9. Maximum Likelihood Estimation

I. Ordinary Least Squares Estimation:

- For a linear model

$$Y_j = \beta_0 + \beta_1 X_{1j} + \cdots + \beta_r X_{rj} + \epsilon_j,$$

the OLS estimator for

$$\beta = \begin{bmatrix} \beta_0 \\ \vdots \\ \beta_r \end{bmatrix} \quad \text{is any} \quad \mathrm{b} = \begin{bmatrix} b_0 \\ \vdots \\ b_r \end{bmatrix}$$

that minimizes the sum of squared residuals

$$Q(\mathrm{b}) = \sum_{j=1}^{n} (Y_j - b_0 - b_1 X_{1j} - \cdots - b_r X_{rj})^2.$$

- The estimating equations (normal equations) are

$$\frac{\partial Q(\mathrm{b})}{\partial b_0} = -2 \sum_{j=1}^{n} (Y_j - b_0 - b_1 X_{1j} \cdots - \beta_r X_{rj}) = 0$$

and

$$\frac{\partial Q(\mathrm{b})}{\partial b_i} = -2 \sum_{j=1}^{n} X_{ij} (Y_j - b_0 - b_1 X_{1j} \cdots - b_r X_{rj}) = 0$$

for $i = 1, 2, \ldots, r$

The matrix form of these equations is

$$(X^T X)\mathrm{b} = X^T \mathrm{Y}$$

and a solution is

$$\mathrm{b} = (X^T X)^- X^T \mathrm{Y}.$$

The OLS estimator for an estimable function $C^T\beta$ is

$$C^T \mathrm{b} = C^T (X^T X)^- X^T \mathrm{Y}$$

for any solution to the normal equations.

- $E(C^T \mathrm{b}) = C^T \beta$
- $Var(C^T \mathrm{b}) = C^T (X^T X)^- X^T \Sigma X [(X^T X)^-]^T C.$ where $\Sigma = Var(\mathrm{Y})$.
- The distribution of Y is not completely specified.

For a Gauss-Markov model with

$$E(\mathrm{Y}) = X\beta \quad \text{and} \quad Var(\mathrm{Y}) = \sigma^2 I$$

the OLS estimator of an estimable function $C^T\beta$ is the unique best linear unbiased estimator (b.l.u.e.) of $C^T\beta$.

- $E(C^T \mathrm{b}) = C^T \beta$
- $Var(C^T \mathrm{b}) = \sigma^2 C^T (X^T X)^- C$ is smaller than the variance of any other linear unbiasd estimator for $C^T\beta$.
- The distribution of Y is not completely specified.

II. Generalized Least Squares Estimation

Consider the Aitken model

$$E(\mathrm{Y}) = X\beta \quad \text{and} \quad Var(\mathrm{Y}) = \sigma^2 V$$

where V is a positive definite symmetric matrix of known constants and σ^2 is an unknown variance parameter.

- A GLS estimator for β is any b that minimizes

  $$Q(\mathrm{b}) = (\mathrm{Y} - X\mathrm{b})^T V^{-1} (\mathrm{Y} - X\mathrm{b})$$

  (from Definition 3.8 with $\Sigma = \sigma^2 V$).

- The estimating equations are

  $$(X^T V^{-1} X)\mathrm{b} = X^T V^{-1} \mathrm{Y}.$$

- A solution is

  $$\mathrm{b}_{GLS} = (X^T V^{-1} X)^- X^T V^{-1} \mathrm{Y}.$$

- For any estimable function $C^T\beta$ the unique b.l.u.e. is

  $$C^T \mathrm{b}_{GLS} = C^T (X^T V^{-1} X)^- X^T V^{-1} \mathrm{Y}$$

  for any solution to the normal equations.

- $E(C^T\mathrm{b}) = C^T\beta$ and $Var(C^T\mathrm{b}) = \sigma^2 C^T (X^T V^{-1} X)^- C.$

- The distribution of Y is not completely specified.

