PBIO 294 HW 25 Oct 2017; Amended on 5 November 2017

MCMC sampler for a linear regression model. Due by midnight on 10 Nov 2017.

Problem 1: ~~This question uses the height vs. weight data from 6 Sep. The data can be found on the class website from that week.~~ Create simulated data to recapture regression parameters. The code to simulate the data can be found on the website from 1 Nov 2017 and is copied here:

```
nReps<-10
x<-seq(from=0,to=100,by=1)
x<-rep(x,nReps)
b0<-1.0; b1<-0.2
ymean<-b0+b1*x
y<-rnorm(n=length(ymean),mean=ymean,sd=1.5)
N<-length(y)
```

a) Fit a linear regression predicting weight from height using your own MCMC sampler. Use a Normal(0,100) prior for both the intercept and slope and a InverseGamma(0.001,0.001) prior for the standard deviation. (Install the invgamma package to access this distribution.) Generate 10,000 samples for each of the model parameters.

b) Use the R coda package to create a trace plot, density plot, autocorrelation plot, and to decide on a burn in lenght.

c) Report the mean and 95% HDI for the intercept, slope, and sd. Show a plot of the fit superimposed on the data