
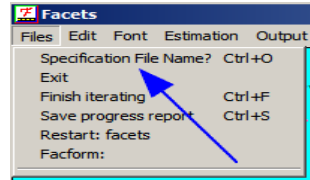
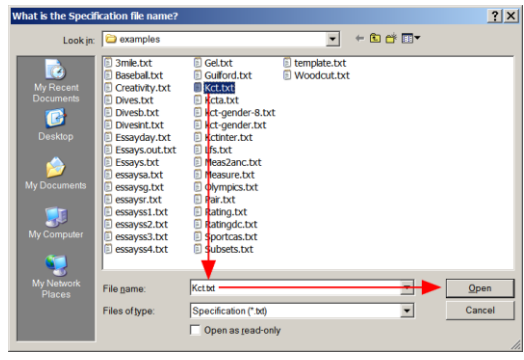
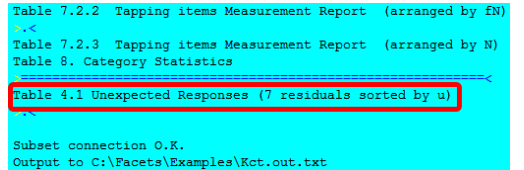
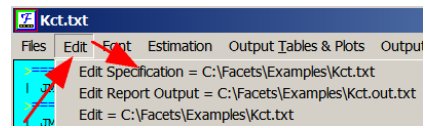
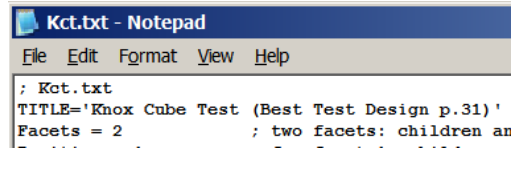
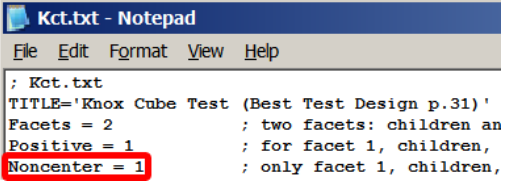
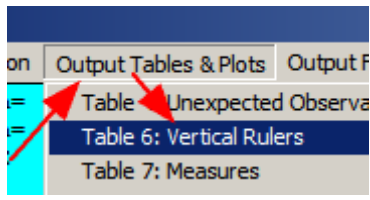
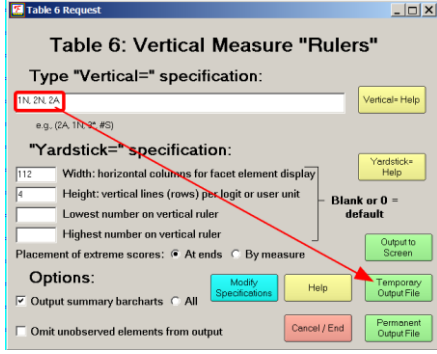
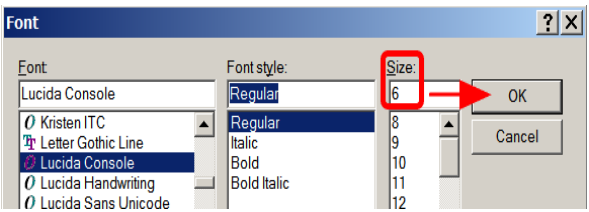
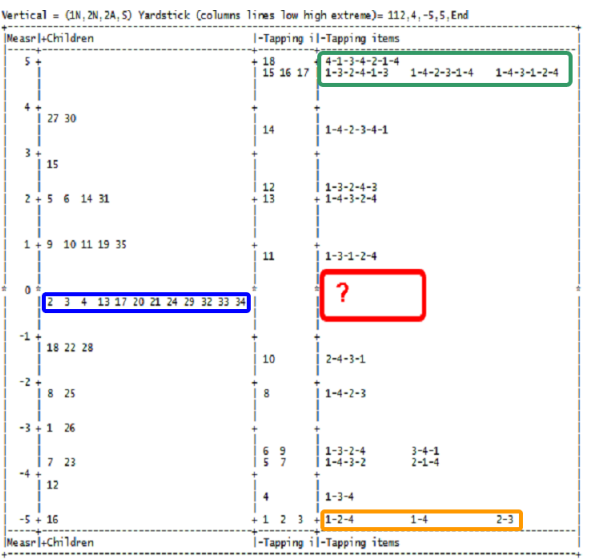




#	<b>Many-Facet Rasch Measurement : Facets Tutorial</b> <b>Mike Linacre - 1/2012</b>	
1.	<b>Tutorial 2. Fit analysis and Measurement Models</b> <i>Welcome back!</i> <ul style="list-style-type: none"> <li>• Observations, expectations and residuals</li> <li>• Quality-control fit statistics elements and observations</li> <li>• Reliability indexes and inter-rater reliability</li> </ul> This tutorial builds on Tutorial 1, so please go back and review when you need to.	
2.	<b>A. Facets Specifications and Data: The Knox Cube Test</b>	
3.	Let's launch <i>Facets</i> again	
4.	To start with we'll look at an analysis that's about as simple as it gets: 2-facets, dichotomous. Click on "Files" Click on "Specification File Name?"	
5.	Click on "Kct.txt" and "Open" or Double-Click on "Kct.txt" "Extra Specifications" - click on "OK" "What is the Report Output file name" - click on "Open"  This is the "Knox Cube Test" data in " <a href="#">Best Test Design</a> " (Wright & Stone, 1979, MESA Press). The Knox Cube Test was devised by Dr. Howard Knox on Ellis Island in New York harbor (next to the Statue of Liberty). It was used to screen immigrants arriving by ship from Europe. It assesses attention-span and short-term memory.	
6.	The Estimation and initial reporting completes. We will be particularly interested in Table 4.1 "Unexpected Responses", but first let's look at what this analysis is all about ....	
7.	Click on "Edit" menu Click on "Edit Specification = ... Kct.txt"	
8.	We know what most of this means: ; starts a comment. I wanted to mention " <i>Kct.txt</i> ", the name of the specification file. <i>TITLE</i> = is title line at the top of each output table Facets = 2 - there are two facets: children and items	

9.	<p><i>Positive = 1</i> - the first facet (children) have positive ability: more score = more measure. The second facet, items has the default setting, negative difficulty, more score = less measure. You have probably realized that the order of the specifications doesn't matter, except that we need to specify <i>Facets=</i> early in the specification file.</p>	 <pre> ; Kct.txt TITLE='Knox Cube Test (Best Test Design p.31)' Facets = 2 ; two facets: children an Positive = 1 ; for facet 1, children, Noncenter = 1 ; only facet 1, children, </pre>
10.	<p><i>Noncenter=1</i></p> <p>Something that must be decided in all measurement is <b>where to measure from</b>. For short distances, we measure length from the end of the tape measure. For mountains, from sea level. For temperature, from freezing point of water for Celsius, but from freezing point of salt water for Fahrenheit. It is the same in Rasch measurement. The measuring convention is that <b>we measure from the center (mean) of the measures for each facet</b>. So item difficulties are measured from the center, <i>the local origin</i>, of the item facet. The average item has a difficulty of 0 logits. Judge severities are measured from the center, <i>the local origin</i>, of the judge facet. The average judge has a severity of 0 logits. We do this for all facets <b>except one</b>, usually the person ability facet. The person abilities are measured from the local origins of all the other facets. If the average ability is high, then the average person has a positive logit measure. If the average ability is low, then the average person has a negative logit measure.</p> <p>So all facets have their local origin at their centers, except one facet.</p> <p><i>Noncenter=1</i> ; the first facet (children) does not have its local origin at its center.</p>	
11.	<p><i>Pt-biserial = Yes</i> - report the point-biserial correlation in the measure tables, Table 7. These may not make much sense if the data are incomplete (so there are missing observations). This dataset is a complete rectangular dataset.</p> <p><i>Vertical</i> = this controls the facets to display in Table 6, the vertical rulers.</p> <p><i>Yard</i> = this controls the size of the display in Table 6.</p> <p>Recommendation: Use the <i>Output Tables</i> pull-down menu to play with different settings of <i>Vertical=</i> and <i>Yard=</i> for Table 6, until you find settings that you like.</p>	<pre> Pt-biserial = Yes ; report Vertical=1*,1A,2N,2A ; show of Yard=112,4 ; Vertical Model = ?,?,D ; element Labels = 1,Children ; Childre 1-17=Boy,,1 ; Pretenc 18-35=Girl,,2 ; Pretenc * ; end of 2,Tapping items ; Items a 1=1-4 ; Items 1 </pre>
12.	<p><i>Model = ?, ?, D</i></p> <p>There is only one model specification, so it can be on the same line as <i>Model=</i>. “?” means “any element of the facet”. The first “?” is for facet 1. The next “?” is for facet 2. So this model specification says: “Any element of facet one can combine with any element of facet two to produce an observation on a D-type scale”. “D” means “dichotomous 0-1 scale”. So <i>Facets</i> expects to see 0's and 1's in the data file. Anything else is treated as a missing value and ignored.</p>	
13.	<p><i>Labels =</i> defines the facet names and the elements in the facets.</p> <p><i>1, Children</i> - the label or name of the first facet is “children”</p> <p><i>1-17=Boy,,1</i> - after the facet name comes the list of elements. In this facet, element numbers 1 to 17 are all labeled “Boy”. They could be given individual labels if desired. “,,1” means “the boys are part of element group 1”. So a measure Table with totals will be produced for the boy group.</p> <p><i>18-35=Girl,,2</i> - element numbers 18 to 35 are all labeled “Girl”. They could be given individual labels if desired. “,,2” means “the girls are part of element group 2”. So a measure Table with totals will also be produced for the girl group.</p> <p>* - element lists end with “*”</p>	



<p><b>20. B. Table 6: The Knox Cube Test Measures</b></p>	
<p><b>21.</b> Let's produce a version of Table 6, the vertical rulers which tells us what we will need to know. In the main <i>Facets</i> window, click on "Output Tables &amp; Plots" click on "Table 6: Vertical Rulers"</p>	
<p><b>22.</b> Kct.txt has: "Vertical=1*,1A,2N,2A" But we want to display the first facet, the children, by element number: 1N the second facet, the items by element number, 2N and by element label (the tapping pattern) 2A so the specification is: <b>Vertical = 1N, 2N, 2A</b> We can specify this by typing 1N, 2N, 2A in the <i>Vertical=</i> box. Then click on "Temporary Output File"</p>	
<p><b>23.</b> Table 6 displays. At first it was too big for my screen and somewhat faint, so (just like in Tutorial 1) I went to the <b>NotePad</b> menu bar, and used the <b>Format</b> pull-down menu to change the <b>Font</b> and the <b>Size</b>. Mine is "Courier New" 8-point. For sizes smaller than 6 points, type the size into the NotePad font size box.</p>	<p>In NotePad:</p> 
<p><b>24.</b> In Table 6, at the extreme left is the measurement scale, "Measr", in logits. It is -5 to +5, a typical range. This is not pre-set. It is estimated from the pattern of the data.</p> <p><b>blue box:</b> The children are shown with many of them just below 0 logits. This is the mode of their distribution. The column heading is "+Children". "+" means "more score ↔ more measure". So the most able (highest scoring) children are at the top. They are 27 and 30.</p> <p>In the third column, the items are shown by element number. The fourth column shows them by label. The column is headed "-Tapping items", so "more score ↔ less measure."</p> <p>The lowest scoring items (least success by the children) are at the top. These are the most difficult items. The easiest items are at the bottom.</p>	

25.	<p>Do you notice any flaws in this version of Knox’s test? Here is one ...</p> <p><b>Red box</b> in #24: There are 12 children in the middle of the range (children 2, 3, 4, 13, ...), but no items at their level. If the test is intended for children like those in this sample, it needs more middle-difficulty items. Can you imagine some extra items that might go in the red box in #24? They will be more difficult than the items below the red box, but easier than the items above the red box.</p> <p><b>green box</b>: Now look at the top of the item columns. There are 4 items that are much too difficult for this sample.</p> <p><b>orange box</b>: And at the bottom there are three items that are somewhat too easy. These extreme items waste everyone’s time, and they may make the children frustrated or over-confident.</p> <p>In their book, “<i>Best Test Design</i>”, Wright &amp; Stone improve this test.</p>																																																																																	
26. <b>C. Table 6: Measures and Expectations</b>																																																																																		
27.	<p>Now look at person 5 in #24. His measure is +2 logits. What do we expect to happen when he encounters item 13, also at +2 logits ? The child has the same ability as the item has difficulty. We don’t know what will happen. The child’s probability of success is 0.5</p>																																																																																	
28.	<p>What about when child 9 of ability +1 logits attempts item 13 of difficulty +2 logits. Child 9 is less able than the item is difficult so the child will probably fail. But what is the child’s exact probability of success? .4, .3, ..?</p>																																																																																	
29.	<p>We can compute the probability of success from the Rasch model for dichotomous observations (Tutorial 1). Let’s fill in the values: <math>B_n = +1</math>, <math>D_i = +2</math></p>	$\log_e(P_{ni}/(1-P_{ni})) = B_n - D_i$ $\log_e(P_{ni}/(1-P_{ni})) = +1 - +2 = -1$																																																																																
30.	<p>Rearrange the algebra. (If you are not sure about “e”, please review Tutorial 1, Appendix 3).</p>	$P_{ni} = e^{-1} / ( 1 + e^{-1} )$ $= 1/2.718 / ( 1 + 1/2.718 )$ $= 0.37 / 1.37 = .27$																																																																																
31.	<p>The probability of success when child 9 of ability +1 logits attempts item 13 of difficulty +2 logits is <math>p = .27</math>.</p>	$= 1 \text{ success in every 4 attempts}$																																																																																
32.	<p><b>Logit-to-Probability Conversion Table</b></p> <p>Here is a Table to guide you when you convert dichotomous logit differences into percents (or probabilities) of success.</p> <p><b>green text</b>: Our difference is -1 logits. Look half-way down the right-hand pair of columns. -1.0 logits is 27% chance of success, which is the same as <math>p=.27</math>.</p> <p>Notice these useful values:</p> <p>1.1 logits difference = 75% chance of success  2.2 logits difference = 90% chance of success  3.0 logits difference = 95% chance of success</p>	<table border="0"> <thead> <tr> <th>Logit</th> <th>diff.</th> <th>% Success</th> <th></th> <th></th> </tr> </thead> <tbody> <tr> <td>5.0</td> <td></td> <td>99%</td> <td>-5.0</td> <td>1%</td> </tr> <tr> <td>4.6</td> <td></td> <td>99%</td> <td>-4.6</td> <td>1%</td> </tr> <tr> <td>4.0</td> <td></td> <td>98%</td> <td>-4.0</td> <td>2%</td> </tr> <tr> <td>3.0</td> <td></td> <td>95%</td> <td>-3.0</td> <td>5%</td> </tr> <tr> <td>2.2</td> <td></td> <td>90%</td> <td>-2.2</td> <td>10%</td> </tr> <tr> <td>2.0</td> <td></td> <td>88%</td> <td>-2.0</td> <td>12%</td> </tr> <tr> <td>1.4</td> <td></td> <td>80%</td> <td>-1.4</td> <td>20%</td> </tr> <tr> <td>1.1</td> <td></td> <td>75%</td> <td>-1.1</td> <td>25%</td> </tr> <tr> <td>1.0</td> <td></td> <td>73%</td> <td>-1.0</td> <td>27%</td> </tr> <tr> <td>0.8</td> <td></td> <td>70%</td> <td>-0.8</td> <td>30%</td> </tr> <tr> <td>0.5</td> <td></td> <td>62%</td> <td>-0.5</td> <td>38%</td> </tr> <tr> <td>0.4</td> <td></td> <td>60%</td> <td>-0.4</td> <td>40%</td> </tr> <tr> <td>0.2</td> <td></td> <td>55%</td> <td>-0.2</td> <td>45%</td> </tr> <tr> <td>0.1</td> <td></td> <td>52%</td> <td>-0.1</td> <td>48%</td> </tr> <tr> <td>0</td> <td></td> <td>50%</td> <td>-0.0</td> <td>50%</td> </tr> </tbody> </table>	Logit	diff.	% Success			5.0		99%	-5.0	1%	4.6		99%	-4.6	1%	4.0		98%	-4.0	2%	3.0		95%	-3.0	5%	2.2		90%	-2.2	10%	2.0		88%	-2.0	12%	1.4		80%	-1.4	20%	1.1		75%	-1.1	25%	1.0		73%	-1.0	27%	0.8		70%	-0.8	30%	0.5		62%	-0.5	38%	0.4		60%	-0.4	40%	0.2		55%	-0.2	45%	0.1		52%	-0.1	48%	0		50%	-0.0	50%
Logit	diff.	% Success																																																																																
5.0		99%	-5.0	1%																																																																														
4.6		99%	-4.6	1%																																																																														
4.0		98%	-4.0	2%																																																																														
3.0		95%	-3.0	5%																																																																														
2.2		90%	-2.2	10%																																																																														
2.0		88%	-2.0	12%																																																																														
1.4		80%	-1.4	20%																																																																														
1.1		75%	-1.1	25%																																																																														
1.0		73%	-1.0	27%																																																																														
0.8		70%	-0.8	30%																																																																														
0.5		62%	-0.5	38%																																																																														
0.4		60%	-0.4	40%																																																																														
0.2		55%	-0.2	45%																																																																														
0.1		52%	-0.1	48%																																																																														
0		50%	-0.0	50%																																																																														

33.	We have the logit measure for every child and every item. They are displayed in Table 6 (pictorially) and Table 7 (numerically). So we can use the Rasch dichotomous model to compute probability of success for every child on every item. These probabilities are the “ <b>expected</b> ” observations.	For dichotomous, 0 or 1, data, probability of success → the expected value of the observation
34.	Think of this in terms of frequency. What would we expect if 100 people of the ability of child 9 attempted item 13?	100 attempts at item 13 by ability of child 9. Logit difference = -1, Percent success = 27% <i>Expect: 27 successes out of 100 attempts</i> Expected value of 1 attempt = $27/100 = .27 =$ Rasch-model probability of success
35.		

