CHAPTER ONE

THE PROPORTIONAL ODDS MODEL: A STATISTICALLY SOUND METHOD TO ANALYZE TURFGRASS RATING DATA

ABSTRACT

A common objective of many turfgrass experiments is to evaluate the

effects of various treatments on turf quality. Analysis of variance (ANOVA) has

traditionally been used to analyze quality rating data. However, many data sets

resulting from turf quality ratings have ordinal outcomes, defined as the ranking

of a set of observed values. These data violate assumptions required for valid

statistical inference from ANOVA since they are not continuous. The

development of the proportional odds model (POM) allows for valid statistical

inference on treatment effects from ordinal rating data. The POM also estimates

treatment parameters and standard errors, making treatment separation tests

and contrasts possible. These options were not available with the traditional

statistical tests appropriate for ordinal data. Unfortunately, to use the POM to its

full potential a researcher had to be an experienced statistical software

programmer, making it unusable for many. The objective of the following work

was to develop a Rating Data Analysis File Package (RDAFP) that (i) analyzes

ordinal rating data in a statistically valid manner using the POM, (ii) outputs

nearly the same amount of information on treatment effects as ANOVA, and (iii)

has an intuitively simple user-interface, from data entry to the production of

output. An example quality rating data set from a 4 x 2 factorial randomized

complete block design was used to demonstrate how the RDAFP analyzes data with the POM and outputs probability distribution charts into MS Excel. Complete analysis of the quality rating data with the POM, comparison of the results to ANOVA, and the production of probability distribution charts were possible with minimal SAS programming knowledge needed.

INTRODUCTION

The majority of turfgrass research is funded by groups interested in improving golf course, lawn care, or athletic field conditions. Therefore, an objective of many turfgrass experiments is to examine the effects of treatments on the functional or aesthetic quality of turf. This objective cannot be addressed without an evaluation of turfgrass quality by the researcher. Historically, quality evaluations have been done by visually rating plots on a scale of 1 to 9, where 1 = dead or brown turf, 6 = minimum acceptable quality (varying depending on the intended use of the turf), and 9 = ideal turf (dark green, dense, and uniform). The 1 to 9 scale was probably first used because of its practicality. Nine rating categories were usually adequate to distinguish quality differences observed among the turf plots, statistical calculations with values from 1 to 9 were relatively simple, and the results presented to the non-scientific community were comprehensible.

Quality rating data have different characteristics from data such as clipping yields that are obtained from an objective measuring device. Rating data resulting from the 1 to 9 scale will only have nine possible values (1, 2, ..., 8, 9). Seventeen values are possible if half steps are used (1, 1.5, ..., 8.5, 9). A typical quality rating may result in less than five unique observed values, whereas a clipping yield measurements usually result in a unique observed value for each experimental unit.

Another property of quality rating values is that they are arbitrary, since the values assigned to turfgrass plots are not from a standardized scale. An

alternate, but equally effective quality rating could be accomplished by using a scale of "A" to "I" where "A" represented ideal turf, "I" represented dead turf and "B through H" represented declining levels of turf quality intermediate to ideal and dead. However, a scale of this sort certainly could not be used to evaluate clipping yields. Clipping yields are measured with a standardized scale. For example, an observed clipping yield of 17.6 grams can be precisely comprehended by any turf researcher. It is obvious that quality rating data are of a different type than clipping yield data.

Classical statistical texts define the type of data resulting from quality ratings as *ordinal* (Freund and Wilson, 1993). Freund and Wilson define ordinal data as, "... a ranking or ordering of a set of observed values. Usually these ranks are assigned integer values starting with '1' for the lowest value, although other representations may be used." In contrast, clipping yield data is *continuous*, meaning that it can take on an infinite number of values within an interval (Freund and Wilson, 1993). Of course, an infinite number of values is limited by the precision of the measuring device.

Analysis of variance is a popular statistical tool because of the relatively large amount of information obtained from the data compared to other statistical analyses. Global hypothesis testing, treatment mean estimation, and treatment mean separation tests can all be accomplished using ANOVA techniques. In contrast, traditional statistical tests appropriate for ordinal data (Kruskal-Wallis test, Friedman test, or Spearman correlation) only test the global hypothesis of treatment equality. The relative weakness of these tests, as well as the better

comprehension of ANOVA calculations by most turf researchers may account for the frequent use of ANOVA for rating data.

