Confidence Intervals with the χ^2_n random variable

↳ pronounced "chi-squared with n degrees of freedom"

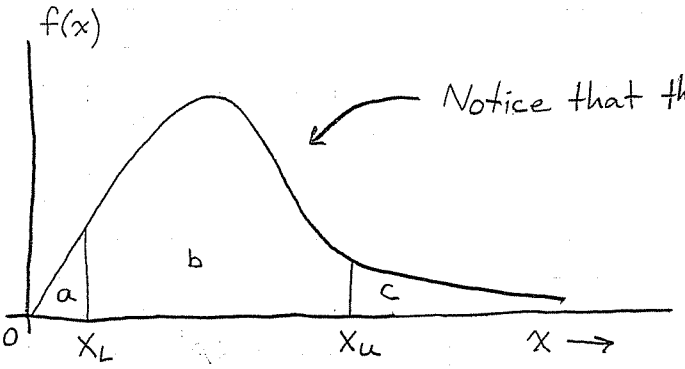

Suppose we want to find the 95% confidence interval. This means that 95% of the area under the curve is found in region b (see above), which leaves 5% of the area under the curve for regions a + c. The problem is that the values of X_L and X_u (the lower and upper boundaries of the confidence interval) are <u>not unique</u>. For example, we could choose region a to have an area of 0.01 and region c to have an area of 0.04. For a χ^2_4 random variable (4 degrees of freedom) this gives

$X_L = 0.297109$ and $X_u = 10.0255$ (see p. 293)

According to Dr. Rowland, the customary thing to do is set the areas of regions a and c equal to each other. For a 95% confidence interval, a and c each have an area of 0.025. For χ^2_4 this gives

$X_L = 0.48442$ and $X_u = 11.1433$ (see again p. 293)

Computing probabilities with the χ^2_n random variable

To find the area under the curve for the χ^2_n PDF you have at least the following options:

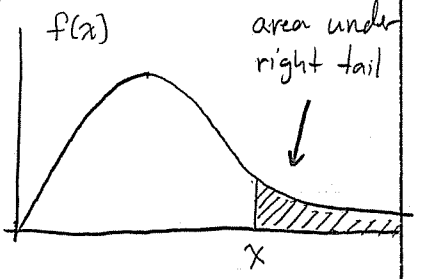

① The table on p. 506, but it is not very detailed. It gives the "right tail" probability (see right)

② Your TI calculator. You need to specify the degrees of freedom, "df", a lower limit, X_L, and an upper limit, X_u.

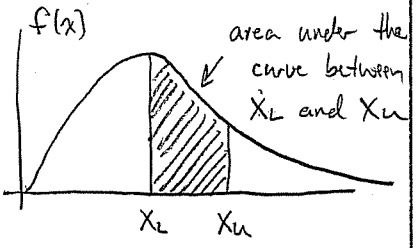

③ Microsoft Excel. There is a function called CHIDIST(X, df) It gives the area under the right tail for the PDF with df degrees of freedom.

④ HP Calculators. For those of us that aren't whimps and know what RPN means, there is a function that computes the area under the right tail also.

⑤ Matlab. Of course, Matlab makes you feel like such a dummy, and would never have a user-friendly function like CHIDIST(). Instead you have to assemble your own CDF out of obscure mathematical functions:

gammainc(X/2, df/2)/gamma(df/2)

returns the actual CDF value, which is 1 − (the area under the right tail)

Go to Wikipedia.org, search: chi squared for an explanation of gammainc() and gamma()

Now that you know how to compute the area under the χ^2_n PDF, and you know how to handle the confidence interval, you run into the problem of finding χ_L and χ_u of the confidence interval

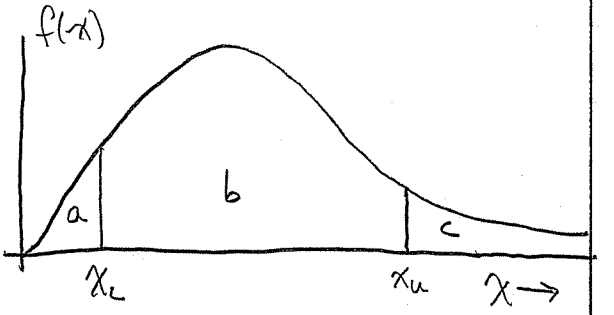

Problem

① Find χ_L such that the area in region a = 0.025

② Find χ_u such that the area in region b = 0.025

Solution

① Read the table on p. 503 "backwards", i.e., find the entries that correspond to 0.025 and 0.975 and read the χ values that correspond to them.

② With your TI calculator, set the lower limit to 0, set "df" to whatever value the problem gives, and adjust the upper limit until the function returns 0.025, you have just found χ_L. To find χ_u, set the upper limit to ∞ and adjust the lower limit (χ_u) until you get an answer of 0.025. If you are a good programmer you can write a little program to do this. Can do this in Matlab and on HP calculator also

③ MY favorite option, MICROSOFT Excel has a function called CHIINV(prob, df)

to find χ_L you use CHIINV(0.975, df)

to find χ_u you use CHIINV(0.025, df)

Easy button option