

Example name	Stents immediate failure
Effect size	Risk ratio
Analysis type	Simple analysis with between-studies variance
Level	Basic
Reference	Figure 2 in Kasapis et al (2009)

#### Synopsis

The goal of this analysis was to assess the use of stenting vs. angioplasty for patients with symptomatic superficial femoral-popliteal artery (SFPA) disease. The researchers looked at several outcomes, and found a difference for some but not others. Here, we focus on the short-term impact only.

The analysis includes eight studies where patients were randomized to receive either stent or angioplasty. Outcome was the proportion of patients that with immediate technical failure. The effect size is the odds ratio.

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 to perform a sensitivity analysis
- How to choose between fixed-effect and random-effects models
- How the statistical model affects the study weights
- How to interpret statistics for effect size
- How to interpret statistics for heterogeneity
- How to create a high-resolution plot

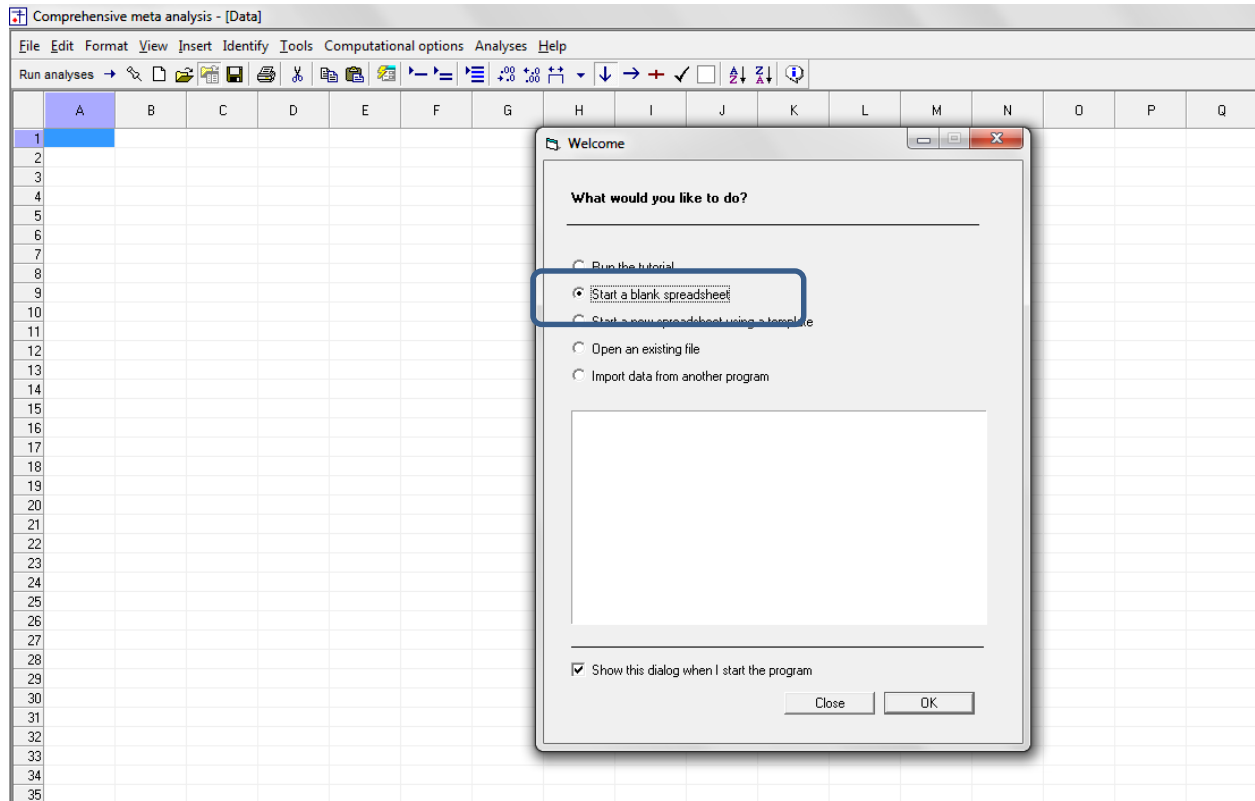
To open a CMA file > [Download and Save file](#) | [Start CMA](#) | [Open file from within CMA](#)

[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

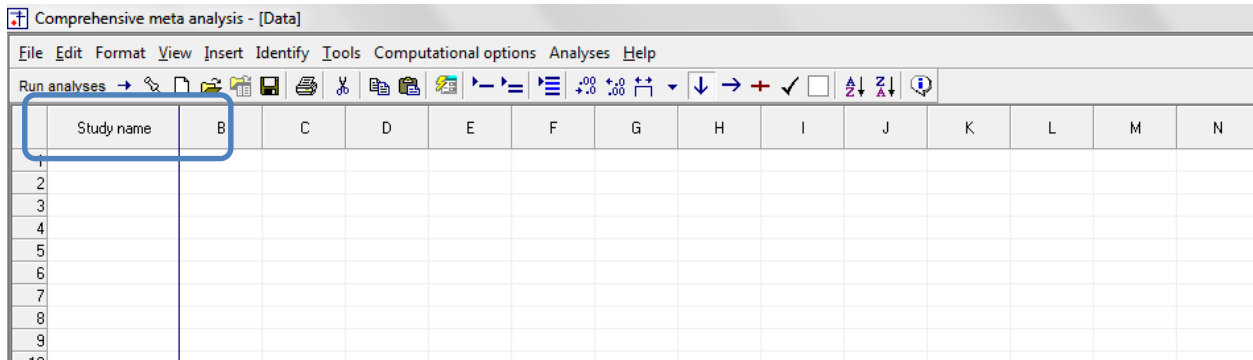
- Select the option [Start a blank spreadsheet]
- Click [OK]




Click Insert > Column for > Study names



The screen should look like this

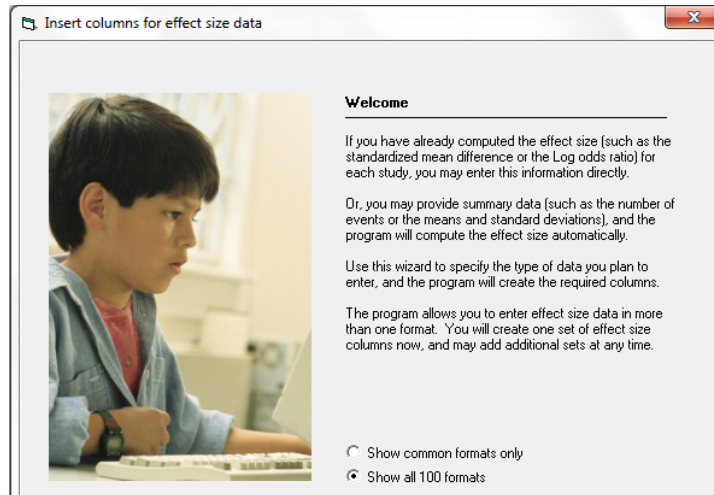


Click Insert > Column for > Effect size data

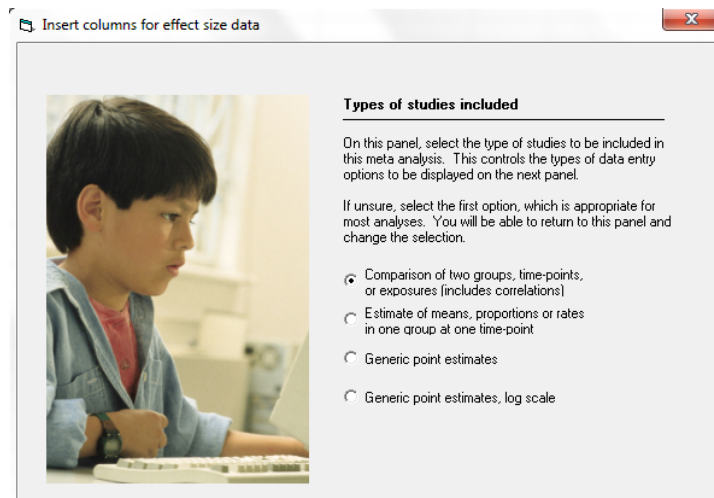


The program displays this wizard

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 > Stent
- Second group > Angio
- Name for events > Problem
- Name for non-events > Ok

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

The screenshot shows the 'Comprehensive meta analysis - [Data]' window. The main data grid has columns for 'Study name', 'Stent Problem', 'Stent Total N', 'Angio Problem', 'Angio Total N', 'Odds ratio', 'Log odds ratio', 'Std Err', 'Variance', and columns J through O. A blue box highlights the first five columns. A 'Group names' dialog box is open, with 'Stent' in the 'Name for first group' field, 'Angio' in the 'Name for second group' field, 'Problem' in the 'Name for events' field, and 'Ok' in the 'Name for non-events' field. The 'Apply' button is highlighted.

