GROUND RULES:

- Print your name at the top of this page. Do not put your name on any other page.
- This is a closed-book and closed-notes exam.
- This exam contains 5 problems, each worth 12 points. This exam is worth **60 points** total.
- You may use a calculator, but this calculator cannot have internet access. You cannot use your phone as a calculator. You cannot share calculators with another student. Show all of your work; use a calculator only to do final calculations or to check your work.
- Each problem contains parts. On each part, there is opportunity for partial credit, so show all of your work and explain all of your reasoning. **Translation:** No work/no explanation means no credit.
- On any problem, you may use the back of the page if you need more space. I also have extra paper if you need it.
- Any discussion or inappropriate communication between you and another examinee, as well as the appearance of any unnecessary material, will result in a declaration of academic dishonesty. Don't risk it!
- You have 75 minutes to complete this exam.

HONOR PLEDGE FOR THIS EXAM:

After you have finished the exam, please read the following statement and sign your name below it.

> *I promise that I did not discuss any aspect of this exam with anyone other than the instructor, that I neither gave nor received any unauthorized assistance on this exam, and that the work presented herein is entirely my own.*

1. The presence of high arsenic (As) levels in water wells in southeastern New Hampshire has been recognized by the United States Environmental Protection Agency since the late 1970's. Civil engineers recently observed a random sample of $n = 32$ water wells in southeastern New Hampshire and recorded the arsenic concentration (in parts per billion, ppb) for each well. The data are shown below:

Well	1	2	3	4	5	6	7	8	9	10	11	12
As (ppb)	45.3	12.6	25.5	97.3	56.8	31.1	29.6	43.3	180.2	33.6	51.1	82.4
Well	13	14	15	16	17	18	19	20	21	22	23	24
As (ppb)	1.8	55.2	6.3	5.7	40.9	11.2	10.1	96.2	43.6	4.8	44.6	11.9
Well	25	26	27	28	29	30	31	32				
As (ppb)	8.5	112.2	137.8	118.9	44.0	34.9	59.1	4.6				

(a) What is the population in this example? Give a reasonable answer.

(b) The sample mean is 48.2 and the sample standard deviation is 43.6.

```
> mean(arsenic)
[1] 48.2
> sd(arsenic)
[1] 43.6
```

What do each of these quantities estimate? Be specific. What are the units attached to these estimates?

Parts (c) and (d) are on the next two pages.

(c) I used R to construct a normal quantile-quantile (qq) plot of the data on the previous page. This plot is shown below:

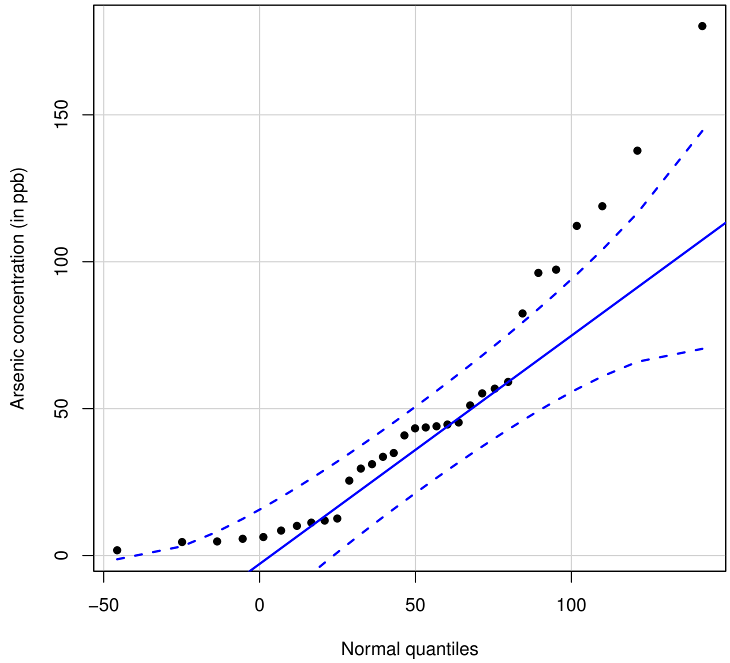

Do you think these data are well represented by a normal population distribution? Explain.

Part (d) is on the next page.

(d) I used R to calculate a 95% confidence interval using the formula

$$\overline{x} \pm t_{n-1,\alpha/2}\frac{s}{\sqrt{n}}.$$

Here is the output:

```
> t.test(arsenic,conf.level=0.95)$conf.int
[1] 32.5 63.9
```

Interpret what this interval means.

2. A manufacturer states their car batteries will last, on average, 4 years with a standard deviation of 1 year. Recent data are consistent with the claim the population mean lifetime is $\mu = 4$ years, but there is concern the claim regarding the population standard deviation $\sigma = 1$ year is not correct.

