Introductory Statistics Tutorial
Chapter 10 – Data on a Continuous Variable

Section A: Paired Comparisons or Two Independent Samples

1. For the experiments described below, state the method of analysis that should be used (choose **either** analysed as paired data **or** analysed as two independent samples).

(a) A consumer group is interested in comparing the performance of two different brands of tyres with respect to tyre wear. One of each brand of tyre was randomly assigned to the rear wheels of eight cars of the same model. These cars were then driven a specified distance by the same driver on the same road. The amount of wear on each tyre was recorded.

(b) The research department of a large concrete products manufacturer is experimenting with a new method of manufacturing concrete blocks. To test whether or not the new method increased the compressive strength of the blocks, ten sample blocks were made by the new method and ten sample blocks were made by the old method. The compressive strength of each block was measured.

(c) A sample of ten students enrolled in a course in German were asked to copy a passage written in German and the number of errors made was recorded. The instructor used a new approach in teaching over the next two weeks. At the end of this time, the same ten students were asked to copy the same passage with the number of errors made being recorded again.

2. Home Shopping Network, Inc., pioneered the idea of merchandising directly to customers through cable TV in America. By watching commercials that run 24 hours per day, viewers can call a number to buy products. Before expanding their services, network managers wanted to know whether this method of direct marketing changed sales, on average. A random sample of 16 viewers was selected for an experiment. All viewers in the sample had recorded how much money they spent shopping during the holiday season of the previous year. The next year these people were given free access to the cable network and were asked to keep a record of their total purchases during the holiday season. The data for the 16 shoppers is shown in the table below.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Previous	334	150	520	95	212	30	1055	300	85	129	40	440	610	208	880	25
Current	405	125	540	100	200	30	1200	265	90	206	18	489	590	310	995	75
Current – Previous	71	-25	20	5	-12	0	145	-35	5	77	-22	49	-20	102	115	50

Summary statistics

	Sample mean	Sample standard deviation
Previous	319.6	309.6
Current	352.4	344.0
Difference (Current – Previous)	32.8	55.75

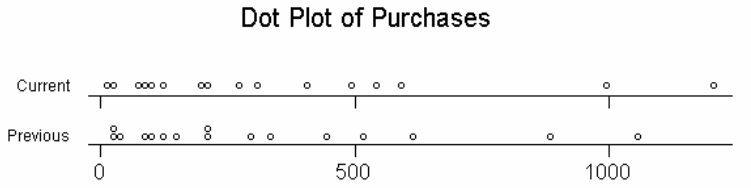

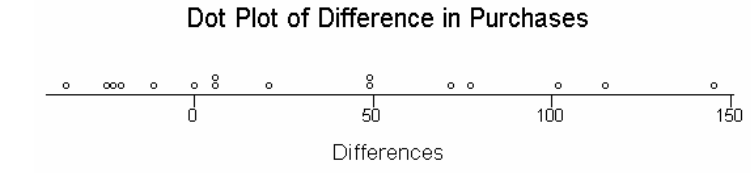

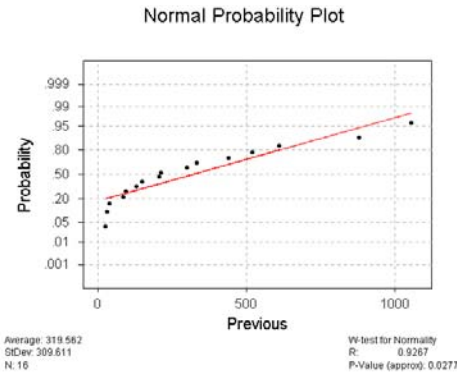

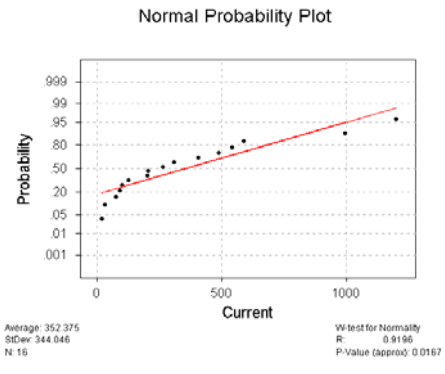

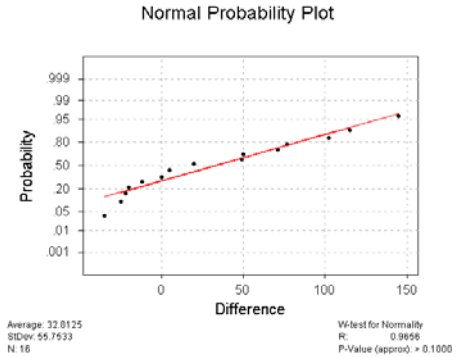

(a) Which test is more appropriate: a *t*-test on the differences or a two independent samples *t*-test? Briefly explain.

