

# Getting a DIF Breakdown with Lertap

Larry Nelson

Curtin University of Technology

Document date: 8 October 2009 (links updated 6 August 2014)

*This document shows how Lertap 5 may be used to look for differences among groups of test takers. Given two groups, say, for example, males and females, did one group out-perform the other, getting higher test scores?*

*Even if the groups appeared to have similar proficiency on the subject matter covered by the test, may there nonetheless have been group differences at the item level? Was there evidence of "DIF", differential item functioning?*


*I'll use a 49-item multiple-choice test, professionally developed by a large-scale test centre, for my examples, showing how the Excel 2007 version of Lertap 5 may be used to answer questions such as these. I'll also demonstrate how to use SPSS to get DIF results.*


*The discussion presupposes some familiarity with Lertap, an Excel-based system. More information about Lertap is available via the URL seen at the top of this page.*

## The data

More than four thousand high-school students sat a 60-item test. The test had ten trial items (not scored), and, as it unfortunately turned out, one item whose stem had a substantial error, not caught before printing. This item was also omitted from scoring, leaving 49 four-option multiple-choice items scored on a right / wrong basis, with one point for each correct answer.

My first concern was to see how the distribution of test scores looked. I of course used Lertap to produce a couple of histograms.





What I particularly wanted to investigate was the number of zero scores. There turned out to be three:

Distribution of "G431-49", as at 4/09/20

z	score	f	%	cf	c%
-3.30	0.00	3	0.1%	3	0.1%
-3.19	1.00	0	0.0%	3	0.1%
-3.07	2.00	0	0.0%	3	0.1%
-2.95	3.00	0	0.0%	3	0.1%
-2.84	4.00	2	0.0%	5	0.1%
-2.72	5.00	9	0.2%	14	0.3%
-2.60	6.00	9	0.2%	23	0.5%
-2.49	7.00	11	0.2%	34	0.7%
-2.37	8.00	19	0.4%	53	1.1%
-2.26	9.00	19	0.4%	72	1.5%
-2.14	10.00	25	0.5%	97	2.1%
-2.02	11.00	28	0.6%	125	2.7%
-1.91	12.00	45	1.0%	170	3.6%
-1.79	13.00	60	1.3%	230	4.9%

To find the Data records corresponding to zero scores, I used Lertap's "Sort" option, found in the collection of icons grouped under the "Basic options" section of Lertap's Excel 2007 ribbon tab:



Records 1024, 1028, and 1067 had missing data for each of the 49 test items. I deleted them from the Data worksheet, and started again (that is, once again ran Lertap's "Interpret" and "Elmillon" options).


## The test

I looked at overall test quality by first going to the Stats1f report, then the bottom of the Stats1b report, after which I used the "Res. Charts" option to generate quintile plots.


### Summary statistics

<b>number of scores (n):</b>	4,712	
<b>lowest score found:</b>	4.00	(8.2%)
<b>highest score found:</b>	49.00	(100.0%)
<b>median:</b>	29.00	(59.2%)
<b>mean (or average):</b>	<u>28.44</u>	<u>(58.0%)</u>
<b>standard deviation:</b>	8.58	(17.5%)
<b>standard deviation (as a sample):</b>	8.58	(17.5%)
<b>variance (sample):</b>	73.63	


<b>number of subtest items:</b>	49	
<b>minimum possible score:</b>	0.00	
<b>maximum possible score:</b>	49.00	
<b>reliability (coefficient alpha):</b>	<u>0.87</u>	
<b>index of reliability:</b>	0.93	
<b>standard error of measurement:</b>	3.07	(6.3%)



This wasn't really my test, but it seemed to have reasonable quality – coefficient alpha was comfortable at 0.87, and the little scatterplot of item discrimination and difficulty, found towards the bottom a Stats1b report, suggests that only one item was sort of an "odd-man out". Here I refer to the only item having a discrimination value less than 0.10, which was item "q32".