(a) The manufacturer offers a warranty benefit for all batteries whose lifetimes are less than 2 years. Assuming the values of $\mu = 4$ and $\sigma = 1$ are correct, and assuming the population distribution of the battery lifetime X is normal, that is, $X \sim \mathcal{N}(4, 1^2)$, approximately what percentage of batteries will qualify for a warranty benefit? *Hint:* Draw a picture.

(b) If the population standard deviation σ is larger than 1 year, how would this affect the percentage of batteries in the population which qualify for the warranty benefit? Would it increase the percentage of batteries which qualify or decrease it? Explain.

Parts (c) and (d) are on the next page.

To investigate the issue regarding the population standard deviation, quality control engineers selected a random sample of $n = 120$ batteries from the population, and the lifetime (X, in years) was measured on each battery. A histogram of the data and a normal qq plot are shown below:

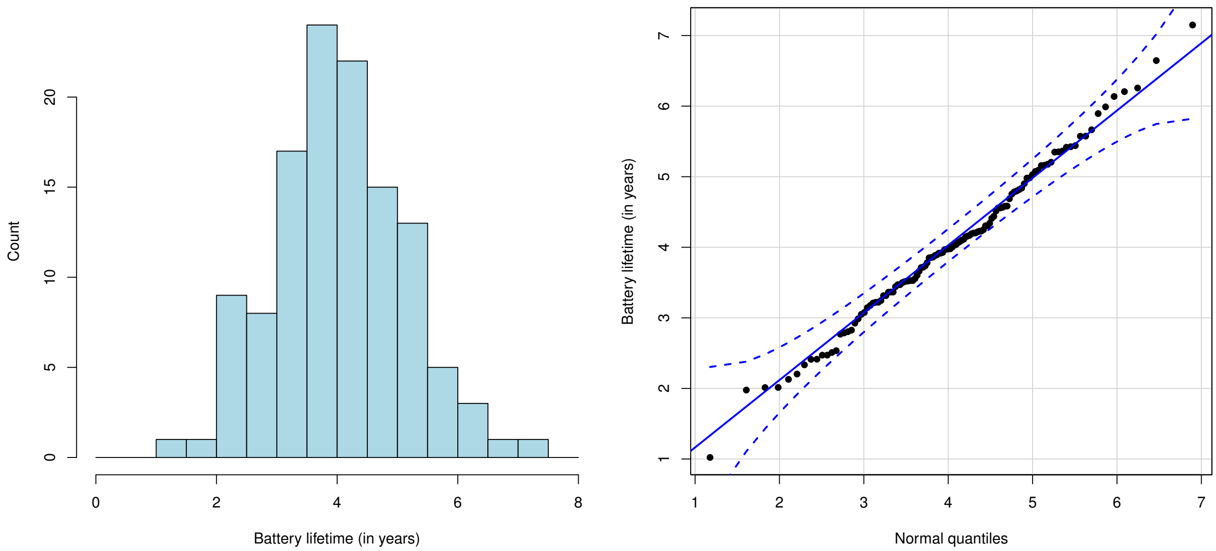

(c) I used R to calculate a 95% confidence interval using the formula

$$\left(\frac{(n-1)s^2}{\chi^2_{n-1,1-\alpha/2}}, \frac{(n-1)s^2}{\chi^2_{n-1,\alpha/2}} \right)$$

Here is the output:

```
> var.ci(batteries,conf.level=0.95) # CI for population variance
[1] 0.93 1.54
```

Is this interval consistent with the manufacturer's claim that the population standard deviation is $\sigma = 1$ year? Or does this interval provide evidence (at the 95% confidence level) the manufacturer's claim is incorrect? Explain.

(d) The confidence interval procedure in part (c) is **not** robust to violations of the normal population distribution assumption. Explain what this means. Use the back of this page if necessary.

3. Nurses who work in hospitals routinely handle substances (e.g., bleach, ammonia, etc.) that could increase the risk of adverse health outcomes, such as asthma. Researchers at Simon Fraser, a university in British Columbia, Canada, recruited two samples of individuals at hospitals throughout the province:

- Group 1: 100 nurses who work in hospitals
- Group 2: 100 office workers who work in hospitals.

After following the 200 individuals for one year, researchers performed lung function tests to assess whether the individuals had developed asthma-related symptoms. None of the 200 individuals had asthma-related symptoms before the study began.

