Potential prior-data conflict when using informative priors

Timothy Mutsvari (Arlenda)
Rosalind Walley (UCB)

Bayes-pharma, Leuven
17 – 20 May 2016

Inspired by patients.
Driven by science.

Arlenda
Advance your business with statistics

Outline

- Introduction to active comparator studies
- Prior predictive distribution
- Case study
- Assessment of power of prior data-conflict test
- Comparison of testing approach with a mixture approach
- Case study results
- Conclusions

Introduction

- Active comparator: a marketed drug presumed to have beneficial effects
- Clinical trial setting: compare new trt vs active comparator (vs placebo):
  - new trt vs active comparator - benchmarking new drug
- High attrition in Phase II studies hence should be executed as quickly and economically as possible
- **Features:**
- Difficult to do a head-to-head if standard of care is a drug cocktail
- Active comparator studies require large sample sizes (usually done in phase 3)
- Possibility of published data or in-house data on performance of marketed drug ➔ informative prior ➔ Bayesian methods
- Risk that an apparent mismatch is observed between the prior and the data

Prior predictive distribution

- Before a new study is run, the uncertainty in the parameter of interest, θ say, is represented by the prior distribution, $p(\theta)$
- The unconditional distribution of data summary e.g. response rate is represented by the "prior predictive distribution", obtained by averaging over $p(\theta)$, to get:

$$p(x) = \int p(x|\theta)p(\theta)d\theta$$

- **Example**
  - Given a beta prior $p(\theta) \sim \text{Beta}(\alpha, \beta)$ and binomial likelihood $r \sim \text{Bin}(n, \theta)$
    - The prior predictive is a beta-binomial: $p(r) \sim \binom{n}{r} \frac{B(r+\alpha, n-r+\beta)}{B(\alpha,\beta)}$
  - Normal likelihood $N(\theta, \sigma^2)$ and normal prior $p(\theta) \sim N(\theta_0, \sigma_0^2)$
    - The prior predictive is a normal: $p(\tilde{y}) \sim N(\theta_0, \sigma_0^2 + \sigma^2/n)$

- Prior predictive distribution provides a pre-study prediction of the data based on the selected prior, hence plausible to use it in the assessment of the compatibility of the data and the prior

Is the prior suitable?

- Assess by comparing the observed study mean with the “prior predictive distribution”
- Declare prior-data conflict if the observed study mean is in the extremes of the prior predictive distribution (5% level)

