# Decision Theory and Bayesian Analysis

**Dr. Vilda Purutcuoglu**

[1]

[1]Edited by Anil A. Aksu based on lecture notes of STAT 565 course by Dr. Vilda Purutcuoglu

# Contents

# LECTURE 1
# Bayesian Paradigm

## 1.1. Bayes theorem for distributions

If $A$ and $B$ are two events,

$$P(A \mid B) = \frac{P(A)P(B \mid A)}{P(B)}. \tag{1.1}$$

This is just a direct consequence of the multiplication law of probabilities that says we can express $P(A \mid B)$ as either $P(A)P(B \mid A)$ or $P(B)P(A \mid B)$. For discrete distributions, if $Z, Y$ are discrete random variables

$$P(Z = z \mid Y = y) = \frac{P(Z = z)P(Y = y \mid Z = z)}{P(Y = y)}. \tag{1.2}$$

- How many distributions do we deal with here?

We can express the denominator in terms of the distribution in the numerator[**1**].

(1.3)

$$P(Y = y) = \sum_z P(Y = y, Z = z) = \sum_z P(Z = z)P(Y = y \mid Z = z).$$

- This is sometimes called the law of total probability

In this context, it is just an expression of the fact that as $z$ ranges over the possible values of $Z$ , the probabilities on the left hand-side of equation 1.2 make up the distribution of $Z$ given $Y = y$, and so they must add up to one. The extension to continuous distribution is easy. If $Z, Y$ are continuous random variable,

$$f(Z \mid Y) = \frac{f(Z)f(Y \mid Z)}{f(Y)}. \tag{1.4}$$

where the denominator is now expressed as an integral:

$$f(Y) = \int f(Z)f(Y \mid Z)\mathrm{d}Z. \tag{1.5}$$

$$f = \begin{cases} continous & name? \\ discrete & name? \end{cases} \tag{1.6}$$

## 1.2. How Bayesian Statistics Uses Bayes Theorem

**Theorem 1.7** (Bayes' theorem)**.**

$$P(A \mid B) = \frac{P(A)P(B \mid A)}{P(B)}$$

*$P(B)$=if we are interested in the event $B$, $P(B)$ is the initial or prior probability of the occurence of event $B$. Then we observe event $A$ $P(B \mid A)$ = How likely $B$ is when $A$ is known to have occurred is the posterior probability $P(B \mid A)$.*

Bayes' theorem can be understood as a formula for updating from prior to posterior probability, the updating consists of multiplying by the ratio $P(B \mid A)/P(A)$. It describes how a probability changes as we learn new information. Observing the occurrence of $A$ will increase the probability of $B$ if $P(B \mid A) > P(A)$. From the law of total probability,

$$P(A) = P(A \mid B)P(B) + P(A \mid B^c)P(B^c). \tag{1.8}$$

where $P(B^c) = 1 - P(B)$.

**Lemma 1.9.**

$$P(A \mid B) - P(A) = \frac{P(A) - P(A \mid B^c)P(B^c)}{1 - P(B^c)} - P(A)$$

**Proof.**

$$P(A \mid B) - P(A) = \frac{P(A) - P(A \mid B^c)P(B^c) - P(A) + P(A)P(B^c)}{P(B)}$$

$$P(A \mid B) - P(A) = \frac{P(B^c)(P(A) - P(A \mid B^c))}{P(B)}$$

$$P(A \mid B) - P(A) = P(B^c)\left(\frac{P(B)P(A \mid B) + P(B^c)P(A \mid B^c)}{P(B)} - \frac{P(A \mid B^c)}{P(B)}\right)$$

$$P(A \mid B) - P(A) = P(B^c)\left(P(A \mid B) - \frac{P(A \mid B^c)(1 - P(B^c))}{P(B)}\right)$$

$$P(A \mid B) - P(A) = P(B^c)(P(A \mid B) - P(A \mid B^c))$$

□

### 1.2.1. Generalization of the Bayes' Theorem

Let $B_1, ..., B_n$ be a set of mutually exclusive events. Then

$$P(B_r \mid A) = \frac{P(B_r)P(A \mid B_r)}{P(A)} = \frac{P(B_r)P(A \mid B_r)}{\sum_{i=1}^{n} P(B_r)P(A \mid B_r)}. \tag{1.10}$$

- Assuming that $P(B_r) > 0$,$P(A \mid B) > P(A)$ if and only if $P(A \mid B) > P(A \mid B^c)$.

- In Bayesian inference we use Bayes' theorem in a particular way.
- $Z$ is the parameter (vector) $\theta$.
- $Y$ is the data (vector) $X$.

So we have

$$f(\theta \mid X) = \frac{f(\theta)f(X \mid \theta)}{f(X)} \tag{1.11}$$

$$f(X) = \int f(\theta)f(X \mid \theta)\mathrm{d}\theta. \tag{1.12}$$

$$f(\theta) = prior. \tag{1.13}$$

$$f(\theta \mid X) = posterior. \tag{1.14}$$

$$f(X \mid \theta) = likelihood. \tag{1.15}$$

### 1.2.2. Interpreting our sense

How do we interpret the things we see, hear, feel, taste or smell?

**Example 1.2.1.** I hear a song on the radio I identify the singer as Robbie Williams. Why do I think it's Robbie Williams?. Because he sounds like that. Formally, $P($ What I hear Robbie Williams $) >> P($What I hear someone else $)$

**Example 1.2.2.** I look out of the window and see what appears to be a tree. It has a big, dark coloured part sticking up out of the ground that branches into thinner sticks and on the ends of these are small green things. Clearly, $P(view \mid tree)$ is high and $P(view \mid car)$ or $P(view \mid Robbie\ \ Williams)$ are very small. But $P(view \mid carboard\ \ cutout\ \ cunningly\ \ painted\ \ to\ \ look\ \ like\ \ a\ \ tree)$ is also very high. Maybe even higher than $P(view \mid tree)$ in the sense that what I see looks almost like a tree.
Does this mean I should now believe that I am seeing a cardboard cut-out cunningly painted to look like a tree? No because it is much less likely to begin with than a red tree.

In statistical terms, consider some data $X$ and some unknown parameter $\theta$. The first step in any statistical analysis is to build a model that links the data to unknown parameters and the main function of this model is to allow us to state the probability of observing any data given any specified values of the parameters. That is the model defines $f(x \mid \theta)$.

When we think of $f(x \mid \theta)$ as a function of $\theta$ for fixed observed data $X$, we call it likelihood function and it by $L(\theta, X)$.

- So how can we combine this with our example?

This perspective underlies the differences between the two main theories of statistical inference.

- Frequentist inference essentially uses only the likelihood, it does not recognize $f(\theta)$.
- Bayesian inference uses both likelihood and $f(\theta)$.

The principal distinguishing feature of Bayesian inference as opposed to frequentist inference is its use of $f(\theta)$.

## 1.3. Prior to Posterior

We refer to $f(\theta)$ as the prior distribution of $\theta$. It represents knowledge about $\theta$ prior to observing the data $X$. We refer to $f(\theta \mid X)$ as the posterior distribution of $\theta$ and it represents knowledge about $\theta$ after observing $X$.

- So we have two sources of information about $\theta$.
- Here $f(x)$ does not depend on $\theta$. Thus $\int f(\theta \mid x)\mathrm{d}\theta = 1$. Since $f(x)$ is a constant within the integral, we can take it outside to get $1 = f^{-1}(x) \int f(\theta) f(x \mid \theta)\mathrm{d}\theta$.
- $f(\theta \mid x) =\propto f(\theta) f(x \mid \theta) \propto f(\theta) L(\theta; x)$ (the posterior is proportional to the prior times the likelihood).
- The constant that we require to scale the right hand side to integrate to 1 is usually called the normalizing constant. If we haven't dropped any constants form $f(\theta)$ or $f(x \mid \theta)$, then the normalising constant is just $f^{-1}(x)$, otherwise it also restores any dropped constants.

## 1.4. Triplot

If for any value of $\theta$, we have either $f(\theta) = 0$ or $f(x \mid \theta) = 0$, then we will also have $f(\theta \mid x) = 0$. This is called the property of zero preservation. So if either:

- the prior information says that this $\theta$ value is impossible
- the data say that this value of $\theta$ is impossible because if it were the true value, then the observed data would have been impossible, then the posterior distribution confirms that this value of $\theta$ is impossible.

**Definition 1.16.** Crowwell's Rule: If either information source completely rules out a specific $\theta$, then the posterior must rule it out too.

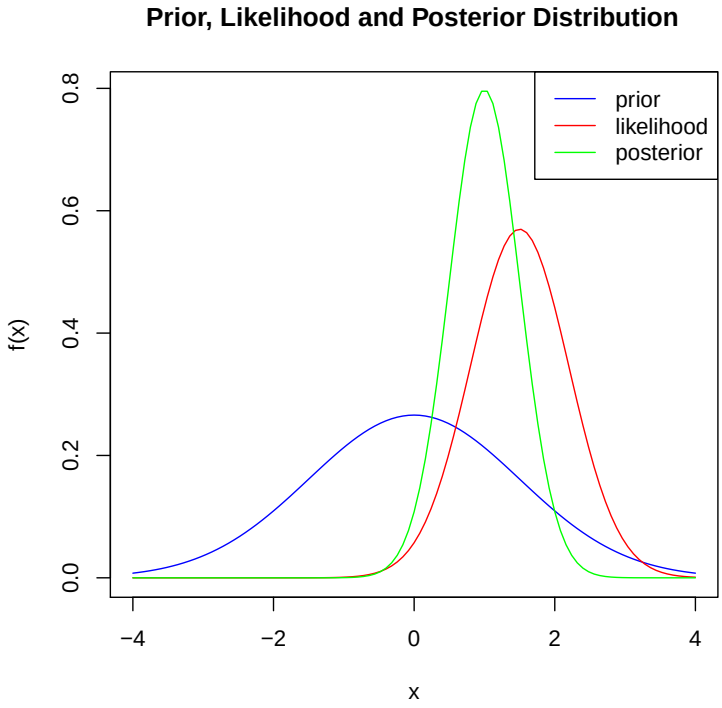


**Figure 1.** Triplot of prior, likelihood and posterior.

This means that we should be very careful about giving zero probability to something unless it is genuinely impossible. Once something has zero probability then no amount of further evidence can cause it to have a non-zero posterior probability.

- More generally, $f(\theta \mid x)$ will be low if either $f(\theta)$ is very small. We will tend to find that $f(x \mid \theta)$ is large when both $f(\theta)$ and $f(x \mid \theta)$ are relatively large, so that this $\theta$ value is given support by both information sources.

When $\theta$ is a scalar parameter, a useful diagram is the triplot, which shows the prior, likelihood and posterior on the same graph. An example is in Figure 1.[1]

A strong information source in the triplot is indicated by a curve that is narrow (and therefore, because it integrates to one, also has a high peak). A narrow curves concentrates on a small range of $\theta$ values, and thereby "rules out" all values of $\theta$ outside that range.

[1]All plots are generated in R, relevant codes are provided in Appendix R Codes

- Over the range $\theta < -1$, the likelihood: low
- Over the range $\theta > 3$,the likelihood: low
- Values of $\theta$ between $-1$ and 3, the likelihood: high
- The maximum value of the posterior at: 1
- The Maximum likelihood estimation (MLE) of $\theta$ is $\approx 2$

### 1.4.1. Normal Mean

For example, suppose that $X_1, X_2, ..., X_n$ are iid $\mathcal{N}(\mu, \sigma^2)$ and $\sigma^2$ is known. Then the likelihood is :

$$f(x \mid \mu) = \prod_{i=1}^{n} f(x_i \mid \mu) = \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2\sigma^2}(x_i - \mu)^2\right) \propto \exp\left(-\frac{1}{2\sigma^2}(x_i - \mu)^2\right). \tag{1.17}$$

As,

$$\begin{aligned}\sum (x_i - \bar{x} + \bar{x} - \mu)^2 &= \sum (x_i - \bar{x})^2 + n(\bar{x} - \mu)^2 + 2(\bar{x} - \mu)\sum (x_i - \bar{x}) \\ &= \sum (x_i - \bar{x})^2 + n(\bar{x} - \mu)^2 \\ &\propto \exp\left(-\frac{1}{2\sigma^2} n(\bar{x} - \mu)^2\right).\end{aligned} \tag{1.18}$$

Note that $2(\bar{x} - \mu)\sum (x_i - \bar{x}) = 0$ as $\sum (x_i - \bar{x}) = 0$. Suppose the prior distribution for $\mu$ is normal:

$$\mu \sim \mathcal{N}(m, v). \tag{1.19}$$

Then applying Bayes' theorem we have:

$$\begin{aligned}f(\mu \mid x) &\propto \underbrace{\exp\left(-\frac{1}{2\sigma^2} n(\bar{x} - \mu)^2\right)}_{f(x|\mu)} \underbrace{\exp\left(-\frac{1}{2v} n(\mu - m)^2\right)}_{f(\mu)} \\ &= \exp\left(-\frac{\theta}{2}\right).\end{aligned} \tag{1.20}$$

Note that

$$\theta = n\sigma^{-2}(\bar{x} - \mu) + v^{-1}(\mu - m)^2 = (v^*)^{-1}(\mu - m^*)^2 + R \tag{1.21}$$

and

$$v^* = (n\sigma^{-2} + v^{-1})^{-1} \tag{1.22}$$

$$m^* = (n\sigma^{-2} + v^{-1})^{-1}(n\sigma^{-2}\bar{x} + v^{-1}m) = a\bar{x} + (1 - a)m \tag{1.23}$$

where $a = n\sigma^-2/(n\sigma^{-2} + v^{-1})$

$$R = (n^{-1}\sigma^2 + v)(\bar{x} - m)^2 \tag{1.24}$$

Therefore,

$$f(\mu \mid x) \propto \exp\left(-\frac{1}{2v} n(\mu - m)^2\right) \tag{1.25}$$

and we have shown that the posterior distribution is normal too: $\mu \mid x \sim \mathcal{N}(m^*, v^*)$

- $m^*$ = weighted average of the mean $m$ and the usual frequentist data-only estimate $\bar{x}$.
  The weights $\propto$:
- Bayes' theorem typically works in this way. We usually find that posterior estimates are compromises between prior estimates and data based estimates and tend to be closer whichever information source is stronger. And we usually find that the posterior variance is smaller than the prior variance.

### 1.4.2. Weak Prior Information

It is the case where the prior information is much weaker that the data. This will occur, for instance, if we do not have strong information about $Q$ before seeing the data, and if there are lots of data. Then in triplot, the prior distribution will be much broader and flatter that the likelihood. So the posterior is approximately proportional to the likelihood.

**Example 1.4.1.** In the normal mean analysis, we get weak prior information by letting the prior precision of $v^{-1}$ become small. Then $m^* \to \bar{x}$ and $v^* \to \sigma^2/n$ so that the posterior distribution of $\mu$ corresponds very closely with standard frequentist theory.

# LECTURE 2
# Some Common Probability Distributions

**Binomial** on $Y \in \{0, 1, ..., n\}$ with parameters $n \in \{1, 2, 3, ...\}$ and $p \in (0, 1)$ is denoted by $Bi(n, p)$ and

$$f(y \mid n, p) = \binom{n}{y} p^y (1-p)^{n-y} \tag{2.1}$$

for $y = 0, 1, ..., n$. The mean is given as:

$$E(y) = np. \tag{2.2}$$

Also the variance is given as:

$$v(y) = np(1-p). \tag{2.3}$$

**Beta** on $Y \in \{0, 1\}$ with parameters $a, b > 0$ is denoted by $Beta(p, q)$ and the density function is:

$$f(y \mid p, q) = \frac{\Gamma(p+q)}{\Gamma(p)\Gamma(q)} y^{p-1}(1-y)^{q-1} \tag{2.4}$$

for $y \in (0, 1)$. The mean is given as:

$$E(y) = \frac{p}{p+q}, \tag{2.5}$$

Also the variance is given as:

$$v(y) = \frac{pq}{(p+q)^2(p+q+1)}. \tag{2.6}$$

$B(p, q) = \int_0^1 y^{p-1}(1-y)^{q-1} \mathrm{d}y$ is the beta function and defined to be the normalizing constant for this density.

- In beta distribution, $p$ and $q$ change the shape of the distribution. Discuss!

