Week 8: The Rise and Fall of HRT
Solutions to Exercises

1. Each of the following examples of a hazard ratio (HR) with confidence interval and P-value from a trial contains a misprint. Explain why we can be sure of this in each case.

   (i) HR 1.52 (95% CI 1.61 to 2.29), P=0.045

The lower limit of the confidence interval (1.61) is higher than the estimate itself (1.52). The true estimated hazard ratio should fall somewhere within the confidence interval.

   (ii) HR 0.61 (95% CI -0.43 to 0.87), P=0.006

The lower limit of the confidence interval (-0.43) is negative. This is impossible: the hazard ratio is a ratio of incidence rates and is always positive.

   (iii) HR 1.14 (95% CI 0.58 to 2.24), P=1.703

The P-value is greater than 1. This is impossible: the P-value is a probability, and is never greater than 1.

   (iv) HR 0.57 (95% CI 0.48 to 0.68), P=0.150

The 95% confidence interval excludes a hazard ratio of 1, which implies a statistically significant result (P<0.05), but here the P-value is greater than 0.05.

2. Suppose you are recruiting for a trial in which subjects are allocated to groups “A” or “B” using block randomisation, not stratified, with a *fixed but unknown* block size. Randomisations are provided to you by an online randomisation service. The first 23 subjects are allocated as follows:

   ABBBAABBAABAAAABBBAABBB

   Can you figure out which group the 24th subject will be allocated to?

This example illustrates how allocation concealment can be compromised when using randomised blocks. In this case, you can guess that the block size is 6: each block of 6 contains exactly 3 As and 3 Bs; no other block size will work in this respect. Having figured this out, you would know that the 24th allocation must be A, to make the last block balance. To get around this problem, trials sometimes use blocks of varying length, where the length of the block is determined randomly.

In questions 3 and 4, decide for each part of the question whether the statement is true or false, and give your reasoning.

3. In a randomised controlled trial to compare a new analgesic with placebo for the control of pain in arthritis, subjects reported less pain while using the analgesic and the difference in pain scores between the two regimes was highly statistically significant (P=0.002).

   We can conclude that:

   a) an important clinical advance has been made.

FALSE – just because the result is statistically significant does not mean the result is clinically important

   b) there is good evidence that the drug reduces pain.

TRUE – there is a low P-value, therefore good evidence for some sort of effect of the drug

   c) the drug is a very effective analgesic.

FALSE – the effect of the drug might be very small even though it is statistically significant

   d) the difference between mean pain scores on the two regimes was 0.002.

FALSE – this is not the interpretation of the P-value

4. In another randomised controlled trial to compare a new analgesic with ibuprofen (the standard treatment) for the control of pain in arthritis, the difference in pain scores between the two regimes was not statistically significant (P>0.05).

   We can conclude that:

   a) the new drug is useless.

FALSE – we haven't found evidence that the drug is effective, but that doesn't prove it is ineffective – the confidence interval will tell us how big the effect of the drug might plausibly be

   b) the trial has failed to demonstrate a difference in analgesia.

TRUE – we have not found evidence that the drug is effective

   c) the difference between the drugs is very small.

FALSE – we don't know from this information how big the effect of the drug might be – we need to see the confidence interval

   d) there are no important differences in the analgesic properties of the drugs.

FALSE – again, the confidence interval might show that clinically important differences are plausible, even though it is also showing that a difference of zero is plausible.