cmrologit	— Rank-ordered logit choice model

Description

cmrologit fits the rank-ordered logistic regression model by maximum likelihood (Beggs, Cardell, and Hausman 1981). This model is also known as the Plackett–Luce model (Marden 1995), as the exploded logit model (Punj and Staelin 1978), and as the choice-based method of conjoint analysis (Hair et al. 2010).

Quick start

Rank-ordered logit model of rankings y on x1, x2, and x3, using cmset data

    cmrologit y x1 x2 x3

Same as above, but interpret the lowest value of y as the best

    cmrologit y x1 x2 x3, reverse

Use Efron's method for handling ties in rankings

    cmrologit y x1 x2 x3, ties(efron)

With cluster–robust standard errors for clustering by levels of cvar

    cmrologit y x1 x2 x3, vce(cluster cvar)

Menu

Statistics > Choice models > Rank-ordered logit model

Syntax

> cmrologit *depvar indepvars* [*if*] [*in*] [*weight*] [, *options*]

options	Description
Model	
<u>inc</u>omplete(*#*)	use *#* to code unranked alternatives; default is incomplete(0)
reverse	reverse the preference order
ties(*spec*)	method to handle ties: exactm, breslow, efron, or none
altwise	use alternativewise deletion instead of casewise deletion
notestrhs	keep right-hand-side variables that do not vary within case
<u>off</u>set(*varname*)	include *varname* in model with coefficient constrained to 1
SE/Robust	
vce(*vcetype*)	*vcetype* may be oim, <u>r</u>obust, <u>cl</u>uster *clustvar*, <u>boot</u>strap, or jackknife
Reporting	
level(*#*)	set confidence level; default is level(95)
display_options	control columns and column formats, row spacing, line width, display of omitted variables and base and empty cells, and factor-variable labeling
Maximization	
maximize_options	control the maximization process; seldom used
<u>coefl</u>egend	display legend instead of statistics

You must cmset your data before using cmrologit; see [CM] **cmset**.

indepvars may contain factor variables; see [U] **11.4.3 Factor variables**.

bootstrap, by, collect, fp, jackknife, and statsby are allowed; see [U] **11.1.10 Prefix commands**.

Weights are not allowed with the bootstrap prefix; see [R] **bootstrap**.

fweights, iweights, and pweights are allowed, except no weights are allowed with ties(efron), and pweights are not allowed with ties(exactm); see [U] **11.1.6 weight**.

coeflegend does not appear in the dialog box.

See [U] **20 Estimation and postestimation commands** for more capabilities of estimation commands.

Options

> [**Model**]

incomplete(*#*) specifies the numeric value used to code alternatives that are not ranked. It is assumed that unranked alternatives are less preferred than the ranked alternatives (that is, the data record the ranking of the most preferred alternatives). It is not assumed that subjects are indifferent between the unranked alternatives. The default is incomplete(0).

reverse specifies that in the preference order, a higher number means a less attractive alternative. The default is that higher values indicate more attractive alternatives. The rank-ordered logit model is not symmetric in the sense that reversing the ordering simply leads to a change in the signs of the coefficients.

ties(*spec*) specifies the method for handling ties (indifference between alternatives) (see [ST] **stcox** for details):

exactm	exact marginal likelihood (default)
breslow	Breslow's method (default if pweights specified)
efron	Efron's method (default if robust VCE)
none	no ties allowed

altwise specifies that alternativewise deletion be used when omitting observations because of missing values in your variables. The default is to use casewise deletion; that is, the entire group of observations making up a case is omitted if any missing values are encountered. This option does not apply to observations that are excluded by the if or in qualifier or the by prefix; these observations are always handled alternativewise regardless of whether altwise is specified.

notestrhs suppresses the test that the independent variables vary within (at least some of) the cases. Effects of variables that are always constant are not identified. For instance, a rater's gender cannot directly affect his or her rankings; it could affect the rankings only via an interaction with a variable that does vary over alternatives.

offset(*varname*); see [R] **Estimation options**.

> SE/Robust

vce(*vcetype*) specifies the type of standard error reported, which includes types that are derived from asymptotic theory (oim), that are robust to some kinds of misspecification (robust), that allow for intragroup correlation (cluster *clustvar*), and that use bootstrap or jackknife methods (bootstrap, jackknife); see [R] *vce_option*.

If ties(exactm) is specified, *vcetype* may be only oim, bootstrap, or jackknife.

> Reporting

level(*#*); see [R] **Estimation options**.

display_options: noci, nopvalues, noomitted, vsquish, noemptycells, baselevels, allbaselevels, nofvlabel, fvwrap(*#*), fvwrapon(*style*), cformat(*%fmt*), pformat(*%fmt*), sformat(*%fmt*), and nolstretch; see [R] **Estimation options**.

> Maximization

maximize_options: iterate(*#*), trace, [no]log, tolerance(*#*), ltolerance(*#*), nrtolerance(*#*), and nonrtolerance; see [R] **Maximize**. These options are seldom used.

The following option is available with cmrologit but is not shown in the dialog box:

coeflegend; see [R] **Estimation options**.

