Class 15 in-class problems, 18.05, Spring 2022

Concept questions

Concept question 1. More Beta

Suppose your prior $f(\theta)$ in the bent coin example is Beta$(6, 8)$. You flip the coin 7 times, getting 2 heads and 5 tails. What is the posterior pdf $f(\theta|x)$?

(a) *Beta(2,5)*

(b) *Beta(11,10)*

(c) *Beta(6,8)*

(d) *Beta(8,13)*

Solution: We saw in the first board question that 2 heads and 5 tails will update a Beta(a, b) prior to a Beta$(a + 2, b + 5)$ posterior.

So the answer is (d), Beta$(8, 13)$.

Concept question 2. Strong priors

Say we have a bent coin with unknown probability of heads θ.

We are convinced that $\theta \leq 0.7$.
Our prior is uniform on $[0, 0.7]$ and 0 from 0.7 to 1.

We flip the coin 65 times and get 60 heads.

Which of the graphs below is the posterior pdf for θ?

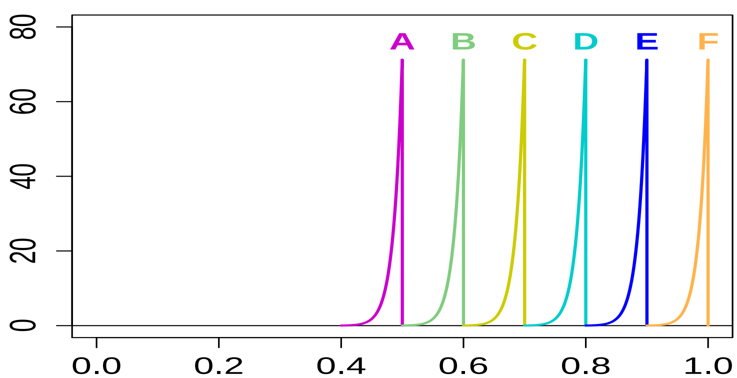

Solution: Graph C, the blue graph spiking near 0.7.

Sixty heads in 65 tosses indicates the true value of θ is close to 1. Our prior was 0 for $\theta > 0.7$. So no amount of data will make the posterior non-zero in that range. That is, we have foreclosed on the possibility of deciding that θ is close to 1. The Bayesian updating puts θ near the top of the allowed range.

Concept question 3. Normal priors, normal likelihood

(a)

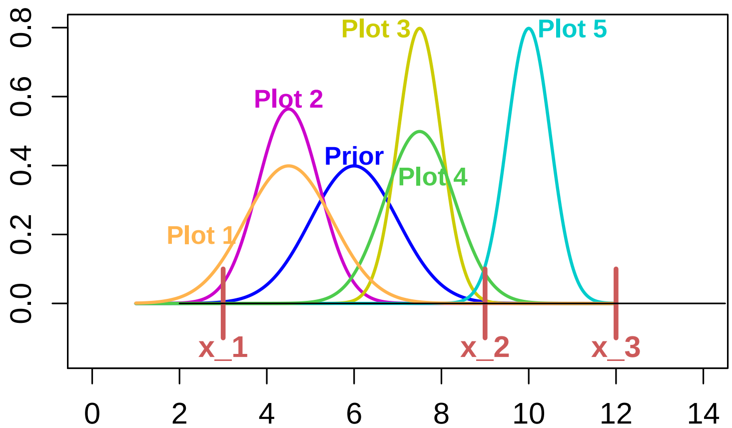

Blue graph = prior,   Red lines = data in order: 3, 9, 12

Which plot is the posterior to just the first data value?

Solution: Plot 2: The first data value is 3. Therefore the posterior must have its mean between 3 and the mean of the blue prior. The only possibilites for this are plots 1 and 2. We also know that the variance of the posterior is less than that of the posterior. Between the plots 1 and 2 graphs only plot 2 has smaller variance than the prior.

