Bayesian random-effects meta-analysis made simple

Christian Röver[1], Beat Neuenschwander[2], Simon Wandel[2], Tim Friede[1]

[1]Department of Medical Statistics, University Medical Center Göttingen, Göttingen, Germany

[2]Novartis Pharma AG, Basel, Switzerland

May 24, 2016

InSPiRe Innovative methodology for small populations research

This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement number FP HEALTH 2013-602144.

Overview

- Meta analysis
  - example
  - the random-effects model
  - the Bayesian approach
- the `bayesmeta` package
  - parameter estimation
  - prediction
- Conclusions

Meta analysis

Example

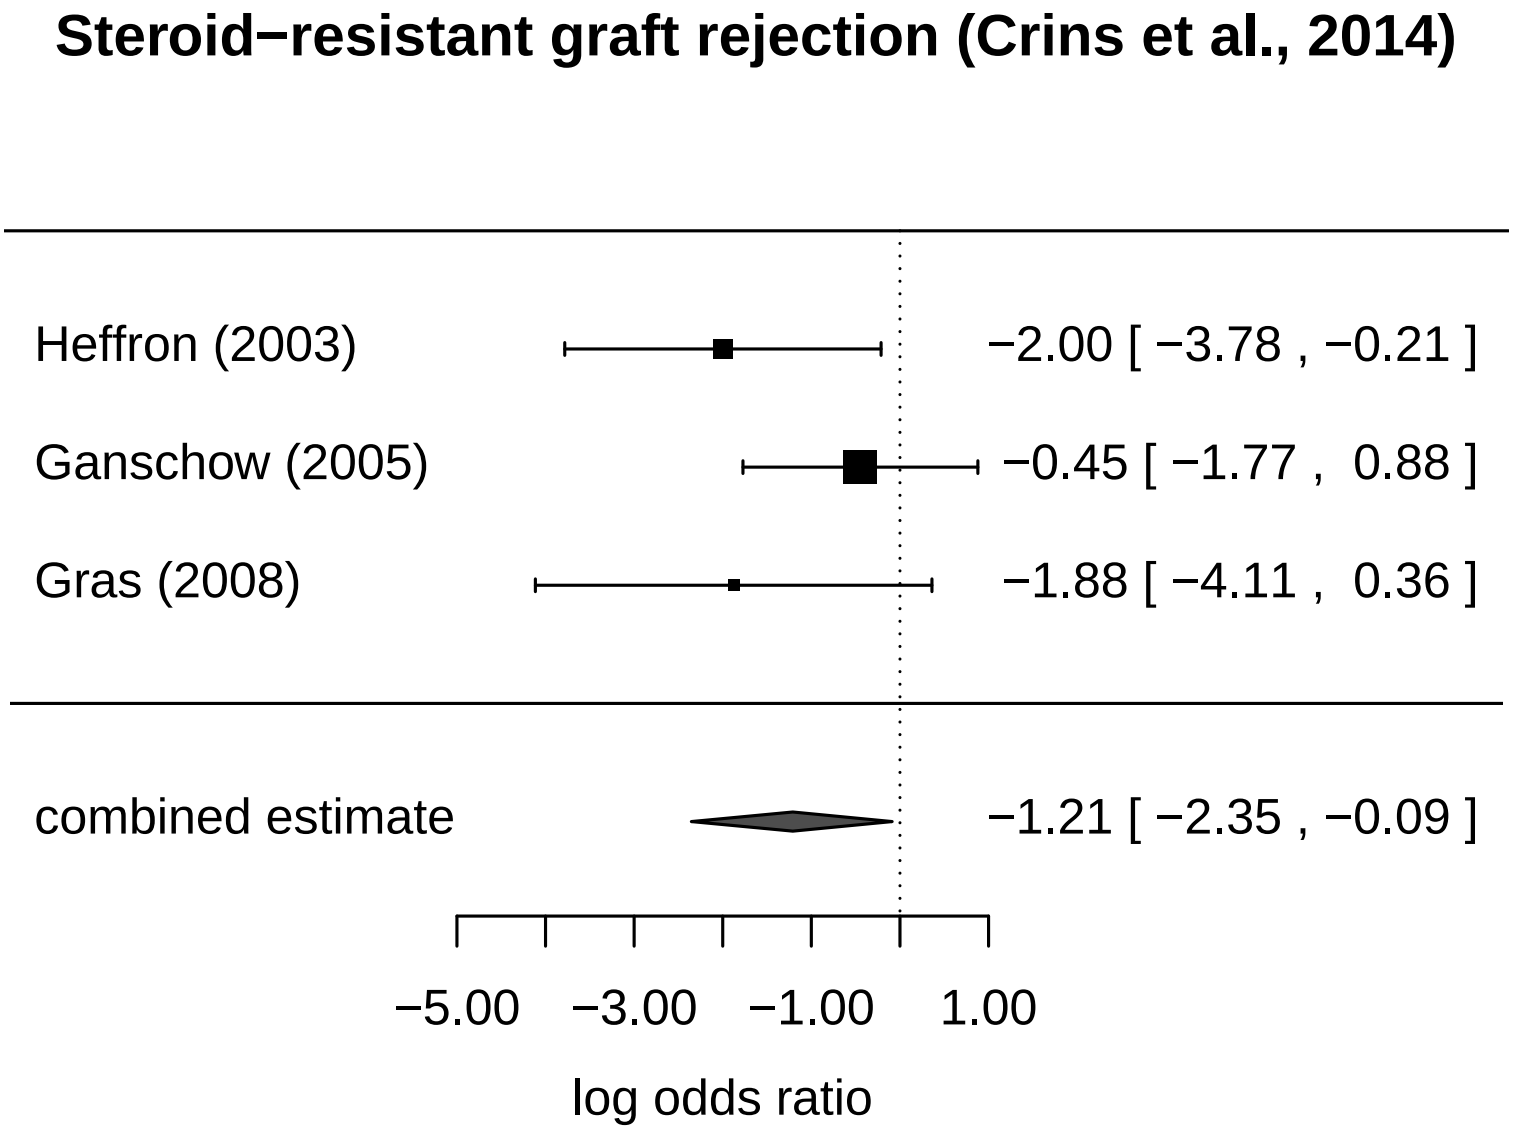

Meta analysis

The random effects model

- assume[1,2]:

$$y_i \sim \text{Normal}(\theta_i, s_i^2), \qquad \theta_i \sim \text{Normal}(\mu, \tau^2)$$

$$\Rightarrow \quad y_i \sim \text{Normal}(\mu, s_i^2 + \tau^2)$$

- model components:

Data:

- estimates y_i
- standard errors s_i

Parameters:

- true parameter value μ
- heterogeneity τ

[1] L. V. Hedges, I. Olkin. *Statistical methods for meta-analysis*. Academic Press, 1985.

[2] J. Hartung, G. Knapp, B. K. Sinha. *Statistical meta-analysis with applications*. Wiley, 2008.

Meta analysis

The random effects model

- assume[1,2]:

$$y_i \sim \text{Normal}(\theta_i, s_i^2), \qquad \theta_i \sim \text{Normal}(\mu, \tau^2)$$

$$\Rightarrow \quad y_i \sim \text{Normal}(\mu, s_i^2 + \tau^2)$$

- model components:

Data:

- estimates y_i
- standard errors s_i

Parameters:

- true parameter value μ
- heterogeneity τ

- $\mu \in \mathbb{R}$ of primary interest ("effect")
- $\tau \in \mathbb{R}^+$ nuisance parameter ("between-trial heterogeneity")

[1] L. V. Hedges, I. Olkin. *Statistical methods for meta-analysis*. Academic Press, 1985.

[2] J. Hartung, G. Knapp, B. K. Sinha. *Statistical meta-analysis with applications*. Wiley, 2008.

Meta analysis

The random effects model

- *normal-normal hierarchical model (NNHM)*
  applicable for many endpoints
- follow Bayesian approach here[3]
- suitable also for few studies (small k)
- consideration of prior information
- propagation of uncertainty
- straightforward interpretation
- computationally more involved, *usually* done via simulation (MCMC, BUGS)

[3] A. J. Sutton, K. R. Abrams. *Bayesian methods in meta-analysis and evidence synthesis*. Statistical Methods in Medical Research, 10(4):277, 2001.

Meta analysis

Prior, posterior

- have:
  - likelihood $p(\vec{y}, \vec{\sigma} \,|\, \mu, \tau)$
  - prior density $p(\mu, \tau) \;=\; p(\mu) \times p(\tau)$
- note:
  - Normal likelihood
  - Normal or (improper) uniform $p(\mu)$
  - $p(\mu)$ and $p(\tau)$ independent
- posterior $p(\mu, \tau \,|\, \vec{y}, \vec{\sigma}) \;\propto\; p(\vec{y}, \vec{\sigma} \,|\, \mu, \tau) \times p(\mu, \tau)$
- integrate out marginal posteriors
  - effect $p(\mu \,|\, \vec{y}, \vec{\sigma}) \;=\; \int p(\mu, \tau \,|\, \vec{y}, \vec{\sigma}) \, \mathrm{d}\tau$
  - heterogeneity $p(\tau \,|\, \vec{y}, \vec{\sigma}) \;=\; \int p(\mu, \tau \,|\, \vec{y}, \vec{\sigma}) \, \mathrm{d}\mu$
- inference: marginal distributions, posterior expectations, medians, quantiles,. . .

Meta analysis

Semi-analytical implementation

- integrals in NNHM may be solved semi-analytically
  - heterogeneity posterior (τ):
    analytical
  - effect posterior (μ):
    conditionally normal $\Rightarrow$ marginal = normal mixture
  - approximation via finite number of mixture components[4]
- method is implemented in `bayesmeta` R package[5]
- provides direct access to posterior densities, cumulative distribution functions, quantiles,...
- numerical accuracy is under control

[4] C. Röver, T. Friede. Discrete approximation of a mixture distribution via restricted divergence. *arXiv preprint 1602.04060* (`http://arxiv.org/abs/1602.04060`)

[5] `http://cran.r-project.org/package=bayesmeta`

Meta analysis

Semi-analytical implementation

- previously investigated general performance and compared to common frequentist methods (bias, coverage,. . .)[6]

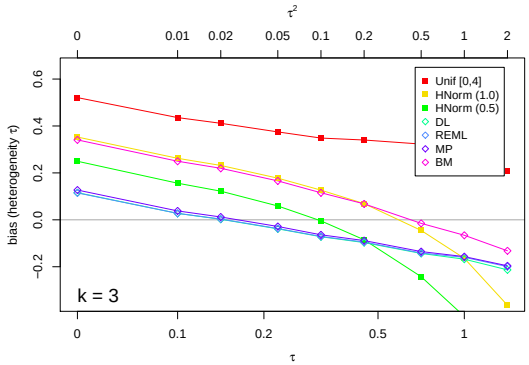

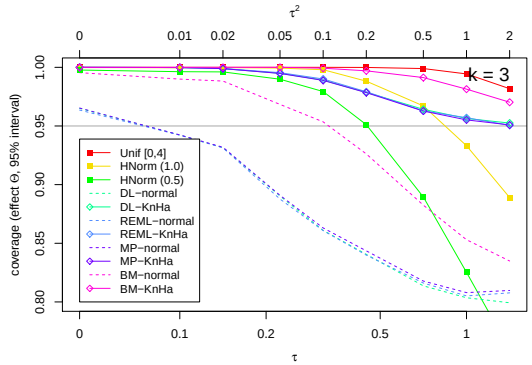

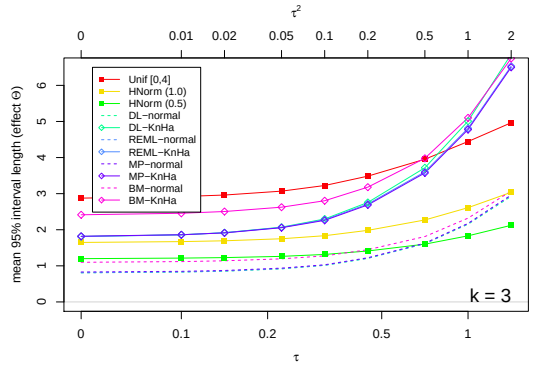

[6] T. Friede, C. Röver, S. Wandel, B. Neuenschwander. Meta-analysis of few small studies in orphan diseases. *Research Synthesis Methods* (in press; `http://arxiv.org/abs/1601.06533`).

