

I. Beta distribution

1. Explore the shape of $\text{Beta}(\alpha, \beta)$ using the applet <https://mathlets.org/mathlets/beta-distribution>. Make sure to look at $\text{Beta}(1,1)$, $\text{Beta}(2,1)$, $\text{Beta}(1,2)$. What can you say about the shape when $\alpha = \beta$, when $\alpha < \beta$, and when $\alpha > \beta$? What happens when $\alpha < 1$ or $\beta < 1$ or both are less than one?
2. If $\pi \sim \text{Beta}(\alpha, \beta)$, then $E(\pi) = \frac{\alpha}{\alpha + \beta}$. Use `rbeta` to check this when $\alpha = 8$ and $\beta = 2$. What is the standard deviation $\sigma(\pi)$ in this case?
3. The PDF for $\text{Beta}(\alpha, \beta)$ is given by

$$f(x) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}$$

Check that the mode of this distribution is $\frac{\alpha-1}{\alpha+\beta-2}$. This is a calculus problem – the mode is the x which maximizes $f(x)$, which happens at a critical point where $f'(x) = 0$.

II. Posterior simulation

(Section 3.5 of Bayes Rules! is helpful here)

Consider a binomial experiment $Y \sim \text{Binom}(10, \pi)$ with 10 trials and an unknown probability π of success. Model our prior understanding with $\pi \sim \text{Unif}(0,1)$.

1. Assume we run the experiment and have 8 successes. What is the theoretical posterior distribution of π ?
2. Simulate 10000 values of π .
3. For each value of π , use `rbinom` to choose one y value.
4. Select just the π values corresponding to $y = 8$. This is a sample from the distribution of $f(\pi|Y = 8)$, which is the posterior we're looking for. Make a density plot of these values and then add the corresponding Beta distribution to the plot.

Hint: In base R, you can use `plot(density(...))` and `curve`. With `ggplot`, you can use `geom_density` and `geom_function`

III. Water and Land

1. As a group, tune a Beta prior to your group's knowledge of what percentage of the Earth's surface is water. What are the mean, mode, and sd for your prior?
2. Use the website <https://www.random.org/geographic-coordinates/> to visit at least 20 locations on Earth and record how many are land and how many are water.
3. Find the posterior Beta distribution for the percentage of water, given your data. What are the mean, mode, and sd?
4. Plot the prior and posterior distributions on the same graph. Use `bayesrules::plot_beta_binomial()` if you have it.