Testing evidence of absence

Tools for Bayesian analysis and equivalence testing

Paul Bürkner

Definitions of Probability

Probability as the limit of relative frequencies – examples:

- "In 40% of the elections 2020, Trump will win"
- "The true parameter value lies within 95% of the confidence intervals"

Probability as the representation of uncertainty – examples:

- "In the election 2020, Trump will win with probability 40%"
- "With 95% probability, the parameter value lies within the credible interval"

Prior and Posterior Uncertainty

- Before the data collection, we have certain **prior** uncertainty about the effects under study.
- After collecting the **data**, we update our uncertainty, which then becomes our **posterior** uncertainty.
- Bayesian inference gets us from prior to posterior uncertainty.

Bayes Theorem

- The posterior probability of the parameters given data is

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)}$$

- Likelihood: $p(y|\theta)$
- Prior: $p(\theta)$
- Marginal likelihood / Evidence: $p(y)$

A Simple Example

- We are interested in measuring the abundance of a specific genetic variant.
- We want to model the rate θ (our parameter) that individuals of the population have the genetic variant of interest.

We need:

- The genetic data for each individual
- The likelihood / generative model (probability of the data given the parameters)
- The prior (probability of the parameters before seeing the data)

The Binomial Likelihood

- We call y_i the outcome of individual i
- The individual may have the genetic variant $(y_i = 1)$ or not $(y_i = 0)$
- The genetic variant occurs with probability θ
- The genetic variant is absent with probability $1 - \theta$
- We have data on a total of N individuals

The Binomial likelihood for the number of individuals with the genetic variant $y := \sum_{i=1}^{N} y_i$:

$$p(y \mid \theta, N) = \binom{N}{y} \theta^y (1 - \theta)^{N-y}$$

Results for Data $y = 4, N = 10$ and a Flat Prior

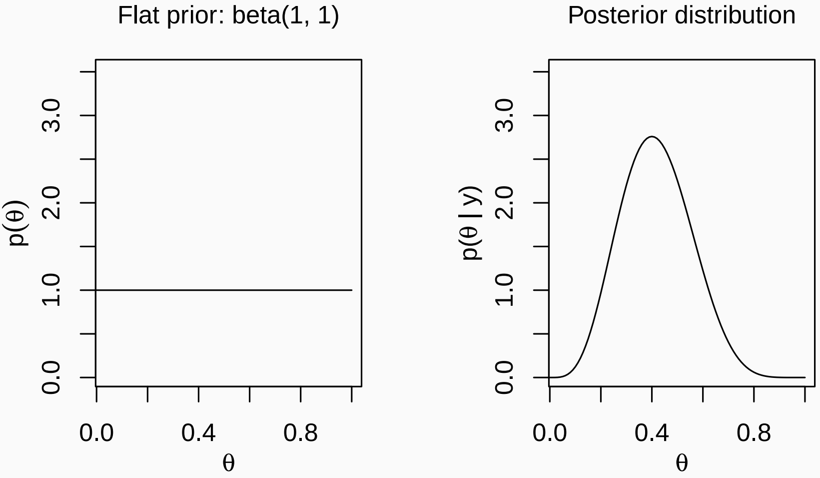

Results for Data $y = 4, N = 10$ and an Informative Prior

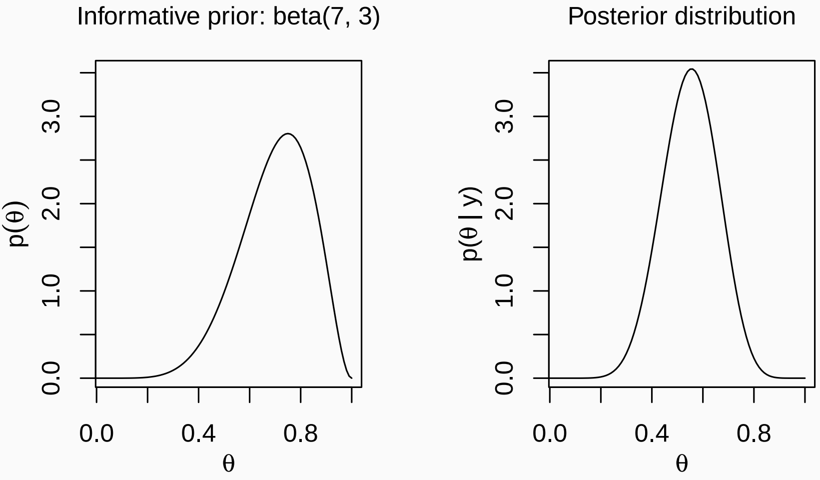

Sampling from the Posterior distribution

Why Do We need Sampling?