There are three options at this point

- Enter the data directly into CMA
- – or – Open the CMA data file “Stents short-term.cma”
- – or – Copy the data from Excel “Stents short-term.xls”


Here, we’ll show how to copy the data from Excel

- Switch to Excel and open the file
- Highlight the rows and columns as shown, and press CTRL-C to copy to clipboard

Study	Stent E	Stent N	Angio E	Angio N
Vroegindeweij, 1997	0	24	3	27
IntraCoil, 2001	25	177	32	175
Cejna, 2001	1	77	12	77
Becquemin, 2003	4	115	17	112
Saxon, 2003	0	15	1	13
Viabahn, 2005	6	97	16	100
Schillinger, 2006	1	51	17	53
Krankenber, 2007	6	123	25	121

Switch back to CMA

- Click in Cell Study name – 1
- Press [CTRL-V] to paste the data into CMA



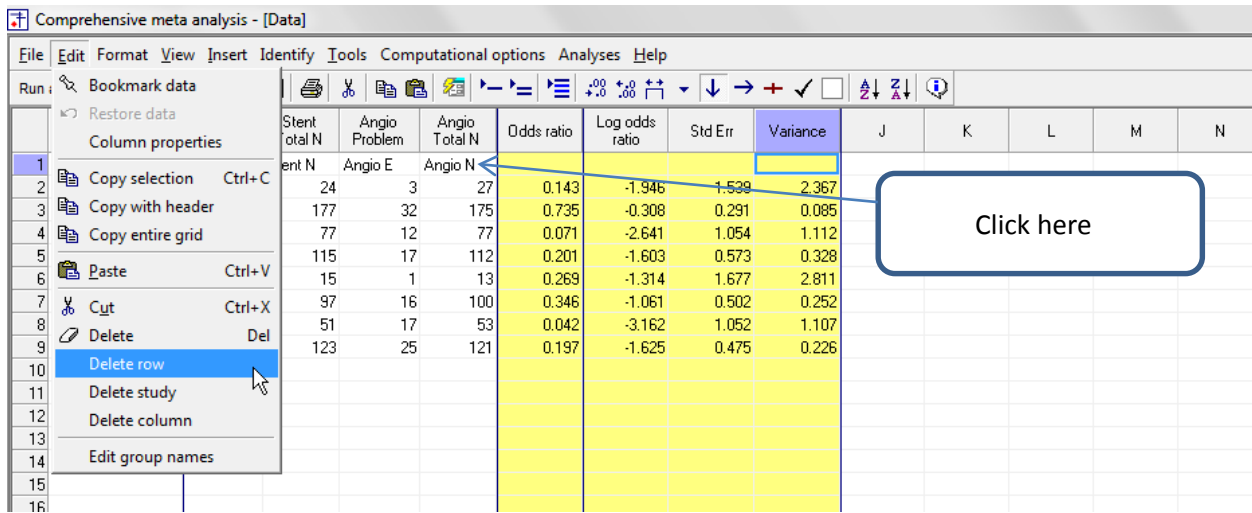
Study name	Stent Problem	Stent Total N	Angio Problem	Angio Total N	Odds ratio	Log odds ratio	Std Err	Variance
Study	Stent E	Stent N	Angio E	Angio N				
Vroegindeweij, 1997	0	24	3	27	0.143	-1.946	1.539	2.367
IntraCoil, 2001	25	177	32	175	0.735	-0.308	0.291	0.085
Cejna, 2001	1	77	12	77	0.071	-2.641	1.054	1.112
Becquemin, 2003	4	115	17	112	0.201	-1.603	0.573	0.328
Saxon, 2003	0	15	1	13	0.269	-1.314	1.677	2.811
Viabahn, 2005	6	97	16	100	0.346	-1.061	0.502	0.252
Schillinger, 2006	1	51	17	53	0.042	-3.162	1.052	1.107
Krankenber, 2007	6	123	25	121	0.197	-1.625	0.475	0.226

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 "Stent E" (for events) is now called "Stent Problem" (since events here represent a problem). 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.

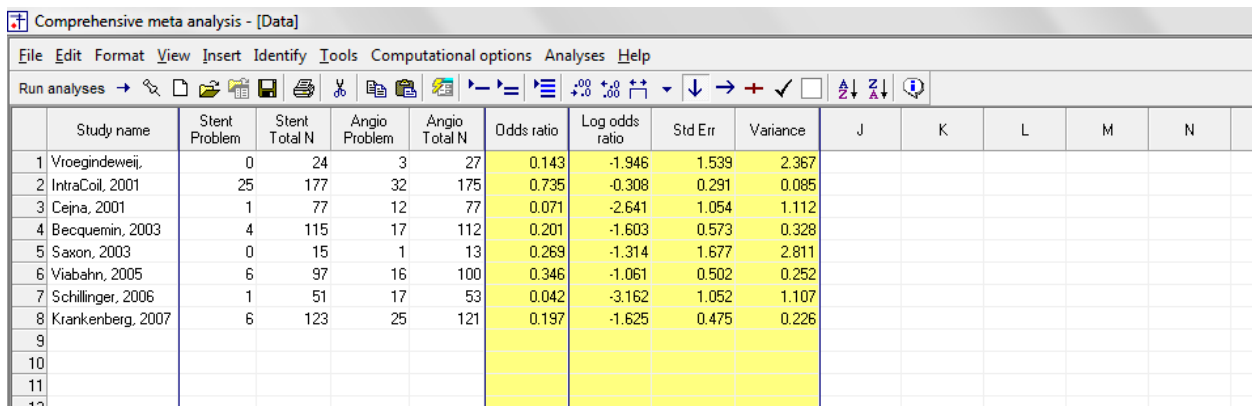
Now, we can remove the first row

- Click in the first row to select it



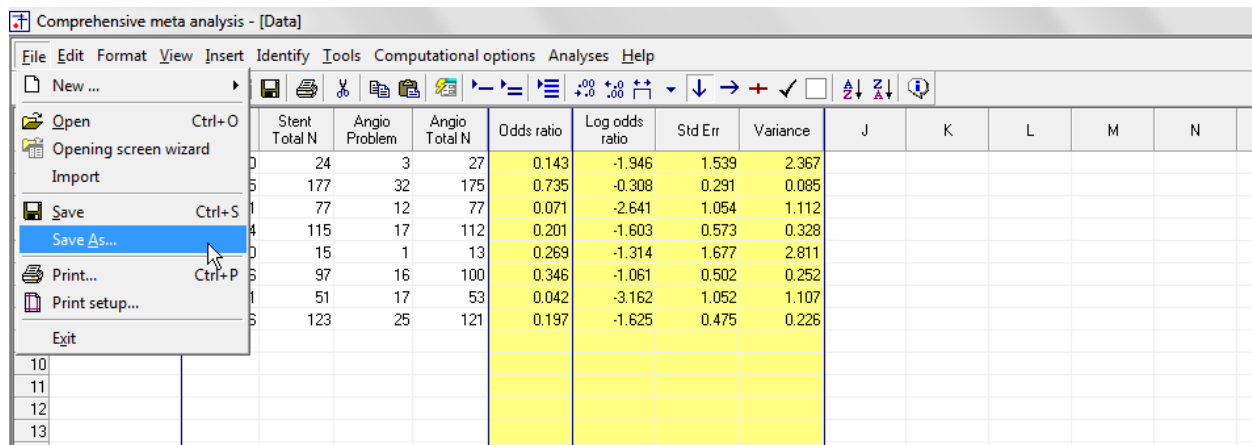
- Click Edit > Delete row and confirm

The screen should look like this



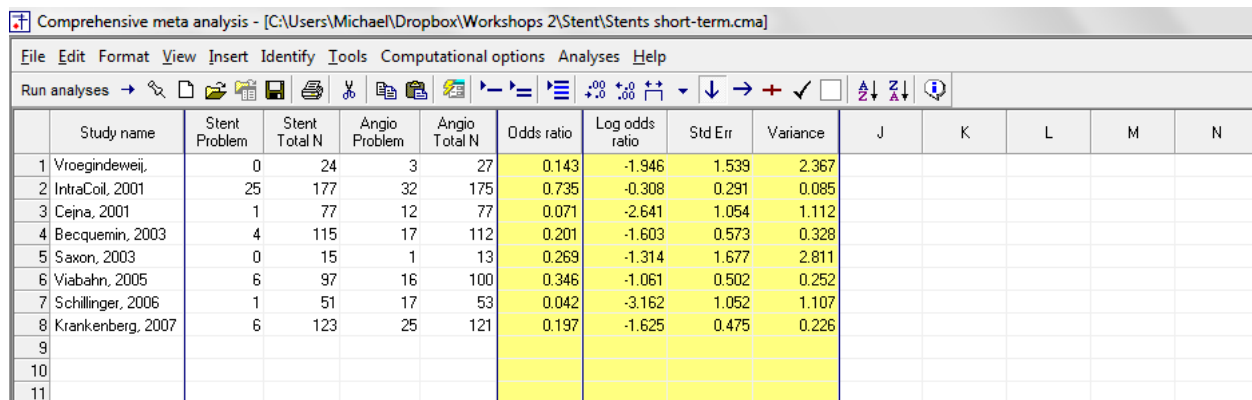
At this point we should save the file

- 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. We want to modify the screen to display the risk ratio.

