

Example name	Guidelines for cardiac arrest
Effect size	Odds ratio
Analysis type	Basic analysis with between-studies variance
Version	Long

Synopsis

This analysis includes nine studies where patients were randomized to be treated for cardiac arrest based on either old guidelines or new guidelines. Outcome was survival (being alive), and we focused on the odds ratio as the effect size.

We use this example to show

- How to enter data from 2x2 tables
- How to get a sense of the weight assigned to each study
- How the study weights are affected by the model
- How to perform a sensitivity analysis
- How to perform a cumulative analysis
- How to interpret statistics for heterogeneity
- How to create a high-resolution plot

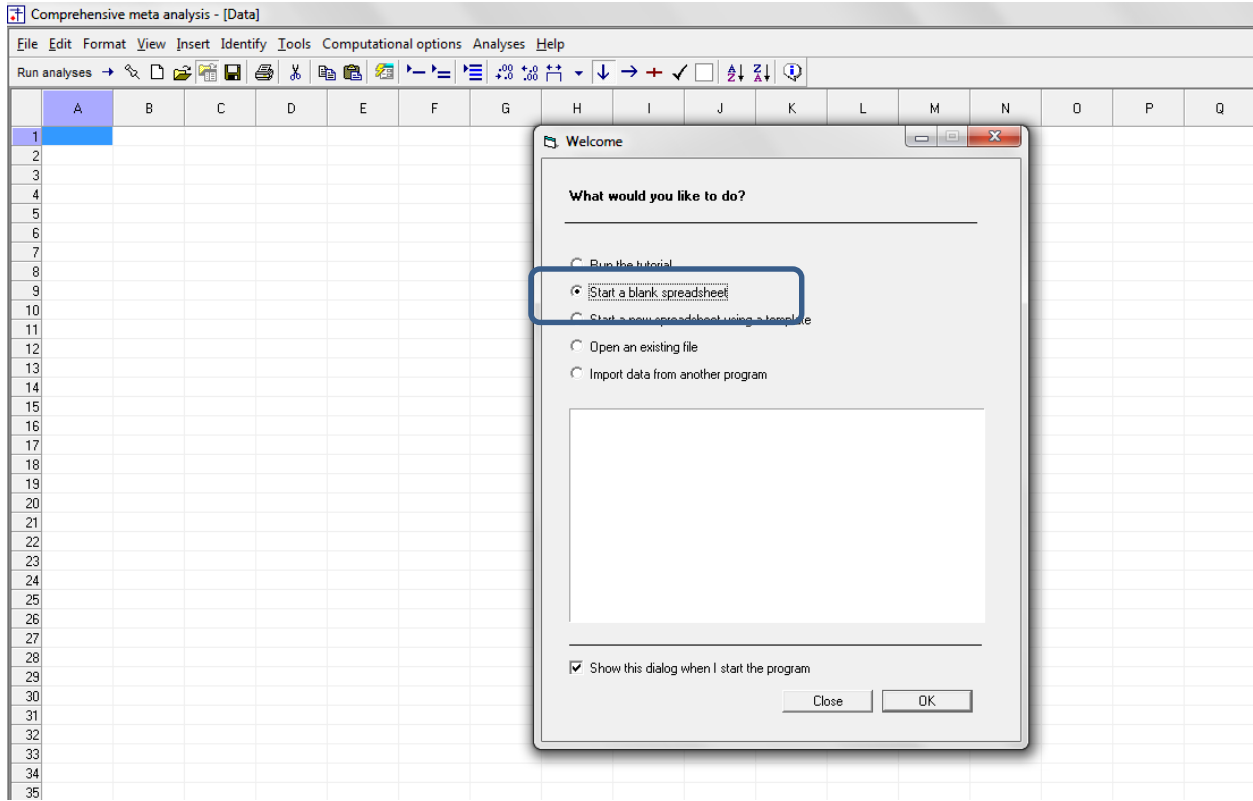
To open a CMA file > [Download and Save file](#) | [Start CMA](#) | [Open file from within CMA](#)

[Download CMA file for computers that use a period to indicate decimals](#)
[Download CMA file for computers that use a comma to indicate decimals](#)

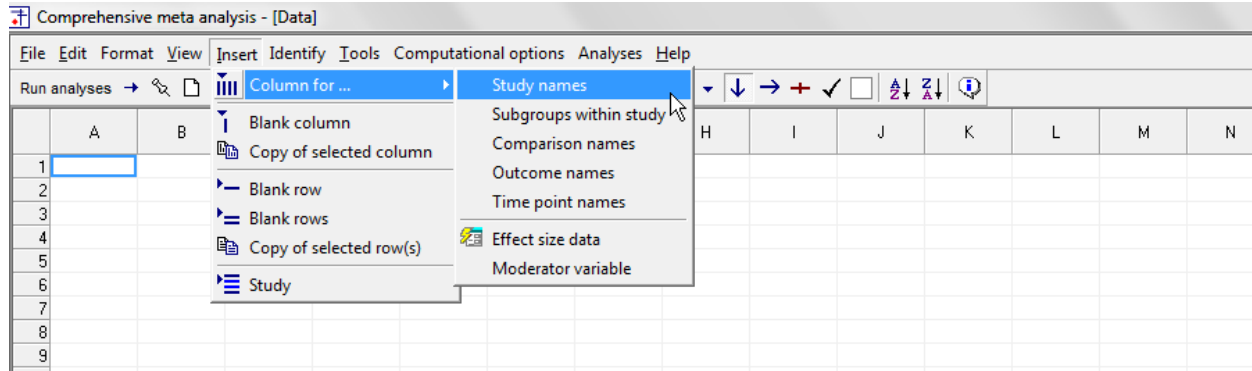
[Download this PDF](#)
[Download data in Excel](#)
[Download trial of CMA](#)

Start the program

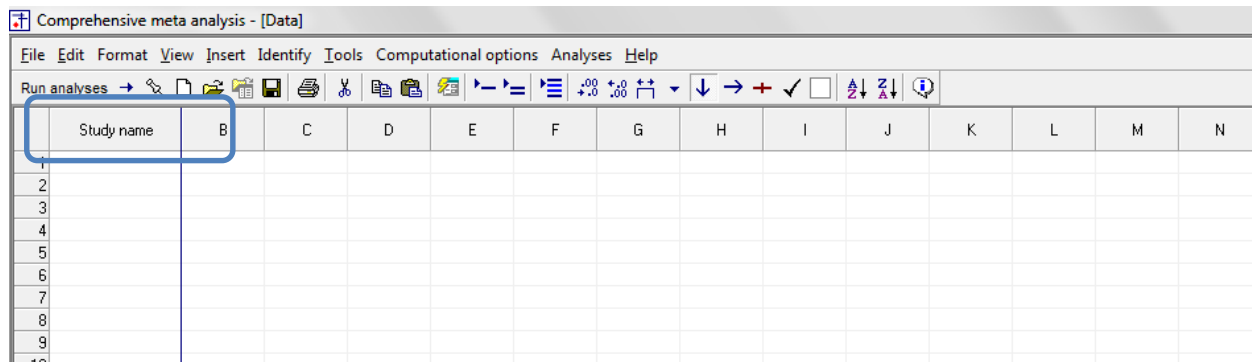
- Select the option [Start a blank spreadsheet]
- Click [OK]



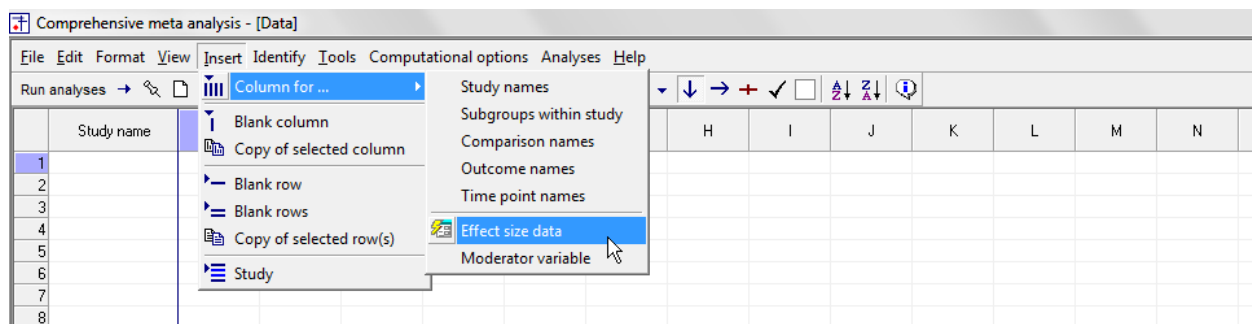
Click Insert > Column for > Study names



The screen should look like this

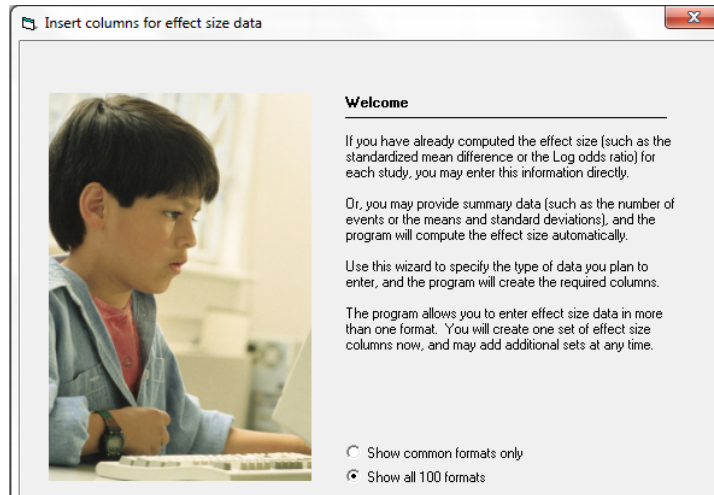


Click Insert > Column for > Effect size data




The program displays this wizard

Select [Show all 100 formats]
Click [Next]




Select [Comparison of two groups...]
Click [Next]



Drill down to

Dichotomous (number of events)
Unmatched groups, prospective ...
Events and sample size in each group



