


Objective Priors for Measures of Evidence

Laura Ventura

(joint work with Elena Bortolato and Monica Musio)

Department of Statistical Sciences
University of Padova, Italy
ventura@stat.unipd.it
homes.stat.unipd.it/lauraventura



Today's -- traditional -- menu

- **Starter:** a taste of background on a posterior distribution & a parameter of interest & a precise null hypothesis
- **First course:** two interesting Bayesian measures of evidence
- **Main course:** two objective matching (median and strong) priors (it depends on your tastes)
- **Dessert:** two little tasters of examples, so as not to overeat before the coffee break
- **Coffee:** conclusions & discussion

It is argued that analytic approximations still have an important role to play in Bayesian inference.


Starter:
a taste of background

The parametric framework

- Random sample $y = (y_1, \dots, y_n)$ from $f(y; \theta)$, d -dimensional parameter θ .
- Let $\theta = (\psi, \lambda)$, with ψ scalar and λ $(d - 1)$ -dimensional nuisance parameters.
- Given a prior density $\pi(\theta) = \pi(\psi, \lambda)$, Bayesian inference for ψ is based on the **marginal posterior density**

$$\pi_m(\psi|y) \propto \int \pi(\psi, \lambda) \mathcal{L}(\psi, \lambda) d\lambda$$

with $\mathcal{L}(\psi, \lambda)$ likelihood function for (ψ, λ) .



The framework of interest

- It is of interest to test the precise (or sharp) null hypothesis

$$\mathcal{H}_0 : \psi = \psi_0 \quad \text{against} \quad \mathcal{H}_1 : \psi \neq \psi_0$$

- Possible examples: ψ is the stress-strength reliability and the null hypothesis is $\mathcal{H}_0 : \psi = 0.5$; ψ is a regression coefficient and the null hypothesis is $\mathcal{H}_0 : \psi = 0$; non-inferiority and superiority testing, with ψ difference of means, or odds ratios, or hazard ratios . . .
- The usual approach relies on the Bayes Factor (BF), which measures the ratio of posterior to prior odds in favor of $\mathcal{H}_0$. But:
 - ▶ undetermined when using improper priors
 - ▶ Jeffreys-Lindley paradox
 - ▶ lacks calibration (its finite sampling distribution is unknown and may depend on nuisance parameters)

The goal

- To overcome these limitations, recent alternative Bayesian measures of evidence for testing precise null hypotheses:
 1. the Full Bayesian Significance Test (FBST) or e -value
 2. the Bayesian Discrepancy measure (BDM)
- We mix these measures of evidence with suitable objective matching priors (to match desirable frequentist properties):
 1. median matching prior, that maintains the invariance of the posterior mode and of a HPD
 2. strong matching prior, so that a frequentist one-sided p -value coincides with a Bayesian posterior survivor probability
- In problems with a large number of nuisance parameters, matching priors have a role to play since they eliminate the need to elicitate on the nuisance parameter.

*First course:
two Bayesian measures of evidence*


1. Full Bayesian Significance Test

- Pereira & Stern (1999) provide a measure of evidence for the **full Bayesian Significance test (FBST)** in favour of $\mathcal{H}_0$, based on a specific loss function. Thus the decision made is the action that minimizes the corresponding posterior risk.
- In practice, from $\pi_m(\psi|y)$, given the set $\mathcal{T}_y(\psi) = \{\psi : \pi_m(\psi|y) \geq \pi_m(\psi_0|y)\}$, the FBST in favour of $\mathcal{H}_0$ is

$$p_\pi = 1 - \mathcal{P}_\pi(\psi \in \mathcal{T}_y(\psi))$$

where $\mathcal{P}_\pi(\cdot)$ denotes posterior probability.

- The null hypothesis $\mathcal{H}_0$ is accepted whenever p_π is large enough.



2. Bayesian Discrepancy Measure

- The BDM (Bertolino et al 2024) is defined as

$$\delta_{\mathcal{H}} = 1 - 2\Delta = 1 - 2 \min \left\{ \int_{-\infty}^{\psi_0} \pi_m(\psi|y) d\psi, 1 - \int_{-\infty}^{\psi_0} \pi_m(\psi|y) d\psi \right\}$$

where Δ can be interpreted as the posterior probability of a "tail" event concerning only $\mathcal{H}_0$, assessing about how "central" the hypothesis $\mathcal{H}_0$ is.

