Math 3081 (Probability and Statistics)

Lecture #26 of 27 ∼ August 18th, 2021

The χ^2 Distribution

- The χ^2 Distribution
- Confidence Intervals and Hypothesis Tests for Variance
- The χ^2 Test for Goodness of Fit

This material represents §5.2.1-5.2.3 from the course notes, and problems 14-17 from WeBWorK 7.

The χ^2 Distribution, I

We now move into the second half of the final chapter of the course: discussing the χ^2 distribution and two different χ^2 tests, which allow us to expand our hypothesis tests to testing statements about the variance (and standard deviation) of a distribution.

- All of our hypothesis tests so far have essentially focused on testing statements about the mean of a distribution.
- However, in certain scenarios, some of which we will discuss now, we might also want to test hypotheses about the variance of a distribution.
- If the underlying distribution is normal, or approximately normal, we can use the χ^2 distribution to construct such tests.

The χ^2 Distribution, II

We have previously discussed (at length) methods for constructing confidence intervals for the mean μ of a normally-distrbuted random variable with (known or unknown) standard deviation σ, given a random sample $x_1, \ldots, x_n$ from this normal distribution.

- Our present goal is to apply the same ideas to construct confidence intervals for the variance σ^2 (or equivalently the standard deviation σ) of the normal distribution.
- Of course, the problem is only interesting when we do not already know σ, which is to say, when we are estimating it from the sample.

The χ^2 Distribution, III

As we have discussed at length, the sample variance $S^2 = \frac{1}{n-1}\left[(x_1 - \overline{x})^2 + \cdots + (x_n - \overline{x})^2\right]$ gives an unbiased estimator for σ^2.

- In order to construct confidence intervals for σ^2, it is enough to write down the underlying distribution of the statistic $\dfrac{(n-1)S^2}{\sigma^2} = \left(\dfrac{x_1 - \overline{x}}{\sigma}\right)^2 + \cdots + \left(\dfrac{x_n - \overline{x}}{\sigma}\right)^2.$

This distribution is essentially given by a χ^2 distribution.

Definition

The χ^2 distribution with k degrees of freedom is the continuous random variable Q_k whose probability density function $p_{Q_k}(x) = \dfrac{1}{2^{k/2}\Gamma(k/2)} \cdot x^{(k/2)-1}e^{-x/2}$ for all real numbers $x > 0$.

The χ^2 Distribution, IV

Examples:

- The χ^2 distribution with 1 degree of freedom has probability density function $p_{Q_1}(x) = \frac{1}{\sqrt{2\pi x}} e^{-x/2}$ for $x > 0$.
- The χ^2 distribution with 2 degrees of freedom has probability density function $p_{Q_1}(x) = \frac{1}{2} e^{-x/2}$ for $x > 0$, which is the exponential distribution with parameter $\lambda = 1/2$.
- The χ^2 distribution with 3 degrees of freedom has probability density function $p_{Q_1}(x) = \frac{\sqrt{x}}{\sqrt{2\pi}} e^{-x/2}$ for $x > 0$.

The χ^2 Distribution, V

Here are some plots of χ^2 pdfs:

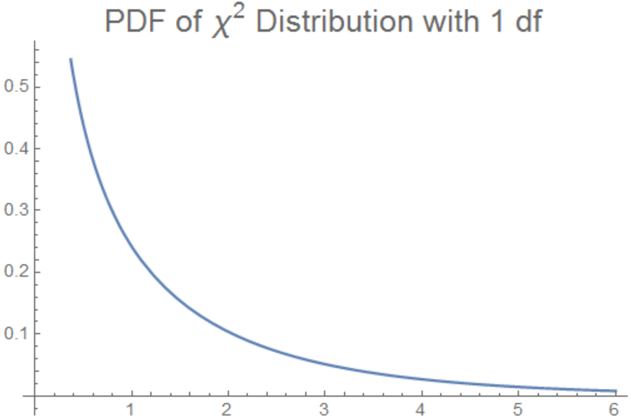

The χ^2 Distribution, VI

Here are some plots of χ^2 pdfs:

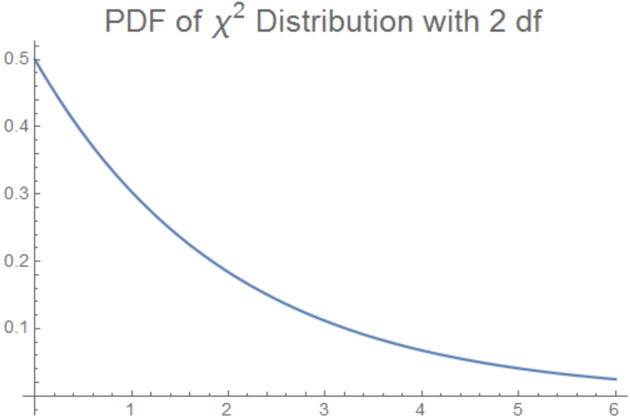

The χ^2 Distribution, VII

Here are some plots of χ^2 pdfs:

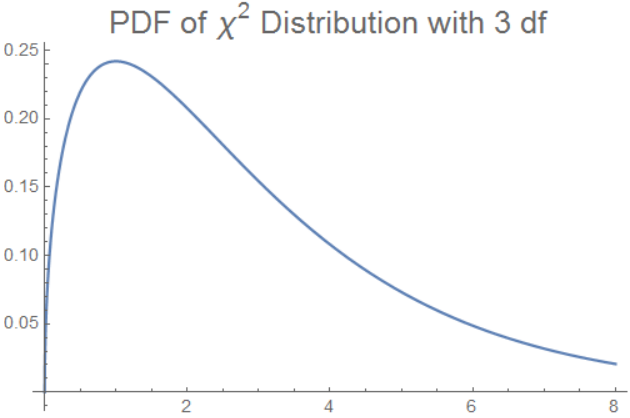

The χ^2 Distribution, VIII

Here are some plots of χ^2 pdfs:

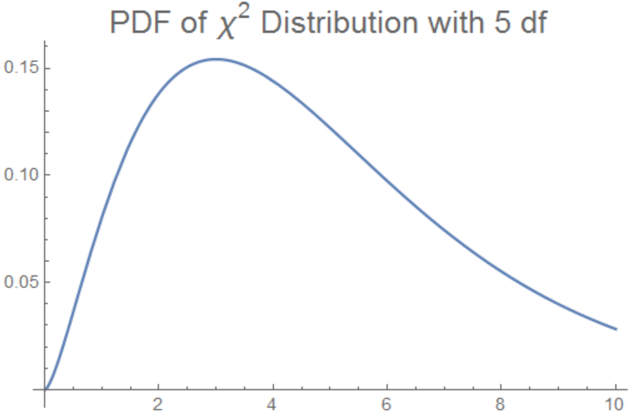

The χ^2 Distribution, IX

Here are some plots of χ^2 pdfs:

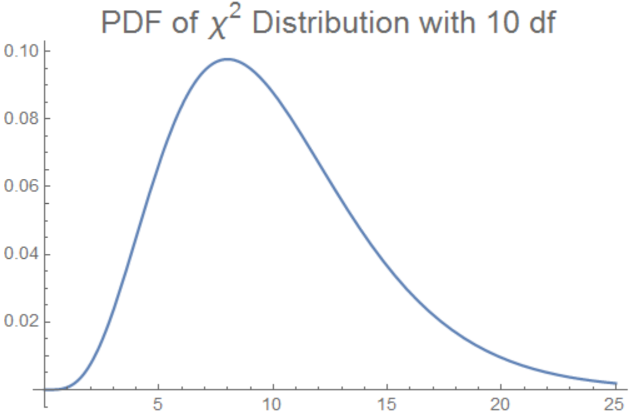

The χ^2 Distribution, X

Here are some plots of χ^2 pdfs:

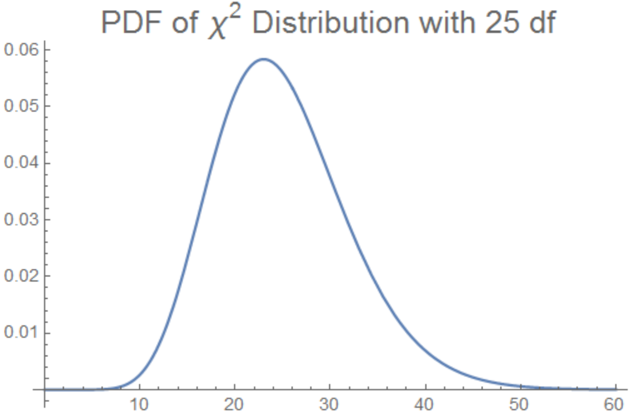

The χ^2 Distribution, XI

Definition

The $\underline{\chi^2 \text{ distribution with } k \text{ degrees of freedom}}$ *is the continuous random variable* Q_k *whose probability density function*
$p_{Q_k}(x) = \frac{1}{2^{k/2}\Gamma(k/2)} \cdot x^{(k/2)-1} e^{-x/2}$ *for all real numbers* $x > 0$.

- It is not hard to show using the probability density function that the χ^2 distribution with k degrees of freedom has mean k and variance $2k$.
- From the plots, we see that the χ^2 distribution, unlike the normal and t distributions, is quite skewed to the right, but the skewness decreases with more degrees of freedom.
- As we now show, the χ^2 distribution with $n-1$ degrees of freedom is the proper model for the test statistic $\frac{(n-1)S^2}{\sigma^2}$.

The χ^2 Distribution, XII

Proposition (χ^2 Distribution From Normals)

If $X_1, \ldots, X_n$ are independent standard normal random variables (i.e., with mean 0 and standard deviation 1), then the random variable $Q_n = X_1^2 + \cdots + X_n^2$ has a χ^2 distribution with n degrees of freedom.

- The proof is a relatively straightforward calculation using the joint pdf of $X_1, \ldots, X_n$ (which is simply the product of the one-variable pdfs, since these variables are independent).
- We then just have to set up and evaluate the appropriate n-dimensional integral to compute the probability density function of $Q_n = X_1^2 + \cdots + X_n^2$.
- The main idea in the computation of the integral is to convert to n-dimensional spherical coordinates.

The χ^2 Distribution, XII

As an easy corollary to the previous result, we see that the χ^2 distribution approaches a normal one as the number of degrees of freedom grows:

Corollary (χ^2 Limit)

As the number of degrees of freedom k increases, the pdf of the χ^2 distribution Q_k approaches the normal distribution with mean k and variance $2k$.

Proof:

- By the proposition, the χ^2 distribution Q_k is obtained by summing k independent, identically-distributed random variables.
- Thus, by the central limit theorem, the normalization $\dfrac{Q_k - k}{\sqrt{2k}}$ approaches the standard normal distribution $N_{0,1}$.

The χ^2 Distribution, XIII

Now we can give our main result about the χ^2 distribution:

Theorem (χ^2 Distribution As Sampling Distribution)

Suppose $n \geq 2$ and that $X_1, X_2, \ldots, X_n$ are independent, identically normally distributed random variables with mean μ and standard deviation σ. If $\overline{X} = \frac{1}{n}(X_1 + \cdots + X_n)$ denotes the sample mean and $S^2 = \frac{1}{n-1}\left[(X_1 - \overline{X})^2 + (X_2 - \overline{X})^2 + \cdots + (X_n - \overline{X})^2\right]$ denotes the sample variance, then the distribution of the test statistic $\frac{(n-1)S^2}{\sigma^2}$ is the χ^2 distribution Q_{n-1} with $n-1$ degrees of freedom.

