

Using the **difR** Package with Lertap5

Larry R Nelson¹
Curtin University, Western Australia
www.Lertap5.com

Document date: 14 April 2020

difR is the name of an R package designed to support “DIF” analyses, where DIF=differential item functioning. It is introduced at [this webpage](#). When I undertook the work described below, I was using difR Version 5, published 2018-05-14.

The use of R packages with Lertap5 is discussed in [this document](#). That document has a few examples of how R packages are installed. Following those examples, the difR package would be installed using

```
> install.packages("difR")
```

Lertap5 has its own DIF analysis capability based on the Mantel-Haenszel (M-H) method. The “[Ibreaks](#)” routine in Lertap5 is the gateway to M-H; a technical paper discussing and exemplifying the use of Lertap5 M-H is [here](#).

The **difR** package has more extensive support for DIF researchers, featuring the M-H method and several others, including logistic regression and SIBTEST (to name just two of the others – there are more).

This paper is for those who have item response data in a Lertap5 workbook and would like to use the data with the **difR** package.

When test results are input to difR, it is expected that the input will consist of item scores and a group variable. Item scores will be zero and one (0,1) – these are often referred to as “binary” scores. The group variable is expected to have two categories². It might, for example, be Gender with codes of F and M, or, as in the example below, Race with codes of B and W. The version of difR used for this paper would allow for missing data among the item scores, but not in the group variable. The codes used for the group variable could be letters or digits³.

To prepare data for difR, Lertap5 users will have to get Lertap5 to make item scores, make sure that their group variable has just two codes, make sure that all data records have a group code (no missing data), and then prepare a “csv” file.

Note at 7 April 2020: the discussion on the next few pages describes how to set up a “csv” Excel workbook suitable for use with the difR program. Another way to set up such a workbook is to use the “[difR1 special macro](#)” – it’s faster and easier, but not as “transparent” – the steps which follow below involve more work but allow greater flexibility in how the csv file’s group variable is brought in. Readers who choose to use the special macro should do so, and then return here, to this document, skipping ahead to the section “**M-H results with difR**”.

As seen below, my suggestion is that the csv file start with columns having the item scores, one column per test item, and end with the group code in the last column. The first row in the csv file will be for variable “headers”. The item headers will typically be, for example, Item1 or Q1.

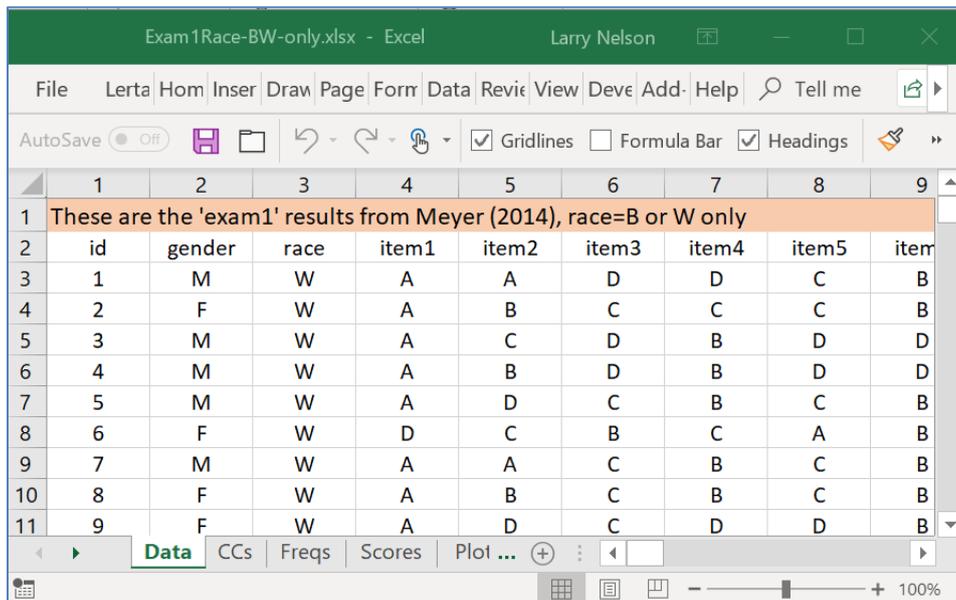
¹ Comments / questions may be sent to l.nelson@curtin.edu.au

² difR supports more than two groups.

³ Lertap5 has a [recode option](#) that might be useful (for example, change {M,F} codes to {0,1})

Creating a csv file from a Lertap5 workbook

I have selected the "exam1" test results for this example. Exam1 is used as a major exhibit in Patrick Meyer's 2014 text "Applied Measurement with jMetrik" (see the Lertap5 [references](#)). I have used these results in my "[GimmeABreak](#)" paper, where they featured in Example 2 (pp. 15-19).