**Uniform distribution** on $Y \in \{l, r\}$ where $-\infty < l < r < \infty$ is denoted by uniform $(l, r)$ and its pdf is:

$$f(y \mid l, r) = \frac{1}{r-l} \tag{2.7}$$

for $y \in \{l, r\}$. The mean is given as:

$$E(y) = \frac{l+r}{2}, \tag{2.8}$$

Also the variance is given as:

$$v(y) = \frac{(r-l)^2}{12}. \tag{2.9}$$

**Poisson distribution** on $Y \in \{0, 1, 2, ..., \}$ with parameter $\theta > 0$ is denoted by $Poisson(\theta)$ and its pdf is:

$$f(y \mid \theta) = \frac{\exp(-\theta)\theta^y}{y!} \tag{2.10}$$

for $y = 0, 1, 2, ...$ The mean and the variance are given as[**4**]:

$$E(y) = v(y) = \theta. \tag{2.11}$$

**Gamma distribution** on $Y > 0$ with shape parameter $\alpha > 0$ and rate parameter $\lambda > 0$ is denoted by $Gamma(\alpha, \lambda)$ and the corresponding density is:

$$f(y \mid \alpha, \lambda) = \frac{\lambda^\alpha}{\Gamma(\alpha)} y^{\alpha-1} \exp(-\lambda y) \tag{2.12}$$

for $y > 0$. The mean is given as:

$$E(y) = \frac{\alpha}{\lambda}, \tag{2.13}$$

Also the variance is given as:

$$v(y) = \frac{\alpha}{\lambda^2}. \tag{2.14}$$

Note that

$$\exp(\lambda) = Gamma(1, \lambda). \tag{2.15}$$

**Univariate normal distribution** on $Y \in \mathbb{R}$ with $\mu \in \mathbb{R}$ and variance $\sigma^2 > 0$ is denoted by $\mathcal{N}(\mu, \sigma^2)$ and its pdf is:

$$f(y \mid \mu, \sigma^2) = \frac{1}{\sigma}\left(\frac{1}{2\pi}\right)^{1/2} \exp\left\{-\frac{1}{2\sigma^2}(y-\mu)^2\right\}. \tag{2.16}$$

The mean is given as:

$$E(y) = \mu, \tag{2.17}$$

Also the variance is given as:

$$v(y) = \sigma^2. \tag{2.18}$$

**K-variate normal distribution** on $Y \in \mathbb{R}^k$ with vector $\mathbf{b} \in \mathbb{R}^k$ and positive definite symmetric (PDS) covariance matrix $\mathbf{C}$ is denoted by $\mathcal{N}_k(\mathbf{b}, \mathbf{C})$ and the corresponding density function is:

$$f(y \mid \mathbf{b}, \mathbf{C}) = \frac{1}{\underbrace{|\mathbf{C}|^{1/2}}_{determinant}} \frac{1}{(2\pi)^{k/2}} \exp\left\{-\frac{1}{2\sigma^2}(y-\mathbf{b})^T\mathbf{C}^{-1}(y-\mathbf{b})\right\}. \tag{2.19}$$

The mean is given as:

$$E(y) = \mathbf{b}, \tag{2.20}$$

And the covariance matrix is given as:

$$Cov(y) = \mathbf{C}. \tag{2.21}$$

## 2.1. Posterior

Not only the beta distributions are the simplest and the most convenient distributions for a random variable confined to $[0, 1]$, they also work very nicely as prior distribution for a binomial observation. If $\theta \sim Be(p, q)$ then

$$f(x \mid \theta) = \binom{n}{x}\theta^x(1-\theta)^{n-x}. \tag{2.22}$$

for $x = 1, 2, ..., n$.

$$f(\theta) = \frac{1}{Be(p,q)}\theta^{p-1}(1-\theta)^{q-1} \tag{2.23}$$

where $0 \leq \theta \leq 1$ and $p, q > 0$.

$$f(x) = \int f(\theta)f(x \mid \theta)\mathrm{d}\theta = \binom{n}{r}\frac{1}{Be(p,q)}\int_0^1 \theta^{p+x-1}(1-\theta)^{q+n-x-1}\mathrm{d}\theta = \binom{n}{r}\frac{Be(p+x, q+n-x)}{Be(p,q)}. \tag{2.24}$$

From

$$f(\theta \mid x) = \frac{f(\theta)f(x \mid \theta)}{f(x)}. \tag{2.25}$$

(2.26)
$$f(\theta \mid x) = \frac{\theta^{p+x-1}(1-\theta)^{q+n-x-1}}{Be(p+x, q+n-x)} \propto \underbrace{\theta^{p-1}(1-\theta)^{q-1}}_{Beta \quad part} \underbrace{\theta^{x}(1-\theta)^{n-x}}_{Binomial \quad part}.$$

So $(\theta \mid x) \propto Beta(p+x, q+n-x)$. The posterior mean is:

$$E(\theta \mid x) = \frac{p+x}{p+q+n} = \frac{p+q}{p+q+n}E(\theta) + \frac{n}{p+q+n}\hat{\theta} \tag{2.27}$$

where $\hat{\theta} = x/n$. The posterior variance is:

$$\begin{aligned} v(\theta \mid x) &= \frac{(p+x)(q+n-x)}{(p+q+n)^2(p+q+n+1)} \\ &= \frac{E(\theta)(1-E(\theta \mid x))}{p+q+n+1} \end{aligned} \tag{2.28}$$

But,

$$v(\theta) = \frac{E(\theta)(1-\theta)}{p+q+1} \tag{2.29}$$

So the posterior has higher relative precision that the prior.

**SPECIAL NOTE:**

The classical theory of estimation regards an estimator as good if it is unbiased and has small variance, or more generally if its mean-square-error is small. The MSE is an average squared error where the error is the difference between $\theta$, i.e. $y$ in previous notation , and the estimate $t$. In accordance with classical theory, the average is taken with respect to the sampling distribution of the estimator.

In Bayesian inferenc, $\theta$ is a random variable and it is therefore appropriate to average the squared error with respect to the posterior distribution of $\theta$. Consider

$$\begin{aligned} E\left\{(t-\theta)^2 \mid x\right\} &= E(t^2 \mid x) - E(2t\theta \mid x) + E(\theta^2 \mid x) \\ &= t^2 - E(2t\theta \mid x) + E(\theta^2 \mid x) \\ &= \{t - E(\theta \mid x)\}^2 + v(\theta \mid x). \end{aligned} \tag{2.30}$$

Therefore the estimate $t$ which minimizes posterior expected square error is $t = E(\theta \mid x)$, the posterior mean. The posterior mean can therefore be seen as an estimate of $\theta$ which is the best in the sense of minimizing expected squared error. This is distinct from, but clearly related to, its more natural role as a useful summary of location of the posterior distribution.

## 2.2. Weak Prior

If we reduce the prior relative precision to zero by setting $p = q = 0$, we get $\theta \mid x \sim Be(x, n - x)$. Then $E(\theta \mid x) = \hat{\theta}$ and $v(\theta \mid x) = \hat{\theta}(1-\hat{\theta})/(n+1)$ results which nicely parallel standard frequentist theory.

- Notice that we are not really allowed to let either parameter of the beta distribution be zero. However, by making $p$ and $q$ extremely small, we get as close to these results as we like. We can think of $p = q = 0$ as a defining limiting (if strictly improper) case of weak prior information.

**Example 2.2.1.** A doctor proposes a new treatment protocol for a certain kind of cancer. With current methods about 40% of patients with this cancer survive six months after diagnosis. After one year of using the new protocol, 15 patients with diagnosis, of whom 6 survived. After two years a further 55 patients have been followed to teh six months mark, of whom 28 survived. So in total we have 34 patients surviving out of 70.
Let $\theta$ be the true success rate of the new treatment protocol, i.e. the true proportion of patients who survive 6 months and we wish to make comparison of $\theta$ with the current survival rate 40%.
Suppose that the doctor in charge has prior information leading her to assign a prior distribution with expectation $E(\theta) = 0.45$, i.e. expects a slight improvement over the existing protocol, form 40% to 45%, however her prior standard deviation is 0.07, $v(\theta) = 0.07^2 = 0.0049$

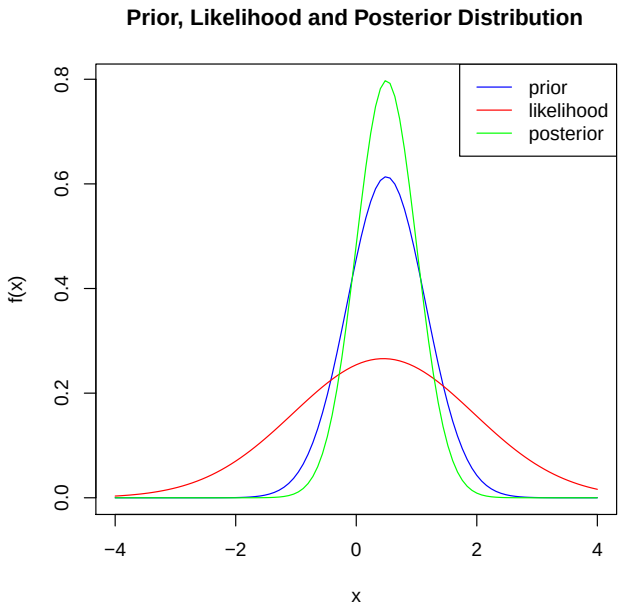


**Figure 1.** The triplot from first year's data.

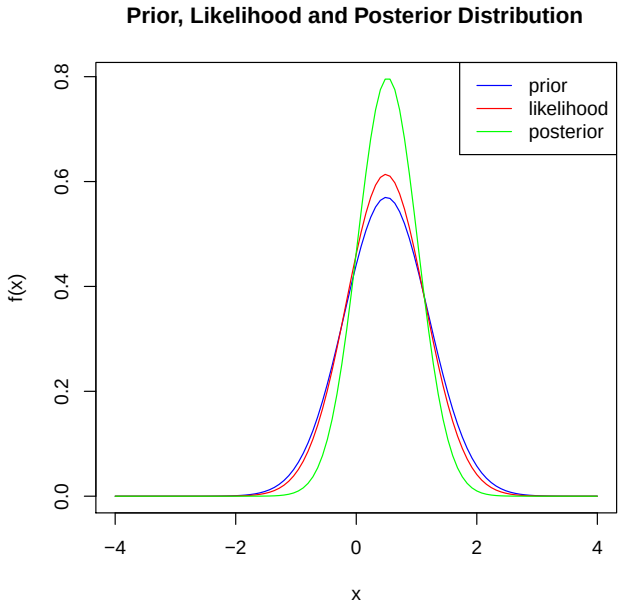


**Figure 2.** The triplot from two years' data.

## 2.3. Sequential Updating

In the last example we pooled the data from the two years and went back to the original prior distribution to use Bayes' theorem. We did not need to do this. A nice feature of Bayes' theorem is the possibility of updating sequentially, incorporating data as they arrive. In this case, consider the data to be just the new patients observed to a six months follow-up during the second year. These comprise 55 patients, of whom 28 had survived. The doctor could consider these as the data $x$ with $n = 55$ and $r = 28$. What would the prior information be? Clearly, the prior distribution should express her information prior to obtaining these new data, i.e. after the first years' data, so her prior for this second analysis is her posterior distribution form the first. This was $Be(28.28, 36.23)$. Combining this prior with the new data gives the same posterior $Be(28.28 + 28, 36.23 + 27) = Be(56.28, 63.23)$ as before. This simply confirms that we can get to the posterior distribution.

- In a single step, combining all the data with a prior distribution representing information available before any of the data were obtained.
- Sequentially, combining each item or block of new data with a prior distribution representing information available just before the new data were obtained (but after getting data previously received).

## 2.4. Normal Sample

Let $X_1, X_2, ..., X_n$ be from $\mathcal{N}(\mu, \sigma^2)$. $\theta = (\mu, \sigma^2) \rightarrow$ unknown parameters. The likelihood is:

$$\begin{aligned} f(x \mid \mu, \sigma^2) &= \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi}\sigma} \exp\left\{-\frac{1}{2\sigma^2}(x_i - \mu)^2\right\} \\ &\propto \sigma^{-n} \exp\left[-\frac{1}{2\sigma^2}\left\{n(\bar{x} - \mu)^2 + S^2\right\}\right] \end{aligned} \tag{2.31}$$

where $S^2 = \sum_{i=1}^{n}(x_i - \bar{x})^2$

## 2.5. NIC distributions

For the prior distribution, we now need a joint distribution for $\mu$ and $\sigma^2$

**Definition 2.32.** The normal-inverse-chi-squared distribution(NIC) has density:

$$f(x \mid \mu, \sigma^2) \propto \sigma^{-(d+3)/2} \exp\left[-\frac{1}{2\sigma^2}\left\{v^{-1}(\mu - m)^2 + a\right\}\right] \tag{2.33}$$

where $a > 0$, $d > 0$ and $v > 0$.

The following facts are easy to derive about $NIC(m, v, a, d)$ distribution.

(a) The conditional distribution of $\mu$ given $\sigma^2$ is $\mathcal{N}(\mu, v\sigma^2)$ so $E(\mu \mid \sigma^2) = m$, $v(\mu \mid \sigma^2) = v\sigma^2$.
(b) The marginal distribution of $\sigma^2$ is such that $a\sigma^{-2} \sim \chi^2_d$. We say that $\sigma^2$ has the inverse-chi-square distribution $IC(a, d)$. We have $E(\sigma^2) = a/(d-2)$ if $d > 2$ and $v(\sigma^2) = 2a^2/\{(d-2)^2(d-4)\}$ if $d > 4$.
(c) The conditional distribution of $\sigma^2$ given $\mu$ is $IC(v^{-1}(\mu-m)^2 + a, d+1)$ and in particular $E(\sigma^2 \mid \mu) = (v^{-1}(\mu-m)^2+a)/(d-1)$ provided $d > 1$.
(d) The marginal distribution of $\mu$ is such that $(\mu-m)\sqrt{d}/\sqrt{av}\mu + d$. We say that $\mu$ has $t$-distribution $t_d(m, av/d)$. We have $E(\mu) = m$ if $d > 1$, and $v(\mu) = av/(d-2)$ if $d > 2$.

## 2.6. Posterior

Supposing then that the prior distribution is $NIC(m, v, a, d)$, we find

$$f(\mu, \sigma^2 \mid x) \propto \sigma^{d+n+3} \exp\left[-\frac{1}{2\sigma^2}\theta\right] \tag{2.34}$$

where $\theta = v^{-1}(\mu-m)^2 + a + n(\bar{x}-\mu) + s^2$ is a quadratic expression in $\mu$. After completing the square, we see that $\mu, \sigma^2 \mid x \propto NIC(m^*, v^*, a^*, d^*)$ where $m^* = (v^{-1}m + n\bar{x})/(v^{-1}+n)$, $v^* = (v^{-1}+n)^{-1}$, $a^* = a + S^2 + (\bar{x}-m)^2/(n^{-1}+v)$, $d^* = d + n$. To interpret these results, note first that the posterior mean of $\mu$ is $m^*$ which is a weighted average of the prior mean $m$ and the usual data only-estimate $\bar{x}$ with weights $v^{-1}$ and $n$.

The posterior mean of $\sigma^2$ is $a^*/(d^*-2)$ which is a weighted average of three terms: the prior mean $a/(d-2)$ with weight $(d-2)$, the usual data-only estimate $S^2/(n-1)$ with weight $(n-1)$ and $(\bar{x}-m)/(n^{-1}+v)$ with weight 1.

## 2.7. Weak prior

We clearly obtain weak prior information about $\mu$ by letting $v$ go to infinity or $v^{-1} \to 0$. Then $m^* = \bar{x}$, $v^* = 1/n$, $a^* = a + S^2$, because the third term disappears.

To obtain weak prior information also about $\sigma^2$, if it is usual to set $a = 0$ and $d = 1$. Then $a^* = S^2$ and $d^* = n - 1$. The resulting inference match the standard frequentist results very closely with these parameters, since we have:

$$\frac{(\mu - \bar{x})\sqrt{n}}{S/\sqrt{n-1}} \propto t_{n-1}, \tag{2.35}$$

$$\frac{S^2}{\sigma^2} \propto \chi^2_{n-1} \tag{2.36}$$

Exactly the same distribution statements underlie standard frequentist inference in this problem.

# LECTURE 3
# Inference

**Summarisation:** Here all inferences are derived form the posterior distribution. In frequentist statistics, there are three kinds of inference:

- Point estimation: unbiasedness, minimum variance estimation
- Interval estimation: Credible intervals in Bayesian inference, confidence interval in frequentist approach
- Hypothesis testing: Signifcance test.

In Bayesian inference, the posterior distribution expresses all that is known about $\theta$. So it uses an appropriate summaries of the posterior distribution to describe the main features of what we now know about $\theta$.

**Plots:** To draw the posterior density

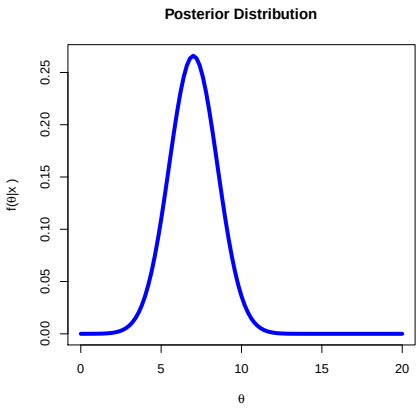


**Figure 1.** A posterior density plot.

For a bivariate parameter, we can still usefully draw the density as a perspective plot or a contour plot.

