## Slides 6: Likelihood ratio tests

- Likelihood ratio
- Nested hypotheses
- Chi-square null distribution
- Chi-squared test of goodness of fit
- A case study

Likelihood ratio test statistic $=-2 \log \frac{\max _{\theta \in \Omega_{0}} L(\theta)}{\max _{\theta \in \Omega} L(\theta)}$
has $\chi_{\mathrm{df}}^{2}$ as an approximate null distribution, with

$$
\mathrm{df}=\operatorname{dim}(\Omega)-\operatorname{dim}\left(\Omega_{0}\right)
$$

Two simple hypotheses
A general method of finding asymptotically optimal tests (having the largest power for a given $\alpha$ ) takes likelihood ratio as the test statistic.

Consider a parametric population distribution with a single parameter $\theta$ and a likelihood function $L(\theta)=L\left(\theta ; x_{1}, \ldots, x_{n}\right)$. For testing

$$
H_{0}: \theta=\theta_{0} \text { against } H_{1}: \theta=\theta_{1},
$$

use the likelihood ratio

$$
\lambda=\frac{L\left(\theta_{0}\right)}{L\left(\theta_{1}\right)}
$$

as a test statistic. Large values of $\lambda$ suggest that $H_{0}$ explains the data set better than $H_{1}$. Therefore, the likelihood ratio test rejects $H_{0}$ for small values of the likelihood ratio.

$$
\text { Likelihood ratio rejection rule is }\left\{\lambda \leq \lambda_{\alpha}\right\}
$$

Neyman-Pearson lemma: the likelihood ratio test is optimal in the case of two simple hypothesis.

Question. How do we find the critical value $\lambda_{\alpha}$ ?

For example, consider $\mathrm{N}(\mu, \sigma)$ model with $\theta=(\mu, \sigma)$. Instead of a pair of two alternative hypotheses $H_{0}: \mu=\mu_{0}$ against $H_{1}: \mu \neq \mu_{0}$, one can think in terms of a pair of nested hypothesis

$$
H_{0}: \mu=\mu_{0}, \quad H: \mu \in(-\infty, \infty)
$$

More generally, consider

$$
H_{0}: \theta \in \Omega_{0}, \quad H: \theta \in \Omega,
$$

where parameter sets $\Omega_{0} \subset \Omega$ are such that $\operatorname{dim}(\Omega)>\operatorname{dim}\left(\Omega_{0}\right)$.
Generalised likelihood ratio

$$
\tilde{\lambda}=\frac{L\left(\hat{\theta}_{0}\right)}{L(\hat{\theta})},
$$

is defined in terms of two maximum likelihood estimates
$\hat{\theta}_{0}=$ maximises the likelihood function $L(\theta)$ over $\theta \in \Omega_{0}$,
$\hat{\theta}=$ maximises the likelihood function $L(\theta)$ over $\theta \in \Omega$.
Question. What is $\mathrm{df}=\operatorname{dim}(\Omega)-\operatorname{dim}\left(\Omega_{0}\right)$ in the example above?

Generalised likelihood ratio test rejects $H_{0}$ for small values of $\tilde{\lambda}$ or equivalently for large values of

$$
-\ln \tilde{\lambda}=\ln L(\hat{\theta})-\ln L\left(\hat{\theta}_{0}\right)
$$

It turns out that the test statistic $-2 \ln \tilde{\Lambda}$ has a nice approximate null distribution

$$
-2 \ln \tilde{\Lambda} \stackrel{H_{0}}{\approx} \chi_{\mathrm{df}}^{2}, \quad \text { where } \mathrm{df}=\operatorname{dim}(\Omega)-\operatorname{dim}\left(\Omega_{0}\right) .
$$

$\chi_{k}^{2}$-distribution is the gamma distribution with $\alpha=\frac{k}{2}, \lambda=\frac{1}{2}$. If independent $Z_{1}, \ldots, Z_{k}$ have the same $\mathrm{N}(0,1)$ distribution, then

$$
Z_{1}^{2}+\ldots+Z_{k}^{2} \sim \chi_{k}^{2}
$$

Question. Consider $\mathrm{N}(\mu, \sigma)$ model with $\theta=(\mu, \sigma)$. With $H_{0}: \mu=\mu_{0}$ against $H_{1}: \mu \neq \mu_{0}$, how would you connect the corresponding likelihood ratio test to the large sample test for the mean?

Chi-squared test of goodness of fit
Suppose that the population distribution is discreet with probabilities $\left(p_{1}, \ldots, p_{J}\right)$. A sample of size $n$ is summarised by the vector of observed counts whose joint distribution is multinomial

$$
\begin{aligned}
& \left(O_{1}, \ldots, O_{J}\right) \sim \operatorname{Mn}\left(n ; p_{1}, \ldots, p_{J}\right) \\
& \mathrm{P}\left(O_{1}=k_{1}, \ldots, O_{J}=k_{J}\right)=\frac{n!}{k_{1}!\cdots k_{J}!} p_{1}^{k_{1}} \cdots p_{J}^{k_{J}}
\end{aligned}
$$

Consider a parametric model for the data

$$
H_{0}:\left(p_{1}, \ldots, p_{J}\right)=\left(v_{1}(\delta), \ldots, v_{J}(\delta)\right)
$$

with unknown $r$-dimensional parameter $\delta=\left(\delta_{1}, \ldots, \delta_{r}\right)$.
To see if the proposed model fits the data, compute $\hat{\delta}$, the maximum likelihood estimate of $\delta$, and then the expected cell counts

$$
E_{j}=n \cdot v_{j}(\hat{\delta})
$$

where "expected" means expected under the null hypothesis model.
Question. What is $\Omega_{0}$ and $\Omega$ in this setting?

