



MetaXL User Guide

version 1.2

Jan J Barendregt & Suhail A Doi



© EpiGear International Pty Ltd
ABN 51 134 897 411
Brisbane, Australia, 2011-2012

Web: www.epigear.com
Email: info@epigear.com

Contents

Contents	3
Introduction.....	5
Meta-analysis in Excel.....	5
Installation.....	5
License	5
Features	5
Design issues.....	6
Operation.....	6
This User Guide	6
Technical note	7
MetaXL functions	8
Introduction.....	8
MAInputTable.....	8
MASubGroups	11
MAPooledEffect	12
MAPooledEffectLCI & MAPooledEffectHCI	12
MAForestName.....	12
MACochranQ.....	12
MAISquare.....	12
MATauSquare.....	12
MAQIndex	13
MANumStudies	13
MetaXL menu and features.....	14
Introduction.....	14
Input templates.....	14
Results.....	14
Options.....	20
MetaXL User Guide.....	20
Examples.....	20
Reset.....	20
About MetaXL	21
Meta-analysis methods.....	22
Introduction.....	22
Various models	22
The random effects model	22
Rationale for the QE model	24
Differences in weighting between models	25
Meta-analysis of Prevalence	27
Introduction.....	27
Prevalence as a variable	27
Transformations	28
Multi-category prevalence	29
The MetaXL implementation.....	30
Conclusion	31
Continuity corrections.....	32
Binary studies.....	32
Pooled prevalence	32
Documentation.....	33



Examples.....	33
Introduction.....	33
ThyroidCancerRR.....	33
Magnesium.....	33
Omega3WMD.....	34
MarshallSMD.....	34
SchizophreniaPrev.....	34
MSMultipleCategoryPrev.....	34
Error messages.....	35
Introduction.....	35
MAInputTable error messages.....	35
Bugs.....	37
Trouble shooting.....	38
Known issues.....	40
MetaXL versions.....	41
Version 1.2.....	41
Version 1.11.....	41
Version 1.1.....	41
Version 1.0.....	42
Technical appendix.....	43
Installation issues.....	43
Statistical and other scientific sources.....	44
Software.....	45
About the authors.....	45
References.....	46

Introduction

Meta-analysis in Excel

MetaXL is a tool for meta-analysis in Microsoft Excel. It extends Excel with several functions for input and output of meta-analysis data, and a menu that offers access to output (both in table and graphical format), options, examples, and this *User Guide*, among other things.

MetaXL employs the same meta-analysis methods that can be accessed in general statistical packages (such as Stata™) and in dedicated meta-analysis software, but makes an additional Quality Effects method available, see the section on Meta-analysis methods below.

Installation

MetaXL has been tested with Excel 2000, 2003, 2007, and 2010¹. Installation is usually straightforward: make sure Excel is not running, and run the MetaXLSetup executable. If all goes well, the MetaXL commands will be present in Excel (in the main Excel menu for Excel version 2003 and earlier, in the MetaXL ribbon tab in version 2007 and later), the MetaXL functions will be visible in the Excel Function Wizard (in the function category 'MetaXL'), and you can open the example spreadsheets without seeing any "#NAME?" errors. See the section on 'Trouble shooting' when this is not the case.

The installation program installs an Excel add-in (called 'MetaXL.xll'), together with this documentation file and a few example spreadsheets. For Excel 2007 and later, an additional add-in (called 'MetaXLRibbon.xlam') is installed. In the Windows Start menu a MetaXL entry gives access to the example spreadsheets, this *User Guide*, and an Uninstall program. For a detailed discussion of the installation issues, see the Technical Appendix below.

License

MetaXL is free, but copyrighted software. MetaXL may be installed on as many computers as the user owns or employs, and can be freely distributed to third parties. However, modification of the installation software and distribution of the software to third parties for a fee are strictly prohibited.

Please contact Epigear (info@epigear.com) if you have any questions about this license arrangement.

Features

MetaXL has a number of features:

1. Fixed effects (inverse variance, Mantel Haenszel, Peto), random effects (DerSimonian & Laird) and quality effects (Doi & Thalib) models.
2. Both binary (relative risk, odds ratio, risk difference) and continuous (weighted mean difference, Cohen's d, Hedges' g, Glass's Δ) methods.
3. Pooled prevalence and rates methods.
4. Heterogeneity statistics: Cochran's Q, I^2 , τ^2 .
5. Sensitivity analysis and study exclusion option.

¹ 32 bit Excel 2010 only. The current version is not compatible with 64 bit Excel 2010. The default Excel 2010 installation is 32 bit.

6. A wide range of input format options.
7. Output in table and graphical (forest and funnel plots) formats, with various ways to copy or save results.

These features are discussed in sections below.

Design issues

MetaXL has been developed with academic research in mind. This implies that transparency of the methods used is deemed important. We have aimed to reference all the algorithms used, or to describe them in sufficient detail in this *User Guide* for algorithms unique to MetaXL. If you have any questions on MetaXL methods left after consulting this *User Guide* and the referenced literature, please contact us.

One of the big advantages of MetaXL being Excel based is that it can very easily be combined with Excel based Monte Carlo simulation software, such as EpiGear's Ersatz. Using Ersatz and MetaXL together makes research into, for example, coverage properties of the various models a breeze, with thousands of meta-analysis calculations done within a minute.

Another important design objective was ease of use. The interface has been kept very simple (and therefore, hopefully, easy to navigate). For users of Excel, MetaXL should pose no hurdles.

Operation

Running MetaXL is done from within Excel. To see how it operates, start Excel, and open one of the example spreadsheets (accessible through the MetaXL menu, which is located in the main Excel menu for Excel version 2003 and earlier, in the MetaXL ribbon tab in version 2007 and later). The example spreadsheets contain tables with study results, and several input table functions (MAInputTable). In addition, there are some functions that report results: MaPooledEffect, and MAPooledEffectLCI and MAPooledEffectHCI for the pooled result and its confidence interval respectively. See the section on MetaXL functions below for a detailed description of these functions and the parameters they take.

Full results are accessed by choosing 'Results' from the MetaXL menu. This will produce a list of choices, one entry for each MAInputTable function in the workbook. Clicking one of these produces a pop-up window with tabs giving access to various outputs such as a table and forest and funnel plots.

Results can be saved to file by choosing 'File|Save results', and can be saved to the clipboard (and then pasted into, for example, Excel or Word) by choosing 'Options|Copy to clipboard', or by right-clicking and choosing 'Copy'.

The MetaXL menu also gives access to input templates for setting up your own meta-analysis, see the discussion of the MAInputTable function below.

This User Guide

Although in the remainder we will discuss the various methods of meta-analysis, this *User Guide* is certainly not meant as an introduction to these issues. The reader is kindly referred to the literature, see the references for some examples. The aim of this document is to make the user familiar with MetaXL and its interface, possibilities, and quirks.

Technical note

Please consult the Technical Appendix below for statistical and other technical details.

MetaXL functions

Introduction

In this section we will describe the functions MetaXL makes available in Excel. All MetaXL functions start with 'MA', and are accessible through Excel's Function Wizard (for Excel 2003 and earlier: InsertFunction in Excel's main menu, for 2007 and later: FormulasInsert Function, or click the icon to the left of the formula bar) in the category 'MetaXL'. Check out the Function Wizard if you don't know it.

Of the functions described in this section, only the MAInputTable is essential: it determines the input and output and the meta-analysis method used.

MAInputTable

MAInputTable(Name, IOType, Method, Table): returns the Name parameter (if all is well) or an error message (if not).

Table 1: parameters of the MAInputTable function

Parameter	Type	Description
Name	text	A unique name for each MAInputTable function in the workbook (and across workbooks if you have more than one workbook open at a time). The Name should be more than 0 and less than 256 characters long.
IOType	text	One of: NumRR, NumOR, NumRD, RRCI, ORCI, RDCI, RDSE, WMD, Cohen, Hedges, Glass, ContCI, ContSE, Prev, NumRate, RateSE (see notes below).
Method	text	One of: IV, RE, QE, MH, Peto (see notes below).
Table	range	An Excel range, containing the study results for the meta-analysis (see notes below).

Notes:

1. The MAInputTable function either returns its 'Name' parameter (indicating that no errors were encountered) or an error message. See the section on Error messages. Please note that you can have multiple MAInputTable functions in an Excel workbook, but that each must have a unique name.
2. Table 1 defines the parameters of the MAInputTable function. Please note that the parameters of type text need to be entered between double quotes.
3. Table 2 gives the values the IOType parameter can take.

Table 2: values of the IOType parameter

Value	Description
NumRR	Binary analysis, input in numbers, output in risk ratios
NumOR	Binary analysis, input in numbers, output in odds ratios
NumRD	Binary analysis, input in numbers, output in risk differences
RRCI	Binary analysis, input in risk ratios and confidence intervals, output in risk ratios
ORCI	Binary analysis, input in odds ratios and confidence intervals, output in odds ratios
RDCI	Binary analysis, input in risk differences and confidence intervals, output in risk differences

Value	Description
RDSE	Binary analysis, input in risk differences and standard errors, output in risk differences
WMD	Continuous analysis, input population size, mean, and standard deviation, output weighted mean difference
Cohen	Continuous analysis, input population size, mean, and standard deviation, output standardised mean difference (Cohen's d)
Hedges	Continuous analysis, input population size, mean, and standard deviation, output standardised mean difference (Hedges' adjusted g)
Glass	Continuous analysis, input population size, mean, and standard deviation, output standardised mean difference (Glass's Δ)
ContCI	Continuous analysis, input either WMD, Cohen's d , Hedges' adjusted g , or Glass's Δ with confidence intervals, output same as input
ContSE	Continuous analysis, input either WMD, Cohen's d , Hedges' adjusted g , or Glass's Δ with standard errors, output same as input
Prev	Pooling prevalences, input population size and number of cases, output proportion.
NumRate	Pooling rates, input person time at risk and number of cases, output rate.
RateSE	Pooling rates, input rate and standard error, output rate

4. The Method parameter stands for the following meta-analysis methods (see the section on Meta-analysis methods below):

Table 3: values of the Method parameter

Value	Description
IV	Inverse variance (fixed effects)
RE	Random effects (DerSimonian & Laird)
QE	Quality effects (Doi & Thalib)
MH	Mantel Haenszel (fixed effects)
Peto	Peto method for odds ratios (fixed effects)

5. Not every IOType is compatible with every Method parameter (and combining incompatible ones will result in an error message to that effect):

Table 4: compatible values of the IOType and Method parameters

IOType\method	IV	RE	QE	MH	Peto
NumRR	✓	✓	✓	✓	
NumOR	✓	✓	✓	✓	✓
NumRD	✓	✓	✓	✓	
RRCI, ORCI, RDCI, RDSE	✓	✓	✓		
WMD, Cohen, Hedges, Glass	✓	✓	✓		
ContCI, ContSE	✓	✓	✓		
Prev, NumRate, RateSE	✓	✓	✓		

6. The number of columns of the Table parameter, and the kind of information in it, depends on the IOType and the Method parameters. Table 5 below specifies

the columns the MAInputTable function expects for the various IOType and Method combinations.