- Alternative definition for $\delta_{\mathcal{H}}$:


Let m_1 be the posterior median and consider the interval defined as [A]

$I_E = (-\infty, \psi_0)$ if $\psi_0 < m_1$ or as [B] $I_E = (\psi_0, +\infty)$ if $m_1 < \psi_0$.

Then, the BDM of the hypothesis $\mathcal{H}_0$ can be computed as

$$\delta_{\mathcal{H}} = 1 - 2 \mathcal{P}(\psi \in I_E | y) = 1 - 2 \int_{I_E} \pi_m(\psi|y) d\psi$$

2. Bayesian Discrepancy Measure (cont.)



- $1 - \delta_{\mathcal{H}}$ gives the posterior probability of an equi-tailed credible interval for ψ that contains ψ_0 .
- High values of $\delta_{\mathcal{H}}$ indicate strong evidence against $\mathcal{H}_0$, while small values suggest data consistency with $\mathcal{H}_0$.
- When $\mathcal{H}_0$ is false $\delta_{\mathcal{H}}$ is more likely to approach 1 than 0. Thresholds can be set to interpret $\delta_{\mathcal{H}}$.
- The BDM is invariant under invertible monotonic reparametrizations.

FBST vs BDM

FBST	BDM
$1 - p_\pi$ represents the posterior probability of a HPD CI for ψ	$\delta_{\mathcal{H}}$ represents the posterior probability of an equi-tailed CI for ψ
asymptotic consistency	asymptotic consistency
p_π is not invariant while the e -value is invariant (but not with respect marginalizations)	$\delta_{\mathcal{H}}$ is invariant
simple extension for $d > 1$	see you in Padova! January 7-9, 2026 <code>skew2026.stat.unipd.it</code>

- The practical computation of the FBST and BDM requires numerical integration and the elicitation on (ψ, λ) .
- Numerical integration can be tackled by resorting to asymptotic approximations.

First-order approximation

- A first-order approximation for $\delta_{\mathcal{H}}$ or $1 - p_{\pi}$ is simply

$$2 \Phi \left(\left| \frac{\psi_0 - \hat{\psi}}{\sqrt{j_p(\hat{\psi})^{-1}}} \right| \right) - 1$$

with $\hat{\psi}$ MLE of ψ and $j_p(\psi) = -\ell_p^{(2)}(\psi)$ profile observed information.

- To first-order, $\delta_{\mathcal{H}}$ agrees numerically with $1 - p$ -value based on the profile Wald statistic and also with $1 - p_{\pi}$.
- In practice, this approximation may be inaccurate, in particular for a small sample size or large number of nuisance parameters, since it fails to account for potential posterior asymmetry and skewness.

Beyond Gaussian: Higher-order asymptotic approximations

- Asymptotic arguments are widely used in Bayesian inference, based on developments of so-called higher-order asymptotics (Ventura & Reid 2014).
- They provides very accurate approximations to posterior distributions and to summary quantities of interest, including tail areas and credible regions. Moreover, they are useful for the **derivation of matching priors**.
- The key tool is the **Laplace approximation**

$$\pi_m(\psi|y) \doteq \frac{1}{\sqrt{2\pi}} |j_p(\hat{\psi})|^{1/2} \exp\{\ell_p(\psi) - \ell_p(\hat{\psi})\} \frac{|j_{\lambda\lambda}(\hat{\psi}, \hat{\lambda})|^{1/2}}{|j_{\lambda\lambda}(\psi, \hat{\lambda}_\psi)|^{1/2}} \frac{\pi(\psi, \hat{\lambda}_\psi)}{\pi(\hat{\psi}, \hat{\lambda})}$$

where $j_{\lambda\lambda}(\psi, \lambda)$ is the (λ, λ) -block of the observed information from the full log-likelihood and $\doteq$ indicates that accuracy is of order $O(n^{-3/2})$.

Beyond Gaussian (cont.)