Lertap's standard quintile plot for q32 clearly indicates a weak test item. For example, compare its plot to that for a nicely-performing item, q27:

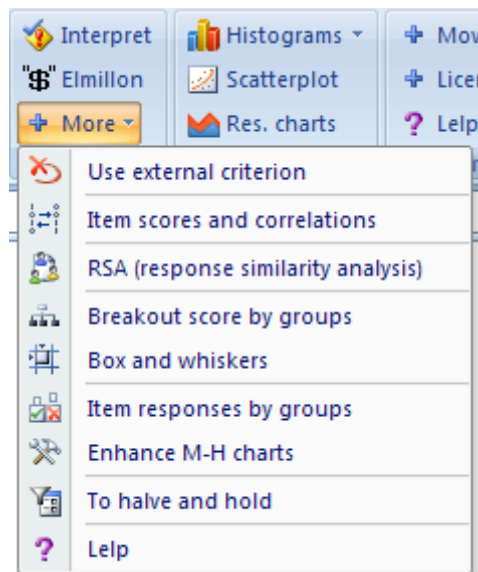


There's more about using quintile plots here:

[larrynelsonstuff.com/Documentation/UsingLertapQuintilePlots.pdf](http://larrynelsonstuff.com/Documentation/UsingLertapQuintilePlots.pdf)

## The groups

Gender information was coded in column 4 of the Data worksheet, with "M" for males, and "F" for females.




Were there gender differences on the test? I used the "Breakout score by groups" option as one means of answering the question. This option prompted Lertap to create two new reports, or worksheets. I copied the following from the "Breaks1" report:

Lertap5 breakout of G431-49 scores by gender			
G431-49	F	M	
n	2,153	2,559	
Min	4.00	4.00	
Median	29.00	29.00	
Mean	28.81	28.13	
Max	49.00	47.00	
s.d.	8.41	8.70	
var.	70.80	75.77	
Range	45.00	43.00	
IQRange	12.00	13.00	
Skewness	-0.24	-0.24	
Kurtosis	-0.42	-0.64	
MinPos	0.00	0.00	
MaxPos	49.00	49.00	
Analysis of variance			
	df	SS	MS
Between	1	541	541
Within	4710	346330	74
Total	4711	346871	
F ratio:	7.35	.01 (<-sig.)	
eta <sup>2</sup> :	0.00		

Not surprisingly, there was a statistically significant result; were I to hypothesize that the population means were equal, I'd reject the hypothesis at the .01 level. But this is not my interest. With more than 2,000 people in each sample, statistical significance would not be my focus, not at all.

The  $\eta^2$  index, a measure of practical significance, indicates that the difference in sample means was nothing to crow about (that is to say, was not at all substantial, hardly even measurable, from a practical standpoint).

Next, I took the "Box and whiskers" option – I wanted a picture of the group results. I found that the boxes and whiskers were about the same for F and M:



I like Lertap's simple histograms. I went back to the Breaks1 report, and then clicked on Lertap's "Histograms" option. Twice I clicked, once for the girls' results, once for the boys'. With a little cutting and pasting, and a scale reduction, I made the following, females on the left, males on the right. The results, as presented below, are hard to read, but I think the gestalt is visible (note that low scores are at the top).

