

Plotting Differences among LSMEANS in Generalized Linear Models

Robin High

University of Nebraska Medical Center, Omaha, NE

Abstract

The effectiveness of visual interpretation of the differences between pairs of LsMeans in a generalized linear model includes the graph's ability to display four inferential and two perceptual tasks. Among the types of graphs which display some or all of these tasks are the forest plot, the mean-mean scatter plot (diffogram), and the mean-mean multiple comparison (MMC) plot. These graphs provide essential visual perspectives for interpretation of the differences among pairs of LsMeans from a Generalized Linear Model (GLM). The diffogram is a graphical option now available through ODS statistical graphics with linear model procedures such as GLIMMIX of SAS®. Through combining ODS output files of the lsmeans and their differences, the SGPLOT procedure can efficiently produce forest and MMC plots.

Visual Displays

Visual displays of distributions are important aids to compare variation in data. For continuous data classified by two or more levels of a categorical variable, standard graphical displays include histograms, boxplots, or dot plots. Summary statistics such as means, medians, variances, and sample sizes may also be displayed in a corresponding table. Figure 1 shows two basic examples of the comparing the distribution of continuous data defined by two groups. The boxplots on the left displays the range and inter-quartile range (IQR) and are best utilized for large data sets. The dot plot on the right is pertinent for small datasets with the mean plotted as a short horizontal line; random horizontal jittering of these individual points minimizes over-strikes. Both visual displays clearly identify the ranges and means of the data for each group level. However, as useful as they are for presenting data distributions by showing similarity with overlapping ranges or dissimilarity, these types of graphs are not designed for or consistently helpful for detecting significant differences in group means.

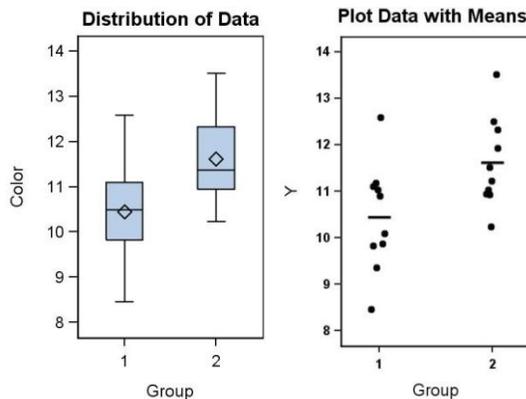


Figure 1. Plots for continuous data classified by one factor with two levels

Visual Displays of LSMeans and their Differences in a Generalized Linear Model

Assume that continuous response data are classified into two groups of equal sample sizes and the ideal conditions of normality, independence, and variance equality are reasonably satisfied. Data collected in this manner are usually analyzed with a two-sample T-test. Graphs of only the LsMeans as bar charts, perhaps with confidence intervals, are all-too-often published with only annotations or explanatory text to indicate which pairs of means are significantly different (reporting pvalue only) rather than display confidence intervals for the differences. However, a more effect plot, even from a T-test, is to plot the two LsMeans connected with the confidence interval for their difference side-by-side as shown in Figure 2. Vertical lines emanating from the two means represent a 95% confidence interval for the means with standard deviations computed from the pooled estimate of the variance. The width of the corresponding 95% confidence interval for their difference depends on the magnitude of the pooled sample variance and the number of observations.

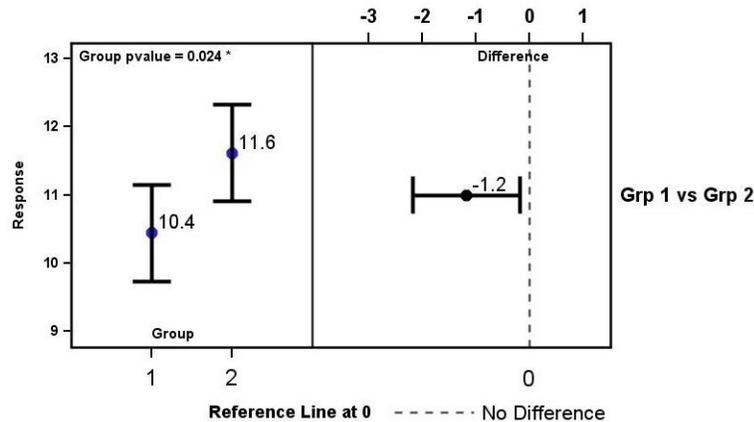


Figure 2. Plot of two LsMeans and a 95% confidence interval for their difference.

The left portion of Figure 2 shows that 95% confidence intervals of the two LsMeans can overlap (which may also be inferred from the considerable overlap of the data distributions in Figure 1) yet still produce a statistically significant difference (i.e., $p < 0.05$, meeting an ethereal dichotomy of success) as shown by the horizontal line of the 95% confidence interval for the difference in the two means in the right portion of the graph (0 not included in the interval).

To illustrate how overlap of LsMean confidence intervals may hinder interpretation of a significant difference in two means, Julious (2004) compares the means from two independent groups with equal sample sizes with the Z-test (variance known). He demonstrates how 84% (or smaller) confidence intervals around the two means that do not overlap ensure that the difference between the means of the two groups will have significance of $\alpha = 0.05$ (or lower) (i.e., a 95% confidence interval for the difference in two means does not contain 0). Continuous data analyzed by the two-sample T-test may also have means with overlapping confidence intervals, yet still display a confidence interval of the difference that is significant.

Interpretations of differences among LsMeans become more problematic when three or more means with confidence intervals are displayed and when these plots are constructed with standard deviations computed from only the data belonging to each level of the classification factor and not with the pooled standard deviation (which is typically the value that produces the output from an ANOVA) from which statistical inferences about the differences in means will be made. Observing that the 95% error bars around the two means in Figure 2 overlap, yet the 95% confidence interval for their difference does not contain 0 indicates why interpreting differences among LsMeans with confidence intervals (computed with the “correct” standard error) either displayed graphically or with the actual numbers is so important.

