## Contents

1 Foundation ..... 2
1.1 Review on MLE ..... 2
1.2 Review on GLM ..... 4
1.2.1 Newly introduced in 711 ..... 5
1.3 Families of distributions, moments / cumulants, and quantiles / percentiles ..... 6
1.4 Transformation of random variables ..... 8
1.5 Others ..... 8
2 Likelihood construction and estimation ..... 8
2.1 Introduction ..... 8
2.2 Likelihood construction ..... 9
2.3 Proportional likelihoods ..... 10
2.4 Empirical distribution function as an MLE ..... 11
2.5 Likelihoods for censored / truncated data ..... 12
2.6 Likelihoods for regression models ..... 12
2.7 Marginal and conditional likelihoods ..... 15
2.8 MLE and information matrix ..... 16
2.8.1 Transformed and modeled parameters ..... 19
2.9 Methods for maximizing the likelihood ..... 19
2.9.1 Why does EM work? ..... 22
2.9.2 Calculating observed info matrix after EM ..... 23
2.10 Uniqueness of MLE ..... 24
3 Likelihood-based tests and confidence regions ..... 26
3.1 Simple null hypothesis ..... 27
3.2 Composite null hypothesis ..... 28
3.3 Confidence interval ..... 31
3.4 Nonstandard hypothesis testing problems ..... 32
4 Bayesian methods ..... 33
4.1 Introduction ..... 33
4.2 Bayesian estimator from decision theory perspective ..... 34
4.3 Credible intervals ..... 35
4.4 Conjugate prior ..... 35
4.5 Noninformative prior ..... 36
4.6 Normal data examples ..... 37
4.6.1 One sample with unknown mean and variance ..... 37
4.6.2 Two samples ..... 38
4.6.3 Normal linear model ..... 38
4.7 Hierarchical Bayes and empirical Bayes ..... 39
4.7.1 James-Stein estimation ..... 39
4.7.2 Meta-analysis applications of hierarchical and empirical Bayes ..... 40
4.8 Monte Carlo estimation of a posterior ..... 41
4.8.1 Noniterative Monte Carlo methods ..... 41
4.9 MCMC methods ..... 42
4.9.1 Substitution sampling ..... 42
4.9.2 Gibbs sampling ..... 43
4.9.3 Metropolis-Hastings algorithm ..... 44
4.9.4 Hybrid forms ..... 45
5 Large sample theory ..... 45
6 M-Estimation (Estimating Equations) ..... 45
6.1 Introduction ..... 45
6.2 Basic approach ..... 46
6.2.1 Estimation for A, B, and V ..... 46
6.3 Delta method via M-estimation ..... 47

## 1 Foundation

### 1.1 Review on MLE

Central limit theorem: Suppose $\left\{X_{1}, X_{2}, \ldots, X_{n}\right\}$ is a sequence of i.i.d. random variables with $E\left(X_{i}\right)=\mu$ and $\operatorname{Var}\left(X_{i}\right)=\sigma^{2}<\infty$ then as $n \rightarrow \infty, \bar{X} \rightarrow N\left(\mu, \frac{\sigma^{2}}{n}\right)$
Score:The partial derivative with respect to $\theta$ of the natural logarithm of the likelihood function is called the score

$$
\begin{gathered}
Z=l^{\prime}=\frac{\partial}{\partial \theta} \log f(X ; \theta) \\
E(Z)=0 \text { and } Z \xrightarrow{d} N\left(0, I\left(\theta_{0}\right)\right)
\end{gathered}
$$

under $\theta_{0}$
Fisher information: The variance of the score is defined to be the Fisher information

$$
\mathcal{I}(\theta)=\mathrm{E}\left[\left.\left(\frac{\partial}{\partial \theta} \log f(X ; \theta)\right)^{2} \right\rvert\, \theta\right]=-\mathrm{E}\left[\left.\frac{\partial^{2}}{\partial \theta^{2}} \log f(X ; \theta) \right\rvert\, \theta\right]
$$

Total information and observed total information

$$
\begin{aligned}
\mathbf{I}_{T}(\boldsymbol{\theta}) & =-E\left\{\frac{\partial^{2}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \log L(\boldsymbol{\theta} \mid \mathbf{Y})\right\} \\
\mathbf{I}_{T}(\mathbf{Y}, \boldsymbol{\theta}) & =-\frac{\partial^{2}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \log L(\boldsymbol{\theta} \mid \mathbf{Y})
\end{aligned}
$$

Information estimation

- Version 1

$$
\begin{aligned}
\overline{\mathbf{I}}(\mathbf{Y}, \boldsymbol{\theta}) & =n^{-1} \sum_{i=1}^{n}\left\{-\frac{\partial}{\partial \boldsymbol{\theta}} \mathbf{s}\left(Y_{i}, \boldsymbol{\theta}\right)\right\} \\
& =n^{-1} \sum_{i=1}^{n}\left\{-\frac{\partial^{2}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \log f\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}
\end{aligned}
$$

- Version 2

$$
\begin{aligned}
\overline{\mathbf{I}}^{*}(\mathbf{Y}, \boldsymbol{\theta}) & =n^{-1} \sum_{i=1}^{n} \mathbf{s}\left(Y_{i}, \boldsymbol{\theta}\right)^{\otimes 2} \\
& =n^{-1} \sum_{i=1}^{n}\left\{\frac{\partial}{\partial \boldsymbol{\theta}^{T}} \log f\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}^{\otimes 2}
\end{aligned}
$$

Property 1. If $\hat{\theta}$ is the MLE estimate of $\theta_{0}$, then it has the following property:

$$
\sqrt{n}\left(\hat{\theta}-\theta_{0}\right) \xrightarrow{d} N\left(0, \frac{1}{I\left(\theta_{0}\right)}\right)
$$

## Ways for estimating MLE:

- Newton-Raphson algorithm

$$
\begin{aligned}
\mathbf{0}=\mathbf{S}(\boldsymbol{\theta}) & \approx \mathbf{S}\left(\boldsymbol{\theta}^{(\nu)}\right)+\left\{\left.\frac{\partial}{\partial \boldsymbol{\theta}} \mathbf{S}(\boldsymbol{\theta})\right|_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(\nu)}}\right\}\left(\boldsymbol{\theta}-\boldsymbol{\theta}^{(\nu)}\right) \\
& =\mathbf{S}\left(\boldsymbol{\theta}^{(\nu)}\right)-\mathbf{I}_{T}\left(\mathbf{Y}, \boldsymbol{\theta}^{(\nu)}\right)\left(\boldsymbol{\theta}-\boldsymbol{\theta}^{(\nu)}\right) \\
\boldsymbol{\theta}^{(\nu+1)} & =\boldsymbol{\theta}^{(\nu)}+\mathbf{I}_{T}\left(\mathbf{Y}, \boldsymbol{\theta}^{(\nu)}\right)^{-1} \mathbf{S}\left(\boldsymbol{\theta}^{(\nu)}\right)
\end{aligned}
$$

1. start with initial $\theta^{(0)}$
2. update from current $\theta^{(v)}$ to obtain $\theta^{(v+1)}$
3. stop if $\left\|\mathbf{S}\left(\boldsymbol{\theta}^{(\nu+1)}\right)\right\|$ or $\left\|\boldsymbol{\theta}^{(\nu+1)}-\boldsymbol{\theta}^{(\nu)}\right\|$ is sufficiently small

Note: Newton-Raphson estimator has local quadratic convergence property

$$
\left\|\boldsymbol{\theta}^{(\nu+1)}-\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right\| \leq c\left\|\boldsymbol{\theta}^{(\nu)}-\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right\|^{2} \text { for some } c>0
$$

This property exists for the situation where the estimator is very close to the true parameter. Thus, at least near the solution, convergence is fast for a Newton method.
More specifically, the local quadratic convergence holds under the following conditions

1. $I_{T}(\mathbf{Y}, \theta) \neq 0$ in a neighborhood of $\hat{\theta}_{\text {MLE }}$
2. $S^{\prime \prime}(\theta)$ is bounded
3. $\theta^{(\nu)}$ is sufficiently close to $\hat{\theta}_{\text {MLE }}$

Proof draft:

$$
\begin{aligned}
0=S\left(\widehat{\theta}_{\mathrm{MLE}}\right)= & S\left(\theta^{(\nu)}\right)-I_{T}\left(\mathbf{Y}, \theta^{(\nu)}\right)\left(\widehat{\theta}_{\mathrm{MLE}}-\theta^{(\nu)}\right) \\
& +\frac{1}{2} S^{\prime \prime}\left(\vartheta^{(\nu)}\right)\left(\widehat{\theta}_{\mathrm{MLE}}-\theta^{(\nu)}\right)^{2} \\
\widehat{\theta}_{\mathrm{MLE}}-\theta^{(\nu+1)}= & \frac{1}{2} \frac{S^{\prime \prime}\left(\vartheta^{(\nu)}\right)}{I_{T}\left(\mathbf{Y}, \theta^{(\nu)}\right)}\left(\widehat{\theta}_{\mathrm{MLE}}-\theta^{(\nu)}\right)^{2}
\end{aligned}
$$

where $\vartheta^{(\nu)}=\theta^{(\nu)}+k *\left(\theta_{\text {MLE }}-\theta^{(\nu)}\right), k \in[0,1]$. This uses the Taylor expansion with Lagrange remainder. One-step estimator: typically, $\widehat{\boldsymbol{\theta}}_{\text {MLE }}-\boldsymbol{\theta}=O_{p}\left(n^{-1 / 2}\right)$. If one starts with $\theta^{(0)}$ such that $\boldsymbol{\theta}^{(0)}-\boldsymbol{\theta}=O_{p}\left(n^{-1 / 2}\right)$, then

$$
\boldsymbol{\theta}^{(1)}-\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}=O_{p}\left(n^{-1}\right)
$$

under regularity conditions. That is $\theta^{(0)}$ and $\hat{\theta_{\text {MLE }}}$ is asymptotically equivalent.

- Fisher scoring algorithm
$\mathbf{I}_{T}\left(\mathbf{Y}, \boldsymbol{\theta}^{(\nu)}\right)$ replaced by its expectation $\mathbf{I}_{T}\left(\boldsymbol{\theta}^{(\nu)}\right)$, which is

$$
\boldsymbol{\theta}^{(\nu+1)}=\boldsymbol{\theta}^{(\nu)}+\mathbf{I}_{T}\left(\boldsymbol{\theta}^{(\nu)}\right) \mathbf{S}\left(\boldsymbol{\theta}^{(\nu)}\right)
$$

- EM algorithm

The basic idea of the EM Algorithm is to view the observed data Y as incomplete, that somehow there is missing data Z that would make the problem simpler if we had it. In some cases Z could truly be missing data, but in others it is just additional data that we wish we had.

1. The first step is to write down the joint likelihood of the "complete" data $(\boldsymbol{Y}, \boldsymbol{Z})$, call it $L_{\mathrm{C}}(\boldsymbol{\theta} \mid \boldsymbol{Y}, \boldsymbol{Z})$. The "E" step of the EM Algorithm is to compute the conditional expectation of $\log L_{\mathrm{C}}(\boldsymbol{\theta} \mid \boldsymbol{Y}, \boldsymbol{Z})$ given $\boldsymbol{Y}$ assuming the true parameter value is $\boldsymbol{\theta}^{(\nu)}$

$$
\begin{aligned}
Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(\nu)}, \boldsymbol{Y}\right) & =\mathrm{E}_{\boldsymbol{\theta}^{(v)}}\left\{\log L_{\mathrm{C}}(\boldsymbol{\theta} \mid \boldsymbol{Y}, \boldsymbol{Z}) \mid \boldsymbol{Y}\right\} \\
& =\int \log L_{\mathrm{C}}(\boldsymbol{\theta} \mid \boldsymbol{Y}, \boldsymbol{z}) f_{Z \mid Y}\left(\boldsymbol{z} \mid \boldsymbol{Y}, \boldsymbol{\theta}^{(v)}\right) d z
\end{aligned}
$$

2. M step: calculate $\boldsymbol{\theta}^{(v+1)}$ that maximizes $Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(\nu)}, \boldsymbol{Y}\right)$ wrt $\theta$

The plausibility of EM

Lemma 1.1. For $\mathbf{Y} \sim f\left(\mathbf{y} ; \boldsymbol{\theta}_{0}\right), E_{\boldsymbol{\theta}_{0}}\{\log f(\mathbf{Y} ; \boldsymbol{\theta})\}$ is maximized at $\boldsymbol{\theta}=\boldsymbol{\theta}_{0}$
Proof: Note that $\psi(x)=-\log (x)$ is a convex function for $x \in(0, \infty)$, since $\psi^{\prime \prime}(x)=x^{-2}>0$. Then by Jensen's inequality we have

$$
-\log \left[E_{\boldsymbol{\theta}_{0}}\left\{\frac{f(\mathbf{Y} ; \boldsymbol{\theta})}{f\left(\mathbf{Y} ; \boldsymbol{\theta}_{0}\right)}\right\}\right] \leq-E_{\boldsymbol{\theta}_{0}}\left[\log \left\{\frac{f(\mathbf{Y} ; \boldsymbol{\theta})}{f\left(\mathbf{Y} ; \boldsymbol{\theta}_{0}\right)}\right\}\right]
$$

and since $E_{\boldsymbol{\theta}_{0}}\left\{\frac{f(y ; \boldsymbol{\theta})}{f\left(y ; \boldsymbol{\theta}_{0}\right)}\right\}=\int \frac{f(y ; \boldsymbol{\theta})}{f\left(y ; \boldsymbol{\theta}_{0}\right)} f\left(y ; \boldsymbol{\theta}_{0}\right) d y=1$ we get $E_{\boldsymbol{\theta}_{0}}\left[\log \left\{\frac{f(\mathbf{Y} ; \boldsymbol{\theta})}{f\left(\mathbf{Y} ; \boldsymbol{\theta}_{0}\right)}\right\}\right] \leq 0$

### 1.2 Review on GLM

some materials: on two forms, canonical forms, GLM and exponential family, Exponential family: Y with distribution with the density of the form

$$
\begin{aligned}
Y & \sim f(y ; \theta, \phi) \\
& =\exp \left\{\frac{y \theta-b(\theta)}{a(\phi)}+c(y, \phi)\right\}
\end{aligned}
$$

where

- $\theta$ is the canonical (natural) parameter - often the parameter of interest
- $\Phi$ is dispersion (scale) parameter - often nuisance parameter

Exponential family also has another commonly used form A good material here

$$
p(x \mid \eta)=h(x) \exp \left\{\eta^{T} T(x)-A(\eta)\right\}
$$

parameter vector $\eta$, often referred to as the canonical parameter for given functions $T$ and $h$. The statistic T(X) is referred to as a sufficient statistic. The function $A(\eta)$ is known as the cumulant function and

$$
A(\eta)=\log \int h(x) \exp \left\{\eta^{T} T(x)\right\} \nu(d x)
$$

with respect to the measure $\nu$
For example, for normal distribution,

$$
\begin{aligned}
f\left(y ; \mu, \sigma^{2}\right) & =\frac{1}{\sqrt{2 \pi \sigma^{2}}} \exp \left\{-\frac{(y-\mu)^{2}}{2 \sigma^{2}}\right\} \\
& =\exp \left\{\frac{y \mu-\mu^{2} / 2}{\sigma^{2}}-\frac{y^{2}}{2 \sigma^{2}}-\frac{1}{2} \log \left(2 \pi \sigma^{2}\right)\right\}
\end{aligned}
$$

then

- $\theta=\mu$
- $\phi=\sigma^{2}$
- $a(\phi)=\phi$
- $b(\theta)=\frac{\theta^{2}}{2}$
- $c(y, \phi)=-\frac{y^{2}}{2 \phi}-\frac{1}{2} \log (2 \pi \phi)$
or

$$
f\left(x ; \mu, \sigma^{2}\right)=\frac{1}{\sqrt{2 \pi} \sigma} \exp \left\{-\frac{1}{2 \sigma^{2}}(x-\mu)^{2}\right\}=\frac{1}{\sqrt{2 \pi}} \exp \left\{\frac{\mu}{\sigma^{2}} x-\frac{1}{2 \sigma^{2}} x^{2}-\frac{1}{2 \sigma^{2}} \mu^{2}-\log \sigma\right\}
$$

then

$$
\begin{aligned}
\eta & =\left[\begin{array}{c}
\mu / \sigma^{2} \\
-1 / 2 \sigma^{2}
\end{array}\right] \\
T(x) & =\left[\begin{array}{c}
x \\
x^{2}
\end{array}\right] \\
A(\eta) & =\frac{\mu^{2}}{2 \sigma^{2}}+\log \sigma=-\frac{\eta_{1}^{2}}{4 \eta_{2}}-\frac{1}{2} \log \left(-2 \eta_{2}\right) \\
h(x) & =\frac{1}{\sqrt{2 \pi}}
\end{aligned}
$$

We also have the following property
Property 1.

$$
\begin{aligned}
E(Y) & =b^{\prime}(\theta) \\
\operatorname{Var}(Y) & =b^{\prime \prime}(\theta) a(\phi)
\end{aligned}
$$

Proof

$$
\ell(\theta)=\log f(Y ; \theta, \phi)=\frac{Y \theta-b(\theta)}{a(\phi)}+c(Y, \phi)
$$

Then the score function for $\theta$ is

$$
U(\theta)=\frac{\partial \ell(\theta)}{\partial \theta}=\frac{Y-b^{\prime}(\theta)}{a(\phi)}
$$

Since $E(U(\theta))=0$, we have

$$
E\left(\frac{Y-b^{\prime}(\theta)}{a(\phi)}\right)=0 \Longrightarrow E(Y)=b^{\prime}(\theta)
$$

Since $E(U(\theta))=0$, we have

$$
\operatorname{Var}(U(\theta))=E\left(\frac{\partial \ell(\theta)}{\partial \theta}\right)^{2}=-E\left(\frac{\partial^{2} \ell(\theta)}{\partial \theta^{2}}\right) \Longrightarrow \frac{\operatorname{Var}(Y)}{a^{2}(\phi)}=\frac{b^{\prime \prime}(\theta)}{a(\phi)} \Longrightarrow \operatorname{Var}(Y)=b^{\prime \prime}(\theta) a(\phi)
$$

### 1.2.1 Newly introduced in 711

Exponential family distributions $\left\{\mathcal{F}_{\boldsymbol{\theta}}, \boldsymbol{\theta} \in \Omega\right\}$ have densities of the general form

$$
f(x ; \boldsymbol{\theta})=h(x) \exp \left\{\sum_{i=1}^{s} g_{i}(\boldsymbol{\theta}) T_{i}(x)-B(\boldsymbol{\theta})\right\}
$$

this representation is not unique. A version with the smallest $s$ is called minimal exponential family under this case both $T_{i}^{\prime} s$ and $g_{i} \theta^{\prime} s$ are free of linear constraints, which means

$$
\sum_{i=1}^{s} c_{i} T_{i}(x)=c_{s+1} \text { a.e. } \nu \Longrightarrow c_{i}=0, i=1, \cdots, s+1
$$

With minimal exponential family $\left(T_{1}, \cdots, T_{s}\right)$ is minimal sufficient. Natural parameter space

$$
\Omega_{0}=\left\{\boldsymbol{\theta}: \int h(x) \exp \left\{\sum_{i=1}^{s} g_{i}(\boldsymbol{\theta}) T_{i}(x)\right\} d \nu(x)<\infty\right\}
$$

Canonical representation

$$
f(x ; \eta)=h(x) \exp \left\{\sum_{i=1}^{s} \eta_{i} T_{i}(x)-A(\boldsymbol{\eta})\right\}
$$

by integration equal to 1 we have

$$
\int h(x) \exp \left\{\sum_{i=1}^{s} \eta_{i} T_{i}(x)\right\} d \nu(x)=\exp [A(\eta)]
$$

with the canonical parameter $\boldsymbol{\eta}=g(\boldsymbol{\theta})$ and accordingly

$$
\Omega_{\eta}=\left\{\eta: \int h(x) \exp \left\{\sum_{i=1}^{s} \eta_{i} T_{i}(x)\right\} d \nu(x)<\infty\right\}
$$

is another form of the natural parameter space.
If $\eta$ is in the interior of $\Omega_{\eta}$ then the m.g.f of $T$ exists in a neighborhood of $u=0$

$$
\begin{aligned}
M_{\mathbf{T}}(\mathbf{u}) & =E e^{u_{1} T_{1}+\cdots+u_{s} T_{s}} \\
& =\int \exp \left\{\sum u_{i} T_{i}\right\} \cdot h(x) \exp \left\{\sum \eta_{i} T_{i}-A(\eta)\right\} d v(x) \\
& =\int h(x) \exp \left\{\sum\left(\eta_{i}+u_{i}\right) T_{i}-A(\eta)\right\} d v(x) \\
& =\int h(x) \frac{\exp \left\{\sum\left(\eta_{i}+u_{i}\right) T_{i}\right\}}{\exp \{A(\eta)\}} d v(x) \\
& =\frac{\exp \{A(\eta+u)\}}{\exp \{A(\eta)\}}=e^{A(\boldsymbol{\eta}+\mathbf{u})-A(\boldsymbol{\eta})}
\end{aligned}
$$

$$
K_{\mathbf{T}}(\mathbf{u})=\log \left\{M_{\mathbf{T}}(\mathbf{u})\right\}=A(\boldsymbol{\eta}+\mathbf{u})-A(\boldsymbol{\eta})
$$

therefore

$$
\frac{\partial A(\boldsymbol{\eta})}{\partial \boldsymbol{\eta}^{T}}=E\{\mathbf{T}(X)\}, \quad \frac{\partial^{2} A(\boldsymbol{\eta})}{\partial \eta \partial \eta^{T}}=\operatorname{Var}\{\mathbf{T}(X)\}
$$

also

$$
\mathbf{I}(\boldsymbol{\eta})=\frac{\partial^{2} A(\boldsymbol{\eta})}{\partial \boldsymbol{\eta} \partial \boldsymbol{\eta}^{T}}=\operatorname{Var}\{\mathbf{T}(X)\}
$$

curved exponential family: $\operatorname{dim}(\eta)>\operatorname{dim}(\theta)$
Completeness: a set of statistics $\mathbf{T}=\left(T_{1}, \cdots, T_{s}\right)$ is complete with respect to the family of their induced distributions indexed by $\theta$ if there are no functions $\Phi(T)$ other than $\Phi=0$ such that $E_{\boldsymbol{\theta}} \phi(\mathbf{T})=0$ for all $\theta \in \Omega$

$$
\text { complete sufficient statistics } \nleftarrow \text { minimal sufficient statistics }
$$

Lehmann-Scheffe theorem: If a statistic is unbiased, complete and sufficient, then it is the unique best unbiased estimator.
Sufficient statistic ( $T_{1}, \cdots, T_{s}$ ) of a canonical exponential family is complete sufficient provided that the family is minimal and that the parameter space contains an s-dimensional rectangle (which ruled out the curve exponential family).

## Distribution of sufficient statistics

$$
\begin{aligned}
& X \sim \text { exponential family with } \boldsymbol{\eta}=(\boldsymbol{\zeta}, \boldsymbol{\psi}): \\
& f(x ; \boldsymbol{\zeta}, \boldsymbol{\psi})=h(x) \exp \left\{\sum_{i=1}^{r} \zeta_{i} U_{i}(x)+\sum_{j=1}^{s} \psi_{j} T_{j}(x)-A(\boldsymbol{\zeta}, \boldsymbol{\psi})\right\}
\end{aligned}
$$

then

1. the distribution of $(U, T)$ is an exponential family
2. marginally $T$ is an exponential family

$$
f(\mathbf{t} ; \boldsymbol{\zeta}, \boldsymbol{\psi})=q(\mathbf{t}) C(\boldsymbol{\zeta}) \exp \left\{\sum_{j=1}^{s} \psi_{j} t_{j}-A(\boldsymbol{\zeta}, \boldsymbol{\psi})\right\}
$$

3. conditional distribution of $U$ given $T=t$ is an exponential family

$$
f(\mathbf{u} \mid \mathbf{t} ; \boldsymbol{\zeta})=q_{\mathbf{t}}(\mathbf{u}) \exp \left\{\sum_{i=1}^{r} \zeta_{i} u_{i}-A_{\mathbf{t}}(\boldsymbol{\zeta})\right\}
$$

### 1.3 Families of distributions, moments / cumulants, and quantiles / percentiles

## Families of distributions

- Parametric model or parametric family: a set of distributions indexed by a finite dimensional parameter

$$
\boldsymbol{\theta}=\left(\theta_{1}, \cdots, \theta_{b}\right)^{T}, \quad\{F(y ; \boldsymbol{\theta}): \boldsymbol{\theta} \in \boldsymbol{\Theta}\}
$$

- Semiparametric model: a set of distributions indexed by a finite dimensional $\theta$ and some infinite dimensional parameters
for example
$\left\{f\left(y ; \mu, f_{0}\right)=f_{0}(y-\mu):-\infty<\mu<\infty, f_{0} \in C\right.$, where $C$ is a class of continuous unimodeal densities $\}$ where $\mu$ belongs to the parametric part and the $f_{0}$ belongs to the nonparametric part
- Nonparametric model: a set of distributions indexed by infinite dimensional parameters for example

$$
Y_{i}=g\left(X_{i}\right)+e_{i}, \quad \text { density of } e \in C
$$

where $g($.$) belongs to a class of functions with certain smoothness$

- jth population moment of random variable $Y$

$$
\mu_{j}^{\prime}=E\left(Y^{j}\right), j=1,2, \cdots
$$

with its sample counterparts as

$$
m_{j}^{\prime}=n^{-1} \sum_{i=1}^{n} Y_{i}^{j}
$$

- jth population central moment of rv Y

$$
\mu_{j}=E\left[\{Y-E(Y)\}^{j}\right]
$$

with its sample counterparts as

$$
m_{j}=n^{-1} \sum_{i=1}^{n}\left(Y_{i}-\bar{Y}\right)^{j}
$$

- moment generating function (MGF)

$$
m(t)=E\{\exp (t Y)\}
$$

a useful trick is to write the MGF as

$$
\begin{aligned}
M(t)=E\left(e^{t X}\right) & =E\left(1+t X+\cdots+t^{r} X^{r} / r!+\cdots\right) \\
& =\sum_{r=0}^{\infty} \mu_{r} t^{r} / r!
\end{aligned}
$$

which implies that

$$
\mu_{r}^{\prime}=M^{(r)}(0)
$$

- cumulant generating function (CGF)

$$
k(t)=\log \{m(t)\}
$$

it's useful to remember that

$$
K(t)=\log M(t)=\sum_{r} \kappa_{r} t^{r} / r!
$$

which is just from the Taylor expansion for $K(t)$ (note: the Taylor expansion around 0 is $f(x)=$ $\left.\sum_{n=0}^{\infty} \frac{f^{(n)}(0)}{n!} x^{n}\right)$ a nice material on this topic

- jth population cumulant

$$
\kappa_{j}=k^{(j)}(0)
$$

for example

$$
\begin{aligned}
& \kappa_{1}=\mu_{1}^{\prime}(\text { mean }) \\
& \kappa_{2}=\mu_{2}^{\prime}-\kappa_{1}^{2}(\text { variance })
\end{aligned}
$$

proof:

$$
\begin{aligned}
\kappa_{1}=\kappa^{\prime}(0) & =\left.\frac{m^{\prime}(t)}{m(t)}\right|_{t=0} \\
& =\left.\frac{E[\exp (t Y) Y]}{E[\exp (t Y)]}\right|_{t=0} \\
& =E[Y]=\mu_{1}^{\prime}
\end{aligned}
$$

and

$$
\begin{aligned}
\kappa^{\prime \prime}(t) & =\left.\frac{m^{\prime \prime}(t) m(t)-\left[m^{\prime}(t)\right]^{2}}{m^{2}(t)}\right|_{t=0} \\
& =E\left(Y^{2}\right)-E(Y)^{2}=\mu_{2}^{\prime}-\kappa_{1}^{2}
\end{aligned}
$$

- skewness

$$
\text { skew }=\mathrm{E}\left[\left(\frac{X-\mu}{\sigma}\right)^{3}\right]=\frac{\mu_{3}}{\sigma^{3}}=\frac{\mathrm{E}\left[(X-\mu)^{3}\right]}{\left(\mathrm{E}\left[(X-\mu)^{2}\right]\right)^{3 / 2}}=\mu_{3} / \mu_{2}^{3 / 2}
$$

which measures symmetry. Note that we are using the central moment in the definition

- kurtosis

$$
\text { kurt }=\mathrm{E}\left[\left(\frac{X-\mu}{\sigma}\right)^{4}\right]=\frac{\mathrm{E}\left[(X-\mu)^{4}\right]}{\left(\mathrm{E}\left[(X-\mu)^{2}\right]\right)^{2}}=\mu_{4} / \mu_{2}^{2}
$$

which measures tail mass. Note that for Normal distribution skew $=0$, kurt $=3$

- coefficient of variation

$$
\mathrm{CV}=\frac{\sigma}{\mu}=\mu_{2}^{1 / 2} / \mu
$$

for non-negative random variable, providing a relative standard deviation

- quantile

$$
\eta_{p} \equiv \inf \{x: F(x) \geq p\}=F^{-1}(p)
$$

the last equality holds only when $F$ is a strictly monotone function and $X$ is continuous

### 1.4 Transformation of random variables

Jacobian method: Let $f_{X_{1} X_{2}}\left(x_{1} x_{2}\right)$ be the value of the joint probability density of the continuous random variables $X_{1}$ and $X_{2}$ at $\left(x_{1}, x_{2}\right)$. If the functions given by $y_{1}=u_{1}\left(x_{1}, x_{2}\right)$ and $y_{2}=u_{2}\left(x_{1}, x_{2}\right)$ are partially differentiable with respect to $x_{1}$ and $x_{2}$ and represent a one-to-one transformation for all values within the range of $X_{1}$ and $X_{2}$ for which $f_{X_{1} X_{2}}\left(x_{1} x_{2}\right) \neq 0$, then, for these values of $x_{1}$ and $x_{2}$, the equations $y_{1}=u_{1}\left(x_{1}, x_{2}\right)$ and $y_{2}=u_{2}\left(x_{1}, x_{2}\right)$ can be uniquely solved for $x_{1}$ and $x_{2}$ to give $x_{1}=w_{1}\left(y_{1}, y_{2}\right)$ and $x_{2}=w_{2}\left(y_{1}, y_{2}\right)$ and for corresponding values of $y_{1}$ and $y_{2}$, the joint probability density of $Y_{1}=u_{1}\left(X_{1}, X_{2}\right)$ and $Y_{2}=u_{2}\left(X_{1}, X_{2}\right)$ is given by

$$
f_{Y_{1} Y_{2}}\left(y_{1} y_{2}\right)=f_{X_{1} X_{2}}\left[w_{1}\left(y_{1}, y_{2}\right), w_{2}\left(y_{1} y_{2}\right)\right] \cdot|J|
$$

where $J$ is the Jacobian of the transformation

$$
J=\left|\begin{array}{cc}
\frac{\partial x_{1}}{y_{1}} & \frac{\partial x_{1}}{y_{2}} \\
\frac{\partial x_{2}}{y_{1}} & \frac{\partial x_{2}}{y_{2}}
\end{array}\right|
$$

all other points of $f_{Y_{1} Y_{2}}\left(y_{1} y_{2}\right)=0$
Distribution function method: suppose we have the distribution function for $Y$ as $F_{Y}(y ; \theta)=P(Y \leq y)$. But we want the distribution function of $X=g(Y)$. Then

$$
F_{X}(x ; \theta)=P\{X \leq x\}=P\{g(Y) \leq x\}
$$

if $g$ is a strictly increasing function which means $g^{-1}$ exists then

$$
F_{X}(x ; \boldsymbol{\theta})=P\{g(Y) \leq x\}=P\left\{Y \leq g^{-1}(x)\right\}=F_{Y}\left\{g^{-1}(x) ; \boldsymbol{\theta}\right\}
$$

if functions are differentiable then

$$
f_{X}(x ; \theta)=f_{Y}\left\{g^{-1}(x) ; \theta\right\} \frac{d g^{-1}(x)}{d x}
$$

other transformation methods

### 1.5 Others

Theorem 1.2. Singular value decomposition
Suppose $A$ is an $n^{*} p$ matrix of rank $r$, where $r \leq \min (n, p)$. There exists orthogonal matrices $U_{p \times p}$ and $V_{n \times n}$ such that

$$
V^{\prime} A U=\left(\begin{array}{cc}
\Delta & 0 \\
0 & 0
\end{array}\right) \longrightarrow A=V D U^{\prime}, \text { where } D=\left(\begin{array}{cc}
\Delta & 0 \\
0 & 0
\end{array}\right)
$$

where $\Delta=\operatorname{diag}\left(\delta_{1}, \ldots, \delta_{r}\right)$ is an $r^{*} r$ diggonal matrix with $\delta_{1} \geq \delta_{2} \ldots \geq \delta_{r}>0$. The $\delta_{i}$ are called the singular values of $A$.