Concept question 4. Normal priors, normal likelihood
(b)

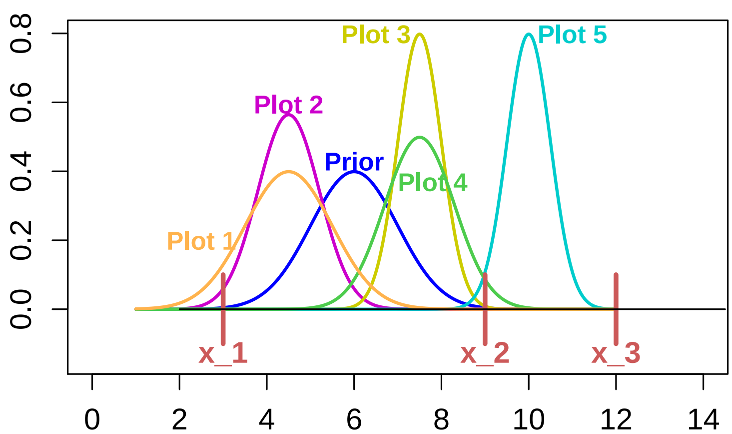

Blue graph = prior,   Red lines = data in order: 3, 9, 12

Which graph is posterior to all 3 data values?

Plot 3: The average of the 3 data values is 8. Therefore the posterior must have mean between the mean of the blue prior and 8. Therefore the only possibilities are the plots 3 and 4. Because the posterior is posterior to the magenta graph (plot 2) it must have smaller variance. This leaves only the Plot 3.

Board questions

Problem 1. Beta priors

Suppose you are testing a new medical treatment with unknown probability of success θ. You don't know θ, but your prior belief is that it's probably not too far from 0.5. You capture this intuition with a Beta(5,5) prior on θ.

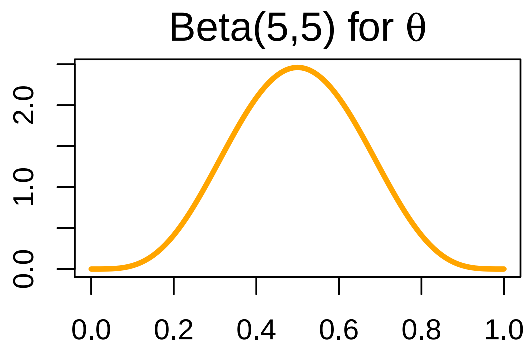

To sharpen this distribution you take data and update the prior.

- *Beta(a, b):* $f(\theta) = \dfrac{(a+b-1)!}{(a-1)!(b-1)!}\theta^{a-1}(1-\theta)^{b-1}$
- *Treatment has prior* $f(\theta) \sim$ *Beta*$(5,5)$

(a) *Suppose you test it on 25 patients and have 20 successes.*
– Find the posterior distribution on θ.
– Identify the type of the posterior distribution.

(b) *Suppose you recorded the order of the results and got*

$$S\,S\,S\,S\,F\,S\,S\,S\,S\,S\,F\,F\,S\,S\,S\,F\,S\,F\,S\,S\,S\,S\,S\,S\,S$$

(20 S and 5 F). Find the posterior based on this data.

(c) *Using your answer to (b) give an integral for the posterior predictive probability of success with the next patient.*

***(d)* Don't do in class.** *Use what you know about pdf's to evaluate the integral without computing it directly*

Solution: (a) Prior pdf is $f(\theta) = \frac{9!}{4!\,4!}\theta^4(1-\theta)^4 = c_1\theta^4(1-\theta)^4$.

hyp.	prior	likelihood	Bayes numer.	posterior
θ	$c_1\theta^4(1-\theta)^4\,d\theta$	$\binom{25}{20}\theta^{20}(1-\theta)^5$	$c_3\theta^{24}(1-\theta)^9\,d\theta$	Beta$(25,10)$
	Beta$(5,5)$	binom. prob.		Beta$(25,10)$

We know the normalized posterior is a Beta distribution because it has the form of a Beta distribution ($c\theta^{a-1}(1-\theta)^{b-1}$ on [0,1]) so by our earlier observation it must be a Beta distribution.