Please be aware that the MAInputTable function is very particular about order and content of the columns. Only when the layout is strictly adhered to will MetaXL function correctly or at all. For this reason, the precise columns are also obtainable through the MetaXL menu 'Input templates' item, see the section on the Menu below).

In the table below N1 and N2 stand for population numbers in active and control arms respectively, each with their Cases and Non-cases (binary) or Mean and Standard deviation (continuous), and Qi stands for a quality index, a number between 0 and 1 (see section on Meta-analysis methods below). Lo and Hi CI stand respectively for the lower and higher limit of the confidence interval as input for the IOType parameters that take a confidence interval as input. The size of the input confidence interval can be set in the Options menu (see the section on the MetaXL menu below) and is independent of the size of the output confidence interval.

7. Please note that within an input Table each study name must be unique: it is an error when there are duplicate study names.

Table 5: columns of the Table parameter

IOTypes / Methods	Table Columns							
NumRR, NumOR, NumRD / IV, RE, MH, Peto	Study name	N1	Cases	Non-cases	N2	Cases	Non-cases	
NumRR, NumOR, NumRD / QE	Study name	N1	Cases	Non-cases	N2	Cases	Non-cases	Qi
ORCI, RRCI, RDCI / IV, RE	Study name	RR/OR /RD	Lo CI	Hi CI				
ORCI, RRCI, RDCI / QE	Study name	RR/OR /RD	Lo CI	Hi CI	Qi			
RDSE / IV, RE	Study name	RD	Standard error					
RDSE / QE	Study name	RD	Standard error	Qi				
WMD, Cohen, Hedges, Glass / IV, RE	Study name	N1	Mean	StDev	N2	Mean	StDev	
WMD, Cohen, Hedges, Glass / QE	Study name	N1	Mean	StDev	N2	Mean	StDev	Qi
ContCI / IV, RE	Study name	Effect size	Lo CI	Hi CI				
ContCI / QE	Study name	Effect size	Lo CI	Hi CI	Qi			
ContSE / IV, RE	Study name	Effect size	Standard error					

IOTypes / Methods	Table Columns							
ContSE / QE	Study name	Effect size	Standard error	Qi				
Prev / IV, RE	Study name	N	Cases					
Prev / QE	Study name	N	Cases	Qi				
NumRate / IV, RE	Study name	Person time	Cases					
NumRate / QE	Study name	Person time	Cases	Qi				
RateSE / IV, RE	Study name	Rate	Standard error					
RateSE / QE	Study name	Rate	Standard error	Qi				

Important note

The remaining MetaXL functions all take a “TableName” parameter. In order for these functions to be associated with the correct MAInputTable function, the “TableName” parameter should be provided by linking it to the output of the corresponding MAInputTable function.

MASubGroups

MASubGroups(Name, TableName, Table): returns the Name parameter (if all is well) or an error message (if not).

Notes:

1. This function allows doing a subgroup analysis of the studies in the MAInputTable function with name TableName. The result will appear in the Results list (see below) as ‘TableName by Name’.
2. The Table parameter is an Excel range that groups the names of the studies in the TableName analysis under appropriate headings. Table 6 shows an example layout with two subgroups, however, there is no limit to the number of subgroups.

Table 6: Layout of the MASubGroup function Table parameter

Subgroup 1 Label	Subgroup 2 Label
Study 1 Name	Study 2 Name
Study 4 Name	Study 3 Name
	Study 5 Name
	Study 6 Name

3. Subgroup results are reported in the forest plot and tabled outcomes (see below).
4. Examples of the use of this function are in the SchizophreniaPrev and MagnesiumSubGroups example files, see the section on the examples below.

MAPooledEffect

MAPooledEffect(TableName): returns the pooled effect size of the studies associated with the MAInputTable function called 'TableName'.

Note:

1. When the meta-analysis is of prevalence with 2 or more categories, this function takes as a second parameter the category number (1,2...). For more information on the meta-analysis of prevalence, see the section on this below.

MAPooledEffectLCI & MAPooledEffectHCI

MAPooledEffectLCI(TableName): returns the lower confidence interval of the pooled effect size of the studies associated with the MAInputTable function called 'TableName'.

MAPooledEffectHCI(TableName): returns the higher confidence interval of the pooled effect size of the studies associated with the MAInputTable function called 'TableName'.

Notes:

1. The size of the confidence interval can be set through the MetaXL menu item 'Options'. Default is 95%.
2. When the meta-analysis is of prevalence with 2 or more categories, these functions take as a second parameter the category number (1,2...). For more information on the meta-analysis of prevalence, see the section on this below.

MAForestName

MAForestName(TableName, Forest plot title): returns the Forest plot title parameter of the forest plot associated with the MAInputTable function called 'TableName'.

Notes:

1. This function can be used to override the default title of the forest plot. Default is the TableName of the linked MAInputTable function. To override the default is particularly useful for multiple category prevalence output, where by default all forest plots get the same name.
2. When the meta-analysis is of prevalence with 2 or more categories, this function takes as a second parameter the category number (1,2...). For more information on the meta-analysis of prevalence, see the section on this below.

MACochranQ

MACochranQ(TableName): returns Cochran's Q statistic for heterogeneity of the studies associated with the MAInputTable function called 'TableName'.

MAISquare

MAISquare(TableName): returns the I^2 statistic for heterogeneity of the studies associated with the MAInputTable function called 'TableName'.

MATauSquare

MATauSquare(TableName): returns the τ^2 statistic for heterogeneity of the studies associated with the MAInputTable function called 'TableName'.

MAQIndex

MAQIndex(TableName): returns Q-index statistic of the Quality Effects model of the studies associated with the MAInputTable function called 'TableName'.

Note:

1. This function will only return the Q-index of the meta-analysis associated with the MAInputTable function if this meta-analysis uses the Quality Effects model. In all other cases, the Q-index is undefined, and this function will return #NUM!.
2. See the section on Meta-analysis methods below for more information on the Quality Effects model and the Q-index.

MANumStudies

MANumStudies(TableName): returns the number of studies associated with the MAInputTable function called 'TableName'.

MetaXL menu and features

Introduction

In this section we will describe the MetaXL menu, and simultaneously explain the various features the menu gives access to. The MetaXL menu is located in the main Excel menu for Excel version 2003 and earlier, and in the MetaXL ribbon tab in version 2007 and later. It has the following entries.

Input templates

This menu item gives access to templates for the Table parameter of the MAInputTable function, see the description of this function above. As discussed there, this function requires a specific layout and content of its input Excel range in order to work correctly. When chosen, this menu item displays the following list of templates:

1. Binary numbers (IOType numOR, numRR, and numRD)
2. OR, RR or RD and confidence interval (IOType RRCI, ORCI, and RDCI)
3. RD and standard error (IOType RDSE)
4. Continuous numbers (IOType WMD, Cohen, Hedges, and Glass)
5. Continuous and confidence interval (IOType ContCI)
6. Continuous and standard error (IOType ContSE)
7. Prevalence numbers (IOType Prev)
8. Rate numbers (IOType NumRate)
9. Rate and standard error (IOType RateSE)

The IOType parameters are described in Table 2 above. The user can choose the appropriate layout template and Copy it to the clipboard, which allows pasting it into Excel using the Paste button or Ctrl-V.

A checkbox determines whether the template will contain a field for a quality index, to be used with the Quality Effects model. Default is on. The template entry for Prevalence has an option to set the number of categories. Default is 1.

Results

While MetaXL has several functions to display selected results of the meta-analysis in Excel itself, full results are only available through this menu item. They are displayed in a pop-up window, with tabs to give access to tabled results, a forest and a funnel plot, the option to exclude one or more studies and recalculate results, and a sensitivity analysis.

When more than one valid MAInputTable function is present in the workbook, choosing this menu item will present a Result list of available analyses. You can have as many results windows simultaneously open as you like.

Tabled results

The results table shows the list of input studies with their effect sizes, confidence intervals, and weight in the meta-analysis, together with the pooled result, the heterogeneity statistics I^2 (with CI) and Cochran's Q (with p-value), and, if the quality effects model is used, the Q-index (see the section on Meta-analysis methods below for an explanation of the Q-index).

The displayed number of decimals (3 is default) and the size of the confidence interval (95% is default) can be set in the Options menu item (see below). Results can be saved to comma-delimited file (*.csv, readable by Excel and many other programs) or copied to the clipboard (right-click, or choose Options) and pasted into Excel or Word. Note that you can adjust column widths and row heights using the mouse. The table below shows the quality effects model results from the Omega3 example workbook.