Interpreting differences among LsMeans with error bars only is more ambiguous and even impossible when complicated analytic features such as unequal sample sizes, unequal variances, random effects, or repeated measures are present. Plots of the LsMeans with confidence intervals do not account multiple comparison adjustments when three or more categorical levels exist nor do they convey information about the strength of the correlations among levels of data collected from within subject factors in repeated measures designs. In essence, confidence intervals surrounding the LsMeans can be inconclusive or misleading and, perhaps with exceptions for situations for considerable overlap (with little or no correlation) or extreme separation, are essentially of little value for interpreting the significance of differences among LsMeans and should not be chosen as a graphical display when the intent is to illustrate differences.

Inferential and Perceptual Tasks

The hypothesis testing framework and evaluating differences among LsMeans with statistical procedures such as PROCs GLM, MIXED, or GLIMMIX has the same objective to test the equality of group means. A visual representation of multiple comparisons should allow an interpreter to make four inferential tasks between pairs of means. The four inferential tasks are grouped into two sets.

The first two inferential tasks examine equality of means by looking for significant differences with a pre-determined Type I error rate, α (e.g., $\alpha = 0.01$ or $\alpha = 0.05$), between groups by either:

- constructing a $100 \cdot (1 - \alpha/2)\%$ confidence interval for $(\mu_i - \mu_j)$ defined by the boundary values L & U where 0 is not contained in (L,U)

- making an accept/reject decision (i.e., $\mu_i < \mu_j$ or $\mu_i > \mu_j$) if the absolute value of the test statistic is greater than the critical value of the test defined by α

And since equality testing is not equivalence testing (i.e., “absence of evidence is not evidence of absence”), setting tolerance limits on how different the means can be from each other and still be considered equivalent leads to two inferential tasks to detect *practical equivalence*:

- construct two one-sided $100 \cdot (1 - \alpha)\%$ confidence intervals for $(\mu_i - \mu_j)$, for which the intersection is completely contained within the tolerance bounds (L, U)
- accept/reject decision is based on $L = -\delta_1 < \mu_i - \mu_j < \delta_2 = U$ with $\delta_1, \delta_2 > 0$ pre-specified

The two-one-sided tests (TOST) available as an option in PROC TTEST with user defined tolerance limits essentially works in this manner to determine the equivalence of two group means.

Intrinsic dependencies among the confidence intervals for the differences in means should also be visible in graphical comparisons. Two perceptual tasks are also important when interpreting graphical displays. First, the elementary contrasts of two means are *additive*:

$$\mu_i - \mu_k = (\mu_i - \mu_j) + (\mu_j - \mu_k)$$

For any two pairs of group means, μ_i and μ_k , with any third mean (μ_j) subtracted from μ_i and also the first mean that appears with the difference with μ_k , the graphical method should center the confidence interval for the difference at the “sum” of the centers of the confidence intervals for $(\mu_i - \mu_j)$ and $(\mu_j - \mu_k)$.

The second perceptual task of a graphical display is the inherent *transitivity* present with significant differences of means; that is, if the multiple comparison method declares $\mu_i > \mu_j$ and also that $\mu_j > \mu_k$, then it must necessarily declare $\mu_i > \mu_k$.

With these inferential and perceptual tasks in mind, according to Heiberger (2006) a graph of LsMeans and their pairwise differences should also display:

1. The means with correct relative distances
2. Point and interval estimates of all pairwise differences
3. Declarations of significance
4. Correct confidence interval widths for unequal sample sizes and unequal variances

Multiple Comparisons of LsMeans

Methods of plotting and interpreting multiple differences in LsMeans are illustrated here with the catalyst data set from Montgomery (1997), the first example from Heiberger and Holland (2006). The continuous response is analyzed through an ANOVA with one fixed factor, catalyst (coded 1, 2, 3, & 4) having four levels A, B, C, & D assigned through a format (the recommended coding method for producing results). The significant F-test provides evidence of unequal mean concentrations for catalyst ($F=9.92, 3, 12, p=.0014$). The graphical techniques discussed here first display the six pairwise differences among the four catalyst means. These data were analyzed with PROC GLIMMIX from SAS/STAT software, V. 9.3, assuming equal variances across the groups. ODS OUTPUT statements save LSMEANS results into SAS datasets which provide all data necessary to produce the graphs:

```
ODS OUTPUT lsmeans=lsm(KEEP= catalyst estimate lower upper)
           diffs=dfs(KEEP= catalyst _catalyst estimate adjlower adjupper adjp);
LSMEANS catalyst / cl diff adjust=tukey;
```

Adjusted confidence intervals for the differences in the LSMEANS are displayed along with the adjusted pvalue to determine the significance of each comparison. Based on the design and to scenario of all possible comparisons of means, adjustments for a one-way ANOVA are computed with the Tukey-Kramer method. Methods for multiple comparisons for other situations and how to implement them are described in detail in Westfall, et. al. (2011).

The Forest Plot

The forest plot is a graphical display which historically has compared results from a meta analysis for odds ratios. This plot is also well suited to display confidence intervals for differences among LsMeans of pairwise combinations of factor levels as part of the ANOVA summary. It is a very simple graph to make with SGPLOT. The ODS file of differences produced with the LSMEANS statement listed above contains all the information needed to plot the adjusted confidence intervals of the pairwise differences displayed as horizontal lines with the value of the difference as the center point of the interval. A significant difference exists when the adjusted 95% confidence interval of the comparison does not intersect a vertical line placed at a difference of 0. The only required modification to the file of differences is to add a text variable which contains the two factor levels involved in the difference of LsMeans, in this example identified by the variables catalyst and _catalyst. The plotted variable for the vertical axis which shows the contents of the comparison is produced with the CAT function in a DATA step. The actual names of the factor levels are entered with a format and the adjusted pvalue can also be added to the row label with a pvalue format. The statements entered into a DATA step which make this label are:

```
LENGTH label $20 ptxt $4;
IF adjp < .001 then ptxt=' , p='; else ptxt=' , p=';
label = CAT(put(catalyst,ct.), ' vs ',
            put(_catalyst,ct.),ptxt,strip(put(adjp,pvalue6.3)));
```

The name of the format of the integer coded catalyst variable is ct; the length of the label must be long enough to contain the contents of the longest expression. Otherwise, the contents of label will be blank. Figure 3 shows a forest plot of the LsMean differences for the example data. SGPLOT also provides the option to print the actual differences at the center of each interval; the values of the endpoints of the intervals can be approximated by vertical lines drawn to the horizontal axis. The labels for the differences can be attached to either the left vertical axis (YAXIS) or to the right vertical axis (Y2AXIS). These comparisons are positioned with equal spacing between the horizontal lines without any visual representation of the values of the two LsMeans for each difference.