(b) The answer is the same. The only change is that the likelihood has a coefficient of 1 instead of a binomial coefficent.

(c) The posterior on θ is Beta(25, 10) which has density

$$f(\theta\,|, \text{data}) = \frac{34!}{24!\,9!}\theta^{24}(1-\theta)^9.$$

The law of total probability says that the posterior predictive probability of success is

$$\begin{aligned} P(\text{success}\,|\,\text{data}) &= \int_0^1 f(\text{success}\,|\,\theta)\cdot f(\theta\,|\,\text{data})\,d\theta \\ &= \int_0^1 \theta\cdot\frac{34!}{24!\,9!}\theta^{24}(1-\theta)^9\,d\theta = \int_0^1 \frac{34!}{24!\,9!}\theta^{25}(1-\theta)^9\,d\theta \end{aligned}$$

(d) We compute the integral in (c) by relating it to the pdf of Beta(26, 10). That pdf is $\frac{35!}{25!\,9!}\theta^{25}(1-\theta)^9$. Since any pdf integrates to 1 we have

$$\int_0^1 \frac{35!}{25!\,9!}\theta^{25}(1-\theta)^9 = 1 \qquad \Rightarrow \qquad \int_0^1 \theta^{25}(1-\theta)^9 = \frac{25!\,9!}{35!}.$$

Thus, we can compute the integral in part (c):

$$P(\text{Success}\,|\,\text{data}) = \int_0^1 \frac{34!}{24!\,9!}\theta^{25}(1-\theta)^9\,d\theta = \frac{34!}{24!\,9!}\cdot\frac{25!\,9!}{35!}. = \boxed{\frac{25}{35}\approx 0.71}.$$

Problem 2. Normal-normal updating

For data $x_1,\dots,x_n$ with data mean $\bar{x} = \frac{x_1+\dots+x_n}{n}$

$$a = \frac{1}{\sigma^2_{prior}},\qquad b = \frac{n}{\sigma^2},\qquad \mu_{post} = \frac{a\mu_{prior}+b\bar{x}}{a+b},\qquad \sigma^2_{post} = \frac{1}{a+b}.$$

Suppose we have one data point $x = 2$ drawn from $N(\theta, 3^2)$

Suppose θ is our parameter of interest with prior $\theta \sim N(4, 2^2)$.

***(a)** Identify μ_{prior}, σ_{prior}, σ, n, and $\bar{x}$.*

***(b)** Make a Bayesian update table, but leave the posterior as an unsimplified product.*

***(c)** Use the updating formulas to find the posterior.*

Solution: (a) $\mu_{\text{prior}} = 4$, $\sigma_{\text{prior}} = 2$, $\sigma = 3$, $n = 1$, $\bar{x} = 2$.

(b)

hypoth.	prior	likelihood	posterior
θ	$f(\theta) \sim \text{N}(4, 2^2)$	$\phi(x\|\theta) \sim \text{N}(\theta, 3^2)$	$f(\theta\|x) \sim \text{N}(\mu_{\text{post}}, \sigma^2_{\text{post}})$
θ	$c_1\exp\left(\frac{-(\theta-4)^2}{8}\right)$	$c_2\exp\left(\frac{-(2-\theta)^2}{18}\right)$	$c_3\exp\left(\frac{-(\theta-4)^2}{8}\right)\exp\left(\frac{-(2-\theta)^2}{18}\right)$

(c) We have $a = 1/4$, $\quad b = 1/9$, $\quad a+b = 13/36$. Therefore

$$\begin{aligned} \mu_{\text{post}} &= (1+2/9)/(13/36) = 44/13 = 3.3846 \\ \sigma^2_{\text{post}} &= 36/13 = 2.7692 \end{aligned}$$

The posterior pdf is $f(\theta|x = 2) \sim \mathrm{N}(3.3846, 2.7692)$.

