SAA for many N level multilevel models

Welcome to an SAA for fitting many model types developed for Stat-JR v1.0.5

Input questions

Firstly on this page you will need to specify the dataset required from the list of available datasets.

Which dataset do you wish to use:	▾

Submit

Next you need to choose many options including the response, estimation method, clustering variables and predictor variables (both continuous and categorical) from the chosen dataset. After choosing these variables the SAA will run and you will see a block of text describing how many observations are to be used at the bottom of this page. The rest of the analysis will appear in pages 2-12.

What estimation method do you want to use:	MCMC
What is the response variable:	attain
What distribution are you going to assume:	Normal
Which higher level classifications do you wish to consider:	pid,sid
Are there any continuous predictors that need including in all models:	No

Are there any categorical predictors that need including in all models:	No
Do you want to include any continuous predictors as candidates for inclusion in the models:	Yes
Which continuous predictors do you want to consider:	vrq,sc
Do you want to include any categorical predictors as candidates for inclusion in the models:	Yes
Which categorical predictors do you want to consider:	sex,fed,choice,med
What selection type do you require:	Forward pass
Do you want to test for random slopes:	No
Do you want to test for interactions:	No
How do you wish to compare models:	Wald
How long to burnin for:	500
How long to then run chains for:	2000
What is the minimum ESS at which to stop (use 0 to just run for number last input):	200

Do you want to use orthogonal parameterisation:	No
What change in DIC denotes a better model:	1

The Analysis Assistant you are currently using is designed to work on complete datasets only and so as a pre-processing step we have to remove any rows that contain missing data in columns used in the analysis that follows. For now the list of columns to be considered is: attain, pid, sid, vrq, sc, sex, fed, choice, med. There are 0 (0.0%) rows that get deleted This results in a dataset of 3435 rows.
On the next page we will look at the shape of the response and, in the case of normal responses, decide whether to log transform.

Exploring the response

We will begin our analysis of the dataset by doing some basic data exploration.

You have chosen attain as your response variable and so a first step is to take a look at this variable and assess its suitability for a normal model. The summary statistics for the variable are in the table below:

Observations	3435
Mean	5.679
Standard Deviation	3.058
Median	5.0

We also look at a histogram of attain to see if it is approximately normally distributed. Although in modelling the response in terms of a set of predictors it is what is unexplained (the model residuals) that need to be normally distributed, it is still useful to look at the response variable as a very skewed variable will often lead to very skewed residuals.

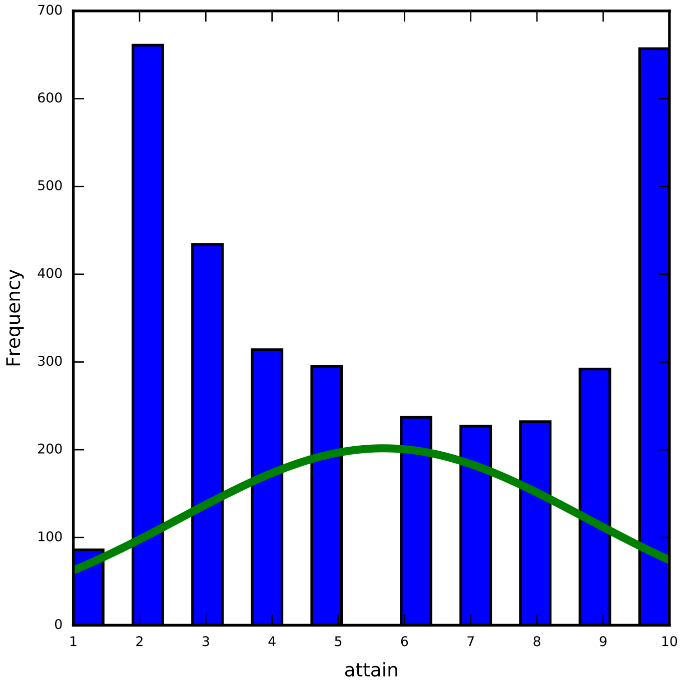

Here the median is smaller than the mean and there is significant skew to the right. The skewness value is 0.158. Here the statistical significance may be to some degree due to the large sample size as from a practical perspective values of skew less than 2 are not considered too big a skew.

There are no obvious outliers in attain.

Do you want to log transform the response variable?:	No

Exploring the predictors individually

We can also look at each of the predictor variables in turn in isolation.

For categorical predictors we are looking at how common each category is in the dataset. In particular we are checking for rare categories which might cause difficulties in modelling and might therefore be usefully merged with other categories (though this would need to be done outside this SAA).

For predictor sex we see the following:

sex	N	Percentage
0	1739	50.626
1	1696	49.374
Total	3435	100

None of the categories of sex have fewer than 5 observations.

