Bayesian Stopping Rules for Meta-analyses

Katherine Daignault

Department of Mathematics and Statistics
McGill University

April 12, 2013

Outline

1. Introduction
2. Methods
   - Bayesian Stopping Rule
   - BIC approximation
3. Results
   - ECMO data
   - Simulation
4. Conclusion

What is a meta-analysis?

- Multiple trials are often run on the same question with different populations.
- It would be practical to use this information to determine an overall result for this question.
- A meta-analysis is a statistical tool that allows the analyst to integrate the findings of these trials.
- Analysts use forest plots to show the result of each study and the combined result from the meta-analysis.

Meta-analysis example

Forest plot for meta-analysis with 5 studies

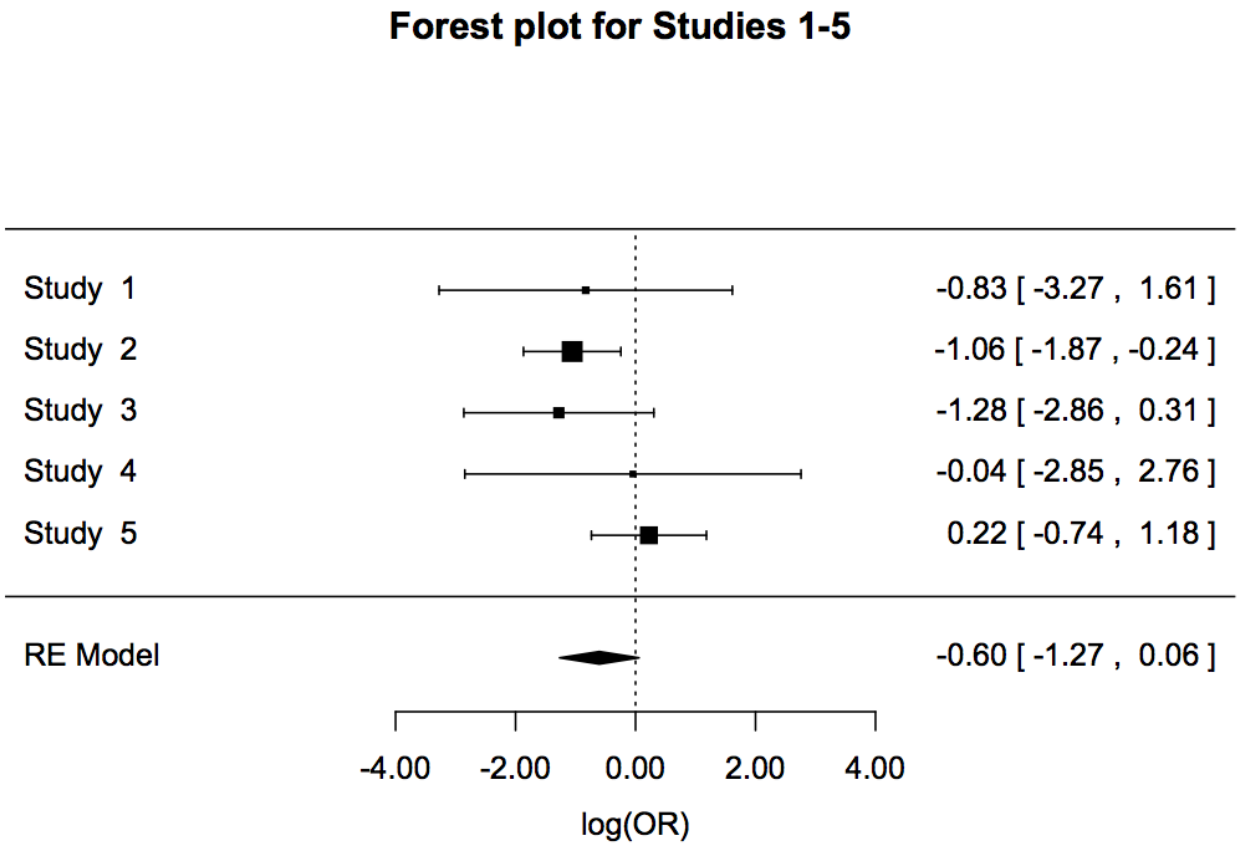

What happens to the overall estimate of log odds ratio if we add 10 more studies to the meta-analysis?