- For simple models, we can compute the marginal likelihood $p(y) = \int p(y|\theta)p(\theta)d\theta$ analytically.

Binomial likelihood with flat prior:

$$p(y) = \int_0^1 \binom{N}{y} \theta^y (1-\theta)^{N-y} \times 1 \; d\; \theta = \frac{1}{N+1}$$

- For a bit more complex models, integration may be done numerically.
- For more than 3 or 4 parameters, numerical computation of the marginal likelihood becomes infeasible
- We need to sample (somehow) from the posterior

Rejection Sampling

- Sample parameter values from the prior
- Sample data from the likelihood based on the sampled parameters
- Only keep those parameter values, which produced data consistent with our observed data
- Repeat this process many times
- The kept parameter values are samples from the posterior

Rejection Sampling: Examples

Hypothetical Sample 1:

- Sampling from the prior yields $\theta_s = 0.7$
- Sampling from the binomial likelihood $p(y \mid \theta = 0.7, N = 10)$ yields $y_s = 6$
- The sample response $y_s = 6$ is different from the observed response $y = 4$ so that $\theta_s = 0.7$ is thrown away

Hypothetical Sample 2:

- Sampling from the prior yields $\theta_s = 0.42$
- Sampling from the binomial likelihood $p(y \mid \theta = 0.42, N = 10)$ yields $y_s = 4$
- The sample response $y_s = 4$ is equal to the observed response $y = 4$ so that $\theta_s = 0.42$ is kept

Using Samples to Approximate Expectations

(Almost) every quantity of interest is an expectation over $p(\theta|y)$:

$$\mathbb{E}_p(h) = \int h(\theta) \, p(\theta \mid y) \, \mathrm{d}\theta$$

Having obtained exact random samples $\{\theta_s\}$ from $p(\theta \mid y)$:

$$\frac{1}{S} \sum_{s=1}^{S} h(\theta_s) \sim \text{Normal}\left(\mathbb{E}_p(h), \sqrt{\frac{\text{Var}_p(h)}{\text{S}}}\right)$$

Evaluation of Hypotheses

Different ways to evaluate hypotheses:

- Estimation with uncertainty intervals
- Posterior probabilities
- Bayes factors

All of them can be used for equivalence testing!

Estimation with Uncertainty

Estimation with Uncertainty

For inference use:

- Point estimates
- Uncertainty intervals (UIs)

Bayesian Uncertainty intervals:

- Credible intervals based on quantiles (CIs)
- Highest posterior density intervals (HDIs)

Computation of Point Estimates in R

Let `posterior` be a vector of posterior samples of θ:

```r
head(posterior)
```

```
## [1] 0.59 0.49 0.42 0.24 0.32 0.43
```

Computation of the posterior mean: `mean(posterior)` = 0.42

Computation of the posterior median: `median(posterior)` = 0.41

Computation of the posterior mode:

- computationally unstable
- rarely sensible

Computation of uncertainty intervals in R

Computation of 95%-CIs:

```
quantile(posterior, probs = c(0.025, 0.975))
```

```
## 0.17 0.69
```

Computation of HDIs may be computationally unstable

Visualization of Uncertainty Intervals

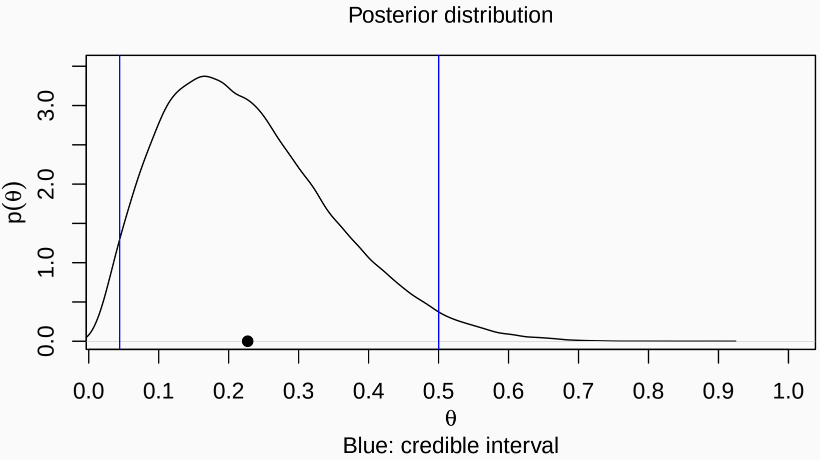

Region of Practical Equivalence (ROPE)

- Define a region that is thought to be practically equivalent to the value being tested.
- Extends the null hypothesis to an interval
- For instance $ROPE = [d = -0.1, d = 0.1]$ in intervention studies