## 2 Likelihood construction and estimation

### 2.1 Introduction

Likelihood is the joint density of the observed data.
For example

- in models for censored or missing data, the likelihood is not the density of the so-called "complete" data that includes the censored or missing values. Rather it is the density of only those components of the data that are observed and used in the statistical analysis
- consider an iid sample $Y_{1}, \ldots, Y_{n}$, and a parametric transformation $h(y, \alpha)$ strictly increasing in $y$ for each $\alpha$. The model assumption is $h\left(Y_{1}, \alpha\right), \ldots, h\left(Y_{n}, \alpha\right)$ are iid with common density $f(y ; \boldsymbol{\theta})$ and distribution function $F(y ; \boldsymbol{\theta})$, where $\alpha$ and $\theta$ are both parameters in the model. To construct the likelihood, we need to find the likelihood of the observed data $Y$ rather than $h(Y)$. The distribution of $Y_{i}$ is (we can also use Jacobian method here)

$$
P\left\{Y_{i} \leq y\right\}=P\left\{h\left(Y_{i}, \alpha\right) \leq h(y, \alpha)\right\}=F\{h(y, \alpha) ; \theta\}
$$

taking derivative wrt $y$

$$
f_{Y}(y ; \boldsymbol{\theta}, \alpha)=f\{h(y, \alpha) ; \boldsymbol{\theta}\} \frac{\partial h(y, \alpha)}{\partial y}
$$

thus the likelihood is

$$
L(\boldsymbol{\theta}, \alpha ; \boldsymbol{Y})=\prod_{i=1}^{n} f\left\{h\left(Y_{i}, \alpha\right) ; \boldsymbol{\theta}\right\}\left\{\left.\frac{\partial h(y, \alpha)}{\partial y}\right|_{y=Y_{i}}\right\}
$$

a common mistake is missing the $\left\{\left.\frac{\partial h(y, \alpha)}{\partial y}\right|_{y=Y_{i}}\right\}$ part here
$\widehat{\boldsymbol{\theta}}_{\text {MLE }}$ of size $b \times 1$ is derived by maximizing

$$
\ell(\boldsymbol{\theta})=\log \{L(\boldsymbol{\theta} \mid \mathbf{Y})\}
$$

under differentiable assumptions we can also solve for

$$
\mathbf{S}(\boldsymbol{\theta})=\left\{\ell^{\prime}(\boldsymbol{\theta})\right\}^{T}
$$

in vector calculus notation

$$
\begin{array}{r}
\ell^{\prime}(\boldsymbol{\theta})=\frac{\partial \ell}{\partial \boldsymbol{\theta}}=\left(\frac{\partial \ell}{\partial \theta_{1}}, \cdots, \frac{\partial \ell}{\partial \theta_{b}}\right) \\
\ell^{\prime}(\boldsymbol{\theta})^{T}=\left(\frac{\partial \ell}{\partial \boldsymbol{\theta}}\right)^{T}=\frac{\partial \ell}{\partial \boldsymbol{\theta}^{T}}=\left(\begin{array}{c}
\frac{\partial \ell}{\partial \theta_{1}} \\
\vdots \\
\frac{\partial \ell}{\partial \theta_{b}}
\end{array}\right)
\end{array}
$$

and

$$
\ell^{\prime \prime}(\boldsymbol{\theta})=\frac{\partial}{\partial \boldsymbol{\theta}}\left(\frac{\partial \ell}{\partial \boldsymbol{\theta}^{T}}\right)=\frac{\partial^{2} \ell}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}}=\left(\begin{array}{ccc}
\frac{\partial^{2} \ell}{\partial^{2} \theta_{1}} & \cdots & \frac{\partial^{2} \ell}{\partial \theta_{b} \partial \theta_{1}} \\
\vdots & \ddots & \vdots \\
\frac{\partial^{2} \ell}{\partial \theta_{1} \partial \theta_{b}} & \cdots & \frac{\partial^{2} \ell}{\partial^{2} \theta_{b}}
\end{array}\right)=\left(\frac{\partial^{2} \ell}{\partial \theta_{j} \partial \theta_{i}}\right)_{i=1, \cdots, b, j=1, \cdots, b}
$$

### 2.2 Likelihood construction

- For discrete IID case, multinomial distribution is widely used cause

1. any discrete distribution with a finite support $\Longrightarrow$ multinomial model
2. any discrete distribution with an infinite support (e.g. Poisson) $\xrightarrow{\text { grouped }}$ multinomial model

$$
N \sim \operatorname{binomial}(n, p), f(N ; p)=\binom{n}{N} p^{N}(1-p)^{n-N}
$$

more generally

$$
\left(N_{1}, \cdots, N_{k}\right) \sim \operatorname{multinomial}\left(n ; p_{1}, \cdots, p_{k}\right), k \geq 2, f\left(N_{1}, \cdots, N_{k} ; p_{1}, \cdots, p_{k}\right)=\frac{n!}{N_{1}!\cdots N_{k}!} p_{1}^{N_{1}} \cdots p_{k}^{N_{k}}
$$ where $\sum_{i=1}^{k} p_{i}=1, \sum_{i=1}^{k} N_{i}=n$, and the model can be interpreted as tossing $n$ balls into $k$ urns. Also

$$
\begin{aligned}
E\left(N_{i}\right) & =n p_{i} \\
\operatorname{var}\left(N_{i}\right) & =n p_{i}\left(1-p_{i}\right) \\
\operatorname{cov}\left(N_{i}, N_{j}\right) & =-n p_{i} p_{j}, \quad i \neq j
\end{aligned}
$$

hint: to prove the covariance, we can write $N_{i}=\sum_{k}^{n} M_{k i}$ where $M_{k i}$ means toss k-th ball into i-th urn.

- Continuous IID case: skipped (nothing very important here)
- Connection Between Discrete and Continuous Likelihoods:2h-method

1. basic formula for continuous variable:

$$
f(y)=\lim _{h \rightarrow 0+} \frac{F(y+h)-F(y-h)}{2 h}=\lim _{h \rightarrow 0+} \frac{P(Y \in(y-h, y+h])}{2 h}
$$

2. bivariate data with both $X$ and $Y$ being continuous

$$
\begin{aligned}
f_{X, Y}(x, y) & =\lim _{h \rightarrow 0+}(2 h)^{-2}\left\{F_{X, Y}(x+h, y+h)-F_{X, Y}(x-h, y+h)\right. \\
& -F_{X, Y}(x+h, y-h)+F_{X, Y}(x-h, y-h) \\
& =\lim _{h \rightarrow 0+} \frac{P(X \in(x-h, x+h], Y \in(y-h, y+h])}{(2 h)^{2}}
\end{aligned}
$$

3. bivaraite data with $X$ discrete and $Y$ continuous
$f_{X, Y}(x, y)=\lim _{h \rightarrow 0+} \frac{P(X \in(x-h, x+h], Y \in(y-h, y+h])}{2 h}=\lim _{h \rightarrow 0+} \frac{P(X=x, Y \in(y-h, y+h])}{2 h}$
note: it is $2 h$ in the denominator
The likelihood based on the $2 h$-method can be summarized as follows
(a) for continuous case

$$
\begin{aligned}
L(\boldsymbol{\theta} \mid \mathbf{Y}) & =\prod_{i=1}^{n} f\left(Y_{i} ; \boldsymbol{\theta}\right) \\
& =\prod_{i=1}^{n} \lim _{h \rightarrow 0+} \frac{F\left(Y_{i}+h ; \boldsymbol{\theta}\right)-F\left(Y_{i}-h ; \boldsymbol{\theta}\right)}{2 h} \\
& =\lim _{h \rightarrow 0+}\left(\frac{1}{2 h}\right)^{n} \prod_{i=1}^{n}\left\{F\left(Y_{i}+h ; \boldsymbol{\theta}\right)-F\left(Y_{i}-h ; \boldsymbol{\theta}\right)\right\}
\end{aligned}
$$

(b) for discrete case

$$
\begin{aligned}
\lim _{h \rightarrow 0^{+}} \prod_{i=1}^{n}\left\{F\left(Y_{i}+h ; \boldsymbol{\theta}\right)-F\left(Y_{i}-h ; \boldsymbol{\theta}\right)\right\} & =\prod_{i=1}^{n}\left\{F\left(Y_{i}^{+} ; \boldsymbol{\theta}\right)-F\left(Y_{i}^{-} ; \boldsymbol{\theta}\right)\right\} \\
& =\prod_{i=1}^{n} f\left(Y_{i} ; \boldsymbol{\theta}\right) \\
& =L(\boldsymbol{\theta} \mid \boldsymbol{Y})
\end{aligned}
$$

(c) for combination of continuous and discrete random variables

$$
L(\boldsymbol{\theta} \mid \boldsymbol{Y})=\lim _{h \rightarrow 0^{+}}\left(\frac{1}{2 h}\right)^{m} \prod_{i=1}^{n}\left\{F_{i}\left(Y_{i}+h ; \boldsymbol{\theta}\right)-F_{i}\left(Y_{i}-h ; \boldsymbol{\theta}\right)\right\}
$$

where $1 \leq m \leq n$ depends on the number of continuous components in the data.

### 2.3 Proportional likelihoods

Likelihoods are equivalent if they are proportional and the constant of proportionality does not depend on unknown parameters.

- transformation of variables

$$
Y_{1}, \cdots, Y_{n} \sim \operatorname{iid} f_{Y}(y ; \boldsymbol{\theta})
$$

we'd like to focus on transformed data $X_{i}=g\left(Y_{i}\right), g(\cdot)$ known, increasing, and continuously differentiable

$$
\begin{aligned}
& f_{X}(x ; \boldsymbol{\theta})=f_{Y}\{h(x) ; \boldsymbol{\theta}\} h^{\prime}(x), \quad h(\cdot)=g^{-1}(\cdot) \\
& L(\boldsymbol{\theta} \mid \mathbf{X})=L(\boldsymbol{\theta} \mid \mathbf{Y}) \prod_{i=1}^{n} h^{\prime}\left(X_{i}\right)=L(\boldsymbol{\theta} \mid \mathbf{Y}) \prod_{i=1}^{n} \frac{1}{g^{\prime}\left(Y_{i}\right)}
\end{aligned}
$$

note that $\prod_{i=1}^{n} \frac{1}{g^{\prime}\left(Y_{i}\right)}$ is constant, thus the estimate of $\theta$ will be the same

- sufficient statistic e.g., $Y_{1}, \cdots, Y_{n} \stackrel{\text { iid }}{\sim} \operatorname{Bernoulli}(p)$

$$
L(p \mid \mathbf{Y})=\prod_{i=1}^{n} p^{Y_{i}}(1-p)^{1-Y_{i}}=p^{S}(1-p)^{n-S}
$$

for $S=\sum_{i=1}^{n} Y_{i}$, which is a sufficient statistic

$$
L(p \mid S)=\binom{n}{S} p^{S}(1-p)^{n-S} \propto L(p \mid \mathbf{Y})
$$

- different sampling plans
the above two examples give the same sampling plan and data for the two likelihoods. Here we give an example that different sampling plan leads to proportional likelihoods.

1. as in the sufficent statitic example

$$
Y_{1}, \cdots, Y_{n} \stackrel{\mathrm{iid}}{\sim} \operatorname{Bernoulli}(p)
$$

which leads to the first likelihood

$$
L_{1}(p \mid S)=\binom{12}{S} p^{S}(1-p)^{12-S}
$$

2. negative binomial: $Y_{i}$ 's are observed until 3 0's appear, which leads to likelihood

$$
L_{2}(p \mid S)=\binom{S+2}{S} p^{S}(1-p)^{3}
$$

(it's $S+2$ cause the last observation must be 0 ) the ratio is

$$
\frac{\binom{12}{S}}{\binom{S+2}{S}}(1-p)^{9-S}
$$

which depends on $p$ except for $S=9$. When $S=9$, the MLE estimate for $p$ will be the same but the inference like hypothesis testing (p-value) is different cause null distribution in the two plans are different

### 2.4 Empirical distribution function as an MLE

$$
\begin{aligned}
Y_{1}, \cdots, Y_{n} & \sim \operatorname{iid} F(y) \\
L_{h}(F \mid \mathbf{Y}) & =\prod_{i=1}^{n}\left\{F\left(Y_{i}+h\right)-F\left(Y_{i}-h\right)\right\} \\
& =\prod_{i=1}^{n} p_{i, h} \text { assuming no ties }
\end{aligned}
$$

for the likelihood here we ignored the $(2 h)^{-m}$ factor and the $p_{i, h}$ here needs to satisfy that $p_{i, h} \geq 0, \sum_{i=1}^{n} p_{i, h} \leq$ 1. Obviously, the likelihood increases as $p_{i, h}$ increases until it reaches that

$$
\sum_{i=1}^{n} p_{i, h}=1
$$

Thus use Lagrange multipliers, we solve for

$$
\log L_{h}(F \mid \mathbf{Y})+\lambda\left(\sum_{i=1}^{n} p_{i, h}-1\right)=\sum_{i=1}^{n} \log p_{i, h}+\lambda\left(\sum_{i=1}^{n} p_{i, h}-1\right)
$$

then

$$
\begin{aligned}
\frac{\partial g}{\partial p_{i, h}} & =\frac{1}{p_{i, h}}+\lambda=0, \quad i=1, \ldots, n \\
\frac{\partial g}{\partial \lambda} & =\sum_{i=1}^{n} p_{i, h}-1=0
\end{aligned}
$$

we can show that

$$
\hat{F}_{h} \xrightarrow{d} \widehat{F}_{\mathrm{EMP}}(y)=n^{-1} \sum_{i=1}^{n} I\left(Y_{i} \leq y\right)
$$

therefore we take the MLE of $F(y)$ as the empirical distribution function

$$
\widehat{F}_{\mathrm{MLE}}(y)=F_{\mathrm{EMP}}(y)=\frac{1}{n} \sum_{i=1}^{n} I\left(Y_{i} \leq t\right)
$$

### 2.5 Likelihoods for censored / truncated data

- fixed censoring (for censoring, we know exactly which individuals are censored) e.g. suppose we have underlying random variable $X \sim \mathrm{~N}\left(\mu, \sigma^{2}\right)$ and we generate $Y$ as follows

$$
\begin{aligned}
Y & =\left\{\begin{array}{cc}
0 & X \leq 0 \\
X & X>0
\end{array}\right. \\
F(y) & =\left\{\begin{array}{cc}
0 & y<0 \\
\Phi\{(y-\mu) / \sigma\} & y \geq 0
\end{array}\right.
\end{aligned}
$$

then the likelihood is

$$
\begin{aligned}
& Y_{1}, \cdots, Y_{n} \sim \text { iid } \\
& L(\mu, \sigma \mid \mathbf{Y})=\left\{\prod_{i: Y_{i}=0} \Phi(-\mu / \sigma)\right\}\left[\prod_{i: Y_{i}>0} \sigma^{-1} \phi\left\{\left(Y_{i}-\mu\right) / \sigma\right\}\right]
\end{aligned}
$$

- truncation (unlike censoring, here we are unaware of the truncated individuals) e.g. suppose income $X \sim F_{X}(x ; \theta)$ and sample $Y_{1}, \ldots, Y_{n}$ comes from incomes above $L_{0}$ (left truncated). Then

$$
\begin{aligned}
P(Y \leq y) & =P\left(X \leq y \mid X>L_{0}\right) I\left(y>L_{0}\right) \\
& =\frac{P\left(X \leq y, X>L_{0}\right)}{P\left(X>L_{0}\right)} I\left(y>L_{0}\right) \\
& =\frac{F_{X}(y ; \boldsymbol{\theta})-F_{X}\left(L_{0} ; \boldsymbol{\theta}\right)}{1-F_{X}\left(L_{0} ; \boldsymbol{\theta}\right)} I\left(y>L_{0}\right) \\
f_{Y}(y) & =\frac{f_{X}(y ; \boldsymbol{\theta})}{1-F_{X}\left(L_{0} ; \boldsymbol{\theta}\right)} I\left(y>L_{0}\right)
\end{aligned}
$$

- random censoring: we have underlying random variables $X$ and $R$ (e.g. $X$ : survival time, $R$ : censoring time) and $X \perp R$. What we observed is $Y=\min (X, R)$ and $\delta=I(X \leq R)$ then

$$
\begin{aligned}
f_{Y, \delta}(y, \delta=1)= & f_{X}(y ; \boldsymbol{\theta})\left\{1-F_{R}(y)\right\} \\
f_{Y, \delta}(y, \delta=0)= & \left\{1-F_{X}(y ; \boldsymbol{\theta})\right\} f_{R}(y) \\
L(\boldsymbol{\theta} \mid \mathbf{Y}, \boldsymbol{\delta})= & \prod_{i=1}^{n} f_{X}\left(Y_{i} ; \boldsymbol{\theta}\right)^{\delta_{i}}\left\{1-F_{R}\left(Y_{i}\right)\right\}^{\delta_{i}} \times\left\{1-F_{X}\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}^{1-\delta_{i}} f_{R}\left(Y_{i}\right)^{1-\delta_{i}} \\
\propto & \prod_{i=1}^{n} f_{X}\left(Y_{i} ; \boldsymbol{\theta}\right)^{\delta_{i}}\left\{1-F_{X}\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}^{1-\delta_{i}} \\
& \text { if } F_{R}(\cdot) \text { is noninformative of } \boldsymbol{\theta}
\end{aligned}
$$

proof: we can use the $2 h$-method to get the density in this case.

$$
\begin{aligned}
\lim _{h \rightarrow 0} \frac{P(Y \in(y-h, y+h], \delta=1)}{2 h} & =\lim _{h \rightarrow 0} \frac{P(x \in(y-h, y+h), R>x)}{2 h} \\
& =\lim _{h \rightarrow 0} \frac{\iint I(x \in(y-h, y+h], r>x) f_{X}(x) f_{R}(r) d x d r}{2 h} \\
& =\lim _{h \rightarrow 0} \frac{\iint\left[I(r>x) f_{R}(r)\right] d r I(x \in(y-h, y+h]) f_{X}(x) d x d r}{2 h} \\
& =\lim _{h \rightarrow 0} \frac{\iint\left[1-F_{R}(x)\right] I(x \in(y-h, y+h]) f_{X}(x) d x d r}{2 h} \\
& =\lim _{h \rightarrow 0} \frac{\int_{y-h}^{y+h}\left\{1-F_{R}(x)\right\} f_{x}(x) d x}{2 h}=\left[1-F_{R}(y)\right] f_{x}(y)
\end{aligned}
$$

### 2.6 Likelihoods for regression models

- normal linear model

$$
\begin{equation*}
Y_{i}=\mathbf{x}_{i}^{T} \boldsymbol{\beta}+e_{i}, \quad i=1, \cdots, n, \quad e_{1}, \cdots, e_{n} \stackrel{\mathrm{iid}}{\sim} N\left(0, \sigma^{2}\right) \tag{1}
\end{equation*}
$$

1. fixed design: $x_{1}, x_{2}, \ldots, x_{n}$ are known nonrandom p -vectors

$$
L\left(\boldsymbol{\beta}, \sigma \mid\left\{Y_{i}, \mathbf{x}_{i}\right\}_{i=1}^{n}\right)=\left(\frac{1}{\sqrt{2 \pi} \sigma}\right)^{n} \exp \left\{-\sum_{i=1}^{n} \frac{\left(Y_{i}-\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)^{2}}{2 \sigma^{2}}\right\}
$$

$$
\widehat{\boldsymbol{\beta}}_{\mathrm{MLE}}=\widehat{\boldsymbol{\beta}}_{\mathrm{LS}}=\left(\mathbf{X}^{T} \mathbf{X}\right)^{-1} \mathbf{X}^{T} \mathbf{Y}
$$

where $\mathbf{X}=\left(\mathbf{x}_{1}, \cdots, \mathbf{x}_{n}\right)^{T}, \mathbf{Y}=\left(Y_{1}, \cdots, Y_{n}\right)^{T}$ and

$$
\widehat{\sigma}_{\mathrm{MLE}}^{2}=n^{-1} \sum_{i=1}^{n} \widehat{e}_{i}^{2}=n^{-1} \sum_{i=1}^{n}\left(Y_{i}-\mathbf{x}_{i}^{T} \widehat{\boldsymbol{\beta}}_{\mathrm{MLE}}\right)^{2}
$$

whereas

$$
\widehat{\sigma}_{\text {unbiased }}=(n-p)^{-1} \sum_{i=1}^{n} \widehat{e}_{i}^{2}
$$

2. random design: $\left\{Y_{i}, \mathbf{x}_{i}\right\}_{i=1}^{n} \stackrel{\text { iid }}{\sim}(Y, \mathbf{x})$ and marginally $\mathbf{x} \sim f_{\mathbf{x}}(\mathbf{x} ; \boldsymbol{\tau})$ then

$$
L\left(\boldsymbol{\beta}, \sigma, \boldsymbol{\tau} \mid\left\{Y_{i}, \mathbf{x}_{i}\right\}_{i=1}^{n}\right)=\text { above } \times \prod_{i=1}^{n} f_{\mathbf{x}}\left(\mathbf{x}_{i} ; \boldsymbol{\tau}\right)=\left(\frac{1}{\sqrt{2 \pi} \sigma}\right)^{n} \exp \left\{-\sum_{i=1}^{n} \frac{\left(Y_{i}-\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)^{2}}{2 \sigma^{2}}\right\} \times \prod_{i=1}^{n} f_{\mathbf{x}}\left(\mathbf{x}_{i} ; \boldsymbol{\tau}\right)
$$

- additive errors nonlinear model: it is very similar to 1 but with the following the modification

$$
Y_{i}=g\left(\mathbf{x}_{i}, \beta\right)+e_{i}
$$

with function $g$ known. Then

$$
\widehat{\sigma}_{\mathrm{MLE}}^{2}=n^{-1} \sum_{i=1}^{n} \widehat{e}_{i}^{2}=n^{-1} \sum_{i=1}^{n}\left(Y_{i}-\mathbf{x}_{i}^{T} \widehat{\boldsymbol{\beta}}_{\mathrm{MLE}}\right)^{2}
$$

but the estimator for $\beta$ tend to have no closed form and need to solve numerically.

- generalized linear model

Exponential family is of the following form:

$$
f(y \mid \theta, \phi)=\exp \left\{\frac{y \theta-b(\theta)}{a(\phi)}+c(y, \phi)\right\}
$$

where $\theta$ is the canonical parameter, $\phi$ is the dispersion parameter, $b(\theta)$ is the cumulative function (different from cumulative generating function) and $a(), b(), c()$ are known functions.
To fit into the regression framework, we have

$$
\log f\left(y_{i} ; \theta_{i}, \phi\right)=\frac{y_{i} \theta_{i}-b\left(\theta_{i}\right)}{a_{i}(\phi)}+c\left(y_{i}, \phi\right)
$$

link function is $g($.$) such that$

$$
g\left(\mu_{i}\right)=\mathbf{x}_{i}^{T} \boldsymbol{\beta} \Longrightarrow \mu_{i}=g^{-1}\left(\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right) \Longrightarrow \theta_{i}=b^{\prime-1}\left\{g^{-1}\left(\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)\right\}
$$

a special case of the link function is the canonical link

$$
g\left(\mu_{i}\right)=\theta_{i}=\mathbf{x}_{i}^{T} \boldsymbol{\beta}
$$

in this setting the log-likelihood becomes

$$
\log L\left(\boldsymbol{\beta}, \phi \mid\left\{Y_{i}, \mathbf{x}_{i}\right\}_{i=1}^{n}\right)=\sum_{i=1}^{n}\left\{\frac{y_{i} \mathbf{x}_{i}^{T} \boldsymbol{\beta}-b\left(\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)}{a_{i}(\phi)}+c\left(y_{i}, \phi\right)\right\}
$$

If $Y \sim f(y ; \theta, \phi)$ and $\phi$ is fixed, then
Property 1. $E(Y)=b^{\prime}(\theta)$
Property 2. $\operatorname{Var}(Y)=b^{\prime \prime}(\theta) a(\phi)$

## Proof

We first prove that

$$
E\left[l^{\prime}(\theta)\right]=0
$$

and

$$
E\left[l^{\prime \prime}(\theta)\right]=-E\left\{\left[l^{\prime}(\theta)\right]^{2}\right\}
$$

$$
\begin{aligned}
& \ell(\theta)=\log f(y ; \theta, \phi)=\frac{y \theta-b(\theta)}{a(\phi)}+c(y, \phi) \\
E & {\left[\frac{\partial}{\partial \theta} \log f(y ; \theta, \phi)\right] } \\
= & \int\left\{\frac{\partial}{\partial \theta} \log f(y ; \theta, \phi)\right\} f(y ; \theta, \phi) d y \\
= & \int f^{\prime}(y ; \theta, \phi) d y \\
& \text { under some regularity conditions } \frac{\partial}{\partial \theta} \int f(y ; \theta, \phi) d y=0
\end{aligned}
$$

and

$$
\begin{aligned}
& E\left[-\frac{\partial^{2}}{\partial \theta^{2}} \log f(y ; \theta, \phi)\right] \\
= & -\int\left\{\frac{f^{\prime \prime}(y ; \theta, \phi) \cdot f(y ; \theta, \phi)-\left[f^{\prime}(y ; \theta, \phi)\right]^{2}}{f(y ; \theta, \phi)^{2}}\right\} \cdot f(y ; \theta, \phi) d y \\
= & -\int f^{\prime \prime}(y ; \theta, \phi) d y+\int\left[\frac{f^{\prime}(y ; \theta, \phi)}{f(y ; \theta, \phi)}\right]^{2} d y \\
= & -\frac{\partial^{2}}{\partial \theta^{2}} \int f(y ; \theta, \phi) d y+E\left[\left[(\log f(y ; \theta, \phi))^{\prime}\right]^{2}\right\} \\
= & E\left\{\left[(\log f(y ; \theta, \phi))^{\prime}\right]^{2}\right\}
\end{aligned}
$$

Then the score function for $\theta$ is

$$
U(\theta)=\frac{\partial \ell(\theta)}{\partial \theta}=\frac{Y-b^{\prime}(\theta)}{a(\phi)}
$$

Since $E(U(\theta))=0$, we have

$$
E\left(\frac{Y-b^{\prime}(\theta)}{a(\phi)}\right)=0 \Longrightarrow E(Y)=b^{\prime}(\theta)
$$

Since $E(U(\theta))=0$, we have

$$
\operatorname{Var}(U(\theta))=E\left(\frac{\partial \ell(\theta)}{\partial \theta}\right)^{2}=-E\left(\frac{\partial^{2} \ell(\theta)}{\partial \theta^{2}}\right) \Longrightarrow \frac{\operatorname{Var}(Y)}{a^{2}(\phi)}=\frac{b^{\prime \prime}(\theta)}{a(\phi)} \Longrightarrow \operatorname{Var}(Y)=b^{\prime \prime}(\theta) a(\phi)
$$

