

ANCHORING VIGNETTES IN R: A (DIFFERENT KIND OF) VIGNETTE

JONATHAN WAND AND GARY KING

ABSTRACT. The `anchors` package in R implements the techniques described in King et al. (2004), King and Wand (2007), and Wand (2007b). The procedures include methods both for evaluating and choosing anchoring vignettes, and for analyzing the resulting data. This document provides a quick introduction to setting up and using `anchors`. A companion article is also available (Wand, King, and Lau, 2007), providing details on the logic of the analysis and results for the same data used in this document. The latest version of this software and related materials are available at the `anchors` website:

<http://wand.stanford.edu/anchors/>.

1. A QUICK OVERVIEW

This section assumes that you have already have `anchors` installed and want a quick introduction/overview. Information on installation, background, and examples of `anchors` are provide in detail in subsequent sections. All examples and objects described in this document assume that you have loaded the package in an R session,

```
> library(anchors)
```

A list of the functions and datasets with help pages can be found using,

```
> help(package = "anchors")
```

For a list of datasets of vignette surveys included in `anchors`, see

```
> data(package = "anchors")
```

For a list of demonstrations of functions, uses of data, and replications of published results,

```
> demo(package = "anchors")
```

The function `anchors()` has two `method=` options

- B non-parametric rank method from Wand (2007a)
- C non-parametric rank method from King et al. (2004) and King and Wand (2007)

Date: April 17, 2011, Version: 3.0-4.

Our thanks to NIA/NIH (grant P01 AG17625-01), NSF (grants SES-0112072 and IIS-9874747), and the Weatherhead Center for International Affairs, and the Robert Wood Johnson Foundation's Health Policy Scholars program for research support.

There are two other key supporting functions that will be discussed in turn: `anchors.order()` and `chopit()`

For methods B and C, one can also specify that all combinations of subsets of vignettes (but retaining the same relative order as submitted in the formula) be analyzed using the option `anchors(..., combn=TRUE)`. The default is `combn=FALSE` since for more than three vignettes, the process requires non-trivial computational time. Details can be found in the later section on vignette selection, and via `help(anchors.combn)`.

Datasets with anchoring vignettes that are made available by the `anchors` package include

<code>chopitsim</code>	Simulated Data for test chopit function
<code>mexchn</code>	China-Mexico political efficacy data
<code>poleff</code>	Simulated Political Efficacy Data
<code>poleffna</code>	Simulated Political Efficacy Data with NA (demo only, don't use)
<code>freedom</code>	Individual freedom of speech data
<code>sleep</code>	Sleep data for china
<code>selfcare</code>	Self-care data for china
<code>table1</code>	Reference from Table 1 of King and Wand (2007)
<code>table1src</code>	Specific response values that have inequalities to create table1

Any of these can be loaded with `data()`, for example,

```
> data(freedom)
```

Demonstration files are available, both to provide examples of the use of functions and as an aid to those who would simply like to re-compute published results that have used versions of the `anchors` package,

<code>anchors.plot</code>	Demo of plotting with anchors
<code>chopit</code>	Demo of chopit: summary, plot
<code>anchors.freedom</code>	Wand et al (2007) rank analysis of freedom
<code>anchors.freedom3</code>	Wand et al (2007) Figure 2 histogram with 3 vignettes
<code>anchors.freedom6</code>	Wand et al (2007) Figure 1 histogram with 6 vignettes
<code>anchors.vign2</code>	King and Wand (2007) Table 1 anchors()
<code>anchors.mexchn</code>	King and Wand (2007) Figure 1 histogram
<code>entropy.mexchn</code>	King and Wand (2007) Figure 2 entropy()
<code>entropy.sleep</code>	King and Wand (2007) Figure 3 entropy()
<code>entropy.self</code>	King and Wand (2007) Figure 4 entropy()
<code>anchors.mexchn2</code>	Repl King et al (2004) Figure 2
<code>chopit.mexchn</code>	King et al (2004) Table 2 (non-linear taus)

Any of these can be invoked with `demo()`, for example,

```
> demo(anchors.freedom)
```

2. GETTING STARTED: INSTALLATION AND THE BASICS

We begin by walking through how to set-up `anchors` on your computer to facilitate the interactive use of the examples that follow. There are many introductions to R available on the R site, <http://www.r-project.org>, and this is only intended

as a brief summary with an emphasis on helping you to specifically get started with **anchors**.

Prior to installing **anchors**, you will need to install the R statistical package available via <http://www.r-project.org>. Use at least R 2.8. For details on installing R the FAQ at <http://cran.r-project.org/faqs.html> are helpful.