Number	Name	MD	LCI 95%	HCI 95%	weight (%)
1	Alekseeva 2000	-0.400	-1.066	0.266	1.590
2	Borkman 1989	-0.150	-0.953	0.653	2.969
3	Connor 1993	-2.200	-9.749	5.349	2.024
4	Goh 1997	-0.860	-1.677	-0.043	4.479
5	Hendra 1990	-1.200	-1.851	-0.549	5.320
6	Jain 2002	-0.390	-0.709	-0.071	8.081
7	Luo 1998	-0.150	-1.066	0.766	5.614
8	McGrath 1996	-0.200	-0.680	0.280	7.252
9	McManus 1996	-0.620	-1.758	0.518	5.055
10	Morgan 1995	-2.250	-5.409	0.909	3.117
11	Mostad 2006	0.010	-0.733	0.753	8.098
12	Pelikanova 1992	0.070	-0.617	0.757	3.319
13	Petersen 2002	-0.500	-1.571	0.571	3.869
14	Puhakainen 1995	-0.580	-2.153	0.993	3.413
15	Schectman 1988	-0.600	-1.270	0.070	3.386
16	Silvis 1990	-0.360	-1.008	0.288	1.622
17	Sirtori 1997	-0.600	-0.824	-0.376	22.419
18	Woodman 2002	-0.170	-0.597	0.257	8.373
	Pooled	-0.525	-0.764	-0.285	100.000
	I-squared	0.542	0.000	50.239	
	Cochran's Q	17.093			
	Chi2, p	0.448			
	Q-Index	46.575			

Forest plot

The figure below shows the forest plot from a subgroup analysis by study size of the quality effects model in the Magnesium example workbook. It shows, in addition to the graph, for each study and the pooled effect the effect size and confidence interval, plus the weight of each study in the pooled effect size.

The forest plot can be saved to disk (as a Windows extended metafile, *.emf) or to the clipboard (right-click and choose 'Copy', or choose 'Copy' from the menu), and then pasted into other applications.

Forest plot options

Right-clicking the graph or choosing 'Options\Plot options' from the menu gives access to the Forest plot options form, where you can set many graph properties. Some of these properties can be saved to become your preferred defaults. Table 7 below displays the various options by section of the Forest plot options window, and indicates which ones can be saved.




Table 7: Forest plot options

Section	Item	Comment	Saved
Colours	Panel	The panel of the graph is the area outside the graph area itself, and contains the title (if shown) and labels. Default is white.	✓
	Back	The colour of the graph area itself. Default is white.	✓
	Square	The studies' squares. Default is blue.	✓


Section	Item	Comment	Saved
	CI Line	The studies' confidence interval lines. Default is blue.	✓
X-axis	Auto	By default, the X-axis is scaled automatically. You need to uncheck this box in order to change the scaling.	
	Minimum, Maximum	By default, the minimum and maximum are set to the lowest and highest values of the confidence intervals of the included studies, respectively. When a confidence interval exceeds your chosen X-axis limit, it will display an arrow head.	
	Ticks	This allows changing the number of ticks. Default is 5.	
Y-axis	Spacing	This allows changing the vertical spacing between the studies in the graph.	✓
	Right margin	The right margin of the plot shows by default the studies' effect size, confidence interval, and weight. This option allows changing the space for this information, which depends on whether you want to display the information, the font size, and your screen resolution.	✓
	Diamond height	The diamond in the graph displays the pooled effect size and confidence intervals. Its width is determined by the confidence intervals, but you can change its height.	✓
Fonts	Titles and labels	All text items in the plot can individually be set to any available font type, style, size, and colour.	
Miscellaneous	Show title	Determines whether the name of the meta-analysis will show as the title of the graph. Default is on.	
	Grid lines	Ticking this box will cause the graph to display grid lines. Default is off.	
	Filled diamond	When ticked, the overall effect diamond is solid, otherwise only the outline shows. Default is on.	✓
	Arrow size	Size of the arrow head that is displayed when confidence interval exceeds the X-axis scale.	✓
	Square size	Sets the size of the study squares (this retains the proportional differences)	✓
	Weights	Determines whether the right margin displays the weight of each study in the pooled effect size. Default is on.	
	Effect size and CI	Determines whether the right margin displays the effect size and confidence interval of each study. Default is on.	

Section	Item	Comment	Saved
	ES ln transform	With the Output option for OR and RR set to log scale (see below), this option determines whether the right margin effect size and confidence interval and tabled results are also on the log scale. Default is off.	✓
	CI line width	Determines the width of the studies' confidence interval width. Default is 2.	✓
Effect size text		Allows setting the displayed effect size name in the graphs and associated tables. Primarily useful when MetaXL displays the generic ES (for effect size).	

When choosing 'File|Save settings' in this form the currently chosen options for which a check mark is in the 'saved' column of table 6 become the new defaults. This menu also allows restoring the original default values.

Funnel plot

The figure below shows the funnel plot from the Inverse variance rate ratio method in the Magnesium example spreadsheet. Again, the funnel plot can be saved to disk or to the clipboard, see the discussion of the forest plot above.



Funnel plot options

Right-clicking the graph or choosing 'Options|Plot options' from the menu gives access to the Funnel plot options form. Table 8 below displays the various options by section of the Funnel plot options window, and indicates which ones can be saved.

Table 8: Funnel plot options

Section	Item	Comment	Saved
Colours	Panel	The panel of the graph is the area outside the graph area itself, and contains the title and labels. Default is white.	✓
	Back	The colour of the graph area itself. Default is white.	✓
	Dot	The studies' dots. Default is red.	✓
X-axis	Ticks	This allows changing the number of ticks. Default is 5.	
Y-axis	Ticks	This allows changing the number of ticks. Default is 5.	
	Scale	There are three options for the Y-axis: Standard error (the default), Precision, and Inverse variance. For the relative merits of each, see Sterne & Egger (Sterne and Egger 2001).	✓
Fonts	Titles and labels	All text items in the plot can individually be set to any available font type, style, size, and colour.	
Miscellaneous	Show title	Determines whether the name of the meta-analysis will show as the title of the graph. Default is on.	
	Grid lines	Ticking this box will cause the graph to display grid lines. Default is off.	
	Tip labels	When ticked, study names are displayed in the graph. Default is off.	
	Dot size	Size of the study dots. Default is 3.	✓
Effect size text		Allows setting the displayed effect size name in the graphs and associated tables. Primarily useful when MetaXL displays the generic ES (for effect size).	

Again, some of the currently chosen options can be saved to become preferred defaults, and the original defaults can be restored.

Exclude

This tab shows the list of studies in the analysis, each with a check box next to it. When the check box is checked, the study is included in the analysis. You can uncheck studies and recalculate results by clicking the button. With this feature you can investigate which studies are the prime determinants of the pooled result, and which are the main source of heterogeneity.

Sensitivity

In this tab a table is displayed from a sensitivity analysis that excludes one by one each study in the analysis. The table shows the pooled effect sizes and the associated heterogeneity statistics. This feature gives a quick indication which study is the prime determinant of the pooled result, and which is the main source of heterogeneity. Note that you can adjust column widths and row heights using the mouse.

Options

The Options menu item is organised in two tabs: Calculation and Output. The Calculation tab lets you set the following values:

1. Size of the confidence interval of the output. Size can range from 50-99%. Default is 95%.
2. Size of the confidence interval of the input for the IOTypes that take a confidence interval as input. Size can range from 50-99%. Default is 95%.
3. The value for the continuity correction (default is 0.5). Choices are, in addition to 0.5, 0.0, 0.25, 0.75, and 1. See the section on Continuity correction below for details.
4. The transformation used when meta-analysing prevalences. Options are 1) no transformation, 2) the logit transformation, and 3) the double arcsine transformation. For more on these transformations, see the section on Meta-analysis of Prevalence below. Default is double arcsine, and users are strongly advised to keep that option.
5. Normalise prevalence. The prevalence transformations may cause a multi-category prevalence pooled result to not add to 1. This option normalises the pooled prevalences such that they do. See the section on Meta-analysis of Prevalence below for more information. Default is on.

The Output tab gives access to the following features:

1. The number of decimals in the results table (default is 3) and the forest plot (default is 2).
2. The size of the font used in the results tables. Default is 8, range is 2-16.
3. The output for relative risks (RR) and odds ratios (OR) can be set to be on the natural scale or the (natural) log scale. Default is the log scale.
4. Graph titles: you can set whether the forest and funnel plots will show the meta-analysis name as the graph title. Default is on. You can override the choice made here through the Options of an individual plot.

All chosen options can be saved to become the preferred defaults by choosing File|Save settings. You can also restore the original defaults through this menu.

MetaXL User Guide

Choosing this item displays this document in a window. You can also consult this document through the Windows start menu MetaXL entry.

Examples

This item displays an Open file dialog where you can choose from the example workbooks that come with the MetaXL installation. You can also open these examples through the Windows start menu MetaXL entry. Please note that the example workbooks are not protected: any changes you make can be saved.

Reset

Clicking this entry makes MetaXL to delete and recalculate all results in the current workbook. This is useful in at least two circumstances:

1. When you rename a worksheet with a MAInputTable function in it, the function will display the “Duplicate MAInputTable function names” error. Clicking “Reset” makes the error go away.



2. After closing a workbook with MAInputTable functions in it, the calculated results remain present in memory and are listed when choosing “Show Results”. Clicking “Reset” deletes the results from memory.

Please note that normally the MetaXL functions recalculate results just like any native Excel function: when Excel’s Recalculation option is in the default “automatic” position, any change in the input data range will immediately be reflected in the results.

About MetaXL

This menu item displays the About box, with details of the MetaXL version and a link to the Epigear website (www.epigear.com).

Meta-analysis methods

Introduction

The aim of this section is to give some background to the special features of MetaXL, including references to the literature, and help the user to make full and appropriate use of these features. We first discuss ways to deal with heterogeneity between studies, and in particular the quality effects model (which is unique to MetaXL) and why we think it is superior to the random effects model. Next we discuss the meta-analysis of prevalence, where (as far as we know) MetaXL is unique in that it provides methods for meta-analysing multiple-category prevalences. And in a third sub-section we discuss the way MetaXL deals with zero counts, the so-called continuity corrections.

One of the first things to realise is that, from a statistical point of view, meta-analysis is a very simple procedure: nothing more than taking the weighted average of a number of study results. The various methods differ in the way they assign individual study weights and how they apply them, but once that is settled, the pooled result is simply the weighted average.

