

Statistics 582, Problem Set 1 Solutions

Wellner; 1/17/2007

1. Suppose that $(X, Y), (X_1, Y_1), \dots, (X_n, Y_n)$ are i.i.d. with bivariate normal distribution $N_2(\mu, \Sigma)$ where $\mu \in R^2$ and

$$\Sigma = \begin{pmatrix} \sigma^2 & \sigma\tau\rho \\ \sigma\tau\rho & \tau^2 \end{pmatrix}$$

where $\sigma^2 > 0$, $\tau^2 > 0$, and $\rho \in (-1, 1)$.

- (a) If we assume that $\mu_1 = \mu_2 \equiv \theta$ and Σ is known, what is the MLE of θ ?
- (b) If we assume that μ is known and $\sigma^2 = \tau^2 \equiv \theta$, what is the MLE of θ ?
- (c) What is the asymptotic distribution of the estimator you found in (b)?
- (d) Under the same assumption as in (b), what is the MLE of ρ ?
- (e) What is the asymptotic distribution of the estimator you found in (d)?

Solution: (a) When $\mu_1 = \mu_2 = \theta$ and Σ is known, then the log-likelihood for one observation is (relabelling $\mu_1 = \mu$, $\mu_2 = \nu$),

$$\log p(x; \theta) = -\frac{1}{2(1-\rho^2)} \left\{ \frac{(x-\theta)^2}{\sigma^2} - 2\rho \frac{(x-\theta)(y-\theta)}{\sigma\tau} + \frac{(y-\theta)^2}{\tau^2} \right\} + \text{constant}.$$

Hence the score for θ for one observation is

$$\begin{aligned} \dot{l}_\theta(x, y) &= \frac{1}{1-\rho^2} \left\{ \frac{(x-\theta)}{\sigma^2} + \frac{(y-\theta)}{\tau^2} - \frac{\rho}{\tau} \frac{(x-\theta)}{\sigma} - \frac{\rho}{\sigma} \frac{(y-\theta)}{\tau} \right\} \\ &= \frac{1}{1-\rho^2} \left\{ \frac{(x-\theta)}{\sigma} \left(\frac{1}{\sigma} - \frac{\rho}{\tau} \right) + \frac{(y-\theta)}{\tau} \left(\frac{1}{\tau} - \frac{\rho}{\sigma} \right) \right\}. \end{aligned}$$

Thus the score equation for θ is

$$0 = \dot{l}_{n\theta}(\theta) = \frac{n}{1-\rho^2} \left\{ \frac{(\bar{X}_n - \theta)}{\sigma} \left(\frac{1}{\sigma} - \frac{\rho}{\tau} \right) + \frac{(\bar{Y}_n - \theta)}{\tau} \left(\frac{1}{\tau} - \frac{\rho}{\sigma} \right) \right\},$$

and hence

$$\begin{aligned} \hat{\theta}_n &= \frac{\bar{X}_n \left(\frac{1}{\sigma^2} - \frac{\rho}{\sigma\tau} \right) + \bar{Y}_n \left(\frac{1}{\tau^2} - \frac{\rho}{\sigma\tau} \right)}{\frac{1}{\sigma^2} - \frac{2\rho}{\sigma\tau} + \frac{1}{\tau^2}} \\ &= a\bar{X}_n + (1-a)\bar{Y}_n \end{aligned}$$

where

$$a = \frac{\frac{1}{\sigma^2} - \frac{\rho}{\sigma\tau}}{\frac{1}{\sigma^2} - \frac{2\rho}{\sigma\tau} + \frac{1}{\tau^2}}.$$

Note that this yields

$$\text{Var}(\hat{\theta}) = \frac{1}{n} \left\{ a^2\sigma^2 + 2a(1-a)\rho\sigma\tau + (1-a)^2\tau^2 \right\}.$$

(b) and (d) If $\sigma^2 = \tau^2 = \theta$ and μ is known, then the log-likelihood for one observation is (again relabelling $\mu_1 = \mu$, $\mu_2 = \nu$),

$$\begin{aligned} \log p(x; \theta, \rho) &= -\log \theta - \frac{1}{2} \log(1 - \rho^2) \\ &\quad - \frac{1}{2(1 - \rho^2)\theta} \{(x - \mu)^2 - 2\rho(x - \mu)(y - \nu) + (y - \nu)^2\} + \text{constant}. \end{aligned}$$