36.	<b>D. Rasch Theory: Observations, Expectations and Residuals: Response-level fit of the data to the Rasch model</b>	
37.	<p>Here is the Knox Cube Test data again:</p> <p>The principles of fit are easier to explain with dichotomous data than with polytomous data, so that is why we are starting here.</p>	<pre>Data = ; no data file nam 1 ,1 ,1 ; child 1 on item 1 ,2-18,1,1,1,1,1,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0 2 ,1-18,1,1,1,1,1,1,0,0,1,1,1,0,0,1,0,0,0,0,0,0,0 3 ,1-18,1,1,1,1,1,1,1,1,1,0,1,0,0,0,0,0,0,0,0,0,0 4 ,1-18,1,1,1,1,1,1,1,1,1,1,0,0,0,0,0,0,0,0,0,0,0 5 ,1-18,1,1,1,1,1,1,1,1,1,1,1,0,0,1,1,0,0,0,0,0,0 6 ,1-18,1,1,1,1,1,1,1,1,1,1,1,1,0,0,0,0,0,0,0,0,0 7 ,1-18,1,1,1,1,0,0,1,0,1,0,0,0,0,0,0,0,0,0,0,0,0 8 ,1-18,1,1,1,1,1,0,1,0,1,1,0,0,0,0,0,0,0,0,0,0,0 9 ,1-18,1,1,1,1,1,1,1,1,1,1,1,0,0,1,0,0,0,0,0,0,0</pre>
38.	Child 9, item 13, is marked in green. The child scored “1”, a success!	<p>the observation: <math>X_{ni} = X_{9,13} = 1</math></p>
39.	We’ve already discovered in #28 that Child n=9 (ability $b = 1$ logit) is less able than item i=13 (difficulty $d = 2$ logits): $b - d = -1$ , $P_{ni} = 0.27$	<p>the expectation = <math>E_{ni} = P_{ni} = e^{b-d} / (1 + e^{b-d})</math> <math>= e^{-1} / (1 + e^{-1}) = 0.27</math></p>
40.	The difference between the <b>observation</b> and its <b>expectation</b> is the “ <b>residual</b> ” (what is left over). This is the part of the observation we did not expect to see ...	<p>the residual: <math>R_{ni} = X_{ni} - E_{ni} = 1 - 0.27 = 0.73</math></p>
41.	<p>We know that what we will see in the KCT data are 0’s and 1’s, but they are <i>not</i> the expected values. The expected values are numbers like .68 and .34. So there are almost always residuals. There is a further question to ask, “Are we <i>surprised</i> about the size of the residual, or is it about the size of the discrepancy we were expecting to see?”</p> <p>There are two aspects to what we expect:</p> <ol style="list-style-type: none"> <li>1. The expected (average) value.</li> <li>2. The expected variation of the observed around its expected value. This is called the “model variance”.</li> </ol> <p>Think of 100 people like child 9 attempting item 13. We expect 27 successes. The expected (average) value is <math>27/100 = .27</math>. But we also expect to see 27 1’s and 73 0’s. So there will be residuals!</p> <p><i>Here is a technical computation:</i></p> <p>the <b>sum-of-squared-residuals</b> = sum of (observation - expectation)<sup>2</sup>  = (count of successes)*(success - expected value)<sup>2</sup> + (count of failure)*(failure - expected value)<sup>2</sup>  = (success count)*(1 - expected value)<sup>2</sup> + (failure count)*(0 - expected value)<sup>2</sup>  = <math>27*(1 - .27)^2 + 73*(0 - .27)^2 = 27 * .73^2 + 73 * .27^2 = 100 * .27 * .73 = 19.71</math></p> <p>the <b>model residual variance</b> = <math>V_{ni} = \text{sum-of-squares} / \text{count of residuals}</math>  <math>V_{ni} = 19.71 / 100 = 0.1971</math></p> <p>the <b>model residual standard-deviation</b> = square-root (variance) = <math>\sqrt{V_{ni}} = \sqrt{0.1971} = 0.44</math>  = the size of the splatter of the observations around their expected values.</p>	
42.	With these numbers, we can calculate how unexpected is our residual, $R_{ni}$ . The <i>standardized</i> residual, $Z_{ni}$ , is as unexpected as the unit <i>normal deviate</i> [see Appendix 1. Unit Normal Deviates of this tutorial].	<p>standardized residual = <math>Z_{ni} = R_{ni} / \sqrt{V_{ni}}</math></p>



43.	In our example, $X_{13,9} = 1$ , the residual $R_{ni} = 0.73$ , the residual S.D. = 0.44, so that the standardized residual, $Z_{ni}$ , is 1.66. This is as unusual as a unit <i>normal deviate</i> of 1.66, $p \approx .10$ (see Table in Appendix 1), but not unusual enough ( $p < .05$ ) to be considered significantly misfitting the Rasch model.	$R_{ni} = 0.73, V_{ni} = 0.1971, \sqrt{(V_{ni})} = 0.44,$ $Z_{ni} = 0.73/0.44 = 1.66; p \approx .10$
44.	Now let's look at the observation I ringed in red in #37: Child 2 on item 14. According to Table 6 (see #24), Child 2 has an ability of about -0.25 logits. Item 14 has a difficulty of about 3.37 logits.	Logit difference (child - item) = $-0.25 - 3.37 = -3.6$ logits Probability of success (Table in #32) = 3%
45.	We observed a success, so $X_{ni} = 1$ . Expectation = 3% success = .03 (we are rounding the computations 2 decimal places for clarity) Now we can compute the residual and the standardized residual. The residual, $R_{ni}$ , is .97 (very large) and the standardized residual, $Z_{ni}$ , is 5.60 (very unexpected), $p < .01$ .	Observation: $X_{ni} = 1$ Expectation: $E_{ni} = P_{ni} = .03$ Residual: $R_{ni} = X_{ni} - E_{ni} = 1 - .03 = 0.97$ Model variance of the observation around its expectation: $V_{ni} = P_{ni} * (1 - P_{ni}) = .03 * .97 = .03$ Standardized residual: $Z_{ni} = R_{ni} / \sqrt{(V_{ni})} = 0.97 / \sqrt{(.03)} = 5.60$
46.	<b>E. Table 4: Unexpected Responses</b>	
47.	We could go through this computation by hand for every observation, but it is easier to have <i>Facets</i> do it for us. Click on the <i>Facets</i> Report output file on your Windows Taskbar (or click on <i>Facets</i> "Edit" menu, click on "Report output file") Scroll down to Table 4. It is the last Table. It shows the unexpected responses (or unexpected observations)	Knox Cube Test (Best Test Design p.31) 4/23/2009 3:12:00 AM Table 4.1 Unexpected Responses (7 residuals sorted by u). <pre> +-----+   Cat  Score  Exp.  Resd  StRes  Nu  Chil  Nu  Tapping items   +-----+-----+   1      1      .0   1.0   6.2   2  Boy  14  1-4-2-3-4-1     0      0      1.0  -1.0  -6.1   2  Boy   7  1-4-3-2       0      0      1.0  -1.0  -6.1   24 Girl 7  1-4-3-2       0      0      1.0  -1.0  -4.8   24 Girl 6  3-4-1         1      1      .1    .9   3.5   24 Girl 12 1-3-2-4-3     1      1      .1    .9   3.5   33 Girl 12 1-3-2-4-3     0      0      .9   -.9  -3.5   28 Girl 5  2-1-4       +-----+-----+   Cat  Score  Exp.  Resd  StRes  Nu  Chil  Nu  Tapping items   +-----+-----+ </pre>
48.	<b>Green box:</b> <i>Facets</i> has done the computation for Child 2 on Item 14 more precisely than I did. It reports that the standardized residual (StRes, $Z_{ni}$ ) is 6.2. This is the most unexpected observation in these data. The observation is unexpectedly high (1) compared with its expected value (.0)	
49.	<b>Red box:</b> Look at the next two observations listed in Table 4. Both are on item 7 and they are equally unexpected StRes = -6.1. The minus - sign means "they did worse than we expected." Both children, 2 and 24, failed on the item when we expected them so succeed. We don't know why, but if we were serious about the children or the instrument we might inquire. The tapping pattern includes the sequence 4-3-2. Perhaps the examiner sped-up unintentionally, or perhaps he didn't clearly tap each cube so the children saw 4-2 instead of 4-3-2.  The list of unexpected responses nearly always contains useful messages about the instrument, the sample, the judges, the dataset, or whatever .....	



**50. F. Table 7: Quality-control fit statistics elements and observations**

**51.** Looking down the list of unexpected responses is somewhat like looking at the pot-holes in a road. You want to pay some attention to them (not too much, usually), but they don't tell you much about the surface of the road as a whole. For that we need to take a wider look.


**52.** Scroll back up the Kct.out.txt Report Output file to Table 7.2.2. - the measure Table for Items in *fit order, descending*, or output a new copy of Table 7 from the "Output Tables" menu.  
**red box:** You will see 4 columns: Infit and Outfit, MnSq and Zstd. These are quality-control fit statistics. They are central to the evaluation of the quality of the data for the construction of measures.

**53.** Table 7.2.2 Tapping items Measurement Report (arranged by fN).

Total Score	Total Count	Obsvd Average	Fair-M Average	Measure	Model S.E.	Infit MnSq	Infit ZStd	Outfit MnSq	Outfit ZStd	Estim. Discrm	Corr. PtBis	Nu Tapping items
31	35	.9	.98	-3.87	.71	1.35	.8	2.30	1.2	.52	.30	7 1-4-3-2
3	35	.1	.03	3.40	.70	1.59	1.2	1.53	1.0	.60	.09	14 1-4-2-3-4-1
30	35	.9	.96	-3.41	.65	1.18	.5	.98	.6	.86	.43	6 3-4-1
6	35	.2	.08	2.25	.55	1.17	.6	1.07	.5	.82	.26	12 1-3-2-4-3
24	35	.7	.80	-1.58	.49	1.07	.3	.84	.0	.96	.46	10 2-4-3-1
12	35	.4	.27	.80	.45	1.08	.4	.79	-.1	.96	.37	11 1-3-1-2-4
31	35	.9	.98	-3.87	.71	1.05	.2	.53	.6	1.00	.47	5 2-1-4
27	35	.8	.90	-2.37	.54	.59	-1.3	.43	-.2	1.35	.65	8 1-4-2-3
7	35	.2	.10	1.96	.53	.70	-1.0	.38	-.2	1.34	.44	13 1-4-3-2-4
32	35	.9	.99	-4.45	.82	.91	.0	.35	.8	1.08	.48	4 1-3-4
30	35	.9	.96	-3.41	.65	.62	-1.0	.21	.0	1.34	.63	9 1-3-2-4
1	35	.0	.01	4.85	1.08	.76	.0	.11	1.4	1.18	.24	15 1-3-2-4-1-3
1	35	.0	.01	4.85	1.08	.76	.0	.11	1.4	1.18	.24	16 1-4-2-3-1-4
1	35	.0	.01	4.85	1.08	.76	.0	.11	1.4	1.18	.24	17 1-4-3-1-2-4
35	35	1.0	1.00	(-6.64	1.86)	Minimum					.00	1 1-4
35	35	1.0	1.00	(-6.64	1.86)	Minimum					.00	2 2-3
35	35	1.0	1.00	(-6.64	1.86)	Minimum					.00	3 1-2-4
0	35	.0	.00	(6.18	1.86)	Maximum					.00	18 4-1-3-4-2-1-4

**54.** How well do the observations of each element fit with the estimate of the its measure?  
 If the fit is good, then we can have confidence that the measure means what it says.  
 If the fit is bad, then the measure could mislead us.  
 If the fit is too good, then perhaps something is constraining the data to be too coherent.  
 If the fit is too bad, then those data could also be damaging the measures of other elements.

**55.** So we have the process of fit evaluation. In *Facets*, most fit statistics are based on summarizing the residuals that we've already thought about one at a time.



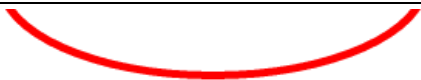

**56.** Imagine that we administer a dichotomous test in which the items are ordered from easy to difficult. What would we expect would happen when a typical person takes the test?  
 Success on the easy items 😊! Failure on the hard items ☹️.  
 And a transition zone 🤔 where the items are about as difficult as the person is able, so we expect to see some successes and some failure.  
 This is what has happened with the top left response-string in the Table in [#Error! Reference source not found.](#):  
 "1110110110100000". And this is also the pattern that the Rasch model predicts: "There is nothing so practical as a good theory" (Kurt Lewin, 1951, p. 169)

Person Responses:  
 Easy -- Items -- Hard  
 1110110110100000

57.	So, how can we verify that this response string does match Rasch expectations? We do this using mean-square fit statistics. A mean-square is a chi-square divided by its degrees of freedom [see <b>Appendix 2. if chi-square sounds like Greek to you ...</b> ]. Let's start with chi-square fit statistics ....	
58.	<p><b>Chi-square fit statistics</b> are very useful for diagnosing the standardized residuals. The standardized residuals are modeled to be unit normal deviates. So when we square them and sum them, we expect their sum will approximate a chi-square distribution with mean equal to the count of the standardized residuals.</p> <p>If the chi-square value is much above the count, then the standardized residuals are further away from 0, on average, than the Rasch model predicts. The observations are farther from their expectations than the Rasch model predicts. The data are too unpredictable, “noisy”. They “underfit” the Rasch model.</p> <p>If the chi-square value is much below the count and so much closer to zero, then the standardized residuals are closer to 0, on average, than the Rasch model predicts. The observations are closer to their expectations than the Rasch model predicts. The data are too predictable. The unexpectedness in the data is “muted”. The data “overfit” the Rasch model.</p> <p>The value of the chi-square, along with its degrees of freedom, enable us to compute how unlikely these data are to be observed by chance when the data fit the Rasch model. When we deem the data too unlikely to have occurred by chance, then we declare that “the data misfit the model”.</p>	
59.	Chi-square statistics are useful for quantifying the fit of the data to the Rasch model, but we can make them even more convenient. The expected mean of a chi-square distribution is its “degrees of freedom”, the number of independent squared unit-normal distributions it represents. If we <b>divide a chi-square value by its degrees of freedom, then we have a mean-square value.</b>	<p>chi-square <math>\chi^2</math> / degrees of freedom (d.f.)  = mean-square (MnSq)  mean (expectation) of MnSq = 1.0  model variance of MnSq = 2 / d.f.  standard deviation of MnSq = <math>\sqrt{(2/d.f.)}</math></p> <p><b>The expected value of a mean-square is 1.0</b></p>
60.	<b>G. Rasch Theory: “The data misfit the model!”</b> 🙄	
61.	<p><i>Are you surprised by that statement?</i> Many statisticians would be. Descriptive statistics are based on summarizing the data efficiently and parsimoniously. The data are considered to be the given (<i>Latin</i> “datum”) truth. The statistical model (regression, ANOVA, etc.) is intended to describe the dataset. So a good <b>descriptive</b> statistical model is one which fits the data. If the model misfits the data, then try a different descriptive model.</p> <p>Rasch is a <b>prescriptive</b> statistical method. The Rasch model gives us what we want (additive measures in a unidimensional framework), so it is our “truth”. The data may, or may not, contain the information that we need. So good data fit our Rasch model. If the data don't fit the model usefully, then the dataset as a whole doesn't support unidimensional measurement. Some part of the dataset may. In fact, usually most of a dataset does, if it is intended to capture one latent variable.</p> <p><i>Thought:</i> Raw scores are the “sufficient statistics” for a Rasch analysis. If the dataset doesn't conform to Rasch analysis, then it doesn't conform to raw-score analysis either ! (But CTT analysts usually do not know this). Raw-score fit analysis tends to be superficial, so the misfit in the dataset to a raw-score Classical Test Theory model is often overlooked.</p>	
62.	We have now had 40 years experience with mean-squares since Wright & Panchapakesan (1969) proposed them for Rasch usage (see Optional Reading at <a href="#">#178</a> ). The following Tables summarizes them from a Rasch measurement perspective.	

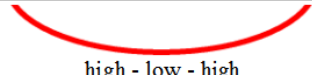
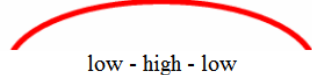
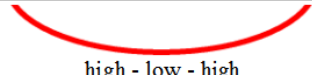
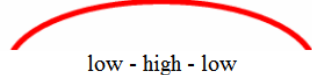
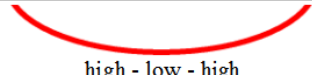
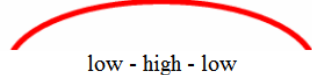
63.	H. Table 7: Interpretation of Element-level Mean-Square Fit Statistics:				
Mean-square	Interpretation				
>2.0	Distorts or degrades the measurement system. (The background noise is starting to drown out the music.)				
1.5 - 2.0	Unproductive for construction of measurement, but not degrading. (The background noise is audible, but not intrusive to the music.)				
0.5 - 1.5	Productive for measurement. (Beautiful music)				
<0.5	Less productive for measurement, but not degrading. May produce misleadingly good reliabilities and separations. (Music too quiet)				

64.	Person Responses: Easy -- Items -- Hard	Diagnosis Pattern	OUTFIT Mean-square	INFIT Mean-square	Point- Measure Correlation
	111 0110110100 000	Modeled/Ideal	1.0	1.1	.62
	111 1111100000 000	Guttman/Deterministic	<b>0.3</b>	<b>0.5</b>	<b>0.87</b>
	000 0000011111 111	Miscode	<b>12.6</b>	4.3	<b>-0.87</b>
	011 1111110000 000	Carelessness/Sleeping	<b>3.8</b>	1.0	0.65
	111 1111000000 001	Lucky Guessing	<b>3.8</b>	1.0	0.65
	101 0101010101 010	Response set/Miskey	<b>4.0</b>	<b>2.3</b>	<b>0.11</b>
	111 1000011110 000	Special knowledge	0.9	<b>1.3</b>	0.43
	111 1010110010 000	Imputed outliers *	<b>0.6</b>	1.0	0.62
	111 0101010101 000	Low discrimination	<b>1.5</b>	<b>1.6</b>	0.46
	111 1110101000 000	High discrimination	<b>0.5</b>	0.7	0.79
	111 1111010000 000	Very high discrimination	<b>0.3</b>	<b>0.5</b>	<b>0.84</b>
	 high - low - high	OUTFIT sensitive to outlying observations	>>1.0 unexpected outliers	>>1.0 disturbed pattern	
	 low - high - low	INFIT sensitive to pattern of inlying observations	<<1.0 overly predictable outliers	<<1.0 Guttman pattern	

65. Look back at the Table above. The mean-squares for our imagined typical respondent are near to 1.0 - good!	<table border="1"> <thead> <tr> <th>Person Responses: Easy -- Items -- Hard</th> <th>Diagnosis Pattern</th> <th>OUTFIT Mean-square</th> <th>INFIT Mean-square</th> </tr> </thead> <tbody> <tr> <td>111 0110110100 000</td> <td>Modelled/Ideal</td> <td>1.0</td> <td>1.1</td> </tr> </tbody> </table>	Person Responses: Easy -- Items -- Hard	Diagnosis Pattern	OUTFIT Mean-square	INFIT Mean-square	111 0110110100 000	Modelled/Ideal	1.0	1.1
Person Responses: Easy -- Items -- Hard	Diagnosis Pattern	OUTFIT Mean-square	INFIT Mean-square						
111 0110110100 000	Modelled/Ideal	1.0	1.1						
66. What about the “lucky guesser” who succeeded on the most difficult item. The OUTFIT mean-square is 3.8, much bigger than 1.0. That lucky guess has degraded the guesser’s measure. It is less secure as a basis for inference. Do we really want “guessing for success”?	<table border="1"> <thead> <tr> <th>Person Responses: Easy -- Items -- Hard</th> <th>Diagnosis Pattern</th> <th>OUTFIT Mean-square</th> <th>INFIT Mean-square</th> </tr> </thead> <tbody> <tr> <td>111 1111000000 01</td> <td>Lucky Guessing</td> <td>3.8</td> <td>1.0</td> </tr> </tbody> </table>	Person Responses: Easy -- Items -- Hard	Diagnosis Pattern	OUTFIT Mean-square	INFIT Mean-square	111 1111000000 01	Lucky Guessing	3.8	1.0
Person Responses: Easy -- Items -- Hard	Diagnosis Pattern	OUTFIT Mean-square	INFIT Mean-square						
111 1111000000 01	Lucky Guessing	3.8	1.0						