- generalized linear mixed model (GLMM)
in this model we extend the canonical-link GLM to accommodate random effects

$$
\theta_{i}=\mathbf{x}_{i}^{T} \boldsymbol{\beta}+\mathbf{z}_{i}^{T} \mathbf{U}, \quad \mathbf{U} \sim f_{\mathbf{U}}(\mathbf{u} ; \boldsymbol{\nu})
$$

then the likelihood is

$$
L\left(\boldsymbol{\beta}, \phi, \boldsymbol{\nu} \mid\left\{Y_{i}, \mathbf{x}_{i}, \mathbf{z}_{i}\right\}_{i=1}^{n}\right)=\prod_{i=1}^{n} \int f_{Y_{i} \mid \mathbf{U}}\left(Y_{i} \mid \mathbf{u}, \mathbf{x}_{i}, \mathbf{z}_{i}, \boldsymbol{\beta}, \phi\right) f_{\mathbf{U}}(\mathbf{u} ; \boldsymbol{\nu}) d \mathbf{u}
$$

- accelerated failure time model

$$
\log T=\mathbf{x}^{T} \beta+\sigma e, \quad T \perp R \mid \mathbf{x}
$$

usually $e \sim \mathrm{~N}(0,1)$
the observed random variable is $Y=\min (T, R), \delta=I(T \leq R)$ then the likelihood is

$$
L\left(\boldsymbol{\beta}, \sigma \mid\left\{Y_{i}, \delta_{i}, \mathbf{x}_{i}\right\}_{i=1}^{n}\right) \propto \prod_{i=1}^{n}\left\{\frac{1}{\sigma} f_{e}\left(\frac{\log Y_{i}-\mathbf{x}_{i}^{T} \boldsymbol{\beta}}{\sigma}\right)\right\}^{\delta_{i}}\left\{1-F_{e}\left(\frac{\log Y_{i}-\mathbf{x}_{i}^{T} \boldsymbol{\beta}}{\sigma}\right)\right\}^{1-\delta_{i}}
$$

### 2.7 Marginal and conditional likelihoods

Suppose $\mathbf{Y} \Longleftrightarrow(\mathbf{W}, \mathbf{V})$ and we have

$$
\left\{\begin{array}{l}
\theta_{1}: \text { parameter of interest }  \tag{2}\\
\theta_{2}: \text { nuisance parameter }
\end{array}\right.
$$

$$
\begin{align*}
f_{\mathbf{Y}}\left(\mathbf{y} ; \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right) & =f_{\mathbf{W}, \mathbf{V}}\left(\mathbf{w}, \mathbf{v} ; \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right)  \tag{3}\\
& =f_{\mathbf{W} \mid \mathbf{v}}\left(\mathbf{w} \mid \mathbf{V} ; \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right) f_{\mathbf{V}}\left(\mathbf{v} ; \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right)
\end{align*}
$$

1. if $f_{\mathbf{V}}\left(\mathbf{v} ; \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right)=f_{\mathbf{V}}\left(\mathbf{v} ; \boldsymbol{\theta}_{1}\right)$, it is a marginal likelihood
2. if $f_{\mathbf{W} \mid \mathbf{V}}\left(\mathbf{w} \mid \mathbf{v} ; \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right)=f_{\mathbf{W} \mid \mathbf{V}}\left(\mathbf{w} \mid \mathbf{v} ; \boldsymbol{\theta}_{1}\right)$, it is a conditional likelihood

Let's illustrate with an example
Neyman-Scott problem

$$
\begin{aligned}
Y_{i j} & \stackrel{i . i . d}{\sim} N\left(\mu_{i}, \sigma^{2}\right), i=1, \cdots, n, j=1,2 \\
\boldsymbol{\theta} & =\left(\sigma^{2}, \mu_{1}, \cdots, \mu_{n}\right)^{T} \operatorname{dim} n+1 \text { (increases with sample size) }
\end{aligned}
$$

then we have

$$
\begin{aligned}
\widehat{\sigma}_{\mathrm{MLE}}^{2} & =(2 n)^{-1} \sum_{i=1}^{n} \sum_{j=1}^{2}\left(Y_{i j}-\widehat{\mu}_{i, \mathrm{MLE}}\right)^{2} \\
& =\frac{n^{-1} \sum_{i=1}^{n}\left(Y_{i 1}-Y_{i 2}\right)^{2}}{4} \\
E\left(\widehat{\sigma}_{\mathrm{MLE}}^{2}\right) & =\frac{\sigma^{2}}{2} \Longrightarrow \widehat{\sigma}_{\mathrm{MLE}}^{2} \xrightarrow{p} \frac{\sigma^{2}}{2} \text { by WLLN }
\end{aligned}
$$

as we can see, consistency no longer exist for MLE in this case cause the number of parameters increases with sample size. Then how to achieve consistency?

$$
\begin{aligned}
V_{i} & =\frac{\left(Y_{i 1}-Y_{i 2}\right)}{\sqrt{2}} \sim N\left(0, \sigma^{2}\right) \\
W_{i} & =\frac{\left(Y_{i 1}+Y_{i 2}\right)}{\sqrt{2}} \sim N\left(\sqrt{2} \mu_{i}, \sigma^{2}\right)
\end{aligned}
$$

1. marginal likelihood

$$
\begin{aligned}
L(\sigma \mid \mathbf{V}) & =\prod_{i=1}^{n} \frac{1}{\sqrt{2 \pi} \sigma} \exp \left(-\frac{V_{i}^{2}}{2 \sigma^{2}}\right) \\
\widehat{\sigma}_{\mathrm{MMLE}}^{2} & =n^{-1} \sum_{i=1}^{n} V_{i}^{2}=\frac{n^{-1} \sum_{i=1}^{n}\left(Y_{i 1}-Y_{i 2}\right)^{2}}{2}
\end{aligned}
$$

in this case, only $V_{i}$ needs to be identified. However, $W_{i}$ may provide a sense of information loss
2. conditional likelihood
the key for this approach is to identify sufficient statistic for nuisance parameter under the assumption that the parameter of interest is known. In our case, $T_{i}=Y_{i 1}+Y_{i 2}$ is sufficient for $\mu_{i} \Longrightarrow\left(Y_{i 1}, Y_{i 2}\right) \mid T_{i}$ doesn't depend on $\mu_{i}$

$$
\begin{aligned}
Y_{i 1} \mid T_{i} & \sim N\left(\frac{T_{i}}{2}, \frac{\sigma^{2}}{2}\right) \\
\widehat{\sigma}_{\mathrm{CMLE}}^{2} & =2 n^{-1} \sum_{i=1}^{n}\left(Y_{i 1}-\frac{T_{i}}{2}\right)^{2} \\
& =\frac{n^{-1} \sum_{i=1}^{n}\left(Y_{i 1}-Y_{i 2}\right)^{2}}{2}
\end{aligned}
$$

Example1: Logistic regression measurement error model

$$
\begin{aligned}
& P(Y=1 \mid X)=\frac{\exp (\alpha+\beta X)}{1+\exp (\alpha+\beta X)} \\
& W=X+U, \quad U \sim N\left(0, \sigma_{U}^{2}\right) \text { with } \sigma_{U}^{2} \text { known }
\end{aligned}
$$

with $U \perp Y$ and $U \perp X, \sigma_{U}^{2}$ is known

- data: $\left(Y_{i}, W_{i}\right), i=1, \cdots, n$
- parameters of interest: $\alpha, \beta$
then

$$
\begin{aligned}
& f\left(Y_{i}, W_{i} \mid \alpha, \beta, X_{i}\right) \\
& \quad=\frac{\exp \left\{Y_{i}\left(\alpha+\beta X_{i}\right)\right\}}{1+\exp \left(\alpha+\beta X_{i}\right)} \frac{1}{\sqrt{2 \pi} \sigma_{U}} \exp \left\{-\frac{\left(W_{i}-X_{i}\right)^{2}}{2 \sigma_{U}^{2}}\right\} \\
& \quad=\frac{1}{\sqrt{2 \pi} \sigma_{U}} \frac{1}{1+\exp \left(\alpha+\beta X_{i}\right)} \times \exp \left\{\frac{\left(W_{i}+Y_{i} \sigma_{U}^{2} \beta\right) X_{i}}{\sigma_{U}^{2}}-\frac{X_{i}^{2}}{2 \sigma_{U}^{2}}+\alpha Y_{i}-\frac{W_{i}^{2}}{2 \sigma_{U}^{2}}\right\}
\end{aligned}
$$

$W_{i}+Y_{i} \sigma_{U U}^{2} \beta$ or equivalently $T_{i}=W_{i}+\left(Y_{i}-1 / 2\right) \sigma_{U}^{2} \beta$ is sufficient for $X_{i}$, assuming $\alpha$ and $\beta$ are known (Sufficient statistic depends on the parameter of interest). Then

$$
P\left(Y_{i}=1 \mid T_{i}\right)=\frac{\exp \left(\alpha+\beta T_{i}\right)}{1+\exp \left(\alpha+\beta T_{i}\right)}
$$

and the conditional likelihood score is based on

$$
\frac{\partial}{\partial(\alpha, \beta)} \log P\left(Y_{i} \mid T_{i}\right)=\left\{Y_{i}-\frac{\exp \left(\alpha+\beta T_{i}\right)}{1+\exp \left(\alpha+\beta T_{i}\right)}\right\}\binom{1}{T_{i}}
$$

and the score is

$$
\left.\sum_{i=1}^{n}\left\{Y_{i}-\frac{\exp \left(\alpha+\beta T_{i}\right)}{1+\exp \left(\alpha+\beta T_{i}\right)}\right\}\binom{1}{T_{i}}\right|_{T_{i}=W_{i}+\left(Y_{i}-1 / 2\right) \sigma_{U}^{2} \beta}=\mathbf{0}
$$

Note: when taking the differentiation, $T_{i}$ is treated as given but not as a function of $\beta$
Example2: exponential families with canonical parameter
$f\left(\mathbf{y} ; \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right)=h(\mathbf{y}) \exp \left\{\sum_{i} \theta_{1 i} W_{i}+\sum_{j} \theta_{2 j} V_{j}-A\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right)\right\} \Longrightarrow \mathbf{W} \mid \mathbf{V} \sim$ exponential family indexed by $\boldsymbol{\theta}_{1}$ only
Example3: conditional logistic regression

$$
\begin{aligned}
\operatorname{logit}\left\{P\left(Y_{i}=1\right)\right\} & =\mathbf{x}_{i}^{T} \boldsymbol{\beta} \\
P\left(Y_{i}=1\right) & =\frac{\exp \left(\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)}{1+\exp \left(\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)} \\
L(\boldsymbol{\beta} \mid \mathbf{Y}, \mathbf{X}) & =\prod_{i=1}^{n} \frac{\exp \left(\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)^{Y_{i}}}{1+\exp \left(\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)}=\frac{\exp \left(\sum_{i=1}^{n} Y_{i} \mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)}{\prod_{i=1}^{n}\left\{1+\exp \left(\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)\right\}}
\end{aligned}
$$

$T_{j}=\sum_{i=1}^{n} x_{i j} Y_{i}$ is sufficient for $\beta_{j}, j=1, \cdots, p$. If $\beta_{k}$ is the parameter of interest then let $W=T_{k}$ and $\mathbf{V}=\left(T_{1}, \cdots, T_{k-1}, T_{k+1}, \cdots, T_{p}\right)^{T}$

$$
P(W \mid \mathbf{V})=\frac{c\left(T_{1}, \cdots, T_{p}\right) \exp \left(\beta_{k} T_{k}\right)}{\sum_{u} c\left(T_{1}, \cdots, T_{k-1}, u, T_{k+1}, \cdots, T_{p}\right) \exp \left(\beta_{k} u\right)}
$$

where $P(W, V)=c\left(T_{1}, \cdots, T_{p}\right) \exp \left(\beta_{k} T_{k}\right)$ and $P(V)=\sum_{w} P(W, V)=\sum_{u} c\left(T_{1}, \cdots, T_{k-1}, u, T_{k+1}, \cdots, T_{p}\right) \exp \left(\beta_{k} u\right)$

### 2.8 MLE and information matrix

Recall that

$$
\begin{aligned}
\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}} & =\arg \max _{\boldsymbol{\theta} \in \boldsymbol{\theta}} L(\boldsymbol{\theta} \mid \mathbf{Y}) \\
\ell(\boldsymbol{\theta}) & =\log L(\boldsymbol{\theta} \mid \mathbf{Y})
\end{aligned}
$$

if $\ell(\boldsymbol{\theta})$ is continuously differentiable, then likelihood score

$$
\mathbf{S}(\boldsymbol{\theta})=\mathbf{S}(\mathbf{Y}, \boldsymbol{\theta})=\frac{\partial \ell(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^{T}}=\left\{\ell^{\prime}(\boldsymbol{\theta})\right\}^{T}
$$

exists and $\hat{\theta_{\text {MLE }}}$ satisfies $\mathbf{S}\left(\widehat{\boldsymbol{\theta}}_{\text {MLE }}\right)=\mathbf{0}$. Under most of the cases, $\theta_{\text {MLE }}$ is unique, or at least there is a principled strategy for choosing a single solution from among the possibly multiple values. But there are also special cases, for example

$$
Y_{1}, \cdots, Y_{n} \stackrel{\mathrm{iid}}{\sim} f(y ; \mu)= \begin{cases}e^{-(y-\mu)} & y>\mu \\ 0 & \text { otherwise }\end{cases}
$$

any $\mu<\min \left(Y_{i}\right)$ leads to likelihood to be 0 .

Property 1. invariant property:

$$
\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}} \text { for } \boldsymbol{\theta} \Longrightarrow \widehat{\boldsymbol{\tau}}_{\mathrm{MLE}}=g\left(\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right) \text { for } \tau=g(\boldsymbol{\theta})
$$

Property 2. asymptotic normality: under regularity conditions

$$
\sqrt{n}\left(\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}-\boldsymbol{\theta}\right) \xrightarrow{d} N\left\{\mathbf{0}, \mathbf{I}(\boldsymbol{\theta})^{-1}\right\}, \quad \text { as } n \rightarrow \infty
$$

where

$$
\begin{aligned}
\mathbf{I}(\boldsymbol{\theta}) & =E\left[\left\{\frac{\partial}{\partial \boldsymbol{\theta}^{T}} \log f(Y ; \boldsymbol{\theta})\right\}^{\otimes 2}\right] \\
& =\operatorname{var}\left\{\frac{\partial}{\partial \boldsymbol{\theta}^{T}} \log f(Y ; \boldsymbol{\theta})\right\} \\
& =-E\left\{\frac{\partial^{2}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \log f(Y ; \boldsymbol{\theta})\right\}
\end{aligned}
$$

How to estimate $I(\theta)$ ?
$I(\theta)$ has 2 components: $I($.$) and \theta$. It is a natural choice for using $\hat{\theta_{\text {MLE }}}$ to estimate $\theta$. But what about $I($.$) ?$

1. theoretical $I(.) \Longrightarrow \mathbf{I}\left(\widehat{\boldsymbol{\theta}}_{\text {MLE }}\right)$ is the MLE for $I(\theta)$ (invariant property)
2. estimated $I$ (.) version 1

$$
\begin{aligned}
\overline{\mathbf{I}}(\mathbf{Y}, \boldsymbol{\theta}) & =n^{-1} \sum_{i=1}^{n}\left\{-\frac{\partial}{\partial \boldsymbol{\theta}} \mathbf{s}\left(Y_{i}, \boldsymbol{\theta}\right)\right\} \\
& =n^{-1} \sum_{i=1}^{n}\left\{-\frac{\partial^{2}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \log f\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}
\end{aligned}
$$

3. estimated $I($.$) version 2$ (usually less efficient)

$$
\begin{aligned}
\overline{\mathbf{I}}^{*}(\mathbf{Y}, \boldsymbol{\theta}) & =n^{-1} \sum_{i=1}^{n} \mathbf{s}\left(Y_{i}, \boldsymbol{\theta}\right)^{\otimes 2} \\
& =n^{-1} \sum_{i=1}^{n}\left\{\frac{\partial}{\partial \boldsymbol{\theta}^{T}} \log f\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}^{\otimes 2}
\end{aligned}
$$

Note: when $Y_{i}$ are independent but not identically distributed. The $f(. ; \theta) \Longrightarrow f_{i}(. ; \theta)$
4. estimated $I($.$) version 1$ : ( $Y_{i}$ are independent but not identically distributed)

$$
\begin{aligned}
\overline{\boldsymbol{I}}(\boldsymbol{Y}, \boldsymbol{\theta}) & =\frac{1}{n} \sum_{i=1}^{n}\left\{-\frac{\partial}{\partial \boldsymbol{\theta}} \boldsymbol{s}_{i}\left(Y_{i}, \boldsymbol{\theta}\right)\right\} \\
& =\frac{1}{n} \sum_{i=1}^{n}\left\{-\frac{\partial^{2}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \log f_{i}\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}
\end{aligned}
$$

the according average expected information matrix (also called the average Fisher information matrix)
$\overline{\boldsymbol{I}}(\boldsymbol{\theta})=\mathrm{E}\{\overline{\boldsymbol{I}}(\boldsymbol{Y}, \boldsymbol{\theta})\}=\frac{1}{n} \sum_{i=1}^{n} \mathrm{E}\left\{-\frac{\partial^{2}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \log f_{i}\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}=\frac{1}{n} \sum_{i=1}^{n} \mathrm{E}\left[\left\{\frac{\partial}{\partial \boldsymbol{\theta}^{T}} \log f_{i}\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}\left\{\frac{\partial}{\partial \boldsymbol{\theta}} \log f_{i}\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}\right]$
in the i.i.d case $\bar{I}(\theta)=I(\theta)$
5. estimated $I($.$) version 2$ : ( $Y_{i}$ are independent but not identically distributed)

$$
\begin{aligned}
\overline{\boldsymbol{I}}^{*}(\boldsymbol{Y}, \boldsymbol{\theta}) & =\frac{1}{n} \sum_{i=1}^{n} s_{i}\left(Y_{i}, \boldsymbol{\theta}\right) s_{i}\left(Y_{i}, \boldsymbol{\theta}\right)^{T} \\
& =\frac{1}{n} \sum_{i=1}^{n}\left\{\frac{\partial}{\partial \boldsymbol{\theta}^{T}} \log f_{i}\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}\left\{\frac{\partial}{\partial \boldsymbol{\theta}} \log f_{i}\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}
\end{aligned}
$$

Actually, we discussed two types of information here

1. the total information:

$$
\mathbf{I}_{T}(\boldsymbol{\theta})=-E\left\{\frac{\partial^{2}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \log L(\boldsymbol{\theta} \mid \mathbf{Y})\right\}
$$

2. the observed total information

$$
\mathbf{I}_{T}(\mathbf{Y}, \boldsymbol{\theta})=-\frac{\partial^{2}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \log L(\boldsymbol{\theta} \mid \mathbf{Y})
$$

Example 1: $N\left(\mu, \sigma^{2}\right)$

$$
\begin{aligned}
\log f(y ; \mu, \sigma) & =\text { constant }-\log \sigma-\frac{1}{2 \sigma^{2}}(y-\mu)^{2} \\
\mathbf{I}(\mu, \sigma) & =\frac{1}{\sigma^{2}}\left(\begin{array}{ll}
1 & 0 \\
0 & 2
\end{array}\right)
\end{aligned}
$$

more specifically

$$
\ell^{\prime}(\theta)=\frac{\partial}{\partial \theta} \log f(y ; \mu, \sigma)=\left(-\frac{1}{\sigma}+\frac{(y-\mu)^{2}}{\sigma^{3}}, \frac{y-\mu}{\sigma^{2}}\right)^{T}
$$

then

$$
\operatorname{var}\left(\ell^{\prime}(\theta)\right)=\left[\begin{array}{ll}
\operatorname{var}\left(-\frac{1}{\sigma}+\frac{(y-\mu)^{2}}{\sigma^{3}}\right) & \operatorname{cov}\left(-\frac{1}{\sigma}+\frac{(y-\mu)^{2}}{\sigma^{3}}, \frac{y-\mu}{\sigma^{2}}\right) \\
\operatorname{cov}\left(-\frac{1}{\sigma}+\frac{(y-\mu)^{2}}{\sigma^{3}}, \frac{y-\mu}{\sigma^{2}}\right) & \operatorname{var}\left(\frac{y-\mu}{\sigma^{2}}\right)
\end{array}\right]
$$

then by the results of normal moments we get the answer

| Order | Non-central moment | Central moment |
| :--- | :--- | :--- |
| 1 | $\mu$ | 0 |
| 2 | $\mu^{2}+\sigma^{2}$ | $\sigma^{2}$ |
| 3 | $\mu^{3}+3 \mu \sigma^{2}$ | 0 |
| 4 | $\mu^{4}+6 \mu^{2} \sigma^{2}+3 \sigma^{4}$ | $3 \sigma^{4}$ |
| 5 | $\mu^{5}+10 \mu^{3} \sigma^{2}+15 \mu \sigma^{4}$ | 0 |
| 6 | $\mu^{6}+15 \mu^{4} \sigma^{2}+45 \mu^{2} \sigma^{4}+15 \sigma^{6}$ | $15 \sigma^{6}$ |
| 7 | $\mu^{7}+21 \mu^{5} \sigma^{2}+105 \mu^{3} \sigma^{4}+105 \mu \sigma^{6}$ | 0 |
| 8 | $\mu^{8}+28 \mu^{6} \sigma^{2}+210 \mu^{4} \sigma^{4}+420 \mu^{2} \sigma^{6}+105 \sigma^{8}$ | $105 \sigma^{8}$ |

Example 2: recall for the normal error regression models

$$
\begin{aligned}
L\left(\beta, \sigma \mid\left\{Y_{i}, \boldsymbol{x}_{i}\right\}_{i=1}^{n}\right) & =\prod_{i=1}^{n} \frac{1}{\sqrt{2 \pi} \sigma} \exp \left\{-\frac{\left(Y_{i}-\boldsymbol{x}_{i}^{T} \boldsymbol{\beta}\right)^{2}}{2 \sigma^{2}}\right\} \\
& =\left(\frac{1}{\sqrt{2 \pi} \sigma}\right)^{n} \exp \left\{-\sum_{i=1}^{n} \frac{\left(Y_{i}-\boldsymbol{x}_{i}^{T} \boldsymbol{\beta}\right)^{2}}{2 \sigma^{2}}\right\}
\end{aligned}
$$

then

$$
\ell=-\log \sqrt{2 \pi}-\log \sigma-\frac{\left(Y_{i}-x_{i}^{\top} \beta\right)^{2}}{2 \sigma^{2}}
$$

it follows that

$$
\ell^{\prime}=\left(\frac{\partial \ell}{\partial \sigma}, \frac{\partial \ell}{\partial \beta}\right)=\left(-\frac{1}{\sigma}+\frac{\left(Y_{i}-x_{i}^{T} \beta\right)^{2}}{\sigma^{3}}, \frac{x_{i}^{\top}\left(Y_{i}-x_{i}^{\top} \beta\right)}{\sigma^{2}}\right)
$$

and

$$
\ell^{\prime \prime}=\left[\begin{array}{ll}
\frac{1}{\sigma^{2}}-3 \frac{\left(Y_{i}-x_{i}^{\top} \beta\right)^{2}}{\sigma^{4}} & -2 \frac{x_{i}^{\top}\left(Y_{i}-x_{i}^{\top} \beta\right)}{\sigma^{3}} \\
-2 \frac{x_{i}^{\top}\left(Y_{i}-x_{i}^{\top} \beta\right)}{\sigma^{3}} & \frac{-x_{i}^{\prime} x_{i}}{\sigma^{2}}
\end{array}\right]
$$

taking expectation and organize we get

$$
\overline{\mathbf{I}}(\beta, \sigma)=\frac{1}{\sigma^{2}}\left(\begin{array}{cc}
\mathbf{X}^{T} \mathbf{X} / n & 0 \\
0 & 2
\end{array}\right)
$$

then invert $\boldsymbol{I}_{\mathrm{T}}(\boldsymbol{\beta}, \sigma)=n \overline{\boldsymbol{I}}(\boldsymbol{\beta}, \sigma)$ we get

$$
\begin{aligned}
& \operatorname{avar}(\hat{\beta})=\left(\boldsymbol{X}^{T} \boldsymbol{X}\right)^{-1} \sigma^{2} \\
& \operatorname{avar}(\hat{\sigma})=\frac{\sigma^{2}}{2 n}
\end{aligned}
$$

Similarly, for $Y_{i}=g\left(\mathbf{x}_{i}, \boldsymbol{\beta}\right)+e_{i}$

$$
\overline{\mathbf{I}}(\boldsymbol{\beta}, \sigma)=\frac{1}{\sigma^{2}}\left(\begin{array}{cc}
\mathbf{G}^{T} \mathbf{G} / n & 0 \\
0 & 2
\end{array}\right)
$$

where $[\mathbf{G}]_{i j}=\partial g\left(\mathbf{x}_{i}, \boldsymbol{\beta}\right) / \partial \beta_{j}$
Example 3: GLM with canonical link $\left(g\left(\mu_{i}\right)=\theta_{i}=\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)$

$$
\begin{aligned}
\mathbf{S}(\boldsymbol{\beta}, \phi) & =\sum_{i=1}^{n} \frac{\left\{Y_{i}-b^{\prime}\left(\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)\right\} \mathbf{x}_{i}}{a_{i}(\phi)}=\sum_{i=1}^{n} \frac{\left(Y_{i}-\mu_{i}\right) \mathbf{x}_{i}}{a_{i}(\phi)} \\
-\frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{S}(\boldsymbol{\beta}, \phi) & =\sum_{i=1}^{n} \frac{b^{\prime \prime}\left(\mathbf{x}_{i}^{T} \boldsymbol{\beta}\right)}{a_{i}(\phi)} \mathbf{x}_{i}^{\otimes 2}=\sum_{i=1}^{n} \frac{\operatorname{var}\left(Y_{i}\right)}{a_{i}(\phi)^{2}} \mathbf{x}_{i}^{\otimes 2} \\
& \equiv \mathbf{X}^{T} \mathbf{V} \mathbf{X}=n \overline{\mathbf{I}}(\mathbf{Y}, \boldsymbol{\beta})=n \overline{\mathbf{I}}(\boldsymbol{\beta})
\end{aligned}
$$

where $\boldsymbol{V}=\operatorname{diag}\left\{\operatorname{Var}\left(Y_{1}\right) / a_{1}(\phi)^{2}, \ldots, \operatorname{Var}\left(Y_{n}\right) / a_{n}(\phi)^{2}\right\}$. In this special case, average Fisher information is equal to average observed information (cause $\mathbf{X}^{T} \mathbf{V} \mathbf{X}$ is irrespective of $Y$ ). When $\Phi$ is unknown, the full average Fisher information is

$$
\overline{\mathbf{I}}(\beta, \phi)=\left(\begin{array}{cc}
\mathbf{X}^{T} \mathbf{V} \mathbf{X} / n & 0 \\
0 & \bar{I}(\phi)
\end{array}\right)
$$

thus estimating $\Phi$ does not increase the variability of $\beta$ estimation asymptotically in canonical link GLM models. Cause the diagonal of $\overline{\mathbf{I}}(\beta, \phi)$ is 0 and therefore we have

$$
\bar{I}^{-1}(\beta, \phi)=\left(\begin{array}{cc}
n\left[X^{\prime} V X\right]^{-1} & 0 \\
0 & \bar{I}^{-1}(\phi)
\end{array}\right)
$$

However, in general

$$
\left(\begin{array}{cc}
I_{11} & \mathbf{I}_{12} \\
& \mathbf{I}_{22}
\end{array}\right)^{-1}=\left(\begin{array}{cc}
\left(I_{11}-\mathbf{I}_{12} \mathbf{I}_{22}^{-1} \mathbf{I}_{21}\right)^{-1} & -I_{11}^{-1} \mathbf{I}_{12}\left(\mathbf{I}_{22}-\mathbf{I}_{21} I_{11}^{-1} \mathbf{I}_{12}\right)^{-1} \\
& \left(\mathbf{I}_{22}-\mathbf{I}_{21} I_{11}^{-1} \mathbf{I}_{12}\right)^{-1}
\end{array}\right)
$$

$\mathbf{I}_{12} \mathbf{I}_{22}^{-1} \mathbf{I}_{21} \geq 0 \Longrightarrow\left(I_{11}-\mathbf{I}_{12} \mathbf{I}_{22}^{-1} \mathbf{I}_{21}\right)^{-1} \geq I_{11}^{-1}$
which means that when adding parameters to a model, the diagonal elements of the inverse information matrix are always no less than the corresponding elements of the simpler model unless $I_{12}=0$. This is called variance inflation

### 2.8.1 Transformed and modeled parameters

If we have $f(y ; \theta)$ where $\theta$ : b-dimensional and therefore $I_{\theta}: b$. We also have $\beta$ :s-dimensional and $s \leq b, \theta=$ $g(\beta)$ then

$$
\begin{aligned}
\ell(\boldsymbol{\beta}) & =\log f\{y ; \mathbf{g}(\boldsymbol{\beta})\} \\
\frac{\partial}{\partial \boldsymbol{\beta}} \ell(\boldsymbol{\beta}) & =\left.\frac{\partial}{\partial \boldsymbol{\theta}} \log f(y ; \boldsymbol{\theta})\right|_{\boldsymbol{\theta}=\mathbf{g}(\boldsymbol{\beta})} \frac{\partial \mathbf{g}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}} \\
\mathbf{I}_{\boldsymbol{\beta}}(\boldsymbol{\beta}) & =\left\{\frac{\partial \mathbf{g}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}}\right\}^{T} \mathbf{I}_{\boldsymbol{\theta}}\{\mathbf{g}(\boldsymbol{\beta})\} \frac{\partial \mathbf{g}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}}
\end{aligned}
$$

