Module 11: Summary statistics for dichotomous outcome data

This module introduces some basic statistics for dealing with dichotomous outcomes. We discuss ways of summarising results and determining treatment effects within single trials, and prepare for meta-analysis.

Learning objectives

- Understand the difference between risk and odds
- Be able to calculate risk ratios and odds ratios from a 2×2 table for a single trial
- Be able to calculate risk differences and numbers needed to treat (NNT) from a 2×2 table for a single trial
- Be able to report and interpret risk ratios, odds ratios, risk differences and NNTs and their confidence intervals
- Be able to choose a suitable summary statistic for the meta-analysis and reporting of dichotomous data

Relevant sections of the *Reviewers' Handbook*

- Sections 8.3.1, 8.4.1

Where does this go in a Cochrane review?

The information in this module will be relevant to many parts of your review

- In the data analysis part of the Methods section of a protocol or review, where you will describe what statistical techniques you are planning to use
- When (or if) you actually perform meta-analyses using the analysis part of RevMan or other software
- In the presentation of results in the Results section of the review
- In the interpretation of results, in the Discussion of your review

Different types of data

There are several different types of data you may encounter when doing a systematic review

There are several different types of data you may come across in your included trials. Some of the more common data types are:

- **Dichotomous data** are data from outcomes that can be divided into two categories (e.g. dead or alive, pregnant or not pregnant), where each participant must be in one or other category, and cannot be in both
- **Counts of events** (for example number of epileptic fits)
- **Short ordinal scales** or scales with a small number of categories where there is a natural order to the categories (for example a pain scale of "none/mild/moderate/severe")
- **Long ordinal scales** or scales with a large number of categories with a natural order (for example the Short Form-36 scale for assessing quality of life or a depression index)
- **Continuous data** which are data from outcomes measured on a continuous scale (for example blood pressure, range of motion of a knee joint)
- **Censored data or survival data** (such as time to recurrence of cancer, where we will not have a measurement on everyone at the end of the study, because some haven't had a recurrence of cancer).

This module, and the next, will discuss issues of dichotomous data, as most Cochrane reviews contain data in this form. If you have continuous outcomes in your review then you will need to complete Additional Module 1 after you have completed Modules 11 and 12.

Often, data measured on different types of scales are converted to dichotomous data for analysis and presentation

Dichotomising data

We often make our own dichotomous data from outcomes that are not truly dichotomous, so that they are easier to manage and understand. For example, converting blood cholesterol (measured on a continuous scale) to 'high cholesterol' or 'not high cholesterol' dichotomised around a clinical threshold above which you would consider the cholesterol to be high; or converting pain measured on a short ordinal scale to 'absent or mild' or 'not absent or mild' (by which we mean moderate or severe). Generally long ordinal scales, or scales with a large number of discrete categories, are treated as continuous data for the purpose of analysis.

Sometimes, censored data are converted into dichotomous data by counting the number of people who have had the event by a particular time (such as the number of people who have a recurrence of cancer within 5 years of an operation). This should only be done when all participants have been followed up to the particular time point.

The benefits of converting non-dichotomous data into dichotomous data relate to ease of analysis and interpretation. Of these, the more important is ease of interpretation. Dichotomous outcomes may be easier for decision makers to understand and make judgements about.

The down side of converting other forms of data to a dichotomous form is that information about the size of the effect may be lost. For example a participant's blood pressure may have lowered when measured on a continuous scale (mmHg), but if it has not lowered below the cut point they will still be in the 'high blood pressure group' and you will not see this improvement. In addition the process of dichotomising continuous data requires the setting of an appropriate clinical point about which to 'split' the data, and this may not be easy to determine.

Summarising dichotomous data

In studies of treatment interventions we aim to describe what is happening to the participants we are studying so we can predict what is likely to happen to others. In order to do this we take an observation about a particular outcome for each participant. If that outcome is dichotomous, then each participant can be in one of two states. For example if we have a new drug thought to save lives in high risk patients, we might test it first in a group of people representative of these high risk patients and observe the number dead or alive at the end of the intervention. There are a couple of ways of summarising the information we get about the whole of our observed sample in a form that can be applied to others.

Risk

Risk can be applied to both a good and a bad event. It means the probability of that event.