67.	I. Table 7: Outfit vs. Infit																																																											
68.	<p>Did you notice that the INFIT mean-square for the lucky-guesser is 1.0, its expected value? What is going on?</p> <p><b>The Outfit statistic is outlier-sensitive.</b> <b>The Infit statistic is sensitive to patterns in the targeted responses. It is inlier-pattern sensitive.</b></p>	<table border="1"> <tr> <td data-bbox="933 231 1258 325">  high - low - high </td> <td data-bbox="1258 231 1511 325">           OUTFIT sensitive to outlying observations </td> </tr> <tr> <td data-bbox="933 325 1258 420">  low - high - low </td> <td data-bbox="1258 325 1511 420">           INFIT sensitive to pattern of inlying observations </td> </tr> </table>	 high - low - high	OUTFIT sensitive to outlying observations	 low - high - low	INFIT sensitive to pattern of inlying observations																																																						
 high - low - high	OUTFIT sensitive to outlying observations																																																											
 low - high - low	INFIT sensitive to pattern of inlying observations																																																											
69.	<p>Take a look at “special knowledge”. Imagine the items are in 4 cluster of difficulty: addition, subtraction, multiplication, division. Then most children will follow the typical Rasch pattern. But those who are taught: addition, multiplication, subtraction, division will have a different pattern: fail on subtraction, succeed on multiplication.</p>	<p><i>Special knowledge or Alternative curriculum</i></p> <p><b>Person Responses:</b> <b>Easy -- Items -- Hard</b></p> <hr/> <p>111;100011110;000</p> <p>Add-Subtract-Multiply-Divide</p>																																																										
70.	<p>The OUTFIT statistic is 0.9 (less than 1.0). The Outfit statistic reports that responses far from the person ability are predictable. The INFIT statistic is 1.3, reporting the patterns in the data are somewhat unpredictable. The Infit statistic detects the unexpected <b>pattern</b> of responses near the person ability.</p>	<table border="1"> <thead> <tr> <th data-bbox="933 777 1144 871">Person Responses: Easy -- Items -- Hard</th> <th data-bbox="1144 777 1315 871">Diagnosis Pattern</th> <th data-bbox="1315 777 1412 871">OUTFIT Mean-square</th> <th data-bbox="1412 777 1511 871">INFIT Mean-square</th> </tr> </thead> <tbody> <tr> <td data-bbox="933 882 1144 934">111;100011110;000</td> <td data-bbox="1144 882 1315 934">Special knowledge</td> <td data-bbox="1315 882 1412 934">0.9</td> <td data-bbox="1412 882 1511 934">1.3</td> </tr> </tbody> </table>	Person Responses: Easy -- Items -- Hard	Diagnosis Pattern	OUTFIT Mean-square	INFIT Mean-square	111;100011110;000	Special knowledge	0.9	1.3																																																		
Person Responses: Easy -- Items -- Hard	Diagnosis Pattern	OUTFIT Mean-square	INFIT Mean-square																																																									
111;100011110;000	Special knowledge	0.9	1.3																																																									
71.	<p>Mathematically, the OUTFIT Mean-square is the conventional statistical chi-square divided by its degrees of freedom. The Infit statistic is an information-weighted mean-square statistic.</p>	<p>For the N observations that we are summarizing in the mean-square statistics:</p> <p>Outfit Mean-square = <math>\sum (R_{ni}^2 / V_{ni}) / N</math></p> <p>Infit Mean-square = <math>\sum (R_{ni}^2) / \sum V_{ni}</math></p>																																																										
72.	<p>Glance back at Table 7.2.2. It should be starting to make more sense to you.</p> <p><b>Red arrow:</b> Item 7 has the biggest Outfit mean-square, MnSq, statistic: 2.25. This is much bigger than the expected mean-square of 1.0. There is more “unmodeled noise” than useful “statistical information” in this item.</p> <p><b>Green arrow:</b> Item 14 has the second biggest Outfit MnSq: 1.48.</p>	<table border="1"> <thead> <tr> <th>Obsvd</th> <th>Obsvd</th> <th>Obsvd</th> <th>Fair-M</th> <th>Model</th> <th>Infit</th> <th>Outfit</th> <th>Estim.</th> <th></th> <th></th> </tr> <tr> <th>Score</th> <th>Count</th> <th>Average</th> <th>Average</th> <th>Measure</th> <th>S.E.</th> <th>MnSq</th> <th>ZStd</th> <th>MnSq</th> <th>ZStd</th> <th>Discrm</th> <th>No Tapping items</th> </tr> </thead> <tbody> <tr> <td>31</td> <td>34</td> <td>.9</td> <td>.97</td> <td>-3.84</td> <td>.71</td> <td>1.34</td> <td>.8</td> <td>2.25</td> <td>1.1</td> <td>.54</td> <td>7 1-4-3-2</td> </tr> <tr> <td>3</td> <td>34</td> <td>.1</td> <td>.03</td> <td>3.37</td> <td>.70</td> <td>1.57</td> <td>1.4</td> <td>1.48</td> <td>.9</td> <td>.62</td> <td>14 1-4-2-3-4-1</td> </tr> <tr> <td>30</td> <td>34</td> <td>.9</td> <td>.96</td> <td>-3.38</td> <td>.64</td> <td>1.18</td> <td>.9</td> <td>.97</td> <td>.6</td> <td>.87</td> <td>6 3-4-1</td> </tr> </tbody> </table>	Obsvd	Obsvd	Obsvd	Fair-M	Model	Infit	Outfit	Estim.			Score	Count	Average	Average	Measure	S.E.	MnSq	ZStd	MnSq	ZStd	Discrm	No Tapping items	31	34	.9	.97	-3.84	.71	1.34	.8	2.25	1.1	.54	7 1-4-3-2	3	34	.1	.03	3.37	.70	1.57	1.4	1.48	.9	.62	14 1-4-2-3-4-1	30	34	.9	.96	-3.38	.64	1.18	.9	.97	.6	.87	6 3-4-1
Obsvd	Obsvd	Obsvd	Fair-M	Model	Infit	Outfit	Estim.																																																					
Score	Count	Average	Average	Measure	S.E.	MnSq	ZStd	MnSq	ZStd	Discrm	No Tapping items																																																	
31	34	.9	.97	-3.84	.71	1.34	.8	2.25	1.1	.54	7 1-4-3-2																																																	
3	34	.1	.03	3.37	.70	1.57	1.4	1.48	.9	.62	14 1-4-2-3-4-1																																																	
30	34	.9	.96	-3.38	.64	1.18	.9	.97	.6	.87	6 3-4-1																																																	

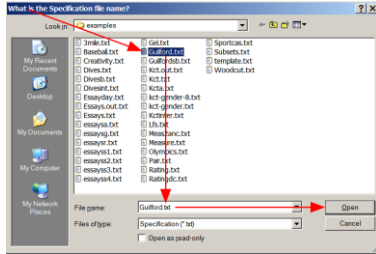
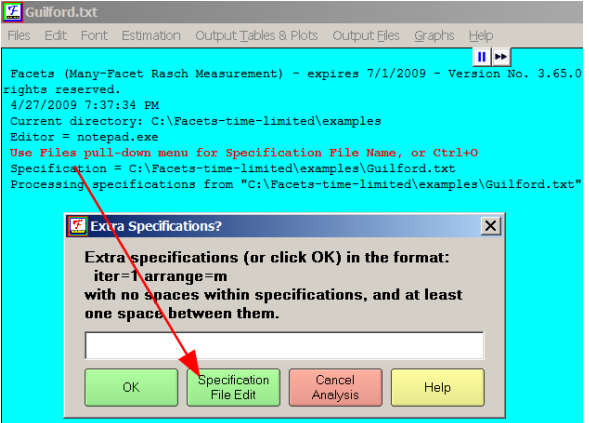

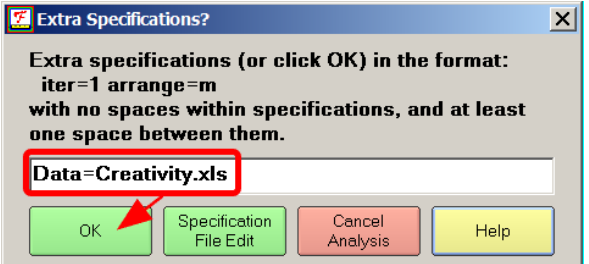
<p>73. <b>Red box:</b> Item 7 (biggest OUTFIT mean-square) has two unexpected responses in Table 4.1  <b>Green box:</b> Item 14 (second biggest Outfit MnSq) has only one response in Table 4.1, but it is the most unexpected response. This is the most outlying response.</p> <p><b>We see that the unexpected responses in Table 4.1 can cause the Outfit MnSq statistics in Table 7 to be large.</b> Start by looking at the Outfit statistics in Table 7 to localize the problem areas in the data. Table 4 tends to be too detailed.</p>	<p>Table 4.1 Unexpected Responses (7 residuals sorted by u).</p> <pre> -----+-----   Cat  Score  Exp.  Resd StRes  Nu Chil Nu Tapping items   -----+-----   1      1      .0   1.0  6.2    2 Boy  14 1-4-2-3-4-1     0      0      1.0  -1.0 -6.1    2 Boy   7 1-4-3-2     0      0      1.0  -1.0 -6.1   24 Girl  7 1-4-3-2     0      0      1.0  -1.0 -4.8   24 Girl  6 3-4-1   -----+----- </pre> <p>In big datasets, Table 4 can be come unmanageably long (which is why it is after Table 8 in the Report output file)</p>
<p>74. Diagnosing misfit:  Large OUTFIT mean-square &gt; 1.5 - Unexpected off-target observations - Look at Table 4  Small OUTFIT mean-square &lt; 0.5 - Off-target observations too predictable - Are there imputed data or other constraints?  Large INFIT mean-square &gt; 1.5 - Unexpected patterns in on-target observations - Very difficult to investigate.  <i>Suggestion:</i> Write out the residual file to Excel. Sort on person (or item, etc.) element number and “logit”. Look at patterns in responses near logit 0.  Small INFIT mean-square &lt; 0.5 - On-target observations too predictable - Are there redundant items or response sets in rated items?</p>	
<p>75. <b>J. Table 7: Misfit: Size vs. Significance: MnSq vs. Zstd</b></p>	
<p>76. We know that a large mean-square statistic flags unexpectedness in the data. But is this an unusual amount of unexpectedness, or merely a reflection of the randomness in the data which the Rasch model requires? The <b>Zstd</b> statistics (mean-squares standardized as z-statistics) answer this. The Outfit and Infit Mean-squares are derived from chi-square statistics with their d.f.. So we know how unlikely we are to observe any particular mean-square value (or worse). This is what the Zstd statistics report.</p>	
<p>77. We could report the probability of the mean-square. Computing the actual d.f. is complicated, so let’s assume the mean-square value is a chi-square with 1 d.f.</p>	<p>Item 7: Outfit MnSq = 2.25  chi-square = 2.25 with d.f. = 1  probability ≈ 0.13</p>
<p>78. Our experience is that small probabilities become long numbers that are often difficult to think with. So instead of reporting the probability, we report the equivalent unit-normal deviate [see Appendix 1], called Zstd, “the mean-square statistic standardized like a Z-score”. This is also a Student’s <i>t</i>-statistic with infinite d.f.</p>	<p>Item 7: Outfit Zstd = 1.1  probability = 0.14</p>
<p>79. Reporting Zstd simplifies interpretation. See Appendix 1 for more Zstd values.</p>	<p> Zstd  ≥ 2.0 are statistically significant   Zstd  ≥ 2.6 are highly significant</p>
<p>80. So, the rule-of-thumb with Outfit and Infit statistics is:  <b>“MnSq size: large enough to be distorting; MnSq &gt; 1.5  Zstd significance: improbable enough to be surprising. Zstd &gt; 2.0”</b></p>	

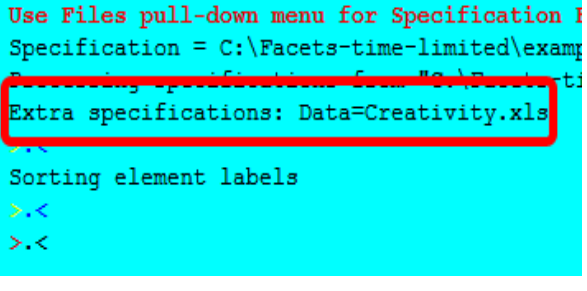
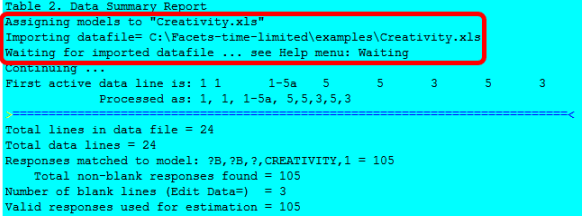
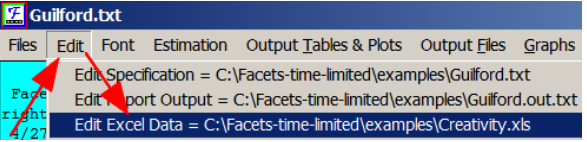
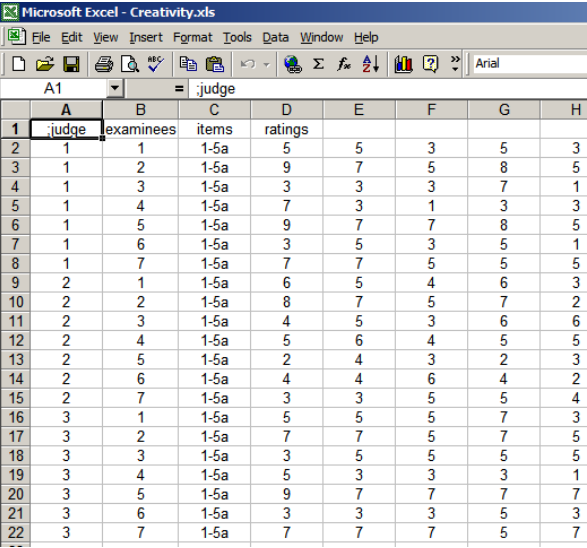
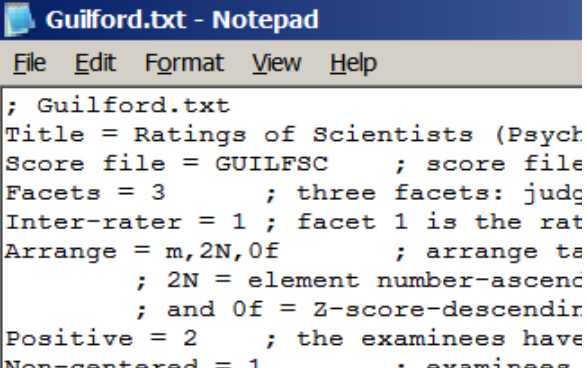
<p><b>81.</b></p>	<p><i>What about unexpectedly low mean-squares?</i></p> <p>These are paradoxical! Here is a Guttman response string, name after psychometrician Louis Guttman. He proposed that the ideal response string is one where someone succeeds on all the easy (for that person) items and fails on all the hard items. The result would be the response string you see here. This is also the best possible response string for Classical Test Theory. It has the highest point-biserial correlation and the highest discrimination indexes, and results in the highest test reliabilities.</p> <p>Guttman’s ideal response-pattern is perfect for <b>ordering</b>, but <b>not for measurement</b>. For measurement we need <b>uncertainty</b> in the responses. <b>The closer to the item difficulty, the more uncertain each person’s responses</b>. This fundamental to Rasch theory.</p>	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 30%;">Person Responses: Easy--Items--Hard</th> <th style="width: 20%;">Diagnosis Pattern</th> <th style="width: 15%;">OUTFIT Mean-square</th> <th style="width: 15%;">INFIT Mean-square</th> </tr> </thead> <tbody> <tr> <td>111 1111100000 000</td> <td>Guttman/Deterministic</td> <td style="color: red;">0.3</td> <td style="color: red;">0.5</td> </tr> </tbody> </table> <p><b>Guttman/deterministic.</b> Louis Guttman, a leading psychometrician around 1950, proposed that the ideal item would be one on which all low performers failed and all high performers succeeded. It would act like a switch. It would have infinitely high discrimination. If you knew the location on the latent variable of the switch for each person, then all person responses would be exactly determined.</p>	Person Responses: Easy--Items--Hard	Diagnosis Pattern	OUTFIT Mean-square	INFIT Mean-square	111 1111100000 000	Guttman/Deterministic	0.3	0.5
Person Responses: Easy--Items--Hard	Diagnosis Pattern	OUTFIT Mean-square	INFIT Mean-square							
111 1111100000 000	Guttman/Deterministic	0.3	0.5							
<p><b>82.</b></p>	<p><i>So what is wrong with a Guttman pattern from a Rasch perspective?</i></p> <p>Rasch proposes that any reasonable subset of items should give statistically the same estimate as the full set. So let’s split the test high-low:</p> <ol style="list-style-type: none"> <li>1. According to the easy items, our respondent is a genius.</li> <li>2. According to the difficult items, our respondent is a dunce.</li> </ol> <p><i>A contradiction!</i></p> <p>The problem is that the Rasch <b>transition zone</b> of uncertain responses to the targeted items is missing. Of course, the Rasch measure has correctly located the respondent between the easy and hard items on the latent variable, but the response string is squeezed together from the Rasch perspective.</p>	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; text-align: center;">         Responses: Easy--Items--Hard   <span style="border: 2px solid green; padding: 2px;">111 11111</span> </td> <td style="width: 50%; text-align: center;"> <b>Genius!</b> </td> </tr> <tr> <td style="width: 50%; text-align: center;">         Responses: Easy--Items--Hard   <span style="border: 2px solid green; padding: 2px;">00000 000</span> </td> <td style="width: 50%; text-align: center;"> <b>Dunce!</b> </td> </tr> </table>	Responses: Easy--Items--Hard  <span style="border: 2px solid green; padding: 2px;">111 11111</span>	<b>Genius!</b>	Responses: Easy--Items--Hard  <span style="border: 2px solid green; padding: 2px;">00000 000</span>	<b>Dunce!</b>				
Responses: Easy--Items--Hard  <span style="border: 2px solid green; padding: 2px;">111 11111</span>	<b>Genius!</b>									
Responses: Easy--Items--Hard  <span style="border: 2px solid green; padding: 2px;">00000 000</span>	<b>Dunce!</b>									



