1.5 The Maximum Likelihood Estimator and the Information Matrix

We have now talked about how to construct likelihoods in a variety of settings, now we can use those constructions to formalize how we make inferences about model parameters.

parameter estimation, hypothesis tests, confidence intervals.

We often restrict attention to likelihoods that are continuously differentiable wrt $\underset{\sim}{\theta}$.

In this case, Recall the score function

$$S(\underset{\sim}{\theta}) = S(\boldsymbol{Y}, \boldsymbol{\theta}) = \begin{pmatrix} \frac{\partial \ell(\underset{\sim}{\theta})}{\partial \theta_1} \\ \vdots \\ \frac{\partial \ell(\underset{\sim}{\theta})}{\partial \theta_b} \end{pmatrix} = \begin{pmatrix} \frac{\partial \log L(\underset{\sim}{\theta} \mid \underset{\sim}{Y})}{\partial \theta_1} \\ \vdots \\ \frac{\partial \log L(\underset{\sim}{\theta} \mid \underset{\sim}{Y})}{\partial \theta_b} \end{pmatrix}$$

This function is random because it depends on the data $\underset{\sim}{Y}$.

?!

Generally, the maximum likelihood estimator $\hat{\boldsymbol{\theta}}_{\text{MLE}}$ is the value of $\boldsymbol{\theta}$ where the maximum (over the parameter space Θ) of $L(\boldsymbol{\theta}|\boldsymbol{Y})$ is attained.

$$\hat{\underset{\sim}{\theta}}_{MLE} = \underset{\underset{\sim}{\theta}}{\arg\max}\, L(\underset{\sim}{\theta} \mid \underset{\sim}{Y}) \iff L(\hat{\underset{\sim}{\theta}}_{MLE} \mid \underset{\sim}{Y}) \geq L(\underset{\sim}{\theta} \mid \underset{\sim}{Y}) \quad \forall\, \underset{\sim}{\theta} \in \Theta.$$

Under the assumption that the log-likelihood is continuously differentiable, then

$$S(\hat{\underset{\sim}{\theta}}_{MLE}) = 0.$$

But not always (?!).

Example (Exponential threshold model): Suppose that $Y_1, \dots, Y_n$ are iid from the exponential distribution with a threshold parameter μ,

strict.

$$f(y;\mu) = \begin{cases} \exp\{-(y-\mu)\} & \mu < y < \infty \\ 0 & \text{otherwise,} \end{cases}$$

for $\infty < \mu < \infty$.

$$L(\mu|\underline{Y}) = \prod_{i=1}^{n} f(Y_i;\mu) = \prod_{i=1}^{n} \exp(-(Y_i-\mu))\, \mathbb{I}(\mu < Y_i).$$

$$= \exp(-n\bar{Y})\exp(n\mu)\underbrace{\prod_{i=1}^{n}\mathbb{I}(\mu<Y_i)}_{\mathbb{I}(\mu<Y_{(1)})}.$$

$\Rightarrow$ the likelihood $=0$ for any value of $\mu \geq Y_{(1)}$, $Y_{(1)} = \min\{Y_1, \dots, Y_n\}$.

Consider the artificial data set $\boldsymbol{y} = [2.47, 2.35, 2.23, 3.53, 2.36]$.