Remarks and examples

Remarks are presented under the following headings:

Overview

The rank-ordered logit model can be applied to analyze how decision makers combine attributes of alternatives into overall evaluations of the attractiveness of these alternatives. The model generalizes a version of McFadden's choice model, one where alternatives are not explicitly identified. It uses information about the comparison of alternatives, namely, how decision makers rank the alternatives rather than just specifying the alternative that they like best.

cmrologit expects the data to be in long form, similar to clogit (see [R] **clogit**), in which each of the ranked alternatives forms an observation. The distinction from a McFadden's choice model fit with clogit is that *depvar* in cmrologit records the rankings of the alternatives, whereas for clogit, *depvar* indicates a single chosen alternative by a value not equal to zero. If your data record only one preferred alternative for each case, cmrologit fits the same model as clogit.

cmrologit interprets equal values of *depvar* as ties. The ranking information may be incomplete "at the bottom" (least preferred alternatives). That is, unranked alternatives may be coded as 0 or as a common value that may be specified with the incomplete() option.

All observations related to an individual are linked together by the case ID variable that you specify in cmset. Alternatives are not explicitly identified by an indicator variable, so cmset is used with the noalternatives option. For example, if id is your case ID variable, to cmset your data before running cmrologit, you type

```
. cmset id, noalternatives
```

For details on cmset, see [CM] **cmset**.

Examples

A popular way to study employer preferences for characteristics of employees is the quasiexperimental "vignette method". As an example, we consider the research by de Wolf (2000) on the labor market position of social science graduates. This study addresses how the educational portfolio (for example, general skills versus specific knowledge) affects short-term and long-term labor-market opportunities.

De Wolf asked 22 human resource managers (the respondents) to rank order the 6 most suitable candidates of 20 fictitious applicants and to rank order these 6 candidates for 3 jobs, namely, 1) researcher, 2) management trainee, and 3) policy adviser. Applicants were described by 10 attributes, including their age, gender, details of their portfolio, and work experience. In this example, we analyze a subset of the data.

To simplify the output, we dropped, at random, 10 nonselected applicants per case. The resulting dataset includes 29 cases, consisting of 10 applicants each. The data are in long form: observations correspond to alternatives (the applications), and alternatives that figured in one decision task are identified by the variable `caseid`. We list the observations for `caseid==7`, in which the respondent considered applicants for a social-science research position.

```
. use https://www.stata-press.com/data/r19/evignet
(Vignet study employer prefs (Inge de Wolf 2000))
. list pref female age grades edufit workexp boardexp if caseid==7, noobs
```

pref	female	age	grades	edufit	workexp	boardexp
0	yes	28	A/B	no	none	no
0	no	25	C/D	yes	one year	no
0	no	25	C/D	yes	none	yes
0	yes	25	C/D	no	internship	yes
1	no	25	C/D	yes	one year	yes
2	no	25	A/B	yes	none	no
3	yes	25	A/B	yes	one year	no
4	yes	25	A/B	yes	none	yes
5	no	25	A/B	yes	internship	no
6	yes	28	A/B	yes	one year	yes

Here 6 applicants were selected. The rankings are stored in the variable `pref`, where a value of 6 corresponds to "best among the candidates", a value of 5 corresponds to "second-best among the candidates", etc. The applicants with a ranking of 0 were not among the best 6 candidates for the job. The respondent was not asked to express his or her preferences among these four applicants, but by the elicitation procedure, it is known that he or she ranks these four applicants below the 6 selected applicants.

The best candidate was a female, 28 years old, with education fitting the job, with good grades (A/B), with 1 year of work experience, and with experience being a board member of a fraternity, a sports club, etc. The profiles of the other candidates read similarly. Here the respondent completed the task; that is, he or she selected and rank ordered the 6 most suitable applicants. Sometimes the respondent performed only part of the task.

```
. list pref female age grades edufit workexp boardexp if caseid==18, noobs
```

pref	female	age	grades	edufit	workexp	boardexp
0	no	25	C/D	yes	none	yes
0	no	25	C/D	no	internship	yes
0	no	28	C/D	no	internship	yes
0	yes	25	A/B	no	one year	no
2	yes	25	A/B	no	none	yes
2	no	25	A/B	no	none	yes
2	no	25	A/B	no	one year	yes
5	no	25	A/B	no	none	yes
5	no	25	A/B	no	none	yes
5	yes	25	A/B	no	none	no

The respondent selected the six best candidates and segmented these six candidates into two groups: one group with the three best candidates and a second group of three candidates that were "still acceptable". The numbers 2 and 5, indicating these two groups, are arbitrary apart from the implied ranking of the groups. The ties between the candidates in a group indicate that the respondent was not able to rank the candidates within the group.

The purpose of the vignette experiment was to explore and test hypotheses about which of the employees' attributes are valued by employers, how these attributes are weighted depending on the type of job (described by variable job in these data), etc. In the psychometric tradition of Thurstone (1927), "value" is assumed to be linear in the attributes, with the coefficients expressing the direction and weight of the attributes. In addition, it is assumed that "valuation" is to some extent a random procedure, captured by an additive random term. For instance, if value depends only on an applicant's age and gender, we would have

$$\text{value}(\texttt{female}_i, \texttt{age}_i) = \beta_1 \texttt{female}_i + \beta_2 \texttt{age}_i + \epsilon_i$$

where the random residual, ϵ_i, captures all omitted attributes. Thus, $\beta_1 > 0$ means that the employer assigns higher value to a woman than to a man.