Thus the scores for θ and ρ are given by

$$\begin{aligned} \dot{l}_\theta(x, y) &= -\frac{1}{\theta} + \frac{1}{2(1 - \rho^2)\theta^2} \{(x - \mu)^2 - 2\rho(x - \mu)(y - \nu) + (y - \nu)^2\}, \\ \dot{l}_\rho(x, y) &= \frac{\rho}{(1 - \rho^2)} - \frac{\rho}{(1 - \rho^2)^2\theta} \{(x - \mu)^2 - 2\rho(x - \mu)(y - \nu) + (y - \nu)^2\} \\ &\quad + \frac{1}{\theta(1 - \rho^2)}(x - \mu)(y - \nu). \end{aligned}$$

Hence the score equations for estimation of θ and ρ are given by

$$0 = \dot{l}_{n\theta}(\theta) = \sum_{i=1}^n \dot{l}_\theta(X_i, Y_i) = -\frac{n}{\theta} + \frac{n}{\theta^2 2(1 - \rho^2)} \{S_{XX} - 2\rho S_{XY} + S_{YY}\},$$

and

$$0 = \dot{l}_{n\rho}(\rho) = \sum_{i=1}^n \dot{l}_\rho(X_i, Y_i) = \frac{n\rho}{1 - \rho^2} - \frac{n\rho}{\theta(1 - \rho^2)^2} \{S_{XX} - 2\rho S_{XY} + S_{YY}\} + \frac{n}{\theta(1 - \rho^2)} S_{XY}$$

where

$$S_{XX} \equiv n^{-1} \sum_{i=1}^n (X_i - \mu)^2, \quad S_{XY} \equiv n^{-1} \sum_{i=1}^n (X_i - \mu)(Y_i - \nu), \quad S_{YY} \equiv n^{-1} \sum_{i=1}^n (Y_i - \nu)^2.$$

Solving the first of these for $\hat{\theta}$ yields

$$\hat{\theta} = \frac{1}{2(1 - \hat{\rho}^2)} \{S_{XX} - 2\hat{\rho}S_{XY} + S_{YY}\};$$

Rewriting the score equation for ρ with a common denominator of $\theta(1 - \rho^2)^2$ yields

$$\theta\rho(1 - \rho^2) - \rho\{S_{XX} - 2\rho S_{XY} + S_{YY}\} + (1 - \rho^2)S_{XY} = 0;$$

and then plugging in the estimator $\hat{\theta}$ of θ yields the equation

$$(1 - \hat{\rho}^2)S_{XY} = \frac{1}{2}\hat{\rho}\{S_{XX} - 2\hat{\rho}S_{XY} + S_{YY}\}.$$

This has the solution

$$\hat{\rho} = \frac{2S_{XY}}{S_{XX} + S_{YY}};$$

plugging this (or more precisely the last form of the equation for $\widehat{\rho}$) into the expression for $\widehat{\theta}$ yields $\widehat{\theta} = (S_{XX} + S_{YY})/2$.

(c) and (e) To find the asymptotic distributions of $\widehat{\theta}$ and $\widehat{\rho}$ we could either (i) proceed directly from first principles (central limit theorems and the delta method), or (b) use theorem 4.1.5 concerning the asymptotic behavior of maximum likelihood estimators. I'll take the second route here. The first step in this direction is to compute the information matrix for (θ, ρ) . Now

$$\ddot{l}_{\theta\theta}(x, y) = \frac{1}{\theta^2} - \frac{1}{(1 - \rho^2)\theta^3} \{ (x - \mu)^2 - 2\rho(x - \mu)(y - \nu) + (y - \nu)^2 \},$$

$$\begin{aligned} \ddot{l}_{\theta\rho}(x, y) &= \frac{2\rho}{2\theta^2(1 - \rho^2)^2} \{ (x - \mu)^2 - 2\rho(x - \mu)(y - \nu) + (y - \nu)^2 \} \\ &\quad - \frac{1}{2\theta^2(1 - \rho^2)^2} 2(x - \mu)(y - \nu), \end{aligned}$$

and

$$\begin{aligned} \ddot{l}_{\rho\rho}(x, y) &= \frac{1}{1 - \rho^2} + \frac{2\rho^2}{(1 - \rho^2)^2} \\ &\quad - \left\{ \frac{1}{(1 - \rho^2)^2} + \frac{4\rho^2}{(1 - \rho^2)^3} \right\} \{ (x - \mu)^2 - 2\rho(x - \mu)(y - \nu) + (y - \nu)^2 \} \\ &\quad + \frac{\rho}{\theta(1 - \rho^2)^2} 2(x - \mu)(y - \nu) + \frac{2\rho}{\theta(1 - \rho^2)^2} (x - \mu)(y - \nu) \\ &= \frac{1 + \rho^2}{(1 - \rho^2)^2} - \frac{1 - 3\rho^2}{\theta(1 - \rho^2)^3} \{ (x - \mu)^2 - 2\rho(x - \mu)(y - \nu) + (y - \nu)^2 \} \\ &\quad + \frac{4\rho}{\theta(1 - \rho^2)^2} (x - \mu)(y - \nu). \end{aligned}$$

