Paper 205-30

Using the Proportional Odds Model for Health-Related Outcomes: Why, When, and How with Various SAS® Procedures

Marc J. Gameroff, New York State Psychiatric Institute, New York, NY

ABSTRACT

Health-related outcomes often possess an intrinsic ordering but fail to meet the assumptions usually needed to perform an ordinary least-squares (OLS) regression. When the distribution of scores is highly non-normal, as occurs when the majority of respondents score at the very bottom or top of the scale, ordinal regression can be more valid, and sometimes more informative, than OLS regression. In SAS/STAT® version 9, PROC LOGISTIC allows one to fit a proportional odds model and to test the validity of interpreting significant predictors as affecting the outcome regardless of where the ordinal outcome might be divided into "higher" vs. "lower" scores. An example is provided based on an analysis of patient scores on the Sheehan Disability Scale, an oft-used outcome measure in mental health research. The model includes various mental disorder diagnoses as primary predictors and several sociodemographic variables as covariates. Results from the proportional odds model are compared to analogous results from an OLS regression obtained with PROC GLM. Other SAS/STAT procedures, such as PROC GENMOD and PROC PROBIT, can also be used to fit proportional odds models, and the differences in assumptions, modeling details, and available output will be described.

INTRODUCTION

The purpose of this paper is to explain the basic features and implementation of the proportional odds model (POM) in SAS. During the past few decades, several methods of ordinal logistic regression have been formalized. In these methods, the ordinality of a response variable having more than two levels can be incorporated into a logistic model. Many outcome variables are difficult, if not impossible, to measure on an interval scale. In biomedical research, for instance, constructs such as self-perceived health can be measured plausibly on an ordinal scale ("very unhealthy," "unhealthy", "healthy", "very healthy"), but numeric values assigned to these levels are at best arbitrary, and may lead to erroneous conclusions when analyzed as if they were equally-spaced points on a continuum (as in ordinary least-squares regression). When faced with this problem, one option is to dichotomize the ordinal outcome and run a binary logistic regression. However, the loss of information and decrease in statistical power are often too high a price to pay. Moreover, the resulting odds ratios may depend on the cut point chosen to dichotomize the outcome, and this choice is often arbitrary.

Ordinal logistic regression overcomes some of these problems. Like in binary and multinomial logistic regression, predictors may be categorical and/or continuous, and the computation of crude or adjusted odds ratios is the typical goal. The unique feature of the POM is that the odds ratio for each predictor is taken to be constant across all possible collapsings of the response variable. When a testable assumption is met, odds ratios in a POM are interpreted as the odds of being "lower" or "higher" on the outcome variable across the entire range of the outcome. The wide applicability and intuitive interpretation of the POM are two reasons for its being considered the most popular model for ordinal logistic regression.

This paper will be most useful to SAS users already familiar with binary and/or multinomial logistic regression as implemented in the LOGISTIC and GENMOD procedures. After a brief description of the conceptual basis of the POM, the syntax required to implement a POM will be described. This will be followed by an example from a study conducted in a primary care clinic, where strategies for building a predictive model of mental health-related disability were not clear-cut. SAS/STAT version 9 has been used for all the examples.

THE PROPORTIONAL ODDS MODEL

The proportional odds model (POM) described by McCullagh (1980) is the most popular model for ordinal logistic regression (Bender & Grouven, 1998). The POM is sometimes referred to as the cumulative logit model, however the latter is actually a more general term. In SAS, three types of cumulative logit models are available: the POM (available in the LOGISTIC, PROBIT, and GENMOD procedures), the *partial* proportional odds model (available in the GENMOD procedure), and the *non*proportional odds model (available in the CATMOD procedure). The hallmark of the POM is that the odds ratio for a predictor can be interpreted as a summary of the odds ratios obtained from separate binary logistic regressions using all possible cut points of the ordinal outcome (Scott et al., 1997). Whereas a binary logistic regression models a single logit, the POM models several cumulative logits. Therefore, if the ordinal outcome has four levels (1, 2, 3, and 4), three logits will be modeled, one for each of the following cut points: 1 vs. 2,3,4; 1,2 vs. 3,4; and 1,2,3 vs. 4. Because the various logits in a single POM are constrained to be equal, the POM has also been referred to as the constrained cumulative logit model (Hosmer & Lemeshow, 2000). Details on the statistical theory behind the POM can be found in several books and articles (see References). Agresti (1996) and Hosmer and Lemeshow (2000) are two sources that offer particularly clear explanations, and a SAS-specific approach can be found in sources such as Allison (2001).

