, the algorithm uses entropy, defined as:

$$H(X) = -\sum_{i=1}^{n} p(x_i) \log_b p(x_i)$$

where x_i is the observed frequency of an outcome over the discrete random variable, X. Intuitively, the more uniform the observed frequencies, the less pure it is (i.e. the entropy is greater). For example, <10, 10, 10> has more entropy than <20, 5, 5>. In the context selecting an attribute, the best node is the one gives the greatest information gain (i.e. Kullback–Leibler divergence):

$$\underset{\theta}{\operatorname{argmax}} = H(X) - H(X \mid \theta)$$

Note, this is a greedy heuristic. Conceivably, something like beam search – which operates over sets of nodes at a time, rather than single ones – would perform better. However, in practice, information gain is a good and performant heuristic (Mitchell 1997, Ch. 3).

Results

To show the usefulness of this methodology, four countries from Wave 6 (2010 - 2014) of the World Value Survey serve as examples: Germany, India, Morocco, and United States.[2] Figure 1 shows the macro-level variation in trust (V24).[3] It plots the proportion of people answering the question most "people can be trusted" affirmatively by country. The opposing response corresponds to "you need to be very careful."

[2] These examples are not cherry-picked. India, Germany, and the US were targets of a larger effort, and thus drove our interest. Morocco was subsequently added to better cover the cultural map for this paper.

[3] WVS variable names documented in parentheses.

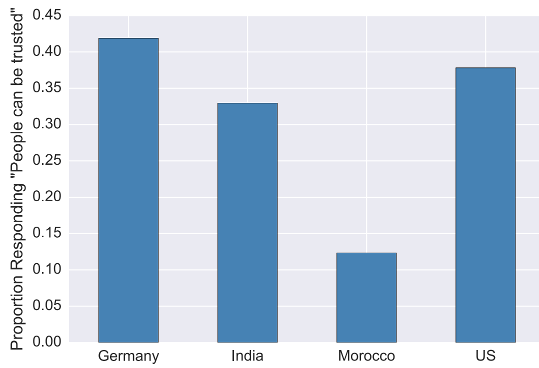

Fig. 1. Trust in Others By Example Country

The cultural map of the world (Inglehart and Welzel 2005) inspired this selection of countries, an illustrative coverage over the reported abstract regions. Said differently, the examples draw from different regions in Tausch's (2015) two-dimensional space. The trees explained in the examples below analyize the with-incountry variance over the same question.

WVS Wave 6, Germany

Figure 2 portrays the decision tree emitted for Germany over the World Value Survey's 6th wave. Below each node, the plot describes the proportion of people answering that they trust others. Next to this number is n, the number of respondents which exist at this point in the tree traversal.

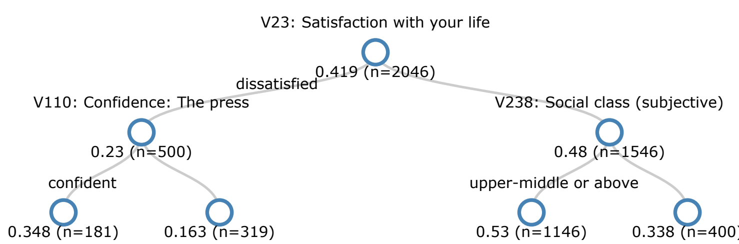

Fig. 2. Germany Decision Tree

This decision tree has an ROC/AUC score of 0.64. ROC/AUC is the Area Under the Curve (AUC) for the Receiver Operating Characteristic (ROC). The ROC plots the true positive rate as a function of the false positive rate. A classifier that performs no better than random would have a AUC/ROC of 0.5, illustrated by a diagonal line. More intuitively, assume you applied your classifier to one random negative and one random positive example. The AUC/ROC measures the probability that the positive example ranks higher than the negative one, according to some ranking function.[4]

Here, the root node partitions the space with respect to "satisfaction with your life" (V23). Respondents answering 'dissatisfied' fall to the left; those satisfied fall to the right.[5] Those satisfied are twice as likely to answer, "most people can be trusted" (i.e. 0.48 : 0.23)

The left child of the root then splits the set over confidence in the press (V110). Of the set remaining, those who are confident in the press are about two times more likely to express trust in most people than those who are not. The right child of the root splits the set over subjective social class, with {Upper, Upper-Middle} branching on one side and {Lower, Working, Lower-Middle} on the other. Those belonging to a higher social class are approximately 50-percent more likely to be trusting than those from the lower classes.