Given this conceptualization of value, it is straightforward to model the decision (selection) among alternatives or the ranking of alternatives: the alternative with the highest value is selected (chosen), or the alternatives are ranked according to their value. To complete the specification of a model of choice and of ranking, we assume that the random residual ϵ_i follows an "extreme value distribution of type I", introduced in this context by Luce (1959). This specific assumption is made mostly for computational convenience.

This model is known by many names. Among others, it is known as the rank-ordered logit model in economics (Beggs, Cardell, and Hausman 1981), as the exploded logit model in marketing research (Punj and Staelin 1978), as the choice-based conjoint analysis model (Hair et al. 2010), and as the Plackett–Luce model (Marden 1995).

The model coefficients are estimated using the method of maximum likelihood. The implementation in cmrologit uses an analogy between the rank-ordered logit model and the Cox regression model observed by Allison and Christakis (1994); see *Methods and formulas*. The cmrologit command implements this method for rankings, whereas clogit deals with the variant of choices; that is, only the most highly valued alternative is recorded. In the latter case, the model is also known as the Luce–McFadden choice model. In fact, when the data record the most preferred (unique) alternative and no additional ranking information about preferences is available, cmrologit and clogit return the same information, though formatted somewhat differently.

Before we can fit our model, we must cmset our data. The argument to cmset is the case ID variable, which must be numeric. For these data, it is the variable caseid, which identifies respondents. Unlike other choice models, alternatives are not specified; that is, there is no variable identifying specific alternatives across respondents. Alternatives simply have characteristics in this model.

```
. cmset caseid, noalternatives
      Case ID variable: caseid
Alternatives variable: <none>
```

We fit our model:

```
. cmrologit pref i.female age i.grades i.edufit i.workexp i.boardexp if job==1

Iteration 0:   Log likelihood =  -95.41087
Iteration 1:   Log likelihood = -71.004807
Iteration 2:   Log likelihood = -68.045946
Iteration 3:   Log likelihood = -67.906223
Iteration 4:   Log likelihood = -67.905657
Refining estimates:
Iteration 0:   Log likelihood = -67.905657
```

| Rank-ordered logit choice model | | | | Number of obs | = | 80 |
| Case ID variable: caseid | | | | Number of cases | = | 8 |

Ties adjustment: No ties in data

Obs per case:
```
                                                    min =         10
                                                    avg =      10.00
                                                    max =         10
```

```
                                        LR chi2(7)        =      55.01
Log likelihood = -67.90566              Prob > chi2       =     0.0000
```

pref	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
female						
yes	-.4554119	.3615579	-1.26	0.208	-1.164052	.2532285
age	-.0851138	.0822504	-1.03	0.301	-.2463216	.0760939
grades						
A/B	3.144718	.6200539	5.07	0.000	1.929434	4.360001
edufit						
yes	.7638799	.3613688	2.11	0.035	.0556102	1.47215
workexp						
internship	1.89448	.6298646	3.01	0.003	.6599679	3.128992
one year	2.9124	.6203927	4.69	0.000	1.696452	4.128347
boardexp						
yes	.8102527	.3972321	2.04	0.041	.031692	1.588813

Focusing only on the variables whose coefficients are significant at the 10% level (we are analyzing 8 respondents only!), we can write the estimated value of an applicant for a job of type 1 (research positions) as

```
value = 3.14*(A/B grades) + 0.76*edufit + 1.89*internship
        + 2.91*(1-year workexp) + 0.81*boardexp
```

Thus, employers prefer applicants for a research position (job==1) whose educational portfolio fits the job, who have better grades, who have more relevant work experience, and who have (extracurricular) board experience. They do not seem to care much about the sex and age of applicants, which is comforting.

Given these estimates of the valuation by employers, we consider the probabilities that each of the applications is ranked first. Under the assumption that the ϵ_i are independent and follow an extreme value type I distribution, Luce (1959) showed that the probability, π_i, that alternative i is valued higher than alternatives $2, \ldots, k$ can be written in the multinomial logit form

$$\pi_i = \Pr\{\text{value}_1 > \max(\text{value}_2, \ldots, \text{value}_m)\} = \frac{\exp(\text{value}_i)}{\sum_{j=1}^{k} \exp(\text{value}_i)}$$

The probability of observing a specific ranking can be written as the product of such terms, representing a sequential decision interpretation in which the rater first chooses the most preferred alternative, and then the most preferred alternative among the rest, etc.