The χ^2 Distribution, XIV

Proof:

- Let $W = \sum_{i=1}^{n} \left[\frac{X_i - \mu}{\sigma}\right]^2$. Then

$$
\begin{aligned}
W &= \sum_{i=1}^{n} \left[\frac{X_i - \mu}{\sigma}\right]^2 = \sum_{i=1}^{n} \left[\frac{(X_i - \overline{X}) + (\overline{X} - \mu)}{\sigma}\right]^2 \\
&= \sum_{i=1}^{n} \left[\frac{X_i - \overline{X}}{\sigma}\right]^2 + 2\sum_{i=1}^{n} \left[\frac{X_i - \overline{X}}{\sigma}\right]\left[\frac{\overline{X} - \mu}{\sigma}\right] + \sum_{i=1}^{n} \left[\frac{\overline{X} - \mu}{\sigma}\right]^2.
\end{aligned}
$$

- The first term is $\frac{(n-1)S^2}{\sigma^2}$, the middle term is zero by evaluating the sum (since $\sum_{i=1}^{n} X_i = \sum_{i=1}^{n} \overline{X}$), and the last term is $n\left[\frac{\overline{X} - \mu}{\sigma}\right]^2 = \left[\frac{\overline{X} - \mu}{\sigma/\sqrt{n}}\right]^2$.

The χ^2 Distribution, XV

Proof (continued):

- So $W = \sum_{i=1}^{n} \left[\frac{X_i - \mu}{\sigma}\right]^2 = \frac{(n-1)S^2}{\sigma^2} + \left[\frac{\overline{X} - \mu}{\sigma/\sqrt{n}}\right]^2.$
- Note that W is the sum of squares of n independent standard normal variables, so it has a χ^2 distribution with n df.
- Also, S and $\overline{X}$ are independent (as we previously noted in our derivation of the properties of the t distribution).
- Since $\left[\frac{\overline{X}-\mu}{\sigma/\sqrt{n}}\right]^2$ is the square of a standard normal variable, and S is independent from it, this means the distribution of $\frac{(n-1)S^2}{\sigma^2} = W - \left[\frac{\overline{X}-\mu}{\sigma/\sqrt{n}}\right]^2$ is given by the sum of squares of $n - 1$ independent standard normal variables.
- This means $\frac{(n-1)S^2}{\sigma^2}$ has a χ^2 distribution with $n - 1$ degrees of freedom, as claimed.

χ^2 Confidence Intervals, I

The point of the theorem is that we can use the χ^2 distribution as a model for the ratio between the sample variance and the population variance, after rescaling appropriately.

- Thus, we can construct confidence intervals for the population variance using χ^2-statistics and the sample variance.
- Specifically, since the statistic $\dfrac{(n-1)S^2}{\sigma^2}$ is modeled by the χ^2 distribution Q_{n-1} with $n-1$ degrees of freedom, we can compute a $100(1-\alpha)\%$ confidence interval using χ^2-statistics in place of the z- and t-statistics that we used for the confidence intervals for the mean of a normally distributed random variable.

χ^2 Confidence Intervals, II

Here is a plot of the underlying χ^2 distribution:

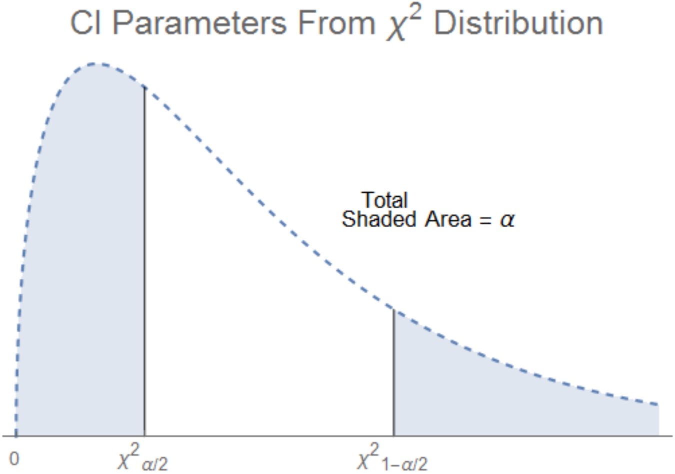

The middle essentially represents the $100(1-\alpha)\%$ CI.

χ^2 Confidence Intervals, III

- We want the parameters $\chi^2_{\alpha/2}$ and $\chi^2_{1-\alpha/2}$ to satisfy $P(Q_{n-1} \leq \chi^2_{\alpha/2}) = \alpha/2 = P(Q_{n-1} \geq \chi^2_{1-\alpha/2})$, so that the total area in each tail of the distribution is α, leaving an area $1-\alpha$ in the middle.
- In other words, we have $P(\chi^2_{\alpha/2} \leq Q_{n-1} \leq \chi^2_{1-\alpha/2}) = 1-\alpha$. Since $\dfrac{(n-1)S^2}{\sigma^2}$ is χ^2-distributed, this is equivalent to saying that $P(\chi^2_{\alpha/2} \leq \dfrac{(n-1)S^2}{\sigma^2} \leq \chi^2_{1-\alpha/2}) = 1-\alpha$.
- We can then rewrite the above equation to get the desired $100(1-\alpha)\%$ confidence interval for σ.