Here

$$E \{ (X - \mu)^2 - 2\rho(X - \mu)(Y - \nu) + (Y - \nu)^2 \} = 2\theta(1 - \rho^2)$$

and

$$E(X - \mu)(Y - \nu) = \rho\theta.$$

Thus we find that

$$\begin{aligned} I_{\theta\theta} &= E(-\ddot{l}_{\theta\theta}(X, Y)) = \theta^{-2}, \\ I_{\theta\rho} &= E(-\ddot{l}_{\theta\rho}(X, Y)) = \frac{\rho}{\theta(1 - \rho^2)}, \end{aligned}$$

and

$$I_{\rho\rho} = E(-\ddot{l}_{\rho\rho}(X, Y)) = \frac{1 + \rho^2}{(1 - \rho^2)^2}.$$

This yields

$$I_{\theta\theta \cdot \rho} = I_{\theta\theta} - I_{\theta\rho} I_{\rho\rho}^{-1} I_{\rho\theta} = \frac{1}{\theta^2} \frac{1}{1 + \rho^2}$$

and

$$I_{\rho\rho \cdot \theta} = I_{\rho\rho} - I_{\rho\theta} I_{\theta\theta}^{-1} I_{\theta\rho} = (1 - \rho^2)^{-2}.$$

Hence it follows from theorem 4.1.5 that

$$\sqrt{n}(\widehat{\theta}_n - \theta) \rightarrow_d N(0, I_{\theta\theta^{-1}}^{-1}) = N(0, \theta^2(1 + \rho^2))$$

while

$$\sqrt{n}(\widehat{\rho}_n - \rho) \rightarrow_d N(0, I_{\rho\rho^{-1}}^{-1}) = N(0, (1 - \rho^2)^2).$$

2. Problem 1, page 117, Ferguson, ACILST. What happens if $\Theta = [1, \infty)$ or $(0, \infty)$?

Solution: (i) Ferguson's problem: here is a verification of the five conditions (a)-(e) when $\Theta = [1, 2]$:

(a) $\Theta = [1, 2]$ is closed and bounded, hence compact.

(b) For fixed $x \leq 1$, $p(x, \theta) = 1/\theta$ is continuous; for $1 < x \leq 2$, $p(x, \theta) = \theta^{-1}1\{\theta \geq x\}$, which is upper semi-continuous. For $x > 2$, $p(x, \theta) = 0$ which is a continuous function of θ . Since log is a continuous function on $(0, \infty)$, these (semi-)continuities carry over to $f(x, \theta) = \log p(x, \theta) - \log p(x, \theta_0)$.

(c) Fix $\theta_0 \in \Theta$. Then

$$\frac{p(x, \theta)}{p(x, \theta_0)} = \frac{\theta_0}{\theta} \frac{1_{[0, \theta]}(x)}{1_{[0, \theta_0]}(x)} = \begin{cases} (\theta_0/\theta), & x \leq 1 \\ (\theta_0/\theta)1_{[x, \infty)}(\theta), & 1 < x \leq \theta_0 \\ \infty, & x > \theta_0, \end{cases}$$

so

$$\sup_{\theta \in \Theta} \frac{p(x, \theta)}{p(x, \theta_0)} \equiv K(x) = \theta_0 1\{x \leq 1\} + \frac{\theta_0}{x} 1_{(1, \theta_0]}(x) + \infty 1_{(\theta_0, \infty)}(x),$$

and

$$\sup_{\theta \in \Theta} \log \frac{p(x, \theta)}{p(x, \theta_0)} \equiv F(x) = (\log \theta_0) 1\{x \leq 1\} + \log(\theta_0/x) 1_{(1, \theta_0]}(x) + \infty 1_{(\theta_0, \infty)}(x),$$

satisfies $E_{\theta_0} F(X) < \infty$.