The word *risk* is fundamental to epidemiology and evidence-based health. Risk is the chance, or probability, of having a specific event. As it relates to clinical trials and systematic reviews, risk is not always of a bad event, we can talk about the 'risk' of a good outcome (such as cure) as well as the 'risk' of a bad outcome (such as death). We can use the word 'risk' to describe the chance of the outcome whether it's good or bad.

Risk is the number of people with the outcome divided by the total number of people.

Given a single group of people, and knowledge of how many have 'the event', we can express the risk of the event by dividing the number with the event by the number of people. For example, of 133 women taking an antibiotic for the treatment of urinary tract infection (UTI), 14 had the event 'still infected' after 6 weeks. The risk of remaining infected was 14/133 = approximately 0.1.

Odds

Odds is the ratio of events to non-events

An alternative measure of describing how likely an event is to happen is called odds. The odds of an event is the ratio of events to non-events. Equivalently (and more formally) it's the risk of having an event divided by the risk of not having it. If we look at the 133 women taking the antibiotic for UTI, the ratio of events (still infected) to non-events (cured) is 14/119 = approximately 0.1. The more formal formula gives 14/133 (risk of having the event) divided by 119/133 (risk of not having the event), which also works out as 14/119 = approximately 0.1.

Risks and odds are similar when the event is rare, but differ when the event is common

In this example the risk and odds are both similar (approximately 0.1), so why bother to have two alternatives? In this example the 133 women taking antibiotics were the treated group in a clinical trial. In this trial there was also a placebo group with another 148 women. Of the 148 receiving placebos, 128 still had a UTI after 6 weeks. So in this group, what's the risk of staying infected? It's 128 (number with the event of 'still infected') /148 (total number in the group) = 0.86. What are the odds? They are 128 (number still infected) /20 (number cured) = 6.4. So in this case the risk and odds are very different.

In fact odds and risk are never identical, but they can be similar. They are similar when they are both small – i.e. when an event is rare. Taking antibiotics, the women rarely stayed infected (the event was rare), so the risk and the odds were similar. Not taking them, they mostly stayed infected (the event was common) so the odds and risk were different.

As values of odds and risks can differ for the same data, it is important to be careful and precise when using statistical summaries of dichotomous outcomes.

Comparing two groups

In the section above we have talked about risk and odds as it applies to a single group of people. In clinical trials, so we can assess the effect of an intervention over and above the natural course of a disease, we usually compare how people respond in an experimental group to how they respond in a control group (i.e. we compare two groups of people). When we are dealing with a dichotomous outcome, we can either compare the risk of having the event between the two groups, or compare odds between the two groups.

Risk ratio is the risk in the experimental group divided by the risk in the control group

Relative Risk or Risk Ratio

Relative risk or risk ratio (they mean the same thing and are both abbreviated as RR) is simply the risk of the event in one group divided by the risk of the event in the other group.

The most common way to go about calculating the risk ratio (and nearly all other statistics from dichotomous data) is to start by presenting your results in a 2x2 table, where each cell in the table contains the number of participants in each category.

Using a 2x2 table

	Event (Still infected)	No event (Not still infected)	Total
Intervention (Antibiotics)	*14*	*119*	*133*
Control (Placebo)	*128*	*20*	*148*

Now, if you think through what you are comparing (risk in the treated group with risk in the control group), the risk ratio is easy to calculate.

RR = risk in the treated group / risk in the control group

$$= \frac{\text{no. with event in treatment group}}{\text{no. in treatment group}} \Big/ \frac{\text{no. with event in control group}}{\text{no. in control group}}$$

= (14/133) / (128/148)
= 0.1/0.86
= 0.12

If an experimental intervention has an identical effect to the control, the risk ratio will be 1. If it reduces the chance of having the event, the risk ratio will be less than 1; if it increases the chance of having the event, the risk ratio will be bigger than 1. The smallest value the risk ratio can take is zero when there are no events in the treated group.

Odds Ratios

Odds ratio is the odds in one group divided by the odds in the other group

Just as odds are an alternative way of expressing how 'likely' events are in a single group, odds ratio is an alternative way of comparing how 'likely' events are between two groups.