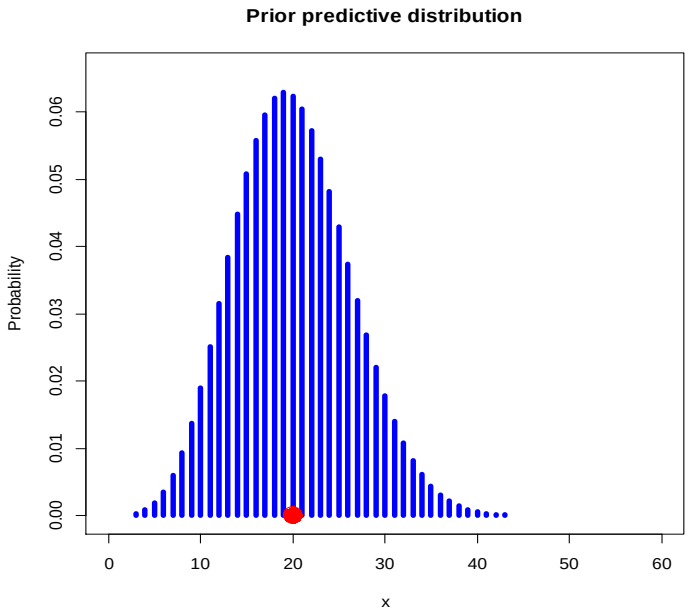

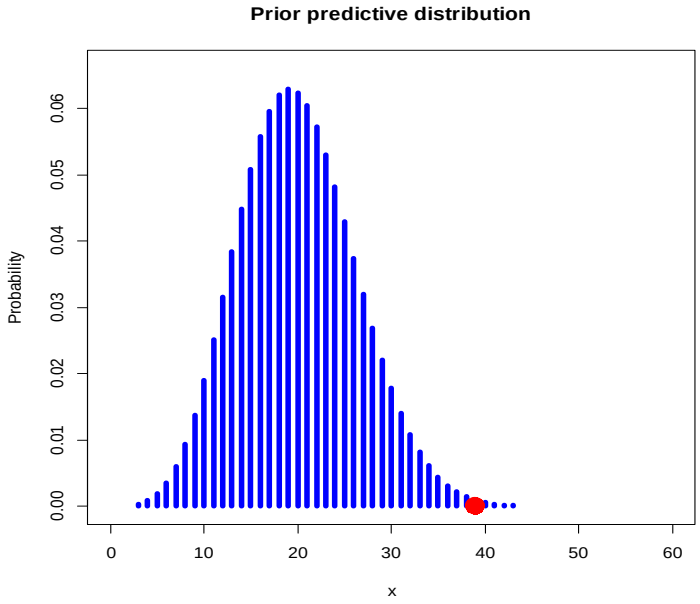

- Assumption: the statistical model used is appropriate as an unsuitable model could also trigger an extreme test statistic

Case Study: UK Medical Research Thiotepa Study

- Study in superficial bladder cancer
- Efficacy end-point: time to first recurrence
- We focus on the comparison of 2 of the 3 treatment arms:
  - control group, n=131
  - immediate installation of thiotepa 30mg, n=126

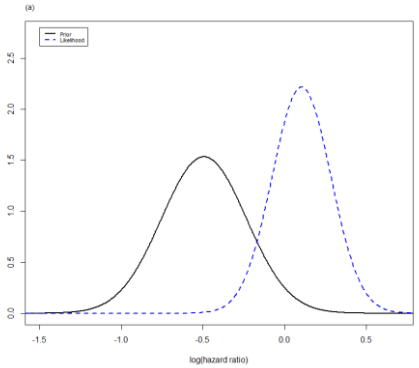

- Concern expressed about apparent mismatch between prior and data. However, prior-data conflict p-value is **0.06**. Raises two issues:
  - How “powerful” is the prior data conflict test?
  - What should you do if you observe apparent prior data conflict?

Simulation Methodology

Assessment of power for prior-data conflict test

- For normal likelihood prior with known variance cut-offs were obtained analytically
- Sample sizes typically encountered in phase II studies (n=10 to 100)
- Number of simulations = 10 000 a range of values of mu [-2, 2]
- Selected prior distribution: $p(\theta)$ ~ N(0, 0.1)
- Cut-offs: 95%