THE PROPORTIONAL ODDS ASSUMPTION

For a POM to be valid, the assumption that all the logit surfaces are parallel must be tested. The standard test is a Score test that SAS labels in the output as the "Score Test for the Proportional Odds Assumption." A nonsignificant test is taken as evidence that the logit surfaces are parallel and that the odds ratios can be interpreted as constant across all possible cut points of the outcome.

There is some debate about the adequacy of the proportional odds test, e.g., it may be too sensitive to sample size (Scott et al., 1997), however it is the test used in standard practice. Whereas a significant test might be misleading, a nonsignificant test is usually taken as sufficient evidence that the POM is valid. If the test is significant, some have proposed additional, often graphical, methods for determining whether the logit surfaces are actually parallel. Other solutions include fitting a *partial* proportional odds model (see below for some general remarks), or dichotomizing the ordinal outcome. However, the latter option should be avoided if possible because of the loss in statistical power and the decreased generality of the analytic conclusions.

SAS/STAT PROCEDURES FOR FITTING A PROPORTIONAL ODDS MODEL

THE LOGISTIC PROCEDURE

PROC LOGISTIC is arguably the easiest and most all-purpose procedure to use in fitting a POM. When the response variable has more than two levels, SAS automatically treats it as an ordinal response and fits the POM. For instance, suppose that self-rated emotional health ("In general, how would you rate your emotional health?") is measured on a 5-point scale, with 4=Excellent, 3=Very good, 2=Good, 1=Fair, and 0=Poor. To model this response as a function of gender (1=female, 0=male), the following syntax is run:

```
proc logistic data=my.sugidata order=internal;
   class gender (param=ref ref='0');
   model emhealth = gender;
run;
```

Turning to the output, the "Response Profile" block shows the ordering of variable `emhealth` used in this analysis:

```
                    Response Profile

        Ordered                              Total
          Value      emhealth            Frequency

              1      Poor [0]                  101
              2      Fair [1]                  283
              3      Good [2]                  228
              4      Very good [3]              64
              5      Excellent [4]              55

   Probabilities modeled are cumulated over the lower Ordered Values.
```

It is important to ensure that the levels of the ordinal response are in fact being modeled ordinally. The default ordering for the response variable is `order=formatted`, and in order to illustrate this point, I applied a custom format to `emhealth` prior to running the syntax above (4="Excellent [4]", 3="Very good [3]", 2="Good [2]", 1="Fair [1]", 0="Poor [0]"). Had I omitted the `order=internal` option, the response levels would have been ordered alphabetically by their formatted values, with Ordered Value 1 being "Excellent [4]" and Ordered Value 5 being "Very good [3]", resulting in a decidedly non-ordinal response and a nonsensical analysis. Alternatively, I could have made sure that `emhealth` was unformatted, in which case the internal values (0, 1, 2, 3, 4) would have determined their order.

Another important matter to consider is the direction of the response levels. The meaning of the statement printed below the list of Ordered Values is that what is being modeled is the probability of being in a *lower* category in the Ordered Value list. Because in this example the lower Ordered Values denote *worse* health, what is being modeled is the probability of having *worse* emotional health. The documentation for PROC LOGISTIC presents several ways to reverse the order of the response levels, one of which would be to add the response variable option `DESCENDING` (alias: `DESC`) to the MODEL statement, in which case the analysis would model the probability of *better* emotional health, like so:

```
proc logistic data=my.sugidata order=internal;
   class gender (param=ref ref='0');
   model emhealth(descending) = gender;
run;
```

An attractive feature of the POM is that reversing the direction of the response levels will change the “direction” of the effects but not their magnitude or significance. Simply, the log-odds will have their signs reversed and the odds ratios will be inverted (i.e., odds ratio → 1/odds ratio).

