

Introduction - Bayesian Statistics

Bayesian inference

Bayesian inference models the process of learning. That is, we start with a certain level of belief, however vague, and through the accumulation of experience, our belief becomes more fine-tuned. Some people take a dislike to Bayesian inference because it is overtly subjective and they like to think of statistics as being objective. We don't agree, because any statistical analysis is necessarily subjective as a result of the need to make assumptions, but also because many approximations are accepted without question, even without warning that they have been made. For that reason we appreciate the extra transparency of Bayesian inference, but it also frequently provides answers where classical statistics cannot. Perhaps more importantly for us as risk analysts, Bayesian inference encourages our clients to think about the level of knowledge *they* have about their problem, and what that means to them.

Bayesian inference is an extremely powerful technique, based on [Bayes' Theorem](#) (sometimes called Bayes' Formula), for using data to improve one's estimate of a parameter. There are essentially three steps involved:

1. Constructing a confidence distribution of the parameter before analyzing the new data set. This is called the prior distribution;
2. Find an appropriate [likelihood function](#) for the observed data; and
3. Modify the prior distribution using the likelihood function to get a revised estimate known as the posterior distribution.

Bayesian inference options

Number of parameters to estimate	1	2	3+
Prior selection:			
	Simple	Uncorrelated	
		Correlated	
Analytical methods:	Markov Chain Monte Carlo		
	Accept / Reject Simulation		
	Construction		
	Conjugate		

This section starts with some introductory topics:

[Introduction to the concept and some simple examples](#)

[How to determine prior distributions](#)

[How to determine likelihood functions](#)

We then turn to the actual execution of a Bayesian inference for risk analysis models with Crystal Ball by looking at various techniques for arriving at the posterior distribution:

[Construction method](#)

[Conjugate prior method](#)

[Simulation with accept/reject method](#)

[Markov Chain Monte Carlo method](#)

Most important of all, we offer a number of worked examples:

Examples of Bayesian inference calculations

General estimation problems

[Identifying a weighted coin](#)

A simple Bayesian inference example using [construction](#)

[Tigers in the jungle](#)

Simple construction model showing the interaction between likelihood functions and informed priors

[Gender of a random sample of people](#)

A simple construction model illustrating the importance of the prior distribution

[Micro-fractures on turbine blades](#)

A model to show how to incorporate hyperparameters by simulation, as well as offering both simulation and construction approaches to determining the posterior distribution

[The Monty Hall problem](#)

A fun model of a classic problem showing the sometimes non-intuitive nature of the Bayesian result

[Using cow pats to estimate infected animals in a herd](#)

A bit gross, but this example shows how you can perform Bayesian inference calculations with very involved stochastic processes (i. e. very complicated likelihood functions) by using Monte Carol simulation with accept/reject criteria.

[Bayesian estimation of a components mean time to failure MTTF](#)

A simple construction example that shows how we use data that describe being above or below a threshold, instead of exact observations

Taylor series approximation to a Bayesian posterior distribution

Showing how Taylor series expansion lets you determine the normal approximation to posterior distributions, and a method for algebraically obtaining the standard deviation

[Normal approximation to the Beta posterior distribution](#)

Example of a Taylor series expansion

Two common statistical problems

[Estimate of the mean of a Normal distribution with unknown standard deviation](#)

A standard statistics problem with the same outcome as the [classical method](#)

[Bayesian estimate of the mean of a Normal distribution with known standard deviation](#)

Another standard statistics problem with the same outcome as the [classical method](#)

Once you have reviewed the material in this section, you might like to test how much you have learned by taking the self-test quiz:

Take

A quiz on Bayesian inference: [quiz](#)