Meta-analysis example

Forest plot for meta-analysis with 15 studies

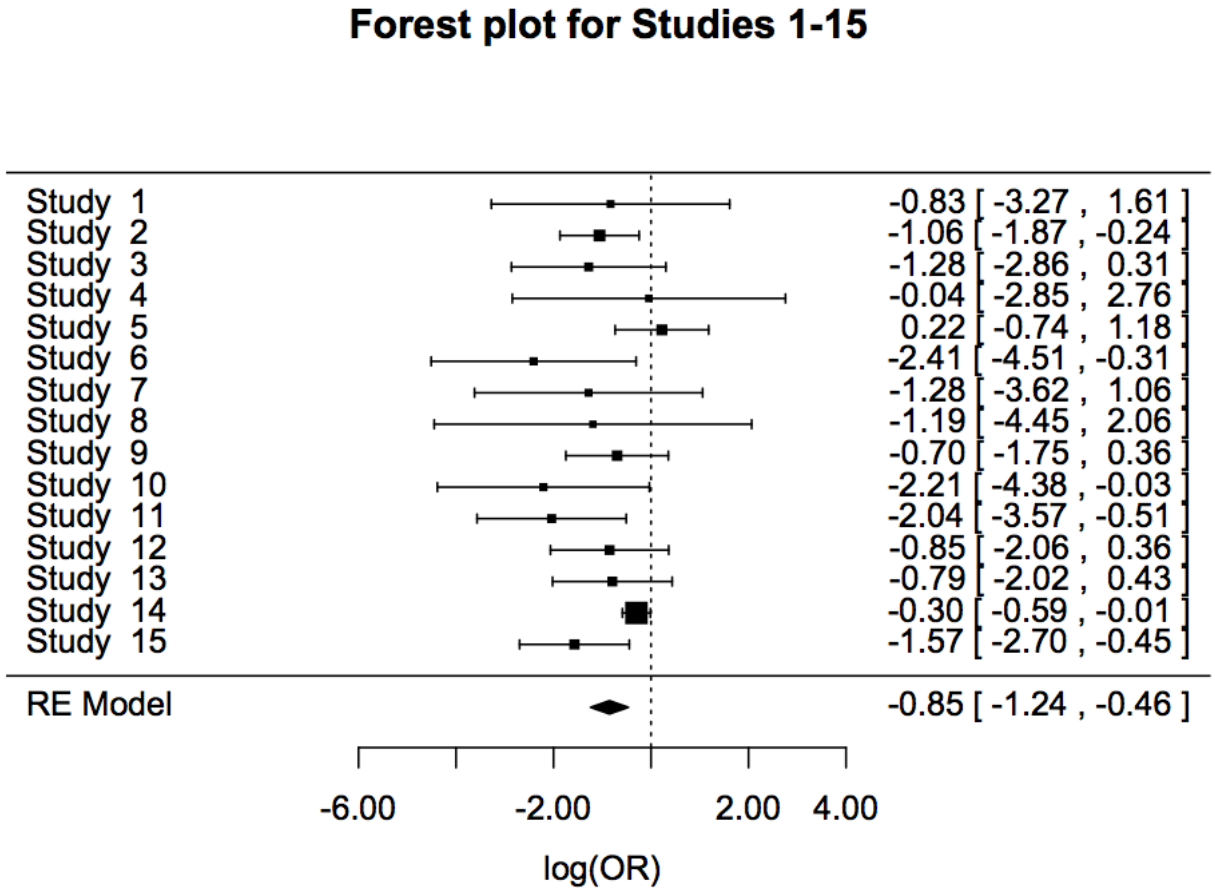

We see that the log odds ratio changed with the addition of 10 more studies. What happens if we add in a mega trial?

Meta-analysis example

Forest plot for meta-analysis with mega trial

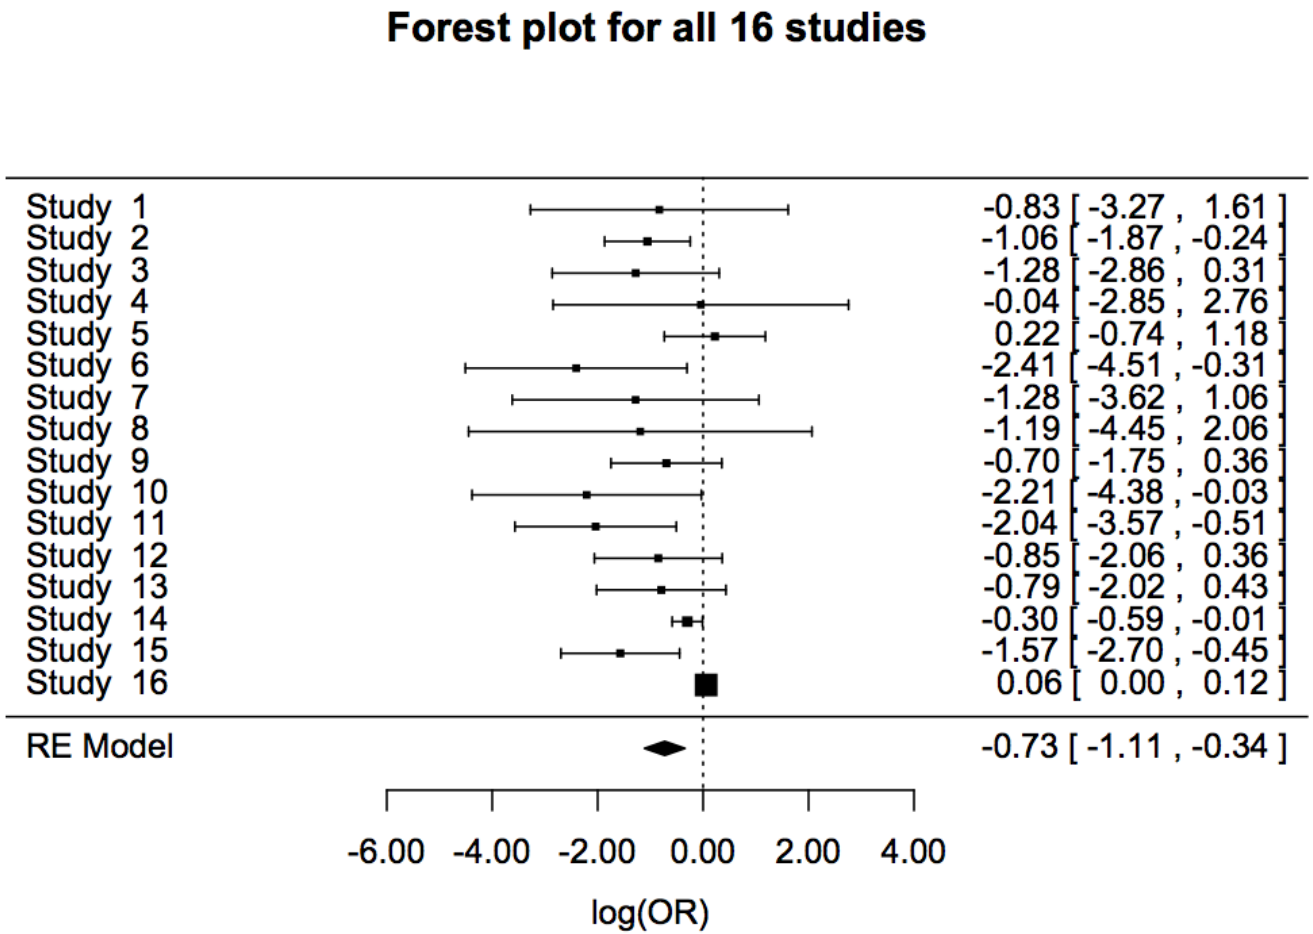

Study 16 had a very large impact on the log odds ratio of the meta-analysis. Should we include it in the analysis?

What is a decision rule?

- In trials, ethical considerations sometimes require the premature termination of a trial, e.g. treatment being tested may be harmful instead of helpful.
- But stopping the trial early might not give an accurate idea of true treatment effect.
- A stopping rule dictates the earliest moment to terminate the trial without changing the overall conclusion of the trial.

