This exam covers primarily chapters 4 and 6. We've also talked about ideas from section 8.1, which deals with adapting the Central Limit Theorem to proportions, so that will be covered on the exam as well. Note that we skipped section 6.4.

Probably the most difficult three ideas on the exam are the Central Limit Theorem, confidence intervals, and hypothesis testing. Before giving you some review problems, here's a review of these three important ideas.

Central Limit Theorem

The Central Limit Theorem describes the sampling distribution for a sample statistic: either a sample mean, or a sample proportion. Recall that the sampling distribution of a statistic shows the distribution which would result from taking a large number of samples (all of the same size) and measuring the statistic of interest for each sample. Thus, the sampling distribution for a statistic gives us an idea of what we can expect: what values of a statistic are common, what values are less common, and what values are all but impossible.

Central Limit Theorem (mean version)

Suppose you take a simple random sample of size n from a population in which the mean of some quantitative variable is μ and the standard deviation of the same variable is σ. Suppose also either that n is fairly large (at least 40 is a good guidline), or that the population can be modeled by the normal curve. Then the sampling distribution of the sample mean ($\bar{x}$) is approximately normal, with a mean of μ and a standard deviation of $\frac{\sigma}{\sqrt{n}}$.

You can find this version of the Central Limit Theorem in section 4.3, in particular on pages 242-244.

Central Limit Theorem (proportion version)

Suppose you take a simple random sample of size n from a population in which the proportion of individuals in some category of a categorical variable is p. Suppose also that n is large enough that we can expect any sample to contain at least 10 individuals in the category, as well as at least 10 individuals outside the category. Then the sampling distribution of the sample proportion ($\hat{p}$) is approximately normal, with a mean of p and a standard deviation of $\sqrt{\frac{p\,(1-p)}{n}}$.

You can find this version of the Central Limit Theorem in section 8.1, on pages 431-432.

Confidence Intervals

The basic idea behind a confidence interval is that we wish to be able to estimate the value of a parameter (pertaining to the entire population) from the value of a corresponding statistic (pertaining only to a single sample). We know that simple random samples will vary, so the mean (or proportion) of our sample may not exactly equal the mean (or proportion) of the population. But because the Central Limit Theorem tells us what the sampling distribution should look like, we can calculate which values of the parameter are most likely to yield the statistic we observe in our sample.

Confidence Intervals for means

Suppose you take a simple random sample of size n, from a population whose mean μ is unknown but whose standard deviation σ is known. If the mean of the sample is $\bar{x}$, then we can be 95% confident that μ is between $\bar{x} - 2\left(\frac{\sigma}{\sqrt{n}}\right)$ and $\bar{x} + 2\left(\frac{\sigma}{\sqrt{n}}\right)$.

For other confidence levels C, we can be C% confident that μ is between $\bar{x} - z^*\left(\frac{\sigma}{\sqrt{n}}\right)$ and $\bar{x} + z^*\left(\frac{\sigma}{\sqrt{n}}\right)$, where the number z^* is chosen to be the z-score which puts C% of the area under the normal curve between $-z^*$ and z^*.

Note that constructing this confidence interval requires us to know, or at least estimate, the population standard deviation σ. It's not very realistic that we would know σ but not μ, so eventually (in chapter 7) we'll want to devise a method to get around this assumption. If n is fairly large (at least 100 is a good guideline), then the sample standard deviation s is a reasonable estimate for σ. For now, we'll stick to situations where σ is known, or when n is large enough to estimate σ with s.

The book discusses confidence intervals in section 6.1; you should understand their summary on page 306. It's also a good idea to make sure you understand the graphic on page 302, which illustrates a large number of samples, each used to construct a confidence interval. Most, but not all, of these confidence intervals contain the true population mean.

Confidence Intervals for proportions

Suppose you take a simple random sample of size n, from a population whose proportion p is unknown. If the proportion of the sample is $\hat{p}$, then we can be 95% confident that p is between $\hat{p} - 2\sqrt{\frac{\hat{p}\,(1-\hat{p})}{n}}$ and $\hat{p} + 2\sqrt{\frac{\hat{p}\,(1-\hat{p})}{n}}$.

For other confidence levels C, we can be C% confident that p is between $\hat{p} - z^*\sqrt{\frac{\hat{p}\,(1-\hat{p})}{n}}$ and $\hat{p} + z^*\sqrt{\frac{\hat{p}\,(1-\hat{p})}{n}}$. As above, the number z^* is chosen to be the z-score which puts C% of the area under the normal curve between $-z^*$ and z^*.

Note that in constructing this confidence interval, we're assuming that $\sqrt{\frac{p\,(1-p)}{n}}$ (the standard deviation for the sampling distribution, which we don't know since we don't know p), is approximately $\sqrt{\frac{\hat{p}\,(1-\hat{p})}{n}}$ (which we can calculate from $\hat{p}$). This assumption is valid so long as our sample is large enough that it contains at least 10 individuals in the category and at least 10 outside.

The book discusses confidence intervals for proportions in section 8.1; see the upper part of the box on page 437.

Hypothesis testing

The basic idea behind hypothesis testing, also called tests of significance, is to assess the evidence that a single sample provides about a claim regarding the population.

The procedure for a test of significance is:

1. Describe in words the parameter you're interested in.

2. Formulate a null hypothesis and an alternative hypothesis about the parameter. Remember that the null hypothesis always specifies a single specific value for the parameter; if the context warrents it this specific value amounts to "no change" or "no difference". The alternative hypothesis will either specify that the parameter is larger, specify that it's smaller, (both of which are one-sided), or will specify that it's not equal to the null hypothesis value (a two-sided alternative).