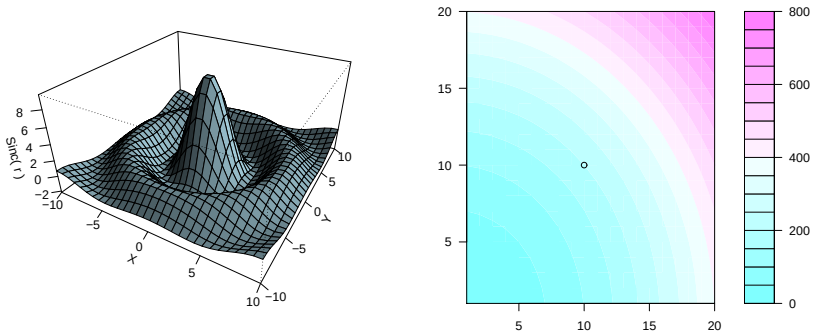


**(a)** Perspective plot **(b)** Contour plot

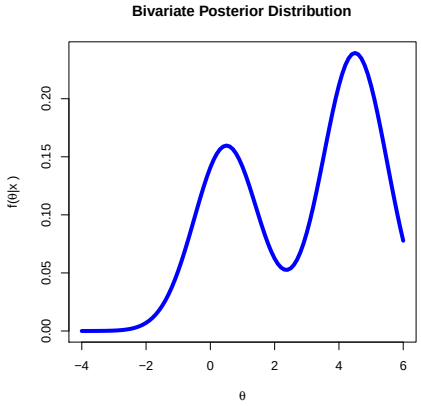


**Figure 3.** Marginal Densities.

## 3.1. Shape

In general, plots illustrate the shape of the posterior distribution. Important features of shape are modes (and antimodes), skewness and kurtosis (peakedness or heavy tails). The quantitative summaries of shape are needed to supplement like the view of mode (antimode).
The first task is to identify turning points of the density, i.e. solutions of $f'(\theta) = 0$. Such points include local maxima and minima of $f(\theta)$ which we call mode and antimode, respectively.

A point $\theta_0$ is characterized as a mode if $f'(\theta_0) = 0$ and $f''(\theta_0) < 0$ whereas it is an antimode if $f'(\theta_0) = 0$ and $f''(\theta_0) > 0$. Any point $\theta_0$ for which $f''(\theta_0) = 0$ is a point of inflection of the density (whether or not $f'(\theta_0) = 0$).

**Example 3.1.1.** Consider the gamma density

$$f(\theta) = \frac{a^b}{\Gamma(b)}\theta^{b-1}e^{-a\theta} \qquad ; \qquad \theta > 0 \tag{3.1}$$

where $a$, $b$ are positive constants.

$$f'(\theta) = \frac{a^b}{\Gamma(b)}\left\{(b-1) - a\theta\right\}\theta^{b-2}e^{-a\theta} \tag{3.2}$$

$$f''(\theta) = \frac{a^b}{\Gamma(b)}\left\{a^2\theta^2 - 2a(b-1)\theta - (b-1)(b-2)\right\}\theta^{b-3}e^{-a\theta} \tag{3.3}$$

So from $f'(\theta)$, the turning point at $\theta = (b-1)/a$. For $b \leq 1$, $f'(\theta) < 0$ for all $\theta \geq 0$, so $f(\theta)$ is monotonic decreasing and the mode is at $\theta = 0$. For $\theta > 1$, $f(\theta) \to 0$ as $\theta \to 0$, so the turning point $\theta = 0$ is not a mode. In this case, $f'(\theta) > 0$ for $\theta < (b-1)/a$ and $f'(\theta) < 0$ for $\theta > (b-1)/a$. Therefore $\theta = (b-1)/a$ is the mode. Looking at $f''(\theta)$, the quadratic expression has roots at $\theta = \frac{b-1}{a} \mp \frac{(b-1)^{1/2}}{a}$. Therefore $b > 1$, these are the points of inflection.

**Example 3.1.2.** Consider the mixture of two normal distributions

$$f(\theta) = \underbrace{\frac{0.8}{\sqrt{2\pi}}\exp\left\{-\frac{1}{2}\theta^2\right\}}_{\sim\mathcal{N} \quad with \quad weight} + \underbrace{\frac{0.2}{\sqrt{2\pi}}\exp\left\{-\frac{1}{2}(\theta-4)^2\right\}}_{\sim\mathcal{N} \quad with \quad weight} \tag{3.4}$$

$$f'(\theta) = -\frac{0.8\theta}{\sqrt{2\pi}}\exp\left\{-\frac{1}{2}\theta^2\right\} - \frac{0.2(\theta-4)}{\sqrt{2\pi}}\exp\left\{-\frac{1}{2}(\theta-4)^2\right\} \tag{3.5}$$

For $\theta \leq 0$, $f'(\theta) > 0$ and for $\theta \geq 0$, $f'(\theta) < 0$, the turning points at $\theta = 0.00034$, 2.46498 and 3.9945.

(3.6)

$$f''(\theta) = -\frac{0.8(\theta^2-1)}{\sqrt{2\pi}}\exp\left\{-\frac{1}{2}\theta^2\right\} - \frac{0.2((\theta^2-8\theta-15)}{\sqrt{2\pi}}\exp\left\{-\frac{1}{2}(\theta-4)^2\right\}$$

This is positive for $\theta \leq -1$, for $1 \leq \theta \leq 3$ and $\theta \geq 5$, confirming that the middle turning point is an antimode. Calculating $f''(\theta)$ at the other points confirms them to be modes. Finally points of inflection are at $\theta = -0.99998$, $\theta = 0.98254$, $\theta = 3.17903$, $\theta = 4.99971$.

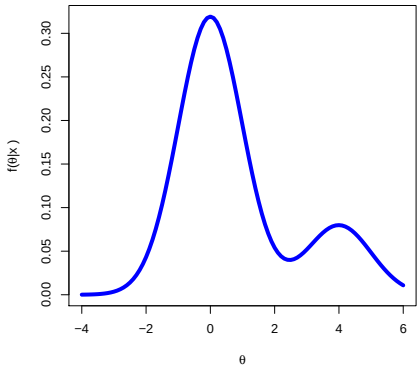


**Figure 4.** Plot of mixture of two normal distributions

## 3.2. Visualizing multivariate densities

**Turning points:**
At a mode, or at a turning pons generally, the gradients of function in all directions are zero. Therefore, the turning points are solutions of the simultaneous equations:

$$0 = \frac{\partial f(\theta,\phi)}{\partial \phi} = \frac{\partial f(\theta,\phi)}{\partial \theta}. \tag{3.7}$$

The turning points may be classified by examining the symmetric matrix $F''(\theta,\phi)$ of the second order partial derivatives. For instance, for a bivariate density $f(\theta,\phi)$,

$$F''(\theta,\phi) = \begin{bmatrix} \frac{\partial^2 f(\theta,\phi)}{\partial \theta^2} & \frac{\partial^2 f(\theta,\phi)}{\partial \theta \partial \phi} \\ \frac{\partial^2 f(\theta,\phi)}{\partial \phi \partial \theta} & \frac{\partial^2 f(\theta,\phi)}{\partial \phi^2} \end{bmatrix} \tag{3.8}$$

$F''(\theta,\phi)$ is known as the Hessian matrix. The second derivative of $f(\theta)$ ,m a direction $t$ is after differentiating. At a mode, this must be negative in all directions. So that the Hessian matrix is negative definite. Similarly, at an antimode, it is positive definite. In the intermediate case, where $F''(\theta,\phi)$ is indefinite, i.e. has both positive and negative eigenvalues, we have a saddle point.
If $F''(\theta,\phi)$ is positive definite, regions are:
$F''(\theta,\phi)$ are negative definite, regions are:
$F''(\theta,\phi)$ is indefinite, indefinite curvature. On the boundaries between these regions, one eigenvalues of $F''(\theta,\phi)$ is zero, so all points on such boundaries are inflection points.
A point of inflection corresponds to the second derivative being zero

in some direction $t$, therefore inflection points are characterized by $F''(\theta, \phi)$ being singular. In fact, in more than two dimensions, we can further subdivide the regions of indefinite curvature according to how many positive eigenvalues of $F''(\theta, \phi)$ has, and all these subregions are also separated by inflection boundaries.

**Location:**

A plot gives a good idea of location, but the conventional location measures for distributions are also useful. These include the mean, mode and median.

**Dispersion:**

The usual dispersion measure is the variance, or for a multivariate distribution the variance-covariance matrix.

**Dependence:**

It is important with multivariate distributions to summarize the dependence between individual parameters. This can be done with correlation coefficients, but plots of regression functions (conditional mean functions) can be more informative.

## 3.3. Informal Inferences

(a) Point estimation: The obvious posterior estimate of $\theta$ is its posterior mean $\hat{\theta} = E(\theta \mid x)$. Modes and medians are also natural point estimates, and they all have intuitively different interpretations. The mean is the expected value, the median is the central value and the mode is the most probable value.

(b) Interval estimation: If asked to provide an interval in which $\theta$ probably lies, we can readily derive such a thing from its posterior distribution. For instance, in the density shown on page 1, there is a probability 0.05 to the left of $\theta = 3.28$ and also 0.05 to the right of $\theta = 11.84$. So the interval $(3.28, 11.84)$ is a 90% posterior probability for $\theta$. We call such an interval a credible interval.

- If a frequentist had found this interval, it means that it would say that if we repeatedly draw samples of data from the same population, and applied the rule that was used to derive this particular interval to each off those datasets, then 90% of those intervals would contain $\theta$.

If a Bayesian approach, there is a posterior probability 0.9 that $\theta$ lies between 3.28 between 11.84.

**Definition 3.9.** A $100(1-\alpha)\%$ credible set for $\theta$ is a subset $C$ such that:

$$1-\alpha \leq P(C \mid y) = \int_C p(\theta \mid y)\mathrm{d}\theta. \tag{3.10}$$

where integration is replaced by summation for discrete components.

**Definition 3.11.** The exact possible coverage of $(1-\alpha)$ can be found by the highest posterior density of HPD credible set as the set:

$$C = \{\theta \in \Theta : p(\theta \mid y) \geq k(\alpha)\}. \tag{3.12}$$

where $k(\alpha)$ is the largest constant satisfying $P(C \mid y) \geq 1-\alpha$.

For 2-sided credible set, we can generally take the $\alpha/2$ and $(1-\alpha/2)$ quantiles of $p(\theta \mid y)$ as our $100(1-\alpha)\%$ credible set for $\theta$. This equal tail credible set will be equal to the HPD credible set if the posterior is symmetric and unimodal, but will be a bit wider otherwise.

(c) Evaluating hypothesis: Suppose we wish to test a hypothesis $H$ which asserts that $\theta$ lies in some region $A$. The Bayesian way to test to the hypothesis is simply to calculate the (posterior) probability that it is true: $P(\theta \in A \mid x)$.

## 3.4. Multivariate inference

All the above treated $\theta$ as a scalar parameter. If we have a vector $\theta$, then in general we can consider inference of the above forms about any scalar function $\phi = g(\theta)$. The inferences are then derived simply from the marginal posterior distribution of $\phi$.

# LECTURE 4
# Formal Inference

Suppose that we want to answer a question that falls neatly into the frequentist point estimation framework, "What is the best estimate of $\theta$?". In the frequentist theory, we need to be explicit about what we would regard as good properties for an estimator in order to identify a best one.
The Bayesian approach also needs to know what makes a good estimate before an answer can be given. This falls into the framework of formal inference.
Formally, we would seek the estimate that minimises expected square error. So the expectation is derived from the posterior distribution. We want to minimise:

$$\begin{aligned} E((\hat{\theta}-\theta)^2 \mid x) &= \hat{\theta}^2 - 2\hat{\theta}E(\theta \mid x) + E(\theta^2 \mid x) \\ &= (\hat{\theta} - E(\theta \mid x))^2 + v(\theta \mid x). \end{aligned} \tag{4.1}$$

So the estimate $\hat{\theta}$ that minimises this expected squared error is $\hat{\theta} = E(\theta \mid x)$.

## 4.1. Utility and decisions

Formal inference aims to obtain optimal answer to inference questions. This is done with reference to a measure of how good or bad the various possible inferences would be deemed to be if we knew the true value of $\theta$. This measure is a utility function. Formally, $u(d, \theta)$ defines the value of inference $d$ if the true value of the parameter is $\theta$. Formal inference casts an inference problem as a decision problem. A decision problem is characterized by:

- a set $\Omega$ of possible parameter values
- a probability distribution for $\theta \in \Omega$
- a set $D$ of possible decisions

- utility function $u(d, \theta)$ for $d \in D$ and $\theta \in \Omega$.

The solution is

$$\boxed{d_{opt} = arg \quad max_d E_\theta(u(d, \theta))} \tag{4.2}$$

Here the distribution of $\theta$ is its prior distribution.

- In inference problems, we generally define a measure of badness of an inference, which we call a loss function $L(d, \theta)$. We can simply define utility to be negative loss, and then the optimal inference is the one which minimises posterior expected loss.

### 4.1.1. Formal Point Estimation

The set $D$ is the set of all possible values of $\theta$. We have seen that if we use squared error loss (which is implicitly the measure used by frequentist in considering mean-squared-error of the variance of an unbiased estimator), formally defining $L(\hat{\theta}, \theta) = (\hat{\theta} - \theta)^2$, then the posterior mean is the optimal estimator. If we use absolute error loss $L(\hat{\theta}, \theta) = \left|\hat{\theta} - \theta\right|$, then the optimal estimator is the posterior median.

### 4.1.2. Formal Interval Estimation

The possible inferences now are interval or, more generally, subsets of possible values of $\theta$. A loss function will penalize an interval if it fails to contain the true value of $\theta$. So the optimal interval is found from this form:

$$\boxed{d_{opt} = \{\theta : f(\theta \mid) \geq t\}} \tag{4.3}$$

where $t$ is chosen to obtain the desired $P(\theta \in d_{opt} \mid x)$. Such a credible interval is called a highest (posterior) density interval (HPD interval or HDI).

**Example 4.1.1.** From the following figure 1 which shows 90% credible interval, it is not HPI. Because here the upper and lower limits should have the same posterior density, and shows that the density at $\theta = 3.28$ is higher that at $\theta = 11.84$. So the 90% highest density interval is actually $(2.78, 11.06)$.

## 4.2. Formal Hypothesis Testing

If we really need to decide whether (to act as if) hypothesis that $\theta \in A$, is true, there are just two inferences. $d_0$ is to say it is true, while $d_1$

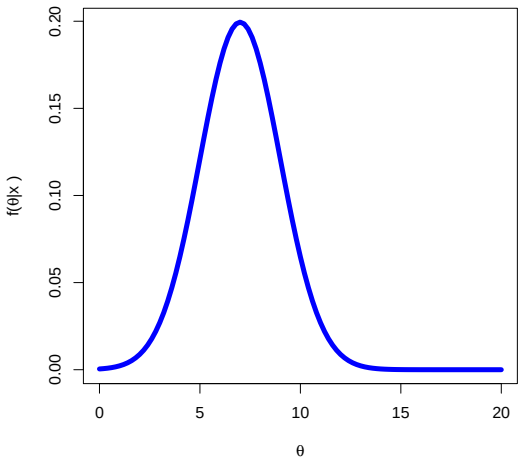


**Figure 1.** A posterior density plot

says it is false. The loss function will take the form:

$$L(d_0, \theta) = 0 \quad if \quad \theta \in A \tag{4.4}$$
$$= 1 \quad if \quad \theta \notin A$$

$$L(d_1, \theta) = k \quad if \quad \theta \in A \tag{4.5}$$
$$= 0 \quad if \quad \theta \notin A$$

where $k$ defines the relative seriousness of the first kind of error relative to the second. Then $E_\theta(L(d_0, \theta)) = P(\theta \notin A \mid x)$, while $E_\theta(L(d_1, \theta)) = kP(\theta \in A \mid x)$. The optimal decision is to select $d_0$ (say that $H$ is true) if its probability $P(\theta \in A \mid x)$ exceeds $\frac{1}{k+1}$. The greater the relative seriousness $k$ of the first kind of error, the more willing we are:

## 4.3. Nuisance Parameter

In any inference problem, the parameter(s) that we wish to make inference about is (are) called the parameter of interest, and the remainder of components of $\theta$ is (are) called nuisance parameters.

**Example 4.3.1.** If we have a sample from $\mathcal{N}(\mu, \sigma^2)$ and we wish to make inference about $\mu$, then nuisance parameter is $\sigma^2$.
If $\theta = (\phi, \psi)$ and $\psi$ is the vector of nuisance parameters, then the inference about $\phi$ is made from marginal value posterior distribution.

## 4.4. Transformation

If $\hat{\theta}$ is an estimate of $\theta$, is $g(\hat{\theta})$ the appropriate estimate of $\phi$?
This depends on the kind of inference being mode. In the particular case of point estimation, then the posterior mean is not invariant in this way.