<p><b>83.</b></p>	<p><b><i>Guttman Patterns and Low Mean-squares &lt; 0.5:</i></b>  Guttman patterns produce low Mean-Squares.  Low mean-squares correspond to persons and items which are too predictable. They are lacking in the uncertainty Rasch needs for constructing measures.  This makes the reported standard errors (measurement precisions) too small and the reported reliabilities (measure reproducibility) too high. In general, however, <b>low mean-squares are not a serious problem.</b></p> <p>Small standard errors (high precision) and high reliability (high measure reproducibility, a consequence of high precision) are good, but only if that level of precision is really supported by the data. Here the reported standard errors (though computed correctly) are too small from a substantive perspective.  A parallel situation arises in physical measurement. Suppose you weigh yourself 100 times. Then your weight will be the average of those weights with precision (standard error of the mean) S.E.M. = S.D. of your 100 weights / 10. But do you believe this high precision about your own weight? No. It is statistically correct, but substantively misleading. You weight varies by more than that S.E.M. during each day. The calculated standard error of your weight is too small, and so may mislead you about how precisely you know your own weight.</p>											
<p><b>84.</b></p>	<p>In general, <b>low mean-squares are not a serious problem, but high mean-squares are.</b>  Low mean-squares rarely lead to incorrect inferences about the meaning of measures, unless they are caused by constraints which invalidate the measures.  So always investigate and remedy high mean-squares, and then re-analyze your data, before investigating low mean-squares. The overall average mean-squares are usually close to 1.0, so high mean-squares force there to be low mean-squares.</p>											
<p><b>85.</b></p>	<p>So what values of the mean-square statistics cause us real concern? Here is my summary table from Winsteps Help “Special Topic” “Misfit Diagnosis ...”</p> <p><i>Here’s a story:</i>  When the mean-square value is around 1.0, we are hearing music! The measurement is <b>accurate</b>  When the mean-square value is less than 1.0, the music is becoming quieter, becoming muted. When the mean-square is less than 0.5, the item is providing only have the music volume (technically “statistical information”) that it should. But mutedness does not cause any real problems. Muted items aren’t efficient. The measurement is less accurate.</p> <p>When the mean-squares go above 1.0, the music level stays constant, but now there is other noise: rumbles, clunks, pings, etc. When the mean-square gets above 2.0, then the noise is louder than the music and starting to drown it out. The measures (though still forced to be additive) are becoming distorted relative to the response strings. So <b>it is mean-square values greater than 2.0 that are of greatest concern. The measurement is inaccurate.</b></p>	<table border="1"> <thead> <tr> <th colspan="2">Interpretation of mean-square fit statistics:</th> </tr> </thead> <tbody> <tr> <td>&gt;2.0</td> <td>Distorts or degrades the measurement system. <i>But be alert, the <b>explosion</b> caused by only one very lucky guess can send a mean-square statistic above 2.0. Eliminate the lucky guess from the data set, and harmony will reign!</i></td> </tr> <tr> <td>1.5 - 2.0</td> <td>Unproductive for construction of measurement, but not degrading.</td> </tr> <tr> <td>0.5 - 1.5</td> <td>Productive for measurement.</td> </tr> <tr> <td>&lt;0.5</td> <td>Less productive for measurement, but not degrading. May produce misleadingly good reliabilities and separations.</td> </tr> </tbody> </table>	Interpretation of mean-square fit statistics:		>2.0	Distorts or degrades the measurement system. <i>But be alert, the <b>explosion</b> caused by only one very lucky guess can send a mean-square statistic above 2.0. Eliminate the lucky guess from the data set, and harmony will reign!</i>	1.5 - 2.0	Unproductive for construction of measurement, but not degrading.	0.5 - 1.5	Productive for measurement.	<0.5	Less productive for measurement, but not degrading. May produce misleadingly good reliabilities and separations.
Interpretation of mean-square fit statistics:												
>2.0	Distorts or degrades the measurement system. <i>But be alert, the <b>explosion</b> caused by only one very lucky guess can send a mean-square statistic above 2.0. Eliminate the lucky guess from the data set, and harmony will reign!</i>											
1.5 - 2.0	Unproductive for construction of measurement, but not degrading.											
0.5 - 1.5	Productive for measurement.											
<0.5	Less productive for measurement, but not degrading. May produce misleadingly good reliabilities and separations.											

<p><b>86.</b></p>	<p>Every Rasch analyst has favorite rules for identifying misfit. The Reasonable Mean-Square Fit Value is from <a href="http://www.rasch.org/rmt/rmt83b.htm">http://www.rasch.org/rmt/rmt83b.htm</a></p> <p>No rules are decisive, but many are helpful.</p>	<table border="1"> <thead> <tr> <th colspan="2">Reasonable Item Mean-square Ranges for INFIT and OUTFIT</th> </tr> <tr> <th>Type of Test</th> <th>Range</th> </tr> </thead> <tbody> <tr> <td>MCQ (High stakes)</td> <td>0.8 - 1.2</td> </tr> <tr> <td>MCQ (Run of the mill)</td> <td>0.7 - 1.3</td> </tr> <tr> <td>Rating scale (survey)</td> <td>0.6 - 1.4</td> </tr> <tr> <td>Clinical observation</td> <td>0.5 - 1.7</td> </tr> <tr> <td>Judged (agreement encouraged)</td> <td>0.4 - 1.2</td> </tr> </tbody> </table>	Reasonable Item Mean-square Ranges for INFIT and OUTFIT		Type of Test	Range	MCQ (High stakes)	0.8 - 1.2	MCQ (Run of the mill)	0.7 - 1.3	Rating scale (survey)	0.6 - 1.4	Clinical observation	0.5 - 1.7	Judged (agreement encouraged)	0.4 - 1.2
Reasonable Item Mean-square Ranges for INFIT and OUTFIT																
Type of Test	Range															
MCQ (High stakes)	0.8 - 1.2															
MCQ (Run of the mill)	0.7 - 1.3															
Rating scale (survey)	0.6 - 1.4															
Clinical observation	0.5 - 1.7															
Judged (agreement encouraged)	0.4 - 1.2															
<p><b>87.</b></p>	<p>Close all open <i>Facets</i> windows</p>															
<p><b>88.</b></p>																

89.	K. Facets Specification and Data: The Guilford Data	
90.	<p>Let's apply what we've learned to some 3-facet rating data.</p> <p>Launch <i>Facets</i>.</p> <p>Click on "Files"</p> <p>Click on "Specification File Name?"</p> <p>Double-click on "<b>Guilford.txt</b>"</p>	
91.	<p>"Extra specifications?"</p> <p><b>Click on "Specification File Edit"</b></p>	
92.	<p>Guilford.txt displays in a NotePad window.</p> <p><b>Scroll down to Data=</b></p> <p>Notice that there are alternative data files. Most are commented out with ";"</p> <p><b>blue box:</b> All these data files contain the same observations. You can see these in option 6.</p> <p><b>red box:</b> We will use the Excel file, "Creativity.xls"</p> <p>Do not edit Guilford.txt - we will make the change using <i>Extra Specifications?</i></p>	<pre> ; This illustrates 6 different ways of accessing your data. ; 5 ways should be commented out with ; so that only one way is active.  ; 1. Data from an SPSS data file Data= Creativity.sav ; SPSS file with 1-5  ; 2. Data from an external text file ; Data = Creativity.txt ; standard text data file  ; 3. Data from an Excel spreadsheet ; Data= Creativity.xls ; Excel file with 1-5  ; 4. Data from an SPSS data file, using dvalues= to simplify formatting ; dvalues = 3, 1-5 ; Data = Guilford.sav ; SPSS file omitting 1-5 for the 3rd facet  ; 5. Data from an Excel data file, using dvalues= to simplify formatting ; dvalues = 3, 1-5 ; Data = Guilford.xls ; Excel file omitting 1-5 for the 3rd facet  ; 6. Data included in the Specification file. You can use , or tab or blank as separators ; Data= ;1,1,1-5,5,3,3,3 ;1,2,1-5,9,7,5,8,5 </pre>
93.	<p>Click on your Facets analysis on the Windows task bar</p>	
94.	<p>Type into the Extra Specifications? box: <b>Data=Creativity.xls</b> <i>with no spaces</i></p> <p>(or copy-and-paste: Ctrl+C Ctrl+V)</p> <p>Click on <b>OK</b></p>	

<p>95. The Extra specification, Data=Creativity.xls is shown in the Facets Analysis window.</p> <p>“What is the Report Output file name?” - click on “<b>Open</b>” to accept the default value: <i>guilford.out.txt</i></p> <p>Analysis begins ....</p>																																																																																																																																																																																																																
<p>96. Notice on your <i>Facets</i> analysis window screen that the “Creativity.xls” is imported.</p> <p>Facets launches Excel to obtain the responses. Excel is sometimes slow, so you may see the “Waiting ....” message.</p>																																																																																																																																																																																																																
<p>97. Take a look at the Guilford data in Creativity.xls <i>Facets</i> menu bar: Click on “Edit” Click on “Edit Excel Data” Excel launches and displays the data ...</p>																																																																																																																																																																																																																
<p>98. The data, from “Psychometric Methods” by J.P. Guilford (1954) are of 3 Senior Scientists (the judges) rating 7 Junior Scientists (the examinees) or 5 items of Creativity. The observed range of the rating scale is 1-9. Guilford omits to tell us what the possible range was. So row 1 of the spreadsheet is: ; judges examinees items ratings the “;” is to tell <i>Facets</i> this row is a comment, not data. Row 2 is the first data row, it says: Judge 1 rated examinee 1 on 5 items, 1 to 5, and the ratings were 5, 5, 3, 5, 3 1-5 means “items 1, 2, 3, 4, 5” “1-5a” is to prevent Excel converting 1-5 into -4. Facets ignores the “a”. There are 21 rows of data, and 105 ratings.</p>	 <table border="1" data-bbox="935 982 1518 1430"> <thead> <tr> <th></th> <th>A</th> <th>B</th> <th>C</th> <th>D</th> <th>E</th> <th>F</th> <th>G</th> <th>H</th> </tr> <tr> <th>1</th> <td>judge</td> <td>examinees</td> <td>items</td> <td>ratings</td> <td></td> <td></td> <td></td> <td></td> </tr> </thead> <tbody> <tr><td>2</td><td>1</td><td>1</td><td>1-5a</td><td>5</td><td>5</td><td>3</td><td>5</td><td>3</td></tr> <tr><td>3</td><td>1</td><td>2</td><td>1-5a</td><td>9</td><td>7</td><td>5</td><td>8</td><td>5</td></tr> <tr><td>4</td><td>1</td><td>3</td><td>1-5a</td><td>3</td><td>3</td><td>3</td><td>7</td><td>1</td></tr> <tr><td>5</td><td>1</td><td>4</td><td>1-5a</td><td>7</td><td>3</td><td>1</td><td>3</td><td>3</td></tr> <tr><td>6</td><td>1</td><td>5</td><td>1-5a</td><td>9</td><td>7</td><td>7</td><td>8</td><td>5</td></tr> <tr><td>7</td><td>1</td><td>6</td><td>1-5a</td><td>3</td><td>5</td><td>3</td><td>5</td><td>1</td></tr> <tr><td>8</td><td>1</td><td>7</td><td>1-5a</td><td>7</td><td>7</td><td>5</td><td>5</td><td>5</td></tr> <tr><td>9</td><td>2</td><td>1</td><td>1-5a</td><td>6</td><td>5</td><td>4</td><td>6</td><td>3</td></tr> <tr><td>10</td><td>2</td><td>2</td><td>1-5a</td><td>8</td><td>7</td><td>5</td><td>7</td><td>2</td></tr> <tr><td>11</td><td>2</td><td>3</td><td>1-5a</td><td>4</td><td>5</td><td>3</td><td>6</td><td>6</td></tr> <tr><td>12</td><td>2</td><td>4</td><td>1-5a</td><td>5</td><td>6</td><td>4</td><td>5</td><td>5</td></tr> <tr><td>13</td><td>2</td><td>5</td><td>1-5a</td><td>2</td><td>4</td><td>3</td><td>2</td><td>3</td></tr> <tr><td>14</td><td>2</td><td>6</td><td>1-5a</td><td>4</td><td>4</td><td>6</td><td>4</td><td>2</td></tr> <tr><td>15</td><td>2</td><td>7</td><td>1-5a</td><td>3</td><td>3</td><td>5</td><td>5</td><td>4</td></tr> <tr><td>16</td><td>3</td><td>1</td><td>1-5a</td><td>5</td><td>5</td><td>5</td><td>7</td><td>3</td></tr> <tr><td>17</td><td>3</td><td>2</td><td>1-5a</td><td>7</td><td>7</td><td>5</td><td>7</td><td>5</td></tr> <tr><td>18</td><td>3</td><td>3</td><td>1-5a</td><td>3</td><td>5</td><td>5</td><td>5</td><td>5</td></tr> <tr><td>19</td><td>3</td><td>4</td><td>1-5a</td><td>5</td><td>3</td><td>3</td><td>3</td><td>1</td></tr> <tr><td>20</td><td>3</td><td>5</td><td>1-5a</td><td>9</td><td>7</td><td>7</td><td>7</td><td>7</td></tr> <tr><td>21</td><td>3</td><td>6</td><td>1-5a</td><td>3</td><td>3</td><td>3</td><td>5</td><td>3</td></tr> <tr><td>22</td><td>3</td><td>7</td><td>1-5a</td><td>7</td><td>7</td><td>7</td><td>5</td><td>7</td></tr> </tbody> </table>		A	B	C	D	E	F	G	H	1	judge	examinees	items	ratings					2	1	1	1-5a	5	5	3	5	3	3	1	2	1-5a	9	7	5	8	5	4	1	3	1-5a	3	3	3	7	1	5	1	4	1-5a	7	3	1	3	3	6	1	5	1-5a	9	7	7	8	5	7	1	6	1-5a	3	5	3	5	1	8	1	7	1-5a	7	7	5	5	5	9	2	1	1-5a	6	5	4	6	3	10	2	2	1-5a	8	7	5	7	2	11	2	3	1-5a	4	5	3	6	6	12	2	4	1-5a	5	6	4	5	5	13	2	5	1-5a	2	4	3	2	3	14	2	6	1-5a	4	4	6	4	2	15	2	7	1-5a	3	3	5	5	4	16	3	1	1-5a	5	5	5	7	3	17	3	2	1-5a	7	7	5	7	5	18	3	3	1-5a	3	5	5	5	5	19	3	4	1-5a	5	3	3	3	1	20	3	5	1-5a	9	7	7	7	7	21	3	6	1-5a	3	3	3	5	3	22	3	7	1-5a	7	7	7	5	7
	A	B	C	D	E	F	G	H																																																																																																																																																																																																								
1	judge	examinees	items	ratings																																																																																																																																																																																																												
2	1	1	1-5a	5	5	3	5	3																																																																																																																																																																																																								
3	1	2	1-5a	9	7	5	8	5																																																																																																																																																																																																								
4	1	3	1-5a	3	3	3	7	1																																																																																																																																																																																																								
5	1	4	1-5a	7	3	1	3	3																																																																																																																																																																																																								
6	1	5	1-5a	9	7	7	8	5																																																																																																																																																																																																								
7	1	6	1-5a	3	5	3	5	1																																																																																																																																																																																																								
8	1	7	1-5a	7	7	5	5	5																																																																																																																																																																																																								
9	2	1	1-5a	6	5	4	6	3																																																																																																																																																																																																								
10	2	2	1-5a	8	7	5	7	2																																																																																																																																																																																																								
11	2	3	1-5a	4	5	3	6	6																																																																																																																																																																																																								
12	2	4	1-5a	5	6	4	5	5																																																																																																																																																																																																								
13	2	5	1-5a	2	4	3	2	3																																																																																																																																																																																																								
14	2	6	1-5a	4	4	6	4	2																																																																																																																																																																																																								
15	2	7	1-5a	3	3	5	5	4																																																																																																																																																																																																								
16	3	1	1-5a	5	5	5	7	3																																																																																																																																																																																																								
17	3	2	1-5a	7	7	5	7	5																																																																																																																																																																																																								
18	3	3	1-5a	3	5	5	5	5																																																																																																																																																																																																								
19	3	4	1-5a	5	3	3	3	1																																																																																																																																																																																																								
20	3	5	1-5a	9	7	7	7	7																																																																																																																																																																																																								
21	3	6	1-5a	3	3	3	5	3																																																																																																																																																																																																								
22	3	7	1-5a	7	7	7	5	7																																																																																																																																																																																																								
<p>99. Now let’s examine the Guilford specification file. You may have it on your Windows task bar, if not ...</p> <p><i>Facets</i> menu bar: Click on “Edit” Click on “Edit Specification”</p> <p><i>You are probably racing ahead of me, but just in case ...</i></p> <p>; starts a comment <b>Title=</b> specifies the title to print at the top of each Table.</p>	 <pre> ; Guilford.txt Title = Ratings of Scientists (Psych Score file = GUILFSC ; score file Facets = 3 ; three facets: judg Inter-rater = 1 ; facet 1 is the rat Arrange = m,2N,Of ; arrange ta ; 2N = element number-ascend ; and Of = Z-score-descendir Positive = 2 ; the examinees have Non-centered = 1 ; examinees </pre>																																																																																																																																																																																																															