(d) The function $\varphi(x, \theta, \rho)$ is given by

$$\varphi(x, \theta, \rho) = \sup_{\theta': |\theta' - \theta| < \rho} \frac{1}{\theta'} 1_{[0, \theta']}(x) = \begin{cases} 1/(\theta - \rho), & x < \theta - \rho \\ 1/x, & |x - \theta| \leq \rho \\ 0, & x > \theta + \rho, \end{cases}$$

which is clearly measurable.

(e) If $p(x, \theta') = p(x, \theta)$ a.e. Lebesgue, then

$$0 = \frac{1}{2} \int |p(x, \theta) - p(x, \theta')| dx = d_{TV}(P_\theta, P_{\theta'}) = 1 - \eta(P_\theta, P_{\theta'}) = 1 - \frac{\theta' \wedge \theta}{\theta' \vee \theta}$$

where the last equality follows by a computation as in the final exam for stat 581, Fall '05. This yields $\theta' \vee \theta = \theta' \wedge \theta$, which implies $\theta = \theta'$.

(ii) When $\Theta = [1, \infty)$ or $\Theta = (0, \infty)$, then Θ is no longer compact, and Wald's theorem does not apply directly. One way to remedy the problem is to compactify the set Θ by some appropriate identification of points in Θ with points on the unit half-circle (as shown in class on 1/9/06). Another way to proceed is to show that the MLE is in a compact set eventually with probability 1; see e.g. van der Vaart's

re-working of Wald's theorem. In this case the MLE is $\hat{\theta}_n = \max_{1 \leq i \leq n} X_i = X_{(n)}$, and it follows that for any $\theta_0 > \delta > 0$

$$P_{\theta_0}(\theta_0 - \delta > \hat{\theta}_n) = \left(\frac{\theta_0 - \delta}{\theta_0} \right)^n,$$

which has a finite sum on n , and hence by the Borel-Cantelli lemma

$$P_{\theta_0}(\hat{\theta}_n < \theta_0 - \delta \text{ infinitely often}) = 0,$$

or, equivalently,

$$P_{\theta_0}(\theta_0 \geq \hat{\theta}_n \geq \theta_0 - \delta \text{ almost always}) = 1.$$

Note that this argument already yields almost sure consistency of the MLE in this case since δ can be chosen to be arbitrarily small.

3. Consider the model introduced in Ferguson, ACILST, problem 17.2, page 117. Show that Theorem 4.3, page 28, of the Chapter 4 notes (or Theorem 17, Ferguson, ACILST, page 114) applies to the MLE of θ in this model.

Solution: We proceed to verify the conditions (a) - (e) of Theorem 4.3 of the course notes. First (a) holds since $[0, 1]$ is compact. (b) holds since for $0 < x < 1$ the function $\theta \mapsto p(x, \theta)$ is continuous (and hence upper-semicontinuous), while for $x = 0$ the function $\theta \mapsto p(0, \theta) = 1_{\{0\}}(\theta)$ is upper-semicontinuous, and similarly for $x = 1$ the function $\theta \mapsto p(1, \theta) = 1_{\{1\}}(\theta)$ is also upper-semicontinuous. To see that (c) holds, consider the two cases $\theta > \theta_0$ and $\theta \leq \theta_0$. When $\theta > \theta_0$,

$$\begin{aligned} f(x, \theta) &= \log p(x, \theta) - \log p(x, \theta_0) \\ &= \begin{cases} \log(\theta_0/\theta) & 0 \leq x \leq \theta_0 < \theta \\ \log(x/(1-x)) - \log(\theta/(1-\theta_0)) & \theta_0 < x \leq \theta \\ \log((1-\theta_0)/(1-\theta)) & \theta < x \leq 1 \end{cases} \\ &\leq \begin{cases} \log(\theta_0/x) & 0 \leq x \leq \theta_0 < \theta \\ \log((1-\theta_0)/(1-x)) & \theta_0 < x \leq \theta \\ \log((1-\theta_0)/(1-x)) & \theta < x \leq 1 \end{cases} \\ &= 1_{[0, \theta_0]}(x) \log(\theta_0/x) + 1_{(\theta_0, 1]}(x) \log((1-\theta_0)/(1-x)) \\ &\equiv F(x). \end{aligned}$$