The program displays this wizard

Enter the following labels into the wizard

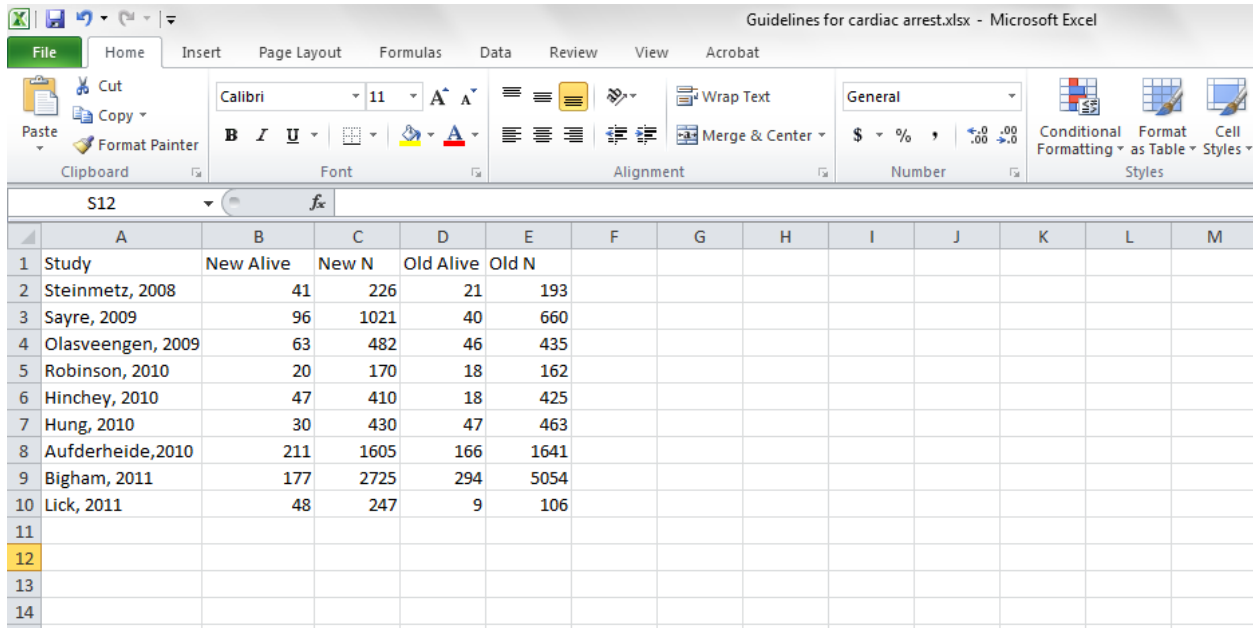
- First group > New
- Second group > Old
- Name for events > Alive
- Name for non-events > Dead

Click [Ok] and the program will copy the names into the grid

The screenshot shows a software window titled "Comprehensive meta analysis - [Data]". The window contains a data grid with columns: Study name, New Alive, New Total N, Old Alive, Old Total N, Odds ratio, Log odds ratio, Std Err, Variance, J, K, L, M, N, O, P. The "New Alive" and "Old Total N" cells are highlighted with a blue box. A "Group names" dialog box is open over the grid. The dialog box has two sections: "Group names for cohort or prospective studies" and "Binary outcome in cohort or prospective studies". The first section has two input fields: "Name for first group (e.g., Treated)" with "New" and "Name for second group (e.g., Control)" with "Old". The second section has two input fields: "Name for events (e.g., Dead)" with "Alive" and "Name for non-events (e.g., Alive)" with "Dead". The dialog box also has "Cancel", "Apply", and "Ok" buttons.

Rather than enter the data directly into CMA we will copy the data from Excel

- Switch to Excel and open the file “Guidelines for cardiac arrest”
- Highlight the rows and columns as shown, and press CTRL-C to copy to clipboard



Guidelines for cardiac arrest.xlsx - Microsoft Excel

	A	B	C	D	E	F	G	H	I	J	K	L	M
1	Study	New Alive	New N	Old Alive	Old N								
2	Steinmetz, 2008	41	226	21	193								
3	Sayre, 2009	96	1021	40	660								
4	Olasveengen, 2009	63	482	46	435								
5	Robinson, 2010	20	170	18	162								
6	Hinchey, 2010	47	410	18	425								
7	Hung, 2010	30	430	47	463								
8	Aufderheide,2010	211	1605	166	1641								
9	Bigham, 2011	177	2725	294	5054								
10	Lick, 2011	48	247	9	106								
11													
12													
13													
14													

Switch back to CMA

- Click in Cell Study name – 1
- Press [CTRL-V] to paste the data into CMA
- Stretch the columns as needed for the text to be fully visible

Click here

Comprehensive meta analysis - [Data]

File Edit Format View Insert Identify Tools Computational options Analyses Help

Run analyses →

	Study name	New Alive	New Total N	Old Alive	Old Total N	Odds ratio	Log odds ratio	Std Err	Variance	J	K	L	M	N
1	Study	New Alive	New N	Old Alive	Old N									
2	Steinmetz, 2008	41	226	21	193	1.815	0.596	0.288	0.083					
3	Sayre, 2009	96	1021	40	660	1.609	0.475	0.195	0.038					
4	Olasveengen, 2009	63	482	46	435	1.272	0.240	0.206	0.043					
5	Robinson, 2010	20	170	18	162	1.067	0.065	0.345	0.119					
6	Hinchey, 2010	47	410	18	425	2.928	1.074	0.286	0.082					
7	Hung, 2010	30	430	47	463	0.664	-0.410	0.244	0.060					
8	Aufderheide, 2010	211	1605	166	1641	1.345	0.296	0.110	0.012					
9	Bigham, 2011	177	2725	294	5054	1.125	0.118	0.098	0.010					
10	Lick, 2011	48	247	9	106	2.600	0.955	0.384	0.147					
11														
12														
13														
14														

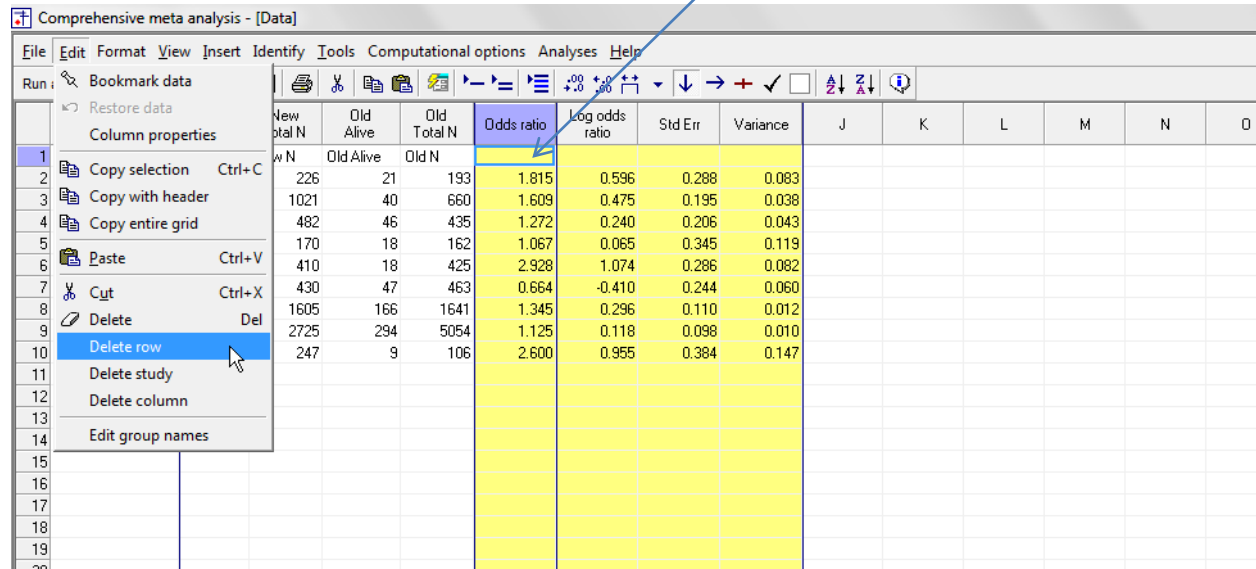
Since we've copied the row with the titles, we can check to ensure that all columns are being used as intended. The column that had been "New Alive" is now called "New Alive". Similarly, all other columns are labeled correctly.

Note that the program has computed the odds ratio for all studies and displayed this (as well as the log odds ratio, its standard error and variance) in the yellow columns.

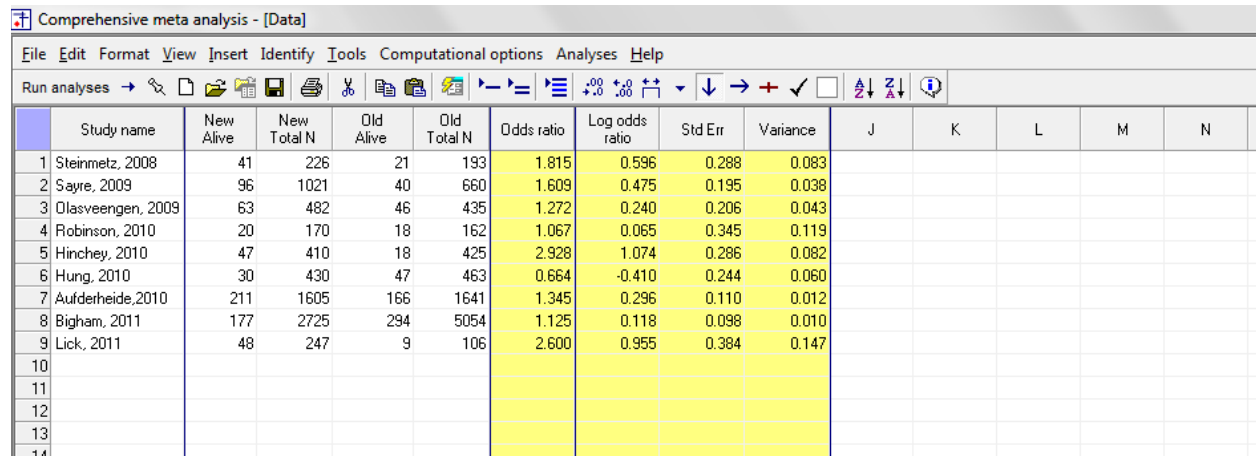
Now, we can remove the first row

Click here

- Click in the first row to select it
- Click Edit > Delete row and confirm



The screen should look like this

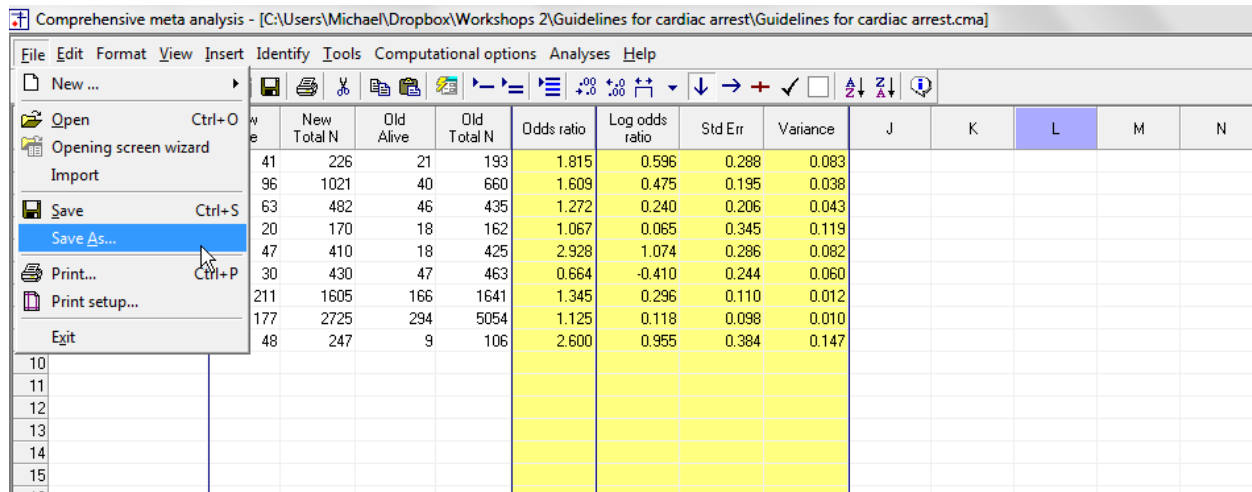


We've followed the convention of putting the treated (new guidelines) group before the control (old guidelines). When we do this, if (a) the treated group does better and (b) the outcome is something good (being alive) the odds ratio will be greater than 1.0.