Once you have R installed, and given you have an active internet connection, the easiest way to install the **anchors** package is from the R command line,

```
> install.packages("anchors", dependencies = TRUE)
```

which will also install the **rgenoud** package if it is not already installed on your system. Alternatively, for *nix systems, you can also install the package manually by

- (1) downloading the source code from the anchors website: `anchors_3.0-4.tar.gz`.
- (2) running from the *nix shell, in the same directory as the downloaded file,

```
% R CMD INSTALL anchors_3.0-4.tgz
```

Once the **anchors** package is installed, and an R session is begun, the package is made available by invoking on the R command-line,

```
> library(anchors)
```

The full list of functions and datasets made available by **anchors** can be found by invoking at any time,

```
> help(package = "anchors")
```

and each individual object with a help file can be queried individually, for example,

```
> help(anchors)
```

3. INTRODUCTION TO ANCHORING VIGNETTES

Consider a survey question along with response categories that is asked as a self-assessment,

How free do you think you are to express yourself without fear
of government reprisal? (1) Completely Free, (2) Very Free, (3)
Moderately Free, (4) Slightly Free, (5) Not Free at All

One key difficulty of analyzing the results from such a survey question is the possibility that individuals apply different standards in the selection of a response category. Researchers have tried to ameliorate the problems of interpersonal and cross-cultural incomparability in survey research with careful question wording, translation (and back translation), focus groups, cognitive debriefing, and other techniques, most of which are designed to improve the survey question. In contrast, anchoring vignettes is a technique that seeks to bring additional data to bear on the problem.

For example, vignettes corresponding to the above political freedom question attempt to describe hypothetical individuals who have different levels of freedom

from government reprisal. The following six vignettes are intended to correspond to distinct levels of political freedom in order of decreasing freedom,

- vign1:** [Kay] does not like many of the government’s policies. She frequently publishes her opinion in newspapers, criticizing decisions by officials and calling for change. She sees little reason these actions could lead to government reprisal.
- vign2:** [Michael] disagrees with many of the government’s policies. Though he knows criticism is frowned upon, he doesn’t believe the government would punish someone for expressing critical views. He makes his opinion known on most issues without regard to who is listening.
- vign3:** [Bob] has political views at odds with the government. He has heard of people occasionally being arrested for speaking out against the government, and government leaders sometimes make political speeches condemning those who criticize. He sometimes writes letters to newspapers about politics, but he is careful not to use his real name.
- vign4:** [Connie] does not like the government’s stance on many issues. She has a friend who was arrested for being too openly critical of governmental leaders, and so she avoids voicing her opinions in public places.
- vign5:** [Vito] disagrees with many of the government’s policies, and is very careful about whom he says this to, reserving his real opinions for family and close friends only. He knows several men who have been taken away by government officials for saying negative things in public.
- vign6:** [Sonny] lives in fear of being harassed for his political views. Everyone he knows who has spoken out against the government has been arrested or taken away. He never says a word about anything the government does, not even when he is at home alone with his family.

After each of these vignettes, a corresponding evaluation question is asked with the same response categories as for the self-assessment.

How free do you think [name] is to express [him/her]self without fear of government reprisal? (1) Completely Free, (2) Very Free, (3) Moderately Free, (4) Slightly Free, (5) Not Free at All

Note: In the case where there are missing values for responses to the self-assessment or the vignettes, it is important that these be coded as ‘0’ (zero), instead of NA or some other missing value if you wish to retain the other (non-missing) responses of an individual in the parametric model to be described shortly (see `chopit`). For all non-parametric analysis that rely on `anchors` or `anchors.order`, cases with missing responses (either NA or zero) must be listwise deleted. We provide a handy function, `replace.value`, that facilitates the alteration of the coding of missing values for subsets of variables.

4. INDEXING NOTATION

Our notation is a generalization of King et al. designed to accommodate our enhancements to the various models. We index survey questions, response categories, and respondents as follows:

- We index *survey questions* by the pair (s, j) , where question set s ($s = 1, \dots, S$) corresponds to the self-assessment question number and refers to the set of questions that includes the self-assessment question (indicated by $j = 0$) and, optionally, one or more vignette questions (indicated by $j = 1, \dots, J_s$).
- We index *response categories* by k ($k = 1, \dots, K_s$) separately for each survey question since they can each have different response categories. Each set of questions (self-assessment and vignettes) must have the same number of choice categories (coded as increasing sequential integers starting with 1). *Missing values* (whether structural, because the question was not asked, or due to nonresponse) should be coded as $k = 0$.
- We index *respondents* by i or ℓ . Respondent i ($i = 1, \dots, n$) is asked all of the self-assessment questions. Respondent ℓ ($\ell = 1, \dots, N$) is asked all of the vignette questions. (Respondents are indexed for self-assessment and vignette questions separately since each could be asked of independent samples; if they are asked of the same individuals, then $i = \ell$ and $n = N$.) If your survey design asks each set of vignette questions in separate samples (and separate from the self-assessment question), then index each set of vignettes according to unique values of ℓ and use the missing value code ($k = 0$) for vignettes that are not asked of a subgroup; in other words, stack the data in block diagonal format.