Here were the sample proportions of individuals who developed asthma-related symptoms by the end of the study:

$$\text{Group 1: } \frac{12}{100} = 0.12 \qquad \text{Group 2: } \frac{6}{100} = 0.06.$$

(a) I used the formula

$$(\widehat{p}_1 - \widehat{p}_2) \pm z_{\alpha/2}\sqrt{\frac{\widehat{p}_1(1-\widehat{p}_1)}{n_1} + \frac{\widehat{p}_2(1-\widehat{p}_2)}{n_2}}$$

to write a 90% confidence interval for the population proportion difference $\Delta = p_1 - p_2$, where

p_1 = population proportion of hospital nurses who would develop symptoms
p_2 = population proportion of hospital office workers who would develop symptoms,

and I got $(-0.006, 0.126)$.

```
> prop.test(c(12,6),c(100,100),conf.level=0.90,correct=FALSE)$conf.int
[1] -0.006  0.126
```

(a) Interpret what this interval means. Then discuss what this interval suggests about how the population proportions compare at the 90% confidence level.

Parts (b) and (c) are on the next page.

(b) At what confidence level do you think the interval on the last page would exclude "0" as a plausible value for the population proportion difference Δ? There is no one right answer (but there are wrong ones).

(c) Suppose you want to plan a larger study that assumes a 95% confidence level ($z_{0.05/2} = 1.96$). Assuming $n_1 = n_2 = n$, find the smallest number of individuals n (per group) that will produce a 95% confidence interval for $\Delta = p_1 - p_2$ with margin of error equal to 0.04. You can use the information in the problem to estimate any parameters needed for this calculation.

4. An experiment was conducted to evaluate the effectiveness of a drug treatment for tapeworm in the stomachs of sheep. A random sample of 24 worm-infected lambs of the same age and health was randomly divided into two groups:

- Group 1: 12 lambs treated for tapeworm with the drug
- Group 2: 12 lambs left untreated.

After a six-month period, the lambs were slaughtered and the following worm counts were recorded:

Drug-treated lambs	18	43	28	50	16	32	13	35	38	33	6	7
Untreated lambs	40	54	26	63	21	37	39	23	48	58	28	39

Side-by-side boxplots of the two samples (treated/untreated) are shown below:

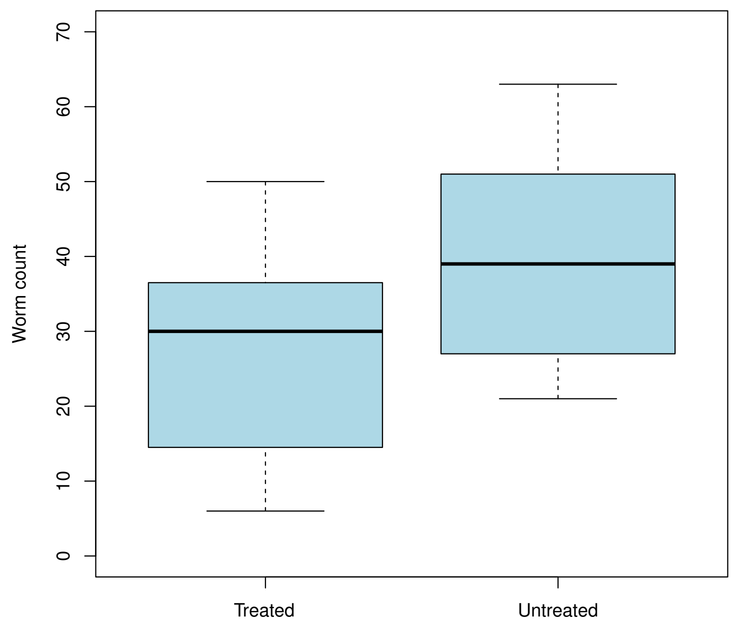

The goal of the experiment is to compare the population mean worm counts for the two groups. I asked R to calculate 95% confidence intervals for the population mean difference $\Delta = \mu_1 - \mu_2$, where μ_1 is the population mean worm count for all treated lambs and μ_2 is the population mean worm count for all untreated lambs:

```
> t.test(treated,untreated,conf.level=0.95,var.equal=TRUE)$conf.int
[1] -25.0  -1.1
> t.test(treated,untreated,conf.level=0.95,var.equal=FALSE)$conf.int
[1] -25.0  -1.1
```

The first interval $(-25.0, -1.1)$ assumes the population variance of the worm counts are equal for the two groups; i.e., $\sigma_1^2 = \sigma_2^2$. The second interval assumes the population variances of the worm counts are not equal. Note that the two intervals are the same.

Parts (a), (b), and (c) are on the next page.

(a) Interpret what the confidence interval means. Then discuss what this interval suggests about how the population mean worm counts compare at the 95% confidence level.

(b) Why do you think the equal-variance and unequal-variance confidence intervals for $\Delta = \mu_1 - \mu_2$ are the same?