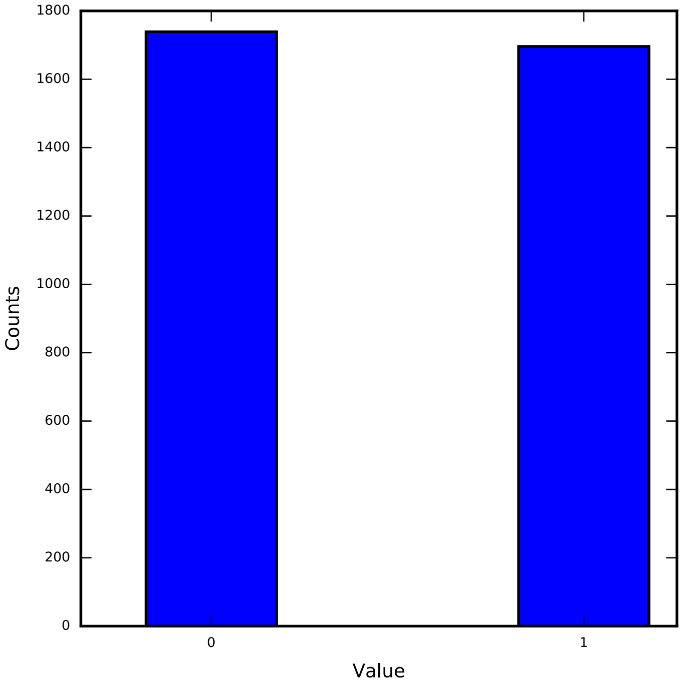

For predictor fed we see the following:

fed	N	Percentage
0	2489	72.46
1	946	27.54
Total	3435	100

None of the categories of fed have fewer than 5 observations.

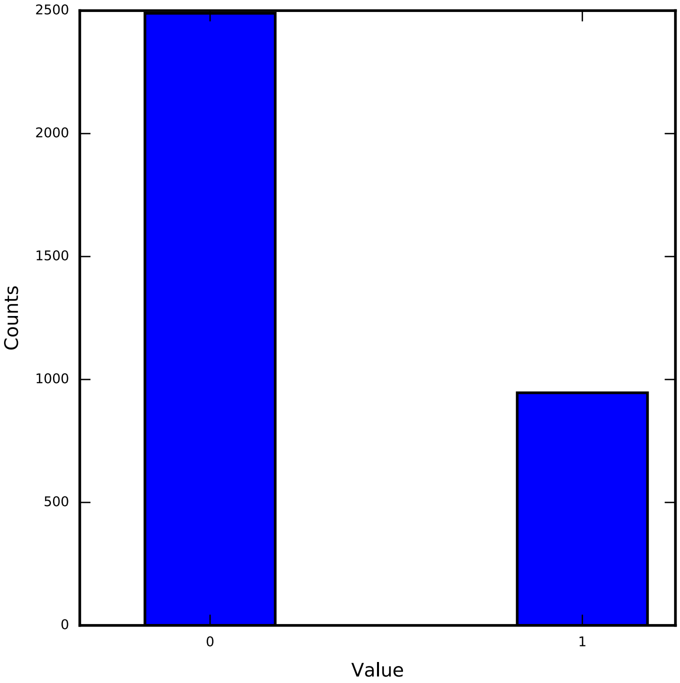

For predictor choice we see the following:

choice	N	Percentage
1	3069	89.345
2	212	6.172
3	3	0.087
4	151	4.396
Total	3435	100

The values 3 are rare categories with fewer than 5 observations.

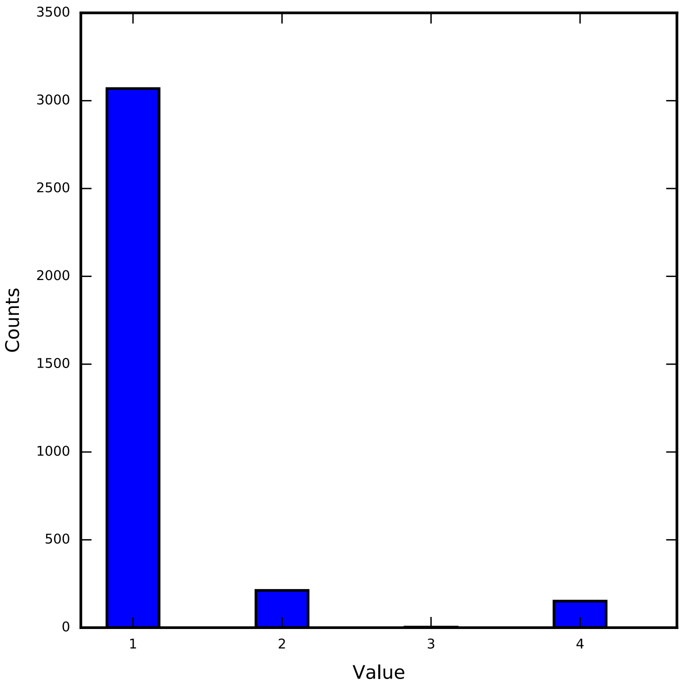

For predictor med we see the following:

med	N	Percentage
0	2260	65.793
1	1175	34.207
Total	3435	100

None of the categories of med have fewer than 5 observations.

