MaxDiff Designer v2

(Updated 7 March 06)

Sawtooth Software, Inc.
Sequim, WA
http://www.sawtoothsoftware.com

Bryan Orme, Editor

About Technical Support

We've designed this manual to teach you how to use our software and to serve as a reference to answer your questions. If you still have questions after consulting the manual, we offer telephone support.

When you call us, please be at your computer and have at hand any instructions or files associated with your problem, or a description of the sequence of keystrokes or events that led to your problem. This way, we can attempt to duplicate your problem and quickly arrive at a solution.

For customer support, contact our Sequim, Washington office at 360/681-2300, email: support@sawtoothsoftware.com, (fax: 360/681-2400).

Outside of the U.S., contact your Sawtooth Software representative for support.

1 The MaxDiff Designer

Introduction

This document describes design software for creating Maximum Difference (Best/Worst) or method of paired comparison (MPC) experiments. It also describes the data preparation procedure if using Sawtooth Software's Latent Class v2.x or CBC/HB software for parameter estimation. Sawtooth Software also offers this design tool integrated within SSI Web for complete web-based interviewing and HB analysis. If you are new to the MaxDiff methodology, we suggest you read the "MaxDiff/Web Technical Paper," available within our technical papers library at www.sawtoothsoftware.com for background information.

Users of the MaxDiff/Web system (within SSI Web) are also entitled to this software as an extension of their MaxDiff/Web license. This stand-alone tool allows MaxDiff/Web users to develop questionnaires outside of SSI Web (such as paper-and-pencil studies) and merge respondent answers into a file appropriate for analysis using our Latent Class v2.x or CBC/HB systems (the .CHO file).

Steps in Using the MaxDiff Designer

1. Create an experimental design using the MaxDiff Designer software.
2. Develop the questionnaire using your other software tools (such as your word-processor or another CAPI or Web-based Interviewing system).
3. Using your own tools, collect the data and prepare the respondent answers as a text-only file (such as .csv, .txt).
4. Use the MaxDiff Designer software to prepare a .CHO file (text-only file containing information needed for computation) for analysis using Latent Class v2.x or CBC/HB software.

Capacities

This design software has the following capacities:

Number of items	500
Items per set	500
Sets per respondent	Limited by prior two inputs which determine possible unique sets, or up to 500
Questionnaire Versions	1000

2 Design Generation

When you click ***File | New*** (or open an existing project), the main Generate Design menu is displayed:

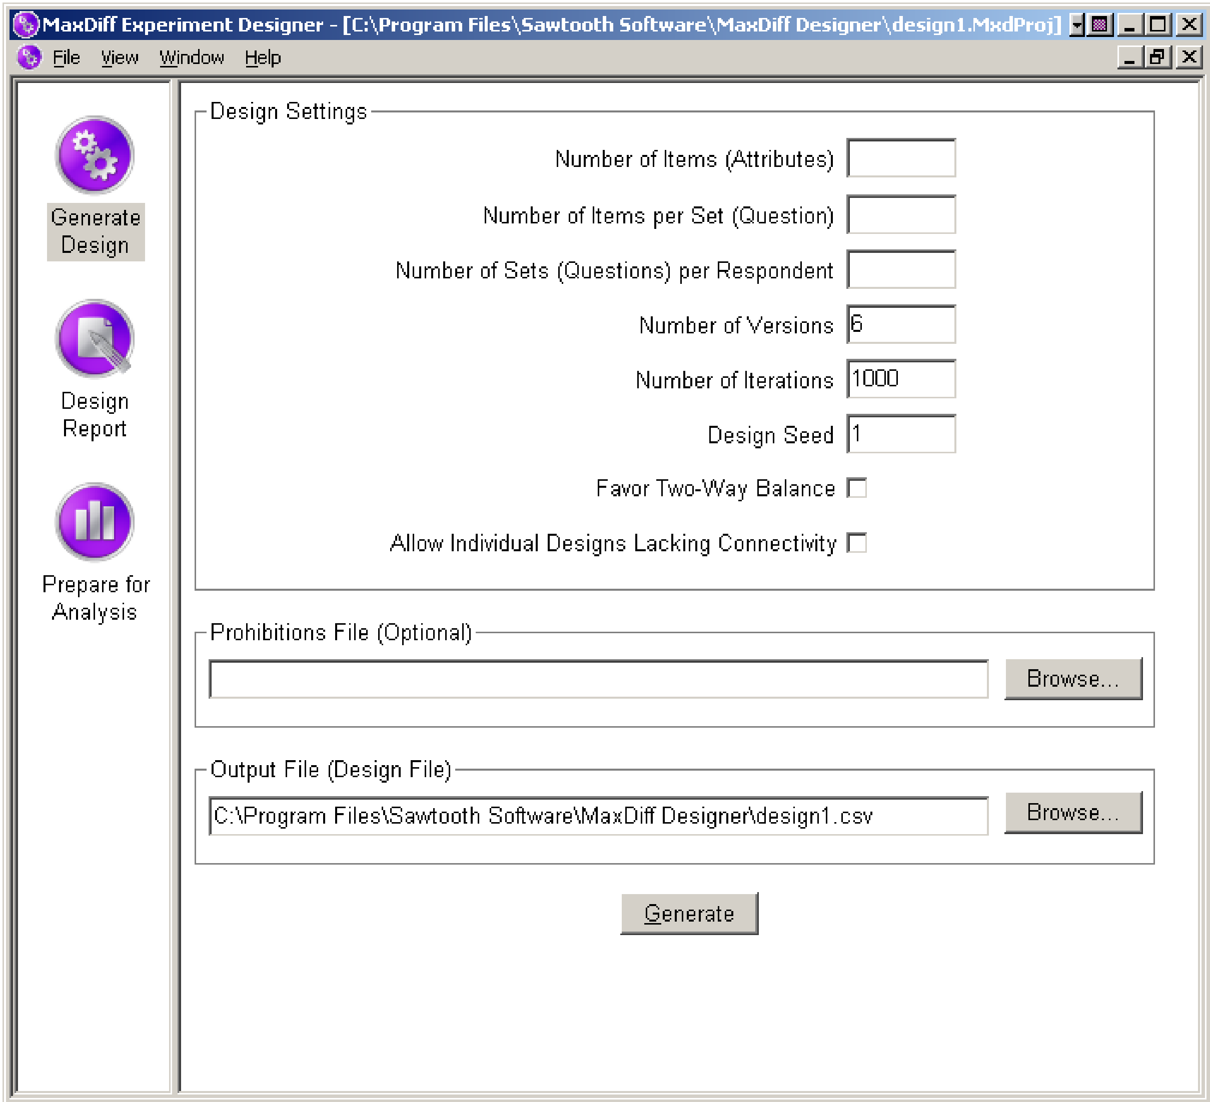

There are three main views (screens) to navigate among: *Generate Design*, *Design Report*, and *Prepare for Analysis*.

When you are at the *Generate Design* screen, this dialog presents a number of fields for you to specify:

Number of Items (Attributes)

In this field, you specify the number of items (i.e. levels or features) to include in the exercise. Depending on your license, you can specify a number up to 500.

Number of Items per Set (Question)

Generally, we recommend displaying either four or five items at a time (per set or question) in MaxDiff questionnaires. However, we do not recommend displaying more than half as many items as there are items in your study. Therefore, if your study has just eight total items, we would not recommend displaying more than four items per set.

If you would like to use the Method of Paired Comparisons approach, then specify just two items per set.

Research using synthetic data suggests that asking respondents to evaluate more than about five items at a time within each set may not be very useful in MaxDiff studies. The gains in precision of the estimates are minimal when using more than five items at a time per set. The small statistical gains from showing even more items may be offset by respondent fatigue or confusion.

Number of Sets (Questions) per Respondent

If using HB to estimate individual-level scores, we generally recommend asking as many sets (questions) per respondent such that each item has the opportunity to appear from three to five times per respondent. For example, consider a study with 20 items where we are displaying four items per set. With 20 total sets, we know that each item will be displayed 4 times (assuming the design is perfectly balanced). This leads to the following decision rule and formula:

For best results (under HB estimation), the suggested number of sets is at least:

3K/k

where **K** is the total number of items in the study, and **k** is the number of items displayed per set.

The software will warn you if the number of sets you have requested leads to each item being displayed fewer than 2 times per respondent.

These recommendations assume HB estimation. It is also possible to analyze the data using Sawtooth Software's Latent Class system, or another analytical program of your choice. When pooling data across respondents using Latent Class or aggregate logit, you may be able to achieve reasonable population-level scores using even fewer sets per respondent than these recommendations.