Simulation Results

Assessment of power for prior-data conflict test

Using 5% cut-offs

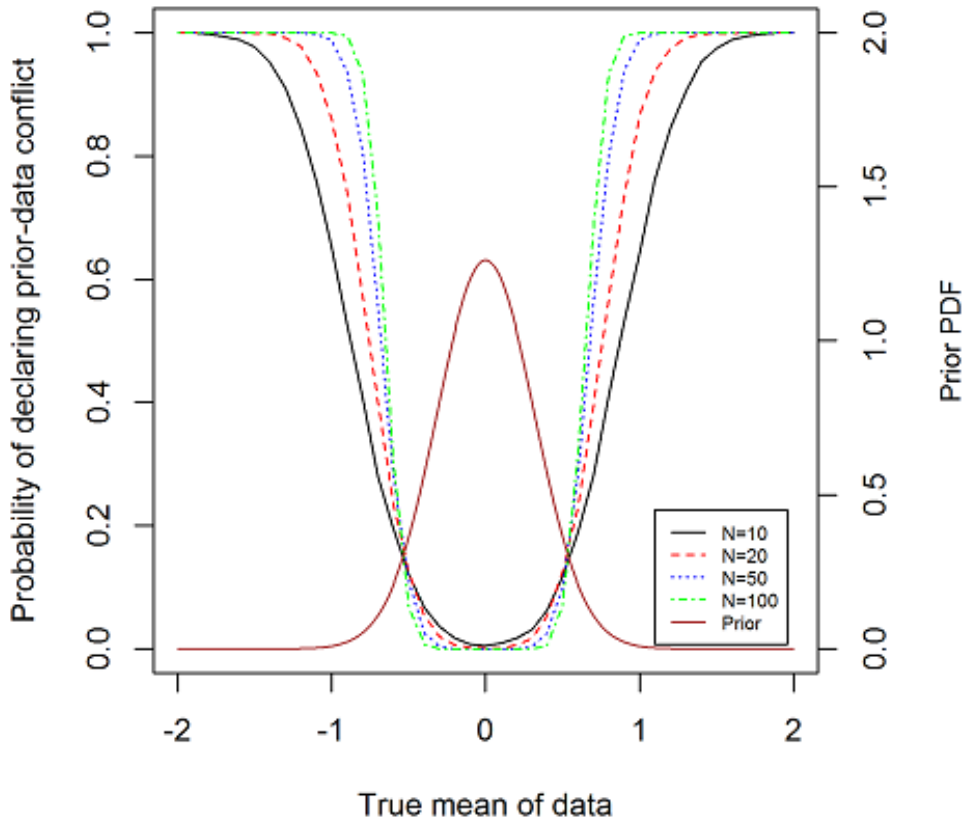

Comparison of Two approaches

Prior-data conflict testing approach

- Construct prior predictive distribution using the informative prior
- Prior-data conflict is declared if the observed mean lies outside the (1-alpha)100 % cut-off points of the prior predictive distribution
- If prior data conflict **is not** declared, continue with informative prior
- If prior data conflict **is** declared, change to an uninformative prior

Comparison of Two Approaches

Mixture prior

- A mixture/robust prior with two components
  - 1st component represents the precise information about the active comparator with high prior probability
  - 2nd component represents a comparatively vague component, illustrating a lack of certainty of knowledge about the parameter with a small prior probability
- For the normal case, if $p_0(\mu)$ and $p_1(\mu)$ represent the precise and vague prior distributions for the treatment mean, respectively., the mixture prior distribution is then

$$p(\mu) = \xi . f(y;\ \mu_0, \sigma_0^2) + (1 - \xi) . f(y;\ \mu_1, \sigma_1^2)$$

- Concept: In presence of prior-data conflict use diffuse prior, otherwise use precise

Simulation Methodology

Head to head comparison of the testing and mixture approaches

Data:

- Normal with mean 0, and known variance, σ^2, sample size n=10 (or 100)

Prior for the mean:

- Normal prior, or mixture of two normal priors with the same mean
- The precise prior → effective sample size of n or 5n. St. dev.= σ_0
- The diffuse prior → effective sample size of 0.01n
- Discrepancy between the true mean of the data and the mean of the priors was set to be 0, $2\sigma_0$ or $4\sigma_0$

Five analyses:

- Precise prior alone
- Diffuse prior alone
- Prior-data conflict testing approach
- Mixture approach with prior probability ξ=0.7 for the precise prior
- Mixture approach with prior probability ξ=0.9 for the precise prior

