

Chapter 25 Summary

Paired Samples and Blocks

What have we learned?

- Pairing can be a very effective strategy.
 - Because pairing can help control variability between individual subjects, paired methods are usually more powerful than methods that compare individual groups.
- Analyzing data from matched pairs requires different inference procedures.
 - Paired t -methods look at pairwise differences.
 - We test hypotheses and generate confidence intervals based on these differences.
 - We learned to *Think* about the design of the study that collected the data before we proceed with inference.

Paired Data

- Data are paired when the observations are collected in pairs or the observations in one group are naturally related to observations in the other group.
- Paired data arise in a number of ways. Perhaps the most common is to compare subjects with themselves before and after a treatment.
 - When pairs arise from an experiment, the pairing is a type of *blocking*.
 - When they arise from an observational study, it is a form of *matching*.
- If you know the data are paired, you can (and must!) take advantage of it.
 - To decide if the data are paired, consider how they were collected and what they mean (check the W 's).
 - There is no test to determine whether the data are paired.
- Once we know the data are paired, we can examine the *pairwise* differences.
 - Because it is the *differences* we care about, we treat them as if they were the data and ignore the original two sets of data.
- Now that we have only one set of data to consider, we can return to the simple one-sample t -test.
- Mechanically, a paired t -test is just a one-sample t -test for the mean of the pairwise differences.
 - The sample size is the number of pairs.

Assumptions and Conditions

- Paired Data Assumption: The data must be paired.
- Independence Assumption: The differences must be independent of each other. Check the:
 - Randomization Condition
- Normal Population Assumption: We need to assume that the population of *differences* follows a Normal model.
 - Nearly Normal Condition: Check this with a histogram or Normal probability plot of the differences.

The Paired t -Test

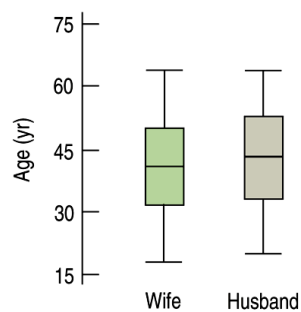
- When the conditions are met, we are ready to test whether the paired differences differ significantly from zero.
- We test the hypothesis $H_0: \mu_d = \Delta_0$, where the d 's are the pairwise differences and Δ_0 is almost always 0.
- We use the statistic $t_{n-1} = \frac{\bar{d} - \Delta_0}{SE(\bar{d})}$ where n is the number of pairs.
- $SE(\bar{d}) = \frac{s_d}{\sqrt{n}}$ is the ordinary standard error for the mean applied to the differences.
- When the conditions are met and the null hypothesis is true, this statistic follows a Student's t -model on $n - 1$ degrees of freedom, so we can use that model to obtain a P-value.

Confidence Intervals for Matched Pairs

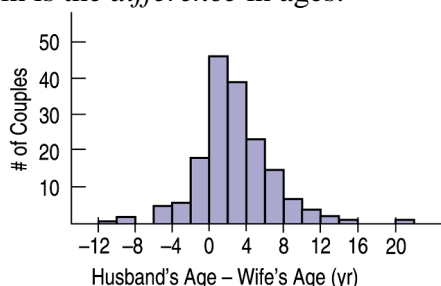
- When the conditions are met, we are ready to find the confidence interval for the mean of the paired differences.
- The confidence interval is $\bar{d} \pm t_{n-1}^* \times SE(\bar{d})$ where the standard error of the mean difference is $SE(\bar{d}) = \frac{s_d}{\sqrt{n}}$
- The critical value t^* depends on the particular confidence level, C , that you specify and on the degrees of freedom, $n - 1$, which is based on the number of pairs, n .

Blocking

- Consider estimating the mean difference in age between husbands and wives.
- The following display is worthless. It does no good to compare all the wives as a group with all the husbands—we care about the paired differences.



- In this case, we have paired data—each husband is paired with his respective wife. The display we are interested in is the *difference* in ages:



- Pairing removes the extra variation that we saw in the side-by-side boxplots and allows us to concentrate on the variation associated with the difference in age for each pair.
- A paired design is an example of blocking.

*The Sign Test Again?

- Because we have paired data, we've been using a simple t -test for the paired differences.
- This suggests that if we want a distribution-free method, a sign test on the paired differences testing whether the *median* of the differences is 0 is appropriate.
- The advantage of the sign test for matched pairs is that we don't require the Nearly Normal Condition for the paired differences.
 - However, if the assumptions of a *paired t-test* are met, the paired t -test is more powerful than the sign test.

What Can Go Wrong?

- Don't use a two-sample t -test for paired data.
- Don't use a paired- t method when the samples aren't paired.
- Don't forget outliers—the outliers we care about now are in the differences.
- Don't look for the difference in side-by-side boxplots.