The screenshot shows an Excel spreadsheet titled "Exam1Race-BW-only.xlsx" with the following data:

	1	2	3	4	5	6	7	8	9
1	These are the 'exam1' results from Meyer (2014), race=B or W only								
2	id	gender	race	item1	item2	item3	item4	item5	item
3	1	M	W	A	A	D	D	C	B
4	2	F	W	A	B	C	C	C	B
5	3	M	W	A	C	D	B	D	D
6	4	M	W	A	B	D	B	D	D
7	5	M	W	A	D	C	B	C	B
8	6	F	W	D	C	B	C	A	B
9	7	M	W	A	A	C	B	C	B
10	8	F	W	A	B	C	B	C	B
11	9	F	W	A	D	C	D	D	B

The screen snapshot above displays some of the rows and columns in the [Lertap5 Data](#) worksheet corresponding to Meyer's "exam1" test results. In this case, the three initial columns were used to house descriptive information; these columns were followed by item responses. There were a total of 56 multiple-choice items; responses were found starting in column 4, ending in column 59.

I will need to have item scores before I'm able to use the difR routines. The screen shot above shows item *responses*, not item *scores*. To get item scores, I'll use the "[Item scores and correlations](#)" option in Lertap5.

The "Item scores and correlations" option adds the "**IStats**" worksheet to the workbook. It has two sections: item scores, followed by item statistics and correlations.

Next, I take the "[BinaryItems](#)"⁴ option in Lertap5. It creates a new workbook with item scores – here's how part of it looks at this stage:

Cautions: readers should make sure that the scores created are indeed {0.00, 1.00}. The difR package accepts only {0.00,1.00,NA} item scores, where NA=missing data). The [Freqs](#) report can be used to reveal the extent of [missing data](#); in Lertap, missing data for cognitive items is usually converted to an item score of 0.00.

⁴ Caution: the BinaryItems option does not check to make sure that cognitive items are in use.

Book8 - Excel Larry Nelson

File Lerta Hom Inser Draw Page Form Data Revi View Deve Ac

AutoSave Off Gridlines

	1	2	3	4	5	6	7	8	9
1	Lertap5 IScores matrix, created: 23/1/19								
2	id	item1	item2	item3	item4	item5	item6	item7	item8
3	1	1.00	0.00	0.00	0.00	1.00	1.00	1.00	1.00
4	2	1.00	0.00	1.00	0.00	1.00	1.00	0.00	1.00
5	3	1.00	0.00	0.00	1.00	0.00	0.00	0.00	1.00
6	4	1.00	0.00	0.00	1.00	0.00	0.00	1.00	1.00
7	5	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
8	6	0.00	0.00	0.00	0.00	0.00	1.00	1.00	1.00
9	7	1.00	0.00	1.00	1.00	1.00	1.00	0.00	1.00
10	8	1.00	0.00	1.00	1.00	1.00	1.00	1.00	1.00
11	9	1.00	1.00	1.00	0.00	0.00	1.00	0.00	1.00

Data CCs

100%

Notice that the BinaryItems option has not copied the Gender and Race information from the original workbook. As I'd like to look at results based on race, I now need to copy the race column from the original workbook and paste it into the new workbook's Data sheet. When I do so, I paste it on the right-hand side of the new workbook:

Book8.xlsx - Excel Larry Nelson

File Home Insert Draw Page Layout Formulas Data Review View Help

AutoSave Off Gridlines

	52	53	54	55	56	57	58
1							
2	item51	item52	item53	item54	item55	item56	race
3	1.00	0.00	1.00	0.00	0.00	1.00	W
4	1.00	0.00	1.00	1.00	1.00	0.00	W
5	1.00	1.00	1.00	0.00	1.00	1.00	W
6	1.00	0.00	0.00	0.00	0.00	0.00	W
7	1.00	0.00	0.00	1.00	0.00	1.00	W
8	0.00	1.00	0.00	1.00	0.00	0.00	W
9	1.00	1.00	1.00	0.00	0.00	1.00	W
10	0.00	0.00	0.00	0.00	1.00	0.00	W
11	1.00	0.00	0.00	0.00	1.00	1.00	W

Data CCs

120%

Next I need to delete the first row so that what shows as row 2 below will become the first row. I will not need the id column, and will delete it.