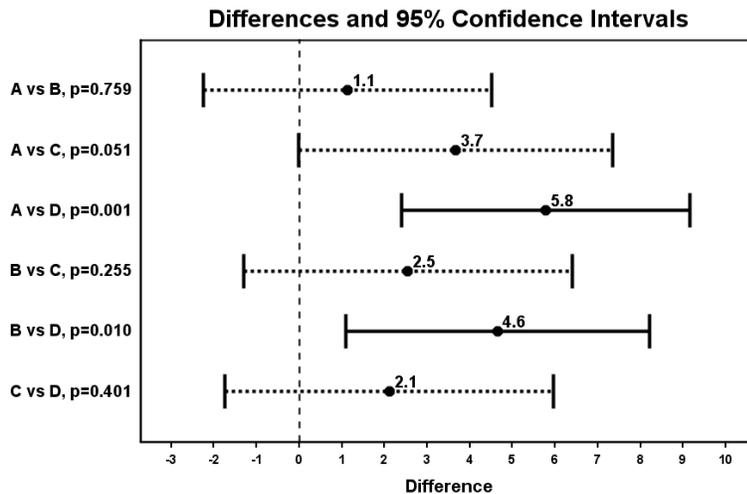


Figure 3. Forest plot of adjusted LsMean differences.

The Mean-Mean Scatter Plot (Diffogram)

A second graphical display for interpretation of differences among the LsMeans available with PROC GLIMMIX is the mean-mean scatter plot (Hsu, 1996), also called a diffogram. The graph is designed to display differences among the LSMEANS when analyzing data having any of the response distributions available with Generalized Linear Models; it is characterized by the following features:

- It is a two-dimensional plot with both the vertical and horizontal axes of the same length and having the same numerical range defined by the largest and smallest LsMean (with a slightly larger range needed to plot the confidence intervals for the differences within the plotting area).
- All possible pairs of the p LsMeans ($LsMean_i$ and $LsMean_j$, where $i=1,2,..,p-1$, $j=2,..,p$, $i \neq j$) are plotted with $LsMean_i$ on the horizontal axis (abscissa) and $LsMean_j$ on the vertical axis (ordinate).

- Within the plotting area the point of intersection of the horizontal and vertical line emanating from both axes for the two LsMeans represents the value of their difference:

$$\text{diff}_{ij} = \text{LsMean}_i - \text{LsMean}_j \quad (i= 1, 2, \dots, p-1, j=2, \dots, p, i \neq j)$$

- The dashed line from the lower left to the upper right corners represents points of equality for two means; that is, points on this line display the difference of two Lsmeans equal to 0. The difference of each LsMean with itself is trivially 0 (note the point of intersection of this diagonal with vertical and horizontal lines assumes equal means).
- A third axis (implied on the graph but not printed by SAS ODS Graphics) exists from the top left to the lower right corners of the plotting area. This hidden axis represents the magnitude of the differences of the LSMEANS defined on the horizontal and vertical axes scaled in such a way [i.e., divided by $\text{SQRT}(2)$] that the confidence interval for the difference crosses the line of equality when the interval contains 0. It also provides values for the endpoints of the confidence intervals necessary for equivalence testing.
- A significant difference between pairs of LsMeans (i.e., the adjusted pvalue is less than $\alpha=0.05$) occurs when the adjusted lower and upper endpoints of the confidence intervals are both positive or both negative; that is, the solid lines sloping at -45 degrees fall completely above or below the line of equality. This line of equality also identifies non-significant differences; when the confidence interval intersects the diagonal (shown as dashed lines), then 0 is included within the confidence interval.

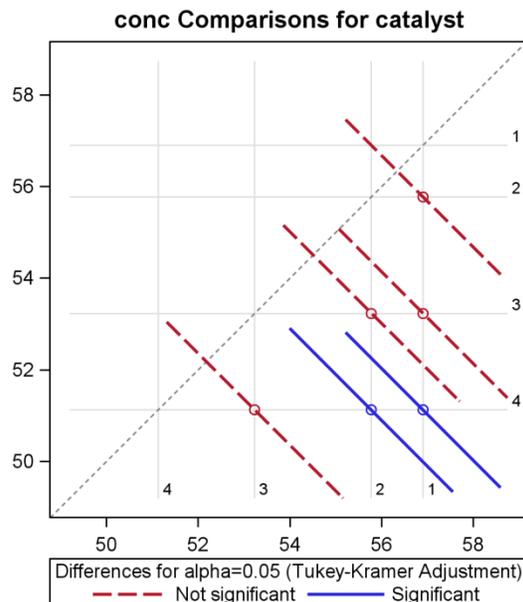


Figure 4. A diffogram (mean-mean scatter plot) produced with PROC GLIMMIX

All six comparisons of the four LsMeans are shown in Figure 4 as the center point of each confidence interval, indicated by a 45 degree downward sloping line. If the difference in the two LsMeans is significant, the solid line for the confidence interval lies completely above or below the upward diagonal dashed line of equality. Non-significant differences are indicated by a dashed line that crosses the line of equality.

