

04 Hierarchical modeling

Shravan Vasishth

September 13, 2021

Example 1: Reading time differences in subject vs object relatives in English

We begin with a relatively simple question from the psycholinguistics literature: are subject relatives easier to process than object relatives? The data come from Experiment 1 in a paper by Grodner and Gibson (2005).

Example 1: Reading time differences in subject vs object relatives in English

Scientific question: Is there a subject relative advantage in reading?

(1a) The *reporter* who the photographer *sent* to the editor was hoping for a good story. (object gap)

(1b) The *reporter* who *sent* the photographer to the editor was hoping for a good story. (subject gap)

Example 1: Reading time differences in subject vs object relatives in English

Load data and reformat

```
library(bcogsci)  
data("df_gg05_rc")
```

Example 1: Reading time differences in subject vs object relatives in English

Experiment design: Latin square and crossed subject and items

Two important properties of these data are worth noticing.

- Latin square design
- Crossed subjects and items

Example 1: Reading time differences in subject vs object relatives in English

Latin-square design

First, the design is the classic repeated measure Latin square set-up. To see what this means, first look at the number of subjects and items, and the number of rows in the data frame:

```
length(unique(df_gg05_rc$subj))
```

```
## [1] 42
```

```
length(unique(df_gg05_rc$item))
```

```
## [1] 16
```

```
dim(df_gg05_rc)[1]
```

```
## [1] 672
```

There are 42 subjects and 16 items. There are $42 \times 16 = 672$ rows in the data frame.

Example 1: Reading time differences in subject vs object relatives in English

Latin-square design

Notice also that each subject sees exactly eight object gap and eight subject gap sentences:

```
head(xtabs(~subj+condition,df_gg05_rc),n=4)
```

```
##      condition
## subj objgap subjgap
##    1      8      8
##    2      8      8
##    3      8      8
##    4      8      8
```

Example 1: Reading time differences in subject vs object relatives in English

Latin-square design

Table 1: The Latin-square design in repeated measures experiments.

item id	group 1	group 2
1	objgap	subjgap
2	subjgap	objgap
3	objgap	subjgap
4	subjgap	objgap
⋮	⋮	⋮
16	subjgap	objgap

Example 1: Reading time differences in subject vs object relatives in English

Latin-square design: Ensuring balance

```
sample(rep(c("order1", "order2"), 11))
```

```
## [1] "order2" "order1" "order2" "order1" "order1" "order2"  
## [9] "order2" "order1" "order1" "order2" "order1" "order1"  
## [17] "order1" "order2" "order1" "order2" "order1" "order2"
```

Latin square designs are used in planned experiments because they are optimal in several ways.

Example 1: Reading time differences in subject vs object relatives in English

Latin-square design: Generating fake data

```
library(MASS)
```

```
##
```

```
## Attaching package: 'MASS'
```

```
## The following object is masked from 'package:dplyr':
```

```
##
```

```
##      select
```

```
nitem <- 16
```

```
nsubj <- 42
```

```
## prepare data frame for two condition in a latin square design
```

```
g1<-data.frame(item=1:nitem,  
               cond=rep(c("objgap","subjgap"),nitem/2))
```

```
g2<-data.frame(item=1:nitem,  
               cond=rep(c("subjgap","objgap"),nitem/2))
```

Example 1: Reading time differences in subject vs object relatives in English

Latin-square design: Generating fake data

```
## assemble data frame in long format:
```

```
gp1<-g1[rep(seq_len(nrow(g1)),  
            nsubj/2),]
```

```
gp2<-g2[rep(seq_len(nrow(g2)),  
            nsubj/2),]
```

```
fakedat<-rbind(gp1,gp2)
```

```
dim(fakedat) ## sanity check
```

```
## [1] 672  2
```

```
fakedat$subj<-rep(1:nsubj,each=nitem) ## add subjects
```

```
fakedat<-fakedat[,c(3,1,2)]
```

```
fakedat$so<-ifelse(fakedat$cond=="objgap",1,-1)
```

Example 1: Reading time differences in subject vs object relatives in English

Latin-square design: Generating fake data

For example, subject 1 sees the following conditions and items:

```
head(fakedat, n=16)
```

```
##      subj item      cond so
## 1      1     1  objgap  1
## 2      1     2 subjgap -1
## 3      1     3  objgap  1
## 4      1     4 subjgap -1
## 5      1     5  objgap  1
## 6      1     6 subjgap -1
## 7      1     7  objgap  1
## 8      1     8 subjgap -1
## 9      1     9  objgap  1
## 10     1    10 subjgap -1
## 11     1    11  objgap  1
## 12     1    12 subjgap -1
```

Example 1: Reading time differences in subject vs object relatives in English

Fully crossed subjects and items

In the data, because of the Latin square design, each subject sees exactly one item in one of the two conditions:

```
xtabs(~subj+item,df_gg05_rc)
```

```
##      item
## subj 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16
##   1  1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
##   2  1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
##   3  1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
##   4  1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
##   5  1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
##   6  1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
##   7  1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
##   8  1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
##   9  1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
##  10  1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
```

Example 1: Reading time differences in subject vs object relatives in English

The implied generative model

The above design implies a particular statistical model that takes us beyond the linear model.

To remind you, a simple linear model of the above data would be:

$$y \sim \text{Normal}(\alpha + \beta * so, \sigma) \quad (1)$$

Here, object gaps are coded +1, subject gaps -1. See Schad et al. (2019) for an explanation of contrast coding.

```
df_gg05_rc$so<-ifelse(df_gg05_rc$condition=="objgap",1,-1)
```

Example 1: Reading time differences in subject vs object relatives in English

The implied generative model

As figure 1 shows, a Normal likelihood doesn't seem well motivated, so we will use the log-normal.




Figure 1: Distribution of reading times in the Grodner and Gibson Experiment 1 data, at the critical region.

Example 1: Reading time differences in subject vs object relatives in English

The implied generative model

$$y \sim \text{LogNormal}(\alpha + \beta * so, \sigma) \quad (2)$$

Example 1: Reading time differences in subject vs object relatives in English

Between subject variability in mean reading time




Figure 2: Between subject variability in mean reading times.

Example 1: Reading time differences in subject vs object relatives in English

Between subject variability in mean reading time

In the linear model, we can express the assumption that the grand mean intercept α needs an adjustment by subject.

Here, $\text{subj}[n]$ =subject id in n-th row of data.

$$y_n \sim \text{LogNormal}(\alpha + u_{\text{subj}[n],1} + \beta * so_n, \sigma) \quad (3)$$

where we now have two sources of variance:

- within subject: σ
- between subject variance in mean reading times:

$$u_{\text{subj}[n],1} \sim \text{Normal}(0, \sigma_{u1})$$

Example 1: Reading time differences in subject vs object relatives in English

Between item variability in mean reading time




Figure 3: Between item variability in mean reading times.

Example 1: Reading time differences in subject vs object relatives in English

Between item variability in mean reading time

For items ranging from $j = 1, \dots, 16$, we can add this assumption to the model:

$$y_n \sim \text{LogNormal}(\alpha + u_{\text{subj}[n],1} + w_{\text{item}[n],1} + \beta * so_n, \sigma) \quad (4)$$

where there are now three variance components:

- $\sigma \sim \text{Normal}(0, \sigma)$
- $u_{\text{subj}[n],1} \sim \text{Normal}(0, \sigma_{u1})$
- between item variability in mean reading time,
 $w_{\text{item}[n],1} \sim \text{Normal}(0, \sigma_{w1})$

This model is called a *varying intercepts model* with crossed varying intercepts for subjects and for items.

Example 1: Reading time differences in subject vs object relatives in English

Between subject and between item variability in objgap cost

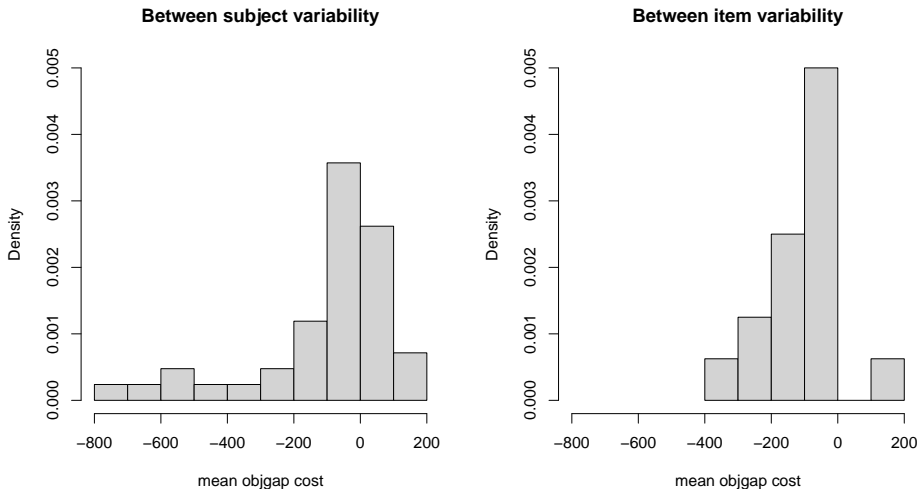


Figure 4: Between subject and item variability in object gap vs subject gap

Example 1: Reading time differences in subject vs object relatives in English

Between subject and between item variability in objgap cost

We can incorporate this assumption into the model by adding adjustments to the β parameter:

$$y_n \sim \text{LogNormal}(\alpha + u_{\text{subj}[n],1} + w_{\text{item}[n],1} + (\beta + u_{\text{subj}[n],2} + w_{\text{item}[n],2}) * SO_n, \sigma) \quad (5)$$

where

- σ
- $u_{\text{subj}[n],1} \sim \text{Normal}(0, \sigma_{u1})$
- $u_{\text{subj}[n],2} \sim \text{Normal}(0, \sigma_{u2})$
- $w_{\text{item}[n],1} \sim \text{Normal}(0, \sigma_{w1})$
- $w_{\text{item}[n],2} \sim \text{Normal}(0, \sigma_{w2})$

This is called the *varying intercepts and slopes* model with *no correlation* between the intercepts and slopes.

Example 1: Reading time differences in subject vs object relatives in English

The “maximal” model

- There is one detail still missing in the model: the adjustments to the intercept and slope are correlated for subjects, and also for items.
- In other words, we have a bivariate distribution for the subject and item random effects:

Example 1: Reading time differences in subject vs object relatives in English

The “maximal” model

$$y_n \sim \text{LogNormal}(\alpha + u_{\text{subj}[n],1} + w_{\text{item}[n],1} + (\beta + u_{\text{subj}[n],2} + w_{\text{item}[n],2}) * SO_n, \sigma) \quad (6)$$

where we have variance components: $\sigma \sim \text{Normal}(0, 1)$ and

$$\Sigma_u = \begin{pmatrix} \sigma_{u1}^2 & \rho_u \sigma_{u1} \sigma_{u2} \\ \rho_u \sigma_{u1} \sigma_{u2} & \sigma_{u2}^2 \end{pmatrix} \quad \Sigma_w = \begin{pmatrix} \sigma_{w1}^2 & \rho_w \sigma_{w1} \sigma_{w2} \\ \rho_w \sigma_{w1} \sigma_{w2} & \sigma_{w2}^2 \end{pmatrix} \quad (7)$$

$$\begin{pmatrix} u_1 \\ u_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma_u \right), \quad \begin{pmatrix} w_1 \\ w_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma_w \right) \quad (8)$$

Example 1: Reading time differences in subject vs object relatives in English

The maximal model

This is a varying intercepts and slopes model with fully specified variance-covariance matrices for the subject and item random effects. It is sometimes called the **maximal model** (Barr et al. 2013).

Example 1: Reading time differences in subject vs object relatives in English

Implementing the model

The above model is simple to implement in the Bayesian framework.

Example 1: Reading time differences in subject vs object relatives in English

Specify and visualize priors

Define some priors first:

- ① $\alpha \sim \text{Normal}(0, 10)$
- ② $\beta \sim \text{Normal}(0, 1)$
- ③ Residual standard deviation: $\sigma \sim \text{Normal}_+(0, 1)$
- ④ All other standard deviations are also: $\text{Normal}_+(0, 1)$
- ⑤ Correlation matrix: $\rho \sim \text{LKJcorr}(2)$.

Example 1: Reading time differences in subject vs object relatives in English

The LKJ prior on the correlation matrix

- In this model, we assume that the vector $\mathbf{u} = \langle u_1, u_2 \rangle$ comes from a bivariate normal distribution with a variance-covariance matrix $\Sigma_{\mathbf{u}}$.
- This matrix has the variances of the adjustment to the intercept and to the slope respectively along the diagonal, and the covariance on the off-diagonals.

Example 1: Reading time differences in subject vs object relatives in English

The LKJ prior on the correlation matrix

- Recall that the covariance $Cov(X, Y)$ between two variables X and Y is defined as the product of their correlation ρ and their standard deviations σ_X and σ_Y , such that, $Cov(X, Y) = \rho\sigma_X\sigma_Y$.

$$\Sigma_{\mathbf{u}} = \begin{pmatrix} \sigma_{u_1}^2 & \rho_u \sigma_{u_1} \sigma_{u_2} \\ \rho_u \sigma_{u_1} \sigma_{u_2} & \sigma_{u_2}^2 \end{pmatrix} \quad (9)$$

Example 1: Reading time differences in subject vs object relatives in English

The LKJ prior on the correlation matrix

The covariance matrix can be decomposed into a vector of standard deviations and a correlation matrix. The correlation matrix looks like this:

$$\begin{pmatrix} 1 & \rho_u \\ \rho_u & 1 \end{pmatrix} \quad (10)$$

Example 1: Reading time differences in subject vs object relatives in English

The LKJ prior on the correlation matrix

In Stan, we write a matrix that has 0's on the off-diagonals as:

$$\text{diag_matrix}(\sigma_{u_1}, \sigma_{u_2}) = \begin{pmatrix} \sigma_{u_1} & 0 \\ 0 & \sigma_{u_2} \end{pmatrix} \quad (11)$$

This means that we can decompose the covariance matrix into three parts:

$$\begin{aligned} \Sigma_{\mathbf{u}} &= \text{diag_matrix}(\sigma_{u_1}, \sigma_{u_2}) \cdot \rho_{\mathbf{u}} \cdot \text{diag_matrix}(\sigma_{u_1}, \sigma_{u_2}) \\ &= \begin{pmatrix} \sigma_{u_1} & 0 \\ 0 & \sigma_{u_2} \end{pmatrix} \begin{pmatrix} 1 & \rho_{\mathbf{u}} \\ \rho_{\mathbf{u}} & 1 \end{pmatrix} \begin{pmatrix} \sigma_{u_1} & 0 \\ 0 & \sigma_{u_2} \end{pmatrix} \end{aligned} \quad (12)$$

Example 1: Reading time differences in subject vs object relatives in English

The LKJ prior on the correlation matrix

So we need priors for the σ_u 's and for ρ_u :

- The basic idea of the LKJ prior is that its parameter (usually called *eta*, η , here it has value 2) increases, the prior increasingly concentrates around the unit correlation matrix (i.e., favors smaller correlation: ones in the diagonals and values close to zero in the lower and upper triangles).
- At $\eta = 1$, the LKJ correlation distribution is uninformative (similar to $Beta(1, 1)$), at $\eta < 1$, it favors extreme correlations (similar to $Beta(a < 1, b < 1)$).

Example 1: Reading time differences in subject vs object relatives in English

Visualize the priors

As always, it is a good idea to visualize these priors. See Figure 5.

Example 1: Reading time differences in subject vs object relatives in English

Visualize the priors

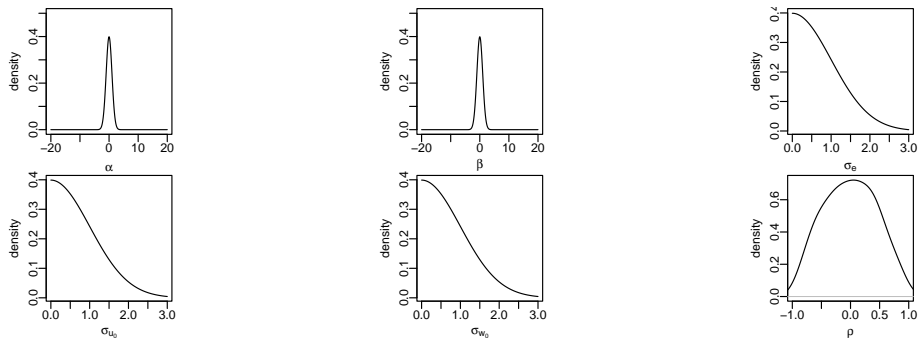


Figure 5: Priors for the Godner and Gibson data.

Example 1: Reading time differences in subject vs object relatives in English

Fit the model using brms

```
priors <- c(set_prior("normal(0, 10)",  
                    class = "Intercept"),  
           set_prior("(0, 1)",  
                    class = "b",  
                    coef = "b_condition"),  
           set_prior("normal(0, 1)",  
                    class = "sd"),  
           set_prior("normal(0, 1)",  
                    class = "sigma"),  
           set_prior("lkj(2)",  
                    class = "cor")  
)
```


Example 1: Reading time differences in subject vs object relatives in English

Fit the model using brms

```
library(tictoc)
tic()
m_gg<-brm(RT~so + (1+so|subj) + (1+so|item),
          df_gg05_rc,family=lognormal())
toc()
```

Example 1: Reading time differences in subject vs object relatives in English

Fit the model using brms



Example 1: Reading time differences in subject vs object relatives in English

Fit the model using brms

Look at the posterior distributions of the parameters on the log ms scale (for the coefficients and standard deviations). Notice that

- The object relative takes longer to read than the subject relative, as predicted. We know this because the parameter b_{so} is positive.
- The largest sources of variance are the subject intercepts, slopes, and the residual standard deviation. Look at the $sd_subject$ parameters, and σ .
- The by-item variance components are relatively small. Look at the sd_item parameters.
- The correlations have very wide uncertainty—the prior is dominating in determining the posteriors as there isn't that much data to obtain accurate estimates of these parameters. Look at the cor parameters.

Example 1: Reading time differences in subject vs object relatives in English

Examine by subject random effects visually

First, extract the posterior samples of the parameters that we will need to compute individual differences.

```
library(bayesplot)

postgg<-posterior_samples(m_gg)
## extract variances:
alpha<-postgg$b_Intercept
beta<-postgg$b_so
cor<-posterior_samples(m_gg,"^cor")
sd<-posterior_samples(m_gg,"^sd")
sigma<-posterior_samples(m_gg,"sigma")
```

Example 1: Reading time differences in subject vs object relatives in English

Examine by subject random effects visually

Example 1: Reading time differences in subject vs object relatives in English

By subject slope adjustments




Figure 6: Variability in subject slope adjustments in the Grodner and Gibson data.

Example 1: Reading time differences in subject vs object relatives in English

By subject slope adjustments

```
## `stat_bin()` using `bins = 30`. Pick better value with `bin
```




Figure 7: Posterior distributions of subject varying intercept and slope correlation parameter in the Grodner and Gibson data.

Example 1: Reading time differences in subject vs object relatives in English

Examine mean and individual differences on the raw ms scale

It is useful to see the effects on the raw ms scale. The log ms scale is difficult to interpret.

Example 1: Reading time differences in subject vs object relatives in English

Mean difference

```
exp(mean(alpha)+mean(beta))-exp(mean(alpha)-mean(beta))
```

```
## [1] 44.10451
```

```
meandiff<- exp(alpha + beta) - exp(alpha - beta)  
(mndiff<-mean(meandiff))
```

```
## [1] 44.59543
```

```
(diffq<-round(quantile(meandiff,prob=c(0.025,0.975)),0))
```

```
## 2.5% 97.5%
```

```
## 7 85
```

Example 1: Reading time differences in subject vs object relatives in English

Mean difference




Figure 8: Mean OR processing cost effect in the Grodner and Gibson data.

Example 1: Reading time differences in subject vs object relatives in English

Individual effects of OR processing cost

Example 1: Reading time differences in subject vs object relatives in English

Individual effects of OR processing cost

```
for(i in 1:42){  
  hist(subjdiff[,i],xlim=c(min(subjdiff),  
                           max(subjdiff)),  
       ylim=c(0,0.02),  
       main=paste("subj",colnames(subjdiff)[i],sep=" "),freq=F)  
  abline(v=mndiff)  
  abline(v=diffq[1],lty=2)  
  abline(v=diffq[2],lty=2)  
  Sys.sleep(.5)  
}
```

Example 1: Reading time differences in subject vs object relatives in English

Individual effects of OR processing cost

This illustrates a point that Blastland and Spiegelhalter (2014) make: “The average is an abstraction. The reality is variation.’”

Example 2: Question-response accuracies (Logistic regression)

The Grodner and Gibson (2005) data also has question-response accuracies: 1 if the response to a question following the sentence was correct, 0 otherwise. We show only the relevant columns below:

```
head(df_gg05_rc[,c(1,2,3,4,6)])
```

```
## # A tibble: 6 x 5
##   subj item condition    RT qcorrect
##   <int> <int> <chr>      <int>    <int>
## 1     1     1  objgap     320        0
## 2     1     2  subjgap     424        1
## 3     1     3  objgap     309        0
## 4     1     4  subjgap     274        1
## 5     1     5  objgap     333        1
## 6     1     6  subjgap     266        1
```

Example 2: Question-response accuracies (Logistic regression)

One could aggregate the accuracy by item, and then just fit a hierarchical linear model:

```
meanp<-with(df_gg05_rc,tapply(qcorrect,
                              IND=list(condition,subj),
                              mean))
q_df<-data.frame(subj=rep(c(1:42),2),
                 so=rep(c(1,-1),each=42),
                 p=c(meanp[1,],meanp[2,]))

head(q_df)
```

```
##   subj so    p
## 1     1  1 0.750
## 2     2  1 0.875
## 3     3  1 1.000
## 4     4  1 1.000
```

Example 2: Question-response accuracies (Logistic regression)

```
library(lme4)
mqlmer<-lmer(p~so+(1|subj),q_df)

## boundary (singular) fit: see ?isSingular

summary(mqlmer)

## Linear mixed model fit by REML ['lmerMod']
## Formula: p ~ so + (1 | subj)
##      Data: q_df
##
## REML criterion at convergence: -97.6
##
## Scaled residuals:
##      Min       1Q   Median       3Q      Max
## -3.03723 -0.77697 -0.07063  0.91823  1.20076
##
```

Example 2: Question-response accuracies (Logistic regression)

Think about the generative process; a 0,1 response is best seen as generated by a Bernoulli distribution with probability of success p :
response $\sim \text{Bernoulli}(p)$. This is the same as a Binomial process, with one trial.

One can therefore model each 0,1 response as being generated from a Bernoulli distribution, which is just a Binomial with a single trial. Thus, what is of interest is the probability of correct responses in subject vs object relatives:

```
round(100*with(df_gg05_rc,  
              tapply(qcorrect,condition,mean)))
```

```
##  objgap  subjgap  
##      88      85
```

Example 2: Question-response accuracies (Logistic regression)

We will transform the probability p of a correct response to a log-odds:

$$\log \frac{p}{1-p} \quad (13)$$

and assume that the log-odds of a correct response is affected by the relative clause type:

$$\log \frac{p}{1-p} = \alpha + \beta * so \quad (14)$$

Example 2: Question-response accuracies (Logistic regression)

This model is called a *logistic* regression because it uses the logistic or logit function to transform p to log odds space. Notice that there is no residual term in this model.

Example 2: Question-response accuracies (Logistic regression)

We can fit the above model easily using brms:

```
m_gg_q1<-brm(qcorrect~so,df_gg05_rc,  
             family=bernoulli(link="logit"))
```

```
## Compiling Stan program...
```

```
## Start sampling
```

```
summary(m_gg_q1)
```

Example 2: Question-response accuracies (Logistic regression)

Obviously, because the question-response data are also repeated measures, we must use a hierarchical linear model, with varying intercepts and slopes for subject and item, as in Example 1:

```
m_gg_q2<-brm(qcorrect~so+(1+so|subj) + (1+so|item),  
             df_gg05_rc,  
             family=bernoulli(link="logit"))
```

```
## Compiling Stan program...
```

```
## Start sampling
```

```
summary(m_gg_q2)
```

This model is not especially good because many of the response accuracies are at ceiling. However, in principle this kind of model is appropriate for binary responses.

Example 2: Question-response accuracies (Logistic regression)

Convert posteriors back to probability space

What is theoretically important is the posterior distribution of the difference between object and subject relative response accuracy. That is on the probability scale. We can go from log-odds space to probability space by solving this equation for p .

Using simple algebra, we can go from:

$$\log \frac{p}{1-p} = \alpha + \beta * so = \mu \quad (15)$$

to:

$$p = \exp(\mu) / (1 + \exp(\mu)) \quad (16)$$

Example 2: Question-response accuracies (Logistic regression)

Convert posteriors back to probability space

For object gap sentences, the factor *so* is coded as 1, so we have $\mu = \alpha + \beta$.

For subject gap sentences, *so* is coded as -1, so we have $\mu = \alpha - \beta$.

Therefore, we just need to plug in the expression for μ for object and subject relatives.

We can now straightforwardly plot the posterior distribution of the difference between object and subject relatives. We see that there isn't any important difference between the two relative clause types.

Example 2: Question-response accuracies (Logistic regression)


Convert posteriors back to probability space

```
postq<-posterior_samples(m_gg_q2)
alpha<-postq$b_Intercept
beta<-postq$b_so
mu_or<-alpha+beta
probor<-exp(mu_or)/(1+exp(mu_or))
mu_sr<-alpha-beta
probsr<-exp(mu_sr)/(1+exp(mu_sr))
```

Example 2: Question-response accuracies (Logistic regression)

Convert posteriors back to probability space

```
hist(probor-probsr, freq=FALSE)  
abline(v=0, lwd=2)
```



References

- Barr, Dale J, Roger Levy, Christoph Scheepers, and Harry J Tily. 2013. “Random Effects Structure for Confirmatory Hypothesis Testing: Keep It Maximal.” *Journal of Memory and Language* 68 (3): 255–78.
- Blastland, Michael, and David Spiegelhalter. 2014. *The Norm Chronicles: Stories and Numbers about Danger and Death*. Basic Books (AZ).
- Grodner, Daniel, and Edward Gibson. 2005. “Consequences of the Serial Nature of Linguistic Input.” *Cognitive Science* 29: 261–90.
- Schad, Daniel J., Sven Hohenstein, Shravan Vasishth, and Reinhold Kliegl. 2019. “How to Capitalize on a Priori Contrasts in Linear (Mixed) Models: A Tutorial.” *Journal of Memory and Language*.