The odds ratio is simply the odds of the event occurring in one group divided by the odds of the event occurring in the other group. If we take the same data from our 2x2 table above,

OR = odds in the treated group / odds in the control group

$$= \frac{\textit{no. with event in treatment group}}{\textit{no. without event in treatment group}} \Big/ \frac{\textit{no. with event in control group}}{\textit{no. without event in control group}}$$

= (14/119)/ (128/20)
= 0.118/6.40
= 0.018

If an experimental intervention has an identical effect to the control, the odds ratio will be 1. If it reduces the chance of having the event, the odds ratio will be less than 1; if it increases the chance of having the event, the odds ratio will be bigger than 1. The smallest value an odds ratio can take is zero.

Swapping the "good" and "bad" outcomes when calculating RR may make a difference. There are effectively two RRs

The difference between good and bad outcomes

Most dichotomised outcomes will be a dichotomy between a good and a bad event. When we describe risk, it can refer to the risk of having a good event or the risk of having the bad event, so 'reducing risk' could be a good or a bad thing. It is important whether we define 'the event' as the good outcome or the bad outcome as the results can change if we swap the good and bad outcomes around.

Taking the UTI example again, suppose we decide to define the event as cure. The risk in the antibiotic group is now 119/133 = 0.895 (i.e. 119 women were no longer infected) and in the placebo group it is 20/148 = 0.135 (i.e. 20 women were no longer infected). The risk ratio is therefore 0.895/0.135 = 6.6

Remember we previously calculated the risk ratio for remaining infected as 0.12. By swapping the good and bad outcomes we have changed the risk ratio from 0.12 to 6.6, but there is no simple relationship between these numbers. This makes it difficult to calculate one from the other without going back to the original data.

This means there are essentially two risk ratios: the risk ratio for a good outcome and the risk ratio for a bad outcome. There is quite a lot of work being done on this issue in the Cochrane Collaboration at the moment, but the general rule is that for outcomes which we aim to **prevent** (e.g. death, recurrence or worsening of symptoms), it is best to report the event as the bad outcome, which is usually the intuitive choice. For outcomes where we are trying to improve health (e.g. healing, resolution of symptoms, clearance of infection), we still do not know which option is best, but you should be very clear in your results section which outcome you are presenting. These rules are based on analysing which statistic is the most consistent – an issue we will discuss in more detail shortly.

Odds ratios of good and bad events are related to each other

For odds ratios, you still need to choose which outcome is the most appropriate to present, but it is easier to convert from 'good' to 'bad' outcomes or vice versa. From the example above you will see that the odds ratio of the "good event" of cured is
OR = odds in the treated group / odds in the control group
=(119/14) / (20/128)
= 8.50 / 0.156
=54

From above we know the odds ratio when using the event 'still infected" is 0.018, and the odds ratio when using 'infection cleared' as the event is 54. These numbers are inversely related (working with more accurate numbers, we find 0/0.01838 = 54.4 and 1/54.4 = 0.01838) and this is always the case with odds ratios. So in some senses it doesn't matter whether we choose good or bad outcomes if we use odds ratios. Whichever we choose, it is vitally important that the results are very clearly reported so that those using the review are clear which outcome they are looking at.

When do risk ratios and odds ratios differ?

It is really important to make clear whether the statistic you are presenting is an odds ratio or a risk ratio. As we saw when we looked at using odds and risks to summarise events in a single group, the risk and odds can be very different. So too with risk ratios and odds ratios.

In general an odds ratio will always be further from the point of no effect (where OR=1, RR=1) than a risk ratio. If the event rate increases in the treatment group, the OR and RR will both be greater than 1, but the OR will be bigger than the RR. If the event rate decreases in the treatment group, both the OR and the RR will be smaller than 1, but the OR will be smaller than the RR.

OR and RR may differ when the event is common.

Odds ratios and risk ratios will be similar when the event is rare, but will differ (often by a lot) when the event is common. In situations of common events, the odds and odds ratio can be misleading, because people tend to interpret an odds ratio as if it were a risk ratio. Trials usually study frequent events, so this is a very real issue.

Later on in this module we will discuss how to chose the appropriate statistic in your review.

Another word of caution about both these measures: because the result is expressed as a proportion of the event rate in the control group, it is not possible to determine the actual number of participants who benefited. For example, a RR of 0.5 can mean a risk is decreased from 40% in one group to 20% in the other, or it can mean a 2% in one group and 1% in the other. In both cases the risk is halved by the intervention, but the actual change in the number of events is very different. Because of this, it may also be useful to express results in absolute terms. One way of doing this is to report a risk difference; another is to report the number needed to treat.

