

Measuring International Socio-Religious Values and Conflict by Coding U.S. State Department Reports

Brian J. Grim¹, Roger Finke², Jaime Harris², Catherine Meyers², Julie VanEerden²
 Pew Forum on Religion & Public Life, and Penn State ARDA¹
 Association of Religion Data Archives (ARDA), Pennsylvania State University²

Abstract

Values, culture and religion are sources of cohesion and conflict. Though some international survey data exist on socio-religious values such as the World Values Survey and ISSP, these are limited in the number of countries covered and at times in the types of sensitive questions that can be asked in the very countries where such questions are of most interest. We overcome these limitations by providing coded measures on socio-religious values and conflict for 195 different countries. Coding of events, transcripts, and reports is used in many disciplines to gauge public values and opinions. Using a “survey questionnaire-type” codebook, we coded up to 250 different variables for 195 countries based on data in the U.S. State Department’s International Religious Freedom Reports (www.state.gov/g/drl/rls/irf/). We provide a detailed account of our coding of the 2001 and 2003 International Religious Freedom Reports.

Keywords: Religion, International, Conflict, Coding, Violence

1. Introduction

One of the chief limitations researchers face when attempting to study religion’s relation to conflict is the lack of empirical measures. Data sets which are available, such as the Minorities at Risk Project, have limited measures that specifically focus on religion. Other cross-national studies, such as the World Values Study which has multiple measures on religion, have few data that specifically relate to religion and conflict. This paper describes our work to address this deficit of cross-national measures on socio-religious conflict.

2. State Department Reports as Data

Our measures are coded from the U.S. State Department’s International Religious Freedom Reports (hereinafter referred to as the “Reports”). These Reports are one of the most comprehensive global treatments of conflict related to religion. Though the Report’s title rightly indicates a focus on religious freedom, the Reports detail many issues related to

religion including situations where religious actors are either perpetrators or victims of conflict and/or violence.

One important concern in using these Reports is whether they are seriously biased due to being authored by the U.S. Government. In this paper we will directly address this concern by comparing the results of our coding with the results of Jonathan Fox and Shmuel Sandler (2003, 2004), who conducted a similar coding project, but used data sources beyond the International Religious Freedom Reports. Our data cover much more than the Fox and Sandler data in that we have measures that specifically focus on socio-religious violence and on the social regulation of religion.

2.1 International Religious Freedom Reports

In fulfillment of U.S. law, each U.S. Embassy prepares an annual Report on the state of religious freedom in their host country. Reporting adheres to a common set of guidelines and training is given to Embassy personnel who investigate the situation and prepare the Reports (see U.S. State Department 2001-2005). Once an Embassy completes a Report, it is then reviewed by various State Department offices with expertise in the affairs of that country and in human rights. The Reports incorporate information from other human rights reports. The U.S. Commission on Religious Freedom also assists the State Department conduct research that feeds into these reports. They are then arranged and vetted under the supervision of the special U.S. Ambassador for International Religious Freedom. They cover the following standard reporting fields for each country: religious demography, legal/policy issues, restrictions of religious freedom, abuses of religious freedom, forced conversions, improvements in respect for religious freedom, and the US Government’s actions.

The Reports are a loosely structured, retrospective, qualitative survey of most countries of the world. The U.S. State Department has been compiling such annual Reports since 1997. In 2001, they took on the reporting format shown in Figure 1.

Figure 1: Report Format (for each country)

<p><i>Introductory Overview</i> [untitled section]</p> <ol style="list-style-type: none"> 1. Religious Demography 2. Status of Religious Freedom <ol style="list-style-type: none"> a. Legal/Policy Framework b. Restrictions on Religious Freedom c. Abuses of Religious Freedom d. Forced Religious Conversion e. Persecution by Terrorist Organizations f. Improvements in Religious Freedom 3. Societal Attitudes 4. U.S. Government Policy
--

Though the Reports are bounded (July 1 to June 30, annually), they include retrospective information on events that have been systematically monitored since 1997. Therefore, the data in these Reports approximate a trend study, which captures both recurring problems and specific problems that occurred during the reporting period. One important note is that with each new Reporting year, the Reports benefit from the information gathered and add to that information, self-correcting errors or inserting data missing from previous Reports. Thus, it is reasonable to expect that the 2003 Reports will be less error-prone than, for example, the 2001 Reports.