Thus, every mathematical symbol in the model could be indexed by s , j , k , and either i or ℓ . In practice, we drop indexes that are constant.

5. A NONPARAMETRIC APPROACH

5.1. Definition. Define C_{is} as the self-assessment *relative* to the corresponding set of vignettes. Let y_i be the self-assessment response and z_{i1}, \dots, z_{iJ} be the J vignette responses, for the i th respondent. For respondents with consistently ordered rankings on all vignettes ($z_{j-1} < z_j$, for $j = 2, \dots, J$), we create the DIF-corrected self-assessment C_i

$$(1) \quad C_i = \begin{cases} 1 & \text{if } y_i < z_{i1} \\ 2 & \text{if } y_i = z_{i1} \\ 3 & \text{if } z_{i1} < y_i < z_{i2} \\ \vdots & \vdots \\ 2J + 1 & \text{if } y_i > z_{iJ} \end{cases}$$

Respondents who give tied or inconsistently ordered vignette responses may have an interval values of C , if the tie/inconsistency results in multiple conditions in equation 1 appearing to be true. A more general definition of C is defined as the minimum to maximum values among all the conditions that hold true in equation 1. Values of C that are intervals, rather than scalar, represent the set of inequalities over which the analyst cannot distinguish without further assumption.

5.2. EXAMPLE CODE: `anchors()`. This example again first loads the library and example dataset, and then `anchors()` calculates C for each individual. In the non-parametric estimation, only *one* self-question and corresponding set of vignettes are analyzed at a time.

```
> library(anchors)
> data(freedom)
> a1 <- anchors(self ~ vign2 + vign3 + vign4 + vign5 +
+   vign6, freedom, method = "C")
> summary(a1)
```

ANCHORS: SUMMARY OF RELATIVE RANK ANALYSIS:

Overview of C-ranks

Number of cases: 1763 with interval value, 1737 with scalar value

Maximum possible C-rank value: 11

Interval on C-scale: Frequency and proportions Cs to Ce

	N	Prop	MinEnt
1 to 1	387	0.111	1
2 to 2	279	0.080	2
3 to 3	336	0.096	3
4 to 4	81	0.023	4
5 to 5	59	0.017	5
6 to 6	28	0.008	6
7 to 7	11	0.003	7
8 to 8	31	0.009	8
9 to 9	22	0.006	9
10 to 10	164	0.047	10
11 to 11	339	0.097	11
1 to 4	16	0.005	1
1 to 5	12	0.003	1
1 to 6	25	0.007	6
1 to 7	5	0.001	6
1 to 8	31	0.009	6
1 to 9	5	0.001	6
1 to 10	32	0.009	6
1 to 11	19	0.005	6
2 to 4	15	0.004	3
2 to 5	11	0.003	3
2 to 6	22	0.006	6
2 to 7	4	0.001	6
2 to 8	51	0.015	6
2 to 9	19	0.005	6
2 to 10	177	0.051	6
2 to 11	91	0.026	6
3 to 6	31	0.009	6

3 to 7	3	0.001	6
3 to 8	93	0.027	6
3 to 9	29	0.008	6
3 to 10	16	0.005	6
3 to 11	11	0.003	6
4 to 6	16	0.005	6
4 to 7	2	0.001	6
4 to 8	94	0.027	6
4 to 9	39	0.011	6
4 to 10	175	0.050	6
4 to 11	39	0.011	6
5 to 8	80	0.023	6
5 to 9	38	0.011	6
5 to 10	9	0.003	6
5 to 11	6	0.002	6
6 to 8	107	0.031	6
6 to 9	61	0.017	6
6 to 10	242	0.069	6
6 to 11	52	0.015	6
7 to 10	1	0.000	10
7 to 11	1	0.000	11
8 to 10	44	0.013	10
8 to 11	39	0.011	11

Note: MinEnt is the rank for the interval that minimizes entropy

Summary of C-ranks with ties/intervals broken:

Distribution of ranks omitting interval cases

1	2	3	4	5	6	7	8	9
0.223	0.161	0.193	0.047	0.034	0.016	0.006	0.018	0.013
10	11							
0.094	0.195							