Risk Difference is the risk in the treated group minus the risk in the control group

Risk Difference

As well as comparing risks in relative terms (i.e. risk in one group divided by the risk in the other), we can also compare them in terms of the *absolute* difference between the two groups (i.e. the risk in one group minus the risk in the other). This we call the risk difference, or absolute risk difference (it means the same thing).

Risk difference is calculated as risk in the experimental group minus risk in the control group. For our example this is:

RD = Risk in antibiotic group – Risk in placebo group
= 0.10 – 0.86
= –0.76

The risk difference describes the absolute change in risk that is attributable to the experimental intervention. If an experimental intervention has an identical effect to the control, the risk difference will be 0. If it reduces risk, the risk difference will be less than 0; if it increases risk, the risk difference will be bigger than 0. The risk difference cannot be above 1 or below –1. Switching between good and bad outcomes for the risk difference causes a change of sign, from + to – or – to +.

Sometimes it may be useful to present figures for 100 times the RD, or 1000 times the RD, which describe how many people have avoided (or incurred) the event for every 100 or 1000 treated, respectively.

Number needed to treat

Another way of looking at the risk difference is the number needed to treat (NNT). Where we are trying to prevent an event, and the risk difference is less than 0 (i.e. the intervention reduces the risk of the event), NNT is the inverse of the risk difference:

NNT = 1 / risk difference

NNT gives us the number of people we need to treat to prevent one extra person from having the event

(Where we drop any minus signs from the risk difference). NNT describes the number of patients you would need to treat with the experimental treatment rather than the control treatment in order to prevent a single event. In other words, if the risk difference is 0.76, that means if we treat 100 people, 76 more will benefit when we use the intervention, who would not have benefited if given control. So how many would we need to treat to help one person? 100/76 or 1.3.

We always round up NNT to the next whole number so in this case we need to treat two women with antibiotics to cure one additional woman (over and above those who would have been cured anyway, i.e. those cured in the control group). It is important with NNT to link it to a time frame, so in this case we would need to treat 2 women with antibiotics *for 6 weeks* to prevent a single *extra* woman from not being cured.

Where the risk difference is greater than 0 (i.e. the risk of the event we are trying to prevent actually increases) the same calculation produces a number known as the NNH – number need to harm. This is the number of participants treated for a length of time for one extra person to have the event.

While NNTs are easy to interpret, making them popular with consumers and clinicians, they cannot be used for performing a meta-analysis because of their mathematical properties. RR, OR and RD are therefore used for meta-analysis, and all may later be converted to NNTs as a way of communicating results in some Cochrane reviews. In later modules, we'll look in more depth at interpreting and applying the results of analyses.

Summary to date.

Here is a reminder of the statistics we have covered so far in this module:

- The **risk** describes the number of participants having the event in a group divided by the total number of participants
- The **odds** describe the number of participants having the event divided by the number of participants not having the event
- The **risk ratio (relative risk)** describes the risk of the event in the intervention group divided by the risk of the event in the control group
- The **odds ratio** describes the odds of the event in the intervention group divided by the odds of the event in the control group
- The **risk difference** describes the absolute change in risk that is attributable to the experimental intervention
- The **number needed to treat** (NNT) gives the number of people you would have to treat with the experimental intervention (compared with the control) to prevent one event.

95% CI tell us how certain we are of the result

Uncertainty

All of these statistics are based on observations in a sample of participants who are randomly split into treatment and control groups. On average randomisation will generate two groups who would have the same event rates if treated identically – so that any observed difference in outcome must be due to the different effects of the treatment and control interventions. However, this comparability is not guaranteed in any particular trial. It is possible that, by chance, the treated group may have a few more people who would naturally do well or badly than the control group, even if they had all received identical treatment.

This means that the observed treatment effect (OR, RR, RD) may actually be an over- or underestimate of the real effect of treatment. A confidence interval (CI) can be calculated as a way of representing the uncertainty in the estimate of treatment effect. The interval contains a range of values above and below the calculated treatment effect within which we can be reasonably certain (usually specified as 95% certain) that the real effect lies. The result is said to be statistically significant if the 95% CI does not include the risk in the two groups being the same (i.e. 1 for risk ratio or odds ratio, 0 for risk difference).