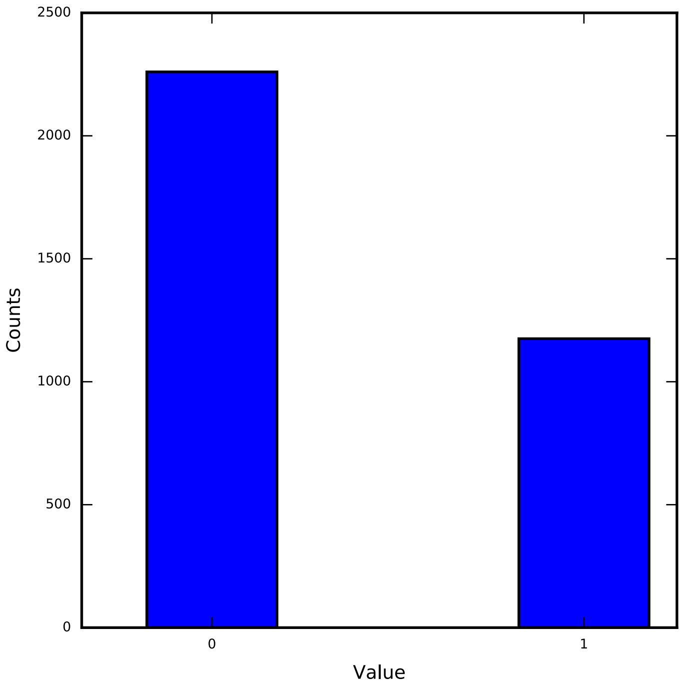

For continuous predictors we are interested in looking at summary statistics, the shape of the distribution and any unusual values. If the distribution is skewed then we might want to transform the variable before fitting it in the model although it is more important to consider transformations of the response variable and remember what is important is whether the relationship between the response and predictor is linear. If there are unusual values we will want to check that the unusual values are correct and not errors and also whether we may want to treat the variable differently. Another possibility for unusual shaped distributions is to instead categorise the variable into ranges of values.

For predictor vrq we see the following:

Name	vrq
Observations	3435
Mean	97.804
Standard Deviation	13.291
Median	98.0

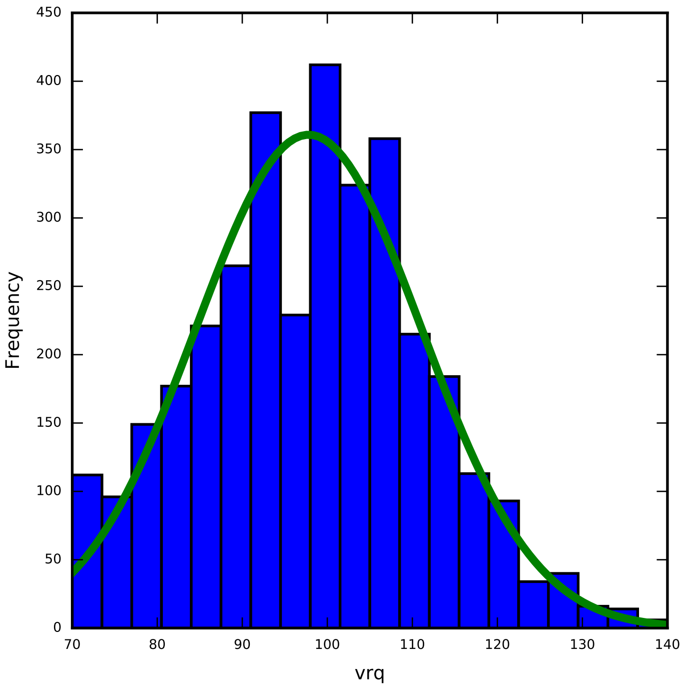

Here the median is smaller than the mean and there is significant skew to the right. The skewness value is 0.098. Here the statistical significance may be to some degree due to the large sample size as from a practical perspective values of skew less than 2 in absolute magnitude are not considered too big a skew.

There are no obvious outliers in vrq.