WVS Wave 6, India

For India, the inferred decision tree (Figure 3), has an ROC/AUC of 0.661. The root of the tree partitions responses by respect for immigrants. Variable V46 prompts, "when jobs are scarce, employers should give priority to people of this country over immigrants." Those who disagree (branching to the right) are more trusting. They are a bit more than twice as likely to answer "most people can be trusted" than those who either agree or don't respond to the original prompt.

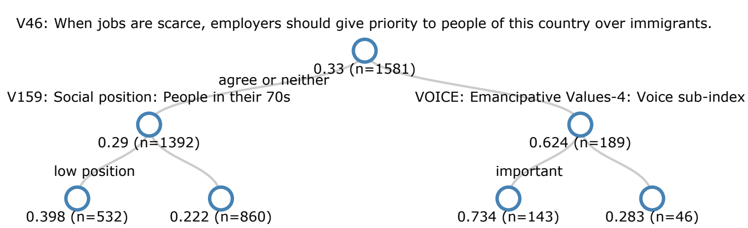

Fig. 3. India Decision Tree

[4] This assumes a function with a real-valued codomain. A decision tree produces a binary response. However, mapping the proportion of 'yes' to 'no' responses in each node allows for computation of AUC/ROC.

[5] Actually, respondents answer on a scale of 1-10, dissatisfied-satisfied. In this case, the decision tree implicitly dichotomized the space according to $x \leq 5$. This is a safe operation over the ordinal values, which are common in WVS.

Traversing the tree down to the left-child of the root, a prompt captures a mixture of respect for authority and traditional values in asking, what is the social position of people in their 70's (V159). Those who believe this cohort has a low position in society are a bit less than twice as likely to trust people than those who lack such respect. On the right-branching child of the root, a compound index which captures the degree to which freedom of expression matters partitions the space (VOICE). Those who value freedom of expression are approximately 2.5 times more likely to trust other people than those who do not.

WVS Wave 6, Morocco

The decision tree for Morocco (Figure 4) achieved an ROC/AUC of 0.658. The root of the tree partitions by the prompt, "how much respect is there for individual human rights nowadays in this country" (V142). Those who believe there is a great deal of respect are about five times more likely to trust other people than everyone else.

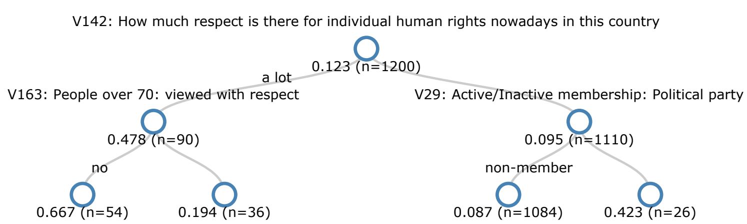

Fig. 4. Morocco Decision Tree

Following the leftward path – those who believe there is a great deal of respect – the child node then decides on the basis of "people over 70 are viewed with respect" (V163). Respondents answering less affirmatively are approximately 3.5 times more likely to trust than those who answer "very likely to be viewed that way." On the right branch of the root, political party membership (V29) discriminates. Those who do not engage in politics by means of party membership are less trusting than those who at least are party members, even if they are inactive.

WVS Wave 6, US

The tree for the United States achieved an AUC/ROC of 0.667. The root node (Figure 5) discriminates based on the respondents trust or distrust of secular institutions (SKEPTICISM). Those who mostly trust these institutions – the courts, the army, and the police – are twice as likely to trust individuals than those who do not.

For those that do trust secular institutions (left branch), neighborhood security (V170) is the next discriminating factor. Those who feel very secure in their neighborhoods are roughly 50% more likely to trust individuals than everyone else. On the right child of the root, satisfaction with the household's financial situation partitions the space (V59). Those who satisfied are approximately twice as likely to trust compared to those dissatisfied.