The Reports are primarily produced by Embassy officials in country capitals and other cities with US Consulates, which limits their scope and may be a potential source of error. The Reports may also be biased by the groups with the loudest national voice. However, these problems are attenuated by the practice of the Reports to incorporate multiple sources of information. The varying length of country Reports should be noted. For example, the 2003 Report for Indonesia contains 12,632 words while the Report for St. Kitts and Nevis contains only 405 words. Rather than view the shorter Reports as a problem of *missing data*, the assumption in this study is that if abuses or restrictions were not reported, then they were negligible or nonexistent.¹ While this assumption may be a source of error in the case of unreported or undiscovered problems, there is reason to believe that

¹ During coding, missing data were coded as system-missing, and beginning with our 2005 coding, “not reported” will be a separate variable for each question in our codebook. In the final data set, system-missing was recoded as whatever value represented the absence of a problem.

any unreported abuses were negligible due to the way the Reports are constructed (as discussed above). The notable exceptions to this approach to missing data are the data coded from the Reports on North Korea, Libya, and Bhutan, where the U.S. State Department did not have an official presence during various reporting periods. In the case of North Korea, the scores calculated for that country cannot be considered reliable due to a lack of verifiable data from any international source. The situation may be worse (or better) in North Korea than the limited information reveals. In spite of such limitations, these Reports are the most comprehensive summary available of the religious conflict in 195 ‘countries’ of the world representing nearly 95 percent of the world’s population. The one notable exception is the United States due to the limitation that the State Department does not report on the situation within the United States.

2.2 Quantitative Coding of the Reports

We employed strict training and coding protocols to ensure that our data quantified from the International Religious Freedom Reports were as objective and reproducible as possible. Coders for this project included senior undergraduate and graduate students. Under the direction of this paper’s lead author, the coders underwent a thorough week-long introduction to the purpose of the project, the data source, and its collection methods. After becoming familiarized with the project, coders were given a copy of the initial coding instrument, which was previously used to code the 2003 Reports, along with a country Report to be coded.² After each person on the team independently coded the Report, each coder described to the group his or her own decision-making processes in arriving at a particular score. Areas of coder discrepancy were discussed at length with no time limit to the discussions. The discrepancies were assessed and reconciled using a combination of methods (Montgomery and Crittendon 1977) including insights from cognitive interviewing (Presser and Blair 1994) and the “Think-Aloud” strategy (e.g., Willis et al. 1999:3). The majority of the instrument is developed from the verbatim responses in the International Religious Freedom Reports, and most coder variation could be solved using the *a posteriori* method of majority rule (Montgomery and Crittendon 1977). Coders presented evidence from the text to justify the logic and detail of their assignment choice, and then a consensus was reached on the correct coding of the data. This method greatly improved the inter-coder

² 2002 Reports were used for training purposes so as not to using training as part of the actual coding.

reliability and reduced marginal and single-cell disagreement (Funkhouser, Ray and Parker 1968) in several practice runs. Some assessments required the coder to assign a code based on a judgment of the entire report rather than on specific mentions or verbatim responses (Type B-2 coding in Montgomery and Crittendon's 1977 terms). In these cases, coders had to utilize their memory of all incidents in the report. Disagreements were subject to thorough discussion and justification until the disagreement was resolved. Using these same procedures, four more country Reports of varying length and complexity were coded and reconciled as part of the training.

During the training process, the instrument was simultaneously being tested for defects. Non-exclusivity of categories was the most common problem in the codification process. This was generally revealed in cases where disagreement over codes could be settled by direct examination of the Reports. Categories that were not mutually exclusive or exhaustive were then modified to achieve exclusivity and objective reproducibility (Crittendon and Hill 1971). This method ensured that in the final coding process, maximum reliability could be achieved without being hampered by unsuitable code categories.

After finalizing the coding instrument and when coder training was completed, the coders began the process of coding the 2001 International Religious Freedom Reports for 195 countries and regions around the world, with the exclusion of the United States of America. Each country's Report was double-blind coded by two coders. Throughout the coding procedure, the instrument was continually monitored for possible defects as new sets of countries were processed. Each coding pair coded for ten hours each week and then met weekly for one hour with the lead author to go over any items they saw as problematic. This means that for every ten hours of coding, one hour was devoted to addressing any further coding issues or questions. This was especially useful as different regions of the world were coded since some regions have particular characteristics, e.g., Shari'a Law does not frequently come up as an issue in the Reports for countries of the Western Hemisphere. For the sake of instrument continuity, only major categorical oversights were grounds for instrument revision, and, then only in such a manner as to not offset all previous coding.

