

Learning outcomes: Chapter 1

1. You should

- know the difference between the classical/frequentist/subjective interpretations of probability;
- be able to give examples of events whose probabilities can be found using the classical/frequentist/subjective interpretations of probability;
- be able to explain if the probabilities of events can be found using the classical/frequentist/subjective interpretations of probability;
- be able to explain the drawbacks of the classical/frequentist/subjective interpretations of probability.

2. You should be able to

- state the result for the *conditional probability* of one event given that another has occurred;
- define a *partition* of the sample space;
- state, and prove, the *Law of Total Probability*;
- state, and prove, *Bayes Theorem*;
- apply Bayes Theorem to simple examples in probability;
- use Bayes Theorem to formulate diagnostic rules given some observed states/symptoms.

3. You should

- know that the *likelihood function* for a parameter θ is simply the joint probability density function after observing data $\mathbf{x}$; assuming i.i.d. observations, this gives

$$f(\mathbf{x}|\theta) = \prod_{i=1}^n f_X(x_i|\theta);$$

- be able to determine the likelihood function for commonly-used statistical models, or indeed *any* given probability density function for data $\mathbf{x}$.

4. You should

- be able to state the *Factorisation Theorem*;
- be able to show that a given statistic $T(\mathbf{X})$ is *sufficient* for a parameter θ by using the Factorisation Theorem;
- be able to use the Factorisation Theorem to determine a sufficient statistic for θ .

Learning outcomes: Chapter 2

1. You should be able to state Bayes' Theorem for distributions, and you should know that

$$\text{posterior} \propto \text{prior} \times \text{likelihood}.$$

2. You should be familiar with the continuous Uniform distribution, the beta distribution and the gamma distribution, although you need *not* memorise their probability density functions, expectations and variances.
3. For a range of models $f(\mathbf{x}|\theta)$ for data $\mathbf{x}$, you should be able to obtain the posterior for θ given a prior $\pi(\theta)$. Examples include:

$X_i \sim$	Prior for θ	Posterior
$Bin(n, \theta)$	Uniform	Beta
	Beta	Beta
$Po(\theta)$	Exp	Gamma
	Gamma	Gamma
$Exp(\theta)$	Gamma	Gamma
$N(\theta, 1/\tau)$	Normal	Normal

You should be able to work with specific scenarios using observed data $\mathbf{x}$, as well as more general scenarios for $\mathbf{X} = X_1, X_2, \dots, X_n$. Examples other than those in the table above may appear in the exam.

4. You should be able to describe how prior beliefs about a parameter of interest θ change in light of some observed data, via
 - prior and posterior numerical summaries;
 - plots of prior and posterior densities.
5. You should understand, and be able to use, the *Bayes linear rule*:

$$E(\theta|\mathbf{x}) = \alpha E(\theta) + (1 - \alpha)f(\mathbf{x}),$$

where f is a function of the data (and usually the likelihood mode).

6. You should understand the role of *sufficient statistics* in Bayesian inference. In particular, you should be able to show that posterior beliefs about a parameter of interest θ having observed the full data $\mathbf{x}$ are exactly the same as having observed only a sufficient statistic.

Learning outcomes: Chapter 3

1. You should know how to obtain parameters for a prior distribution for θ given some suggested prior summaries for θ . For example,
 - parameters for the gamma distribution for θ given $E(\theta)$ and $Var(\theta)$;
 - parameters for the beta distribution for θ given the mode for θ and a probability for θ .
2. You should be able to describe the *trial roulette* and *bisection* methods of prior elicitation. You should also be able to:
 - Interpret output from the *MATCH* elicitation tool;
 - explain the processes of *feedback* and *refinement*;
 - explain the role of *feedback percentiles*, and be able to transpose these into suitable questions to ask a non-statistician.
3. You should be able to explain how we might operate Bayes' theorem with *vague prior knowledge*, and you should be able to obtain posterior distributions in circumstances when our prior knowledge about θ is indeed vague.
4. You should be able to explain *prior ignorance*; you should know how to construct ignorance priors for θ , including Jeffreys priors.
5. Given a model $f(\mathbf{x}|\theta)$, you should be able to obtain the asymptotic posterior distribution for θ .

Learning outcomes: Chapter 4

1. You should understand the difference between a *point estimate* and an *interval estimate*.
2. You should be able to define a *frequentist confidence interval*, a *Bayesian confidence interval* and a *highest density interval*, and you should understand the differences in interpretation between these intervals.
3. You should be able to determine a highest density interval for θ in the following simple situations:
 - When the posterior for θ is a beta distribution and the lower/upper bounds of the interval are clearly 0/1;
 - When the posterior for θ is symmetric, e.g. a Normal distribution.
4. You should be able to write simple functions in R to obtain a highest density interval for any unimodal posterior distribution.
5. You should be able to use the following result to obtain the predictive probability density function for a future observation y :

$$f(y|\mathbf{x}) = \int_{\Theta} f(y|\theta) \pi(\theta|\mathbf{x}) d\theta.$$

6. You should be able to obtain prediction intervals in simple examples for discrete θ , or when the predictive distribution is of known form.