Meta analysis

Semi-analytical implementation

- previously investigated general performance and compared to common frequentist methods (bias, coverage,. . .)[6]

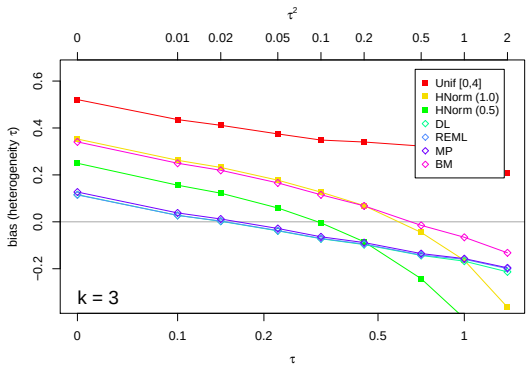

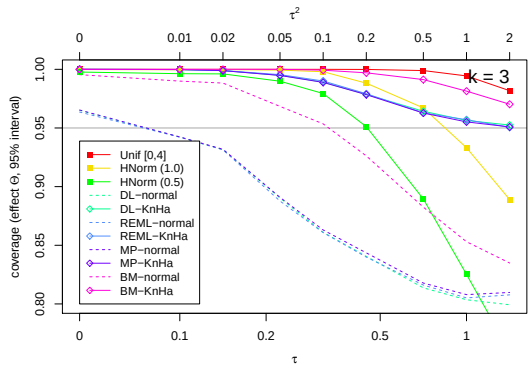

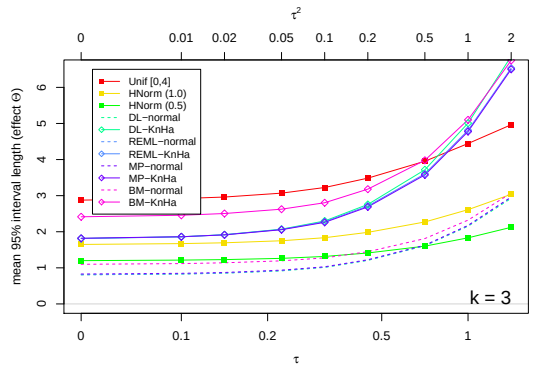

- how to carry out analysis in practice?

[6] T. Friede, C. Röver, S. Wandel, B. Neuenschwander. Meta-analysis of few small studies in orphan diseases. *Research Synthesis Methods* (in press; `http://arxiv.org/abs/1601.06533`).

Example

Crins et al. (2014) data

- Example data set: studies on steroid-resistant graft rejection in pediatric liver transplantation[7].
  $k = 3$ estimates (log-ORs) and standard errors

	treatment	control	effect size (log-OR)	
study	(events/total)	(events/total)	estimate (y_i)	std.err. (σ_i)
Heffron (2003)	2 / 61	4 / 20	-1.998	0.911
Ganschow (2005)	4 / 54	6 / 54	-0.446	0.676
Gras (2008)	1 / 50	4 / 34	-1.877	1.142

[7] N.D. Crins et al. Interleukin-2 receptor antagonists for pediatric liver transplant recipients: A systematic review and meta-analysis of controlled studies. *Pediatric Transplantation* 18(8):839–850, 2014.

[8] R.M. Turner et al. Predicting the extent of heterogeneity in meta-analysis, using empirical data from the Cochrane Database of Systematic Reviews. *International Journal of Epidemiology* 41(3):818–827, 2012.

E. Kontopantelis et al. A re-analysis of the Cochrane Library data: The dangers of unobserved heterogeneity in meta-analyses. *PLoS ONE* 8(7):e69930, 2013.

Example

Crins et al. (2014) data

- Example data set: studies on steroid-resistant graft rejection in pediatric liver transplantation[7].
  $k = 3$ estimates (log-ORs) and standard errors

	treatment	control	effect size (log-OR)	
study	(events/total)	(events/total)	estimate (y_i)	std.err. (σ_i)
Heffron (2003)	2 / 61	4 / 20	-1.998	0.911
Ganschow (2005)	4 / 54	6 / 54	-0.446	0.676
Gras (2008)	1 / 50	4 / 34	-1.877	1.142

- $k = 2$ to 3 studies is a common scenario
  (*majority* of meta analyses in Cochrane Database[8])

[7] N.D. Crins et al. Interleukin-2 receptor antagonists for pediatric liver transplant recipients: A systematic review and meta-analysis of controlled studies. *Pediatric Transplantation* 18(8):839–850, 2014.

[8] R.M. Turner et al. Predicting the extent of heterogeneity in meta-analysis, using empirical data from the Cochrane Database of Systematic Reviews. *International Journal of Epidemiology* 41(3):818–827, 2012.

E. Kontopantelis et al. A re-analysis of the Cochrane Library data: The dangers of unobserved heterogeneity in meta-analyses. *PLoS ONE* 8(7):e69930, 2013.