Example 1: in the normal setting, let $\boldsymbol{\theta}=(\mu, \sigma)^{T}$ and $\boldsymbol{\beta}=\left(\mu, \sigma^{2}\right)^{T}$ then

$$
\begin{aligned}
\frac{\partial \mathbf{g}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}} & =\operatorname{diag}\{1,1 /(2 \sigma)\} \\
\mathbf{I}_{\boldsymbol{\beta}}(\boldsymbol{\beta}) & =\left(\begin{array}{cc}
1 & 0 \\
0 & 1 /(2 \sigma)
\end{array}\right) \frac{1}{\sigma^{2}}\left(\begin{array}{ll}
1 & 0 \\
0 & 2
\end{array}\right)\left(\begin{array}{cc}
1 & 0 \\
0 & 1 /(2 \sigma)
\end{array}\right) \\
& =\frac{1}{\sigma^{2}}\left(\begin{array}{cc}
1 & 0 \\
0 & 1 /\left(2 \sigma^{2}\right)
\end{array}\right)
\end{aligned}
$$

### 2.9 Methods for maximizing the likelihood

1. Profile likelihood: maximize in a sequential way and therefore achieve dimension reduction

$$
\max _{\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}} L\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right)=\max _{\boldsymbol{\theta}_{1}} L\left\{\boldsymbol{\theta}_{1}, \widehat{\boldsymbol{\theta}}_{2}\left(\boldsymbol{\theta}_{1}\right)\right\}, \quad \widehat{\boldsymbol{\theta}}_{2}\left(\boldsymbol{\theta}_{1}\right)=\arg \max _{\boldsymbol{\theta}_{2} \mid \boldsymbol{\theta}_{1}} L\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}\right)
$$

Example 1: $Y_{1}, \cdots, Y_{n} \stackrel{\text { iid }}{\sim} \operatorname{Gamma}(\alpha, \beta)$

$$
\begin{aligned}
f_{Y}(y ; \alpha, \beta)= & \frac{1}{\Gamma(\alpha) \beta^{\alpha}} y^{\alpha-1} e^{-y / \beta} \\
\ell(\alpha, \beta)= & -n \log \Gamma(\alpha)-n \alpha \log \beta+(\alpha-1) \sum_{i} \log Y_{i}-\frac{\sum_{i} Y_{i}}{\beta} \\
\frac{\partial}{\partial \beta} \ell(\alpha, \beta)= & -\frac{n \alpha}{\beta}+\frac{\sum_{i} Y_{i}}{\beta^{2}} \Longrightarrow \widehat{\beta}(\alpha)=\bar{Y} / \alpha \\
\ell\{\alpha, \widehat{\beta}(\alpha)\}= & -n \log \Gamma(\alpha)-n \alpha(\log \bar{Y}-\log \alpha) \\
& +(\alpha-1) \sum_{i} \log Y_{i}-n \alpha
\end{aligned}
$$

2. Newton methods

$$
\begin{aligned}
\mathbf{0}=\mathbf{S}(\boldsymbol{\theta}) & \approx \mathbf{S}\left(\boldsymbol{\theta}^{(\nu)}\right)+\left\{\left.\frac{\partial}{\partial \boldsymbol{\theta}} \mathbf{S}(\boldsymbol{\theta})\right|_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(\nu)}}\right\}\left(\boldsymbol{\theta}-\boldsymbol{\theta}^{(\nu)}\right) \\
& =\mathbf{S}\left(\boldsymbol{\theta}^{(\nu)}\right)-\mathbf{I}_{T}\left(\mathbf{Y}, \boldsymbol{\theta}^{(\nu)}\right)\left(\boldsymbol{\theta}-\boldsymbol{\theta}^{(\nu)}\right) \\
\boldsymbol{\theta}^{(\nu+1)} & =\boldsymbol{\theta}^{(\nu)}+\mathbf{I}_{T}\left(\mathbf{Y}, \boldsymbol{\theta}^{(\nu)}\right)^{-1} \mathbf{S}\left(\boldsymbol{\theta}^{(\nu)}\right)
\end{aligned}
$$

(a) start with initial value $\boldsymbol{\theta}^{(0)}$
(b) update from current $\boldsymbol{\theta}^{(\nu)}$ to obtain $\boldsymbol{\theta}^{(\nu+1)}$
(c) stop if $\left\|\mathbf{S}\left(\boldsymbol{\theta}^{(\nu+1)}\right)\right\|$ or $\left\|\boldsymbol{\theta}^{(\nu+1)}-\boldsymbol{\theta}^{(\nu)}\right\|$ is sufficiently small. Otherwise keep updating


Figure 1: Newton-Raphson method (source)
If the first derivative is not well behaved in the neighborhood of a particular root, the method may overshoot. To prevent overshooting: At each step, ensure improvement of $\boldsymbol{\theta}^{(\nu+1)}$ over $\boldsymbol{\theta}^{(\nu)}$ by repeatedly halving the step size.
Under certain conditions, Newton methods have local quadratic convergence, which means that for some $c>0$

$$
\left\|\boldsymbol{\theta}^{(\nu+1)}-\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right\| \leq c\left\|\boldsymbol{\theta}^{(\nu)}-\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right\|^{2}
$$

proof: (although I am not totally agree with this proof)

$$
\begin{aligned}
& 0=S\left(\widehat{\theta}_{\mathrm{MLE}}\right)=S\left(\theta^{(\nu)}\right)-I_{T}\left(\mathbf{Y}, \theta^{(\nu)}\right)\left(\widehat{\theta}_{\mathrm{MLE}}-\theta^{(\nu)}\right)+\frac{1}{2} S^{\prime \prime}\left(\vartheta^{(\nu)}\right)\left(\widehat{\theta}_{\mathrm{MLE}}-\theta^{(\nu)}\right)^{2} \\
& \widehat{\theta}_{\mathrm{MLE}}-\theta^{(\nu+1)}=\frac{1}{2} \frac{S^{\prime \prime}\left(\vartheta^{(\nu)}\right)}{I_{T}\left(\mathbf{Y}, \theta^{(\nu)}\right)}\left(\widehat{\theta}_{\mathrm{MLE}}-\theta^{(\nu)}\right)^{2}
\end{aligned}
$$

Thus, the local quadratic convergence holds under the following conditions:

- $I_{T}(\mathbf{Y}, \theta) \neq 0$ in a neighborhood of $\hat{\theta}_{\text {MLE }}$
- $S^{\prime \prime}(\theta)$ is bounded
- $\theta^{(\nu)}$ is sufficiently close to $\hat{\theta}_{\text {MLE }}$ (so that we can eliminate $S\left(\theta^{(\nu)}\right)$ )

3. Fisher scoring: $\mathbf{I}_{T}\left(\mathbf{Y}, \boldsymbol{\theta}^{(\nu)}\right)$ in the above method replaced by its expectation $\mathbf{I}_{T}\left(\boldsymbol{\theta}^{(\nu)}\right)$
4. One-step estimator: typically $\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}-\boldsymbol{\theta}=O_{p}\left(n^{-1 / 2}\right)$ (from asymptotic normality). If one starts with $\boldsymbol{\theta}^{(0)}$ such that $\boldsymbol{\theta}^{(0)}-\boldsymbol{\theta}=O_{p}\left(n^{-1 / 2}\right)$, then

$$
\boldsymbol{\theta}^{(1)}-\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}=O_{p}\left(n^{-1}\right)
$$

under regularity conditions (by the local quadratic convergence property).
5. EM algorithm (nice reading): view observed data $Y$ as incomplete, with $Z$ missing. Write log joint likelihood of "complete" data as $\ell_{C}(\boldsymbol{\theta} \mid \mathbf{Y}, \mathbf{Z})$
(a) E step: calculate

$$
\begin{aligned}
Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(\nu)}, \mathbf{Y}\right) & =E_{\boldsymbol{\theta}^{(\nu)}}\left\{\ell_{C}(\boldsymbol{\theta} \mid \mathbf{Y}, \mathbf{Z}) \mid \mathbf{Y}\right\} \\
& =\int \ell_{C}(\boldsymbol{\theta} \mid \mathbf{Y}, \mathbf{z}) f_{\mathbf{Z} \mid \mathbf{Y}}\left(\mathbf{z} \mid \mathbf{Y}, \boldsymbol{\theta}^{(\nu)}\right) d \mathbf{z}
\end{aligned}
$$

(b) maximize $Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(\nu)}, \mathbf{Y}\right)$ with respect to $\boldsymbol{\theta}$ to obtain $\boldsymbol{\theta}^{\nu+1}$

Example 1: 2-component mixtures

$$
Y_{1}, \cdots, Y_{n} \stackrel{\mathrm{iid}}{\sim} f(y ; \boldsymbol{\theta})=p f_{1}\left(y ; \mu_{1}, \sigma_{1}\right)+(1-p) f_{2}\left(y ; \mu_{2}, \sigma_{2}\right)
$$

where $f_{1}, f_{2}$ are normal densities

$$
\begin{aligned}
\boldsymbol{\theta} & =\left(\mu_{1}, \sigma_{1}, \mu_{2}, \sigma_{2}, p\right)^{T} \\
\ell(\boldsymbol{\theta}) & =\sum_{i=1}^{n} \log \left\{p f_{1}\left(Y_{i} ; \mu_{1}, \sigma_{1}\right)+(1-p) f_{2}\left(Y_{i} ; \mu_{2}, \sigma_{2}\right)\right\}
\end{aligned}
$$

it's hard to maximize directly. Therefore, turn to the EM algorithm

$$
\begin{aligned}
Y_{i} & =Z_{i} X_{1 i}+\left(1-Z_{i}\right) X_{2 i} \\
X_{11}, \cdots, X_{1 n} & \stackrel{\text { iid }}{\sim} N\left(\mu_{1}, \sigma_{1}^{2}\right) \\
X_{21}, \cdots, X_{2 n} & \stackrel{\text { iid }}{\sim} N\left(\mu_{2}, \sigma_{2}^{2}\right) \\
Z_{1}, \cdots, Z_{n} & \stackrel{\text { iid }}{\sim} \operatorname{Bernoulli}(p)
\end{aligned}
$$

join likelihood of complete data ( $Y, Z$ )

$$
\begin{aligned}
L_{C}(\boldsymbol{\theta} \mid \mathbf{Y}, \mathbf{Z}) & =\prod_{i=1}^{n}\left\{p f_{1}\left(Y_{i} ; \mu_{1}, \sigma_{1}\right)\right\}^{Z_{i}} \times\left\{(1-p) f_{2}\left(Y_{i} ; \mu_{2}, \sigma_{2}\right)\right\}^{\left(1-Z_{i}\right)} \\
\ell_{C}(\boldsymbol{\theta} \mid \mathbf{Y}, \mathbf{Z}) & =\sum_{i=1}^{n}\left\{Z_{i} \log f_{1}\left(Y_{i} ; \mu_{1}, \sigma_{1}\right)+\left(1-Z_{i}\right) \log f_{2}\left(Y_{i} ; \mu_{2}, \sigma_{2}\right)+Z_{i} \log p+\left(1-Z_{i}\right) \log (1-p)\right\}
\end{aligned}
$$

cause

$$
\begin{aligned}
& p(y, z=1)=p(y \mid z=1) p(z=1)=p f_{1} \\
& p(y, z=0)=p(y \mid z=0) p(z=0)=(1-p) f_{2}
\end{aligned}
$$

- E-step

$$
\begin{aligned}
Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(\nu)}, \mathbf{Y}\right) & =E_{\boldsymbol{\theta}^{(\nu)}}\left\{\ell_{C}(\boldsymbol{\theta} \mid \mathbf{Y}, \mathbf{Z}) \mid \mathbf{Y}\right\} \\
& =\sum_{i=1}^{n}\left\{w_{i}^{(\nu)} \log f_{1}\left(Y_{i} ; \mu_{1}, \sigma_{1}\right)+\left(1-w_{i}^{(\nu)}\right) \log f_{2}\left(Y_{i} ; \mu_{2}, \sigma_{2}\right)+w_{i}^{(\nu)} \log p+\left(1-w_{i}^{(\nu)}\right) \log (1-p)\right\}
\end{aligned}
$$

where

$$
\begin{aligned}
w_{i}^{(\nu)} & =E_{\boldsymbol{\theta}^{(\nu)}}\left(Z_{i} \mid Y_{i}\right) \\
& =\frac{p^{(\nu)} f_{1}\left(Y_{i} ; \mu_{1}^{(\nu)}, \sigma_{1}^{(\nu)}\right)}{p^{(\nu)} f_{1}\left(Y_{i} ; \mu_{1}^{(\nu)}, \sigma_{1}^{(\nu)}\right)+\left(1-p^{(\nu)}\right) f_{2}\left(Y_{i} ; \mu_{2}^{(\nu)}, \sigma_{2}^{(\nu)}\right)}
\end{aligned}
$$

substituting normal densities

$$
\begin{aligned}
Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(\nu)}, \mathbf{Y}\right)= & \sum_{i=1}^{n}\left[w_{i}^{(\nu)}\left\{-\log \sigma_{1}-\frac{\left(Y_{i}-\mu_{1}\right)^{2}}{2 \sigma_{1}^{2}}\right\}\right. \\
+ & \left(1-w_{i}^{(\nu)}\right)\left\{-\log \sigma_{2}-\frac{\left(Y_{i}-\mu_{2}\right)^{2}}{2 \sigma_{2}^{2}}\right\} \\
& \left.+w_{i}^{(\nu)} \log p+\left(1-w_{i}^{(\nu)}\right) \log (1-p)\right] \\
+ & \text { const. }
\end{aligned}
$$

- M-step: maximizing $Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(\nu)}, \mathbf{Y}\right)$ is simple. When the "complete" data likelihood has the form of an exponential family, M-step is straightforward.

Example 2: right-censored data
Let $\left(Y_{1}, \delta_{1}\right), \cdots,\left(Y_{n}, \delta_{n}\right)$ be i.i.d where $Y_{i}=\min \left(X_{i}, R_{i}\right)$ and $\delta_{i}=I\left(X_{i} \leq R_{i}\right)$.

$$
\begin{aligned}
X_{i} & \sim f(x ; \boldsymbol{\theta})=\sigma^{-1} \exp (-x / \sigma) \\
L(\boldsymbol{\theta} \mid \mathbf{Y}, \boldsymbol{\delta}) & \propto \prod_{i=1}^{n} f\left(Y_{i} ; \boldsymbol{\theta}\right)^{\delta_{i}}\left\{1-F\left(Y_{i} ; \boldsymbol{\theta}\right)\right\}^{1-\delta_{i}} \\
& =\sigma^{-n_{u}} \exp \left(-\sum_{i=1}^{n} Y_{i} / \sigma\right)
\end{aligned}
$$

cause $F\left(Y_{i} ; \theta\right)=1-\exp (-x / \sigma)$, where $n_{\mu}$ is the number of uncensored observations. The likelihood is easy to maximize. Nonetheless, let $Z_{i}=X_{i}$

$$
\begin{aligned}
\ell(\boldsymbol{\theta} \mid \mathbf{Y}, \boldsymbol{\delta}, \mathbf{X}) & =\sum_{i=1}^{n} \log f\left(X_{i} ; \boldsymbol{\theta}\right) \\
& =\sum_{i=1}^{n_{u}} \log f\left(X_{i} ; \boldsymbol{\theta}\right)+\sum_{i=n_{u}+1}^{n} \log f\left(X_{i} ; \boldsymbol{\theta}\right)
\end{aligned}
$$

- E-step

$$
\begin{aligned}
Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(\nu)}, \mathbf{Y}, \boldsymbol{\delta}\right) & =\sum_{i=1}^{n_{u}} \log f\left(Y_{i} ; \boldsymbol{\theta}\right)+\sum_{i=n_{u}+1}^{n} E_{\boldsymbol{\theta}^{(\nu)}}\left\{\log f\left(X_{i} ; \boldsymbol{\theta}\right) \mid Y_{i}, \delta_{i}\right\} \\
& =-n \log \sigma-\sigma^{-1} \sum_{i=1}^{n_{u}} Y_{i}-\sigma^{-1} \sum_{i=n_{u}+1}^{n} E_{\sigma^{(\nu)}}\left(X_{i} \mid X_{i}>Y_{i}\right) \\
& =-n \log \sigma-\sigma^{-1} \sum_{i=1}^{n} Y_{i}-\sigma^{-1}\left(n-n_{u}\right) \sigma^{(\nu)}
\end{aligned}
$$

things are simplied here by assuming $R_{i}$ is not random

- M-step

$$
\sigma^{(\nu+1)}=n^{-1}\left\{\sum_{i=1}^{n} Y_{i}+\left(n-n_{u}\right) \sigma^{(\nu)}\right\}
$$

as $\nu \rightarrow \infty$, we expect $\sigma^{(\nu+1)} \rightarrow \widehat{\sigma}_{\text {MLE }}$ and $\sigma^{(\nu)} \rightarrow \widehat{\sigma}_{\text {MLE }}$ thus

$$
\widehat{\sigma}_{\mathrm{MLE}}=n^{-1}\left\{\sum_{i=1}^{n} Y_{i}+\left(n-n_{u}\right) \widehat{\sigma}_{\mathrm{MLE}}\right\} \Longrightarrow \widehat{\sigma}_{\mathrm{MLE}}=\sum_{i=1}^{n} Y_{i} / n_{u}
$$

### 2.9.1 Why does EM work?

- Jensen's inequality: for a random variable $X$ and convex function $\psi(\cdot), \psi\{E(X)\} \leq E\{\psi(X)\}$
- for $\mathbf{Y} \sim f\left(\mathbf{y} ; \boldsymbol{\theta}_{0}\right), E_{\boldsymbol{\theta}_{0}}\{\log f(\mathbf{Y} ; \boldsymbol{\theta})\}$ is maximized at $\theta=\theta_{0}$ proof: $\psi(x)=-\log (x)$ is a convex function for $x \in(0, \infty)$. Therefore, by Jensen's inequality

$$
-\log \left[E_{\boldsymbol{\theta}_{0}}\left\{\frac{f(\mathbf{Y} ; \boldsymbol{\theta})}{f\left(\mathbf{Y} ; \boldsymbol{\theta}_{0}\right)}\right\}\right] \leq-E_{\boldsymbol{\theta}_{0}}\left[\log \left\{\frac{f(\mathbf{Y} ; \boldsymbol{\theta})}{f\left(\mathbf{Y} ; \boldsymbol{\theta}_{0}\right)}\right\}\right]
$$

note that

$$
E_{\theta_{0}}\left\{\frac{f(Y ; \theta)}{f\left(Y ; \theta_{0}\right)}\right\}=\int \frac{f(Y ; \theta)}{f\left(Y ; \theta_{0}\right)} f\left(Y ; \theta_{0}\right) d y=1
$$

therefore the left side is 0

$$
\begin{aligned}
& -E_{\theta_{0}}\left[\log \left\{\frac{f(Y ; \theta)}{f\left(Y ; \theta_{0}\right)}\right\}\right] \geqslant 0 \Rightarrow E_{\theta_{0}}\left[\log \left\{\frac{f\left(Y ; \theta_{0}\right)}{f(Y ; \theta)}\right\}\right] \leq 0 \\
& \Rightarrow E_{\theta_{0}}[\log f(Y ; \theta)] \leqslant E_{\theta_{0}}\left[\log f\left(Y ; \theta_{0}\right)\right] \Rightarrow E_{\boldsymbol{\theta}_{0}}\{\log f(\mathbf{Y} ; \boldsymbol{\theta})\} \text { is maximized at } \boldsymbol{\theta}=\boldsymbol{\theta}_{0}
\end{aligned}
$$

- now go back to the EM algorithm

$$
\ell(\boldsymbol{\theta} \mid \mathbf{Y}) \equiv \log f(\mathbf{Y} \mid \boldsymbol{\theta})=\log f(\mathbf{Y}, \mathbf{Z} \mid \boldsymbol{\theta})-\log f(\mathbf{Z} \mid \mathbf{Y}, \boldsymbol{\theta})
$$

cause $f(y, z ; \theta)=f(z \mid y ; \theta) f(y ; \theta)$. Taking expectation on both sides treating $Z$ as a random variable with density $f\left(\mathbf{Z} \mid \mathbf{Y}, \boldsymbol{\theta}^{(\nu)}\right)$

$$
\ell(\boldsymbol{\theta} \mid \mathbf{Y})=Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(\nu)}, \mathbf{Y}\right)-E_{\boldsymbol{\theta}^{(\nu)}}\{\log f(\mathbf{Z} \mid \mathbf{Y}, \boldsymbol{\theta}) \mid \mathbf{Y}\}
$$

where based on the previous step $E_{\boldsymbol{\theta}^{(\nu)}}\{\log f(\mathbf{Z} \mid \mathbf{Y}, \boldsymbol{\theta}) \mid \mathbf{Y}\}$ is maximized at $\boldsymbol{\theta}^{(\nu)}$. In performing the EM algorithm, we ensure that $Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(\nu)}, \mathbf{Y}\right) \geq Q\left(\boldsymbol{\theta}^{(\nu)}, \boldsymbol{\theta}^{(\nu)}, \mathbf{Y}\right)$ by maximizing $Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(\nu)}, \mathbf{Y}\right)$. Therefore $\ell\left(\boldsymbol{\theta}^{(\nu+1)} \mid \mathbf{Y}\right) \geq \ell\left(\boldsymbol{\theta}^{(\nu)} \mid \mathbf{Y}\right)$.

### 2.9.2 Calculating observed info matrix after EM

The general way is

$$
\begin{aligned}
\mathbf{S}_{\mathbf{Y}}(\mathbf{Y}, \boldsymbol{\theta}) & =\frac{\partial}{\partial \boldsymbol{\theta}^{T}} \log f_{\mathbf{Y}}(\mathbf{Y} ; \boldsymbol{\theta})=\frac{f_{\mathbf{Y}}^{\prime}(\mathbf{Y} ; \boldsymbol{\theta})^{T}}{f_{\mathbf{Y}}(\mathbf{Y} ; \boldsymbol{\theta})} \\
\mathbf{I}_{\mathbf{Y}}(\mathbf{Y}, \boldsymbol{\theta}) & =-\frac{\partial^{2}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \log f_{\mathbf{Y}}(\mathbf{Y} ; \boldsymbol{\theta}) \\
& =-\frac{f_{\mathbf{Y}}^{\prime \prime}(\mathbf{Y} ; \boldsymbol{\theta})}{f_{\mathbf{Y}}(\mathbf{Y} ; \boldsymbol{\theta})}+\mathbf{S}_{\mathbf{Y}}(\mathbf{Y}, \boldsymbol{\theta})^{\otimes 2}
\end{aligned}
$$

Actually, there is a better way than direct computation. Note for the complete data $f_{\mathbf{Y}, \mathbf{Z}}(\mathbf{y}, \mathbf{z} ; \boldsymbol{\theta})$

$$
\begin{align*}
\mathbf{S}_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{Z}, \boldsymbol{\theta}) & =\frac{f_{\mathbf{Y}, \mathbf{Z}}^{\prime}(\mathbf{Y}, \mathbf{Z} ; \boldsymbol{\theta})^{T}}{f_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{Z} ; \boldsymbol{\theta})}  \tag{4}\\
\mathbf{I}_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{Z}, \boldsymbol{\theta}) & =-\frac{f_{\mathbf{Y}, \mathbf{Z}}^{\prime \prime}(\mathbf{Y}, \mathbf{Z} ; \boldsymbol{\theta})}{f_{\mathbf{Y}, \mathbf{Z}}(\mathbf{y} ; \boldsymbol{\theta})}+\mathbf{S}_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{Z}, \boldsymbol{\theta})^{\otimes 2}
\end{align*}
$$

given that the complete density is related to the density we are interested in, we have and based on 4 we have

$$
\begin{aligned}
f_{\mathbf{Y}}(\mathbf{y} ; \boldsymbol{\theta})= & \int f_{\mathbf{Y}, \mathbf{Z}}(\mathbf{y}, \mathbf{z} ; \boldsymbol{\theta}) d \mathbf{z} \\
\mathbf{S}_{\mathbf{Y}}(\mathbf{Y}, \boldsymbol{\theta})= & \frac{\partial}{\partial \boldsymbol{\theta}^{T}} \log \int f_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{z} ; \boldsymbol{\theta}) d \mathbf{z} \\
= & \frac{\int f_{\mathbf{Y}, \mathbf{Z}}^{\prime}(\mathbf{Y}, \mathbf{z} ; \boldsymbol{\theta})^{T} d \mathbf{z}}{\int f_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{z} ; \boldsymbol{\theta}) d \mathbf{z}} \\
= & \frac{\int \mathbf{S}_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{z}, \boldsymbol{\theta}) f_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{z} ; \boldsymbol{\theta}) d \mathbf{z}}{\int f_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{z} ; \boldsymbol{\theta}) d \mathbf{z}} \\
= & E_{\boldsymbol{\theta}}\left\{\mathbf{S}_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{Z}, \boldsymbol{\theta}) \mid \mathbf{Y}\right\} \\
\mathbf{I}_{\mathbf{Y}}(\mathbf{Y}, \boldsymbol{\theta})= & -\frac{\partial^{2}}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}} \log \int f_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{z} ; \boldsymbol{\theta}) d \mathbf{z} \\
= & -\frac{\int f_{\mathbf{Y}, \mathbf{Z}}^{\prime \prime}(\mathbf{Y}, \mathbf{z} ; \boldsymbol{\theta}) d \mathbf{z}}{\int f_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{z} ; \boldsymbol{\theta}) d \mathbf{z}}+\mathbf{S} \mathbf{Y}(\mathbf{Y}, \boldsymbol{\theta})^{\otimes 2} \\
= & E_{\boldsymbol{\theta}}\left\{\mathbf{I}_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{Z}, \boldsymbol{\theta}) \mid \mathbf{Y}\right\} \\
& -E_{\boldsymbol{\theta}}\left\{\mathbf{S}_{\mathbf{Y}, \mathbf{Z}}(\mathbf{Y}, \mathbf{Z}, \boldsymbol{\theta})^{\otimes 2} \mid \mathbf{Y}\right\} \\
& +\mathbf{S}_{\mathbf{Y}}(\mathbf{Y}, \boldsymbol{\theta})^{\otimes 2}
\end{aligned}
$$

where $\mathbf{I}_{\mathbf{Y}, \mathbf{Z}}$ and $\mathbf{Z}_{\mathbf{Y}, \mathbf{Z}}$ only need to be computed at the last iteration of the EM procedure where $\mathbf{S}_{\mathbf{Y}}(\mathbf{Y}, \widehat{\boldsymbol{\theta}})=\mathbf{0}$.

### 2.10 Uniqueness of MLE

Converge to boundary: sequence $\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}, \ldots$ in $\Theta$ is said to converge to boundary $\partial \Theta$ if for every compact set (If a closed set $A$ is bounded, then $A$ is a compact set) $K \subset \Theta$, there exists $k_{0} \geq 1$ such that $\boldsymbol{\theta}^{(k)} \notin K \quad \forall k \geq k_{0}$ (if $\Theta=R^{b}$, then this definition is equivalent to $\lim _{k \rightarrow \infty}\left\|\boldsymbol{\theta}^{(k)}\right\|=\infty$ )
Constant on boundary: a real-valued function $f$ defined on $\Theta$ is said to be constant on boundary $\partial \Theta$ if $\lim _{k \rightarrow \infty} f\left(\boldsymbol{\theta}^{(k)}\right)=c$ for every sequence $\boldsymbol{\theta}^{(k)}$ in $\Theta$ converging to $\partial \Theta ; c$ could be $\pm \infty$. Written as $\lim _{\boldsymbol{\theta} \rightarrow \partial \Theta} f(\boldsymbol{\theta})=c$.