- Starting from the Laplace approximation, the $O(n^{-3/2})$ approximation to the tail area probability is

$$\int_{\psi_0}^{\infty} \pi_m(\psi|y) d\psi \doteq \Phi(r_B^*(\psi_0))$$

with $r_p(\psi) = \text{sign}(\hat{\psi} - \psi)[2(\ell_p(\hat{\psi}) - \ell_p(\psi))]^{1/2}$ profile likelihood root and

$$r_B^*(\psi) = r_p(\psi) + \frac{1}{r_p(\psi)} \log \frac{q_B(\psi)}{r_p(\psi)},$$

$$q_B(\psi) = \ell'_p(\psi) |j_p(\hat{\psi})|^{-1/2} \frac{|j_{\lambda\lambda}(\psi, \hat{\lambda}_\psi)|^{1/2}}{|j_{\lambda\lambda}(\hat{\psi}, \hat{\lambda})|^{1/2}} \frac{\pi(\hat{\psi}, \hat{\lambda})}{\pi(\psi, \hat{\lambda}_\psi)}.$$

- Note that using $r_B^*(\psi)$ an $(1 - \alpha)$ equi-tailed credible interval for ψ can be computed as $CI = \{\psi : |r_B^*(\psi)| \leq z_{1-\alpha/2}\}$, with $z_{1-\alpha/2}$ quantile of the standard normal distribution, and in practice it can reflect asymmetries of the posterior.
- Moreover, the marginal posterior median can be computed as the solution in ψ of the estimating equation $r_B^*(\psi) = 0$.

Beyond Gaussian (cont.)

- Using the tail area approximation, third-order approximations for the FBST and the BDM can be derived.
- For the FBST, assuming without loss of generality that ψ_o is smaller than the MAP, we have (Cabras et al 2015)

$$p_\pi \doteq 1 - \Phi(r_B^*(\psi_o)) + \Phi(r_B^*(\psi_o^*))$$

with ψ_o^* such that $\pi_m(\psi_o^*|y) = \pi_m(\psi_o|y)$.

- For the BDM we have (Bortolato et al 2025)

$$\delta_{\mathcal{H}} \doteq 2\Phi(|r_B^*(\psi_o)|) - 1$$

- When $\pi_m(\psi|y)$ is symmetric we have

$$p_\pi \doteq 2(1 - \Phi(r_B^*(\psi_o))) \doteq 1 - \delta_{\mathcal{H}}$$

- Note that these approximations do not call for any condition on the prior $\pi(\psi, \lambda)$.

*Main course:
the mix of matching priors with $FBIT$ and BDM*

Matching priors

- In the presence of nuisance parameters, starting from the Laplace's approximation it is possible to define a general posterior distribution for ψ of the form

$$\pi^*(\psi|y) \propto \pi^*(\psi) \mathcal{L}_p(\psi)$$

where $\pi^*(\psi)$ is a prior distribution on ψ only.

- Bayesian inference based on pseudo-likelihoods has been widely discussed and motivated in the literature.
- Especially when the dimension of λ is large, there are two advantages in using $\pi^*(\psi|y)$:
 1. the elicitation over λ is not necessary
 2. the computation of integrals is circumvented
- Here we focus on matching priors for ψ .
- Examples of matching priors are for posterior quantiles, for credible regions and for prediction.
- Here we consider median and strong matching priors.

1. Median matching prior for the FBST

- The FBST and the e -value in the presence of nuisance parameters are not invariant.
- We want to find a matching prior $\pi^*(\psi)$ that ensures the invariance of the MAP. As a consequence, the invariance extends to HPDs, as well as the FBST.
- To this end we focus on the median modification of the profile score equation, whose solution respects equivariance under monotone reparameterizations.
- The median modification of the profile score does not rely on finiteness of the MLE, thereby effectively preventing infinite estimates.
- Thus, for Bayesian inference with $\pi^*(\psi)$ we have a proper posterior also in the case of monotone likelihoods.

1. Median matching prior for the FBST (cont.)

- In practice, we impose that the MAP of $\pi^*(\psi|y)$ coincides with a refined version of the MLE, obtained as the solution of the median modified score function.
- The modified version of the profile score function is

$$t_p(\psi) = \ell'_p(\psi) + m(\psi, \hat{\lambda}_\psi)$$

where $m(\psi, \lambda)$ is a suitable correction term of order $O(1)$, given by

$$m(\psi, \lambda) = -\kappa_{1\psi} + \frac{\kappa_{3\psi}}{6\kappa_{2\psi}}$$

where $\kappa_{1\psi}$, $\kappa_{2\psi}$ and $\kappa_{3\psi}$ are the first three cumulants of $\ell'_p(\psi)$.

- For the estimator $\tilde{\psi}_p$, defined as the solution of $t_p(\psi) = 0$, parameterization equivariance holds under interest respecting reparameterizations.