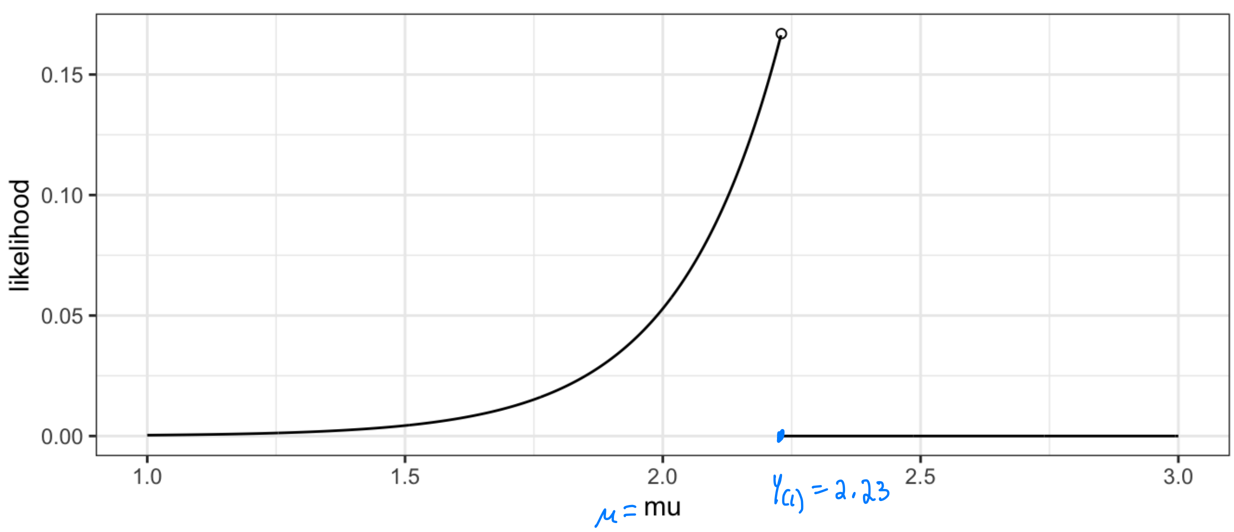

$\hat{\mu}_{MLE} = 2.23$, right? $L(2.23|\underline{Y}) = 0 \Rightarrow (*)\ L(\hat{\mu}_{MLE}|Y) \ngeq L(\mu|Y)\ \forall \mu \in \mathbb{R}$. AND

yes.

$S(\hat{\mu}_{MLE}) \neq 0$ because $\ell(\hat{\mu}_{MLE})$ not differentiable here.

If replace $\mu < y$ with $\mu \leq y$ in $f(y;\mu)$ then $(*)$ will hold by not score equation.

to see this consider maximizing $L_h(\mu|\underline{Y}) = \left(\frac{1}{2h}\right)^n \prod_{i=1}^{n} \{F_Y(Y_i+h;\mu) - F_Y(Y_i-h;\mu)\}$ for "small enough" value of h.

Then maximize $\lim_{h\to 0^+} L_h(\mu|\underline{Y})$. $\Rightarrow \hat{\mu}_{MLE} = Y_{(1)}$.

Rest of this section: assume support doesn't depend on the parameter value.

1.5.1 The Fisher Information Matrix

$Y_i \overset{iid}{\sim} f(y;\underline{\theta})$

The Fisher information matrix $\boldsymbol{I}(\boldsymbol{\theta})$ is defined as the $b \times b$ matrix where (dimension of $\underline{\theta}$)

$$I_{ij}(\boldsymbol{\theta}) = E\left[\left\{\frac{\partial}{\partial\theta_i}\log f(Y_1;\underline{\theta})\right\}\left\{\frac{\partial}{\partial\theta_j}\log f(Y_1;\underline{\theta})\right\}\right]$$

Is this random? No! it's an expectation!

Notice: this is the "information" in one observation. (note Y_1).

In matrix form,

$$\boldsymbol{I}(\boldsymbol{\theta}) = E\left[\underbrace{\left(\frac{\partial}{\partial\underline{\theta}^T}\log f(Y_1;\underline{\theta})\right)}_{\text{column vector.}}\underbrace{\left(\frac{\partial}{\partial\underline{\theta}}\log f(Y_1;\underline{\theta})\right)}_{\text{row vector}}\right]$$

||

Let $s(y;\underline{\theta}) = \left\{\frac{\partial}{\partial\underline{\theta}}\log f(y;\underline{\theta})\right\}^T$ ← column vector.

score contribution.