The probabilities for alternatives to be ranked first are conveniently computed by `predict`.

```
. predict p if e(sample)
(option pr assumed; conditional probability that alternative is ranked first)
(210 missing values generated)
. sort caseid pref p
. list pref p grades edufit workexp boardexp if caseid==7, noobs
```

pref	p	grades	edufit	workexp	boardexp
0	.0021934	C/D	yes	none	yes
0	.0043086	C/D	no	internship	yes
0	.0051824	A/B	no	none	no
0	.0179498	C/D	yes	one year	no
1	.0403597	C/D	yes	one year	yes
2	.0226441	A/B	yes	none	no
3	.2642474	A/B	yes	one year	no
4	.0322894	A/B	yes	none	yes
5	.1505625	A/B	yes	internship	no
6	.4602626	A/B	yes	one year	yes

There clearly is a positive relation between the stated ranking and the predicted probabilities for alternatives to be ranked first, but the association is not perfect. In fact, we would not have expected a perfect association, because the model specifies a (nondegenerate) probability distribution over the possible rankings of the alternatives. These predictions for sets of 10 candidates can also be used to make predictions for subsets of the alternatives. For instance, suppose that only the last three candidates listed in this table would be available. According to parameter estimates of the rank-ordered logit model, the probability that the last of these candidates is selected equals $0.460/(0.032 + 0.151 + 0.460) = 0.715$.

Comparing respondents

The `cmrologit` model assumes that all respondents, HR managers in large public-sector organizations in The Netherlands, use the *same* valuation function; that is, they apply the same decision weights. This is the substantive interpretation of the assumption that the β's are constant between the respondents. To probe this assumption, we could test whether the coefficients vary between different groups of respondents. For a metric characteristic of the HR manager, such as `firmsize`, we can consider a trend model in the valuation weights,

$$\beta_{ij} = \alpha_{i0} + \alpha_{i1}\texttt{firmsize}_j$$

and we can test that the slopes α_{i1} of `firmsize` are zero.

```
. generate firmsize = employer

. cmrologit pref i.edufit i.grades i.workexp
> c.firmsize#(1.edufit 1.grades 1.workexp 2.workexp 1.boardexp)
> if job==1, nolog
```

Rank-ordered logit choice model			Number of obs	=	80
Case ID variable: caseid			Number of cases	=	8

Ties adjustment: No ties in data	Obs per case:		
		min =	10
		avg =	10.00
		max =	10
	LR chi2(9)	=	57.84
Log likelihood = -66.49266	Prob > chi2	=	0.0000

pref	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
edufit						
yes	1.163862	1.134384	1.03	0.305	-1.059489	3.387213
grades						
A/B	6.657894	2.401567	2.77	0.006	1.950908	11.36488
workexp						
internship	2.346752	1.908085	1.23	0.219	-1.393026	6.08653
one year	2.78297	1.794268	1.55	0.121	-.7337314	6.299672
edufit#						
c.firmsize						
yes	-.0118484	.0708709	-0.17	0.867	-.1507529	.127056
grades#						
c.firmsize						
A/B	-.217949	.1277188	-1.71	0.088	-.4682733	.0323753
workexp#						
c.firmsize						
internship	-.0353952	.107952	-0.33	0.743	-.2469773	.1761869
one year	.0078179	.1125715	0.07	0.945	-.2128183	.228454
boardexp#						
c.firmsize						
yes	.0426636	.0235837	1.81	0.070	-.0035596	.0888868

```
. testparm c.firmsize#(i.edufit i.grades i.workexp i.boardexp)

 ( 1)  1.edufit#c.firmsize = 0
 ( 2)  1.grades#c.firmsize = 0
 ( 3)  1.workexp#c.firmsize = 0
 ( 4)  2.workexp#c.firmsize = 0
 ( 5)  1.boardexp#c.firmsize = 0

          chi2(  5) =     7.48
        Prob > chi2 =   0.1871
```

The Wald test that the slopes of the interacted `firmsize` variables are jointly zero provides no evidence upon which we would reject the null hypothesis; that is, we do not find evidence against the assumption of constant valuation weights of the attributes by firms of different size. We did not enter `firmsize` as a predictor variable. Characteristics of the decision-making agent do not vary between alternatives. Thus, an additive effect of these characteristics on the valuation of alternatives does not affect the agent's ranking of alternatives and his or her choice. Consequently, the coefficient of `firmsize` is

not identified. `cmrologit` would in fact have diagnosed the problem and dropped `firmsize` from the analysis. Diagnosing this problem can slow the estimation considerably; the test may be suppressed by specifying the `notestrhs` option.

Incomplete rankings and ties

`cmrologit` allows incomplete rankings and ties in the rankings as proposed by Allison and Christakis (1994). `cmrologit` permits rankings to be incomplete only "at the bottom"; namely, that the ranking of the least attractive alternatives for subjects may not be known—do not confuse this with the situation that a subject is indifferent between these alternatives. This form of incompleteness occurred in the example discussed here because the respondents were instructed to select and rank only the top six alternatives. It may also be that respondents refused to rank the alternatives that are very unattractive.

`cmrologit` does not allow other forms of incompleteness, for instance, data in which respondents indicate which of four cars they like best, and which one they like least, but not how they rank the two intermediate cars. Another example of incompleteness that cannot be analyzed with `cmrologit` is data in which respondents select the three alternatives they like best but are not requested to express their preferences among the three selected alternatives.

`cmrologit` also permits ties in rankings. `cmrologit` assumes that if a subject expresses a tie between two or more alternatives, he or she actually holds one particular strict preference ordering, but with all possibilities of a strict ordering consistent with the expressed weak ordering being equally probable.