To check that things are working as planned let's use the last study. The proportion surviving is roughly 20% (48/247) in the New group and around 10% (9/106) in the Old group. The odds ratio (2.600) is indeed greater than 1. In the analysis, odds ratio greater than 1 should be labeled "Favors New" while odds ratios less than 1 should be labeled "Favors Old". We need to apply these labels manually.

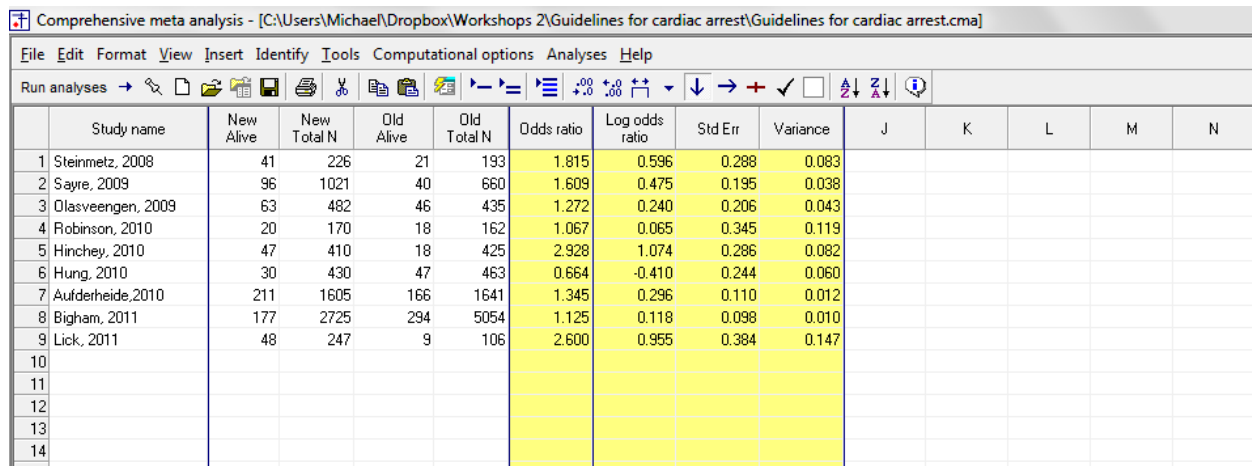
At this point we should save the file

- Click File > Save As ...



Note that the file name is now in the header.

- [Save] will over-write the prior version of this file without warning
- [Save As...] will allow you to save the file with a new name



By default the program displays the odds ratio.

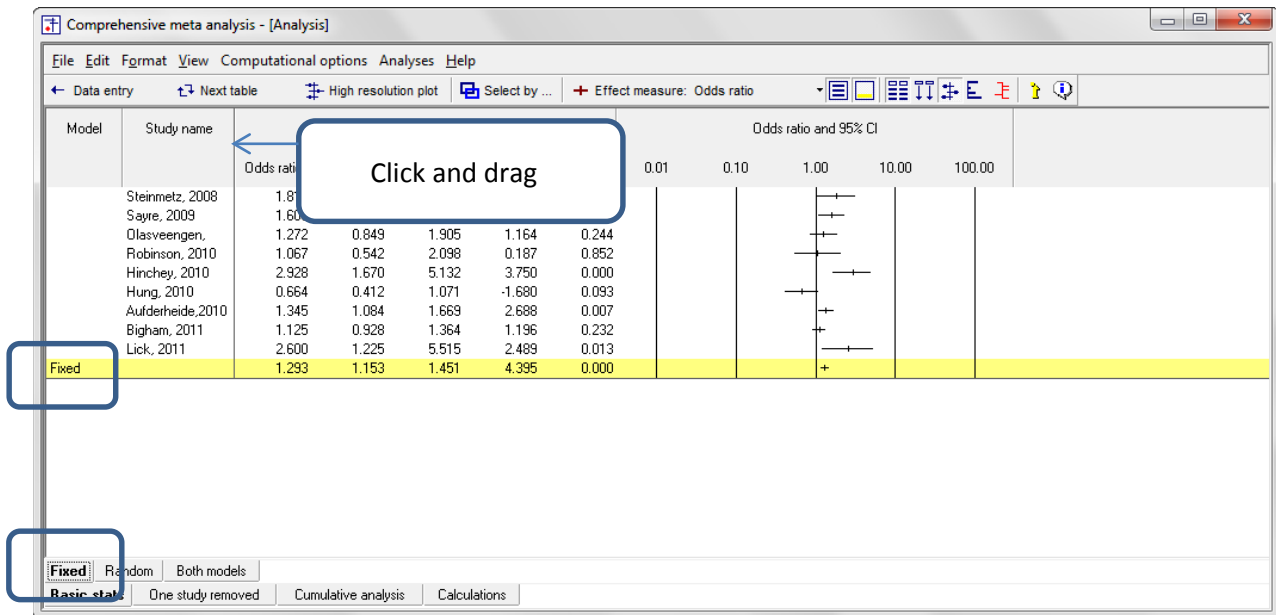
This is what we want to use in the analysis, so no modification is needed.

- To run the analysis, click [Run analysis]

This is the basic analysis screen

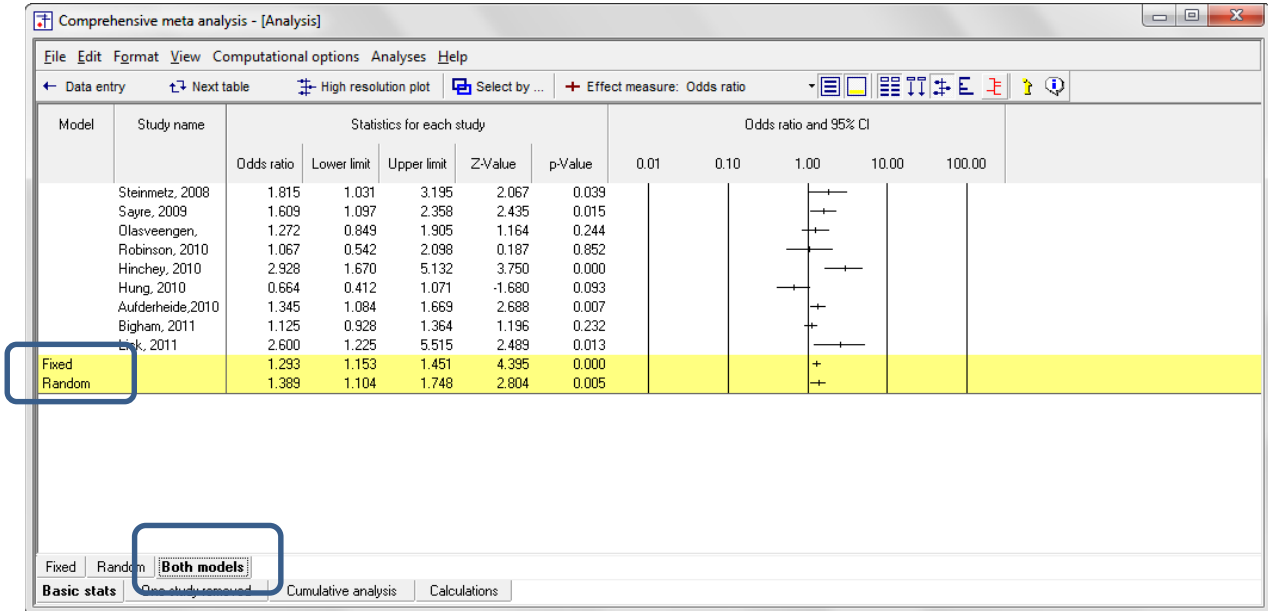
Stretch the Study name column so the full name displays

Initially, the program displays the fixed-effect analysis. This is indicated by the tab at the bottom and the label in the plot.



Click [Both models]

The program displays results for both the fixed-effect and the random-effects analysis.



The fact that the two results are not identical tells us that the weights are different, which means that the effect size varies from study to study. (This means that T^2 , the estimate of between-study variance in true effects is non-zero. It is *not* a test of statistical significance).

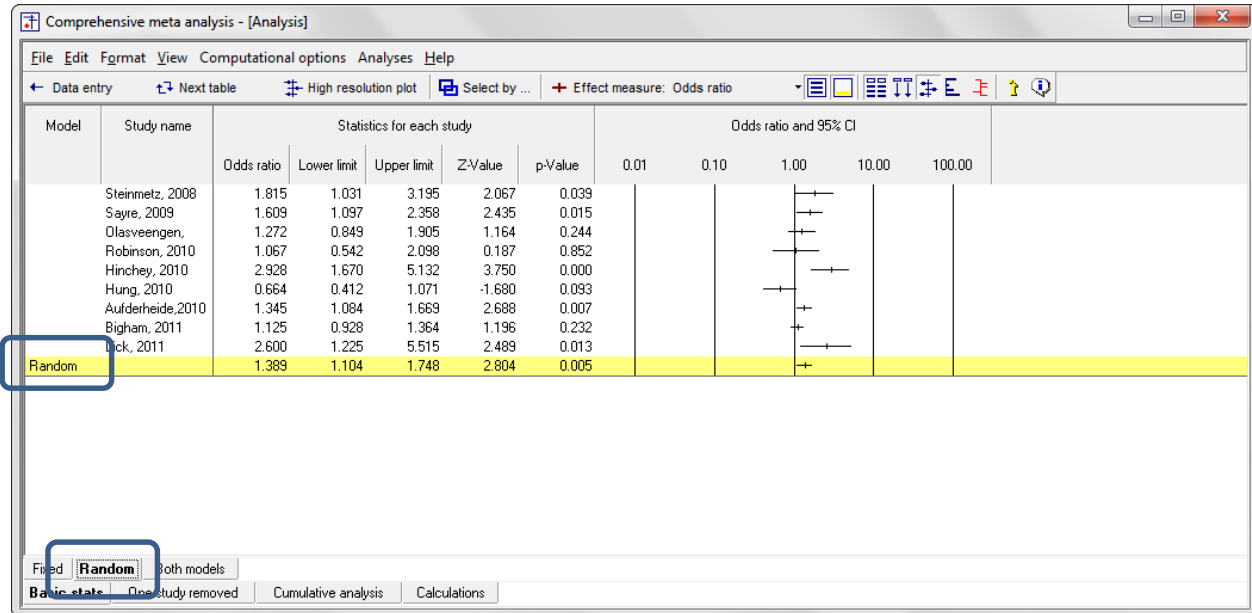
The confidence interval is wider for random-effects than for fixed-effects. This will always be the case when T^2 is non-zero.

The random-effects estimate of the effect size is slightly stronger (further from 1.0) than the fixed-effect estimate. The FE and RE estimates will almost always differ from each other when T^2 is non-zero, but the difference can be in either direction.

In any event, the random-effects model is a better fit for the way the studies were sampled, and therefore that is the model we will use in the analysis.