Theorem 2.1. for $\Theta \subset R^{b}, b \geq 1$ which is a connected open set. If

1. $\ell(\boldsymbol{\theta})$ is twice continuously differentiable with $\lim _{\boldsymbol{\theta} \rightarrow \partial \Theta} \ell(\boldsymbol{\theta})=c$, where constant $c$ is either a real number or $-\infty$
2. $\mathbf{I}_{T}(\mathbf{Y}, \boldsymbol{\theta})=-\partial^{2} \ell(\boldsymbol{\theta}) / \partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^{T}$ (observed total information matrix) is positive definite at every point $\boldsymbol{\theta} \in \Theta$ for which $\partial \ell(\boldsymbol{\theta}) / \partial \boldsymbol{\theta}^{T}=\mathbf{0}$
then
3. the critical point is unique and is the MLE
4. $\ell(\boldsymbol{\theta})>c, \forall \boldsymbol{\theta} \in \Theta$

Recall: Note a matrix $A$ is positive definite if $\mathbf{x}^{T} A \mathbf{x}>0, \forall x . A$ is positive definite if and only if

$$
a_{11}>0 ;\left|\begin{array}{ll}
a_{11} & a_{12} \\
a_{21} & a_{22}
\end{array}\right|>0 ;\left|\begin{array}{lll}
a_{11} & a_{12} & a_{13} \\
a_{21} & a_{22} & a_{23} \\
a_{31} & a_{32} & a_{33}
\end{array}\right|>0 ; \cdots ;|A|>0
$$

Note: the constancy on boundary is not really necessary for the scalar parameter but is need for other cases

- with a scalar parameter, the second conditional that minus Hessian is positive definite at all stationary points ( $x$ at which $f^{\prime}(x)=0$ ) $\Rightarrow$ absence of local minima $\Rightarrow$ absence of multiple local maxima (multiple local maxima cannot occur without local minima)
- counter-example showing the necessity of "constancy on boundary"

$$
\begin{aligned}
& \ell=g(x, y)=-e^{-2 y}-e^{-y} \sin x \\
& \frac{\partial}{\partial x} g(x, y)=-e^{-y} \cos x \\
& \frac{\partial}{\partial y} g(x, y)=2 e^{-2 y}+e^{-y} \sin x
\end{aligned}
$$

solution: $x=(2 k+1.5) \pi, k=1,2, \cdots$, and $y=\log 2$

$$
\begin{aligned}
-\nabla^{2} g(x, y) & =\left(\begin{array}{cc}
-e^{-y} \sin x & -e^{-y} \cos x \\
-e^{-y} \cos x & 4 e^{-2 y}+e^{-y} \sin x
\end{array}\right) \\
& \text { at stationary points }\left(\begin{array}{cc}
0.5 & 0 \\
0 & 0.5
\end{array}\right)
\end{aligned}
$$

which is positive definite $\Rightarrow$ all solutions corresponding to maxima, which is the solutions are all maxima and there are countably infinite of them cause the "constancy on boundary" is violated

Theorem 2.2. for $\Theta \subset R^{b}, b \geq 1$ which is a connected open set and $\ell(\boldsymbol{\theta})$ is twice continuously differentiable. if

1. $\partial \ell(\boldsymbol{\theta}) / \partial \boldsymbol{\theta}^{T}=\mathbf{0}$ has at least one solution
2. $\mathbf{I}_{T}(\mathbf{Y}, \boldsymbol{\theta})$ is positive definite $\forall \boldsymbol{\theta} \in \Theta$
then
3. $\ell(\boldsymbol{\theta})$ is concave
4. the solution is unique and is the MLE

Example 1: normal location-scale model

$$
Y_{1}, \cdots, Y_{n} \sim \operatorname{iid} N\left(\mu, \sigma^{2}\right), \Theta=(-\infty, \infty) \times(0, \infty)
$$

$$
\begin{aligned}
\ell(\mu, \sigma) & =\text { const }-n \log \sigma-\frac{1}{2 \sigma^{2}} \sum_{i=1}^{n}\left(Y_{i}-\mu\right)^{2} \\
\frac{\partial}{\partial \boldsymbol{\theta}^{T}} \ell(\mu, \sigma) & =\binom{\frac{1}{\sigma^{2}} \sum_{i=1}^{n}\left(Y_{i}-\mu\right)}{-\frac{n}{\sigma}+\frac{1}{\sigma^{3}} \sum_{i=1}^{n}\left(Y_{i}-\mu\right)^{2}} \\
\mathbf{I}_{T}(\mu, \sigma) & =\left(\begin{array}{cc}
\frac{n}{\sigma^{2}} & \frac{2}{\sigma^{3}} \sum_{i=1}^{n}\left(Y_{i}-\mu\right) \\
\frac{2}{\sigma^{3}} \sum_{i=1}^{n}\left(Y_{i}-\mu\right) & -\frac{n}{\sigma^{2}}+\frac{3}{\sigma^{4}} \sum_{i=1}^{n}\left(Y_{i}-\mu\right)^{2}
\end{array}\right)
\end{aligned}
$$

only solution to $\partial \ell(\mu, \sigma) / \partial \boldsymbol{\theta}^{T}$ is $\widehat{\mu}=\bar{Y}, \widehat{\sigma}=s_{n}$. Use 2.1 we can show that 1$) \mathbf{I}_{T}\left(\bar{Y}, s_{n}\right)=\operatorname{diag}\left(n / s_{n}^{2}, 2 n / s_{n}^{2}\right)$ is positive definite, and 2) $\lim _{\boldsymbol{\theta} \rightarrow 2 \Theta} \ell(\boldsymbol{\theta})=-\infty$ (HW3 2.58)
2.2 is not applicable cause $\mathbf{I}_{T}(\mu, \sigma)$ is not positive definite everywhere. $\left|\mathbf{I}_{T}(\mu, \sigma)\right|=n^{2} \sigma^{-6}\left\{3 s_{n}^{2}-\sigma^{2}-(\bar{Y}-\mu)^{2}\right\}$ could be negative.
Example 2: Exponential threshold model (combine with the profile likelihood approach)

$$
f(y ; \mu, \sigma)=\frac{1}{\sigma} \exp \left(-\frac{y-\mu}{\sigma}\right), \quad \mu \leq y<\infty
$$

with $Y_{i}^{\prime} s$ are i.i.d and $\Theta=(-\infty, \infty) \times(0, \infty)$

$$
L(\mu, \sigma \mid \mathbf{Y})=\sigma^{-n} \exp \left(-\sum_{i=1}^{n} \frac{Y_{i}-\mu}{\sigma}\right) \prod_{i=1}^{n} I\left(\mu \leq Y_{i}\right)
$$

which is not differentiable in $\mu$ everywhere. By using the profile likelihood method. For each $\sigma, L(\mu, \sigma \mid \mathbf{Y})$ is maximized at $\hat{\mu}=Y_{(1)} \Rightarrow$ log-profile likelihood $\ell^{*}(\sigma)=-n \log \sigma-\sigma^{-1} \sum_{i}\left(Y_{i}-Y_{(1)}\right)$

$$
\begin{gathered}
\frac{\partial}{\partial \sigma} \ell^{*}(\sigma)=-n \sigma^{-1}+n \sigma^{-2}\left(\bar{Y}-Y_{(1)}\right) \Longrightarrow \widehat{\sigma}=\bar{Y}-Y_{(1)} \text { if } \bar{Y}>Y_{(1)} \\
-\frac{\partial^{2}}{\partial \sigma^{2}} \ell^{*}(\sigma)=-n \sigma^{-2}+2 n \sigma^{-3}\left(\bar{Y}-Y_{(1)}\right)
\end{gathered}
$$

is not always positive thus 2.2 does not apply. Nonetheless, 2.1 does

1. $-\partial^{2} \ell^{*}(\sigma) /\left.\partial \sigma^{2}\right|_{\sigma=\widehat{\sigma}}=n \widehat{\sigma}^{-2}>0$
2. $\lim _{\sigma \rightarrow 0} \ell^{*}(\sigma)=\lim _{\sigma \rightarrow \infty} \ell^{*}(\sigma)=-\infty$

Example 3: example to show that existence and uniqueness of the MLE is not always a given

$$
Y_{1}, \cdots, Y_{n} \stackrel{i . i . d}{\sim} \text { mixture of normals }
$$

$$
f(y ; \mu, \sigma, p)=p \phi(y-\mu)+(1-p) \frac{1}{\sigma} \phi\left(\frac{y-\mu}{\sigma}\right)
$$

the log-likelihood assuming $\mu=Y_{1}$

$$
\ell=\log \left\{p \phi(0)+(1-p) \frac{1}{\sigma} \phi(0)\right\}+\sum_{i=2}^{n} \log \left\{p \phi\left(Y_{i}-Y_{1}\right)+(1-p) \frac{1}{\sigma} \phi\left(\frac{Y_{i}-Y_{1}}{\sigma}\right)\right\}
$$

as $\sigma \rightarrow 0$ the first term $\rightarrow \infty$ and the rest are bounded. Thus, MLE does not exist in the strict sense. Furthermore, there exist multiple local maxima.
Nonetheless, local maxima satisfying the likelihood score tend to behave well. EM algorithm can find at least one of them.
Uniqueness of the MLE for exponential family
$\mathrm{X} \sim$ minimal exponential family

$$
f(x ; \boldsymbol{\theta})=h(x) \exp \left\{\sum_{i=1}^{s} g_{i}(\boldsymbol{\theta}) T_{i}(x)-B(\boldsymbol{\theta})\right\}=h(x) \exp \left\{\sum_{i=1}^{k} \eta_{i} T_{i}(\eta)-A(\eta)\right\}
$$

where $g(\boldsymbol{\theta})$ is 1-to-1, twice differentiable in $\Theta$ and $\Theta$ is an open subset of $R^{s}$. Then if there is at least one solution to the transformed likelihood equation $E_{\boldsymbol{\theta}}\{\mathbf{T}(X)\}=\mathbf{T}(x)$ then with 2.2 the solution is unique and is the MLE.
With the canonical representation, the likelihood score is

$$
\frac{\partial}{\partial \eta^{T}} A(\boldsymbol{\eta})-\mathbf{T}(x)=0 \Rightarrow \frac{\partial}{\partial \eta^{T}} A(\boldsymbol{\eta}) \equiv E_{\boldsymbol{\eta}} \mathbf{T}(X)=\mathbf{T}(x)
$$

the minus Hessian is

$$
\frac{\partial^{2}}{\partial \boldsymbol{\eta} \partial \boldsymbol{\eta}^{T}} A(\boldsymbol{\eta}) \equiv \mathbf{I}(\boldsymbol{\eta})=\operatorname{Var}_{\boldsymbol{\eta}}\{\mathbf{T}(X)\}
$$

which is positive definite as $T$ is affinely independent. Then based on invariant property so is $\widehat{\boldsymbol{\theta}}_{\text {MLE }}=$ $g^{-1}\left(\widehat{\boldsymbol{\eta}}_{\text {MLE }}\right)$
Families with truncation or threshold parameters

$$
\begin{gathered}
x_{1}, \cdots, x_{n} \stackrel{i . i . d}{\sim} f\left(x ; \boldsymbol{\theta}, \mu_{1}, \mu_{2}\right) \\
f\left(x ; \boldsymbol{\theta}, \mu_{1}, \mu_{2}\right)=c\left(\mu_{1}, \mu_{2}, \boldsymbol{\theta}\right) d(x, \boldsymbol{\theta}) \quad \mu_{1}<x<\mu_{2}
\end{gathered}
$$

for fixed $\theta,\left(X_{(1)}, X_{(n)}\right)$ is minial sufficent for $\left(\mu_{1}, \mu_{2}\right)$. Conditional on $\left(X_{(1)}, X_{(n)}\right)=\left(x_{(1)}, x_{(n)}\right)$, sample values between $X_{(1)}$ and $X_{(n)}$ say

$$
z_{1}, \cdots, z_{n-2} \stackrel{i . i . d}{\sim} q(z, \boldsymbol{\theta})=\frac{d(z, \boldsymbol{\theta}) I\left(x_{(1)}<z<x_{(n)}\right)}{\int_{x_{(1)}}^{x_{(n)}} d(z, \boldsymbol{\theta}) d z}
$$

if $d(x, \boldsymbol{\theta})$ has an exponential family form then so does $q(z, \boldsymbol{\theta})$
Similar results for

$$
\begin{array}{ll}
f(x ; \boldsymbol{\theta}, \mu)=c_{1}(\mu, \boldsymbol{\theta}) d_{1}(x, \boldsymbol{\theta}) & \mu<x \\
f(x ; \boldsymbol{\theta}, \mu)=c_{2}(\mu, \boldsymbol{\theta}) d_{2}(x, \boldsymbol{\theta}) & x<\mu
\end{array}
$$

Example 1:

$$
f(x ; \theta, \mu)=\theta e^{-\theta(x-\mu)} I(\mu<x)
$$

conditional on $X_{(1)}=x_{(1)}$, observations larger than $x_{(1)}$ have density

$$
q(z, \theta)=\theta e^{-\theta\left(z-x_{(1)}\right)} I\left(x_{(1)}<z\right)=e^{-\theta z+\theta x_{(1)}+\log (\theta)} I\left(x_{(1)}<z\right)
$$

## 3 Likelihood-based tests and confidence regions

Let's first look at the simplest scalar parameter case

$$
H_{0}: \theta=\theta_{0} \text { vs. } H_{a}: \theta \neq \theta_{0}
$$

- Wald test

$$
T_{\mathrm{W}}=\frac{\left(\widehat{\theta}_{\mathrm{MLE}}-\theta_{0}\right)^{2}}{\left\{I_{\mathrm{T}}\left(\widehat{\theta}_{\mathrm{MLE}}\right)\right\}^{-1}}
$$

- Likelihood ratio test

$$
T_{\mathrm{LR}}=-2\left\{\ell\left(\theta_{0}\right)-\ell\left(\widehat{\theta}_{\mathrm{MLE}}\right)\right\}
$$

- Score test

$$
T_{\mathrm{S}}=\frac{S\left(\theta_{0}\right)^{2}}{I_{\mathrm{T}}\left(\theta_{0}\right)}
$$

Under $H_{0}$, asymptotically they are all $\chi_{1}^{2}$. Under local alternatives, they have identical asymptotic non-central $\chi^{2}$ distribution.

The idea behind those tests are graphically illustrated below

Figure 2: Graphical representation of the relationships between Wald, Score and Likelihood Ratio test statistics


Parameter $\theta$

The Likelihood Ratio test statistic is a multiple of the difference, $\Delta_{1}$; the Wald test statistic is a multiple of the squared difference, $\Delta_{2}^{2}$; and the Score test statistic is a multiple of the squared slope $\beta^{2}$ (this figure is from the textbook)

### 3.1 Simple null hypothesis

By "simple" we mean don't have nuisance parameter which is we are interested in the whole parameter vector.

$$
H_{0}: \boldsymbol{\theta}=\boldsymbol{\theta}_{0} \text { vs. } H_{a}: \boldsymbol{\theta} \neq \boldsymbol{\theta}_{0} \quad b \text {-dimensional }
$$

- Wald test

$$
T_{\mathrm{W}}=\left(\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}-\boldsymbol{\theta}_{0}\right)^{T} \mathbf{I}_{\mathrm{T}}\left(\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right)\left(\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}-\boldsymbol{\theta}_{0}\right)
$$

where $\mathbf{I}_{\mathrm{T}}\left(\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right)$ may be replace by $\mathbf{I}_{\mathrm{T}}\left(\boldsymbol{\theta}_{0}\right)$ or $\mathbf{I}_{\mathrm{T}}\left(\mathbf{Y}, \widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right)$

- Likelihood ratio test

$$
T_{\mathrm{LR}}=-2\left\{\ell\left(\boldsymbol{\theta}_{0}\right)-\ell\left(\hat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right)\right\}
$$

By taylor expansion
$\ell\left(\boldsymbol{\theta}_{0}\right)=\ell\left(\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right)+\mathbf{S}\left(\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right)^{T}\left(\boldsymbol{\theta}_{0}-\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right)-\frac{1}{2}\left(\boldsymbol{\theta}_{0}-\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right)^{T} \mathbf{I}_{\mathrm{T}}\left(\mathbf{Y}, \widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right)\left(\boldsymbol{\theta}_{0}-\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right)+$ residual here $\mathbf{S}\left(\widehat{\boldsymbol{\theta}}_{\text {MLE }}\right)^{T}=0$ cause this is how MLE is derived. Therefore

$$
T_{\mathrm{W}}+\delta_{n} \text { with } \delta_{n} \xrightarrow{p} 0
$$

where $\delta_{n}$ is the residual

- Score test

$$
T_{\mathrm{S}}=\mathbf{S}\left(\boldsymbol{\theta}_{0}\right)^{T} \mathbf{I}_{\mathrm{T}}\left(\boldsymbol{\theta}_{0}\right)^{-1} \mathbf{S}\left(\boldsymbol{\theta}_{0}\right)
$$

Under $H_{0}, E\left\{\mathbf{S}\left(\boldsymbol{\theta}_{0}\right)\right\}=0$ and $\operatorname{var}\left\{\mathbf{S}\left(\boldsymbol{\theta}_{0}\right)\right\}=\mathbf{I}_{\mathrm{T}}\left(\boldsymbol{\theta}_{0}\right)$. By central limit theorem and continuous mapping theorem $T_{\mathrm{S}} \xrightarrow{d} \chi_{b}^{2}$.

Example 1:

$$
\begin{aligned}
& Y_{1}, \cdots, Y_{n} \stackrel{\text { i.i.d }}{\sim} N(\mu, 1) \cdot H_{0}: \mu=\mu_{0} \\
& \ell(\mu)=-\frac{n}{2} \log (2 \pi)-\frac{1}{2} \sum_{i}\left(Y_{i}-\mu\right)^{2} \\
& S(\mu)=\frac{\partial}{\partial \mu} \ell(\mu)=\sum_{i}\left(Y_{i}-\mu\right) \Longrightarrow \widehat{\mu}=\bar{Y} \\
& I_{\mathrm{T}}(\mathbf{Y}, \mu)=-\frac{\partial}{\partial \mu} S(\mu)=n
\end{aligned}
$$

we got the test statistics as follows

$$
\begin{aligned}
T_{\mathrm{W}} & =\left(\bar{Y}-\mu_{0}\right)(n)\left(\bar{Y}-\mu_{0}\right)=n\left(\bar{Y}-\mu_{0}\right)^{2} \\
T_{\mathrm{S}} & =\left\{\sum_{i}\left(Y_{i}-\mu_{0}\right)\right\} n^{-1}\left\{\sum_{i}\left(Y_{i}-\mu_{0}\right)\right\}=T_{\mathrm{W}} \\
T_{\mathrm{LR}} & =-2\left\{-\frac{1}{2} \sum_{i}\left(Y_{i}-\mu_{0}\right)^{2}+\frac{1}{2} \sum_{i}\left(Y_{i}-\bar{Y}\right)^{2}\right\}=T_{\mathrm{W}}
\end{aligned}
$$

Example 2: $Y_{1}, \cdots, Y_{n} \stackrel{\text { iid }}{\sim} \operatorname{Bernoulli}(p) . H_{0}: p=p_{0}$. Let $X=\sum_{i} Y_{i}$

$$
\begin{aligned}
L & =\prod_{i=1}^{n} p^{y_{i}}(1-p)^{\left(1-y_{i}\right)} \\
\ell & =\sum_{i=1}^{n} y_{i} \log (p)+\left(1-y_{i}\right) \log (1-p)=X \log (p)+(n-X) \log (1-p) \\
S(p) & =\frac{X}{p}-\frac{n-X}{1-p}=\frac{X-n p}{p(1-p)} \Longrightarrow \widehat{p}=\frac{X}{n} \Longrightarrow X=n \hat{p} \\
I_{\mathrm{T}}(\mathbf{Y}, p) & =\frac{X}{p^{2}}+\frac{n-X}{(1-p)^{2}} \quad I_{\mathrm{T}}(p)=\frac{n}{p(1-p)} \\
T_{\mathrm{W}} & =\left(\widehat{p}-p_{0}\right) \frac{n}{\widehat{p}(1-\widehat{p})}\left(\widehat{p}-p_{0}\right)=\frac{n\left(\widehat{p}-p_{0}\right)^{2}}{\widehat{p}(1-\widehat{p})} \\
T_{\mathrm{S}} & =S\left(p_{0}\right) I_{\mathrm{T}}\left(p_{0}\right)^{-1} S\left(p_{0}\right)=\frac{n\left(\widehat{p}-p_{0}\right)^{2}}{p_{0}\left(1-p_{0}\right)} X \text { is replaced by } n \hat{p} \\
T_{\mathrm{LR}} & =-2\left[X \log \left(p_{0} / \widehat{p}\right)+(n-X) \log \left\{\left(1-p_{0}\right) /(1-\widehat{p})\right\}\right]
\end{aligned}
$$

### 3.2 Composite null hypothesis

In this case, $\boldsymbol{\theta}$ is partitioned and we are only interested in part of it.

$$
\underset{b \times 1}{\boldsymbol{\theta}}=\left(\begin{array}{c}
\boldsymbol{\theta}_{1} \\
r \times 1 \\
\boldsymbol{\theta}_{2} \\
(b-r) \times 1
\end{array}\right)
$$

$H_{0}: \boldsymbol{\theta}_{1}=\boldsymbol{\theta}_{10}$ vs. $H_{a}: \boldsymbol{\theta}_{1} \neq \boldsymbol{\theta}_{10}$, with $\boldsymbol{\theta}_{2}$ as nuisance.

- Wald test

$$
\mathbf{I}_{\mathrm{T}}\left(\widehat{\boldsymbol{\theta}}_{\mathrm{MLE}}\right) \equiv \widehat{\mathbf{I}}_{\mathrm{T}}=\left(\begin{array}{ll}
\widehat{\mathbf{I}}_{\mathrm{T}, 11} & \widehat{\mathbf{I}}_{\mathrm{T}, 12} \\
\widehat{\mathbf{I}}_{\mathrm{T}, 21} & \widehat{\mathbf{I}}_{\mathrm{T}, 22}
\end{array}\right)
$$

Under $H_{0}$, Avar $\left(\widehat{\boldsymbol{\theta}}_{1}-\boldsymbol{\theta}_{10}\right)$ is the upper $(1,1)$ element of $\widehat{\mathbf{I}}_{\mathrm{T}}^{-1}$, given by

$$
\left(\widehat{\mathbf{I}}_{\mathrm{T}, 11}-\widehat{\mathbf{I}}_{\mathrm{T}, 12} \widehat{\mathbf{I}}_{\mathrm{T}, 22}^{-1} \widehat{\mathbf{I}}_{\mathrm{T}, 21}\right)^{-1}
$$

therefore

$$
T_{\mathrm{W}}=\left(\widehat{\boldsymbol{\theta}}_{1}-\boldsymbol{\theta}_{10}\right)^{T}\left(\widehat{\mathbf{I}}_{\mathrm{T}, 11}-\widehat{\mathbf{I}}_{\mathrm{T}, 12} \widehat{\mathbf{I}}_{\mathrm{T}, 22}^{-1} \widehat{\mathbf{I}}_{\mathrm{T}, 21}\right)\left(\widehat{\boldsymbol{\theta}}_{1}-\boldsymbol{\theta}_{10}\right) \xrightarrow{d} \chi_{r}^{2}
$$

be careful here, don't write the middle part as $\left(\widehat{\mathbf{I}}_{\mathrm{T}, 11}\right)^{-1}$

- Likelihood ratio test

$$
T_{\mathrm{LR}}=-2 \log \left\{\frac{\sup _{\boldsymbol{\theta} \in H_{0}} L(\boldsymbol{\theta} \mid \mathbf{Y})}{\sup _{\boldsymbol{\theta} \in \boldsymbol{\Theta}} L(\boldsymbol{\theta} \mid \mathbf{Y})}\right\}=-2\{\ell(\widetilde{\boldsymbol{\theta}})-\ell(\widehat{\boldsymbol{\theta}})\}
$$

where $\sup _{\boldsymbol{\theta} \in H_{0}} L(\boldsymbol{\theta} \mid \mathbf{Y})$ is the restricted MLE and $\sup _{\boldsymbol{\theta} \in \boldsymbol{\Theta}} L(\boldsymbol{\theta} \mid \mathbf{Y})$ is the unrestricted MLE

- Score test

First derive the $H_{0}$-restricted MLE

$$
\begin{gathered}
\widetilde{\boldsymbol{\theta}}=\binom{\boldsymbol{\theta}_{10}}{\widetilde{\boldsymbol{\theta}}_{2}} \quad \widetilde{\boldsymbol{\theta}}_{2}=\arg \max _{\boldsymbol{\theta}_{2}} \ell\left(\boldsymbol{\theta}_{10}, \boldsymbol{\theta}_{2}\right) \\
\mathbf{I}_{\mathrm{T}}\left(\tilde{\boldsymbol{\theta}}_{\mathrm{MLE}}\right) \equiv \tilde{\mathbf{I}}_{\mathrm{T}} \\
\mathbf{S}(\boldsymbol{\theta})=\binom{\mathbf{S}_{1}(\boldsymbol{\theta})}{\mathbf{S}_{2}(\boldsymbol{\theta})}=\binom{\frac{\partial}{\partial \boldsymbol{\theta}_{1}^{T}} \ell(\boldsymbol{\theta})}{\frac{\partial}{\partial \boldsymbol{\theta}_{2}^{T}} \ell(\boldsymbol{\theta})} \\
T_{\mathrm{S}}=\mathbf{S}(\widetilde{\boldsymbol{\theta}})^{T} \widetilde{\mathbf{I}}_{\mathrm{T}}^{-1} \mathbf{S}(\widetilde{\boldsymbol{\theta}}) \\
=\left(\mathbf{S}_{1}(\widetilde{\boldsymbol{\theta}})^{T}, \mathbf{0}^{T}\right)\left(\begin{array}{cc}
\widetilde{\mathbf{I}}_{\mathrm{T}, 11} & \widetilde{\mathbf{I}}_{\mathrm{T}, 12} \\
\widetilde{\mathbf{I}}_{\mathrm{T}, 21} & \widetilde{\mathbf{I}}_{\mathrm{T}, 22}
\end{array}\right)^{-1}\binom{\mathbf{S}_{1}(\widetilde{\boldsymbol{\theta}})}{\mathbf{0}} \\
=\mathbf{S}_{1}(\widetilde{\boldsymbol{\theta}})^{T}\left(\widetilde{\mathbf{I}}_{\mathrm{T}, 11}-\widetilde{\mathbf{I}}_{\mathrm{T}, 12} \widetilde{\mathbf{I}}_{\mathrm{T}, 22}^{-1} \widetilde{\mathbf{I}}_{\mathrm{T}, 21}\right)^{-1} \mathbf{S}_{1}(\widetilde{\boldsymbol{\theta}})
\end{gathered}
$$

where $\mathbf{S}_{2}(\widetilde{\boldsymbol{\theta}})=0$ by definition of MLE
Composite null hypotheses of general form $H_{0}: \underset{r \times 1}{\mathbf{h}(\boldsymbol{\theta})}=\mathbf{0}$

$$
\underset{r \times b}{\mathbf{H}}(\boldsymbol{\theta})=\frac{\partial}{\partial \boldsymbol{\theta}} \mathbf{h}(\boldsymbol{\theta}), \quad r \leq b
$$

which needs to be of full rank
Below are some examples to help understand the definition of $h$ and $H$ Example 1: bi-variate normal data

$$
\begin{gathered}
\boldsymbol{\theta}=\left(\mu_{1}, \mu_{2}, \sigma_{1}, \sigma_{2}, \rho\right)^{T} \\
H_{0}: \mu_{1}=\mu_{2} \\
h(\boldsymbol{\theta})=\mu_{1}-\mu_{2} \\
\mathbf{H}(\boldsymbol{\theta})=(1,-1,0,0,0)
\end{gathered}
$$

Example 2: linear hypotheses $\mathbf{K}^{T} \beta=\mathbf{m}$

$$
\mathbf{h}(\beta)=\mathbf{K}^{T} \beta-\mathbf{m}
$$

Example 3: partitioned-vector hypothesis $H_{0}: \boldsymbol{\theta}_{1}=\boldsymbol{\theta}_{10}$

$$
\mathbf{h}(\boldsymbol{\theta})=\boldsymbol{\theta}_{1}-\boldsymbol{\theta}_{10}
$$

Now let's jump in to the test statistics

- Wald test

$$
\begin{aligned}
\widehat{\boldsymbol{\theta}} & \sim A N\left(\boldsymbol{\theta}, \mathbf{I}_{\mathrm{T}}(\boldsymbol{\theta})^{-1}\right) \\
\mathbf{h}(\widehat{\boldsymbol{\theta}}) & \sim A N\left\{\mathbf{h}(\boldsymbol{\theta}), \mathbf{H}(\boldsymbol{\theta}) \mathbf{I}_{\mathrm{T}}(\boldsymbol{\theta})^{-1} \mathbf{H}(\boldsymbol{\theta})^{T}\right\} \\
T_{\mathrm{W}} & =\mathbf{h}(\widehat{\boldsymbol{\theta}})^{T}\left\{\mathbf{H}(\widehat{\boldsymbol{\theta}}) \mathbf{I}_{\mathrm{T}}(\widehat{\boldsymbol{\theta}})^{-1} \mathbf{H}(\widehat{\boldsymbol{\theta}})^{T}\right\}^{-1} \mathbf{h}(\widehat{\boldsymbol{\theta}})
\end{aligned}
$$

from the first line to second line we used Delta method one problem is that the test statistic varies with reparameterization and choice of h, e.g., $h\left(\mu_{1}, \mu_{2}\right)=\mu_{1}-\mu_{2}$ and $h\left(\mu_{1}, \mu_{2}\right)=\mu_{1} / \mu_{2}-1$.

- Score test

$$
T_{\mathrm{S}}=\mathbf{S}(\widetilde{\boldsymbol{\theta}})^{T} \widetilde{\mathbf{I}}_{\mathrm{T}}^{-1} \mathbf{S}(\widetilde{\boldsymbol{\theta}})
$$

where $\tilde{\boldsymbol{\theta}}$ maximizes the likelihood subject to $\mathbf{h}(\boldsymbol{\theta})=\mathbf{0}$.
Actually a better way to do this is using the Lagrange multiplier for

$$
\max \ell(\boldsymbol{\theta}) \text { subject to } \mathbf{h}(\boldsymbol{\theta})=\mathbf{0}
$$

we will focus on $\ell(\boldsymbol{\theta})-\mathbf{h}(\boldsymbol{\theta})_{r \times 1}^{T} \boldsymbol{\lambda}$

$$
\begin{aligned}
\mathbf{S}(\boldsymbol{\theta})-\mathbf{H}(\boldsymbol{\theta})^{T} \boldsymbol{\lambda} & =\mathbf{0} \\
\mathbf{h}(\boldsymbol{\theta}) & =\mathbf{0}
\end{aligned}
$$

Denote the solution by $\tilde{\boldsymbol{\theta}}$ and $\tilde{\boldsymbol{\lambda}}$ then we get $\mathbf{S}(\tilde{\boldsymbol{\theta}})=\mathbf{H}(\tilde{\boldsymbol{\theta}})^{T} \tilde{\boldsymbol{\lambda}}$ and the score test statistic turns into

$$
T_{\mathrm{S}}=\tilde{\boldsymbol{\lambda}}^{T} \mathbf{H}(\widetilde{\boldsymbol{\theta}}) \tilde{\mathbf{I}}_{\mathrm{T}}^{-1} \mathbf{H}(\widetilde{\boldsymbol{\theta}})^{T} \tilde{\boldsymbol{\lambda}}
$$

Below is a comparison of the three method
invariant to re-parameterization and the choice of $\mathbf{h}(\cdot) \quad$ computing MLE's
unrestricted
restricted

| Wald | N | Y | N |
| :--- | :--- | :--- | :--- |
| Score | Y | N | Y |
| LR | Y | Y | Y |

Nonetheless, LR can be convenient for nested models. Typically, Wald does not have as good type I error as score and LR. And Wald and score are easier to adjust in the case of model misspecification.
Example 1: Normal location-scale model $Y_{1}, \cdots, Y_{n} \stackrel{\text { iid }}{\sim} N\left(\mu, \sigma^{2}\right)$ and $H_{0}: \mu=\mu_{0} \quad$ vs. $\quad H_{a}: \mu \neq \mu_{0} \quad \sigma$ unrestricted

$$
\begin{aligned}
\ell(\mu, \sigma) & =\text { constant }-n \log \sigma-\frac{1}{2 \sigma^{2}} \sum_{i}\left(Y_{i}-\mu\right)^{2} \\
\mathbf{S}(\mu, \sigma) & =\sum_{i}\left[\begin{array}{c}
\sigma^{-2}\left(Y_{i}-\mu\right) \\
\sigma^{-3}\left\{\left(Y_{i}-\mu\right)^{2}-\sigma^{2}\right\}
\end{array}\right] \Rightarrow \widehat{\mu}=\bar{Y} \quad \widehat{\sigma}^{2}=s_{n}^{2}=n^{-1} \sum_{i}\left(Y_{i}-\bar{Y}\right)^{2}, \mathbf{I}_{\mathrm{T}}(\mu, \sigma)=\operatorname{diag}\left(n / \sigma^{2}, 2 n / \sigma^{2}\right) \\
\widetilde{\mu} & =\mu_{0} \quad \tilde{\sigma}^{2}=n^{-1} \sum_{i}\left(Y_{i}-\mu_{0}\right)^{2}=s_{n}^{2}+\left(\bar{Y}-\mu_{0}\right)^{2} \\
T_{\mathrm{w}} & =\left(\bar{Y}-\mu_{0}\right) \frac{n}{s_{n}^{2}}\left(\bar{Y}-\mu_{0}\right)=\frac{n\left(\bar{Y}-\mu_{0}\right)^{2}}{s_{n}^{2}}=\frac{n}{n-1} \frac{n\left(\bar{Y}-\mu_{0}\right)^{2}}{n s_{n}^{2} /(n-1)}=\frac{n}{n-1} t^{2} \text { where } t \text { means t-distribution with } d f=1 \\
T_{\mathrm{S}} & =\left\{\frac{1}{\widetilde{\sigma}^{2}} \sum_{i}\left(Y_{i}-\mu_{0}\right)\right\} \frac{\tilde{\sigma}^{2}}{n}\left\{\frac{1}{\widetilde{\sigma}^{2}} \sum_{i}\left(Y_{i}-\mu_{0}\right)\right\}=\frac{n\left(\bar{Y}-\mu_{0}\right)^{2}}{\tilde{\sigma}^{2}}=\frac{n T_{\mathrm{W}}}{n+T_{\mathrm{W}}} \\
T_{\mathrm{LR}} & =n \log \frac{\tilde{\sigma}^{2}}{\hat{\sigma}^{2}}=n \log \left\{1+\frac{\left(\bar{Y}-\mu_{0}\right)^{2}}{s_{n}^{2}}\right\}=n \log \left\{1+\frac{T_{\mathrm{W}}}{n}\right\}
\end{aligned}
$$

Since $\frac{x}{1+x} \leq \log (1+x) \leq x$ for $x>-1$

$$
T_{\mathrm{S}} \leq T_{\mathrm{LR}} \leq T_{\mathrm{W}}
$$

Using exact distributions, they are all equivalent to $t$ test (can be transformed to $t$ test). Using the asymptotic $\chi_{1}^{2}$ critical values, $T_{\mathrm{W}}$ is more liberal.