<p><b>100</b></p>	<p><b>Score file</b>= specifies the file names to use for writing out score files for each facet. The score files contain summary statistics for each element in each facet. See <a href="#">Facets Help</a> for exact details</p>	<p><i>The Score file for Facet 1:</i></p> <pre> GUI\FSC.1.txt - Notepad File Edit Format View Help 1 Senior scientists   T.Score   T.Count   Obs.Avge   Fai   171.00    35.00     4.89   156.00    35.00     4.46   181.00    35.00     5.17 </pre>
<p><b>101</b></p>	<p><b>Facets</b> = specifies the number of facets in the analysis. We have 3 facets: judges, examinees and items.</p> <p><b>Inter-rater</b> = specifies the facet number of the rater or judge facet. This instruct <i>Facets</i> to compute rater-relevant statistics for this facet. For us, facet 1 is the judge facet, the “Senior Scientists”.</p>	<p><i>Inter-rater = will produce this in Table 7:</i></p> <pre> -----   Exact Agree.     Obs %  Exp %   N Senior scientists   -----   21.4  25.2   2 Brahe   35.7  25.8   1 Avogadro   37.1  25.3   3 Cavendish ----- </pre>
<p><b>102</b></p>	<p><b>Arrange</b> = tells <i>Facets</i> in what order to arrange the elements when they are displayed in Table 7. “Arrange = m” means “Arrange in measure order descending” so that the highest measure appears first in the Table. This is done for all the facets. “Arrange = m, 2N” means: “after doing Arrange = m, then <b>red boxes</b>: output a copy of Table 7 for facet 2 with the elements in numerical order ascending” so that element 1 is displayed first.</p>	<p><i>Arrange= will produce this:</i> (arranged by <b>2N</b>).</p> <pre> ----- it      Outfit  Estim.   q ZStd  MnSq  ZStd Discrm   N Junior Scientists   ----- 4 -3.2  .23 -3.2  1.48   1 Anne 1 -1.1  .60 -1.2  1.30   2 Betty 3  .4  1.22  .7  .84   3 Chris 0  .9  1.37  1.0  .87   4 David 4  2.2  1.94  2.2  .34   5 Edward 9  -.8  .77  -.5  .93   6 Fred 5  -.3  .84  -.4  1.37   7 George ----- </pre>
<p><b>103</b></p>	<p><b>Positive</b>= defines which facets are oriented so that more score implies more measure. In this analysis, we have chosen to do that for facet 2, the “Junior Scientists” who are the examinees.</p>	<p><i>Positive and negative facets in Table 6:</i> Table 6.0 All Facet Vertical "Rulers".</p> <pre> Vertical = (2N, 3A, 2*, 1A, 1A, S) Yardstick +-----+  Measr + junior Scientists -traits +-----+ </pre>
<p><b>104</b></p>	<p><b>Non-centered</b>= specifies which facet does not have a local origin, but is measured relative to the origins of the other facets.</p> <p>Example (not the Guilford analysis): <b>Red box</b>: shows the subjects non-centered. Other red arrows: the Raters and Items are centered</p> <p>Guilford.txt is a study of rater behavior, so we have chosen facet 1, the “Senior Scientist” judges, to be non-centered, so that they “float” relative to the other facets.</p>	<p><i>Table 6 from a large analysis:</i></p> <pre> +-----+  Measr +Subjects   +Raters   +Items   R +-----+ 5 +-----+ +-----+ +-----+ +-----+ (3)     *         4 + *             *             *         3 + *             *             *         2 + *             *             *         1 + *             *             *         0 + *             *         -1 + *         </pre>
<p><b>105</b></p>	<p>More specifications in Guilford.txt</p> <p>The next two specifications are for Table 4 of “Unexpected Responses” (or observations)</p>	<pre> Unexpected = 2 ; report ratings if s Usort = (1,2,3), (3,1,2), (2,3) ; sor Vertical = 2N,3A,2*,1L,1A ;defi Zscore = 1,2 ;report biases greate Pt-biserial = measure ; point-measure </pre>

106	<p><b>Unexpected</b>= says how unexpected? “2” means “report responses with a standardized residual, StRes, of 2 or bigger in Table 4”.</p>	<p>Table 4: Unexpected = 2:</p> <table border="1"> <thead> <tr> <th>Cat</th> <th>Score</th> <th>Exp.</th> <th>Resd</th> <th>StRes</th> </tr> </thead> <tbody> <tr> <td>6</td> <td>6</td> <td>2.9</td> <td>3.1</td> <td>2.4</td> </tr> <tr> <td>2</td> <td>2</td> <td>6.0</td> <td>-4.0</td> <td>-2.7</td> </tr> </tbody> </table>	Cat	Score	Exp.	Resd	StRes	6	6	2.9	3.1	2.4	2	2	6.0	-4.0	-2.7																																													
Cat	Score	Exp.	Resd	StRes																																																										
6	6	2.9	3.1	2.4																																																										
2	2	6.0	-4.0	-2.7																																																										
107	<p><b>Usort</b>= specifies how Table 4 is sorted. See <a href="#">Facets Help</a>.</p>	<p>Table 4.2 Unexpected Responses (4 residuals sorted by 3,1,2).</p> <table border="1"> <thead> <tr> <th>Cat</th> <th>Score</th> <th>Exp.</th> <th>Resd</th> <th>StRes</th> <th>N</th> <th>Senior</th> <th>sc</th> <th>N</th> <th>Junior</th> <th>N</th> <th>Traits</th> </tr> </thead> <tbody> <tr> <td>2</td> <td>2</td> <td>6.0</td> <td>-4.0</td> <td>-2.7</td> <td>2</td> <td>Brahe</td> <td>5</td> <td>Edward</td> <td>1</td> <td>Attack</td> <td></td> </tr> <tr> <td>6</td> <td>6</td> <td>2.9</td> <td>3.1</td> <td>2.4</td> <td>2</td> <td>Brahe</td> <td>6</td> <td>Fred</td> <td>3</td> <td>Clarity</td> <td></td> </tr> <tr> <td>2</td> <td>2</td> <td>6.1</td> <td>-4.1</td> <td>-2.7</td> <td>2</td> <td>Brahe</td> <td>5</td> <td>Edward</td> <td>4</td> <td>Daring</td> <td></td> </tr> <tr> <td>6</td> <td>6</td> <td>2.9</td> <td>3.1</td> <td>2.4</td> <td>2</td> <td>Brahe</td> <td>3</td> <td>Chris</td> <td>5</td> <td>Enthusiasm</td> <td></td> </tr> </tbody> </table>	Cat	Score	Exp.	Resd	StRes	N	Senior	sc	N	Junior	N	Traits	2	2	6.0	-4.0	-2.7	2	Brahe	5	Edward	1	Attack		6	6	2.9	3.1	2.4	2	Brahe	6	Fred	3	Clarity		2	2	6.1	-4.1	-2.7	2	Brahe	5	Edward	4	Daring		6	6	2.9	3.1	2.4	2	Brahe	3	Chris	5	Enthusiasm	
Cat	Score	Exp.	Resd	StRes	N	Senior	sc	N	Junior	N	Traits																																																			
2	2	6.0	-4.0	-2.7	2	Brahe	5	Edward	1	Attack																																																				
6	6	2.9	3.1	2.4	2	Brahe	6	Fred	3	Clarity																																																				
2	2	6.1	-4.1	-2.7	2	Brahe	5	Edward	4	Daring																																																				
6	6	2.9	3.1	2.4	2	Brahe	3	Chris	5	Enthusiasm																																																				
108	<p><b>Vertical</b>= defines the layout for the Table 6 “vertical rulers” <b>Zscore</b>= is for the interaction/bias analysis in Tutorial 3</p>	<p>Vertical = 2N,3A,2*,1L,1A Zscore = 1,2 ;report bias</p>																																																												
109	<p><b>Pt-biserial</b>= is the point-biserial correlation or the point-measure correlation.</p> <p><b>PtMea</b> is the observed point-measure correlation <b>PtExp</b> is the expected value of the point-measure correlation when the data fit the Rasch model. When possible, both the observed and the expected values of the correlations are reported.</p>	<p>Table 7:</p> <table border="1"> <thead> <tr> <th></th> <th>Correlation</th> </tr> </thead> <tbody> <tr> <td>PtMea</td> <td>PtExp</td> </tr> <tr> <td></td> <td>.49</td> </tr> <tr> <td></td> <td>.56</td> </tr> <tr> <td></td> <td>.67</td> </tr> <tr> <td></td> <td>.62</td> </tr> <tr> <td></td> <td>.51</td> </tr> </tbody> </table>		Correlation	PtMea	PtExp		.49		.56		.67		.62		.51																																														
	Correlation																																																													
PtMea	PtExp																																																													
	.49																																																													
	.56																																																													
	.67																																																													
	.62																																																													
	.51																																																													
110	<p><b>Model</b>= specifies the model. The B’s are for the interaction/bias analysis in Tutorial 3. The ?’s mean “any element”. “Creativity” is the name of a user-defined rating scale.</p> <p><b>Rating scale</b>= defines the rating scale. It is called “Creativity”, and it is “R9”, a rating scale with highest category 9. 1=lowest is the number and name of a category, * ends the category list</p>	<pre>Model = ?B,?B,?,Creativity ; A bias/interaction ; senior scientists ( ; log(PnijK/PnijK-1) ; Bn = ability n, Di ; PnijK = probability Rating scale = Creativity,R9 1 = lowest ; name of low 5 = middle ; no need to 9 = highest ; name of hig *</pre>																																																												
111	<p>In a model specification, "?" or "\$" means "any element of this facet". "#" means any element of this facet, and each element has its own rating scale. So, F<sub>k</sub> will look like: Models = ?, ?, ?, R F<sub>jk</sub> will look like: Models = ?, #, ?, R ; assuming that "j" is the second facet.</p>	<p>Allowing each judge to have a unique (partial credit) rating scale: Model = #, ?, ?, Creativity</p> <p>? does not work correctly in some versions of Windows. You can use \$ instead.</p>																																																												



112	<p><b>Labels=</b> specifies the facet and element names and numbers.</p> <p>If you look at the Junior Scientists, you can see that the element numbers can be jumbled.</p>	<pre>Labels=           ;to 1,Senior scientists ;na 1=Avogadro       ;na 2=Brahe          ;th 3=Cavendish * 2,Junior Scientists 2=Betty 5=Edward</pre>																
113	<p><b>Data=</b> specifies the data.</p> <p>The data always have the same layout, but they can be files of different types. Here are the current options:</p> <table border="1" data-bbox="175 655 880 1087"> <thead> <tr> <th>Suffix</th> <th>Data= Format</th> </tr> </thead> <tbody> <tr> <td>.txt</td> <td>text file (MS-DOS or Windows)</td> </tr> <tr> <td>.xls .xlsx</td> <td>Excel workbook: first or only worksheet</td> </tr> <tr> <td>.rda</td> <td>R statistics data file</td> </tr> <tr> <td>.sdata</td> <td>SAS data file</td> </tr> <tr> <td>.sav</td> <td>SPSS data file</td> </tr> <tr> <td>.dta</td> <td>STATA data file</td> </tr> <tr> <td>(other)</td> <td>text file (MS-DOS or Windows)</td> </tr> </tbody> </table>	Suffix	Data= Format	.txt	text file (MS-DOS or Windows)	.xls .xlsx	Excel workbook: first or only worksheet	.rda	R statistics data file	.sdata	SAS data file	.sav	SPSS data file	.dta	STATA data file	(other)	text file (MS-DOS or Windows)	<pre>; 1. Data from an SPSS data file Data= Creativity.sav ; SPSS file with 1-5  ; 2. Data from an external text file ; Data = Creativity.txt ; standard text data file  ; 3. Data from an Excel spreadsheet ; Data= Creativity.xls ; Excel file with 1-5  ; 4. Data from an SPSS data file, using dvalues= to simplify ; dvalues = 3, 1-5 ; Data = Guilford.sav ; SPSS file omitting 1-5 for the 3rd f  ; 5. Data from an Excel data file, using dvalues= to simplif ; dvalues = 3, 1-5 ; Data = Guilford.xls ; Excel file omitting 1-5 for the 3rd  ; 6. Data included in the Specification file. You can use , ; Data= ;1,1,1-5,5,5,3,5,3 ;1,2,1-5,9,7,5,8,5 ;1,3,1-5,3,3,7,1</pre>
Suffix	Data= Format																	
.txt	text file (MS-DOS or Windows)																	
.xls .xlsx	Excel workbook: first or only worksheet																	
.rda	R statistics data file																	
.sdata	SAS data file																	
.sav	SPSS data file																	
.dta	STATA data file																	
(other)	text file (MS-DOS or Windows)																	
114	<p><b>Dvalues=</b> some facet information in the data file can be specified once, instead of in every data line. This is time-saving. <i>More to come about this ...</i></p>	<pre>; 5. Data from an Excel data file, ; dvalues = 3, 1-5 ; Data = Guilford.xls ; Excel file</pre>																

115.	L. Facets Output Tables: The Guilford Report Output File	
116.	<p>On the Windows task bar, click on “Guilford.txt.out” or on the <i>Facets</i> menu bar, click on “Edit” click on “Edit Report Output”</p> <p><b>Table 1</b> reports the specifications for this analysis. The crucial details to check are the numbers of elements in each facet. If incorrect, modify your specification file.</p>	<pre>Table 1. Specifications from file "C:\Facets-time-limited\examples\Guilf Title = Ratings of Scientists (Psychometric Methods p.282 Guilford 1954) Data file = Creativity.xls Output file = C:\Facets-time-limited\examples\Guilford.out.txt  ; Data specification Facets = 3 Non-centered = 1 Positive = 2 Labels = 1,Senior scientists ; (elements = 3) 2,Junior scientists ; (elements = 7) 3,Traits ; (elements = 5) Model = ?B,?B,?,CREATIVITY,1 Rating (or other) scale = CREATIVITY,R9,General,Ordinal</pre>
117.	<p><b>Table 2</b> reports what happened to the data. We have 3 judges x 7 examinees x 5 traits = 105 observations. We expect them all to match our measurement model. <i>They do. Great!</i></p>	<pre>Table 2. Data Summary Report.  Assigning models to "creativity.xls" Total lines in data file = 22 Total data lines = 22 Responses matched to model: ?B,?B,?,CREATIVITY,1 = 105</pre>
118.	<p><b>Table 3</b> reports the estimation process. We expect the last iteration to have very small numbers.</p> <p><b>red box:</b> for the largest raw-score difference, less than .5</p> <p><b>blue box:</b> for the biggest logit change, less than .01. We have these. Great!</p> <p>For big data sets, the maximum raw-score residual can be considerably larger without affecting the accuracy of the estimates.</p>	<pre>Table 3. Iteration Report.  +-----+-----+-----+-----+-----+   Iteration   Max. Score Residual   Max. Logit Change                 Elements  % Categories Elements Steps   +-----+-----+-----+-----+-----+   PROX      1             21.0444  17.5      24.7419   .7405       JMLE      2             -4.1066  -3.4      -.7658    -.0648  -2.2989     JMLE      3             -1.4643  -1.2      .8815     -.0244  -.0930     JMLE      4             -1.2319  -1.0      .3762     -.0198  -.0517     JMLE      5             -.8471   -.7       .3664     -.0134  -.0214     JMLE      6             -.6849   -.5       .2855     -.0102  -.0218     JMLE      7             -.5835   -.4       .2359     .0081   -.0171     JMLE      8             -.4905   -.3       .1932     .0067   -.0142     JMLE      9             -.4082   -.3       .1592     .0055   -.0117     JMLE     10             -4082    -3       .1592     .0055   -.0096   +-----+-----+-----+-----+-----+  If you want to stop the iterative process early, press you Ctrl+F keys together.</pre>
119.	<p>“Subset connection O.K.” so that the measures of all the elements belong to one cohesive structure. <i>We will discuss this in Tutorial 4.</i></p>	<p style="text-align: center;">Subset connection O.K.</p>
120.	<p><b>Table 4</b> appears after Table 8</p>	
121.	<p><b>Table 5</b> shows some global summary statistics. For each observation: <b>Cat</b> (category) is the observation <b>Score</b> is the category after it has been recounted <b>Exp.</b> is the expected value of the Score <b>Resd.</b> is the residual = Score - Expectation <b>StRes</b> is the standardized residual <b>Mean</b> (average) is the average for the observations. Count is the number of observations in the analysis. <b>S.D. (Population)</b> is the standard deviation if the elements are all possible elements for the facet <b>S.D. (Sample)</b> is the (larger) standard deviation if the elements are a sample of all possible elements (the population) for the facet.</p>	<pre>Table 5. Measurable Data Summary.  +-----+-----+-----+-----+-----+   Cat Score Exp. Resd StRes   +-----+-----+-----+-----+-----+   4.84 4.84 4.84 .00 .01   Mean (Count: 105)     1.88 1.88 1.18 1.44 1.00   S.D. (Population)     1.89 1.89 1.19 1.45 1.00   S.D. (Sample)   +-----+-----+-----+-----+-----+  red box: When estimation has been successful we expect the mean residual (Resd) and the mean standardized residual (StRes) to be 0.0, green box: and the S.D. of the standardized residual (StRes) to be 1.0.</pre>

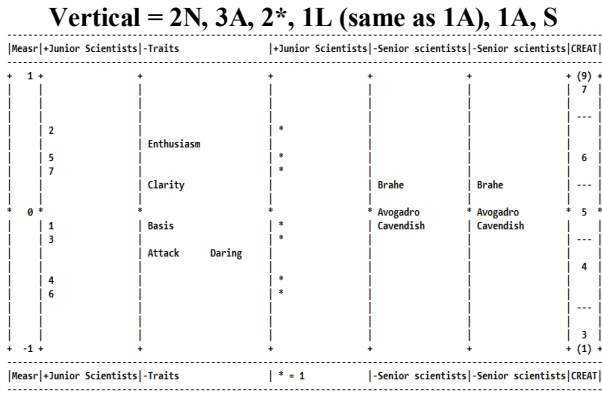
**122.** An approximate global fit statistic, a log-likelihood chi-square is shown. Its degrees of freedom, d.f., are roughly (number of responses - number of elements).  
The Rasch model is a model of perfection, so **we always expect to see significant misfit** to the model in empirical data, as we do here:  $p=.0000$

Data log-likelihood chi-square = 331.4227  
Approximate degrees of freedom = 85  
Chi-square significance prob. = .0000


**123.** *Red box:* Part of the variance in the data is explained by the Rasch measures, and, as the Rasch model predicts, part is unexplained. In these data, 41% is explained by the Rasch measures, a usual amount - even though it 41% looks low!

	Count	Mean	S.D.	Params
Responses used for estimation	= 105	4.84	1.88	20
Count of measurable responses	= 105.00			
Raw-score variance of observations	= 3.53	100.00%		
Variance explained by Rasch measures	= 1.45	41.02%		
Variance of residuals	= 2.08	58.98%		

**124.** **Table 6** shows the measures graphically. We can see that there is a noticeable spread among the Junior Scientists (examinees) and the Traits (items) which we want. There is also a smaller spread among the Senior Scientists (judges) which we don't usually want, but the Rasch measures have adjusted for.  
The rating scale, "CREAT" is shown to the right.  
*Which is the most lenient judge?* The column heading "-Senior Scientist" tells us. The most lenient judge will give the highest ratings. "-" means "high score implies low measure", so Cavendish is the most lenient judge.



**125.** Table 6.1 is a graphical representation of the measures we see in Table 7. It is useful when we need to picture the statistics for large samples. M represents the mean, S=1 standard deviation, and Q=2 standard deviations. The numbers represent elements. The numbers match Table 7.  
*Red box:* In the bottom distribution for a much larger dataset, there are 28 elements at "M", the mean. Read the numbers vertically.



**126.** **M. Table 7. Measure Tables**

**127.** **Table 7** shows the scores and measures. Measures are often reported in logits or other units unfamiliar to our audience. They often ask, "but what do they mean in terms of the scores I'm familiar with?"  
**Green box:** This is what the "Fair Average" does. It takes the measures and shows what they imply as ratings for a standard person rated by standard judge on a standard item. "Standard" means an imaginary element with the average measure of the elements of the facet. In this example, the data are complete, so the Observed Average rating is close to the Fair Average rating. But when there are missing data, the Fair Average adjusts for the missing data but the Observed Average does not.

Obsvd Score	Obsvd Count	Obsvd Average	Fair-M Average	Model Measure	Model S.E.
156	35	4.5	4.39	.24	.12
171	35	4.9	4.86	.04	.11
181	35	5.2	5.17	-.09	.11
169.3	35.0	4.8	4.81	.06	.12
10.3	.0	.3	.32	.13	.00
12.6	.0	.4	.39	.16	.00

**128.** In a practical assessment situation, different people may be administered different tasks and rated on different items by different judges. You encounter the difficult tasks and the severe judges. I encounter the easy tasks and the lenient judges. *Fine!* Your ability measure and my ability measure adjust for this.