Problem 2. Normal/normal

Question. *On a basketball team the average career free throw percentage over all players follows a $N(75, 6^2)$ distribution. In a given year individual players free throw percentage is $N(\theta, 4^2)$ where θ is their career average.*

This season Sophie Lie made 85 percent of her free throws. What is the posterior expected value of her career percentage θ?

Solution: This is a normal/normal conjugate prior pair, so we use the update formulas.

Parameter of interest: θ = career average.

Data: $x = 85$ = this year's percentage.

Prior: $\theta \sim \mathrm{N}(75, 36)$

Likelihood $x \sim \mathrm{N}(\theta, 16)$. So $\phi(x|\theta) = c_1 \mathrm{e}^{-(x-\theta)^2/2\cdot 16}$.

The updating weights are

$$a = 1/36, \quad b = 1/16, \quad a + b = 52/576 = 13/144.$$

Therefore

$$\mu_{\text{post}} = (75/36 + 85/16)/(13/144) \approx 81.9, \qquad \sigma^2_{\text{post}} = 144/13 \approx 11.1.$$

The posterior pdf is $f(\theta|x = 85) \sim \mathrm{N}(81.9, 11.1)$.

In class examples and discussion

1. Likelihood principle

Suppose the prior has been set. Let x_1 and x_2 be two sets of data. Which of the following are true?

(a) *If the likelihoods $\phi(x_1|\theta)$ and $\phi(x_2|\theta)$ are the same then they result in the same posterior.*

(b) *If the likelihoods $\phi(x_1|\theta)$ and $\phi(x_2|\theta)$ are proportional (as functions of θ) then they result in the same posterior.*

(c) *If two likelihood functions are proportional then they are equal.*

Solution: a: true;

b: true, scale factors don't matter

c: false

Extra problems

Extra 1. Conjugate priors

Which of the following likelihood/prior pairs are conjugate?

	hypothesis	data	prior	likelihood
(a) Exponential/Normal	$\theta \in [0, \infty)$	x	$\mathrm{N}(\mu_{\text{prior}}, \sigma^2_{\text{prior}})$	$\exp(\theta)$
	θ	x	$c_1 \exp\left(-\frac{(\theta-\mu_{\text{prior}})^2}{2\sigma^2_{\text{prior}}}\right)$	$\theta \mathrm{e}^{-\theta x}$
(b) Exponential/Gamma	$\theta \in [0, \infty)$	x	$\text{Gamma}(a, b)$	$\exp(\theta)$
	θ	x	$c_1 \theta^{a-1} \mathrm{e}^{-b\theta}$	$\theta \mathrm{e}^{-\theta x}$
(c) Binomial/Normal	$\theta \in [0, 1]$	x	$\mathrm{N}(\mu_{\text{prior}}, \sigma^2_{\text{prior}})$	$\text{binomial}(N, \theta)$
(fixed N)	θ	x	$c_1 \exp\left(-\frac{(\theta-\mu_{\text{prior}})^2}{2\sigma^2_{\text{prior}}}\right)$	$c_2\, \theta^x (1-\theta)^{N-x}$

1. none	*2. a*	*3. b*	*4. c*
5. a,b	*6. a,c*	*7. b,c*	*8. a,b,c*

Solution: (b) is the only conjugate pair.

We have a conjugate prior if the posterior *as a function of* θ has the same form as the prior.

(a) Exponential/Normal posterior:

$$f(\theta|x) = c_1 \theta \mathrm{e}^{-\frac{(\theta-\mu_{\text{prior}})^2}{2\sigma^2_{\text{prior}}} - \theta x}$$

The factor of θ before the exponential means this is not the pdf of a normal distribution. Therefore it is not a conjugate prior.

(b) Exponential/Gamma posterior: Note, we have never learned about Gamma distributions, but it doesn't matter. We only have to check if the posterior has the same form:

$$f(\theta|x) = c_1 \theta^a \mathrm{e}^{-(b+x)\theta}$$

The posterior has the form $\text{Gamma}(a+1, b+x)$. This is a conjugate prior.

(c) Binomial/Normal: It is clear that the posterior does not have the form of a normal distribution.