

1.5 Sequential Bayesian Updating

- Last year, and so far this year: Bayes Theorem to update beliefs about θ after observing a single sample $\mathbf{x}$
 - Real-world: New samples of data coming online continuously!
 - First sample $\mathbf{x}^{(1)} + \pi(\theta) \longrightarrow \pi(\theta|\mathbf{x}^{(1)})$
 - New sample $\mathbf{x}^{(2)}$?
 - Use $\pi(\theta|\mathbf{x}^{(1)})$ as the “new prior”
 - Update once again to obtain the “new posterior” $\pi(\theta|\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$

1.5 Sequential Bayesian Updating

- Last year, and so far this year: Bayes Theorem to update beliefs about θ after observing a single sample $\mathbf{x}$
- Real-world: New samples of data coming online continuously!
 - First sample $\mathbf{x}^{(1)} + \pi(\theta) \longrightarrow \pi(\theta|\mathbf{x}^{(1)})$
 - New sample $\mathbf{x}^{(2)}$?
 - Use $\pi(\theta|\mathbf{x}^{(1)})$ as the “new prior”
 - Update once again to obtain the “new posterior” $\pi(\theta|\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$

1.5 Sequential Bayesian Updating

- Last year, and so far this year: Bayes Theorem to update beliefs about θ after observing a single sample $\mathbf{x}$
- Real-world: New samples of data coming online continuously!
- First sample $\mathbf{x}^{(1)} + \pi(\theta) \longrightarrow \pi(\theta|\mathbf{x}^{(1)})$
- New sample $\mathbf{x}^{(2)}$?
 - Use $\pi(\theta|\mathbf{x}^{(1)})$ as the “new prior”
 - Update once again to obtain the “new posterior” $\pi(\theta|\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$

1.5 Sequential Bayesian Updating

- Last year, and so far this year: Bayes Theorem to update beliefs about θ after observing a single sample $\mathbf{x}$
- Real-world: New samples of data coming online continuously!
- First sample $\mathbf{x}^{(1)} + \pi(\theta) \longrightarrow \pi(\theta|\mathbf{x}^{(1)})$
- New sample $\mathbf{x}^{(2)}$?
 - Use $\pi(\theta|\mathbf{x}^{(1)})$ as the “new prior”
 - Update once again to obtain the “new posterior” $\pi(\theta|\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$

1.5 Sequential Bayesian Updating

In other words, we might replace the usual mantra

$$\text{Posterior} \propto \text{Prior} \times \text{Likelihood}$$

with

$$\text{Revised} \propto \text{Current} \times \text{New likelihood},$$

with the ability to update our likelihood function with each new sample of data. Why does this work?

1.5 Sequential Bayesian Updating

In other words, we might replace the usual mantra

$$\text{Posterior} \propto \text{Prior} \times \text{Likelihood}$$

with

$$\text{Revised} \propto \text{Current} \times \text{New likelihood},$$

with the ability to update our likelihood function with each new sample of data. Why does this work?

1.5 Sequential Bayesian Updating

Suppose we have data assumed drawn from a probability distribution with parameter θ

1. Before we make any observations at all, our prior beliefs about θ are represented by $\pi(\theta)$. We might refer to this as our “original prior”
2. Now we observe a sample $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$; our posterior is then

$$\pi(\theta|\mathbf{x}) \propto \pi(\theta) \times f(\mathbf{x}|\theta) = \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta),$$

assuming our observations are a random sample

1.5 Sequential Bayesian Updating

Suppose we have data assumed drawn from a probability distribution with parameter θ

1. Before we make any observations at all, our prior beliefs about θ are represented by $\pi(\theta)$. We might refer to this as our “original prior”
2. Now we observe a sample $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$; our posterior is then

$$\pi(\theta|\mathbf{x}) \propto \pi(\theta) \times f(\mathbf{x}|\theta) = \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta),$$

assuming our observations are a random sample

1.5 Sequential Bayesian Updating

Suppose we have data assumed drawn from a probability distribution with parameter θ

1. Before we make any observations at all, our prior beliefs about θ are represented by $\pi(\theta)$. We might refer to this as our “original prior”
2. Now we observe a sample $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$; our posterior is then

$$\pi(\theta|\mathbf{x}) \propto \pi(\theta) \times f(\mathbf{x}|\theta) = \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta),$$

assuming our observations are a random sample

1.5 Sequential Bayesian Updating

3. Now suppose we make a new observation x_{n+1} assumed from the same probability distribution as $\mathbf{x}$
- Beforehand our *current* beliefs about θ are represented by $\pi(\theta|\mathbf{x})$; re-label this $\pi_n(\theta)$, our “new prior”
 - Assume the *new likelihood* has the same form as the likelihood for $\mathbf{x}$
 - Then

$$\begin{aligned}\text{Revised} &\propto \text{Current} \times \text{New likelihood} \\ \pi_{n+1}(\theta|x_{n+1}) &\propto \pi_n(\theta) \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta|\mathbf{x}) \times f(x_{n+1}|\theta) \\ &\propto \left\{ \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta) \right\} \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta) \times \prod_{i=1}^{n+1} f(x_i|\theta) \\ &\propto \text{original prior} \times \text{full likelihood}\end{aligned}$$