Analysis of variance is only valid on continuous data, and only if the data: 1) result from a linear combination of the treatment effects and random error, 2) error values are random and from a Gaussian distribution with mean = 0 and variance = σ^2, and 3) data values are from independent and random samples (Freund and Wilson, 1993). In addition to violating the continuous data stipulation, rating data often violate the second assumption of ANOVA. Since visual quality ratings usually lead to few unique outcomes (typical rating data may have a minimum value of "5" and maximum value of "8"), the error values do not approximate a Gaussian distribution well. Furthermore, the analysis used to analyze rating data should accommodate whatever rating scale is used by the researcher. It would be impossible to use ANOVA if an "A" to "I" scale was used to rate quality. Despite these statistical flaws, ANOVA has been used to analyze turf rating data for decades.

McCullagh and Nelder (1980) described POMs capable of predicting ordinal responses from independent variables. These models yield nearly the same amount of treatment information as ANOVA. However, calculations of treatment effects and standard errors are complex, and typically require programming of statistical software.

Recently, Schabenberger et al. (2000) authored SAS® macros that produce global hypothesis tests, treatment comparisons, and contrasts that resemble ANOVA output. The complex SAS environment and macro

programming language may deter many turf researchers from using the macro. The development of a simplified user-interface for this SAS macro may result in more turf researchers using it to analyze rating data.

The objective of the following work was to develop a Rating Data Analysis File Package (RDAFP) that (i) analyzes ordinal rating data in a statistically valid manner using the POM, (ii) outputs nearly the same amount of information on treatment effects as ANOVA, and (iii) has an intuitively simple user-interface, from data entry to the production of output.

STATISTICAL METHODS

Ordinal data are traditionally analyzed by non-parametric methods that only test global hypotheses of treatment equality. Logistic regression models, first used in the 1940's to analyze bioassay data (McCullagh and Nelder, 1989) , estimate the probability of a response based on predictor variables. Because the model estimates probabilities rather than mean rating values, it is applicable regardless of the rating scale used by the researcher. Logistic regression models have gained popularity in the last 20 years, paralleling the refinement of mathematical techniques used in their calculation. Kleinbaum (1994) presents an overview of logistic regression models with applied examples in a format palatable to the non-statistician.

McCullagh (1980) described the POM, specialized for the analysis of ordinal data. The POM involves parallel logistic regressions that estimate the probabilities of an observation to fall into the ordered response categories, based

on the values of independent predictor variables. The RDAFP discussed later uses the POM to analyze rating data.

The POM estimates a value, ranging from $-\infty$ to ∞, for each parameter in the model. For rating data, the model parameters consist of the independent treatment variables, their interactions, and all observed rating categories. For example, a completely randomized design with treatment factors A (with 2 levels) and B (with 3 levels), and observed rating values of "5", "6", "7", and "8" would have values estimated for the following parameters: α_1, α_2, β_1, β_2, β_3, $\alpha_1\beta_1$, $\alpha_1\beta_2$, $\alpha_1\beta_3$, $\alpha_2\beta_1$, $\alpha_2\beta_2$, $\alpha_2\beta_3$, π_5, π_6, π_7, and π_8. Most software packages will estimate treatment and category effects as differences from a reference level. Therefore, the parameter estimates for the first treatment levels (α_1, β_1, and $\alpha_1\beta_1$) and the highest ranking rating category level (π_8) will be zero.

Parameter estimates are calculated by maximum likelihood techniques (Shenton and Bowman, 1977). Maximum likelihood calculations result in parameter estimates that best predict the observed values in the data set. Calculations involve iterative, re-weighted, differentiation of likelihood functions and become very complex with few model parameters. However, with the development of powerful PC processors, maximum likelihood calculations have become commonplace.

A latent variable, *Z*, represents a linear combination of the parameter estimates for the treatment and rating category combination of interest. For example, if a researcher was interested in the probability of a turf plot receiving level 1 of factor A and level 3 of factor B being rated *at best* a 7, then $Z = \alpha_1 + \beta_3$

+ $\alpha_1\beta_3 + \pi_7$. Since the parameter estimates range from $-\infty$ to ∞, Z must also share this range. However, probabilities are constrained between zero and one.

A logit-link function is used to transform Z values into probability predictions by the following equation: $1 / (1 + e^{-z})$, where e is Euler's number (2.178). This function has a range of zero to one, regardless of the value of Z. Plugging the sum of the treatment parameter estimates of interest into the logit-link function will result in a cumulative probability, the probability to be rated, at best, in a given category. Individual category probabilities are calculated by differencing cumulative probabilities for two adjacent rating categories. For example, the probability of the treatment described above to be rated *exactly* a "7" is calculated by:

$$\{ [1 / (1 + e^{-(\alpha_1 + \beta_3 + \alpha_1\beta_3 + \pi_7)})] - [1 / (1 + e^{-(\alpha_1 + \beta_3 + \alpha_1\beta_3 + \pi_6)})] \}$$

A variance-covariance matrix for the parameter estimates can be produced by maximum likelihood calculations in computer software programs. From this matrix and the parameter estimates, statistical tests can be performed on the equality between any combination of treatment levels. These tests give information similar to the mean comparison tests and pre-planned contrasts often used with ANOVA.