χ^2 Confidence Intervals, IV

Proposition (χ^2 Confidence Intervals)

A $100(1-\alpha)\%$ confidence interval for the unknown variance σ^2 of a normal distribution with unknown mean and standard deviation is given by $\left(\dfrac{(n-1)S^2}{\chi^2_{1-\alpha/2,n}}, \dfrac{(n-1)S^2}{\chi^2_{\alpha/2,n}}\right)$ *where n sample points $x_1, \ldots, x_n$ are taken from the distribution, $\hat{\mu} = \frac{1}{n}(x_1 + \cdots + x_n)$ is the sample mean, $S = \sqrt{\frac{1}{n-1}[(x_1-\hat{\mu})^2 + \cdots + (x_n-\hat{\mu})^2]}$ is the sample standard deviation, and $\chi^2_{\alpha/2,n}$ and $\chi^2_{1-\alpha/2,n}$ are the constants satisfying $P(Q_{n-1} \le \chi^2_{\alpha/2,n-1}) = \alpha/2 = P(Q_{n-1} \ge \chi^2_{1-\alpha/2,n-1})$ where Q_{n-1} is χ^2-distributed with $n-1$ degrees of freedom.*

For σ, take the square root: $\left(\sqrt{\dfrac{n-1}{\chi^2_{1-\alpha/2,n-1}}}S, \sqrt{\dfrac{n-1}{\chi^2_{\alpha/2,n-1}}}S\right)$.

χ^2 Confidence Intervals, V

In order to compute the necessary χ^2 statistics, we must (as with the normal distribution or *t* distribution) either use a table of values or a computer to evaluate the inverse cumulative distribution function.

We need to compute both $\chi^2_{\alpha/2,n}$ and $\chi^2_{1-\alpha/2,n}$, since the χ^2 distribution is not symmetric.

χ^2 Confidence Intervals, VI

Here is a small table of such values:

Inverse-CDF entries give $\chi^2_{\beta,n}$ such that $P(Q_n < \chi^2_{\beta,n}) = \beta$.										
df	0.005	0.01	0.025	0.05	0.1	0.9	0.95	0.975	0.99	0.995
1	0.0000	0.0002	0.0010	0.0039	0.0158	2.7055	3.8415	5.0239	6.6349	7.8794
2	0.0100	0.0201	0.0506	0.1026	0.2107	4.6052	5.9915	7.3778	9.2103	10.597
3	0.0717	0.1148	0.2158	0.3518	0.5844	6.2514	7.8147	9.3484	11.344	12.838
4	0.2070	0.2971	0.4844	0.7107	1.0636	7.7794	9.4877	11.143	13.276	14.860
5	0.4117	0.5543	0.8312	1.1455	1.6103	9.2364	11.071	12.833	15.086	16.750
6	0.6757	0.8721	1.2373	1.6354	2.2041	10.644	12.592	14.449	16.811	18.548
7	0.9893	1.2390	1.6899	2.1673	2.8331	12.017	14.067	16.012	18.475	20.277
8	1.3444	1.6465	2.1797	2.7326	3.4895	13.361	15.507	17.534	20.090	21.955
9	1.7349	2.0879	2.7004	3.3251	4.1682	14.683	16.919	19.022	21.666	23.589
10	2.1559	2.5582	3.2470	3.9403	4.8652	15.987	18.307	20.483	23.209	25.188
15	4.6009	5.2293	6.2621	7.2609	8.5468	22.307	24.996	27.488	30.578	32.801
20	7.4338	8.2604	9.5908	10.851	12.443	28.412	31.410	34.170	37.566	39.997

χ^2 Confidence Intervals, VII

Example: A normal distribution is sampled six times yielding values -3, 1, 5, -2, 7, and 8. Find 80%, 90%, and 99% confidence intervals for the standard deviation of the distribution.

χ^2 Confidence Intervals, VII

Example: A normal distribution is sampled six times yielding values -3, 1, 5, -2, 7, and 8. Find 80%, 90%, and 99% confidence intervals for the standard deviation of the distribution.

- We first compute the sample mean $\mu = 2.6667$ and sample standard deviation $S = 4.6762$.
- Since there are 6 values, the number of degrees of freedom for the underlying χ^2 statistics is 5.

χ^2 Confidence Intervals, VII

Example: A normal distribution is sampled six times yielding values -3, 1, 5, -2, 7, and 8. Find 80%, 90%, and 99% confidence intervals for the standard deviation of the distribution.

- We first compute the sample mean $\mu = 2.6667$ and sample standard deviation $S = 4.6762$.
- Since there are 6 values, the number of degrees of freedom for the underlying χ^2 statistics is 5.
- For the 80% confidence interval, the required values are $\chi^2_{0.9,5} = 9.2364$ and $\chi^2_{0.1,5} = 1.6103$, and so the confidence interval for σ is
$$\left(\sqrt{\frac{5}{9.2364}} \cdot 4.6762, \sqrt{\frac{5}{1.6103}} \cdot 4.6762\right) = (3.4405, 8.2400).$$

χ^2 Confidence Intervals, VIII

Example: A normal distribution is sampled six times yielding values -3, 1, 5, -2, 7, and 8. Find 80%, 90%, and 99% confidence intervals for the standard deviation of the distribution.

- $\mu = 2.6667$, $S = 4.6762$, $df = 5$.

χ^2 Confidence Intervals, VIII

Example: A normal distribution is sampled six times yielding values -3, 1, 5, -2, 7, and 8. Find 80%, 90%, and 99% confidence intervals for the standard deviation of the distribution.

- $\mu = 2.6667$, $S = 4.6762$, $df = 5$.
- For the 90% confidence interval, the required values are $\chi^2_{0.95,5} = 11.0705$ and $\chi^2_{0.05,5} = 1.1455$, and so the confidence interval for σ is $\left(\sqrt{\dfrac{5}{11.0705}} \cdot 4.6762, \sqrt{\dfrac{5}{1.1455}} \cdot 4.6762\right) =$ $(3.1426, 9.7697)$.

χ^2 Confidence Intervals, VIII

<u>Example</u>: A normal distribution is sampled six times yielding values -3, 1, 5, -2, 7, and 8. Find 80%, 90%, and 99% confidence intervals for the standard deviation of the distribution.