z	score	f	%	cf	c%
-2.95	4.00	1	0.0%	1	0.0%
-2.83	5.00	7	0.3%	8	0.4%
-2.71	6.00	5	0.2%	13	0.6%
-2.59	7.00	4	0.2%	17	0.8%
-2.47	8.00	6	0.3%	23	1.1%
-2.35	9.00	6	0.3%	29	1.3%
-2.24	10.00	8	0.4%	37	1.7%
-2.12	11.00	7	0.3%	44	2.0%
-2.00	12.00	18	0.8%	62	2.9%
-1.88	13.00	20	0.9%	82	3.8%
-1.76	14.00	26	1.2%	108	5.0%
-1.64	15.00	25	1.2%	133	6.2%
-1.52	16.00	40	1.9%	173	8.0%
-1.40	17.00	46	2.1%	219	10.2%
-1.28	18.00	57	2.6%	276	12.8%
-1.17	19.00	48	2.2%	324	15.0%
-1.05	20.00	54	2.5%	378	17.6%
-0.93	21.00	50	2.3%	428	19.9%
-0.81	22.00	81	3.8%	509	23.6%
-0.69	23.00	75	3.5%	584	27.1%
-0.57	24.00	84	3.9%	668	31.0%
-0.45	25.00	63	3.2%	737	34.2%
-0.33	26.00	86	4.0%	823	38.2%
-0.21	27.00	38	1.8%	861	40.0%
-0.10	28.00	34	1.6%	895	41.7%
0.02	29.00	30	1.4%	925	43.0%
0.14	30.00	100	4.6%	1025	47.5%
0.26	31.00	104	4.8%	1129	50.3%
0.38	32.00	81	3.8%	1210	54.1%
0.50	33.00	84	3.9%	1294	58.0%
0.62	34.00	98	4.6%	1392	62.5%
0.74	35.00	91	4.2%	1483	66.8%
0.85	36.00	73	3.4%	1556	70.2%
0.97	37.00	74	3.4%	1630	73.6%
1.09	38.00	70	3.3%	1700	76.9%
1.21	39.00	65	3.0%	1765	80.3%
1.33	40.00	53	2.5%	1818	83.3%
1.45	41.00	50	2.3%	1868	86.1%
1.57	42.00	33	1.5%	1901	88.2%
1.69	43.00	24	1.1%	1925	89.3%
1.81	44.00	20	0.9%	1945	90.2%
1.92	45.00	17	0.8%	1962	91.0%
2.04	46.00	8	0.4%	1970	91.4%
2.16	47.00	10	0.5%	1980	91.9%
2.28	48.00	2	0.1%	1982	92.0%
2.40	49.00	1	0.0%	1983	92.0%

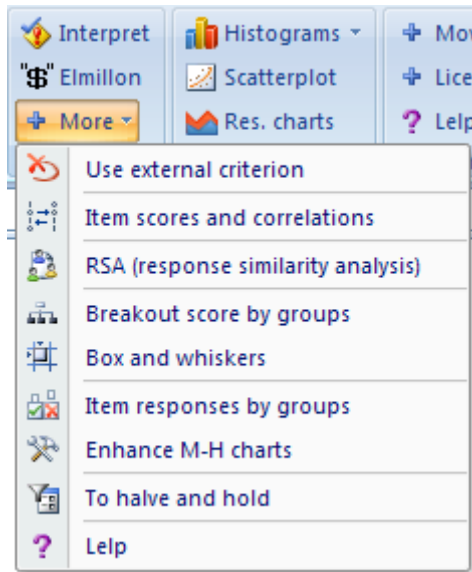
f	%	cf	c%
1	0.0%	1	0.0%
2	0.1%	3	0.1%
4	0.2%	7	0.3%
7	0.3%	14	0.5%
13	0.5%	27	1.1%
13	0.5%	40	1.6%
17	0.7%	57	2.2%
21	0.8%	78	3.0%
27	1.1%	105	4.1%
40	1.6%	145	5.7%
33	1.3%	178	7.0%
58	2.3%	236	9.2%
64	2.5%	300	11.7%
51	2.0%	351	13.7%
63	2.5%	414	16.2%
51	2.0%	465	18.2%
75	2.9%	540	21.1%
79	3.1%	619	24.2%
81	3.2%	700	27.4%
92	3.6%	792	30.9%
74	2.9%	866	33.8%
79	3.1%	945	36.9%
85	3.3%	1030	40.3%
105	4.1%	1135	44.4%
118	4.6%	1253	49.0%
104	4.1%	1357	53.0%
34	1.3%	1451	57.7%
117	4.6%	1568	61.3%
110	4.3%	1678	65.6%
113	4.4%	1791	70.0%
94	3.7%	1885	73.7%
107	4.2%	1992	77.8%
92	3.6%	2084	81.4%
88	3.4%	2172	84.9%
79	3.1%	2251	88.0%
71	2.8%	2322	90.7%
68	2.7%	2391	93.4%
48	1.9%	2439	95.3%
40	1.6%	2479	96.9%
32	1.3%	2511	98.1%
17	0.7%	2528	98.8%
14	0.5%	2542	99.3%
8	0.3%	2550	99.6%
3	0.1%	2553	100.0%

The negative skewing is a bit more noticeable among the males' scores (seen on the right), but then there is just a bit of an extra blip at the very low end of the females' scores (seven girls got a score of 5.00 on the test, compared to two boys).

