

glmmEP: fast and accurate binary response mixed model analysis via expectation propagation

J.C.F. YU AND M.P. WAND

2nd February, 2018

1 Introduction

The `glmmEP` package supports binary response mixed model analysis based on an *expectation propagation* approximation to the log-likelihood. Full details of the methodology, as well as theoretical back-up, are in the article by Hall, Johnstone, Ormerod, Wand & Yu (2018).

The package's central function is `glmmEP()`. In this vignette we explain how to set up the input data matrices for `glmmEP()` and then obtain inferential summaries of the fit. We first use a simulated data set, corresponding to a simulation study described in Hall *et al.* (2018). We then conduct some analyses involving a dataset that arose for a contraception use study.

2 Nature of the Computations and Some Caveats

Exact likelihood-based inference for binary response mixed models involves numerical integration with dimension matching that of the random effects vectors. Typically this dimension is a low number between 1 and 5, with 1 (corresponding to random intercept models) and 2 (corresponding to random slope models) being the most common. The essence of the expectation propagation approach is to replace each multivariate integral by a fixed-point iterative algorithm which, in certain cases, has closed formed updates. This has the attraction of circumventing the numerical integration requirement. On the other hand, fixed-point iterative algorithms are susceptible to breakdown. **NEXT MENTION START VALUES. SAY THAT THOROUGHLY TESTED ON SIMULATED DATA..**

Get 'em to e-mail us.
Add limitations.

3 Illustration for Simulated Data

In an R session the `glmmEP` package is loaded via the command:

```
library(glmmEP)
```

3.1 Generation and Format of the Simulated Data

The next chunk of code obtains simulated data corresponding to the simulation study in Section 4.1.2. of Hall *et al.* (2018):

```
dataObj <- glmmSimData(seed=54321) ; y <- dataObj$y ; idNum <- dataObj$idNum  
Xfixed <- dataObj$Xfixed ; Xrandom <- dataObj$Xrandom
```

The dimension for these data are:

$$\begin{aligned} m &= \text{number of groups} = 2,500, \\ n_i &= \text{number of measurements in the } i\text{th group} \\ &= \text{a randomly generated integer between 20 and 30,} \\ \sum_{i=1}^m n_i &= \text{total number of measurements} = 6,229, \\ d^F &= \text{fixed effects dimension} = 6 \\ \text{and } d^R &= \text{random effects dimension} = 2. \end{aligned}$$

3.2 Probit Mixed Model Analysis of the Data

The appropriate fitting command is:

```
fitSimulRanIntAndSpl <- glmmEP(y,Xfixed,Xrandom,idNum)
```

and takes about 20–30 seconds to fit on typical 2018 computers.

An inferential summary of the fit is obtained via:

```
summary(fitSimulRanIntAndSpl)
```

and leads to

	95% C.I low	estimate	95% C.I upp
intercept	0.13272044	0.30246606	0.4722118
x1	0.69789090	0.88043814	1.0629854
x2	-0.54163761	-0.41337685	-0.2851161
x3	-0.03726938	0.09050807	0.2182855
x4	-1.45392531	-1.31994887	-1.1859724
x5	1.04426639	1.17685492	1.3094436
sigma1	0.60352982	0.70879985	0.8324305
sigma2	0.76176324	0.94008393	1.1601476
rho12	-0.60952776	-0.44659755	-0.2474686

The first six rows of this summary table are estimates and 95% confidence intervals for the fixed effects parameters, corresponding to the intercept and the predictors and the predictors x_1, \dots, x_5 . The last three rows are estimates and 95% confidence intervals for the parameters σ_1, σ_2 and ρ_{12} within the random effects covariance matrix:

$$\begin{bmatrix} u_{i0} \\ u_{i1} \end{bmatrix} \stackrel{\text{ind.}}{\sim} N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_1^2 & \rho_{12} \\ \rho_{12} & \sigma_2^2 \end{bmatrix} \right)$$