- Right-click anywhere in the yellow columns
- Click [Customize computed effect size display]







- In the drop-down menu, select [Risk ratio]

Comprehensive meta analysis - [C:\Users\Michael\Dropbox\Workshops 2\Stent\Stents short-term.cma]

File Edit Format View Insert Identify Tools Computational options Analyses Help

Run analyses → [Icons]

	Study name	Stent Problem	Stent Total N	Angio Problem	Angio Total N	Odds ratio	Log odds ratio	Std Err	Variance	J	K	L	M	N	O
1	Vroegindewij	0	24	3	27	0.143	-1.946								
2	IntraCoil, 2001	25	177	32	175	0.735	-0.308								
3	Cejna, 2001	1	77	12	77	0.071	-2.641								
4	Becquemin, 2003	4	115	17	112	0.201	-1.603								
5	Saxon, 2003	0	15	1	13	0.269	-1.314								
6	Viabahn, 2005	6	97	16	100	0.346	-1.061								
7	Schillinger, 2006	1	51	17	53	0.042	-3.162								
8	Krankenber, 2007	6	123	25	121	0.197	-1.625								
9															
10															
11															
12															
13															
14															
15															
16															
17															
18															
19															
20															
21															
22															
23															
24															
25															
26															
27															
28															
29															
30															
31															
32															
33															
34															

Effect size indices

Use the following as the primary index

- Risk ratio
- Odds ratio
- Log odds ratio
- Peto odds ratio
- Log Peto odds ratio
- Risk ratio**
- Log risk ratio
- Risk difference

Std diff in means

- Log Peto odds ratio
- Risk ratio
- Log risk ratio
- Risk difference
- Std diff in means
- Hedges's g
- Difference in means
- Std Paired Difference
- Correlation
- Fisher's Z
- Rate ratio
- Log rate ratio
- Rate difference
- Hazard ratio

Also show standard error

Also show variance

Show the primary index only

Show all selected indices

Ok  
Cancel

- Un-check the box for [Odds ratio]
- Un-check the box for [Log odds ratio]
- Click [OK]

Comprehensive meta analysis - [C:\Users\Michael\Dropbox\Workshops 2\Stent\Stents short-term.cma]

File Edit Format View Insert Identify Tools Computational options Analyses Help

Run analyses → [Icons]

	Study name	Stent Problem	Stent Total N	Angio Problem	Angio Total N	Odds ratio	Log odds ratio	Std Err	Variance	J	K	L	M	N
1	Vroegindewei	0	24	3	27	0.143	-1.946							
2	IntraCoil, 2001	25	177	32	175	0.735	-0.308							
3	Cejna, 2001	1	77	12	77	0.071	-2.641							
4	Becquemin, 2003	4	115	17	112	0.201	-1.603							
5	Saxon, 2003	0	15	1	13	0.269	-1.314							
6	Viabahn, 2005	6	97	16	100	0.346	-1.061							
7	Schillinger, 2006	1	51	17	53	0.042	-3.162							
8	Krankenber	6	123	25	121	0.197	-1.625							
9														
10														
11														
12														
13														
14														
15														
16														
17														
18														
19														
20														
21														
22														
23														
24														
25														
26														
27														
28														
29														
30														
31														
32														
33														
34														

Effect size indices

Use the following as the primary index

Risk ratio

Display columns for these indices

- Odds ratio
- Log odds ratio
- Peto odds ratio
- Log Peto odds ratio
- Risk ratio
- Log risk ratio
- Risk difference
- Std diff in means
- Hedges's g
- Difference in means
- Std Paired Difference
- Correlation
- Fisher's Z
- Rate ratio
- Log rate ratio
- Rate difference
- Hazard ratio

Also show standard error

Also show variance

Show the primary index only

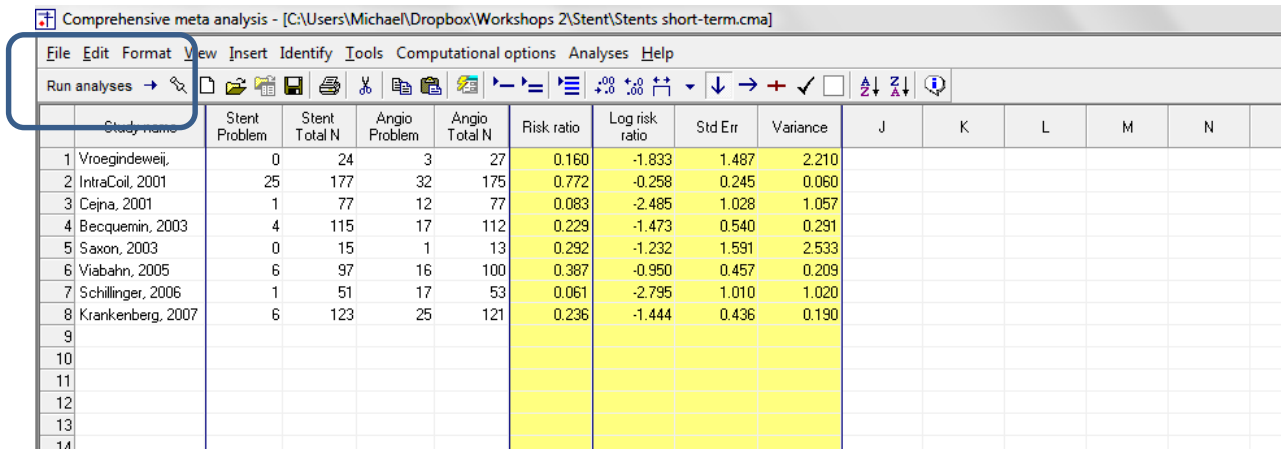
Show all selected indices

Ok

Cancel

The screen should look like this

- To run the analysis, click [Run analysis]



Comprehensive meta analysis - [C:\Users\Michael\Dropbox\Workshops 2\Stent\Stents short-term.cma]

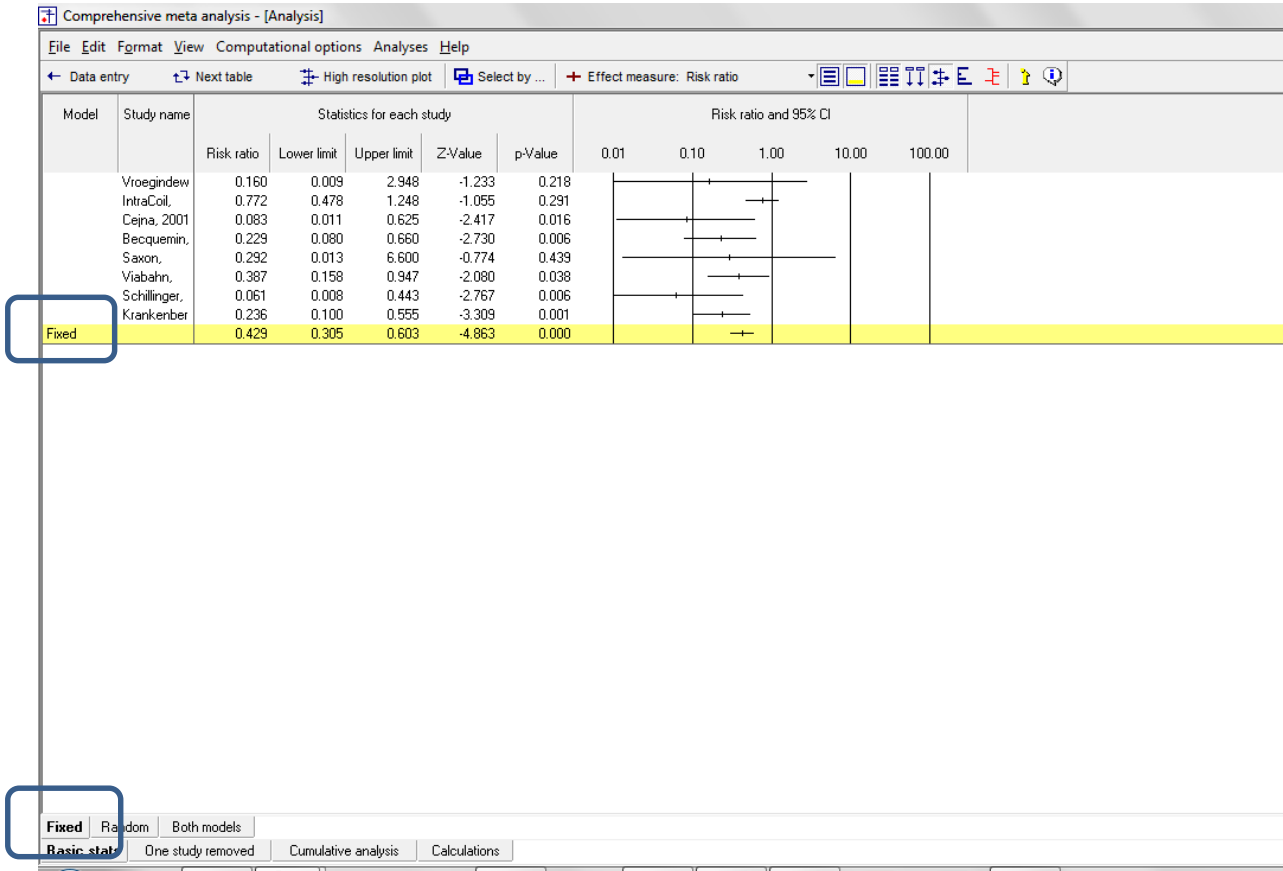
File Edit Format View Insert Identify Tools Computational options Analyses Help

Run analyses → [Icons]

Study name	Stent Problem	Stent Total N	Angio Problem	Angio Total N	Risk ratio	Log risk ratio	Std Err	Variance	J	K	L	M	N
1 Vroegindewei,	0	24	3	27	0.160	-1.833	1.487	2.210					
2 IntraCoil, 2001	25	177	32	175	0.772	-0.258	0.245	0.060					
3 Cejna, 2001	1	77	12	77	0.083	-2.485	1.028	1.057					
4 Becquemin, 2003	4	115	17	112	0.229	-1.473	0.540	0.291					
5 Saxon, 2003	0	15	1	13	0.292	-1.232	1.591	2.533					
6 Viabahn, 2005	6	97	16	100	0.387	-0.950	0.457	0.209					
7 Schillinger, 2006	1	51	17	53	0.061	-2.795	1.010	1.020					
8 Krankenberg, 2007	6	123	25	121	0.236	-1.444	0.436	0.190					
9													
10													
11													
12													
13													
14													