Three possible outcomes of the hypothesis:

- ROPE and UI do not intersect: Reject the null hypothesis
- UI is completely within ROPE: Accept the null hypothesis
- Else: Evidence is inconclusive

Visualization of ROPEs: Reject the Null Hypothesis

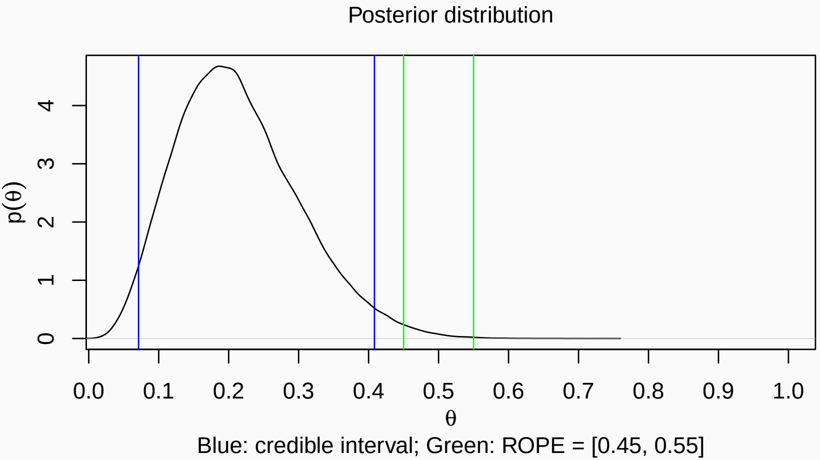

Visualization of ROPEs: Accept the Null Hypothesis

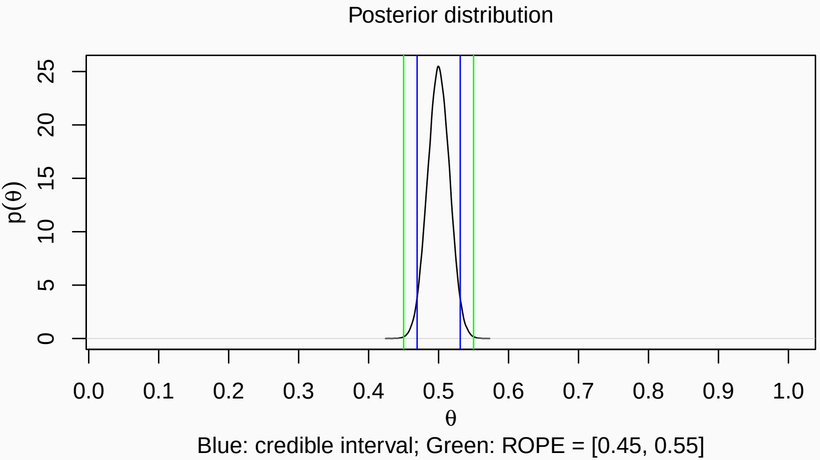

Visualization of ROPEs: Inconclusive

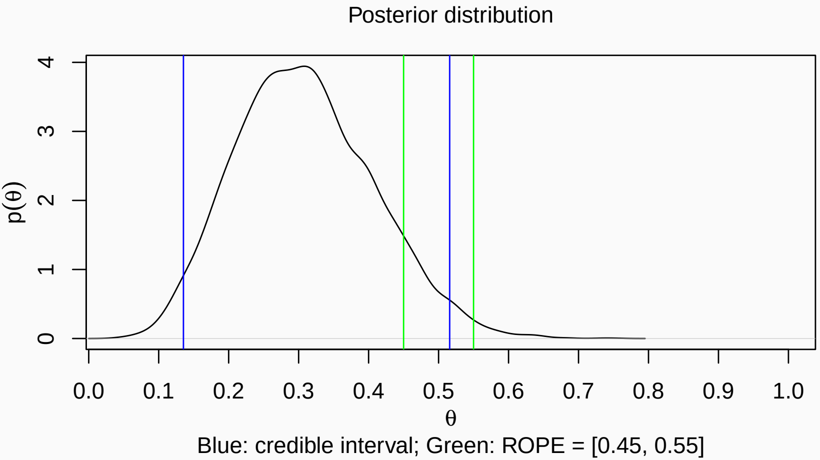

Posterior Probabilities

Posterior Probabilities

Applicable to interval hypotheses – examples:

If $H : \theta > 0.5$ then

$$P(H) = P(\theta > 0.5) = \frac{1}{S} \sum_{s=1}^{S} 1_{>0.5}(\theta_s)$$

If $H : \theta \in [0.4, 0.6]$ then

$$P(H) = P(\theta \in [0.4, 0.6]) = \frac{1}{S} \sum_{s=1}^{S} 1_{[0.4,0.6]}(\theta_s)$$