RATING DATA ANALYSIS FILE PACKAGE

Excel version 95 or later (Microsoft, 1995) and SAS release 6.12 (SAS Institute, 1996) must be installed on the user's PC to use the file package described in this paper. The files needed to analyze, output, and graph rating

data are bundled in an installation program called "RDAFP.exe". This program can be downloaded from "Rating Data Analysis File Package" web page at the URL, http://www.msu.edu/~karcherd/ratings (Karcher, 2000).

Running "RDAFP.exe" will create a directory called "Ordinal Analysis" on the C drive of the user's PC. The files, "PropOddsModel.sas" (Schabenberger et al., 2000), "turfrate.sas" (Appendix A), "Rating Charts.xlt", and "readme.txt", are all placed in the "Ordinal Analysis" directory. Additionally, a shortcut to "Rating Charts.xlt" is placed on the PC Desktop during installation.

"Rating Charts.xlt" is an MS Excel template that produces probability distribution charts from data output by SAS. The "readme.txt" file is a text file containing detailed instructions for RDAFP and covers installation through interpretation of results. The "turfrate.sas" file (Appendix A) was created by the author to run the Schabenberger et al. (2000) macro from easy to use web based forms and the downloaded MS Excel template.

Once the installation program is completed, a data file needs to be created. Although data files can be created in either MS Excel or SAS, using MS Excel simplifies the analysis process. If the data file is created in MS Excel, variable names must be in the first row of the spreadsheet and the data values must begin in the second row (Figure 1). There cannot be any blank rows within the data when using MS Excel. Additionally, an MS Excel data file must be saved in an Excel 95 file format (Figure 1). Finally, the MS Excel file must be closed during analysis since the data cannot be imported into SAS if the file is left open.
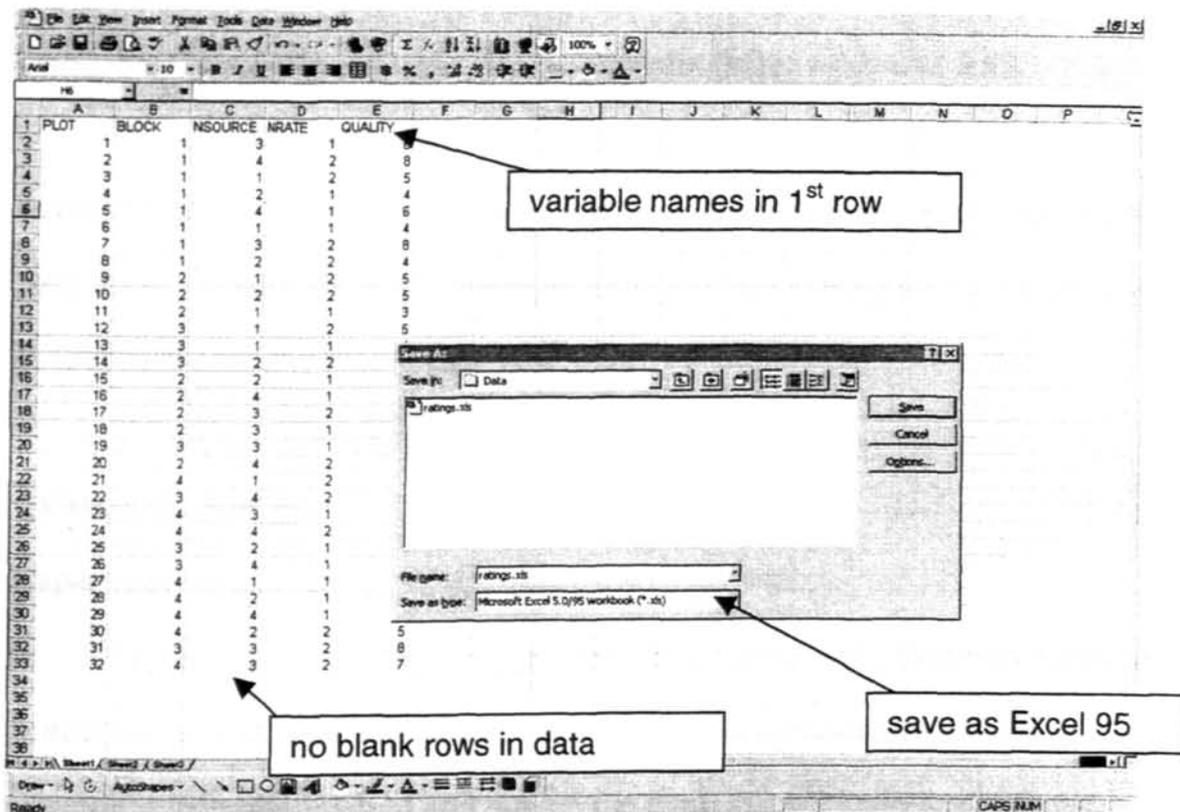
Figure 1. Rating data file created in MS Excel.