**Example 4.4.1.** If $\phi = \theta^2$), then

$$E(\phi \mid x) = E(\theta^2 \mid x) = v(\theta \mid x) + E(\theta \mid x)^2 \geq E(\theta \mid x)^2 \tag{4.6}$$

The mode is not invariant to transformations but the median is invariant, at least to $1-1$ transformations.

Interval estimates are also invariant to $1-1$ transformations in the sense that if $[a, b]$ is a 90% interval, say for $\theta$, then $[g(a), g(b)]$ is a 90% interval for $\phi$ if g is a monotone increasing function. If $[a, b]$ is a 90% HPD interval for $\theta$, then $[g(a), g(b)]$ is an HPD interval for $\phi$?

## 4.5. The prior distribution

The nature of probability:

(a) Frequency probability: Frequentist statistics uses the familiar idea that the probability of an event is the limiting relative frequency with which that event would occur in an infinite sequence of repetitions. For this definition of probability to apply, it is necessary for the event to be, at least in principle, repeatable.
(b) Personal probability: Bayesian statistics is based on defining the probability of a proposition to be a measure of a person's degree of belief is the truth of that proposition.

In the Bayesian framework, wherever there is uncertainty there is probability.

In particular, parameters have probability distributions.

## 4.6. Subjectivity

The main critic to Bayesian methods is the subjectivity due to the prior density.

If the data are sufficiently strong, the remaining element of personal judgement will not matter, because all priors based on reasonable interpretation of the prior information will lead to effectively the same posterior inferences. Then we can claim robust conclusion on the basis of the synthesis of prior information and data.

If the data are not that strong, then we do not yet have enough scientific evidence to reach an objective conclusion. Any method which

claims to produce a definitive answer in such a situation is misleading, so this is actually a strength of the Bayesian approach.

## 4.7. Noninformative Priors

The basis of this is that if we have a completely flat prior distribution such that $f(\theta)$ is a constant, then the posterior density if proportional to the likelihood and inferences will be based only on the data. If we can do this, we can get the other benefits of Bayesian analysis, such as having more meaningful inferences that actually answer the question, but without the supposed disadvantage of subjectivity.

The main problem with this neat solution is that it can not be applied consistently.

**Example 4.7.1.** $f(\theta) = 1$ for all $\theta \in [0, 1]$ uniform distribution represents complete ignorance about $\theta$.

If $\theta$ is ignored, then $\phi = \theta^2$, which also takes values in $[0, 1]$, is also completely ignored.

But the implied distribution $f(\phi) = 1$ is not consistent with the previous specification of $f(\theta) = 1$. The uniform prior distribution for $\theta$ implies that $\phi = \theta^2$ should have the density $f(\phi) = \frac{1}{2\sqrt{\phi}}$. Conversely, if $\phi$ has a uniform prior distribution, then the implied prior for $\theta$ has density $f(\theta) = 2\theta$.

In general, a uniform prior for $\theta$ translates into a non-uniform prior for any function of $\theta$. Another complication is that if the range of possible values of $\theta$

Another complication is that if the range of possible values of $\theta$ is bounded then we can not properly give it a uniform distribution. For instance, if $\theta \in [0, \infty)$ and we try to define a prior distribution $f(\theta) = c$ for some constant $c$, then there is no value of $c$ that will make this density integrate to 1. For $c = 0$, it integrates to 0, and for any positive $c$ it integrates to infinity. In these situations, we appeal to proportionality and simply write $f(\theta) \propto 1$

A distribution expressed as $f(\theta) \propto h(\theta)$ when there can not be any proportionality constant that would make this into a proper density function, is called an improper distribution. This arises whenever the integral of $h(\theta)$ over the range of possible values of $\theta$ does not change.

**Example 4.7.2.** For a parameter $\theta \in [0, 1]$, three favourite recommendations are $f(\theta) = 1$, $f(\theta) = \pi^{-1}\theta^{-1/2}(1-\theta)^{-1/2}$ and $f(\theta) \propto \theta^{-1}(1-\theta)^{-1}$, the last of these being improper. We can identify these as the $Be(1, 1)$, $Be(1/2, 1/2)$ and $Be(0, 0)$ distributions.

Improper distributions are not in fact usually much of a problem, since we can appeal to proportionality. That is, the absence of a well-defined proportionality constant is ignored and assumed to cancel in the proportionality constant of Bayes' theorem [**2**]. In effect, we are obtaining the limit of the posterior distribution as we go through a range of increasingly flat priors towards the uniform limit.

(a) The posterior distribution may also be improper. In this case, technically, the limit of the above process is not well-defined. Improper prior distribution should never be used when the resulting posterior distribution is improper, so it is important to verify propriety of the posterior.
(b) When comparing different models for the data, improper distributions always lead to undefined model comparisons. This is an area outside the scope of this course, but very important in practice.

## 4.8. Informative Priors

So in specifying an informative prior distribution, (a) we specify values for whatever summaries best express the features of the prior information, then (b) we simply choose any conventional $f(\theta)$ that has those summaries.

**Example 4.8.1.** I wish to formulate my prior beliefs about number $N$ of students who will turn up to one of my lectures. I first ask myself what my best estimate would be, and I decide on 38, so I set $E(N) = 38$. I next ask myself how far wrong this estimate might be. I decide that the actual number could be as high as 48 or as low as 30, but I think the probability of the actual number being outside that range is small, maybe only 10%. Now a convenient prior distribution that matches these summaries is the Poisson distribution.

So it has mean 38 and $P(30 \leq N \leq 48) = 0.87$, which seems a good enough fit to my specified summaries.

## 4.9. Prior Choices

There are several alternatives to overcome the problem of improper prior.

(a) Jeffrey's Prior:
Let $I(\theta)$ be the Fisher information:

$$I(\theta) = -E\left\{\frac{\partial^2 \log f(x \mid \theta)}{\partial \theta^2}\right\}. \tag{4.7}$$

In the case of a vector parameter, $I(\theta)$ is the matrix formed as minus the expectation of the matrix of second order partial derivatives of $\log f(x \mid \theta)$. The Jeffrey's prior distribution it then:

$$\boxed{f_o(\theta) \propto |I(\theta)|^{1/2}} \tag{4.8}$$

**Example 4.9.1.** If $x_1, x_2, ..., x_n$ are normally distributed with mean $\theta$ and known variance $v$, then

$$f(x \mid \theta) \propto \exp\left\{-\frac{n(\bar{x}-\theta)^2}{2v}\right\} \tag{4.9}$$

What is the Jeffrey's prior for this distribution?
**Solution:**

$$\log f(x \mid \theta) = -\frac{n(\bar{x}-\theta)^2}{2v} \tag{4.10}$$

$$\frac{\mathrm{d}^2}{\mathrm{d}\theta^2} \log f(x \mid \theta) = -\frac{n}{2v} \tag{4.11}$$

Therefore,

$$I(\theta) = -E(\frac{\mathrm{d}^2}{\mathrm{d}\theta^2} \log f(x \mid \theta)) = \frac{n}{2v} \tag{4.12}$$

As a result,

$$f_0(\theta) = \sqrt{I(\theta)} = \sqrt{\frac{n}{2v}}. \tag{4.13}$$

**Example 4.9.2.** If $x_1, x_2, ..., x_n$ are distributed as $\mathcal{N}(\mu, \sigma^2)$ with $\theta = (\mu, \sigma^2)$, then

$$f(x \mid \theta) \propto \sigma^{-n} \exp\left\{-\frac{-n(s + (\bar{x}-\theta)^2)}{2\sigma^2}\right\} \tag{4.14}$$

where $s = \frac{\sum(x_i - \bar{x})^2}{n}$, Then what is the Jeffrey's prior of $f(\mu, \sigma^2)$?
**Solution:**

**Example 4.9.3.** If $x_1, x_2, ..., x_n$ are normally distributed with known mean $m$ and variance $\theta$, then

$$f(x \mid \theta) \propto \theta^{-n/2} \exp\left\{-\frac{-s}{(2\theta)^2}\right\} \tag{4.15}$$

where $s = \sum (x_i - m)^2$, then what is the Jeffrey's prior for $\theta$?
**Solution:**

$$\log f(x \mid \theta) = \tag{4.16}$$

A number of objections can made to the Jeffrey's prior, the most important of which is that it depends on the form of the data. The prior distribution should only represent the prior information, and not be influenced by what data are to be collected.

(b) Maximum Entropy:
The entropy $H(f) = -\int_{-\infty}^{\infty} f(\theta) \log f(\theta) \mathrm{d}\theta$ of the density $f(\theta)$ can be thought of a measure of how uninformative $f(\theta)$ is about $\theta$. For if we try to convert our information about $\theta$ as a general form of inference in the scoring rule framework, $H(f)$ is the lowest obtainable expected loss. If $H(f)$ is high, then the best decision is still poor. Now to represent prior ignorance we could use the prior density $f(\theta)$ which maximizes the entropy.

**Example 4.9.4.** Suppose that $\theta$ is discrete with possible values $\theta_1, \theta_2, ..., \theta_k$. The prior distribution with maximum entropy will then maximize $\sum_{i=1}^{k} p_i \log p_i + \lambda \sum p_i$, where $\lambda$ is a Lagrange multiplier, $\partial F / \partial p_i = \log p_i + 1 + \lambda$. Equating this to zero yields the solution $p_i = k^{-1}$, $i = 1, 2, 3, ..., k$. That is the maximum entropy prior is the uniform distribution.

The primary criticism of this approach is that it is not invariant under change of parametrization, the problem which the Jeffrey's prior was designed to avoid. In general, unrestricted maximization of entropy leads to a uniform prior distribution, which was shown to be sensitive to parametrization.

(c) Reference Prior: The expected amount of information provided by observing $x$ is given by

$$H\{f(\theta)\} - E\left[H\{f(\theta \mid x)\}\right], \tag{4.17}$$

where the expectation is over the preposterior distribution of $f(x)$ of $x$. If the experiment yielding $x$ were to be repeated, giving a new observation independent of $x$ given $\theta$ and with the same distribution, the posterior distribution would be expected to show a further reduction in entropy, representing the expected information in the second observation. If this were

repeated indefinitely we would eventually learn $\theta$ exactly and so remove all the entropy in the original prior distribution.

In the case of discrete $\theta$ taking a finite of possible values, this process reduces ??? maximizing prior entropy, and so gives the uniform distribution. This is not the case for continuous $\theta$. It is shown that under appropriate regularity conditions, the reference prior distribution is the Jeffrey prior.

# LECTURE 5
# Structuring Prior Information

Independence: Suppose that $x_1, x_2, ..., x_n$ are a sample from the $\mathcal{N}(\mu, 1)$ distribution, which we write formally as:

$$x_i \mid \mu \sim \mathcal{N}(\mu, 1) \tag{5.1}$$

, independent. From a Bayesian perspective, what is meant here is conditional dependence. That is the $x_i$'s are independent given $\mu$.

Exchangaability: It is same as the independence in frequentist approach.

**Definition 5.2.** Random variables $x_1, x_2, ..., x_m$ are said to be exchangeable if their joint distribution is unaffected by permitting the order of the $x_i$'s.

So first consider the (marginal) distribution of $x_1$. The definition says that every one of the $x_i$'s must have the same marginal distribution, because we can permute them so that any desired $x_i$ comes into the first position in the sequence. So one implication of exchangeability is that the random variables in question are identically distributed.

next consider the joint distribution of $x_1$ and $x_2$ Exchangeability means that every pair $(x_i, x_j)$ (for $i \neq j$) has the same bivariate distribution as $(x_1, x_2)$. In particular, the correlation between any pair of random variables is the same.

And so it goes to higher order joint distributions. The joint distributions $(x_1, x_2, ...., x_k)$ is the same as that of any other collection of $k$ distinct $x_i$'s. This is the meaning of exchageability.

- In general, suppose that $x_i$'s have a common distribution $g(x \mid \theta)$, and are independent, given $\theta$. Then the joint density is:

$$\begin{aligned} f(x_1, x_2, ..., x_m) &= \int f(x_1, x_2, ..., x_m, \theta) \mathrm{d}\theta \\ &= \int f(x_1, x_2, ..., x_m) f(\theta) \mathrm{d}\theta \\ &= \int \prod_{i=1}^{m} g(x_i \mid \theta) f(\theta) \mathrm{d}\theta. \end{aligned} \tag{5.3}$$

  which is unaffected by permuting the $x_i$'s. Their common marginal density is:

$$f(x) = \int g(x \mid \theta) f(\theta) \mathrm{d}\theta, \tag{5.4}$$

  and the common distribution of any pair of the $x_i$'s is

$$f(x, y) = \int g(x \mid \theta) g(y \mid \theta) f(\theta) \mathrm{d}\theta, \tag{5.5}$$

This is generally what a frequentist means by the $x_i$'s being iid, or being a random sample from the distribution $g(x \mid \theta)$. From the Bayesian perspective all frequentist statements are conditional on the parameters. So exchangeability is the Bayesian concept that corresponds very precisely with the frequentist idea of a random sample.

## 5.1. Binary Exchangeability

**Theorem 5.6** (De Finetti, 1937)**.** *Let $x_1, x_2, x_3, ...$ be an infinite sequence of exchangeable binary random variables. Then their joint distribution is characterised by a distribution $f(\theta)$ for a parameter $\theta \in [0, 1]$ such that the $x_i$'s are independent of $\theta$ with $P(x_i = 1 \mid \theta) = 0$.*

De Finetti's theorem says that if we have a sequence of binary variables, so that each $x_i$ takes the value 1 (success) or 0 (failure), then we can represent them as independent Bernouilli trials with probability $\theta$ of success in each trial.

## 5.2. Exchangeable Parameters

Consider the simple one-way analysis of variance model:

$$y_{ij} \sim \begin{cases} \mathcal{N}(\mu_i, \sigma^2), & i = 1, 2, ..., k \\ & j = 1, 2, ..., n_i \quad , \textit{independent} \end{cases} \tag{5.7}$$

where $\theta = (\mu_1, \mu_2, ..., \mu_k, \sigma^2)$. This says that we have $k$ independent normal samples, where the $i$-th sample has mean $\mu_i$ and size $n_i$, and where all the observations have a common variance $\sigma^2$.

For this model, we need to specify a joint prior distribution for $\mu_1, \mu_2, ..., \mu_k$ and $\sigma^2$. Now when it comes to formulating a joint prior distribution for many parameters, it is much easier if we can regard them as independent. Then we could write:

$$f(\mu_1, \mu_2, ..., \mu_k, \sigma^2) = f(\sigma^2) \prod_{i=1}^{k} f(\mu_i), \tag{5.8}$$

and we would only need to specify the prior distribution of each parameter separately. Unfortunately, this is unlikely to be the case in practice with this model.

The model is generally used when the samples are from a related or similar populations. For instance, they might be weight gains of pigs given $k$ different diets, and $\mu_i$ is the mean weight gain with diet $i$. In this sort of situation, the $\mu_i$'s would not be independent. Then we could add to the "prior model"

$$\mu_i \mid \xi, \tau^2 \sim \mathcal{N}(\xi, \tau^2) \tag{5.9}$$

, independent, which says that the $\mu_i$'s are drawn form a normal (assumed) population with unknown mean $\xi$ and unknown variance $\tau^2$.

The prior distribution could then be completed by specifying a joint prior distribution $f(\xi, \sigma^2, \tau^2)$ for the remaining parameters.

## 5.3. Hierarchical Models

The kind of modelling seen in the previous section is called "hierarchical". In general, we can consider a model of the form:

- Data model: $f(x \mid \theta)$
- First level of prior: $f(\theta \mid \phi)$
- Second level of prior: $f(\phi)$

We often refer to $\phi$ as the hyperparameter(s)

- If we are only interested in the parameter $\theta$ of the original data model,

$$f(\theta) = \int f(\theta \mid \phi) f(\phi) \mathrm{d}\phi. \tag{5.10}$$

- Let us be actually interested in the hyperparameters $\phi$,

$$f(x \mid \phi) = \tag{5.11}$$

$$f(\theta, \phi) = \tag{5.12}$$

$$f(\theta, \phi \mid x) \propto \tag{5.13}$$

$$f(\theta \mid x) = \tag{5.14}$$

Shrinkage: It means that the posterior distribution and posterior estimates of these parameters will generally be closer together than their corresponding data estimates. This is a phenomenon known as "shrinkage". Let $w = (\xi, \sigma^2, \tau^2)$ and $\mu = (\mu_1, \mu_2, ..., \mu_k)$ and

$$\begin{aligned} f(\mu \mid x, w) &\propto f(x \mid \mu, \sigma^2) f(\mu \mid w) \\ &= \prod_{i=1}^{k} \left\{ \prod_{j=1}^{n_i} \frac{1}{\sqrt{2\pi}\sigma} \exp\left[ -\frac{(y_{ij} - \mu_i)^2}{2\sigma^2} \right] \right\} \exp\left[ -\frac{(\mu_i - \xi)^2}{2\tau^2} \right] \\ &= \prod_{i=1}^{k} f(\mu_i \mid x, w) \end{aligned} \tag{5.15}$$

where each of the $f(\mu_i \mid x, w)$ comes from the analysis of a single normal sample given in Lecture 1. That is conditional on $w$, the $\mu_i$'s are independent $\mathcal{N}(m_i^*, v_i^*)$, where:

$$v_i^* = (n_i \sigma^{-2} + \tau^{-2})^{-1} \tag{5.16}$$

$$m_i^* = (n_i \sigma^{-2} + \tau^{-2})^{-1} (n_i \sigma^{-2} \bar{y}_i + \tau^{-2} \xi) \tag{5.17}$$

We can already see the shrinkage in the model, because the posterior mean of each $\mu_i$ is a weighted average of its own data estimate $\bar{y}_i$ and the common value $\xi$. So they are shrunk towards this common value.