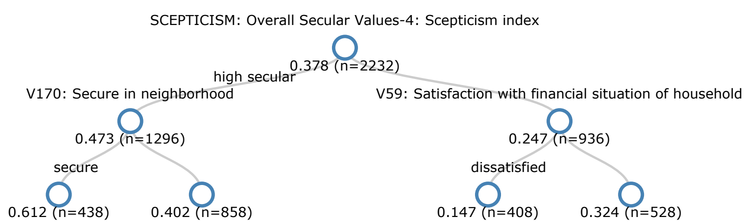

Fig. 5. US Decision Tree

Discussion

Stylized Facts

The emitted decision trees comport to the expectations of the extant literature. Categorizing the node according to Tausch (2015)'s space is subjective, but does not stretch credulity. Looking at the decision trees, many stylized facts are apparent. The following examples are a small sample of those available.

The root node for the US tree – that is, the most discriminatory node, according to information gain – splits based on whether the respondent trusts in secular institutions. This aligns well with the idea that, among Western nations, America is markedly religious (Pew Research Center 2011). Those who are religious (traditional) tend to devalue the secular. As theorized in the literature, institutional trust and social trust are different things. But, in the US, they are more closely correlated than in other countries – so much so that it was not a selected variable in any of the other trees.[6]

The root node for Morocco is almost counter-intuitive. As per the world culture map, Moroccans emphasize traditional and survival values. Yet, the tree learner

[6] We dropped highly-correlated trust variables during the data cleaning phase. Partially, this required reading the code book and striking variables that interrogated trust with an alternate wording. But, as a backup, the software calculated correlations. We retained the SCEPTICISM variable after observing a correlation lower than expected.

selected a question that concerns respect for individual rights. However, given the type of data – microdata survey – this makes sense. If you are in the minority of people who value individual rights, such a value must especially important. Therefore, given the hierarchal structure of values, your propensity to trust is likely to be high. In the context of the decision tree, this split is likely to result in relatively pure subsets.

The case of India is particularly interesting. The root node splits those who favor their in-group (i.e. not immigrants) when economic opportunity is scarce. This may obliquely capture elements of xenophobia, but it should also reflect a high-value on traditionalism. Given the realities of a country with rich ethnic/racial heterogeneity and a high population growth rate, this seems to be a critical survey question. And, the decision tree learner uncovered the importance immediately.

Finally, the case of Germany provides evidence for frequent assertion: in general, the more satisfied you are, the greater your propensity to trust. The between-country comparisons for trust reflect the plausibility of the same assertion – richer countries are more trusting. From a classic hierarchy of needs perspective this makes sense. Higher level needs such as esteem and self-actualization rise to the top, given met lower-level needs (Maslow 1943). In service of these needs, the individual places a heavy weight on autonomy, which emphasizes the role of trust.

Decision Trees and Exploratory Analysis

The results of this exploration adds to the evidence in favor of extant theories on trust and values. Yet, the design is not especially novel. The literature is replete with studies that use the same data to arrive at similar conclusions by means of different methods. Consequently, the reported results do not add much weight to the scales of evidence. However, the use of decision tree learning as a exploratory tool proved especially pragmatic.

The algorithms for inferring decision trees are fast and robust. As a consequence, we were able to design an IPython (Pérez and Granger 2007) interface that ran analyses in real-time, given modeler input. This allowed for on-demand country/wave-specific runs. In the context of trust, this is important for discovery. Between-country differences in trust dwarf within-country ones (Inglehart and Welzel 2010). Said differently, between-country differences explain more variance. But, these drivers are more sociological. Assuming an interest in psychological drivers, the within-country factors are more relevant.

Ignoring performant computation, the emitted trees are easy to interpret. This matters for the modeler, for whom time binds inelastically. Compare the visual portrayal of a decision tree to the tabular output of something like a logistic regression. The latter requires effort interpreting explained variance to deduce relative importance of each factor. In the decision tree, the nodes location portrays this information quickly. And, this ignores the costs of the more expensive operation – pruning the factors included in the regression to combat the problem of collinearity. The decision tree is self-pruning by the nature of the inference algorithm. Given collinearity, the selected node makes the collinear one unimportant from an information gain perspective. It simply drops out of the analysis, without any numerical issues.

