Matthew Draper | POLI 271 | Homework #2 | April 15th, 2019

1a. $f(x) = \{ \theta e^{-\theta x}$ if $x > 0$

$\{ 0$ if $x \leq 0$

Support (X) = $[0, \infty)$

1b. n observations $(X = x_1 \ldots x_n) \sim \exp(\theta)$

$f_e\ (x, \theta) = P(x_1 = x; 0)$

$\mathcal{L}(\theta) = \Pi\ f_e\ (x_i, \theta)$

$\mathcal{L}(\theta) = \Pi(\theta \exp^{-\theta x})$

$\log \mathcal{L}(\theta) = \Sigma(\log(\theta \exp^{-\theta x}))$

$= \Sigma(\log(\theta)) + \Sigma(\log(\exp^{-\theta xi})$

$= n(\log(\theta)) - \theta\Sigma(x_i)$ (log likelihood)

1c. $\hat{\theta} = \text{argmax}(\log \mathcal{L}(\theta))$

$\partial/\partial\theta \log \mathcal{L}(\theta) = n(1/\theta) - \Sigma(xi)(1)$ (first derivative)

$0 = n/\Sigma(x_i) = 1/\text{sample mean}$ (extrema)

$\partial^2/\partial\theta^2 \log \mathcal{L}(\theta) = n(-\theta^{-2})$ (second derivative)

This expression is < 0, so the first derivative was a maximum.

1d. set.seed(2) : Sets the seed of R's random number generator, which allows for reproducible (quasi-) randomness.

x <- rexp(1000, 3) : Generates 1000 random observations from an exponential distribution with a parameter (rate) of 3. Exponential distributions assume a constant rate at which the event occurs.

2a. $X_n \sim N(\mu, \sigma^2)$

$f(x, \mu, \sigma^2) = 1/\sqrt{(2\pi\, \sigma^2)} * ((\exp(-x-\mu)^2)/2\, \sigma^2)$

Assuming μ is fixed ($X = x_1 \ldots x_n$):

$\mathcal{L}(\sigma^2) = \hat{\Pi}(1/\sqrt{(2\pi\, \sigma^2)} * ((\exp(-x_i-\mu)^2)/2\, \sigma^2))$

$\log(\mathcal{L}(\sigma^2)) = \Sigma \log(1/\sqrt{(2\pi\, \sigma^2)}) - \Sigma(-x_i-\mu)^2/2\, \sigma^2$ (log-likelihood)

So:

$\log(\mathcal{L}(\sigma^2)) = \Sigma \log(2\pi\, \sigma^2)^{-1/2} - 1/(2\pi\, \sigma^2) * \Sigma(x_i-\mu)^2$

$= -½\, n \log(2\pi\, \sigma^2) - 1/(2\pi\, \sigma^2) * \Sigma(x_i-\mu)^2$

$\partial/\partial\sigma^2 \log(\mathcal{L}(\sigma^2)) = (-n/2) * (1/(2\pi\, \sigma^2)) * (2\,\pi) + ½ * \Sigma(x_i-\mu)^2 * (\sigma^2)^{-2}$

(first derivative)

$0 = (-n/2) * (1/\sigma^2) + ½\, \Sigma(x_i-\mu)^2 + (1/\sigma^2)^2$

$0 = (-n\sigma^2 + \Sigma(x_i-\mu)^2)/(2(\sigma^2)^2)$

$0 = -n\sigma^2 + \Sigma(x_i-\mu)^2$

$\Sigma(x_i-\mu)^2 = n\sigma^2$

$\hat{\sigma}^2 = (\Sigma(x_i-\mu)^2)/n$ (maxima)

$\partial^2/\partial(\sigma^2)^2 = -1$

This expression is < 0, so the first derivative was a maximum.