**Example 5.3.1.** Let's consider a simple regression situation in which we have observations $y_1, y_2, ..., y_n$ at values $x_1 < x_2 < .. < x_n$ of the explanatory variables. The usual linear regression model specifies

$$y_i \mid \alpha, \beta, \sigma^2 \sim \mathcal{N}(\alpha + \beta x_i, \sigma^2) \tag{5.18}$$

Instead we can create the hierarchical model:

- Data model:
- First level of prior:
- Second level of prior:

# LECTURE 6
# Sufficiency and Ancillary

What happens if $f(x \mid \theta)$ does not depend on $\theta$? If $f(x \mid \theta)$ does not depend on $\theta$, the data are completely uninformative.

**Definition 6.1.** $t(x)$ is sufficient if, for any given $\theta$, $f(x \mid \theta)$ is a function only of $t(x)$, apart froma multiplicative factor that can be any function of $x$

It means that we only need to know $t(x)$ in order to obtain the posterior distribution. It is sufficient to know $t(x)$. Therefore the posterior distribution is the same as if we only had observed $t(x)$ rather than the whole of $x$.

Suppose that $s(x)$ represents all other information in $x$, $x = (t(x), s(x))$. Then

$$f(x \mid \theta) = f(t(x) \mid \theta) f(s(x) \mid t(x), \theta). \tag{6.2}$$

$t(x)$ is sufficient $f(s(x) \mid t(x), \theta)$ must not depend on $\theta$. So once we know $t(x)$ there is no information in $s \mid x$.

**Definition 6.3.** $s(x)$ is ancillary if $f(s(x) \mid \theta)$ does not depend on $\theta$.

**Example 6.0.1.** Let $x_i \mid\sim \mathcal{N}(\mu, \sigma^2)$ with $\sigma^2$ known. Then $x_i - x_j \sim \mathcal{N}(0, 2\sigma^2)$ is ancillary for any $i \neq j$.

- Let $t(x)$ be the rest of the information in $x$, so that again we have $x = (t(x), s(x))$. Then

$$f(x \mid \theta) = f(s(x) \mid \theta) f(t(x) \mid s(x), \theta), \tag{6.4}$$

  but now this implies that $f(x \mid \theta) \propto f(t(x) \mid s(x), \theta)$

## 6.1. The Likelihood Principle

The likelihood principle asserts that inference should be based only on the likelihood.

**Example 6.1.1.** Let $x \sim Be(1, \theta)$ and the results of a fixed number $n$ of independent Bernoullie trials.

$$f_x(x \mid \theta) = \theta^r (1-\theta)^{n-r} \tag{6.5}$$

where $r$ is the number of success. If $x \sim Bi(n, \theta)$, then $f_R(r \mid \theta) = \binom{n}{r}\theta^r(1-\theta)^{n-r} \propto \theta^r(1-\theta)^{n-r}$.

If we keep observing independent Bernoulli trials until we get a fixed number $r$ of success.

$$f_y(y \mid \theta) = \theta^r (1-\theta)^{n-r} \tag{6.6}$$

where $r$-th success on the $n$-th trial. If the distribution is negative binomial

$$f_N(n) = \binom{n-1}{r-1}\theta^r(1-\theta)^{n-r} \propto \theta^r(1-\theta)^{n-r} \tag{6.7}$$

## 6.2. Identifiability

Let $\theta = (y(\theta), h(\theta))$ and $f(x \mid \theta)$ depend only on $g(\theta)$.

$$\begin{aligned} f(\theta \mid x) &\propto f(x \mid \theta) f(\theta) = f(x \mid g(\theta)) f(\theta) \\ &= f(x \mid g(\theta)) f(g(\theta)) f(h(\theta) \mid g(\theta)) \\ &\propto f(g(\theta) \mid x) f(h(\theta) \mid g(\theta)). \end{aligned} \tag{6.8}$$

This says that the posterior distribution of $\theta$ is made up of the distribution of $g(\theta)$ and the <u>prior</u> distribution of $h(\theta)$ given $g(\theta)$. So it is the <u>conditional</u> posterior distribution of $h(\theta)$ given $g(\theta)$ that is the same as the prior. That is

$$f(h(\theta) \mid x, g(\theta)) = f(h(\theta) \mid g(\theta)). \tag{6.9}$$

We say that $h(\theta)$ is <u>not identifiable</u> from these data. No matter how much data we get, we can not learn exactly what $h(\theta)$ is. With sufficient data we can learn $g(\theta)$, but not $h(\theta)$.

## 6.3. Asymptotic Theory

Suppose that we have a sequence of iid observations $x_1, x_2, x_3, ...$ and suppose that $\mathbf{x}_n = (x_1, x_2, ..., x_n)$ comprises the first $n$ observations.

We can now consider a sequence of posterior distributions $f(\theta \mid x_1)$, $f(\theta \mid x_2)$, $f(\theta \mid x_3)$,...,$f(\theta \mid x_n)$,$f(\theta \mid x_{n+1})$. We wish to know how the posterior distribution $f(\theta \mid x_n)$ behaves as $n \to \infty$.

So as we get more data, we expect that the posterior will in some sense converge to the true value. Also, the weight the posterior gives to the data increases, and therefore we can expect that in the limit the posterior will be insensitive to the prior.

Subject to some regularity conditions. Regularity conditions:

(a) The whole $\theta$ needs to be identifiable.
(b) Prior possibility condition: the prior probability does not give zero probability to the true value of $\theta$.
(c) Continuity condition: we need a continuity condition for $\theta$.

## 6.4. Preposterior Properties

Let $X$ and $Y$ be any two random variables. Then

$$E(Y) = E\left\{E(Y \mid X)\right\}, \tag{6.10}$$

$$v(Y) = E\left\{v(Y \mid X)\right\} + v\left\{E(Y \mid X)\right\} \tag{6.11}$$

Let's replace $Y$ by the parameter vector $\theta$ and $X$ by the data vector $X$.

**Remark.** If we use the posterior mean $E(\theta \mid x)$ to estimate $\theta$, the its expected bias 0.

## 6.5. Conjugate Prior Forms

The conjugacy is a joint property of the prior and the likelihood function that provides a posterior from the same distributional family as the prior.

**Example 6.5.1.** Conjugacy in exponential specifications.

$$E(x \mid \theta) = \theta \exp\left\{-\theta x\right\} \tag{6.12}$$

where $0 \leq x$, $0 < \theta$. If $\theta \sim Gamma(\alpha, \beta)$, then

$$f(\theta \mid \alpha, \beta) = \frac{1}{\Gamma(\alpha)} \beta^{\alpha} \theta^{\alpha-1} \exp\left\{-\beta\theta\right\} \tag{6.13}$$

where $\theta, \alpha, \beta > 0$

Suppose we now observe $x_1, x_2, ..., x_n \sim iid$. The likelihood is

$$L(\theta \mid x) = \prod_{i=1}^{n} \theta e^{-\theta x_i} = \theta^n \exp\left\{-\theta \sum x_i\right\} \tag{6.14}$$

Thus,

$$\begin{aligned} \pi(\theta \mid x) &\propto E(x \mid \theta) L(\theta \mid x) \\ &= \theta^n \exp\left\{-\theta \sum x_i\right\} \frac{1}{\Gamma(\alpha)} \beta^{\alpha} \theta^{\alpha-1} \exp\left\{-\beta\theta\right\} \\ &\propto \theta^{\alpha+n-1} \exp\left\{-\theta(\sum x_i + \beta)\right\} \end{aligned} \tag{6.15}$$

This is the kernel of a $Gamma(\alpha + n, \sum x_i + \beta)$ and therefore the gamma distribution is shown to be conjugate to the exponential likelihood function.

**Table 1.** Conjugate Prior Distribution Table

| Likelihood Form | Conjugate Prior Distribution | Hyperparameters |
|---|---|---|
| Bernouilli | Beta | $\alpha > 0$, $\beta > 0$ |
| Binomial | Beta | $\alpha > 0$, $\beta > 0$ |
| Multinomial | Dirichlet | $\theta_j > 0$, $\sum \theta_j = \theta_0$ |
| Negative Binomial | Beta | $\alpha > 0$, $\beta > 0$ |
| Poisson | Gamma | $\alpha > 0$, $\beta > 0$ |
| Exponential | Gamma | $\alpha > 0$, $\beta > 0$ |
| Gamma (ind $\chi^2$) | Gamma | $\alpha > 0$, $\beta > 0$ |
| Normal for $\mu$ | Normal | $\mu \in \mathbb{R}$, $\sigma^2 > 0$ |
| Normal for $\sigma^2$ | Inverse Gamma | $\alpha > 0$, $\beta > 0$ |
| Pareto for $\alpha$ | Gamma | $\alpha > 0$, $\beta > 0$ |
| Pareto for $\beta$ | Pareto | $\alpha > 0$, $\beta > 0$ |
| Uniform | Pareto | $\alpha > 0$, $\beta > 0$ |

# LECTURE 7
# Tackling Real Problems

There are various computational tools that are widely used in practical Bayesian statistics, the most well-know one is Markov Chain Monte Carlo or MCMC. The basic idea is that we randomly draw a very large sample $Q^{(1)}, Q^{(2)}, ...$ from the posterior distribution. Given such a sample, we can compute any inference we wish. If we want to make inference about some derived parameter $\phi = g(\theta)$ then $\phi^{(1)} = g(Q^{(1)})$, $\phi^{(2)} = g(Q^{(2)})$,..., is a sample from its posterior distribution $f(\theta \mid x)$. The sample mean $\bar{\phi}$ is an estimate of $E(\phi \mid x)$. In principle, we could draw such a sample using simple Monte Carlo sampling. That is each $Q^{(j)}$ is independently drawn from $f(\theta \mid x)$ There are algorithms for efficiently drawing random samples from a wide variety of standard distributions.

## 7.1. What is a Markov Chain?

"A stochastic process" is a consecutive set of random quantities defined on some known state space $Q$, indexed so that the order is known. $\left\{Q^{[t]}, t \in T\right\}$. Here the state space (which is parameter space for us) is just the allowable range of values for the random vector of interest. The state space $Q$ is either discrete or continuous depending on how the variable of interest is measured.

A Markov chain is a stochastic process with the property that at time $t$ in the series, the probability of making a transition to any new state is dependent only on the current state of the process.

$$p(Q^{[t]} \in A \mid Q^{[0]}, Q^{[1]}, ..., Q^{[t-2]}, Q^{[t-1]}) \tag{7.1}$$

where $A$ is the identified set on the complete state space.

A fundamental concern is the transition process that defines the probabilities of moving to other points in the state space, given the current location of the chain. This structure is defined via the transition kernel $K$ as a general mechanism for describing the probabilities of moving to some other specified state based on the current chain status. When the state space is discrete, then $K$ is a matrix, $k \times k$ for $k$ discrete elements in $A$, where each cell defines the probability of a state transition from the first term in the parentheses to all possible states:

$$P_A = \begin{bmatrix} p(\theta_1, \theta_1) & ... & p(\theta_1, \theta_k) \\ \vdots & & \\ p(\theta_k, \theta_1) & ... & p(\theta_k, \theta_k) \end{bmatrix} \tag{7.2}$$

The row and columns indicate:

- 
- 

An important feature of the transition kernel is that the transition probabilities between two selected states for arbitrary numbers of steps $m$ can be calculate multiplicative:

(7.3)

$$p^m(\theta_j^{[m]} = y \mid \theta_i^{[0]} = x) = \underbrace{\sum_{\theta_1} \sum_{\theta_2} .. \sum_{\theta_{m-1}}}_{all \quad possible \quad paths} \underbrace{p(\theta_i, \theta_1) p(\theta_1, \theta_2) ... p(\theta_{m-2}, \theta_{m-1})}_{transition \quad products}$$

So $p^m(\theta_j^{[m]} = y \mid \theta_i^{[0]} = x)$ is also a stochastic transition matrix.

**Example 7.1.1.**

$$\underbrace{P}_{current \quad period} = \begin{matrix} \theta_1 \\ \theta_2 \end{matrix} \begin{bmatrix} 0.8 & 0.2) \\ 0.6 & 0.4 \end{bmatrix} \tag{7.4}$$

Let the starting point $S_0 = [0,5 \quad 0,5]$, to get the first state

$$S_1 = \tag{7.5}$$

To get the second state,

$$S_2 = \tag{7.6}$$

$$S_3^{=} \tag{7.7}$$

$$S_4 = \tag{7.8}$$

So the choice proportions are converging to $[0,75 \quad 0,25]$ since the transition matrix is pushing toward a steady state or stationary distribution of the proportions. So when we reach this distribution, all future states are constant, that is stationary.

## 7.2. The Chapman-Kolmogorov Equations

These equations specify how successive events are bound together probabilistically. If we abbreviate the hand side of expression 7.3

$$p^{m_1+m_2}(x,y) = \sum_{all \quad z} p^{m_1}(x,z)p^{m_2}(z,y) \quad discrete \quad case. \tag{7.9}$$

$$p^{m_1+m_2}(x,y) = \int_{range \quad z} p^{m_1}(x,z)p^{m_2}(z,y)\mathrm{d}z \quad continuouscase. \tag{7.10}$$

This is also equal to:

$$p^{m_1+m_2} = p^{m_1}p^{m_2} = p^{m_1}p^{m_2-1}p = p^{m_1}p^{m_2-2}p^2 = ... \tag{7.11}$$

For discrete case. Thus iterative probabilities can be decomposed into segmented products in any way we like, depending on the interim step.

## 7.3. Marginal Distributions

The marginal distributions at some step $m$-th from the transition kernel is found via

$$\underbrace{\pi^m(\theta)}_{\text{current value of the chain}} = \underbrace{[p^m(\theta_1), p^m(\theta_2), ..., p^m(\theta_k)]}_{\text{the row of the transition kernel for the } m\text{-th step}} \tag{7.12}$$

So the marginal distribution at the first step of the Markov chain is given by

$$\boxed{\pi^1(\theta) = \pi^0(\theta)p^1}. \tag{7.13}$$

where $\pi^0$ = initial starting value assigned to the chain, $p^1 = P$ = simple transition matrix, Thus

$$\pi^n = \tag{7.14}$$

$$\pi^m(\theta_j) = \tag{7.15}$$

## 7.4. General Properties of Markov Chains

(a) Homogeneity: A homogeneous Markov chain at step $n$-th transition probability that does not depend on the value of so the decision to move at this step is independent of this being the current point in time.

(b) Irreducibility: A Markov chain is irreducible if every point or collection of points can be reached from every other point or collection of points. So irreducibility implies the existence of a path between any two points in the subspace

(c) Recurrence:If a subspace is closed, finite, and irreducible, then all states within the supspace are recurrent An irreducible Markov chain is called recurrent with regard to a given state, A, which is a single point or a defined collection of point, if the probability that the chain occupies A infinitely often over unbounded time is non-zero.

(d) Stationarity: Let $\pi(\theta)$ =stationary distribution of the Markov chain for $\theta$ of the state space.

$p(\theta_i, \theta_j)$ = the probability that the chain will move from $\theta_j$ to $\theta_j$ at some arbitrary step $t$ from the transition kernel.

$\pi^t(\theta)$ = the marginal distribution, thus the stationary distribution is defined as:

$$\sum_{\theta_i} \pi^t(\theta_i) p(\theta_i, \theta_j) = \pi^{t+1}(\theta_j) \quad \text{discrete case} \tag{7.16}$$

$$\int \pi^t(\theta_i) p(\theta_i, \theta_j) \mathrm{d}\theta_i = \pi^{t+1}(\theta_j) \quad \text{continuous case} \tag{7.17}$$

That is $\pi = \pi p$. The marginal distribution remains fixed when the chain reaches the stationary distribution.

Once the chain reaches its stationary distribution (also its invariant distribution, equilibrium distribution or limiting distribution), it stays in this distribution and moves about or "mixes" throughout the subspace according to marginal distribution, $\pi(\theta)$, forever.