1. Median matching prior for the FBST (cont.)

- We look for a matching prior $\pi^*(\psi)$ such that

$$\frac{\partial \log \pi^*(\psi)}{\partial \psi} = m(\psi, \hat{\lambda}_\psi)$$

Thus $\pi^*(\psi)$ is known through its first derivative.

- The posterior based on the median matching prior can be written as

$$\pi^*(\psi|y) \propto \exp \left(\ell_p(\psi) + \int m(\psi, \hat{\lambda}_\psi) d\psi \right)$$

- The invariant FBST is

$$p_\pi^* = 1 - \mathcal{P}_\pi(\psi \in \mathcal{T}_y^*(\psi))$$

where $\mathcal{T}_y^*(\psi) = \{\psi : \pi^*(\psi|y) \geq \pi^*(\psi_0|y)\}$.

- Computation of p_π^* with simulation approaches: ABC with $\tilde{\psi}_p$ or $t_p(\psi)$ as summary statistics (see Ruli et al 2020), or Manifold MCMC methods (see Bortolato & Ventura 2024).

2. Strong matching priors for BDM

- A so-called strong-matching prior ensures that a frequentist one-sided p -value coincides with a Bayesian posterior survivor probability to a high degree of approximation, in the marginal posterior density.
- When using a strong matching prior, the marginal posterior density can be written as

$$\pi_m^m(\psi|y) \propto \exp\left(-\frac{1}{2}r_p^*(\psi)^2\right) \left|\frac{s_p(\psi)}{r_p(\psi)}\right|$$

where $s_p(\psi) = \ell_p^{(1)}(\psi)/j_p(\hat{\psi})^{1/2}$ is the profile score statistic and $r_p^*(\psi) = r_p(\psi) + \frac{1}{r_p(\psi)} \log \frac{q_p(\psi)}{r_p(\psi)}$ is the frequentist modified profile likelihood root, with $q_p(\psi)$ suitably defined correction term.

2. Strong matching priors for BDM (cont.)

- The tail area can be approximated to third-order as

$$\mathcal{P}_m(\psi \geq \psi_o | y) \doteq \Phi(r_p^*(\psi_o))$$

- A remarkable advantage of $\pi_m^m(\psi | y)$ and $\Phi(r_p^*(\psi_o))$ is that the expressions automatically include the matching prior, without requiring its explicit computation.
- With a strong matching prior, an asymptotic equi-tailed credible interval for ψ can be computed as $CI = \{\psi : |r_p^*(\psi)| \leq z_{1-\alpha/2}\}$, i.e., as a confidence interval for ψ based on $r_p^*(\psi)$, and the posterior median can be computed as the solution of $r_p^*(\psi) = 0$.

2. Strong matching priors for BDM (cont.)

- A third-order approximation of the BDM is

$$\delta_{\mathcal{H}}^* \doteq 1 - 2 \min\{\Phi(r_p^*(\psi_0)), 1 - \Phi(r_p^*(\psi_0))\} = 2\Phi(|r_p^*(\psi_0)|) - 1$$

- In this case

$$\delta_{\mathcal{H}}^* = 1 - p_r^*$$

where p_r^* is the p -value based on $r_p^*(\psi)$.

- Thus, when using strong matching priors there is an agreement between Bayesian and frequentist testing hypothesis, point and interval estimation.
- From a practical point of view, the computation of $\delta_{\mathcal{H}}^*$ can be easily performed in practical problems using the `likelihoodAsy` package of the statistical software R.

Dessert:
two little tasters of examples

1. FBST: Multivariate regression model

- Regression model


$$\mathcal{Y}_{ij} = \beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + \epsilon_{ij}, \quad i = 1, \dots, n, \quad j = 1, 2,$$

with $\epsilon_i \sim \mathcal{N}_2(0, \Sigma)$, $\Sigma = \sigma^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$ positive definite matrix.

- The null hypothesis is $\mathcal{H}_0 : \rho = 0.9$, with $\rho_0 = 0.95$ and $n = 20$.
- Comparison of the posterior distributions based on the median matching prior, the predictive matching prior and an inverse-Wishart prior for the covariance matrix Σ with one degree of freedom, identity position, uniform prior on the regression parameters.
- The expressions of the modified profile estimating functions are obtained from Bortolato et al (2023). Hence, the Manifold MCMC method can be used to obtain the implied posteriors, whose approximation is comparable to that of any MCMC sampler. We used 20000 iterations.