The SAS code needed to analyze and generate probability charts from the rating data can be generated by the RDAFP web page (Karcher, 2000). This is a form based web page where the user answers a few questions regarding the experimental design and treatment structure of the study that generated the rating data. After answering all questions, clicking the "Generate SAS Code" button will generate code in a separate window titled, "SAS Code for Ordinal Data Analysis" that is ready for pasting into the Program Editor window of SAS.

Copying the text from the web window into SAS is relatively simple: 1) activate the output window by clicking it with the mouse pointer, 2) drag over the code with the mouse to select, 3) press Ctrl+C to copy all of the code, 4) activate

the Program Editor window of SAS by clicking it with the mouse pointer, and 5) press Ctrl+V to paste the code into the Program Editor window of SAS.

If probability distribution charts created in MS Excel are desired, "Shortcut to Rating Charts" should be opened from the PC's Desktop prior to executing the SAS code. The SAS code is executed by activating the Program Editor window in SAS and pressing the F8 key.

EXAMPLE USAGE

Experimental Design

The data used in this example are from a quality rating taken on a nitrogen fertilization study. The objective of the study was to compare the effects of nitrogen application method and rate on the quality of a 'Penncross' creeping bentgrass (*Agrostis palustris* Huds.) putting green. The application methods included nitrogen injection using high pressure water injection cultivation (WIC) (Murphy and Rieke, 1994) and traditional surface applications. The study was 4 x 2 factorial with four replications in a randomized complete block design. The first factor was application method, having four levels: 1) surface sprayed N, no WIC, 2) surface sprayed N, followed by WIC, 3) N applied via WIC to a 7.5 cm depth, and 4) N applied via WIC to a 15 cm depth. The second factor was N rate, having two levels: 1) 2.4 and 2) 4.8 g N m^{-2} application^{-1}.

Generating SAS Code

Figure 2 shows the RDAFP web page for downloading the file package and generating SAS code for analysis. Step #1 on the page instructs the user to

click the installation icon with the mouse pointer and download the file package. After downloading and running the installation program, the form on the web page can be used to generate SAS code.

Steps #2 through #6 on the form **must** be completed in order to generate the proper SAS code. Step #2 defines the path on the researcher's PC to the MS Excel file containing the rating data, which in this case was "C:\DATA FILES\RATINGS.xls" (Figure 1). Step #3 defines the treatment factors as they are named in the MS Excel data file. The treatment factors in the data file created for this study were named NSOURCE and NRATE (Figure 1). Step #4 defines the blocking factor, if present, as it is named in the MS Excel data file, which was BLOCK in this case (Figure 1). Step #5 defines the response variable, as named in the MS Excel data file. The response variable was named QUALITY in this situation (Figure 1).

Step #6 gives the user an opportunity to label the observed data values from the rating. Possible labels for a typical quality rating scale are (1=dead, 2=mostly dead, 3=severely flawed, 4=flawed, 5=slightly flawed, 6=acceptable, 7=good, 8=excellent, 9=ideal). For this study, "3", "4", "5", "6", "7", and "8" were the only observed quality rating values. They were labeled here as "SEVFLAWED", "FLAWED", "SLIFLAWED", "ACCEPTABLE", "GOOD", AND "EXCELLENT". Labels should not contain any spaces or special characters and be relatively short in order to accommodate output. Labels provide the researcher with an opportunity to describe the basis for rating the turf and are typically more informative than arbitrary numbers.

Figure 2. SAS code generating form, completed with information from nitrogen application method study. This form is from the "Rating Data Analysis" web page at "http://www.msu.edu/~karcherd/ratings".