For example, the expectation propagation-approximate estimate of the fixed effect associated with x_1 is

$$\hat{\beta}_{\sim 1} = 0.88043814 \text{ with corresponding 95\% confidence interval } (0.69789090, 1.0629854).$$

The estimate of the standard deviation of the random slope is

$$\hat{\sigma}_{\sim 2} = 0.94008393 \text{ with corresponding 95\% confidence interval } (0.76176324, 1.16701476).$$

The code

```
uHat <- fitSimulRanIntAndSpl$randomEffects
plot(uHat[,1],uHat[,2],col="dodgerblue",xlab="random intercepts predicted values",
     ylab="random slopes predicted values",bty="n",lwd=2,cex.lab=1.5,cex.axis=1.5)
abline(v=0,col="slateblue",lwd=2) ; abline(h=0,col="slateblue",lwd=2)
```

leads to the plot shown in Figure 1 concerning the expectation propagation-approximate best predictions of the random effects.

3.3 Controlling the Convergence Parameters

FILL IN HERE.

4 Analysis of Data from a Contraception Use Study

Data from the 1988 Bangladesh Fertility Survey are stored in the data frame `Contraception` within the R package `mlmRev` (Bates, Maechler and Bolker, 2014). Steele, Diamond and Amin (1996) contains details of the study and some multilevel analyses. Variables in the `Contraception` data frame include:

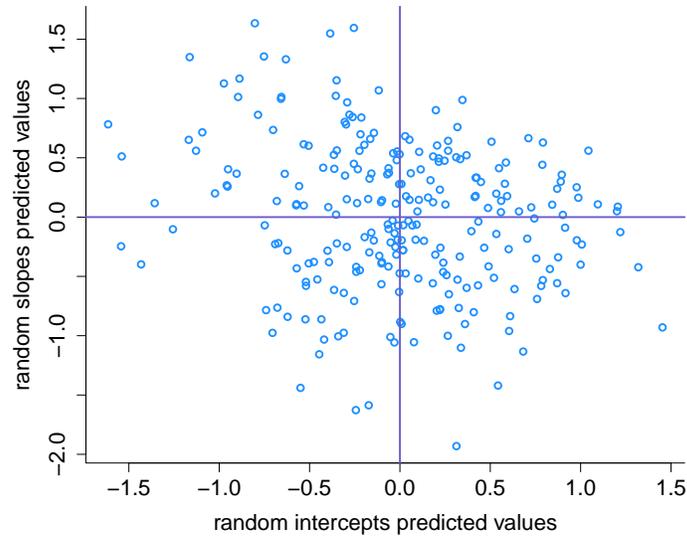


Figure 1: *Scatterplot of the expectation propagation-approximate best predictions of the random slopes and corresponding random intercepts, which are part of the fitSimulRanIntAndSlp fit object.*

use a two-level factor variable indicating whether a woman is a user of contraception at the time of the survey, with levels Y for use and N for non-use.

age age of the woman in years at the time of the survey, centred about the average age of all women in the study.

district a multi-level factor variable that codes the district, out of 60 districts in total, in which the woman lives,

urban a two-level factor variable indicating whether or not the district in which the woman lives is urban, with levels Y for urban dwelling and N for rural dwelling.

livch a four-level factor variable that indicates the number of living children of the woman, with levels 0 for no children, 1 for one child, 2 for two children and 3+ for three or more children.

The following code leads to visualisation of the data shown in Figure 2:

```
library(lattice)
colourVec <- c("forestgreen","sienna")
ContraceptionHiLivCh <- Contraception[Contraception$livch=="3+",]
figRaw <- xyplot(jitter(as.numeric(use),factor=0.5)
  ~ age|district,groups=district,data=ContraceptionHiLivCh,
  layout=c(6,10),
  xlab=list(label="age (years) centred about average",cex=1.35),
  ylab=list(label="indicator of contraception use (jittered)",cex=1.35),
  scales=list(cex=1.25),strip=FALSE,as.table=TRUE,
  key=list(title="subset of data for women with three or more living children",
  columns=2, points=list(pch=rep(1,2),col=colourVec[1:2])),
  text=list(c("rural district","urban district"),cex=1.55)),
  panel=function(x,y,subscripts,groups)
  {
    colourInd <- 3 - as.numeric(ContraceptionHiLivCh$urban[subscripts[1]])
    panel.superpose(x,y,subscripts,groups,col=colourVec[colourInd],
      pch=1,cex=0.5)
  }
```

```

    })
print(figRaw)