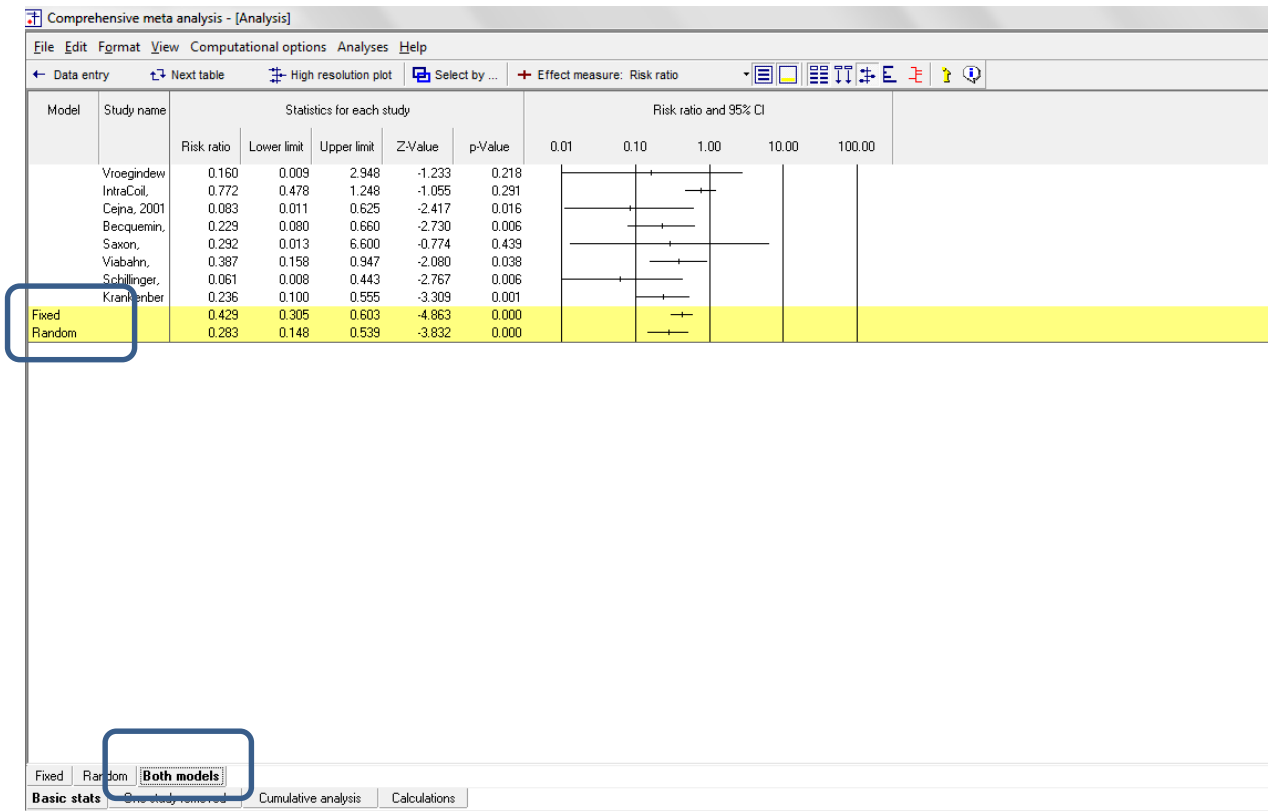
This is the basic analysis screen

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).

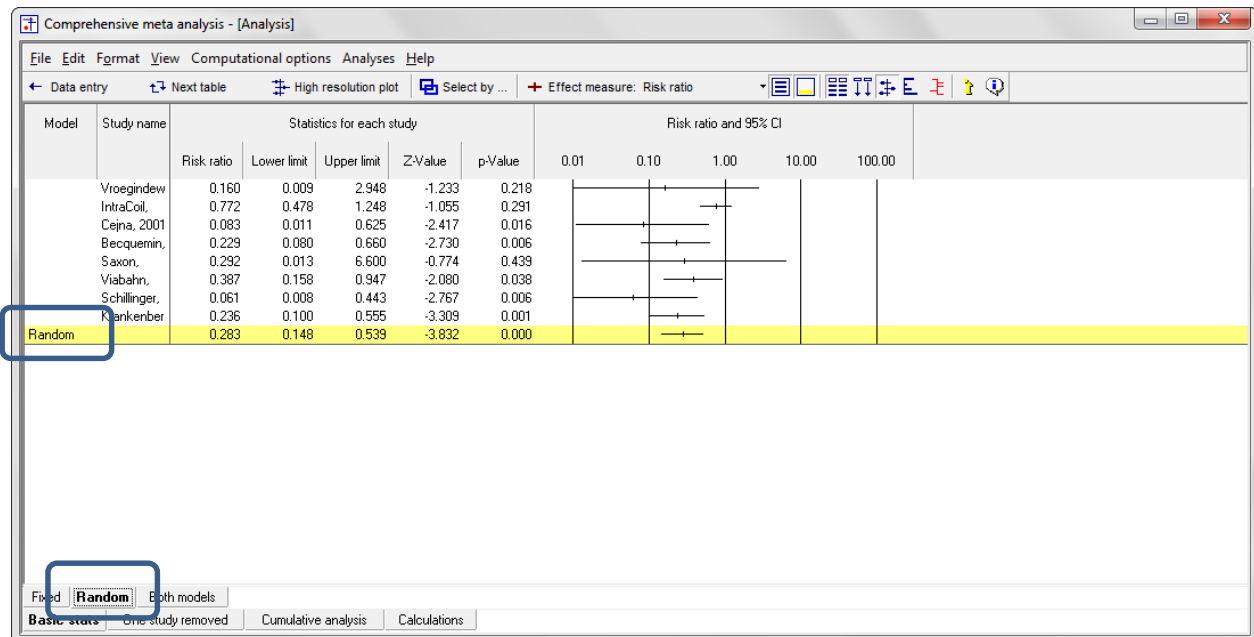
The confidence interval is wider for random-effects than for fixed-effects. This will always be the case when  $T^2$  is non-zero.

The random-effects estimate of the effect size is 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.



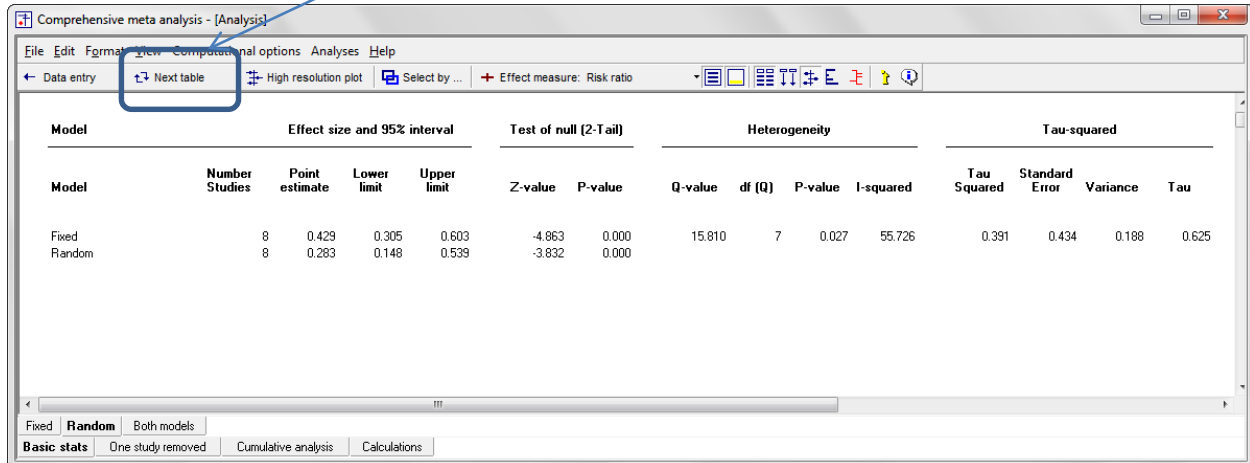
A quick view of the plot suggests the following

- Each of the studies suggests an advantage for stent over angioplasty
- In most studies this effect falls in the range of 0.10 to 0.80, though the CI for single studies extend well beyond that range
- The summary effect is 0.283 with a CI of 0.148 to 0.539. Thus, the mean effect is clearly in the clinically important range.
- The summary effect has a Z-value  $-3.832$  and a  $p$ -value of  $< 0.001$ . Thus we can reject the null hypotheses that the true risk 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 second 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



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	8	0.429	0.305	0.603	-4.863	0.000	15.810	7	0.027	55.726	0.391	0.434	0.188	0.625
Random	8	0.283	0.148	0.539	-3.832	0.000								