Were there gender differences? If yes, they're slight, very slight. Very.

### The groups at an item level

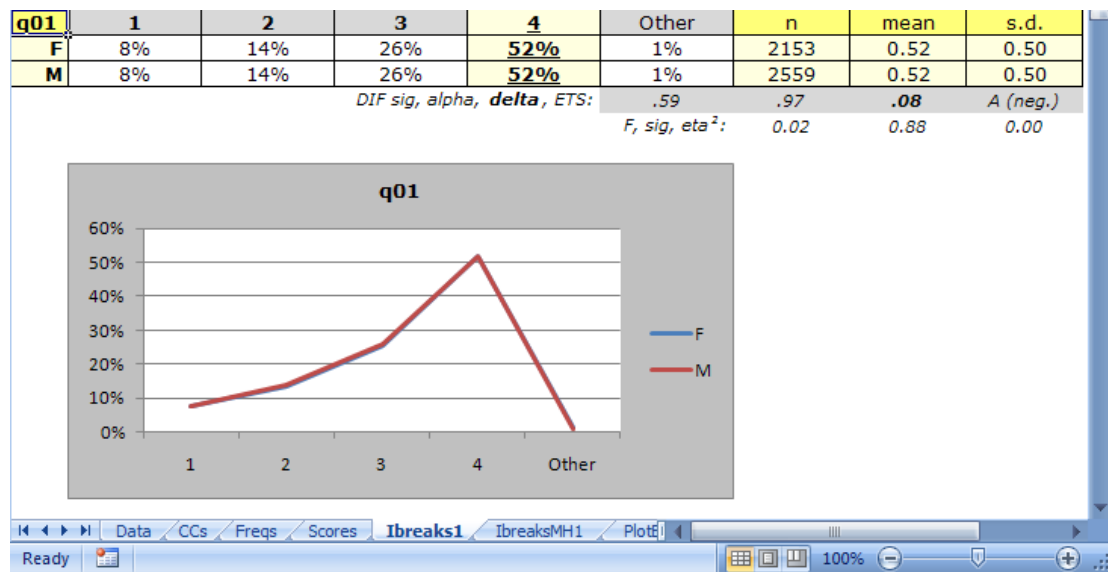
On the overall test there were no meaningful gender differences. But, what if I were to drill deeper, down to the level of results for individual items? Might there be gender differences there?



To answer this, I click on the "Item responses by groups" option.

When you try this option, you'll note that Lertap can present quite a number of questions, asking for instructions, especially when there are just two groups. This reflects Lertap's support for users who may be looking for "DIF", differential item functioning.