2b. In this case, we don't observe bias, but that's because we held μ fixed. If μ had not been fixed, we would have lost degrees of freedom (n – k – 1). Our expected value would look like this:

$E\, [(\Sigma(x_i-\boldsymbol{\mu})^2)/n] = 1/n * E(x_i - \mu)^2 = 1/n * \Sigma\, \sigma^2 = \sigma^2$ (holding μ fixed)

$E\, [(\Sigma(x_i-\mathbf{xbar})^2)/n] = 1/n * E(x_i - \mu)^2 = 1/n * \Sigma\, \sigma^2 = \sigma^2$ (variable μ)

Draper HW 2

Matthew Draper

April 15, 2019

```
## Question 1d
set.seed(2)
x <- rexp(1000, 3)
exp.ll <- function(theta,x)
{n <- length(x)
  l <- n*(log(theta))-(theta*(sum(x)))
  l}
theta <- seq(1,5,by=0.001)
l<-1:length(theta)
for (e in 1:(length(theta))) {
  l[e] <- exp.ll(theta[e],x)
}
```

```
plot(theta,l,type='l',ylab='log-likelihood')
```

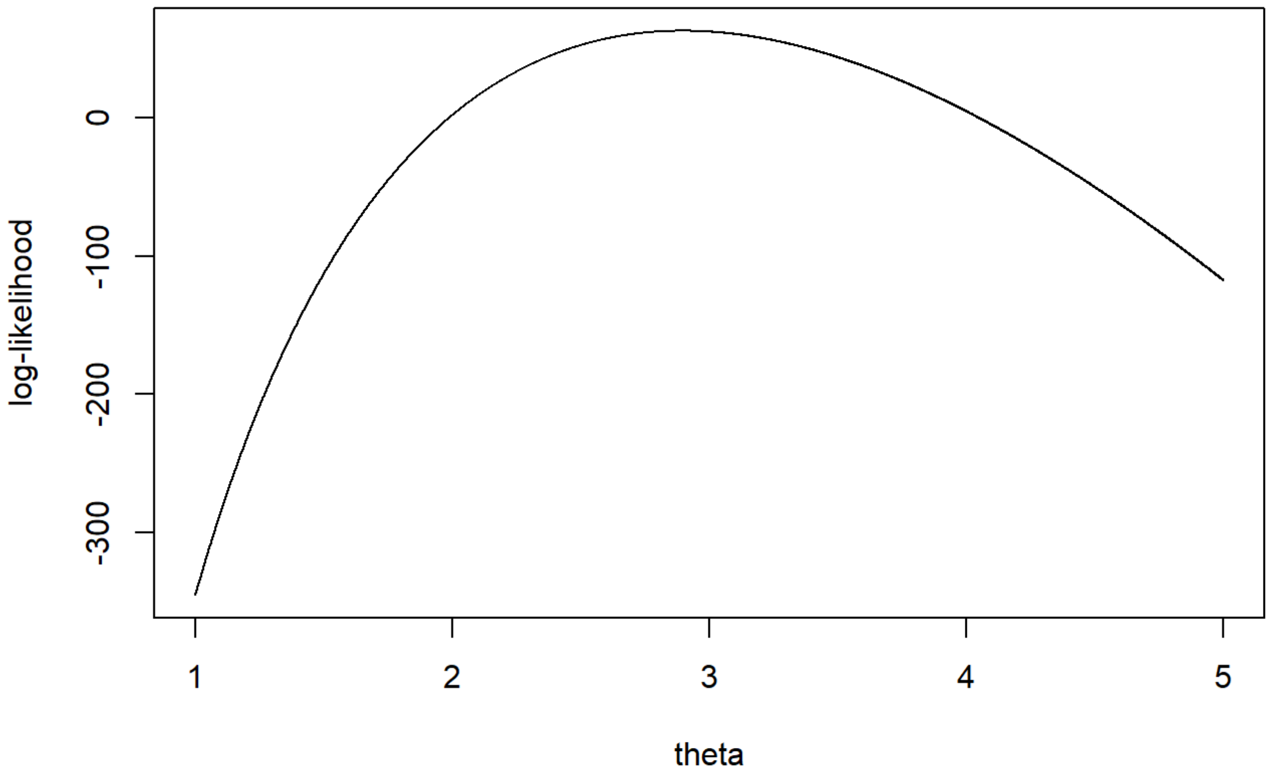


```
thetahat <- length(x)/sum(x)
print("The value of theta hat (given x) is:")
```

```
## [1] "The value of theta hat (given x) is:"
```

```
print(thetahat)
```

```
## [1] 2.895693
```