HW2, STAT453/553, due 2/8/17

1. Let X be exponential with mean 1, and let Y be exonential with mean 2, where X and Y are independent. Let $U = X + Y$ and let $V = X - Y$.

(a) Find the the joint density $f_{U,V}(u, v)$. Be sure to specify the density for all real values of u and v.

(b) Find the marginal densities for U and V. Note that U does not have a gamma distribution.

2. Let X and Y be i.i.d. gamma random variables with parameters a and b (use the form of the density in Casella and Berger). Let $U = \frac{X}{X+Y}$ and let $V = X$. Find the joint density of U and V and show that the marginal distribution of U follows a beta distribution (and find the parameters of the beta distribution).

3. Let $X_1, \ldots X_n$ be an i.i.d. sample from the density

$$f(x) = \frac{k}{x^3} \ \ I(x > 1)$$

(a) find the constant k that makes this a density (i.e., it should integrate to 1).

(b) Find the density for $\min_i X_i = X_{(1)}$

(c) Find the density for $\max_i X_i = X_{(n)}$

(d) Assuming that n is an odd number, find the density for the median.

4. To illustrate Jensen's inequality, we'll look at a couple cases. The first is

$$E\left[\frac{U}{1-U}\right] \text{ versus } \frac{E[U]}{1-E[U]}$$

where U has a uniform(0,1) distribution. Jensen's inequality states that if $g(x)$ is convex (concave up) then $E[g(X)] \geq g(E[X])$, while if $g(x)$ is concave down, then $E[g(X)] \leq g(E[X])$.

(a) Determine whether $g(x) = \frac{x}{1-x}$ is concave up or concave down (for example, using the second derivative) and therefore whether $E\left[\frac{U}{1-U}\right]$ should be larger or less than $\frac{E[U]}{1-E[U]}$. To review this topic, look up Jensen's inequality in chapter 4 of the book.

(b) Confirm the results in (a) by simulating $U/(1-U)$ over many iterations and taking the average. To do this in R, you can use the following code for I iterations:

```
I <- 10000
u <- unif(I)
mean(u/(1-u))
```

What is your average value of $u/(1-u)$? This should approximate $E\left[\frac{U}{1-U}\right]$. Is this larger or less than $\frac{E[U]}{1-E[U]}$? Also make a histogram of the distribution of $U/(1-U)$ in R using `hist(u/(1-u))`. Can you explain why the histogram looks like this?

(c) For the second illustration of Jensen's inequality, we'll compare $\sqrt{E(S^2)}$ and $E(\sqrt{S^2}) = E(S)$. In this case, we need many iterations of a sample. Suppose the sample size is $n = 10$. Then we need I iterations of samples of size $n = 10$.

To generate the data, we can use the following code in R (you can use another language such as Matlab if you want)

```
I <- 10000
n <- 10
myxbar <- 1:I
mys <- 1:I
mys2 <- 1:I
for(i in 1:I) {
  x <- rexp(10,0,1) # sample of size 10 from standard normal
  myxbar[i] <- mean(x) # sample mean
  mys2[i] <- var(x) # sample variance
  mys[i] <- sqrt(var(x)) # sample standard deviation
}
print(mean(myxbar))
print(mean(mys2))
print(mean(mys))
```

The above code is for samples of size 10 from standard normal.

(a) Comment on whether the sample means, variances, and standard deviations tend to be good approximations to the population (i.e., theoretical) mean, variance, and standard deviation, or whether they appear to over or underestimate them. If you are not sure, then a suggestion is to get a confidence interval for the expected value of the statistic. For example if you use `t.test(mys)` this should give you a plausible interval for $E(S)$, you might check whether this interval includes the theoretical standard deviation. (For example, for a uniform(0,1), the theoretical variance is 1/12 and the theoretical standard deviation is $\sqrt{1/12} = 1/(2\sqrt{3})$.

(b) Repeat the experiment for a uniform(0,1) distribution, a beta(2,2) distribution, and an exponential(1) distribution. To simulate from these distributions, use the functions `runif()`, `rbeta()`, and `rexp()`.

(c) For all distributions, also calculate the sample correlation between $\overline{x}$ and s^2. This can be done in R using

```
cor(myxbar,mys2)
```

Here the idea is that you have 10000 values of $\overline{X}$ and 10000 values of S^2, and you are just computing one correlation for these 10000 pairs. You might also make a scatterplot of these 10000 pairs to get an idea of how $\overline{X}$ and S^2 are related. For the uniform, it is interesting to try this when the sample size is $n = 2$.