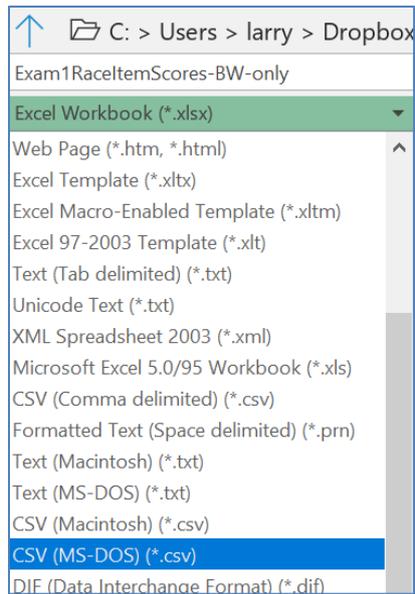
	1	2	3	4	5	6	7	8	9
1	Lertap5 IScores matrix, created: 23/1/19 h								
2	id	item1	item2	item3	item4	item5	item6	item7	item8
3	1	1.00	0.00	0.00	0.00	1.00	1.00	1.00	1.00

I now tell Excel to save this new workbook, giving it an appropriate name: "Exam1RaceItemScores-BW-only.xlsx".

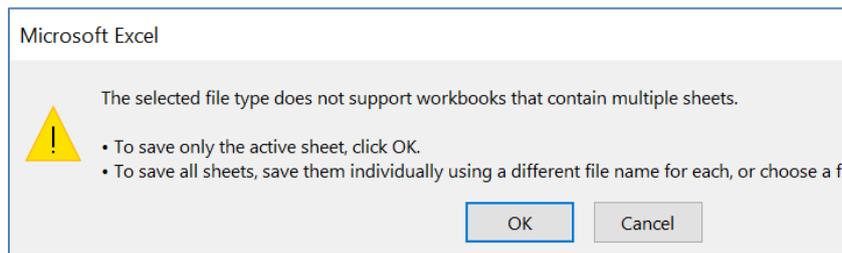
	1	2	3	4	53	54	55	56	57
1	item1	item2	item3	item4	item53	item54	item55	item56	race
2	1.00	0.00	0.00	0.00	1.00	0.00	0.00	1.00	W
3	1.00	0.00	1.00	0.00	1.00	1.00	1.00	0.00	W
4	1.00	0.00	0.00	1.00	1.00	0.00	1.00	1.00	W
5	1.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	W
6	1.00	1.00	1.00	1.00	0.00	1.00	0.00	1.00	W
7	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	W
8	1.00	0.00	1.00	1.00	1.00	0.00	0.00	1.00	W
9	1.00	0.00	1.00	1.00	0.00	0.00	1.00	0.00	W
10	1.00	1.00	1.00	0.00	0.00	0.00	1.00	1.00	W
11	0.00	0.00	1.00	1.00	0.00	0.00	0.00	0.00	W
12	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	W
13	1.00	1.00	1.00	0.00	1.00	0.00	0.00	1.00	W

The snapshot above shows that the new workbook has **item1** scores in the first column, then more item scores, and, at the end, **race** in column 57. This exemplifies the format for a workbook that will best serve my needs when using **difR**: columns with item scores followed by the "group" variable for DIF (which is 'race' in the case).

Now, with the focus on the Data sheet (as in the snapshot above), I ask Excel to save as a "CSV (MS-DOS)" file, highlighted in blue below:

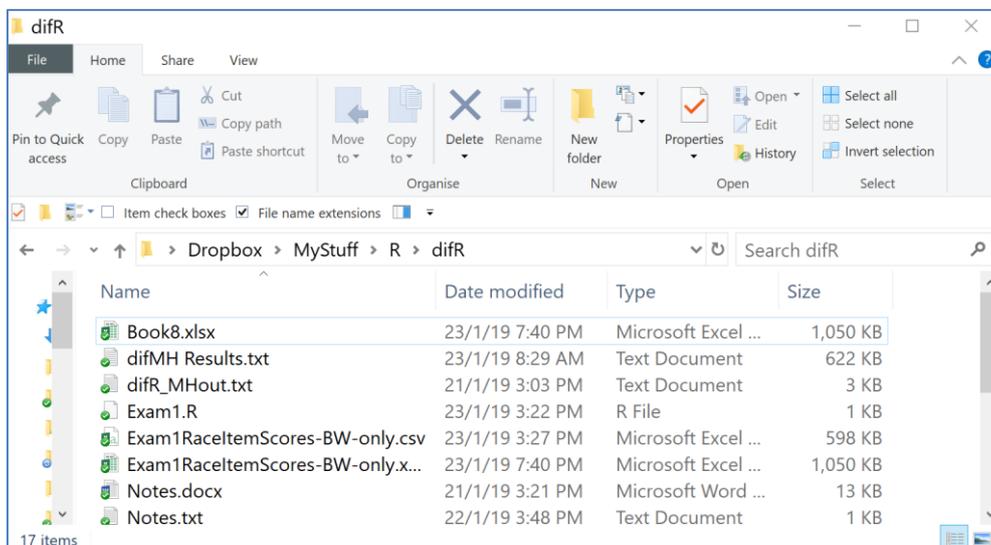


Excel sounds a warning:



I click OK and close the workbook.