Here for the middle term we used

$$\begin{aligned} \log x - \log(1-x) + \log(1-\theta_0) + \log(1/\theta) &\leq \log x - \log(1-x) + \log(1-\theta_0) + \log(1/x) \\ &= \log((1-\theta_0)/(1-x)) \end{aligned}$$

since $1/\theta \leq 1/x$ on the set $\theta_0 < x \leq \theta$. Note that

$$E_{\theta_0} F(X) = 2 \int_0^{\theta_0} \frac{x}{\theta_0} \log\left(\frac{\theta_0}{x}\right) dx + 2 \int_{\theta_0}^1 \frac{1-x}{1-\theta_0} \log\left(\frac{1-\theta_0}{1-x}\right) dx < \infty.$$

Similarly, when $\theta \leq \theta_0$,

$$\begin{aligned}
f(x, \theta) &= \log p(x, \theta) - \log p(x, \theta_0) \\
&= \begin{cases} \log(\theta_0/\theta) & 0 \leq x \leq \theta \leq \theta_0 \\ \log((1-x)/x) - \log((1-\theta)/\theta_0) & \theta < x \leq \theta_0 \\ \log((1-\theta_0)/(1-\theta)) & \theta_0 < x \leq 1 \end{cases} \\
&\leq \begin{cases} \log(\theta_0/x) & 0 \leq x \leq \theta_0 < \theta \\ \log((1-\theta_0)/(1-x)) & \theta_0 < x \leq \theta \\ \log((1-\theta_0)/(1-x)) & \theta \leq \theta_0 < x \leq 1 \end{cases} \\
&= 1_{[0, \theta_0]}(x) \log(\theta_0/x) + 1_{(\theta_0, 1]}(x) \log((1-\theta_0)/(1-x)) \\
&= F(x),
\end{aligned}$$

so the same envelope function works in this case. To verify (d) note that $\theta \mapsto p(x, \theta)$ is a continuous function of θ for $0 < x < 1$, and two indicator functions of θ when $x \in \{0, 1\}$, and hence the supremum involved in the condition is measurable. Finally, the identifiability condition (e) holds easily: note that

$$\begin{aligned}
\rho(P_{\theta_0}, P_\theta) &= \int_0^1 \sqrt{p_{\theta_0}(x)p_\theta(x)} dx \\
&= \frac{\theta^2}{\sqrt{\theta\theta_0}} + \frac{2}{\sqrt{\theta_0(1-\theta)}} \int_\theta^{\theta_0} \sqrt{x(1-x)} dx + \frac{(1-\theta_0)^2}{\sqrt{(1-\theta_0)(1-\theta)}} \\
&= 1 \quad \text{if and only if} \quad \theta = \theta_0.
\end{aligned}$$

Thus $p_{\theta_0}(x) = p_\theta(x)$ a.e. Lebesgue implies that $\theta = \theta_0$. Thus the conditions (a) - (e) all hold and we conclude that the MLE $\hat{\theta}_n$ of θ is (almost surely) consistent: $\hat{\theta}_n \rightarrow_{a.s.} \theta$.

4. Suppose that X, X_1, \dots, X_n are i.i.d. Weibull(α_0, β_0) (if X has the Weibull(θ) distribution where $\theta = (\alpha, \beta)$, then $1 - F_\theta(x) = P_\theta(X > x) = \exp(-(x/\alpha)^\beta)$ for $x \geq 0$). Recall that the MLE $\hat{\alpha}$ of α is given by

$$\hat{\alpha} = \left\{ \frac{1}{n} \sum_{i=1}^n X_i^{\hat{\beta}} \right\}^{1/\hat{\beta}}$$

where $\hat{\beta}$ is the MLE of β . As a simpler alternative to maximum likelihood, I propose to use the alternative estimator $\bar{\beta}_n$ of β obtained from the slope of an ordinary least squares fit of a Weibull Q-Q plot, and then estimate α by

$$\bar{\alpha}_n = \left\{ \frac{1}{n} \sum_{i=1}^n X_i^{\bar{\beta}_n} \right\}^{1/\bar{\beta}_n}.$$

(a) Suppose that $\bar{\beta}_n \rightarrow_p \beta_0$ is known. Show that $\bar{\alpha}_n \rightarrow_p \alpha_0$. [Hint: use a uniform strong law of large numbers.]

(b) Show that $\bar{\alpha}_n$ is a “pseudo-MLE” in the sense that $\bar{\alpha}_n$ maximizes $l_n(\alpha, \bar{\beta}_n)$.