- S = Number of posterior samples
- θ_s = Posterior sample number s of parameter θ
- $1_I(x) = 1$ if x is in the interval I and $1_I(x) = 0$ otherwise

Bayes Factors

Marginal Likelihoods of Models

Marginal likelihood of model M:

$$p(y|M) = \int p(y|\theta, M) p(\theta|M) d\theta$$

- This is the probability of the data given the model
- Can be considered as a measure of model fit
- Depends heavily on the prior $p(\theta|M)$

Bayes Factors

- Used to compare two models M_1 and M_2:

$$BF_{12} = \frac{p(y|M_1)}{p(y|M_2)}$$

- Closely related to the posterior Odds:

$$\frac{p(M_1|y)}{p(M_2|y)} = \frac{p(M_1)}{p(M_2)} BF_{12}$$

- $p(M_1)$ and $p(M_2)$ are the prior probabilites of the models M_1 and M_2
- Usually $p(M_1) = p(M_2) = 1/2$

The Savage-Dickey Ratio

- Computation of the evidence is complicated and so is the computation of the BF
- Assume we are testing $M_1 : \theta = \theta_0$ against $M_2 : \theta \neq \theta_0$
- (We could use the word 'hypothesis' instead of 'models')
- Then the Bayes factor can be computed as

$$BF_{12} = \frac{p(\theta_0 | y, M_2)}{p(\theta_0 | M_2)}$$

Bayes Factors: Visualization

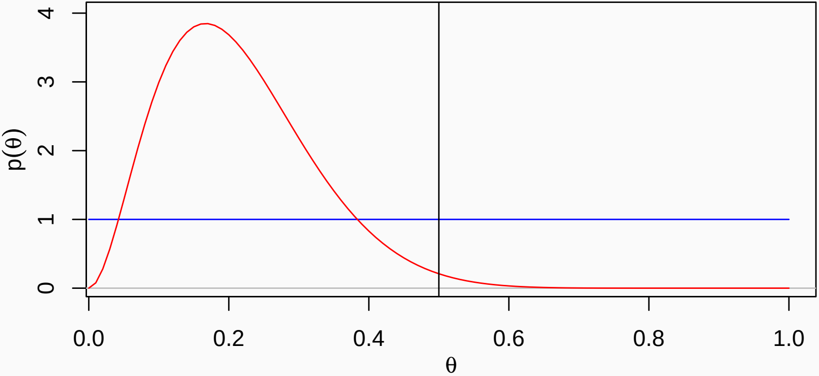

Blue: Prior; Red: Posterior

Bayes Factors: Example

- We assume a flat prior $\mathit{beta}(1, 1)$
- We observed $y = 2$ for $N = 12$.
- The resulting posterior (computed analytically) is $\mathit{beta}(3, 11)$
- We are interested in the BF at $\theta_0 = 0.5$

```
dbeta(0.5, 3, 11) / dbeta(0.5, 1, 1)
```

```
## [1] 0.2094727
```

Bayes factor: Influence of Priors

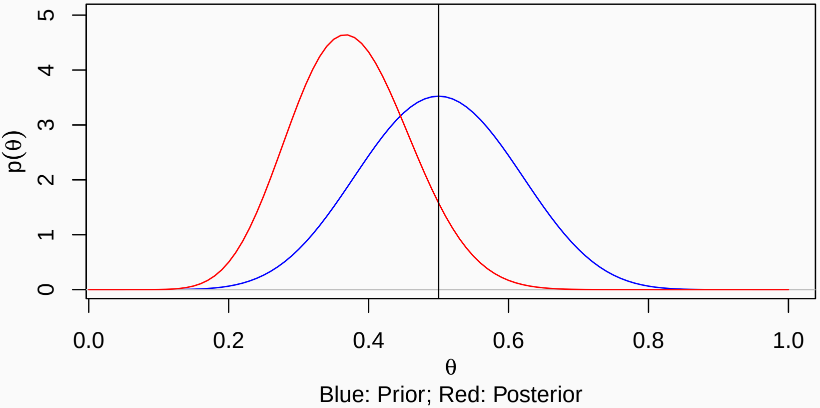

Bayes Factors: Confirming the Null Hypothesis

- We again assume a flat prior $beta(1, 1)$
- We observed $y = 49$ for $N = 100$.
- The resulting posterior (computed analytically) is $beta(50, 52)$
- We are interested in the BF at $\theta_0 = 0.5$

```
dbeta(0.5, 50, 52) / dbeta(0.5, 1, 1)
```

```
## [1] 7.880895
```

Bayes Factors: Confirming the Null Hypothesis