The last step in the process was to reconcile final differences between coders for each country. Printouts of the variables where a discrepancy existed were provided to the coding teams for each country. Under the supervision of the lead author, the coders then went

back to each country's Report to decide on the best code for each disputed item. Most of the discrepancies were differences between levels on a scale rather than discrepancies on whether a particular item was present or absent. This last step provides a single set of codes for each country.

3. Reliability, Error/Bias, Validity

We will now turn to an assessment of our data. Specifically, we will (a) analyze the *reliability* of several key coded measures, (b) look at some indicators that can evaluate whether *error and bias* in the State Department Reports adversely affect these measures, and (c) briefly consider whether there is some evidence that these measures *validly* measure the constructs they attempt to represent. The key coded measures we will analyze for reliability are in two groups.

3.1 Reliability

3.1.1 Grim/Finke Indexes

The three Grim/Finke indexes (Grim and Finke 2006) are the Government Regulation of Religion Index (GRI), the Government Favoritism of Religion Index (GFI), and the Social Regulation of Religion Index (SRI) that they calculated using coded data from the 2003 International Religious Freedom Reports. We have calculated these same indexes using our coding of the 2001 International Religious Freedom Reports.

The GRI operationalizes the concept of *government regulation of religion*, which Grim and Finke define as "the restrictions placed on the practice, profession, or selection of religion by the official laws, policies, or administrative actions of the state" (2006a:7). The GRI index, which is composed of six variables, was coded with very good overall reliability (.892). These six measures were each coded very reliably, with inter-rater correlations at .842 and above. Such high reliability can partly be attributed to the standardized nature of the language used in the Reports. For example, each Report uses the term "generally respects" decidedly (see US State Department 2001-2005), and the coders can key in on such language.

The second Index, Government Favoritism of Religion (GFI), is composed of five variables, the fifth being a funding index constructed from six additional variables. The variables operationalize the concept of *government favoritism of religion*, which Grim and Finke define as "subsidies, privileges, support, or favorable sanctions provided by the state to a select religion or a small group of religions" (2006a:7-8).

The GFI was coded with a good overall reliability of .836. The first four variables were each coded very reliably, all with inter-rater correlations at .853 and above. The Government Funding of Religion Index was more difficult to code reliably (.651), primarily because the coders had to make inferences based on less standardized reporting language. This is especially the case where, for example, a Report may state that there is an official religion and that religion is taught in the public schools, but then not report on all the types of funding that such a relationship brings. It might be reasonable to infer that government-supported public teaching of religion implies all of the other types of funding as well. The Reports, however, do not catalogue each type of funding, so the lower reliability stems from coders making different ratings based on an inference or leaving the items un-coded.

The Social Regulation of Religion Index (SRI) provides a measure of the social regulation that moves beyond government actions. Grim and Finke define *social regulation* as “the restrictions placed on the practice, profession, or selection of religion by other religious groups, associations, or the culture at large” (2006a:8). The variables related to the SRI were coded with a passable overall reliability of .762. The first two variables were each coded with good reliability (.847 and above), and the remaining three variables were coded with acceptable reliability, though the variable measuring whether established religions try to shut out new religions was coded with a inter-rater reliability of .660, which is towards the lower end of acceptability. This item was more difficult to code because the Reports do not specifically use this language, and thus the coders must make inferences based on the situations that are reported. We should note, however, that recognizing a problem is the pathway to a solution. Having such reliability statistics allows us to focus on strategies to improve the coding of low reliability items on the next round of coding (i.e., for the 2005 Reports).

3.1.2 Socio-religious Violence and Conflict Measures

The second set of measures initially developed by Grim (2004, 2005) includes three different indicators of socio-religious violence. The first measure is of the number of human beings adversely affected by socio-religious violence and conflict. The second scale is of the highest type or level of socio-religious violence and conflict. The third measure is of the extent of the socio-religious violence and conflict. These are unique measures that focus on a specific form of social conflict—conflict that specifically involves religion and/or religious adherents.