Distribution of ranks allocating interval cases uniformly

1	2	3	4	5	6	7	8	9	10	11
0.116	0.1	0.125	0.07	0.07	0.09	0.079	0.091	0.06	0.09	0.107

Distribution of ranks allocating interval cases via cpolr
and conditioning on observed ranks

1	2	3	4	5	6	7	8	9	10
0.118	0.103	0.142	0.051	0.045	0.138	0.025	0.155	0.017	0.095
11									
0.110									

Allocating cases to their MinEnt values produces

1	2	3	4	5	6	7	8	9	10
0.119	0.080	0.103	0.023	0.017	0.472	0.003	0.009	0.006	0.060

11
0.108

The names of vignettes must be passed to the function in the same order as the direction of the responses. In the example, `vign2` is in the same (highest) direction as the response category 1, while the `vign6` is in the same direction (lowest) as the response category 5. (We drop `vign1` here for space reason when printing the summary—with the different combinations of intervals of C can be numerous.)

If `anchors` produces many ties you should check that you passed the vignettes in the correct order, but we also offer a function that investigates the ordering of vignettes in detail.

5.3. EXAMPLE CODE: `anchors.order()`. The function `anchors.order()`, and the associated methods `summary.anchors.order` and `barplot.anchors.order` investigate the relationship between vignette responses *without* reference to the self-assessment question.

```
> vo1 <- anchors.order(~vign2 + vign3 + vign4 +
+   vign5 + vign6, freedom)
> summary(vo1, top = 10, digits = 3)
```

ANCHORS: SUMMARY OF VIGNETTE ORDERING

Treatment of ties: represent as sets

Number of cases with at least two distinct vignette responses: 3223
and with no violations of natural ordering: 1178
and with no more than 1 violation of natural ordering: 1959
and with no more than 2 violation of natural ordering: 2621

Proportion of cases a vignette (row) is less than another (column):

	<1	<2	<3	<4	<5
1	NA	0.663	0.732	0.707	0.754
2	0.121	NA	0.457	0.363	0.575
3	0.080	0.138	NA	0.183	0.374
4	0.068	0.198	0.339	NA	0.495
5	0.070	0.081	0.100	0.103	NA

Upper tri = $p_{\{ij\}} - p_{\{ji\}}$ (negative values suggest misorderings)

Lower tri = $1 - p_{\{ij\}} - p_{\{ji\}}$ (big numbers means many ties)


	1	2	3	4	5
1	NA	0.542	0.652	0.639	0.684
2	0.215	NA	0.320	0.165	0.494
3	0.188	0.440	NA	-0.156	0.275
4	0.405	0.477	0.345	NA	0.392
5	0.225	0.176	0.526	0.402	NA

Top 10 orderings (out of 262 unique orderings):

	Frequency	Proportion	Ndistinct	Nviolation
1,{2,3,4,5}	309	0.0883	2	0
{1,2,3,4,5}	277	0.0791	1	0
1,2,{3,4,5}	221	0.0631	3	0
1,{2,4},{3,5}	212	0.0606	3	1
1,2,{3,4},5	178	0.0509	4	0
1,{2,3,4},5	177	0.0506	3	0
1,4,{2,3,5}	157	0.0449	3	2
1,2,4,{3,5}	106	0.0303	4	1
1,{2,4},3,5	96	0.0274	4	1
1,4,2,{3,5}	73	0.0209	4	2

```
> barplot(vo1)
```

Treatment of ties: represent as sets



Details of how to interpret and use the output of the summary are provided in Wand, King, and Lau (2007), where it is discussed in detail how **vign6** is given the highest response almost half the time, however **vign4** is more often given the highest response than **vign5**.

In light of this it is worth reestimating C using the consensus ordering of the vignettes,

```
> a2 <- anchors(self ~ vign2 + vign3 + vign5 + vign4 +
+   vign6, freedom, method = "C")
> summary(a2)
```

ANCHORS: SUMMARY OF RELATIVE RANK ANALYSIS:

Overview of C-ranks

Number of cases: 1654 with interval value, 1846 with scalar value

Maximum possible C-rank value: 11

Interval on C-scale: Frequency and proportions Cs to Ce

	N	Prop	MinEnt
1 to 1	387	0.111	1
2 to 2	279	0.080	2
3 to 3	336	0.096	3
4 to 4	81	0.023	4
5 to 5	59	0.017	5
6 to 6	80	0.023	6
7 to 7	38	0.011	7
8 to 8	61	0.017	8
9 to 9	22	0.006	9
10 to 10	164	0.047	10
11 to 11	339	0.097	11
1 to 4	16	0.005	1
1 to 5	12	0.003	1
1 to 6	20	0.006	6
1 to 7	1	0.000	6
1 to 8	39	0.011	6
1 to 9	6	0.002	6
1 to 10	32	0.009	6
1 to 11	19	0.005	6
2 to 4	15	0.004	3
2 to 5	11	0.003	3
2 to 6	31	0.009	6
2 to 7	6	0.002	6
2 to 8	51	0.015	6
2 to 9	8	0.002	6
2 to 10	177	0.051	6
2 to 11	91	0.026	6
3 to 6	63	0.018	6
3 to 7	19	0.005	6
3 to 8	67	0.019	6
3 to 9	7	0.002	6
3 to 10	16	0.005	6
3 to 11	11	0.003	6
4 to 6	59	0.017	6
4 to 7	17	0.005	6
4 to 8	60	0.017	6
4 to 9	15	0.004	6
4 to 10	175	0.050	6

4 to 11	39	0.011	6
5 to 8	28	0.008	6
5 to 9	11	0.003	6
5 to 10	9	0.003	6
5 to 11	6	0.002	6
6 to 8	107	0.031	6
6 to 9	31	0.009	6
6 to 10	158	0.045	6
6 to 11	50	0.014	6
7 to 10	3	0.001	10
7 to 11	1	0.000	11
8 to 10	126	0.036	10
8 to 11	41	0.012	11

Note: MinEnt is the rank for the interval that minimizes entropy

Summary of C-ranks with ties/intervals broken:

Distribution of ranks omiting interval cases

1	2	3	4	5	6	7	8	9	10
0.21	0.151	0.182	0.044	0.032	0.043	0.021	0.033	0.012	0.089
11									
0.184									

Distribution of ranks allocating interval cases uniformly

1	2	3	4	5	6	7	8	9	10
0.116	0.1	0.127	0.073	0.068	0.096	0.072	0.09	0.057	0.093
11									
0.107									

Distribution of ranks allocating interval cases via cpolr and conditioning on observed ranks

1	2	3	4	5	6	7	8	9	10
0.118	0.103	0.144	0.056	0.042	0.147	0.037	0.120	0.016	0.107
11									
0.110									

Allocating cases to their MinEnt values produces

1	2	3	4	5	6	7	8	9	10
0.119	0.080	0.103	0.023	0.017	0.431	0.011	0.017	0.006	0.084
11									
0.109									

Changing the assumed ordering of the vignettes increased the number of cases without any order violation by 60 percent. With respect to the top sets of types of ordering,

The analysis of vignettes is useful both at the stage of evaluating a pilot study of survey instruments, as well at the stage of choosing how (and whether) to use

particular vignettes. The results of non-parametric anchoring vignettes analysis using C are entirely dependent on which vignettes are included and the order in which they are specified.

5.4. Example Code: Subsets of vignettes. Calculating entropy for subsets of vignettes as suggested by Wand and King (2007) is straightforward. The `anchors(...,combn=TRUE)` calculates statistics of interest, including entropy measures, for every ordered combination of vignettes. For more details, please see `help(anchors.combn)` in R and King and Wand (2007).

```
> data(freedom)
> fo <- list(self = self ~ 1, vign = cbind(vign1,
+     vign3, vign6) ~ 1, cpolr = ~as.factor(country) +
+     sex + age + educ)
> ent <- anchors(fo, data = freedom, method = "C",
+     combn = TRUE)
> summary(ent, digits = 3)
```

ANCHORS: SUMMARY OF RELATIVE RANK ANALYSIS:

Overview of C-ranks

Number of cases: 522 with interval value, 2925 with scalar value

Maximum possible C-rank value: 7

Interval on C-scale: Frequency and proportions Cs to Ce

	N	Prop	MinEnt
1 to 1	496	0.144	1
2 to 2	225	0.065	2
3 to 3	492	0.143	3
4 to 4	236	0.068	4
5 to 5	497	0.144	5
6 to 6	489	0.142	6
7 to 7	490	0.142	7
1 to 4	22	0.006	3
1 to 5	1	0.000	5
1 to 6	28	0.008	5
1 to 7	12	0.003	5
2 to 4	39	0.011	3
2 to 5	10	0.003	5
2 to 6	124	0.036	5
2 to 7	31	0.009	5
3 to 6	9	0.003	5
3 to 7	9	0.003	5
4 to 6	193	0.056	5
4 to 7	44	0.013	5

Note: MinEnt is the rank for the interval that minimizes entropy

Summary of C-ranks with ties/intervals broken:

Distribution of ranks omiting interval cases

1	2	3	4	5	6	7
0.17	0.077	0.168	0.081	0.17	0.167	0.168

Distribution of ranks allocating interval cases uniformly

1	2	3	4	5	6	7
0.147	0.082	0.161	0.108	0.179	0.175	0.148

Distribution of ranks allocating interval cases via cpolr and conditioning on observed ranks