3. Now we ask "what would happen if the null hypothesis were true?" In particular, we want to know how unusual our sample would be if the null hypothesis were true. To answer this, we calculate the test statistic, which is just the z-score that our sample would receive in the sampling distribution, if the null hypothesis were true.

4. Once we've calculated the test statistic, we can find the P-value: the probability of getting a sample at least as extreme as ours, if the null hypothesis were true. This probability will be an area under the normal curve, corresponding either to one "tail" (if the alternative hypothesis is one-sided) or to both "tails" (if the alternative is two-sided).

5. Once we've calculated the P-value, we compare it to the significance level we're using as a criterion to decide whether or not to reject the null hypothesis. If the P-value is small enough (usually this means either less than 5% or less than 1%), then we reject the null hypothesis because the assumption that the null hypothesis is true leads us to conclude that our sample is a very unusual one.

The formula we use to calculate the test statistic for a mean is on page 329, but note that it's just a z-score coming from the normal curve which models the sampling distribution. Don't worry too much about the lower part of the box on page 329, just make sure you understand that the three pictures show the samples which are at least as extreme as the one we have. The corresponding formula and pictures for proportions is on page 437.

I will give you the two versions of the Central Limit Theorem on the exam. Note that you can get formulas for both confidence intervals and tests of significance from these, so I won't put those on the exam.

Sample Problems

1. A college committee consists of six people, of whom four are men and two are women. Two committee members were chosen as officers to carry out most of the committee's administrative work. As it turned out, the two women were the two who were selected. A colleague wonders how likely it would be for the two women to be chosen if the selection had been made randomly.

a. What is the sample space of possible pairs of officers from the committee? (For convenience, suppose the six people are named Alice, Bonnie, Carl, Daniel, Evan, and Francisco.)

b. What is the probability that, if the choice of officers was made randomly, Alice and Bonnie would be the two people chosen?

c. What is the probability that, if the choice of officers was made randomly, that the officers would both be men?

2. Between July 26 and August 8, 1999, 190 Major League Baseball games were played. The difference between the winning team's score and the losing team's score in each game is listed below. (So, for example, in 38 of the games, the winning team scored one more run than the losing team.)

Difference in Score	1	2	3	4	5	6	7	8	9	10	11
Number of Games	38	37	27	23	12	16	13	14	1	5	4

The average of these 190 differences is 3.93, and the standard deviation is 2.66.

a. Find a 95% confidence interval for the average margin of victory in all 1999 Major League Baseball games.

b. How many of the 190 games in our sample have a margin of victory that falls within this confidence interval? Is it close to 95%? Should it be? Explain.

c. Observe that the data in our sample are skewed to the right. Does this fact cast doubt on the validity of the technique you used to calculate a confidence interval? Explain briefly.

3. A company owns and services a fleet of cars for its travelling employees. It finds that the lifetime of the brake pads on these cars varies according to a normal distribution with a mean of 55,000 miles and a standard deviation of 4,500 miles.

a. If the new brand of brake pads has the same lifetime distribution as the previous type, what is the distribution of the sample mean lifetimes for 8 cars? Sketch a graph of the sampling distribution.

b. What is the probability that the brake pads on these 8 cars will have an average lifetime of 51,800 miles or less?

4. A study examines the number of hours that first-year college students spend working at paid jobs. A simple random sample of 1,000 first-year students at public universities is chosen, and the students in the sample are found to work an average of 12.2 hours per week, with a standard deviation of 10.5 hours. A simple random sample of 1,000 first-year students at private colleges and universities has an average of 9.2 hours per week, with a standard deviation of 9.9 hours.

a. Formulate a null hypothesis and an alternative hypothesis suitable for assessing whether these two groups of students work a different number of hours, on average.

b. Conduct a test of significance, and report the P-value.

c. Is there in fact a difference between these populations? Explain carefully what conclusion you can draw from the test of significance.

5. Researchers studying Vietnam veterans found that in a random sample of 2101 male veterans, 777 had been divorced at least once.

a. Use this information to form a 90% confidence interval for the percentage of divorced men among all male Vietnam veterans.

b. How would a 99% confidence interval differ from the 90% interval you just calculated? Would it contain larger values or smaller values? Would it be narrower or wider? You need not calculate the 99% confident interval, just describe how it's different.

c. Would it be legitimate to use this interval as an estimate of the divorce rate among all middle-aged American men? Explain.

6. A researcher studying sleeping habits of college students asks a random sample of 5,000 students how much sleep they got the night before. (The question was asked always on a weekday.) He finds that the students in the sample averaged 6.75 hours of sleep, with a standard deviation of 1.5 hours.

a. The researcher wishes to know whether the average college student sleeps less than 7 hours per weeknight. Formulate appropriate hypotheses, and conduct a test of significance to address this question.

b. Find a 95% confidence interval for the average amount of sleep that all college students get on a typical weeknight.

c. Which of the following conclusions best describes the conclusions you can draw about the average amount of sleep college students get on weeknights? Explain your answer.

- The sample provides strong evidence that college students sleep less than 7 hours per night.
- The sample provides evidence that college students sleep much less than 7 hours per night.
- Neither of the previous two conclusions fits the results above.

7. Suppose that 50 different survey organizations take simple random samples of 400 people age 25 and over from a certain town. Each organization finds the average education level of its sample, and calculates a 95% confidence interval. These confidence intervals are plotted below as lines. (Each organization's confidence interval is plotted at a different height, so that all of them can be seen.)

a. Why do these confidence intervals have different centers?

b. Why do they have different widths?

c. About how many of these confidence intervals would you expect to contain the population mean? Explain.

d. How many of them do in fact contain the population mean of 12 years, marked by the vertical line?