- $\mu = 2.6667$, $S = 4.6762$, $df = 5$.
- For the 90% confidence interval, the required values are $\chi^2_{0.95,5} = 11.0705$ and $\chi^2_{0.05,5} = 1.1455$, and so the confidence interval for σ is $\left(\sqrt{\dfrac{5}{11.0705}} \cdot 4.6762, \sqrt{\dfrac{5}{1.1455}} \cdot 4.6762\right) =$ $(3.1426, 9.7697)$.
- For the 99% confidence interval, the required values are $\chi^2_{0.995,5} = 16.7496$ and $\chi^2_{0.005,5} = 0.4117$, and so the confidence interval for σ is $\left(\sqrt{\dfrac{5}{16.7496}} \cdot 4.6762, \sqrt{\dfrac{5}{0.4117}} \cdot 4.6762\right) =$ $(2.5549, 16.2962)$.

χ^2 Hypothesis Tests, I

We can also adapt our characterization to give a procedure for doing a hypothesis test about the unknown variance of a normal distribution based on an independent sampling of the distribution yielding n values $x_1, x_2, \ldots, x_n$: this is the χ^2 test for variance.

- As usual with hypothesis tests, we first select appropriate null and alternative hypotheses and a significance level α.
- Our null hypothesis will be of the form H_0: $\sigma^2 = c$ for some constant c, with an appropriate one-sided or two-sided alternative hypothesis.
- We take the test statistic $\chi^2 = \dfrac{(n-1)S^2}{c}$, where S is the sample standard deviation.
- From our results about the χ^2 distribution, the test statistic is χ^2-distributed with $n-1$ degrees of freedom.

χ^2 Hypothesis Tests, II

If the test is one-sided, we can calculate the p-value based on the alternative hypothesis.

- If the hypotheses are $H_0 : \sigma^2 = c$ and $H_a : \sigma^2 > c$, then the p-value is $P(Q_{n-1} \geq \chi^2)$.
- If the hypotheses are $H_0 : \sigma^2 = c$ and $H_a : \sigma^2 < c$, then the p-value is $P(Q_{n-1} \leq \chi^2)$.
- If the hypotheses are $H_0 : \sigma^2 = c$ and $H_a : \sigma^2 \neq c$, then it is not as obvious how to compute a p-value because of the asymmetry of the χ^2 distribution. We will take the convention of doubling the appropriate one-sided tail probability (as we did with z tests and t tests).

We then compare the p-value to the significance level and then either reject or fail to reject the null hypothesis, as usual.

χ^2 Hypothesis Tests, III

Example: A normal distribution is sampled six times yielding values -3, 1, 5, -2, 7, and 8.

1. Test at the 10% and 1% significance levels that the variance is greater than 16.
2. Test at the 10% and 1% significance levels that the variance is less than 225.

- We calculated the sample standard deviation $S = 4.6762$ earlier, and the number of degrees of freedom is still 5.

χ^2 Hypothesis Tests, IV

Example: A normal distribution is sampled six times yielding values -3, 1, 5, -2, 7, and 8.

1. Test at the 10% and 1% significance levels that the variance is greater than 16.

χ^2 Hypothesis Tests, IV

Example: A normal distribution is sampled six times yielding values -3, 1, 5, -2, 7, and 8.

1. Test at the 10% and 1% significance levels that the variance is greater than 16.

- Our hypotheses are $H_0 : \sigma^2 = 16$ and $H_a : \sigma^2 > 16$, since in fact the sample variance is greater than 16.
- Our test statistic is $\chi^2 = \dfrac{(n-1)S^2}{\sigma^2} = \dfrac{5 \cdot 4.6762^2}{16} = 6.8333$, and so the p-value is $P(Q_5 > 6.8333) = 0.2333$.
- Since the p-value is greater than both significance levels, we fail to reject the null hypothesis in both cases.
- This result is reasonable, since the sample variance is not that much greater than 16. We can also see that $\sigma = 4$ lies well inside the 80% confidence interval we computed earlier.

χ^2 Hypothesis Tests, V

Example: A normal distribution is sampled six times yielding values -3, 1, 5, -2, 7, and 8.

2. Test at the 10% and 1% significance levels that the variance is less than 225.

χ^2 Hypothesis Tests, V

Example: A normal distribution is sampled six times yielding values -3, 1, 5, -2, 7, and 8.

2. Test at the 10% and 1% significance levels that the variance is less than 225.

- Our hypotheses are $H_0 : \sigma^2 = 225$ and $H_a : \sigma^2 < 225$, since in fact the sample standard deviation is less than 225.
- Our test statistic is $\chi^2 = \dfrac{(n-1)S^2}{\sigma^2} = \dfrac{5 \cdot 4.6762^2}{225} = 0.4859$, and so the p-value is $P(Q_5 < 0.4859) = 0.00737$.
- Since the p-value is less than both significance levels, we fail reject the null hypothesis in both cases.
- This result is also reasonable, since the sample variance is quite a bit less than 225. We can also see that $\sigma = 15$ lies well outside the 80% confidence interval we computed earlier, but it is inside the 99% confidence interval (corresponding to the fact that the p-value is greater than 0.005).

χ^2 Hypothesis Tests, VI

Example: An investor wants to determine whether a particular mutual fund is a good investment and also a stable investment, so they measure the yearly rate of return over four different years, obtaining 5.1%, 4.5%, 7.8%, and 4.6% returns.

1. Find a 95% confidence interval for the average return.
2. Test at the 2% significance level whether the average rate of return is above 4%.
3. Find a 95% confidence interval for the standard deviation in the rate of return.
4. Test at the 4% significance level whether the standard deviation in the rate of return is below 2.5%.

χ^2 Hypothesis Tests, VI

Example: An investor wants to determine whether a particular mutual fund is a good investment and also a stable investment, so they measure the yearly rate of return over four different years, obtaining 5.1%, 4.5%, 7.8%, and 4.6% returns.

