An Introduction to Using WinBUGS for Cost-Effectiveness Analyses in Health Economics

Dr. Christian Asseburg

Centre for Health Economics
University of York, UK

ca505@york.ac.uk

Part 2

Meta-Analyses and Mixed Treatment Comparisons

Meta-Analyses and Mixed Treatment Comparisons

- Meta-Analyses
  - Baseline and treatment effects
  - Fixed versus random baseline
  - Fixed versus random treatment effects
  - Choice of scales
  - WinBUGS examples
- Mixed treatment comparisons
- Hierarchical models

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk

Evidence synthesis

- Often data from different sources must be combined to inform a cost-effectiveness model.
- The evidence synthesis component of a model relates the various pieces of data to the relevant unknown parameters.
- Different data often share the same form – an evidence model that combines these data into estimates of an underlying parameter is called a **meta-analysis**.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 1

- Want to model a disease that can a) infect a patient and b) be dormant or break out in an infected patient.
- Evidence on
  - Rates of infection
  - Probability of symptoms in infected patients
  - Overall incidence of symptoms in population
- Need a hierarchical model that can include all three types of data

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

- Want to model a treatment agent that can cure a disease
- Evidence on
  - Risk ratios from *n* RCTs that compare the agent to placebo
  - Rate of healing in untreated target population
- Need a meta-analysis of the RCTs and combine this with external baseline

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Patient-level trial data from a RCT that compared five different treatment agents
- Evidence on
  - Success of treatment
  - Patient-level covariates
- Need a regression model to explain treatment success by patient covariates
- Could extract overall odds ratios for the treatments, as this is a RCT

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Model design

- Bayesian models are very flexible – all the previous examples can be addressed.
- There is usually more than one “correct” model
- Model selection is covered in the next session. In this talk I introduce OpenBUGS models that could be applied to the previous settings.
- For illustrative purposes, I write equations in OpenBUGS syntax

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 1

- Probability of infection `p`
- Probability of symptoms given infection `s`
- Probability of symptoms `p*s`

Imagine a treatment that reduces infection rate – how will it impact on symptoms?

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 1

- Probability of infection `p`
- Probability of symptoms given infection `s`
- Probability of symptoms `p*s`

Data: Let's assume that we have the following.

Screening study: 550 out of 1421 were infected
Symptoms given infection: 200 in screening study
Incidence of symptoms: logOR mean -1.7, sd 0.1 (independent study)

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 1

Data: Let's assume that we have the following.

Screening study: 550 out of 1421 were infected
Symptoms given infection: 200 in screening study
Incidence of symptoms: logOR mean -1.7, sd 0.1 (independent study)

```
# First attempt
model {
  nI ~ dbin(p, N)
  nS ~ dbin(s, nI)
  lor ~ dnorm(orm, orp)
  lor <- log(p*s/(1-p*s))
}
```

```
# Data
list(
  nI=550, nS=200, N=1421,
  orm= -1.7, orp=100
)
```


Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 1

When I hit "Compile" OpenBUGS says:

multiple definitions of node **lor**

Reason: **lor** is on the LHS twice, we can only define it once.

```
# First attempt
model {
   nI ~ dbin(p, N)
   nS ~ dbin(s, nI)
   lor ~ dnorm(orm, orp)
   lor <- log(p*s/(1-p*s))
}
```

```
# Data
list(
   nI=550, nS=200, N=1421,
   orm= -1.7, orp=100
)
```


Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

"Directed Acyclical Graphs"

OpenBUGS cannot sample posteriors for every Bayesian model. The model has to be in a class of models that are called **"directed acyclical graphs" (DAG)**.

The code below is not a DAG because the node **`lor`** is defined twice. However, in this case there is a work-around because the normal distribution is symmetric.

```
# First attempt
model {
   nI ~ dbin(p, N)
   nS ~ dbin(s, nI)
   lor ~ dnorm(orm, orp)
   lor <- log(p*s/(1-p*s))
}
```

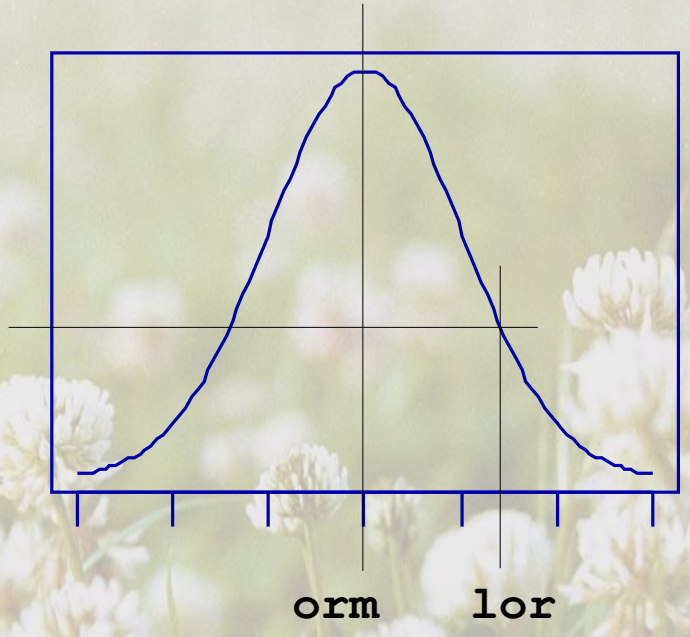

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 1

Reason: **`lor`** is on the LHS twice, we can only define it once.

Fix: We can "flip" the first definition of **`lor`**.

```
orm ~ dnorm(lor, orp)
```

```
# Second attempt
model {
  nI ~ dbin(p, N)
  nS ~ dbin(s, nI)
  orm ~ dnorm(lor, orp)
  lor <- log(p*s/(1-p*s))
}
```

```
# Data
list(
  nI=550, nS=200, N=1421,
  orm= -1.7, orp=100
)
```


Dr Christian Asseburg
University of York, UK ca505@york.ac.uk

Example Setting 1

When I hit "Compile", OpenBUGS complains

made use of undefined node **p**

Reason: **I forgot the priors**!

```
# Second attempt
model {
   nI ~ dbin(p, N)
   nS ~ dbin(s, nI)
   orm ~ dnorm(lor, orp)
   lor <- log(p*s/(1-p*s))
}
```

```
# Data
list(
   nI=550, nS=200, N=1421,
   orm= -1.7, orp=100
)
```


Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 1

Reason: **I forgot the priors**!

Fix: Make up priors.

If there is information from some other relevant study (maybe on a similar disease), we could derive vague priors from it.

Otherwise, use non-informative priors and test that they are non-informative.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 1

Fix: Use non-informative priors and test that they are non-informative.

```
# Third attempt
model {
  # priors
  p ~ dbeta(1,1)
  s ~ dbeta(1,1)
  # model
  nI ~ dbin(p, N)
  nS ~ dbin(s, nI)
  orm ~ dnorm(lor, orp)
  lor <- log(p*s/(1-p*s))
}
```

```
# Data
list(
  nI=550, nS=200, N=1421,
  orm= -1.7, orp=100
)
```


Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 1

Model compiled!

I set monitor nodes on **p**, **s** and **lor**.

```
# Third attempt
model {
  # priors
  p ~ dbeta(1,1)
  s ~ dbeta(1,1)
  # model
  nI ~ dbin(p, N)
  nS ~ dbin(s, nI)
  orm ~ dnorm(lor, orp)
  lor <- log(p*s/(1-p*s))
}
```

These are the results:

lor mean -1.77, sd 0.06.
95%-CI [-1.89, -1.65]

p mean 0.39, sd 0.012.
95%-CI [0.37, 0.42]

s mean 0.37, sd 0.017.
95%-CI [0.34, 0.41]

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

- 25 RCTs compare agent X to placebo (in patient populations that are not identical to target population)
- Each RCT reports
  - Number of patients in treatment arm and in control arm (ITT).
  - Outcome “patient cured after 2 months”.
- From an observational study of the target population we have an estimate of probability of cure after 2 months.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Design decisions:

- What to use the RCT data for... Treatment effect estimate only? Both treatment and baseline?
  - If we think that the RCT baselines do not apply to the target population, we can extract just the treatment effects.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Design decisions:

- RCT data: Fixed or random treatment effect?
  - With 25 RCTs there is enough information to model random treatment effects. Do we consider the populations in each trial to be distinct?
  - If we model a random treatment effect, then how does this apply to our target population?
    - We could use the “average” random treatment effect...
    - ...or we could treat the target equivalent to a RCT population and apply a random effect (“predictive distribution”).

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Design decisions:

- RCT data: Fixed or random baseline?
  - Again, how distinct to we consider the populations in the individual RCTs to be from each other?
  - In terms of the baseline, if we decide to extract only the treatment effect information from the trials, it still matters what uncertainty model we choose for the baseline!

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Design decisions:

- Relation external baseline to RCT baselines?
  - Should we ignore the baseline information from the RCTs completely?
  - Or should we maybe calculate a weighted average between the “average” RCT baseline and the baseline found in the study of the target population?

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Design decisions:

- Scale on which treatment effect is modelled?
  - Probabilities and treatment effects concerning probabilities are often modelled on the **log-odds scale**, but this is more out of convenience than for any “hard” reason.
  - Bayesian models do not require normality assumptions in the way classical inference would.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Design decisions:

- What to use the RCT data for... Treatment effect estimate only? Both treatment and baseline?
- RCT data: Fixed or random treatment effect?
- RCT data: Fixed or random baseline?
- Relation external baseline to RCT baselines?
- Scale on which treatment effect is modelled?

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Design decisions made:

- Use RCTs for treatment effect only
- RCTs random treatment effect, use PD
- RCTs random baseline
- Use only external baseline
- Treatment effect on log-odds probability scale

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Let's write the RCT bit first.

```
# First attempt
model {
  # RCT priors
  basesd ~ dunif(0,2)
  tesd   ~ dunif(0,2)
  baseprec<-pow(basesd,-2)
  teprec  <-pow(tesd,  -2)
  B ~ dnorm(0,.0001)
  T ~ dnorm(0,.0001)
...
```

- RCTs: Random baseline, random treatment effects
- Probabilities modelled on the log-odds scale

Here I write the priors... the rest of the model is shown on the next slide.

Remember to use precision with **dnorm**, not variance!

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Let's write the RCT bit first.

```
# First attempt
mo
...
  # RCT model
  for (i in 1:N) {
   b[i]~dnorm(B,baseprec)
   t[i]~dnorm(T,teprec)
   for (j in 1:2) {
    mu[i,j]<-b[i]+equals(j,2)*t[i]
    logit(p[i,j])<-mu[i,j]
    r[i,j]~dbin(p[i,j],n[i,j])
   }
  }
...
```

I am using two-dimensional arrays in this example to show you how to specify array data.

Note a useful function **`equals(a,b)`**:
- =1 if a=b
- =0 otherwise

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Let's write the RCT bit first.

```
# First attempt
mo
...
  # model
...
  # target population
  oC ~ dbin(pbase,nC)
  logit(pt)<-logit(pbase)+tstar
  tstar~dnorm(T,teprec)
  # target priors
  pbase ~ dbeta(1,1)
  }
...
```

Next we need the external baseline, and we need to predict the treatment effect in the target population.

From an observational study we have 369 cured out of 747.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

The model is getting complex – and difficult to view on the screen all at once!

```
# First attempt
model {
  # RCT priors
  basesd ~ dunif(0,2)
  tesd   ~ dunif(0,2)
  baseprec<-pow(basesd,-2)
  teprec  <-pow(tesd,  -2)
  B ~ dnorm(0,.0001)
  T ~ dnorm(0,.0001)
  # RCT model
  for (i in 1:N) {
   b[i]~dnorm(B,baseprec)
   t[i]~dnorm(T,teprec)
   for (j in 1:2) {
    mu[i,j]<-b[i]+equals(j,2)*t[i]
    logit(p[i,j])<-mu[i,j]
    r[i,j]~dbin(p[i,j],n[i,j])
   }
  }
```

```
  # target population
  oC ~ dbin(pbase,nC)
  logit(pt)<-logit(pbase)+tstar
  tstar~dnorm(T,teprec)
  # target priors
  pbase ~ dbeta(1,1)
}
```

Example Setting 2

Now we need to enter the data.

```
#Data
list(N=25, oC=369, nC=747)

r[,1]  r[,2]  n[,1]  n[,2]
13     20     72     53
26     21     72     73
...
11     19     77     77
END

# put an empty line after "END"
```

We can enter array-shaped data in this syntax. Note that the row headers must be unique and that the first index (row index) must be empty. In another example you might have had something like

```
age[]   sex[]   duration[]   cured[]
```


Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Now we need to enter the data.

```
#Data
list(N=25, oC=369, nC=747)

r[,1]  r[,2]  n[,1]  n[,2]
13     20     72     53
26     21     72     73
...
11     19     77     77
END

# put an empty line after "END"
```

To load these data, note that you have to click the "Load data" button twice. Once you highlight the keyword **`list`** and click "Load data", then you highlight the header row of the table and click "Load data" again.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

Here's an equivalent way of entering the array-shaped data.

```
#Data
list(N=25, oC=369, nC=747)

r[,1]  r[,2]  n[,1]  n[,2]
13     20     72     53
26     21     72     73
...
```

The first version is called "rectangular", the second version is called "S-Plus" data.

```
#Data
list(N=25, oC=369, nC=747,
r=structure(.Data=c(13,20,26,21,...),
            .Dim =c(25,2)),
n=structure(.Data=c(72,53,72,73,...),
            .Dim =c(25,2))
)
```

Note: In S-Plus, the structure on the left would result in a different array because OpenBUGS fills the array row by row.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 2

In the second practical, let's explore some of the design choices that were made at the beginning of this example.

There, you will also learn how to specify initial values for the chains (instead of using the “auto init” button).

For now, here are the results of fitting this model.

	mean	sd	MC_error	val2.5pc	median	val97.5pc
B	-1.098	0.06528	5.344E-4	-1.228	-1.098	-0.9692
T	0.3634	0.08637	6.927E-4	0.1889	0.3641	0.5292
basesd	0.2204	0.06784	5.986E-4	0.09961	0.2165	0.368
pbase	0.4939	0.01828	1.547E-4	0.4586	0.4939	0.5301
pt	0.5824	0.07438	6.348E-4	0.4244	0.5856	0.7268
tesd	0.277	0.09401	8.774E-4	0.1024	0.2719	0.4763

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Patient-level trial data from a RCT that compared five different treatment agents
- Evidence on
  - Success of treatment (yes/no)
  - Patient-level covariates
- Need a regression model to explain treatment success by patient covariates
- Could extract overall odds ratios for the treatments, as this is a RCT

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Design decisions:
  - "Baseline" and treatment effects, or simply model each treatment individually?
  - Include some or all covariates?
  - Include covariate/treatment interactions?
  - Type of regression model: Logistic regression?
- More complicated models require more data.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Let's start by looking at the data.

```
#Data
list(N=1500, NT=5)
t[]   age[] sex[] HbA1c[]  BMI[]    BP[]     s[]
2     -1.31    1  1.343    0.5683   0.5552   1
2     -1.483   0  0.9778   0.2377   -1.2     1
2     1.389    0  -0.213   0.2433   -1.397   0
1     0.9683   1  -0.3772 -0.4553 -0.736   0
5     -0.9591  1  0.5816   -1.324   0.9554   1
2     0.7234   0  -1.323   -0.1137 0.6394   0
.        .     .     .         .        .      .
.        .     .     .         .        .      .
.        .     .     .         .        .      .
.        .     .     .         .        .      .
```


Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Let's start by looking at the data.
- We have 1500 patients (rows of data), in particular we have 300 patients for each treatment.
- It seems feasible to employ 5 covariates to explain 300 observations in each treatment group, so let's model treatment/covariate interactions... we can always prune the model if some of these seem too insignificant.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Should we pick one treatment as “baseline” and model treatment effects relative to it?
  - Advantage: All patients inform the baseline estimate
- Or should we model each treatment group individually?
  - Advantage: Simpler model
- Other approaches are possible, e.g. we could model treatment effects for all 5 treatments and use some kind of average as baseline.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- For now, let's pick treatment 1 as a baseline treatment.
- We'll model P(treatment success) as a fixed-effect logistic regression on the covariates.
- Treatment effects are modelled as additive to the linear coefficients (on the logit scale) – i.e. multiplicative on the odds ratios.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- First, let's code the baseline prediction.

```
model {
  for (i in 1:N) {
    logit(p[i])<-I+
            a[1]*age[i]+
            a[2]*sex[i]+
            a[3]*HbA1c[i]+
            a[4]*BMI[i]+
            a[5]*BP[i]
    s[i]~dbern(p[i])
  }
...
```

dbern(p) is the Bernoulli distribution.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Second, the regression coefficients depend on the treatment indicator, so add indices.

```
model {
  for (i in 1:N) {
    logit(p[i])<-I[t[i]]+
            a[1,t[i]]*age[i]+
            a[2,t[i]]*sex[i]+
            a[3,t[i]]*HbA1c[i]+
            a[4,t[i]]*BMI[i]+
            a[5,t[i]]*BP[i]
    s[i]~dbern(p[i])
  }
...
```


Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Now, we link the regression coefficients to the baseline (treatment 1).

```
...
  for (j in 1:NT) {
    I[j]<-bI+tI[j]
    for (k in 1:5) {
      a[k,j]<-ba[k]+ta[k,j]
    }
  }
...
```

bI is the baseline intercept, **tI** is the treatment effect on this, resulting in intercept **I**.
ba is the baseline coefficient, **ta** is the associated treatment effect, resulting in **a**.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Finally, we need some priors.

```
...
  tI[1]<-0
  for (k in 1:5) {
    ta[k,1]<-0
    ba[k]~dnorm(0,0.1)
    for (j in 2:NT) {
      ta[k,j] ~ dnorm(0,0.1)
    }
  }
  bI~dnorm(0,0.0001)
  for (j in 2:NT) {
    tI[j]~dnorm(0,0.0001)
  }
}
```

Note that in terms of the log-odds treatment effect, we would consider a value of 5 quite extreme – so the prior variance should not be 10.000 like for the base-line log odds.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Here's the whole code on one page.

```
model {
  for (i in 1:N) {
    logit(p[i])<-I[t[i]]+
            a[1,t[i]]*age[i]+
            a[2,t[i]]*sex[i]+
            a[3,t[i]]*HbA1c[i]+
            a[4,t[i]]*BMI[i]+
            a[5,t[i]]*BP[i]
    s[i]~dbern(p[i])
  }
  for (j in 1:NT) {
    I[j]<-bI+tI[j]
    for (k in 1:5) {
      a[k,j]<-ba[k]+ta[k,j]
    }
  }
  for (k in 1:5) {
    ta[k,1]<-0
    ba[k]~dnorm(0,0.1)
    for (j in 2:NT) {
      ta[k,j] ~ dnorm(0,0.1)
    }
  }
  bI~dnorm(0,0.0001)
  tI[1]<-0
  for (j in 2:NT) {
    tI[j]~dnorm(0,0.0001)
  }
}
```


Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Some observations:
  - Running the sampler takes much longer, for two reasons...
    - There's a lot of data.
    - The model is quite complex.
  - Auto-correlation within the chains is also quite high.
  - If you set “extreme” priors (like a variance of 10,000 for the treatment effect log-odds), some of the OpenBUGS calculations can fail and result in **trap messages**.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Anyway, here's some output.
- Because the sampler is so slow, I've calculated only 500 draws (+200 discarded for burn-in).
- In a real-life application I would leave OpenBUGS running overnight or simplify the model.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- First, the intercept coefficients...

	mean	sd	MC_error	val2.5pc	median	val97.5pc
bl	5.17	0.7778	0.06345	3.749	5.118	6.866
tl[2]	6.103	1.73	0.09475	2.565	6.119	9.411
tl[3]	-1.082	0.9616	0.07168	-3.027	-1.015	0.6832
tl[4]	1.463	1.287	0.09108	-1.093	1.471	3.949
tl[5]	6.239	1.672	0.1091	3.304	6.188	9.896

It seems that treatments 2 and 5 have a significantly positive log-odds effects in the intercept coefficient.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- Now, the covariate coefficients for the baseline:

	mean	sd	MC_error	val2.5pc	median	val97.5pc
ba[1]	-11.97	1.395	0.1246	-14.83	-11.98	-9.277
ba[2]	0.2867	0.7306	0.04471	-1.166	0.2857	1.652
ba[3]	1.097	0.3551	0.01799	0.3961	1.096	1.82
ba[4]	0.3997	0.4073	0.01963	-0.3846	0.3889	1.211
ba[5]	0.06609	0.3849	0.01168	-0.6301	0.0571	0.8241

Treatment success in baseline (i.e. treatment 1) decreases with “Age” and increases with “HbA1c” value.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- ...and the covariate/treatment interactions...

	mean	sd	MC_error	val2.5pc	median	val97.5pc
ta[1,2]	-5.738	2.412	0.1319	-10.15	-5.71	-0.7207
ta[1,3]	7.12	1.53	0.1331	4.117	7.131	10.34
ta[1,4]	-0.4302	2.036	0.1605	-4.21	-0.4426	3.463
ta[1,5]	-5.672	2.267	0.1447	-10.4	-5.585	-1.489
ta[2,2]	0.6078	1.255	0.04478	-1.845	0.6125	3.209
ta[2,3]	-1.356	0.8883	0.0503	-3.11	-1.346	0.3319
ta[2,4]	0.605	1.119	0.0555	-1.428	0.5618	2.916
ta[2,5]	-1.245	1.094	0.04953	-3.327	-1.252	0.9531
ta[3,2]	-0.1029	0.7076	0.02315	-1.474	-0.109	1.328
ta[3,3]	-0.7986					
ta[3,4]	-0.3381					
ta[3,5]	0.5762					
ta[4,2]	0.7662					
ta[4,3]	-0.2042					
...						

If you read the whole table, you can see that the covariate/treatment interactions are significant for “age” in T2, “age” and “BP” in T3, “BP” in T4, and “age” and “BMI” in T5.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- How to interpret these results:
  - For an overall effect of each covariate on treatment success, we need to combine the baseline (T1) and treatment effect parameters – so the previous table was a bit misleading.
  - From a Bayesian point of view, it is irrelevant whether a parameter is significantly different from zero. All parameters contribute to the predicted probability of success, and even if they seem insignificant, their contribution is not equal to zero.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Example Setting 3

- In a real-life example, we could explain these results to physicians and see if they agree that e.g. age should have an influence on treatment success, or if treatment 5 really should perform better in patients with high BMI.
- Experts may also be able to confirm that some covariates are believed to be irrelevant to some treatments.
- We could simplify the model in accordance with clinical opinion.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Choice of scales

- In the previous example, we transformed a probability onto the log-odds scale, using the logistic link function

$$y = \log\left(\frac{p}{1-p}\right)$$

- *y* can take any value on the real line. It is also usually assumed to be normally distributed.
- These simplifying assumptions are not required in Bayesian models, so we could model probabilities on other scales.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Choice of scales

For example, instead of the previous model:

$$y = \log\left(\frac{p}{1-p}\right) = a_0 + a_1 \cdot x_1 + a_2 \cdot x_2 + \ldots + a_n \cdot x_n$$

we could just as well have written a linear regression directly on the probability *p*:

$$p = \min\left(1, \max\left(b_0 + b_1 \cdot x_1 + b_2 \cdot x_2 + \ldots + b_n \cdot x_n\right)\right)$$

Note that min and max are required to ensure 0≤p≤1.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Mixed treatment comparisons

- A mixed treatment comparison (MTC) is an evidence synthesis model in which more than two treatments are compared.

- In some trials, more than two therapies are compared directly.
- In other MTCs, a “network of evidence” is created by bringing together many trials, each of which compares some therapies.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Mixed treatment comparisons

- To enter the trial data we need to specify not only the trial results for each arm, but also indicate which treatment was given in each arm.
- For example, we can enter the data like this:

s[]	m[]	t[]	r[]	n[]	b[]
1	1	1	16	66	1
1	2	4	34	70	1
2	1	2	19	56	2
2	2	3	30	54	2
2	3	4	18	72	2
.	.	.	.	.	.

s - study number
m - arm number
t - treatment on arm
r - trial outcome
n - sample size
b - baseline indicator

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Mixed treatment comparisons

- If all trials are 2-arm trials, we could enter a random-baselines, random-effects model like this:

```
model {
  for (i in 1:N) {
    logit(p[i]) <- mu[s[i]]+
      delta[i]*(1-equals(t[i],b[i]))
    r[i] ~ dbin(p[i],n[i])
    delta[i] ~ dnorm(md[i],tau)
    md[i] <- d[t[i]]-d[b[i]]
  }
  d[1] <- 0
  # add priors for:
  #   study baselines
  #   treatment effects
  #   variances
```

```
s
1
1
2
2
2
.
```

s - study number
m - arm number
t - treatment on arm
r - trial outcome
n - sample size
b - baseline indicator

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Mixed treatment comparisons

- If some arms are 3-arm trials, we have to add a variance adjustment.

s	- study number
m	- arm number
t	- treatment on arm
r	- trial outcome
n	- sample size
b	- baseline indicator

```
model {
  sw[1] <- 0
  for (i in 1:N) {
    logit(p[i]) <- mu[s[i]]+
      delta[i]*(1-equals(t[i],b[i]))
    r[i] ~ dbin(p[i],n[i])
    delta[i] ~ dnorm(md[i],taud[i])
    taud[i] <- tau*(1+equals(m[i],3)/3)
    md[i] <- d[t[i]]-d[b[i]]+
      equals(m[i],3)*sw[i]
  }
  for (i in 2:N) {
    sw[i] <- (delta[i-1]-d[t[i-1]]+
      d[b[i-1]])/2
  }
```

```
s
1
1
2
2
2
.
```

Hierarchical models

- A hierarchical model is a model in which the unknown parameters have a hierarchical order.
- For example, in a fixed-effects, fixed-baseline model, there are two parameters (mean baseline and mean effect). They share the same level in the hierarchy – both are used to explain trial data directly.
- But in a random-effects model, the study treatment effects in each trial are random and the underlying treatment effect is also random. **The underlying treatment effect defines the random study effects, i.e. it is above the study effects in the hierarchy.**

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Hierarchical models

- Random-effects models are a case of hierarchical models.
- If you think of all the variables with probability distributions assigned to them:
  - Prior distributions contain no unknown variables, they are at level “1” in the hierarchy.
  - Unknown variables that are defined with reference to variables with priors are one level up in the hierarchy.
  - Some variables do not fall neatly into a particular level in the hierarchy.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Hierarchical models

- Random-effects models are a case of hierarchical models.
- In some models, some variables do not fall neatly into a particular level in the hierarchy.
- Usually it's of no practical relevance whether a model is hierarchical.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics

Summary

- In example 1 we modelled two probabilities: infection and symptoms given infection.
- In example 2 we combined evidence from 25 RCTs on the effectiveness of a treatment agent compared to placebo.
- In example 3 we studied how patient-level covariates interact with the success of 5 treatments in a RCT.
- Finally we saw an example of a mixed treatment comparison.
- OpenBUGS can be used with many different models.

Dr Christian Asseburg
University of York, UK ca505@york.ac.uk
CHE Centre For Health Economics