Then $\underline{I}(\underline{\theta}) = E\left[s(Y_1;\underline{\theta})\, s(Y_1;\underline{\theta})^T\right]$.

Again this depends on 1 observation (not n of them).

Fisher information facts:

1. The Fisher information matrix is the variance of the score contribution.

Why? $E[S(Y_i, \underline{\theta})] = 0$

Fact 1 from GLM section.

Big Result

2. If regularity conditions are met,

defined wrt a single observation.

$$\underbrace{\sqrt{n}(\hat{\boldsymbol{\theta}}_{\text{MLE}} - \boldsymbol{\theta})}_{\text{based on } n \text{ observations.}} \xrightarrow{d} \mathrm{N}_b(\mathbf{0}, \boldsymbol{I}(\boldsymbol{\theta})^{-1}).$$

unbiased inverse Fisher information.

"approximately distributed"

$\Rightarrow$ if n is large $\sqrt{n}(\hat{\underline{\theta}}_{MLE} - \underline{\theta}) \overset{\cdot}{\sim} N(\underline{0}, I(\underline{\theta})^{-1})$

or, $\hat{\underline{\theta}}_{MLE} - \underline{\theta} \overset{\cdot}{\sim} N(\underline{0}, \frac{1}{n} I(\underline{\theta})^{-1})$

$\hat{\underline{\theta}}_{MLE} \overset{\cdot}{\sim} N(\underline{\theta}, \{n I(\underline{\theta})\}^{-1})$ $(*)$.

We will prove this result for $b=1$. (later).

3. If $b = 1$, then any unbiased estimator must have variance greater than or equal to $\{nI(\boldsymbol{\theta})\}^{-1}$

↑ Cramer-Rao lower bound.

If $b \neq 1$: If Σ is the asymptotic cov matrix of any other consistent estimator, then

$$\Sigma - I(\theta)^{-1} \text{ is positive definite.}$$

4. The information matrix is related to the $\underbrace{\text{curvature}}_{\text{Hessian}}$ of the log-likelihood contribution.

$$I(\underline{\theta}) = E\left[\left(\frac{\partial}{\partial \underline{\theta}^T} \log f(Y_1; \underline{\theta})\right)\left(\frac{\partial}{\partial \underline{\theta}} \log f(Y_1; \underline{\theta})\right)\right]$$

$$= E\left[-\frac{\partial^2}{\partial \underline{\theta}\, \partial \underline{\theta}^T} \log f(Y_1; \underline{\theta})\right]$$ ← assuming ℓ is twice differentiable and using fact ② from GLM section.

$$= E\left[-\frac{\partial}{\partial \underline{\theta}} s(Y_1; \underline{\theta})\right]$$ (writing another way).

1.5.2 Observed Information

The information matrix is not random, but it is also not observable from the data.

Let $Y_1, \ldots, Y_n$ be iid with density $f_Y(y_i; \boldsymbol{\theta})$. The log likelihood is defined as

taking two derivatives and dividing by n results in

Definition: The matrix $n\bar{I}\,(Y;\hat{\boldsymbol{\theta}}_{\mathrm{MLE}})$ is called the sample information matrix, or the *observed information matrix.*

Why use $I(\boldsymbol{\theta})=\mathrm{E}\left[-\frac{\partial^2}{\partial\boldsymbol{\theta}\partial\boldsymbol{\theta}^\top}\log f(Y_1;\boldsymbol{\theta})\right]$ as the basis for an estimator, rather than $I(\boldsymbol{\theta})=\mathrm{E}\left[\left\{\frac{\partial}{\partial\boldsymbol{\theta}^\top}\log f(Y_1;\boldsymbol{\theta})\right\}\left\{\frac{\partial}{\partial\boldsymbol{\theta}}\log f(Y_1;\boldsymbol{\theta})\right\}\right]$?

Now let's prove the asymptotic normality of the MLE (in the scalar case).