Number of Versions

Though it is possible to estimate scores relatively efficiently using just a single questionnaire version for all respondents, there is practical benefit to using multiple versions (sometimes called "blocks") of the questionnaire. With multiple versions of the questionnaire, different respondents see different series of questions. Across respondents, this dramatically increases the variation in the way items are combined within sets, which can reduce potential context biases (which are usually modest).

The default is 6 questionnaire versions, which assumes paper-and-pencil interviewing and leads to a modest degree of combinatorial variation in the sets. More versions would be better. However, if conducting paper-and-pencil studies, it usually is difficult to manage multiple versions. Therefore, often just a few versions (such as 4 to 6) are adequate to obtain precise estimates of item scores and to reduce order and context effects.

Number of Iterations

By default, the MaxDiff designer repeats its algorithm 1000 times and returns the best design found across those passes. The quality of the design is based on the one-way, two-way, and positional balance (in that order of precedence). For many designs, 1000 iterations can be done within a few seconds. We encourage you to increase the number of passes, as it may enable you to find a slightly better design. Up to 1,000,000 iterations may be specified.

Design Seed

Specify a value from 1 to 999,999,999. This integer (seed) is used for determining the starting point for the design algorithm. Different seeds will yield different designs, all having approximately the same overall design efficiency. You can try different starting seeds to see if you obtain slightly better designs.

Favor Two-Way Balance

The MaxDiff designer by default puts primary emphasis on balancing the number of times each item is included (one-way balance), with secondary emphasis on two-way balance (how often each paired combination of levels appears together). If you have specified prohibitions, these two goals are at odds with one another. The software puts more emphasis toward balancing the two-way frequencies if you check this box. Recognize, however, that with prohibitions it may be impossible to achieve near-perfect balance in either the one-way or two-way frequencies.

We strongly suggest checking this box when displaying two items at a time within sets (pairwise presentation). The design algorithm will almost always produce one-way frequencies that differ very little from when one-way balance is favored. But, the two-way frequency balance will usually be far superior.

Allow Individual Designs Lacking Connectivity

Connectivity is the property that the items cannot be divided into two sets wherein no comparison is made between any item in one set and another item from the other. Even if creating choice sets randomly, if respondents receive enough sets, it would be difficult for the items to lack connectivity for any given respondent. However, if you select a design with relatively many items and relatively few tasks per respondent, many individual versions of the questionnaire could lack connectivity. The default for the software is to reject any version (questionnaire for one respondent) that lacks connectivity. But in the case of many items and few sets per respondent, it may become impossible to satisfy connectivity. If you would like to permit versions lacking connectivity (which may be the case if you ask few questions of any one respondent, probably leading to aggregate analysis), then you can check this box. Make sure to specify many questionnaire versions. To ensure stable estimation, advanced users may wish to generate dummy (random) response data, prepare the data for estimation under aggregate logit, and examine the size of the standard errors to ensure reasonable precision.

Prohibitions File Name

Optionally, you can specify a text-only (.txt) file (e.g. constran.txt) that includes prohibitions (e.g. item 4 never should appear in the same set as item 5). The format is one prohibition per line, where the item numbers are separated by commas, such as:

```
4,5
5,8
9,10
```

Above is an example of three prohibited combinations (items 4 and 5; 5 and 8; and 9 and 10). Prohibitions decrease design efficiency, though we have found design efficiency to be quite robust when a modest pattern and number of prohibitions is specified.

When you compute the ***Generate*** button, if the software returns warnings that the questionnaire versions lack connectivity, this is clear sign that your prohibitions are excessive (after ruling out the possibility that you are not displaying enough sets per respondent).

Output File (Design File) Name

Browse to a folder and type a file name, such as Design1.csv. The file will be written in comma-separated text format (.csv) which may be automatically opened in Excel. For convenience with working with Excel, we suggest you specify a file with .csv extension. Alternatively, you may also specify a file with .txt extension if you wish to open the document in a word processing program.

Generate

When you are pleased with your settings, click ***Generate Design*** to produce a design and write the information to the Output File.

The design algorithm follows these guidelines:

1. Create a design that features item **frequency balance**. A balanced design is one in which the one-way frequencies are nearly equivalent (how many times each level appears across the entire design) and two-way frequencies are also nearly equivalent (how many times each pair of items appears within the same set across the entire design). When one-way and two-way frequencies are balanced, this is termed **orthogonality**.