For predictor sc we see the following:

Name	sc
Observations	3435
Mean	6.845
Standard Deviation	10.886
Median	0.0

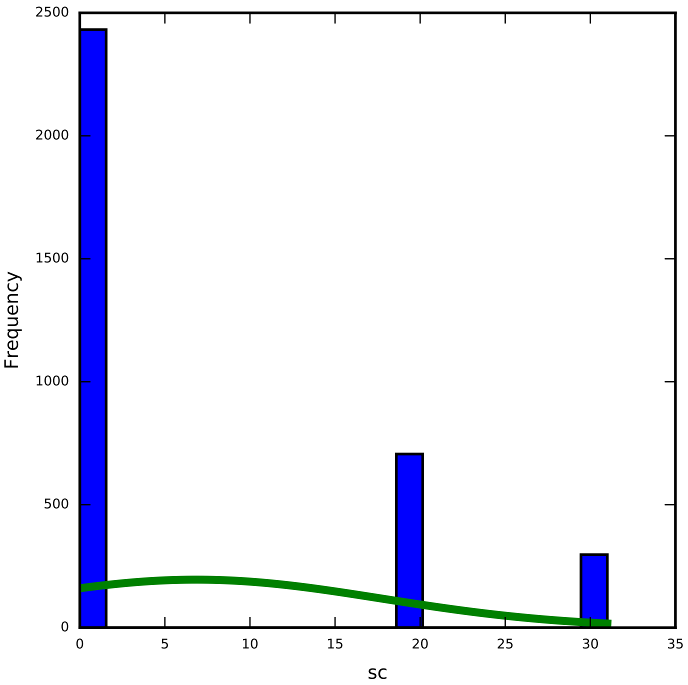

The variable has only 4 unique values but is being treated as continuous.

Here the median is smaller than the mean and there is significant skew to the right. The skewness value is 1.136. Here the statistical significance may be to some degree due to the large sample size as from a practical perspective values of skew less than 2 in absolute magnitude are not considered too big a skew.

There are no obvious outliers in sc.

Assessing the relationship between the response and individual predictors

Once we are happy with our response variable and our set of predictors we now want to have a preliminary look at them together before progressing to the univariable modelling.

For the categorical predictors it is worth looking at the mean value of the response in each category to assess if there are differences. We can then formally test this with a t-test for binary predictors or an ANOVA for predictors with more than 2 categories.

Here is a tabulation of the response, attain for predictor sex with category 1 having the largest mean and category 0 the smallest.

Category	N	Mean	Standard Deviation	Median
0	1739	5.422	3.076	5.0
1	1696	5.942	3.017	6.0

The formal test is as follows:

There are two groups in the data:
The first group has 1739 observations with mean 5.422 standard deviation 3.077.
The second group has 1696 observations with mean 5.942 standard deviation 3.018.
We are trying to test a hypothesis as to whether the two groups differ in their (population) means by a statistically significant amount. Statistical significance is related to how likely a result is to be a chance occurance. Here we are trying to differentiate between a real difference (no matter how small) and a difference that may have occurred due to the samples we have chosen.
The mean difference is 0.521 with the second group having the larger sample mean.
We need to quantify if this difference is large relative to the variability in the data. To do this we calculate the standard error of the difference. This is a function of the variabilities in the samples from group A and group B combined with their sample sizes. The bigger the 2 variabilities the larger the standard error, whilst the smaller the variability the smaller the standard error.
For our data the standard error of the mean difference is 0.104 and we divide our observed difference by this standard error to give a test statistic with value 5.007.

This test statistic is then compared to a t distribution with degrees of freedom equal to the sum of the sample sizes in each group (3435) - 2. In this case a t distribution with 3433. This t table has values of 1.961 for p=0.05 and 2.577 for p=0.01.

As 5.007 > 2.577 our p value is less than 0.01 and we have strong evidence to reject the null hypothesis (at the p=0.01 level).

The p-value is in fact less than 0.0001.

The t test assumes that the distribution of the response in each group follows a Normal distribution. We could check this by looking at histograms of the variable in each group. If we were concerned about the normality assumption then we could instead use a Mann Whitney (MW) test.

A Mann Whitney test works simply on the order (or ranks) of the responses across the two groups. So the response variable is firstly sorted and then each value is ranked. The ranks for each group are then summed and the value that is larger is compared with what would be expected if there was no difference between the groups.

In this case the MW U statistic is 1324016 which for samples of size 1739 and 1696 corresponds to a p value of less than 0.0001.