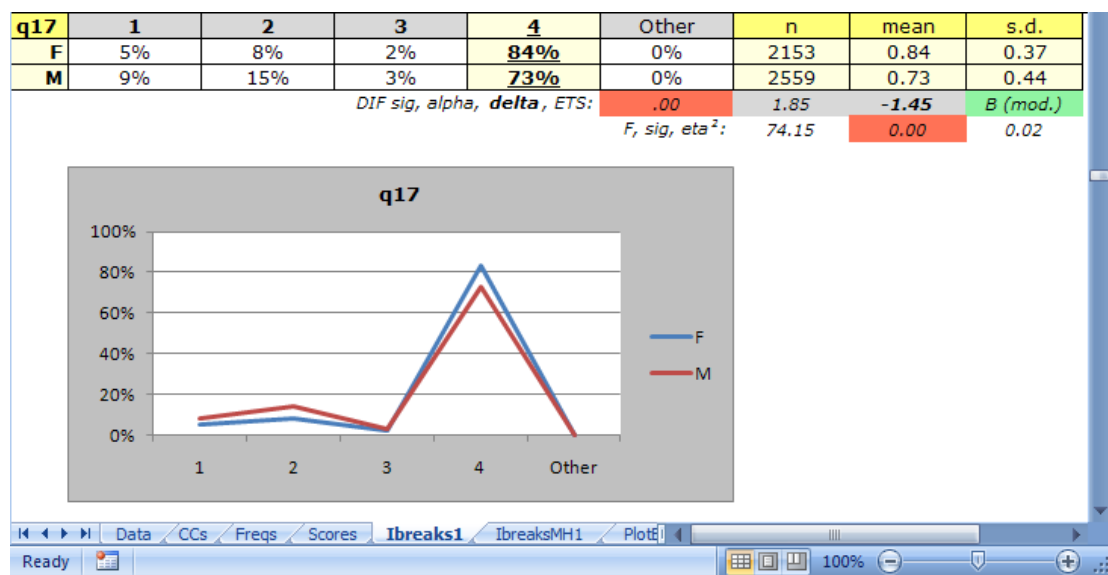
In this case I was interested in a DIF analysis, and I answered Lertap's questions accordingly.



I get two new reports, "Ibreaks1", and "IbreaksMH1".

The first of these produces the small table and graph shown above. On the item labelled q01, there were no differences at all between the two groups. The two trace lines, one for F, one for M, overlap each other, making it seem there's just a single trace.

The item labelled q17 displayed a different picture:



Now I've got a difference between group means, 0.84 versus 0.73, and although it strikes me as slight, there may indeed be something to it, something in this item which may, for some reason, favour the girls.



The two lines of statistics which appear below the table, with colours, are used to answer questions commonly found in data analysis of this sort.

The *F*, *sig*, *eta*<sup>2</sup> results are entirely analogous to those seen earlier in the analysis of variance table of the "Breaks1" report. Are the differences between the sample means statistically significant? Yes, the *F* ratio of 74.15 is significant at the 0.00 level. But once again I tread lightly here; given the large sample sizes, even very small differences are likely to be statistically significant – what may be more meaningful is the *eta*<sup>2</sup> figure. At 0.02 it is not large at all, but then, in the scheme of things, maybe this has some meaning.

If I want my test items to show no favour, to have equal challenge for each gender, how much difference will I tolerate?

To be honest, I don't know. This is a question for the test developers, not me.

What I do know is that Lertap has more information for me to study. It turns out that the **B (mod.)** outcome seen above is a flag, waving away, urging me to have a look at the second report, "IbreaksMH1". I turn to it.

Lertap5 Mantel-Haenszel results based on score levels from G431-49, grouped by gender.											
Score levels->	4	5	6	7	8	9	10	11	12	13	
<b>F (r)</b>	1	7	5	4	6	6	8	7	18	20	
<b>M (f)</b>	1	2	4	7	13	13	17	21	27	40	
<b>q17</b>											
<b>F diff</b>	.00	.43	.20	.25	.50	.50	.50	.14	.50	.30	
<b>M diff</b>	.00	.00	.00	.43	.31	.15	.41	.24	.19	.40	
<b>odds ratio-&gt;</b>	.00	.00	.00	.44	2.25	5.50	1.43	.53	4.40	.64	
MH chi-sq: 59.25 Prob: <b>.00</b> MH alpha: 1.85 MH D-DIF: -1.45 ETS level: <b>B (mod.)</b>											
<b>q18</b>											
<b>F diff</b>	.00	.00	.20	.00	.33	.00	.13	.14	.17	.05	
<b>M diff</b>	.00	.00	.25	.00	.08	.15	.18	.29	.07	.10	
<b>odds ratio-&gt;</b>	.00	.00	.75	.00	6.00	.00	.67	.42	2.50	.47	
MH chi-sq: 3.66 Prob: .06 MH alpha: .88 MH D-DIF: .30 ETS level: A (neg.)											