1.5 Sequential Bayesian Updating

3. Now suppose we make a new observation x_{n+1} assumed from the same probability distribution as $\mathbf{x}$
- Beforehand our *current* beliefs about θ are represented by $\pi(\theta|\mathbf{x})$; re-label this $\pi_n(\theta)$, our “new prior”
 - Assume the *new likelihood* has the same form as the likelihood for $\mathbf{x}$
 - Then

$$\begin{aligned}\text{Revised} &\propto \text{Current} \times \text{New likelihood} \\ \pi_{n+1}(\theta|x_{n+1}) &\propto \pi_n(\theta) \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta|\mathbf{x}) \times f(x_{n+1}|\theta) \\ &\propto \left\{ \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta) \right\} \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta) \times \prod_{i=1}^{n+1} f(x_i|\theta) \\ &\propto \text{original prior} \times \text{full likelihood}\end{aligned}$$

1.5 Sequential Bayesian Updating

3. Now suppose we make a new observation x_{n+1} assumed from the same probability distribution as $\mathbf{x}$
- Beforehand our *current* beliefs about θ are represented by $\pi(\theta|\mathbf{x})$; re-label this $\pi_n(\theta)$, our “new prior”
 - Assume the *new likelihood* has the same form as the likelihood for $\mathbf{x}$
 - Then

$$\begin{aligned}\text{Revised} &\propto \text{Current} \times \text{New likelihood} \\ \pi_{n+1}(\theta|x_{n+1}) &\propto \pi_n(\theta) \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta|\mathbf{x}) \times f(x_{n+1}|\theta) \\ &\propto \left\{ \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta) \right\} \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta) \times \prod_{i=1}^{n+1} f(x_i|\theta) \\ &\propto \text{original prior} \times \text{full likelihood}\end{aligned}$$

1.5 Sequential Bayesian Updating

3. Now suppose we make a new observation x_{n+1} assumed from the same probability distribution as $\mathbf{x}$
- Beforehand our *current* beliefs about θ are represented by $\pi(\theta|\mathbf{x})$; re-label this $\pi_n(\theta)$, our “new prior”
 - Assume the *new likelihood* has the same form as the likelihood for $\mathbf{x}$
 - Then

$$\begin{aligned}\text{Revised} &\propto \text{Current} \times \text{New likelihood} \\ \pi_{n+1}(\theta|x_{n+1}) &\propto \pi_n(\theta) \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta|\mathbf{x}) \times f(x_{n+1}|\theta) \\ &\propto \left\{ \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta) \right\} \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta) \times \prod_{i=1}^{n+1} f(x_i|\theta) \\ &\propto \text{original prior} \times \text{full likelihood}\end{aligned}$$

1.5 Sequential Bayesian Updating

3. Now suppose we make a new observation x_{n+1} assumed from the same probability distribution as $\mathbf{x}$
- Beforehand our *current* beliefs about θ are represented by $\pi(\theta|\mathbf{x})$; re-label this $\pi_n(\theta)$, our “new prior”
 - Assume the *new likelihood* has the same form as the likelihood for $\mathbf{x}$
 - Then

$$\begin{aligned}\text{Revised} &\propto \text{Current} \times \text{New likelihood} \\ \pi_{n+1}(\theta|x_{n+1}) &\propto \pi_n(\theta) \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta|\mathbf{x}) \times f(x_{n+1}|\theta) \\ &\propto \left\{ \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta) \right\} \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta) \times \prod_{i=1}^{n+1} f(x_i|\theta) \\ &\propto \text{original prior} \times \text{full likelihood}\end{aligned}$$

1.5 Sequential Bayesian Updating

3. Now suppose we make a new observation x_{n+1} assumed from the same probability distribution as $\mathbf{x}$
- Beforehand our *current* beliefs about θ are represented by $\pi(\theta|\mathbf{x})$; re-label this $\pi_n(\theta)$, our “new prior”
 - Assume the *new likelihood* has the same form as the likelihood for $\mathbf{x}$
 - Then

$$\begin{aligned}\text{Revised} &\propto \text{Current} \times \text{New likelihood} \\ \pi_{n+1}(\theta|x_{n+1}) &\propto \pi_n(\theta) \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta|\mathbf{x}) \times f(x_{n+1}|\theta) \\ &\propto \left\{ \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta) \right\} \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta) \times \prod_{i=1}^{n+1} f(x_i|\theta) \\ &\propto \text{original prior} \times \text{full likelihood}\end{aligned}$$