2. Ensure that each version of the design (respondent questionnaire) features **connectivity** (meaning that all items are linked either directly or indirectly), unless the *Allow Individual Designs Lacking Connectivity* option is checked. Without connectivity it becomes very difficult for HB to scale the items properly relative to one another for each respondent.

3. After the designs have been generated based on the guidelines in 1 and 2, swap the order of the items within each set so that each item appears approximately an equal number of times in each position. **Positional balance** reduces or virtually eliminates order bias. Finally, randomize the order of the tasks within each version.

The design process is repeated 1000 separate times, and the replication that demonstrates the best one-way balance is selected. If multiple designs have the same degree of one-way balance, then we select among those designs based on the best two-way balance. If multiple designs have the same degree of one-way and two-way balance, then we select among those designs based on the best positional balance. With small to moderate sized designs and no prohibitions, this usually happens within a few seconds.

3 The Design Report

After the design has been generated, a report window appears showing the one-way and two-way frequencies, along with positional frequencies. If you move off this report window, you can access the design report again for the current project by clicking the *Design Report* icon in the navigation panel on the left.

Here is the report for an example study of just 8 items, for just 2 questionnaire versions:

```
The following design report is for the MaxDiff design file ('C:\Program
Files\Sawtooth Software\MaxDiff Designer\design1.txt')
which was created on Friday, March 03, 2006 at 7:56:23 AM

Number of Items (Attributes): 8
Number of Items per Set: 4
Number of Sets per Respondent: 10
Number of Versions: 2
Number of Iterations: 1000
Random Number Seed: 999
Iteration 634 was chosen


One Way Frequencies:
 Item  Times Used
-----------------
    1      10
    2      10
    3      10
    4      10
    5      10
    6      10
    7      10
    8      10

Mean = 10.000000
Std Dev. = 0.000000


Two Way Frequencies:
 Item|       1       2       3       4       5       6       7       8
-----+-----------------------------------------------------------------
    1|      10       4       4       4       5       4       5       4
    2|       4      10       4       4       5       5       4       4
    3|       4       4      10       4       4       4       5       5
    4|       4       4       4      10       4       4       5       5
    5|       5       5       4       4      10       5       3       4
    6|       4       5       4       4       5      10       4       4
    7|       5       4       5       5       3       4      10       4
    8|       4       4       5       5       4       4       4      10

Off Diagonal Non-prohibited Elements
Mean = 4.285714
Std Dev. = 0.524891
```

```
Positional Frequencies:
          | Pos.       1       2       3       4
----------+------------------------------------
Item     1|            3       2       2       3
         2|            2       3       3       2
         3|            3       2       2       3
         4|            3       2       3       2
         5|            3       3       2       2
         6|            2       3       2       3
         7|            2       2       3       3
         8|            2       3       3       2

Mean = 2.500000
Std Dev. = 0.500000
```

In this example, the one-way frequencies are perfectly balanced, and the off-diagonal two-way frequencies nearly so. It is not necessary to achieve *exact* balance in a design to have a very satisfactory design, but balance is a desirable property. Methods such as Logit, Latent Class and HB do not require perfect balance to achieve unbiased estimates of parameters. The Positional Frequencies report how many times each item appears in the first, second, third, and so on positions.

A standard deviation is reported for each table. Lower standard deviations are better, with a standard deviation of zero reflecting perfect balance. With most design specifications (number of items, items per set, sets, and versions), it is impossible to achieve exact balance in all three tables. Again, exact balance is not required for near-optimal efficiency and unbiased estimation of parameters. However, the researcher interested in achieving slightly better designs should try different starting seeds and/or more iterations, and compare the results.