- An unbiased estimator for σ^2 in the Aitken model is

  $$\hat{\sigma}^2_{GLS} = \frac{\mathrm{Y}^T \left[V^{-1} - V^{-1} X (X^T V^{-1} X)^- X^T V^{-1}\right] \mathrm{Y}}{n - rank(X)}$$

  $$= \frac{(\mathrm{Y} - X\mathrm{b}_{\mathrm{GLS}})^T V^{-1} (\mathrm{Y} - X\mathrm{b}_{\mathrm{GLS}})}{n}$$

- In practice, V may not be known. Then b_{GLS} and σ^2_{GLS} can be approximated by replacing V with a consistent estimator:
  - The estimator for $C^T\beta$ is not b.l.u.e.
  - The estimator for σ^2 is not unbiased.
  - Both estimators are consistent.

III. Maximum Likelihood Estimation

The model must include a specification of the joint distribution of the observations.

Example: Normal theory Gauss-Markov model:

$$Y_j = \beta_0 + \beta_1 X_{1_j} + \cdots + \beta_r X_{r_j} + \epsilon_j$$

where

$$\epsilon_j \sim \text{NID}(0, \sigma^2), \qquad i = 1, \ldots, n$$

or

$$\mathrm{Y} = \begin{bmatrix} Y_1 \\ \vdots \\ Y_n \end{bmatrix} \sim N(X\beta, \sigma^2 I)$$

- Find the parameter values that maximize the "likelihood" of the observed data.

For the normal-theory Gauss-Markov model, the likelihood function is

$$L(\beta, \sigma^2; Y_1, \ldots, Y_n) = \frac{1}{(2\pi)^{n/2}\sigma^n} e^{-\frac{1}{2\sigma^2}(\mathrm{Y}-X\beta)^T(\mathrm{Y}-X\beta)}$$

Find values of β and σ^2 that maximize this likelihood function.

- This is equivalent to finding values of β and σ^2 that maximize the log-likelihood.

$$\begin{aligned}
&\ell(\beta, \sigma^2; Y_1, \ldots, Y_n) \\
&= \log\left(L(\beta, \sigma^2; Y_1, \ldots, Y_n)\right) \\
&= -\frac{n}{2}\log(2\pi) - \frac{n}{2}\log(\sigma^2) \\
&\qquad -\frac{1}{2\sigma^2}(\mathrm{Y} - X\beta)^T(\mathrm{Y} - X\beta) \\
&= -\frac{n}{2}\log(2\pi) - \frac{n}{2}\log(\sigma^2) \\
&\qquad -\frac{1}{2\sigma^2}\underline{\sum_{j=1}^{n}(Y_j - \beta_0 - \cdots - \beta_r X_{r_j})^2}
\end{aligned}$$

↗ this is minimized by an OLS estimator for β regardless of the value of σ^2

Solve the likelihood equations:

$$\begin{aligned}
0 &= \frac{\partial \ell(\beta, \sigma^2; \mathrm{Y})}{\partial \beta_0} \\
&= \frac{1}{\sigma^2}\sum_{j=1}^{n}(Y_j - \beta_0 - \cdots - \beta_r X_{r_j}) \\
0 &= \frac{\partial \ell(\beta, \sigma^2; \mathrm{Y})}{\partial \beta_i} \\
&= \frac{1}{\sigma^2}\sum_{j=1}^{n} X_{ij}(Y_j - \beta_0 - \cdots - \beta_r X_{r_j}) \\
&\qquad \text{for } i = 1, 2, \ldots, r \\
0 &= \frac{\partial \ell(\beta, \sigma^2; \mathrm{Y})}{\partial \sigma^2} = -\frac{n}{2\sigma^2} \\
&\qquad + \frac{1}{2(\sigma^2)^2}\sum_{j=1}^{n}(Y_j - \beta_0 - \cdots - \beta_r X_{r_j})^2
\end{aligned}$$

Solution:

$$\hat{\beta} = \mathrm{b}_{\mathbf{OLS}} = (X^T X)^- X^T \mathrm{Y}$$

and

$$\begin{aligned}\hat{\sigma}^2 &= \frac{1}{n}\sum_{j=1}^{n}(Y_j - \hat{\beta}_0 - \cdots - \hat{\beta}_r X_{r_j})^2 \\ &= \frac{1}{n}\,\mathrm{Y}^T(I - P_X)\mathrm{Y} \\ &= \frac{1}{n}\,\mathbf{SSE}\end{aligned}$$

↗

- This is a biased estimator for σ^2.
- $\frac{1}{n-\mathbf{rank}(X)}SSE$ is an unbiased estimator for σ^2.
- $\frac{1}{n}SSE$ and $\frac{1}{n-\mathbf{rank}(X)}SSE$ are asymptotically equivalent.