(b) Below is MINITAB output for two different *t*-tests. State the hypotheses of the more appropriate of these two tests. Use the appropriate output to interpret your test. Remember to interpret the test and the confidence interval.

Two Sample T-Test and Confidence Interval

```
Two sample T for Current vs Previous

           N      Mean     StDev   SE Mean
Current   16       352       344        86
Previous  16       320       310        77

95% CI for mu Current - mu Previous: ( -204,  269)
T-Test mu Current = mu Previous (vs not =): T = 0.28  P = 0.78  DF = 29
```

Paired T-Test and Confidence Interval

```
Paired T for Current - Previous

                 N      Mean     StDev   SE Mean
Current         16     352.4     344.0      86.0
Previous        16     319.6     309.6      77.4
Difference      16      32.8      55.8      13.9

95% CI for mean difference: (3.1, 62.5)
T-Test of mean difference = 0 (vs not = 0): T-Value = 2.35  P-Value = 0.033
```

(c) Comment on the appropriateness of using a *t*-test. Briefly explain.

(d) Carry out a Sign test on the differences. Use the MINITAB output to help find the *P-value*.

Hypotheses:

Signs of the differences:

***P-value*:**

Interpretation:

MINITAB output: **Cumulative Distribution Function**

```
Binomial with n = 15 and p = 0.500000

         x    P( X <= x )
      5.00         0.1509
     10.00         0.9408

Binomial with n = 16 and p = 0.500000

         x    P( X <= x )
      5.00         0.1051
     10.00         0.8949
```

(e) Which test is more appropriate: the parametric *t*-test on the differences or the nonparametric Sign test on the differences? Briefly explain.

3. What is useful about a nonparametric test?

4. The manufacturers of a rechargeable hedge trimmer claimed that the trimmer would operate for an average of 1.8 hours before a recharge was required. After receiving customers' complaints saying that their trimmers needed to be recharged much sooner than the claimed 1.8 hours of operation, it was decided to check this claim. The number of hours of operation before a recharge was required was recorded for eleven different hedge trimmers as follows:

1.5, 2.2, 0.4, 1.7, 2.0, 1.6, 1.8, 1.5, 2.0, 1.3, 1.7.

Summary statistics:

Sample mean = 1.609 Sample median = 1.7 Sample standard deviation = 0.478

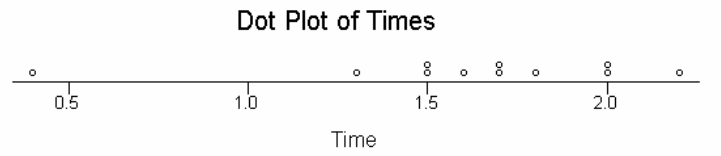

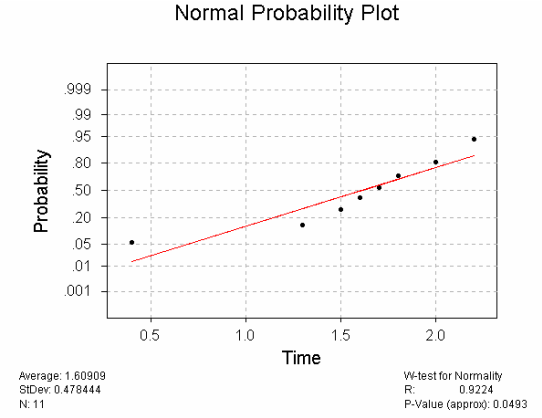

Sign Test for Median

```
Sign test of median = 1.800 versus  <  1.800

              N  Below  Equal  Above       P     Median
Time         11      7      1      3  0.1719      1.700
```

Sign Confidence Interval

```
Sign confidence interval for median

                          Achieved
              N   Median  Confidence    Confidence interval   Position
Time         11    1.700      0.9346  (    1.500,    2.000)         3
                              0.9500  (    1.484,    2.000)       NLI
                              0.9883  (    1.300,    2.000)         2
```


(a) Explain why a one-tailed test is appropriate here.

(b) Explain why the Sign test is more appropriate than a one-sample *t*-test.

(c) State the hypotheses used in the Sign test shown in the MINITAB output above.

(d) What is the *P-value* of this Sign test?

(e) Interpret the Sign test result.

(f) Explain how the *P-value* was obtained by making specific reference to a distribution.

Section B: More Than Two Independent Samples

1. (a) When do we consider using an *F*-test?

(b) What are the hypotheses for a four independent samples *F*-test?

(c) List the assumptions for an *F*-test and describe how we check each of them.