But then the examination authorities say “Rasch measures are great, but when we publish the results, we want them expressed as ratings on the original rating scale!”

So we have to go from Rasch measures back to the rating scale in a way that is fair - as though you and I encountered the same judges and performed the same tasks. *Facets* does this for us by computing the ratings we would have received (according to the Rasch measures) if you and I had both performed a task of average difficulty and we were both rated by judges of average severity. This gives a “Fair Average” rating.

**129.** **Red box:** In Table 7, the “Model S.E.” is the **precision** of the measure. This indicates how fuzzy is the location of the element measure on the latent variable.

In everyday speech, the words “**precision**” and “**accuracy**” often mean the same thing, but for us they are different.

Imagine arrows being shot at a target. If the arrows form a close group, then the archery is precise. If the arrows are in the neighborhood of the center of the target, the archery is accurate. When the arrows all hit the bull’s eye, the archery is accurate and precise.

						Precision	Accuracy			
Obsvd	Obsvd	Obsvd	Fair-M	Model	S.E.	Infit	Outfit			
Score	Count	Average	Average	Measure		MnSq	ZStd	MnSq	ZStd	
131	28	4.7	4.60	.07	.14	1.35	1.3	1.34	1.2	
148	28	5.3	5.28	-.23	.13	.95	-.1	.98	.0	
150	28	5.4	5.36	-.27	.13	.62	-1.6	.60	-1.7	

**Measurement Precision:** how exact is the location on the latent variable?

**Measurement Accuracy:** is it the correct location?

**Estimation Precision** (decimal places): how closely does our estimate match the estimation criteria?

Statistics are often reported with 6 decimal places (high estimation precision) even though they are reporting only a few data points (low substantive precision).

**130.** **Precision** means “how reproducible is the location of the measure on the latent variable with data like these”. It is like the gradations on measurement scale. It is internal to the measuring system, and is quantified in the standard error of measurement, S.E.

The more observations of an element, the more precise will be the estimate. As carpenters say, “*Measure twice, cut once!*”

**Accuracy** means “how well does the measure correspond to an external standard”. In our case, the external standard is the Rasch-model ideal of invariant measure additivity.

If the data fit the Rasch model, then the parameter estimates accurately reflect the ideal additive measurement framework. For us, accuracy is quantified in the quality-control fit statistics, Infit and Outfit.

**131.** We can obtain **higher precision** for an element’s measure by:

- 1. More observations** of the element, e.g., a person takes a longer test or is rated by more judges.
- 2. Better targeting** of the element, e.g., a person takes a test that is not too easy or too high.
- 3. More categories** in the rating scale, e.g., a 5-category rating scale instead of a 3-category scale, but beware of over-categorization .... which we will soon meet!

132 **N. Table 7: Fit Statistics**

133 **Green box:** Guilford.out.txt Table 7.1.1 shows the measures for the judges.  
**Blue box:** Brahe (lowest Total Score) is the slightly most severe judge.  
**Red box:** But, more importantly, look at the fit statistics. Brahe is the most misfitting (Mean-squares > 1.0). The other two judges have about the same fit.  
**Orange box:** The average of the mean-squares is usually near 1.0, so a misfitting judge, like Brahe, forces the other judges, Avogadro and Cavendish to be reported as overfitting.

Total Score	Total Count	Model Measure	Infit S.E.	Outfit MnSq	ZStd	N Senior scientists
156	35	.24	.12	1.42	1.7	2 Brahe
171	35	.04	.11	.84	-1.6	1 Avogadro
181	35	-.09	.11	.66	-1.6	3 Cavendish
169.3	35.0	.06	.12	.97	-2	Mean (Count: 3)
10.3	.0	.13	.00	.33	1.4	S.D. (Population)
12.6	.0	.16	.00	.40	1.7	S.D. (Sample)

**Always investigate underfit (high mean-squares) before overfit (low mean-squares).** Often the overfit disappears when the underfit is eliminated from the data.

134 Notice also that the Infit and Outfit columns are similar. This is usual with long rating scales (9 categories here) so that the operational range of each item is very wide.  
 Under these circumstances, my choice is only to report **Outfit**, because it is the conventional statistical chi-square (divided by its d.f.) which is familiar to most statisticians, but please do report both if your audience expects to see them.  
 Polytomous mean-square statistics have the same characteristics as dichotomous ones, **#Error!** Reference source not found., but are much harder to diagnose by eye.

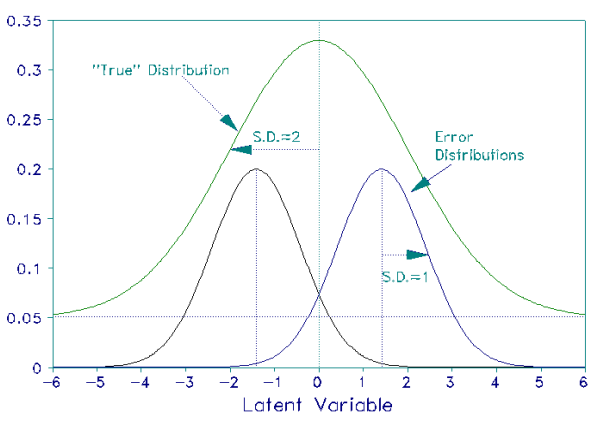
Polytomous Mean-square Fit Statistics

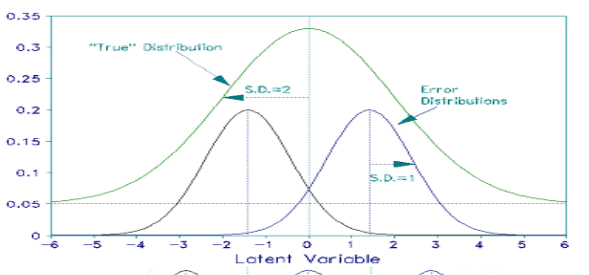
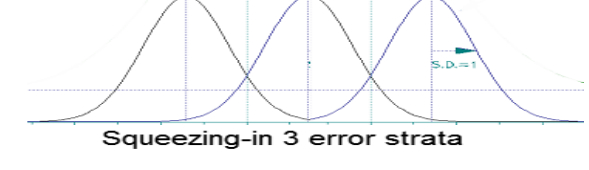
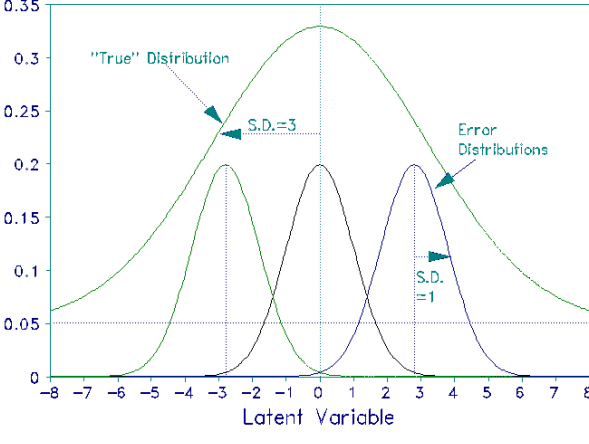
Response String Easy.....Hard	INFIT MnSq	OUTFIT MnSq	RPM Corr.	Diagnosis
<b>I. modelled:</b>				
33333132210000001011	.98	.99	.78	<i>Stochastically monotonic in form, strictly monotonic in meaning</i>
31332332321220000000	.98	1.04	.81	
33333331122300000000	1.06	.97	.87	
33333331110010200001	1.03	1.00	.81	
<b>II. overfitting (muted):</b>				
33222222221111111100	.18	.22	.92	Guttman pattern high discrimination low discrimination tight progression
333332222211111100000	.31	.35	.97	
32222222221111111110	.21	.26	.89	
32323232121212101010	.52	.54	.82	
<b>III. limited categories:</b>				
33333333332222222222	.24	.24	.87	high (low) categories central categories only 3 categories
22222222221111111111	.24	.34	.87	
33333322222222221111	.16	.20	.93	
<b>IV. informative-noisy:</b>				
3222222201111111130	.94	1.22	.55	noisy outliers erratic transitions noisy progression extreme categories
33233332212333000000	1.25	1.09	.77	
33133330232300101000	1.49	1.40	.72	
33333333330000000000	1.37	1.20	.87	
<b>V. non-informative:</b>				
22222222222222222222	.85	1.21	.00	one category central flip-flop rotate categories extreme flip-flop random responses
12121212121212121212	1.50	1.96	-.09	
01230123012301230123	3.62	4.61	-.19	
03030303030303030303	5.14	6.07	-.09	
03202002101113311002	2.99	3.59	-.01	
<b>VI. contradictory:</b>				
11111122233222111111	1.75	2.02	.00	folded pattern central reversal high reversal Guttman reversal extreme reversal
11111111112222222222	2.56	3.20	-.87	
22222222233333333333	2.11	4.13	-.87	
00111111112222222233	4.00	5.58	-.92	
00000000033333333333	8.30	9.79	-.87	

135	O. Table 7: Inter-rater Statistics															
136	<p><b>Inter-rater</b>= has instructed <i>Facets</i> to compute some rater agreement statistics.</p> <p><b>Green box:</b> The “Exact Agreement Observed %” report what percent of the ratings by this rater agree exactly with the ratings made by another rater. the “Exact Agreement Expected %” reports the agreement that would be seen if the data fit the Rasch model perfectly.</p>	<table border="1"> <thead> <tr> <th colspan="2">Exact Agree.</th> <th rowspan="2">N Senior sci</th> </tr> <tr> <th>Obs %</th> <th>Exp %</th> </tr> </thead> <tbody> <tr> <td>21.4</td> <td>25.2</td> <td>2 Brahe</td> </tr> <tr> <td>35.7</td> <td>25.8</td> <td>1 Avogadro</td> </tr> <tr> <td>37.1</td> <td>25.3</td> <td>3 Cavendish</td> </tr> </tbody> </table>	Exact Agree.		N Senior sci	Obs %	Exp %	21.4	25.2	2 Brahe	35.7	25.8	1 Avogadro	37.1	25.3	3 Cavendish
Exact Agree.		N Senior sci														
Obs %	Exp %															
21.4	25.2	2 Brahe														
35.7	25.8	1 Avogadro														
37.1	25.3	3 Cavendish														
137	<p>For Brahe the observed agreement is 21.4%. Is this good or bad? We would tend to expect a much higher agreement. But <i>Facets</i> provides a reference point. It reports that for these raters, examinees and items, the “Exact Agreement Expected %” for Brahe is 25.2%. Usually the observed agreement is <i>slightly higher</i> than the expected agreement, because most raters try to be “agreeable” with each other. Look at Avogadro and Cavendish, their observed agreement %’s (35.7%, 37.1%) are much higher than expected (25.8%, 25.3%). They are agreeing together against Brahe.</p>															
138	<p>Under Table 7.1.1, the agreement statistics are summarized. In these data, the observed “exact agreement” is 31.4%, but the expected agreement is 25.4%. <i>The judges are agreeing too well!!</i> <b>Something is wrong!</b></p>	<p>Inter-Rater agreement opportunities: 105 Exact agreements: 33 = 31.4% Expected: 26.7 = 25.4%</p>														
139	<p><i>Facets</i> models the raters to be “<b>independent experts</b>”. These would produce an “exact agreement” percent, which is the same or slightly higher than the “expected agreement” percent. But many raters are trained to behave like “<b>rating machines</b>”. Agreement is encouraged among the raters, and disagreements are penalized. For these raters we expect the “exact agreement” percent to be <i>much higher</i> than the “expected agreement” percent. When the “exact agreement” approaches 100%, the raters are behaving the same way as optical scanners do for “bubble sheets”. The raters have become part of the data-collection mechanism, they are no longer a facet of the measurement situation.</p>															

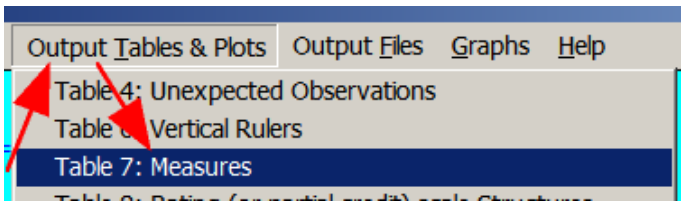
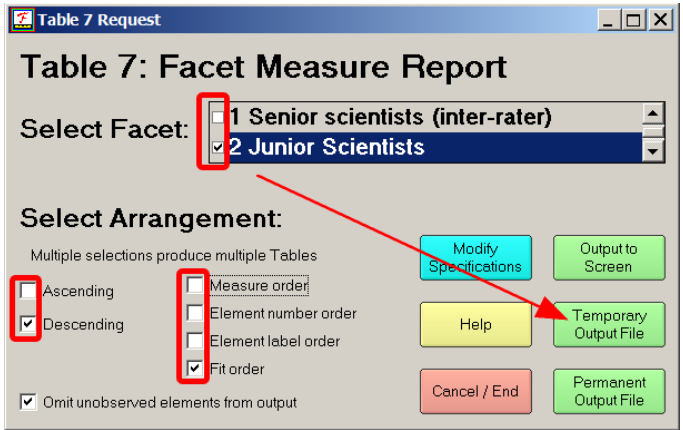
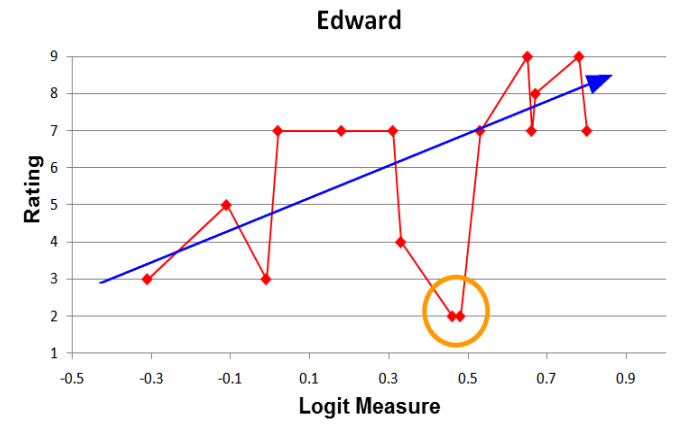


140	<b>P. Table 7: Reliabilities and Separations</b>																																																							
141	Model, Populn: RMSE .12 Adj (True) S.D. .07 Separation .60 Strata 1.13 Reliability (not inter-rater) .26 Model, Sample: RMSE .12 Adj (True) S.D. .12 Separation 1.02 Strata 1.69 Reliability (not inter-rater) .51																																																							
142	<p>Under each Table 7 is a set of reliability statistics. These show the reliability of the differences between the measures in the facet. They indicate the <b>reproducibility</b> of the measures, <b>not the accuracy</b> of the measures. These reliabilities are not <i>inter-rater reliability</i> statistics (which show the rater similarity).  <b>“Reproducible”</b> - we can expect the same number if we repeated the same data collection. A stopped clock is highly reproducible, so it is highly <b>reliable</b>. <i>Of course, it is reliably wrong!</i>  <b>“Accuracy”</b> - the current number is near the “true” number.</p>																																																							
143	<p><b>“Model”</b> means “assuming all misfit in the data is due to the randomness predicted by the Rasch model”  <b>“Real”</b> (when shown) means “assuming all misfit in the data contradicts the Rasch model”  <b>“Population”</b> means “assuming this set of elements is the entire population.”  <b>“Sample”</b> means “assuming this set of elements is a random sample from the population of interest”  <b>“RMSE”</b> means “root mean-square error”, a statistical average of the standard errors of the measures.  <b>“Adj (True) SD”</b> means “the standard deviation of the measures, (Adj=) adjusted for measurement error”, also called the <i>“True” standard deviation</i>.  <b>“Separation”</b> is the True SD / RMSE. It indicates how many measurement strata could be statistically distinguishable among the measures, if the tails of the measure distribution are conceptualized to be caused <i>by outlying random noise</i>.  <b>“Reliability”</b> is the ratio of the “True” variance of the measures to the observed variance.  <b>“Strata”</b> is <math>(4 * \text{Separation} + 1) / 3</math>. It indicates how many measurement strata could be statistically distinguishable among the measures, if the tails of the measure distribution are conceptualized to be caused <i>by outlying “true” measures</i>.</p>																																																							
144	<p>This table shows the relationship between measurement variance, measurement error, and reliability.</p> <table border="1" data-bbox="175 1161 1513 1470"> <thead> <tr> <th>Error RMSE</th> <th>True SD</th> <th>True Variance = True SD<sup>2</sup></th> <th>Observed Variance = RMSE<sup>2</sup> + True Variance</th> <th>Signal-to-Noise Ratio</th> <th>Separation = True SD / RMSE</th> <th>Reliability = True Variance / Observed Variance</th> <th>Strata = <math>(4 * \text{Separation} + 1) / 3</math></th> </tr> </thead> <tbody> <tr> <td>1</td> <td>0</td> <td>0</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>0.3</td> </tr> <tr> <td>1</td> <td>1</td> <td>1</td> <td>2</td> <td>1</td> <td>1</td> <td>0.5</td> <td>1.7</td> </tr> <tr> <td>1</td> <td>2</td> <td>4</td> <td>5</td> <td>2</td> <td>2</td> <td>0.8</td> <td>3</td> </tr> <tr> <td>1</td> <td>3</td> <td>9</td> <td>10</td> <td>3</td> <td>3</td> <td>0.9</td> <td>4.3</td> </tr> <tr> <td>1</td> <td>4</td> <td>16</td> <td>17</td> <td>4</td> <td>4</td> <td>0.94</td> <td>5.7</td> </tr> </tbody> </table>								Error RMSE	True SD	True Variance = True SD <sup>2</sup>	Observed Variance = RMSE <sup>2</sup> + True Variance	Signal-to-Noise Ratio	Separation = True SD / RMSE	Reliability = True Variance / Observed Variance	Strata = $(4 * \text{Separation} + 1) / 3$	1	0	0	1	0	0	0	0.3	1	1	1	2	1	1	0.5	1.7	1	2	4	5	2	2	0.8	3	1	3	9	10	3	3	0.9	4.3	1	4	16	17	4	4	0.94	5.7
Error RMSE	True SD	True Variance = True SD <sup>2</sup>	Observed Variance = RMSE <sup>2</sup> + True Variance	Signal-to-Noise Ratio	Separation = True SD / RMSE	Reliability = True Variance / Observed Variance	Strata = $(4 * \text{Separation} + 1) / 3$																																																	
1	0	0	1	0	0	0	0.3																																																	
1	1	1	2	1	1	0.5	1.7																																																	
1	2	4	5	2	2	0.8	3																																																	
1	3	9	10	3	3	0.9	4.3																																																	
1	4	16	17	4	4	0.94	5.7																																																	