Bayesian Stopping Rules for trials

- Stopping rules were originally a frequentist approach.
- Intuitively, it would be reasonable to utilize any existing information concerning the problem at hand (why would we start from nothing if we already know something?).
- Frequentist methods do not incorporate any previous knowledge, but Bayesian methods will.

Bayesian Statistics

Let $\pi(\theta)$ be a probability distribution that incorporates prior information about our parameter of interest θ. Let $L(\theta|x) = \prod_{i=1}^{n} f(x|\theta)$ be the likelihood of the data. Then the posterior distribution is

$$\pi(\theta|x) = \frac{L(\theta|x)\pi(\theta)}{\int L(\theta|x)\pi(\theta)d\theta} = \frac{L(\theta|x)\pi(\theta)}{m_\pi(x)}$$

where $m_\pi(x)$ is the marginal likelihood of the data.

Greenhouse and Wasserman paper

- Greenhouse and Wasserman(1995) introduce a robust Bayesian stopping rule for clinical trials.
- Robustness here refers to the use of a class of prior distributions with similar properties instead of begin restricted to just one distribution.
- Due to this robustness, their approach is useful for meta-analysis since it requires only a class of priors to be specified, reducing the difficulty in determining an acceptable prior for meta-analysis.

Greenhouse and Wasserman paper

- The ϵ-contaminated class of priors, $\Gamma_\epsilon = \{\pi = (1-\epsilon)\pi + \epsilon q : q \in Q\}$ where Q is some class of prior and $\epsilon \in (0,1)$, allows the analyst to control the contribution of the prior on posterior inference.
- This prior allows us to obtain upper and lower bounds on our inference in the form

$$\overline{E}_{\pi(\theta|x)}(g(\theta)|x) = \sup_\theta \frac{(1-\epsilon)m_\pi(x)E_\pi(g(\theta)|x) + \epsilon L(\theta|x)g(\theta)}{(1-\epsilon)m_\pi(x) + \epsilon L(\theta|x)}$$

$$\underline{E}_{\pi(\theta|x)}(g(\theta)|x) = \inf_\theta \frac{(1-\epsilon)m_\pi(x)E_\pi(g(\theta)|x) + \epsilon L(\theta|x)g(\theta)}{(1-\epsilon)m_\pi(x) + \epsilon L(\theta|x)}$$

Greenhouse and Wasserman paper

The form of the posterior bounds also provide nice graphical interpretability.

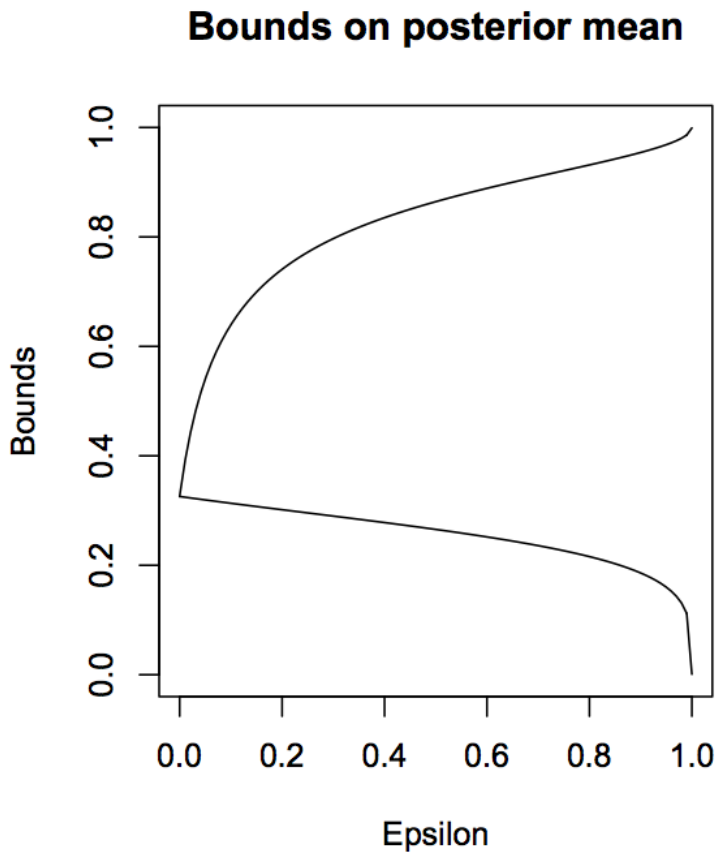

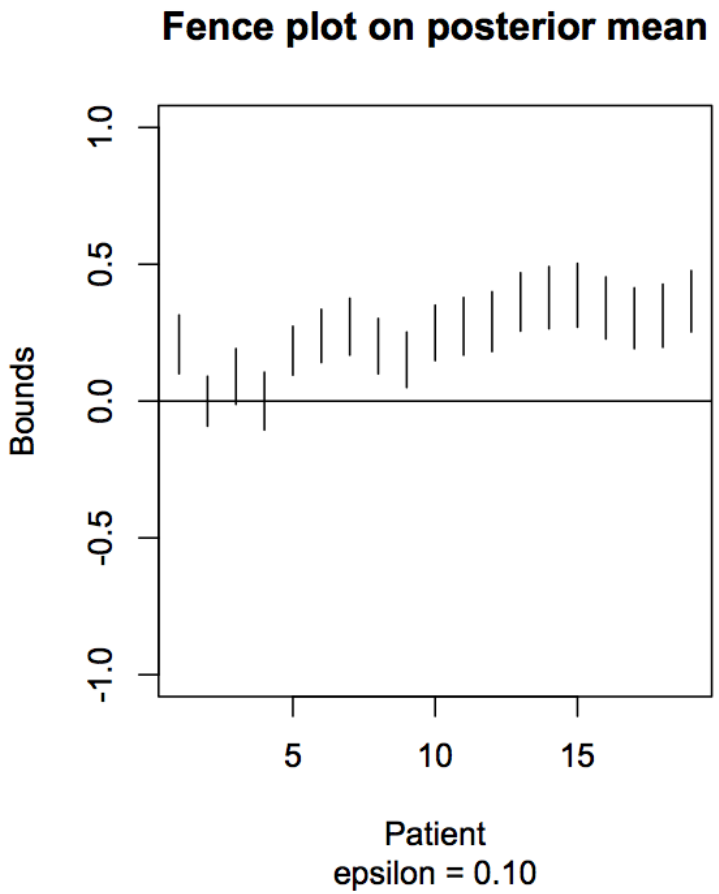

Problems

- The graphical features of the bounds make a wonderful tool for determining stopping rules.
- The largest obstacle in the calculation of the bounds is that the marginal likelihood, $m_\pi(x)$ may not always be in closed form and therefore not analytically integrable.
- Many computational procedures exist for approximating integrals. But how badly will they affect the bounds on our stoping rule?