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

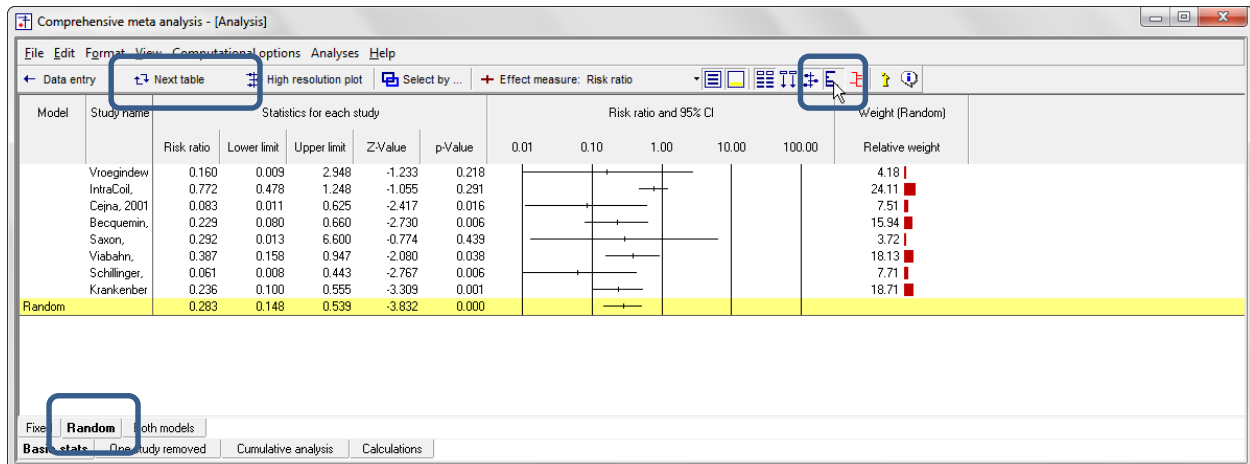
- Under the random-effects model the Risk ratio is 0.283 with a 95% confidence interval of 0.148 to 0.539. The test of the null (that the true risk ratio is 1.0) yields a Z-value of -3.823 and a corresponding p-value of < 0.001.
- The statistics at the upper right relate to the dispersion of effect sizes across studies.
- The Q-value is 15.810 with df=7 and p=0.027. 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 7). Here, Q exceeds that value, and exceeds it by enough that the difference is probably not due to chance. The p-value of 0.027 allow 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 0.391.  $T$  is the estimate of the between-study standard deviation in true effects. This estimate is 0.625. 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 risk 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 I2 value is 55.726, which tell us that about half the variance 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 half.

- 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 3.72% to a high of 24.11%. So, no single study dominated the analysis. Four studies seem to account for about 80% of the weights, with the other four sharing the remaining 20%.

Do any studies appear to be outliers?

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

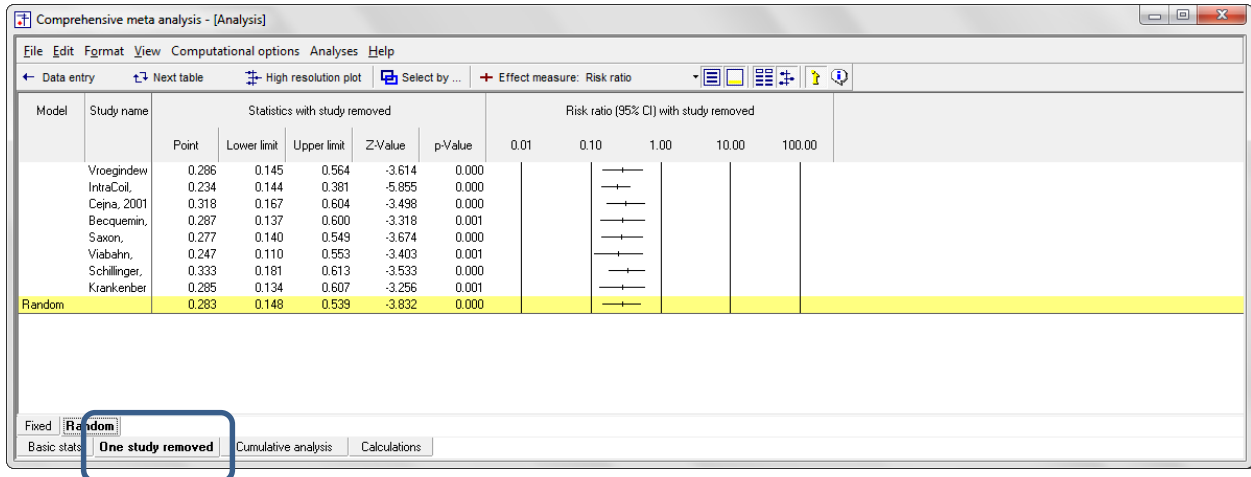
One study has a standardized residual of 1.72 which falls within the generally expected range. Also, this study's effect was closer to 1.0 than the summary effect, so it would have diminished the summary effect size rather than increased it.



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.



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 a risk ratio of 0.28. If we remove any one study the effect size stays in the range of 0.23 to 0.33, which is probably not a clinically important difference.

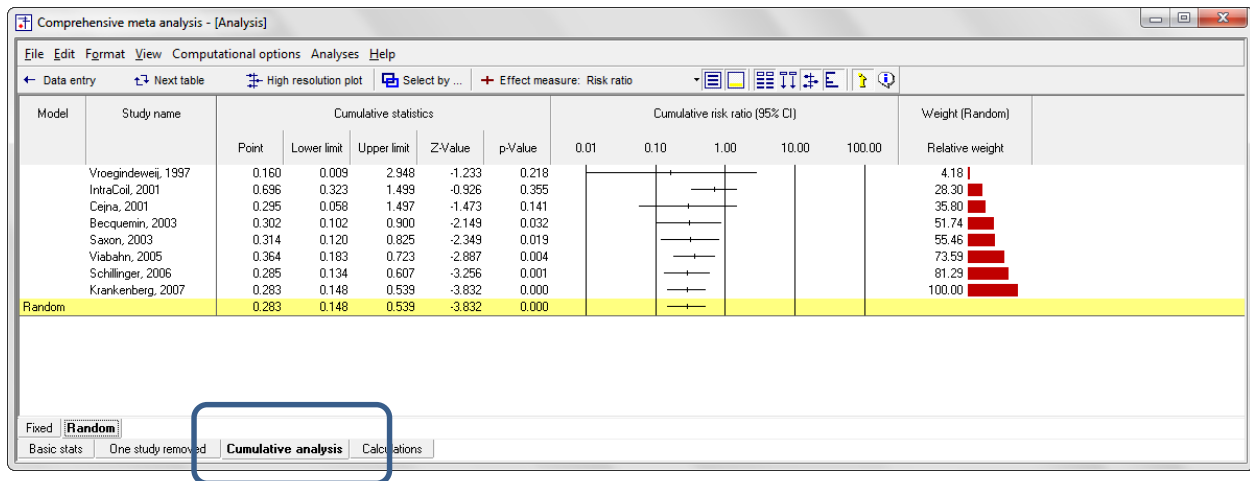
For those concerned about significance tests, the p-value with all studies is < 0.001, and it stays in that range with any single study removed.

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



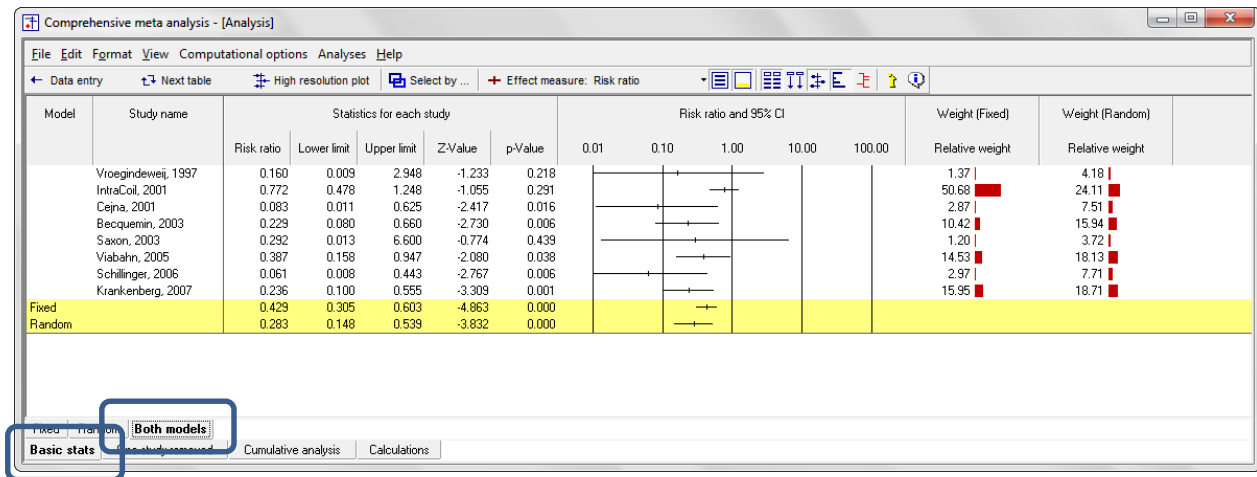
The risk ratio would have been estimated at .160 after the one study, at 0.696 after two studies, and at 0.295 after three studies. From that point on, the estimate is remarkably stable to the end.

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 0.027, 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 risk ratio moves to the left (a larger effect) when we move from fixed-effect weights to random-effects weights. The columns with the weights show why this happens. The second study happens to have an unusually small effect size relative to the others. This study also happens to have the largest sample size, by a wide margin. Under the FE model this study gets 50% of the weight, and pulls the summary effect to the right. Under the RE model this study gets only 24% of the weight. It still pulls the summary effect to the right, but not as much.