The "MH" part of this report's label stands for Mantel-Haenszel, named after authors who years ago developed a statistical method now frequently used for representing the extent to which the responses of two groups might differ. The method is based on the "odds ratio", an index of how one group, say the "reference" group, might be favoured to get an item correct when compared to the second group, often called the "focal" group. An odds ratio greater than 1.00 indicates that the odds favour the reference group (they're more likely to get the item right); a ratio less than 1.00 favours the focal group.

Lertap's IbreaksMH tables of item results are broken into score levels, which are seen at the top of the report. These levels generally begin with the lowest score found, 4 in this case, and range out to the highest score found, which in this case was 49. In the tables showing above, I can see that one person in the F group had a score of 4, as did one person in the M group.

The *diff* figures indicate the proportion of people in each group, at each score level, who correctly answered an item. For q17, 43% of the 7 female students with a score of 5 correctly answered the item; of the 2 males with a score of 5, none answered q17 correctly.


These tables are usually quite wide, with one column required to capture all the action at each score level.

The statistics which appear for each item, immediately below the row of odds ratios, are those often used in an M-H analysis. You'll find more about them


here:

[larrynelsonstuff.com/HTMLHelp/HTML/index.html?dif.htm](http://larrynelsonstuff.com/HTMLHelp/HTML/index.html?dif.htm)

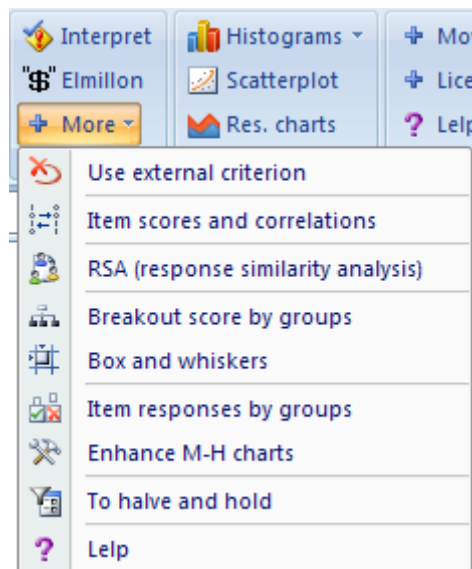
Lertap has an option to make line graphs. Its icon is found in the "Basic options" section of the Excel 2007 Lertap ribbon tab:




Have a look at a line graph of the two **diff** rows for item q17, made, in this case, by selecting cells starting at score level 7 (**diffs** of .25 and .43), going out to a score level of 41, and then using Lertap's "Line" option:



At this point, I can make good use of the option to "Enhance M-H charts" shown below:



After clicking on the Enhance option, my chart improves quite a bit:




Now it's fairly easy to see that there's quite a range of score levels where the girls (F) outscored the boys (M). The difference isn't uniform in the score levels range I've used here – at the lower end there are a few levels where the advantage went to the boys. However, the number of cases in the lower score levels is generally not great; the M-H statistic is a weighted average of an item's odds ratios, so those few levels where the boys were stronger, coming in a relatively sparse region of the score distribution, will be washed out.

### Non-uniform DIF

Lertap uses Mantel-Haenszel methods for its DIF analysis. While these methods generally receive good support in the literature, it is recognised that M-H procedures are not capable of uncovering non-uniform DIF, that is, the case where an item may favour one group in one area of the scores range, but, in another part of the scores range, favour the other group.

As an example, consider item "I15" from the test mentioned in Lertap help:

[larrynelsonstuff.com/HTMLHelp/HTML/index.html?dif.htm](http://larrynelsonstuff.com/HTMLHelp/HTML/index.html?dif.htm)



MH alpha for I15 was 2.39, with MH D-DIF at -2.04, an ETS level of C (large), indicating DIF in favour of the reference group, Nat.

For most of the scores range (seen along the x-axis), there seems to be something about I15 which works to favour the Nats. However, as shown in the small table below, there is an area of the range, corresponding to test scores of 25, 26, and 27, where the odds ratio is in favour of the focal group, Ing. And, there are quite a few students at these three score levels: 82 Nats and 91 Ings (respective total group sizes were 876 and 844, so we're talking about approximately 10% of each group; the odds ratio at a score level of 7 was .25, but there were few students with this score, only 3 in each group).