(e) Ergodicity: If the chain is irreducible, positive Harris recurrent (i.e. recurrence under unbounded continuous state space), and aperiodic, the we call it ergodic. Ergodic Markov chains have the property:

$$\lim_{n \to \infty} p^n(\theta_i, \theta_j) = \pi(\theta_j) \tag{7.18}$$

for all $\theta_i$, and $\theta_j$ in the subspace Therefore in the limit, the marginal distribution at one step is identical to the marginal distributions at all other steps. The ergodic theorem is analogous to the strong law of large numbers bot for Markov chains. Thus suppose that $\theta_{i+1}, ..., \theta_{i+n}$ are $n$ values from a Markov chain that has reached its ergodic distribution. A statistic of interest, $h(\theta)$, can be calculated empirically:

$$h(\theta_i) = \frac{1}{n} \sum_{j=i+1}^{i+n} h(\theta_j) \approx h(\theta) \tag{7.19}$$

For a given empirical estimator $\hat{h}(\theta_i)$ with bounded limiting variance, we get the central limit theorem results

$$\sqrt{n}\frac{\hat{h}(\theta_i) - h(\theta)}{\sqrt{v(\hat{h}(\theta_i))}} \to_{n\to\infty} \mathcal{N}(0,1). \tag{7.20}$$

## 7.5. Noniterative Monte Carlo Methods

(a) Direct Methods:
The most basic definition of Monte Carlo integration is that $\theta \sim h(\theta)$ and we seek $\gamma = E[f(\theta)] = \int f(\theta)h(\theta)\mathrm{d}\theta$. Then if $\theta_1, ..., \theta_N \sim^{iid} h(\theta)$, we have

$$\hat{\gamma} = \frac{1}{N}\sum_{j=1}^{N} f(\theta_j) \tag{7.21}$$

with converges to $E[f(\theta)]$ with probability 1 as $N \to \infty$ by the strong law of large numbers. In our case, $h(\theta)$ is a posterior distribution and $\gamma$ is the posterior mean of $f(\theta)$. Hence the computation of posterior expectations require only a sample of size $N$ from the posterior distribution.

(b) Indirect methods:
If we can't directly sample from distribution? In this case, we can use one of the following methods

- Importance sampling: Suppose we wish to approximate a posterior expectation, say

$$E(f(\theta) \mid y) = \frac{\int f(\theta)L(\theta)\pi(\theta)\mathrm{d}\theta}{\int L(\theta)\pi(\theta)\mathrm{d}\theta} \tag{7.22}$$

where $f$ is the function of interest, $L$ is the likelihood or the data $y$. By defining the weight function $w(\theta) = \frac{L(\theta)\pi(\theta)}{g(\theta)}$, we have

$$\begin{aligned} E(f(\theta) \mid y) &= \frac{\int f(\theta)w(\theta)g(\theta)\mathrm{d}\theta}{\int w(\theta)g(\theta)\mathrm{d}\theta} \\ &\approx \frac{\frac{1}{N}\sum_{j=1}^{N} f(\theta_j)w(\theta_j)}{\frac{1}{N}\sum_{j=1}^{N} w(\theta_j)} \end{aligned} \tag{7.23}$$

where $\theta_j \sim^{iid} g(\theta)$. Here $g(\theta)$ is called the importance function, how closely it resembles $L(\theta)\pi(\theta)$ controls how good the approximation in the equation is. If $g(\theta)$ is a good approximation, the weights will all be roughly equal,

which in turn will minimize the variance of the numerator and denominator.

If $g(\theta)$ is a poor approximation, many of the weights will be close to zero, and thus a few $\theta_j$'s will dominate the sums, producing an inaccurate approximation.

Thus in importance sampling, one chooses a known density function $g(\theta)$ that is easy to sample. The procedure works best if $g(\theta)$ is similar in shape to the known kernel of the posterior $L(\theta)\pi(\theta)$ with tails that do not decay more rapidly than the tails of the posterior.

- Refection sampling: Here instead of trying to approximate the normalized posterior:

$$h(\theta) = \frac{L(\theta)\pi(\theta)}{\int L(\theta)\pi(\theta)\mathrm{d}\theta}, \tag{7.24}$$

we try to "blanket" it. That is suppose there exists an identifiable constant $\mu > 0$ and a smooth density $g(\theta)$, called the envelope function, such that $L(\theta)\pi(\theta) < \mu g(\theta)$ for all $\theta$.

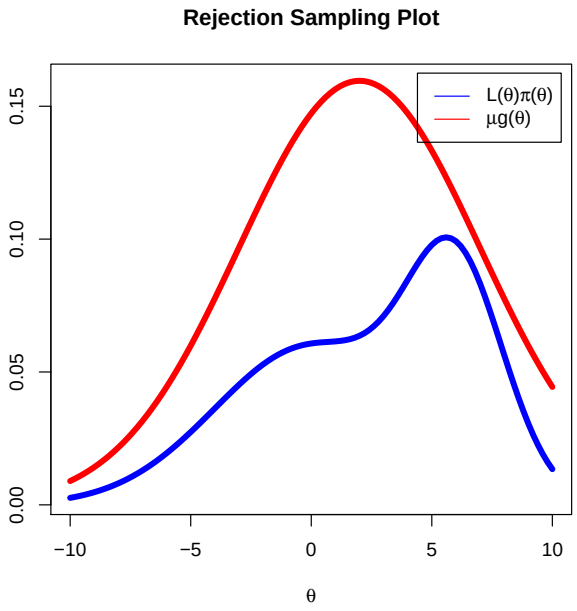


**Figure 1.** Plot of $L(\theta)\pi(\theta)$ and $\mu g(\theta)$

**Example 7.5.1.**

The rejection method proceeds as follows:

(a) Generate $\theta_j \sim g(\theta)$
(b) Generate $U \sim Uniform(0,1)$
(c) If $\mu u g(\theta_i) < L(\theta_j)\pi(\theta_j)$, accept $\theta_j$, otherwise reject $\theta_j$
(d) Return to step $(i)$ and repeat until the desired sample $\{\theta_j, j = 1, ..., N\}$ is obtained. The members of this sample will then be random samples from $h(\theta)$. Like an importance sampling density, the envelope density $g$ should be similar to the posterior in general appearance, but with heavier tails and sharper infinite peaks, in order to assure that there are sufficiently many rejection candidates available across its entire domain. Also $\mu g$ is actually an "envelope" for the unnormalized posterior $L\pi$.

- Weighted bootstrap: Suppose an $\mu$ appropriate for the rejection method is not readily available, but that we do have a sample $\theta_1, ..., \theta_N$ from same approximating density $g(\theta)$, Define:

$$w_i = \frac{L(\theta_i)\pi(\theta_i)}{g(\theta_i)} \tag{7.25}$$

and

$$q_i = \frac{w_i}{\sum_{j=1}^{N} w_j} \tag{7.26}$$

Now draw $\theta^*$ from the discrete distribution over $\{\theta_1, , ..., \theta_N\}$ which places mass at $\theta_i$. Then

$$\theta^* \sim h(\theta) = \frac{L(\theta)\pi(\theta)}{\int L(\theta)\pi(\theta)\mathrm{d}\theta} \tag{7.27}$$

with the approximation improving as $N \to \infty$. This is a weighted boottrap, since instead of resampling from the set $\{\theta_1, , ..., \theta_N\}$ with equally likely probabilities of selection , we are resampling some points more often that others due to the unequal weighting.

NOTE: Importance sampling, rejection sampling and the weighted bootstrap are all "one-off", or non-iterative methods: they draw a sample of size $N$, and stop. Hence there is no notion of the algorithm "converging " we simply require $N$ sufficiently large. But for many problems, especially high-dimensional ones, it may be difficult or impossible to find an importance sampling density (or envelope function)

which is an acceptably accurate approximation to the lay posterior, but still easy to sample from.

Solution: MCMC based methods

- Gibb's sampling
- Metropolis Algorithm
- Metropolis-Hastling Algorithms

# LECTURE 8

- There is a superficial resemblance between MCMC and the frequentist technique of bootstrapping.

  An MCMC sampler is a Markov chain in $\theta$ space. We start the chain with an arbitrary $\theta^{(1)}$, which is therefore not random. Each subsequent $\theta^{(j)}$ drawn from a distribution $q(\theta^{(j)} \mid \theta^{(j-1)})$, so that it depends only on the previous point $\theta^{(j-1)}$ and not on the history of the chain up to that point. This is the Markov property. The distribution $q(. \mid .)$ defines the transition probabilities of the chain.
- Let $\theta^{(2)}$ be the distribution given the initial $\theta^{(1)}$ so, just $q(\theta^{(2)} \mid \theta^{(1)})$. The distribution of $\theta^{(3)}$ is:

$$f(\theta^{(3)} \mid \theta^{(1)}) = \int q(\theta^{(3)} \mid \theta^{(2)}) q(\theta^{(2)} \mid \theta^{(1)}) \mathrm{d}\theta^{(2)}. \tag{8.1}$$

  The distribution of $\theta^{(j)}$ can then be obtained in principle by iterating this convolution using

$$f(\theta^{(j)} \mid \theta^{(1)}) = \int q(\theta^{(j)} \mid \theta^{(j-1)}) q(\theta^{(j-1)} \mid \theta^{(1)}) \mathrm{d}\theta^{(j-1)}. \tag{8.2}$$

  The Markov chain theory then says that, subject to some conditions, there is a unique limiting distribution $p(\theta)$ such that for all sufficiently large $j$ we have $f(\theta^{(j)} \mid \theta^{(1)}) \approx p(\theta^{(j)})$, and this limiting distribution is independent of the arbitrary starting value $\theta^{(1)}$.

## 8.1. The Gibbs Sampler

Procedure:

(a) Choose starting values: $\theta^{[0]} = \left[\theta_1^{[0]}, \theta_2^{[0]}, ..., \theta_k^{[0]}\right]$

(b) At the $j$-th starting at $j = 1$ complete the single cycle by drawing values from the $k$ distributions given by:

$$\begin{aligned}
\theta_1^{[j]} &\approx \pi(\theta_1 \mid \theta_2^{[j-1]}, \theta_3^{[j-1]}, \theta_4^{[j-1]}, ..., \theta_{k-1}^{[j-1]}, \theta_k^{[j-1]}) \\
\theta_2^{[j]} &\approx \pi(\theta_2 \mid \theta_1^{[j]}, \theta_3^{[j-1]}, \theta_4^{[j-1]}, ..., \theta_{k-1}^{[j-1]}, \theta_k^{[j-1]}) \\
\theta_3^{[j]} &\approx \pi(\theta_3 \mid \theta_1^{[j]}, \theta_2^{[j]}, \theta_4^{[j-1]}, ..., \theta_{k-1}^{[j-1]}, \theta_k^{[j-1]}) \\
&\vdots \\
&\vdots \\
\theta_{k-1}^{[j]} &\approx \pi(\theta_{k-1} \mid \theta_1^{[j]}, \theta_2^{[j]}, \theta_3^{[j]}, ..., \theta_{k-2}^{[j]}, \theta_k^{[j-1]}) \\
\theta_k^{[j]} &\approx \pi(\theta_k \mid \theta_1^{[j]}, \theta_2^{[j]}, \theta_3^{[j]}, ..., \theta_{k-2}^{[j]}, \theta_{k-1}^{[j]})
\end{aligned}$$

(c) Increment $j$ and repeat until convergence.

**Example 8.1.1.** $\theta_1 \mid \theta_2 \sim \mathcal{N}(\begin{bmatrix}0\\0\end{bmatrix}, \begin{bmatrix}1 & \rho\\ \rho & 1\end{bmatrix})$

$$\theta_1^{[j]} \mid \theta_2^{[j-1]} \sim \mathcal{N}(\rho\theta_2^{[j-1]}, 1-\rho^2) \tag{8.3}$$

$$\theta_2^{[j]} \mid \theta_1^{[j]} \sim \mathcal{N}(\rho\theta_1^{[j]}, 1-\rho^2) \tag{8.4}$$

**Example 8.1.2.** Let $x_1, x_2, ..., x_n$ be a series of count data where there exists the possibility of a changepoint at some period $k$, along the series. Therefore there are two Poisson data-generating processes:

$$x_i \mid \lambda \sim P(\lambda) \qquad i = 1, 2, ..., k \tag{8.5}$$

$$x_i \mid \phi \sim P(\phi) \qquad i = k+1, ..., n \tag{8.6}$$

The parameters to be estimated are $\lambda$, $\phi$ and $k$. Also the three independent priors applied to this model are:

$$\begin{aligned}
\lambda &\sim Gamma(\alpha, \beta), \\
\phi &\sim Gamma(\gamma, \delta), \\
k &\sim \text{Discrete uniform on } [1, 2, ..., n]\,.
\end{aligned}$$

So the joint posterior is

$$\begin{aligned}
\pi(\lambda, \phi, k \mid y) &\propto L(\lambda, \phi, k \mid y)\pi(\lambda \mid \alpha, \beta)\pi(\phi \mid \gamma, \delta)\pi(k) \\
&= (\prod_{i=1}^{k} \frac{e^{-\lambda}\lambda^{y_i}}{y_i!})(\prod_{i=k+1}^{n} \frac{e^{-\phi}\phi^{y_i}}{y_i!})(\frac{\beta^\alpha}{\Gamma(\alpha)}\lambda^{\alpha-1}e^{-\beta\lambda})(\frac{\delta^\gamma}{\Gamma(\gamma)}\phi^{\gamma-1}e^{-\delta\phi})\frac{1}{n} \qquad (8.7) \\
&\propto \lambda^{\alpha-1+\sum_{i=1}^{k} y_i}\phi^{\gamma-1+\sum_{i=k+1}^{n} y_i}e^{-(k+\beta)\lambda-(n-k+\delta)\phi}.
\end{aligned}$$

So

$$\lambda \mid \phi, k \sim Gamma(\alpha + \sum_{i=1}^{k} y_i, \beta + k),$$

$$\phi \mid \lambda, k \sim Gamma(\gamma + \sum_{i=k+1}^{n} y_i, \delta + n - k).$$

Let $\lambda$ and $\phi$ be fixed,

$$\begin{aligned} p(y \mid k, \lambda, \phi) &= (\prod_{i=1}^{k} \frac{e^{-\lambda}\lambda^{y_i}}{y_i!})(\prod_{i=k+1}^{n} \frac{e^{-\phi}\phi^{y_i}}{y_i!}) \\ &= (\prod_{i=1}^{k} \frac{1}{y_i!}) e^{k(\phi-\lambda)} e^{-n\phi} \lambda^{\sum_{i=1}^{k} y_i} (\prod_{i=k+1}^{n} \phi^{y_i}) \\ &= (\prod_{i=k+1}^{n} \frac{e^{-\phi}\phi^{y_i}}{y_i!})(e^{k(\phi-\lambda)}(\frac{\lambda}{\phi})^{\sum_{i=1}^{k} y_i}) \\ &= f(y, \phi) L(y \mid k, \lambda, \phi). \end{aligned} \tag{8.8}$$

**Listing 8.1.** MCMC Function in R

```
1  bcp<- function(theta.matrix,y,a,b,g,d){
2
3  n<- length(y);
4  k.prob <- rep(0,length=n);
5
6  for (i in 1:nrow(theta.matrix)){
7
8  lambda <- rgamma(1,a+sum(y[theta.matrix[(i-1),3]:n])
9  ,b+theta.matrix[(i-1),3]);
10
11 phi <- rgamma(1,g+sum(y[theta.matrix[(i-1),3]:n])
12 ,d+length(y)-theta.matrix[(i-1),3]);
13
14 for(j in 1:n) k.prob[j] <- ...
       exp(j*(phi-lambda))*(lambda/phi)^sum(y[1:j]);
15
16 k.prob <- k.prob/sum(k.prob);
17 k <- sample(1:n,,size=1,prob=k.prob);
18 theta.matrix[i,] <-c(lambda,phi,k);
19 }
20  return(theta.matrix)
21 }
```

**Example 8.1.3.** The time series stored in the tile gives the # of British cool mining disasters per year over the period $1851 - 1962$.