This observation is certainly not meant to belittle meta-analysis: it is an invaluable and increasingly important tool, and a prime source of information for evidence based medicine and health economic evaluation. But the hard work in meta-analysis is in the validly collecting, digesting, and rating of often very heterogeneous studies. Once that is done, the statistical pooling of the results is almost a breeze.

Various models

While meta-analysis is essentially taking a weighted average of study results, there are different ways to take a weighted average. In the literature the main distinction is between fixed-effects and random-effects models (Deeks, Altman et al. 2001). The fixed-effects model is sub-divided into three methods: inverse variance, Mantel-Haenszel, and Peto's method for odd ratios. The random-effects model (see below) is a modification of the inverse variance fixed-effects model. For these methods and models, MetaXL implements the equations as outlined in the chapter by (Deeks, Altman et al. 2001).

MetaXL is unique in that it also implements a quality-effects model. The reason for this is that the random-effects model is an unsatisfactory answer to the problem it tries to address. We therefore propose the quality-effects model as a better alternative to the random-effects model. As the random-effects model, the quality-effects model is a modification of the fixed-effects inverse variance method. Below we first describe the random-effects model and the way it tries to address the heterogeneity problem, next discuss the quality-effects model, and finally compare the two.

The random effects model

In conducting a meta-analysis, you must conventionally choose between a fixed- or random-effects model. The two models make very different assumptions about the data. The fixed-effects model assumes that all of the studies are estimating a common population effect-size and that the differences between the studies are simply a function of sampling error (Lipsey and Wilson 2001; Senn 2007). This is rarely plausible with the possible exception of a meta-analysis of pure replications.

A random-effects model assumes two sources of variability in effects, one from sampling error and one from study level differences. However the latter are based on the variability in effect sizes across the group of studies and it has been taken to imply that the random-effects model assumes a distribution of true population effects from which the observed studies are sampled (Lipsey and Wilson 2001; Senn 2007). The problem with this assumption is that a more likely cause of differences in effect size between trials is systematic error, and such sources of bias are not addressed appropriately by the random effects meta-analysis model (Conn and Rantz 2003). Obviously, if a random variable is inserted to inflate the variance based on heterogeneity, it is not clear what aspect of between-trial differences is being assessed and fails to take into account meaningful differences between the individual studies. Also, because of this limitation of the random effects model, when used in a meta-analysis of badly designed studies, it will still result in bad statistics even though there is some statistical adjustment for heterogeneity. Furthermore, if we look critically at computational aspects of the random model, we immediately realize that the random effects meta-analysis is simply a process whereby the inverse variance weighting of the fixed model is reversed (to a variable extent) thus moving the weighted mean effect size back towards an unweighted mean. The extent of this reversal is solely dependent on two factors (Al Khalaf, Thalib et al. 2011):

1. Heterogeneity of precision (study size): The extent of the spread of precision of the studies involved as indicated by the maximum minus minimum inverse variance weights
2. Heterogeneity of effect size (τ^2): How many times bigger than the average variance of studies within the meta-analysis is the value of τ^2 . (τ^2 alone is not comparable across different meta-analyses).

This model reduces the extreme diminution of the effect size of the smaller studies so that the pooled effect size now moves towards an intermediate value between trials with extremes of precision (Helfenstein 2002). This movement backwards towards an unweighted mean, however, is based on penalizing larger studies based on sample size and effect size (τ^2) heterogeneity. Is the random model analysis now any more valid than a fixed effects analysis when heterogeneity is present? Unfortunately the answer is *no* because there is no reason to automatically assume that a larger variability in study sizes or effect sizes (meta-analysis τ^2) automatically indicates a faulty larger study or more reliable smaller studies. Indeed, there is no reason why the conclusiveness of a meta-analysis should be associated with the random model method of reversal of the inverse variance weighting process of the included studies. As such the changes in weight introduced by this model to each study have no statistical or probabilistic interpretation and thus bear no relationship to what the studies actually have to offer.

Heterogeneity between trials can occur for a multitude of reasons some of which include chance, different definitions of treatment effects, credibility related heterogeneity (quality), and finally unexplainable and real differences (Bailey 1987). Credibility related heterogeneity (quality) is an important quantifiable difference and this refers to the likelihood of the trial design to generate unbiased results that are sufficiently precise and allow application in clinical practice (Verhagen, de Vet et al. 2001). Naturally, the flaws in the design of individual studies can be expected to create heterogeneity between trials as well as influence the magnitude of the meta-analysis results. If the quality of the primary material is inadequate, this may falsify the conclusions of the review, regardless of the presence or absence of effect size

heterogeneity. This need for addressing heterogeneity in trials via study specific assessment has been obvious for a long time and the solution involves much more than the random model strategy of just inserting a random term based on effect size heterogeneity (Senn 2007). We thus developed the QE model that is implemented in MetaXL.

Rationale for the QE model

In a group of homogenous trials, it is assumed that because the effect sizes are homogenous, the studies are all estimating the same target effect (we can call this a type A trial). In this situation, the inverse variance weights of Woolf (Woolf 1955) will minimize the variance since:

$$MSE = Expected(Estimate - True)^2 = Variance + (Bias)^2$$

Bias is zero if the underlying true effect size's are equal and thus minimizing variance is optimal and the weighted MSE = Variance. It is thought that the inverse variance weighted analysis tests the null hypothesis that all studies in the meta-analysis are identical and show no effect of the intervention under consideration regardless of homogeneity (Senn 2007). This of course requires the assumption that trials are exchangeable so that if one large trial is null and multiple small trials show an effect, the large trial essentially confirms that the null hypothesis is true. Exchangeability however is a big assumption and therefore if we do not believe the trials are exchangeable then in this situation we have two alternatives: Either the trials have been affected by bias even though underlying true effects are identical (we can call these type B trials) or the trials represent different underlying true effects (we can call these type C trials). In the former case, the trial effect size from a biased trial might seem like it is coming from a different underlying true effect, thus giving the same impression as type C trials: that the trials represent different underlying true effects. In both types B & C trials, inverse variance weights will not minimize the variance as it will just exaggerate and create gross bias in these situations. Furthermore, any set of weights in a type A situation estimates the same target, but in a type B or C situation each set of weights estimates a different target. Thus inverse variance weights in the latter situation just increases bias and is not optimal for type B or C trials. Thus in type B or C trials we would want to use situation specific weights.

One such situation specific weight that has been suggested for type B trials is weighting according to the probability (Q_i) of credibility (internal validity or quality) (Tritchler 1999; Doi and Thalib 2008; Doi and Thalib 2009; Doi, Barendregt et al. 2011) of the studies making up the meta-analysis. Although this can correct for distortions due to systematic error, it can also introduce errors of another type. For example a study of a small sample, which is not representative of the underlying population, may get undeserved weight from a large underlying population and this can skew the data. It might thus be informative to weigh according to precision and then re-distribute the weights according to situation specific requirements. In this case importance of smaller studies will only get upgraded if the larger or more precise study is deemed poor by its situation specific weight.

This line of thought is not new as this is precisely what the random effects model attempts to do (DerSimonian and Laird 1986). The unfortunate thing however is that the "situation specific" weight used in this particular model is an index of the variability of the effect sizes across trials and the same "situation specific" weight is

applied to all trials. This creates two problems, first that these weights are meaningless and second that they are not study specific. In order to rectify this situation, an alternative approach has been proposed in 2008 (Doi and Thalib 2008; Doi and Thalib 2009). In the QE model, both of the draw-backs of the random model type of re-distribution have been addressed.

One further consideration, in type C trials which deal usually with burden of disease where true differences across populations are expected, is the fractional population represented by the effect size in each study. Of course, a study of, for example, 1000 respondents is equally useful in examining the mortality in a country with 10 million inhabitants as it would be in a country with a population of only 1 million. Without weighting, any figures that combine data for two or more countries would over-represent smaller countries at the expense of larger ones. So a population size weight is needed to make an adjustment to ensure that each country risk is represented in the pooled estimate proportional to its population size. Although such weighting has been attempted previously, it has been improperly applied (Batham, Gupta et al. 2009). The best method is to assign a proportional weight between zero and 1 of each study in relation to the largest based on the underlying population size. The population size weight (Pweight) is thus the proportional weight $Psize_i/Psize_{max}$. This corrects for the fact that most studies may have very similar sample sizes, no matter how large or small their underlying population, These weights can be multiplied by internal validity weights (also between zero and 1) to yield modified internal validity weights that can then be utilized in the QE model.

Differences in weighting between models

The standard approach frequently used in meta-analysis in clinical research is termed the 'inverse variance method' or FE model based on Woolf (Woolf 1955). The average effect size across all studies is computed, whereby the weights are equal to the inverse variance of each study's effect estimator. Larger studies and studies with less random variation are given greater weight than smaller studies. The weights (w) allocated to each of the studies are then inversely proportional to the square of the standard error (se), thus for the i th study

$$w_i = \frac{1}{se_i^2}$$

which gives greater weight to those studies with smaller standard errors.

As can be seen above the variability within each study is used to weight each studies effect in the current approach to combining them into a weighted average as this minimizes the variance (assuming each study is estimating the same target). So if a study reports a higher variance for its effect size estimate it would get lesser weight in the final combined estimate and vice versa. This approach however does not take into account the innate variability that exists between the studies arising from differences inherent to the study such as their protocols and how well they were executed and conducted. This major limitation has been well recognized and it gave rise to the random effects (RE) model approach (DerSimonian and Laird 1986). Here, a constant is generated from the homogeneity statistic Q and using this and other study parameters a random effects variance component (τ^2) is generated. The inverse of the sampling variance plus this constant that represents the variability across the population effects is then used as the weight

$$w_i^* = \frac{1}{se_i^2 + \tau^2}$$

where w_i^* is the random effects weight for the i th study.