- Click Random on the tab at the bottom

The plot now displays the random-effects analysis alone.



A quick view of the plot suggests the following

- Almost all studies suggest an advantage for the new guidelines over the old
- In most studies this effect falls in the range of 1.067 to 2.928, though the CI for single studies extend well beyond that range. One study suggests that the old guidelines did better, but the effect is not statistically significant.
- The summary effect is 1.389 with a CI of 1.104 to 1.748. Thus, the mean effect is likely in the clinically important range.
- The summary effect has a Z-value 2.804 and a p-value of 0.005. Thus we can reject the null hypotheses that the true odds ratio is 1.0.
- It's clear from the plot that there is variation in the *observed* effects. It also looks like the observed variation exceeds the amount that we would expect to see based solely on within-study sampling error. For example, the CI for the Hinchey study has little (if any) overlap with the CI for the summary effect.
- We also know that there must be some true between-study variance in the sample, by virtue of the fact that the RE estimate is different than the FE estimate.
- To have a closer look at this variance we turn to the next table.

Click [Next table]

Click here

The screenshot shows the 'Comprehensive meta analysis - [Analysis]' window. The 'Next table' button is highlighted with a blue box and a callout. The main table displays results for two models: Fixed and Random. The table has columns for Model, Number Studies, Point estimate, Lower limit, Upper limit, Z-value, P-value, Q-value, df (Q), P-value, I-squared, Tau Squared, Standard Error, Variance, and Tau.

Model	Effect size and 95% interval				Test of null (2-Tail)		Heterogeneity			Tau-squared				
	Number Studies	Point estimate	Lower limit	Upper limit	Z-value	P-value	Q-value	df (Q)	P-value	I-squared	Tau Squared	Standard Error	Variance	Tau
Fixed	9	1.293	1.153	1.451	4.395	0.000	24.015	8	0.002	66.687	0.071	0.063	0.004	0.267
Random	9	1.389	1.104	1.748	2.804	0.005								

The statistics at the left duplicate those we saw on the prior screen.

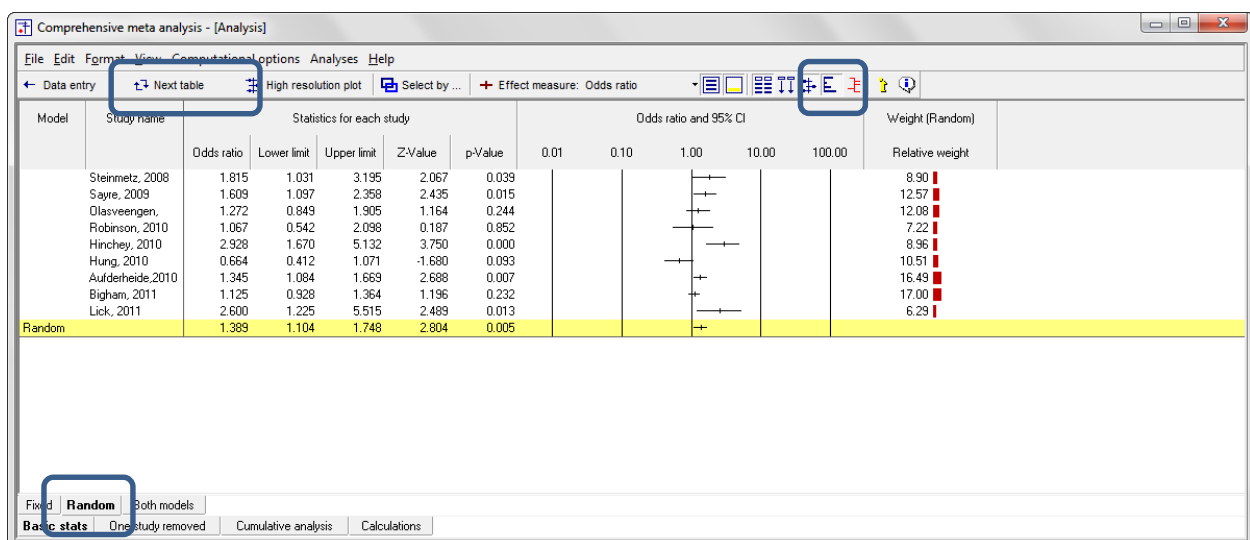
- Under the random-effects model the odds ratio is 1.389 with a 95% confidence interval of 1.104 to 1.748. The test of the null (that the true risk ratio is 1.0) yields a Z-value of 2.804 and a corresponding p-value of 0.005.
- The statistics at the upper right relate to the dispersion of effect sizes across studies.
- The Q-value is 24.015 with $df=8$ and $p=0.002$. Q reflects the distance of each study from the mean effect (weighted, squared, and summed over all studies). Q is always computed using FE weights (which is the reason it is displayed on the “Fixed” row, but applies to both FE and RE analyses).
- If all studies actually shared the same true effect size, the expected value of Q would be equal to df (which is 8). Here, Q exceeds that value, and exceeds it by enough that the difference is probably not due to chance. The p-value of 0.002 allows us to reject the null hypothesis that all studies share the same true effect size, and conclude that the true effect size actually varies from study to study.
- T^2 is the estimate of the between-study variance in true effects. This estimate is 0.071. T is the estimate of the between-study standard deviation in true effects. This estimate is 0.267. Note that these values are in log units. Therefore, to use these estimates to compute confidence intervals or prediction intervals we would need to convert all values into log units, perform the computations, and convert the values back into odds ratios. (This is handled automatically by the program.)
- The variance in effect sizes includes both sampling error and variance in the true effect size from study to study. The I^2 value is 66.687, which tells us that about two-thirds of the *observed* variance in effect sizes reflects differences in *true* effect sizes. This means that if each of the

studies had a huge sample size (so that the observed effect closely mirrored the true effect size for that study's population) the observed effects would fall closer to each other than they do now, but would not align exactly. The variance of the observed effects would drop by about a third.

Click [Next table] to return to this screen

It's always important to see how much each study contributed to the summary effect. For example, if the summary effect was dominated by one or two studies, we would want to be aware of this and consider if it affects the generalizability of the results.

- Click the [Weights] button on the toolbar to display the relative weight assigned to each study.

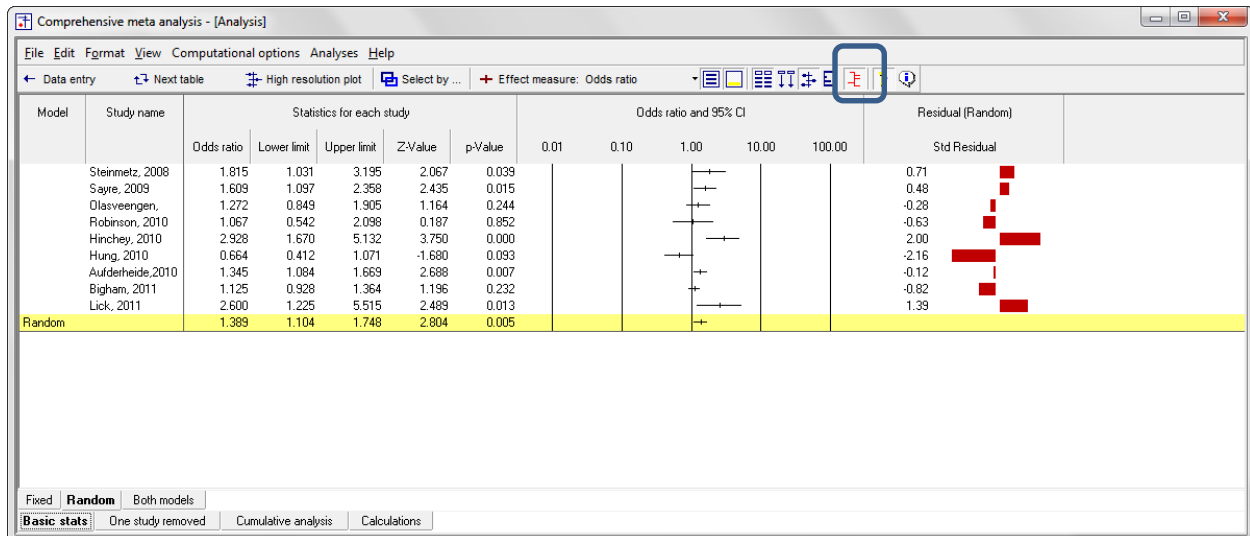


We can see that the relative weights ranged from a low of 3.29% to a high of 17.00%. So, no single study dominated the analysis. Five studies have relative weights in the range of 10% to 17% while the remaining four have weights in the range of 6% to 9%.

Do any studies appear to be outliers?

- Click the residual button on the toolbar to display the Residual column.

One study has a standardized residual of +2.00 while another has a standardized residual of -2.16. We may want to have a closer look at these and their impact on the analysis. It's likely that they offset each other's impact on the mean effect, while increasing the estimate of between-study variance.



We might wonder how the summary effect would have shifted if any one study had been excluded from the analysis. To address this question we can run the analysis eight times, each time with a single study excluded.

The program allows us to run these eight analyses in one step.

- Click [One study removed] at the bottom.

Model	Study name	Statistics with study removed					Odds ratio (95% CI) with study removed				
		Point	Lower limit	Upper limit	Z-Value	p-Value	0.01	0.10	1.00	10.00	100.00
	Steinmetz, 2008	1.354	1.061	1.728	2.440	0.015			++		
	Sayre, 2009	1.365	1.057	1.762	2.387	0.017			++		
	Olasveengen,	1.414	1.089	1.834	2.604	0.009			++		
	Robinson, 2010	1.422	1.112	1.817	2.811	0.005			++		
	Hinchey, 2010	1.281	1.046	1.569	2.393	0.017			++		
	Hung, 2010	1.487	1.205	1.835	3.702	0.000			++		
	Aulderheide,2010	1.416	1.058	1.894	2.341	0.019			++		
	Bigham, 2011	1.459	1.105	1.926	2.661	0.008			++		
	Lick, 2011	1.330	1.059	1.670	2.450	0.014			++		
Random		1.389	1.104	1.748	2.804	0.005			++		

The program displays this screen. The yellow line at the bottom is the analysis with all eight studies. The first row is what the yellow line would look like if the first study was excluded from the analysis. The second row is what the yellow line would look like if the second study was excluded from the analysis. And so on.

It seems that the effect size would remain pretty-much unchanged if any one study was removed – in other words, the conclusions are not dependent on any single study. The summary effect is an odds ratio of 1.389. If we remove any one study the effect size stays in the range of 1.281 to 1.489, which is probably not a clinically important difference.