1. Find a 95% confidence interval for the average return.
2. Test at the 2% significance level whether the average rate of return is above 4%.
3. Find a 95% confidence interval for the standard deviation in the rate of return.
4. Test at the 4% significance level whether the standard deviation in the rate of return is below 2.5%.

- The sample mean is $(5.1\% + 4.5\% + 7.8\% + 4.6\%)/4 = 5.5\%$.
- The sample standard deviation is $\sqrt{\frac{1}{3}[(5.1\% - 5.5\%)^2 + \cdots + (4.6\% - 5.5\%)^2]} = 1.5556\%$.

χ^2 Hypothesis Tests, VII

Example: An investor wants to determine whether a particular mutual fund is a good investment and also a stable investment, so they measure the yearly rate of return over four different years, obtaining 5.1%, 4.5%, 7.8%, and 4.6% returns.

1. Find a 95% confidence interval for the average return.

χ^2 Hypothesis Tests, VII

Example: An investor wants to determine whether a particular mutual fund is a good investment and also a stable investment, so they measure the yearly rate of return over four different years, obtaining 5.1%, 4.5%, 7.8%, and 4.6% returns.

1. Find a 95% confidence interval for the average return.

- We have $\hat{\mu} = 5.5\%$ and $S = 1.5556\%$.
- Since the standard deviation is unknown, we use a t confidence interval. For a 95% CI with $df = 3$, the t-statistic $t_{\alpha/2,df} = 3.1824$.
- Then the confidence interval is $5.5\% \pm 3.1824 \cdot 1.5556\%/\sqrt{4} = (3.02\%, 7.97\%)$.

χ^2 Hypothesis Tests, VIII

Example: An investor wants to determine whether a particular mutual fund is a good investment and also a stable investment, so they measure the yearly rate of return over four different years, obtaining 5.1%, 4.5%, 7.8%, and 4.6% returns.

2. Test at the 2% significance level whether the average rate of return is above 4%.

χ^2 Hypothesis Tests, VIII

<u>Example</u>: An investor wants to determine whether a particular mutual fund is a good investment and also a stable investment, so they measure the yearly rate of return over four different years, obtaining 5.1%, 4.5%, 7.8%, and 4.6% returns.

2. Test at the 2% significance level whether the average rate of return is above 4%.

- We have $\hat{\mu} = 5.5\%$ and $S = 1.5556\%$.
- This is a one-sample t test. Our hypotheses are $H_0 : \mu = 4\%$ and $H_a : \mu > 4\%$, since the average rate in the sample was above 4%.
- The test statistic is $\dfrac{\hat{\mu} - \mu}{S/\sqrt{n}} = \dfrac{5.5\% - 4\%}{1.5556\%/\sqrt{4}} = 1.9285$.
- Thus, the p-value is $P(T_3 \geq 1.9285) = 0.07469$.
- At the 2% significance level, we fail to reject the null hypothesis: we do not have strong enough evidence to conclude the average rate of return exceeds 4%.

χ^2 Hypothesis Tests, IX

Example: An investor wants to determine whether a particular mutual fund is a good investment and also a stable investment, so they measure the yearly rate of return over four different years, obtaining 5.1%, 4.5%, 7.8%, and 4.6% returns.

3. Find a 95% confidence interval for the standard deviation in the rate of return.

χ^2 Hypothesis Tests, IX

Example: An investor wants to determine whether a particular mutual fund is a good investment and also a stable investment, so they measure the yearly rate of return over four different years, obtaining 5.1%, 4.5%, 7.8%, and 4.6% returns.

3. Find a 95% confidence interval for the standard deviation in the rate of return.

- We have $\hat{\mu} = 5.5\%$ and $S = 1.5556\%$.
- Here, we must use a χ^2 confidence interval. For a 95% CI with $df = 3$, the two χ^2-statistics we need are $\chi^2_{\alpha/2,df} = 0.2158$ and $\chi^2_{1-\alpha/2,df} = 9.3484$.
- Then the desired confidence interval is $\left(\sqrt{\dfrac{3}{9.3484}} \cdot 1.5556\%, \sqrt{\dfrac{3}{0.2158}} \cdot 1.5556\%\right) =$ $(0.881\%, 5.800\%)$.

χ^2 Hypothesis Tests, X

Example: An investor wants to determine whether a particular mutual fund is a good investment and also a stable investment, so they measure the yearly rate of return over four different years, obtaining 5.1%, 4.5%, 7.8%, and 4.6% returns.

4. Test at the 4% significance level whether the standard deviation in the rate of return is below 2.5%.

χ^2 Hypothesis Tests, X

Example: An investor wants to determine whether a particular mutual fund is a good investment and also a stable investment, so they measure the yearly rate of return over four different years, obtaining 5.1%, 4.5%, 7.8%, and 4.6% returns.

4. Test at the 4% significance level whether the standard deviation in the rate of return is below 2.5%.

- We have $\hat{\mu} = 5.5\%$ and $S = 1.5556\%$.
- This is a χ^2 test for variance. Our hypotheses are $H_0 : \sigma = 2.5\%$ and $H_a : \sigma < 2.5\%$, since the sample standard deviation was below 2.5%.
- The test statistic is $\chi^2 = \dfrac{(n-1)S^2}{\sigma^2} = \dfrac{3 \cdot 1.5556^2}{2.5^2} = 1.1616$.
- Thus, the p-value is $P(Q_3 < 1.1616) = 0.2378$.
- At the 4% significance level, we fail to reject the null hypothesis: we do not have strong enough evidence to conclude the standard deviation is below 2.5%.

χ^2 Goodness of Fit, I

We often have reasons to believe that sample data should adhere to a particular shape or distribution. However, in many cases, we need to verify whether a particular model actually fits the data set we have collected.