However, because of the limitations of the RE model, when used in a meta-analysis of badly designed studies, it will still result in biased estimates even though there is statistical adjustment for effect size heterogeneity (Senn 2007). Furthermore, such adjustments, based on an artificially inflated variance, lead to a widened confidence interval, supposedly to reflect effect size “uncertainty”, but do not have much clinical relevance (Senn 2007; Al Khalaf, Thalib et al. 2011). Therefore, a new method for adjusting for heterogeneity can use situation specific parameters and in this case we use study specific scores. These are then rescaled to a probability Q_i , the judgement of the probability (0 to 1) that study i is credible or the proportion (0-1) that study i patients make up of the total population of interest and so on. From Q_i a study specific composite called \hat{Q}_i is generated that takes into consideration study specific information and its relationship to other studies to re-distribute inverse variance weights. As studies increase in quality or are more representative of the population, re-distribution becomes progressively less and ceases when all studies are of perfect quality or representation. This is totally different from the direct adjustment for quality of fixed effects weights reported by several authors (Berard and Bravo 1998; Trichler 1999; Leeflang, Reitsma et al. 2007). With the QE model, rather than use study specific scores directly, the composite \hat{Q}_i is generated (Doi and Thalib 2009).

The QE weight, w_i^δ , is then:

$$w_i^\delta = \frac{\hat{Q}_i}{se_i^2} = \hat{Q}_i w_i$$

Detailed aspects of computation of this model have been published with an important update for overdispersion using a correction factor for individual study variances (Doi, Barendregt et al. 2011).

The extent of redistribution of these weights due to non-credibility can be assessed via a Q-index defined as follows:

$$\bar{Q} = \sum_1^n \left(\frac{(1 - Q_i) \times w_i}{\sum_1^n w_i} \right)$$

The rationale for this index is that it tells us the probability that the information contained within the studies in the meta-analysis is non-credible. The higher the Q-index, the more likely that the information contained is non-credible and thus, as Q-index increases, the weights in the meta-analysis come together and the confidence interval of the pooled estimate increases.

Meta-analysis of Prevalence

Based on a draft article of the same name, authored by: Jan J Barendregt, Suhail A Doi, Yong Yi Lee, Rosana Norman, Theo Vos.

Introduction

The large majority of meta-analyses are devoted to establishing the effects of interventions, and therefore aim to get a pooled estimate of effect size, be that relative risk, odds ratio, risk difference, or weighted or standardised mean difference. However, meta-analysis methods can be useful as well to get a more precise estimate of disease frequency, such as disease incidence rates and prevalence proportions. This section looks at the meta-analysis of prevalence, which is a rather less developed field than that of intervention effects. MetaXL introduces some new methods, which we describe in this section. We first discuss the specific properties of prevalence as a variable and the options to deal with these in the meta-analysis, next look at the multiple category case, and discuss the implementation of these methods in the MetaXL software.

Prevalence as a variable

While just about all epidemiologists will habitually speak of ‘prevalence rate’, prevalence is defined as a proportion: the number of cases in a population, divided by the population number. This definition implies that 1) prevalence is always between 0 and 1 (inclusive), and 2) the sum over categories always equals 1.

The definition of prevalence is the same as the definition of the binomial distribution (number of successes in a sample), and therefore the standard assumption is that prevalence follows a binomial distribution. With the main meta-analysis methods based on the inverse variance method (or modifications thereof), the binomial equation for variance (expressed as a proportion) can be used to obtain the individual study weights:

$$Var(p) = \frac{p(1-p)}{N} \quad (1)$$

Where p is the prevalence proportion, and N the population size.

With the variance of the individual studies nailed down, the pooled prevalence estimate P then becomes (according to the inverse variance method):

$$P = \frac{\sum_i \frac{p_i}{Var(p_i)}}{\sum_i \frac{1}{Var(p_i)}} \quad (2)$$

With standard error:

$$SE(P) = \sqrt{\sum_i \frac{1}{Var(p_i)}} \quad (3)$$

The confidence interval of the pooled prevalence can then be obtained by:

$$CI_{\gamma}(P) = P \pm Z_{\alpha/2} SE(P) \quad (4)$$

where $Z_{\alpha/2}$ denotes the appropriate factor from the standard Normal distribution for the desired confidence percentage (e.g. $Z_{0.025}=1.96$).

Transformations

While this works fine for prevalence proportions around 0.5, increasing problems arise when the proportions get closer to the limits of the 0..1 range. The first problem is mostly cosmetic: the equation for the confidence interval does not preclude confidence limits outside the 0..1 range. While this is annoying, the second problem is much more substantial: when the proportion becomes small or big, the variance of the study is squeezed towards 0 (see eq 1). As a consequence, in the inverse variance method, the study gets a large weight. A meta-analysis of prevalence according to the method described above therefore puts undue weight on the studies at the extreme of the 0..1 range.

The way to deal with these problems is to transform the prevalence to a variable that is not constrained to the 0..1 range, has an approximately normal distribution, and by being unconstrained avoids the squeezing of variance effect. The meta-analysis can be carried out on the transformed proportions, using the inverse of the variance of the transformed proportion as study weight. For final presentation, the pooled transformed proportion and its confidence interval are back transformed to a proportion. We discuss two transformations that have been used: the logit and the double arcsine.

Logit

The logit transformation is given by (Lipsey and Wilson 2001):

$$t = \ln\left(\frac{p}{1-p}\right) \quad (5)$$

with variance:

$$Var(t) = \frac{1}{Np} + \frac{1}{N(1-p)} \quad (6)$$

The back transformation to a proportion is done using:

$$p = \frac{\exp(t)}{\exp(t) + 1} \quad (7)$$

However, while the logit transformation solves the problem of estimates falling outside the 0..1 limits, unfortunately, it does not succeed in stabilizing the variance. Rather, there is a reversal of the variance instability of the non-transformed proportions: studies with proportions close to 0 or 1 get their variance estimates grossly magnified, while the variance for proportions around 0.5 is squeezed. So a large study with a small prevalence gets completely outweighed by a small study with

prevalence around 0.5. The variance instability that plagues non-transformed proportions thus persists even after logit transformation.

Double arcsine

The double arcsine transformation is given by (Freeman and Tukey 1950):

$$t = \sin^{-1} \sqrt{\frac{n}{N+1}} + \sin^{-1} \sqrt{\frac{n+1}{N+1}} \quad (8)$$

with n the number of people in the category. The variance of t is given by:

$$\text{Var}(t) = \frac{1}{N+0.5} \quad (9)$$

The back transformation to a proportion is done using (Miller 1978):

$$p = 0.5 \left\{ 1 - \text{sgn}(\cos t) \left[1 - \left(\sin t + \frac{(\sin t - \frac{1}{\sin t})}{N} \right)^2 \right]^{0.5} \right\} \quad (10)$$

A much simpler but less accurate alternative back transformation is:

$$p = \left(\sin \frac{t}{2} \right)^2 \quad (11)$$

The double arcsine transformation addresses both the problem of confidence limits outside the 0..1 range and that of variance instability, and it is therefore preferred over the logit transformation.

Multi-category prevalence

The discussion so far has implicitly been about a single category (which of course implies a second category of those not in the first). But in some instances k -category prevalences may be meta-analysed where $k > 2$ (eg mild, moderate and severe disease), and this complicates matters.

The first thing to mention is that a multi-category prevalence from a study is a single (vector-valued) observation, which suggests that a single study weight across all categories applies. This is no problem in the case of the double arcsine transformation, because its variance depends on the population size (N) only. But for the untransformed and logit transformation options the variance of both p and $\text{logit}(p)$ depends on p itself, and this implies that the same study would get a different weight in each category, which seems hard to justify.

A second issue is the calculation of the heterogeneity statistic Cochran's Q . With multiple categories, the standard calculation for Q will return a different value for each category, again something that is hard to justify. Moreover, Q is used for further calculations, in the random effects model for τ^2 , and in the quality effects model for the overdispersion correction (see the section on the quality effects model above).

The third issue is that because of the non-linearities introduced by the logit and double arcsine transformations, the back-transformed proportions will not necessarily add up to 1. We discuss solutions to these three issues in the next section.

The MetaXL implementation

MetaXL implements three variants of meta-analysis of prevalence: untransformed, logit transformed, and double arcsine transformed prevalence, with the double arcsine transformation as the default option. The user can specify as many categories as desired, and can choose between fixed effects (inverse variance only), random effects, and quality effects models.

In some cases equation 10 can become numerically unstable, and to avoid this, the double arcsine back transformation is implemented as follows.

Let:

$$\bar{s} = \left(\sin \frac{\bar{t}}{2} \right)^2 \quad (12)$$

Where \bar{t} is the pooled t . Then:

$$\bar{P} = \begin{cases} \bar{s} & , \text{if } \frac{\bar{s}}{\bar{v}} < 2 \\ \bar{s} & , \text{if } \left(\frac{1 - \bar{s}}{\bar{v}} \right) < 2 \\ 0.5 \left\{ 1 - \text{sgn}(\cos \bar{t}) \left[1 - \left(\sin \bar{t} + \frac{\left(\sin \bar{t} - \frac{1}{\sin \bar{t}} \right)}{\frac{1}{\bar{v}}} \right)^2 \right]^{0.5} \right\} & , \text{otherwise} \end{cases} \quad (13)$$

Where \bar{P} is the pooled prevalence and \bar{v} is the pooled variance. Instead of the harmonic mean, as suggested by Miller (Miller 1978), we use $1/\bar{v}$ because our simulation studies suggest that this is a more stable estimate for N than the harmonic mean. This works because the variance of a double arcsine transformed proportion is approximately $1/N$

The lower (LLC) and upper (ULC) limit are given by:

$$LLC = \begin{cases} 0 & , \text{if } \frac{\bar{P}}{\bar{v}} < 2 \\ 0.5 \left\{ 1 - \text{sgn}(\cos \bar{t}) \left[1 - \left(\sin \bar{t} + \frac{\left(\sin \bar{t} - \frac{1}{\sin \bar{t}} \right)}{\frac{1}{\bar{v}}} \right)^2 \right]^{0.5} \right\} & , \text{otherwise} \end{cases} \quad (14)$$

and

$$ULC = \begin{cases} 1 & , \text{if } \left(\frac{1 - \bar{P}}{\bar{v}} \right) < 2 \\ 0.5 \left\{ 1 - \text{sgn}(\cos \bar{t}) \left[1 - \left(\sin \bar{t} + \frac{\left(\sin \bar{t} - \frac{1}{\sin \bar{t}} \right)}{\frac{1}{\bar{v}}} \right)^2 \right]^{0.5} \right\} & , \text{otherwise} \end{cases} \quad (15)$$

Where $\bar{t} = \bar{t} \pm Z_{\alpha/2} \sqrt{\bar{v}}$.