Solution: Fix $\delta > 0$ (small). The family of functions $\mathcal{F} = \{f(x, \beta) = x^\beta : \beta \in [\beta_0 - \delta, \beta_0 + \delta]\}$ are indexed by the compact set $[\beta_0 - \delta, \beta_0 + \delta]$, are continuous in β for every $x \geq 0$, and are bounded by

$$\sup_{\beta \in [\beta_0 - \delta, \beta_0 + \delta]} |f(x, \beta)| = x^{\beta_0 + \delta} \vee x^{\beta_0 - \delta} \leq x^{\beta_0 + \delta} + x^{\beta_0 - \delta} \equiv F(x)$$

which satisfies $E_0 F(X) < \infty$ if $\delta < 2\beta_0$. Thus by theorem 4.4.1 (of the section 4 revision) the uniform strong law of large numbers holds for \mathcal{F} :

$$\sup_{\beta: |\beta - \beta_0| \leq \delta} |\mathbb{P}_n f(\cdot, \beta) - P_0 f(\cdot, \beta)| \rightarrow_{a.s.} 0.$$

If $\bar{\beta}_n \rightarrow_{a.s.} \beta_0$, $\bar{\beta}_n \in [\beta_0 - \delta, \beta_0 + \delta]$, with probability 1 for n sufficiently large, and it follows from the uniform strong law of large numbers (Theorem 1, section 4.4 revision) together with continuity of $\mu(\beta) \equiv E_0 f(X, \beta)$ that

$$\begin{aligned} \bar{\alpha}_n^{\bar{\beta}_n} &= \frac{1}{n} \sum_{i=1}^n X_i^{\bar{\beta}_n} \\ &\rightarrow_{a.s.} E_0 f(X, \beta_0) = \alpha_0^{\beta_0}. \end{aligned}$$

(If instead $\bar{\beta}_n \rightarrow_p \beta_0$, then for and given $\epsilon > 0$ and $n \geq N_{\epsilon, \delta}$ large, $P_{\theta_0}(\bar{\beta}_n \in [\beta_0 - \delta, \beta_0 + \delta]) > 1 - \epsilon$ and we can simply argue on this set.) But now

$$\bar{\alpha}_n = \{\bar{\alpha}_n^{\bar{\beta}_n}\}^{1/\bar{\beta}_n} = g(\bar{\alpha}_n^{\bar{\beta}_n}, \bar{\beta}_n)$$

where $g(u, v) \equiv u^{1/v}$ is continuous and $(\bar{\alpha}_n^{\bar{\beta}_n}, \bar{\beta}_n) \rightarrow_{a.s.} (\alpha_0^{\beta_0}, \beta_0)$. Hence by the continuous mapping theorem

$$\bar{\alpha}_n = g(\bar{\alpha}_n^{\bar{\beta}_n}, \bar{\beta}_n) \rightarrow_{a.s.} g(\alpha_0^{\beta_0}, \beta_0) = \alpha_0.$$

B. The log-likelihood is

$$l_n(\alpha, \beta) = n \log(\beta/\alpha) + (\beta - 1) \sum_{i=1}^n \log(X_i/\alpha) - \sum_{i=1}^n \left(\frac{X_i}{\alpha}\right)^\beta,$$

and hence

$$\begin{aligned} l_n(\alpha, \bar{\beta}_n) &= n \log(\bar{\beta}_n/\alpha) + (\bar{\beta}_n - 1) \sum_{i=1}^n \log(X_i/\bar{\alpha}) - \sum_{i=1}^n \left(\frac{X_i}{\alpha}\right)^{\bar{\beta}_n} \\ &= -n \bar{\beta}_n \log \alpha - \frac{\sum_{i=1}^n X_i^{\bar{\beta}_n}}{\alpha^{\bar{\beta}_n}} + \text{constant in } \alpha \\ &= -n \log \eta - \frac{\sum_{i=1}^n X_i^{\bar{\beta}_n}}{\eta} + \text{constant in } \alpha \text{ and } \eta \end{aligned}$$

where $\eta \equiv \alpha^{\bar{\beta}_n}$. This is easily seen to be maximized by

$$\bar{\eta} \equiv \frac{1}{n} \sum_{i=1}^n X_i^{\bar{\beta}_n}$$

and hence

$$\bar{\alpha}_n = \left\{ \frac{1}{n} \sum_{i=1}^n X_i^{\bar{\beta}_n} \right\}^{1/\bar{\beta}_n}$$

as claimed. Thus $\bar{\alpha}_n$ is a pseudo-MLE of α .