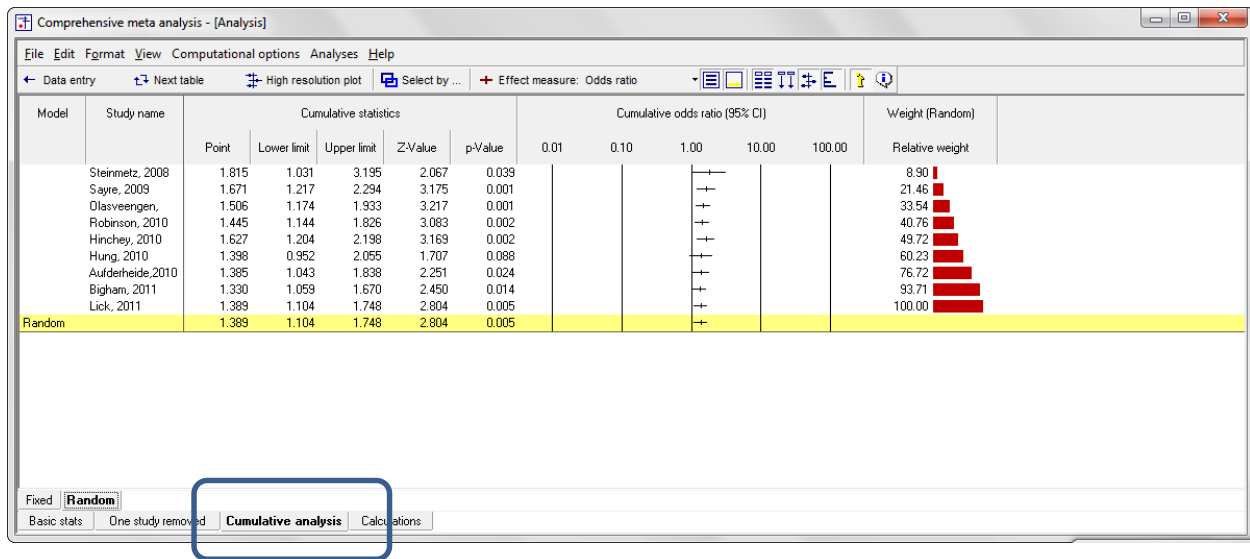
For those concerned about significance tests, the p-value with all studies is 0.005, and it stays at < 0.02 with any single study removed.

We might wonder how the weight of the evidence has shifted over time. In other words, what would a meta-analysis have shown if we had performed it after the first study, after the first two studies, and so on.

To run this analysis we need to ensure that the studies are sorted by year on the data-entry screen. In this case, they are, and so we can proceed.

- Click [Cumulative analysis] on the bottom
- Click the tool for relative weights on the menu

The program displays this screen



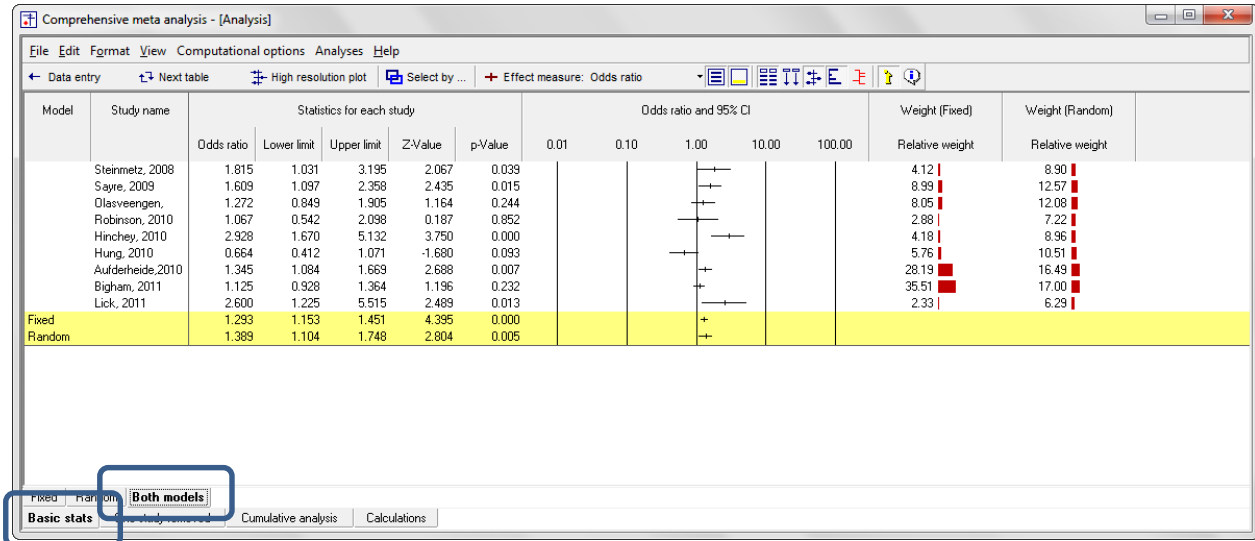
The risk ratio would have been estimated at 1.815 after the one study, at 1.671 after two studies, and at 1.506 after three studies. From that point on, the estimate is pretty stable to the end.

While the estimate of the effect size changes little beyond the first three studies, the confidence interval narrows with each additional study, which means that we have a more precise estimate of the effect size. It also means that the p-value moves closer to zero.

This does not mean that the additional studies were unimportant. They may have provided important information about the robustness of the effect. They did provide information about other outcomes (not reported here).

Please note that the cumulative analysis shown here is intended only as a look-back. It would be a very bad idea to repeat a meta-analysis every time a new study was added to the literature, with the goal of stopping when the p-value hits 0.05. If the goal is to repeat the analysis every time a study is added, then adjustments must be made to the p-value and confidence interval.

- Click [Basic stats] to return to this screen
- Click [Both models] to display both Fixed-effect and Random-effects estimates
- Click [Weights] to display the relative weights assigned under both models



As noted earlier, we use the random-effects model because it matches the way the studies were sampled. (As it happens, the test of homogeneity also yielded a p-value of 0.002, which tells us that the fixed-effect model is empirically false. However, we would have selected random-effects even if this p -value was not statistically significant).

In that context, this screen is presented purely for instructional purposes, to show how the weights affect the model.

The plot shows that the odds ratio moves slightly to the right (a larger effect, since we are working with values above 1.0) when we move from fixed-effect weights to random-effects weights. The columns with the weights show why this happens. The Bigham study happens to have an unusually small effect size relative to the others. This study also happens to have a large sample size. Under the FE model this study gets 35% of the weight, and pulls the summary effect to the left. Under the RE model this study gets only 17% of the weight. It still pulls the summary effect to the left, but not as much. The weight for other studies shifts as well, with the overall impact of the shifts being for the RE estimate to be slightly higher.

Next, we want to create a high-resolution plot.

When we click [Hi-Res], the screen will be recreated in a high-resolution plot. Therefore, *before* moving to that screen we need to determine which columns to include (or exclude). The basic idea is to exclude as many columns as possible, to ensure maximum clarity on the plot.

- Click [Weights] to hide the weights. In the high-res version these will be reflected in the size of each study.
- Click [Random] at the bottom.

Right-click anywhere in the section called [Statistics for each study]

The program opens a wizard as shown here

- Check the boxes for odds-ratio and p-value
- Uncheck all other boxes
- Click Ok

Right-click here

The screenshot shows the 'Comprehensive meta analysis - [Analysis]' window. The main data table is as follows:

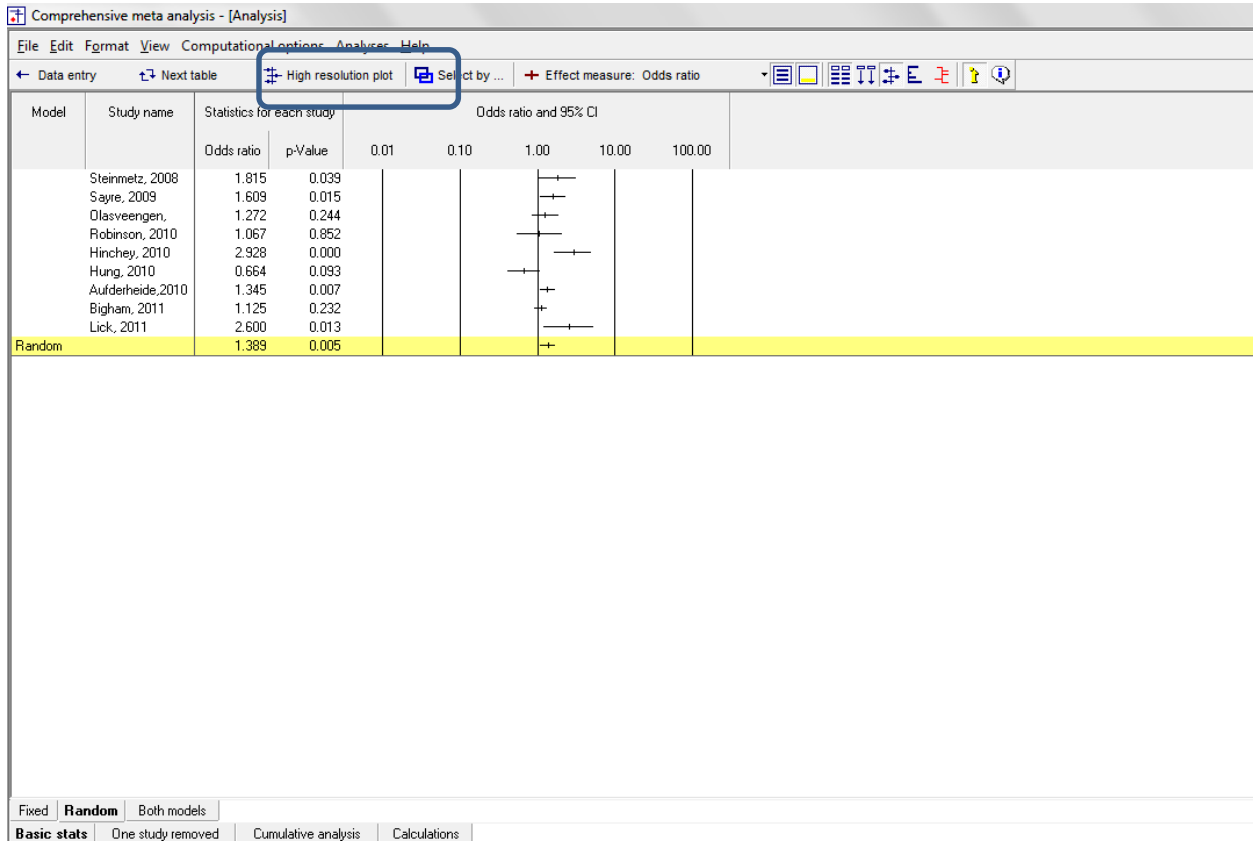
Model	Study name	Statistics for each study					Odds ratio and 95% CI				
		Odds ratio	Lower limit	Upper limit	Z-Value	p-Value	0.01	0.10	1.00	10.00	100.00
	Steinmetz, 2008	1.815	1.031	3.195	2.067	0.039					
	Sayre, 2009	1.609	1.097	2.358	2.435	0.015					
	Olasveengen,	1.272	0.849	1.905	1.164	0.244					
	Robinson, 2010	1.067	0.542	2.098	0.187	0.852					
	Hinchey, 2010	2.928	1.670	5.132	3.750	0.000					
	Hung, 2010	0.664	0.412	1.071	-1.680	0.093					
	Aufderheide,2010	1.345	1.084	1.669	2.688	0.007					
	Bigham, 2011	1.125	0.928	1.364	1.196	0.232					
	Lick, 2011	2.600	1.225	5.515	2.489	0.013					
Random		1.389	1.104	1.748	2.804	0.005					

The 'Customize display' dialog box is open, showing the following settings:


Show	Decimals	Alignment
<input checked="" type="checkbox"/> All columns in this block	[Auto]	[Auto]
<input checked="" type="checkbox"/> Odds ratio	[Auto]	[Auto]
<input type="checkbox"/> Standard error	[Auto]	[Auto]
<input type="checkbox"/> Variance	[Auto]	[Auto]
<input type="checkbox"/> Lower limit	[Auto]	[Auto]
<input type="checkbox"/> Upper limit	[Auto]	[Auto]
<input type="checkbox"/> Z-Value	[Auto]	[Auto]
<input checked="" type="checkbox"/> p-Value	[Auto]	[Auto]

At the bottom of the window, the 'Random' model is selected, and the 'Basic stats' tab is active.

