MS Module 22: χ^2 tests (overview 3rd edition)

(The attached PDF file has better formatting.)

(Readings from the third 3rd edition of the Devore, Berk, and Carlton text.)

Reading §13.1 Goodness-of-Fit Tests, including χ^2 for Completely Specified Probability Distributions," and part of "Goodness-of-Fit Tests for Composite Hypotheses" (until the end of example13.6 on human blood types)

This module covers χ^2 tests, along with related *p* values.

The null hypotheses may specify the probabilities or a relation among the probabilities. If the study examines the probabilities of dominant and recessive alleles (versions of genes), we may test whether the dominant allele has a specified probability or whether the assumed allele structure follows Mendel's laws.

Binomial experiments use the statistical tests for proportions in earlier modules. Multinomial experiments use χ^2 tests, explained in this module.

The textbook uses illustrations from biology and from sports. Actuarial science also offers many illustrations.

! If mortality rates follow mathematical curves, the percentage of deaths by age and sex track the expected.
! If accident frequencies follow mathematical curves, the accidents by age, sex, and territory track the expected.

A χ^2 test helps decide if differences of actual from expected are random fluctuations or if the expected values are incorrect.

Know the test statistic in equation 13.1 and that it has a χ^2 distribution. Chapter 6 of the text describes the χ^2 distribution; if you are not familiar with this topic, review the relevant sections of that chapter. You will be tested on the use of the χ^2 distribution for hypothesis testing, not on the mathematics of the χ^2 distribution.

Know example 13.1. Both this example and the final exam problems assume you are familiar with Mendel's laws of inheritance at a high school level. The experiment crosses AaBb with AaBb.

p_2= dominant alleles for both characteristics: AABB, AaBB, aABB, AABb, AAbB, AaBb, AabB, aABb, aAbB,
p_2= a dominant allele for the first characteristic, and no dominant allele for the second characteristic: AAbb, Aabb, aAbb
p_2= no dominant allele for the first characteristic, and a dominant allele for the second characteristic: aaBB, aaBb, aabB
p_4= no dominant allele for either characteristic: aabb

Of the 16 genotypes, the phenotypic probabilities are 9/16, 3/16, 3/16, 1/16

From the section "χ^2 for Completely Specified Probability Distributions," know example 13.2. Remember that p_{10} means $p_{1\text{ nought}}$, not p_{ten}, and p_{50} means $p_{5\text{ nought}}$, not p_{fifty}. Skip the material beginning "The χ^2 test can also be used to test whether a sample comes from a specific underlying continuous distribution" *until* the sub-section "Goodness-of-Fit Tests for Composite Hypotheses." Know expression 13.3, example 13.5, and example 13.6.

This section covers χ^2 tests when observed probabilities are functions of other parameters. The probabilities may be phenotypes, and the parameter may be the percentage of alleles of a certain type in the population.

An actuary may model mortality for a class as a combination of high mortality and low mortality policyholders or accident frequencies for a class as a combination of high risk and low risk insureds. The relevant question is whether the model is reasonable, not whether the parameters are correctly specified. Example 13.5-6 uses a simple distribution so that the maximum likelihood fitting is easy. Final exam problems use the same type of distribution but with different observed values.

Final exam problems may give the expected values for (k-1) levels. The total probability is one, so you derive the probability of the last layer. As the textbook says: “The fact that df = k – 1 in the preceding theorem is a consequence of the restriction $\Sigma N_i = n$: although there are k observed counts, once any k – 1 are known, the remaining one is uniquely determined. That is, there are only k – 1 “freely determined” cell counts, and thus k – 1 df.”

Pricing actuaries use χ^2 tests to judge whether accident frequency has a Poisson distribution or whether risk classifications are homogeneous. The actuary is not testing a specific Poisson distribution, but only testing whether a Poisson distribution is a reasonable model for accident frequency. In general, Poisson distributions are reasonable only for homogeneous classes, so the actuary is testing whether the class is homogeneous.

The final exam problems uses allele probabilities, for which the arithmetic can be done without a computer. Some examples using Poisson distributions and professional sports playoff games cannot be done by pencil and paper.

Review end of chapter exercises 1, 3, 5, 6, 7, and 13.