5. (Profile likelihood) [For nice plots to accompany this exercise, see pages 41 - 43 of Cox, D. R. and Oakes, D. (1984); *Analysis of Survival Data*, Chapman and Hall.] As in problem 1.3, consider the Weibull family of example 3.2.5: $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$ with $\Theta \subset \mathbb{R}^{+2}$ given by the (Lebesgue) densities

$$p_\theta(x) = \frac{\beta}{\alpha} \left(\frac{x}{\alpha}\right)^{\beta-1} \exp\left(-\left(\frac{x}{\alpha}\right)^\beta\right) 1_{[0,\infty)}(x)$$

where $\theta \equiv (\alpha, \beta) \in (0, \infty) \times (0, \infty) \subset \mathbb{R}^2$.

(a) For a sample of n observations from p_θ , we know that, for each fixed value of β the value of α which maximizes the likelihood as a function of α is

$$\hat{\alpha}(\beta) = \left\{ \frac{1}{n} \sum_{i=1}^n X_i^\beta \right\}^{1/\beta}.$$

Use this to compute the *profile likelihood* $l_{\text{profile}}(\beta) = l_{\text{profile}}(\beta|\underline{X})$ defined by

$$l_{\text{profile}}(\beta) = l(\hat{\alpha}(\beta), \beta) = l(\hat{\alpha}(\beta), \beta|\underline{X}).$$

(b) Use what we know from problem 10.3 to show that the profile likelihood is strictly concave and hence has a unique maximum. Show that maximizing the profile likelihood as a function of β yields the maximum likelihood estimate: i.e. that $(\hat{\alpha}, \hat{\beta}) = (\hat{\alpha}_{\text{profile}}, \hat{\beta}_{\text{profile}})$.

(c) What is the relationship of the score function for β from the profile likelihood, $\dot{l}_{\beta, \text{profile}}$ to the (efficient) score for β from the full likelihood? Prove or disprove my claim: the profile score for β (based on n observations) is asymptotically equivalent to the sum of efficient scores for β over the sample in the sense that their difference divided by \sqrt{n} converges to 0 in probability.

(d) What is the relationship of the observed information from the profile likelihood $-\ddot{l}_{\beta, \text{profile}}$ to information quantities from the full likelihood?

Solution: (a) The log-likelihood is

$$l(\alpha, \beta) = n \log(\beta/\alpha) + (\beta - 1) \sum_{i=1}^n \log\left(\frac{X_i}{\alpha}\right) - \sum_{i=1}^n \left(\frac{X_i}{\alpha}\right)^\beta$$

and for fixed β the value of α which maximizes this is

$$\hat{\alpha}(\beta) = \left(\frac{1}{n} \sum_{i=1}^n X_i^\beta\right)^{1/\beta}.$$

Thus the profile log-likelihood is

$$l_{profile}(\beta) = l(\hat{\alpha}(\beta), \beta) = n \log \beta - n \log \left(\sum_{i=1}^n X_i^\beta \right) + (\beta - 1) \sum_{i=1}^n \log X_i + n \log n - n.$$

(b) It follows that the score function for β corresponding to the profile log-likelihood is

$$\dot{\mathbf{l}}_{profile, \beta}(\underline{X}) = \frac{n}{\beta} - n \frac{\sum_{i=1}^n X_i^\beta \log X_i}{\sum_{i=1}^n X_i^\beta} + \sum_{i=1}^n \log X_i,$$

and the observed information is

$$\begin{aligned} -\ddot{\mathbf{l}}_{profile, \beta}(\underline{X}) &= \frac{n}{\beta^2} + n \left\{ \frac{\sum_{i=1}^n X_i^\beta (\log X_i)^2}{\sum_{i=1}^n X_i^\beta} - \left(\frac{\sum_{i=1}^n X_i^\beta \log X_i}{\sum_{i=1}^n X_i^\beta} \right)^2 \right\} \\ &> 0 \end{aligned}$$

since the term in brackets is a variance, and hence is positive. Thus the profile likelihood is strictly concave and its maximum is unique.