<p><b>145</b> What does this table of separations and reliabilities mean? <b>Here is a picture of Separation = 2.</b>  <b>Green curve:</b> The larger curve is the conceptual “true” distribution.  <b>Black and blue curves:</b> The smaller curves are the error distributions for individual measures.  The x-axis on the graph locates the person measures on the latent variable.  The y-axis is on the graph is the local density, i.e., what proportion of the sample we expect at the x-axis location (larger curve) or what proportion of observations of a measure with error we expect at an x-axis value (smaller curves).</p>	 <p style="text-align: center;"><b>Separation = 2 implies 2 error strata</b></p>
<p><b>146</b> So the question becomes “<i>How many numerically different measures can we reasonably discriminate within the true distribution?</i>”  For Separation = 2, we can see that two measures at about -1.5 and +1.5, together with their error distributions (the fuzziness of the measurement), matches the “true” distribution of the measures. The splatter around only two measures (shown by the two smaller curves) covers the whole reasonable range of the upper curve.</p>	
<p><b>147</b> <b>A physical analogy:</b> Imagine I have a classroom of children and must report their heights. The report is overdue, so I measure their heights quickly by eye. This will yield imprecise measurements with high uncertainty. The measures of height would have big standard errors.</p> <p><i>Question 1. How far apart must two heights be for me to be reasonably sure that the children's heights are different?</i>  Answer: Roughly three standard errors. we are comparing two measures and both have standard errors. So, assuming the standard errors are approximately the same, the statistical value is (two error distributions)*(p&lt;.05)*(RMSE) = <math>\sqrt{2} * 1.96 * SE \approx 3 * S.E.</math></p> <p><i>Question 2. In the observed distribution of heights, how many "reasonably sure" height-difference strata are there, assuming there are no unusually short or tall children?</i>  Answer: This is the "Separation", which is (True S.D. / Standard error).</p> <p>Practical example:  My children have only two "true" heights: half are 1.5 meters high, and half are 1.7 meters high, so the true S.D. of their height measures is .1 meters.  But when I record their heights by eye, the 1.5-meter-children have a range from 1.4 meters to 1.6 meters, and the 1.7 meter children have a range from 1.6 meters to 1.8 meters. So the standard error of my height measurements is about .05 meters  The "separation" of my children is (true S.D./S.E.) = (.1 / .05) = 2. Two strata - exactly right.  If my standard error had been larger, the two height ranges would have overlapped and the separation would have dropped. I would not have been able to distinguish tall children from short children.  But if my standard error had been smaller, there would have been a gap between the two error distributions. The separation would have been bigger, perhaps 3 strata, alerting me that I could have distinguished a third strata of children of height 1.5 meters, if there had been any.</p> <p>It is the same with Rasch measures and their standard errors.</p>	

<p><b>148</b></p>	<p>For Separation = 2, notice that the peaks of the S.E. curves are 3 units apart.</p> <p><b>Strata:</b> If we needed to discriminate 3 “strata”, we could squeeze them in. The very top (with curve peak at 3), against the middle (with curve peak at 0), against the very bottom (with curve peak at -3).</p> <p>So, a very high performer can be discriminated from a middle performer, can be discriminated from a very low performer. <i>But we would need to look at the empirical distribution of measures (in the “rulers” in Table 6) to see if the distribution does have long tails where those very high and very low performers would be located.</i></p>	<p><b>Separation = 2 = 2 error distributions</b></p>  <p><b>Squeezing-in 3 error strata</b></p>  <p><b>Strata = 3</b></p>
<p><b>149</b></p>	<p>Here’s the same thing for <b>reliability = 0.9</b>, separation = 3. We can see how the narrower error distributions allow for more different measures to be squeezed into the “true” distribution.</p> <p>Separation (True SD / Error SD) is more useful than reliability when reliabilities get much above 0.9. The maximum reliability is 1.0 so changes in reliability are not noticeable. Changes in the equivalent separation are always identifiable.</p> <p><i>Strata:</i> we could squeeze in another S.E. distribution curve, by placing the peaks at -4.5, -1.5, +1.5, -4.5</p>	 <p><b>Separation = 3 implies 3 error strata</b></p>
<p><b>150</b></p>	<p>Decision-makers say, “<i>We are going to use this instrument to discriminate x levels of performance.</i>” We might respond, “<i>Fine! based on the sample measure distribution, and the separation pictures, this instrument can do that.</i>” Or we might say, “<i>This test really does not have the discriminating power for x many levels of performance, so there will be a lot of mis-classification.</i>”</p> <p>In language testing it used to be common (maybe it still is) to try to discriminate 10 or so performance-levels based on a test that only has the statistical power to discriminate 3 or so levels. No one computed the standard errors for individual measures, so no one knew how arbitrary the classification of examinees was.</p>	
<p><b>151</b></p>	<p><b>How much reliability or separation do we need?</b> It depends what our purpose is, but we nearly always want to separate high performers from low performers, so a person separation of 2, reliability of 0.8, is often the benchmark for practical use. But reliabilities are always computed for the current sample, despite the convention of calling them the “test reliabilities”. Next time the sample will be differently distributed, so the separation may be different .....</p>	

152	<b>Measure summary chi-square statistics</b>
153	<p>There are two questions we may ask ourselves about the elements of a facet:</p> <ol style="list-style-type: none"> <li>1. <i>Are the measures of the elements in a facet all statistically the same, except for measurement error?</i> This particularly applies to raters. We want them to have the same leniency. This hypothesis is tested with the <b>“fixed (all same) chi-square.”</b></li> <li>2. <i>Are the measures a random sample from a normal distribution?</i> This particularly applies to large samples of persons. If they are, we can conveniently summarize them with a mean measure and a standard deviation. This hypothesis is tested with the <b>“Random (normal) chi-square”</b>.</li> </ol> <p><b>Model, Fixed (all same) chi-square: 39.1 d.f.: 6 significance (probability): .00</b>  <b>Model, Random (normal) chi-square: 5.2 d.f.: 5 significance (probability): .39</b></p> <p>In this example, the hypothesis that the elements have the same measure, apart from measurement error, has significance <math>p=.00</math>, so this hypothesis is <i>rejected</i>.  The hypothesis that the measures are a random sample from a normal distribution has significance <math>p=.39</math>, so this hypothesis is <i>not rejected</i>.</p>

<p>154</p>	<p style="text-align: center;"><b>Rasch Fit Statistics: Are the Measures Accurate and Effective?</b></p>																																																																																		
<p>155</p>	<p>Let's produce Table 7 for "Junior Scientists", facet 2 in fit order, ascending:</p> <p>Facets analysis window for the Guilford.txt Click on "Output Tables &amp; Plots" menu Click on "Table 7: Measures"</p>																																																																																		
<p>156</p>	<p>"Table 7 Request" dialog box: Click on "All" to uncheck-mark it. Scroll the list ... Click on "2 Junior Scientists" to check-mark it Click on "Measure" to uncheck-mark it Click on "Fit order" check-mark it Click on "Temporary Output File"</p>	<table border="1"> <thead> <tr> <th>Infit</th> <th colspan="2">Outfit</th> <th colspan="2"> Estim.</th> <th colspan="2"> Correlation</th> <th colspan="2"></th> </tr> <tr> <th>MnSq</th> <th>ZStd</th> <th>MnSq</th> <th>ZStd</th> <th>Discrm</th> <th>PtMea</th> <th>PtExp</th> <th>N</th> <th>Junior</th> </tr> </thead> <tbody> <tr> <td>1.94</td> <td>2.2</td> <td>1.94</td> <td>2.2</td> <td>.34</td> <td>.51</td> <td>.48</td> <td>5</td> <td>Edward</td> </tr> <tr> <td>1.31</td> <td>.9</td> <td>1.37</td> <td>1.0</td> <td>.87</td> <td>.20</td> <td>.44</td> <td>4</td> <td>David</td> </tr> <tr> <td>1.13</td> <td>.4</td> <td>1.22</td> <td>.7</td> <td>.84</td> <td>.16</td> <td>.46</td> <td>3</td> <td>Chris</td> </tr> <tr> <td>.85</td> <td>-.3</td> <td>.84</td> <td>-.4</td> <td>1.37</td> <td>.30</td> <td>.48</td> <td>7</td> <td>George</td> </tr> <tr> <td>.70</td> <td>-.8</td> <td>.77</td> <td>-.5</td> <td>.93</td> <td>.40</td> <td>.43</td> <td>6</td> <td>Fred</td> </tr> <tr> <td>.61</td> <td>-1.1</td> <td>.60</td> <td>-1.2</td> <td>1.30</td> <td>.85</td> <td>.47</td> <td>2</td> <td>Betty</td> </tr> <tr> <td>.24</td> <td>-3.2</td> <td>.23</td> <td>-3.2</td> <td>1.48</td> <td>.81</td> <td>.47</td> <td>1</td> <td>Anne</td> </tr> </tbody> </table> <p><b>Always investigate underfit (high mean-squares) before overfit (low mean-squares).</b> Often the overfit disappears when the underfit is eliminated from the data.</p>	Infit	Outfit		Estim.		Correlation				MnSq	ZStd	MnSq	ZStd	Discrm	PtMea	PtExp	N	Junior	1.94	2.2	1.94	2.2	.34	.51	.48	5	Edward	1.31	.9	1.37	1.0	.87	.20	.44	4	David	1.13	.4	1.22	.7	.84	.16	.46	3	Chris	.85	-.3	.84	-.4	1.37	.30	.48	7	George	.70	-.8	.77	-.5	.93	.40	.43	6	Fred	.61	-1.1	.60	-1.2	1.30	.85	.47	2	Betty	.24	-3.2	.23	-3.2	1.48	.81	.47	1	Anne
Infit	Outfit		Estim.		Correlation																																																																														
MnSq	ZStd	MnSq	ZStd	Discrm	PtMea	PtExp	N	Junior																																																																											
1.94	2.2	1.94	2.2	.34	.51	.48	5	Edward																																																																											
1.31	.9	1.37	1.0	.87	.20	.44	4	David																																																																											
1.13	.4	1.22	.7	.84	.16	.46	3	Chris																																																																											
.85	-.3	.84	-.4	1.37	.30	.48	7	George																																																																											
.70	-.8	.77	-.5	.93	.40	.43	6	Fred																																																																											
.61	-1.1	.60	-1.2	1.30	.85	.47	2	Betty																																																																											
.24	-3.2	.23	-3.2	1.48	.81	.47	1	Anne																																																																											
<p>157</p>	<p>Here is the Table in a NotePad file. Let us think about what this means ...</p> <p><b>Orange box:</b> Edward has mean-squares of 1.94, much larger than the expected 1.0. The ratings of Edward underfit the Rasch model. They are <b>too unpredictable from the Rasch measures or "noisy"</b>.</p> <p><b>Blue circle:</b> Edward has a high correlation - this usually means "predictable" - why?</p> <p><b>Red box:</b> Anne has mean-squares of .24 and .23, much lower than the expected 1.0. The ratings of Anne <b>overfit</b> the Rasch model. They are <b>too predictable from the Rasch measures or "muted"</b>.</p>																																																																																		
<p>158</p>	<p><b>Blue circle:</b> our investigation! Hre is a plot of Edward's ratings and the logit measures that are modeled to produce them. If you have some skill with Excel, Appendix 3. Excel plots from the Residual file explains how to make this plots for yourself. <b>Blue line:</b> a strong trend = high correlation. <b>Orange circle:</b> two observations of "2" are surprising. Overall, Edward's the ratings are much less predictable (from the Rasch measures) than the Rasch model expects.</p>																																																																																		

**159 Edward:**

**1. “Outfit MnSq = 1.94”, “Infit MnSq=1.94”.**

The “Outfit mean-square” reports primarily about observations where the combined (summed) measures are far from zero.

The “Infit mean-square” reports primarily about patterns of observations where the combined (summed) measures are near to zero.

In Guilford.txt the rating scale is so long (9 categories) that the operational range of the rating scale for each item is much wider than the spread of the measures. Accordingly Outfit and Infit report essentially the same results. I prefer to report only Outfit, but some reviewers prefer Infit or both Infit and Outfit.

**2. “MnSq = 1.94”**

The mean-square is much greater than 1.0, so these ratings are too unpredictable. They underfit the Rasch model. They twice as much randomness as the model predicts.

Ben Wright explained fit like an old phonograph record.

When the mean-square is close to 1.0, the music can be heard clearly.

When the mean-square is much less than 1.0, the music is muted, muffled. It loses its rich tones.

When the mean-square is much greater than 1.0, the music is there, but so are the pops, rumbles due to scratches and surface noise. When the mean-square is above 2.0, the noise is starting to overwhelm the music.

From the plot, we can that the data do not concur about Edward’s performance. The ratings in the orange circle say that Edward is a low performer, but other ratings say that he is a high performer. Whichever is correct, the estimated measure is a compromise, so it is an inaccurate estimate of Edward’s “true” measure.

**3. “Zstd = 2.2”** in #157 - this is reporting the result of a statistical hypothesis test: *“These ratings conform to the Rasch model.”*

**4. “Zstd = +...”** - indicates that the ratings underfit (too much noise) the Rasch model

**5. “= +2.2”** - this value is a unit-normal deviate indicating the probability that these ratings conform to the Rasch model. It is unlikely ( $p < .05$  in Appendix 1) that these ratings are the chance outcomes of a Rasch process based on the estimated measures.

**6. “Do these data fit the Rasch model or not?”** - the hypothesis test of fit to the Rasch model reports “They do not!”. They underfit the model: the mean-squares says the misfit is big, and the Zstd says that the misfit is unlikely to have happened by chance.

**7. “What action do we take?”** This depends on the circumstances.

A. The data aren’t perfect - but we expected that.


B. These data underfit the model. They are too unpredictable. Is that a cause of concern for us? Yes, the measure of Edward’s performance (based on these data) is inaccurate for practical purposes.

C. If this is our first look at the data, always examine high mean-squares (underfit) before low mean-squares (overfit). This is because the average mean-square is usually forced to be close to 1.0. So investigate Edward (MnSq = 1.94) before Anne (MnSq=0.24).

D. If we consider that the ratings in the orange circle are not representative of Edward’s general performance, we might omit Edward from the analysis, or omit those ratings. Then Edward’s idiosyncrasies won’t impact other aspects of the analysis, such as Anne’s mean-square. A later tutorial will show us how we can anchor (fix) the other measures at their good values, and measure Edward with all his ratings.

E. If this is a diagnostic test, then the orange-circled ratings may be the most important ones. They tell us where to focus our remedial action for Edward to improve his performance.





**161 Anne:**

**2. “MnSq = 0.24”**  
 The mean-square is much less than 1.0, so these ratings are too predictable. They overfit the Rasch model. They only contain one-quarter of the randomness that the model predicts. In Classical Test Theory (CTT) this would be considered good. In Rasch this indicates that these ratings contain only 24% of the measurement information that they should. These ratings are inefficient and will cause the reported standard errors to be too small and the reported reliabilities to be too high. But the measure of Anne’s performance (based on these data) is accurate.

**3. “Zstd = -3.2”** - this is reporting the result of a statistical hypothesis test: “These ratings conform to the Rasch model.”

**4. “Zstd = -”** - indicates that the ratings overfit.

**5. “= -3.2”** - this value is a unit-normal deviate indicating the probability that these ratings conform to the Rasch model. It is extremely unlikely ( $p < .01$  in Appendix 1) that these ratings are the chance outcomes of a Rasch process based on the estimated measures.

**6. “Do these data fit the Rasch model or not?”** - the hypothesis test of fit to the Rasch model reports “They do not!”. They overfit the model, highly statistically significantly.

**7. “What action do we take?”** This depends on the circumstances.

A. The data aren’t perfect - but we expected that.

B. These data overfit the model. They are too predictable. Is that a cause of concern for us? Yes, if it is a roulette wheel. But usually No if it is the performance of a child on an educational test. The measure of Anne’s performance (based on these data) is **accurate for practical purposes**.

C. If this is a standard testing situation, then overfit slightly stretches the measures (increases their range), inflates their reliability and reduces their standard errors. These are technical issues usually not of concern to anyone other than psychometricians. So it would require very strong external motivation to omit or alter this set of ratings.

D. If this a rater training situation, low mean-squares are typical of raters “playing it safe” by exhibiting central tendency or trying to agree with the ratings they think the other raters will give. This is often the result of training which emphasizes “if you disagree with the other raters too much, you will be fired!” So, before being concerned about the individual, review the training material and the instructions given to the raters. Are they explicitly or implicitly being told to agree with each other?

*At the Olympic Ice-Skating, the organizers think that “rater agreement = more credibility”, but to psychometricians, “excessive rater agreement = psychological pressure to agree = loss of objectivity and fairness”.*

162 **Q. Table 8: Rating Scale Structure**

163 Table 8 tells us about the 9-category rating scale of Creativity. It is packed with useful information about the success of our data collection - information which J.P. Guilford completely overlooked when he wrote the chapter on rating scales in his book, *“Psychometric Methods”*.

164

Model = ?B,?B,?,CREATIVITY  
Rating (or partial credit) scale = CREATIVITY,R9,G,0

DATA				QUALITY CONTROL			STEP		EXPECTATION		MOST	.5 Cumul.	Cat	Response
Category	Counts	Cum.		Avge	Exp.	OUTFIT	CALIBRATIONS	Measure	at	PROBABLE	Probabil.	PEAK	Category	
Score	Used	%	%	Meas	Meas	MnSq	Measure	S.E.	Category	-0.5	from	at	Prob	Name
1	4	4%	4%	-.86	-.74	.8			(-2.70)		low	low	100%	lowest
2	4	4%	8%	-.11	-.58	2.7	-.64	.53	-1.65	-2.21		-1.75	17%	
3	25	24%	31%	-.36*	-.40	.9	-2.32	.39	-.93	-1.26	-1.48	-1.39	48%	
4	8	8%	39%	-.43*	-.22	.5	.83	.25	-.41	-.66		-.46	11%	
5	31	30%	69%	-.04	-.03	.8	-1.48	.24	.02	-.19	-.32	-.29	39%	middle
6	6	6%	74%	-.46*	-.17	4.1	1.71	.25	.44	-.23		.34	9%	
7	21	20%	94%	.45	.34	.6	-1.00	.26	.94	-.68	.35	.47	47%	
8	3	3%	97%	.75	.50	.5	2.36	.44	1.62	1.24		1.37	16%	
9	3	3%	100%	.77	.62	.8	.54	.60	( 2.69)	2.17	1.45	1.70	100%	highest
									(Mean)			(Modal)		
												(Median)		

165 **Blue box:** There are 9 categories, 1 to 9. Look at how they have been “Used”, the category frequency counts. Do you notice anything conspicuous? *Yes you do!* Only categories 3, 5 and 7 have large counts. Perhaps the judges, the Senior Scientists, could only discriminate 3 levels of Creativity, but were told to use a 9-category scale. **Over-categorization** leads to artificially reduced standard errors, inflated reliabilities and poor fit to the Rasch model.

166 **Red box:** We can see evidence of poor fit in the “Average Measure” column. The rating scale is intended to represent a series of qualitative advances along the latent variable. Each category is assumed to be a quantitative advance (of a size yet to be determined) beyond the previous category. So, **higher categories should imply higher measures, and higher measures should be observed as higher categories. But did this happen?**


167 The “Average Measures” are the averages of the measures that combined to produce the observations in the category. We expect them to advance with category number.