For non-normal responses identified with the `distribution=<>` option on the MODEL statement of GLIMMIX, the differences displayed on the diffogram are on the scale of the specified or default link function (for example, the defaults are logits for binary data and logs for the poisson or negative binomial distributions).

Mean-Mean Multiple Comparison Plot

The Mean-Mean Multiple Comparison (MMC) plot is described by Heiberger and Holland (2006) as a generalization and refinement of the features of a mean-mean scatter plot (diffogram) presented above. They also summarize the geometric relationship between the two graphs. All the characteristics listed above for the diffogram apply to the MMC plot. The only difference between the two graphs is the MMC plot tilts the contents of the diffogram 45 degrees to the left, so that it balances on the lower left corner where the horizon and vertical axes intersect. In its essential details,

the MMC plot conveys exactly the same information as the diffogram, yet with two distinct advantages: it is much easier to produce with SGPLOT and it is much easier to interpret, since the values of the Lsmeans, factor identifiers, and their differences are read directly from the left vertical and bottom horizontal axes.

The same confidence intervals for differences in LsMeans and their associated pvalues as shown on the forest plot (Figure 3) and the diffogram (Figure 4) are presented in the MMC plot (Figure 5). The values of the LsMeans and their associated factor levels are placed relative to their magnitude on the left vertical axis. The values of the differences, not available from the GLIMMIX diffogram, are now easily read from the lower horizontal axis. The confidence interval for each difference is positioned at the average of the two LsMeans from which it is computed. The names of the two factor levels involved each comparison are printed on the right vertical axis.

The SGPLOT procedure produces the MMC plot from the saved output of the LSMEANS statement with two files produced with ODS OUTPUT statements with either PROCs MIXED (normal distribution) or GLIMMIX (any of the available distribution choices). The next step to make a data set to plot the graph requires some programming, in particular a DATA step to merge in the average of the two LsMeans that produced the difference with the file of differences (this task is greatly simplified when integer codes for the ANOVA factor beginning with 1 define the levels). The average value is the vertical coordinate at which the horizontal confidence interval for each difference is placed on the graph. Computation of this average value leads to one disadvantage of both the diffogram and MMC plot in that intervals in close proximity or even overlapping with each other are possible for two or more comparisons that have the same average of their component LsMeans (or values close to each other), even though the individual Lsmeans may be very different. Since a format is applied to this average value to print the comparison on the vertical axis, they must remain distinct. Possible solutions when this occurs is to add a small offset to these average values (which is possible with the MCC plot rather than with the diffogram) or increase the vertical height of the plot for average values with close proximity. If these approaches do not work consider computing a contrast of the LsMeans involved or make a forest plot of the subset of differences and place them in a separate plot next to the MMC plot.

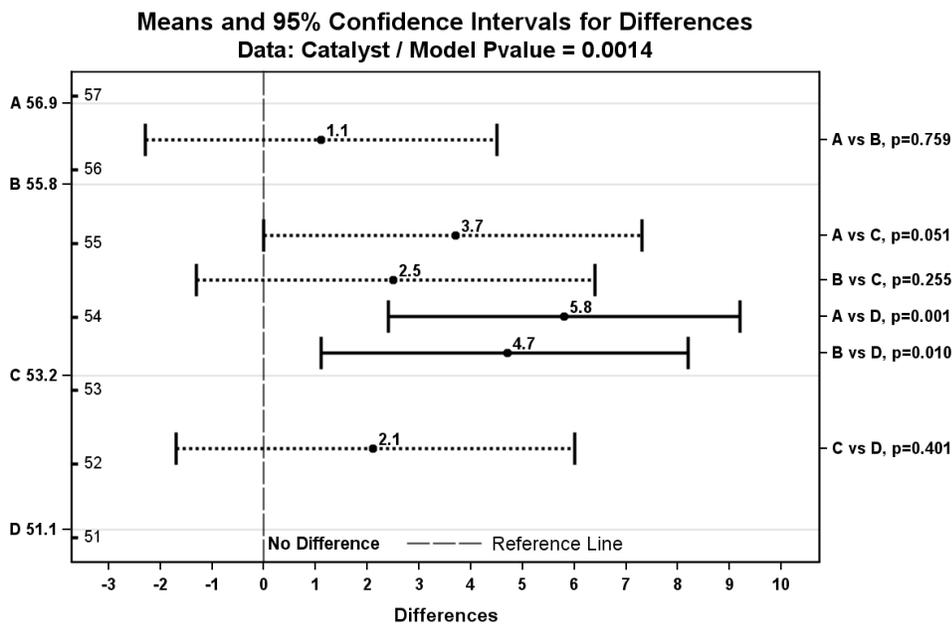


Figure 5. MMC plot produced with PROC SGPLOT

The MMC plot provides the same interpretation of LsMeans and their differences as provided by the diffogram, but in a layout which generally is much easier for viewers to interpret. One important advantage of the MMC plot is the factor levels and their respective LsMeans can be printed directly on the outside of the left vertical axis. This feature is accomplished here by making a format with the file of LsMeans and the SPLOT plotting options on the YAXIS statement (see code in the Appendix). The sequential values covering a relevant range of means are placed inside the plotting area with an annotation data set to minimize overstrikes that may occur with the factor mean values and labels. Also, as observed on the forest plot, SGPLOT provides the option to print the differences at the center of each confidence interval; the values of the interval endpoints can be approximated by vertical lines to the horizontal axis.

Statistical Significance

These graphs allow direct conclusions to be made about the statistical significance of LsMean comparisons. All three plots readily identify significant differences by line type: a solid line is chosen for a significant difference and a dashed line for a difference that is not significant (i.e., the adjusted confidence interval includes 0). Also significance is determined with the pvalue printed in the comparison label with the forest and MMC plots based on the choice of the investigator's level of a Type I error α .