1	2	3	4	5	6	7
0.148	0.075	0.165	0.094	0.187	0.183	0.148

Allocating cases to their MinEnt values produces

1	2	3	4	5	6	7
0.144	0.065	0.160	0.068	0.278	0.142	0.142

Summary of entropy and intervals by subsets of vignettes:


	Vignettes	Estimated entropy	Minimum entropy
1	123	1.903	1.844
4	12	1.450	1.430
3	13	1.490	1.465
2	23	1.527	1.489
7	1	0.795	0.795
6	3	0.959	0.959
5	2	1.044	1.044

	Interval Cases	Span avg.	Max. rank
1	522	1.471	7
4	220	1.151	5
3	195	1.137	5
2	418	1.285	5
7	0	1.000	3
6	0	1.000	3
5	0	1.000	3

One important feature to be noted about including `cpolr=` variables is that cases with any missing value in the covariates will be listwise deleted for both the estimated and minimum entropy calculations to ensure a common basis for comparisons. As such, the minimum entropy values may change as a function of what variables (if any) are included in `cpolr=`.

The `plot()` method is described in `help(plot.anchors.rank)`, and an example is given here,

```
> plot(ent)
```



6. PARAMETRIC MODEL

This section describes the Compound Hierarchical Ordered Probit (chopit) model.

6.1. Self-assessment component. Figure 1 summarizes the self-assessment component of the model.

The *actual* level for respondent i is μ_i , a continuous unidimensional variable (with higher values indicating more freedom, better health, etc., defined by the order of the vignettes). Respondent i perceives μ_i only with random normal error so that

$$(2) \quad Y_{is}^* \sim N(\mu_i, \sigma_s^2)$$

is respondent i 's unobserved *perceived* level. The actual level is a linear function of observed covariates X_i , the first column of which can be a constant term (if it is not needed for identification) and an independent normal random effect η_i :

$$(3) \quad \mu_i = X_i\beta + \eta_i$$

with parameter β and

$$(4) \quad \eta_i \sim N(0, \omega^2).$$

The *reported* survey response category is y_{is} and is generated by the model via this observation mechanism:

$$(5) \quad y_{is} = k \quad \text{if } \tau_{is}^{k-1} \leq Y_{is}^* < \tau_{is}^k$$




FIGURE 1. Self-Assessment Component: All levels vary over observations i . Each solid arrow denotes a deterministic effect; a squiggly arrow denotes the addition of normal random error with variance indicated at the arrow's source.




FIGURE 2. Vignette Component for question set s ($s = 1, \dots, S'$, $S' \leq S$). All levels vary over observations ℓ . Each solid arrow denotes a deterministic effect; a squiggly arrow denotes the addition of normal random error with variance indicated at the arrow's source.

with a vector of thresholds τ_{is} (where $\tau_{is}^0 = -\infty$, $\tau_{is}^{K_s} = \infty$, and $\tau_{is}^{k-1} < \tau_{is}^k$, with indexes for categories $k = 1, \dots, K_s$ and self-assessment questions $s = 1, \dots, S$) that vary over the observations as a function of a vector of covariates, V_i (the first column of which can be a constant term), and unknown parameter vectors γ_s (with elements the vector γ_s^k):

$$(6) \quad \begin{aligned} \tau_{is}^1 &= \gamma_s^1 V_i \\ \tau_{is}^k &= \tau_{is}^{k-1} + e^{\gamma_s^k V_i} \quad (k = 2, \dots, K_s - 1) \end{aligned}$$

6.2. Vignette Component. Figure 2 summarizes the vignette component of the model for question set s ($s = 1, \dots, S$). Under the model, one or more of the self-assessment questions have corresponding vignettes.

The actual level for vignette j is θ_j ($j = 1, \dots, J_s$), measured on the same scale as μ_i and the τ 's. Respondent ℓ perceives θ_j with random normal error so that

$$(7) \quad Z_{\ell sj}^* \sim N(\theta_j, \sigma_{sj}^2)$$

represents respondent ℓ 's unobserved assessment of the level of vignette j for question set s .