Let $l^\#(\beta) = l_{profile}(\beta | \underline{X})$. Then, by the chain rule,

$$\dot{\mathbf{l}}_\beta^\#(\underline{X}) = \dot{\mathbf{l}}_{n\alpha} |_{\hat{\alpha}(\beta)} \dot{\alpha}(\beta) + \dot{\mathbf{l}}_{n\beta} |_{\hat{\alpha}(\beta)} = \dot{\mathbf{l}}_{n\beta} |_{\hat{\alpha}(\beta)} \quad (0.1)$$

since

$$\dot{\mathbf{l}}_{n\alpha} |_{\hat{\alpha}(\beta)} = 0. \quad (0.2)$$

Hence solving the profile score equation $\dot{\mathbf{l}}_\beta^\#(\underline{X}) = 0$ yields a solution of the likelihood equations $\dot{\mathbf{l}}_{n\alpha} = 0$ and $\dot{\mathbf{l}}_{n\beta} = 0$.

(c) As we have seen in (b), (0.1) and (0.2) hold by definition of profile (log-)likelihood. Expanding (0.2) about α yields

$$0 = \dot{\mathbf{l}}_{n\alpha} |_{\hat{\alpha}(\beta)} = \dot{\mathbf{l}}_{n\alpha} + \ddot{\mathbf{l}}_{n\alpha\alpha}(\alpha^*)(\hat{\alpha}(\beta) - \alpha),$$

and hence

$$\hat{\alpha}(\beta) - \alpha = -\ddot{\mathbf{l}}_{n\alpha\alpha}^{-1}(\alpha^*) \dot{\mathbf{l}}_{n\alpha}. \quad (0.3)$$

Expanding (0.1) about α yields

$$\dot{\mathbf{l}}_\beta^\#(\underline{X}) = \dot{\mathbf{l}}_{n\beta} |_{\hat{\alpha}(\beta)} = \dot{\mathbf{l}}_{n\beta} + \ddot{\mathbf{l}}_{n\beta\alpha}(\alpha^{**})(\hat{\alpha}(\beta) - \alpha). \quad (0.4)$$

Substitution of (0.3) into (0.4) yields

$$\begin{aligned} \dot{\mathbf{l}}_\beta^\#(\underline{X}) &= \dot{\mathbf{l}}_{n\beta} + \ddot{\mathbf{l}}_{n\beta\alpha}(\alpha^{**})(\hat{\alpha}(\beta) - \alpha) \\ &= \dot{\mathbf{l}}_{n\beta} - \ddot{\mathbf{l}}_{n\beta\alpha}(\alpha^{**}) \ddot{\mathbf{l}}_{n\alpha\alpha}^{-1}(\alpha^*) \dot{\mathbf{l}}_{n\alpha}. \end{aligned}$$

It follows easily from this relationship (under the conditions A0 - A4, the Cramér conditions) that

$$\begin{aligned}
& \frac{1}{\sqrt{n}} \{\mathbf{i}_\beta^\#(\underline{X}) - \mathbf{I}_{n,\beta}^*\} \\
&= \frac{1}{\sqrt{n}} \sum_{i=1}^n \{I_{\beta,\alpha} I_{\alpha\alpha}^{-1} \mathbf{i}_\alpha(X_i) - (-\ddot{\mathbf{I}}_{n\beta\alpha}(\alpha^{**}))(-\ddot{\mathbf{I}}_{n\alpha\alpha}^{-1}(\alpha^*)) \mathbf{i}_\alpha(X_i)\} \\
&= o_p(1).
\end{aligned}$$

(d) Differentiation across (0.2) with respect to β yields

$$\ddot{\mathbf{I}}_{n\alpha\alpha} \dot{\alpha}(\beta) + \ddot{\mathbf{I}}_{n\alpha\beta} = 0,$$

and hence

$$\dot{\alpha}(\beta) = -\ddot{\mathbf{I}}_{n\alpha\beta} \ddot{\mathbf{I}}_{n\alpha\alpha}^{-1}. \quad (0.5)$$

Similarly, differentiation of (0.1) with respect to β and then plugging in (0.5), yields

$$\ddot{\mathbf{I}}_{\beta\beta}^\# = \ddot{\mathbf{I}}_{n\beta\beta} + \ddot{\mathbf{I}}_{n\beta\alpha} \dot{\alpha}(\beta) = \ddot{\mathbf{I}}_{n\beta\beta} - \ddot{\mathbf{I}}_{n\beta\alpha} \ddot{\mathbf{I}}_{n\alpha\alpha}^{-1} \ddot{\mathbf{I}}_{n\alpha\beta}.$$

With some added minus signs, this becomes

$$-\ddot{\mathbf{I}}_{\beta\beta}^\# = \hat{I}_{n\beta\beta \cdot \alpha} |_{\hat{\alpha}(\beta)}.$$