```

In Figure 2 each panel corresponds to a different district and each point is the `age/use` pair for a woman in that district. The `use` data are re-coded as 0 if the woman is a non-user of contraception and 1 if the woman uses contraception. Jittering has been added to these vertical axes data to aid visualisation. Lastly, Figure 2 is restricted to the subset for which `livch=3+`, namely women with three or more living children.

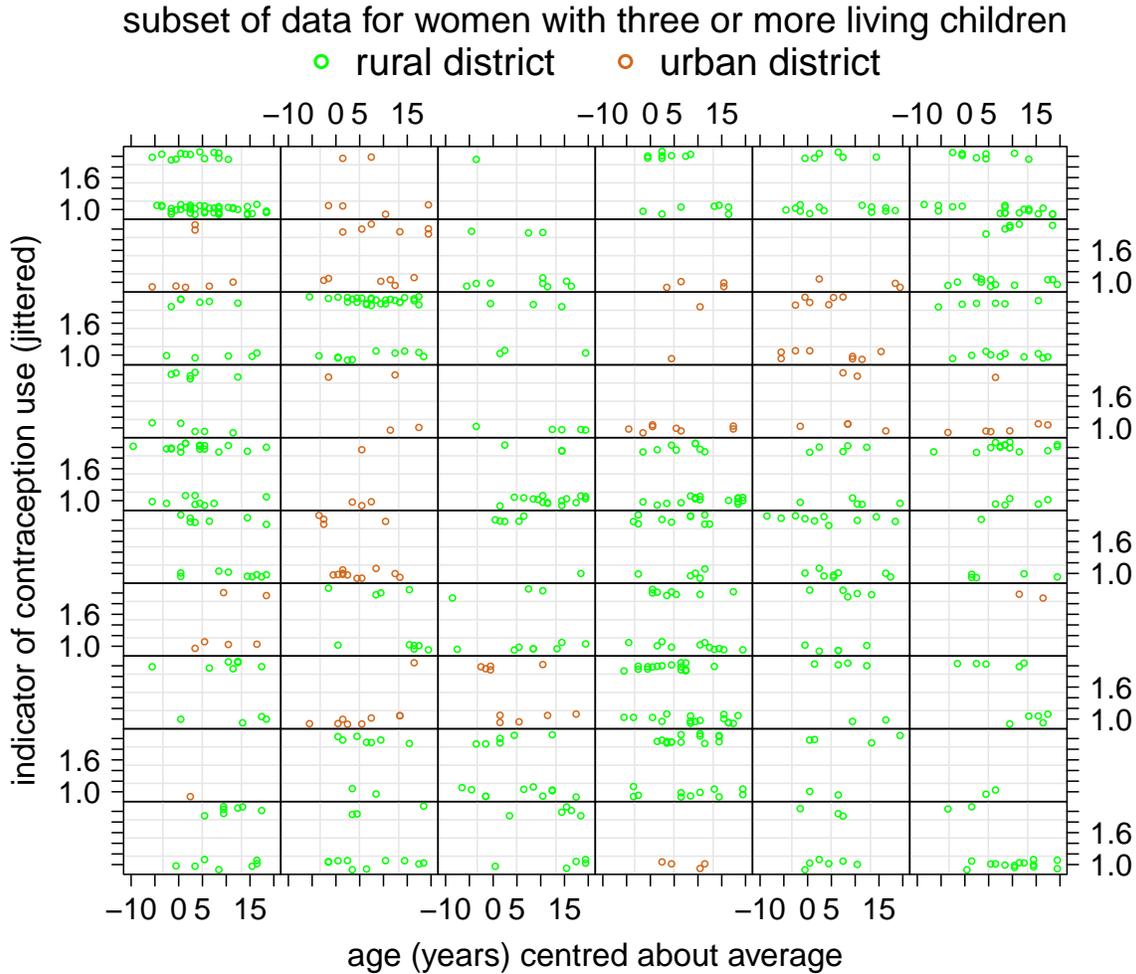


Figure 2: Visualisation of the *Contraception* data frame for the `livch=3+` subset, corresponding to women with three or more living children. Each panel is for a different district with colour-coding of the points according to whether the district is rural or urban. The indicators of contraception use values have been jittered to aid visualisation.