Simulation Results

Compare the testing and mixture approaches

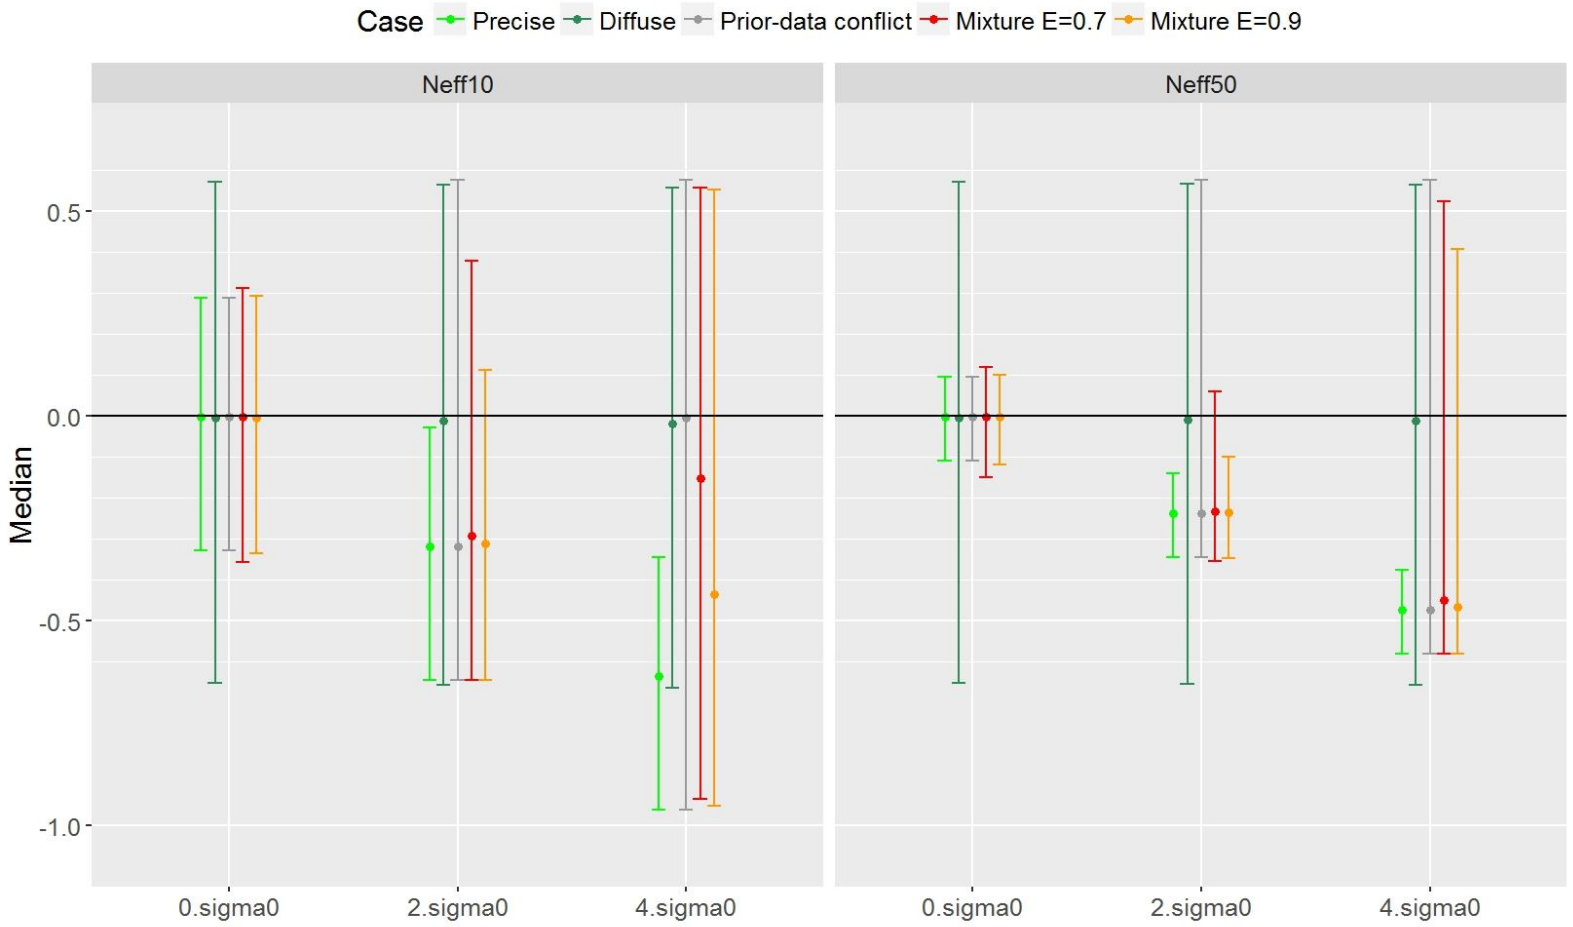

Sigma0 = sd of the precise prior

Simulation Results

Compare the testing and mixture approaches

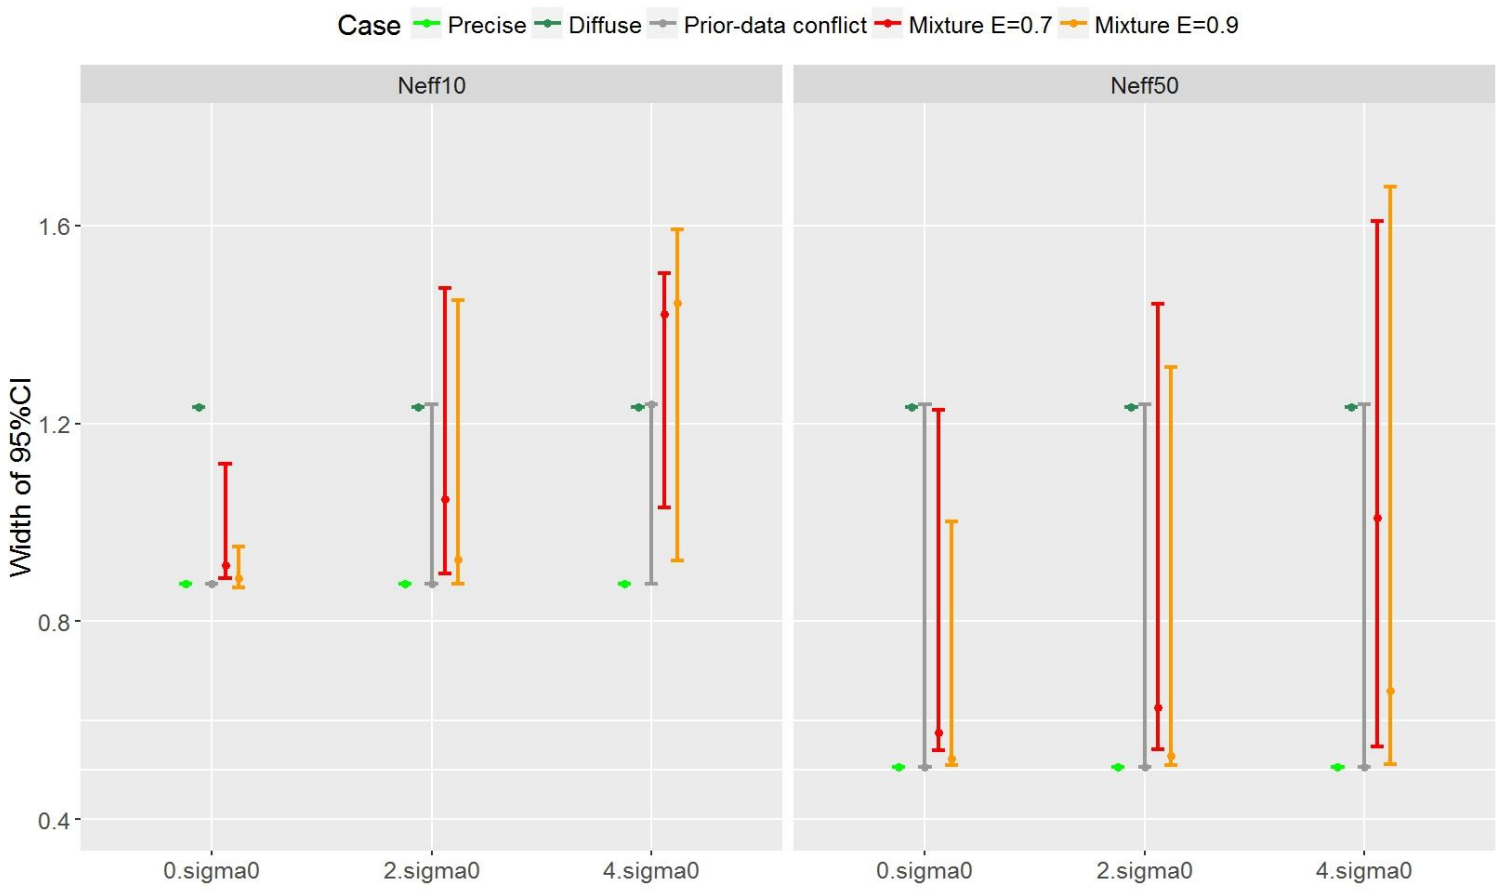

Sigma0 = sd of the precise prior

Simulation Results

Compare the testing and mixture approaches

Effective sample size	10 (n)			50 (5n)		
Discrepancy (as a multiple of σ_0)	0	2	4	0	2	4
P(prior data conflict declared)	0.00	0.22	0.89	0.03	0.10	0.35
P(precise prior | ξ=0.7)	0.94	0.86	0.25	0.95	0.94	0.85
P(precise prior | ξ=0.9)	0.98	0.96	0.56	0.99	0.98	0.96