Another way of thinking about CIs is that it gives us an estimate of the range in which the estimate would fall a fixed percentage of times if we repeated the study many times. Picking a 95% CI means that in 5% of all possible trials the effect estimate would fall outside the 95% CI (2.5% above and 2.5% below). In some situations, you may want your CI to include more of the possible trial results, to be more sure that you are quoting an interval that contains the real effect. You can choose, for example, a 99% CI. The interval you come out with then will be wider than for a 95% CI, making your interpretation more conservative.

Putting these statistics in words

To make all these numbers useful to decision makers, we have to be able to express them in words. Using the example from the previous section, here are some suggestions of how to express your results.

The RR of 'still being infected' on antibiotics relative to no antibiotics was 0.12. We can express this as

- The risk of still being infected on antibiotics was about 12% of the risk on control
- Treatment reduced the risk to 12% of what it would have been
- Treatment reduced the risk by 88% of what it was in the control group

We could reasonably exchange the word 'risk' for 'chance' or 'probability', as they are commonly used to mean the same thing.

Odds ratios are harder. The OR in this example is 0.02. You could express this as:

- Antibiotics reduced the odds of still being infected to about 2% of what they would have been
- Treatment reduced the odds by 98% of what they were in the control group

Note that we have to use the word odds and must not use words like chance, risk or probability.

The risk difference in this example was –0.76. This is best expressed as:

- Antibiotics reduced the risk of still being infected by 76 percentage points.

It is important to be clear about how you express the reduction. If we said the reduction was 76%, it is difficult to know if this reduction is 76% of the risk without treatment (i.e. the control group risk), which in this case was 0.86, or it is a reduction of 76 percentage points. In most cases RD should be expressed as percentage points, as these are the units we need to calculate number needed to treat.

Most people find risk ratio and risk difference (and NNT) easier to interpret than odds ratios. However, communication of the results of your review is only one factor to consider when choosing the best statistic to use.

Choosing an effect measure

From what we have seen to date we know that if a randomised trial measures dichotomous outcomes, we can compare the event rates in the two groups using several different summary statistics:

- A risk ratio (of either the good and bad outcome)
- The risk difference
- The NNT
- The odds ratio.

Unfortunately there is no easy way to decide which statistic to use in your review. You will need to make two decisions,

- Which statistic to choose for the analysis
- Which statistic to choose to present the results.

Within RevMan, the choice you make for expressing the results of an individual study will also apply to the meta-analysis if you choose to do one. Whatever statistic you choose for your analysis, you always have the option of re-interpreting the results using another measure in the text of your review. For example, you might perform a meta-analysis by selecting risk ratio in RevMan, then interpret the results by converting it to an NNT in the results section of your review. Remember that NNT is useful for presenting results, but not for analysis purposes. The other three statistics (OR, RR, RD) can be used for either. But also remember that we have two quite different RRs depending on which way the events are coded.

There are three principal issues to consider when choosing a summary statistic:

- Communication, i.e. a straightforward and clinically useful interpretation
- Consistency of the statistic across different studies
- Reasonable mathematical properties.

Communication
We have seen in the section on putting statistics into words that it can be quite hard to explain an odds ratio. Most clinicians and consumers have less difficulty with understanding a risk ratio, however we need to be careful to also give some idea of the absolute difference between the two groups, as relative measures can be misleading. Take the example of buying two lottery tickets instead of one. We could say you are doubling your chances of winning, or we could say your chances of winning have gone up by 1 in 400,000. Both versions give you incomplete information because neither tells us clearly what the chance of winning is in the first place. The statements are likely to be interpreted differently, because many people would think an increase of 1 in 400,000 sounds a lot less attractive than a doubling of the chance of winning.

NNT is a very useful way to express effect in clinical terms (for example, “I have to treat x number of my patients with this treatment for y number of weeks in order to help 1 patient who would not have got better anyway”), as is risk difference (for example, “for every 100 treated, x% will benefit”).

Consistency
The results of your review are probably drawn from many trials, and will be applied in many populations, so it is desirable that the statistic you chose is consistent, i.e. stays nearly the same, or is stable, when applied in different places. In any meta-analysis there is likely to be variation in event rates between trials. The risk ratio, risk difference and odds ratio all vary to some extent in different situations, but one may be more stable than the others. Let’s have a look at a hypothetical example.