*Our estimation process in Tutorial 1:*  
 $B_n - D_i - R_r - S_s - \{F_{kj}\} \rightarrow X_{nirs}$   
 Average Measure for category “j” =  
 Average  $(B_n - D_i - R_r - S_s)$  for all  $X_{nirs} = j$


168 In Table 8, we can see that the Average Measure for category 1 is **-.86**. Then for category 2 it is **-.11**. *Good so far*, categories and average measures are advancing together. But the Average Measure for Category 3 is **-.36**. **The Average Measure has gone backwards, and so is flagged with “\*”**. *This contradicts our theory about the rating scale.*  
**Green box:** the “Expected Measure” column shows what the Average Measures would be if the data fit the Rasch model. We can see big differences, particularly for Category 6 (-.46 vs. .17). Something is seriously wrong. *What is it?*

169 **Orange box:** Look at the category-level “Outfit MnSq” column. We expect these to be 1.0 or less due to dependency among the categories of the rating scale. Most mean-squares are in this range. But Category 2, mean-square 2.7, and Category 6, mean-square 4.1, are showing considerably unpredictability. *Another symptom that something has gone seriously wrong with the functioning of the rating scale!*

**170** Let's look at the rating scale from a graphical perspective:  
 On the *Facets* menu bar,  
 Click "Graphs"



**171** The Rasch model has made as much sense of the rating scale as it can.  
 The x-axis is the latent variable, drawn relative to the difficulty of the item.  
 The y-axis is the probability of observing each category of the 9-category rating scale.  
 Categories 3, 5, 7 are observed more often, so they have higher probability curves. Categories 1 and 9 are the extreme categories, so the Rasch model extrapolates that they are the categories most probable to be observed outside of the range of the data. The rest of the categories have low probability of being observed.




If you want to know which category a curve represents, click on the curve.  
 (Click on the plot background if the plot is not redrawn correctly.)

**172** On the "Graphs" window,  
 Click on "Prob+Empirical Cat. Curves"

Prob = Probability (as predicted by the Rasch model)

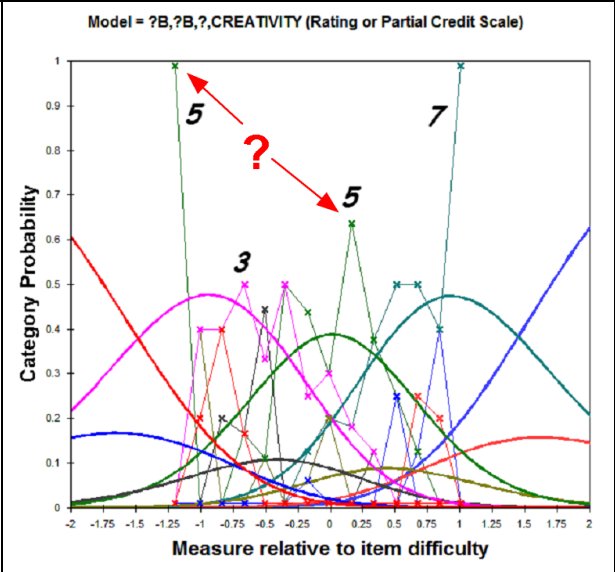
Empirical = Observed (as summarized from the dataset)


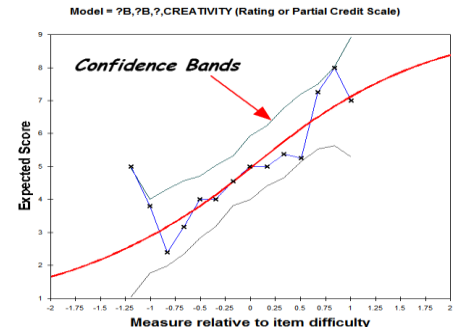
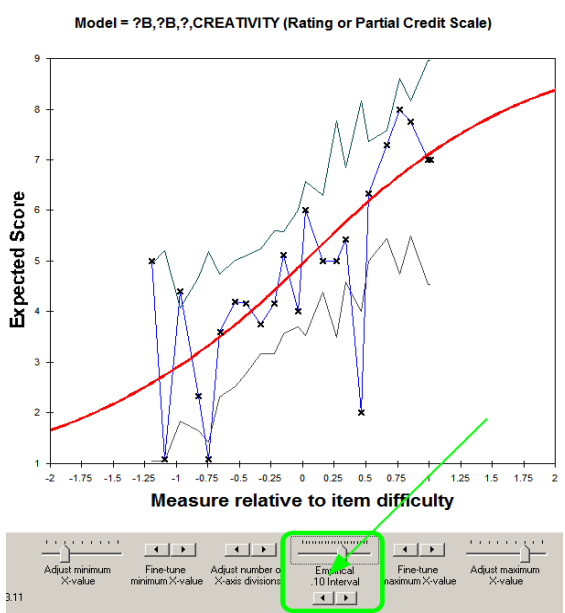



**173** Prob+Empirical Cat. Curves:

The thinner lines with x's are the empirical category frequency lines, summarizing how the rating scale categories were used. Their colors match the Rasch-model smooth curves.

Do you see that categories 3, 5, 7 are the only high-frequency categories? **Category 5** peaks in the center, where it should, but also down at the bottom - *weird!*




<p><b>174</b> Click on “Exp+ Empirical ICC”. This will display the expected (Rasch-model) and empirical (what the data say) Item Characteristic Curves, ICCs. These show the average functioning of the rating scale along the latent variable.</p>	
<p><b>175</b> The solid <b>red line</b> is the Rasch-Model ideal ICC for these data. The <b>blue jagged line</b> is what the data say. The <b>green 95% Confidence Bands</b> are the statistical limits of the divergence of the empirical from the ideal, as predicted by the Rasch model. We can see that the data only just remain within the lines. There is a problem at the bottom end of the empirical ICC, matching the problem with category 5 which we saw in its empirical category curve.</p>	
<p><b>176</b> <b>Green box and arrow:</b> Move the slider below the plot to make the empirical summarizing-interval 0.10 logits. You will see that now the empirical blue line crosses over the confidence bands, which are two-sided 95% confidence intervals.</p> <p>Even at this level of summarization, the misfit in the rating data are apparent, suggesting that the misfit should be investigated in greater detail in other Tables, such as Table 4. Something is seriously wrong with this Guilford dataset. J.P. Guilford did not notice it himself, but we will discover exactly what it is in the next Tutorial.</p> <p>Play with the “Graphs” screen, clicking different buttons and different slider settings. Do you see anything intriguing or diagnostically useful for you?</p>	
<p><b>177</b></p>	
<p><b>178</b> <i>Optional Reading:</i>  <b>#14</b> - Knox’s “Cube Imitation” Test - <a href="http://www.rasch.org/rmt/rmt133j.htm">http://www.rasch.org/rmt/rmt133j.htm</a>  <b>#62</b> - Wright &amp; Panchapakesan (1969) “A Procedure for Sample-Free Item Analysis” - <a href="http://www.rasch.org/memo46.htm">http://www.rasch.org/memo46.htm</a>      For a conceptual summary of what we have done so far, and also a glance ahead, please read “A Facets Model for Judgmental Scoring” - <a href="http://www.rasch.org/memo61.htm">http://www.rasch.org/memo61.htm</a></p>	<p style="text-align: center;"><b>A Facets Model for Judgmental Scoring</b></p> <p><small>Abstract: An extension to the Rasch model for fundamental measurement is proposed in which there is parameterization not only for examinee ability and item difficulty but also for judge severity. Several variants of this model are discussed. Its use and characteristics are explained by an application of the model to an empirical testing situation.</small></p> <p><small>Key-words: Rater, Rasch, Judging, Latent Trait, IRT</small></p> <p><small>Authors:      John M. Linacre, University of Chicago      Benjamin D. Wright, University of Chicago      Mary E. Lunz, American Society of Clinical Pathologists      MESA Memo 61, written in July 1990 and accepted for a special issue of <i>Applied Measurement in Education</i>, but not published due to lack of space.</small></p> <p><small>I Introduction:      This century has seen efforts to remove subjectivity from the measurement of examinee ability, aptitude or knowledge. There are still many areas, however, in which performance ratings depend on assessments made by judges. Artistic skills, essay writing and science projects are but a few of the many such areas in education. In measuring professional performance, “By far the most widely used performance measurement techniques are judgmental ones” (Landy and Farr, 1983 p.57).      Four factors dominate the rating given to an examinee’s performance: the ability of the examinee, the difficulty of the task performed, the severity of the judge and the way in which the judge applies the rating scale. In a competitive diving competition, each diver performs a series of dives, the test “items”, and is rated on each dive by several judges, who probably differ in levels of severity, as well as their application of the rating scale. The aim is to determine the diver of the highest skill, regardless of which dives are actually performed, and which judges happen to rate them.</small></p>
<p><b>179</b> Close all windows.</p>	

## Appendix 1. Unit Normal Deviates


The “normal” distribution is fundamental to statistics. It describes what happens when events happen “normally”, purely by chance. The Figure shows the probability of different numbers of “heads” when a coin is tossed 15 items in the red bars:

<http://mathworld.wolfram.com/NormalDistribution.html>

We can see that the overall pattern follows a bell-shaped curve the continuous black line. This pattern gets closer to a smooth line, the more coins we toss. The black continuous line for an infinite number of tosses is the “normal distribution”.



We are interested in a special case of the normal distribution. We want the one when its mean is zero, and its standard deviation is 1.0. This is called the “unit normal distribution”, abbreviated  $N(0,1)$ . Statisticians use the Greek letter mu,  $\mu$ , for the mean or average, and the Greek letter sigma,  $\sigma$ , for the standard deviation or spread, so the general normal distribution is  $N(\mu, \sigma)$ . Look at the plot, the values along the x-axis at labeled “z”, these are unit normal deviates. The area under the red curve indicates the probability of observing those values. <http://faculty.vassar.edu/lowry/ch6pt1.html> 68% of the area under the is within 1 S.D. of the mean, so we expect about 2/3 of the values we observe to be statistically close to the mean.



We are usually concerned about values far away from the mean on either side (2-sided). The Figure it says that 2.28% of the area under the curve is to the right of +2, and 2.28% is less than -2. So, when we sample from random behavior modeled this way, we expect to encounter values outside of  $\pm 2.0$  only  $2.28\% + 2.28\% = 4.56\%$  of the time. This is less than the 5% (in other words,  $p < .05$ ) that are conventionally regarded as indicating statistical significance.

The precise value of probability  $< .05$  is

$z > |\pm 1.96|$  for  $p < .05$

and for probability  $< .01$  is

$z > |\pm 2.58|$  for  $p < .01$

Handy table of unit normal deviates (z) and probabilities (p) for a “two sided z-test”, also called a “two-sided t-test with infinite degrees of freedom”

*Zstd values also use this probability table:*

z >	p <
<b>±2.58</b>	<b>0.01</b>
±2.33	0.02
±2.17	0.03
<b>±1.96</b>	<b>0.05</b>
±1.64	0.10
±1.28	0.20
±1.04	0.30
±0.84	0.40
±0.67	0.50

But, remember, just because a value is statistically significant doesn't mean that it is wrong. We do expect to see those values occasionally. The question to ask ourselves is ***"Why now?"***

What if we don't have a unit-normal distribution? We can often approximate it by taking our set of numbers, our data, subtracting from them their mean (arithmetic average) and dividing them by their standard deviation)

$$\frac{(\text{the data} - \text{their mean})}{(\text{their standard deviation})} \rightarrow N(0,1)$$


Residuals from our data,  $\{R_{ni}\}$ , have a mean of zero, and a modeled standard deviation of  $V_{ni}^{0.5}$  so the standardized residuals  $\{Z_{ni}\}$  should approximate  $N(0,1)$

$$\{R_{ni} / V_{ni}^{0.5}\} = \{Z_{ni}\} \rightarrow N(0,1)$$



## Appendix 2. Chi-square, mean-square and degrees of freedom


We talked about the unit-normal distribution in Appendix 1. And have discovered that the standardized residuals  $\{Z_{ni}\}$  approximate  $N(0,1)$ , the unit-normal distribution. So, what happens when we accumulate them?



Add two unit-normal distributions:

$$N(0,1) + N(0,1) = N(0, 2)$$

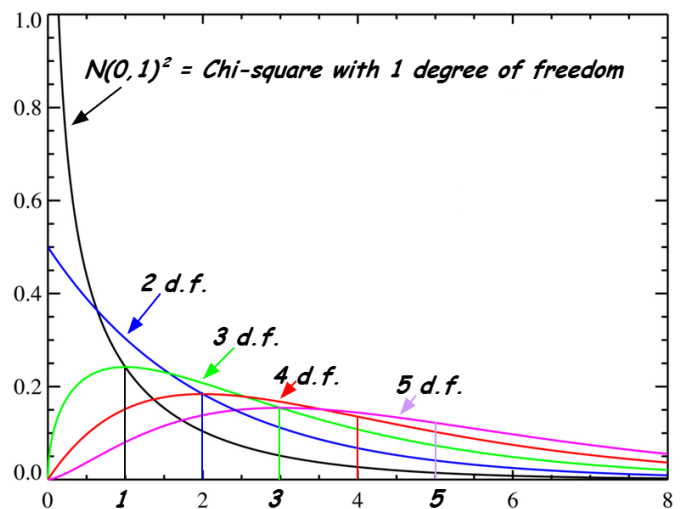
The average stays the same, but they spread out more. The combined distribution has twice the variance of a unit-normal distribution:



But what if we square the unit-normal distribution?  $N(0,1)^2$  is called the “chi-square distribution with 1 degree of freedom”, shortened to  $\chi^2_1$ . It is the black curved line on the plot. Its mean is its degrees of freedom, indicated by the black vertical line going up from 1.

We can add two of these  $N(0,1)^2 + N(0,1)^2 = \chi^2_2$ . This has two degrees of freedom, d.f., and is shown by the blue curve on the plot.

We can keep adding more. So, when we have added “k” squared (unit normal distributions) we have a chi-square distribution with k d.f.,  $\chi^2_k$ . It has a mean of k and a variance of 2k, so a standard deviation of  $\sqrt{2k}$ .



Since the mean of chi-square statistic is its d.f., it is convenient to divide the chi-square by its d.f., so that its value can be compared with 1.0. This makes scanning a Table of fit statistics much easier than when chi-square statistics with their d.f. are reported.

$$\text{Mean-square} = \chi^2_k / k$$

Mean-square  $\ll 1$  is over-fit, dependency, over-parameterization, over-predictability

Mean-square  $\gg 1$  is under-fit, noise, misfit, lack of predictability

*Facets* reports the significance (probability) of a mean-square as a unit-normal deviate (Zstd).

ZStd = Wilson-Hilferty (mean-square, d.f.)  
see <http://www.rasch.org/rmt/rmt162g.htm>

### Appendix 3. Excel plots from the Residual file

Here is how you can produce Excel plots from the Facets “Residuals file”. This needs some skill with Excel.

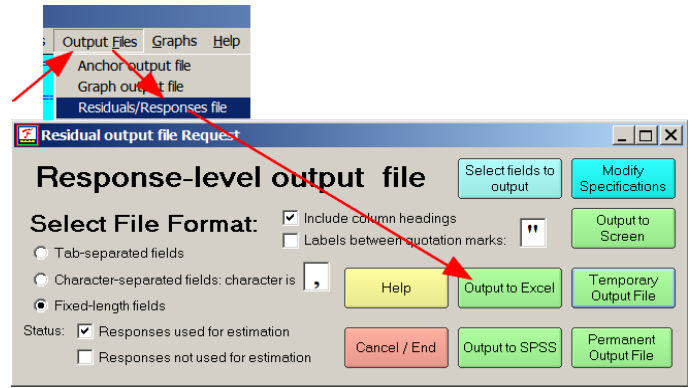
Facets Analysis window

Click on “Output Files”

Click on “Residuals/Responses file”

“Residual output file Request” dialog box

Click on “Output to Excel”



Excel worksheet:

Click on the worksheet

“Select All” Ctrl+A

“Sort and Filter”

“Custom Sort”

Top row is headings

Sort fields:

The facet number you want, ascending, e.g., “2”

The x-axis value you want, ascending, e.g., “Logit”

OK

The worksheet is sorted

1	Obs	Exp	Res	Var	StRes	Wt	LProb	Logit	1	2	3	Senior scie	Junior Scie	Traits	
2	5	5	5.34	-0.34	2.41	-0.22	1	-0.9704	0.16	1	1	1	Avogadro	Anne	Attack
3	5	5	5.03	-0.03	2.42	-0.02	1	-0.9465	0.03	1	1	2	Avogadro	Anne	Basis
4	3	3	4.23	-1.23	2.25	-0.82	1	-1.104	-0.31	1	1	3	Avogadro	Anne	Clarity
5	5	5	5.39	-0.39	2.4	-0.25	1	-0.9782	0.18	1	1	4	Avogadro	Anne	Daring
6	3	3	3.6	-0.6	1.98	-0.43	1	-0.8364	-0.6	1	1	5	Avogadro	Anne	Enthusiasm
7	9	9	6.87	2.13	1.85	1.56	1	-1.9162	0.87	1	2	1	Avogadro	Betty	Attack
8	7	7	6.63	0.37	1.97	0.26	1	-0.7822	0.74	1	2	2	Avogadro	Betty	Basis
9	5	5	5.9	-0.9	2.29	-0.59	1	-1.1168	-0.4	1	2	3	Avogadro	Betty	Clarity
10	8	8	6.91	1.09	1.82	0.8	1	-2.2179	0.89	1	2	4	Avogadro	Betty	Daring
11	5	5	5.2	-0.2	2.42	-0.13	1	-0.9543	0.1	1	2	5	Avogadro	Betty	Enthusiasm
12	3	3	4.92	-1.92	2.41	-1.23	1	-1.5647	-0.02	1	3	1	Avogadro	Chris	Attack
13	3	3	4.61	-1.61	2.36	-1.05	1	-1.34	-0.14	1	3	2	Avogadro	Chris	Basis
14	3	3	3.84	-0.84	2.09	-0.58	1	-0.9219	-0.49	1	3	3	Avogadro	Chris	Clarity
15	7	7	4.97	2.03	2.41	1.31	1	-1.6426	0.01	1	3	4	Avogadro	Chris	Daring
16	1	1	3.27	-2.27	1.81	-1.68	1	-2.1605	-0.78	1	3	5	Avogadro	Chris	Enthusiasm
17	7	7	4.4	2.6	2.31	1.71	1	-2.1951	-0.23	1	4	1	Avogadro	David	Attack
18	3	3	4.11	-1.11	2.21	-0.75	1	-1.045	-0.36	1	4	2	Avogadro	David	Basis


Scatterplot:

“Series name” is the element label you want in facet 2.

x-axis values are the “Logits” (or whatever) for the element you want in facet 2

y-axis values are the “Obs(ervations)” (or whatever) for the element you want in facet 2

Excel produces a plot



Use Excel tools to customize your plot.

*Be exuberant! With a little time and talent, this becomes fun.*