Case Study Results

UK Medical Research Thiotepa Study in superficial bladder cancer

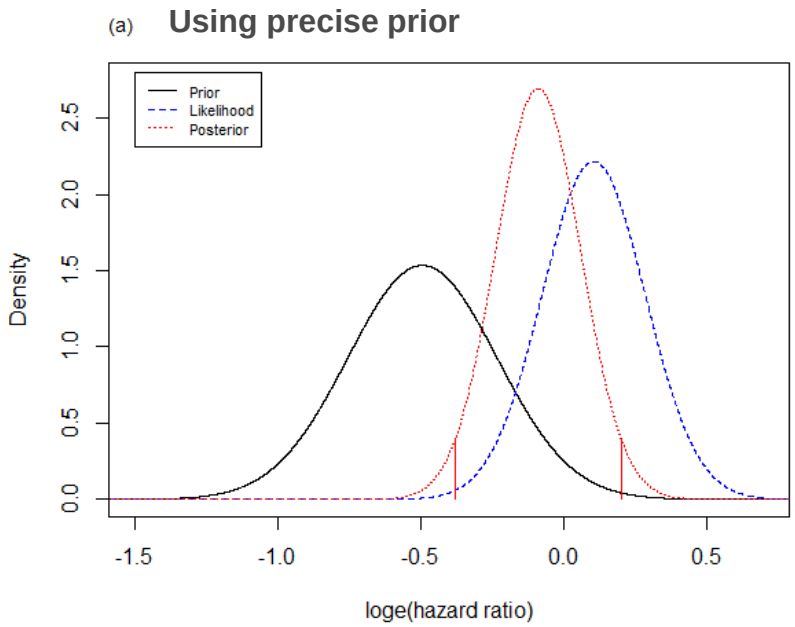

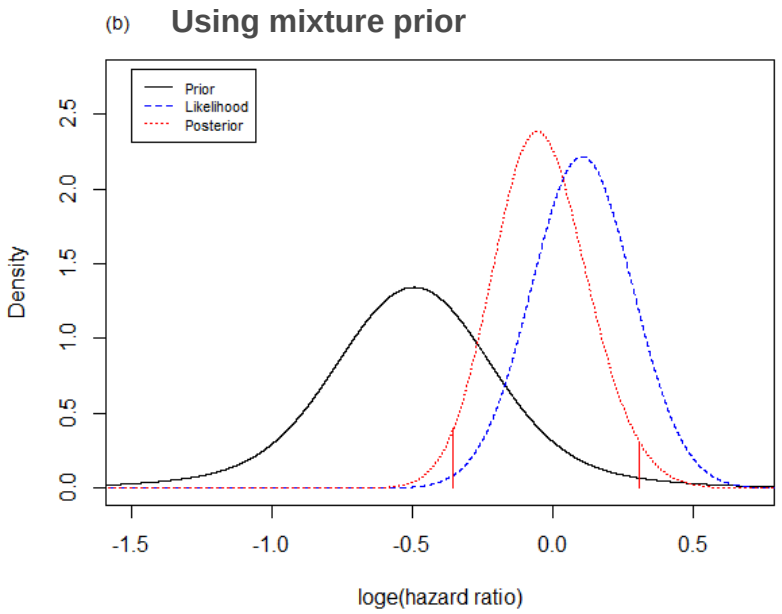

Posterior = $N(-0.09, 0.15^2)$

Posterior = $0.61*N(-0.09, 0.15^2) + 0.39*N(0.04, 0.17^2)$

Conclusions

- Whenever an informative prior is used, prior-data conflict is a potential issue
- Prior data conflict test is a way to assess this, but:
  - May be underpowered for phase II studies
  - Not obvious what to do if prior data conflict is declared. We focused on reverting to a noninformative prior

- When the observed discrepancy is small:
  Prior-data conflict testing ≈ Mixture analysis ≈ standard analysis with informative prior
- When the observed discrepancy is larger both proposed methods have some attractive features over the standard analysis. For example, wider posterior credible intervals than when there is no observed discrepancy.
- Whichever approach is selected, it is vital that at the design stage the operating characteristics of analysis should be assessed and discussed with all stakeholders

References

Prior data conflict:

- *Evans & Moshonov (Bayesian Analysis 2006; **1**(4):893–914)*

Mixture models: Bolstat

- *John Wiley and Sons: New Jersey, 2007, pp. 261–270*

Simulation:

- *Mutsvari, Walley & Tytgat (Pharmaceutical Statistics 2015; **15**(1): 28-36)*

Thiotepa Study:

- *UK Medical Research (Journal of Urology 1994; **73**(6):632–638)*

Backup