(d) What formula is given in the formulae sheet for calculating the value of the *F*-test statistic, f_0?

(e) What are the formulae for df_1 and df_2?

(f) What are s_B^2 and s_W^2 called and what do they measure?

2. In a recent test of the effectiveness of a new sleeping pill, 75 patients were randomly assigned to three groups of 25. The first group was given the new drug, the second group was given a placebo and the third group was given no treatment at all. The number of minutes it took for the patient to fall asleep was measured and the results of the experiment are shown in the table below.

Drug	12	11	18	37	12	10	63	19	8	17	5	27	50
	27	4	22	28	34	42	2	32	21	33	41	29	
Placebo	44	30	3	13	56	28	24	4	67	32	22	12	27
	32	22	9	13	28	12	42	54	37	22	20	42	
Neither	32	12	55	44	38	89	6	56	59	33	15	67	60
	49	63	8	28	21	14	72	36	66	23	9	39	

Summary statistics

Treatment	Sample size	Sample mean	Sample standard deviation
Drug	25	24.16	15.22
Placebo	25	27.80	16.40
Neither	25	39.76	23.14

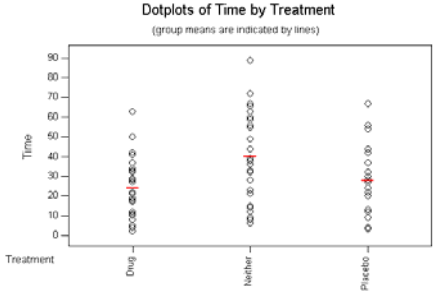

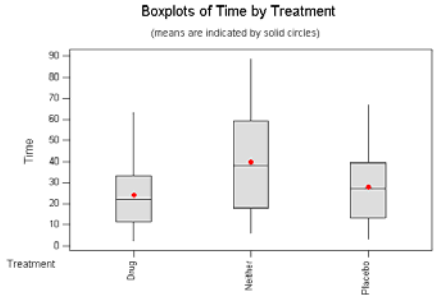


```
Tukey's pairwise comparisons

    Family error rate = 0.0500
Individual error rate = 0.0195

Critical value = 3.38

Intervals for (column level mean) - (row level mean)

           Drug        Neither

 Neither      -28.16
               -3.04

 Placebo      -16.20       -0.60
                8.92       24.52
```


(a) Comment on what the box plots and dot plots reveal about the data.

(b) Comment on the validity of the assumptions for the *F*-test.

(c) State the hypotheses for the *F*-test.

(d) Complete the ANOVA table below:

	DF	SS	MS	F	P
Treatment	_____	3330	1665	_____	0.011
Error	_____	24868	345		
Total	_____	28198			

(e) Interpret the test.

(f) Use an interval to interpret the difference between the underlying mean times to fall asleep for people taking the placebo and people taking the drug.

(g) **(i)** Are we able to determine a single level of treatment that has the longest underlying mean time to fall asleep? If so, name the level of treatment.

(ii) Are we able to determine a single level of treatment that has the shortest underlying mean time to fall asleep? If so, name the level of treatment.

3. **(a)** What should you always do first to any dataset?

(b) Why?

4. When carrying out hypothesis tests (both parametric and nonparametric) we use sample data. What assumption is made about this sample data?

5. From the list below, select the assumptions required for the stated test.

(i) Underlying distribution(s) is (are) Normal.
(ii) Independence between samples.
(iii) Independence between pairs of observations.
(iv) Independence between observations.
(v) Equal population standard deviations.
(vi) Equal sample sizes.
(vii) Equal sample means.

(a) One sample *t*-test.

(b) One sample Sign test.

(c) One sample *t*-test on the differences (from paired-comparison data).

(d) Two sample *t*-test.

(e) *F*-test.

Section C: Identifying Appropriate Type of Analysis

For each of the situations in Questions 1 and 2, the aim is for you to correctly identify the appropriate form of analysis you should use for investigating the problem posed. Below is a table of the types of analyses you have studied.

Code	Form of analysis
A	One-sample *t*-test on a mean
B	One-sample *t*-test on a proportion
C	One-sample *t*-test on a mean of differences
D	Two-sample *t*-test on a difference between two means
E	Two-sample *t*-test on a difference between two proportions
F	*F*-test for one-way analysis of variance

1. The following data set is from a random sample of recent First Semester Stage 2 Statistics students (STATS 20x). Below is some information on the variables collected.