Returning to the output, the test that assesses the validity of the proportional odds assumption is given:

```
          Score Test for the Proportional Odds Assumption

                Chi-Square        DF      Pr > ChiSq

                    2.3152         3          0.5096
```

The observed p-level does not lead to rejection of the proportional odds assumption.

The “Analysis of Maximum Likelihood Estimates” block includes an intercept for each logit. The intercept parameters quantify the “shift” in location between the four cumulative logits being modeled. The intercept parameters, however, are seldom of practical importance.

```
                                                Standard       Wald
       Parameter           DF      Estimate      Error     Chi-Square     Pr > ChiSq

Intercept Poor (0)          1      -2.3960      0.1653      210.1040        <.0001
Intercept Fair (1)          1      -0.4323      0.1377        9.8526        0.0017
Intercept Good (2)          1       1.1367      0.1454       61.1427        <.0001
Intercept Very good (3)     1       2.0188      0.1734      135.5770        <.0001
gender                      1       0.7124      0.1556       20.9610        <.0001
```

The odds ratio for gender (female vs. male) is exp(0.7124) = 2.039 and is statistically significant ($p<.0001$). The 95% confidence interval for the odds ratio is 1.503–2.766 (output not shown).

Because the proportional odds assumption was not rejected, the conclusion is that in this sample, females are approximately twice as likely to report worse emotional health than males. Importantly, this interpretation holds across the entire range of emotional health, from “poor” to “excellent.”

THE PROBIT PROCEDURE

Like the LOGISTIC procedure, the PROBIT procedure fits a POM by default when the response variable has more than two levels. However, when deciding which procedure to use, PROC PROBIT does not appear to have a clear advantage over PROC LOGISTIC unless the nature of the data is better suited to the unique capabilities of PROC PROBIT, specifically, the modeling of quantal response data. Still, by specifying the logistic distribution function in the model statement (`DISTRIBUTION=LOGISTIC`; aliases: `D` / `DIST` = `LOGISTIC`), PROC PROBIT produces results similar to the PROC LOGISTIC results above:

```
   proc probit data=my.sugidata order=internal;
      * Response variable must be listed in the CLASS statement;
      class emhealth;
      model emhealth = gender /distribution=logistic;
   run;
```

For gender, the parameter estimate (0.7124) is identical in PROC PROBIT and PROC LOGISTIC, and the standard error and chi-square are nearly identical (respectively, 0.1565 and 20.7317 with PROC PROBIT, and 0.1556 and 20.9610 with PROC LOGISTIC).

THE GENMOD PROCEDURE

PROC GENMOD can also be used to fit a POM. In order to do this, you must specify the multinomial distribution with the `DIST=MULTINOMIAL` option (aliases: `D` / `ERROR` / `ERR` = `MULT`), and the link function must be specified as `CUMLOGIT` (alias: `CLOGIT`).

```
proc genmod data=my.sugidata order=internal;
   class emhealth
         gender (param=ref ref='0');
   model emhealth = gender / dist=multinomial link=cumlogit;
run;
```

The parameter estimate, standard error, and chi-square for gender from this PROC GENMOD analysis are identical to those obtained with the PROC PROBIT analysis above.

Perhaps the most compelling reason to choose PROC GENMOD over PROC LOGISTIC is if the test of the proportional odds assumption is significant (i.e., the assumption of common slopes for all of the cumulative logits is not valid), and you decide to fit a *partial* proportional odds model (PPOM), which can only be done with PROC GENMOD. The PPOM loosens the constraint that *all* the predictors have a common parameter across the response logits. To implement a PPOM in PROC GENMOD, you must reconfigure the original data set so that each observation is expanded into *n* observations, where *n* is the number of logits being modeled. Further details are beyond the scope of this paper, but may be found, e.g., in Stokes et al. (2000).

