

Stat 414 – Day 4
Likelihood Ratio Tests (LRT), Restricted Maximum Likelihood Estimation (REML)

Last Time: Maximum Likelihood Estimation

- Alternative method for estimating slope coefficients
- Maximizes the (log) likelihood of the observed data

$$L(\mu, \sigma; y_i) = 1/(\sigma\sqrt{2\pi})^n e^{-\sum(y_i-\mu)^2/(2\sigma^2)}$$

$$L(\beta_0, \beta_1, \sigma; y_i, x_i) = -\frac{n}{2} \ln 2\pi - \frac{n}{2} \ln \sigma^2 - \frac{1}{2\sigma^2} \sum_{i=1}^n (Y_i - (\beta_0 + \beta_1 X_i))^2$$

- Regression coefficient estimates same as OLS
- Can compare models by comparing the (log) likelihood values
 - $L = -\frac{n}{2} \ln(2\pi) - \frac{n}{2} \ln \hat{\sigma}^2 - \frac{1}{2} \left(\frac{SSE}{\hat{\sigma}^2}\right)$
 - $AIC = -2 \times \log\text{-likelihood} + 2p = \text{deviance} + 2p$
 - May be reported differently by different software packages
 - Advice: Compute variety of measures, identify models favored by multiple criteria

Example 1: Recall the example predicting airfares from distances.

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	214.99440	69.81197	3.08	0.0116 *
distance	0.14246	0.03384	4.21	0.0018 **

Analysis of Variance Table

		Df	Sum Sq	Mean Sq	F value	Pr(>F)
Residual standard error: 83.46 on 10 degrees of freedom						
distance	1	123463	123463	17.723	0.0018	
Residuals	10	69662	6966			

Multiple R-squared: 0.6393, Adjusted R-squared: 0.6032
 F-statistic: 17.72 on 1 and 10 DF, p-value: 0.0018

(a) What is $\hat{\sigma}_{OLS}$? What are the degrees of freedom for $\hat{\sigma}_{OLS}$? What is the value of SSE? What is the likelihood based on these model estimates? What is the “deviance”?

(b) What is the likelihood of this model compared to the null model? How do we compare these?

Definition: The likelihood ratio test compares nested models by using the statistic $-2(L_0 - L_1)$ which asymptotically follows a chi-square distribution with df = difference in number of parameters in the two model

In R: `1-pchisq(2*(75.144 - 69.023), 1)`

F-tests vs. Likelihood ratio tests for nested models

- Full and Reduced models ($H_0: \beta_j = \dots = \beta_{j+k} = 0$)
- $F = [(\text{drop in SSE})/k] / \text{MSE (full)}$ with df k and $n - p - 1$
- $\chi^2 = 2 \ln(L_{\text{full}} / L_{\text{reduced}})$ with df = k

Example 2: Recall our election data from Day 1. Fit two gls models, one with the interaction terms and one without. Compare the two models with a likelihood ratio test, what do you conclude?

Example 1 cont.

(c) But what about $\hat{\sigma}_{MLE}$?

Maximum Likelihood Estimators for the variance are biased because they do not take into account the estimation of the “nuisance” parameters (β_0, β_1).

An alternative proposed in the 1930s is *restricted* (or *residual*) maximum likelihood, REML.

- Maximizes a different likelihood function (special matrix multiplication, some rows constrained) that doesn't depend on β 's (residuals, removes the fixed effects from the model)
- With simple regression models, the slope coefficient estimates are the same
- The parameter estimate of the variances differ
- $E(\hat{\sigma}_{REML}^2) = \sigma^2$

Definition: *Restricted Maximum Likelihood Estimators (REML)* are MLEs but reflect the number of parameters estimated and are unbiased. REMLs have much better properties especially for estimating variances.

(d) Use restricted maximum likelihood estimation for model 2.

- What is the residual standard error?
- How do the slope and intercept estimates change?
- How does the achieved log likelihood value change?
- How does the AIC value change?

(e) What happens if you try to compare this to the null model?

To compare models with different fixed effects, you must use maximum likelihood estimation. To compare models with different random effects, you can use either ML or REML. REML provides better standard error estimates.

Example 3: Recall the squid data from Day 2/3.

(a) Model 1 assumed $\text{Var}(\epsilon_i) = \sigma^2$. Run model1 using gls.

- Is it using ML or REML?

- How many parameters are being estimated?

(b) Model 2 assumed DML was a variance covariate: $\text{Var}(\epsilon_i) = \sigma^2 / \text{DML}_i$. Run the model. How many parameters are being estimated?

(c) Model 3 assumed the variance covariate was the month, $\text{Var}(\epsilon_{ij}) = \sigma^2_j$ where j was the month number. This actually estimates one month's variance and then all the other month variances are multiples of that. Run model 3, how many parameters are being estimated?

(d) What do you learn from the anova comparing the 3 models?

(e) Run model 4 which allows the power on DML to change differently for each month. $\text{Var}(\epsilon_{ij}) = \sigma^2 / \text{DML}_i^k$. How many parameters? Can we compare model3 to model 4?