Example 2: Score test for multinomial data $=$ Pearson $\chi^{2}$ test

$$
\begin{aligned}
& \left(N_{1}, \cdots, N_{k}\right) \sim \text { Multinomial }\left(n, p_{1}, \cdots, p_{k}\right) \\
& \mathbf{p}=\left(p_{1}, \cdots, p_{k-1}\right)^{T} \quad p_{k}=1-p_{1}-\cdots-p_{k-1} \\
& f\left(n_{1}, \cdots, n_{k}\right)=\frac{n!}{n_{1}!\cdots n_{k}!} p_{1}^{n_{1}} \cdots p_{k}^{n_{k}} \\
& \ell(\mathbf{p})=\sum_{i=1}^{k-1} N_{i} \log p_{i}+N_{k} \log \left(1-\sum_{i=1}^{k-1} p_{i}\right)+\text { const. } \\
& \mathbf{S}(\mathbf{p})=\left(\frac{N_{1}}{p_{1}}-\frac{N_{k}}{p_{k}}, \cdots, \frac{N_{k-1}}{p_{k-1}}-\frac{N_{k}}{p_{k}}\right)^{T} \\
& \mathbf{I}_{\mathrm{T}}(\mathbf{p})=n \operatorname{diag}\left(\frac{1}{p_{1}}, \cdots, \frac{1}{p_{k-1}}\right)+\frac{n}{p_{k}} \mathbf{1}^{\otimes 2} \\
& \mathbf{I}_{\mathrm{T}}(\mathbf{p})^{-1}=\left\{\operatorname{diag}(\mathbf{p})-\mathbf{p}^{\otimes 2}\right\} / n \\
& T_{\mathrm{S}}=\mathbf{S}(\widetilde{\mathbf{p}})^{T} \mathbf{I}_{\mathrm{T}}(\widetilde{\mathbf{p}})^{-1} \mathbf{S}(\widetilde{\mathbf{p}}) \\
& =n^{-1} \sum_{i=1}^{k}\left(\frac{N_{i}}{\widetilde{p}_{i}}-\frac{N_{k}}{\widetilde{p}_{k}}\right)^{2} \widetilde{p}_{i} \\
& -n^{-1}\left\{\sum_{i=1}^{k}\left(\frac{N_{i}}{\widetilde{p}_{i}}-\frac{N_{k}}{\widetilde{p}_{k}}\right) \widetilde{p}_{i}\right\} \\
& =n^{-1} \sum_{i=1}^{k}\left\{\frac{N_{i}}{\widetilde{p}_{i}}-\frac{N_{k}}{\widetilde{p}_{k}}-\left(n-\frac{N_{k}}{\widetilde{p}_{k}}\right)\right\}^{2} \widetilde{p}_{i} \\
& =\sum_{i=1}^{k} \frac{\left(N_{i}-n \widetilde{p}_{i}\right)^{2}}{n \widetilde{p}_{i}}
\end{aligned}
$$

in the last step, the variance equality

$$
\sum a_{i}^{2} p_{i}-\left(\sum a_{i} p_{i}\right)^{2}=\sum\left(a_{i}-\sum a_{i} p_{i}\right)^{2} p_{i} \text { is used. }
$$

Example 3: Testing for Hardy-Weinberg equilibrium
Multinomial $k=3: p_{A A}, p_{A a}, p_{a a}$
Under the equilibrium, $p_{A A}=p_{A}^{2}, p_{A a}=2 p_{A}\left(1-p_{A}\right)$,
$p_{a a}=\left(1-p_{A}\right)^{2}$
restricted MLE: $\widetilde{p}_{A}=\left(2 N_{A A}+N_{A a}\right) /(2 n)$

$$
\begin{aligned}
T_{\mathrm{S}}= & \frac{\left(N_{A A}-n \widetilde{p}_{A}^{2}\right)^{2}}{n \widetilde{p}_{A}^{2}}+\frac{\left\{N_{A a}-2 n \widetilde{p}_{A}\left(1-\widetilde{p}_{A}\right)\right\}^{2}}{2 n \widetilde{p}_{A}\left(1-\widetilde{p}_{A}\right)} \\
& +\frac{\left\{N_{a a}-n\left(1-\widetilde{p}_{A}\right)^{2}\right\}^{2}}{n\left(1-\widetilde{p}_{A}\right)^{2}}
\end{aligned}
$$

the degree of freedom is 1 cause

$$
h\left(P_{A A}, P_{A a}\right)=P_{A a}-2 P_{A A}^{1 / 2}\left(1-P_{A A}^{1 / 2}\right) \Rightarrow r=1
$$

### 3.3 Confidence interval

$$
\begin{aligned}
\underset{b \times 1}{\boldsymbol{\theta}} & =\left(\begin{array}{c}
\boldsymbol{\theta}_{1} \\
r \times 1 \\
\boldsymbol{\theta}_{2} \\
(b-r) \times 1
\end{array}\right) \\
T\left(\boldsymbol{\theta}_{1}\right) & =\text { test stat with } \boldsymbol{\theta}_{1} \text { as the null value (actually the true value) } \\
C_{1-\alpha} & =\left\{\boldsymbol{\theta}_{1}: T\left(\boldsymbol{\theta}_{1}\right) \leq \chi_{r}^{2}(1-\alpha)\right\}
\end{aligned}
$$

Asymptotically, (random) confidence region contains the true parameter value $\theta_{10}$ with probability $1-\alpha$ if $\boldsymbol{\theta}_{10}$ is the true value. If $\boldsymbol{\theta}_{10} \notin C_{1-\alpha}$ then we reject the null hypothesis

$$
P\left(\boldsymbol{\theta}_{10} \in C_{1-\alpha}\right)=P\left\{T\left(\boldsymbol{\theta}_{10}\right) \leq \chi_{r}^{2}(1-\alpha)\right\}
$$

Example 1: Binomial model
We previously derived the test statistics for the binomial model as $T_{\mathrm{W}}=\frac{n(\widehat{p}-p)^{2}}{\widehat{p}(1-\widehat{p})}$ the solving

$$
T_{\mathrm{W}}=\frac{n(\widehat{p}-p)^{2}}{\widehat{p}(1-\widehat{p})} \leq \chi_{1}^{2}(1-\alpha)=z_{1-\alpha / 2}^{2} \Longrightarrow\left(\widehat{p}-z_{1-\alpha / 2} \sqrt{\frac{\widehat{p}(1-\widehat{p})}{n}}, \widehat{p}+z_{1-\alpha / 2} \sqrt{\frac{\widehat{p}(1-\widehat{p})}{n}}\right)
$$

### 3.4 Nonstandard hypothesis testing problems

If $\left.\frac{\partial}{\partial \boldsymbol{\theta}} \ell(\boldsymbol{\theta} \mid \mathbf{Y})\right|_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}_{\text {MLE }}} \neq \mathbf{0}$ in probability, typically the asymptotic normality of $\hat{\boldsymbol{\theta}}_{\text {MLE }}$ is no longer true and the 3 likelihood-based tests do not have a limiting $\chi_{r}^{2}$ null distribution. Example 1: Exponential threshold model

$$
Y_{1}, \cdots, Y_{n} \stackrel{\mathrm{iid}}{\sim} f(y ; \mu)= \begin{cases}e^{-(y-\mu)} & y>\mu \\ 0 & \text { otherwise }\end{cases}
$$

and the MLE is $\widehat{\mu}_{\text {MLE }}=Y_{(1)}$. In this case, the score equation is not 0 .
Null hypotheses on the boundary of the parameter space

$$
\begin{aligned}
& Y_{1}, \cdots, Y_{n} \stackrel{\text { iid }}{\sim} N(\mu, 1) \quad \mu \geq \mu_{0} \\
& \widehat{\mu}_{\mathrm{MLE}}=\max \left(\bar{Y}, \mu_{0}\right) \\
& H_{0}: \mu=\mu_{0} \quad \text { vs. } \quad H_{a}: \mu>\mu_{0} \\
& T_{\mathrm{W}}=n\left(\widehat{\mu}_{\mathrm{MLE}}-\mu_{0}\right)^{2} \\
& = \begin{cases}n\left(\bar{Y}-\mu_{0}\right)^{2} & \bar{Y} \geq \mu_{0} \\
0 & \text { otherwise }\end{cases} \\
& T_{\mathrm{S}}=n\left(\bar{Y}-\mu_{0}\right)^{2} \sim \chi_{1}^{2} \text { under } H_{0} \\
& T_{\mathrm{LR}}=T_{\mathrm{W}}
\end{aligned}
$$

When a null hypothesis value, say $\boldsymbol{\theta}_{0}$ lies on the boundary of the parameter space, then maximum likelihood estimators are often truncated at that boundary because by definitionb $\widehat{\boldsymbol{\theta}}_{\text {MLE }}$ must lie in the parameter space of $\boldsymbol{\theta}$. Thus $\widehat{\boldsymbol{\theta}}_{\text {MLE }}$ equal to the boundary value $\boldsymbol{\theta}_{0}$ with positive probability and correspondingly $T_{L R}$ zero for those cases. The result is that the limiting distribution of $T_{L R}$ is a mixture of a point mass at zero and a chi-squared distribution.
In this case, $\bar{Y}$ used to be greater than $\mu_{0}$ with probability $\frac{1}{2}$. Therefore $T_{\mathrm{w}}$ or $T_{\mathrm{LR}}$ has a null distribution that is an equal mixture of a point mass at 0 and a $\chi_{1}^{2}$ :

$$
Z^{2} I(Z>0), \quad Z \sim N(0,1)
$$

with this modified limiting distribution. The confidence interval shall be modified accordingly: level- $\alpha$ test: reject $H_{0}$ with critical value being $1-2 \alpha$ quantile of $\chi_{1}^{2}$

$$
\begin{aligned}
P\left\{Z^{2} I(Z>0) \geq \chi_{1}^{2}(1-2 \alpha)\right\} & =P\left\{Z^{2} \geq \chi_{1}^{2}(1-2 \alpha), Z>0\right\} \\
& =\frac{1}{2} P\left\{Z^{2} \geq \chi_{1}^{2}(1-2 \alpha)\right\}=\alpha
\end{aligned}
$$

Although the score test statistics still follows the chi-square distribution. By following standard approach we can reject if

$$
\bar{Y}<\mu_{0}-\sqrt{\chi_{1}^{2}(\alpha) / n} \text { or } \bar{Y}>\mu_{0}+\sqrt{\chi_{1}^{2}(\alpha) / n}
$$

however, the $\bar{Y}<\mu_{0}-\sqrt{\chi_{1}^{2}(\alpha) / n}$ part doesn't make sense cause the parameter space is $\left[\boldsymbol{\theta}_{0}, \infty\right)$ thus a natural solution is to reject when $\sqrt{n}\left(\bar{Y}-\mu_{0}\right)>z_{\alpha}$
For cases that are not in the same form as this example, we can try to transfer the test statistics to be like that in this example.
Example 2: Isotonic regression
There are $k$ independent normal samples of size $n_{1}, \cdots, n_{k}$, each i.i.d with means $\mu_{1}, \cdots, \mu_{k}$, respectively, and common variance $\sigma^{2}$, satisfying

$$
\mu_{1} \leq \mu_{2} \leq \cdots \leq \mu_{k}
$$

The MLE minimizes $\sum_{i=1}^{k} n_{i}\left(\bar{Y}_{i}-\mu_{i}\right)^{2}$ subject to $\mu_{1} \leq \mu_{2} \leq \cdots \leq \mu_{k}$. We'd like to test
$H_{0}: \quad \mu_{1}=\cdots=\mu_{k}$
$H_{a}: \quad \mu_{1} \leq \cdots \leq \mu_{k}$ with at least 1 strict inequality
we can reparameterize as

$$
\begin{gathered}
\mu_{1} \\
\Delta \mu_{2}=\mu_{2}-\mu_{1} \\
\vdots \\
\Delta \mu_{k}=\mu_{k}-\mu_{k-1}
\end{gathered}
$$

Then, this results in null hypothesis on the boundary:
$H_{0}: \quad \Delta \mu_{2}=\cdots=\Delta \mu_{k}=0$
$H_{a}: \Delta \mu_{2} \geq 0, \cdots, \Delta \mu_{k} \geq 0$ with at least 1 strict inequality

## 4 Bayesian methods

$\star$ this part will not be tested

* this part of the note is combined with notes from BIS 567 Bayesian Statistics at Yale to enhance understanding


### 4.1 Introduction

- Frequentist approach
- the unknown parameter $\theta$ is assume to be constant
- data $\mathbf{Y} \sim f(\mathbf{y} ; \boldsymbol{\theta})$ is considered random
- estimation can be via MLE, M-estimation, ...
- hypothesis testing can be likelihood-based,...
- confidence interval is derived by inverting a test statistic
- Bayesian approach
- the unknown parameter $\theta$ is from $\boldsymbol{\theta} \sim \pi(\boldsymbol{\theta})$, the prior distribution
- data is from $\mathbf{Y} \sim f(\mathbf{y} \mid \boldsymbol{\theta})$
- estimation is via the posterior density

$$
\pi(\boldsymbol{\theta} \mid \mathbf{Y}=\mathbf{y})=\frac{f(\mathbf{y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta})}{\int f(\mathbf{y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d \boldsymbol{\theta}}
$$

point estimation can use ,e.g., posterior mean

- $1-\alpha$ credible region: parameter space with posterior probability $1-\alpha$

Marginal density (prior predictive density) of $Y$

$$
\mathbf{Y}, m(\mathbf{y})=\int f(\mathbf{y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d \boldsymbol{\theta}
$$

Posterior predictive density

$$
\begin{aligned}
m\left(\mathbf{y}_{\text {new }} \mid \mathbf{Y}\right) & =\int f\left(\mathbf{y}_{\text {new }} \mid \boldsymbol{\theta}, \mathbf{Y}\right) \pi(\boldsymbol{\theta} \mid \mathbf{Y}) d \boldsymbol{\theta} \\
& =\int f\left(\mathbf{y}_{\text {new }} \mid \boldsymbol{\theta}\right) \pi(\boldsymbol{\theta} \mid \mathbf{Y}) d \boldsymbol{\theta} \text { when } \mathbf{Y}_{\text {new }} \mathbf{Y}
\end{aligned}
$$

Where does the prior $\pi(\boldsymbol{\theta})$ come from?

- Subjective Bayesian: personal uncertainty about $\boldsymbol{\theta}$
- From previous knowledge: previous information about $\boldsymbol{\theta}$ (Bayesian analysis is used to combine previous info with current data)
- For convenience: convenient technical density to employ the Bayesian machinery

$$
\begin{aligned}
f(y \mid p) & =\binom{n}{p} p^{y}(1-p)^{n-y} \quad y=0,1, \cdots, n \\
\pi(p) & =\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha) \Gamma(\beta)} p^{\alpha-1}(1-p)^{\beta-1}(\text { Beta }) \\
m(y) & =\int f(y \mid p) \pi(p) d p=\frac{B(y+\alpha, n-y+\beta)}{B(\alpha, \beta)}\binom{n}{y} \\
\pi(p \mid \mathbf{Y}) & =\frac{f(y \mid p) \pi(p)}{m(y)}=\frac{\binom{n}{p} p^{y}(1-p)^{n-y} \cdot \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha) \Gamma(\beta)} p^{\alpha-1}(1-p)^{\beta-1}}{m(y)} \\
& =\frac{\binom{n}{p} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha) \Gamma(\beta)} p^{y+\alpha-1}(1-p)^{n-y+\beta-1}}{m(y)} \sim \operatorname{beta}(Y+\alpha, n-Y+\beta)
\end{aligned}
$$

the posterior mean is

$$
E(p \mid \mathbf{Y})=\frac{Y+\alpha}{\alpha+\beta+n}=\frac{\alpha+\beta}{\alpha+\beta+n} \frac{\alpha}{\alpha+\beta}+\frac{n}{\alpha+\beta+n} \frac{Y}{n}
$$

which shrinks the MLE toward prior mean. Bayesian estimator may have smaller MSE than the MLE.
When $\alpha=\beta=0$, the prior is improper

$$
\Longrightarrow \text { posterior beta }(Y, n-Y) \text { is proper if } 1 \leq Y \leq n-1 \Longrightarrow \text { posterior mean } Y / n
$$

Example 2: $\operatorname{Normal}\left(\theta, \sigma_{Y}^{2}\right)$

$$
\begin{aligned}
& Y_{1}, \cdots, Y_{n} \stackrel{\mathrm{iid}}{\sim} N\left(\theta, \sigma_{Y}^{2}\right) \quad \sigma_{Y}^{2} \text { known } \\
& \text { prior } \theta \sim N\left(\mu_{0}, \sigma_{0}^{2}\right) \quad \mu_{0}, \sigma_{0}^{2} \text { known } \\
& \text { posterior } \theta \left\lvert\, \mathbf{Y} \sim N\left(\frac{\tau_{0} \mu_{0}+\tau_{n} \bar{Y}}{\tau_{0}+\tau_{n}}, \frac{1}{\tau_{0}+\tau_{n}}\right)\right.
\end{aligned}
$$

where precisions $\tau_{0}=1 / \sigma_{0}^{2}$ and $\tau_{n}=n / \sigma_{Y}^{2}$
$95 \%$ credible interval

$$
\left(\frac{\tau_{0} \mu_{0}+\tau_{n} \bar{Y}}{\tau_{0}+\tau_{n}}-\frac{1.96}{\sqrt{\tau_{0}+\tau_{n}}}, \frac{\tau_{0} \mu_{0}+\tau_{n} \bar{Y}}{\tau_{0}+\tau_{n}}+\frac{1.96}{\sqrt{\tau_{0}+\tau_{n}}}\right)
$$

A nice feature of independent data is that one can sequentially update the prior for each additional datum, or all at once. For example, suppose $Y_{1}, \ldots, Y_{n}$ are independent with respective densities $f_{i}\left(y_{i} \mid \boldsymbol{\theta}\right), i=1, \ldots, n$ Then, the posterior is

$$
\begin{aligned}
\pi(\boldsymbol{\theta} \mid \boldsymbol{Y}) & =\frac{\pi(\boldsymbol{\theta}) \prod_{i=1}^{n} f_{i}\left(Y_{i} \mid \boldsymbol{\theta}\right)}{\int \pi(\boldsymbol{\theta}) \prod_{i=1}^{n} f_{i}\left(Y_{i} \mid \boldsymbol{\theta}\right) d \boldsymbol{\theta}} \\
& =\frac{\pi\left(\boldsymbol{\theta} \mid Y_{1}\right) \prod_{i=2}^{n} f_{i}\left(Y_{i} \mid \boldsymbol{\theta}\right)}{\int \pi\left(\boldsymbol{\theta} \mid Y_{1}\right) \prod_{i=2}^{n} f_{i}\left(Y_{i} \mid \boldsymbol{\theta}\right) d \boldsymbol{\theta}} \\
& =\frac{\pi\left(\boldsymbol{\theta} \mid Y_{1}, \ldots Y_{k}\right) \prod_{i=k+1}^{n} f_{i}\left(Y_{i} \mid \boldsymbol{\theta}\right)}{\int \pi\left(\boldsymbol{\theta} \mid Y_{1}, \ldots, Y_{k}\right) \prod_{i=k+1}^{n} f_{i}\left(Y_{i} \mid \boldsymbol{\theta}\right) d \boldsymbol{\theta}}
\end{aligned}
$$

Here $\pi\left(\theta \mid Y_{1}\right)$ is the posterior from using the prior and $Y_{1}$. It then is used as the prior for the remaining data. Sufficient statistics often make calculations easier. Because of the factorization theorem, if a sufficient statistic for $\theta$ exists, then the posterior depends on the data only through the sufficient statistic.
Estimation and inference: Frequentist uses sampling distribution, whereas Bayesian uses posterior density.

### 4.2 Bayesian estimator from decision theory perspective

Loss function $L\{\boldsymbol{\theta}, \boldsymbol{\delta}(\mathbf{Y})\}$ : non-negative function of the true parameter and an estimator $\delta(\mathbf{Y})$. For example:

1. squared error loss: $L\{\boldsymbol{\theta}, \boldsymbol{\delta}(\mathbf{Y})\}=\|\boldsymbol{\theta}-\boldsymbol{\delta}(\mathbf{Y})\|^{2}$
2. absolute error loss: $L\{\theta, \delta(\mathbf{Y})\}=|\theta-\delta(\mathbf{Y})|$
3. 0-1 loss: $L\{\boldsymbol{\theta}, \boldsymbol{\delta}(\mathbf{Y})\}= \begin{cases}0 & \text { if } \boldsymbol{\theta}=\boldsymbol{\delta}(\mathbf{Y}) \\ 1 & \text { otherwise }\end{cases}$

Risk: is the average loss over $Y$

$$
R(\boldsymbol{\theta}, \boldsymbol{\delta})=\int L\{\boldsymbol{\theta}, \boldsymbol{\delta}(\mathbf{y})\} f(\mathbf{y} \mid \boldsymbol{\theta}) d \mathbf{y}
$$

which is not a single number, but rather a function of $\theta$ Several standard frequentist approaches regarding the risk:

1. minimize $R(\boldsymbol{\theta}, \boldsymbol{\delta})$, with squared error loss, over the class of unbiased estimators $\Longrightarrow$ if exists, minimum variance unbiased estimator (MVUE)
2. minimax estimator

$$
\inf _{\boldsymbol{\delta}} \sup _{\boldsymbol{\theta}} R(\boldsymbol{\theta}, \boldsymbol{\delta})
$$

maximize the risk over $\boldsymbol{\theta}$ first then seek an estimator that minimizes the maximum risk
Bayesian approach (take another average over $\theta$ )

$$
R_{\text {Bayes }}(\pi, \boldsymbol{\delta})=\int R(\boldsymbol{\theta}, \boldsymbol{\delta}) \pi(\boldsymbol{\theta}) d \boldsymbol{\theta}
$$

Bayesian estimator $\delta_{\text {Bayes }}$ minimizes $R_{\text {Bayes }}(\pi, \boldsymbol{\delta})$
We can also minimize the posterior risk instead

$$
\rho\{\pi, \boldsymbol{\delta}(\mathbf{Y})\}=\int L\{\boldsymbol{\theta}, \boldsymbol{\delta}(\mathbf{Y})\} \pi(\boldsymbol{\theta} \mid \mathbf{Y}) d \boldsymbol{\theta}
$$

But typically $\arg \min \rho\{\pi, \boldsymbol{\delta}(\mathbf{Y})\} \equiv \arg \min R_{\text {Bayes }}(\pi, \boldsymbol{\delta})$

$$
\begin{aligned}
R_{\text {Bayes }}(\pi, \boldsymbol{\delta}) & =\iint L\{\boldsymbol{\theta}, \boldsymbol{\delta}(\mathbf{y})\} f(\mathbf{y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d \mathbf{y} d \boldsymbol{\theta} \\
= & \iint L\{\boldsymbol{\theta}, \boldsymbol{\delta}(\mathbf{y})\} \pi(\boldsymbol{\theta} \mid \mathbf{y}) m(\mathbf{y}) d \mathbf{y} d \boldsymbol{\theta} \\
= & \int \rho\{\pi, \boldsymbol{\delta}(\mathbf{y})\} m(\mathbf{y}) d \mathbf{y}
\end{aligned}
$$

thus minimize $\rho\{\pi, \boldsymbol{\delta}(\mathbf{Y})\}$ usually minimize $R_{\text {Bayes }}(\pi, \boldsymbol{\delta})$
Different loss function leads to different parameter estimate

- squared error loss $\Longrightarrow$ posterior mean
- absolute error loss $\Longrightarrow$ posterior median
- 0-1 loss $\Longrightarrow$ posterior mode

Bayes estimators with proper priors are generally not unbiased in the frequentist sense. However, they typically have good risk behavior in the frequentist sense. Interestingly, a key technique for finding minimax estimators starts with a Bayes estimator (see Lehmann and Casella, 1998, Ch. 5). Moreover, an admissible estimator (an estimator not uniformly larger in risk compared to any other estimator) must be a Bayes estimator or the limit of Bayes estimators. Thus, Bayes estimators are not only good in terms of Bayes risk but are often of interest to frequentists willing to sacrifice unbiasedness.