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 risk-ratio and p-value
- Uncheck all other boxes
- Click Ok

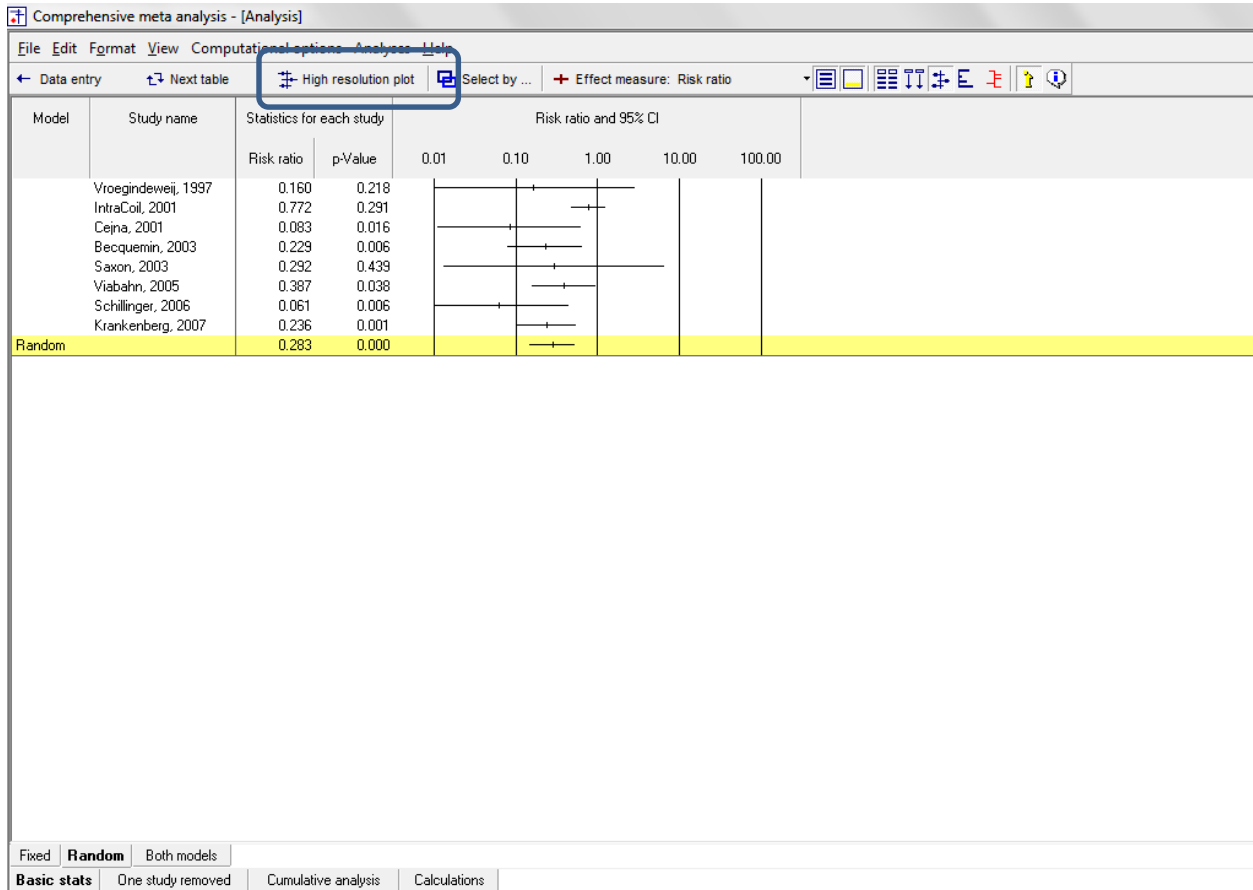
Right-click here

The screenshot shows the 'Comprehensive meta analysis - [Analysis]' window. The main area displays a forest plot with columns for 'Model', 'Study name', 'Risk ratio', 'Lower limit', 'Upper limit', 'Z-Value', 'p-Value', and 'Risk ratio and 95% CI'. The 'Random' model is highlighted in yellow. A 'Customize display' dialog box is open, allowing users to select which columns to show in the high-resolution plot. The 'Show' column has checkboxes for 'All columns in this block', 'Risk ratio', 'Standard error', 'Variance', 'Lower limit', 'Upper limit', 'Z-Value', and 'p-Value'. The 'Decimals' and 'Alignment' columns have dropdown menus. The 'Risk ratio' and 'p-Value' checkboxes are checked, while all others are unchecked. The 'Apply' button is highlighted.


Model	Study name	Statistics for each study					Risk ratio and 95% CI				
		Risk ratio	Lower limit	Upper limit	Z-Value	p-Value	0.01	0.10	1.00	10.00	100.00
	Vroegindewei, 1997	0.160	0.009	2.948	-1.233	0.218					
	IntraCoil, 2001	0.772	0.478	1.248	-1.055	0.291					
	Cejna, 2001	0.083	0.011	0.625	-2.417	0.016					
	Becquemini, 2003	0.229	0.080	0.660	-2.730	0.006					
	Saxon, 2003	0.292	0.013	6.600	-0.774	0.439					
	Viabahn, 2005	0.387	0.158	0.947	-2.080	0.039					
	Schillinger, 2006	0.061	0.008	0.443	-2.767	0.006					
	Krankenber, 2007	0.236	0.100	0.555	-3.309	0.001					
Random		0.283	0.148	0.539	-3.832	0.000					



- Click the tool for Hi-res plot

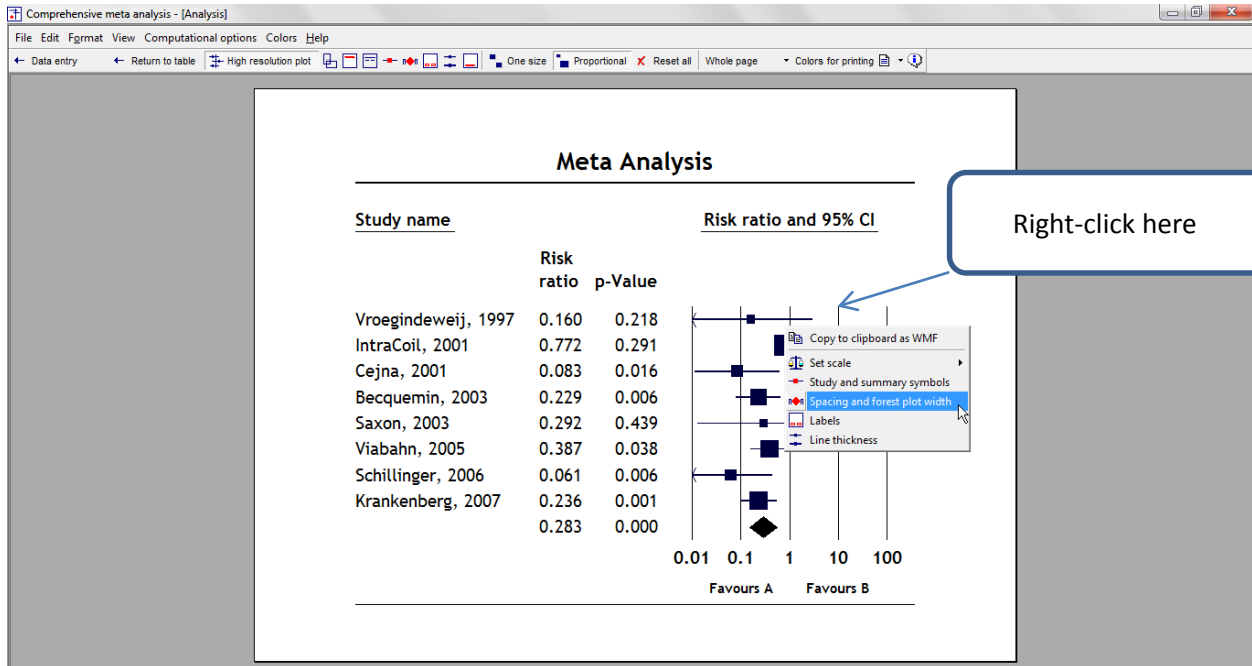


- On this screen click [Reset all]
- The screen should look like this




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]