The perception of respondent ℓ about the level of vignette j elicited via a survey question s with the same K_s ordinal categories as the corresponding self-assessment question. Thus, the respondent turns the continuous $Z_{\ell sj}^*$ into a categorical answer to the survey question $z_{\ell sj}$ via this observation mechanism:

$$(8) \quad z_{\ell sj} = k \quad \text{if } \tau_{\ell s}^{k-1} \leq Z_{\ell sj}^* < \tau_{\ell s}^k$$

with thresholds determined by the same γ_s coefficients as in (6) for y_{i1} , and the same explanatory variables but with values measured for units ℓ , V_ℓ :

$$(9) \quad \begin{aligned} \tau_{\ell 1}^1 &= \gamma_s^1 V_\ell \\ \tau_{\ell 1}^k &= \tau_{\ell s}^{k-1} + e^{\gamma_s^k V_\ell} \quad (k = 2, \dots, K_1 - 1). \end{aligned}$$

6.3. Identification. The model as specified above has an infinite number of equivalent maximum likelihood solutions. To identify the model, two choices must be made:

- (1) The mean of the actual level must be set, by choosing one point. This can be done by setting the constant term $\beta_0 = 0$ (in which case be aware of your choice of the scale of the variables in X), or one of the θ 's.
- (2) The variance of the actual level must also be set. This can be done by setting all the self-assessment variances (such as $\sigma_s^2 = 1$, for all s) or by setting another point among β_0 or the θ 's.

Two common parameterizations are as follows:

- (1) The ordinal probit parameterization is useful for comparing chopit to this simpler model. Set $\beta_0 = 0$ and $\sigma_1^2 = \dots = \sigma_S^2 = 1$.
- (2) Another option is parameterization defined by the extreme vignettes. Let $\theta_1 = 0$ and $\theta_{J_s} = 1$. This lets estimates of μ be interpreted on the scale of the vignettes, with 0 being the level of the lowest vignette and 1 the level of the highest. Note that μ can still be higher than 1 or lower than 0, but the units are easily interpretable.

6.4. EXAMPLE CODE: chopit(). The `chopit()` function provided by `anchors` at its most basic simply requires specifying the formula's defining ys , zs , and τs . For example, using variables from the `data(freedom)` dataset, we have the *named* list.

```
> fo <- list(self = self ~ sex + age + educ + factor(country),
+   vign = cbind(vign1, vign2, vign3, vign4, vign5,
+   vign6) ~ 1, tau = ~sex + age + educ +
+   factor(country))
```

The names `self=`, `vign=`, and `tau=` as written, are required. On the LHS of the equality signs are the variables of the dataset that specify the details of the models as for other models (e.g., `lm()`).

The self-assessment variable `self` is modeled to have a mean that is a linear additive function of `sex`, `age`, `educ` and `country` dummies. The vignettes are specified as a vector of outcomes `cbind(vign1,vign2,vign3,vign4,vign5,vign6)` as a function of only an intercept ' ~ 1 '. This is both a simple and accurate way to describe the model of θ s which are the mean locations of the vignettes. The τ cutpoints shared by the self-assessment and the vignettes are specified as their own formula without a LHS variable.

Beyond the formula and data, the rest will be set by default in the basic invocation,

```
> out <- chopit(fo, data = freedom)
```

which can be summarized by the `summary` method,

```
> summary(out)
```

ANCHORS: SUMMARY OF RELATIVE CHOPIT ANALYSIS:

Model formula:

`$self`

`self ~ sex + age + educ + factor(country)`

`$vign`

`cbind(vign1, vign2, vign3, vign4, vign5, vign6) ~ 1`

`$tau`

`~sex + age + educ + factor(country)`

`$cpolr`

`~1`

`<environment: 0x7bb8a8>`

Coefficients:

	coeff	se
gamma.cut1.(Intercept)	-1.6697	0.0774
gamma.cut1.sex	0.0570	0.0228
gamma.cut1.age	-0.0028	0.0007
gamma.cut1.educ	0.0109	0.0068
gamma.cut1.factor(country)Eurasia	0.0447	0.0504
gamma.cut1.factor(country)Oceania	-0.1262	0.0309
gamma.cut2.(Intercept)	0.6655	0.0388
gamma.cut2.sex	-0.0426	0.0205
gamma.cut2.age	0.0013	0.0006
gamma.cut2.educ	-0.0140	0.0061
gamma.cut2.factor(country)Eurasia	-0.0286	0.0449
gamma.cut2.factor(country)Oceania	0.0260	0.0274
gamma.cut3.(Intercept)	0.7068	0.0319
gamma.cut3.sex	-0.0211	0.0167
gamma.cut3.age	-0.0001	0.0005