Steps #7 through #13 are extra analysis options. Step #7 gives an option to the user to define a reduced model. By default, a full model is used that contains all treatment factors and all possible interactions. When the number of treatment factors and observed rating categories is large relative to the number

of observations in the data set, a full model might result in errors during

maximum likelihood calculations. The following message (Figure 3) appears in

the Log window of SAS when maximum likelihood errors occur:

```
WARNING: There is possibly a quasicomplete separation in the sample points.  The
         maximum likelihood estimate may not exist.
WARNING: The LOGISTIC procedure continues in spite of the above warning. Results
         shown are based on the last maximum likelihood iteration. Validity of the
                    model fit is questionable.
```

Figure 3. Warning message that appears in the Log window of SAS when errors occur during
maximum likelihood calculation.

Checking "Yes" in Step #7 will cause a text prompt to appear upon clicking

the "Generate SAS Code" button (Figure 4). The user may define a reduced

model in this text prompt if maximum likelihood errors occur when analyzing the

full model.

The example data set had 32 observations and 24 (4 BLOCK, 4

NSOURCE, 2 NRATE, 8 NSOURCE x NRATE, and 6 QUALITY) parameter

estimates in the full model. The full model resulted in maximum likelihood errors,

causing the error message in Figure 3 to be printed in the Log window of SAS.

Therefore, a reduced model was used by dropping the NSOURCE X NRATE

interaction term from the full model (Figure 4). No error messages resulted from

analyzing the reduced model.
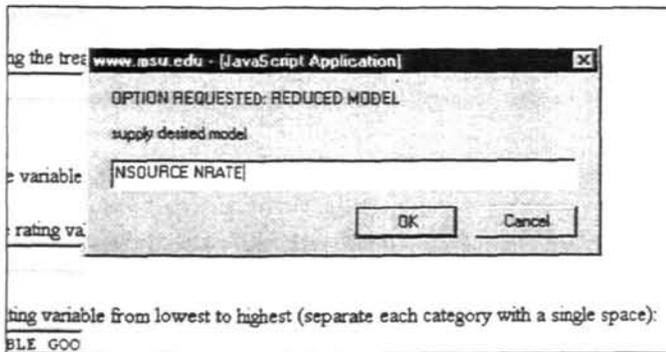
Figure 4. Textbox used to supply a reduced model. In this case, the NSOURCE x NRATE interaction term has been dropped from the full model.

Step #8 defines if probability distributions for treatment main effects are to be graphed using the "Rating Charts" template in MS Excel. The default is "Yes" and the template must be opened prior to executing code in SAS if "Yes" is checked. Step #9 defines if probability distributions for an interaction term are to be graphed in the "Rating Charts" template. The default for this option is "No". Interaction distributions can only be graphed after main effects analysis has been executed in SAS. Following main effects analysis, if "Yes" is checked in Step #9, a text prompt will appear for the user to define the interaction term for which probability distributions are to be graphed in MS Excel.

Step #10 allows the user to compare results produced by POM analysis with results obtained from ANOVA. The default for Step #10 is "No". Checking Step #10 will produce side-by-side tests of fixed effects and mean separation tests from the POM and ANOVA.

Steps #11 and #12 give the user an opportunity to test treatment contrasts and slice interaction terms by user-defined effects. These are the only options that require SAS programming knowledge by the user. Checking these steps will

produce text prompts where the user must provide the proper "contrast" or "slice" statement, using syntax identical to that of SAS's glm procedure.

Clicking the "Generate SAS Code" button with the mouse pointer will generate a new window that contains SAS code (Figure 5). The code shown in Figure 5 is for the reduced model. This code must be pasted into the SAS Program Editor window before POM analysis can take place. This is accomplished by: 1) selecting the code below the horizontal rule with the mouse, 2) pressing Ctrl+C to copy the code, 3) clicking inside the Program Editor window of SAS with the mouse pointer, and 4) pressing Ctrl+V to paste the code. After pasting the code into the SAS Program Editor window, pressing the F8 key will execute the "turfrate.sas" macro that uses the POM and all other files in the RDAFP to generate output.
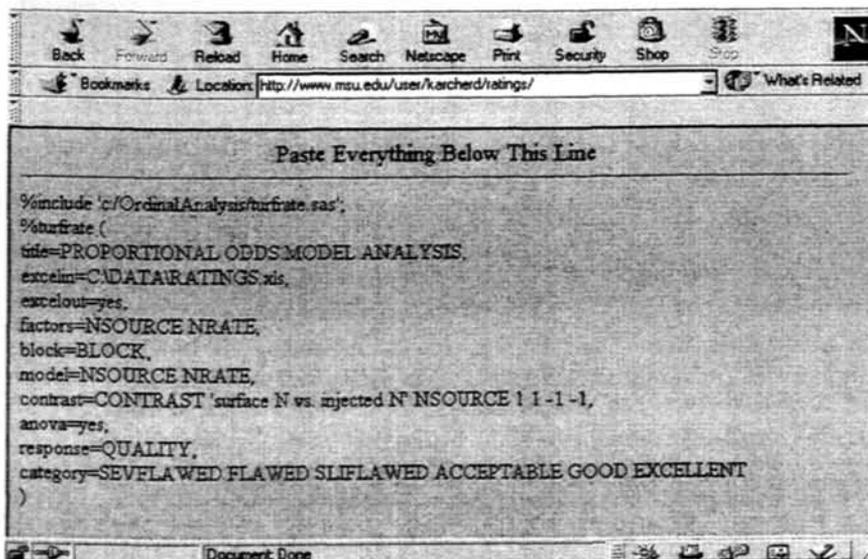
Figure 5. Output generated from the "Rating Data Analysis" web page form. This code can simply be pasted into SAS v. 6.12 for expedient analysis using the POM.

