I. 1D grid approximation

1. Let's model the annual number of burglaries on SLU's campus with a Gamma-Poisson model:

$$Y_i|\lambda \sim \text{Pois}(\lambda)$$
$$\lambda \sim \text{Gamma}(s, r)$$

   Choose values of s and r that match your prior ideas about campus crime.

2. In the last three years, there were 17, 27, and 22 burglaries. What is the posterior distribution of λ given this data?

3. Approximate the posterior using a grid with 100 values of λ between 10 and 30. Plot it.

II. 2D grid approximation

1. Model body temperature for women using the normal model with mean and variance as parameters:

$$Y_i|\mu, \sigma \sim \text{Normal}(\mu, \sigma^2)$$
$$\mu \sim \text{Normal}(\theta, \tau^2)$$
$$\sigma \sim \text{Gamma}(s, r)$$

   Choose reasonable values for the hyperparameters θ, τ, s, and r.

2. Create a 50x50 grid that covers the interesting part of parameter space. Compute the prior at each grid point. Make a plot using color (or alpha) at each grid point to show the prior distribution.

3. Load the `normtemp.csv` file from our course web page, and take your data to be the 65 measurements of female body temperature.

4. Compute the grid approximation to the posterior. Plot it using a dot at each grid point, then also use `geom_contour` to overlay contour height lines.

III. MCMC with rstan

1. Using the female temperature data from Part II, perform a posterior simulation using `rstan`. The default RStudio stan program will work in this setting.

2. Use `extract` to extract the simulated μ and σ values. Make a scatterplot and add contour lines to the plot with `geom_density_2d()`. Compare with the grid approximation in part II.

3. Make a trace plot of the chains.

---

4. Using the model and burglary data from Part I, perform a posterior simulation using `rstan`.

5. Use `extract` to extract the simulated λ values. Make a density plot. Compare with the grid approximation in part I.

6. Make a trace plot of the chains.