The three issues caused by generalising to more than two categories are dealt with as follows. To obtain a single study weight across categories for the untransformed and logit transformed prevalences, MetaXL uses the average of the category weights.

For the second issue, to obtain a single estimate for Cochran's Q, we have chosen to take the maximum category Q. This because if we believe that the effect size variations across studies in one category of proportions is not independent of variations in the other categories, the maximum category Q value would be the best and most conservative estimate of a common Q.

For the random effects model, a single τ^2 is calculated in the usual way from this common Q, and this τ^2 is then added to the variance of each study in each category to obtain the category standard error, while it is added to the average category variance to obtain a single study weight.

In the quality effects model, the redistribution of inverse variance weights is done using a quality parameter individualized to each study rather than the common τ^2 parameter across studies as is done with the random effects model (Doi and Thalib 2008; Doi and Thalib 2009). In addition to this, the quality effects model applies an overdispersion correction based on Cochran's Q chi-squared statistic, resulting in a more conservative confidence interval in the presence of heterogeneity of effect sizes (Doi, Barendregt et al. 2011). Note that, unlike in the random effects model, this increase does not affect the study weights. For the multi-category prevalence analysis, the overdispersion correction is applied to each of the categories separately, thus boosting the CIs in case of heterogeneity.

The third and final issue is that, due to the transformations, the back transformed proportions may not sum to 1. The difference is usually small, and MetaXL deals with it by offering the option to normalise the prevalences in each category after pooling and back transformation. The confidence intervals, however, are not adjusted.

Conclusion

From the discussion above, we conclude that the double arcsine transformation is the preferred way to deal with the meta-analysis of prevalence. Consequently, this transformation is the default one, but the user can change this through the Options menu. Because this transformation may lead to multi-category prevalences that do not add up to 1, the normalisation option is also the default.

Continuity corrections

Binary studies

In some instances the outcome of interest may be very rare (or in some circumstances very common) leading to zero events occurring in either or both arms of a study. In such cases, the log-odds ratio and the log risk ratio become undefined (as are their variances) and hence this causes problems with the meta-analysis computations. To overcome this problem, a so called ‘continuity correction’ factor, of 0.5, is often added to each cell of the 2×2 table for the studies with zero events in either arm. Using a continuity correction for studies with zero counts allows the log-odds ratio or log risk ratio to be estimated, and hence allows synthesis via the meta-analysis methods used in MetaXL.

Yates first used the term correction for continuity in 1934 (Yates 1934). The premise for the use of 0.5 as a continuity correction is through the approximation of a discrete distribution by a continuous one (Sweeting, Sutton et al. 2004). In standard 2×2 tables, it has been argued that the use of 0.5 makes the average approximation of the error equal to zero (Cox 1970).

The default continuity correction for binary studies in MetaXL is 0.5. This correction is only applied to studies with 0 counts, other studies in the analysis are not affected. A continuity correction is not applied to Peto’s method: a zero count poses no problem in that method.

The options menu of MetaXL lets you choose the size of the continuity correction, including setting it to 0 (see the Options section above). If you do that and MetaXL encounters a study with zero counts, the MAInputTable function gives an error message, see the section on Error messages below.

Pooled prevalence

In the case of pooling prevalences (IOType Prev), a continuity correction is applied automatically for studies where the observed prevalence is either 0 (adjusted to 0.0005) or 1 (adjusted to 0.9995). The setting of the continuity correction option has no bearing on the correction used when pooling prevalences.

Documentation

Examples

Introduction

The MetaXL installation comes with a number of example Excel workbooks. All example workbooks are accessible through the MetaXL menu and the MetaXL entry in the Windows Start menu. Where they are located on your harddisk is difficult to say: it depends on your Windows version, your privileges on the PC, and possible changes to the proposed default directories that you may have made during installation. Details on this can be found in the section Installation issues of the Technical appendix of this guide.

Please note that the example workbooks are not protected: any changes that you make can be saved. If you want to preserve the original example workbooks but also experiment with them, you should first make copies of them to experiment with, for example by using 'Save as' to copy them to your standard Excel workbook directory. Of course you can always revert to the original example workbooks by re-installing MetaXL. Below is a brief description of each of the example workbooks.

ThyroidCancerRR

This meta-analysis is about high dose radio-active iodine (^{131}I) versus low dose after thyroidectomy in patients with differentiated thyroid cancer. The study looked at the risk of non-ablation, and the outcome is that high-dose patients have a 20-30% lower risk of non-ablation.

File name: ThyroidCancerRR.xls.

Magnesium

This binary numbers example gives all the methods available for this kind of data: odds ratios according to five different methods, and risk ratios and differences according to four methods each. In addition, there are sheets with the same studies but then taking effect sizes and their 95% confidence intervals as inputs, and for the risk difference effect sizes and standard error as inputs, all for the IV, RE and QE methods. Note that in the Options menu you can set the size of the input confidence interval independently from the output confidence interval. Default is 95%. Also note that the results are identical between the various input formats.

The workbook summarises results of a number of trials of intravenous magnesium after suspected acute myocardial infarction. There has been much controversy about this topic, with the smaller trials showing a beneficial effect, but the large ISIS trial showing a small but not statistically significant detrimental effect (Nuesch and Juni 2009).

In fixed effects analyses the large ISIS trial weighs-in heavily, and as a consequence the pooled effect sizes are not significant. The random effects model assigns a smaller weight to large studies in case of heterogeneity, resulting in significant pooled effect sizes. Interestingly, the quality effects model also finds significant beneficial effects, mainly because the ISIS trial does not score very well on quality.

The heterogeneity statistics Cochran's Q and I^2 indicate that there is considerable heterogeneity. In a subgroup analysis of large (ISIS and MAGIC) versus smaller (all the rest) studies the heterogeneity within the subgroups disappears. An interesting finding, but not so easy to interpret.

File names: Magnesium.xls and MagnesiumSubGroups.xls.

Omega3WMD

A weighted mean difference example. The workbook contains sheets with input in numbers (population sizes, means, and standard deviations for both arms), and with effects sizes and either confidence intervals or standard errors as inputs. In all cases results are calculated according to the IV, RE, and QE methods, and are identical across input formats (as should be).

File name: Omega3WMD.xls.

MarshallSMD

This example is taken from the chapter in Egger et al where most of the methods implemented in MetaXL are based on (Deeks, Altman et al. 2001). Three trials comparing assertive community treatment with standard care for severely mental ill. The trials used different scoring systems, and therefore a standardised mean difference was used to combine them. Note that the scale used by the Lehman study runs opposite to the other two, the mean is therefore entered as a negative number. Results are presented for all three methods for standardised mean difference: Cohen's d , Hedges' adjusted g , and Glass's Δ . For these studies, results are very similar across SMD methods.

The workbook gives results for the IV, RE, and QE methods, and for input formats with numbers, confidence intervals, and standard errors.

File name: MarshallSMD.xls.

SchizophreniaPrev

A pooled prevalence estimate of acute schizophrenia in the population with schizophrenia. This example also gives the example of a quality scoring list, that runs on a scale from 0 to 11, and how to convert the score to the 0..1 range expected by MetaXL. There is also a subgroup analysis of the studies by the income status of the countries where they were carried out.

File name: SchizophreniaPrev.xls.

MSMultipleCategoryPrev

Also a pooled prevalence example, this time of three multiple sclerosis categories "mild", "moderate", and "severe". The three categories should always add up to 1, but in case of the logit and double arcsine transformations do so only when the "Normalize prevalence" option in the options menu is set.

This example illustrates the use of the MAForestName function to override the default forest plot title which would otherwise be the same for each category. And finally, it illustrates the use of an additional parameter for the MAPooledEffect, MAPooledEffectLCI, MAPooledEffectHCI, and MAForestName functions to designate the category number the function should refer to.

File name: MSMultipleCategoryPrev.xls.

Error messages

Introduction

Generally in software, there are two types of error messages: planned ones and un-planned ones. The former kind anticipates user mistakes and tries to give an informative message, pointing the user in the right direction. The latter kind pops-up when un-anticipated circumstances occur, and the error messages of the compiler used (in this case Borland's Delphi), host application (Microsoft's Excel), or operating system (Microsoft's Windows) kick in. This kind of error "message" includes crashing of the application, for an Excel add-in like MetaXL this usually means crashing of Excel itself as well. Software developers refer to this kind of error messages as "bugs".

While we have tried to anticipate all user mistakes, experience learns that software users are far more creative in making mistakes than software developers will ever fathom. So it is likely that you will encounter both kinds of error messages. Below we first give the messages for the mistakes we anticipated, and then for the ones we did not. The former will be returned by the MAInputTable function and when you try to get results from an analysis that shows such an error message.

MAInputTable error messages

Error: unknown IOType parameter

The MAInputTable function did not recognise the IOType parameter name you specified. Please check your spelling, and table 2 above for the IOType parameters that are defined. Note that the parameter name has to be enclosed in double quotes (ie "NumOR"). The parameter name is not case-sensitive.

Error: unknown Method parameter

The MAInputTable function did not recognise the Method parameter name you specified. Please check your spelling, and table 3 above for the Method parameters that are defined. Note that the parameter name has to be enclosed in double quotes (ie "IV"). The parameter name is not case-sensitive.

Error: incompatible input and method parameters

Both the IOType and Method parameters are recognised, but not all combinations of them are valid. Please consult table 4 above to see which combinations of parameter values are compatible.

Error: too many (few) columns in table

Given the IOType and the Method parameters you specified, the MAInputTable function expects a defined number of columns in the input range parameter. This error occurs when there are too many (few) columns in this table. Please consult table 5 above to see which columns MAInputTable expects, and what should be in each. Use the Input Templates entry from the MetaXL Excel menu to set up a correct input table.