EXAMPLE

Data for this example come from a mental health survey conducted in a primary care clinic serving a low-income community. Patients in the clinic's waiting room were systematically sampled over the course of several months. Further details of the study can be found in Olfson et al. (2000). In order to simplify the presentation, the analyses presented here are based on the subsample of participants who identified as Hispanic or Latino (n=731; 73% of the full sample). Participants had the choice of Spanish or English versions of all assessments. The goal is to model mental health-related disability (henceforth, "disability") in an effort to address the question "Which demographic and clinical characteristics are uniquely related to disability?"

Mental-health related disability was assessed with the Sheehan Disability Scale (SDS; Sheehan et al., 1996). The SDS has three 11-point items, and the respondent is asked to rate, on a numerical scale, the extent to which emotional problems have disrupted her/his *work*, *social life*, and *family life/home responsibilities* over the last month. Each item is rated from 0 to 10, with all numbers evenly spaced across the page. Text descriptions appear above the numbers. Written above 0 is "Not at all"; centered above 1, 2, and 3 is "Mildly"; centered above 4, 5, and 6 is "Moderately"; centered above 7, 8, and 9 is "Markedly"; and written above 10 is "Extremely." Scores for the three items are summed for a possible score of 0 to 30. Because only a small portion of the sample was gainfully employed, the *work* item was usually left blank and was therefore not used in the analyses. The modified version of the SDS used the *social life* and *family life/home responsibilities* items, for a possible total score of 0 to 20.

As can be seen in the Figure, the distribution of SDS scores was severely skewed. More than half the subjects had a total score of 0 (i.e., no disability). The floor effect commonly seen with the SDS has been discussed by McQuaid et al. (1999). Floor and ceiling effects for measures such as this are not uncommon. They are likely to occur with health-related and clinical measures, for instance, when the sample is not selected on the basis of having at least some degree of the construct being measured. A measure of *physical* health, for example, might be less skewed in a primary care sample such as this because most, if not all, individuals seeking primary care might be expected to have at least some physical health problems.

FIGURE. Scores on the Modified Sheehan Disability Scale (N=731)