(c) The analysis on the previous page assumes we have two independent samples of lambs (i.e., 12 treated; 12 untreated). An investigator asks you if it would have been possible to design this study as a matched-pairs experiment. Do you think this would be possible?

- If so, carefully explain how you would redesign the experiment using the same 24 lambs.
- If not, carefully explain why it is not possible to perform a matched-pairs experiment.

Use the back of this page if necessary.

5. Industrial engineers at a company which stamps gaskets out of sheets of cork, rubber, and plastic want to compare the mean number of gaskets per hour for the three types of material. The engineers sample 10 hourly production numbers at random for the three materials. Here are the data (number of gaskets per hour):

Sample 1	Cork	4395	4464	3843	4105	4270
		3557	3792	4035	4031	4308
Sample 2	Rubber	3710	3831	3629	3454	4072
		4245	4120	3220	3820	3863
Sample 3	Plastic	3831	3807	3652	3665	3752
		3924	3827	3550	3838	3475

Side-by-side boxplots of the three samples (cork/rubber/plastic) are shown below:

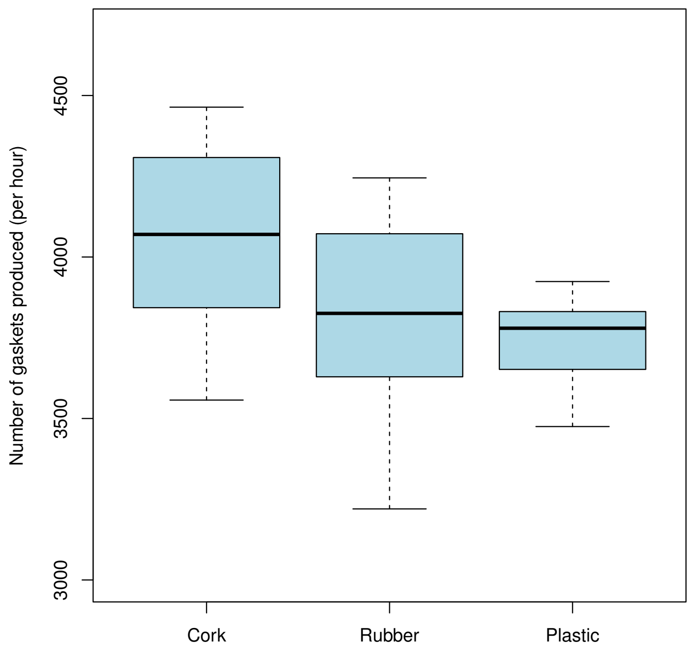

The first objective was to test whether the population mean number of gaskets produced (per hour) was the same for all three materials, that is, testing

$$H_0 : \mu_1 = \mu_2 = \mu_3$$
versus
$$H_1 : \text{the population means } \mu_i \text{ are not all equal.}$$

Parts (a), (b), (c), and (d) are on the next two pages.

(a) Here is the R output to perform the hypothesis test on the last page:

```
> fit = lm(Gaskets ~ Material)
> anova(fit)

Analysis of Variance Table
          Df  Sum Sq  Mean Sq  F value  Pr(>F)
Material   2  685326   342663    5.116   0.013
Residuals 27 1808537    66983
```

The overall F statistic is 5.116 with p-value = 0.013. Sketch a graph of the sampling distribution of F when H_0 is true, identify what the sampling distribution is, and show on your graph how the p-value is determined. I'm looking for three things here.

(b) Here is the R output to do a follow-up Tukey analysis:

```
> TukeyHSD(aov(fit),conf.level=0.95)

Tukey multiple comparisons of means
95% family-wise confidence level

                 diff     lwr    upr  p adj
plastic-cork   -347.9  -634.9  -60.9  0.015
rubber-cork    -283.6  -570.8    3.4  0.053
rubber-plastic   64.3  -222.7  351.3  0.845
```

Interpret these results. Use the back of this page if necessary.

Parts (c) and (d) are on the next page.

(c) Had the engineers not done the overall ANOVA and just focused on groups 2 (rubber) and 1 (cork) to begin with, a 95% confidence interval for $\mu_2 - \mu_1$ based on the two-independent sample and equal population variance assumptions would have been $(-565.9, -1.3)$.

```
> t.test(rubber,cork,conf.level=0.95,var.equal=TRUE)$conf.int
[1] -565.9   -1.3
```

Why is this interval (and its conclusion) different than the interval for $\mu_2 - \mu_1$ in part (b)?

(d) Are you concerned about the underlying statistical assumptions for a one-way classification analysis? If so, which one(s), and why are you concerned? *Hint:* Look at the boxplots again.