Lertap5 Mantel-H									
Score levels->	22	23	24	25	26	27	28	29	
Nat. (r)	39	27	45	37	29	34	19	16	
Ing. (f)	51	23	27	33	27	36	28	28	
<b>I15</b>									
Nat. diff	.85	.93	.87	.89	.83	.91	.79	1.00	
Ing. diff	.63	.57	.85	.88	.85	.94	.93	.89	
odds ratio->	3.27	9.62	1.13	1.14	.83	.61	.29	.00	

I myself would not have singled out I15 as displaying non-uniform DIF, but a logistic regression program did, finding "moderate" DIF for I15, saying that it "exhibited non-uniform DIF too".

Perhaps, given a Lertap IbreaksMH table such as that seen here, a *simple* index of non-uniform DIF could be made using the odds ratio row, counting the number of times it flips from below 1.00 to above 1.00. An item with uniform DIF might never flip at all. An item with no DIF could be expected to have lots of flips, and an item with non-uniform DIF might flip just once. (Any reader thinking this idea to possibly have a smidgen of merit might contact me; it shouldn't be too difficult to compute. We could call it the "OR Flip Flop" index. To test it we'd need a few datasets, some known to have non-uniform DIF.)

## SPSS and DIF

It's possible to get **SPSS** ([www.spss.com](http://www.spss.com)) to calculate Mantel-Haenszel statistics. Lertap can be used to prepare data for SPSS (which of course is just as well as all my data are in Lertap to start with).

Here's how to go from Lertap to SPSS.

Starting with Lertap:

Make sure the categorical variable in the Lertap Data sheet has a numeric code. If it was in a column which used {M, F} for codes, for example, use Lertap's recoder to map these to {1, 2}. The "Recode a Data column" is an option under the Move+ option under the 'Others menu'.

Copy the categorical variable over to the Lertap Scores sheet. Use the "Copy a Data column to the Scores worksheet" option, also found under the Move+ menu.

Create a Lertap IStats worksheet by using the "Item scores and correlations" option found under the More+ option on the 'Run menu'. Then, from IStats, select the columns and rows with the 1.00 / 0.00 item scores (they always start in row 3), and copy them to columns on the Scores worksheet, to the right of the column with the categorical variable. Yes, this can be a big copy and paste when there are thousands of test takers.

Make sure you have the Excel referencing style set to "A1": this is simple to do – just use the "Ref. style" option found under "Excel" in the 'Basic options' menu, to the right of the "Spread" option. Click it, and the columns in the Scores worksheet should change from having numbers as labels to letters, starting, naturally, with the letter A.

Note the cells where the scores begin; most likely this will be A3. Note where they end; on a sample job I ran, with 40 items and 1,720 test takers, the scores ended in cell AQ1722. And note: this cell will **not** be at the bottom of the Scores worksheet, just in the row where the scores end. And another note: getting stuff from Scores into SPSS is not difficult, but it can be made just a bit easier if you delete the first Scores row so that "variable labels" come up to reside there, in the first row. If you don't want to fiddle with the Scores worksheet in this manner, use Excel to make a copy of it, and then delete the first row in the copy. Then, once again note the cell where the scores begin, and that where they end; A2 to AQ1721, for example.

Now, close the workbook. No need to exit Excel, but you do need to close the workbook to avoid a potential file sharing violation when getting into SPSS.

Going over to SPSS:

Use the File menu to / Open / Data

Tell SPSS you're looking for Files of type Excel (\*.xls, \*.xlsx, \*.xlsm)

Find and Open the Lertap Excel workbook you just closed.

If you did delete the first Scores row, then leave a tick in the little box which says "Read variable names from the first row of data".

Pick out the worksheet in the workbook which has your scores. Put in the range, including the first row if you have the variable names in it (in my example, this would be A1:AQ1721).

That should do it. SPSS should now have a dataset with a tab for "Data View" and another for "Variable View".