Example

Crins et al. (2014) data

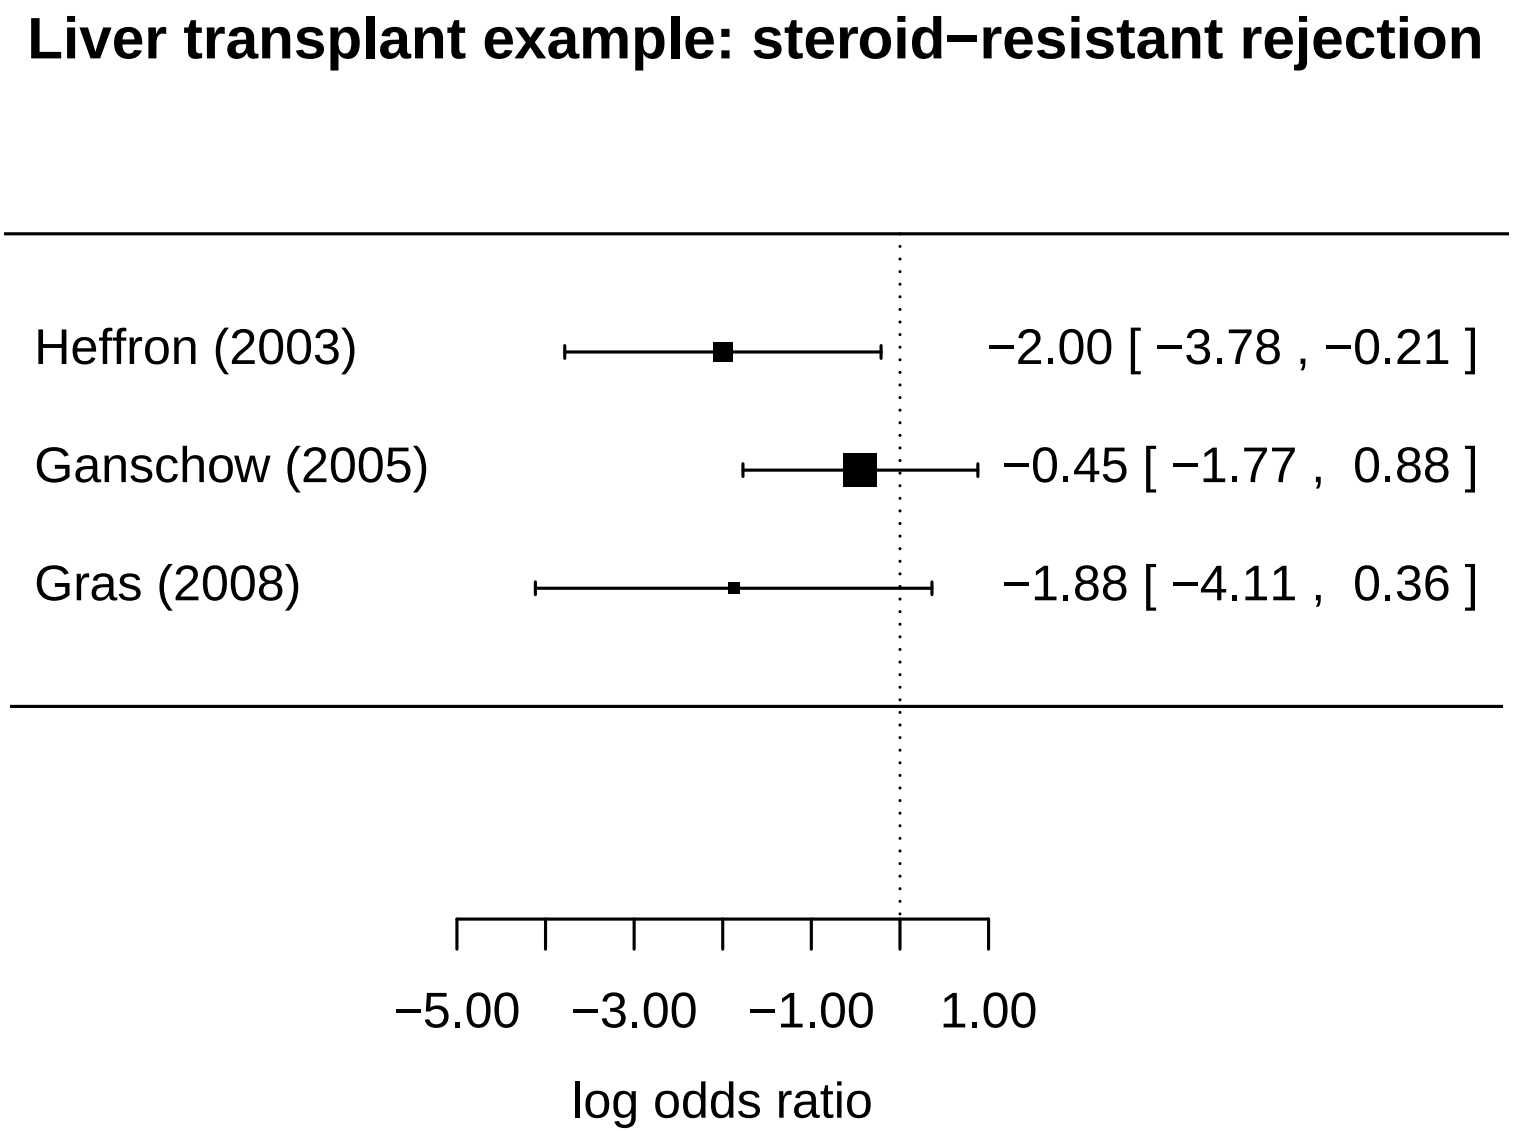

Example

Prior specification

- prior specification - - two unknowns;
  for example:

  - effect: $\mu \sim \text{Normal}(\mu = 0, \sigma = 10)$

    (vague prior)

  - heterogeneity $\tau \sim \text{half-Normal}(\sigma = 0.5)$

    prior distribution:

quantile	heterogeneity τ	$\exp(\tau)$
2.5 %	0.016	1.016
50.0 %	0.337	1.401
97.5 %	1.121	3.067

    (spans range from *homogeneity* to "*substantial heterogeneity*"[9])

[9] D.J. Spiegelhalter, K.R. Abrams, J.P. Myles. *Bayesian approaches to clinical trials and health-care evaluation*. Wiley, 2004.

Example

Data preparation

- data (counts):

```
> CrinsEtAl2014[c(1,4,6),c(1,11,12,14,15)]
      publication exp.SRR.events exp.total cont.SRR.events cont.total
1  Heffron (2003)              2        61               4         20
4 Ganschow (2005)              4        54               6         54
6     Gras (2008)              1        50               4         34
```

Example

Data preparation

- data (counts):

```
> CrinsEtAl2014[c(1,4,6),c(1,11,12,14,15)]
      publication exp.SRR.events exp.total cont.SRR.events cont.total
1  Heffron (2003)              2        61               4         20
4 Ganschow (2005)              4        54               6         54
6     Gras (2008)              1        50               4         34
```

- compute effect sizes (log-ORs):

```
> library("metafor")
> effsize <- escalc(measure="OR",
                    ai=exp.SRR.events,  n1i=exp.total,
                    ci=cont.SRR.events, n2i=cont.total,
                    slab=publication, data=CrinsEtAl2014[c(1,4,6),])
> effsize[,c(1,16,17)]
      publication      yi     vi
1  Heffron (2003) -1.9981 0.8294
2 Ganschow (2005) -0.4463 0.4575
3     Gras (2008) -1.8769 1.3037
```

Example

Computation

- perform analysis:

```
> library("bayesmeta")
> bm01 <- bayesmeta(effsize,
                    mu.prior.mean=0.0, mu.prior.sd=10.0,
                    tau.prior=function(x){dhalfnormal(x,scale=0.5)})
```

(specify effect prior via moments, heterogeneity prior via density)

Example

Computation

- perform analysis:

```
> library("bayesmeta")
> bm01 <- bayesmeta(effsize,
                    mu.prior.mean=0.0, mu.prior.sd=10.0,
                    tau.prior=function(x){dhalfnormal(x,scale=0.5)})
```

  (specify effect prior via moments, heterogeneity prior via density)

- may also specify data (y, `sigma`) and labels individually:

```
> bm01 <- bayesmeta(y=as.vector(effsize$yi),
                    sigma=sqrt(effsize$vi),
                    labels=effsize$study,
                    mu.prior.mean=0.0, mu.prior.sd=10.0,
                    tau.prior=function(x){dhalfnormal(x,scale=0.5)})
```

Example

Results

- print default output:

```
> bm01
 'bayesmeta' object.

3 estimates:
Heffron (2003), Ganschow (2005), Gras (2008)

tau prior:
function(x){dhalfnormal(x,scale=0.5)}

mu prior:
normal(mean=0, sd=10)

ML and MAP estimates:
                       tau        mu
ML joint     1.771042e-04 -1.160107
ML marginal  5.077174e-01        NA
MAP joint    9.862966e-05 -1.157224
MAP marginal 0.000000e+00 -1.194867

marginal posterior summary:
                 tau          mu
mode      0.0000000 -1.19486744
median    0.3380733 -1.20525311
mean      0.3920815 -1.21188547
sd        0.2881225  0.57387301
95% lower 0.0000000 -2.34756473
95% upper 0.9406362 -0.08617254
```

Example

Results

- the `bayesmeta()` function returns the main result (current example: stored in "`bm01`" object).
- `list` object; elements:
  - `...$dposterior`: posterior density function (μ, τ or joint)
  - `...$pposterior`, `...$qposterior`: cumulative distribution function, quantile function
  - `...$post.interval`: function to determine credibility intervals
  - `...$summary`: table of essential summary statistics
  - ...

Example

Results

- show posterior density of effect μ:

```
mu <- seq(from=-3, to=1, length=100)
plot(mu, bm01$dposterior(mu=mu), type="l",
     col="blue", xlab="effect (log-OR)", ylab="probability density")
```

Example

Results

- show posterior density of effect μ:

```
mu <- seq(from=-3, to=1, length=100)
plot(mu, bm01$dposterior(mu=mu), type="l",
     col="blue", xlab="effect (log-OR)", ylab="probability density")
```

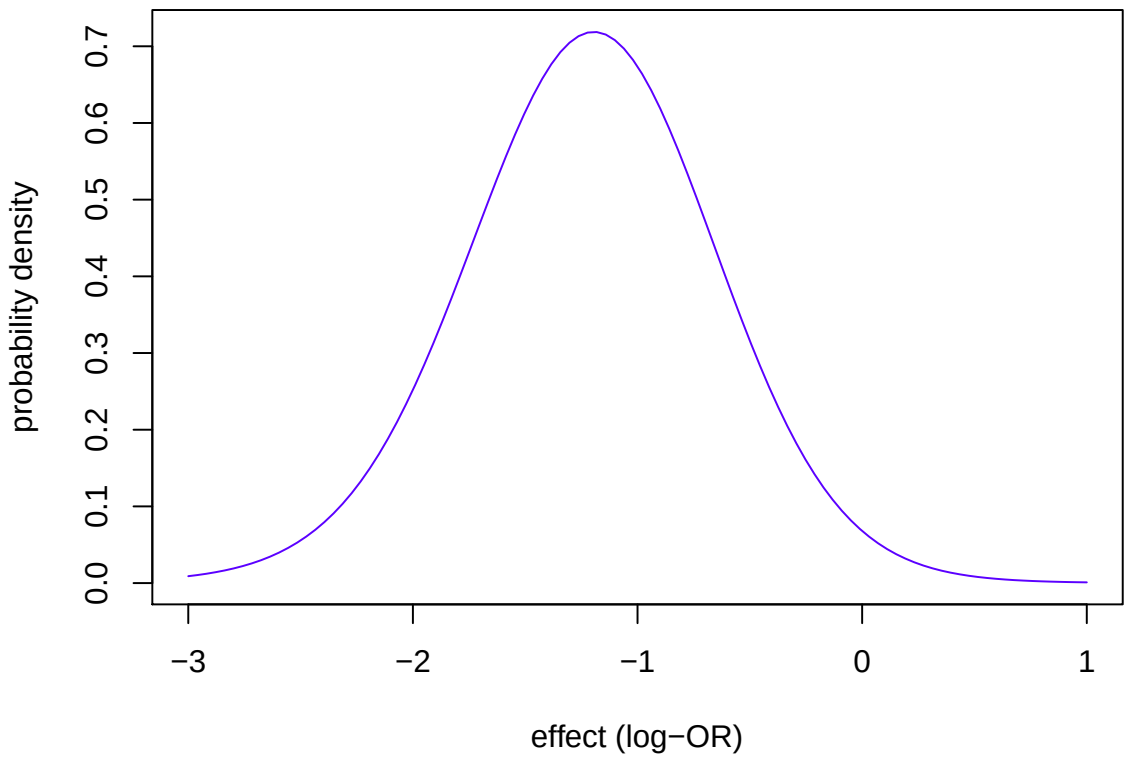

Example

Results

- show posterior density of effect μ:

```
mu <- seq(from=-3, to=1, length=100)
plot(mu, bm01$dposterior(mu=mu), type="l",
     col="blue", xlab="effect (log-OR)", ylab="probability density")

lines(mu, bm01$dprior(mu=mu), lty="dashed", col= "blue")
abline(h=0, v=0)
```

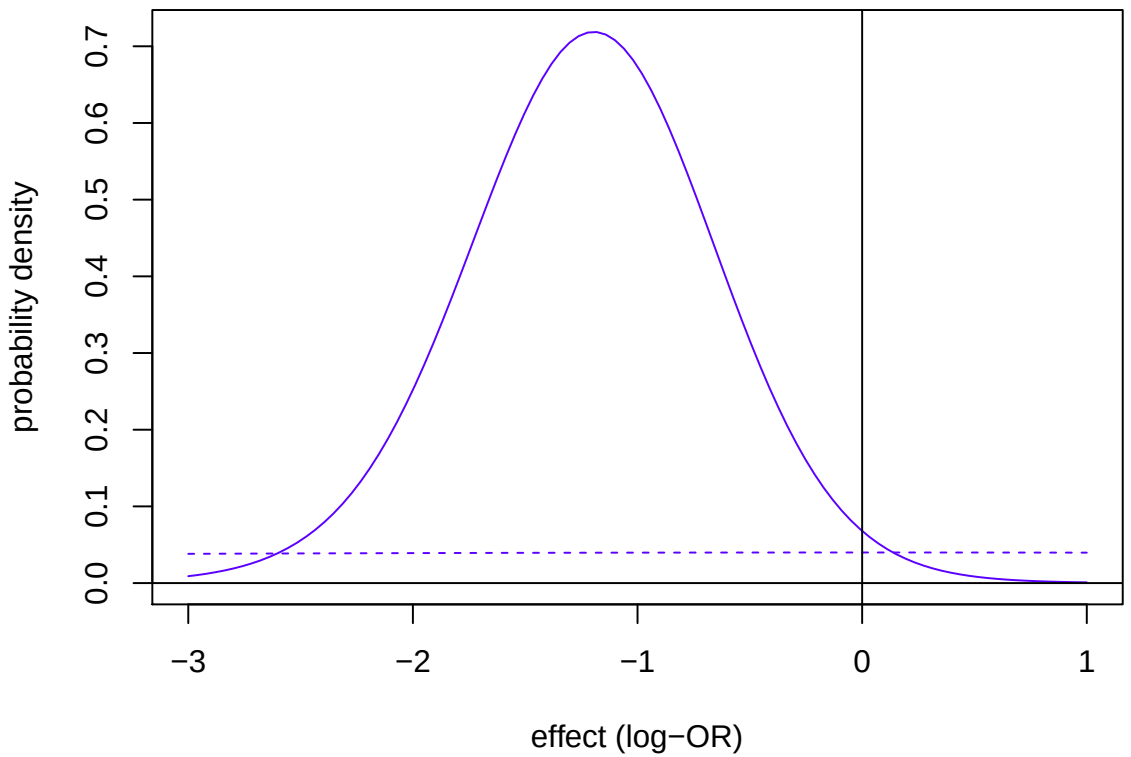

Example

Results

- what is the probability of a beneficial effect (P($\mu \leq 0$))?
  $\rightarrow$ evaluate cumulative distribution function:

Example

Results

- what is the probability of a beneficial effect (P($\mu \leq 0$))?
  $\rightarrow$ evaluate cumulative distribution function:

```
> bm01$pposterior(mu=0)
[1] 0.9833152
```

Example

Results

- what is the probability of a beneficial effect (P($\mu \leq 0$))?
  $\rightarrow$ evaluate cumulative distribution function:

```
> bm01$pposterior(mu=0)
[1] 0.9833152
```

- what is the 95% upper limit on the log-OR?
  $\rightarrow$ evaluate quantile function:

```
> bm01$qposterior(mu=0.95)
[1] -0.2859446
```

Example

Results

- what is the probability of a beneficial effect (P($\mu \leq 0$))?
  $\rightarrow$ evaluate cumulative distribution function:

```
> bm01$pposterior(mu=0)
[1] 0.9833152
```

- what is the 95% upper limit on the log-OR?
  $\rightarrow$ evaluate quantile function:

```
> bm01$qposterior(mu=0.95)
[1] -0.2859446
```

- analogous for τ...

```
> bm01$pposterior(tau=1.0)
[1] 0.9632259
> bm01$qposterior(tau=0.95)
[1] 0.9406362
```

Example

Results

- posterior density of heterogeneity τ:

```
tau <- seq(from=0, to=2.5, length=100)
plot(tau, bm01$dposterior(tau=tau), type="l",
     col="green", xlab="heterogeneity (tau)", ylab="probability densi
abline(h=0, v=0)
```

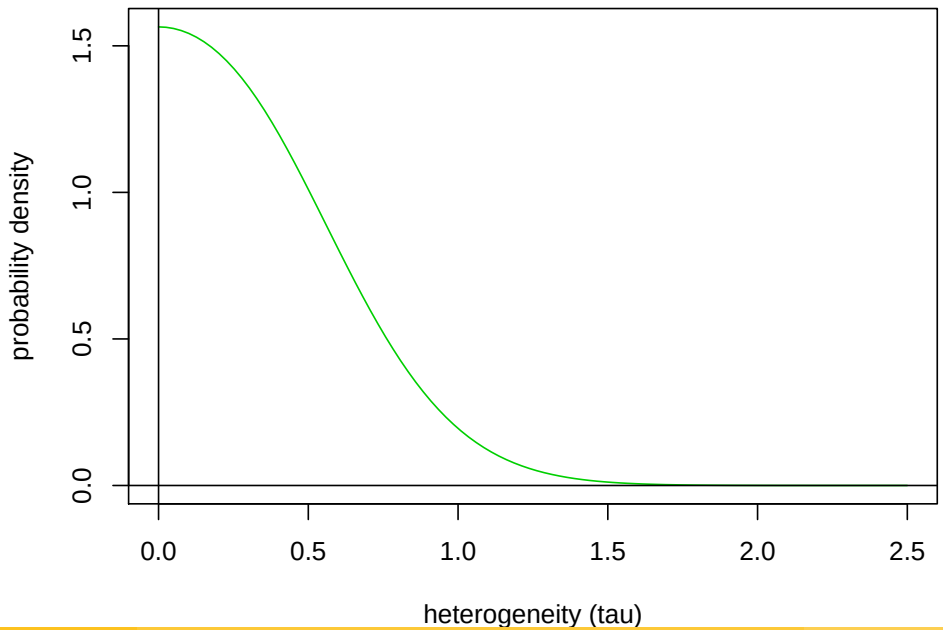

Example

Results

- posterior density of heterogeneity τ:

```
tau <- seq(from=0, to=2.5, length=100)
plot(tau, bm01$dposterior(tau=tau), type="l",
     col="green", xlab="heterogeneity (tau)", ylab="probability densi
abline(h=0, v=0)
```

- 95% credibility interval (default: shortest interval):

```
> bm01$post.interval(tau.level=0.95)
[1] 0.0000000 0.9406362
> abline(v=bm01$post.interval(tau.level=0.95), col="red", lty=2)
```

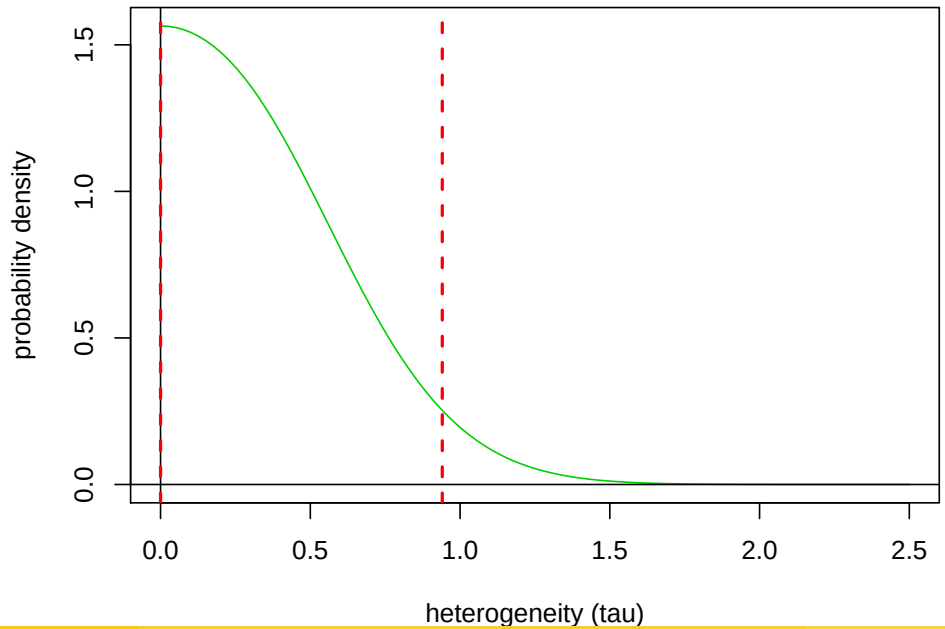

Example

Results

- predictive distributions –
  distribution of "new" study's true effect θ_{k+1}

```
> bm01$qposterior(mu=c(0.025, 0.975))
[1] -2.3634190 -0.1011219
> bm01$qposterior(mu=c(0.025, 0.975), predict=TRUE)
[1] -2.7717423  0.2576317
```

- provides *meta-analytic-predictive (MAP) prior*[10]

[10] B. Neuenschwander, G. Capkun-Niggli, M. Branson, and D.J. Spiegelhalter. Summarizing historical information on controls in clinical trials. *Clinical Trials* 7(1):5-18, 2010.

H. Schmidli et al. Robust meta-analytic-predictive priors in clinical trials with historical control information. *Biometrics* 70(4):1023-1032, 2014.

Example

Results

- quick sensitivity checks
  (uniform effect prior, *very* wide heterogeneity prior):

```
bm01 <- bayesmeta(effsize,
                  mu.prior.mean=0.0, mu.prior.sd=10.0,
                  tau.prior=function(x){dhalfnormal(x,scale=0.5)})
bm02 <- bayesmeta(effsize,
                  tau.prior=function(x){dhalfnormal(x,scale=1.0)})
```

Example

Results

- quick sensitivity checks
  (uniform effect prior, *very* wide heterogeneity prior):

```
bm01 <- bayesmeta(effsize,
                  mu.prior.mean=0.0, mu.prior.sd=10.0,
                  tau.prior=function(x){dhalfnormal(x,scale=0.5)})
bm02 <- bayesmeta(effsize,
                  tau.prior=function(x){dhalfnormal(x,scale=1.0)})
```

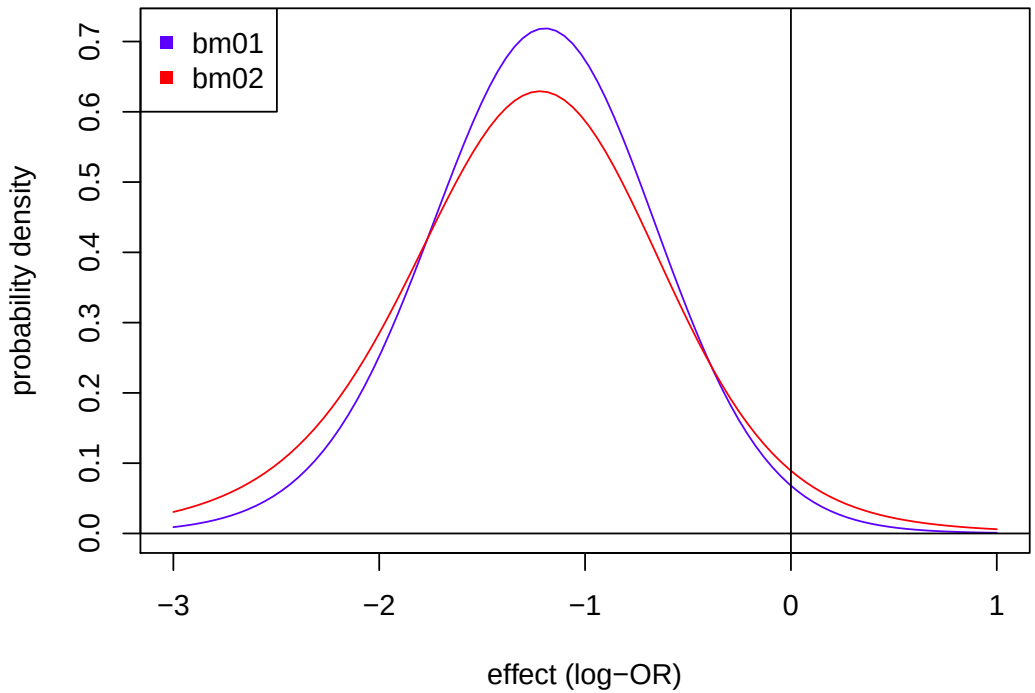

Example

Results

- forest plot:

```
> forest(bm01, main="Steroid-resistant rejection (SRR)")
```

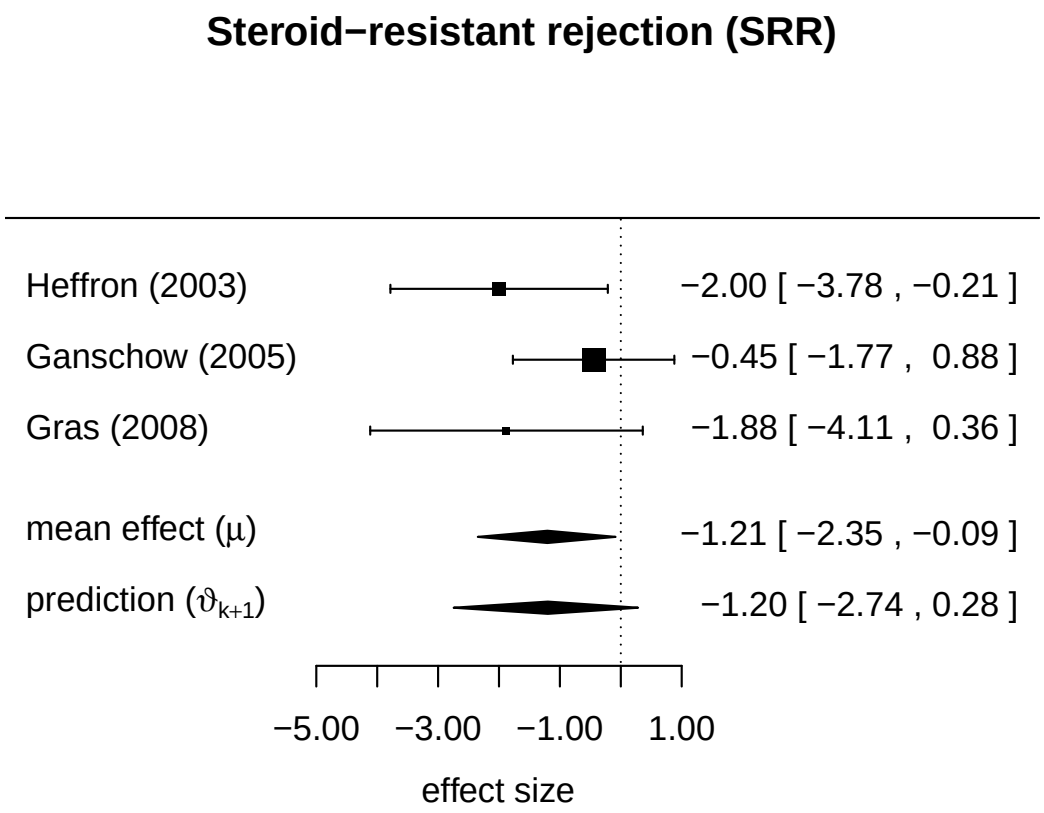

Conclusions

- random-effects meta-analysis model covers wide range of cases
- semi-analytical integration simplifies Bayesian meta-analysis (esp.: no MCMC sampling necessary)
- R implementation is straightforward to use
- flexible prior specification
- quick sensitivity analyses
- includes predictive distributions

- `bayesmeta` package available on CRAN[11]

[11] `http://cran.r-project.org/package=bayesmeta`