For instance, suppose that a respondent ranks alternative 1 highest. He or she prefers alternatives 2 and 3 over alternative 4 and is indifferent between alternatives 2 and 3. We assume that this respondent has the strict preference ordering either $1 > 2 > 3 > 4$ or $1 > 3 > 2 > 4$, with both possibilities being equally likely. From a psychometric perspective, it may actually be more appropriate to also assume that the alternatives 2 and 3 are close; for instance, the difference between the associated valuations (utilities) is less than some threshold or minimally discernible difference. Computationally, however, this is a more demanding model.

Clustered choice data

We have seen that applicants with work experience are in a relatively favorable position. To test whether the effects of work experience vary between the jobs, we can include interactions between the type of job and the attributes of applicants. Such interactions can be obtained using factor variables.

Because some HR managers contributed data for more than one job, we cannot assume that their selection decisions for different jobs are independent. We can account for this by specifying the `vce(cluster` *clustvar*`)` option. By treating choice data as incomplete ranking data with only the most preferred alternative marked, `cmrologit` may be used to estimate the model parameters for clustered choice data.

```
. cmrologit pref job##(i.female i.grades i.edufit i.workexp),
> vce(cluster employer) nolog
note: 2.job 3.job omitted because of no within-case variance.
```

Rank-ordered logit choice model	Number of obs	=	290
Case ID variable: caseid	Number of cases	=	29

```
Ties adjustment: efron                  Obs per case:
                                                     min =           10
                                                     avg =        10.00
                                                     max =           10
                                        Wald chi2(15)    =        93.03
Log pseudolikelihood =   -296.08        Prob > chi2      =       0.0000
```

(Std. err. adjusted for 22 clusters in employer)

pref	Coefficient	Robust std. err.	z	P>|z|	[95% conf. interval]	
job						
managemen..	0	(omitted)				
policy ad..	0	(omitted)				
female						
yes	-.2249088	.2396552	-0.94	0.348	-.6946244	.2448068
grades						
A/B	2.847569	.8087366	3.52	0.000	1.262475	4.432664
edufit						
yes	.6759788	.2318277	2.92	0.004	.2216049	1.130353
workexp						
internship	1.450487	.5518405	2.63	0.009	.3689	2.532075
one year	2.495849	.72438	3.45	0.001	1.07609	3.915608
job#female						
managemen.. #						
yes	.0227426	.4749959	0.05	0.962	-.9082322	.9537174
policy ad.. #						
yes	.1138279	.3559085	0.32	0.749	-.58374	.8113958
job#grades						
managemen.. #						
A/B	-2.381741	.9589772	-2.48	0.013	-4.261302	-.50218
policy ad.. #						
A/B	-1.92171	.8563522	-2.24	0.025	-3.60013	-.2432907
job#edufit						
managemen.. #						
yes	-.2427363	.4041935	-0.60	0.548	-1.034941	.5494684
policy ad.. #						
yes	-.2942018	.3578419	-0.82	0.411	-.9955589	.4071553
job#workexp						
managemen.. #						
internship	-1.183344	.7275346	-1.63	0.104	-2.609286	.2425972
managemen.. #						
one year	-1.495627	.7763694	-1.93	0.054	-3.017283	.0260292
policy ad.. #						
internship	-.6353285	.8809524	-0.72	0.471	-2.361963	1.091306
policy ad.. #						
one year	-.9522599	1.02415	-0.93	0.352	-2.959556	1.055036

The parameter estimates for the first job type are very similar to those that would have been obtained from an analysis isolated to these data. Differences are due only to an implied change in the method of handling ties. With clustered observations, `cmrologit` uses Efron's method. If we had specified the `ties(efron)` option with the separate analyses, then the parameter estimates would have been identical to the simultaneous results.

Another difference is that `cmrologit` now reports robust standard errors, adjusted for clustering within respondents. These could have been obtained for the separate analyses, as well by specifying the `vce(robust)` option. In fact, this option would also have forced `cmrologit` to switch to Efron's method as well.

Given the combined results for the three types of jobs, we can test easily whether the weights for the attributes of applicants vary between the jobs, in other words, whether employers are looking for different qualifications in applicants for different jobs. A Wald test for the equality hypothesis of no difference can be obtained with the `testparm` command:

```
. testparm job#(i.female i.grades i.edufit i.workexp)

 ( 1)   2.job#1.female = 0
 ( 2)   3.job#1.female = 0
 ( 3)   2.job#1.grades = 0
 ( 4)   3.job#1.grades = 0
 ( 5)   2.job#1.edufit = 0
 ( 6)   3.job#1.edufit = 0
 ( 7)   2.job#1.workexp = 0
 ( 8)   2.job#2.workexp = 0
 ( 9)   3.job#1.workexp = 0
 (10)   3.job#2.workexp = 0

           chi2( 10) =    18.01
         Prob > chi2 =    0.0548
```

We find only mild evidence that employers look for different qualities in candidates according to the job for which they are being considered.