### 4.3 Credible intervals

$1-\alpha$ credible interval/region: a region of $\boldsymbol{\theta}$ of posterior prob $1-\alpha$
highest posterior density (HPD) region: the region with minimized volume
region with equal tail probability is often used for a scalar parameter

### 4.4 Conjugate prior

Conjugate prior: when the data $Y$ has density $f(y \mid \theta)$ and the prior and posterior are from the same family of densities, we say that the prior is conjugate.
$\pi(\boldsymbol{\theta})$ governed by fixed hyperparameters
$\gamma_{\text {prior }} \pi(\boldsymbol{\theta} \mid \mathbf{Y})$ has updated hyperparameters $\gamma_{\text {post }}$, via a known function of $\gamma_{\text {prior }}$ and $\mathbf{Y}$
Some examples:

- beta $(\alpha, \beta)$ prior for binomial $(n, p)$ data
- $N\left(\mu_{0}, \sigma_{0}^{2}\right)$ prior for $N\left(\theta, \sigma_{Y}^{2}\right)$ data (known $\left.\sigma_{Y}^{2}\right)$
- beta $(\alpha, \beta)$ prior for negative binomial data
- gamma prior for Poisson data
- gamma prior for gamma data
- Pareto prior for Uniform $(0, \theta)$ data


## Example 1:

$$
\begin{aligned}
\mathbf{Y} & \sim \text { multinomial }\left(n ; p_{1}, \cdots, p_{k}\right) \\
\pi\left(p_{1}, \cdots, p_{k}\right) & \sim \operatorname{Dirichlet}\left(\alpha_{1}, \cdots, \alpha_{k}\right) \\
\pi(\mathbf{p}) & =\frac{\prod_{i=1}^{k} p_{i}^{\alpha_{i}-1}}{B(\boldsymbol{\alpha})}, \quad B(\boldsymbol{\alpha})=\frac{\prod_{i=1}^{k} \Gamma\left(\alpha_{i}\right)}{\Gamma\left(\sum_{i=1}^{k} \alpha_{i}\right)} \\
\pi(\mathbf{p} \mid \mathbf{Y}) & \propto \prod_{i=1}^{k} p_{i}^{N_{i}} \prod_{i=1}^{k} p_{i}^{\alpha_{i}-1} \\
\Longrightarrow \pi(\mathbf{p} \mid \mathbf{Y}) & \sim \operatorname{Dirichlet}\left(\alpha_{1}+N_{1}, \cdots, \alpha_{k}+N_{k}\right)
\end{aligned}
$$

Any data density having a sufficient statistic of fixed dimension $\forall n$ has a conjugate prior.
Example 2: exponential family canonical form

$$
\begin{aligned}
f(\mathbf{y} ; \boldsymbol{\eta}) & =h(\mathbf{y}) \exp \left\{\sum_{i=1}^{s} \eta_{i} T_{i}(\mathbf{y})-A(\boldsymbol{\eta})\right\} \\
\text { conjugate } \pi(\boldsymbol{\eta} \mid \boldsymbol{\gamma}, \lambda) & =K(\boldsymbol{\gamma}, \lambda) \exp \left\{\sum_{i=1}^{s} \gamma_{i} \eta_{i}-\lambda A(\boldsymbol{\eta})\right\} \\
\text { posterior } & \sim \pi(\boldsymbol{\eta} \mid \boldsymbol{\gamma}+\mathbf{T}(\mathbf{y}), \lambda+1)
\end{aligned}
$$

### 4.5 Noninformative prior

Truly noninformative may not be possible. For example, assigning prior probability to events $A, B$, and $C=C_{1} \cup C_{2}$. The choice between the following two cases reflects some kind of information

$$
\begin{aligned}
& P(A)=P(B)=P(C)=1 / 3 \text { or } \\
& P(A)=P(B)=P\left(C_{1}\right)=P\left(C_{2}\right)=1 / 4 ?
\end{aligned}
$$

For the case of a location parameter $\mu$ taking values on $(-\infty, \infty)$ : improper prior $\pi(\mu)=1$ giving equal weight to all values can be justified on a variety of grounds. An improper prior does not have a finite integral. However, there seems to be no philosophical problem with improper priors as long as they lead to proper posteriors.
For the case of a scale parameter $\sigma$ taking values in $(0, \infty)$ : Jeffreys' suggestion $\pi(\sigma)=1 / \sigma$ which has an invariance argument. . His invariance argument is that any power transformation of $\sigma$, say $\gamma=\sigma^{a}$ has via a change-of-variables, the improper density

$$
\gamma=\sigma^{a} \Longrightarrow \pi(\gamma)=\frac{1}{\gamma^{1 / a}}\left|\frac{\gamma^{1 / a-1}}{a}\right|=\frac{1}{a \gamma}
$$

which is similar in form to $1 / \sigma$
Combining these last two improper priors, a location-scale family with $(\mu, \sigma) \in(-\infty, \infty) \times(0, \infty)$, suggests using the improper prior

$$
\pi(\mu, \sigma) \propto \frac{1}{\sigma}
$$

Now, moving to the case of general parameters on continuous parameter spaces. For Jeffreys prior

- Scalar parameter $\pi(\theta) \propto I(\theta)^{1 / 2}$. For transformed parameter $\gamma=g(\theta)$,

$$
\pi(\gamma) \propto \frac{1}{\left|g^{\prime}\left\{g^{-1}(\gamma)\right\}\right|} I\left\{g^{-1}(\gamma)\right\}^{1 / 2}=I(\gamma)^{1 / 2}
$$

Example 1: for $\operatorname{binomial}(n, p) . I(p)=n /\{p(1-p)\}$

$$
\pi(p) \propto\{p(1-p)\}^{-1 / 2} \sim \operatorname{beta}(1 / 2,1 / 2)
$$

- Vector parameter $\pi(\boldsymbol{\theta}) \propto|\mathbf{I}(\boldsymbol{\theta})|^{1 / 2} ;|\cdot|$ is determinant

Example 1: for multinomial $\left(n ; p_{1}, \cdots, p_{k}\right)$

$$
\begin{aligned}
\mathbf{I}_{T}(\mathbf{p}) & =n\left\{\operatorname{diag}\left(1 / p_{1}, \cdots, 1 / p_{k-1}\right)+11^{T} / p_{k}\right\} \\
\pi(\mathbf{p}) & \propto \frac{1}{\sqrt{p_{1} \cdots p_{k}}} \sim \operatorname{Dirichlet}(1 / 2, \cdots, 1 / 2)
\end{aligned}
$$

For $N\left(\mu, \sigma^{2}\right)$ case with $\boldsymbol{\theta}=(\mu, \sigma), \mathbf{I}(\boldsymbol{\theta})=\operatorname{diag}\left(\sigma^{-2}, 2 \sigma^{-2}\right) \Longrightarrow \pi(\mu, \sigma) \propto \sigma^{-2}$ which is different from $\sigma^{-1}$ given earlier. Thus, Jeffreys modified his original proposal in the presence of location parameters say $\mu_{1}, \ldots, \mu_{k}$ to

$$
\pi\left(\mu_{1}, \ldots, \mu_{k}, \boldsymbol{\theta}\right) \propto|\boldsymbol{I}(\boldsymbol{\theta})|^{1 / 2}
$$

where $\mu_{i}^{\prime} s$ are held as fixed.

### 4.6 Normal data examples

### 4.6.1 One sample with unknown mean and variance

Data: $Y_{1}, \cdots, Y_{n} \stackrel{\text { iid }}{\sim} N\left(\mu, \sigma^{2}\right), \boldsymbol{\theta}=\left(\mu, \tau=1 / \sigma^{2}\right)$ unknown

$$
\begin{aligned}
& \text { Prior: } \pi(\mu \mid \tau) \sim N\left(\mu_{0},\left(\tau n_{0}\right)^{-1}\right) \\
& \quad \pi(\tau) \sim \operatorname{gamma}\left(\alpha_{0}, 1 / \beta_{0}\right), \text { mean }=\alpha_{0} / \beta_{0}
\end{aligned}
$$

a rough way to estimate the prior variance is treating $\tau=\alpha_{0} / \beta_{0} \Longrightarrow$ prior variance of $\mu \approx\left(\beta_{0} / \alpha_{0}\right) n_{0}^{-1}$ $\mu \sim t$ with center $\mu_{0}$, scale $e^{2}=\left(\beta_{0} / \alpha_{0}\right) / n_{0}$, and df $2 \alpha_{0}$
then the likelihood is

$$
\left(2 \pi \sigma^{2}\right)^{-n / 2} \exp \left[-\frac{1}{2 \sigma^{2}}\left\{\sum\left(Y_{i}-\bar{Y}\right)^{2}+n(\bar{Y}-\mu)^{2}\right\}\right]
$$

then the posterior is

$$
\begin{aligned}
& \pi(\mu, \tau \mid \mathbf{Y}) \propto \tau^{1 / 2} \exp \left\{-\frac{\tau n^{\prime}}{2}\left(\mu-\mu^{\prime}\right)^{2}\right\} \tau^{\alpha^{\prime}-1} e^{-\tau \beta^{\prime}} \\
& \mu^{\prime}=\frac{n_{0} \mu_{0}+n \bar{Y}}{n_{0}+n} \\
& n^{\prime}=n_{0}+n \\
& \alpha^{\prime}=n / 2+\alpha_{0} \\
& \beta^{\prime}=\beta_{0}+\frac{\sum\left(Y_{i}-\bar{Y}\right)^{2}}{2}+\frac{1}{2}\left(\frac{n_{0} n}{n_{0}+n}\right)\left(\mu_{0}-\bar{Y}\right)^{2} \\
& \Longrightarrow \mu, \tau \mid \mathbf{Y} \sim \mathrm{N}\left(\mu^{\prime},\left(\tau n^{\prime}\right)^{-1}\right) \times \operatorname{gamma}\left(\alpha^{\prime}, 1 / \beta^{\prime}\right) \text { (conjugate prior) } \\
& \mu \mid \mathbf{Y} \sim t \text { with center } \mu^{\prime}, \text { scale } e^{2}=\left(\beta^{\prime} / \alpha^{\prime}\right) / n^{\prime}, \text { and df } 2 \alpha^{\prime} .
\end{aligned}
$$

posterior predictive density

$$
m\left(y_{n+1} \mid Y_{1}, \cdots, Y_{n}\right)=\int f\left(y_{n+1} \mid \boldsymbol{\theta}\right) \pi\left(\boldsymbol{\theta} \mid Y_{1}, \cdots, Y_{n}\right) d \boldsymbol{\theta}
$$

which is a $t$ distribution with mean $\mu^{\prime}$, scale $e^{2}=\left(\beta^{\prime} / \alpha^{\prime}\right)\left(1+1 / n^{\prime}\right)$, and $d f=2 \alpha^{\prime}$.
Letting $\alpha_{0} \rightarrow-1 / 2, \beta_{0} \rightarrow 0$, and $n_{0} \rightarrow 0$ s.t. $\left(n_{0} / \beta_{0}\right)^{1 / 2} \rightarrow 1$
$\Longrightarrow \pi(\mu, \tau) \rightarrow 1 / \tau$, an improper prior
$\Longrightarrow \mu \mid \mathbf{Y} \sim t$ with mean $\bar{Y}$, scale ${ }^{2}=s_{n-1}^{2} / n$, and $\operatorname{df} n-1$
$\Longrightarrow$ Bayes estimator (with squared error loss) is $\bar{Y}$ credible interval $\equiv$ usual frequentist $t$ interval

### 4.6.2 Two samples

$$
X_{1}, \cdots, X_{m} \stackrel{\text { iid }}{\sim} N\left(\mu_{1}, \sigma^{2}\right), \quad Y_{1}, \cdots, Y_{n} \stackrel{\text { iid }}{\sim} N\left(\mu_{2}, \sigma^{2}\right), \boldsymbol{\theta}=\left(\mu_{1}, \mu_{2}, \tau=1 / \sigma^{2}\right)^{T}
$$

$\Delta=\mu_{1}-\mu_{2}$ is of interest
prior

$$
\begin{aligned}
\pi\left(\mu_{1} \mid \tau\right) & \sim N\left\{\mu_{10},\left(\tau m_{0}\right)^{-1}\right\} \\
\pi\left(\mu_{2} \mid \tau\right) & \sim N\left\{\mu_{20},\left(\tau n_{0}\right)^{-1}\right\} \\
\pi(\tau) & \sim \operatorname{gamma}\left(\alpha_{0}, 1 / \beta_{0}\right)
\end{aligned}
$$

with similar derivations as in the last example, we can get the posterior

$$
\begin{aligned}
& \theta \mid \mathbf{X}, \mathbf{Y} \sim \mathrm{N}\left\{\mu_{1}^{\prime},\left(\tau m^{\prime}\right)^{-1}\right\} \times \mathrm{N}\left\{\mu_{2}^{\prime},\left(\tau n^{\prime}\right)^{-1}\right\} \times \operatorname{gamma}\left(\alpha^{\prime}, 1 / \beta^{\prime}\right) \\
& m^{\prime}=m_{0}+m \mu_{1}^{\prime}=\left(m_{0} \mu_{10}+m \bar{X}\right) /\left(m_{0}+m\right) \\
& n^{\prime}=n_{0}+n \mu_{2}^{\prime}=\left(n_{0} \mu_{20}+n \bar{Y}\right) /\left(n_{0}+n\right) \\
& \alpha^{\prime}=\alpha_{0}+m / 2+n / 2 \\
& \beta^{\prime}=\beta_{0}+\frac{\sum\left(X_{i}-\bar{X}\right)^{2}}{2}+\frac{1}{2}\left(\frac{m_{0} m}{m_{0}+m}\right)\left(\mu_{10}-\bar{X}\right)^{2}+\frac{\sum\left(Y_{i}-\bar{Y}\right)^{2}}{2}+\frac{1}{2}\left(\frac{n_{0} n}{n_{0}+n}\right)\left(\mu_{20}-\bar{Y}\right)^{2} \\
& \Longrightarrow \Delta, \tau \mid \mathbf{X}, \mathbf{Y} \sim \mathrm{N}\left\{\Delta^{\prime}=\mu_{1}^{\prime}-\mu_{2}^{\prime}, 1 /\left(\tau m^{\prime}\right)+1 /\left(\tau n^{\prime}\right)\right\} \times \operatorname{gamma}\left(\alpha^{\prime}, 1 / \beta^{\prime}\right) \\
& \Longrightarrow \Delta \mid \mathbf{X}, \mathbf{Y} \sim t \text { with center } \Delta^{\prime}, \text { scale }{ }^{2}=\left(\beta^{\prime} / \alpha^{\prime}\right)\left(m^{\prime}+n^{\prime}\right) /\left(m^{\prime} n^{\prime}\right), \text { and df } 2 \alpha^{\prime}
\end{aligned}
$$

Letting $\alpha_{0} \rightarrow-1, \beta_{0} \rightarrow 0, m_{0} \rightarrow 0, n_{0} \rightarrow 0$ such that $\left(m_{0} n_{0}\right)^{1 / 2} / \beta_{0} \rightarrow 1$
$\longrightarrow$ improper prior $\pi(\boldsymbol{\theta})=1 / \tau$ (Jeffrey's prior)
$\Longrightarrow \Delta \mid \mathbf{X}, \mathbf{Y} \sim t$ with mean $\bar{X}-\bar{Y}$, scale ${ }^{2}=s_{p}^{2}(1 / m+1 / n)$, and $d f=m+n-2$, where $s_{p}^{2}$ is the usual pooled estimated variance.

### 4.6.3 Normal linear model

$$
\mathbf{Y}_{(n \times 1)}=\mathbf{X}_{(n \times p)} \boldsymbol{\Delta}_{(p)}+\mathbf{e}_{(n \times 1)}, \quad e_{1}, \cdots, e_{n} \stackrel{\text { i.i.d }}{\sim} N\left(0, \sigma^{2}\right)
$$

Normal-gamma prior:

$$
\boldsymbol{\Delta} \mid \tau=1 / \sigma^{2} \sim N\left(\boldsymbol{\Delta}_{0}, \boldsymbol{\Sigma}_{0}^{-1} / \tau\right), \quad \tau \sim \operatorname{gamma}\left(\alpha_{0}, 1 / \beta_{0}\right)
$$

Posterior is also Normal-gamma:

$$
\boldsymbol{\Delta}\left|\tau, \mathbf{Y} \sim N\left(\boldsymbol{\Delta}^{\prime},\left(\mathbf{X}^{T} \mathbf{X}+\boldsymbol{\Sigma}_{0}\right)^{-1} / \tau\right), \quad \tau\right| \mathbf{Y} \sim \operatorname{gamma}\left(\alpha^{\prime}, 1 / \beta^{\prime}\right)
$$

where

$$
\begin{aligned}
\Delta^{\prime} & =\left(\mathbf{X}^{T} \mathbf{X}+\boldsymbol{\Sigma}_{0}\right)^{-1}\left(\mathbf{X}^{T} \mathbf{Y}+\boldsymbol{\Sigma}_{0} \boldsymbol{\Delta}_{0}\right) \\
\alpha^{\prime} & =\alpha_{0}+n / 2 \\
\beta^{\prime} & =\beta_{0}+\frac{1}{2}\left\{\left(\mathbf{Y}-\mathbf{X} \boldsymbol{\Delta}^{\prime}\right)^{T} \mathbf{Y}+\left(\boldsymbol{\Delta}_{0}-\boldsymbol{\Delta}^{\prime}\right)^{T} \boldsymbol{\Sigma}_{0} \boldsymbol{\Delta}_{0}\right\}
\end{aligned}
$$

$\Delta^{\prime}$ is a weighted average of the OLS estimator $\left(\mathbf{X}^{T} \mathbf{X}\right)^{-1} \mathbf{X}^{T} \mathbf{Y}$ and prior mean $\Delta_{0}$
$\Delta \mid \mathbf{Y}$ is multivariate $t$ with center $\Delta^{\prime}$, scale matrix $\left(\beta^{\prime} / \alpha^{\prime}\right)\left(\mathbf{X}^{T} \mathbf{X}+\boldsymbol{\Sigma}_{0}\right)^{-1}$, and df $2 \alpha^{\prime}$. For a subset of $\boldsymbol{\Delta}$, the marginal posterior is also a multivariate $t$.
let $\Delta_{0}=\mathbf{0}, \boldsymbol{\Sigma}_{0}=d \mathbf{I}_{p}$ with constant $d \Longrightarrow \Delta^{\prime}=\left(\mathbf{X}^{T} \mathbf{X}+d \mathbf{I}_{p}\right)^{-1} \mathbf{X}^{T} \mathbf{Y}$, the ridge regression estimator.
Letting $\boldsymbol{\Sigma}_{0} \rightarrow \mathbf{0}, \alpha_{0} \rightarrow-p / 2$, and $\beta_{0} \rightarrow 0$
$\Longrightarrow \pi(\Delta, \tau)=1 / \tau$
$\Longrightarrow \Delta^{\prime}=\left(\mathbf{X}^{T} \mathbf{X}\right)^{-1} \mathbf{X}^{T} \mathbf{Y}$, and $\boldsymbol{\Delta} \mid \mathbf{Y}$ is a $p$-dimensional $t$ with df $n-p$ and scale matrix $s^{2}\left(\mathbf{X}^{T} \mathbf{X}\right)^{-1}$, where $s^{2}=(n-p)^{-1} \sum\left(Y_{i}-\mathbf{x}_{i}^{T} \Delta^{\prime}\right)^{2}$

### 4.7 Hierarchical Bayes and empirical Bayes

Previously we discussed that case where prior $\pi(\boldsymbol{\theta})$ has been specified fully, say, $\pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha})$ with hyperparameter $\boldsymbol{\alpha}=\boldsymbol{\alpha}_{0}$ given. What if unsure about $\boldsymbol{\alpha}_{0}$ ? We turn to Hierarchical Bayes.
Hierarchical Bayes: Apart from $\pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha})$, we also specify a hyperprior $h(\boldsymbol{\alpha})=h\left(\boldsymbol{\alpha} \mid \gamma_{0}\right)$ with $\gamma_{0}$ given. The resulting prior

$$
\begin{aligned}
& \pi(\boldsymbol{\theta})=\pi\left(\boldsymbol{\theta} \mid \gamma_{0}\right)=\int \pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha}) h\left(\boldsymbol{\alpha} \mid \gamma_{0}\right) d \boldsymbol{\alpha} \\
& m(\mathbf{y} \mid \boldsymbol{\alpha})=\int f(\mathbf{y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha}) d \boldsymbol{\theta} \\
& \pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha}, \mathbf{Y})=\frac{f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha})}{m(\mathbf{Y} \mid \boldsymbol{\alpha})} \\
& \pi(\boldsymbol{\alpha} \mid \mathbf{Y})=\frac{m(\mathbf{Y} \mid \boldsymbol{\alpha}) h(\boldsymbol{\alpha})}{\iint f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha}) h(\boldsymbol{\alpha}) d \boldsymbol{\alpha} d \boldsymbol{\theta}} \\
& \pi(\boldsymbol{\theta} \mid \mathbf{Y})=\frac{\int f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha}) h(\boldsymbol{\alpha}) d \boldsymbol{\alpha}}{\iint f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha}) h(\boldsymbol{\alpha}) d \boldsymbol{\alpha} d \boldsymbol{\theta}} \\
&=\int \pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha}, \mathbf{Y}) \pi(\boldsymbol{\alpha} \mid \mathbf{Y}) d \boldsymbol{\alpha}
\end{aligned}
$$

Empirical Bayes: use MLE $\widehat{\boldsymbol{\alpha}}$ from marginal likelihood $m(\mathbf{Y} \mid \boldsymbol{\alpha})$ (other estimators may also be used).

$$
\pi(\boldsymbol{\alpha} \mid \mathbf{Y}) \text { highly peaked at } \widehat{\boldsymbol{\alpha}} \Longrightarrow \pi(\boldsymbol{\theta} \mid \mathbf{Y}) \approx \pi(\boldsymbol{\theta} \mid \widehat{\boldsymbol{\alpha}}, \mathbf{Y}) \text { (taking } \pi(\hat{\boldsymbol{\alpha}} \mid \mathbf{Y}) \approx 1 \text { ) }
$$

Empirical Bayes posterior has a similar mean, but smaller variability than the full Bayes.
Example 1: one-way normal random effects model

$$
Y_{i j}, i=1, \cdots, k ; j=1, \cdots, n_{i}
$$

Frequentist methods: fixed and random effects ANOVA.
Bayesian with fixed effects: earlier normal linear model. A Bayesian analogue of the random effects model:

$$
\begin{aligned}
Y_{i j} \mid \theta_{i}, \sigma_{e}^{2} & \sim N\left(\theta_{i}, \sigma_{e}^{2}\right) \text { given } \theta_{1}, \cdots, \theta_{k}, \sigma_{e}^{2}, Y_{i j} \text { 's are mutually independent } \\
\theta_{1}, \cdots, \theta_{k} \mid \boldsymbol{\alpha}=\left(\mu, \sigma_{a}\right) & \text { i.i.d } N\left(\mu, \sigma_{a}^{2}\right) \\
\pi\left(\sigma_{e}^{2}\right) & \propto 1 / \sigma_{e}^{2}(\text { Jeffrey's } \\
h(\mu) & \propto 1 \\
h\left(\sigma_{a}\right) & \propto 1
\end{aligned}
$$

Of interest:

1. random effects population: $\mu$ and $\sigma_{a}$
2. individual $\theta_{i}$

For the empirical Bayes approach, consider a simplified model with $\sigma_{e}^{2}$ known. Obtain $\widehat{\boldsymbol{\alpha}}$ from the marginal likelihood $m(\mathbf{y} \mid \boldsymbol{\alpha})$. The posterior is $\pi(\boldsymbol{\theta} \mid \widehat{\boldsymbol{\alpha}}, \mathbf{Y})$.

### 4.7.1 James-Stein estimation

$Y_{1}, \cdots, Y_{b}$ are independent: $Y_{i} \sim N\left(\theta_{i}, \sigma_{0}^{2}\right)$ with $\sigma_{0}^{2}$ known.
While the parameters are unconnected in the model, the inferences regarding them are connected through the squared error loss $\|\boldsymbol{\theta}-\widehat{\boldsymbol{\theta}}(\mathbf{Y})\|^{2}$. Consider $\hat{\boldsymbol{\theta}}_{\mathrm{MLE}}=\mathbf{Y}, R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}})=b \sigma_{0}^{2}$.
Stein (1955) proved the remarkable result that for $b \geq 3, \widehat{\boldsymbol{\theta}}_{\text {MLE }}$ is inadmissible. James and Stein (1961) provided a dominating estimator,

$$
\widehat{\boldsymbol{\theta}}_{\mathrm{JS}}=\left\{1-\frac{(b-2) \sigma_{0}^{2}}{\sum_{i=1}^{b} Y_{i}^{2}}\right\} \mathbf{Y}
$$

shrinking $\mathbf{Y}$ toward $\mathbf{0}$.
Expected squared error loss

$$
\sum_{i=1}^{b} E\left(\widehat{\theta}_{\mathrm{JS}, i}-\theta_{i}\right)^{2}=b \sigma_{0}^{2}-(b-2)^{2} \sigma_{0}^{4} E\left(\frac{1}{\sum_{i=1}^{b} Y_{i}^{2}}\right)
$$

which is less than the risk for $\widehat{\boldsymbol{\theta}}_{\text {MLE }}$ provided $b \geq 3$
Empirical Bayes interpretation:

$$
\text { prior: } \boldsymbol{\theta} \sim \operatorname{MN}\left(\mathbf{0}, \sigma_{a}^{2} \mathbf{I}_{b}\right)
$$

$\Longrightarrow$ posterior: $\boldsymbol{\theta} \mid \mathbf{Y} \sim$ normal with mean $\left(1-\frac{\sigma_{0}^{2}}{\sigma_{0}^{2}+\sigma_{a}^{2}}\right) \mathbf{Y} \quad$ marginal $\mathbf{Y} \sim M N\left(\mathbf{0},\left(\sigma_{0}^{2}+\sigma_{a}^{2}\right) \mathbf{I}_{b}\right) \sigma_{a}^{2}$ is unknown, but can be estimated from the fact

$$
E\left\{\frac{(b-2) \sigma_{0}^{2}}{\sum_{i=1}^{b} Y_{i}^{2}}\right\}=\frac{\sigma_{0}^{2}}{\sigma_{0}^{2}+\sigma_{a}^{2}}
$$

Thus, $\hat{\boldsymbol{\theta}}_{\text {JS }}$ may be viewed as an empirical Bayes estimator (Efron and Morris, 1972).
Alternative prior

$$
\boldsymbol{\theta} \sim M N\left(\mu, \sigma_{a}^{2} \mathbf{I}_{b}\right)
$$

Bayes posterior mean $B \mu+(1-B) \mathbf{Y}, B=\frac{\sigma_{0}^{2}}{\sigma_{0}^{2}+\sigma_{a}^{2}}$

$$
\text { Marginal } \mathbf{Y} \sim \sim \mathrm{N}\left(\mu,\left(\sigma_{0}^{2}+\sigma_{a}^{2}\right) \mathbf{I}_{b}\right)
$$

unbiased estimators $\widehat{\mu}=\bar{Y} \quad \widehat{B}=(b-3) \sigma_{0}^{2} / \sum_{i=1}^{b}\left(Y_{i}-\bar{Y}\right)^{2}$
James-Stein estimator $\widehat{B} \widehat{\mu}+(1-\widehat{B}) \mathbf{Y}$,
shrinking towards sample mean $\bar{Y}$ and having expected squared error loss less than $\mathbf{Y}$ if $b \geq 4$.

### 4.7.2 Meta-analysis applications of hierarchical and empirical Bayes

Meta analysis: analysis of a group of studies related to the same question of interest. Each study has a point effect estimate and its associated standard error.
Hierarchical Bayes: $k$ studies. $i$ th study - effect parameter $\theta_{i}$ and data $Y_{i}$. Ignoring first-stage nuisance parameters, 3 levels of the model are

$$
\begin{gathered}
f(\mathbf{y} \mid \boldsymbol{\theta})=f\left(y_{1}, \cdots, y_{k} \mid \theta_{1}, \cdots, \theta_{k}\right)=\prod_{i=1}^{k} f\left(y_{i} \mid \theta_{i}\right) \\
\pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha})=\prod_{i=1}^{k} f\left(\theta_{i} \mid \boldsymbol{\alpha}\right) \\
h(\boldsymbol{\alpha})
\end{gathered}
$$

The 2 nd and 3 rd stages $\Longrightarrow$ prior for $\boldsymbol{\theta}$ :

$$
\pi(\boldsymbol{\theta})=\int \prod_{i=1}^{k} f\left(\theta_{i} \mid \boldsymbol{\alpha}\right) h(\boldsymbol{\alpha}) d \boldsymbol{\alpha}
$$

Empirical Bayes:
marginal density of $\mathbf{Y} \Longrightarrow$ estimator $\widehat{\boldsymbol{\alpha}}$ for $\boldsymbol{\alpha}$
$\Longrightarrow$ estimated prior $\pi(\boldsymbol{\theta} \mid \widehat{\boldsymbol{\alpha}})$
$\Longrightarrow$ estimated posterior $\pi(\boldsymbol{\theta} \mid \widehat{\boldsymbol{\alpha}}, \mathbf{Y})$
$\Longrightarrow$ empirical Bayes estimate close to posterior mean from the hierarchical Bayes, but the variance is smaller. Example 1: Normal models with known variance
Data: $\left(Y_{1}, V_{1}\right), \cdots,\left(Y_{k}, V_{k}\right)$, where $Y_{i}$ is approximately normal with variance estimated by $V_{i}$ (but treated as known in analysis)
Frequentist fixed effect approach: all $Y_{i}$ 's are approximately unbiased estimators of true effect $\mu \Longrightarrow$ weighted average $\sum_{i=1}^{k} V_{i}^{-1} Y_{i} / \sum_{i=1}^{k} V_{i}^{-1}$ is approximately optimal. Assumption is not quite realistic!
Frequentist random effect approach: $\theta_{i} \sim N\left(\mu, \sigma_{a}^{2}\right)$. The likelihood of $\mathbf{Y}$ or other methods can be used for the estimation of $\mu$ and $\sigma_{a}$.
Hierarchical Bayesian:

$$
\begin{aligned}
& Y_{i} \mid \theta_{i} \sim N\left(\theta_{i}, V_{i}\right) \\
& \theta_{i} \mid \mu, \sigma_{a} \sim N\left(\mu, \sigma_{a}^{2}\right) \\
& h\left(\mu, \sigma_{a}\right) \propto 1 \text { (noninformative prior) }
\end{aligned}
$$

Empirical Bayes: use estimator of ( $\mu, \sigma_{a}$ ) (same as those from the frequentist random effect approach) without the hyperprior.
Full Bayes: with the hyperprior, MCMC is needed for the estimation.

### 4.8 Monte Carlo estimation of a posterior

(most materials from Carlin and Louis, 1996)
Main technical problem in Bayesian analysis: obtaining the posterior density and computing summary quantities such as posterior mean, standard deviation, and quantiles.
Which means we need to calculate integrals! As soon as we move away from conjugate priors and/or to hierarchical models.
One approach is to resort to asymptotic results. More focus is on Monte Carlo methods.