Summary of Output

Executing the code in Figure 5 produced the output shown in Figure 6
through Figure 10. Figure 6 shows the comparison of fixed effects tests between
the POM and ANOVA. The output shows the degrees of freedom and computed
chi-square and F values used to determine the respective *P*-values for each
statistical test. In this example, nitrogen application method and nitrogen rate
significantly affected turf quality. The two statistical analyses produced
remarkably similar *P*-values for NSOURCE ($P = 0.0001$) and NRATE ($P = 0.0003$) effects.

Figure 6. Tests of fixed effects produced by RDAFP. Both application method (NSOURCE) and
nitrogen rate (NRATE) effects were highly significant when analyzed by the POM and
ANOVA.

Predicted probabilities for each treatment level to be rated into each
quality category and a comparison of mean separation tests between the POM
and ANOVA for NSOURCE are shown in Figure 7. Treatment #3, which
corresponded to nitrogen injected to a 7.5 cm depth, had the highest probability

(41%) to be rated as excellent. Conversely, Treatment #2, which corresponded

to surface applications of nitrogen plus WIC, had the highest probability (6%) to

be rated as severely flawed. From the treatment separation tests following

analysis by the POM , treatments #3 and #4 (followed by A's) were significantly

different from treatments #1 and #2 (followed by B's). Examination of the

category probabilities for the treatment levels reveals that treatments #3 and #4

produced significantly higher quality than treatments #1 and #2. Similar results

were calculated by a post ANOVA LSD test.

PROPORTIONAL ODDS MODEL
PROPORTIONAL ODDS MODEL ANALYSIS

Predicted Category Probabilities [P = RATING CATEGORY]
for NSOURCE effects

OBS	NSOURCE	SEVFLAWED	FLAWED	SLIFLAWED	ACCEPTABLE	GOOD
1	3	0.000024	0.00129	0.02001	0.04531	0.52714
2	4	0.000125	0.00659	0.09417	0.16792	0.61393
3	1	0.026858	0.56843	0.36577	0.02706	0.01177
4	2	0.057575	0.71054	0.21401	0.01235	0.00525

OBS	EXCELLENT	POM LS Mean	ANOVA LS Mean	POM Mean Groups	ANOVA Mean Groups
1	0.40623	-4.67000	7.25000000	A	A
2	0.11727	-3.03108	6.50000000	A	A
3	0.00061	2.34905	4.75000000	B	B
4	0.00027	3.16295	4.25000000	B	B

Figure 7. Probability of each NSOURCE treatment to be rated into each quality rating category, as well as mean separation tests from the POM and ANOVA.

Category probabilities are easier to compare among treatments using a

probability distribution chart. The probability distribution chart shown in Figure 8

was created automatically in the "Rating Charts.xlt" MS Excel template.