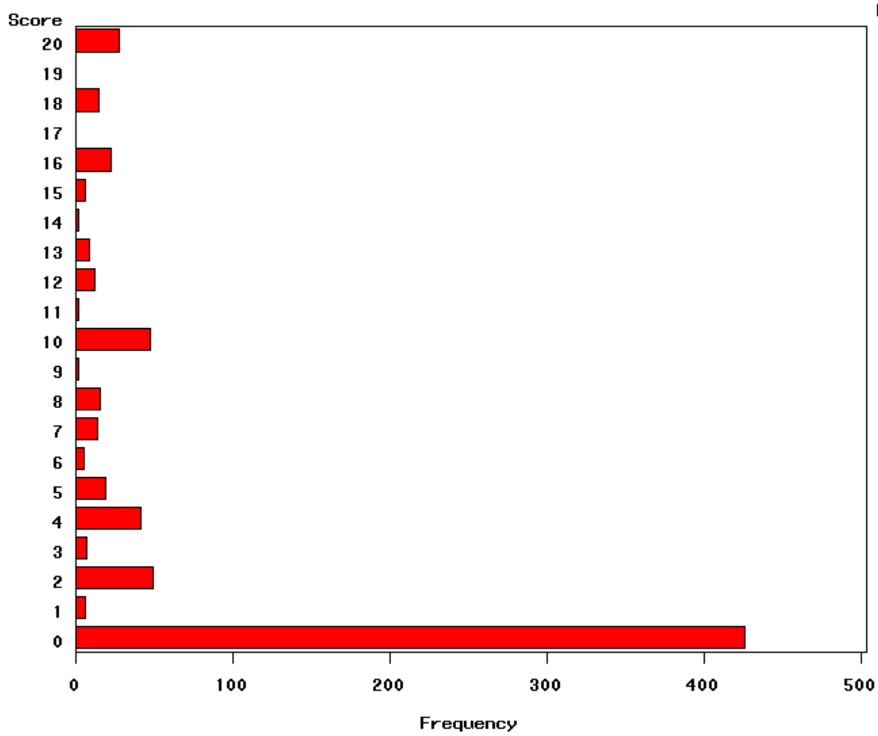

Score	FREQ.	CUM. FREQ.	PCT.	CUM. PCT.
20	28	28	3.83	3.83
19	0	28	0.00	3.83
18	15	43	2.05	5.88
17	0	43	0.00	5.88
16	23	66	3.15	9.03
15	6	72	0.82	9.85
14	2	74	0.27	10.12
13	9	83	1.23	11.35
12	12	95	1.64	13.00
11	2	97	0.27	13.27
10	48	145	6.57	19.84
9	2	147	0.27	20.11
8	16	163	2.19	22.30
7	14	177	1.92	24.21
6	5	182	0.68	24.90
5	19	201	2.60	27.50
4	42	243	5.75	33.24
3	7	250	0.96	34.20
2	49	299	6.70	40.90
1	6	305	0.82	41.72
0	426	731	58.28	100.00

The first potential problem with treating this variable as interval-level and modeling it with ordinary least-squares regression is the extreme skew. Obviously, there is no way to satisfactorily transform these scores to arrive at a normal distribution because 58% of the scores will always have the same (transformed) value. The second problem came to light upon inspection of the individual item distributions, which revealed that the numbers 2, 5, and 8 were much more commonly chosen than the numbers flanking them (i.e., 1, 3, 4, 6, 7, and 9). Because 2, 5, and 8 were the numbers *directly* below the text descriptions “mildly,” “moderately,” and “markedly,” it seemed clear that participants had taken notice of the words and that many were drawn more to them than to the numerical scale per se, at least in the 1-9 score range. These response patterns provided further evidence that SDS scores might best be modeled as an ordered categorical variable rather than an interval-level one.

I decided to transform the “continuous” version (`sds`) into an ordered categorical variable with five levels (`sds_5`):

```
if            sds = 0 then sds_5=0;
else if  1<=sds<= 6 then sds_5=1;
else if  7<=sds<=12 then sds_5=2;
else if 13<=sds<=18 then sds_5=3;
else if 19<=sds<=20 then sds_5=4;
```

The following is one of the basic models I ran during the model-building process. The first four class variables code the presence (vs. absence) of a positive screening result for that mental or addictive disorder (1=present):

```
proc logistic data=my.sugidata descending;
   class mdd_2 (param=ref ref='0')  /* Major depressive disorder (0,1)*/
         pan_2 (param=ref ref='0')  /* Panic disorder (0,1)*/
         gad_2 (param=ref ref='0')  /* Generalized anxiety disorder (0,1)*/
         adud_2 (param=ref ref='0') /* Alcohol or drug use disorder (0,1)*/
         gender (param=ref ref='Male');
   model sds_5 = mdd_2 pan_2 gad_2 adud_2 gender age
                 /expb parmlabel clodds=wald;
   units age=-10; /* Will give odds ratio for each decrement of 10 years */
run;
```

Selected output follows:

```
                         The LOGISTIC Procedure

                            Model Information

              Data Set                         MY.SUGIDATA
              Response Variable                sds_5
              Number of Response Levels        5
              Number of Observations           731
              Model                            cumulative logit
              Optimization Technique           Fisher's scoring


                             Response Profile

                    Ordered                          Total
                      Value          sds_5       Frequency

                          1              4              28
                          2              3              55
                          3              2              94
                          4              1             128
                          5              0             426

    Probabilities modeled are cumulated over the lower Ordered Values.
```

Score Test for the Proportional Odds Assumption

Chi-Square	DF	Pr > ChiSq
23.9739	18	0.1559

Model Fit Statistics

Criterion	Intercept Only	Intercept and Covariates
AIC	1766.981	1552.826
SC	1785.359	1598.771
-2 Log L	1758.981	1532.826