By using the Windows file manager I can confirm that the workbook has indeed been saved as a file named "Exam1RaceItemScores-BW-only.csv"⁵.



I now have a csv file that can readily be used with the **difR** package.

⁵ If the [difR1_special_macro](#) has been used, the name of the csv file will always be "difR-IScores.csv".

M-H results with difR

I have both R and RStudio installed on my laptop, as mentioned in [this paper](#). Were I to continue using **difR**, and well I might, I'd probably develop a special script for use with RStudio. For the moment, however, I'll just run with R.

The "Exam1.R" script file shows in the file list above; its contents are shown below:

```
# Made to work with the Exam1 dataset, 24 January 2019
# By Larry Nelson, Curtin University, Western Australia
# l.nelson@curtin.edu.au
setwd("C:/Users/larry/Dropbox/MyStuff/R/difR/Exam1")
library(difR)
dir()
AAitems <-read.csv(file="Exam1RaceItemScores-BW-only.csv")
difItems<-AAitems[colnames(AAitems)]
head(difItems)
difMH(difItems, group="race", focal.name="W", save.output=TRUE)
```

When R finishes running the code lines, it provides this message:

```
Output was captured and saved into file
'C:/Users/larry/Dropbox/MyStuff/R/difR/Exam1/out.txt'
```

Some of the output is shown below:

Mantel-Haenszel Chi-square statistic:

	Stat.	P-value	
item1	19.3862	0.0000	***
item2	14.5016	0.0001	***
item3	7.1053	0.0077	**
item4	21.0385	0.0000	***
item5	1.4150	0.2342	
item6	17.1696	0.0000	***
item7	0.1694	0.6806	
item8	0.9321	0.3343	
item9	6.5306	0.0106	*
item10	0.0695	0.7920	

Items detected as exhibiting DIF:

```
item1
item2
item3
item4
item6
item9
item16
```

(this is a partial listing; there were 25 items in the whole list, all with P-values less than .05)

Effect size (ETS Delta scale):

```
Effect size code:
'A': negligible effect
'B': moderate effect
'C': large effect
```

	alphaMH	deltaMH	
item1	0.7198	0.7726	A
item2	0.7791	0.5865	A
item3	0.7722	0.6076	A
item4	0.7396	0.7090	A
item5	0.9179	0.2014	A
item6	1.4001	-0.7909	A
item7	1.0295	-0.0683	A
item8	0.8904	0.2728	A
item9	0.8404	0.4085	A
item10	1.0327	-0.0756	A
item37	2.0179	-1.6499	C
item45	1.8638	-1.4632	B

Only two items, item37 and item45, evidenced practical significance as indexed by the deltaMH statistic.

As I mentioned earlier, difR can look for possible DIF using numerous methods – here I’ve used only one, Mantel-Haenszel, the method used in Lertap5, and a method with good support in the literature.

difR is able to make graphical summaries of DIF.

Have a look at these URLs:

<https://cran.r-project.org/web/packages/difR/difR.pdf>

https://eeecon.uibk.ac.at/psychoco/2011/slides/Magis_hdt.pdf

The following Appendix has other examples of running difR.

Appendix

```
# Made to work with the Exam1 dataset, 13 February 2019
# By Larry Nelson, Curtin University, Western Australia
# l.nelson@curtin.edu.au
setwd("C:/Users/larry/Dropbox/MyStuff/R/difR/Exam1")
library(difR)
dir()
AAitems <- read.csv(file="Exam1RaceItemScores-BW-only.csv")
difItems <- AAitems[colnames(AAitems)]
head(difItems)
# difMH(difItems, group="race", focal.name="W", save.output=TRUE)
# difMH(difItems, group="race", focal.name="W")
# dichDif(difItems, group="race", focal.name="W",
# method = c("TID", "MH", "Std", "Logistic", "SIBTEST"), correct =
FALSE,
# thrSTD = 0.08, thrTID = 1, purify = TRUE, save.output = TRUE,
output = c("dichodif1", "default"))
# dichDif(difItems, group="race", focal.name="W",
# method = c("TID", "MH", "Std", "Logistic", "SIBTEST"), correct =
FALSE,
# thrSTD = 0.08, thrTID = 1, purify = FALSE, save.output = TRUE,
output = c("dichodif1", "default"))
dichDif(difItems, group="race", focal.name="W",
method = c("MH", "Std", "Logistic", "SIBTEST"), correct = FALSE,
thrSTD = 0.08, thrTID = 1, purify = FALSE, save.output = TRUE,
output = c("dichodif1", "default"))
```