Save this SPSS dataset.

Still in SPSS, go to Analyze / Descriptive Statistics / Crosstabs

Put the categorical variable into the Row(s) box, all the items into the column(s) box, and the actual test score into the Layer 1 of 1 box. Under the Statistics option, tick "Cochran's and Mantel-Haenszel statistics".

Click Continue / OK

SPSS grinds out lots of results. Here are its calculations for I15:

#### Tests of Conditional Independence

	Chi-Squared	df	Asymp. Sig. (2-sided)
Cochran's	56.769	1	.000
Mantel-Haenszel	55.051	1	.000

#### Mantel-Haenszel Common Odds Ratio Estimate

Estimate				.419
ln(Estimate)				-.870
Std. Error of ln(Estimate)				.118
Asymp. Sig. (2-sided)				.000
Asymp. 95% Confidence Interval	Common Odds Ratio	Lower Bound	.333	
		Upper Bound	.528	
	ln(Common Odds Ratio)	Lower Bound	-1.101	
		Upper Bound	-.639	

The Mantel-Haenszel common odds ratio estimate is asymptotically normally distributed under the common odds ratio of 1.000 assumption. So is the natural log of the estimate.

Compare these results with I15 data given in an IbreaksMH report from Lertap:

I15										
Nat. diff	.00	.00	.00	.33	.46	.18	.50	.61	.59	.63
Ing. diff	.00	.00	.25	.67	.07	.19	.17	.21	.22	.41
odds ratio->	.00	.00	.00	.25	12.00	.96	4.75	6.00	5.24	2.35
		MH chi-sq: 55.05		Prob: .00	MH alpha: 2.39		MH D-DIF: -2.04		ETS level: C (large)	

The Mantel-Haenszel Chi-square values agree exactly (55.05), but the SPSS estimate of the Mantel-Haenszel common odds ratio is .419, compared to an MH alpha of 2.39 from IbreaksMH. What appears to be a major discrepancy has come about because I took little care when using Lertap to recode the categorical variable from Nat and Ing to numeric codes: I let Ing become 1, while Nat became 2. SPSS effectively takes the lowest code as the reference group, so now

I need to divide SPSS' .419 into 1 (that is, 1.00 divided by .419) to reverse this coding error. What do I get? Agreement: 2.39.

## **Other M-H DIFfers: SAS and jMetrik**

I know that **SAS** ([www.sas.com](http://www.sas.com)) is eminently capable of reporting Mantel-Haenszel results, but that's all I can say; I don't know how to use it.

SIBTEST (from [www.assess.com](http://www.assess.com) and other sites) used to support those interested in applying Mantel-Haenszel methods, but this is no longer the case.

**jMetrik** will do Mantel-Haenszel stuff, and it's free. Look for it at [www.itemanalysis.com](http://www.itemanalysis.com). Authored by Professor Patrick Meyer, jMetrik has apparently taken over from the same site's ItemQual program.

I used ItemQual as a counter-check when I first started to develop the Mantel-Haenszel support now seen in Lertap<sup>1</sup>.

## **Back to Lertap?**

Of all the systems mentioned in this little paper, I suspect Lertap is the only one which makes it fairly simple to get graphs of results, a fact which stems from its Excel base, and the fairly extensive charting facilities offered by Excel. If jMetrik follows in the footsteps of ItemQual, it may very well eventually come to have the ability to display DIF graphs, empirical and theoretical.

Lertap may be downloaded from:

[http://larrynelsonstuff.com/HTMLHelp/HTML/index.html?how\\_to\\_get\\_lertap.htm](http://larrynelsonstuff.com/HTMLHelp/HTML/index.html?how_to_get_lertap.htm)

---

Larry Nelson, PhD

School of Education  
Curtin University of Technology  
Western Australia

College of Research Methodology and Cognitive Science  
Burapha University  
Chonburi, Thailand

Questions welcomed. Please send to: [larry@lertap.com](mailto:larry@lertap.com)

---

<sup>1</sup> Thanks to Professor Meyer for his assistance.