Testing Global Null Hypothesis: BETA=0

Test	Chi-Square	DF	Pr > ChiSq
Likelihood Ratio	226.1544	6	<.0001
Score	207.0935	6	<.0001
Wald	201.2175	6	<.0001

Type 3 Analysis of Effects

Effect	DF	Wald Chi-Square	Pr > ChiSq
mdd_2	1	58.0898	<.0001
pan_2	1	17.1458	<.0001
gad_2	1	15.7358	<.0001
adud_2	1	0.8144	0.3668
gender	1	2.9129	0.0879
AGE	1	21.9964	<.0001

Analysis of Maximum Likelihood Estimates

Parameter		DF	Estimate	Standard Error	Wald Chi-Square	Pr > ChiSq	Exp(Est)
Intercept	4	1	-3.2372	0.4387	54.4617	<.0001	0.039
Intercept	3	1	-1.8482	0.4055	20.7715	<.0001	0.158
Intercept	2	1	-0.6570	0.3942	2.7769	0.0956	0.518
Intercept	1	1	0.4013	0.3927	1.0442	0.3068	1.494
mdd_2	1	1	1.5414	0.2022	58.0898	<.0001	4.671
pan_2	1	1	1.0482	0.2531	17.1458	<.0001	2.852
gad_2	1	1	0.9053	0.2282	15.7358	<.0001	2.473
adud_2	1	1	0.2625	0.2908	0.8144	0.3668	1.300
gender	Female	1	0.3306	0.1937	2.9129	0.0879	1.392
AGE		1	-0.0295	0.00629	21.9964	<.0001	0.971

```
                 Analysis of Maximum Likelihood Estimates

                   Parameter         Label

                   Intercept 4       Intercept: sds_5=4
                   Intercept 3       Intercept: sds_5=3
                   Intercept 2       Intercept: sds_5=2
                   Intercept 1       Intercept: sds_5=1
                   mdd_2      1      MDD (0-1) 1
                   pan_2      1      PAN (0-1) 1
                   gad_2      1      GAD (0-1) 1
                   adud_2     1      adud_2 1
                   gender     Female Gender Female
                   AGE               Age (yrs)


            Wald Confidence Interval for Adjusted Odds Ratios

 Effect                        Unit    Estimate    95% Confidence Limits

 mdd_2  1 vs 0               1.0000       4.671       3.142        6.943
 pan_2  1 vs 0               1.0000       2.852       1.737        4.685
 gad_2  1 vs 0               1.0000       2.473       1.581        3.867
 adud_2 1 vs 0               1.0000       1.300       0.735        2.299
 gender Female vs Male       1.0000       1.392       0.952        2.035
 AGE                       -10.0000       1.343       1.188        1.520
```

Let's step through the output. Because for variable `sds_5` (as for variable `sds`), higher scores denote greater disability, I included the `DESCENDING` option on the PROC LOGISTIC statement so that SAS would model the probability of greater (rather than less) disability. The score test for the proportional odds assumption is not statistically significant (p=.16), so I decide not to reject the assumption of parallel slopes. Looking at the "Type 3 Analysis of Effects" and the "Wald Confidence Interval for Adjusted Odds Ratios" blocks, the variables in this model that predict greater self-reported disability over and above the other variables are (younger) age, major depressive disorder, panic disorder, and generalized anxiety disorder.

Perhaps not surprisingly, these four variables are the same ones that are statistically significant when running the same multivariable model in PROC GLM, but using `sds` (i.e., the original, "continuous" version) as the response.