The MMC plot combines features of the forest plot and the diffogram (the same plot except oriented as a horizontal rectangle, not as a square tilted at a 45 degree angle). The means on the diffogram are read from where the two lines from each difference intersect each axis; they are printed on the graph itself, as observed on left axis of the MMC plot. On the MMC plot, the confidence interval for the difference of any two LsMeans is positioned halfway between the two left vertical axis labels for the means; the difference between the levels with the two largest means will always be the top horizontal line. On the forest plot, this confidence interval may appear anywhere.

Results from the MMC plot in Figure 5 can be compared with the forest plot of Figure 3 and the diffogram of Figure 4. The MCC plot clearly indicates how the two largest means, one from catalyst A (56.9) and the other from catalyst B (55.8), are both significantly different from the mean catalyst D (51.1). Both significant differences are displayed with solid lines for these two confidence intervals do not intersect 0. Pairs of means for catalysts AB, AC, BC, and CD are not significant from each other (their dashed lines all include a difference of 0).

Practical Equivalence

The length and position of the confidence intervals in all three plots determine whether practical equivalence of the two factor levels exists, that is, how tightly the confidence interval of the difference is placed around 0. In the diffogram two additional parallel lines at the specified upper and lower tolerance limits where the difference is considered trivial, could be drawn perpendicular to the upward diagonal axis of equality (not illustrated here). Without the implied third axis printed, equivalence is not directly observable with the diffogram currently produced with SAS ODS Graphics. With the scale of differences printed on the lower horizontal axis, both the forest and MMC plots allow vertical lines to easily be added as appropriate limits for equivalence. These values in both directions from the vertical line of equality at 0, such that a confidence interval for the difference in the two LsMeans is completely contained within it are the values where one could state the two factor levels are equivalent.

Perceptual Tasks: Additivity and Transitivity

With the forest plot the perceptual tasks are not readily identified. The diffogram and MMC plot allow interpretation of additivity for all of the comparisons, whether they are significant or not, display the relative position of each other in an additive sense. It is noteworthy to see how additivity is present here by examining the difference between any two factor means. Choosing catalysts levels A and B, their observed difference is 1.1. The differences of in means for catalysts A and B with a third mean, say level D, are 5.8 and 4.7 respectively. The difference between these two differences is also 1.1 ($=5.8-4.7$). For transitivity, the example data analysis only produces two pairs of Lsmeans that show significant differences, so demonstrating transitivity as defined above is not feasible.

Enhancements to the Forest and MMC Plots

The forest and MMC plots as described above are constructed to display confidence intervals and significance for differences in LsMeans as their primary objective. For those situations where a visual display of the LsMeans with confidence intervals is requested, an enhancement to these graphs requires relatively little extra effort. Rather than showing means with vertical bars, the means may be displayed as lines showing the 95% confidence intervals. SGPLOT is able to partition the plot area into two separate regions so that both the LsMeans with vertical confidence intervals and horizontal confidence intervals for their differences can be placed into separate sections of one graph. Specifications and options available in the XAXIS/YAXIS and X2AXIS/Y2AXIS statements allow these plots achieve this divided layout.

This enhanced graph requires data from both the LsMeans and the difference files to be appended (forced) together in a block diagonal layout with all data set to missing for variables from one section which are not needed for the other section. To make this data file variables with the same name in both files need to be renamed. With the two output files from the LSMEANS statement, "Estimate" is common to both files so it must be given a new name in at least one file, e.g., rename "estimate" as "lsmean" or "estimate" as "difference" respectively (since adjusted endpoints

for the confidence interval on the differences are recommended, the lower and upper values for the differences in the diff file are not needed; also, the factor names (i.e., catalyst and _catalyst) in the file of differences should be dropped after the axis labels are produced). Figure 6 shows an enhanced forest plot with the LsMeans and their 95% confidence intervals plotted vertically in the left portion of the graph and the forest plot for the 95% confidence intervals of the adjusted pairwise LsMean differences in the right portion:

Means and 95% Confidence Intervals for LsMean Differences

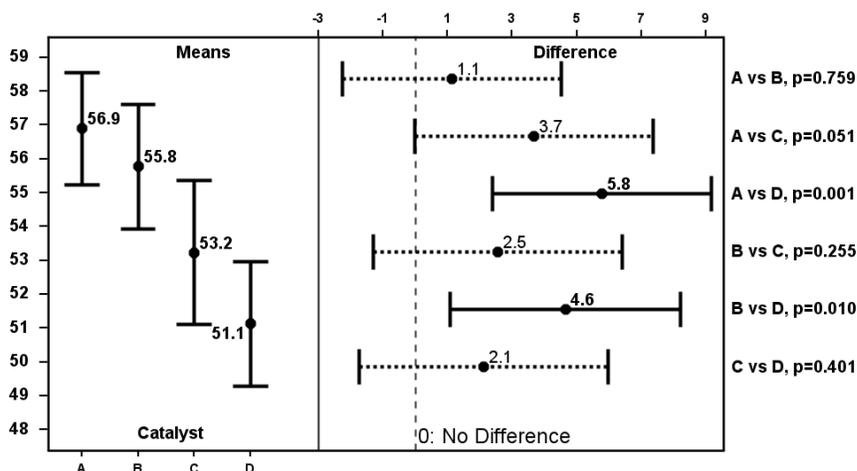


Figure 6. LsMeans (left) with forest plot of adjusted differences (right).

The same process to append files can be applied to produce an enhanced MMC plot with the LsMeans in the left portion and their pairwise differences in the right. A potential problem with this graph is the extra vertical space required for the confidence intervals of the largest and smallest LsMeans implies less vertical space will exist to plot the confidence intervals for the differences, which increases the likelihood of overlapping lines. If this situation causes overlap, a separate graph of the LsMeans could be made or resort to the forest plot which due to the nature of how it is made will not have overlapping intervals.