Normal-theory Aitken model

$$\mathrm{Y} = X\beta + \epsilon$$

where $\epsilon \sim N(0, \sigma^2 V)$ and V is a known positive definite matrix.

The multivariate normal likelihood function is

$$L(\beta; \mathrm{Y}) = \frac{1}{(2\pi\sigma^2)^{n/2}|V|^{1/2}}\; e^{-\frac{1}{2\sigma^2}(\mathrm{Y}-X\beta)^T V^{-1}(\mathrm{Y}-X\beta)}$$

The log-likelihood function is

$$\begin{aligned}\ell(\beta; \mathrm{Y}) &= -\frac{n}{2}\log(2\pi) - \frac{1}{2}\log(|V|) \\ &\quad -\frac{n}{2}log(\sigma^2) \\ &\quad -\frac{1}{2\sigma^2}(\mathrm{Y} - X\beta)^T V^{-1}(\mathrm{Y} - X\beta)\end{aligned}$$

For any value of σ, the log-likelihood is maximized by finding a β that minimizes

$$(\mathrm{Y} - X\beta)^T V^{-1}(\mathrm{Y} - X\beta)$$

The estimating equations are

$$(X^T V^{-1} X)\beta = X^T V^{-1}\mathrm{Y}$$

Solutions are of the form

$$\hat{\beta} = \mathrm{b}_{\mathtt{GLS}} = (X^T V^{-1} X)^-\; X^T V^{-1}\mathrm{Y}$$

When V is known the mle for β is also the generalized least squares estimator.

The additional estimating equation corresponding to σ^2 is

$$0 = \frac{\partial \ell(\beta, \sigma^2; \mathrm{Y})}{\partial \sigma^2} = -\frac{n}{2\sigma^2} + \frac{1}{2(\sigma^2)^2}(\mathrm{Y} - X\beta)^T V^{-1}(\mathrm{Y} - X\beta)$$

Substituting the solution to the other estimating equations for β, the solution is

$$\hat{\sigma}^2 = \frac{1}{n}(\mathrm{Y} - X\mathrm{b}_{\mathrm{GLS}})^T V^{-1}(\mathrm{Y} - X\mathrm{b}_{\mathrm{GLS}})$$

↗ This is a biased estimator for σ^2.

When V contains unknown parameters:

- You could maximize the log-likelihood

$$\ell(\beta, \Sigma; \mathrm{Y}) = -\frac{n}{2}\log(2\pi) - \frac{1}{2}\log(|V|) - \frac{n}{2}\, log(\sigma^2) - \frac{1}{2\sigma^2}(\mathrm{Y} - X\beta)^T\ V^{-1}(\mathrm{Y} - X\beta)$$

with respect to β, σ^2 and the parameters in V.

- There may be no algebraic formulas for the solutions to the joint likelihood equations.
- The MLE's for σ^2 and the parameters in V are usually biased (too small).
- REML estimates are often used.

General Properties of MLE's

Regularity Conditions:

(i) The parameter space has finite dimension, is closed and compact, and the true parameter vector is in the interior of the parameter space.

(ii) Probability distributions defined by any two different values of the parameter vector are distinct (an identifiability condition).

(iii) First three partial derivatives of the log-likelihood function, with respect to the parameters

(a) exist
(b) are bounded by a function with a finite expectation.

(iv) The expectation of the negative of the matrix of second partial derivatives of the log-likelihood is

(a) finite
(b) positive definite

in a neighborhood of the true value of the parameter vector. This matrix is called the *Fisher information matrix*.