gamma.cut3.educ	0.0112	0.0051
gamma.cut3.factor(country)Eurasia	0.0250	0.0374
gamma.cut3.factor(country)Oceania	-0.0985	0.0218
gamma.cut4.(Intercept)	0.5937	0.0294
gamma.cut4.sex	0.0436	0.0159
gamma.cut4.age	0.0007	0.0005
gamma.cut4.educ	0.0163	0.0049
gamma.cut4.factor(country)Eurasia	0.0605	0.0365
gamma.cut4.factor(country)Oceania	0.0166	0.0211
sigma.random.effect	1.0000	NaN
sigma.self	1.0000	NaN
sigma.vign1	0.7951	0.0183
sigma.vign2	0.9974	0.0239
sigma.vign3	0.7546	0.0173
sigma.vign4	0.8336	0.0208
sigma.vign5	0.7246	0.0171
sigma.vign6	1.3307	0.0420
theta.vign1	-1.0863	0.0721
theta.vign2	-1.2051	0.0734
theta.vign3	-0.2478	0.0706
theta.vign4	0.1660	0.0715
theta.vign5	-0.0562	0.0706
theta.vign6	0.9519	0.0820
beta.(Intercept)	0.0000	NaN
beta.sex	0.1434	0.0388
beta.age	-0.0019	0.0012
beta.educ	-0.0569	0.0117
beta.factor(country)Eurasia	0.4600	0.0897
beta.factor(country)Oceania	-0.7019	0.0517

-Log-likelihood of CHOPIT: 32421.69

Partition of CHOPIT -Log-likelihood by question:

	-LL	N
Self (self)	5154.965	3447
vign1	5032.314	3447
vign2	5207.052	3447
vign3	4766.234	3447
vign4	4340.710	3447
vign5	4485.543	3447
vign6	3434.877	3447

Number of cases that contribute at least partially to likelihoods:

- a) in self-responses: 3447
- b) in vign-responses: 3447

The default invocation uses the the ordinal probit normalization, which identifies/normalizes the model by omitting the intercept in μ , and setting $\sigma_1 = 1$ (the

variance of the first self-assessment question). If one specified the explanatory variables of `self=` to include an intercept, then that intercept parameter would be constrained to be zero as will be `beta.(Intercept)` in this example.

7. LIST OF FUNCTIONS

Here is a list of function available in `anchors`, and help files are available for each

	<code>allequal.test</code>	all.equal with expected outcome test
	<code>anchors</code>	Analysis of surveys with anchoring vignettes
	<code>anchors.chopit</code>	Compound Hierarchical Ordered Probit (CHOPIT)
	<code>anchors.combn</code>	Calculate known minimum or estimated entropy for subsets
	<code>anchors.data</code>	Organized data from surveys with anchoring
	<code>anchors.options</code>	Set or query <code>anchors()</code> parameters
	<code>anchors.order</code>	Calculate frequency of vignette orderings
	<code>anchors.rank</code>	Non-parametric analysis of surveys with
	<code>convert</code>	Convert factor variables into integers
	<code>cpolr</code>	Censored ordered probit
of them.	<code>fitted.anchors.cpolr</code>	Conditional and unconditional prediction for cpolr
	<code>fitted.anchors.rank</code>	Fitted values of non-parametric models
	<code>fitted.cpolr</code>	Conditional and unconditional prediction for cpolr
	<code>insert</code>	Insert DIF-corrected variable into dataframe
	<code>barplot.anchors.order</code>	Plot frequency of vignette orderings
	<code>barplot.anchors.rank</code>	Plot distribution of non-parametric ranks
	<code>replace.list</code>	Updating contents of one list using a second
	<code>replace.value</code>	Replaces occurrences of a value with another
	<code>summary.anchors.chopit</code>	Summary of CHOPIT Analysis
	<code>summary.anchors.order</code>	Calculate frequency of vignette orderings
	<code>summary.anchors.rank</code>	Summary of non-parameteric anchors analysis
	<code>trim.data</code>	Trim a dataset to match <code>anchors.data</code>

REFERENCES

- King, Gary, C.J.L. Murray, J.A. Salomon, and A. Tandon. 2004. "Enhancing the Validity and Cross-Cultural Comparability of Measurement in Survey Research." *American Political Science Review* 98 (01): 191–207.
- King, Gary, and Jonathan Wand. 2007. "Comparing Incomparable Survey Responses: New Tools for Anchoring Vignettes." *Political Analysis* 15: 46–66.
- Wand, Jonathan. 2007a. "Credible Comparisons Using Interpersonally Incomparable Data: Ranking self-evaluations relative to anchoring vignettes or other common survey questions." Stanford University.
URL: http://wand.stanford.edu/anchors/wand_anchors.pdf
- Wand, Jonathan. 2007b. "Surveys with Anchoring Vignettes: Theory and Practice." Stanford University, mimeo.
- Wand, Jonathan, Gary King, and Olivia Lau. 2007. "Anchors: Software for Anchoring Vignette Data." *Journal of Statistical Software*. Submitted for review.