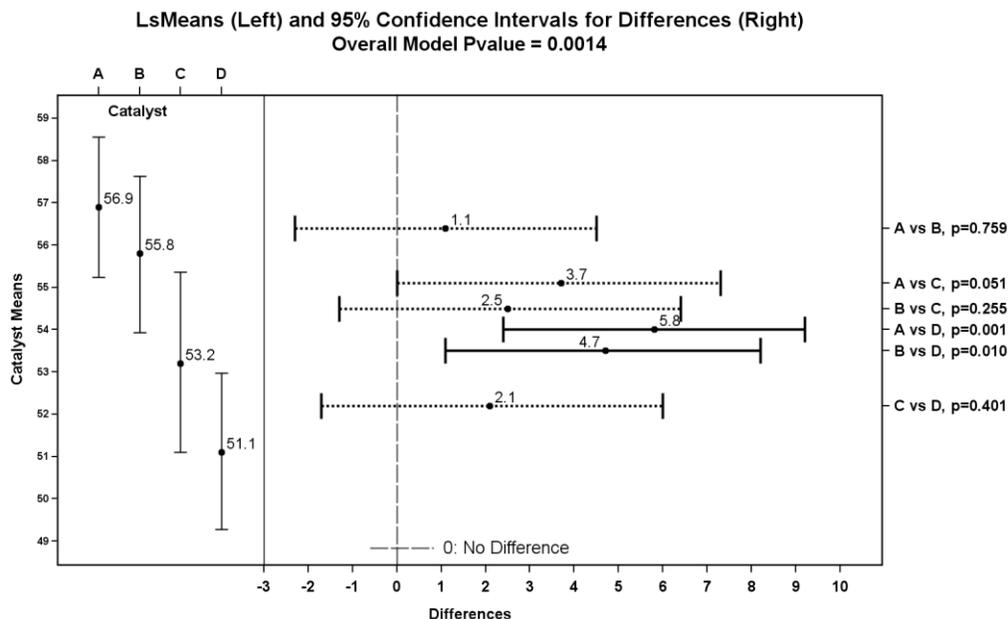


Figure 7. LsMeans (left) with MMC plot of adjusted differences (right).

Summary and Conclusions

Three plots are described which allow interpretation of differences in LsMeans in Generalized Linear Models according to four inferential and two perceptual tasks. These types of graphs have rarely been used to visually display confidence intervals from statistical analyses. The complexity of the design does not influence how they are produced, as the techniques remain the same from ANOVAs with independent groups to multiple factors with repeated measures or random effects. The magnitude of the standard error of a difference for the LsMeans, the divisor to compute p-values and form confidence intervals, depends on the choice of a repeated measures covariance matrix or effects treated as random. If only a subset of all possible comparisons is needed, the LSMESTIMATE statement can provide the relevant output file of differences with appropriately adjusted confidence intervals.

The diffogram is automatically produced with ODS statistical graphics, but is limited in its interpretability to the two inferential and two perceptual tasks, plus the graph itself may be confusing to first time viewers. It is possible to make an enhanced diffogram to avoid some of these difficulties. However, the effort required to do so is not justified when realizing the contents and interpretation of the MMC plot are the same as the diffogram and the MMC plot also clearly displays labels for means and differences. The MCC plots provide a visual method to examine all four inferential and two perceptual tasks, an essential characteristic of a graphical display of confidence intervals for LsMean differences.

In the presence of significant interactions of two factors, the forest and MMC plots can also be made by slicing one factor and then examining the simple effects of the other factor. That is, for each level of one factor, plot the confidence intervals for differences in LsMeans for all pairs of levels from the other factor. The diffogram displays adjusted confidence intervals for the difference in means for all possible pairs of the two factors, which even with relatively small numbers of levels, add up to a very large number.

In conclusion, these graphs allow interpretation of LsMean differences and could become a more common graph to summarize differences from ANOVA. The degree of difficulty to make them varies, as the forest plot is the easiest to produce requiring only the file of LsMean differences, whereas the MMC plot requires some basic DATA step manipulations.

REFERENCES

1. Heiberger, Richard M. and Burt Holland, (2006) *Mean-Mean Multiple Comparison Displays for Families of Linear Contrasts* Journal of Computational and Graphical Statistics, Volume 15, Number 4, pp. 937-955.
2. Hsu, Jason, Mario Peruggia, (1994) *Graphical Representations of Tukey's Multiple Comparison Method*, Journal of Computational and Graphical Statistics, Volume 3, No. 2, pp. 143-161.
3. Hsu, Jason (1996) *Multiple Comparisons Theory and Methods*, Chapman & Hall/CRC: Boca Raton.
4. Julious, Steven A. (2004) "Using confidence intervals around individual means to assess statistical significance between two means," *Pharmaceutical Statistics*, Vol. 3: 217-222.
5. Montgomery, D. C. (1997), *Design and Analysis of Experiments* (4th ed.), New York: Wiley.
6. Westfall, Peter H., Tobias, Randall D., and Wolfinger, Russell D., (2011). *Multiple Comparisons and Multiples Tests Using SAS®*, Cary, NC: SAS Institute Inc.

CONTACT INFORMATION

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

Robin High
Statistical Coordinator
College of Public Health
Department of Biostatistics
University of Nebraska Medical Center
984375 Nebraska Medical Center
Omaha, NE 68198-4375
Phone: (402) 559-2929
email: rhigh@unmc.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.