❑ Technical note

Allison (1999) stressed that the comparison between groups of the coefficients of logistic regression is problematic, especially in its latent-variable interpretation. In many common latent-variable models, only the regression coefficients divided by the scale of the latent variable are identified. Thus, a comparison of logit regression coefficients between, say, men and women is meaningful only if one is willing to argue that the standard deviation of the latent residual does not differ between the sexes.

The rank-ordered logit model is also affected by this problem. While we formulated the model with a scale-free residual, we can actually think of the model for the value of an alternative as being scaled by the standard deviation of the random term, representing other relevant attributes of alternatives. Again, comparing attribute weights between jobs is meaningful to the extent that we are willing to defend the proposition that "all omitted attributes" are equally important for different kinds of jobs.

❑

Comparison of cmrologit and clogit

The rank-ordered logit model also has a sequential interpretation. A subject first chooses the best among the alternatives. Next, he or she selects the best alternative among the remaining alternatives, etc. The decisions at each of the subsequent stages are described by a conditional logit model, and a subject is assumed to apply the same decision weights at each stage.

Some authors have expressed concern that later choices may well be made more randomly than the first few decisions. A formalization of this idea is a heteroskedastic version of the rank-ordered logit model in which the scale of the random term increases with the number of decisions made (for example, Hausman and Ruud [1987]). This extended model is currently not supported by `cmrologit`. However, the hypothesis that the same decision weights are applied at the first stage and at later stages can be tested by applying a Hausman test.

First, we fit the rank-ordered logit model on the full ranking data for the first type of job.

```
. cmrologit pref age i.female i.edufit i.grades i.boardexp if job==1, nolog
Rank-ordered logit choice model              Number of obs    =          80
Case ID variable: caseid                     Number of cases  =           8

Ties adjustment: No ties in data             Obs per case:
                                                          min =          10
                                                          avg =       10.00
                                                          max =          10

                                             LR chi2(5)       =       26.15
Log likelihood = -82.33537                   Prob > chi2      =      0.0001
```

pref	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
age	-.1160121	.0767137	-1.51	0.130	-.2663682	.0343439
female						
yes	-.0791295	.3354377	-0.24	0.814	-.7365753	.5783163
edufit						
yes	.2475177	.3186609	0.78	0.437	-.3770462	.8720817
grades						
A/B	1.866381	.4542095	4.11	0.000	.9761463	2.756615
boardexp						
yes	-.0418455	.3194106	-0.13	0.896	-.6678788	.5841878

Second, we save the estimates for later use with the `estimates` command.

```
. estimates store Ranking
```

Third, to estimate the decision weights on the basis of the most preferred alternatives only, we create a variable, `best`, that is 1 for the best alternatives and 0 otherwise. The `by` prefix is useful here.

```
. by caseid (pref), sort: generate best = (pref == pref[_N]) if job==1
(210 missing values generated)
```

By specifying `(pref)` with `by caseid`, we ensured that the data were sorted in increasing order on `pref` within `caseid`. Hence, the most preferred alternatives are last in the sort order. The expression `pref == pref[_N]` is true (1) for the most preferred alternatives, even if the alternative is not unique, and false (0) otherwise.

If the most preferred alternatives were sometimes tied, we could still fit the model for the based-alternatives-only data via `cmrologit`, but `clogit` would yield different results because it deals with ties in a less appropriate way for continuous valuations. To ascertain whether there are ties in the selected data regarding applicants for research positions, we can combine `by` with `assert`:

```
. by caseid (pref), sort: assert pref[_N-1] != pref[_N] if job==1
```

There are no ties. We can now fit the model on the choice data by using either `clogit` or `cmrologit`.

```
. cmrologit best age i.edufit i.grades i.boardexp if job==1, nolog
Rank-ordered logit choice model                 Number of obs    =          80
Case ID variable: caseid                        Number of cases  =           8

Ties adjustment: No ties in data                Obs per case:
                                                              min =          10
                                                              avg =       10.00
                                                              max =          10

                                                LR chi2(4)       =        4.32
Log likelihood = -16.25952                      Prob > chi2      =      0.3641
```

best	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
age	-.0552291	.1887951	-0.29	0.770	-.4252607	.3148024
edufit						
yes	-.049965	.7530442	-0.07	0.947	-1.525905	1.425975
grades						
A/B	1.505808	1.11493	1.35	0.177	-.6794136	3.69103
boardexp						
yes	.995195	.8461853	1.18	0.240	-.6632977	2.653688

```
. estimates store Choice
```

The same results, though with a slightly different formatted header, would have been obtained by using `clogit` on these data.

```
. clogit best age i.edufit i.grades i.boardexp if job==1, group(caseid) nolog
Conditional (fixed-effects) logistic regression     Number of obs =        80
                                                    LR chi2(4)    =      4.32
                                                    Prob > chi2   =    0.3641
Log likelihood = -16.259518                         Pseudo R2     =    0.1173
```

best	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
age	-.0552291	.1887951	-0.29	0.770	-.4252607	.3148024
edufit						
yes	-.049965	.7530442	-0.07	0.947	-1.525905	1.425975
grades						
A/B	1.505808	1.11493	1.35	0.177	-.6794136	3.69103
boardexp						
yes	.995195	.8461853	1.18	0.240	-.6632977	2.653688