In the current setting, the likelihood ratio test statistic $-2 \log \tilde{\lambda}$ is approximated by the so-called chi-squared test statistic

$$
\mathrm{X}^{2}=\sum_{j=1}^{J} \frac{\left(O_{j}-E_{j}\right)^{2}}{E_{j}}
$$

The approximate null distribution of $\mathrm{X}^{2}$ is $\chi_{\mathrm{df}}^{2}$ with $\mathrm{df}=J-1-r$, since

$$
\operatorname{dim}\left(\Omega_{0}\right)=r \quad \text { and } \quad \operatorname{dim}(\Omega)=J-1,
$$

where dim stands for dimension or the number of independent parameters. A mnemonic rule for the number of degrees of freedom:

$$
\mathrm{df}=(\text { number of cells })-1
$$

- (number of independent parameters estimated from the data).

Since the chi-squared test is approximate, all expected counts are recommended to be at least 5 . If not, then you should combine small cells in larger cells and recalculate the number of degrees of freedom df.

A 1889 study made in Germany recorded the numbers of boys $\left(x_{1}, \ldots, x_{n}\right)$ for $n=6115$ families with 12 children each. The general model is described by a vector $\theta=\left(p_{0}, p_{1}, \ldots, p_{12}\right)$ such that

$$
p_{j}=\mathrm{P}(X=j), \quad j=0,1, \ldots, 12 .
$$

We first test a simple null hypothesis claiming that $X \sim \operatorname{Bin}(12,0.5)$, or

$$
H_{0}: p_{j}=\binom{12}{j} \cdot 2^{-12}, \quad j=0,1, \ldots, 12 .
$$

The expected cell counts

$$
E_{j}=6115 \cdot\binom{12}{j} \cdot 2^{-12}, \quad j=0,1, \ldots, 12
$$

are summarised in the table below. The chi-squared test statistic

$$
\mathrm{X}^{2}=\sum_{j=0}^{12} \frac{\left(O_{j}-E_{j}\right)^{2}}{E_{j}}
$$

has the observed value $\mathrm{X}^{2}=249.2$. We have $\mathrm{df}=13-1-0=12$. Since $\chi_{12}^{2}(0.005)=28.3$, we reject $H_{0}$ at $0.5 \%$ level.

| cell $j$ | $O_{j}$ | Model 1: $E_{j}$ | $\operatorname{and} \frac{\left(O_{j}-E_{j}\right)^{2}}{E_{j}}$ | Model 2: $E_{j}$ | $\operatorname{and} \frac{\left(O_{j}-E_{j}\right)^{2}}{E_{j}}$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 7 | 1.5 | 20.2 | 2.3 | 9.6 |
| 1 | 45 | 17.9 | 41.0 | 26.1 | 13.7 |
| 2 | 181 | 98.5 | 69.1 | 132.8 | 17.5 |
| 3 | 478 | 328.4 | 68.1 | 410.0 | 11.3 |
| 4 | 829 | 739.0 | 11.0 | 854.2 | 0.7 |
| 5 | 1112 | 1182.4 | 4.2 | 1265.6 | 18.6 |
| 6 | 1343 | 1379.5 | 1.0 | 1367.3 | 0.4 |
| 7 | 1033 | 1182.4 | 18.9 | 1085.2 | 2.5 |
| 8 | 670 | 739.0 | 6.4 | 628.1 | 2.8 |
| 9 | 286 | 328.4 | 5.5 | 258.5 | 2.9 |
| 10 | 104 | 98.5 | 0.3 | 71.8 | 14.4 |
| 11 | 24 | 17.9 | 2.1 | 12.1 | 11.7 |
| 12 | 3 | 1.5 | 1.5 | 0.9 | 4.9 |
| Total | 6115 | 6115 | $\mathrm{X}^{2}=249.2$ | 6115 | $\mathrm{X}^{2}=110.5$ |

Consider next a more flexible model $X \sim \operatorname{Bin}(12, \delta)$. Model 2 leads to a composite null hypothesis

$$
H_{0}: p_{j}=\binom{12}{j} \cdot \delta^{j}(1-\delta)^{12-j}, \quad j=0, \ldots, 12, \quad 0 \leq \delta \leq 1
$$



Estimate $\delta$ using the maximum likelihood estimate of the proportion of boys in a family

$$
\hat{\delta}=\frac{\text { number of boys }}{\text { number of children }}=\frac{1 \cdot 45+2 \cdot 181+\ldots+12 \cdot 3}{6115 \cdot 12}=0.481
$$

The expected cell counts

$$
E_{j}=6115 \cdot\binom{12}{j} \cdot \hat{\delta}^{j} \cdot(1-\hat{\delta})^{12-j}
$$

are given in the table and the graph above.
The observed chi-squared test statistic for Model 2

$$
X^{2}=110.5
$$

is much smaller than that for Model 1 . However, with $r=1, \mathrm{df}=11$, and the table value $\chi_{11}^{2}(0.005)=26.76$, we reject even Model 2 at $0.5 \%$ level. We see that what is needed is an even more flexible model addressing large variation in the observed cell counts.

Suggestion for Model 3: allow the probability of a male child $\delta$ to differ from family to family. Namely, assume that for each family the value $\delta$ is generated by a beta-distribution $\operatorname{Beta}(a, b)$.

Question. What is dimension $r$ for the suggested Model 3?