- Click the tool for Hi-res plot

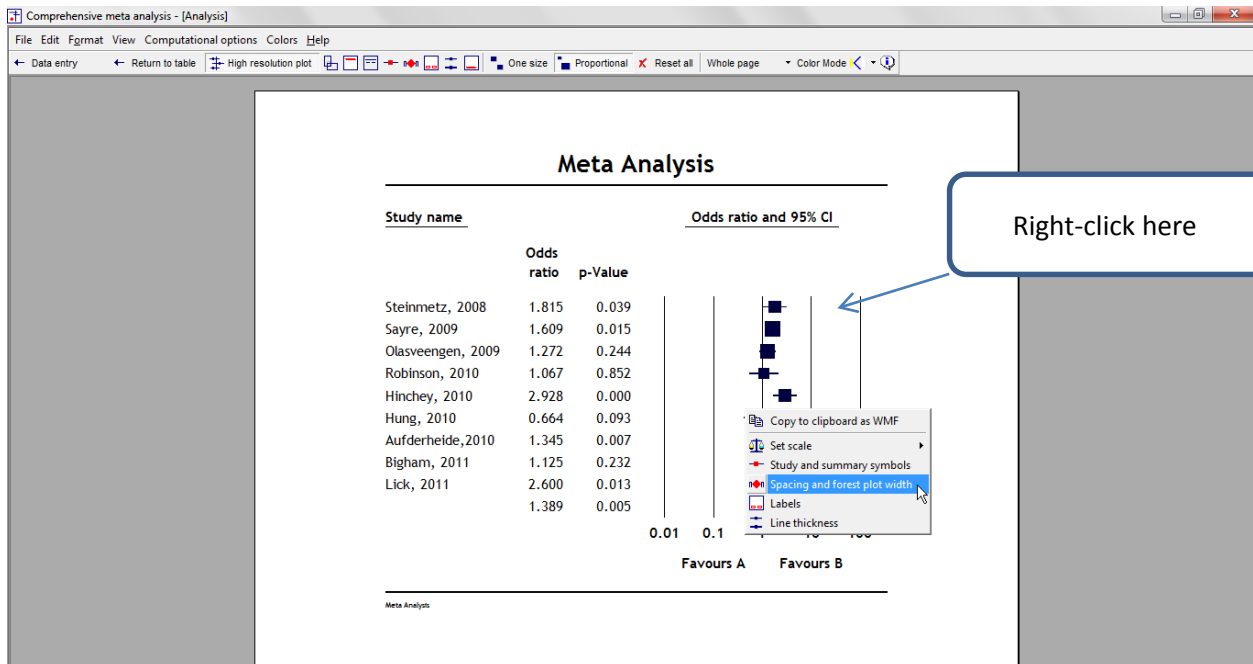


- On this screen click [Reset all]
- The screen should look like this

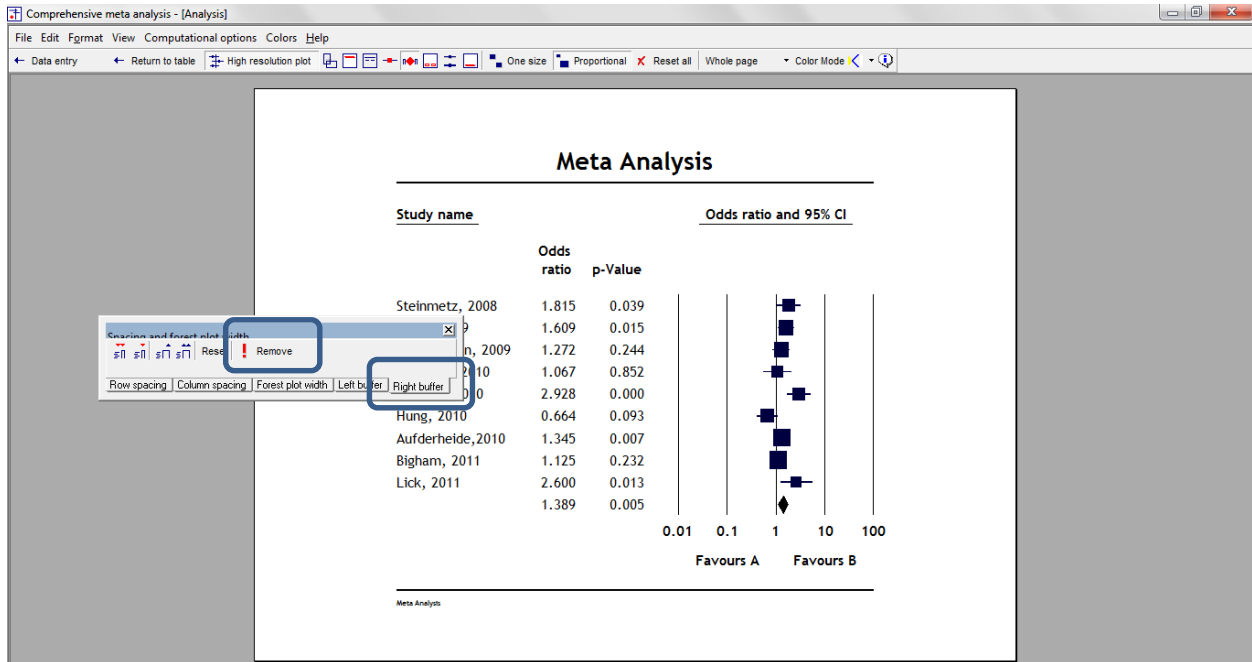


At this point we can allocate more or less space to each element of the plot

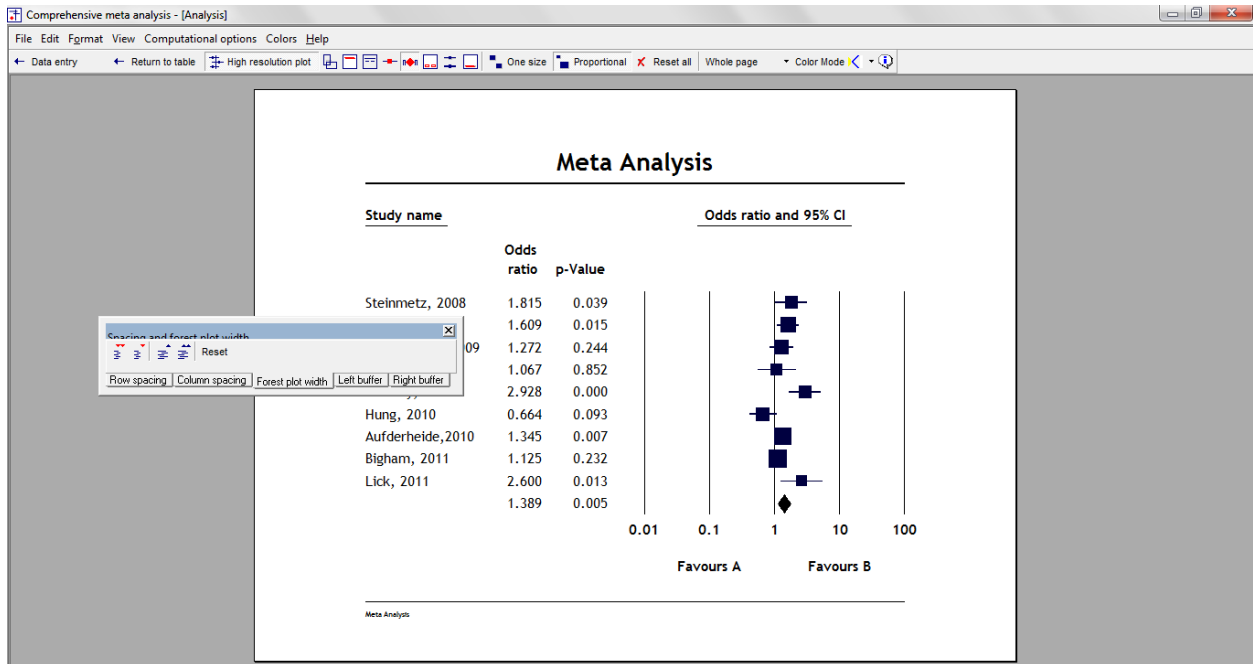
- Right click on the forest plot
- Select [Spacing and forest plot width]



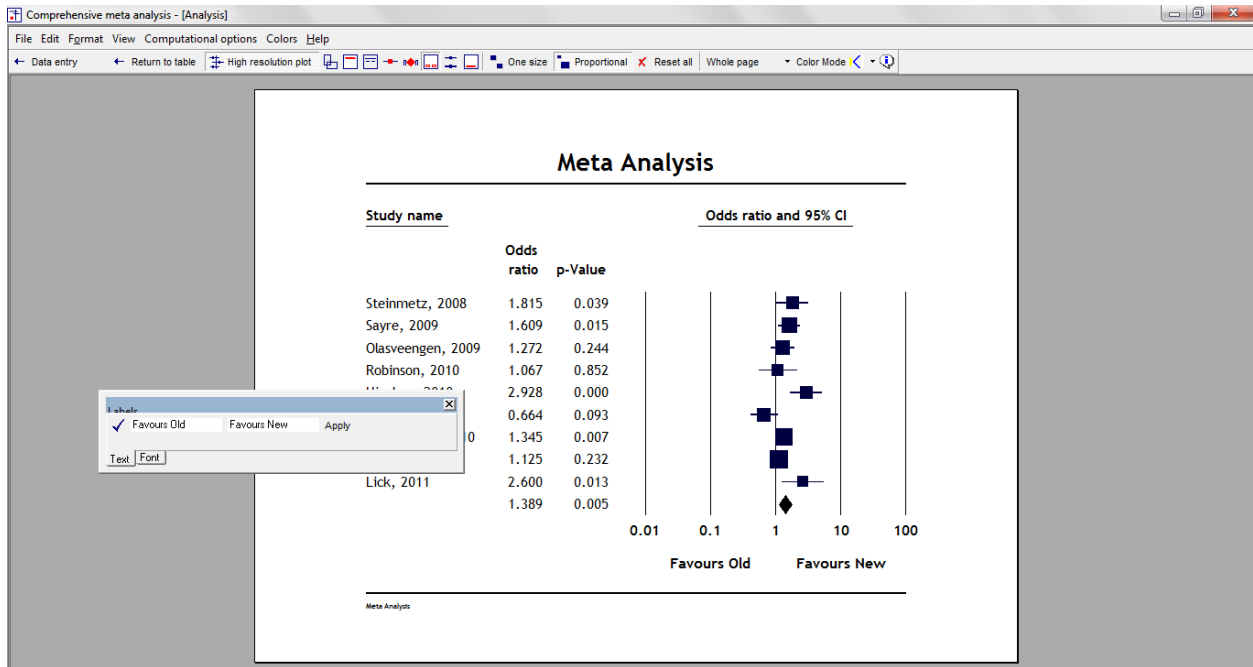
- If using Windows 7 you may need to drag the bottom edge of this (and other) toolbars to make it fully visible.
- Select [Right buffer] and then [Remove]
- This removes the space to the right of the forest plot (we don't need it since we have no columns to the right of the forest plot)




- Click [Forest plot width]
- Click the ↑↑ button. This will increase the width of the plot section.