These measures were coded with a consistently high degree of reliability, owing to the fact that the Reports give priority attention to documenting specific incidents of violence and abuse that occur due to a person’s religious affiliation. A strength of the Reports, and thus a benefit to the coding endeavor, is that the Reports avoid making unsubstantiated allegations about the abuse or displacement of persons due to religion. The Reports are in fact careful to make a clear distinction between whether a particular violent incident was simply, for example, abuse that happened upon a religious person, or whether the abuse was committed *because* of the person’s religious identity.

Table 1: Coding Reliability of Conflict Variables

Inter-rater Reliability Analysis	
Alpha	Socio-religious Conflict & Violence Variables
0.908	Considering the entire Report, estimate the following: The number of people who were physically abused or displaced due to religion in this country is: 0=none 1= < 10 2= 10 - 200 3= 201-1000 4= 1001 - 10,000 5= > 10,000
0.916	What is the highest level of religious violence reported? 0=none 1=anti-religious brand graffiti 2=vandalism to religious brand property 3=bombing or burning or religious brand property 4=beating, rape or physical assault of person(s) due to religious brand 5=torture or killing of person(s) due to religious brand 6=massacre of and/or war between religious brands
0.887	To what extent is there religiously-related violence in the nation (where religion is the victim and/or perpetrator)? 0=Absent 1=Isolated acts of religiously related violence 2=Widespread acts or covering several regions with religiously related violence 3=Ongoing war with religiously related violence

3.2 Error and Bias

We will not attempt to measure the error in the Reports directly; nevertheless, we can get an idea whether the level of error is decreasing over time by comparing how the data we have coded from the two different Report years correlate with other measures. As mentioned above, the Reports are produced annually, and information that was omitted from (or in error in) one year’s country Report can be—and often is—corrected in subsequent years. For example, the omission of a long-standing missionary presence in the UAE in the 2001 Report was corrected in the UAE’s 2003 Report. If this self-correcting process is the case, then it is reasonable to expect that measures from 2001 will be less strongly correlated with external data with which they should theoretically be correlated than the 2003 measures.

The first two columns of correlations in Table 5 show just this (see Appendix for Table 5). For simplicity of presentation, we combined the three main indexes described above (GRI, GFI, and SRI) into a composite

measure by adding their weighted totals. Both composite measures (2001 and 2003) correlate strongly and significantly with other measures of low freedom, such as Freedom House's religious freedom scale (.755 and .794), their civil liberty scale (.663 and .684), as well as Heritage Foundation's economic freedom scale (.585 and .607). All three of these scales have low freedom having a higher scale value, so the correlations are in the expected directions. The important point to note is that *in every case*, the 2003 Grim/Finke Composite Index³ is more strongly correlated with these measures. The same holds true for the remaining comparison variables in Table 2 (in Appendix). This provides rudimentary evidence that the Reports have less error as the years go by. Certainly as readers of the Reports, we find this to be the case.

[See Table 2 at end.]

Aside from the problem of error in the Reports contaminating the coded data, we must look at whether our results are biased because they rely on Reports compiled by the U.S. Government, albeit drawing on multiple sources themselves. A plausible way to test whether there is a bias that affects the data is by comparing the two Grim/Finke composite indexes with a composite measure created by Fox and Sandler (2003, 2005).

Fox and Sandler created a composite coded measure on the separation of religion and state. This measure is extremely useful in our analysis because their coding project was designed and carried out in a very similar way as our study, *except* that they coded their data from multiple sources. They did not rely exclusively on the State Department Reports, as we did. Fox and Sandler's composite variable [all2002] scores the lack of separation (akin to what we call regulation and favoritism) high. So, we should expect that the relationships will all go in the same direction. Also,

³ To create this weighted composite index, we treated three Indexes as being predicted by a common factor (e.g., religious freedom). Using a confirmatory factor analysis structural equation model, each index was weighted by the regression weight predicted by the model (all regression weights significant at $p < .001$, two-tailed). For 2001: SRI=1.000 (designated constant), GRI=1.117, GFI=.734; for 2003: SRI=1.000, GRI=.836, GFI=.570; for a Composite of 2001-03: SRI=1.000, GRI=.977, GFI=.627. This combination also makes the measures more closely comparable to the Fox/Sandler measure that is described below. See Grim and Finke (2006) for a detailed explanation of the index calculations.

though the variables are measuring slightly different concepts, they are closely related and should behave very similarly when compare to other theoretically related data.