APPENDIX

```
* Basic SAS code assumes coding categorical factors as integers, 1, 2, 3, .. ;
* Apply a format to attach factor level descriptions;

PROC FORMAT ;
VALUE ct 0=' ' 1='A' 2='B' 3='C' 4='D' 5=' ';
run;

DATA ctlyst;
label catalyst='Catalyst' conc='Concentration';
input catalyst conc @@;
cards;
1 58.2 1 57.2 1 58.4 1 55.8 1 54.9
2 56.3 2 54.5 2 57.0 2 55.3
3 50.1 3 54.2 3 55.4
4 52.9 4 49.9 4 50.0 4 51.7
;

ODS OUTPUT lsmeans=lsmdif(KEEP=catalyst estimate lower upper)
           diff=dfs(KEEP= catalyst _catalyst estimate adjlower adjupper adjp);

ODS GRAPHICS ON / reset = all height= 4 in width=4 in border= off ;
ODS LISTING image_dpi=200 gpath="u:\sas\plots";
ODS GRAPHICS / imagename = "diffogram" ; * diffogram file;

PROC GLIMMIX DATA=ctlyst order=internal;
CLASS catalyst;
MODEL conc = catalyst / solution;
LSMEANS catalyst / cl diff adjust=tukey PLOT=diff(noabs center);
FORMAT catalyst ct. ;
run;

ODS GRAPHICS off;

* Forest Plot of Differences ;

* with the file of differences
  add a text variable called label describing each comparison;

DATA dfs2;
SET dfs(rename =( estimate=mndif1 adjlower=adjlower1 adjupper=adjupper1));
DROP ptxt;
LENGTH label $30 ptxt $4;
IF adjp < .001 then ptxt=', p'; else ptxt=', p=';
label = CAT(put(catalyst,ct.),' vs ',
           put(_catalyst,ct.),ptxt,STRIP(put(adjp,pvalue6.3)));
* new variables for non-significant differences, set existing to missing;
IF (adjp GE 0.05) then do;
  mndif2 = mndif1; adjlower2 = adjlower1; adjupper2 =adjupper1;
  mndif1 = . ; adjlower1 = . ; adjupper1 = . ;
end;
RUN;

PROC PRINT DATA=dfs2 NOObs;
RUN;
```

```

ODS GRAPHICS ON / reset = all height= 4 in width=4 in border= off ;
ods listing image_dpi=200 gpath="u:\sas\plots";
ods graphics / imagename = "forest_diffs" ;

PROC SGPLOT DATA=dfs2 noautolegend ;
REFLINE 0 / axis=x lineattrs=(color=black pattern=2 thickness=1) transparency=0;
SCATTER y=label x=mndif1 / xerrorlower=adjlower1 xerrorupper=adjupper1
           errorbarattrs=(color=black pattern=1 thickness=2)
           markerattrs=(color=black symbol= circlefilled size=6)
           datalabel=mndif1 datalabelattrs=(color=black weight=bold size=7);

/* print the differences at the center of each confidence interval */
SCATTER y=label x=mndif2 / xerrorlower=adjlower2 xerrorupper=adjupper2
           errorbarattrs=(color=black pattern=2 thickness=2)
           markerattrs=(color=black symbol= circlefilled size=6)
           datalabel=mndif2 datalabelattrs=(color=black weight=bold size=7);
XAXIS offsetmin=0.05 offsetmax=0.05 label='Difference' values=(-3 to 10 by 1);
YAXIS offsetmin=0.12 offsetmax=0.12 display=(Nolabel Noticks) reverse ;
FORMAT mndif1 mndif2 5.1;
TITLE "Differences and Adjusted 95% Confidence Intervals";
RUN;

ODS GRAPHICS off;

* MMC plot ;

* need values for plotting rounded to the smallest number of decimals
  that allow graphical resolution, in this case round to nearest 1 decimal;
* remove attached formats on catalyst and _catalyst ;

* For plotting and assigning formats, round LsMeans
  to resolution of graph or smaller, if necessary, so they remain distinct;

%LET rnd = .1 ; * round lsmean to the nearest .1 decimal value;

DATA lsm; set lsm;
estimate=round(estimate,&rnd.);

PROC DATASETS NOLIST;
MODIFY lsm; FORMAT catalyst 3.0 estimate 5.1;
RUN; quit;
PROC PRINT DATA=lsm NOObs; run;

PROC DATASETS nolist;
MODIFY dfs; FORMAT catalyst _catalyst 4.0;
run; quit;
proc print data=dfs NOObs; run;

DATA lsm2; SET lsm;
KEEP catalyst estimate start;
start=estimate;

PROC SORT DATA=lsm2;
BY estimate; RUN;

* Labels for means placed the left vertical axis;

DATA Lfmt; SET lsm2;
DROP y cmpr catalyst ;
LENGTH label $15; * NEED length long enough to contain format and mean ;
RETAIN fmtname "L1b1" type 'N';
label = CATX(' ',put(catalyst,ct.),put(estimate,best12.));
run;

```

```

proc format cntlin=Lfmt; run;

* test the format;

PROC FREQ DATA=ls2 order=formatted;
TABLES estimate / nocum ;
FORMAT estimate Llbl. ;
RUN;

* place the lsmeans into a macro variable from smallest to largest
  (needed for values=() option YAXIS statement);

PROC SQL NOprint;
SELECT distinct estimate
      into: mns separated by ' '
FROM ls2
ORDER BY estimate;
quit;

* the right vertical axis Y2AXIS, plotting values;
* get low to high values of lsmeans transposed into macro variable;

PROC TRANSPOSE DATA=ls2 out=tlsm (drop=_name_ _label_) prefix=_mn;
VAR estimate; ID catalyst;
proc print data=tlsm NOObs; run;

DATA dfs2; SET dfs;
IF _n_ = 1 then set tlsm; * read transposed means;
DROP _mn1 _mn2 _mn3 _mn4 ptxt catalyst _catalyst;
LENGTH ylabel $35 ptxt $4;
FORMAT y2 6.2 ;
ARRAY mns{4} _mn1 _mn2 _mn3 _mn4 ;
y = (mns{catalyst} + mns{_catalyst})/2 ;
y2=y; * copy of y for yaxis2;
IF adjp < .001 then ptxt=' p'; else ptxt=' p=';
ylabel = CAT(strip(PUT(catalyst,ct.)), ' vs ', strip(PUT(_catalyst,ct.)),
            strip(ptxt), strip(PUT(adjp,pvalue6.3))); * values to label y2axis;
grp=(adjp < .05);
* For non-significant differences and confidence intervals;
IF grp=0 then
  DO; estimate2=estimate; adjupper2=adjupper; adjlower2=adjlower;
      estimate = . ; adjupper = . ; adjlower = . ;
  END;
RUN;

proc print data=dfs2 NOObs;
var ylabel Estimate adjLower adjUpper adjp grp estimate2 adjupper2 adjlower2 y; ;
format estimate adjlower adjupper 6.1 adjp pvalue6.3 ;
run;

* place average of pairs of lsmeans into a macro variable
  from smallest to largest for Y2AXIS;

PROC SQL NOprint;
SELECT distinct y2
      into: amns separated by ' '
FROM dfs2
ORDER BY y2;
quit;

* the value of y2 must be different for each comparison for the format to work;
* for equal values, need to either refine the rounding or offset values;