### 4.8.1 Noniterative Monte Carlo methods

## Direct sampling:

Suppose that $\boldsymbol{\theta} \sim f(\boldsymbol{\theta})$ and $\gamma=E\{h(\boldsymbol{\theta})\}=\int h(\boldsymbol{\theta}) f(\boldsymbol{\theta}) d \boldsymbol{\theta}$ is of interest. Generate $\boldsymbol{\theta}_{1}, \cdots, \boldsymbol{\theta}_{N} \stackrel{\text { iid }}{\sim} f(\boldsymbol{\theta})$ and obtain

$$
\widehat{\gamma}=N^{-1} \sum_{j=1}^{N} h\left(\boldsymbol{\theta}_{j}\right)
$$

## Indirect sampling:

1. Importance sampling

Consider a posterior expectation

$$
E\{h(\boldsymbol{\theta}) \mid \mathbf{Y}\}=\frac{\int h(\boldsymbol{\theta}) f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d \boldsymbol{\theta}}{\int f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d \boldsymbol{\theta}}
$$

Suppose that we can roughly approximate the normalized likelihood times prior, $c f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta})$, by some density $g(\boldsymbol{\theta})$ from which we can easily sample. Define weight function $w(\boldsymbol{\theta})=f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) / g(\boldsymbol{\theta})$

$$
E\{h(\boldsymbol{\theta}) \mid \mathbf{Y}\}=\frac{\int h(\boldsymbol{\theta}) w(\boldsymbol{\theta}) g(\boldsymbol{\theta}) d \boldsymbol{\theta}}{\int w(\boldsymbol{\theta}) g(\boldsymbol{\theta}) d \boldsymbol{\theta}} \approx \frac{N^{-1} \sum_{j=1}^{N} h\left(\boldsymbol{\theta}_{j}\right) w\left(\boldsymbol{\theta}_{j}\right)}{N^{-1} \sum_{j=1}^{N} w\left(\boldsymbol{\theta}_{j}\right)}
$$

where $\boldsymbol{\theta}_{j} \stackrel{\text { iid }}{\sim} g(\boldsymbol{\theta})$, the importance function.
The performance depends on how close $g(\boldsymbol{\theta})$ resembles $c f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta})$.
2. Rejection sampling

Consider posterior sampling

$$
\pi(\boldsymbol{\theta} \mid \mathbf{Y})=\frac{f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta})}{\int f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d \boldsymbol{\theta}}
$$

Suppose $\exists M>0$ and a smooth density $g(\boldsymbol{\theta})$ - envelope function - s.t. $f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta})<M g(\boldsymbol{\theta})$ then
(a) Generate $\boldsymbol{\theta}_{j} \sim g(\boldsymbol{\theta})$
(b) Generate $U \sim \operatorname{Uniform}(0,1)$
(c) If $M U g\left(\boldsymbol{\theta}_{j}\right)<f\left(\mathbf{Y} \mid \boldsymbol{\theta}_{j}\right) \pi\left(\boldsymbol{\theta}_{j}\right)$, accept $\boldsymbol{\theta}_{j}$; otherwise reject $\boldsymbol{\theta}_{j}$.
(d) Return to step (i) and repeat until a desired size is obtained. The members of this sample is an iid sample from $\pi(\boldsymbol{\theta} \mid \mathbf{Y})$

Figure 3: Graphical representation of the rejection sampling method


The third condition ensures that the density of the sampled $\theta$ is within $[A, B]$. Therefore, more likely to be from the posterior distribution
3. Weighted bootstrap

Again consider the posterior density. Suppose that we have a sample $\boldsymbol{\theta}_{1}, \cdots, \boldsymbol{\theta}_{N}$ from approximating
density $g(\boldsymbol{\theta})$.
Define

$$
w_{i}=\frac{f\left(\mathbf{Y} \mid \boldsymbol{\theta}_{i}\right) \pi\left(\boldsymbol{\theta}_{i}\right)}{g\left(\boldsymbol{\theta}_{i}\right)}, \quad q_{i}=\frac{w_{i}}{\sum_{j=1}^{N} w_{j}}
$$

Draw $\boldsymbol{\theta}^{*}$ from the discrete distribution over $\left\{\boldsymbol{\theta}_{1}, \cdots, \boldsymbol{\theta}_{N}\right\}$ which places mass $q_{i}$ at $\boldsymbol{\theta}_{i}$.
In these indirect sampling methods, the prior $\pi(\boldsymbol{\theta})$, if proper, can play a role.
Example 1: Suppose $Y_{1}, \cdots, Y_{n} \stackrel{\text { iid }}{\sim} N\left(\theta, \sigma^{2}\right)$ and $\pi(\theta)=$ Cauchy $(\mu, \tau)$ with known $\sigma^{2}, \mu$, and $\tau$. Likelihood $f(\mathbf{Y} \mid \theta)$ is maximized at $\widehat{\theta}=\bar{Y} \Longrightarrow M=f(\mathbf{Y} \mid \widehat{\theta})$ in the rejection method, and $g(\theta)=\pi(\theta)$. However, $\pi(\boldsymbol{\theta})$ is often very flat relative to $f(\mathbf{Y} \mid \theta) \Rightarrow$ quite inefficient.

### 4.9 MCMC methods

### 4.9.1 Substitution sampling

Consider the 3 -stage hierarchical model

$$
\begin{aligned}
& \text { likelihood } p(Y \mid \theta) \\
& \text { prior } p(\theta \mid \eta) \\
& \text { hyperprior } p(\eta)
\end{aligned}
$$

Assuming the prior is conjugate with the likelihood $\Longrightarrow$ marginal distribution $p(Y \mid \eta)=\int p(Y \mid \theta) p(\theta \mid \eta) d \theta$ easily computed $\Longrightarrow$ closed-form posterior $p(\theta \mid Y, \eta)=\frac{p(Y \mid \theta) p(\theta \mid \eta)}{p(Y \mid \eta)}$
Also assuming the hyperprior is conjugate with the prior $\Longrightarrow$ closed-form $p(\eta \mid \theta)$ For inference, we seek the marginal posterior,

$$
p(\theta \mid Y)=\int p(\theta \mid Y, \eta) p(\eta \mid Y) d \eta
$$

and we also have

$$
p(\eta \mid Y)=\int p(\eta \mid \theta) p(\theta \mid Y) d \theta
$$

Switching to generic notation, we have a system of two linear integral equations

$$
\begin{aligned}
& p(x)=\int p(x \mid y) p(y) d y \\
& p(y)=\int p(y \mid x) p(x) d x
\end{aligned}
$$

where $p(x \mid y)$ and $p(y \mid x)$ are known, and we seek $p(x)$. This is a fixed point system:

$$
\begin{aligned}
p(x) & =\int p(x \mid y) \int p\left(y \mid x^{\prime}\right) p\left(x^{\prime}\right) d x^{\prime} d y \\
& =\int h\left(x, x^{\prime}\right) p\left(x^{\prime}\right) d x^{\prime}
\end{aligned}
$$

where $h\left(x, x^{\prime}\right)=\int p(x \mid y) p\left(y \mid x^{\prime}\right) d y$. Sampling-based algorithm:
Draw $X^{(0)} \sim p_{0}(x)$
Draw $Y^{(1)} \sim p\left(y \mid X^{(0)}\right) \sim p_{1}(y)=\int p(y \mid x) p_{0}(x) d x$
Draw $X^{(1)} \sim p\left(x \mid Y^{(1)}\right) \sim p_{1}(x)=\int h\left(x, x^{\prime}\right) p_{0}\left(x^{\prime}\right) d x^{\prime}$
Repeat this process:

$$
X^{(i)} \xrightarrow{d} X \sim p(x) \quad Y^{(i)} \xrightarrow{d} Y \sim p(y)
$$

A variant: multiple $Y_{1}^{(i)}, \cdots, Y_{m}^{(i)} \stackrel{\text { iid }}{\sim} p\left(y \mid X^{(i-1)}\right)$ and single

$$
X^{(i)} \sim \widehat{p}_{i}(x)=\frac{1}{m} \sum_{j=1}^{m} p\left(x \mid Y_{j}^{(i)}\right)
$$

$\Longrightarrow$ automatically produce a smooth estimate of $p(x)$
Parallel sampling with multiple chains: Marginally independent replicates $\Longrightarrow$ presumably better $\widehat{p}_{i}(x)$. But wasteful.
Ergodic sampling with a single chain: Continue for an additional $m-1$ iterations after convergence at iteration $i$.

$$
\widehat{p}_{i}(x)=\frac{1}{m} \sum_{j=1}^{m} p\left(x \mid Y_{j}^{(i+j-1)}\right)
$$

To reduce high autocorrelation, retain only every $k$ th iteration:

$$
\widehat{p}_{i}(x)=\frac{1}{m} \sum_{j=1}^{m} p\left(x \mid Y_{j}^{(i+(j-1) k)}\right)
$$

Although $X$ and $Y$ can conceivably be vectors, sampling from complex multivariate distributions is difficult $\Longrightarrow$ Require a $K$-variate extension. Consider $K=3$ :

$$
\begin{aligned}
p(x) & =\int p(x, z \mid y) p(y) d y d z \\
p(y) & =\int p(y, x \mid z) p(z) d x d z \\
p(z) & =\int p(z, y \mid x) p(x) d x d y
\end{aligned}
$$

Closed-form bivariate conditional distributions are unlikely available. With univariate distributions only:

$$
\begin{aligned}
& p(x)=\int p(x \mid z, y) p(z \mid y) p(y) d y d z \\
& p(y)=\int p(y \mid x, z) p(x \mid z) p(z) d x d z \\
& p(z)=\int p(z \mid y, x) p(y \mid x) p(x) d x d y
\end{aligned}
$$

6 of them!
General $K$-dimensional problem $\Longrightarrow K(K-1)$ distributions $\Longrightarrow$ impractical for large $K$.

### 4.9.2 Gibbs sampling

Gibbs sampler can be viewed as a special case of Metropolis-Hastings. And Metropolis-Hastings algorithm is a generalization of the Metropolis algorithm.
Under mild conditions, between full or complete conditional distributions and joint distribution

$$
\left\{p_{i}\left(U_{i} \mid U_{j \neq i}\right), i=1, \cdots, K\right\} \quad \Longleftrightarrow \quad p\left(U_{1}, \cdots, U_{K}\right)
$$

Suppose these full conditional distributions are available for sampling. Given an arbitrary set of starting values $\left(U_{1}^{(0)}, \cdots, U_{K}^{(0)}\right)$, the algorithm proceeds as follows:

$$
\begin{aligned}
& \text { Draw } U_{1}^{(1)} \sim p_{1}\left(U_{1} \mid U_{2}^{(0)}, \cdots, U_{K}^{(0)}\right) \\
& \operatorname{Draw} U_{2}^{(1)} \sim p_{2}\left(U_{2} \mid U_{1}^{(1)}, U_{3}^{(0)}, \cdots, U_{K}^{(0)}\right) \\
& \operatorname{Draw} U_{K}^{(1)} \sim p_{K}\left(U_{K} \mid U_{1}^{(1)}, \cdots, U_{K-1}^{(1)}\right)
\end{aligned}
$$

completing one iteration of the scheme. After $t$ such iterations, we obtain $\left(U_{1}^{(t)}, \cdots, U_{K}^{(t)}\right)$.

$$
\left(U_{1}^{(t)}, \cdots, U_{K}^{(t)}\right) \xrightarrow{d}\left(U_{1}, \cdots, U_{K}\right) \sim p\left(U_{1}, \cdots, U_{K}\right)
$$

Hierarchical models with conjugate priors and hyperpriors $\Longrightarrow$ closed form $\left\{p_{i}\left(U_{i} \mid U_{j \neq i}, \mathbf{Y}\right), i=1, \cdots, K\right\}$ $\Longrightarrow$ joint posterior $p\left(U_{1}, \cdots, U_{K} \mid \mathbf{Y}\right)$ by Gibbs sampler.
Marginal posterior $p\left(U_{i} \mid \mathbf{Y}\right)$ can be obtained the same manner as before. For example, with parallel sampling,

$$
\widehat{p}_{t}\left(U_{i} \mid \mathbf{Y}\right)=\frac{1}{m} \sum_{j=1}^{m} p\left(U_{i} \mid U_{1, j}^{(t)}, \cdots, U_{i-1, j}^{(t)}, U_{i+1, j}^{(t)}, \cdots, U_{K, j}^{(t)}, \mathbf{Y}\right)
$$

which is less variable than a kernel-smoothed estimate.
What about non-conjugate priors? One may consider, for example, an indirect sampling method. But such a solution is not ideal.

### 4.9.3 Metropolis-Hastings algorithm

Target: joint posterior $p(u)$ for (possibly vector-valued) $U$
Candidate or proposal density: $q(v, u)$ such that $q(\cdot, u)$ is a $\operatorname{pdf} \forall u$, and $q(u, v)=q(v, u) \forall u, v$.
The Metropolis algorithm:

1. Draw $v \sim q\left(\cdot, U^{(t-1)}\right)$
2. Compute the density ratio $r=p(v) / p(u)$
3. If $r \geq 1$, set $U^{(t)}=v$

If $r<1$, set $U^{(t)}=\left\{\begin{array}{l}v \text { with prob } r \\ u \text { with prob } 1-r\end{array}\right.$
Mild conditions $\Longrightarrow U^{(t)} \xrightarrow{d} U \sim p(\cdot)$.
$p(\cdot)$ is needed only up to proportionality constant:

$$
\pi(\boldsymbol{\theta} \mid \mathbf{Y}) \propto f(\mathbf{Y} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta})
$$

in Bayesian applications.
Why symmetric $q$ ?
Consider finite state spaces, where the transition kernel can be represented by matrix $\mathbf{P} ; P_{i j}$ is the prob of moving from state $i$ to $j$.
The Markov chain has equilibrium distribution

$$
\mathbf{d}=\left(d_{1}, \cdots, d_{k}\right)^{T}
$$

if and only if

$$
\mathbf{d}^{T} \mathbf{P}=\mathbf{d}^{T}
$$

Symmetric $q \Longrightarrow$

$$
\begin{aligned}
d_{i} P_{i j} & =d_{i}\left[\min \left(1, \frac{d_{j}}{d_{i}}\right) Q_{i j}\right]=\min \left(d_{i}, d_{j}\right) Q_{i j} \\
& =\min \left(d_{i}, d_{j}\right) Q_{j i}=d_{j} P_{j i}
\end{aligned}
$$

$\Longrightarrow$ chain is reversible $\Longrightarrow$

$$
\left(\mathbf{d}^{T} P\right)_{j}=\sum_{i=1}^{k} d_{i} P_{i j}=\sum_{i=1}^{k} d_{j} P_{j i}=d_{j} \sum_{i=1}^{k} P_{j i}=d_{j}
$$

For the continuous parameter settings, a convenient choice for $q$ is a $N\left(\boldsymbol{\theta}^{(i-1)}, \widetilde{\boldsymbol{\Sigma}}\right)$. In theory, any positivedefinite $\widetilde{\Sigma}$ suffices. Care is required in practice:

- too large $\Longrightarrow$ large jumps $\Longrightarrow$ many candidates far from posterior support and rejected $\Longrightarrow$ tendency to "get stuck"
- too small $\Longrightarrow$ "baby-stepping"

A simple but important generalization of the Metropolis algorithm was due to Hastings, by dropping the symmetry requirement of $q$ and redefining

$$
r=\frac{p(v) q(u, v)}{p(u) q(v, u)}
$$

## From Yale BIS 567

Metropolis Algorithm
Define $h(\theta) \equiv f(y \mid \theta) f(\theta)$ and $q\left(\boldsymbol{\theta}^{*} \mid \boldsymbol{\theta}^{(t-1)}\right)$ is a candidate or proposal density. It must satisfy

1. valid density function for every possible value of $\boldsymbol{\theta}^{(t-1)}$
2. be symmetric $q\left(\boldsymbol{\theta}^{*} \mid \boldsymbol{\theta}^{(t-1)}\right)=q\left(\boldsymbol{\theta}^{(t-1)} \mid \boldsymbol{\theta}^{*}\right)$. For example

- $q\left(\boldsymbol{\theta}^{*} \mid \boldsymbol{\theta}^{(t-1)}\right)=$ Uniform $\left(\boldsymbol{\theta}^{(t-1)}-\delta, \boldsymbol{\theta}^{(t-1)}+\delta\right)$
- $q\left(\boldsymbol{\theta}^{*} \mid \boldsymbol{\theta}^{(t-1)}\right)=\operatorname{Normal}\left(\boldsymbol{\theta}^{(t-1)}, \delta^{2}\right)$
$\delta$ is chosen to make the algorithm run more efficiently. A poor choice of $\delta$ leads to high correlation in the Markov chain.
(a) $\delta$ too big: Low acceptance of a new $\theta$ and high correlation between samples
(b) $\delta$ too small: High acceptance of a new $\theta$ and high correlation between samples

Generally, acceptance around $20 \%$ to $50 \%$ is desired. Pilot runs are often required to properly tune the algorithm and select a reasonable $\delta$.

The algorithm: for $t=1, \ldots, T$ :

1. draw $\theta^{*}$ from $q\left(. \mid \theta^{(t-1)}\right)$
2. compute $r=h\left(\boldsymbol{\theta}^{*}\right) / h\left(\boldsymbol{\theta}^{(t-1)}\right)=\exp \left[\log \left\{h\left(\boldsymbol{\theta}^{*}\right)\right\}-\log \left\{h\left(\boldsymbol{\theta}^{(t-1)}\right)\right\}\right]$
3. if $r \geq 1$, set $\boldsymbol{\theta}^{(t)}=\boldsymbol{\theta}^{*}$; if $r<1$, set

$$
\boldsymbol{\theta}^{(t)}= \begin{cases}\boldsymbol{\theta}^{*} & \text { with probability } r \\ \boldsymbol{\theta}^{(t-1)} & \text { with probability } 1-r\end{cases}
$$

Under mild conditions (similar to those for the Gibbs sampler), a draw $\boldsymbol{\theta}^{(t)}$ converges in distribution to a draw from the true posterior density $f(\boldsymbol{\theta} \mid \boldsymbol{Y}=\boldsymbol{y})$.
Intuition:
Suppose we have a working set of posterior samples $\left\{\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(s)}\right\}$. We want to add a new sample to the set, $\theta^{*}$, which is nearby $\theta^{(t)}$. If $f\left(\boldsymbol{\theta}^{*} \mid \boldsymbol{Y}=\boldsymbol{y}\right)>f\left(\boldsymbol{\theta}^{(s)} \mid \boldsymbol{Y}=\boldsymbol{y}\right)$, then we want to add $\theta^{*}$ to the set. If $f\left(\boldsymbol{\theta}^{*} \mid \boldsymbol{Y}=\boldsymbol{y}\right)<f\left(\boldsymbol{\theta}^{(s)} \mid \boldsymbol{Y}=\boldsymbol{y}\right)$, then we don't necessarily want to add it to the set. So we accept it with probability $r$

$$
r=\frac{f\left(\boldsymbol{\theta}^{*} \mid \boldsymbol{Y}=\boldsymbol{y}\right)}{f\left(\boldsymbol{\theta}^{(s)} \mid \boldsymbol{Y}=\boldsymbol{y}\right)}=\frac{f\left(y \mid \theta^{*}\right) f\left(\theta^{*}\right)}{f(y)} \frac{f(y)}{f\left(y \mid \theta^{(s)}\right) f\left(\theta^{(s)}\right)}=\frac{f\left(y \mid \theta^{*}\right) f\left(\theta^{*}\right)}{f\left(y \mid \theta^{(s)}\right) f\left(\theta^{(s)}\right)}
$$

### 4.9.4 Hybrid forms

Several MCMC algorithms may be combined in a single problem, to take advantage of the strengths of each. Markov kernels $P_{1}, \cdots, P_{m}$ all have stationary distribution $p$; they may correspond to, say, one Gibbs sampler and $m-1$ Metropolis algorithms.
Mixer: At each iteration, kernel $P_{i}$ is chosen with prob $\alpha_{i}$, where $\sum_{i} \alpha_{i}=1$.
Cycle: Each kernel $P_{i}$ is used in a prespecified order.
Metropolis within Gibbs: In the case that, say, all full conditional distributions are available in closed form except for one, a Metropolis subalgorithm may be embedded in the Gibbs sampler.

## 5 Large sample theory

## 6 M-Estimation (Estimating Equations)

### 6.1 Introduction

$$
\begin{aligned}
\sum_{i=1}^{n} \boldsymbol{\Psi}\left(Y_{i} ; \boldsymbol{\theta}\right) & =\mathbf{0} \\
Y_{1}, \cdots, Y_{n} & : \text { independent but not necessarily identically distributed } \\
\boldsymbol{\theta} & : b \times 1 \text { parameter } \\
\boldsymbol{\Psi}: b & \times 1 \text { function }
\end{aligned}
$$

Example 1. the mean estimator $\widehat{\theta}=n^{-1} \sum_{i=1}^{n} Y_{i}$ is from

$$
\sum_{i=1}^{n}\left(Y_{i}-\theta\right)=0
$$

the deviance from the mean $\widehat{\theta}_{1}=n^{-1} \sum_{i=1}^{n}\left|Y_{i}-\bar{Y}\right|$ is from

$$
\sum_{i=1}^{n}\binom{\left|Y_{i}-\theta_{2}\right|-\theta_{1}}{Y_{i}-\theta_{2}}=\mathbf{0}
$$

### 6.2 Basic approach

Suppose $Y_{1}, \cdots, Y_{n} \stackrel{\text { iid }}{\sim} F$ and the true parameter $\theta_{0}$ is defined as a solution to

$$
E_{F} \boldsymbol{\Psi}\left(Y_{1}, \boldsymbol{\theta}\right)=\int \boldsymbol{\Psi}(y, \boldsymbol{\theta}) d F(y)=\mathbf{0}
$$

if $\Psi$ is suitably smooth then by Taylor expansion of

$$
\mathbf{G}_{n}(\boldsymbol{\theta})=n^{-1} \sum_{i=1}^{n} \boldsymbol{\Psi}\left(Y_{i}, \boldsymbol{\theta}\right)
$$

we have

$$
\begin{aligned}
0=\mathbf{G}_{n}(\widehat{\boldsymbol{\theta}}) & =\mathbf{G}_{n}\left(\boldsymbol{\theta}_{0}\right)+\mathbf{G}_{n}^{\prime}\left(\boldsymbol{\theta}_{0}\right)\left(\widehat{\boldsymbol{\theta}}-\boldsymbol{\theta}_{0}\right)+o_{p}\left(n^{-1 / 2}\right) \\
\sqrt{n}\left(\widehat{\boldsymbol{\theta}}-\boldsymbol{\theta}_{0}\right) & =\left\{-\mathbf{G}_{n}^{\prime}\left(\boldsymbol{\theta}_{0}\right)\right\}^{-1} \sqrt{n} \mathbf{G}_{n}\left(\boldsymbol{\theta}_{0}\right)+o_{p}(1)
\end{aligned}
$$

by WLLN

$$
-\mathbf{G}_{n}^{\prime}\left(\boldsymbol{\theta}_{0}\right)=-n^{-1} \sum_{i=1}^{n} \boldsymbol{\Psi}^{\prime}\left(Y_{i}, \boldsymbol{\theta}_{0}\right) \xrightarrow{p} E_{F}\left\{-\boldsymbol{\Psi}^{\prime}\left(Y_{1}, \boldsymbol{\theta}_{0}\right)\right\} \equiv \mathbf{A}\left(\boldsymbol{\theta}_{0}\right)
$$

by CLT

$$
\sqrt{n} \mathbf{G}_{n}\left(\boldsymbol{\theta}_{0}\right) \xrightarrow{d} N\left(\mathbf{0}, \mathbf{B}\left(\boldsymbol{\theta}_{0}\right)\right) \text { where } \mathbf{B}\left(\boldsymbol{\theta}_{0}\right)=E_{F}\left\{\boldsymbol{\Psi}\left(Y_{1}, \boldsymbol{\theta}_{0}\right)^{\otimes 2}\right\}
$$

then by Slutsky's theorem

$$
\widehat{\boldsymbol{\theta}} \sim A N\left(\boldsymbol{\theta}_{0}, \mathbf{V}\left(\boldsymbol{\theta}_{0}\right) / n\right) \text { where } \mathbf{V}\left(\boldsymbol{\theta}_{0}\right)=\mathbf{A}\left(\boldsymbol{\theta}_{0}\right)^{-1} \mathbf{B}\left(\boldsymbol{\theta}_{0}\right)\left\{\mathbf{A}\left(\boldsymbol{\theta}_{0}\right)^{-1}\right\}^{T}
$$

More generally, the estimating equation does not have to be exactly 0 , for example

$$
\sum_{i=1}^{n} \boldsymbol{\Psi}\left(Y_{i}, \boldsymbol{\theta}\right)=o_{p}\left(n^{1 / 2}\right) \Rightarrow \frac{1}{n} \sum_{i=1}^{n} \boldsymbol{\Psi}\left(Y_{i}, \boldsymbol{\theta}\right)=o_{p}\left(n^{-1 / 2}\right)
$$

The asymptotic distribution of $\hat{\theta}$ remains the same.

### 6.2.1 Estimation for $A, B$, and $V$

Empirical estimator of $\mathbf{A}\left(\boldsymbol{\theta}_{0}\right)$ :

$$
\mathbf{A}_{n}(\mathbf{Y}, \widehat{\boldsymbol{\theta}})=-n^{-1} \sum_{i=1}^{n} \boldsymbol{\Psi}^{\prime}\left(Y_{i}, \widehat{\boldsymbol{\theta}}\right)
$$

Empirical estimator of $\mathbf{B}\left(\boldsymbol{\theta}_{0}\right)$ :

$$
\mathbf{B}_{n}(\mathbf{Y}, \widehat{\boldsymbol{\theta}})=n^{-1} \sum_{i=1}^{n} \boldsymbol{\Psi}\left(Y_{i}, \widehat{\boldsymbol{\theta}}\right)^{\otimes 2}
$$

then

$$
\mathbf{V}_{n}(\mathbf{Y}, \widehat{\boldsymbol{\theta}})=\mathbf{A}_{n}(\mathbf{Y}, \widehat{\boldsymbol{\theta}})^{-1} \mathbf{B}_{n}(\mathbf{Y}, \widehat{\boldsymbol{\theta}})\left\{\mathbf{A}_{n}(\mathbf{Y}, \widehat{\boldsymbol{\theta}})^{-1}\right\}^{T}
$$

A special case is the MLE, where

$$
\mathbf{A}\left(\boldsymbol{\theta}_{0}\right)=\mathbf{B}\left(\boldsymbol{\theta}_{0}\right)=\mathbf{I}\left(\boldsymbol{\theta}_{0}\right) \quad \mathbf{V}\left(\boldsymbol{\theta}_{0}\right)=\mathbf{I}\left(\boldsymbol{\theta}_{0}\right)^{-1}
$$

Example 2. Sample mean and variance

$$
\begin{aligned}
\boldsymbol{\Psi}\left(Y_{i}, \boldsymbol{\theta}\right) & =\binom{Y_{i}-\theta_{1}}{\left(Y_{i}-\theta_{1}\right)^{2}-\theta_{2}} \\
\mathbf{A}\left(\boldsymbol{\theta}_{0}\right) & =E\left\{-\boldsymbol{\Psi}^{\prime}\left(Y_{1}, \boldsymbol{\theta}_{0}\right)\right\}=E\left(\begin{array}{cc}
1 & 0 \\
2\left(Y_{1}-\theta_{10}\right) & 1
\end{array}\right)=\left(\begin{array}{ll}
1 & 0 \\
0 & 1
\end{array}\right) \\
\mathbf{B}\left(\boldsymbol{\theta}_{0}\right) & =E\left\{\boldsymbol{\Psi}\left(Y_{1}, \boldsymbol{\theta}_{0}\right)^{\otimes 2}\right\}=\left(\begin{array}{cc}
\sigma^{2} & \mu_{3} \\
\mu_{3} & \mu_{4}-\sigma^{4}
\end{array}\right) \\
\mathbf{V}\left(\boldsymbol{\theta}_{0}\right) & =\mathbf{B}\left(\boldsymbol{\theta}_{0}\right) \\
\mathbf{B}_{n}(\mathbf{Y}, \widehat{\boldsymbol{\theta}}) & =n^{-1} \sum_{i=1}^{n}\left(\begin{array}{cc}
\left(Y_{i}-\bar{Y}\right)^{2} & \left(Y_{i}-\bar{Y}\right)^{3} \\
\left(Y_{i}-\bar{Y}\right)^{3} & \left\{\left(Y_{i}-\bar{Y}\right)^{2}-s_{n}^{2}\right\}^{2}
\end{array}\right)
\end{aligned}
$$

the estimated $\widehat{\boldsymbol{\theta}}=\left(\bar{Y}, s_{n}^{2}\right)^{T}$ is the MLE for the normal density $f(y ; \boldsymbol{\theta})=\frac{1}{\sqrt{2 \pi \theta_{2}}} \exp \left\{-\frac{\left(y-\theta_{1}\right)^{2}}{2 \theta_{2}}\right\}$

## Example 3. Ratio estimator

For i.i.d $\left(Y_{1}, X_{1}\right), \cdots,\left(Y_{n}, X_{n}\right)$ and $\mu_{X} \neq 0$

$$
\begin{aligned}
\Psi\left(Y_{i}, X_{i}, \theta\right) & =Y_{i}-\theta X_{i} \\
A\left(\theta_{0}\right) & =E\left(X_{1}\right)=\mu_{X} \\
B\left(\theta_{0}\right) & =E\left\{\left(Y_{1}-\theta_{0} X_{1}\right)^{2}\right\} \\
V\left(\theta_{0}\right) & =E\left\{\left(Y_{1}-\theta_{0} X_{1}\right)^{2}\right\} / \mu_{X}^{2} \\
A_{n}(\mathbf{Y}, \widehat{\theta}) & =\bar{X} \\
B_{n}(\mathbf{Y}, \widehat{\theta}) & =\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\frac{\bar{Y}}{\bar{X}} X_{i}\right)^{2} \\
V_{n}(\mathbf{Y}, \widehat{\theta}) & =\frac{1}{\bar{X}^{2}} \frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\frac{\bar{Y}}{\bar{X}} X_{i}\right)^{2}
\end{aligned}
$$

The $\Psi$ functions are not unique, it can even be of different dimensions

$$
\begin{aligned}
\boldsymbol{\Psi}\left(Y_{i}, X_{i}, \boldsymbol{\theta}\right) & =\left(\begin{array}{c}
Y_{i}-\theta_{1} \\
X_{i}-\theta_{2} \\
\theta_{1}-\theta_{3} \theta_{2}
\end{array}\right) \\
\mathbf{A}\left(\boldsymbol{\theta}_{0}\right) & =\left(\begin{array}{ccc}
1 & 0 & 0 \\
0 & 1 & 0 \\
-1 & \theta_{30} & \theta_{20}
\end{array}\right) \\
\mathbf{B}\left(\boldsymbol{\theta}_{0}\right) & =\left(\begin{array}{ccc}
\sigma_{Y}^{2} & \sigma_{Y X} & 0 \\
\sigma_{Y X} & \sigma_{X}^{2} & 0 \\
0 & 0 & 0
\end{array}\right) \text { singular } \\
\mathbf{V}\left(\boldsymbol{\theta}_{0}\right) & =\mathbf{A}\left(\boldsymbol{\theta}_{0}\right)^{-1} \mathbf{B}\left(\boldsymbol{\theta}_{0}\right)\left\{\mathbf{A}\left(\boldsymbol{\theta}_{0}\right)^{-1}\right\}^{T} \\
v_{33} & =\frac{1}{\theta_{20}^{2}}\left(\sigma_{Y}^{2}-2 \theta_{30} \sigma_{Y X}+\theta_{30}^{2} \sigma_{X}^{2}\right)
\end{aligned}
$$

### 6.3 Delta method via M-estimation