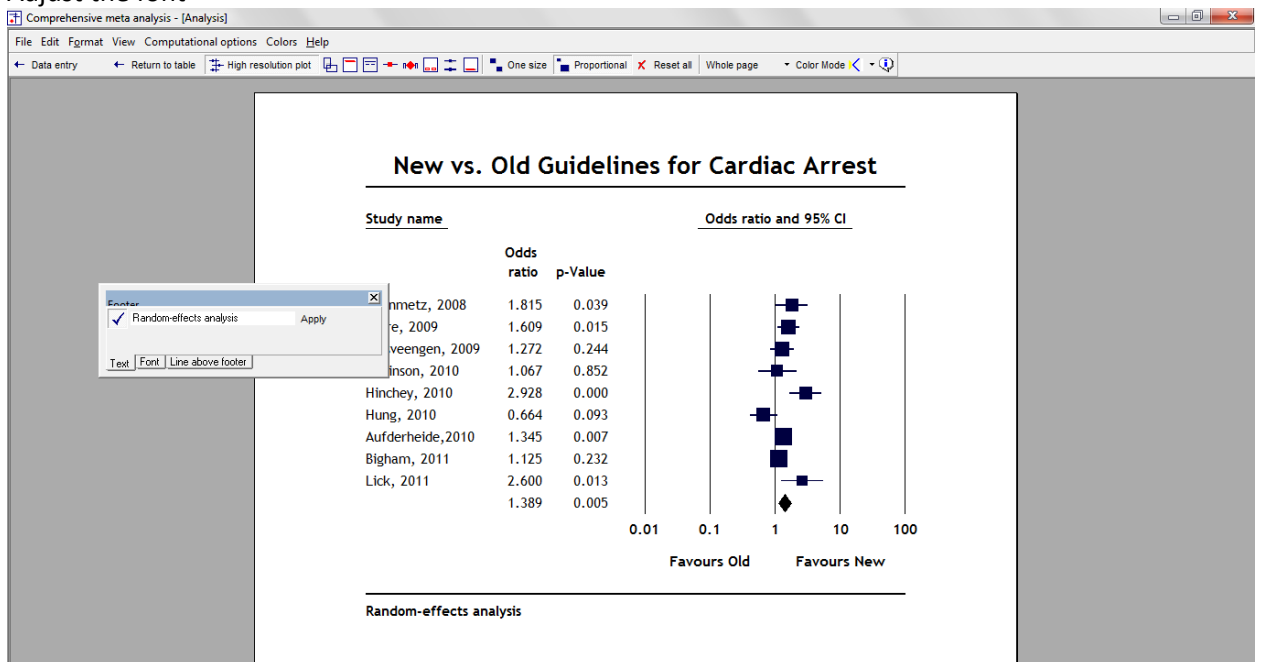
- Right-click on the words “Favours A” or “Favours B” to open a dialog
- Change these to “Favors Old” and “Favors New”
- Click [Apply]
- Click on [Font]
- Click ↑↑ to increase the font size




- Right-click on the title
- Set the title to “New vs. Old Guidelines for Cardiac Arrest” | Click [Apply]



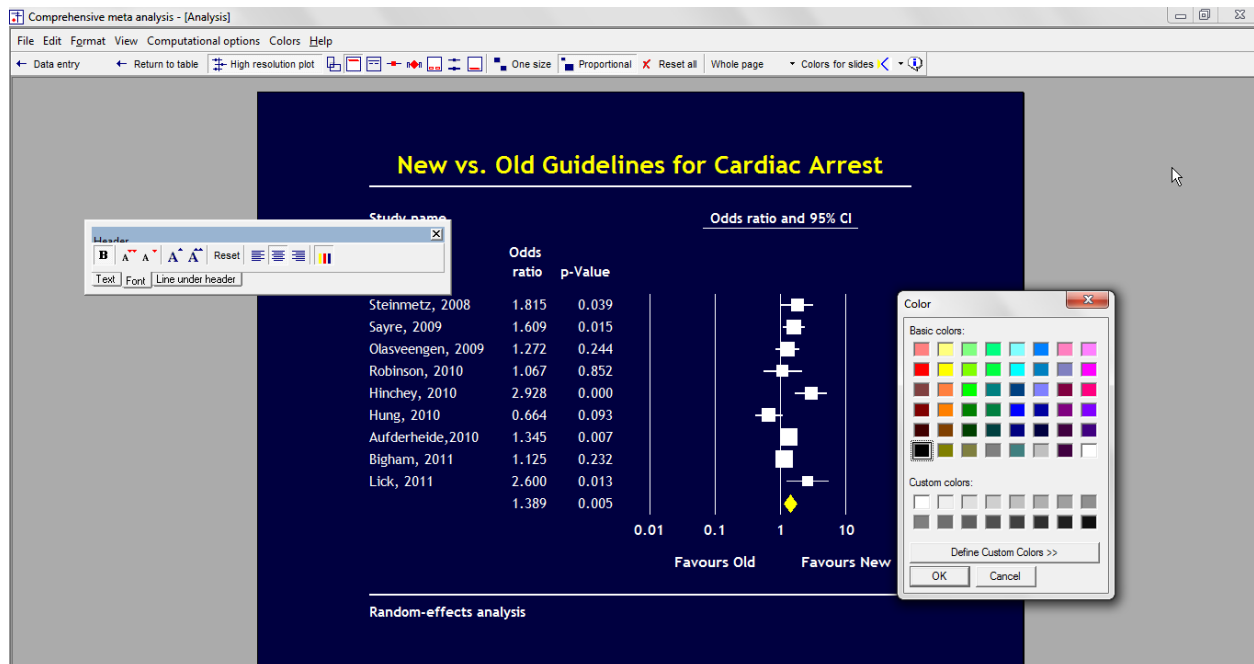
- Click Format > Footer
- Enter the text “Random-effects analysis”
- Depress the check-box at the left of this toolbar
- Adjust the font



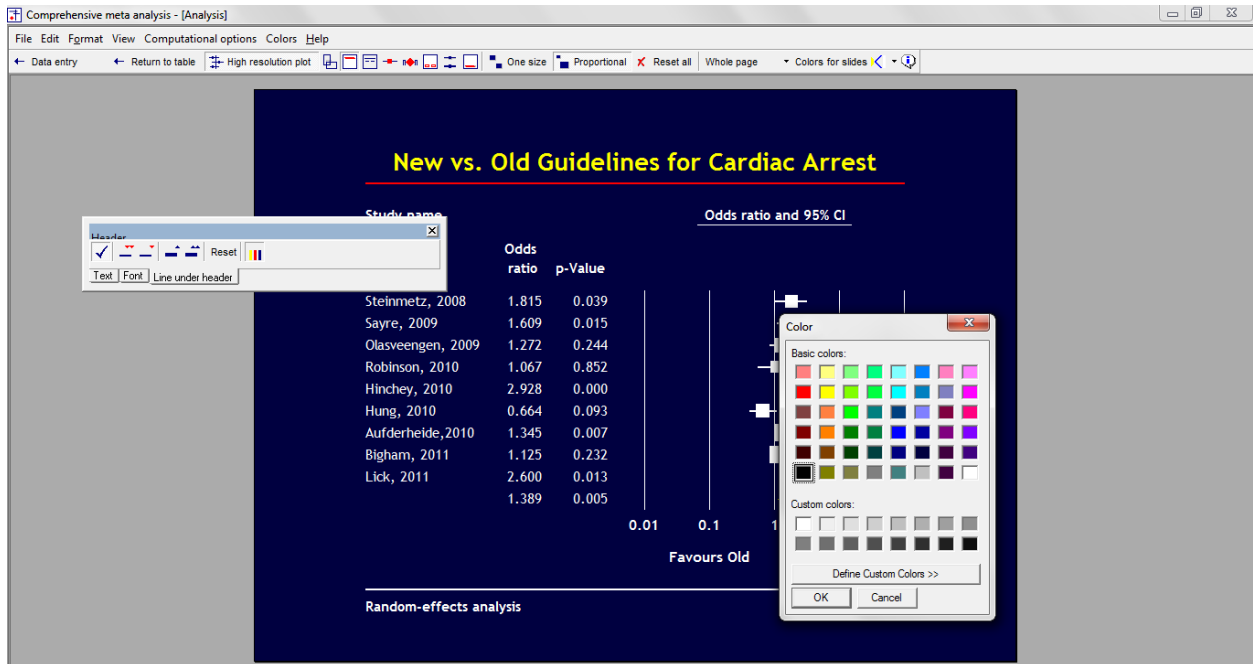
- Click Color for Slides



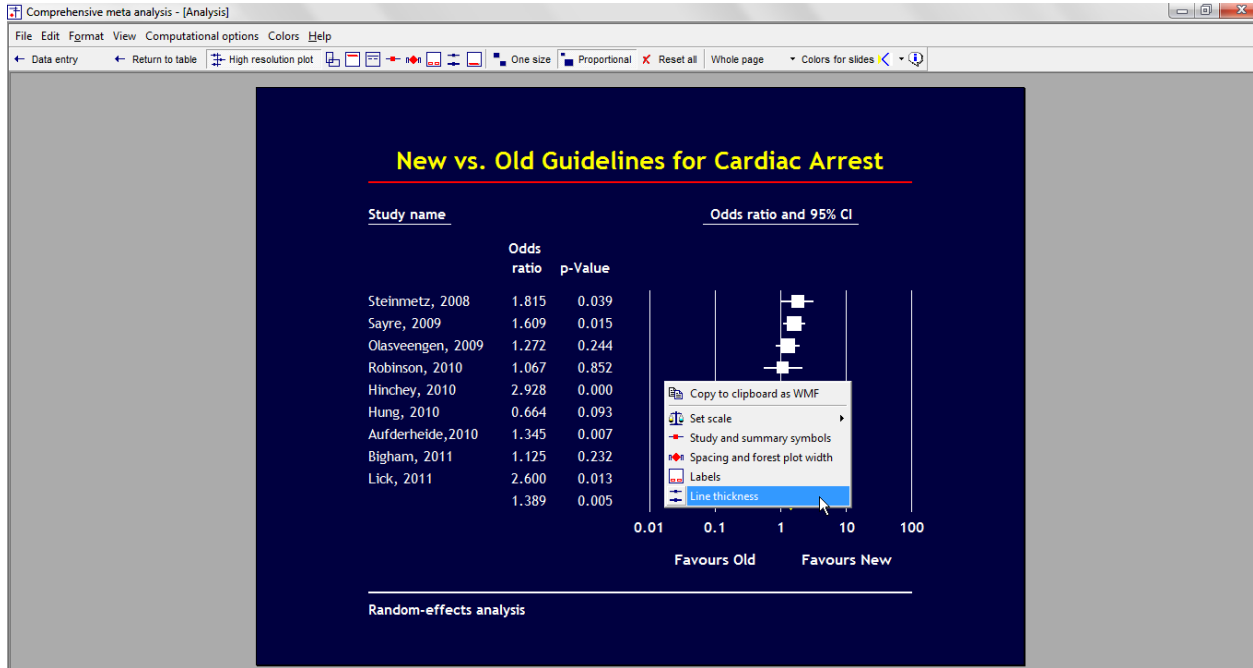
- Right-click on the title
- Click Font | Click the Color button
- Select a color | Click Ok