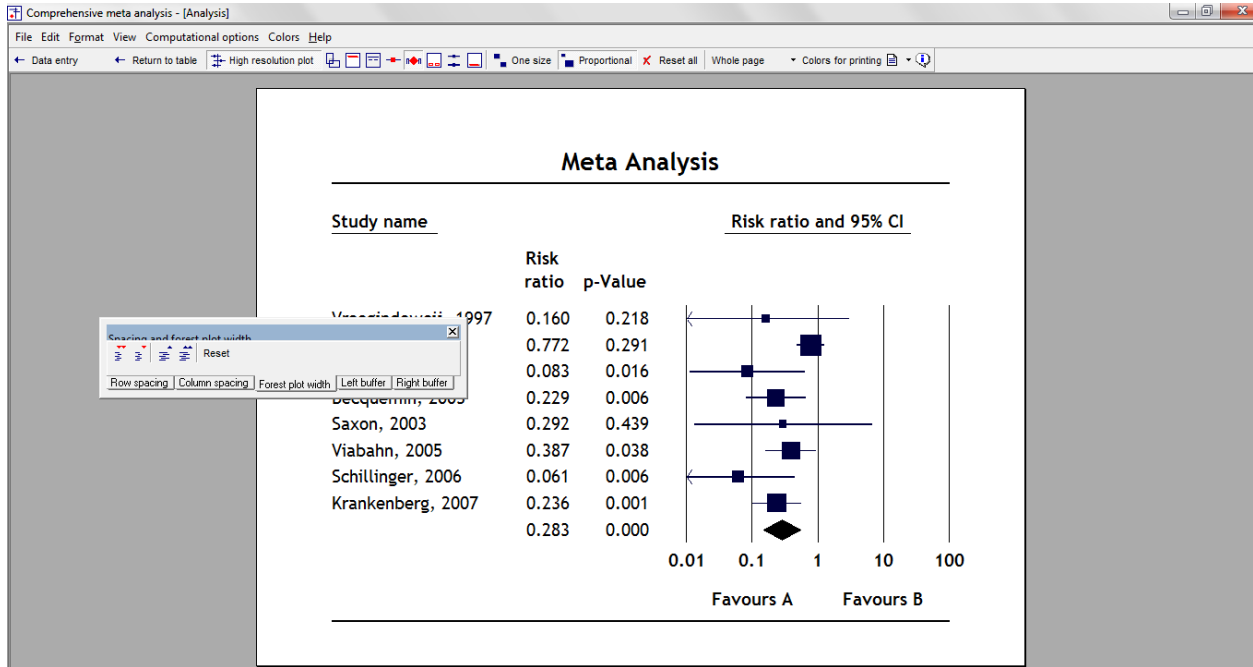


- If using Windows 7 you may need to drag the bottom edge of this (and other) toolbars to make it fully visible.
- Select [Right buffer] and then [Remove]
- This removes the space to the right of the forest plot (we don't need it since we have no columns to the right of the forest plot)




- Click [Forest plot width]

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



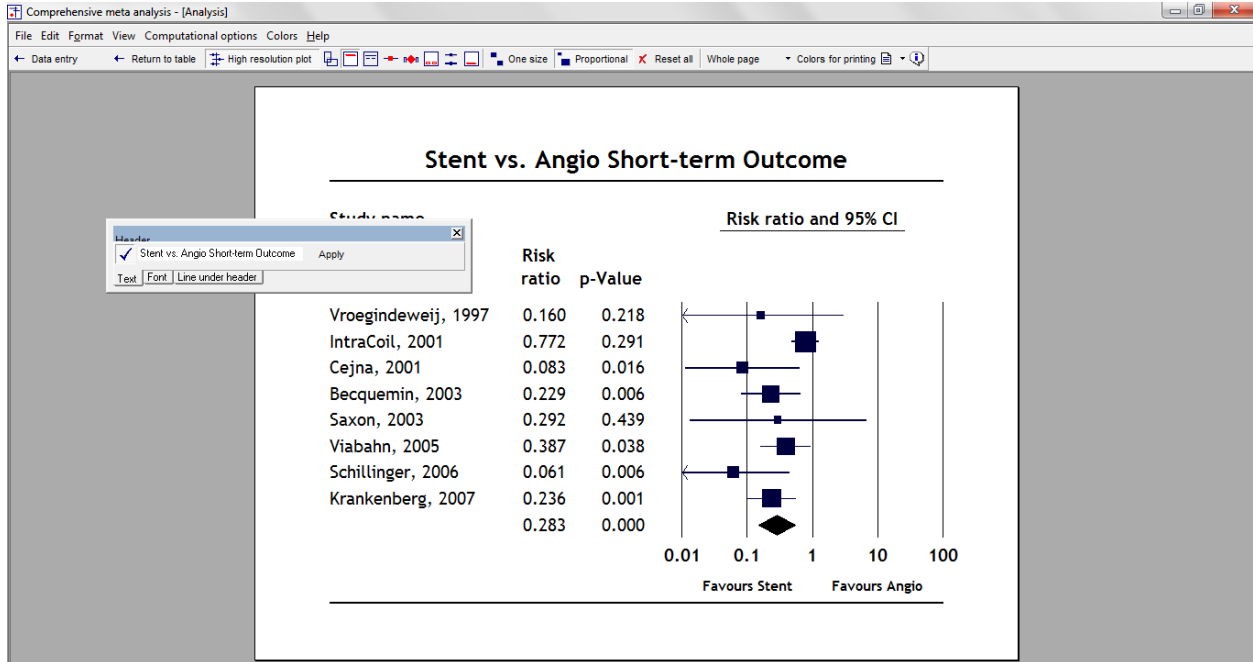
Right-click on the words “Favours A” or “Favours B” to open a dialog

- Change these to “Favors Stent” and “Favors Angio”
- If in the UK, Change these to “Favours Stent” and “Favours Angio”
- Click [Apply]
- Click on [Font]
- Click ↑↑ to increase the font size




Right-click on the title

- Set the title to “Stent vs. Angio Short-term Outcome”
- Click [Apply]




Click Format > Footer






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



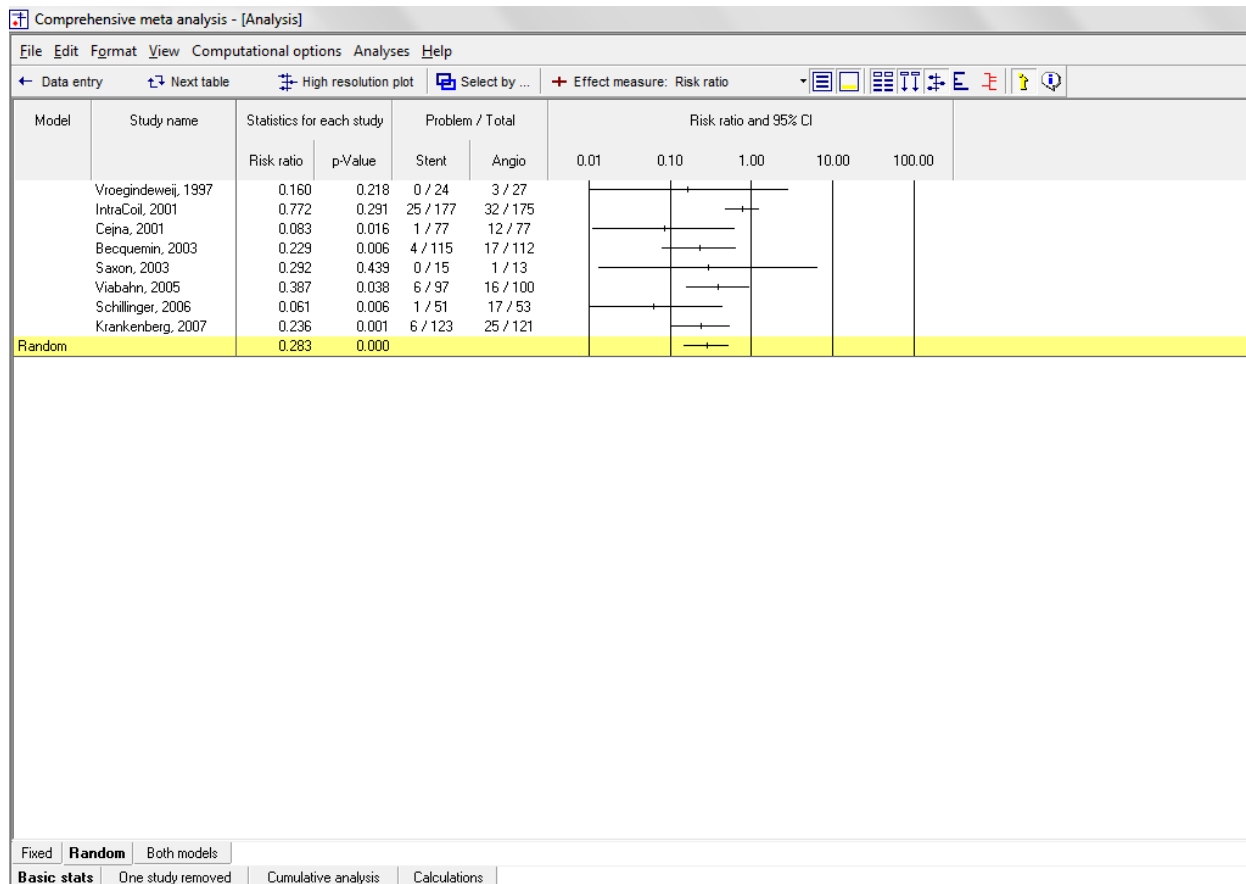


Click Return to table




We might want to include the number of events and patients in each study

- Click the button for [Show/Hide counts]
- The screen should look like this.
- This helps explain why the second study received more weight than the others – it has 57 events and 352 total N, which is more than any other study.
- Click High-res plot




Now, this additional information is included in the plot.


- The easiest way to reformat the plot may be to reset all parameters and then make adjustments from there.
- Click reset all




- Right-click on the plot
- Select Right-buffer > remove



Click Column spacing →→ to increase columns widths




Edit the names of the title, outcomes, and footer




It's always a good idea to double-check the labels, and ensure that the studies classified as "Favors Stent" did indeed have the lower event rate if the event is a bad outcome, as it is here – or a higher event rate if the event is a good outcome.