Variable	Type
Grade	The student's final grade for the course (A, B, C, D)
Pass	Whether the student passed the course (Yes, No)
Exam	The student's mark in the final exam (out of 100)
Degree	The degree the student is enrolled for (BA, BCom, BSc, Other)
Gender	The student's gender (Female, Male)
Attend	Whether the student regularly attended class (Yes, No)
Assign	The student's assignment mark (out of 20)
Test	The student's test mark (out of 20)
Diff	The difference in a student's assignment and test mark (Assign – Test)
Stage1	The student's Stage 1 grade (A, B, C)
Repeat	Whether the student was repeating the paper (Yes, No)

The following five different scenarios were of interest to the teaching staff in the Statistics department.

Scenario 1: *Does the exam mark depend on class attendance*?

(i) Write down the name of the variable(s) under investigation and its type (e.g. qualitative or quantitative).

(ii) Can the scenario be viewed as investigating a relationship? HINT: Use the chart at the end of Lecture 3/2 in the Lecture Workbook.

(iii) What tool(s) should you use to begin to investigate the scenario?

(iv) Which form of analysis (from the table above) should be used in the investigation of the scenario?

Scenario 2: *Historically the pass rate for STATS 20x is 76%. Is the underlying pass rate for this semester the same as the historical value?*

(i) Write down the name of the variable(s) under investigation and its type.

(ii) Can the scenario be viewed as investigating a relationship?

(iii) What tool(s) should you use to begin to investigate the scenario?

(iv) Which form of analysis (from the table above) should be used in the investigation of the scenario?

Scenario 3: *Do we expect a student's assignment mark to be higher than their test mark*?

(i) Write down the name of the variable(s) under investigation and its type.

(ii) Can the scenario be viewed as investigating a relationship?

(iii) What tool(s) should you use to begin to investigate the scenario?

(iv) Which form of analysis (from the table above) should be used in the investigation of the scenario?

Scenario 4: *Does the exam mark depend on the degree the student is enrolled for*?

(i) Write down the name of the variable(s) under investigation and its type.

(ii) Can the scenario be viewed as investigating a relationship?

(iii) What tool(s) should you use to begin to investigate the scenario?

(iv) Which form of analysis (from the table above) should be used in the investigation of the scenario?

Scenario 5: *Does the pass rate depend on class attendance*?

(i) Write down the name of the variable(s) under investigation and its type.

(ii) Can the scenario be viewed as investigating a relationship?

(iii) What tool(s) should you use to begin to investigate the scenario?

(iv) Which form of analysis (from the table above) should be used in the investigation of the scenario?

2. For each of the situations below correctly identify the appropriate form of analysis.

As a guideline to determining the correct form of analysis to use:

- Identify the response variable (that is, the variable the researcher is interested in).
- Identify the treatment variable(s) you would use to investigate the response variable.
- Identify the type of hypothesis you would be testing.

(a) In the early 1990s, aspirin, well known as a mild painkiller, was recognised to have blood-thinning properties. If a small dosage of aspirin is taken regularly, the risk of a heart attack caused by blood clotting can be reduced substantially. A drug company in the United States is presently interested in testing a new drug, which it believes has similar properties to aspirin. A total of 19,000 people will take part in the study. 9500 subjects were randomly selected and given a regular dosage of aspirin, the remaining subjects were given a regular dosage of the new drug. The experiment used double blinding methods in order to prevent any possible biases by the subjects or the people administering the treatments. At the end of a 10-year period, the number of deaths due to heart attacks was recorded for each group.

(b) An orchardist in Napier is concerned about the presence of a species of (hungry!) bugs that have infested the apple trees in his orchard. There are four different sprays available that will eliminate this pest and the orchardist wants to test the effectiveness of these to see which one, if any, is worth using. There are 80 apple trees that are infested with the bug. The orchardist randomly assigned the trees into five groups with each of the four sprays being randomly allocated to one of the groups, and the remaining group acting as a control and being left untreated. Two weeks after spraying, the orchardist took a random sample of 100 leaves from each of the 80 trees and counted the number of these leaves that have bugs on them.

(c) A market research company has been contracted to carry out a large survey. The survey involves asking a detailed series of questions that usually takes between 15 to 25 minutes to answer. The company is aware that the longer it takes to answer the questionnaire, the less likely people are to complete the survey. They conduct a small pilot study to investigate which of the two questionnaire designs will take the least time to answer. 50 subjects were randomly allocated into two groups of 25; one group is interviewed using the first questionnaire design, and the other group is interviewed using the second questionnaire design. For each subject, the time taken to complete the questionnaire is recorded.

(d) The quality control manager of a pharmaceutical company is investigating a new method of measuring the levels of acetaminophen in the cold remedies produced. The current method used is known to be accurate, but is more expensive than the new method. Therefore, the manager needs to know if there is any evidence that the new method is less accurate than the current method. The company produces a wide range of cold remedies at varying strengths and twenty-four (24) different batches of cold remedies are sampled for testing. The level of acetaminophen is measured in each batch using both the current and new method of measuring.