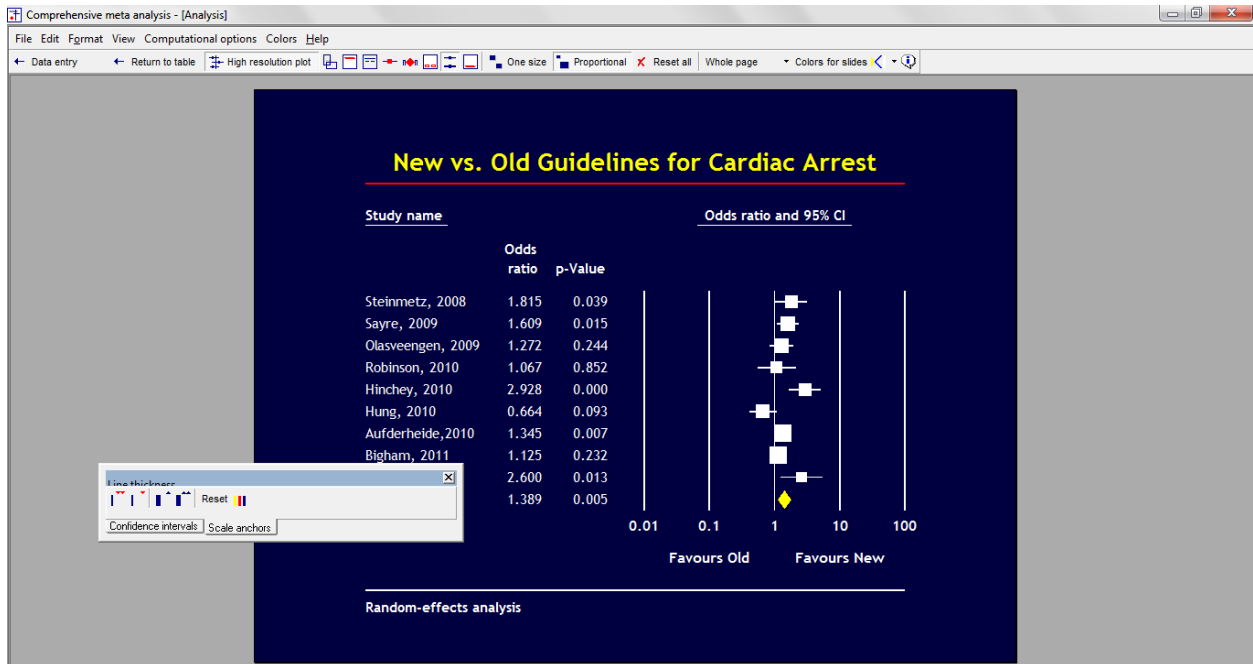
- Click Line Under Header
- Click the color tool | Select a color
- Click Ok




- Right-click on the plot
- Select Line thickness



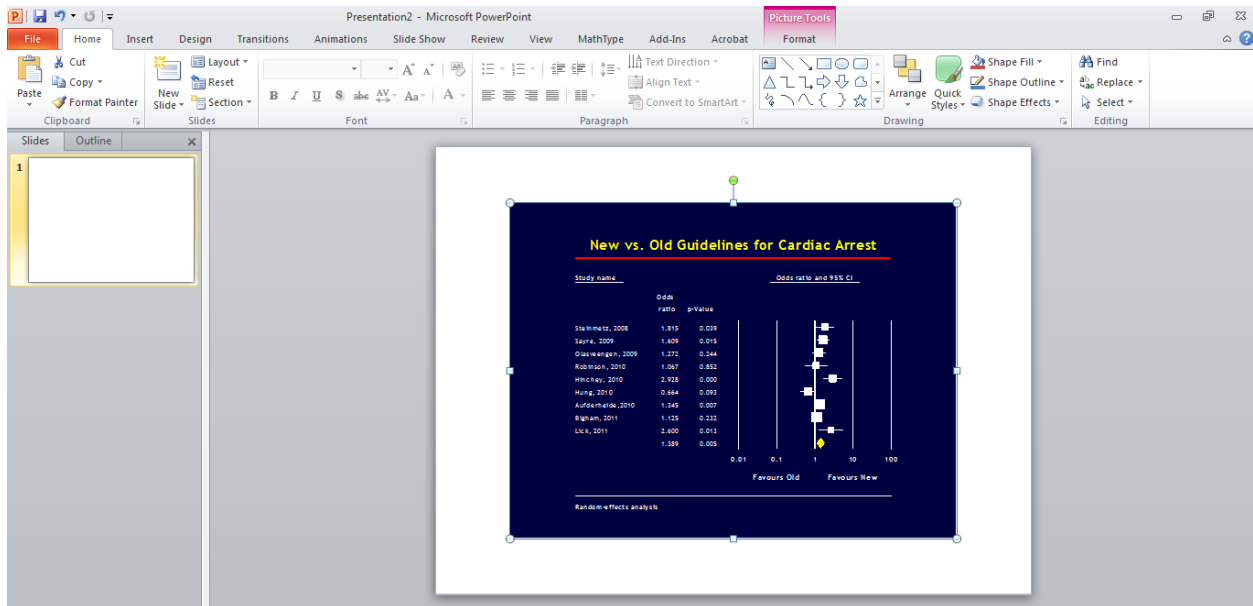
Select Scale Anchors and increase the size. This ensures that the vertical lines will be visible in PowerPoint.



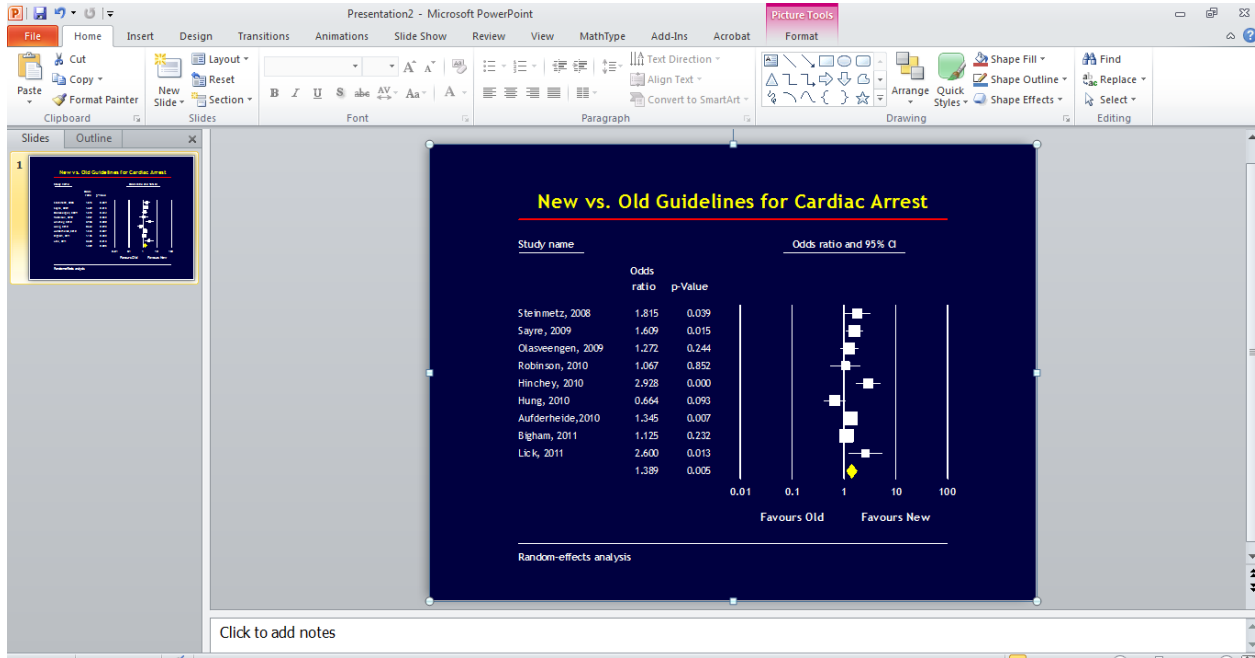
Click File > Export to PowerPoint



Look for the PowerPoint icon on your Windows toolbar and click



Depending on which version of PowerPoint you're using, you may need to stretch the plot to cover the full slide



Summary

This analysis includes nine studies where patients were randomized to be treated for cardiac arrest based on either old guidelines or new guidelines. Outcome was survival (being alive), and we used the odds ratio as the effect size.

Do the guidelines affect the likelihood of survival?

The mean odds ratio is 1.389, which means that being assigned to the newer guidelines resulted in a higher chance of surviving.

These studies were sampled from a universe of possible studies defined by certain inclusion/exclusion rules as outlined in the full paper. The confidence interval for the odds ratio is 1.104 to 1.748, which tell us that the mean odds ratio in the universe of studies could fall anywhere in this range. This range does not include an odds ratio of 1.0, which tells us that the mean odds ratio is probably not 1.0.

Similarly, the Z-value for testing the null hypothesis (that the mean odds ratio is 1.0) is 2.004, with a corresponding *p*-value of 0.005. We can reject the null that the likelihood of survival is the same in both groups, and conclude that the likelihood of survival is higher in the new-guidelines group.

Does the effect size vary across studies?

The *observed* effect size varies somewhat from study to study, but a certain amount of variation is expected due to sampling error. We need to determine if the observed variation falls within the range that can be attributed to sampling error (in which case there is no evidence of variation in true effects), or if it exceeds that range.

The *Q*-statistic provides a test of the null hypothesis that all studies in the analysis share a common effect size. If all studies shared the same effect size, the expected value of *Q* would be equal to the degrees of freedom (the number of studies minus 1).

The *Q*-value is 24.015 with 8 degrees of freedom and the corresponding *p*-value is 0.002. Thus, we can reject the null hypothesis that the true odds ratio is the same in all studies, and all the variation in observed effects is due to sampling error.

The I^2 statistic tells us what proportion of the observed variance reflects differences in true effect sizes rather than sampling error. I^2 is 66.689, which means that about two-thirds of the observed variance reflects variance in true effects. Put another way, if we could plot the true effects rather than the observed effects, the variance of the new plot would shrink by about a third.

T^2 is the variance of true effect sizes (in log units). Here, T^2 is 0.071 in log units. T is the standard deviation of true effects (in log units). Here, T is 0.267 in log units.