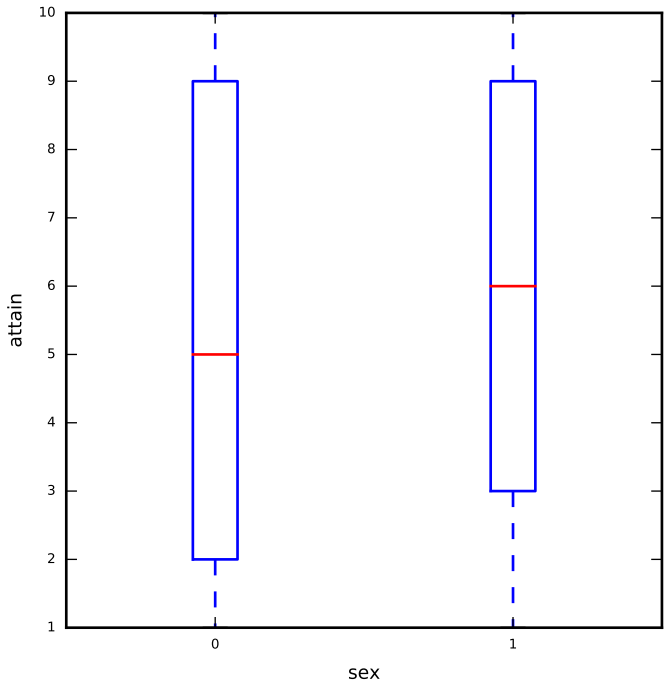

Here is a tabulation of the response, attain for predictor fed with category 1 having the largest mean and category 0 the smallest.

Category	N	Mean	Standard Deviation	Median
0	2489	5.419	3.012	5.0
1	946	6.362	3.074	7.0

The formal test is as follows:

There are two groups in the data:
The first group has 2489 observations with mean 5.419 standard deviation 3.012.
The second group has 946 observations with mean 6.362 standard deviation 3.076.
We are trying to test a hypothesis as to whether the two groups differ in their (population) means by a statistically significant amount. Statistical significance is related to how likely a result is to be a chance occurance. Here we are trying to differentiate between a real difference (no matter how small) and a difference that may have occurred due to the samples we have chosen.
The mean difference is 0.942 with the second group having the larger sample mean.
We need to quantify if this difference is large relative to the variability in the data. To do this we calculate the standard error of the difference. This is a function of the variabilities in the samples from group A and group B combined with their sample sizes. The bigger the 2 variabilities the larger the standard error, whilst the smaller the variability the smaller the standard error.
For our data the standard error of the mean difference is 0.117 and we divide our observed difference by this standard error to give a test statistic with value 8.068.
This test statistic is then compared to a t distribution with degrees of freedom equal to the sum of the sample sizes in each group (3435) - 2. In this case a t distribution with 3433. This t table has values of 1.961 for p=0.05 and 2.577 for p=0.01.

As 8.068 > 2.577 our p value is less than 0.01 and we have strong evidence to reject the null hypothesis (at the p=0.01 level).

The p-value is in fact less than 0.0001.

The t test assumes that the distribution of the response in each group follows a Normal distribution. We could check this by looking at histograms of the variable in each group. If we were concerned about the normality assumption then we could instead use a Mann Whitney (MW) test.

A Mann Whitney test works simply on the order (or ranks) of the responses across the two groups. So the response variable is firstly sorted and then each value is ranked. The ranks for each group are then summed and the value that is larger is compared with what would be expected if there was no difference between the groups.

In this case the MW U statistic is 969272 which for samples of size 2489 and 946 corresponds to a p value of less than 0.0001.

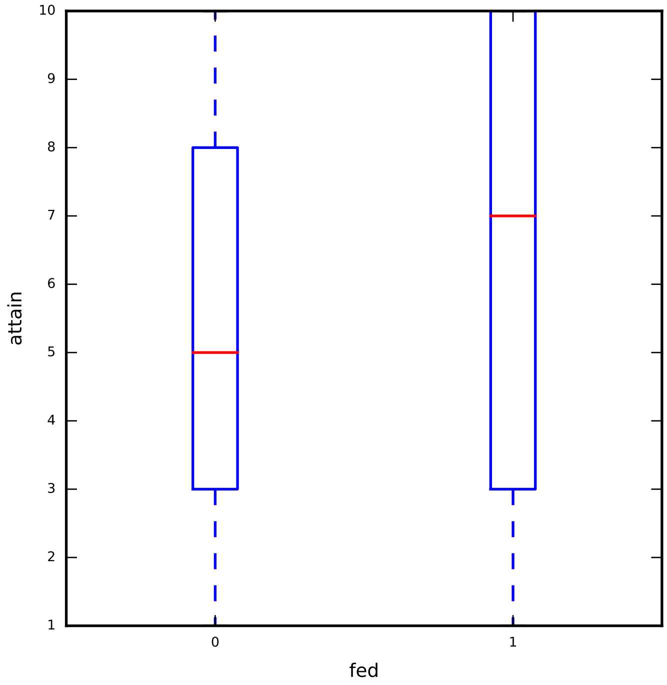

Here is a tabulation of the response, attain for predictor choice with category 3 having the largest mean and category 4 the smallest.