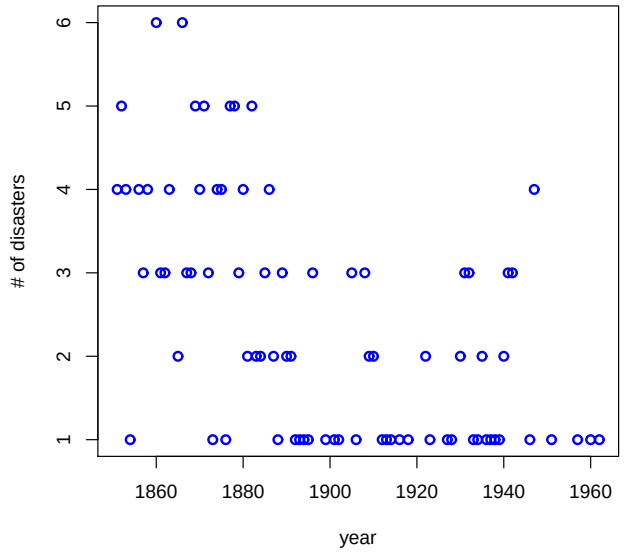


**Figure 1.** Plot of # of disasters over $1851 - 1962$

There has been a reduction in the rate of disasters over the period. Let $y_i$ denote the # of disasters in year $i = 1, ..., 12$ (relabelling the years by numbers from 1 to $n = 112$). A model that has been proposed in the literature has the form:

$$\begin{aligned} y_i &\sim Poisson(\theta), \quad i = 1, ..., k \\ y_i &\sim Poisson(\lambda), \quad i = k+1, ..., n \end{aligned}$$

Let

$$\begin{aligned} \theta &\sim Gamma(a_1, b_1) \\ \lambda &\sim Gamma(a_2, b_2) \\ k &\sim \text{discrete uniform over } \{1, ..., n\} \\ b_1 &\sim Gamma(c_1, d_1) \\ b_2 &\sim Gamma(c_2, d_2) \end{aligned}$$

So

$$\begin{aligned} \pi(\theta, \lambda, k, b_1, b_2 \mid y) \propto\ & e^{-\theta k} \theta^{\sum_{i=1}^{k} y_i} e^{-\theta(n-k)} \lambda^{\sum_{i=k+1}^{n} y_i} b_1^{a_1} \theta^{a_1 - 1} e^{-b_1 \theta} \\ & \times b_1^{c_1 - 1} e^{-d_1 b_1} b_2^{a_2} \lambda^{a_2 - 1} e^{-b_2 \lambda} b_2^{c_2 - 1} e^{-d_2 b_2} I\left[k \in \{1, 2, ..., n\}\right]. \end{aligned}$$

$$\theta \mid y, \lambda, b_1, b_2, k \sim Gamma(a_1 + \sum_{i=1}^{k} y_i, b_1 + k)$$

$$\lambda \mid y, \theta, b_1, b_2, k \sim Gamma(a_2 + \sum_{i=k+1}^{n} y_i, b_2 + n - k)$$

$$b_1 \mid y, \theta, \lambda, b_2, k \sim Gamma(c_1 + a_1, d_1 + \theta)$$

$$b_2 \mid y, \theta, \lambda, b_1, k \sim Gamma(c_2 + a_2, d_2 + \lambda)$$

and

$$p(k \mid y, \theta_1, \lambda, b_1, b_2) = \frac{e^{(\lambda-\theta)k}(\theta/\lambda)^{\sum_{i=1}^{k} y_i} I\left[k \in \{1, 2, ..., n\}\right]}{\sum_{j=1}^{n} \left\{ e^{(\lambda-\phi)j} (\frac{\theta}{\lambda})^{\sum_{i=1}^{k} y_i} \right\}}$$

Burn-in period: The observations obtained after the chain has settled down to the posterior will be more useful in estimating probabilities and expectations for $p(\theta)$. If we throw out the early observations, taken while the process was settling down, the remainder of the process should be a very close approximation to one in which every observation is sampled from the posterior. Dropping the early observations is referred to as using a burn-in period.

Thinning is a process used to make the observations more nearly independent, hence more nearly random sample from the posterior distribution. Frankly, after a burn-in, there is not much point in thinnings unless the correlations are extremely large. If there is a lot of correlation between adjacent observations, a larger overall MC sample size is needed to achieve reasonable numerical accuracy, in addition to needing a much larger burn-in.

# LECTURE 9
# Summary of the properties of Gibbs sampler

(a) Since the Gibbs sampler conditions on values from the last iteration of its chain values, it clearly constitutes a Markov chain.
(b) The Gibbs sampler has the true posterior distribution of parameter vector as its limiting distribution.
(c) The Gibbs sampler is a homogeneous Markov chain, the consecutive probabilities are independent of $n$, the current length of the chain.
(d) The Gibbs sampler converges at a geometric rate: the total variation distance between an arbitrary time and the point of convergence decreases at a geometric rate in time $(t)$.
(e) The Gibbs sampler is an ergodic Markov chain.

## 9.1. Metropolis Algorithm

The Metropolis algorithm[**3**] is another type of accept-reject algorithm. It requires a candidate generating distribution, sometimes referred to as the proposal distribution. The algorithm begins with an initial value $\theta^1$. At the $k$-th iteration we have $(\theta^1, \theta^2, ..., \theta^k)$. The $(k+1)$st iteration first generates $\theta^*$ from a proposal density $h(\theta^* \mid \theta^k)$ This density should mimic the actual posterior distribution in some sense, but in theory, it can be any distribution with the same support as the posterior. Define:

$$\alpha(\theta^*, \theta^k) = \min\left\{1, \frac{p(\theta^*)h(\theta^k \mid \theta^k)}{p(\theta^k)h(\theta^* \mid \theta^k)}\right\} \cong \alpha. \tag{9.1}$$

We then simulate $u \sim u\,[0,1]$ and we select $\theta^{k+1} = \theta^*$ if $u \leq \alpha$ and otherwise take $\theta^{k+1} = \theta^k$. Thus

$$\theta^{k+1} = \begin{cases} \theta^* & \text{with probability} \quad \alpha(\theta^*, \theta^k) \\ \theta^k & \text{with probability} \quad 1 - \alpha(\theta^*, \theta^k) \end{cases} \tag{9.2}$$

Here $\alpha$ only uses the ratio of two values of $p(.)$, so it is enough to know the kernel of the posterior density. The acceptance ration can be also written as:

$$\alpha = \frac{p(\theta^* \mid y)/h_t(\theta^* \mid \theta^{t-1})}{p(\theta^{t-1} \mid y)/h_t(\theta^{t-1} \mid \theta^*)}. \tag{9.3}$$

The ratio $\alpha$ is always defined, because a jump from $\theta^{t-1}$ to $\theta^*$ can only occur if both $p(\theta^{t-1} \mid y)$ and $h_t(\theta^* \mid \theta^{t-1})$ are nonzero.

So here the proposal density is asymmetric.

## 9.2. Metropolis Algorithm

The original Metropolis algorithm assumes that $h(\theta^k \mid \theta^*) = h(\theta^* \mid \theta^k)$ so that $\alpha(\theta^*, \theta^k) = \min\left\{p(\theta^*)/p(\theta^k)\right\}$. This is called the random walk.

Various suggestions are made about how to choose $h(\theta^* \mid \theta^k)$. Often it is taken as a $\mathcal{N}(\theta^k, \sum^k)$ distribution with various suggestions for $\sum^k$.

The Metropolis algorithm is an adaptation of a random walk that uses an acceptance/rejection rule to converge to the specified target distribution. The algorithm proceeds as follows:

(a) Draw a starting $\theta^0$, for which $p(\theta^0 \mid y) > 0$ from a starting distribution $p_0(\theta)$. The starting distribution might, for example, be based on an approximation or we may simply choose starting values dispersed around a crude approximate estimate.

(b) For $t = 1, 2, ...$

- Sample a proposal $\theta^*$ from a jumping distribution (or proposal distribution) at time $t$, $R_t(\theta^* \mid \theta^{t-1})$. For the Metropolis algorithm, the jumping distribution must be symmetric, satisfying the condition $h_t(\theta_a \mid \theta_b) = h_t(\theta_b \mid \theta_a)$ for all $\theta_a$, $\theta_b$ and $t$.
- Calculate the ratio of the densities,

$$\alpha = \frac{p(\theta^* \mid y)}{p(\theta^{t-1} \mid y)} \tag{9.4}$$

- Set

$$\theta^t = \begin{cases} \theta^* & \text{with probability} \quad \min(\alpha, 1) \\ \theta^{t+1} & \text{otherwise} \end{cases} \tag{9.5}$$

Given the current value $\theta^{t-1}$, the transition distribution $h_t(\theta^t \mid \theta^{t-1})$ of the Markov chain is thus a mixture of a point mass at $\theta^t = \theta^{t-1}$, and a weighted version of the jumping distribution, $h_t(\theta^t \mid \theta^{t-1})$, that adjusts for the acceptance rate. The algorithm requires the ability to calculate the ratio $\alpha$ for all $(\theta, \theta^*)$, and to draw $\theta$ from the jumping distribution $h_t(\theta^* \mid \theta)$ for all $\theta$ and $t$ In addition, step (c) requires the generation of a uniform random number.

If $\theta^t = \theta^{t-1}$, that is the jump is not accepted, this counts as an iteration in the algorithm.

Interpretation of the Gibbs sampler as a special case of the Metropolis-Hastling algorithm:

Gibbs sampling can be viewed as a special case of the Metropolis-Hasting's algorithm in the following way. We first define iteration $t$ to consist of a series of $d$ steps, with step $j$ of iteration $t$ corresponding to an update of the subvector $\theta_j$ conditional on all other elements of $\theta$. Then the jumping distribution $h_{j,t}(. \mid .)$ at step $j$ of iteration $t$ only jumps along the $j$-th subvector, and does so with the conditional posterior density of $\theta_j$ density $\theta^{t-1}_{-j}$:

$$h^{Gibbs}_{j,t}(\theta^* \mid \theta^{t-1}) = \begin{cases} p(\theta^*_j \mid \theta^{t-1}_{j-1}, y) & \text{if} \quad \theta^*_{-j} = \theta^t_{-j} \\ 0 & \text{otherwise} \end{cases} \tag{9.6}$$

The only possible jumps are to parameter vectors that match. $\theta^{t-1}$ on all components other than $j$-th. Under this jumping distribution, the ratio at the $j$-th step of iteration $t$ is:

$$\begin{aligned} \alpha &= \frac{p(\theta^* \mid y)/h^{Gibbs}_{j,t}(\theta^* \mid \theta^{t-1})}{p(\theta^{t-1} \mid y)/h^{Gibbs}_{j,t}(\theta^{t-1} \mid \theta^*)} \\ &= \frac{p(\theta^* \mid y)/p(\theta^*_j \mid \theta^{t-1}_{-j}, y)}{p(\theta^{t-1}_j \mid y)/p(\theta^{t-1}_{-j} \mid \theta^*, y)} \\ &= \frac{p(\theta^{t-1}_{-j} \mid y)}{p(\theta^{t-1}_{-j} \mid y)} \\ &= 1 \end{aligned} \tag{9.7}$$

and thus every jump is accepted.

## 9.3. Data Augmentation

It is a technique that can be helpful in making problems amenable to Gibbs sampling.

(a) In the real world, some of our data can be missing.
(b) The likelihood function is not fractable for one reason or another but conditional on a collection of unobserved random variables, the likelihood becomes easy to handle.

**Example 9.3.1.** Suppose $Y_0, Y_1, ...Y_n$ is a time series of random variables defined by $Y_0 = 0$ and for each $i = 1, ..., n$, $Y_i = Y_{i-1} + s_i$ where $s_i \sim Beta(\theta, \theta)$, $\theta > 0$. Therefore:

$$Y_i \mid Y_0, Y_1, ..., Y_{i-1} \sim Y_{i-1} + s_i \tag{9.8}$$

Likelihood of $\theta$

(9.9)

$$f(y_0, ..., y_n \mid \theta) = f(y_0 \mid \theta) \prod_{i=1}^{n} f(y_i \mid y_0, ..., y_{i-1}, \theta)$$

$$= \prod_{i=1}^{n} f(y_i \mid y_{i-1}, \theta)$$

$$= \prod_{i=1}^{n} \frac{\Gamma(2\theta)}{\{\Gamma(\theta)\}^2} (y_i - y_{i-1})^{\theta-1} \{1 - (y_i - y_{i-1})\}^{\theta-1} I\left[0 < y_i - y_{i-1} < 1\right].$$

However suppose that an observation $y_i^*$ is missing then the likelihood is no longer available in closed form.

**Solution:**

$z$ additional variables to be included in the model. $f(y \mid \theta)$ is not tractable but $f(y, z \mid \theta)$ is tractable. The posterior distribution of $(\theta, z)$ is proportional to:

$$\pi(\theta, z \mid y) \propto f(y, z \mid \theta)\pi(\theta). \tag{9.10}$$

$y = (y_0, y_i^*, y_n)$ where $z = y_i^*$.

(9.11)

$$f(y, y_i^* \mid \theta) = f(y_0, ..., y_n \mid \theta)$$

$$= \prod_{i=1}^{n} \frac{\Gamma(2\theta)}{\{\Gamma(\theta)\}^2} (y_i - y_{i-1})^{\theta-1} \{1 - (y_i - y_{i-1})\}^{\theta-1} I\left[0 < y_i - y_{i-1} < 1\right].$$

Therefore the posterior density of $\theta$ given $(y, y_i^*)$ is explicit:

$$\Gamma(\theta \mid y, y_i^*) \propto f(y, y_i^* \mid \theta)\pi(\theta). \tag{9.12}$$

To complete the Gibbs sampler, we also need to sample from the conditional posterior distribution of $y_i^*$. This has density:

(9.13)

$$f(y_i^*, y \mid \theta) \propto f(y, y_i^* \mid \theta)$$
$$\propto \left[(y_i^* - y_{i-1}^*)\left\{1 - (y_i^* - y_{i-1}^*)\right\}(y_{i+1}^* - y_i^*)\left\{1 - (y_{i+1}^* - y_i^*)\right\}\right]^{\theta-1}$$

on the region $y_i^* \in (y_{i-1}^*, y_{i-1}^* + 1) \cap (y_{i+1}^* - 1, y_{i+1}^*)$. This sampling can be carried out by rejection sampling.

# GIBBS SAMPLING

$$
\begin{aligned}
y_i &\sim Poisson(\theta), \quad i = 1, ..., k \\
y_i &\sim Poisson(\lambda), \quad i = k+1, ..., n \\
\theta &\sim Gamma(a_1, b_1) \\
\lambda &\sim Gamma(a_2, b_2) \\
k &\sim \text{discrete uniform over } \{1, ..., n\} \\
b_1 &\sim Gamma(c_1, d_1) \\
b_2 &\sim Gamma(c_2, d_2)
\end{aligned}
$$

Likelihoods:

$$f(y_I \mid \theta) = \prod_{i=1}^{k} f(y_i \mid \theta)$$

$$f(y_J \mid \lambda) = \prod_{j=k+1}^{n} f(y_j \mid \lambda)$$

$$f(y_i \mid \theta) = \prod_{i=1}^{k} \frac{e^{-\theta}\theta^{y_i}}{y_i!} = \frac{e^{-k\theta}\theta^{\sum_{i=1}^{k} y_i}}{\prod_{i=1}^{k} y_i!}$$

$$f(y_j \mid \theta) = \prod_{j=k+1}^{n} \frac{e^{-\theta}\theta^{y_j}}{y_j!} = \frac{e^{-(n-k)\theta}\lambda^{\sum_{i=k+1}^{n} y_i}}{\prod_{i=k+1}^{n} y_i!}$$

Priors:

$$\pi(\theta) = \frac{1}{b_1^{a_1}\Gamma(a_1)}\theta^{a_1-1}e^{-\theta/b_1}$$

$$\pi(\lambda) = \frac{1}{b_2^{a_2}\Gamma(a_2)}\theta^{a_2-1}e^{-\theta/b_2}$$

$$\pi(b_1) = \frac{1}{d_1^{c_1}\Gamma(c_1)}b_1^{c_1-1}e^{-b_1/d_1}$$

$$\pi(b_2) = \frac{1}{d_2^{c_2}\Gamma(c_2)}b_2^{c_2-1}e^{-b_2/d_2}$$

Since

$$g(x;\alpha,\beta) = \begin{cases} \frac{1}{\beta^\alpha \Gamma(\alpha)} x^{\alpha-1} e^{-x/\beta} & \text{for} \quad x > 0 \\ 0 & \text{elsewhere} \end{cases}$$

where $\alpha > 0$ and $\beta > 0$ for Gamma Distribution $(\alpha, \beta)$.