1. FBST: Multivariate regression model (cont.)

- The FBST are 0.25 with the median matching prior, 0.36 for the predictive matching prior and 0.60 with the inverse-Wishart prior.
- The BDM provides $1 - \delta_{fi} = 0.12$ with the median matching prior, 0.04 for the predictive matching prior and 0.87 with the inverse-Wishart prior.
- The only invariant results are those with the median matching prior.



2. BDM: Heteroscedastic one-way random effects model

- Analysis of data from inter-laboratory study with the one-way random effects model with heteroscedastic error variance.
- There are m laboratories, with n_j observations at the j -th laboratory, for $j = 1, \dots, m$. The model is


$$y_{ij} = \mu + b_j + \varepsilon_{ij},$$

where y_{ij} denotes the i -th observation at the j -th laboratory, and b_j and ε_{ij} are independent random variables with distribution $b_j \sim \mathcal{N}(0, \sigma^2)$ and $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma_j^2)$, respectively.

- Typically, the parameter of interest is the consensus mean μ , which is also the mean of y_{ij} , $i = 1, \dots, n_j$ and $j = 1, \dots, m$.
- The remaining $(m + 1)$ parameters of the model, i.e., within-laboratory variances $(\sigma_1^2, \dots, \sigma_m^2)$ and the between laboratory variability σ^2 , are nuisance parameters.

Example 1: Heteroscedastic one-way random effects model

Data on the Ki-67 protein on adrenocortical tumors, coming from an inter-laboratory study. It is of interest to carry out inference on the mean of the Ki-67 level (on logarithmic scale).



For $\mathcal{H}_0 : \mu = 0$ vs $\mathcal{H}_1 : \mu \neq 0$, we obtain $\delta_{\mathcal{H}}^* = 1 - p_r^* = 1 - 0.062 = 0.938$.

In this case, FBST also coincides, up to four decimal places, to $1 - \delta_{\mathcal{H}}$, in view of the symmetry of the marginal posterior of μ .

The first-order approximation gives $\delta_{\mathcal{H}} = 1 - p_w = 1 - 0.002 = 0.998$.

Hence, first-order results suggest that there is a strong evidence against $\mathcal{H}_0$.

Coffee:
conclusions & discussion

Coffee: conclusions & discussion

- Approximate analytical tools have still a role to play in the modern era of Bayesian statistics, where high computational power allows the use of stochastic simulation techniques to obtain exact (i.e., simulation consistent) answers.
- In problems with a large number of nuisance parameters, approximate Bayesian computations provide important quantities of the posterior distribution with very little computational effort, in a fraction of the time required for a full simulation approach. Moreover, sensitivity and influence analyses may also be carried out quickly within this framework.
- Advantages of the proposed approximations are that no elicitation on the nuisance parameters is required.
- Although the approximations described in this paper are derived from asymptotic considerations, they perform extremely well in moderate or even small sample situations.
- When using objective Bayesian procedures based on strong matching priors and higher-order asymptotics, there is an agreement between Bayesian and frequentist point and interval estimation, and also in the significance measure BDM. This is not true in general with the FSBT.

Some (personal) references

- Bortolato, Bertolino, Musio, Ventura (2025) Bayesian discrepancy measure: Higher-order and skewed approximations, *Entropy*
- Bertolino, Manca, Musio, Racugno, Ventura (2024) A new Bayesian discrepancy measure, *SMA*
- Bortolato, Ventura (2024) Objective priors for invariant e -values in the presence of nuisance parameters, *Entropy*
- Bortolato, Kenne Pagui (2023) Bias reduction and robustness in gaussian longitudinal data analysis, *JSCS*
- Ruli, Ventura (2021) Can Bayesian, confidence distribution and frequentist inference agree? *SMA*
- Ruli, Sartori, Ventura (2020) Robust approximate Bayesian inference, *JSPI*
- Cabras, Racugno, Ventura (2015) Higher-order asymptotic computation of Bayesian significance tests for precise null hypotheses in the presence of nuisance parameters, *JSCS*
- Ventura, Reid (2014) Approximate Bayesian computation with modified log-likelihood ratios, *Metron*
- Ventura, Sartori, Racugno (2013), Objective Bayesian higher-order asymptotics in models with nuisance parameters, *CSDA*

Thanks for your time and Hope you enjoyed your meal!