4 The Design File

A design is written to the output file specified in the *Output File Name* field. The design tells you the item numbers that need to be presented to respondents in each choice set. An example is shown below, for two questionnaire versions. If you use a .csv extension when you specify your output file, you can directly open the file in Excel. The file might look something like this when opened in Excel:

Version	Set	Item1	Item2	Item3	Item4
1	1	8	5	4	1
1	2	4	7	3	2
1	3	2	4	7	1
1	4	1	3	8	7
1	5	5	1	6	3
1	6	6	2	5	7
1	7	4	6	2	5
1	8	3	6	8	4
1	9	7	8	2	3
1	10	7	5	1	8
2	1	1	3	5	2
2	2	3	5	7	6
2	3	5	8	2	4
2	4	4	7	3	1
2	5	2	1	6	8
2	6	6	2	8	3
2	7	8	4	7	6
2	8	3	8	4	5
2	9	5	2	1	6
2	10	1	6	4	7

There are just two versions of the questionnaire in this example. Each version of this questionnaire contains 10 sets. And, the four items to show in each set are specified. For example, the first set of the first questionnaire version includes items 8, 5, 4 and 1.

5 Preparing for Analysis (Creating .CHO file)

Creating the File of Respondent Answers

Once you have collected respondent data, you need to prepare a text-only file containing respondent answers. Such a file can be created with Excel (saving the file in .csv, .txt, or .prn format). Or, you can create this file using a text editor (saving as text-only .txt format).

The file of respondent answers should contain one respondent per row. The data should be organized as:

Respondent#, Version#, b1, w1, b2, w2, ... and so on.

The data may be separated by spaces, tabs, comas or semicolons.

Where b1 is the answer to the "best" from set1, w1 is the answer to the "worst" from set1, etc. These answers may either reflect the item# itself ("Answers Are Item Numbers"), or the position in the set ("Answers Are Alternatives". For example, consider a set including items 8, 5, 4 and 1. Assume the respondent chooses item 1 as best and item 5 as worst. Answers may be coded as "item numbers" (b1=1, w1=5) or as "Alternatives" (b1=4, w1=2).

If using method of paired comparisons (MPC, showing only 2 items per set), then only a "best" answer is recorded for each set.

Preparing Files for Analysis

When you click the Prepare for Analysis icon, the following screen is displayed:

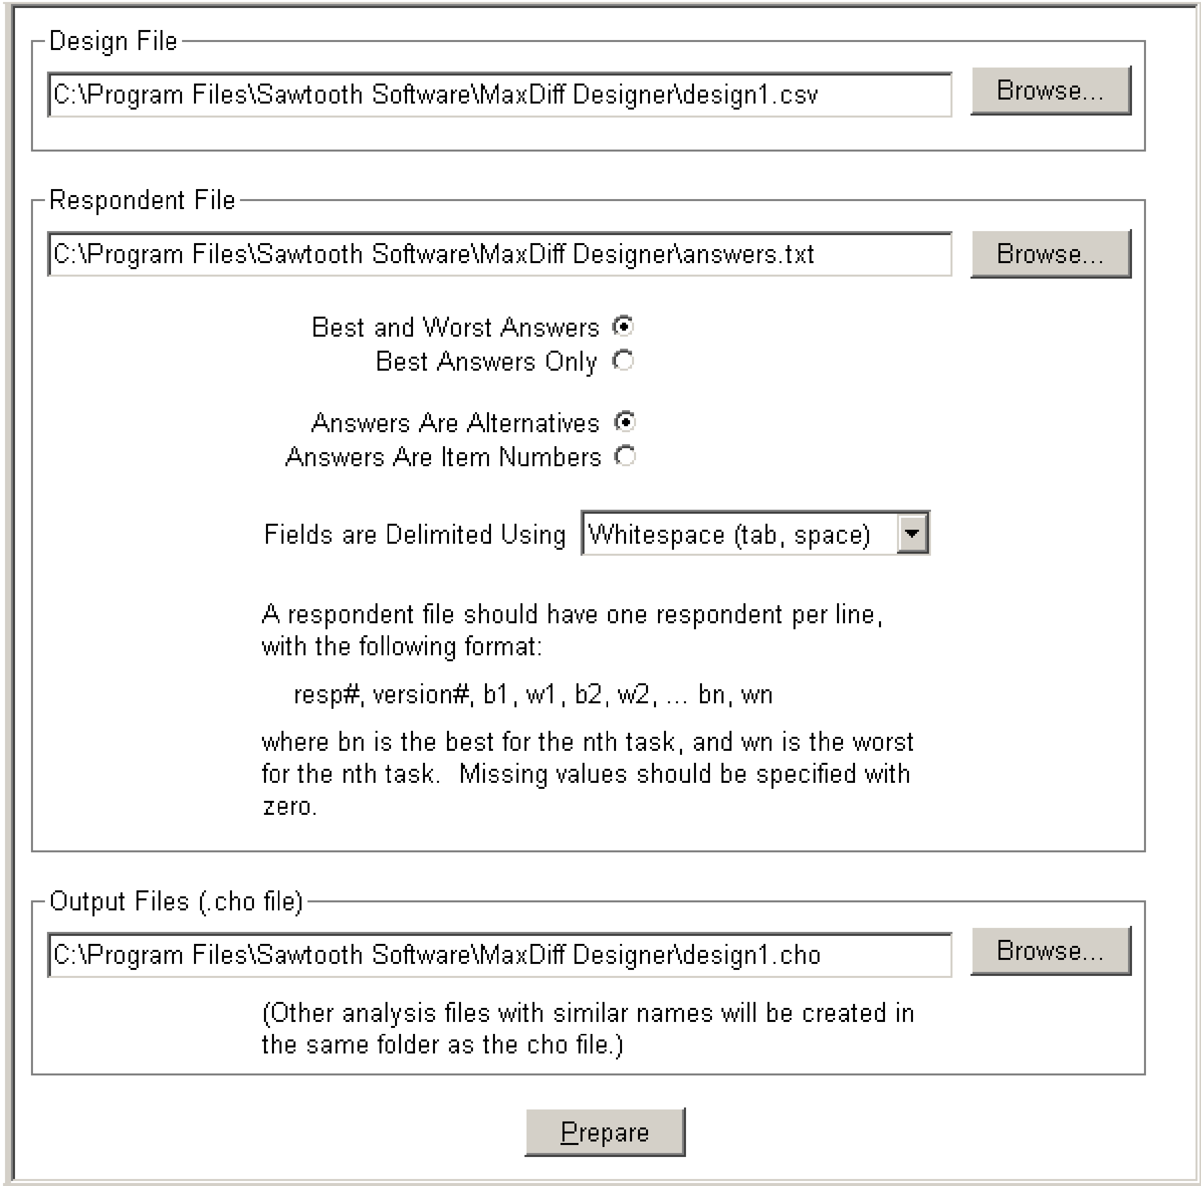

The name of the design file you previously generated for this project should be displayed in the **Design File** field.

Browse to the file of respondent answers you prepared (as described above). After you browse for and select that file, it should appear within the **Respondent File** field (e.g. C:\Program Files\Sawtooth Software\MaxDiff Designer\answers.txt)

Specify whether the respondents answers are "Alternatives" or "Item Numbers" (see description above for clarification).

Select the appropriate Delimiter for your Respondent File.

Browse to the folder and type the name of the file (.cho) in which to write the output (e.g. C:\Program Files\Sawtooth Software\MaxDiff Designer\design1.cho)

When you click ***Prepare***, the following files are prepared (within the same folder) for use with Latent

Class v2.5 or CBC/HB software:

yourfilename.cho (choice data file)
yourfilename.att (attribute labels file)
yourfilename.val (attribute level values file)
yourfilename.eff (effects file)
yourfilename.mtrx (prior covariance matrix file, for use in CBC/HB only)

6 Output Files for Latent Class and CBC/HB Estimation

This section is for the interested reader, and describes the layout of various files required for Latent Class v2.x and CBC/HB estimation.

The MaxDiff Designer software automatically creates the following files when ***Prepare*** is clicked from the *Prepare for Analysis* screen.

Studyname.ATT (Attribute level labels file)
Studyname.EFF (Effects file)
Studyname.VAL (Level values file)
Studyname.CHO (Choice data file)
Studyname.MTRX (Prior covariance file for HB estimation)

Additionally, the following two files are required for Latent Class v2.5 estimation. You can create these files with a text editor:

Names.CBC (Study name file, *required only for Latent Class*)
Studyname.PAR (Parameters file, *required only for Latent Class*)

For Latent Class v2.5 estimation, you must place all files in the same folder as your Lbuild.EXE, Lclass.EXE, and Userid.lcl files.

Studyname.ATT is also a text-only file, containing labels for the items (attributes) in your best/worst study (remember, if you have 8 items in total, you'll be using n-1, or 7 labels in the .ATT file). To conform to the rule in standard Latent Class or CBC/HB that each attribute have at least two levels, we represent each best/worst item as a two-level attribute. For example, your bwexamp.ATT file should contain:

Red
(use a single blank line here to separate the levels)
Red
(use two blank lines here to separate attributes, etc.)