The parameters of the ranking and choice models look different, but the standard errors based on the choice data are much larger. Are we estimating parameters with the ranking data that are different from those with the choice data? A Hausman test compares two estimators of a parameter. One of the estimators should be efficient under the null hypothesis, namely, that choosing the second-best alternative is determined with the same decision weights as the best, etc. In our case, the efficient estimator of the decision weights uses the ranking information. The other estimator should be consistent, even if the null hypothesis is false. In our application, this is the estimator that uses the first-choice data only.

```
. hausman Choice Ranking
                   ─── Coefficients ───
                     (b)            (B)            (b-B)       sqrt(diag(V_b-V_B))
                   Choice        Ranking        Difference          Std. err.
         age     -.0552291      -.1160121         .060783            .1725068
   1.edufit       -.049965       .2475177       -.2974828            .6822982
   1.grades       1.505808       1.866381       -.3605724            1.018215
 1.boardexp        .995195      -.0418455        1.037041            .7835857

              b = Consistent under H0 and Ha; obtained from cmrologit.
       B = Inconsistent under Ha, efficient under H0; obtained from cmrologit.
Test of H0: Difference in coefficients not systematic
    chi2(4) = (b-B)'[(V_b-V_B)^(-1)](b-B)
            =    2.37
Prob > chi2 = 0.6672
```

We do not find evidence for misspecification. We have to be cautious, though, because Hausman-type tests are often not powerful, and the number of observations in our example is very small, which makes the quality of the method of the null distribution by a χ^2 test rather uncertain.

On reversals of rankings

The rank-ordered logit model has a property that you may find unexpected and even unfortunate. Compare two analyses with the rank-ordered logit model, one in which alternatives are ranked from "most attractive" to "least attractive", the other a reversed analysis in which these alternatives are ranked from "most unattractive" to "least unattractive". By unattractiveness, you probably mean just the opposite of attractiveness, and you expect that the weights of the attributes in predicting "attractiveness" to be minus the weights in predicting "unattractiveness". This is, however, *not* true for the rank-ordered logit model.

The assumed distribution of the random residual takes the form $F(\epsilon) = 1 - \exp\{\exp(-\epsilon)\}$. This distribution is right skewed. Therefore, slightly different models result from adding and subtracting the random residual, corresponding with high-to-low and low-to-high rankings. Thus, the estimated coefficients will differ between the two specifications, though usually not in an important way.

You may observe the difference by specifying the `reverse` option of `cmrologit`. Reversing the rank order makes rankings that are incomplete at the bottom become incomplete at the top. Only the first kind of incompleteness is supported by `cmrologit`. Thus, for this comparison, we exclude the alternatives that are not ranked, omitting the information that ranked alternatives are preferred over excluded ones.

```
. cmrologit pref grades edufit workexp boardexp if job==1 & pref!=0
  (output omitted)
. estimates store Original
. cmrologit pref grades edufit workexp boardexp if job==1 & pref!=0, reverse
  (output omitted)
. estimates store Reversed
```

```
. estimates table Original Reversed, stats(aic bic)
```

Variable	Original	Reversed
grades		
A/B	1.9842801	−1.2118836
edufit		
yes	−.16397461	−.11698355
workexp		
internship	.50190232	−.70220848
one year	2.3208572	−2.1518673
boardexp		
yes	.32930412	−.19747927
aic	97.817471	101.24092
bic	107.17348	110.59692

Thus, although the weights of the attributes for reversed rankings are indeed mostly of opposite signs, the magnitudes of the weights and their standard errors differ. Which one is more appropriate? We have no advice to offer here. The specific science of the problem will determine what is appropriate, though we would be surprised indeed if this helps here. Formal testing does not help much either because the models for the original and reversed rankings are not nested. The model-selection indices, such as the AIC and BIC, however, suggest that you stick to the rank-ordered logit model applied to the original ranking rather than to the reversed ranking.

Stored results

cmrologit stores the following in e():

Scalars
e(N)	number of observations
e(N_case)	number of cases
e(N_ic)	N for Bayesian information criterion (BIC)
e(N_clust)	number of clusters
e(df_m)	model degrees of freedom
e(ll)	log likelihood
e(ll_0)	log likelihood of the null model ("all rankings are equiprobable")
e(chi2)	χ^2
e(r2_p)	pseudo-R^2
e(p)	p-value for model test
e(code_inc)	value for incomplete preferences
e(alt_min)	minimum number of alternatives
e(alt_avg)	average number of alternatives
e(alt_max)	maximum number of alternatives
e(rank)	rank of e(V)
e(converged)	1 if converged, 0 otherwise