Of course, an experienced researcher may find cause for concern in this methodology. It is positivist to the point of being a-theoretical. And, given the opportunity of point-and-click analysis, it is easy to find results that are seem novel but are truly arbitrary. This is true, and it warrants caution. But, as a tool for preliminary analysis, it strikes a good balance. In any early-stage research project, there is a period of pure exploration. The researcher becomes more familiar with the data by forming expectations and testing them. When done in a perfunctory manner, this process is prone to error. Consequently, cursory tests are less decisive than they could be, so clarity emerges slowly. Given the robustness and obviousness of decision trees, such tests over expectations are more robust. And, assuming commensurate effort or time, it should leave the researcher with more a refined understanding of the target system.

Conclusion

Research is messy; fraught with dead ends; and, riddled with bad paths. Decision trees allow for pragmatic exploration. As familiarity with the problem domain increases, such a tool may cease to be useful. But, during the early stages, it allows for parsimonious explanations while favoring large effects. That is stable ground upon which to build.

We used a decision tree learner to explore the question of how individuals' values affect their propensity to trust. From this exploration, we validated the extant literature. That is, using a different methodology, we generated evidence in agreement with the prior theory. Moreover, the trees made country-specific factors immediately salient. Perhaps, these factors would not surprise a cultural anthropologist with domain expertise on the specific country. But, they are easy to miss when generalizing over all countries at once, in deference to methodological convenience. With decision trees, the data acts as a tour guide, allowing for a more nuanced perspective, even in absence of domain expertise. Consequently, they grant the modeler a more intimate understanding of the problem.

Bibliography

Cioffi-Revilla, Claudio. 2013. *Introduction to Computational Social Science: Principles and Applications*. Springer Science & Business Media.

Coleman, James S. 1990. *Foundations of Social Theory*. Cambridge, MA: Belknap.

Devos, Thierry, Dario Spini, and Shalom H Schwartz. 2002. “Conflicts Among

Human Values and Trust in Institutions." *British Journal of Social Psychology* 41 (4). Wiley Online Library: 481–94.

Inglehart, Ronald, and Christian Welzel. 2005. *Modernization, Cultural Change, and Democracy: The Human Development Sequence*. Cambridge University Press.

———. 2010. "Changing Mass Priorities: The Link Between Modernization and Democracy." *Perspectives on Politics* 8 (02). Cambridge Univ Press: 551–67.

Jen, Min Hua, Erik R Sund, Ron Johnston, and Kelvyn Jones. 2010. "Trustful Societies, Trustful Individuals, and Health: An Analysis of Self-Rated Health and Social Trust Using the World Value Survey." *Health & Place* 16 (5). Elsevier: 1022–9.

Maslow, Abraham Harold. 1943. "A Theory of Human Motivation." *Psychological Review* 50 (4). American Psychological Association: 370.

Mitchell, Tom M. 1997. *Machine Learning*. Burr Ridge, IL: McGraw Hill.

Morselli, Davide, Dario Spini, and Thierry Devos. 2012. "Human Values and Trust in Institutions Across Countries: A Multilevel Test of Schwartz's Hypothesis of Structural Equivalence." In *Survey Research Methods*, 6:49–60. 1.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, et al. 2011. "Scikit-Learn: Machine Learning in Python." *Journal of Machine Learning Research* 12: 2825–30.

Pew Research Center. 2011. "The American-Western European Values Gap." Report. *Pew Global Attitudes Project*.

Pérez, Fernando, and Brian E. Granger. 2007. "IPython: A System for Interactive Scientific Computing." *Computing in Science and Engineering* 9 (3). IEEE Computer Society: 21–29.

Putnam, Robert D, Robert Leonardi, and Raffaella Y Nanetti. 1994. *Making Democracy Work: Civic Traditions in Modern Italy*. Princeton university press.

Schwartz, Shalom H. 1992. "Universals in the Content and Structure of Values: Theoretical Advances and Empirical Tests in 20 Countries." *Advances in Experimental Social Psychology* 25 (1). Orlando, FL: 1–65.

Tausch, Arno. 2015. "Towards New Maps of Global Human Values, Based on World Values Survey (6) Data." Based on World Values Survey (6) Data (March 31, 2015). Available at SSRN: http://ssrn.com/abstract=2587626.

World Values Survey Association. 2016. "WORLD VALUES SURVEY Wave 6 2010-2014 OFFICIAL AGGREGATE V.20150418." Aggregate File Producer: Asep/JDS, Madrid Spain.