I. Poisson and Gamma distributions

1. Plot the PDF of the Poisson distribution for $\lambda = 10.5$. What is its mode?

2. Explore the Gamma(s,r) distribution using the Desmos applet on our course page. Note that they call the parameters $s = \alpha$ and $r = \beta_{\text{rate}}$. For a fixed rate $r = 1$, what happens as the shape passes from $s > 1$ through $s = 1$ to $s < 1$?

3. The Gamma distribution has mean $\frac{s}{r}$, mode $\frac{s-1}{r}$, and variance $\frac{s}{r^2}$.
   (a) Find Gamma(s,r) with mean 5 and sd 1/2. Check by simulating with `rgamma`.
   (b) Find Gamma(s,r) with mode 6 and mean 8.

II. Posterior simulation

Consider a Poisson variable $Y \sim \text{Pois}(\lambda)$ with an unknown rate λ of success. Model our prior understanding with $\lambda \sim$ Gamma(1,1). We make two observations and find $Y_1 = 5, Y_2 = 3$.

1. What is the theoretical posterior distribution of λ?

2. Simulate 1,000,000 values of λ.

3. For each value of λ, use `rpois` to choose a Y_1 and a Y_2 value.

4. Select just the λ values corresponding to $Y_1 = 5, Y_2 = 3$. This is a sample from the posterior distribution of $f(\lambda | Y_1 = 5, Y_2 = 3)$. How many of the 1,000,000 values did you get in your sample?

5. Make a density plot of these values of λ and then add the theoretical posterior distribution to the plot.

III. Soccer goals

Pick your favorite soccer team, and think about their rate λ of goals scored per game.

1. Tune a Gamma prior to your group's knowledge of λ. What are the mean, mode, and sd for your prior?

2. Make a list of how many goals your team scored in each of their last ten games.

3. Find the posterior Gamma distribution for λ. What are the mean, mode, and sd?

4. Plot the prior and posterior distributions on the same graph. Use `bayesrules::plot_gamma_poisson()` if you have it.

IV. Maximum likelihood

Consider the Gamma-Poisson model with data $y_1, \ldots, y_n$. The likelihood function is given (up to a constant) by

$$L(\lambda) \propto \lambda^{\sum y_i} e^{-n\lambda}$$

Show that the maximum of the likelihood function occurs when $\lambda = \bar{y} = \frac{1}{n} \sum_{i=1}^{n} y_i$