1.5 Sequential Bayesian Updating

3. Now suppose we make a new observation x_{n+1} assumed from the same probability distribution as $\mathbf{x}$
- Beforehand our *current* beliefs about θ are represented by $\pi(\theta|\mathbf{x})$; re-label this $\pi_n(\theta)$, our “new prior”
 - Assume the *new likelihood* has the same form as the likelihood for $\mathbf{x}$
 - Then

$$\begin{aligned}\text{Revised} &\propto \text{Current} \times \text{New likelihood} \\ \pi_{n+1}(\theta|x_{n+1}) &\propto \pi_n(\theta) \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta|\mathbf{x}) \times f(x_{n+1}|\theta) \\ &\propto \left\{ \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta) \right\} \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta) \times \prod_{i=1}^{n+1} f(x_i|\theta) \\ &\propto \text{original prior} \times \text{full likelihood}\end{aligned}$$

1.5 Sequential Bayesian Updating

3. Now suppose we make a new observation x_{n+1} assumed from the same probability distribution as $\mathbf{x}$
- Beforehand our *current* beliefs about θ are represented by $\pi(\theta|\mathbf{x})$; re-label this $\pi_n(\theta)$, our “new prior”
 - Assume the *new likelihood* has the same form as the likelihood for $\mathbf{x}$
 - Then

$$\begin{aligned}\text{Revised} &\propto \text{Current} \times \text{New likelihood} \\ \pi_{n+1}(\theta|x_{n+1}) &\propto \pi_n(\theta) \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta|\mathbf{x}) \times f(x_{n+1}|\theta) \\ &\propto \left\{ \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta) \right\} \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta) \times \prod_{i=1}^{n+1} f(x_i|\theta) \\ &\propto \text{original prior} \times \text{full likelihood}\end{aligned}$$

1.5 Sequential Bayesian Updating

3. Now suppose we make a new observation x_{n+1} assumed from the same probability distribution as $\mathbf{x}$
- Beforehand our *current* beliefs about θ are represented by $\pi(\theta|\mathbf{x})$; re-label this $\pi_n(\theta)$, our “new prior”
 - Assume the *new likelihood* has the same form as the likelihood for $\mathbf{x}$
 - Then

$$\begin{aligned}\text{Revised} &\propto \text{Current} \times \text{New likelihood} \\ \pi_{n+1}(\theta|x_{n+1}) &\propto \pi_n(\theta) \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta|\mathbf{x}) \times f(x_{n+1}|\theta) \\ &\propto \left\{ \pi(\theta) \times \prod_{i=1}^n f(x_i|\theta) \right\} \times f(x_{n+1}|\theta) \\ &\propto \pi(\theta) \times \prod_{i=1}^{n+1} f(x_i|\theta) \\ &\propto \text{original prior} \times \text{full likelihood}\end{aligned}$$

Example 1.15

Table 1.6 provides an update to table 1.1, giving the number of cases of foodbourne botulism in England and Wales between 2006 and 2010. Use your solution to Example 1.1, and sequential Bayesian updating, to obtain your new posterior distribution for θ .

Year	2006	2007	2008	2009	2010
Cases	1	0	0	1	1

Table: Number of cases of foodbourne botulism in England and Wales, 2006–2010

Solution

...

Example 1.16

Devise a generic sequential Bayesian updating plan for successive random samples from a Poisson distribution, assuming a conjugate prior specification.

Solution

...

1.6 Learning objectives

By the end of this chapter, you should be able to:

1. Determine the likelihood function using a random sample from **any** distribution
2. Combine this likelihood function with **any** prior distribution to obtain the posterior distribution
3. Name the posterior distribution if it is a “standard” distribution listed in these notes or on the exam paper – this list may well include distributions that are standard within the subject but which you have not met before. If the posterior is not a “standard” distribution then it is okay just to give its density (or probability function) up to a constant.
4. Do all the above for a particular data set or for a general case with random sample $x_1, \dots, x_n$

5. Describe the different levels of prior information; determine and use conjugate priors and vague priors
6. Determine the asymptotic posterior distribution
7. Determine the predictive distribution, particularly when having a random sample from any distribution and a conjugate prior via Candidate's formula
8. Describe and calculate the confidence intervals, HDIs and prediction intervals
9. Operate sequential Bayesian updating schemes