Posterior Distribution

$$\pi(\theta, \lambda, k, b_1, b_2 \mid y) \propto f(y_i \mid \theta) f(y_j \mid \lambda) \pi(\lambda \mid a_2, b_2) \pi(\theta \mid a_1, b_1) \pi(b_1) \pi(b_2) \pi(k)$$

Explicitly,

$$\pi(\theta, \lambda, k, b_1, b_2 \mid y) \propto (e^{-k\theta} \theta^{\sum_{i=1}^{k} y_i})(e^{-(n-k)\theta} \lambda^{\sum_{i=k+1}^{n} y_i})(d_1^{-c_1} b_1^{c_1 - 1} e^{-b_1/d_1})$$
$$(d_2^{-c_2} b_2^{c_2 - 1} e^{-b_2/d_2}) I\left[k \in \{1, 2, ..., n\}\right] (b_2^{-a_2} \theta^{a_2 - 1} e^{-\theta/b_2})(b_1^{-a_1} \theta^{a_1 - 1} e^{-\theta/b_1})$$

So the conditional distributions are:

(a)

$$\begin{aligned} \pi(\theta \mid y, \lambda, k, b_1, b_2) &\propto \int \pi(\theta, \lambda, k, b_1, b_2 \mid y) \mathrm{d}\lambda \mathrm{d}\theta \mathrm{d}b_1 \mathrm{d}b_2 \\ &\propto \int (e^{-k\theta} \theta^{\sum_{i=1}^{k} y_i} \theta^{a_1 - 1} e^{-\theta/b_1}) \text{constant} \mathrm{d}\theta \\ &\propto \int (e^{-(k+1/b_1)\theta} \theta^{\sum_{i=1}^{k} y_i + a_1 - 1} \mathrm{d}\theta \\ &= Gamma(a_1 + \sum_{i=1}^{k} y_i, b_1 + k) \end{aligned}$$

(b)

$$\begin{aligned} \pi(\lambda \mid y, \theta, k, b_1, b_2) &\propto \int \pi(\theta, \lambda, k, b_1, b_2 \mid y) \mathrm{d}\theta \mathrm{d}b_1 \mathrm{d}b_2 \mathrm{d}\theta \\ &\propto \int (e^{-(n-k)\lambda} \theta^{\sum_{i=k+1}^{n} y_i} \theta^{a_2 - 1} e^{-\theta/b_2}) \text{constant} \mathrm{d}\lambda \\ &\propto \int (e^{-((n-k)+b_2)\lambda} \theta^{\sum_{i=1}^{k} y_i + a_2 - 1} \mathrm{d}\lambda \\ &= Gamma(a_2 + \sum_{i=k+1}^{n} y_i, b_2 + (n-k)) \end{aligned}$$

(c)

$$\begin{aligned}\pi(b_1 \mid y, \lambda, \theta, k, b_1) &\propto \int \pi(\theta, \lambda, k, b_1, b_2 \mid y)\mathrm{d}y\mathrm{d}\theta\mathrm{d}b_2\mathrm{d}\theta \\ &\propto \int b_1^{c_1-1} e^{-b_1/d_1} \theta^{a_1-1} e^{-\theta/b_1} \mathrm{constant}\mathrm{d}b_1 \\ &\propto \int b_1^{c_1+a_1-1} e^{-b_1/(d_1+\theta)}\mathrm{d}b_1 \\ &= Gamma(a_1 + c_1, d_1 + \theta)\end{aligned}$$

(d)

$$\begin{aligned}\pi(b_2 \mid y, \lambda, \theta, k, b_1) &\propto \int \pi(\theta, \lambda, k, b_1, b_2 \mid y)\mathrm{d}y\mathrm{d}\theta\mathrm{d}b_1\mathrm{d}k \\ &\propto \int b_2^{c_2-1} e^{-b_2/d_2} \lambda^{a_2-1} e^{-\lambda/b_2} \mathrm{constant}\mathrm{d}b_2 \\ &\propto \int b_2^{c_2+a_2-1} e^{-b_2/(d_2+\lambda)}\mathrm{d}b_2 \\ &= Gamma(a_2 + c_2, d_2 + \lambda)\end{aligned}$$

(e)

$$\begin{aligned}\pi(k \mid y, \lambda, \theta, b_1, b_2) &\propto \int \pi(\theta, \lambda, k, b_1, b_2 \mid y)\mathrm{d}y\mathrm{d}\theta\mathrm{d}b_1\mathrm{d}b_2 \\ &= \frac{e^{(\lambda-\theta)k}(\theta/\lambda)^{\sum_{i=1}^{k} y_i}}{\sum_{j=1}^{n} \left\{ e^{(\lambda-\theta)j}(\theta/\lambda)^{\sum_{i=1}^{j} y_i} \right\}} I\left[k \in \{1, 2, ..., n\}\right]\end{aligned}$$

As the conditional distribution of $k$ is discrete, thereby characterized by a probability mass function.

# DATA AUGMENTATION

Suppose $Y_0, Y_1, ..., Y_n$ is a time series of random defined by $Y_0 = 0$, $i = 1, ..., n$, $Y_i = Y_{i-1} + S_i$ where $S_i \sim Beta(\theta, \theta)$, $\theta > 0$. Therefore

$$Y_i \mid Y_0, Y_1, ..., Y_{i-1} \sim Y_{i-1} + S_i$$

Likelihood for the observations $(y_0, y_1, ..., y_n)$ of $\theta$:

$$f(y_0, ..., y_n \mid \theta) = f(y_0 \mid \theta) \prod_{i=1}^{n} f(y_i \mid y_0, ..., y_{i-1}, \theta) = \prod_{i=1}^{n} f(y_i \mid y_{i-1}, \theta)$$

$$= \prod_{i=1}^{n} \frac{\Gamma(2\theta)}{\{\Gamma(\theta)\}^2} (y_i - y_{i-1})^{\theta-1} \{1 - (y_i - y_{i-1})\}^{\theta-1} I[0 < y_i - y_{i-1} < 1].$$

However suppose that an observation $y_i^*$ is missing. So since the likelihood is no longer in closed form, we assume that $Y_1, ..., Y_n$ are iid data from the mixture density:

$$f(y_i \mid \theta) = \frac{1}{2} \left( \frac{1}{(2\pi)^{1/2}} e^{-y_i^2/2} + \frac{1}{(2\pi)^{1/2}} e^{-(y_i-\theta)^2/2} \right).$$

and

$$L(y_i \mid \theta) = \prod_{i=1}^{n} f(y_i \mid \theta) \left( e^{-y_i^2/2} + e^{-(y_i-\theta)^2/2} \right).$$

Let

- z=the additional variables included in the model(z may be just a single variable or a vector containing several variables).
- $\theta$ = original parameters in the model with prior $\pi(\theta)$.
- $y$ = vector of observations.

So the posterior distribution of $(\theta, z)$ is proportional to

$$\pi(\theta, z \mid y) \propto f(y, z \mid \theta)\pi(\theta)$$

Data augmentation proceeds by carrying out Gibbs sampling to sample successively form $\theta$ and $z$ to produce a sample from this joint distribution. The marginal distribution of $\theta$ therefore the posterior distribution of interest.

- $y = (y_0, ..., y_n)$ excluding $y_i^*$, the missing observation $z = y_i^*$

$$f(y_i^*, y \mid \theta) \propto f(y, y_i^* \mid \theta)$$

$$\propto \left[(y_i^* - y_{i-1}^*)\left\{1 - (y_i^* - y_{i-1}^*)\right\}(y_{i+1}^* - y_i^*)\left\{1 - (y_{i+1}^* - y_i^*)\right\}\right]^{\theta-1}$$

Therefore, the posterior density of $\theta$ given $(y, y_i^*)$ is

$$\pi(\theta, y_i^* \mid y) \propto f(y, y_i^* \mid \theta)\pi(\theta)$$

the conditional posterior distribution of $y_i^*$ is:

$$f(y_i^*, y \mid \theta) \propto f(y, y_i^* \mid \theta)$$

$$\propto \left[(y_i^* - y_{i-1}^*)\left\{1 - (y_i^* - y_{i-1}^*)\right\}(y_{i+1}^* - y_i^*)\left\{1 - (y_{i+1}^* - y_i^*)\right\}\right]^{\theta-1}$$

here $y_i^* \in (y_{i-1}^*, y_{i-1}^* + 1) \cap (y_{i+1}^* - 1, y_{i+1}^*)$.

- Here $z$ is a sequence of $n$ heads on tails, with one element per iteration. $z = (z_1, ..., z_n)$ and

$$z_i = \begin{cases} 1 & \text{if i-th observation is head} \\ 2 & \text{if i-th observation is fail} \end{cases}$$

Suppose the prior of $\theta \sim \mathcal{N}(0, 1)$. Then

$$f(y_i, z_i \mid \theta) = f(y_i \mid z_i, \theta)P(z_i).$$

where

$$f(y_i \mid z_i, \theta) = \begin{cases} e^{-y_i^2/2} & \text{if} \quad z_i = 1 \\ e^{-(y_i-\theta)^2/2} & \text{if} \quad z_i = 2 \end{cases}$$

and $P(z_i = 1) = P(z_i = 2) = 1/2$

Using that

$$\pi(\theta, z \mid y) \propto \left\{\prod_{i=1}^{n} f(y_i \mid z_i, \theta)P(z_i)\right\}\pi(\theta)$$

$$\propto \left(e^{-\frac{1}{2}\sum_{i:z_i=z}(y_i-\theta)^2}\right)\left(e^{-\theta^2}\right) \sim \mathcal{N}\left(\frac{\sum_{i:z_i=z} y_i}{1+n_2}, \frac{1}{1+n_2}\right)$$

where $n_2 = \#$ of observations for which $z_i = 2$

$$P(z_i = 1 \mid \theta, y) = \frac{P(z_i = 1)}{P(z_i = 1) + P(z_i = 2)} = \frac{e^{-y_i^2/2}}{e^{-(y_i-\theta)^2/2} + e^{-y_i^2/2}}$$

$$P(z_i = 2 \mid \theta, y) = \frac{P(z_i = 1)}{P(z_i = 1) + P(z_i = 2)} = \frac{e^{-(y_i-\theta)^2/2}}{e^{-(y_i-\theta)^2/2} + e^{-y_i^2/2}}$$

Then a Gibbs sampler is used to simulate the posterior distribution of $(\theta, z_1, ..., z_n)$.

**Example .0.2.** A genetic linkage $y = (y_1, y_2, y_3, y_4) = (125, 18, 20, 34)$ with cell probabilities $(\frac{2+\theta}{4}, \frac{1}{4}(1-\theta), \frac{1}{4}(1-\theta), \frac{\theta}{4})$; $0 \leq \theta \leq 1$. Prior of $\theta \sim Uniform(0,1)$, so the posterior density of $\theta$:

$$\pi(\theta \mid y) \propto f(y \mid \theta)\pi(\theta) \propto (2+\theta)^{y_1}(1-\theta)^{y_2+y_3}\theta^{y_4} I\left[\theta \in (0,1)\right]$$

Then

(a) Sample the posterior distribution of $\theta$ directly (e.g. via rejection sampling)

(b) Use data augmentation

Augment the observed data $(y_1, y_2, y_3, y_4)$ by dividing the first cell into two partitions, with respective probabilities proportional to $\theta$ and 2. That is,

$$z \mid y, \theta \sim Binomial(y_1, \frac{\theta}{2+\theta})$$

Then the likelihood function is:

$$f(y, z \mid \theta) = f(y \mid \theta)\pi(z \mid y, \theta)$$

$$\propto (2+\theta)^{y_1}(1-\theta)^{y_2+y_3}\theta^{y_4}\binom{y_1}{z}(\frac{\theta}{2+\theta})^{z}(\frac{2}{2+\theta})^{y_1-z}$$

$$= (1-\theta)^{y_2+y_3}\theta^{y_4}\theta^2 z^{y_1-z}\binom{y_1}{z}$$

The conditional posterior of $\theta$ is:

$$(\theta \mid y, z) = \int f(y, z \mid \theta)\mathrm{d}z\mathrm{d}y \propto \theta^{y_4+z}(1-\theta)^{y_2+y_3} I\left[\theta \in (0,1)\right]$$

$$= Beta(y_4 + z + 1, y_2 + y_3 + 1)$$

To complete the Gibbs sampler, we also generate draws from the conditional posterior distribution of $z$ which is

$$z \mid y, \theta \sim Binomial(y_1, \frac{\theta}{2+\theta})$$

# R Codes

**Listing 1.** Triplot Code in R

```
#####################################################
#                                                   #
#      A Sample Triplot by Anil Aksu                  #
#   It is developed to show some basics of R        #
#                                     #
#####################################################

## the range of sampling
x=seq(-4,4,length=101)
## this function gets numbers from console
prior=dnorm(x, mean = 0.5, sd = 0.7, log = FALSE)
likelihood=dnorm(x, mean = 0.49, sd = 0.65, log = FALSE)
posterior=dnorm(x, mean = 0.52, sd = 0.5, log = FALSE)


## let's plot them
plot(range(x), range(c(likelihood,prior,posterior)), ...
    type='n', xlab=expression(paste(theta)), ...
    ylab=expression(paste("f(", theta, " )")))
lines(x, prior, type='l', col='blue')
lines(x, likelihood, type='l', col='red')
lines(x, posterior, type='l', col='green')

title("Prior, Likelihood and Posterior Distribution")
legend(
  "topright",
  lty=c(1,1,1),
  col=c("blue", "red", "green"),
  legend = c("prior", "likelihood","posterior")
)
```

**Listing 2.** Inference Plots Code in R

```
#####################################################
```

```
#                                                  #
#      Posterior, Perspective and Contour plots    #
#                  by Anil Aksu                    #
#                                     #
#####################################################

## the range of sampling
x=seq(0,20,length=101)
## this function gets numbers from console
posterior=dnorm(x, mean = 7, sd = 1.5, log = FALSE)


## let's plot them
plot(range(x), range(c(posterior)), type='n', ...
    xlab=expression(paste(theta)), ...
    ylab=expression(paste("f(", theta, "|x )")))

lines(x, posterior, type='l', col='blue',lwd=5)

title("Posterior Distribution")
  legend = c("posterior")

## perspective plot
x <- seq(-10, 10, length= 30)
y <- x
f <- function(x, y) { r <- sqrt(x^2+y^2); 10 * ...
    sin(r)/r }
z <- outer(x, y, f)
z[is.na(z)] <- 1
op <- par(bg = "white")
persp(x, y, z, theta = 30, phi = 30, expand = 0.5, ...
    col = "lightblue")
persp(x, y, z, theta = 30, phi = 30, expand = 0.5, ...
    col = "lightblue",
       ltheta = 120, shade = 0.75, ticktype = "detailed",
       xlab = "X", ylab = "Y", zlab = "Sinc( r )"
) -> res
round(res, 3)

# contour plot
a <- expand.grid(1:20, 1:20)
b <- matrix(a[,1]^2 + a[,2]^2, 20)
filled.contour(x = 1:20, y = 1:20, z = b,
               plot.axes = { axis(1); axis(2); ...
                   points(10, 10) })


## bivariate posterior sampling
```

```
44
45 ## the range of sampling
46 x=seq(-4,6,length=101)
47 ## this function gets numbers from console
48 posterior=0.8*dnorm(x, mean = 0, sd = 1, log = ...
       FALSE)+0.2*dnorm(x, mean = 4, sd = 1, log = FALSE)
49
50 ## let's plot them
51 plot(range(x), range(c(posterior)), type='n', ...
       xlab=expression(paste(theta)), ...
       ylab=expression(paste("f(", theta, "|x )")))
52
53 lines(x, posterior, type='l', col='blue',lwd=5)
54
55 # title("Bivariate Posterior Distribution")
56   legend = c("posterior")
57
58 ## credible interval posterior plot
59
60 ## the range of sampling
61 x=seq(0,20,length=101)
62 ## this function gets numbers from console
63 posterior=dnorm(x, mean = 7, sd = 2, log = FALSE)
64
65 ## let's plot them
66 plot(range(x), range(c(posterior)), type='n', ...
       xlab=expression(paste(theta)), ...
       ylab=expression(paste("f(", theta, "|x )")))
67
68 lines(x, posterior, type='l', col='blue',lwd=5)
69
70 # title("Bivariate Posterior Distribution")
71   legend = c("posterior")
```

**Listing 3.** Rejection Sampling Plots Code in R

```
1  ######################################################
2  #                                                    #
3  #            Rejection sampling plots                  #
4  #                  by Anil Aksu                        #
5  #                                      #
6  ######################################################
7
8
9  require(SMPracticals)
10 ## rejection sampling
11
```

```
## the range of sampling
x=seq(-10,10,length=101)
## this function gets numbers from console
posterior=0.6*dnorm(x, mean = 0, sd = 4, log = ...
    FALSE)+0.4*dnorm(x, mean = 6, sd = 2, log = FALSE)
envelope=2*dnorm(x, mean = 2, sd = 5, log = FALSE)
## let's plot them
plot(range(x), range(c(posterior,envelope)), ...
    type='n', xlab=expression(paste(theta)), ylab="")
lines(x, posterior, type='l', col='blue',lwd=5)
lines(x, envelope, type='l', col='red',lwd=5)

title("Rejection Sampling Plot")
legend("topright", legend=c(expression(paste("L(", ...
    theta, ")",pi,"(", theta, ...
    ")")),expression(paste(mu,"g(", theta, ")"))),
    lty=1, col=c('blue', 'red'),inset = .02)

## British Coal Mining accidents
data(coal)
# years of coal mining accidents
years <- unique(as.integer(coal$date))
# the number of accidents in each year
accident <- integer(length(years))
for (i in 1:length(years)){
accident[i]<-sum(as.integer(coal$date) == years[i])
}

plot(years ,accident, col='blue',lwd=2, xlab="year", ...
    ylab="# of disasters")
#rug(coal$date)
```

# BIBLIOGRAPHY


1. Allen B. Dawney. *Think Bayes: Bayesian Statistics in Python*. O'REILLY, 2013.
2. D. Gamerman and F. L. Herbert. *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*. Chapman and Hall/CRC, 2006.
3. Peter D. Hoff. *A First Course in Bayesian Statistical Methods*. Springer Publishing Company, Incorporated, 1st edition, 2009.
4. Sheldon Ross. *Introduction to Probability Models*. Academic Press, Boston, 2014.