**September 13, 2015**
**Solutions to HW #1**

1. Question 1

   (a) The relevant R–code is

```
> library(smoothmest)
> n <- 10000
> x <- rnorm(n,mean=0,sd=1)
> y <- rdoublex(n,mu=0,lambda=1/sqrt(2))
> plot(density(x),xlim=c(-4,4),
+      col='red',ylim=c(0,.8),
+      xlab="x",ylab="Density",main="")
> lines(density(y),col="black",lty=2)
> legend("topleft",
+         c("Normal","Double Exponential"),
+         col=c("red","black"),
+         lty=1:2)
```

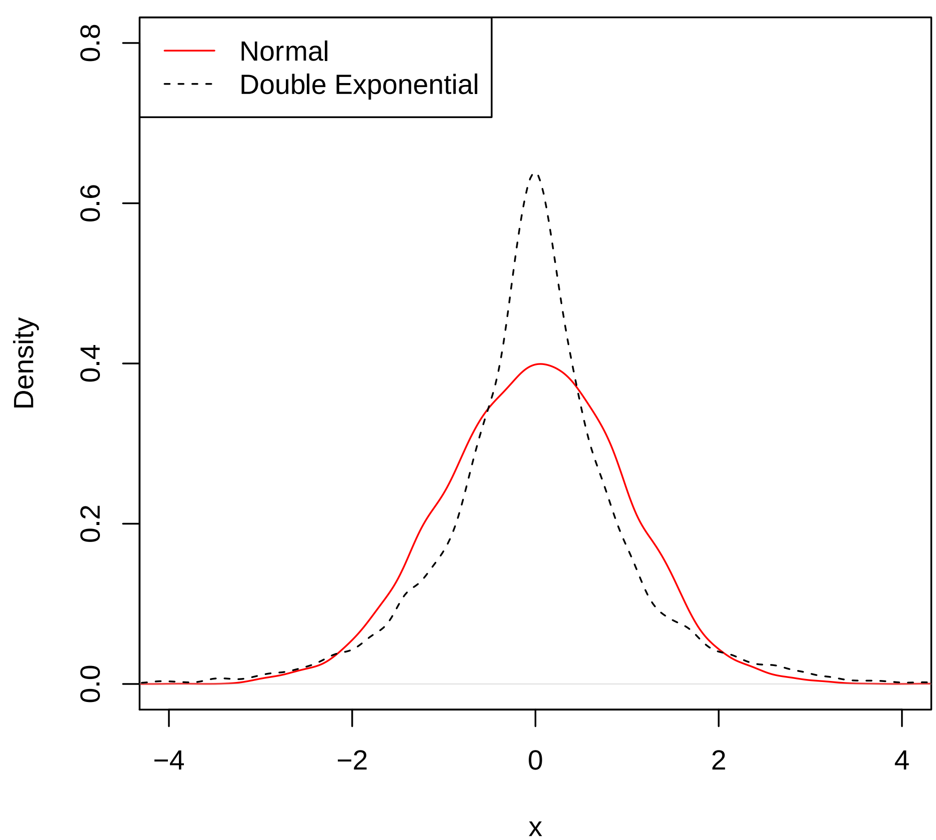


The double exponential is more concentrated around 0 but with heavier tails.

(b) Note that $g(x) = x\frac{1}{2\lambda}e^{-|x|/\lambda}$ is an odd function ie $g(x) = -g(-x)$. Therefore

$$\begin{aligned}
\mathbb{E}[X] &= \int_{-\infty}^{\infty} x\frac{1}{2\lambda}e^{-|x|/\lambda}dx \\
&= \int_{-\infty}^{0} x\frac{1}{2\lambda}e^{-|x|/\lambda}dx + \int_{0}^{\infty} x\frac{1}{2\lambda}e^{-|x|/\lambda}dx \\
&= 0
\end{aligned}$$

Since the expectation is 0, Var $(X) = \mathbb{E}[X^2]$. With the final equality following from integration by parts twice, we have

$$\begin{aligned}
\text{Var } (X) &= \mathbb{E}[X^2] \\
&= \int_{-\infty}^{\infty} \frac{x^2}{2\lambda}e^{-|x|/\lambda}dx \\
&= \frac{1}{\lambda}\int_{0}^{\infty} x^2 e^{-x/\lambda}dx \\
&= 2\lambda^2
\end{aligned}$$

Thus the standard deviation is $\lambda\sqrt{2}$.

