



Bayesian Statistics for Decision-Makers

Training: Priors, Posteriors and Probabilistic Forecasts



Bayesian Statistics for Decision-Makers Training: Priors, Posteriors and Probabilistic Forecasts

Introduction:

Bayesian statistics for decision-makers, covering priors, posteriors and probabilistic forecasts, is a 5-day course for risk, reliability, product and operations analysis teams, ending with a Bayesian Decision Plan and Uncertainty Report. Many units still treat a small trial, a handful of failures or a short A/B test as either proof or noise, and decisions stall or proceed on false certainty. Nominees already analyse data in spreadsheets or code and now need to combine expert judgement with evidence, taught through a modelling build in Python or R. CoreConcept Training Center delivers this Bayesian statistics course.

Course Objectives:

- Translate a business question into a prior, a likelihood and a posterior quantity that answers it
- Elicit defensible priors from subject experts and document how sensitive conclusions are to them
- Fit conjugate and MCMC models in Python or R and confirm convergence before using the results
- Report credible intervals, posterior probabilities and predictive ranges that managers can act on
- Compare competing models with cross-validation and Bayes factors and justify the model retained
- Prepare a Bayesian Decision Plan and Uncertainty Report that sets decision thresholds from expected loss

Target Audience:

- Risk analysis staff who estimate the likelihood and impact of losses, incidents and claims
- Reliability and maintenance engineering staff who judge equipment from sparse failure records
- Product, digital and marketing teams who run A/B tests and pricing or campaign trials
- Quality and supplier management staff who accept or reject lots and vendors on small samples
- Operations and planning analysts who must state uncertainty in projections given to management

Course Outline:

Day 1: Bayesian Reasoning for Decisions Under Uncertainty

- Bayes' Theorem Applied to Diagnostic and Risk Questions
 - Frequentist P-Values Versus Posterior Probability Statements Compared
 - Likelihood Functions for Counts, Rates and Continuous Measurements
 - Base-Rate Neglect and Probability Calibration Exercises for Managers
 - Current Decision Process Audit for Hidden Uncertainty Assumptions
-

Day 2: Prior Elicitation, Conjugate Models and the Bayesian Workflow

- Bayesian Workflow From Model Specification to Posterior Checking
- Structured Prior Elicitation With Domain Experts and Quantile Matching
- Beta-Binomial Model for Conversion, Defect and Approval Rates
- Gamma-Poisson Model for Failure, Incident and Claim Counts
- Normal-Normal Model for Small-Sample Mean Estimation

Day 3: MCMC Computation and Posterior Summaries in Python or R

- PyMC and ArviZ Model Build in Python Notebooks
- Equivalent brms and Stan Model Build in R
- NUTS Sampler Settings With R-hat and Effective Sample Size
- Highest Posterior Density and Equal-Tailed Credible Interval Reporting
- Posterior Predictive Checks Against Observed Operational Data

Day 4: Hierarchical Models, Model Comparison and Problem Cases

- Hierarchical Partial Pooling Across Branches, Sites and Products
- LOO Cross-Validation and WAIC for Competing Model Comparison
- Bayes Factors and Prior Sensitivity Analysis Limits
- Weibull Reliability Model for Sparse Equipment Failure Data
- Divergent Transitions and Prior-Data Conflict Diagnosis

Day 5: Modelling Build of a Bayesian Decision Plan for a Case Organisation

- Bayesian A/B Test With Expected Loss Stopping Rule
- Small-Sample Supplier Defect Evaluation Using Beta-Binomial Updating
- Posterior Predictive Forecast of Next-Quarter Incident Counts
- Decision Threshold Setting From Posterior Expected Cost
- Bayesian Decision Plan and Uncertainty Report Completion

Skills You Will Gain:

- Prior Elicitation
 - Posterior Inference
 - Probabilistic Programming
 - MCMC Convergence Diagnosis
 - Credible Interval Interpretation
 - Hierarchical Modelling
 - Bayesian Model Comparison
 - Expected Loss Analysis
-

Why Attend This Course:

- Deliver a Bayesian Decision Plan and Uncertainty Report to the department head and the risk or investment committee that approves the decision
- Decide when an A/B test, supplier trial or reliability review has gathered enough evidence to act
- Avoid costly reversals caused by acting on a single small sample or ignoring what experts already know
- Coach colleagues to read credible intervals and posterior probabilities correctly in management reports

Conclusion:

Back at work, the participant gives the department head and the relevant risk or investment committee a Bayesian Decision Plan and Uncertainty Report for one live decision, such as launching a product variant, accepting a supplier or extending an asset's service interval. Managers use the posterior probabilities and expected loss figures to approve, delay or stop the decision with a stated level of confidence. After the first decision cycle, the unit should compare outcomes with the posterior predictive ranges, revisit the priors and update the model with the new data.

Frequently Asked Questions FAQ:

What should participants know before a Bayesian statistics course for decision-makers?

Participants should be comfortable with averages, proportions and basic probability and should have run simple analyses in spreadsheets, Python or R. No prior Bayesian experience is expected. Bringing an anonymised small dataset or a pending decision from their own unit makes the modelling build more useful.

How does Bayesian statistics for decision-makers differ from a general business statistics course?

It replaces p-values and confidence intervals with priors, posterior probabilities and expected loss. General business statistics courses cover descriptive measures, hypothesis tests and regression broadly, while this course goes deeper into combining expert judgement with sparse data and turning posteriors into decision thresholds.

Why is Bayesian statistics useful for decisions with small samples?

Bayesian statistics lets a unit combine a small new sample with documented prior knowledge, so estimates stay stable and uncertainty is stated honestly. The posterior answers the direct question managers ask, such as the probability that option B beats option A, rather than a test statistic.

What do participants take back to work from the Bayesian statistics for decision-makers course?

Participants take back a Bayesian Decision Plan and Uncertainty Report for a real decision, with the prior rationale, fitted model, convergence checks, credible intervals and an expected loss threshold. They also keep the Python or R notebooks used in class as reusable templates.