The R code above displays some of my experiments with difR. *All lines that begin with the # character are just comments.* The last of dichDif code lines above are repeated below (for convenience)

```
dichDif(difItems, group="race", focal.name="W",
method = c("MH", "Std", "Logistic", "SIBTEST"), correct = FALSE,
thrSTD = 0.08, thrTID = 1, purify = FALSE, save.output = TRUE,
output = c("dichodif1", "default"))
```

and, when run, resulted in this output from difR:

Comparison of DIF detection results using 4 methods

Methods used:

- Mantel-Haenszel
- Standardization
- Logistic regression
- Crossing-SIBTEST

Matching variable: test score

No set of anchor items was provided

Parameters:

- Significance level: 0.05
- Standardization threshold: 0.08
- Mantel-Haenszel DIF statistic: Chi-square statistic
- Mantel-Haenszel continuity correction: No
- Type of Mantel-Haenszel test: asymptotic test
- Weights for standardized P-DIF statistic: based on the focal group
- Logistic regression DIF statistic: LRT statistic

DIF effect(s) tested by logistic regression: both DIF effects
Item purification: No

No p-value adjustment for multiple comparisons

Comparison of DIF detection results:

	M-H	Stand.	Logistic	CSIBTEST	#DIF
item1	DIF	NoDIF	DIF	DIF	3/4
item2	DIF	NoDIF	DIF	DIF	3/4
item3	DIF	NoDIF	DIF	DIF	3/4
item4	DIF	NoDIF	DIF	DIF	3/4
item5	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item6	DIF	NoDIF	DIF	DIF	3/4
item7	NoDIF	NoDIF	DIF	DIF	2/4
item8	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item9	DIF	NoDIF	DIF	DIF	3/4
item10	NoDIF	NoDIF	DIF	NoDIF	1/4
item11	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item12	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item13	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item14	NoDIF	NoDIF	DIF	NoDIF	1/4
item15	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item16	DIF	NoDIF	DIF	DIF	3/4
item17	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item18	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item19	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item20	DIF	NoDIF	DIF	DIF	3/4
item21	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item22	NoDIF	NoDIF	DIF	NoDIF	1/4
item23	DIF	NoDIF	DIF	DIF	3/4
item24	NoDIF	NoDIF	DIF	NoDIF	1/4
item25	DIF	NoDIF	DIF	NoDIF	2/4
item26	DIF	NoDIF	DIF	NoDIF	2/4
item27	DIF	NoDIF	DIF	DIF	3/4
item28	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item29	DIF	NoDIF	DIF	DIF	3/4
item30	DIF	NoDIF	DIF	DIF	3/4
item31	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item32	DIF	NoDIF	DIF	DIF	3/4
item33	NoDIF	NoDIF	DIF	NoDIF	1/4
item34	DIF	NoDIF	DIF	DIF	3/4
item35	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item36	DIF	NoDIF	DIF	DIF	3/4
item37	DIF	DIF	DIF	DIF	4/4
item38	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item39	DIF	NoDIF	NoDIF	DIF	2/4
item40	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item41	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item42	DIF	NoDIF	DIF	DIF	3/4
item43	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item44	DIF	NoDIF	DIF	DIF	3/4
item45	DIF	DIF	DIF	DIF	4/4
item46	DIF	NoDIF	DIF	DIF	3/4
item47	NoDIF	NoDIF	DIF	NoDIF	1/4
item48	NoDIF	NoDIF	DIF	NoDIF	1/4
item49	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item50	NoDIF	NoDIF	NoDIF	DIF	1/4
item51	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item52	DIF	NoDIF	NoDIF	DIF	2/4
item53	DIF	NoDIF	DIF	DIF	3/4
item54	NoDIF	NoDIF	NoDIF	NoDIF	0/4

item55	NoDIF	NoDIF	NoDIF	NoDIF	0/4
item56	NoDIF	NoDIF	NoDIF	DIF	1/4

The following line draws on the "Standardization DIF method" in difR:

```
plot(difStd(difItems,group="race",focal.name = "w"))
```

and when run resulted in this plot:

