Bayesian Statistics

Debdeep Pati Florida State University

April 9, 2013

◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 臣 のへぐ

Historical Background

The Reverend Thomas Bayes, began the objective Bayesian theory, by solving a particular problem

- Suppose X is Binomial (n,p); an 'objective' belief would be that each value of X occurs equally often.
- The only prior distribution on p consistent with this is the uniform distribution.
- Along the way, he codified Bayes theorem.
- Alas, he died before the work was finally published in 1763.



REV. T. BAYES

Historical Background

The real inventor of Objective Bayes was Simon Laplace (also a great mathematician, astronomer and civil servant) who wrote *Théorie Analytique des Probabilité* in 1812

- He established the 'central limit theorem' showing that, for large amounts of data, the posterior distribution is asymptotically normal (and the prior does not matter).
- He virtually always utilized a 'constant' prior density (reasons: CLT; parameter choice; robustness).
- He solved very many applications, especially in physical sciences.
- He had numerous methodological developments, e.g., a version of the Fisher exact test.
- Later in his life he invented frequentist statistics.



What's in a name, part I

- It was called *probability theory* until 1838.
- From 1838-1950, it was called *inverse* probability, apparently so named by Augustus de Morgan.
- From 1950 on it was called *Bayesian* analysis (as well as the other names); for why, see Fienberg (2006).



Augustus de Morgan

- Stigler (1983) attributes it to Saunderson (1683-1739), a blind professor of Optics
- The first deduction of the least square method made by Gauss (1795) using Bayesian methods

◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 臣 の�?

Motivating example

- Assess whether a selected population for growth rate has a higher growth rate than a control population.
- Classical statistics: the hypothesis to be tested is that there is no difference between the two treatments
- Before making the experiment, the error of rejecting this hypothesis when it is actually true is fixed at a level of 5%
- ► Repeating the experiment an infinite number of times difference between the averages of these samples (x
 ₁ - x
 ₂)

▶ true value of the difference between selected and control populations $(m_1 - m_2)$

Motivation

If our sample lies in the shadow area,

- There is no difference between treatments, our sample is a rare one
- ▶ The treatments are different, and repeating an infinite number of times the experiment, $(\bar{x}_1 \bar{x}_2)$ will not be distributed around zero but around an unknown value different from zero.



- Natural to find the most probable value of a parameter based on our data rather than to find which value of this parameter, if it would be the true value, would produce our data with a highest probability.
- To make probability statements based on our data we need some prior information and it is not clear how to introduce this prior information in our analysis or how to express lack of information using probability statements.

Apply Bayes Theorem!

► A, B are 2 events

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

- Interested in assessing effect of a drug on growth rate of a rabbit population
- Selected group of rabbits and a control group in which growth rate has been measured.
- S : Effect of the drug the selected group, C: Effect of the control group, Interested in assessing (S − C).
- ► Want to find the probabilities of all possible values of (S C) according to the information provided by our data.
- This can be expressed as $P(S C \mid y)$

The posterior distribution

$$P(S-C \mid y) = \frac{P(y \mid S-C)P(S-C)}{P(y)}$$

- ► P(y | S C) : distribution of the data for given value of the unknown, often known or assumed to be known from reasonable hypotheses.
- ► P(S C) : Prior probability of the difference between selected and control group independent of data.

• P(y): the probability of the sample.

 Information about the parameters we want to estimate that exists before we perform our experiment.

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > <

- It is almost impossible to do this formally, with some exceptions. We will distinguish three scenarios:
 - 1. exact prior information
 - 2. vague prior information
 - 3. No prior information describing ignorance

Exact prior information



◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 三臣 - のへで

Vague / No prior information

- Toss a coin n times want to estimate the probability of head
- Three states of beliefs were tested
- Prior 3 called objective or non-informative by Bayesian statisticians



- General set up: $y_i \sim f(y \mid \theta), \theta \sim \Pi(\theta)$ &
- Obtain posterior distribution $\Pi(\theta \mid y_1, \dots, y_n)$ as

$$\Pi(\theta \mid y_1, \ldots, y_n) = \frac{\prod_{i=1}^n f(y_i \mid \theta) \Pi(\theta)}{\int_{\theta} \prod_{i=1}^n f(y_i \mid \theta) \Pi(\theta)}$$

1877 Algorithm: Normal Prior-Posterior



- $X \sim Bin(n, p), p \sim Beta(a, b)$
- $p \mid X \sim \text{Beta}(a + X, n X + b)$

◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 臣 のへぐ

- Measure of Discrepancy $R = E_{\theta, y_1, \dots, y_n} L(\hat{\theta}(y_1, \dots, y_n), \theta)$
 - 1. Posterior mean: minimizes R with squared error L
 - 2. Posterior median: minimizes R with L as absolute deviation

◆□▶ ◆□▶ ◆三▶ ◆三▶ 三三 のへぐ

3. Posterior mode minimizes R with L as the 0-1 loss.



◆□ > ◆□ > ◆ □ > ◆ □ > → □ = → ○ < ⊙

Loss functions

- Mean: 2-fold inconvenience: penalizes high errors, this risk function is not invariant to transformations
- Mode: signifies the most probable value, easier to calculate in the pre-MCMC era - may not be representative
- Median: true value has a 50% of probability of being higher or lower than the median. Attractive loss - invariant to one-to-one transformations





◆□ > ◆□ > ◆ □ > ◆ □ > → □ = → ○ < ⊙

Posterior mean for the Beta-Binomial problem

Posterior mean:

$$E(p \mid X) = \frac{a + X}{a + b + n}$$

= $\frac{X}{n} \frac{n}{\alpha + \beta + n} + \frac{\alpha}{\alpha + \beta} \frac{\alpha + \beta}{\alpha + \beta + n}$

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > <

 MLE * (precision of MLE) (precision of MLE + prior precision) + prior mean * (prior precision) (precision of MLE + prior precision)

- Confidence interval: How often the interval contains the true value if the samples are generated according to the truth
- Credible intervals: Given the data, how often does the interval contains the true parameter
- $P(\theta \in [L(y), U(y)] \mid y) = 0.95$
- Can find the shortest interval with a 95% probability of containing the true value (what is called the Highest posterior density interval at 95%).

Highest posterior density credible interval



Shortest interval with P=0.95



Symmetric interval with P=0.95

◆□▶ ◆□▶ ◆三▶ ◆三▶ 三三 のへぐ