Macros
e(cmd)	cmrologit
e(cmdline)	command as typed
e(depvar)	name of dependent variable
e(caseid)	name of case ID variable
e(wtype)	weight type

e(wexp)	weight expression
e(marktype)	casewise or altwise, type of markout
e(key_N_ic)	cases, key for N for Bayesian information criterion (BIC)
e(title)	title in estimation output
e(clustvar)	name of cluster variable
e(offset)	linear offset variable
e(chi2type)	Wald or LR; type of model χ^2 test
e(reverse)	reverse, if specified
e(ties)	breslow, efron, exactm
e(vce)	*vcetype* specified in vce()
e(vcetype)	title used to label Std. err.
e(properties)	b V
e(predict)	program used to implement predict
e(marginsok)	predictions allowed by margins
e(marginsnotok)	predictions disallowed by margins
e(marginsdefault)	default predict() specification for margins
e(asbalanced)	factor variables fvset as asbalanced
e(asobserved)	factor variables fvset as asobserved

Matrices
e(b)	coefficient vector
e(V)	variance–covariance matrix of the estimators
e(V_modelbased)	model-based variance

Functions
e(sample)	marks estimation sample

In addition to the above, the following is stored in r():

Matrices
r(table)	matrix containing the coefficients with their standard errors, test statistics, p-values, and confidence intervals

Note that results stored in r() are updated when the command is replayed and will be replaced when any r-class command is run after the estimation command.

Methods and formulas

Allison and Christakis (1994) demonstrate that maximum likelihood estimates for the rank-ordered logit model can be obtained as the maximum partial-likelihood estimates of an appropriately specified Cox regression model for waiting time ([ST] **stcox**). In this analogy, a higher ranking of an alternative is formally equivalent to a higher hazard rate of failure. cmrologit uses stcox to fit the rank-ordered logit model based on such a specification of the data in Cox terms. A higher stated preference is represented by a shorter waiting time until failure. Incomplete rankings are dealt with via censoring. Moreover, decision situations (subjects) are to be treated as strata.

Finally, as proposed by Allison and Christakis, ties in rankings are handled by the marginal-likelihood method, specifying that all strict preference orderings consistent with the stated weak preference ordering are equally likely. The marginal-likelihood estimator is available in stcox via the exactm option. The methods of the marginal likelihood due to Breslow and Efron are also appropriate for the analysis of rank-ordered logit models. Because in most applications the number of ranked alternatives by one subject will be fairly small (at most, say, 20), the number of ties is small as well, and so you rarely will need to turn to methods to restrict computer time. Because the marginal-likelihood estimator in stcox does not support the cluster adjustment or pweights, you should use the Efron method in such cases.

This command supports the clustered version of the Huber/White/sandwich estimator of the variance using vce(robust) and vce(cluster *clustvar*). See [P] **_robust**, particularly *Maximum likelihood estimators* and *Methods and formulas*. Specifying vce(robust) is equivalent to specifying vce(cluster *caseid*), where *caseid* is the variable that identifies the cases.

Acknowledgment

References

Allison, P. D. 1999. Comparing logit and probit coefficients across groups. *Sociological Methods and Research* 28: 186–208. https://doi.org/10.1177/0049124199028002003.

Allison, P. D., and N. Christakis. 1994. "Logit models for sets of ranked items". In *Sociological Methodology*, edited by P. V. Marsden, vol. 24: 199–228. Oxford: Blackwell. https://doi.org/10.2307/270983.

Beggs, S., S. Cardell, and J. A. Hausman. 1981. Assessing the potential demand for electric cars. *Journal of Econometrics* 17: 1–19. https://doi.org/10.1016/0304-4076(81)90056-7.

de Wolf, I. 2000. *Opleidingsspecialisatie en arbeidsmarktsucces van sociale wetenschappers*. Amsterdam: ThelaThesis.

Hair, J. F., Jr., W. C. Black, B. J. Babin, and R. E. Anderson. 2010. *Multivariate Data Analysis*. 7th ed. Upper Saddle River, NJ: Pearson.

Hausman, J. A., and P. A. Ruud. 1987. Specifying and testing econometric models for rank-ordered data. *Journal of Econometrics* 34: 83–104. https://doi.org/10.1016/0304-4076(87)90068-6.

Luce, R. D. 1959. *Individual Choice Behavior: A Theoretical Analysis*. New York: Dover.

Marden, J. I. 1995. *Analyzing and Modeling Rank Data*. London: Chapman and Hall. https://doi.org/10.1201/b16552.

McCullagh, P. 1993. "Permutations and regression models". In *Probability Models and Statistical Analysis for Ranking Data*, edited by M. A. Fligner and J. S. Verducci, 196–215. New York: Springer. https://doi.org/10.1007/978-1-4612-2738-0_11.

Plackett, R. L. 1975. The analysis of permutations. *Journal of the Royal Statistical Society*, C ser., 24: 193–202. https://doi.org/10.2307/2346567.

Punj, G. N., and R. Staelin. 1978. The choice process for graduate business schools. *Journal of Marketing Research* 15: 588–598. https://doi.org/10.1177/002224377801500408.

Thurstone, L. L. 1927. A law of comparative judgment. *Psychological Reviews* 34: 273–286. https://doi.org/10.1037/h0070288.

Yellott, J. I., Jr. 1977. The relationship between Luce's choice axiom, Thurstone's theory of comparative judgment, and the double exponential distribution. *Journal of Mathematical Psychology* 15: 109–144. https://doi.org/10.1016/0022-2496(77)90026-8.

Also see

For suggested citations, see the FAQ on citing Stata documentation.