Error: duplicate MAInputTable function names

This error occurs when two or more MAInputTable functions are present with the same Name parameter. The Name parameter needs to be unique, so the remedy is

simple: rename the duplicate. Please note that this error will also occur when the MAInputTable functions with the same name are in different Excel workbooks that are open simultaneously.

This error also occurs when you rename a worksheet with MAInputTable functions in it. Remedy: go to the MetaXL menu and click 'Reset'.

Error: duplicate study names

Each study included in the MAInputTable range parameter needs to have a unique name, and this error is reported when there is at least one duplicate study name. The remedy is to make sure study names are different.

Error: invalid or missing study name

MetaXL expects studies included in the MAInputTable range parameter to have valid names. A missing name means the cell that should contain the name is empty, an invalid name is one with at least one character that offends. This error can occur when the range parameter includes an empty row. Remedy: make sure your range parameter refers to the correct range, fill empty Name cells, and do not use characters outside the alphanumeric range.

Error: continuity correction needed

This error can only occur when you have changed the default MetaXL setting of the continuity correction. See the discussion of the continuity correction option setting above.

Error: invalid negative or 0 values in input table

This error can occur with the IOType parameters ORCI and RRCI. With these IOType parameters, the MAInputTable function expects either RR or OR effect sizes with their confidence intervals (see table 5 above and the Input Templates entry from the MetaXL Excel menu for the correct inputs for this IOType). By definition, all these input fields must be > 0 , and this error occurs when one of them is not.

This error also occurs when in a continuous analysis population size numbers or standard deviations are less than or equal to 0. And when an IOType with standard error input is used with a standard error that is not more than 0.

Error: invalid negative values in input table

This error is returned by the pooled prevalence IOType Prev method when either the population number or number of cases entered is less than 0. And the binary methods (IOType NumRR, NumOR, and NumRD) return this error when the numbers entered are less than 0.

Error: quality weights must be between 0 and 1 (inclusive)

The quality weights input MetaXL expects to be in the 0..1 range. If you use a different quality scoring system, you will have to rescale it.

An unknown error occurred

An error occurred, and MetaXL did catch it, but it is at a loss what caused it. It should not be one of the errors listed above, but that is of little help. If you get this error and it is reproducible (i.e. if it occurs predictably with some workbook or action), we are



very interested in receiving a description of the problem, if possible with a copy of the offending workbook. Please send email to info@epigear.com.

Bugs

No software is without bugs, and that is undoubtedly true for MetaXL (and for Excel, for that matter). Things can go wrong, and once they have done so, the software may have become unstable and start to produce error messages whatever you do. In such cases it is often advisable to quit Excel altogether and start anew.

In particular, when MetaXL has given incomprehensible error messages like ‘Access violation’, ‘Invalid floating point operation’, or ‘Range check error’, or when Excel crashes, things have seriously gone wrong and a restart is often required. If the problem is reproducible (i.e. if it occurs predictably with some workbook or action), we are very interested in receiving a description of the problem, if possible with a copy of the offending workbook. Please send email to info@epigear.com.

Trouble shooting

Sometimes the combination of MetaXL and Excel will not behave as the user expects, with a frustrated user as a consequence. In this section we try to pre-empt some of the situations where this might occur by explaining what is going on and offering a solution. Most of this is based on user experiences in our own work environment (including ourselves). We are happy to hear of any other potential conundrums: if you have any, please email them to info@epigear.com and we will either try to solve the issue or add them to this list.

1. The MetaXL functions are not recognised by Excel (they show the #NAME? error).

If you see this error on just a single or a few MetaXL functions, while other MetaXL functions are fine, you most likely made a typo in the function name. If you are unsure how a particular function name is spelled, use the Excel Function Wizard: all MetaXL functions are available in the category 'MetaXL'.

If all MetaXL functions show this error, and you are confident that they are spelled correctly, most likely MetaXL is not properly (or not at all) installed. Check whether the MetaXL add-in is listed and active (Excel 2003 and earlier: Tools\Add-ins; Excel 2007 and later: Office Button\Excel Options\Add-ins\Go). Make sure the add-in is listed and the checkbox next to its name is checked.

If the add-in is not listed, you can try to use the 'Browse' button of the Excel Add-in Manager to add it to the list, but most likely something went wrong during installation (e.g. you had Excel running while installing MetaXL). Remedy: quit Excel, and install MetaXL.

2. The MetaXL functions are recognised by Excel, but not correctly (they show the #VALUE! error).

If you see this error on just a single or a few MetaXL functions, while other MetaXL functions are fine, you most likely made an error in the number of parameters the function takes. If you are unsure about the parameters of a particular function, use the Excel function wizard: all MetaXL functions are available in the category 'MetaXL'.

This error also occurs when you forgot to put double quotes around the text input parameters of the MAInputTable function.

Another occasion this error occurs (but this is an Excel and not a MetaXL issue) is when you use the wrong parameter separator in Excel functions: when your system or Excel is configured to have a comma as the decimal separator (as in many European countries), the parameter separator is a semicolon (;), when the decimal separator is a period, the parameter separator is a comma.

3. A MetaXL function returns #NUM!.

When the MAInputTable function returns an error, all MetaXL function linked

to it will return this error. Remedy: fix the problem with the MAInputTable function.

If this is not the case, you most likely made an error in the value of the parameters the function takes. If you are unsure about the values the parameters of a particular function can take, use the Excel function wizard: all MetaXL functions are available in the category 'MetaXL', and help is given on the parameter values they can take. Or look up the function in the MetaXL Function section above.

Known issues

1. Excel keeps focus when a MetaXL pop-up window should have it: this hampers using the keyboard to navigate the MetaXL windows.
2. When you rename a worksheet with MAInputTable functions in it, the functions report the “duplicate MAInputTable function names” error. Remedy: go to the MetaXL menu, and click ‘Reset’.
3. The “ticks” forest plot option for the x-axis scaling sometimes gives unpredictable results.

MetaXL versions

A list of the updates of MetaXL, most recent first.

Version 1.2

May, 2012.

A minor update.

1. This version introduces full support for the Excel 2007+ ribbon user interface, with a MetaXL tab in the ribbon.
2. The MetaXL menu in Excel 2003 and earlier is now positioned at the end, in an attempt to accommodate non-English versions of Excel.

Version 1.11

April, 2012.

A minor update.

1. An additional function: MAForestName, to override the default title of the forest plot.
2. With multiple category prevalence output the various output windows now cascade instead of overlap.
3. An additional example: MSMultipleCategoryPrev.
4. The installation detects whether Excel is running, and quits after announcing it should not.
5. The installation detects whether Excel 64-bit is installed, and quits after announcing that this version of MetaXL is not compatible with 64-bit Excel.

Version 1.1

February, 2012.

A major update, with several enhancements.

1. Subgroup analysis: MASubGroups function and the ability of the forest plot to display subgroups.
2. Other enhancements of the forest plot, such as options to determine the width of the confidence interval lines.
3. Additional functions to display outcomes in the Excel spreadsheet: MATauSquare and MAQIndex.
4. More options can now be saved to become user-preferred defaults.
5. Several enhancements to the meta-analysis of prevalence:
 - a. Introduction of three different options for transformations: none, logit, and double arcsine.
 - b. Introduction of multi-category prevalence meta-analysis.
 - c. A 'normalise' option to make sure multiple category pooled results sum to 1.
 - d. An additional parameter for the MaPooledEffect, MaPooledEffectLCI, and MAPooledEffectHCI functions to determine the category number.
 - e. Multiple output windows in case of multiple category prevalence.
 - f. Addition of a section to this *User Guide* on the *Meta-analysis of Prevalence*.
6. A complete rewrite of the methods used to exchange data with Excel, in preparation of planned support for 64-bit Excel.



Version 1.0

January, 2011.

The first public release.

Technical appendix

Installation issues

Where are your files?

Installation is usually painless, the most important thing to remember is that Excel should not be running when you install MetaXL. If Excel is running, you will get an error message during installation that the add-in could not be installed. Remedy: from now on read the messages installation programs display, quit Excel and try again. Once you've installed MetaXL, where are the files? This depends. For one thing, the installation program gives you a choice. But even if you go with the defaults, it still depends, in particular on the user rights you have on the PC. If you have administrator rights, the MetaXL xll add-in will go into Program Files, while the example spreadsheets and such go into a subdirectory of Documents and Setting\All Users\MetaXL.

If you do not have the rights to install software on your PC, you nevertheless can install MetaXL. MetaXL will then install all files under Documents and Setting\Your User Name\MetaXL. In this case you may get the message during installation that "You may have to install the add-in manually". In practice this usually turns out to be not the case: if you open one of the example spreadsheets and do not get #NAME? errors, you're fine. If you do get those errors, you will have to go to Tools\Add-ins (Excel 2003 and earlier) or Office Button\Excel Options\Add-ins\Go (Excel 2007) and browse where the file 'MetaXL.xll' is located (now you know why it is important to know where your files are).

The Excel development team at Microsoft seriously dropped the ball with respect to support for the 2007+ ribbon interface in the C application programming interface (C-API): there is none whatsoever. This forced us to take the messy route of XLM and VBA to support the ribbon interface. Yuck! The Excel 2007 and later installation therefore has a second add-in (called MetaXLRibbon.xlam) installed that takes care of the ribbon issues.

To minimise the pain of finding the MetaXL files, all supplementary files (documentation and examples) are accessible through the Start\Programs\MetaXL menu that is created by the installation program.

An unhelpful Excel offer

Sometimes Excel will offer to move the add-in file 'MetaXL.xll' to its 'AddIns' folder. Never, repeat never, accept this. If you do, it will cause all kinds of mischief, in particular after upgrading to a newer version of MetaXL. It is one of those features of Excel where you really wonder what they were thinking at Microsoft.

If you have already accepted the offer, here are the steps to undo the harm:

1. Go to the AddIns directory (you can find where it is by going through the steps outlined in the previous section to manually add the add-in: when you click the Browse button, Excel will start in the AddIn folder).
2. Delete the MetaXL.xll file. You cannot do so when Excel is running, so you have to quit Excel first.
3. Start Excel. It will complain that it cannot find the MetaXL add-in. When you go to the Excel add-ins list (see above), it will ask you whether the MetaXL add-in should be removed from its list. Click 'yes' to that.
4. Re-install MetaXL.