Category	N	Mean	Standard Deviation	Median
1	3069	5.71	3.061	5.0
2	212	6.434	2.906	6.0
3	3	8.667	1.886	10.0
4	151	3.921	2.502	3.0

The formal test is as follows:

	df	SS	MS	F
Between groups	3	617.5	205.8	22.41
Within groups	3431	31510.0	9.183	
Total	3434	32120.0	9.354	
Pooled within-group S.D.	3.03			
Between-group variance component	0.876			

For the ANOVA we are testing whether there are differences in the means of the response variable between the different groups. As shown in the table above this is done by constructing an ANOVA table that compares how much of the variability in the data is within the groups compared to between the groups. This results in a test statistic that follows an F distribution with 3 and 3431 degrees of freedom. This F table has values of 3.12 for p=0.05 and 4.287 for p=0.01.

As 22.415 > 4.287 our p value is less than 0.01 and we have strong evidence to reject the null hypothesis (at the p=0.01 level).

The p-value is in fact less than 0.0001.

Category	N	Mean	S.E.M.
1	3069	5.71	0.0547
2	212	6.434	0.208
3	3	8.667	1.75
4	151	3.921	0.247

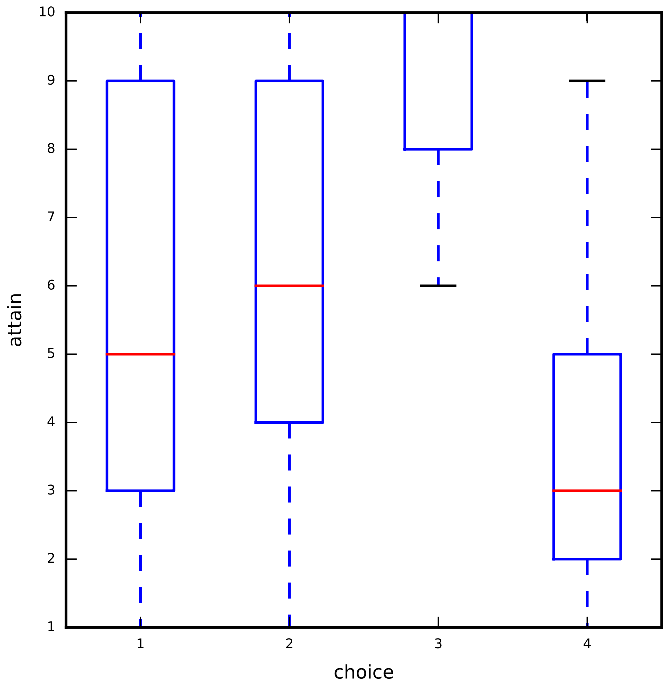

Here is a tabulation of the response, attain for predictor med with category 1 having the largest mean and category 0 the smallest.

Category	N	Mean	Standard Deviation	Median
0	2260	5.404	2.992	5.0
1	1175	6.207	3.113	6.0

The formal test is as follows:

There are two groups in the data:
The first group has 2260 observations with mean 5.404 standard deviation 2.993.
The second group has 1175 observations with mean 6.207 standard deviation 3.115.
We are trying to test a hypothesis as to whether the two groups differ in their (population) means by a statistically significant amount. Statistical significance is related to how likely a result is to be a chance occurance. Here we are trying to differentiate between a real difference (no matter how small) and a difference that may have occurred due to the samples we have chosen.
The mean difference is 0.803 with the second group having the larger sample mean.
We need to quantify if this difference is large relative to the variability in the data. To do this we calculate the standard error of the difference. This is a function of the variabilities in the samples from group A and group B combined with their sample

sizes. The bigger the 2 variabilities the larger the standard error, whilst the smaller the variability the smaller the standard error.

For our data the standard error of the mean difference is 0.111 and we divide our observed difference by this standard error to give a test statistic with value 7.263. This test statistic is then compared to a t distribution with degrees of freedom equal to the sum of the sample sizes in each group (3435) - 2. In this case a t distribution with 3433. This t table has values of 1.961 for p=0.05 and 2.577 for p=0.01.

As 7.263 > 2.577 our p value is less than 0.01 and we have strong evidence to reject the null hypothesis (at the p=0.01 level).

The p-value is in fact less than 0.0001.

The t test assumes that the distribution of the response in each group follows a Normal distribution. We could check this by looking at histograms of the variable in each group. If we were concerned about the normality assumption then we could instead use a Mann Whitney (MW) test.

A Mann Whitney test works simply on the order (or ranks) of the responses across the two groups. So the response variable is firstly sorted and then each value is ranked. The ranks for each group are then summed and the value that is larger is compared with what would be expected if there was no difference between the groups.

In this case the MW U statistic is 1129544 which for samples of size 2260 and 1175 corresponds to a p value of less than 0.0001.