Two hypothetical trials with varying event rates with consistent OR, RRx2 and RD

Trial	Relation to Trial 1	Control	Treatment	OR	Treatment group RR(E)	RD	RR(NE)
1		24/100	16/100	0.60	0.67	-0.08	1.11
2A	Same OR	42/100	30/100	0.60	0.71	-0.12	1.21
2B	Same RR(E)	42/100	28/100	0.54	0.67	-0.14	1.24
2C	Same RD	42/100	34/100	0.71	0.81	-0.08	1.14
2D	Same RR(NE)	42/100	36/100	0.78	0.86	-0.06	1.11

In this table we take trial 1 as the reference trial. In trial 2 the events are more common – 36% of all patients have events compared to 20% in trial 1. Despite this, in trial 2A, the effect size is the same as trial 1 if we used *odds ratio* as the summary statistic. But if we were to use risk ratio (either of the event or the non-event) or if we used risk difference, the effect measure would not be the same. Now consider trial 2B. Here the effect is the same in the 2 trials if we choose *risk ratio of the event* as the effect measure, but it varies when the other measures are chosen. Trial 2C and 2D show the same phenomenon, with risk difference and risk ratio of the non-event respectively being consistent.

So, as event rates vary, OR, RR and RD may vary to different degrees. When we are choosing which one to use, it would be helpful to choose the one which is most likely to be consistent across the event rates in the studies we have, as it is also the one most likely to be consistent in clinical practice. To try to help make this decision, some researchers looked at many meta-analyses in Cochrane reviews, calculating the results with OR, RR and RD. Of these, RD varied most across trials included in the meta-analyses, and OR and RR less so. This suggests that the relative measures are more likely to be consistent in Cochrane reviews than RD. The RR of the bad event varied less than the RR of the good event.

Mathematical Properties

In order for a summary statistic to be able to be used in a meta-analysis it needs one mathematical property. That is the ability to reliably estimate its variance. This is because the way in which we assign weight to studies within the meta-analysis is inversely proportional to variance (we will cover this in more detail in the next module). We cannot use NNT in a meta-analysis because we don't have a usable estimate of its variance.

There are two other properties, which are not essential but are mathematically desirable. We've already seen that it is easier to switch between odds ratios of 'good' and 'bad' outcomes, than it is with risk ratios. This is sometimes argued to be a helpful mathematical property of odds ratios.

Another issue that arises when applying results is called 'bounding', as we can get predictions outside the bounds of possibility. For example, we calculated earlier on that the risk ratio of cure of UTI for antibiotic use in pregnant women was 6.6. What would happen if we tried to apply this to a group of women where we thought half would get better without antibiotics?

Risk without treatment × risk ratio = risk with treatment

$0.5 \times 6.6 = 3.3$

This result is nonsense as it predicts that 330% (i.e. more than all) of the women will be cured!

A similar thing can happen with risk difference. The risk difference we calculated for risk of not being cured of UTI with antibiotics was –0.76. Let's try to apply that in a situation where we think that, without antibiotics, 30% of people would not be cured.

Risk without treatment + risk difference = risk with treatment

$0.3 + -0.76 = -0.46$

Again this is nonsense as it predicts that using antibiotics will mean that –46% of women will not be cured.

In practice, this is less important than it may appear, as what we are doing in these examples is applying a result from one situation to a *very* different situation. In reality, we would not expect to apply results from very high risk populations to very low risk populations or vice versa.

Summary

In summary, risk ratio and odds ratio are better for meta-analysis than the risk difference. Risk ratios are easier to understand than odds ratios, but require some care in choosing whether to analyse the risk ratio of the event or non-event.

This overall view is summarized in the following table:

	OR	RR	RD
Communication	-	+	++
Consistency	+	+	-
Mathematics	++	-	-

Once the analysis is performed the results require careful interpretation. Odds ratios must not be interpreted as if they were risk ratios unless events are very rare. It will be helpful to a reader if any relative effects are re-expressed as absolute effects (RD or NNT) – maybe working out these figures for a range of possible scenarios.

Some review groups will prefer you to use a particular measure to give uniformity across their reviews, and you should check with them or on *The Cochrane Library* to see if this is the case. If, after reading this module, you don't agree with their policy, you could always challenge them to justify it.

In the next module we will discuss how results from individual trials are combined to give an overall estimate of treatment effect.