

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.