The Fox/Sandler variable is in the third correlation column of Table 5, and comparing it with the first two Grim/Finke columns shows our expectation to be exactly the case. In fact, across a whole range of questions from the World Values Survey,⁴ the direction and strength of the correlations are consistent. For example, the mean response of those surveyed about whether "children should learn tolerance and respect" is negatively correlated with religious regulation, which can be considered a form of religious intolerance. The Grim/Finke composite indexes correlated in the same direction as did the Fox/Sandler composite index (-.389 and -.433 versus -.321). Also as expected, the mean response of those surveyed about whether "children should learn religious faith" is positively correlated with religious regulation, which is logical since those who feel that religion should be taught would plausibly expect that government and society would assist in enforcing that value. Again, the Grim/Finke indexes work in the same positive direction as the Fox/Sandler index (.304 and .369 versus .378). The remaining variables offer no surprises: religious beliefs and traditional values are positively associated with religious regulation and/or state involvement in religion. Likewise, low importance of religion in life is negatively correlated with both Grim/Finke indexes (-.295 and -.372) and the Fox/Sandler index (-.314). The fact that the Grim/Finke and Fox/Sandler correlations run in the same direction and are fairly similar in magnitude gives some evidence that the State Department data are not biased. If they are biased, it is a shared bias between Grim/Finke and Fox/Sandler across all the measures considered.

3.3 Validity

The question of validity is simply: Do the measures accurately reflect the concepts being measured? To this question, we can confidently say that our measures correlate significantly, consistently and predictably with other related measures. Specifically, as shown in Table 5, our 2001 and 2003 composite indexes correlate significantly, strongly and logically with related religious regulation measures (with Freedom House at .755 and .794), other freedom measures, and with related measures from the World Values Survey.

⁴ We aggregated by country the weighted mean responses to questions from Waves 3 and 4 of the World Values Survey.

Also, the Grim/Finke composite index for 2001-2003 correlates with the Fox/Sandler composite index at .817. Such correlations are indications of validity, though not proofs.

To dig a bit deeper, we can look at our three socio-religious violence and conflict measures. The last three columns of Table 5 show the correlation of these three variables with the same external measures. These continue to behave in the expected directions, with low freedoms being associated with more violence, while liberal attitudes are associated with less religion-specific violence. For example, the mean response of those surveyed who expressed a “low importance of religion in life” is negatively associated with all three violence scales (-.358, -.320, -.356). Those with higher secular values are also negatively associated with the three socio-religious violence scales (-.468, -.450, and -.504). Conversely, the higher level of religious values and beliefs, the greater the association with *religious violence*. For example, God being important in life, belief that churches give answers to moral problems, and belief in hell are all positively and significantly correlated with all three measures of socio-religious violence and conflict. It is especially important to note that the correlation is with religious violence, not crime or other forms of violence. The critical point is that such correlations would be expected if our measures are actually valid measures of *socio-religious conflict and violence*. As mentioned above, these are unique measures that focus on conflict that specifically involves religion and/or religious adherents. The proposition is that the conflict being measured is intertwined with religion. The data support that this proposition is valid.

As other direct indicators of socio-religious conflict and violence become available to us, we plan to further test the validity of our measures.

4. Conclusions

Cross-national data on religion and socio-religious conflict are deficient. One way to obtain data in a relatively cost- and time-effective way is to code existing government reports, such as the State Department’s International Religious Freedom Reports. This paper summarizes one such project. We were able to code 250 fairly sophisticated measures with a high degree of inter-rater reliability for 195 countries in less than nine months. The resulting data not only reliably make quantitative summaries of the Reports, but our data also relate to other external data sources in ways that suggest that they are valid indicators of the constructs underlying the measures.

We have used the *Reports* due to advantages they offer above other sources of data. First, the State Department officials who compile the Reports rely on a wide array of sources: Embassy personnel and other government and State Department employees, the U.S. Commission on Religious Freedom, journalists, human rights organizations, religious groups, local government, and academics. Rather than relying entirely on local government reports or scattered journalist accounts, the reports are a systematic collection of information, taken from diverse sources, reported in a standardized format, and completed for 195 countries. Second, embassies follow a standardized format with similar information included for each country. Third, embassy representatives receive training in completing the forms and gathering the information. Fourth, the representatives assembling the data live in the country but are not representatives of the local government or long-term residents, giving a positive balance between *nearness* and *remoteness*. Fifth, these Reports are vetted by various State Department offices with expertise in regional affairs and human rights. Sixth, the Reports are getting richer over time. And most importantly, the State Department Reports allow coding of social regulation of religion and socio-religious conflict and violence—these are important measures that are absent in the Fox/Sandler data.