Green

Green

Blue

Blue

Studyname.EFF is a text-only file, containing the attribute numbers to be used in analysis. If you have 8 items in total in your best/worst study, you need to follow a dummy-coding procedure in which you estimate n-1 or 7 parameters (the utility of the omitted parameter is set to zero, and all other utilities are scaled relative to the omitted item). Assuming 8 items, specify the following on the first line of your bwexamp.EFF file:

#1 #2 #3 #4 #5 #6 #7

The "pound notation" is specialized for use when the analyst wishes to override the typical effects-coding procedure that Latent Class v2 and CBC/HB uses for conjoint data.

Studyname.VAL is also required whenever using "pound notation," and is a text-only file with two lines per item in your study. Again, assuming an 8 item study, you are only estimating 7 parameters, and your bwexamp.VAL file should have 14 total lines (2 for each of the 7 parameters). The first six lines look like:

```
"Attribute 1 Label"
1 2
"Attribute 2 Label"
1 2
"Attribute 3 Label"
1 2   (etc.)
```

Although we are only estimating a single parameter for each attribute, we list the level values as "1" and "2." This is just to satisfy our software (which was designed for conjoint experiments, where every attribute has at least two levels), and to ensure that the scaling of data in the text report that we produce is proper.

Studyname.CHO What sets best/worst data apart from traditional conjoint/choice data is that each set is coded twice: once to represent the item chosen as "best" and once for the item selected "worst." For a questionnaire involving 10 sets per respondent, we code these as 20 separate sets. Each respondent's data occupies multiple lines in the file, and the next respondent follows on the line directly beneath the previous (NO blank lines between respondents).

Assume respondent number 1001 received items 7, 8, 3, and 2 in the first of ten sets (and the study involved 8 total items). Further assume that item 7 was selected as this respondent's "best" and item 3 as the "worst." The first few lines of this file representing the coding for respondent 1001's first set, should look something like:

```
1001  0   7   20   0
4  1
 0   0   0   0   0   0   1
 0   0   0   0   0   0   0
 0   0   1   0   0   0   0
 0   1   0   0   0   0   0
1  99
4  1
 0   0   0   0   0   0  -1
 0   0   0   0   0   0   0
 0   0  -1   0   0   0   0
 0  -1   0   0   0   0   0
3  99
(etc. for 9 more sets)
```

The exact spacing of this file doesn't matter. Just make sure it is in text-only format and that you have arranged the data on separate lines, and that the values are separated by at least one space. We describe each line as follows:

Line 1

1001	0	7	20	0
Respondent #1001	No segmentation variables	7 attributes	20 total sets	No "None"

Lines 2 through 7 reflect the information for the "best" item chosen from set #1.

Line 2

4	1
Next follows a set with 4 items	One selection from this set

Line 3

0 0 0 0 0 0 1
Dummy codes representing the first item in the first set (item 7 in our example). Each value represents an item (less the last item, which is omitted in the coding). The dummy codes are "0" if the item is not present, and "1" if the item is present. Since this row represents item 7, the 7th value is specified as 1.

Line 4

0 0 0 0 0 0 0
Dummy codes representing item 8. In our study, if the 8th item is present, all seven values are at 0.

Lines 5 and 6 follow the same formatting rules for dummy coding, to code items 3 and 2 in our example. Next follows the line in which the respondent's "best" item is indicated.

Line 7

1	99
The item in row 1 of this set is best	A filler value (time) of 99 to be compatible with Latent Class

Lines 8 through 13 reflect the information for the "worst" item chosen from set #1.

Line 8

4	1
Here follows a set with 4 items	One selection from this set

Lines 9 through 12 reflect the dummy codes (inverted) for the items shown in set one, considered with respect to the "worst" item selected. All values that were "1" in the previous task are now "-1".

Line 13

3	99
The item in row 3 of this set is worst	A filler value (time) of 99 to be compatible with Latent Class

Names.CBC is a text-only file that contains the studyname on the first line (must be 8 characters or fewer). For our examples, we'll assume a studyname called "bwexamp."

Using a text editor like "Notepad" or "Wordpad," create a text-only file called **Names.CBC** that has the text **bwexamp** on the first line. Make sure your text editor doesn't put a ".TXT" extension on the end of your filename, or Lbuild.exe and Lclass.exe will not be able to recognize the study names file. (Same warning applies to all text files you create.)