(c) The maximum likelihood estimator is

$$\begin{aligned}
\widehat{\mu}_{MLE} &= \underset{\mu}{\text{argmax}}\, f(\vec{m}|\mu, \lambda) \\
&= \underset{\mu}{\text{argmax}} \prod \frac{1}{2\lambda}e^{-|m_i-\mu|/\lambda} \\
&= \underset{\mu}{\text{argmax}} \sum_{i=1}^{n} \log\left(\frac{1}{2\lambda}e^{-|m_i-\mu|/\lambda}\right)
\end{aligned}$$

The maximum does not depend on $\lambda$ in any way so

$$\widehat{\mu}_{MLE} = \underset{\mu}{\text{argmax}} \sum_{i=1}^{n} -|m_i - \mu| = \underset{\mu}{\text{argmin}} \sum_{i=1}^{n} |m_i - \mu| = median(m_i).$$

(d) We could also estimate $\mu$ using the mean of the $m_i$. We call this estimator $\bar{\mu}$ where

$$\bar{\mu} = \frac{1}{n}\sum_{i=1}^{n} m_i$$

We compare $\widehat{\mu}_{MLE}$ to $\bar{\mu}$ using a simulation study. The simulation study will provide some idea as to the form of the sampling distribution. The estimator with a sampling distribution that is tighter around the true $\mu$ (for example, has a lower MSE), is the better estimator. In order to have a standard deviation of 0.3, $\lambda = 0.3/\sqrt{2}$. The relevant code is

```
> N <- 1000
> n <- 100
> m <- matrix(rdoublex(n*N,mu=18,lambda=0.3/sqrt(2)),nrow=N)
> mean_est <- apply(m,1,mean)
> median_est <- apply(m,1,median)
> plot(density(mean_est),xlim=c(17.9,18.1),
+      col='red',ylim=c(0,20),
+      xlab="estimate",ylab="Density",main="")
> lines(density(median_est),col="black",lty=2)
> legend("topleft",
+         c("Mean","Median"),
+         col=c("red","black"),
+         lty=1:2)
```

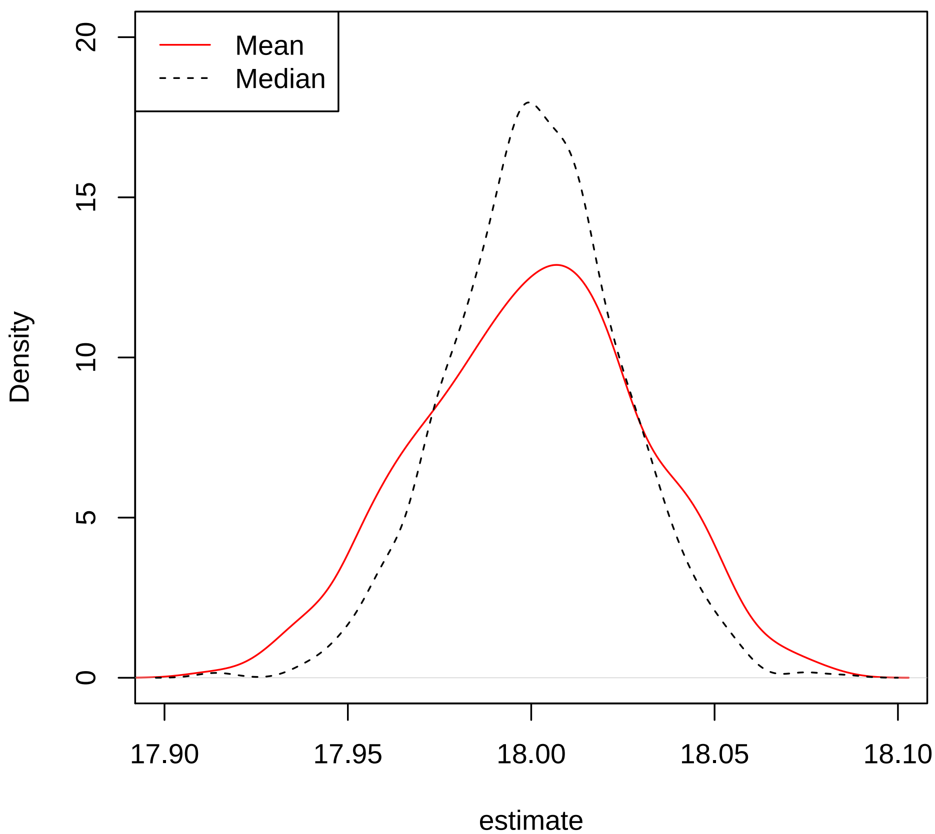


From the plot, it is clear that the sampling distribution of the median is more concentrated around the true value of 18 than the mean. Thus the median is a better estimator. You could quantify the improvement in the estimator through the mean squared error. The mean squared error for the two estimators are

```
> mean((mean_est - 18)^2)
```

```
[1] 0.0009165382
> mean((median_est - 18)^2)
[1] 0.0005312188
```

