AMS 206 – Homework Assignment #6
due Friday, March 2

Required Problem

1. The time to download web site homepages varies from city to city. Keynote Systems measured download times for selected web sites in various cities and assigned each city an average download time. The data file `download.txt` gives a sample of 17 cities and their web page download ratings (in seconds) for the week of December 19, 1999 to December 26, 1999.

   (a) Read the data into R. (Here is an example of a real-world dataset that doesn't load into R on your first try. You'll have to figure out why not.)

   (b) Make a histogram of the data. Do they look normally distributed?

   (c) We'll use a gamma likelihood for this dataset, with both parameters unknown. Use exponential priors for both parameters (i.e., gamma priors with first parameter equal to 1). Since you probably don't have much intuition about these parameters, we'll use somewhat vague priors—pick both priors to be exponential with mean 10.

      i. Fit this model by drawing a posterior sample using the Metropolis-Hastings algorithm. Be sure to show your work for computing the acceptance probabilities and to report your acceptance rates (and to adjust your proposals if your rates are not good).
      ii. Show your trace plots for both parameters and comment on the mixing.
      iii. Plot histograms for both marginal posterior distributions.
      iv. Compute the posterior mean and variance for both parameters.
      v. Plot autocorrelations for both parameters.

   (d) The second parameter can be sampled using a Gibbs step (but the first one will still require a Metropolis-Hastings step).

      i. Fit this model by drawing a posterior sample using Gibbs sampling for the second parameter, and the Metropolis-Hastings algorithm for the first. Be sure to show your work for computing the complete conditional distribution of interest and to report your acceptance rate for the other parameter.
      ii. Show your trace plots for both parameters and comment on the mixing.
      iii. Plot histograms for both marginal posterior distributions.
      iv. Compute the posterior mean and variance for both parameters.
      v. Plot autocorrelations for both parameters.
      vi. Compare your results to those from part (c) above in terms of both the resulting posteriors and the mixing.

Optional Problem

2. Assess the sensitivity of your results above to the choice of your prior by choosing two other very different priors and repeating parts (d)-iii and (d)-iv for each.