Suppose $\mathrm{Y}_1, \ldots, \mathrm{Y}_n$ are independent vectors of observations, with

$$\mathrm{Y}_j = \begin{bmatrix} Y_{1j} \\ \vdots \\ Y_{pj} \end{bmatrix},$$

and the density function (or probability function) is

$$f(\mathrm{Y}_j; \theta)$$

Then, the joint likelihood function is

$$L(\theta; \mathrm{Y}_1, \ldots, \mathrm{Y}_n) = \prod_{j=1}^{n} f(\mathrm{Y}_j; \theta)$$

The log-likelihood function is

$$\begin{aligned} \ell(\theta; \mathrm{Y}_1, \ldots, \mathrm{Y}_n) &= \log\left(L(\theta; \mathrm{Y}_1, \ldots, \mathrm{Y}_n)\right) \\ &= \sum_{j=1}^{n} \log\left(f(\mathrm{Y}_j; \theta)\right). \end{aligned}$$

The score function

$$\mathrm{u}(\theta) = \begin{bmatrix} u_1(\theta) \\ \vdots \\ u_r(\theta) \end{bmatrix} = \begin{bmatrix} \frac{\partial \ell(\theta; \mathrm{Y}_1, \ldots, \mathrm{Y}_n)}{\partial \theta_1} \\ \vdots \\ \frac{\partial \ell(\theta; \mathrm{Y}_1, \ldots, \mathrm{Y}_n)}{\partial \theta_r} \end{bmatrix}$$

is the vector of first partial derivatives of the log-likelihood function with respect to the elements of

$$\theta = \begin{bmatrix} \theta_1 \\ \vdots \\ \theta_r \end{bmatrix}.$$

The likelihood equations are

$$u(\theta; \mathrm{Y}_1, \ldots, \mathrm{Y}_n) = 0$$

The maximum likelihood estimator (MLE)

$$\hat{\theta} = \begin{bmatrix} \hat{\theta}_1 \\ \vdots \\ \hat{\theta}_r \end{bmatrix}$$

is a solution to the likelihood equations, that maximizes the log-likelihood function.

Fisher information matrix:

$$\begin{aligned} i(\theta) &= \mathrm{Var}(u(\theta; \mathrm{Y}_1, \ldots, \mathrm{Y}_n)) \\ &= E\left(u(\theta; \mathrm{Y}_1 \ldots, \mathrm{Y}_n) \right. \\ &\qquad \left. [u(\theta; \mathrm{Y}_1, \ldots, \mathrm{Y}_n)]^T\right) \\ &= -E\left(\left[\frac{\partial \ell(\theta; \mathrm{Y}_1, \ldots, \mathrm{Y}_n)}{\partial \theta_r \partial \theta_k}\right]\right) \end{aligned}$$

Let

- θ denote the parameter vector
- $i(\theta)$ denote the Fisher information matrix
- $\hat{\theta}$ denote the MLE for θ.

Then, if the Regularity Conditions are satisfied, we have the following results:

Result 9.1: $\hat{\theta}$ is a consistent estimator.

$$Pr\left\{(\hat{\theta}-\theta)^T(\hat{\theta}-\theta) > \epsilon\right\} \to 0$$

as $n \to \infty$, for any $\epsilon > 0$.

Result 9.2: Asymptotic normality

$$\sqrt{n}(\hat{\theta}-\theta) \stackrel{\text{dist'n}}{\longrightarrow} N\left(0, \lim_{n\to\infty} n[i(\theta)]^{-1}\right)$$

as $n \to \infty$.

With a slight abuse of notation we may express this as

$$\hat{\theta} \stackrel{\bullet}{\sim} N\left(\theta, [i(\theta)]^{-1}\right)$$

for "large" sample sizes.

Result 9.3: If $\hat{\theta}$ is the mle for θ, then the mle for $g(\theta)$ is $g(\hat{\theta})$ for any function $g(\ )$.

References:

Anderson, T.W. (1984). An Introduction to Multivariate Statistical Analysis, (2nd ed.), Wiley, New York.

Cox, C. (1984). American Statistician, 38, pp. 283–287.

Cox, D.R. and Hinkley, D.V. (1974). Theoretical Statistics, Chapman & Hall, London (Chapters 8 and 9).

Rao, C.R. (1973). Linear Statistical Inference, Wiley, New York (Chapter 5).