At this point you cannot say that the median is always better than the mean because we only did the simulation at one standard deviation and one mean. You can simulate at more values to get an understanding of what happens at different parameters. A better approach is to use math. It is possible to show using the central limit theorem and asymptotic results for the median, that whenever the data is generated from a double exponential distribution, then median will be a better estimator of $\mu$. The result is somewhat counter–intuitive because $\mu$ is the population mean, but the sample mean is not the best estimator for it.

(e) We repeat the simulation but draw $m$ from a Gaussian with mean 18 and standard deviation 0.3.

```
> N <- 1000
> n <- 100
> m <- matrix(rnorm(n*N,mean=18,sd=0.3),nrow=N)
> mean_est <- apply(m,1,mean)
> median_est <- apply(m,1,median)
> plot(density(mean_est),xlim=c(17.9,18.1),
+     col='red',ylim=c(0,15),
+     xlab="estimate",ylab="Density",main="")
> lines(density(median_est),col="black",lty=2)
> legend("topleft",
+        c("Mean","Median"),
+        col=c("red","black"),
+        lty=1:2)
```

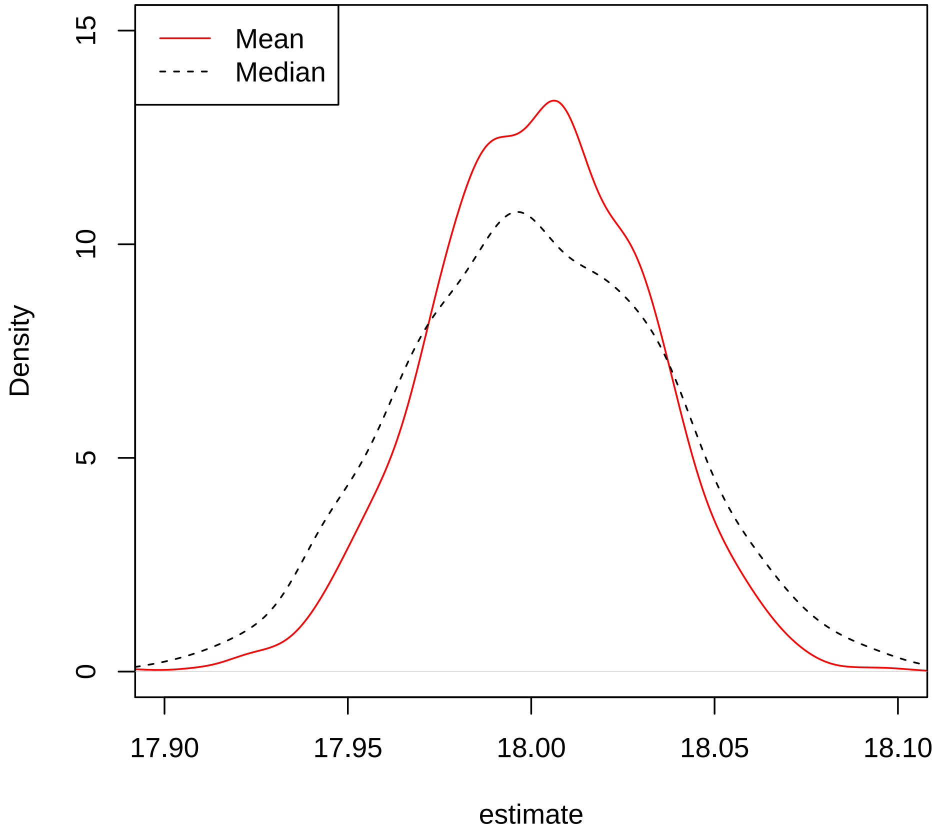


We now see that the mean is a better estimator. This is also reflected in the MSE

```
> mean((mean_est - 18)^2)
[1] 0.0008264765
> mean((median_est - 18)^2)
[1] 0.001302012
```