Studyname.PAR is an optional file for use only with Latent Class. This file has the same options and format as specified in the Latent Class documentation. We strongly encourage you to use a .PAR file so that you can tell Latent Class to perform multiple replicates, selecting the run with the highest fit. Without a .PAR file, the default is to run all solutions between 1-group and 5 groups, with just one replication per solution. The file may look something like this:

```
1  5  0  0  0  1  1  100  .01  0  1
```

The values are typed on a single line, separated by a space. There are 11 values in the file, meaning:

Position	**Default**	**Information**
1	1	minimum number of groups
2	5	maximum number of groups
3	0	weight respondents (0=no or 1=yes)
4	0	report standard erors (0=no or 1=yes)
5	0	tabulate all pairs of solutions (0=no or 1=yes)
6	1	save respondent probabilities (0=no or 1=yes)
7	1	display re-scaled utilities (0=no or 1=yes)
8	100	maximum number of iterations
9	.01	convergence limit for log likelihood
10	0	random seed (0 to use time of day)
11	1	number of replications for each solution

7 Running Latent Class and CBC/HB

Running Latent Class

Double-click the **lbuild.exe** file (or type **lbuild** and press ENTER from the DOS prompt). A report is shown specifying how many choices are made for each category, and the total number of respondents. Press any key to continue.

Double-click the **lclass.exe** file (or type **lclass** and press ENTER from the DOS prompt). Based on the settings in the optional .PAR file, the software produces studyname.LCS and studyname.LCU reports. The utility weights for the 7 items are written to studyname.LCU, for each group. Please refer to the Latent Class documentation for details. Although it is not explicitly given in the report, **the utility of the last omitted item (due to dummy coding) is 0**. This is the "reference level" and the utility values of the "explicit" levels are estimated *relative* to this value.

Running CBC/HB

CBC/HB can also be used to estimate parameters for MaxDiff experiments, and it results in individual-level estimates rather than group estimates. As with latent class analysis, **the utility of the last omitted item (due to dummy coding) is 0**. This is the "reference level" and the utility values of the "explicit" levels are estimated *relative* to this value.

We should note that when using HB to estimate parameters for many items under dummy coding, the estimates of the parameters (relative to the reference "0" level) can sometimes be distorted downward, often quite severely when there are many items in the study and when the number of questions asked of any one respondent is relatively few. (This distortion of course makes it appear as if the "omitted" item's estimate is "too high" relative to the others.) To see if this difficulty is appearing for your data set, you might try coding a different level as the "omitted" item and compare the results. Or, you could take the step described below to avoid the problem.

To avoid potential problems when dealing with dummy coding under HB and very sparse data conditions, you can specify a more appropriate "user-defined prior covariance matrix" for CBC/HB v3.2 or later. (Even though the software automatically creates "proper" prior covariance matrices if you select dummy coding or effects coding from CBC/HB's dialog, since with MaxDiff designs you are "commandeering" the design matrix with the "#" notation usage in the .eff file, the software cannot know what prior covariance structure is appropriate.) The MaxDiff designer automatically creates a **studyname.mtrx** file, where "STUDYNAME" is the name of your study. This file should contain an n x n matrix of values, where n is equal to the number of parameters estimated in your model. If you wish to specify a prior variance of 1.0 (a typical default), the matrix is composed of "2"s across the main diagonal, and "1"s in the off-diagonal positions, such as:

```
2 1 1 1 . . .
1 2 1 1 . . .
1 1 2 1 . . .
1 1 1 2 . . .
. . . . . . .
. . . . . . .
. . . . . . .
```

To specify a different prior variance, multiply all values in the prior covariance matrix above by the desired variance constant.

Transforming Weights to the 0-100 Scale

To convert the raw weights to the 0-100 point scale, perform the following steps for each respondent:

1. Take the antilog of each weight (including for the "omitted" item, which has a weight of zero). In Excel, this is accomplished using the "=exp(A1)" function, where in this example the weight to be transformed is located in cell A1.
2. For each respondent, sum the transformed weights from step 1.
3. Divide each of the weights in step 2 by the sum, and multiply each by 100. The weights should now range from 0 to 100, and the sum of the weights should be exactly 100 for each respondent.