18

Treatments with larger white bars were poorer in quality than treatments with larger dark gray and black bars. Figure 8 demonstrates that analysis by the POM yields a greater amount of information regarding treatment effects than the arbitrary mean rating values produced by ANOVA.
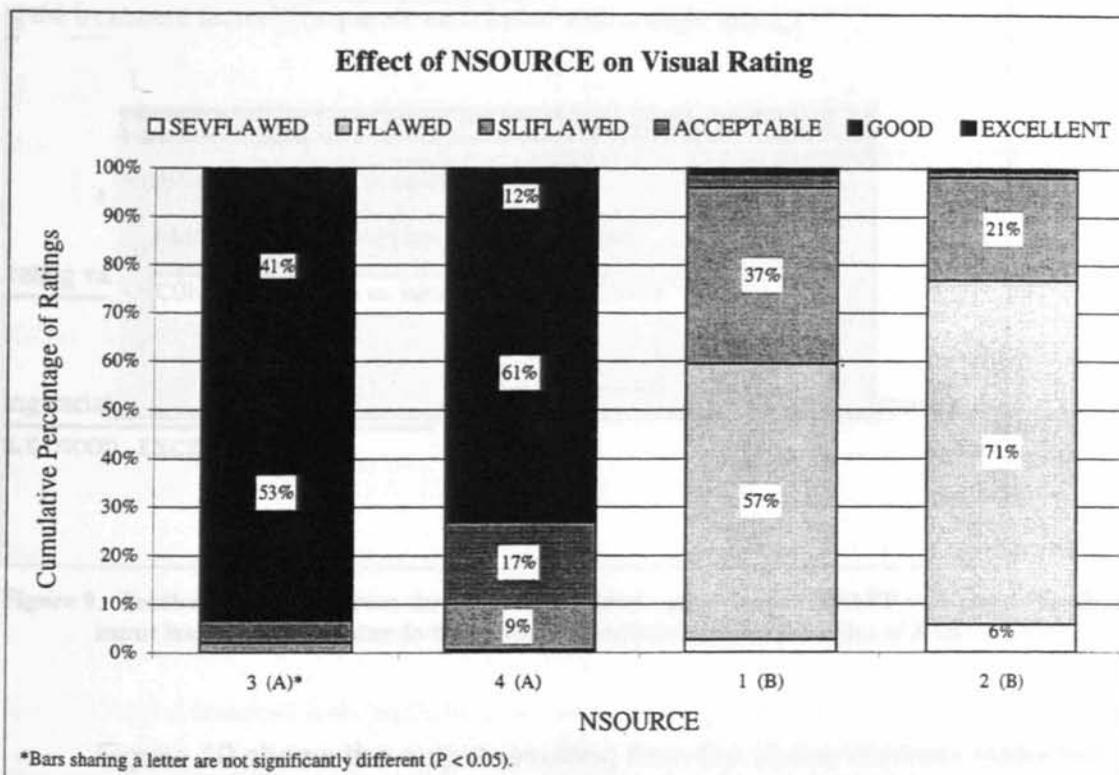
Figure 8. Probability distribution chart created automatically by RDAFP. Cumulative probabilities are shown on the y-axis, whereas individual category probabilities (greater than 5%) are labeled within each bar section.

Treatments #1 and #2 corresponded to turf receiving surface applications of nitrogen, whereas treatments #3 and #4 correspond to turf injected with nitrogen. A contrast testing equality between the treatment groups would test the effects of injecting nitrogen vs. surface applications of nitrogen on turf quality. A hypothesis test comparing treatments #1 and #2 vs. #3 and #4 was accomplished by checking "Yes" in Step #7 on the RDAFP web page and

inputting the appropriate contrast statement (Figure 9). The contrast statement

in the textbox has identical syntax to the contrast statement used in SAS's proc

glm (minus a semi-colon).

g the treatment factor(s)(separate each factor with a single space):

variable

rating va

www.msu.edu – [JavaScript Application]

OPTION REQUESTED: CONTRAST TEST

supply contrast statement [see SAS documentation]

CONTRAST 'surface vs. injected nitrogen' NSOURCE 1 1 -1 -1

OK   Cancel

ing variab   space):

LE GOOD EXCELLENT
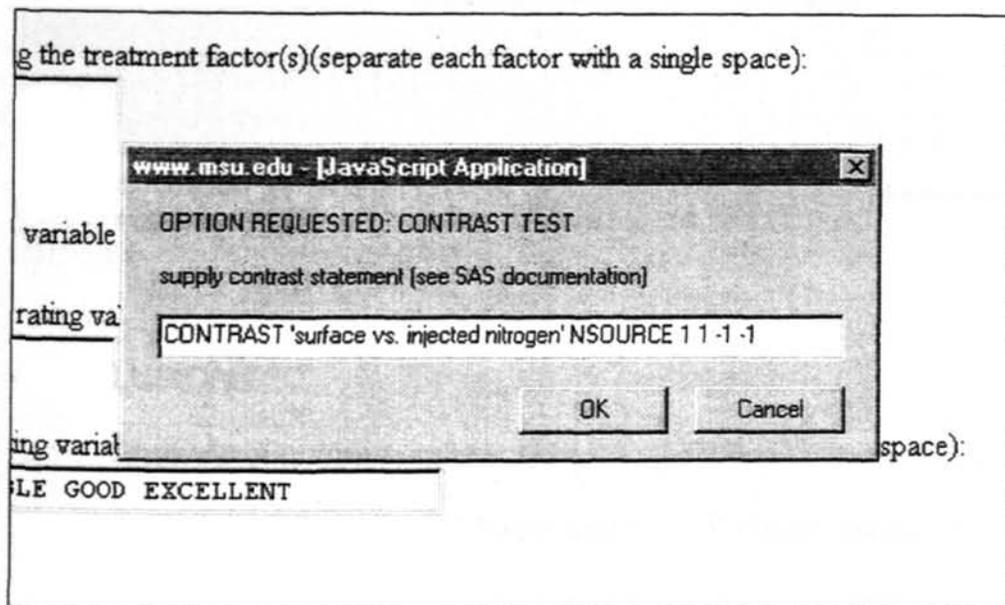
Figure 9. Textbox generated from checking the contrast option on the RDAFP web page. Textbox input has identical syntax to the contrast statement used in proc glm of SAS.