While we are not making any substantive or predictive claims in this paper, the types of measures we have developed beg for further analysis. How do different religions and religious regulation relate to violence? Does the relationship change according to the type of violence? Do some types of religious beliefs lead to more violence? Is it true that people fight for what they hold dear and believe in? Why are some societies plagued with socio-religious violence (e.g., Iraq and Israel), while others quite nearby are generally peaceful (e.g., the UAE and Kazakhstan)? With better data, such questions can be empirically explored.

Acknowledgements

Supported by a grant from the John Templeton Foundation. The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the John Templeton Foundation. The authors gratefully acknowledge the valuable research assistance of Dr. Gail Johnson-Ulmer, Daniel McKenrick, Melissa Grim, and Laura Tach. The 2005 reports are currently being coded. These data will be freely available at the Association of Religion Data Archives (www.TheARDA.com) housed at Penn State University.

References

- Crittendon, Kathleen S. and Richard J. Hill. 1971. "Coding Reliability and Validity of Interview Data." *American Sociological Review* 36: 1073-1080.
- Fox, Jonathan, and Shmuel Sandler. 2003. "Quantifying Religion: Toward Building More Effective Ways of Measuring Religious Influence on State-Level Behavior." *Journal of Church and State* 45:559-588.
- _____. 2005. "Separation of Religion and State in the 21st Century: Comparing the Middle East and Western Democracies." *Comparative Politics* 37:317-335.
- Funkhouser, G. Ray and Edwin B. Parker. 1968. "Analyzing Coding Reliability: The Random Systematic-Error Coefficient." *The Public Opinion Quarterly* 32:122-128.
- Grim, Brian J. 2004. "Behavioral Measures for Religious Violence." Pew Forum on Religion & Public Life, invited lecture for the U.S. State Department. November 17, 2004.
- _____. 2005. "Religious Regulation's Impact on Religious Persecution: The Effects of De Facto and De Jure Religious Regulation." Ph.D. Dissertation in Sociology. University Park, PA: The Pennsylvania State University.
- Grim, Brian J. and Roger Finke. 2006. "International Religion Indexes: Government Regulation, Government Favoritism, and Social Regulation of Religion." *Interdisciplinary Journal of Research on Religion* 2 (Article 1). www.ReligJournal.com
- Montgomery, Andrew, and Kathleen S. Crittendon. 1977. "Improving Coding Reliability for Open-ended Questions." *The Public Opinion Quarterly* 41:235-243.
- Presser, Stanley and Johnny Blair. 1994. "Survey Pretesting: Do Different Methods Produce Different Results?" *Sociological Methodology* 24:73-104.
- U.S. State Department. 2001-2005. *International Religious Freedom Report*, online, accessed from <http://www.state.gov/g/drl/rls/irf/>.
- Willis, Gordon B. et al. 1999. "Cognitive Interviewing: A 'How to' Guide." Research Triangle Institute. Accessed 10-2003. <http://appliedresearch.cancer.gov/areas/cognitive/interview.pdf>
- Freedom House. Data downloaded November 2005. <http://www.freedomhouse.org/>
- Grim, Brian J. and Roger Finke. Data from www.TheARDA.com and from the working data set described in this paper.
- Heritage Foundation. Data downloaded November 2005. <http://www.heritage.org/research/features/index/>
- United Nations. Human Development Report data downloaded November 2005. <http://hdr.undp.org/statistics/>
- U.S. State Department. 2001-2005. *International Religious Freedom Report*, online, accessed from <http://www.state.gov/g/drl/rls/irf/>.
- World Values Survey. Data set with Waves 3 and 4 downloaded November 2005. <http://www.worldvaluessurvey.org/>

Data Sources

- Fox, Jonathan, and Shmuel Sandler. Data obtained from Jonathan Fox April 2006: foxjon@mail.biu.ac.il