```

```

* Difference labels for the right vertical axis;
DATA Rfmt;
SET dfs2(keep=y2 ylabel rename=(y2=start ylabel=label));
RETAIN fmtname "Rlbl" type 'N';
run;

proc format cntlin=Rfmt; run;

* test the format;
proc freq data=dfs2 order=formatted;
tables y2 / nocum ;
format y2 Rlbl. ;
run;

* find the min and max values derived from both of these PROC MEANS;
proc means data=dfs2 min max maxdec=3; var y y2; run;
proc means data=lsm2 min max maxdec=3; var estimate ; run;

* Enter boundary values for lsmeans from MEANS output;

%LET ymn=51; * minimum for both vertical axes;
%LET ymx=57; * maximum for both vertical axes;

* annotate data set for YAXIS scale placed inside graph;
DATA LsmT;
DROP yt;
LENGTH function $4 label $2;
RETAIN linecolor 'black' linestyle 'solid' linesize 1
      y1space "datavalue" x1space "wallpercent"
      y2space "datavalue" x2space "wallpercent"
      textcolor 'black' textsize 7 ;
* tick mark for LsMeans (left axis);
DO yt = &ymn. to &ymx. by 1;
  function= 'line'; label= ' ';
  x1 = 0 ; y1 = yt;
  x2 = x1 + 1; y2 = yt; OUTPUT;
  function= "text"; label= strip(put(yt,3.0));
  x1=x2 + 2; x2=x1 + 5; OUTPUT;
end;

proc print data=LsmT NOObs; run;

proc means data=dfs2 min max maxdec=2;
var adjlower adjupper adjlower2 adjupper2;
run;

* Enter boundary values for differences (horizontal axis) from PROC MEANS ;

%LET dmn= -3; * minimum difference;
%LET dmx= 10; * maximum difference;
%LET dcr= 1; * increment for differences on lower horizontal axis;

```

```

ODS GRAPHICS on / reset = all height= 4 in width=6 in border= off;
ODS LISTING image_dpi=200 gpath="u:\sas\plots";
ODS GRAPHICS / imagename = "mmc_diffs" ;

PROC SGPLOT DATA=dfs2 noautolegend sganno=lsmT;

* plot the differences, solid for significant, dashed for non-significant;
SCATTER x=estimate y=y / y2axis xerrorlower=adjlower xerrorupper=adjupper
        errorbarattrs=(color=black pattern=1 thickness=2)
        markerattrs=(color=black symbol= circlefilled size=5)
        datalabel=estimate datalabelattrs=(color=black weight=bold size=7);
SCATTER x=estimate2 y=y / y2axis xerrorlower=adjlower2 xerrorupper=adjupper2
        errorbarattrs=(color=black pattern=2 thickness=2)
        markerattrs=(color=black symbol= circlefilled size=5)
        datalabel=estimate2 datalabelattrs=(color=black weight=bold size=7);
Y2AXIS offsetmin=0.05 offsetmax=0.05
        min=&ymn. max=&ymx. values=( &mn ) VALUESHINT valueattrs=(weight=bold size=7)
        tickvalueformat=data display=(Nolabel);
* min and max values needed determined by range of lsmeans;

XAXIS offsetmin=0.05 offsetmax=0.05
        values=(&dmn. to &dmx. by &dcr. ) valueattrs=(weight=bold size=7)
        label="Differences" labelattrs=(weight=bold size=8);
REFLINE 0 / axis=x LINEATTRS=(thickness=1 color=black pattern=21)
        transparency=.2 name="zr";
KEYLEGEND "zr" / across=1 down=1 noborder location=inside position=bottom
        title='No Difference' TITLEATTRS=(Color=black Size=8 Weight=Bold);

* place the formatted label for means on the Left Vertical axis (does not plot);
SCATTER x=estimate y=y2 / markerattrs=(color=white size=0);
YAXIS offsetmin=0.05 offsetmax=0.05
        min=&ymn. max=&ymx. values=( &mn ) VALUESHINT valueattrs=(weight=bold size=7)
        grid tickvalueformat=data display=(nolabel);

FORMAT estimate: 5.1 y2 Llbl. y Rlbl. ;
TITLE1 'Means and Adjusted 95% Confidence Intervals for Differences';
TITLE2 "Catalyst Data" ;
run;

```

```
ods graphics off;
```

Important points about the options entered for the YAXIS and Y2AXIS statements:

- * offsetmin= and offsetmax= values need to match on both statements
- * min= and max= must be the same values on both statements
 - as they specify both left and right axis to have the same length
- * min= and max= have no effect if you specify the VALUES= option
 - and you do not also specify valueshint
- * valueshint: the minimum and maximum axis values are determined independently
 - of the values specified in the VALUES= option

The SGPLOT code to make the enhanced forest and MMC plots looks very much like two examples above for the differences only and is not shown here. To make the required data set, the lsmeans and diffs files are appended together and variable name estimate must be renamed in at least one file (e.g., estimate=LsMean). Another SCATTER statement for the LsMeans with corresponding x=catalyst and y=LsMean for the coordinates plus the yerrorlower=lower and yerrorupper=upper with variable names is added to the SGPLOT code associated with a X2AXIS statement. Also, partition the division of the plotting space into two sections with the offsetmin= and offsetmax= options on both the XAXIS and X2AXIS statements. The XAXIS and X2AXIS statements may be interchanged for plotting the means and differences, if desired.