Figure 10 shows the output resulting from the above contrast statement.

Whether nitrogen was applied on the surface or injected significantly affected turf

quality ($P < 0.001$). Caution must be exercised when interpreting contrast

results. A *positive* chi-square value means that treatments corresponding to

negative coefficients in the contrast statement had higher ratings. This results

because probabilities calculated from the logit-link function increase as Z values

decrease. In this example, the negative coefficients correspond to treatments #3

and #4, which correspond to injected nitrogen. Since the chi-square value was

positive (20.59), these treatments had significantly higher ratings.

```
                    PROPORTIONAL ODDS MODEL ANALYSIS

                            Test of Contrasts
                            Data Set = data


                                                      Degrees
                                         Chi-Square      of        P >
        OBS     Source of Variation        Value      Freedom   Chi-Square


         1    SURFACE N VS. INJECTED N     20.5878        1      .0000056959
```

Figure 10. Contrast test from the RDAFP. Here, treatments receiving surface nitrogen were significantly different ($P < 0.001$) than turf injected with nitrogen.

CONCLUSIONS

Analysis of turfgrass quality data with the RDAFP was a simple process. The data file was created in MS Excel and the SAS code needed to run the RDAFP was generated from an intuitive web based form. After generating SAS code from the web and pasting it into the SAS Program Editor window, pressing the F8 key executed the "turfrate.sas" macro. This macro accessed the other files in the RDAFP to import the data from an MS Excel data file, analyze the data using the POM, perform treatment separation tests, and output probability distributions to the "Rating Charts" template in MS Excel. This occurred without the need to program any SAS code. The only SAS knowledge needed was how to paste code into the Program Editor window, and then press the F8 key to execute the pasted code.

The RDAFP has potential to be a valid, user-friendly data analysis tool for researchers in other agricultural sciences when data is acquired from subjective, qualitative rankings. Examples include, but are not restricted to, disease ratings

21

on potatoes, insect damage ratings on tree leaves, and color brilliance ratings on flowers. Several applications of the RDAFP also exist in the non-agricultural sciences.

During the initial phase of the RDAFP creation, version 6.12 was the latest release of SAS. Since then, versions 7.0 and 8.0 have been released, each containing procedures (tlogisitic and genmod) capable of proportional odds model analysis. An updated version of the RDAFP is under development that will work with these procedures in the later versions of SAS.

BIBLIOGRAPHY

Freund, R. J., and W. J. Wilson. 1993. Statistical Methods. Academic Press, Inc., San Diego, CA.

Karcher, D. K. 2000. Rating data analysis file package web page. Available at http://www.msu.edu/~karcherd/ratings.

Kleinbaum, D. G. 1994. Logistic regression; a self learning text. Springer-Verlag., New York, NY.

McCullagh P. 1980. Regression models for ordinal data. Journal of the Royal Statistical Society, Series B 42:109142, London.

McCullagh P., and J. A. Nelder. 1989. Generalized linear models. 2nd ed. Chapman and Hall. London.

Microsoft. 1995. Excel 95 for Windows. Microsoft Corp., Seattle, WA.

Murphy, J. A., and P. E. Rieke. 1994. High pressure water injection and core cultivation of a compacted putting green. Agron J. 86:719-724.

SAS Institute. 1996. The SAS system for Windows. Release 6.12. SAS Inst., Cary, NC.-

Schabenberger, O., D. E. Karcher, P. E. Rieke, J. D. Kelly, and J. C. Stier. 2000. Treatment effects and comparisons in experiments with ordered outcomes. Submitted to Crop Science.

Shenton, L. R., and K. O. Bowman. 1977. Maximum likelihood estimation in small samples. Macmillan, New York, NY.