A BRIEF INTRODUCTION TO

Bayesian methods

Bayesian statistics is an inferential framework rooted in the work of 18th-century mathematician Thomas Bayes. The approach applies a mathematical method *(Bayes' Theorem)* to update the probability of an effect as new evidence is observed.

DID YOU KNOW?

Bayesian analysis uses the same data, outcome measures, and analytical designs as conventional frequentist methods, serving as a complementary approach to evaluate impacts and quantify uncertainty. It can be implemented with statistical software (e.g. Python, R, Stata).

BAYESIAN

A framework that uses observed data to estimate the probability of an effect, updating prior assumptions in light of that data.

HYPOTHESIS

Compares the prior distribution to the posterior distribution. Treats the effect as uncertain and asks: **Which values of the effect look most plausible given the observed data?**

HOW IT WORKS

Explores the range of possible effect sizes using simulation algorithms (e.g. Markov Chain Monte Carlo), running thousands of cycles to weight each possibility by how well it fits the observed data and prior distribution. Constructs a posterior distribution.

Examples: Bayesian regression, Bayesian adaptive trials

VS

FREQUENTIST

A framework that uses observed data to estimate an effect, evaluating how unusual those data would be under a given hypothesis or assumption.

HYPOTHESIS

Tests the null hypothesis. Assumes there is one true, fixed effect and asks: **How surprising is the data if the true effect is zero?**

HOW IT WORKS

Calculates an estimate from the observed sample, deriving uncertainty from the expected variation of that statistic across hypothetical repeated samples from the same population.

Examples: regression, difference in means, ANOVA

Key terminology

Prior: what is believed about the effect before seeing the data (can be deliberately vague or draw on previous research)
Likelihood: how well each possible effect size fits the data observed
Posterior: what is believed about the effect once the data has been taken into account; the prior updated by the likelihood
Credible interval: a range that contains the true effect with a specified probability

Null hypothesis: The assumption that there is no true effect.
Estimate: A number calculated from sample data that quantifies an unknown effect.
p-value: A measure of how surprising the observed result would be if the null hypothesis were true.
Confidence interval: The range of effects compatible with the observed data.

BAYESIAN METHODS

Implementation

BAYESIAN

INPUTS

Data, model specification, prior distribution, simulation settings (e.g. number of cycles)

OUTPUTS

Posterior distribution, point estimate (mean or median), credible interval

INTEPRETATION

The probability that the true effect lies within a specific range (“X% likely that the true effect is between A and B”), or the probability it cleared a specific threshold (“X% likely to have exceeded Y”)

VS

FREQUENTIST

INPUTS

Data, model specification, null hypothesis, significance threshold (e.g. 5%), target confidence level (e.g. 95%)

OUTPUTS

Point estimate, p-value, confidence interval

INTERPRETATION

How surprising the observed difference is if the data was generated by chance (“data this extreme would arise X% of the time under the null hypothesis”; “compatible with impacts between A and B”)