Click [Colors for slides]

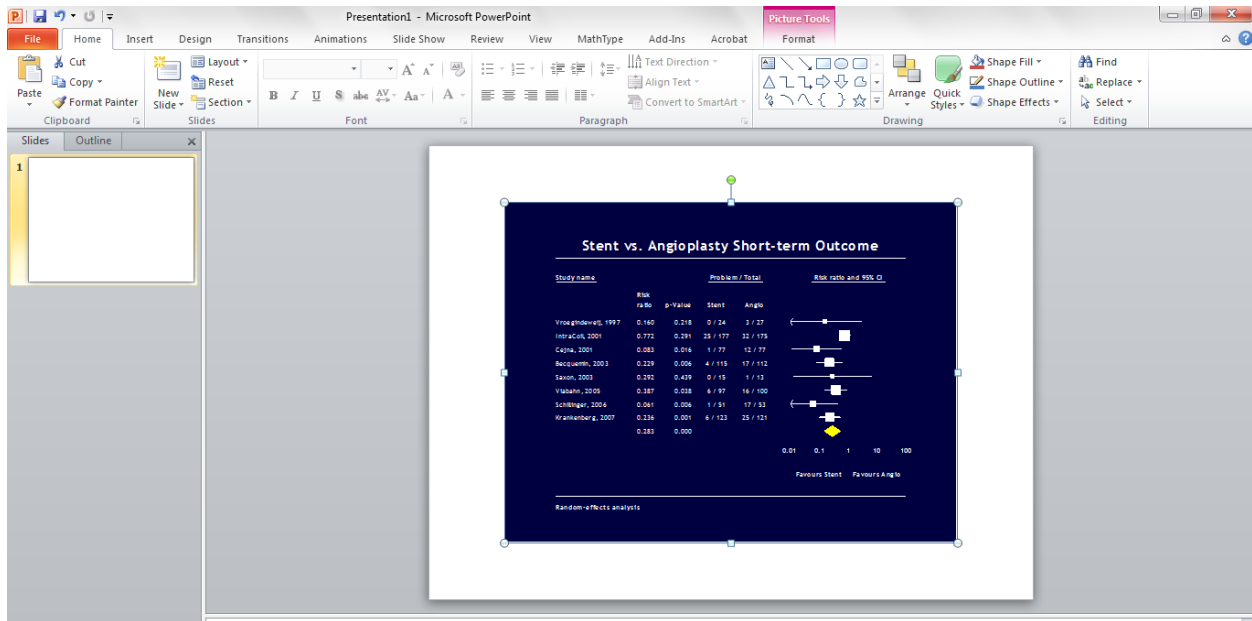


Click File > Export to PowerPoint

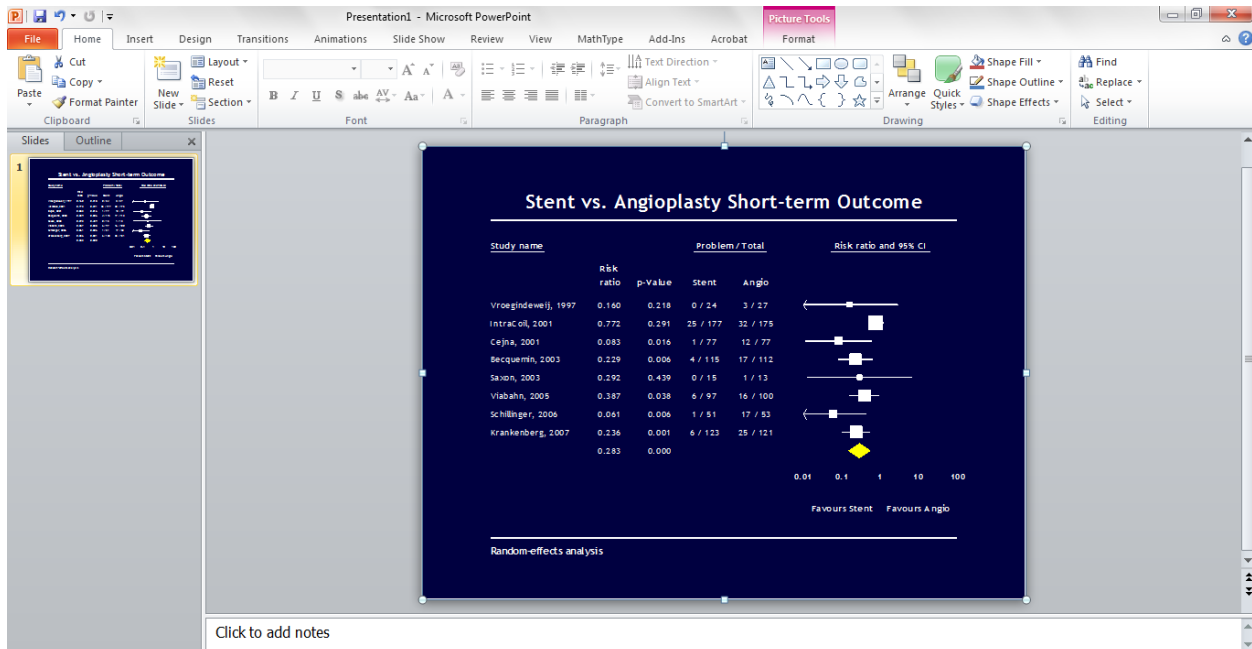


On your Windows toolbar, look for the PowerPoint icon

You should find that CMA has opened PowerPoint and inserted a copy of the plot



Depending on which version of PowerPoint you're using, you may need to stretch the plot to cover the whole slide. Stretch it from the corners, not the sides, to maintain the ratios.

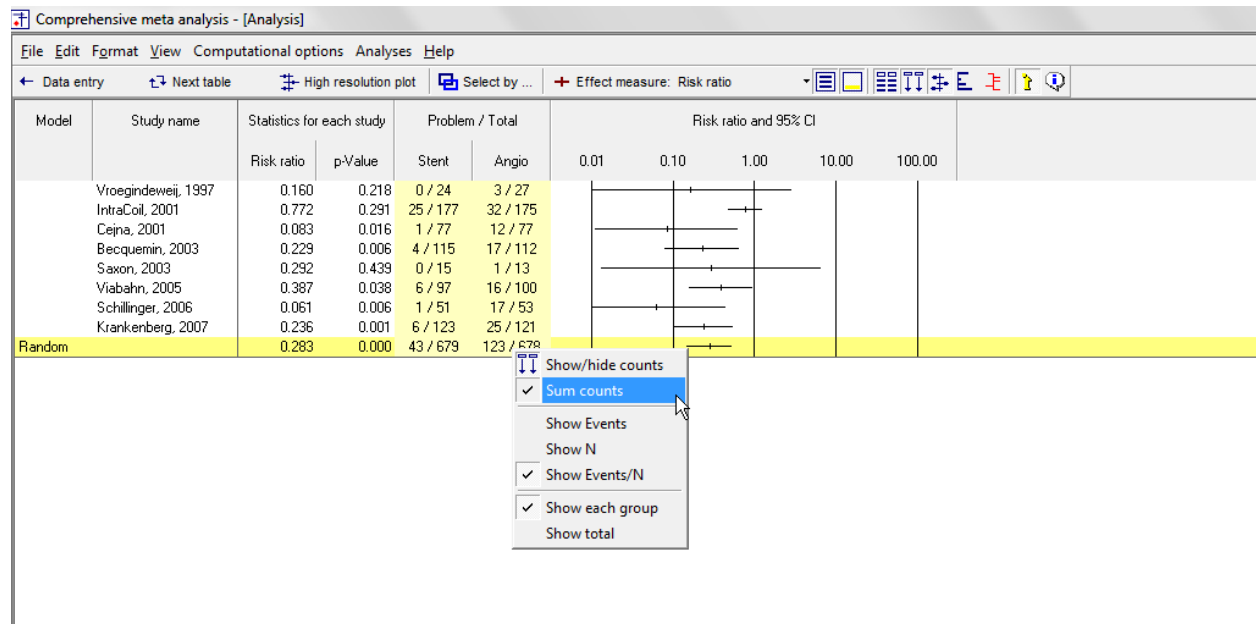




Click Next table to return here

You may have noticed that the program displayed the counts for each study but not the total. In general, this is a good idea because if we display the total counts some people might assume that we can use these counts to compute the risk ratio, when in fact we cannot. In this case the risk ratio based on counts of 43/679 vs 123/678 would be 0.349, while the correct value (using random-effect weights) is 0.283.

However, you can tell the program to display the counts by right-clicking on the counts and checking [Sum counts]. In this case, of course, the total would display on the high-resolution version as well.



## Summary

The analysis includes eight studies where patients were randomized to receive either stent or angioplasty. Outcome was the proportion of patients that required an immediate repeat procedure. The effect size is the risk ratio.

### Is treatment (Stent vs. Angioplasty) related to risk of immediate technical failure?

The mean risk ratio is 0.283, which means that stent was 72% less likely than angioplasty to result in immediate technical failure.

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 risk ratio is 0.148 to 0.539, which tells us that the mean risk ratio in the universe of studies could fall anywhere in this range. This range does not include a risk ratio of 1.0, which tells us that the mean risk ratio is probably not 1.0.

Similarly, the Z-value for testing the null hypothesis (that the mean risk ratio is 1.0) is -3.832, with a corresponding *p*-value is < 0.001. We can reject the null that the risk of immediate technical failure is the same in both groups, and conclude that the risk is lower in the stent group as compared with the angioplasty 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 15.810 with 7 degrees of freedom and *p* = 0.027. We reject the null hypothesis that the true effect size is the same in all studies. The variance in effect sizes is quantified by the following statistics.

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 55.726, which tells us that about 55% of the variance in the observed effects is due to variance in the true effects.

$T^2$  is the variance of true effect sizes (in log units). Here,  $T^2$  is 0.391.  $T$  is the standard deviation of true effects (in log units). Here,  $T$  is 0.625.