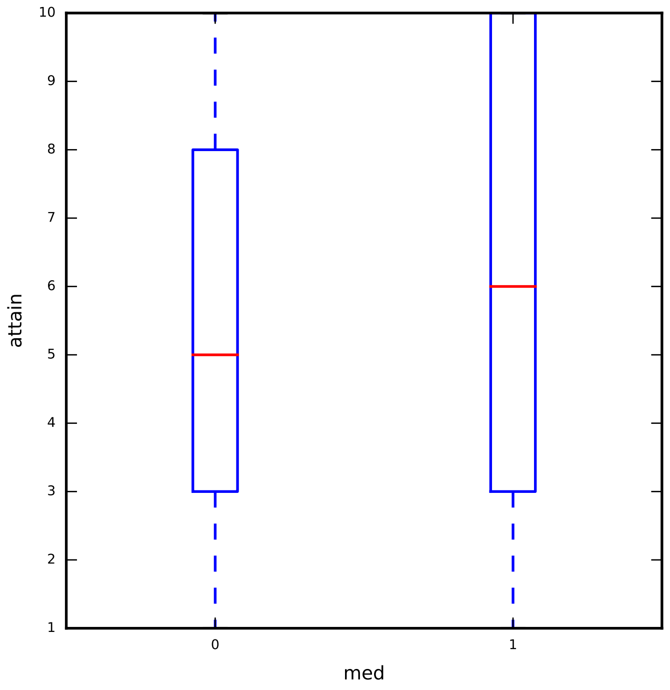

For the continuous predictors we can look at correlations with the response and scatterplots to see if there is a linear relationship.

Predictor : vrq

The Pearson correlation between attain and vrq is 0.719 (p value < 0.001).

The Spearman rank correlation between attain and vrq is 0.734 (p value < 0.001).

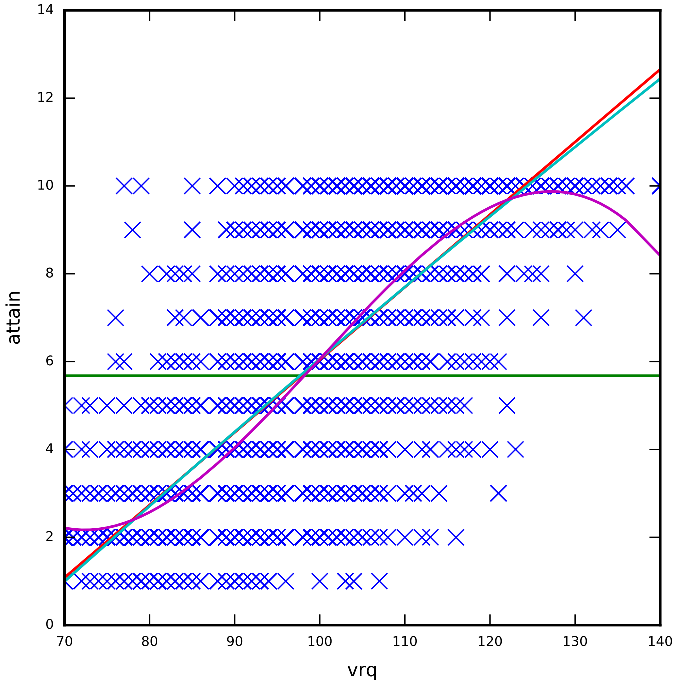

The graph includes best fitting curves for a constant, linear, quadratic and cubic relationship between attain and vrq. In this case a linear relationship is most appropriate.

Predictor : sc

The Pearson correlation between attain and sc is 0.249 (p value < 0.001).

The Spearman rank correlation between attain and sc is 0.323 (p value < 0.001).

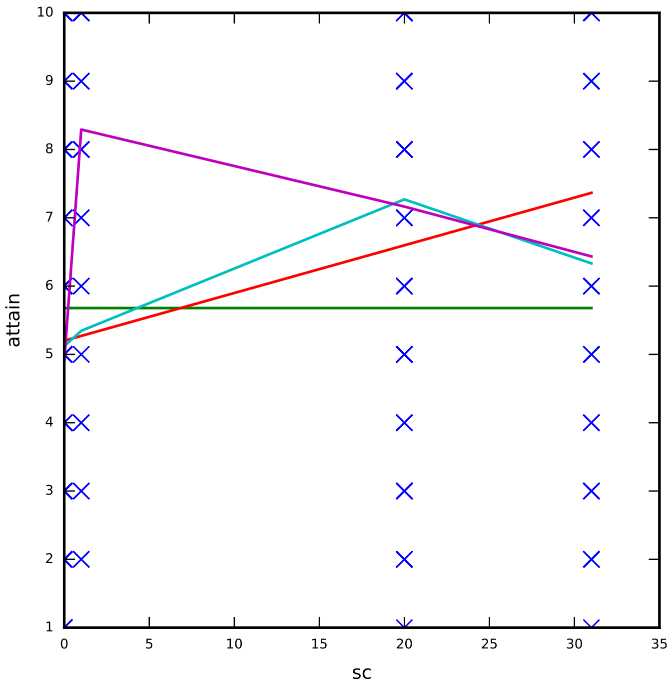

The graph includes best fitting curves for a constant, linear, quadratic and cubic relationship between attain and sc. In this case a polynomial of cubic or higher order is the most appropriate and you might consider including polynomial terms in your predictor list.