- In situations where we have a single variable of interest, we can often use the hypothesis tests we have already developed to test the reasonableness of a model.
- For example, our z-test for unknown proportion is testing whether a particular Bernoulli random variable is a good model for the observed data set (i.e., the collection of successes and failures observed in a sequence of Bernoulli trials).
- However, most situations have a wider array of data values that we will want to compare to a prediction, and the hypothesis tests we have previously developed are not suitable for that more complicated task.

χ^2 Goodness of Fit, II

For example, we might want to test whether a die is fair by rolling it many times and tabulating the number of times each of the outcomes 1-6 is observed.

- Of course, when we roll the die, we do not expect to get a proportion of precisely $1/6$ for each possible outcome.
- Indeed, the distribution of the number of each roll will be binomially distributed.
- What we want is a way to combine these results into a single test statistic to determine whether all of the results are collectively reasonable or unreasonable.

χ^2 Goodness of Fit, III

The following theorem of Pearson gives a χ^2 test statistic for precisely this type of scenario where values are drawn from a discrete random variable:

Theorem (χ^2 Goodness of Fit)

Suppose that a discrete random variable E has outcomes $e_1, e_2, \ldots, e_k$ with respective probabilities $p_1, p_2, \ldots, p_k$. If we sample this random variable n times, obtaining the respective outcomes $e_1, e_2, \ldots, e_k$ a total of $x_1, x_2, \ldots, x_k$ times, then as $n \to \infty$ the random variable $D = \dfrac{(x_1 - np_1)^2}{np_1} + \dfrac{(x_2 - np_2)^2}{np_2} + \cdots + \dfrac{(x_k - np_k)^2}{np_k}$ is χ^2-distributed with $k - 1$ degrees of freedom.

We will not prove this theorem, as the actual details are quite technical (the proof relies on using moment-generating functions).

χ^2 Goodness of Fit, IV

A few remarks about the ingredients of the theorem:

- Note that each individual total $x_1, x_2, \ldots, x_k$ is binomially distributed (n trials, success probability p_i). The precise joint distribution of all of these totals is called a multinomial distribution.
- Thus, the quantity np_i represents the expected number of times we would expect to see the outcome e_i if we sample the random variable n times.
- As a practical matter, the approximation will be good whenever the expected frequencies np_i are all at least 5 or so.

χ^2 Goodness of Fit, V

Some very brief motivation of where the theorem comes from:

- Since x_i is binomially distributed, that means it is approximately normally distributed with mean np_i and standard deviation $\sqrt{np_i(1-p_i)}$.
- Equivalently, that means $\dfrac{x_i - np_i}{\sqrt{np_i}}$ is approximately normally distributed with mean 0 and standard deviation $\sqrt{1-p_i}$, and so the quantity $(1-p_i)\dfrac{(x_1 - np_1)^2}{np_1}$ is approximately χ^2-distributed with 1 degree of freedom.
- Summing over all of the random variables and noting that $(1-p_1)+(1-p_2)+\cdots+(1-p_n) = n-1$ shows that D is "almost" the sum of $n-1$ χ^2-distributed variables each with 1 degree of freedom, which is equivalent to saying that it is a χ^2-distributed variable with $n-1$ degrees of freedom.

χ^2 Goodness of Fit, VI

Using this theorem, we can give a hypothesis testing procedure for analyzing the goodness of fit of a model:

- We take our test statistic as

$$\begin{aligned} d &= \frac{(x_1 - np_1)^2}{np_1} + \frac{(x_2 - np_2)^2}{np_2} + \cdots + \frac{(x_k - np_k)^2}{np_k} \\ &= \sum_{\text{data}} \frac{[\text{Observed} - \text{Expected}]^2}{\text{Expected}}. \end{aligned}$$

- Our hypotheses are usually $H_0 : d = 0$ and $H_a : d > 0$, since the value $d = 0$ means the model is perfect and a positive value of d indicates deviation from the model.

χ^2 Goodness of Fit, VII

In order to apply Pearson's result above, we must verify that most of the predicted observation sizes np_i are at least 5.

- We will adopt the convention that at least 80% of the entries should be at least 5 or larger. (Another option is to combine some of these small entries into groups that have a predicted size greater than 5.)
- If the hypotheses are satisfied, then the test statistic is χ^2-distributed with $k-1$ degrees of freedom, and we can calculate the p-value as $P(Q_{k-1} \geq d)$.

We then compare the p-value to the significance level and then either reject or fail to reject the null hypothesis, as usual.

χ^2 Goodness of Fit, VIII

Example: To test for fairness, a six-sided die is rolled 2000 times, yielding the results below. Test at the 10%, 3%, and 0.4% significance levels whether the die is fair.

Outcome	1	2	3	4	5	6
Observed	354	347	318	312	333	336

χ^2 Goodness of Fit, VIII

Example: To test for fairness, a six-sided die is rolled 2000 times, yielding the results below. Test at the 10%, 3%, and 0.4% significance levels whether the die is fair.

Outcome	1	2	3	4	5	6
Observed	354	347	318	312	333	336

- If the die is fair, we would expect each outcome to occur with probability $1/6$, meaning that the expected totals are $2000/6 = 333.\overline{3}$ for each of the six possibilities.
- We can tabulate the test statistic more conveniently by adding a few additional rows to the table.

χ^2 Goodness of Fit, IX

Example: To test for fairness, a six-sided die is rolled 2000 times, yielding the results below. Test at the 10%, 3%, and 0.4% significance levels whether the die is fair.