BIC approximation

- The Bayesian Information Criterion (BIC) is usually used for comparing models in model selection problems.
- The standard expression for the BIC is

$$\text{BIC} = -2 \ln L(x|\theta) + k \ln n$$

where $L(x|\theta)$ is the maximized likelihood, k is the number of parameters to be estimated, and n is the sample size.

BIC approximation

- There is an alternative form for the BIC that may be used to approximate the marginal likelihood.
- We may write

$$\text{BIC} = -2 \ln L(x|\theta) + k \ln n \approx -2 \ln m_\pi(x)$$

where $m_\pi(x)$ is the marginal likelihood

- However, the biggest problem with this approach is that it does not handle random effects well.

Harvard ECMO clinical trial data

- Extracorporeal membrane oxygenation (ECMO) compared to conventional medical therapy (CMT) to treat persistent pulmonary hypertension of newborns.
- Patients were randomly assigned to treatment; as soon as 4 deaths occurred in either arm, randomization stopped.
- 19 patients enrolled; 9/9 survived in ECMO arm, 6/10 survived in CMT.

Harvard ECMO clinical trial data

- Our likelihood is binomial, so we have conjugate beta priors on the probabilities of success.
- Let p_E and p_C be the probability of survival for each arm. We put independent uniform priors on each which can be expressed in terms of a Beta distribution
- Here, we can obtain the marginal likelihood in closed form.
- We are in a fixed effect setting, so we would expect our approximations to be comparable, depending on the starting values of the function.

Harvard ECMO clinical trial data

Here we see the difference in posterior probabilities under both methods.

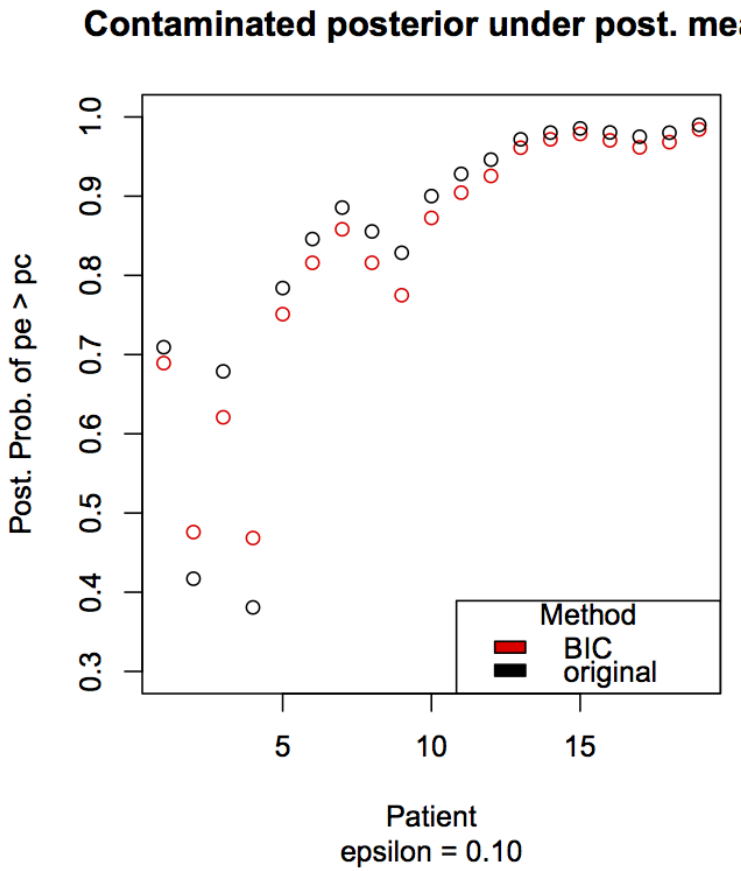

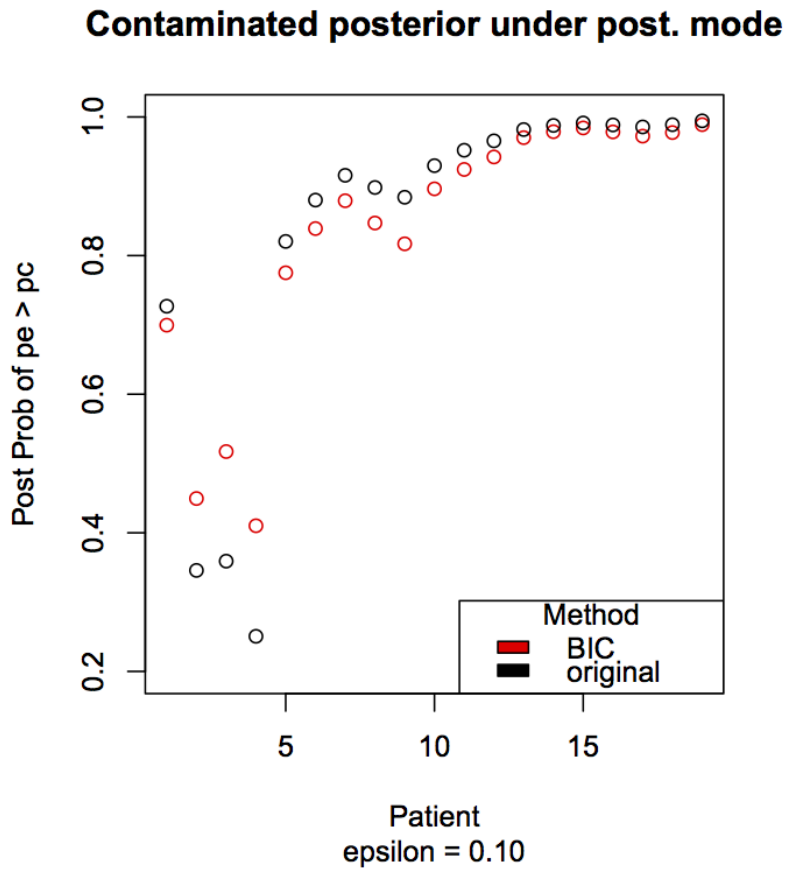

Harvard ECMO clinical trial data

Now we can see the effect each method would have on the stopping decision in the trial.

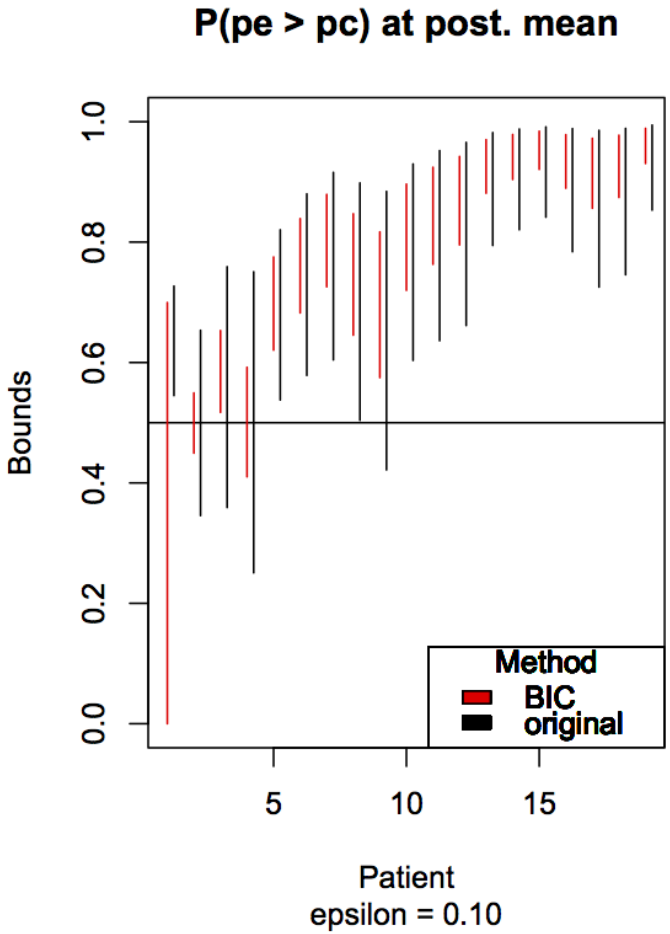

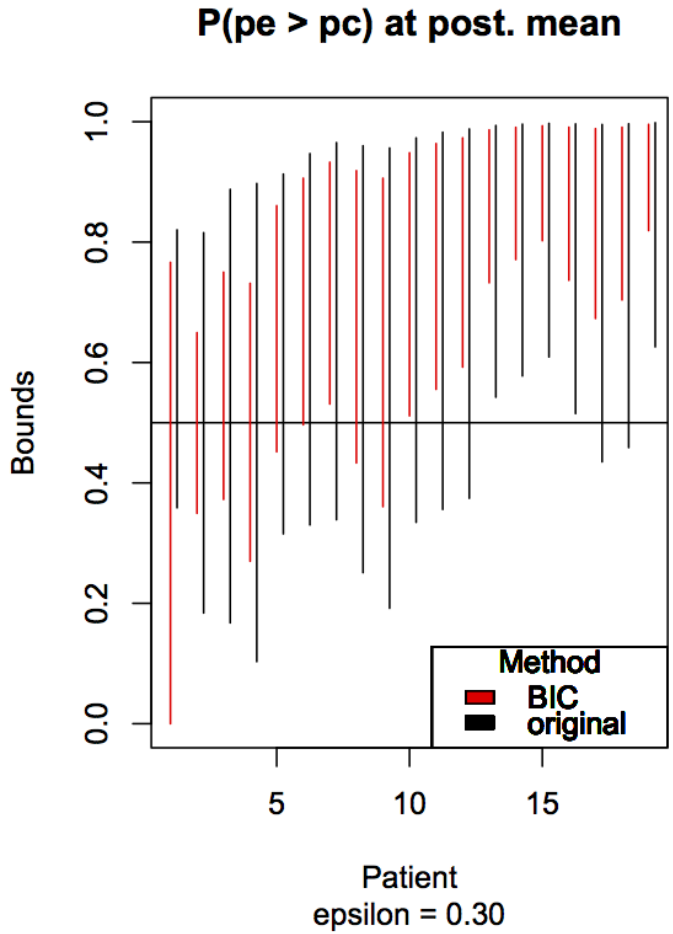

Harvard ECMO clinical trial data

- The BIC approximation of the marginal likelihood for trial data seems to underestimate the bounds of the posterior probabilities.
- If we are quite confident in our choice of prior (ϵ small), then we seem to agree on a stopping point with the closed form method.
- However, as we become more uncertain of the validity of the prior, the underestimation of the BIC becomes problematic.
- Small sample size could be to blame; so what happens if we apply this technique to a meta-analysis?

Meta-analysis simulation study

- Simulate 16 studies of size 200 each, with 100 subjects in each of the treatment and control arms.
- The number of cases in each arm were generated from a Binomial distribution where

$$\log(p_p) \sim U(-3.665, -0.995)$$

$$\log(p_t) = \log(p_p) + \theta$$

  where θ is the treatment effect.
- This allows us to have a common baseline among the studies but a study-specific treatment effect.

Meta-analysis simulation study

Forest plot for the simulated data.

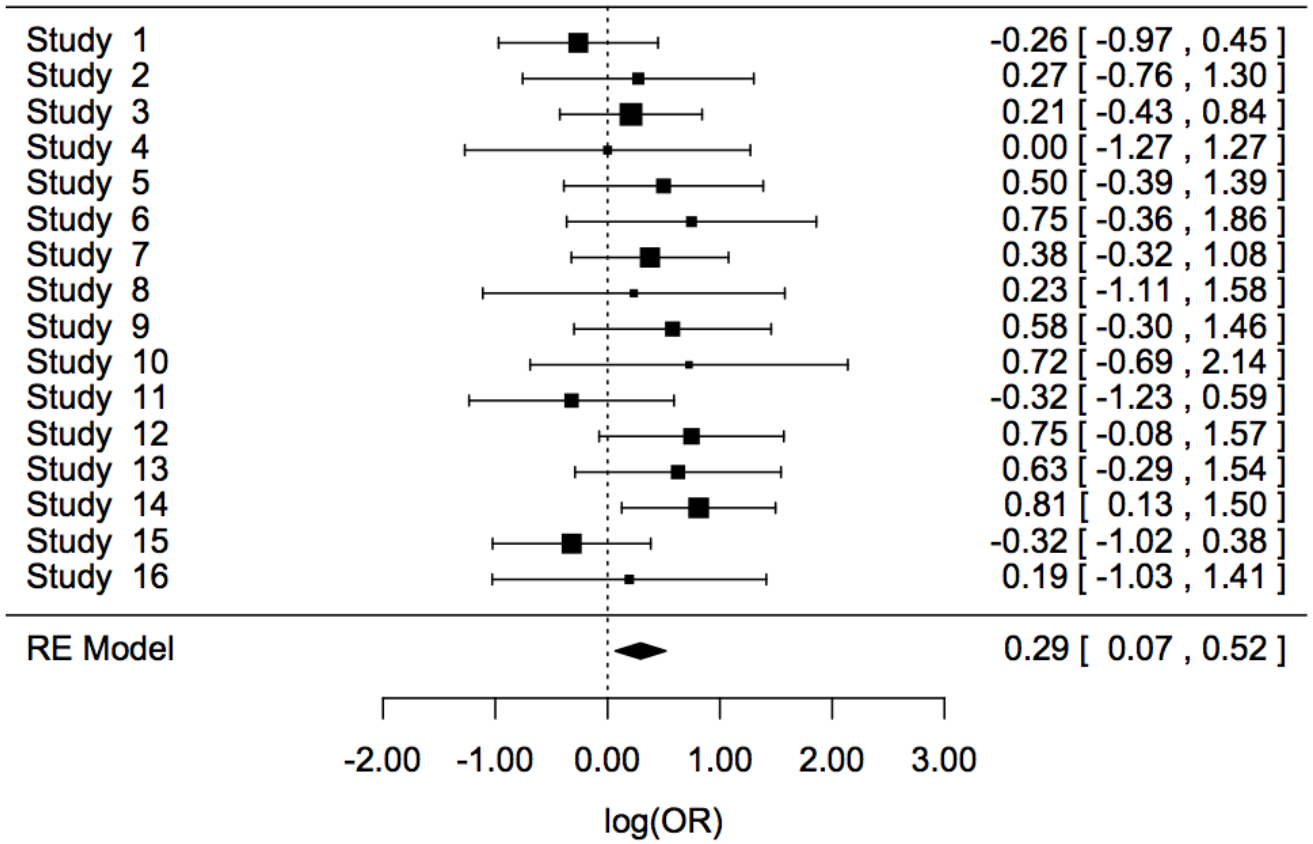

Meta-analysis simulation study

- While allowing for a random effect is far more realistic, we can no longer employ our previous methods to these data.
- The marginal likelihood can no longer be obtained in closed form.
- The BIC approximation is for fixed effect models only where it does a questionable job already.
- So what can we do?

Meta-analysis simulation study

- Linear mixed models (LM) can incorporate a random effect in the model and provide us with an estimate of the BIC.
- Random effects models are far more realistic as they allow modelling of both between- study and within-study variability.
- Fitting the simulated data with "lmer" in R, we can extract an estimate of the marginal likelihood under a random effects model and a value for the log likelihood, which we can use in our calculations of the posterior probabilities.

Meta-analysis simulation study

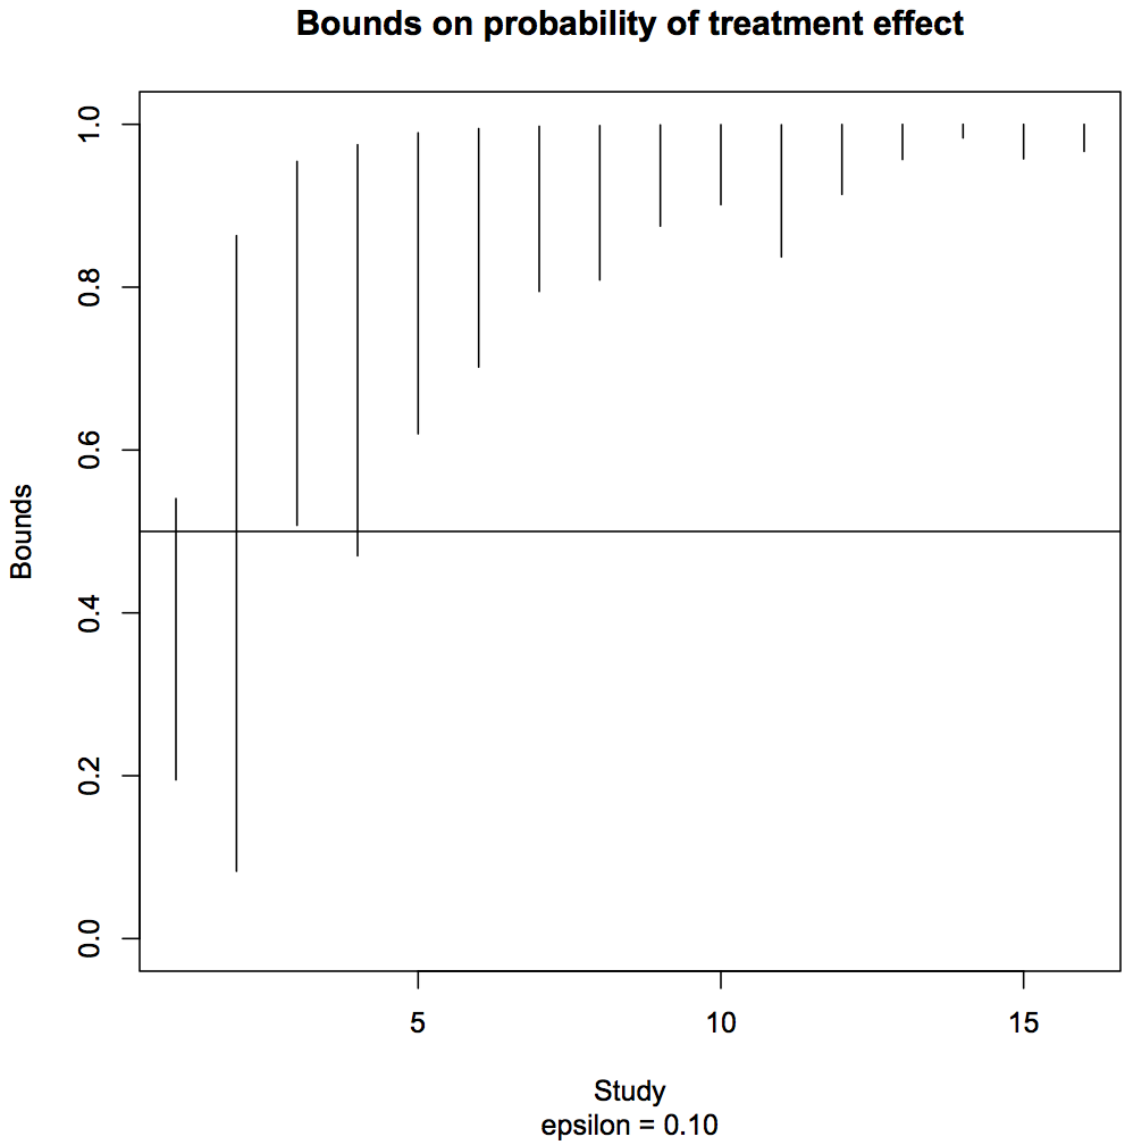

Meta-analysis simulation study

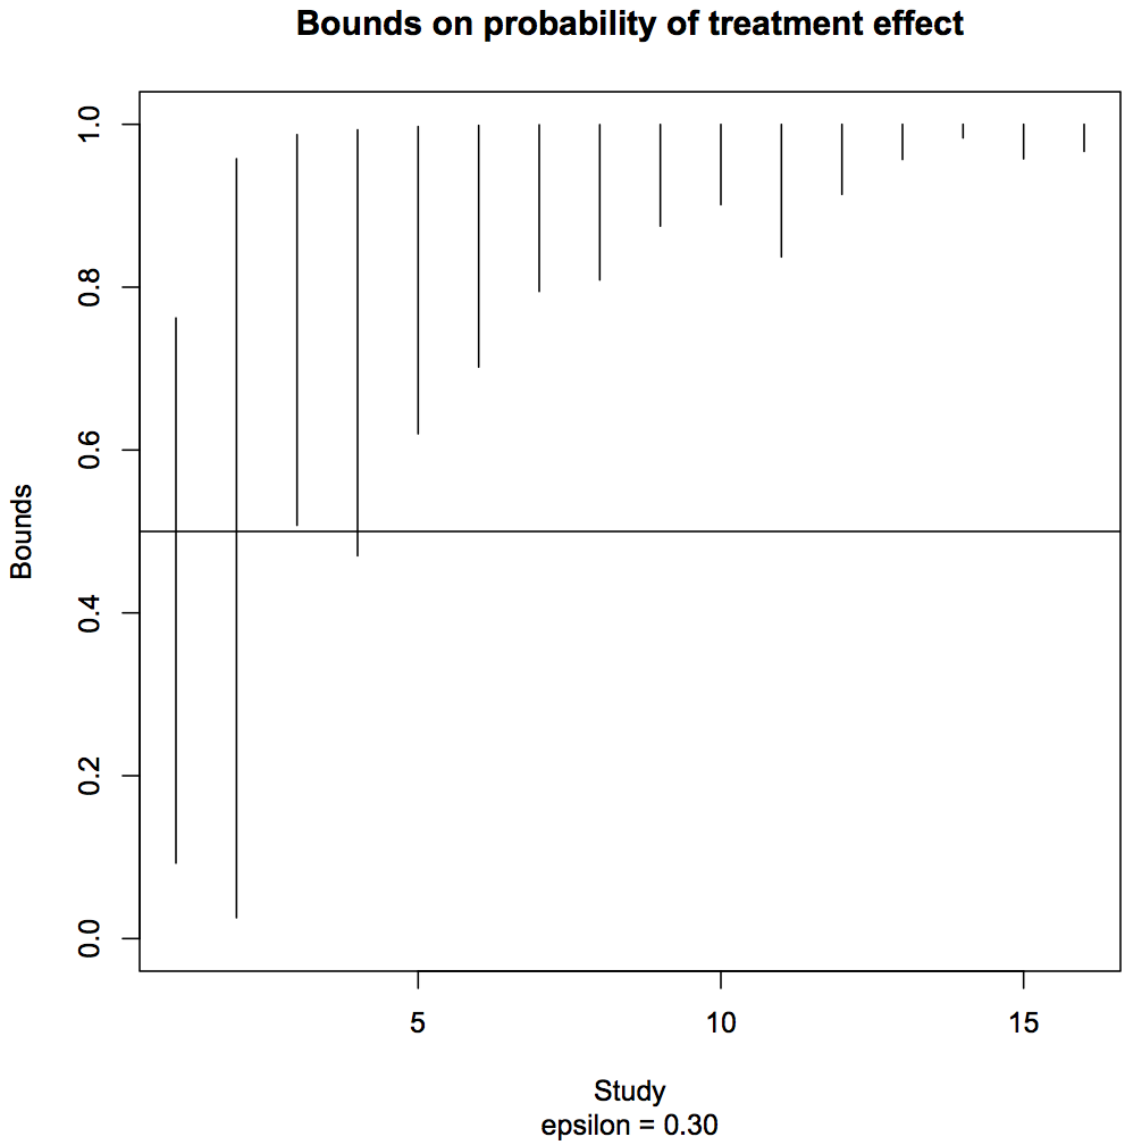

Meta-analysis simulation study

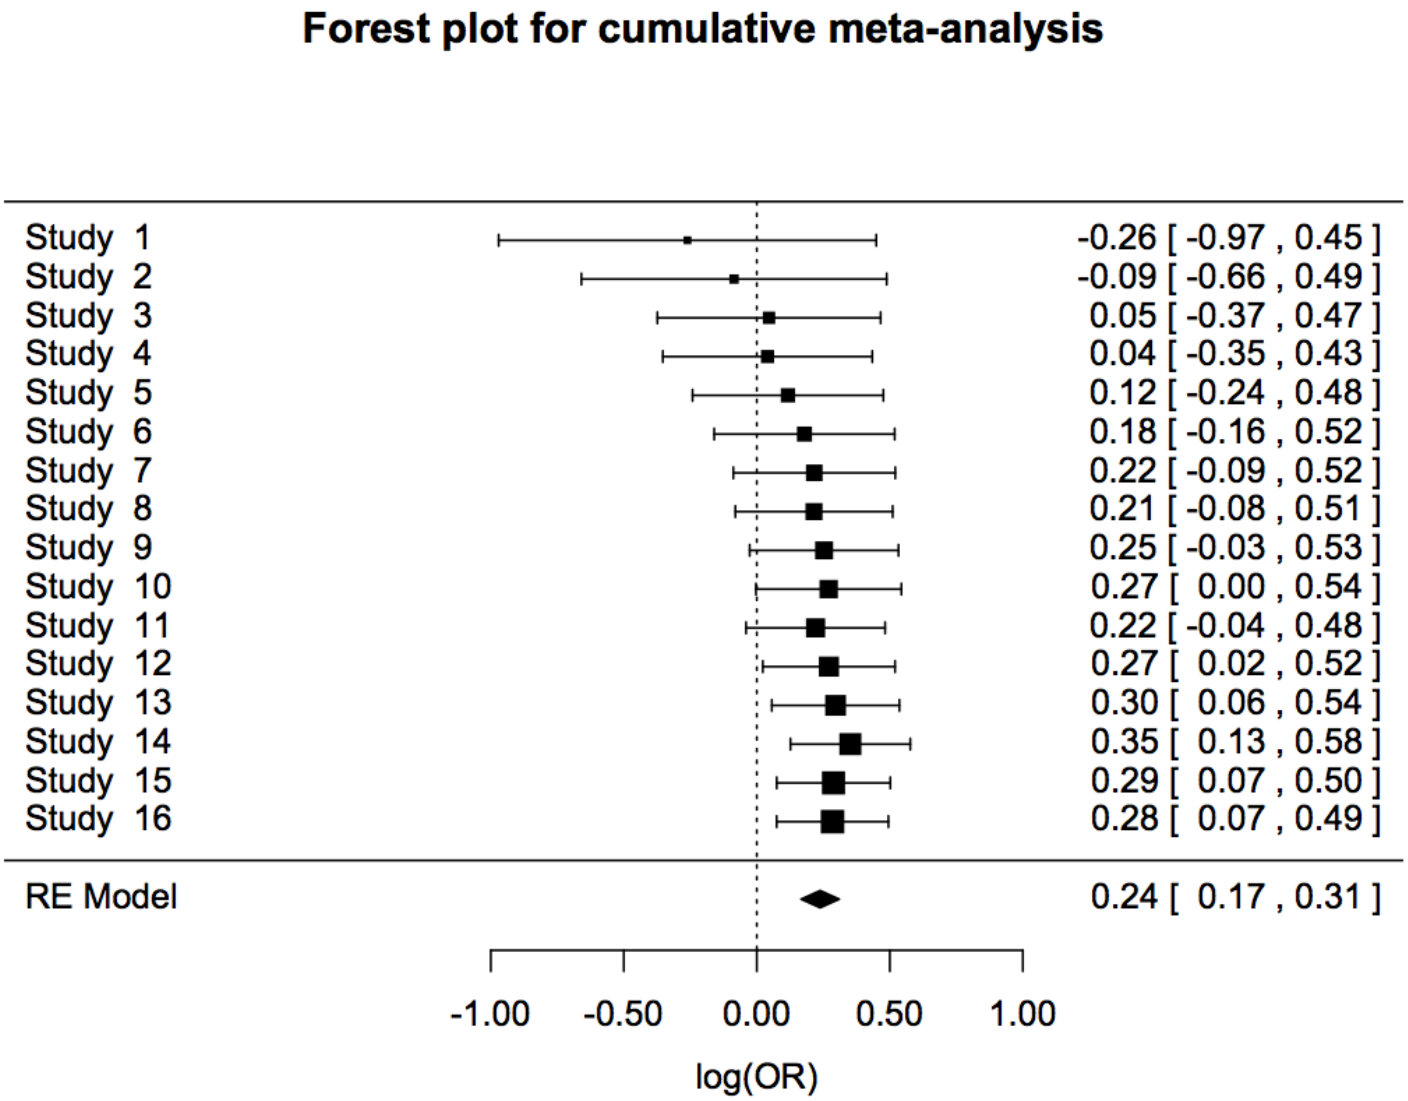

Meta-analysis simulation study

- Through the bounds that we obtain on our posterior probability of treatment effect, we see that even under the worst choice of prior, we would need a very high cutoff for probability of treatment effect in order to conclude that all 16 studies were necessary.
- We notice that as we increase our epsilon, we do not notice much of a change in the bounds.
- Likely caused by the size of each study: at 5 studies, we have a cumulative sample size of 500 subjects in each arm.

Conclusion

- For fixed effect models, we can use the naive BIC approximation for our marginal likelihood and it performs adequately.
- For random effect models, the problem becomes more complicated, but it appears that the BIC value from fitting a linear mixed model can be used.
- The $\epsilon-$ contaminated approach provides a nice technique for managing the effect of our prior distribution on the posterior inference.
- Added advantage: nice graphical methods for determining stopping rules.

Future work

- Use path sampling method to approximate the marginal likelihood.
- Path sampling is a Monte Carlo computation approach, frequently use in Bayesian statistics as an integration technique used to estimate marginal likelihoods.
- Also investigate better asymptotic approximations based on ML estimates as discussed in Pauler et al. (1999).