4.1 Random Intercepts Model

Our first analysis of these data using `glmmEP()` involves a probit mixed model with the response variable being the indicator of contraception use and predictors urban versus rural status, age and number of living children. The district is the grouping variable. The following code sets up the input data for fitting via `glmmEP()` (since the current release of `glmmEP()` does not support factor-type variables we are required to use indicator variable coding for categorical variables):

```

y <- as.numeric(as.character(Contraception$use=="Y"))
age <- Contraception$age

```

```

isUrban <- as.numeric(as.character(Contraception$urban)=="Y")
livchFactor <- Contraception$livch
livChEq1 <- as.numeric(as.character(Contraception$livch=="1"))
livChEq2 <- as.numeric(as.character(Contraception$livch=="2"))
livChGe3 <- as.numeric(as.character(Contraception$livch=="3+"))
Xfixed <- cbind(1,isUrban,age,livChEq1,livChEq2,livChGe3)
colnames(Xfixed) <- c("intercept","isUrban","age",
                    "livChEq1","livChEq2","livChGe3")
idNum <- as.numeric(as.character(Contraception$district))
Xrandom <- as.matrix(rep(1,length(y)))
colnames(Xrandom) <- "intercept"

```

Fitting is then achieved via:

```
fitContraRanInt <- glmmEP(y,Xfixed,Xrandom,idNum)
```

and the inferential summary of the model parameters is produced from the command:

```
summary(fitContraRanInt)
```

resulting in the output:

	95% C.I low	estimate	95% C.I upp
intercept	-0.9683392	-0.8077013	-0.6470634
isUrban	0.3066033	0.4491163	0.5916293
age	-0.8477430	-0.5374353	-0.2271276
livChEq1	0.4845205	0.6701782	0.8558358
livChEq2	0.6292168	0.8348053	1.0403938
livChGe3	0.6043650	0.8147953	1.0252256
sigma	0.2031284	0.2825052	0.3929001

We see from this output that each of the fixed parameters are statistically significant. For example, the coefficient of the indicator of the district being urban has an estimate of 0.4491163 and a corresponding 95% confidence interval of (0.3066033, 0.5916293). The random intercept standard deviation corresponds to the row labelled `sigma` and its estimate is 0.2825052 with a 95% confidence interval of (0.2031284, 0.3929001), indicating a significant amount of within-district correlation.

The following code:

```

hist(uHat,xlab="random intercepts predicted values",probability=TRUE,
     col="dodgerblue",breaks=15,main="",cex.lab=1.5)
abline(v=0,col="slateblue",lwd=2)

```

leads to the histogram shown in Figure 3. This is a visual summary of the expectation propagation-approximate best predictions of the random intercepts.

Lastly, we embellish Figure 2 by computing the estimated probability of contraception use curves. The code for this is:

```

ng <- 101
ageg <- seq(min(age),max(age),length=ng)
probRanIntg <- vector("list",60)
idNumOrig <- idNum
idNum <- match(idNumOrig,unique(idNumOrig))
distSttInds <- c(1,(1:length(y))[diff(idNum)==1] + 1)
for (i in 1:60)
  probRanIntg[[i]] <- pnorm(betaHat[1]+uHat[i]+betaHat[2]*ageg
                          + betaHat[3]*isUrban[distSttInds[i]]+betaHat[6])
figFitRanInt <- xyplot(jitter((as.numeric(use)-1),factor=0.5)
                    ~ age/district,groups=district,data=ContraceptionHiLivCh,
                    layout=c(6,10),
                    xlab=list(label="age (years) centred about average",cex=1.35),
                    ylab=list(label="indicator of contraception use (jittered)",cex=1.35),

```

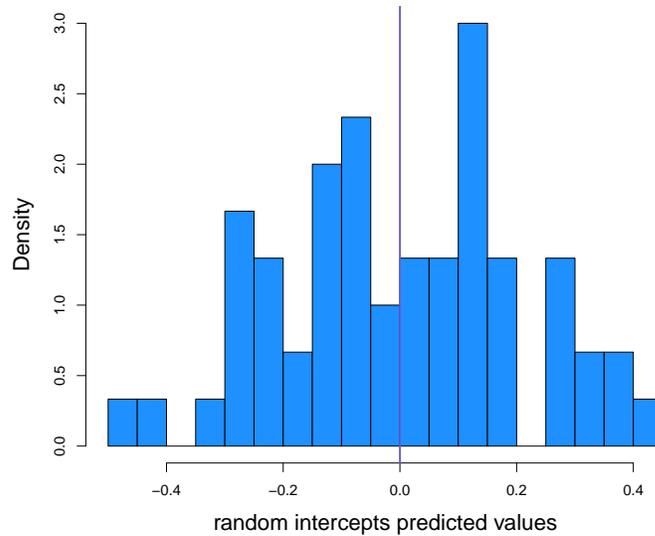


Figure 3: Histogram of the expectation propagation-approximate best predictions of the random intercepts, which are part of the `fitContracRanInt` fit object.

```
scales=list(cex=1.25),strip=FALSE,as.table=TRUE,
key=list(title="subset of data for women with three or more living children",
         columns=2,
         points=list(pch=rep(1,2),col=colourVec[1:2],lwd=rep(2,2)),
         text=list(c("rural district","urban district"),cex=1.55)),
panel=function(x,y,subscripts,groups)
{
  iDistrict <- panel.number() ; panel.grid()
  colourInd <- 3 - as.numeric(ContraceptionHiLivCh$urban[subscripts[1]])
  panel.superpose(x,y,subscripts,groups,
                 col=colourVec[colourInd],pch=1,cex=0.5)
  panel.xyplot(ageg,probRanIntg[[iDistrict]],col="blue",lwd=2,type="l")
}
```

Note that the calculation of the ordinate vectors in `probRanIntg` is simplified by the fact that Figure 2 is restricted to the subset of women with three or more living children. The plot that results from this code is shown in Figure 4. It shows that the estimated probability of contraception use increases steeply with age about 5 years either side of the average age. Differences between the districts and those with rural and urban status is difficult to discern visually. However the confidence intervals given earlier in this subsection show that there are, indeed, significant differences.

4.2 Random Intercepts and Slopes Model

We now extend the model to allow the slope of the urban district indicator to be random. The only change in the design matrix set-up is that `Xrandom` is now:

```
Xrandom <- cbind(1,isUrban)
colnames(Xrandom) <- c("intercept","isUrban")
```

With this new version of `Xrandom` we call `glmmEP()` as before using:

```
fitContracRanIntAndSlp <- glmmEP(y,Xfixed,Xrandom,idNum)
```

The inferential summary from the command:

```
summary(fitContracRanIntAndSlp)
```

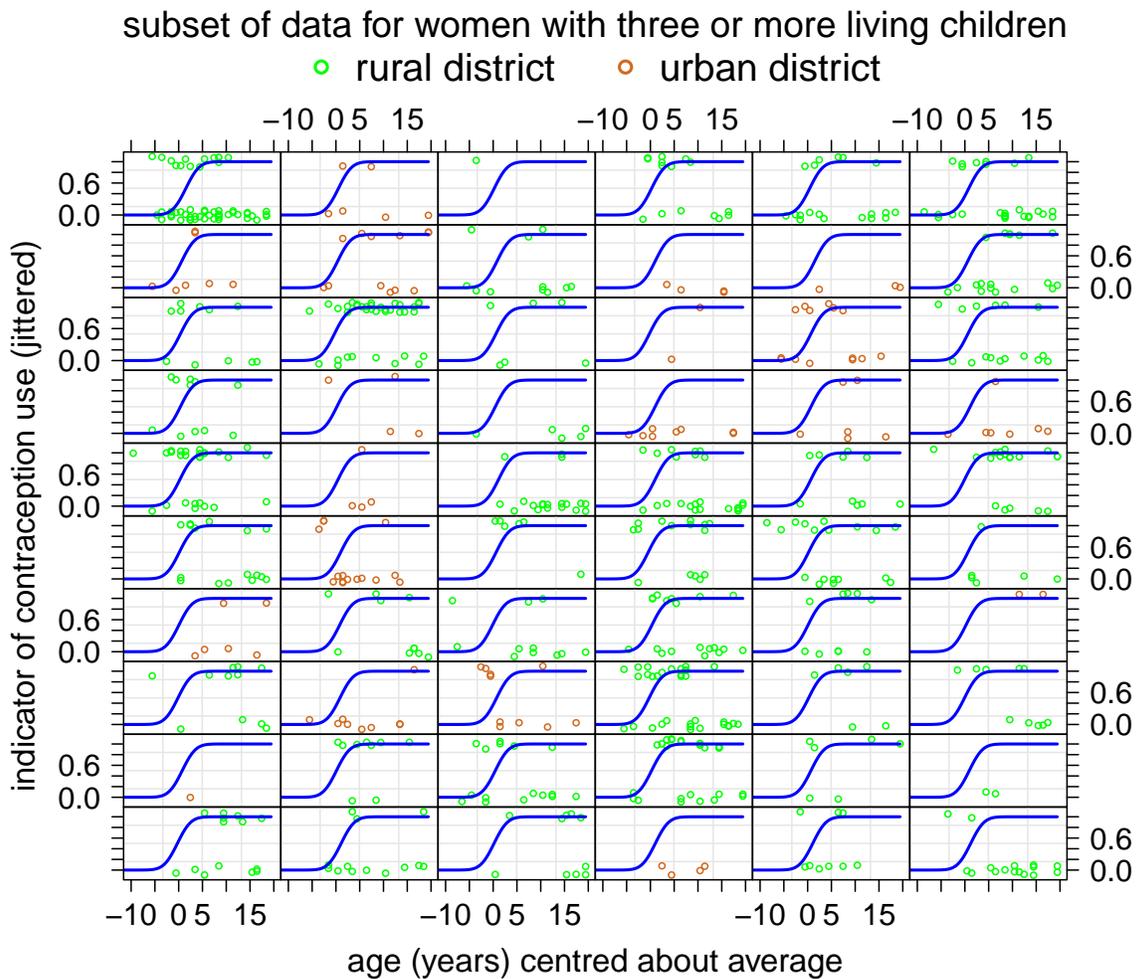


Figure 4: *The data from Figure 2 with the addition of the estimated probability of contraception use curves, according to the expectation propagation random intercepts model fit within the fitContracRanInt fit object.*

resulting in the output:

	95% C.I low	estimate	95% C.I upp
intercept	-1.22775470	-1.04180458	-0.855854477
isUrban	0.29587366	0.50025510	0.704636522
age	-0.02587946	-0.01634977	-0.006820457
livChEq1	0.49356594	0.68154005	0.869514157
livChEq2	0.62253242	0.83058450	1.038636581
livChGe3	0.61062229	0.82445626	1.038290261
sigma1	0.27557767	0.37853878	0.519968122
sigma2	0.31078854	0.49648071	0.793121574
rho12	-0.93599174	-0.79843144	-0.449328947

Note that the estimate of the random slope coefficient has estimate

$$\hat{\sigma}_2 = 0.49648071 \text{ with corresponding 95\% confidence interval } (0.31078854, 0.793121574).$$

The tight confidence interval well away from zero verifies significant variability in the random slopes associated with the indicator of a district being urban.

The scatterplot shown in Figure 5 is a visual summary of the expectation propagation-approximate best predictions of the bivariate random intercepts and slopes, and is produced from the following code:

```
plot(uHat[,1],uHat[,2],col="dodgerblue",lwd=2,xlab="random intercepts predicted values",
     ylab="random slopes predicted values",bty="l",cex.lab=1.5,cex.axis=1.5)
abline(v=0,col="slateblue",lwd=2) ; abline(h=0,col="slateblue",lwd=2)
```

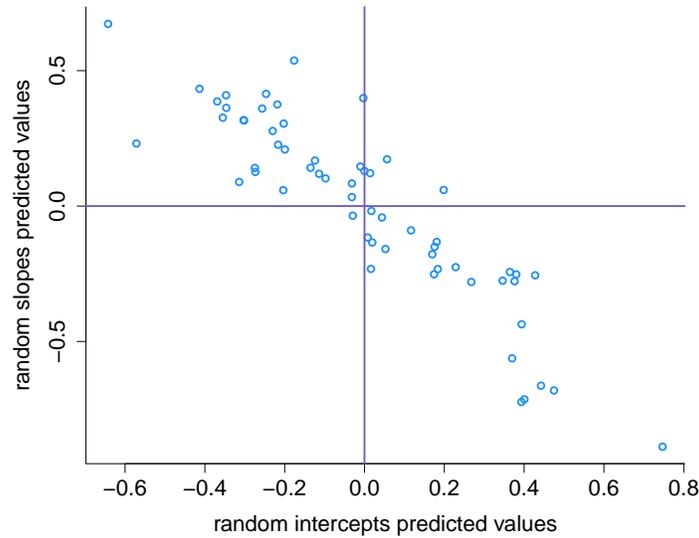


Figure 5: Scatterplot of the expectation propagation-approximate best predictions of the random slopes and corresponding random intercepts, which are part of the `fitContracRanIntAndSlp` fit object.

Our final plot is the analogue of Figure 4 for the random intercepts and slopes model, which is shown in Figure 6
Figure 6 is produced using:

```
probRanIntAndSlpg <- vector("list",60)
for (i in 1:60)
  probRanIntAndSlpg[[i]] <- pnorm(betaHat[1]+uHat[i,1]+betaHat[2]*ageg
    +(betaHat[3]+uHat[i,2])*isUrban[distSttInds[i]]
    +betaHat[6])
figFitRanIntAndSlp <- xyplot(jitter((as.numeric(use)-1),factor=0.5)
  ~ age|district,groups=district,data=ContraceptionHiLivCh,
  layout=c(6,10),
  xlab=list(label="age (years) centred about average",cex=1.35),
  ylab=list(label="indicator of contraception use (jittered)",cex=1.35),
  scales=list(cex=1.25),strip=FALSE,as.table=TRUE,
  key=list(title="subset of data for women with three or more living children",
    columns=2,
    points=list(pch=rep(1,2),col=colourVec[1:2],lwd=rep(2,2)),
    text=list(c("rural district","urban district"),cex=1.55)),
  panel=function(x,y,subscripts,groups)
  {
    iDistrict <- panel.number()
    panel.grid()
    colourInd <- 3 - as.numeric(ContraceptionHiLivCh$urban[subscripts[1]])
    panel.superpose(x,y,subscripts,groups,
      col=colourVec[colourInd],pch=1,cex=0.5)
    panel.xyplot(ageg,probRanIntAndSlpg[[iDistrict]],col="blue",lwd=2,type="l")
  })
```

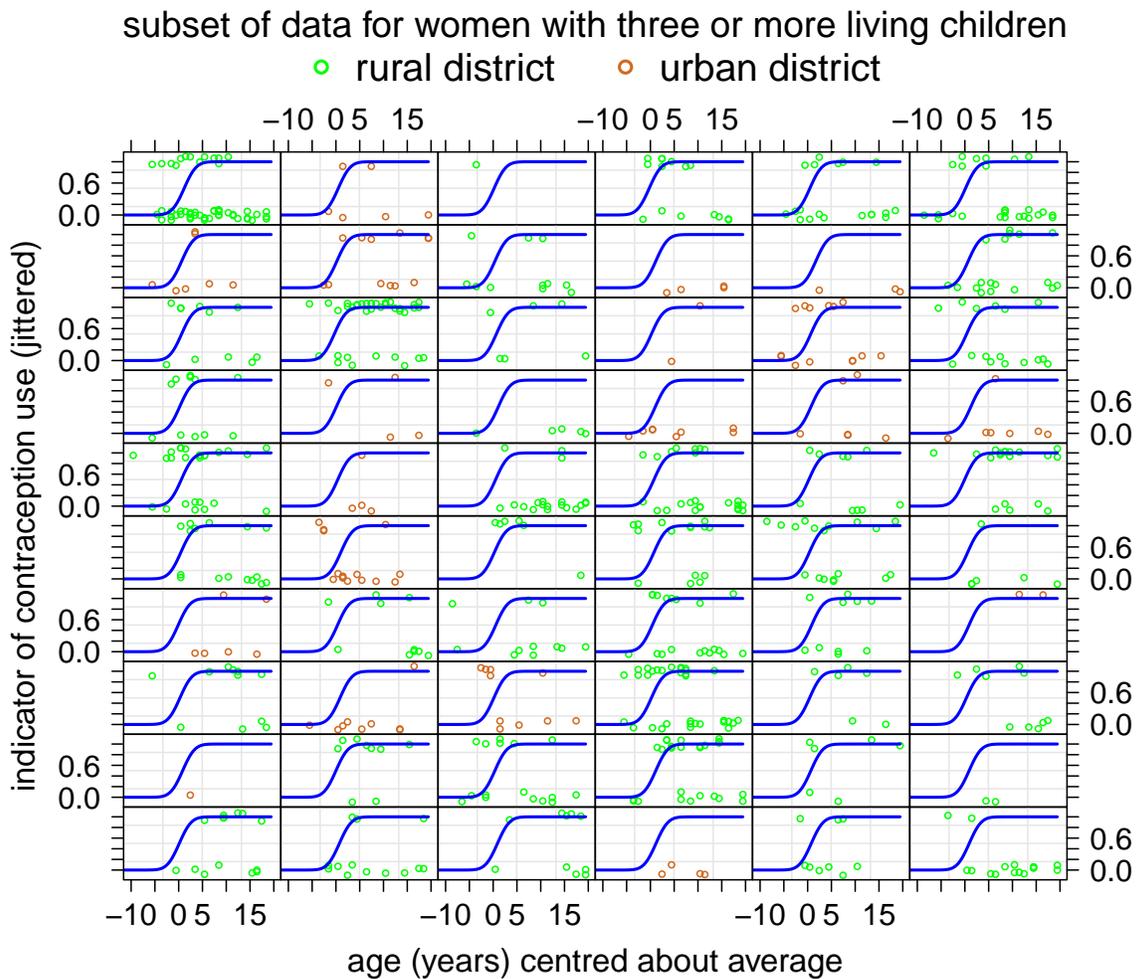


Figure 6: The data from Figure 2 with the addition of the estimated probability of contraception use curves, according to the expectation propagation random intercepts and slopes model fit within the `fitContracRanIntAndSlp` fit object.

References

- Bates, D., Maechler, M. and Bolker, B. (2014). `mlmRev`: Examples from multilevel modelling software review. R package version 1.0. <http://cran.r-project.org>.
- Steele, F., Diamond, I. and Amin, S. (1996). Immunization uptake in rural Bangladesh: a multilevel analysis. *Journal of the Royal Statistical Society, Series A*, **159**, 289–299.
- Hall, P., Johnstone, I.M., Ormerod, J.T., Wand, M.P. and Yu, J.C.F. (2018). Fast and accurate binary response mixed model analysis via expectation propagation. Submitted for publication. <http://matt-wand.utsacademics.info/statsPapers.html>

FINAL TO DO LIST

- * Add bit about `glmmEP.control()`
- * At the end make sure about consistency of axis label sizes.
- * Rethink no. of signif. values in vignette and package.