

Stat 414 - Day 12

Likelihood Ratio Tests, Restricted Maximum Likelihood

Last Time: Maximum Likelihood Estimation

- An alternative to least squares for fitting the parameter estimates in our regression model that simultaneously estimates slope coefficients (β_s) and variance components (σ) by maximizing the “likelihood” of the observed data.

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

- The MLE estimate of σ^2 is biased - does not account for the other parameters being estimated
- Comparisons of model fit can include penalties for number of parameters in the model
 - e.g., want small AIC = $-2 \times \log\text{-likelihood} + 2p$; BIC = $-2 \times \log\text{-likelihood} + p \times \ln(n)$
- Like a drop in SSEror test, we can compare “nested” models by looking at the change in the log likelihood.
 - $2(L_1 - L_0)$ follows a χ^2 distribution with df = difference in # of parameters in the two models

Example: Recall our data on 768 male squid.

#new trick, just show the first few rows of the "t-table"

```
summary(model1)$coeff[1:5,]
```

```
sum(residuals(model1)^2); sigma(model1) #residual standard error
```

(a) How many parameters are being estimated in this model?

(b) Show how to calculate the residual standard error ($\hat{\sigma}_{LS}$).

(c) Show how to calculate $\hat{\sigma}_{ML}$.

(d) Fit the model using maximum likelihood and verify your answer to (c).

```
library(nlme)
```

```
model1ML <- gls(Testisweight ~ DML * fMONTH, data = Squid, method = "ML")
```

How do the regression coefficients change? What is the achieved log likelihood? How many parameters are being estimated in this model?

Because of the biased nature of the estimate of σ , an alternative proposed in the 1930s is **restricted (residual) maximum likelihood**, REML. Some properties:

- Maximizes a different likelihood function (special matrix multiplication, some rows constrained), that doesn't depend on the slope coefficients (residuals, removes the “fixed” effects from model).
- Is often an “iterative” estimation process (e.g., estimate slope coefficients, then sigma, then slope coefficients, etc.)
- With simple regression models, the estimated slope coefficients are the same, but the parameter estimate of the variances differ.
- $E(\hat{\sigma}_{REML}^2) = \sigma^2$ (so really unbiased for variance not standard deviation)
- Overall, REMLs are better to use for estimating variances

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.

(e) Fit the model with restricted maximum likelihood estimation. Does the residual standard error change? Do the regression coefficients change? Does the “achieved likelihood” change?

The main thing to keep in mind is that if you want to compare models, use the same estimation procedure (e.g., both use REML or both use ML). If the only differences in your model are the “fixed effects” (regression slopes and intercept) but the same random components (e.g., variance assumptions), probably recommend ML, but REML if variance estimation is involved...

The flexibility gained by using maximum likelihood allows us to explore very different variance-covariance structures for our errors. Previously for this dataset, we used weighted least squares to allow the variance in y increase with DML .

Model 2: $Var(\epsilon_i) = \sigma^2 \times DML_i$ for observation i .

```
model2REML = gls(Testisweight ~ DML * fMONTH, data=Squid, weights = varFixed(~DML)
```

(f) Does this model “fit” better? Are the model assumptions sufficiently met?

(g) How do the likelihoods of model 1 and model 2 compare? What does this tell you? Can you carry out a likelihood ratio test?

We also note pretty different variances in y across the different months.

Model 3: weighted regression allowing the variances to differ by month $Var(\epsilon_{ij}) = \sigma_j^2$ for observation i and month j .

```
weights = varIdent(form= ~ 1 | fMONTH), method="REML")
```

(h) How does the likelihood of model 3 compare? What does this tell you? DF? Can you carry out a likelihood ratio test? How are the model assumptions?

Now we can go really crazy, we can let the variances increase with DML , in perhaps a different way (different power) for each month. **Model 4:** $Var(\epsilon_{ij}) = \sigma^2 DML_i^{\delta_j}$

```
vfbymonth <- varPower(form = ~ DML | fMONTH ) #different powers of DML, per month
```

(i) How many parameters are being estimated in this model? What do the parameter estimates at the bottom represent? (*Hint:* Look back at the $Var(\epsilon)$ expression.) Is it worth it? (Has the residual plot improved? Are the additional parameter estimates statistically significant?)

Notes

- The last likelihood ratio test isn't quite appropriate (why?) but not a horrible idea
- The `varPower` used in model 4 should not be used with a quantitative predictor that equal zero.
- Finding the right variance structure for a study like this would be largely trial-and-error, using tools like AIC to compare the models. Better yet, use subject-matter knowledge/past information to help inform your choice of model.
- Notice if we fail to reject a model when compared to the “basic” model that saw we can assume homogeneity in the residuals; this is another (some say better) “test for heterogeneity” (e.g., vs. Bartlett's test, Levene's test). In particular, likelihood ratio tests work when you have violations of normality.