While the conclusions from the PROC LOGISTIC and PROC GLM analyses are similar, let's examine how we might be able to phrase the conclusions. For simplicity, I will restrict my remarks to the adjusted effect of major depressive disorder (`mdd_2`) in each analysis:

SAS Procedure Used	**Outcome Variable**	**Conclusion (version A)** **(phrased in terms of participants)**	**Conclusion (version B)** **(phrased in terms of the disorder)**
PROC GLM	SDS ("continuous" with possible range of 0-20)	Participants with major depressive disorder, compared to those without, scored on average 4.37 (95% CI, 3.35-5.38) points higher on the Sheehan Disability Scale.	The presence of major depressive disorder was associated with an expected increase of 4.37 points (95% CI, 3.35-5.38) on the Sheehan Disability Scale.
PROC LOGISTIC (Proportional odds model)	SDS_5 (5 ordered categories)	Participants with major depressive disorder, compared to those without, were 4.67 (95% CI, 3.14-6.94) times as likely to report greater disability.	The presence of major depressive disorder was associated with a 4.67-fold (95% CI, 3.14-6.94) increased odds of reporting greater disability.

As a side point, it should be noted that the proximity of the numbers 4.37 (points) and 4.67 (odds ratio) is entirely coincidental. For instance, had the models included a different set of covariates, or if the possible total score range was narrower or wider than 0-20, these two numbers might be vastly different from each other.

Although the above results from PROC GLM and PROC LOGISTIC are not contradictory, the magnitude of 4–5 points on a 20-point disability scale is difficult to put into real-life terms, whereas the magnitude of 4–5 times the likelihood of having more disability is perhaps easier to apprehend and apply in a real-life context, especially as this statement is true across the entire range of disability (from “none” to “extreme”).

CONCLUSION

The proportional odds model is a convenient and easy-to-implement model available in SAS software for modeling ordered categorical outcomes that many not be appropriate for ordinary least-squares regression. Not covered in this paper are alternative ordinal regression models available in SAS, such as the adjacent categories model and the continuation-ratio model, which are more appropriate than the proportional odds model for certain types of analyses. SAS/STAT documentation should be consulted by anyone interested in fitting their first POM or other type of ordinal regression model in SAS.

REFERENCES

Agresti A (1996), *An Introduction to Categorical Data Analysis,* New York: Wiley.

Allison PD (2001), *Logistic Regression Using the SAS System: Theory and Application*, Cary, NC: SAS Publishing, BBU Press.

Bender R, Grouven U (1998), “Using binary logistic regression models for ordinal data with non-proportional odds,” *Journal of Clinical Epidemiology*, 51(10), 809-816.

Hosmer DW and Lemeshow S (2000), *Applied Logistic Regression* (2nd ed.), New York: Wiley.

McCullagh P (1980), “Regression models for ordinal data (with discussion),” Journal of the Royal Statistical Society, Series B, 42(2), 109-142.

McQuaid JR, Stein MB, Laffaye, C, McCahill, ME (1999), “Depression in a primary care clinic: the prevalence and impact of an unrecognized disorder,” *Journal of Affective Disorders*, 55, 1-10.

Olfson M, Shea S, Feder A, Fuentes M, Nomura Y, Gameroff M, Weissman MM (2000), “Prevalence of anxiety, depression, and substance use disorders in an urban general medical practice,” *Archives of Family Medicine*, 9(9), 876-883.

Scott SC, Goldberg MS, May NE (1997), “Statistical assessment of ordinal outcomes in comparative studies,” *Journal of Clinical Epidemiology*, 50(1), 45-55.

Sheehan DV, Harnett-Sheehan K, Raj BA (1996), “The measurement of disability,” *International Clinical Psychopharmacology*, 11(suppl 3), 89-95.

Stokes ME, Davis CS, Koch GG (2000), *Categorical Data Analysis Using the SAS System* (2nd ed.), Cary, NC: SAS Publishing, BBU Press and John Wiley Sons Inc.

ACKNOWLEDGMENTS

I would like to thank Dr. Myrna Weissman and Dr. Mark Olfson for allowing me to use the data from their study for the examples in this paper.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Marc Gameroff, Ph.D.
New York State Psychiatric Institute
1051 Riverside Drive – Unit 24
New York, NY 10032
Work Phone: (212) 543-5849
Fax: (212) 937-2142
Email: gameroff@childpsych.columbia.edu

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies.