Workshop 8 - Hierarchical Data

JPS

Introduction

We are going to work with the European Social Survey (ESS), where there is a clear two-level hierarchy. Level-1 observations are represented by the individuals that participated in the survey, with countries as level-2 observations (or clusters). We will learn the different approaches that can be used to adjust for this hierarchical structure (taking within cluster correlations as a nuisance), but we will also learn how to control for such hierarchical structures using fixed effects, and how to interpret them using multilevel modelling.

Exercise 1: In the first part of this exercise we are going to assess the effect of leaving the within cluster correlations present in the ESS unadjusted (known as naive analysis). We will do so by comparing results from such a naive analysis with adjustments based on the sandwich estimator and fixed effects models. As we learn how to adjust for within cluster correlations we will also try to estimate the specific association between household income and reported trust on other people. In the second part of the exercise we will employ multilevel models to explore two additional research questions: Is the variability in trust between countries larger than between individuals within the same countries? And, is the association between income and trust uniform across countries?

Exercise 1. Social trust in Europe

Let's access the ESS and trim it down so we keep just those variables that we will be using today. These are: gender ('gndr'), age ('agea'), how often meet with other people ('sclmeet'), trust in other people ('ppltrst'), household's total net income ('hinctnta'), and the country where each interview took place ('cntry').

```
options(scipen=999, digits=5) #This is to remove scientific notation.
library(foreign)
ess = read.dta("ESS9e01_1.dta")
vars = c("gndr","agea","sclmeet","ppltrst","hinctnta","cntry")
ess = ess[vars]
```

We can move on to run a quick exploratory analysis for the three main variables of interest, 'ppltrst', 'hinctnta', and 'cntry'. As we did in Workshop 1 for the case of self-reported happiness, we will also proceed to recode 'ppltrst' so it can be expressed as a numeric variable.

```
plot(ess$ppltrst)
```

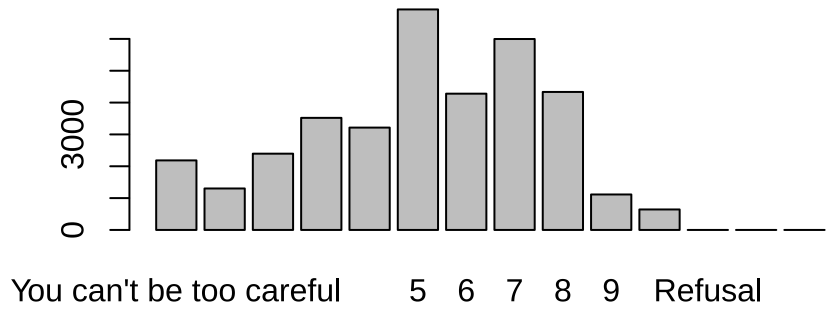


```
table(ess$ppltrst, useNA="ifany") #We need to turn this into a continuous variable.
#To do so we can follow the same procedure we used in Workshop 1 to transform happiness.
ess$ppltrst_rec = ifelse(ess$ppltrst=="You can't be too careful",0,
                  ifelse(ess$ppltrst=="Most people can be trusted",10,
                  as.numeric(as.character(ess$ppltrst))))
table(ess$ppltrst_rec, useNA="ifany")  #The recoding process went ok.
table(ess$cntry)  #Looking at the distribution of level-2 clusters.
table(ess$hinctnta, useNA="ifany")  #7303 of cases are missing, ideally we would
#impute this, but we do not have time to cover that today.
class(ess$hinctnta) #Also, the variable is a factor, but we can express it more simply
#as a continuous variable.
ess$hinctnta = as.numeric(ess$hinctnta)
table(ess$hinctnta, useNA="ifany")  #The transformation worked as expected.
```

Now, to start exploring the association between income and trust we can start with a linear model.

```
naive = lm(ppltrst_rec ~ hinctnta+agea+gndr+sclmeet, data=ess)
summary(naive)
```

```
##
## Call:
## lm(formula = ppltrst_rec ~ hinctnta + agea + gndr + sclmeet,
##     data = ess)
##
## Residuals:
##    Min     1Q Median     3Q    Max
##  -6.30  -1.61   0.18   1.79   6.97
##
## Coefficients:
##                             Estimate Std. Error t value             Pr(>|t|)
## (Intercept)                 2.678002   0.124263   21.55 < 0.0000000000000002
## hinctnta                    0.136553   0.005288   25.82 < 0.0000000000000002
## agea                        0.004864   0.000807    6.03         0.0000000017
## gndrFemale                 -0.024813   0.028250   -0.88                 0.38
## sclmeetLess than once a month  0.701946   0.118066    5.95         0.0000000028
## sclmeetOnce a month         1.252843   0.117773   10.64 < 0.0000000000000002
```

```
## sclmeetSeveral times a month   1.432728   0.113406   12.63 < 0.0000000000000002
## sclmeetOnce a week             1.723626   0.113979   15.12 < 0.0000000000000002
## sclmeetSeveral times a week    1.921902   0.112622   17.07 < 0.0000000000000002
## sclmeetEvery day               1.677348   0.116442   14.40 < 0.0000000000000002
##
## (Intercept)                   ***
## hinctnta                      ***
## agea                          ***
## gndrFemale
## sclmeetLess than once a month ***
## sclmeetOnce a month           ***
## sclmeetSeveral times a month  ***
## sclmeetOnce a week            ***
## sclmeetSeveral times a week   ***
## sclmeetEvery day              ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 2.37 on 28498 degrees of freedom
##   (7507 observations deleted due to missingness)
## Multiple R-squared:  0.0546, Adjusted R-squared:  0.0543
## F-statistic:  183 on 9 and 28498 DF,  p-value: <0.0000000000000002
```

This is an ok model since we have controlled for some relevant confounders. Ideally you need to think more carefully about what other potential confounding factors could be lurking out there and try to control for them, but we do not have enough time to do so here. Gender is not significant, so we could drop it in future models. The main coefficient of interest, 'hinctnta', is positive and significant. However, we are completely disregarding the hierarchical dimension of the data (interviewees clustered within countries), and with it the potentially non-negligible violation of the assumption of independence. We normally refer to such models, where an assumption has been flagrantly violated, as *naive* models. To assess whether that is the case we will use the sandwich estimator, which can be called using the *coeftest* command within the *sandwich* package.

```r
library(lmtest)
library(sandwich)
```

```
## Warning: package 'sandwich' was built under R version 4.4.2
```

```r
#Below we are requesting to re-estimate the model 'naive' using a variance covariance
#empirically determined (i.e. invoking no assumptions of independence or heteroskedasticity)
coeftest(naive, vcov = vcovHC(naive, type="HC1"))
```

```
##
## t test of coefficients:
##
##                                Estimate Std. Error t value             Pr(>|t|)
## (Intercept)                    2.678002   0.136188   19.66 < 0.0000000000000002
## hinctnta                       0.136553   0.005287   25.83 < 0.0000000000000002
## agea                           0.004864   0.000813    5.98         0.0000000022
## gndrFemale                    -0.024813   0.028199   -0.88                 0.38
## sclmeetLess than once a month  0.701946   0.131951    5.32         0.0000001047
## sclmeetOnce a month            1.252843   0.131169    9.55 < 0.0000000000000002
## sclmeetSeveral times a month   1.432728   0.126796   11.30 < 0.0000000000000002
## sclmeetOnce a week             1.723626   0.127046   13.57 < 0.0000000000000002
## sclmeetSeveral times a week    1.921902   0.125947   15.26 < 0.0000000000000002
## sclmeetEvery day               1.677348   0.130200   12.88 < 0.0000000000000002
##
```

```
## (Intercept)                   ***
## hinctnta                      ***
## agea                          ***
## gndrFemale
## sclmeetLess than once a month ***
## sclmeetOnce a month           ***
## sclmeetSeveral times a month  ***
## sclmeetOnce a week            ***
## sclmeetSeveral times a week   ***
## sclmeetEvery day              ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Notice how the size of the coefficients are identical in each model, but the standard errors across in the second model are slightly bigger. This difference is the result of the naive model assuming independence and therefore underestimating the standard errors and any other measures of uncertainty derived from them. The sandwich estimator is a handy technique to add to our toolbox. To be considered when the precision in the estimation process is of paramount importance and we suspect the assumptions of independence or heteroskedasticity do not hold.

Since there are only 19 countries we could also consider a fixed effects (FE) model introducing each level-2 unit in the model as dummy variables. This should help to partially adjust the standard errors by modelling average between cluster differences in trust directly into the fixed part of the model. In addition, this approach can be quite useful to control for unobserved country-level differences that might be confounding (biasing) the effect of 'hinctnta' estimated in our previous two models.

```r
FE = lm(ppltrst_rec ~ hinctnta+agea+gndr+sclmeet+cntry, data=ess)
summary(FE)
```

```
##
## Call:
## lm(formula = ppltrst_rec ~ hinctnta + agea + gndr + sclmeet +
##     cntry, data = ess)
##
## Residuals:
##    Min     1Q Median     3Q    Max
##  -7.13  -1.46   0.20   1.53   7.68
##
## Coefficients:
##                                Estimate Std. Error t value             Pr(>|t|)
## (Intercept)                    3.390197   0.127760   26.54 < 0.0000000000000002
## hinctnta                       0.101158   0.005050   20.03 < 0.0000000000000002
## agea                           0.005770   0.000761    7.58   0.0000000000000348
## gndrFemale                     0.010230   0.026559    0.39              0.70012
## sclmeetLess than once a month  0.692380   0.111143    6.23   0.0000000004740213
## sclmeetOnce a month            1.083391   0.111229    9.74 < 0.0000000000000002
## sclmeetSeveral times a month   1.236610   0.107453   11.51 < 0.0000000000000002
## sclmeetOnce a week             1.394726   0.108186   12.89 < 0.0000000000000002
## sclmeetSeveral times a week    1.540479   0.107227   14.37 < 0.0000000000000002
## sclmeetEvery day               1.470290   0.110933   13.25 < 0.0000000000000002
## cntryBE                       -0.418349   0.073901   -5.66   0.0000000152020243
## cntryBG                       -1.808345   0.072052  -25.10 < 0.0000000000000002
## cntryCH                        0.391473   0.080976    4.83   0.0000013420448802
## cntryCY                       -1.726076   0.100996  -17.09 < 0.0000000000000002
## cntryCZ                       -0.552549   0.074315   -7.44   0.0000000000001073
## cntryDE                       -0.171366   0.069366   -2.47              0.01350
```

```
## cntryEE                        0.160963   0.071842    2.24               0.02507
## cntryFI                        1.306254   0.073906   17.67 < 0.0000000000000002
## cntryFR                       -0.902905   0.071992  -12.54 < 0.0000000000000002
## cntryGB                       -0.267271   0.071474   -3.74               0.00018
## cntryHU                       -0.702562   0.087923   -7.99   0.0000000000000014
## cntryIE                        0.065461   0.074826    0.87               0.38167
## cntryIT                       -0.652472   0.075462   -8.65 < 0.0000000000000002
## cntryNL                        0.544160   0.077787    7.00   0.0000000000027014
## cntryNO                        1.152692   0.080075   14.40 < 0.0000000000000002
## cntryPL                       -1.208786   0.089198  -13.55 < 0.0000000000000002
## cntryRS                       -1.647160   0.075356  -21.86 < 0.0000000000000002
## cntrySI                       -0.968134   0.082202  -11.78 < 0.0000000000000002
##
## (Intercept)                   ***
## hinctnta                      ***
## agea                          ***
## gndrFemale
## sclmeetLess than once a month ***
## sclmeetOnce a month           ***
## sclmeetSeveral times a month  ***
## sclmeetOnce a week            ***
## sclmeetSeveral times a week   ***
## sclmeetEvery day              ***
## cntryBE                       ***
## cntryBG                       ***
## cntryCH                       ***
## cntryCY                       ***
## cntryCZ                       ***
## cntryDE                       *
## cntryEE                       *
## cntryFI                       ***
## cntryFR                       ***
## cntryGB                       ***
## cntryHU                       ***
## cntryIE
## cntryIT                       ***
## cntryNL                       ***
## cntryNO                       ***
## cntryPL                       ***
## cntryRS                       ***
## cntrySI                       ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 2.22 on 28480 degrees of freedom
##   (7507 observations deleted due to missingness)
## Multiple R-squared:  0.168,  Adjusted R-squared:  0.167
## F-statistic:  212 on 27 and 28480 DF,  p-value: <0.0000000000000002
```

We can see that the standard errors are practically identical to the naive model, so in this case the FEs model did not really adjust them, possibly because the assumption of independence held relatively well anyway. However, we can see that the estimates of the regression coefficients are not exactly the same. For the case of 'hinctnta' the regression coefficient is around 30% smaller now. This is probably due to the country-level confounding effects that we are now controlling for (e.g. differences in cultural, educational, or institutional differences, economic inequality, etc.).

Coming back to the comparison of the three models that we have used (naive, sandwich and fixed-effects) and considering that our aim is the estimation of the effect of 'hhinctnta' as accurately and precisely as possible, I would choose to report the FEs model. This is because:

1. The bias in FEs standard errors compared to those estimated using the sandwich estimator is almost negligible; and

2. There is a significant difference in the regression coefficient for 'hinctnta', compared to either the naive model or the model resorting on the sandwich estimator.

Beyond, answering our first research question, the results of our FEs model are also quite illustrative of the substantial between country variability in reported levels of trust. To visualise this more clearly we can use the *plot_model* command available in the *sjPlot* library.

```
library(sjPlot)
```

```
## Warning: package 'sjPlot' was built under R version 4.4.2
```

```
plot_model(FE)
```

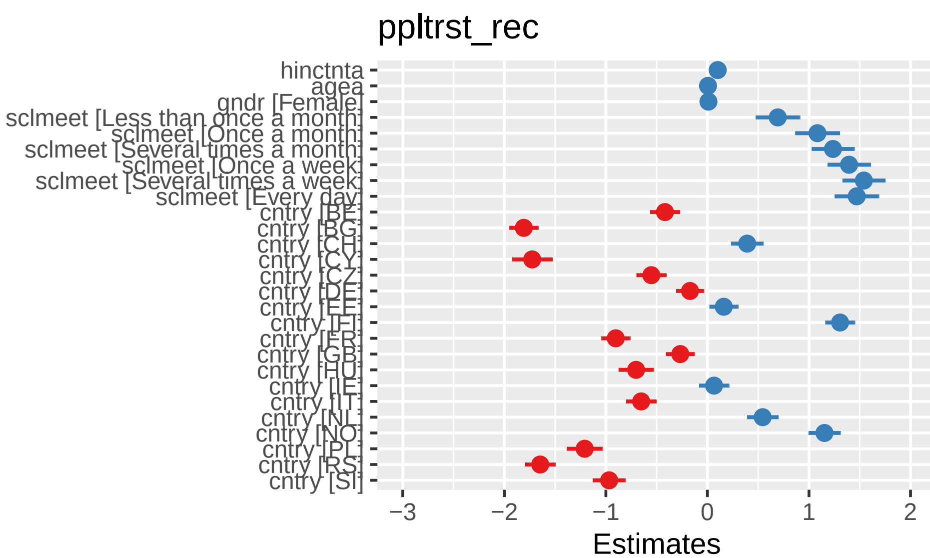

People in Bulgaria ('BG') appear to be really mistrusting whereas Finnish (FI) lie on the opposite side of the spectrum. Note Austria ('AT') is not there as it is being used a reference category. If you want to find out what labels are used for other countries you can find them in this website.

In the second part of this exercise we proceed to explore the two substantial research questions:

1. Is the variability in trust between countries larger than between individuals within the same countries?

2. Is the effect of income on trust uniform across countries?

The between country variability observed in the FEs model signals the relevance of these questions. However, the FEs model is only giving us country specific estimates of trust, which is not exactly a measure of the overall variability in trust between countries, and it definitely is not telling us anything about any potential variability in the effect of 'hinctnta' on trust between countries.

To explore those research questions we will use multilevel modelling. For example, to estimate the overall variability in trust across countries we can expand our naive analysis by adding a random intercepts term u_j, so the model becomes, $Y_{ji} = (\beta_0 + u_j) + \beta_k X_{kji} + \epsilon_{ji}$. To estimate this multilevel model we will use the *lmer* command from the *lme4* package. The code required is identical to the standard *lm* model with the one difference that now we will need to specify the variable capturing the cluster units as follows '(1|cluster)', see below.

```
library(lme4)
RI = lmer(ppltrst_rec ~ hinctnta+agea+gndr+sclmeet+(1|cntry), data=ess)
summary(RI)
```

```
## Linear mixed model fit by REML ['lmerMod']
## Formula: ppltrst_rec ~ hinctnta + agea + gndr + sclmeet + (1 | cntry)
##    Data: ess
##
## REML criterion at convergence: 126574
##
## Scaled residuals:
##    Min     1Q Median     3Q    Max
## -3.207 -0.656  0.090  0.691  3.453
##
## Random effects:
##  Groups   Name        Variance Std.Dev.
##  cntry    (Intercept) 0.79     0.889
##  Residual             4.94     2.222
## Number of obs: 28508, groups:  cntry, 19
##
## Fixed effects:
##                              Estimate Std. Error t value
## (Intercept)                  2.998911   0.235437   12.74
## hinctnta                     0.101315   0.005050   20.06
## agea                         0.005768   0.000761    7.58
## gndrFemale                   0.010039   0.026559    0.38
## sclmeetLess than once a month 0.692353   0.111141    6.23
## sclmeetOnce a month          1.084206   0.111225    9.75
## sclmeetSeveral times a month  1.237731   0.107447   11.52
## sclmeetOnce a week           1.396510   0.108179   12.91
## sclmeetSeveral times a week   1.542725   0.107218   14.39
## sclmeetEvery day             1.471910   0.110923   13.27
##
## Correlation of Fixed Effects:
##             (Intr) hnctnt agea   gndrFm sLtoam sclOam scStam sclOaw scStaw
## hinctnta    -0.124
## agea        -0.222  0.210
## gndrFemale  -0.070  0.103 -0.001
## sclmtLstoam -0.402 -0.045  0.028  0.003
## sclmtOncamn -0.408 -0.061  0.052  0.000  0.851
## sclmtSvrtam -0.423 -0.078  0.057 -0.002  0.882  0.891
## sclmtOncawk -0.423 -0.075  0.069 -0.003  0.876  0.886  0.923
## sclmtSvrtaw -0.432 -0.070  0.088 -0.002  0.885  0.895  0.934  0.931
## sclmtEvrydy -0.424 -0.056  0.111 -0.001  0.856  0.866  0.903  0.900  0.914
```

Notice how the output from our 'RI' model includes a *random effects* part. Here we can find the estimated standard deviation of the level-2 and level-1 residuals, $\hat{\sigma}_u = 0.89$ and $\hat{\sigma}_\epsilon = 2.22$, which can be interpreted as the unexplained differences (after controlling for the four variables included in our model) between countries and between individuals. As we anticipated from the results of our FE model, the variability stemming at the country level is substantial. However, since the variability stemming from the individual level is more than twice bigger, we can determine that individual characteristics are more important at determining differences in trust than country characteristics. To be more precise in answering the research question we can calculate the intracluster correlation coefficient, $ICC = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_\epsilon^2}$.

```
0.79 / (0.79+4.94)
```

So, only 13.8% of the unobserved variability stems from country differences, the rest is due to individual differences. Another common interpretation of the ICC is as the correlation (similarity) of level-1 units within clusters, in our cases individuals within countries. In relation to our specific research question, we could interpret this findings as: although there are significant average differences in trust across countries, in determining individual trust, individual characteristics are more meaningful than country characteristics.

If we want to obtain the specific random intercept for each country ($\hat{u}_j$) we can use *ranef*. Results are very similar to those from the FE model, with Finland leading and Bulgaria lagging behind. Notice however, that unlike for the FE model these estimates are obtained from the residual (random) part of the model. As such, they are not really controlling for any potential country level confounders.

```
ranef(RI)
```

```
library(glmmTMB) #This is to be able to use the option "re" in plot_model
```

```
## Warning: package 'glmmTMB' was built under R version 4.4.2
```

```
plot_model(RI, type="re")
```

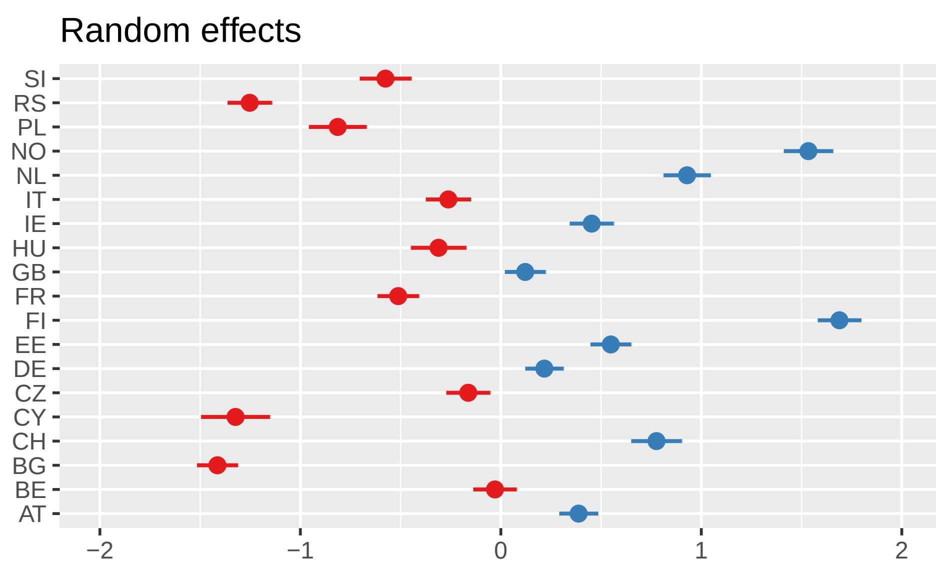

If you want to obtain the specific intercept for each country ($\beta_0 + u_j$) you can use *coef*. We can also create a linechart putting in context both the variability of these random intercepts and the effect of 'hinctnta' using *plot_model*.

```
coef(RI)
```

```
#To drop the legend use, 'show.legend=FALSE'
plot_model(RI, type="pred", terms=c("hinctnta","cntry"), pred.type="re", colors="gs", ci.lvl=NA)
```

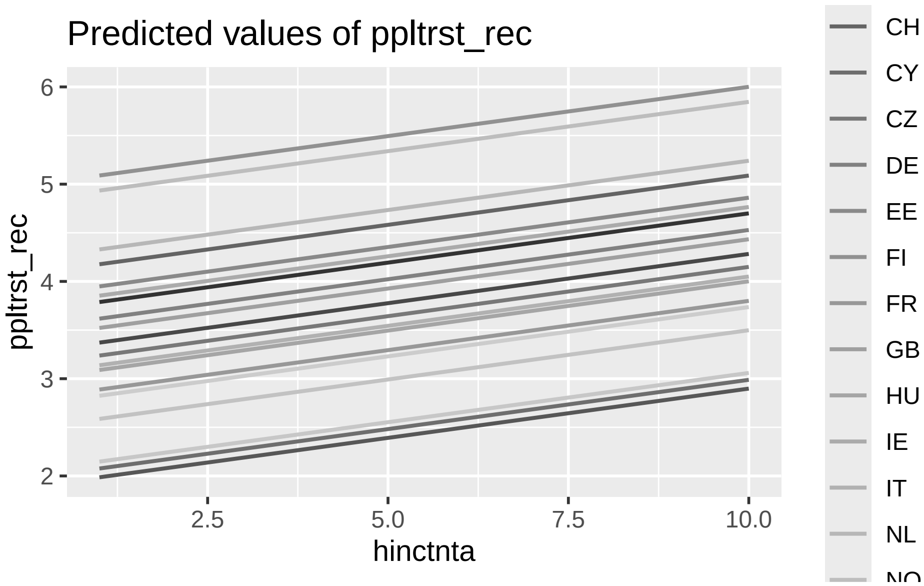

Notice how the random intercepts model that we have estimated, and plotted in the last graph, assumes that the association between income and trust is uniform across countries, what changes is the intercept (the overall level of trust), but not the slope. We can relax this assumption using a random slopes model. However, random slopes models are much more complex, and in many instances cannot be estimated reliably, especially when the number of level-2 units is low (as it is the case here), and the number of explanatory variables is large. In those instances you might get a a warning saying that 'the model failed to converge'.

This problem can be resolved using different estimation process, for example we could use Bayesian methods based on Markov chain Monte Carlo (McMC), but that takes us far beyond the scope of today's session. Instead, what I recommend that you do is first of all to try and reduce the range of your variables, perhaps by standardising them (as we saw in Workshop 2), or by taking log-transformations if you have really strong outliers on the right end of the distribution. You can also try dropping explanatory variables, but that might affect your results if those variables were controlling for important confounders (as we saw in Workshop 3). Another solution, is to forget about multilevel modelling and explore the variability of your variable of interest across level-2 units by including interactions. This is just an expansion of the fixed effects approach that we have used above.

```
#Notice how the interaction between cntry and hinctnta is specified using '*'.
Interact = lm(ppltrst_rec ~ hinctnta*cntry+agea+sclmeet, data=ess)
summary(Interact)
```

Now we have that the coefficient of 'hinctnta' refers to the association between income and trust for Austria (the reference category in 'cntry'). All the other 'hinctnta:cintry' coefficients can be interpreted as how different the association between income and trust is in all other countries compared to Austria. If we take a look at it we can see that there are some significant differences. For example, we can note how for Austria, Cyprus, or Norway, the association between income and trust is non-significant, which contrasts with what we see for example for Slovenia or Hungary, where for every additional decile of household income the average trust increases in more than 0.1 points.

As we saw before for the case of fixed effects, this is very useful, but we do not have an overall measure of how much the association between income and trust varies across countries, or how uniform it is. To get that we really need to use a random slopes model. To tell R that we want both random intercepts and a random slopes term for 'hinctnta' we need to include 'hinctnta' in the random part of the model, as shown below.

```
RS = lmer(ppltrst_rec ~ hinctnta+agea+gndr+sclmeet+(1+hinctnta|cntry), data=ess)
summary(RS)
```

Notice how in the random effects part of the output we now get a standard deviation for 'hinctnta' (0.030), that is the random slope representing the between country variability in the association between income and trust. Notice as well that now we have a correlation term. This is indicating the correlation between the random intercepts and random slopes ($\hat{\rho}_{u_{0j}u_{1j}}$). The correlation is negative and quite strong, which points at how, on average, in those countries where average trust is higher, the positive association between income and trust is weaker, and vice versa, in countries with average low trust, the association between income and trust is stronger (more positive).

However, to confirm this conclusion formally, we need to test whether the addition of this random slope (together with its correlation with the random intercepts) is significant. To test that we can run a likelihood ratio test using the command *lrtest* within the *lmtest* package.

```
library(lmtest)
#We test whether the random slopes model provides a significant improvement to the
#random intercepts model goodness of fit.
lrtest(RI, RS)
```

```
## Likelihood ratio test
##
## Model 1: ppltrst_rec ~ hinctnta + agea + gndr + sclmeet + (1 | cntry)
## Model 2: ppltrst_rec ~ hinctnta + agea + gndr + sclmeet + (1 + hinctnta |
##     cntry)
##   #Df LogLik Df Chisq Pr(>Chisq)
## 1  12 -63287
## 2  14 -63279  2  16.2     0.0003 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
#'Pr(>Chisq)' represents the p-value for this test, which is below 0.05.
```

The extension of the random intercepts to a random slopes model is statistically significant. Therefore, we can conclude that the association between income and trust is not uniform across countries. However, we need to make an additional effort at interpreting our results to ascertain not only statistical significance but also substantive significance (effect size). In our case, whether the variability in the association between income and trust seems meaningful. We should always do this, but especially when using samples as big as the one we have here, since the bigger the sample the more likely that we will find results to be statistically significant even when the effect is really small, and therefore not that meaningful.

To assess the effect size of these random slopes ($\hat{u}_{1j}$), we can put them in context by comparing them with the coefficient of 'hinctnta' from the fixed part of the model ($\hat{\beta}_{hinctnta}$). One simple way to do so is by calculating the range of associations between income and trust across countries after taking into account the country-specific random slopes.

```
coef(summary(RS))  #Estimates from the fixed part of the random slopes model.
#These represent average effects across the whole sample, like in a standard linear
#model.
coef(summary(RS))[2,1]  #0.104, the average association between income and trust.
coef(RS)$cntry     #Country-specific estimates, notice how each country has a
#different estimate for hinctnta, since we used random slopes and allow these to
#vary.
min(coef(RS)$cntry[,2]) #0.046, the weakest association between income and trust,
#Austria.
max(coef(RS)$cntry[,2]) #0.151, the strongest association, Hungary.
```

For example, we have that the average association between income and trust, as estimated in the random slopes model, is 0.104, but that can vary across countries, from as low as 0.046 in Austria to as high as 0.151 in Hungary. In conclusion, yes, we can confirm that there is statistically significant variability in the association between income and trust, AND that this variability is moderately meaningful. This result illustrates very well the relevance of our research question, and the importance of context and culture when modelling behavioural and attitudes data. As for the case of the random intercepts model, we can plot the specific random slopes using *plot_model*.

```
plot_model(RS, type="pred", terms=c("hinctnta","cntry"), pred.type="re", colors="gs", ci.lvl=NA)
```

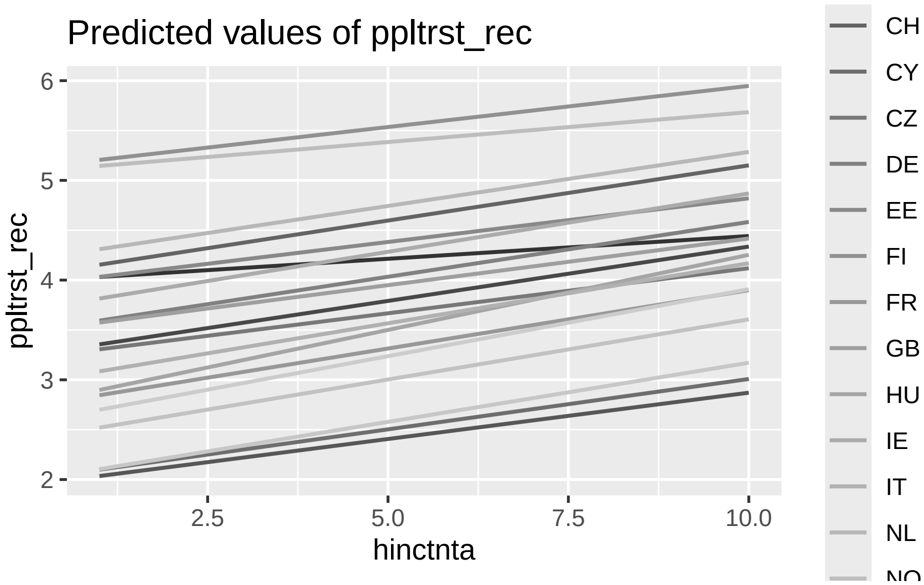

To recap, when presented with hierarchical data we can opt for one of three approaches, robust standard errors using the sandwich estimator, fixed effects models including the cluster units as dummy variables, or multilevel modelling. If we care most about precision in the estimation of the standard errors we should use the sandwich estimator. This approach is also interesting because (unlike multilevel modelling) it does not rely on additional assumptions, and because it adjusts for lack of independence in our residuals (within cluster correlation, autocorrelation, etc.), but also for heteroskedasticty. If we want to control for potential cluster level confounders, and in so doing obtain more accurate estimates of our regression coefficients we should use fixed effects. However, this approach might end up rendering an overfitted model if the number of cluster units is too big. The advantage of multilevel modelling lies on its capacity to model the unexplained cluster-level variability, which as we have just seen can help us explore some important research questions. Incidentaly, if properly specified, multilevel models can adjust for problems of within cluster correlation affecting the assumption of independence.

Preparation for next week's workshop

In next week's workshop we will explore growth curve models, which are nothing more than the application of the same multilevel models that we have seen today, only applied for the analysis of longitudinal data. As such, you can see next week as a continuation of what we did today, which will helps us settle some of the various new concepts and modelling approaches that we saw today. Unlike today, next week's exercise will be largely unguided, requiring you to replicate some of the models used here for a different dataset, where we will be exploring judges' patterns of sentencing across their careers. As always, you are required to prepare this exercise in advance, try to reuse some of the procedures that we have used today, and bring any questions that might arise. See you next week.