After having gone through this, you will probably remember to decline Excel's offer in the future ☺.

Statistical and other scientific sources

In the development of MetaXL, we have adopted the Stata™ implementation of meta-analysis as the gold standard. There are two reasons for this. The first is that we have access to Stata (through our jobs at the School of Population Health of the University of Queensland), and therefore it is easy to check the MetaXL results against those of Stata. But the second reason is more important: most of the statistical input for MetaXL is from the excellent book edited by Matthias Egger et al, in particular the chapter by Deeks, Altman, and Bradburn (Deeks, Altman et al. 2001). Various authors of the book's chapters, most of which were first published in the BMJ, were also involved in the implementation of Stata's metan command. If you find any discrepancies between Stata and MetaXL results, we are eager to hear them (info@epigear.com).

So if MetaXL is designed to reproduce the Stata results, then why produce it in the first place? Again for two reasons. The first is that not everybody has access to Stata, and buying a Stata license just for doing the odd meta-analysis is an expensive option. For the many people who have Excel installed, MetaXL offers free access to meta-analysis.

But again the second reason is more important. The overarching issue in meta-analysis is how to deal with heterogeneity between studies. Currently the standard approach is to reach for the random effects model when the heterogeneity statistics (Cochran's Q and I^2) indicate that heterogeneity exceeds a certain threshold, making the various fixed effects models inappropriate.

We think the random effects model is a poor solution to the problem of heterogeneity. The assumption of the random effects model is that the observed heterogeneity is due to real differences in effect sizes: the different effect sizes themselves have a Normal distribution, and thus the meta-analysis aims to estimate the mean underlying true effect.

It is debatable whether the variation in underlying true effect is more common than variations due to systematic errors between studies. The latter is certainly more likely for the bulk of studies in Medicine and thus the random effects model is not an adequate solution, unless you accept its basic premise, which we think results in a random effects summary that lacks interpretation. Indeed, Peto even called use of random-effects meta-analysis "wrong" because it answers a question that is "abstruse and uninteresting" (Peto 1987).

While we are prepared to accept that the real effect size might have a Normal distribution, we are convinced that the studies in a meta-analysis are neither exchangeable nor that their observed heterogeneity is overwhelmingly driven by differences in the underlying effect size. We are also convinced that we are correct in our assumption that sources of bias are the main drivers of heterogeneity, and thus it is much better to deal with heterogeneity by explicitly assessing and weighting for study quality.

So that is where MetaXL comes in: no other meta-analysis software offers the quality effects model that is implemented in MetaXL². By making MetaXL freely available, we hope and expect that explicitly dealing with study quality in meta-analysis will be

² MIX version 1.7 does, but it is no longer supported, and MIX version 2 has dropped the quality effects model.

encouraged, and that the inappropriate use of the random effects model for this purpose will become obsolete.

Software

MetaXL was developed in Object Pascal, using Borland's Delphi 7 programming environment. While many people seem to think C++ is essential for this kind of project, Object Pascal in the Borland implementation is actually just as powerful, easier to use, and beats all C++ compilers by its blazing speed.

Writing xll add-ins for Excel has been described as a black art. While we would hesitate to call it that, it certainly requires lots of stamina. The main reasons are that the Microsoft Excel xll software development kits are rather short on detail, and that Excel proves to be a wilful environment to program for, with a tendency to either crash or sulk if anything is wrong rather than be explicit about it.

Fortunately, where Microsoft left a lot to be desired, some people have stepped in. We are indebted to David Bolton whose article "Writing (Non Com) Excel Add-ins in Delphi" (which originally appeared in Delphi Magazine but is now floating around on the Web), despite some strange errors, gave the necessary heads up in our first steps on this road. An invaluable resource is also Steve Dalton's "Financial applications using Excel add-in development in C/C++" (Dalton 2007). But as the title suggests, you need to be able to understand C(++) code to get the most out of this book.

Excel 2007 is a major upgrade from previous versions. Among other things, it sports a 'big grid', i.e. many more rows and columns than previous versions. This required a change in its basic data type, the XLOPER, to the XLOPER12. These special data types have the flexibility to describe the different things Excel works with: text, numbers, ranges, etc. Excel 2007 is largely backward compatible with add-ins written for earlier versions, but MetaXL actively supports the new features of Excel 2007 by implementing a dual interface, and presenting the applicable one depending on which version of Excel is running. A not yet supported feature of Excel 2007 is multi-threaded recalculation, which is planned for MetaXL version 2.0.

Tip: when you use MetaXL with Excel 2007, make sure you have the latest Office service pack installed. The first release of Excel 2007 was very slow, and while with SP 2 Excel 2007 is still slower than earlier versions, recalculation is about five times as fast as the first release.

The MetaXL installation program was written using Inno Setup 5 (www.jrsoftware.org). This is one of those amazing things: it is freeware, but beats commercial installation software such as InstallShield hands down on features as ease of use, power, and flexibility. Recommended.

About the authors

Jan Barendregt is Associate Professor of Epidemiological Modelling at the Centre for Burden of Disease and Cost-Effectiveness of the School of Population Health, University of Queensland, Brisbane, Australia. He is also the founder and owner of EpiGear International (www.epigear.com), which makes free and low-cost software. Email: j.barendregt@sph.uq.edu.au or jan@epigear.com.

Suhail Doi is Associate Professor of Clinical Epidemiology at the School of Population Health, University of Queensland, Brisbane, Australia. Email: s.doi@sph.uq.edu.au or sardo@gmx.net.

References

- Al Khalaf, M. M., L. Thalib, et al. (2011). "Combining heterogeneous studies using the random-effects model is a mistake and leads to inconclusive meta-analyses." *J Clin Epidemiol* **64**: 119-23.
- Bailey, K. R. (1987). "Inter-study differences: how should they influence the interpretation and analysis of results?" *Stat Med* **6**(3): 351-60.
- Batham, A., M. A. Gupta, et al. (2009). "Calculating prevalence of hepatitis B in India: using population weights to look for publication bias in conventional meta-analysis." *Indian J Pediatr* **76**(12): 1247-57.
- Berard, A. and G. Bravo (1998). "Combining studies using effect sizes and quality scores: application to bone loss in postmenopausal women." *J Clin Epidemiol* **51**(10): 801-7.
- Conn, V. S. and M. J. Rantz (2003). "Research methods: managing primary study quality in meta-analyses." *Res Nurs Health* **26**(4): 322-33.
- Cox, D. R. (1970). "The continuity correction." *Biometrika* **57**: 217-219.
- Dalton, S. (2007). *Financial Applications Using Excel Add-in Development in C/C++*. Chichester, Wiley.
- Deeks, J. J., D. G. Altman, et al. (2001). Statistical methods for examining heterogeneity and combining results from several studies in meta-analysis. *Systematic reviews in health care: Meta-analysis in context*. M. Egger, G. Davey Smith and D. G. Altman. London, BMJ Publishing Group: 285-312.
- DerSimonian, R. and N. Laird (1986). "Meta-analysis in clinical trials." *Control Clin Trials* **7**(3): 177-88.
- Doi, S. A., J. J. Barendregt, et al. (2011). "Meta-analysis Of Heterogeneous Clinical Trials: An Empirical Example." *Contemporary Clinical Trials* **32**(2): 288-98.
- Doi, S. A. and L. Thalib (2008). "A quality-effects model for meta-analysis." *Epidemiology* **19**(1): 94-100.
- Doi, S. A. and L. Thalib (2009). "An alternative quality adjustor for the quality effects model for meta-analysis." *Epidemiology* **20**(2): 314.
- Freeman, M. F. and J. W. Tukey (1950). "Transformations Related to the Angular and the Square Root." *The Annals of Mathematical Statistics* **21**(4): 607-611.
- Helfenstein, U. (2002). "Data and models determine treatment proposals--an illustration from meta-analysis." *Postgrad Med J* **78**(917): 131-4.
- Leeflang, M., J. Reitsma, et al. (2007). "Impact of adjustment for quality on results of meta-analyses of diagnostic accuracy." *Clin Chem* **53**(2): 164-72.
- Lipsey, M. and D. B. Wilson (2001). *Practical meta-analysis*. Thousand Oaks, Sage Publications.
- Lipsey, M. W. and D. B. Wilson (2001). "The way in which intervention studies have "personality" and why it is important to meta-analysis." *Eval Health Prof* **24**(3): 236-54.
- Miller, J. J. (1978). "The Inverse of the Freeman-Tukey Double Arcsine Transformation." *The American Statistician* **32**(4): 138.
- Nuesch, E. and P. Juni (2009). "Commentary: Which meta-analyses are conclusive?" *Int J Epidemiol* **38**(1): 298-303.
- Peto, R. (1987). "Why do we need systematic overviews of randomized trials?" *Stat Med* **6**: 233-44.
- Senn, S. (2007). "Trying to be precise about vagueness." *Stat Med* **26**(7): 1417-30.

- Sterne, J. A. C. and M. Egger (2001). "Funnel plots for detecting bias in meta-analysis: Guidelines on choice of axis." Journal of Clinical Epidemiology **54**: 1046–1055.
- Sweeting, M. J., A. J. Sutton, et al. (2004). "What to add to nothing? Use and avoidance of continuity corrections in meta-analysis of sparse data." Stat Med **23**(9): 1351-75.
- Titchler, D. (1999). "Modelling study quality in meta-analysis." Stat Med **18**(16): 2135-45.
- Verhagen, A. P., H. C. de Vet, et al. (2001). "The art of quality assessment of RCTs included in systematic reviews." J Clin Epidemiol **54**(7): 651-4.
- Woolf, B. (1955). "On estimating the relation between blood group and disease." Ann Hum Genet **19**(4): 251-3.
- Yates, F. (1934). "Contingency tables involving small numbers and the chi-squared test." Journal of the Royal Statistical Society **1**: 217-235.