Outcome	1	2	3	4	5	6
Observed	354	347	318	312	333	336
Expected	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$
$(O-E)^2/E$	$1.281\overline{3}$	$0.560\overline{3}$	$0.705\overline{3}$	$1.365\overline{3}$	$0.000\overline{3}$	$0.021\overline{3}$

χ^2 Goodness of Fit, IX

Example: To test for fairness, a six-sided die is rolled 2000 times, yielding the results below. Test at the 10%, 3%, and 0.4% significance levels whether the die is fair.

Outcome	1	2	3	4	5	6
Observed	354	347	318	312	333	336
Expected	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$
$(O-E)^2/E$	$1.281\overline{3}$	$0.560\overline{3}$	$0.705\overline{3}$	$1.365\overline{3}$	$0.000\overline{3}$	$0.021\overline{3}$

- Then the test statistic is just given by summing the bottom row: it comes out as $d = 3.934$.
- There are 6 outcomes hence $6 - 1 = 5$ degrees of freedom.
- Thus, the p-value is $P(Q_5 \geq 3.934) = 0.5590$. Since this is well above each of our significance levels, we fail to reject the null hypothesis in each case.
- Remark: The values were obtained by actually simulating a fair die roll, so it is not surprising that the p-value is large!

χ^2 Goodness of Fit, IX

Our test is set up so that data perfectly fitting the model are not rejected, only data that are far away from the prediction.

- However, in some situations, we may instead want to test whether a model is "too good to believe" (e.g., if we are investigating whether it is reasonable to think that the data have been falsified or altered to adhere too closely to a model).
- In those situations we would instead want the hypotheses to be $H_0 : d = c$ and $H_a : d < c$ for (an arbitrary) positive c, and we would compute the p-value instead as $P(Q_{k-1} \leq d)$.

χ^2 Goodness of Fit, X

Example: To test for fairness, a six-sided die is rolled 2000 times, yielding the results below.

Outcome	1	2	3	4	5	6
Observed	332	334	333	334	333	334

1. Test at the 0.1% significance level against the alternative hypothesis $H_a : \chi^2 > c$ and interpret the result of the test.
2. Test at the 0.1% significance level against the alternative hypothesis $H_a : \chi^2 < c$ and interpret the result of the test.

χ^2 Goodness of Fit, X

Example: To test for fairness, a six-sided die is rolled 2000 times, yielding the results below.

Outcome	1	2	3	4	5	6
Observed	332	334	333	334	333	334

1. Test at the 0.1% significance level against the alternative hypothesis $H_a : \chi^2 > c$ and interpret the result of the test.
2. Test at the 0.1% significance level against the alternative hypothesis $H_a : \chi^2 < c$ and interpret the result of the test.

- As before, the expected entry in each cell is $333.\overline{3}$, so we just add rows to the table like before.

χ^2 Goodness of Fit, XI

Example: To test for fairness, a six-sided die is rolled 2000 times, yielding the results below.

Outcome	1	2	3	4	5	6
Observed	332	334	333	334	333	334
Expected	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$
$(O-E)^2/E$	$0.005\overline{3}$	$0.001\overline{3}$	$0.000\overline{3}$	$0.001\overline{3}$	$0.000\overline{3}$	$0.001\overline{3}$

1. Test at the 0.1% significance level against the alternative hypothesis $H_a : \chi^2 > c$ and interpret the result of the test.

χ^2 Goodness of Fit, XI

Example: To test for fairness, a six-sided die is rolled 2000 times, yielding the results below.

Outcome	1	2	3	4	5	6
Observed	332	334	333	334	333	334
Expected	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$
$(O-E)^2/E$	$0.005\overline{3}$	$0.001\overline{3}$	$0.000\overline{3}$	$0.001\overline{3}$	$0.000\overline{3}$	$0.001\overline{3}$

1. Test at the 0.1% significance level against the alternative hypothesis $H_a : \chi^2 > c$ and interpret the result of the test.

- The test statistic is just given by summing the bottom row: it comes out as $d = 0.01$, with $df = 5$.
- With $H_a : d > c$, the p-value is $P(Q_5 \geq 0.01) = 0.9999995$, so we fail to reject the null hypothesis.
- Our interpretation of this result is that we have basically no evidence suggesting the data don't fit the model.

χ^2 Goodness of Fit, XII

Example: To test for fairness, a six-sided die is rolled 2000 times, yielding the results below.

Outcome	1	2	3	4	5	6
Observed	332	334	333	334	333	334
Expected	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$
$(O-E)^2/E$	$0.005\overline{3}$	$0.001\overline{3}$	$0.000\overline{3}$	$0.001\overline{3}$	$0.000\overline{3}$	$0.001\overline{3}$

2. Test at the 0.1% significance level against the alternative hypothesis $H_a : \chi^2 < c$ and interpret the result of the test.

χ^2 Goodness of Fit, XII

Example: To test for fairness, a six-sided die is rolled 2000 times, yielding the results below.

Outcome	1	2	3	4	5	6
Observed	332	334	333	334	333	334
Expected	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$	$333.\overline{3}$
$(O-E)^2/E$	$0.005\overline{3}$	$0.001\overline{3}$	$0.000\overline{3}$	$0.001\overline{3}$	$0.000\overline{3}$	$0.001\overline{3}$

2. Test at the 0.1% significance level against the alternative hypothesis $H_a : \chi^2 < c$ and interpret the result of the test.

- The test statistic is $d = 0.01$, with $df = 5$.
- With $H_a : d < c$, the p-value is $P(Q_5 \leq 0.01) = 0.0000005$, so we reject the null hypothesis.
- Our interpretation of this result is that we have extremely strong evidence that the data fit the model too well, suggesting that something very suspicious is occurring.

Summary

We introduced the χ^2 distribution and characterized its properties as a sampling distribution.

We discussed confidence intervals and hypothesis tests for the population variance and standard deviation using χ^2-statistics.

We introduced the χ^2 test for goodness-of-fit.

Next lecture: The χ^2 tests for goodness-of-fit and independence