Choosing appropriate random classifications

We begin this section by deciding which of the possible random classifications to include in the modelling.

This is done by fitting all possible combinations and picking the model with the lowest DIC. All models are displayed along with their DIC values in the table below:

Higher-level classifications	DIC
None	17431.23
pid	17080.59
sid	17309.98
pid,sid	17048.06

The best model based on the DIC has classifications: pid,sid

As this is a multilevel modelling SAA we will also want to look at how the response is distributed across the levels of the model.

For this we will use the best model chosen above and look at how the variance is distributed across levels.

Variable	Coefficient	SE	ESS
Intercept	5.509	0.192	206
sid Variance	0.42	0.22	1155
pid Variance	1.145	0.211	1225
Level 1 Variance	8.12	0.197	3972

Here we see that the VPC for sid = 0.42/9.685 = 0.0434, so we see that sid effects explain 4.337% of the variability in attain.

Here we see that the VPC for pid = 1.145/9.685 = 0.118, so we see that pid effects explain 11.82% of the variability in attain.

Performing univariable modelling

Our next step in modelling now that we have a set of potential predictors is to consider models for each predictor in turn along with a random intercept at each chosen classification from the best model in the last section. In the fixed part these models simply contain an intercept and the particular predictor and so for continuous predictors will be multilevel linear regressions and for categorical predictors will be multilevel generalisations of ANOVAs. In the table below we summarise the modelling by showing the coefficients for each predictor along with the p value comparing the model with that predictor with a Null model. This Univariable modelling step will identify a set of candidate predictors to be taken forward into the next stage of modelling.

Variable	Coefficient	SD	ESS	p value	Significance
vrq	0.16	0.00283	2215	< 0.001	***
sc	0.054	0.00462	3374	< 0.001	***
sex_1	0.501	0.0985	4767	< 0.001	***
fed_1	0.741	0.112	3589	< 0.001	***
choice_2	0.796	0.215	2814	< 0.001	***
choice_3	2.481	1.678	4136		
choice_4	-1.895	0.243	3713		
med_1	0.671	0.105	3742	< 0.001	***

Which predictors we consider for the next stage of analysis will depend on their significance in the above table (but may in practice also depend on the size the effect and substantive interest of the variable though this is hard to automate). We will use a threshold on the p values associated with the predictors to decide whether to include the predictors in the next stage. Here we are currently using a threshold of 0.05. so the predictors to carry forward are: med, vrq, fed, choice, sc, and sex.

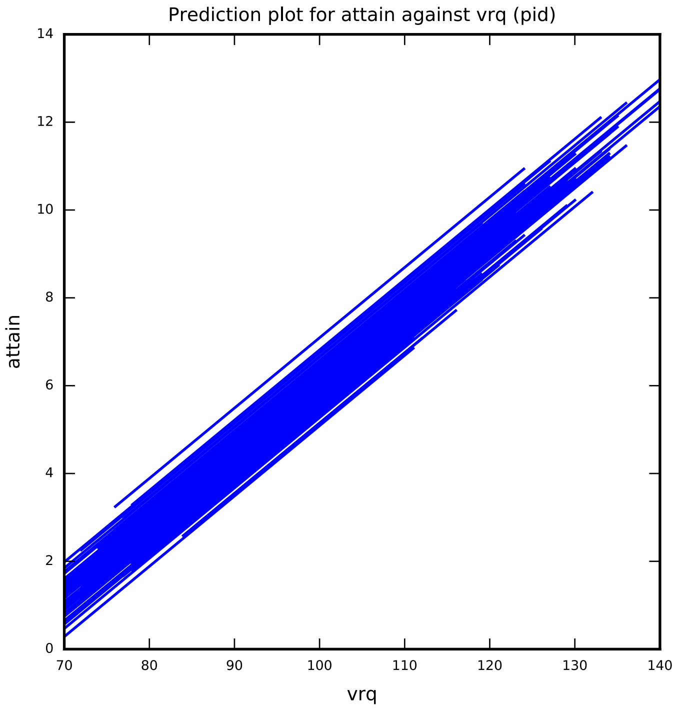

Prediction plot for attain against vrq (pid)
attain
vrq
0
2
4
6
8
10
12
14
70
80
90
100
110
120
130
140

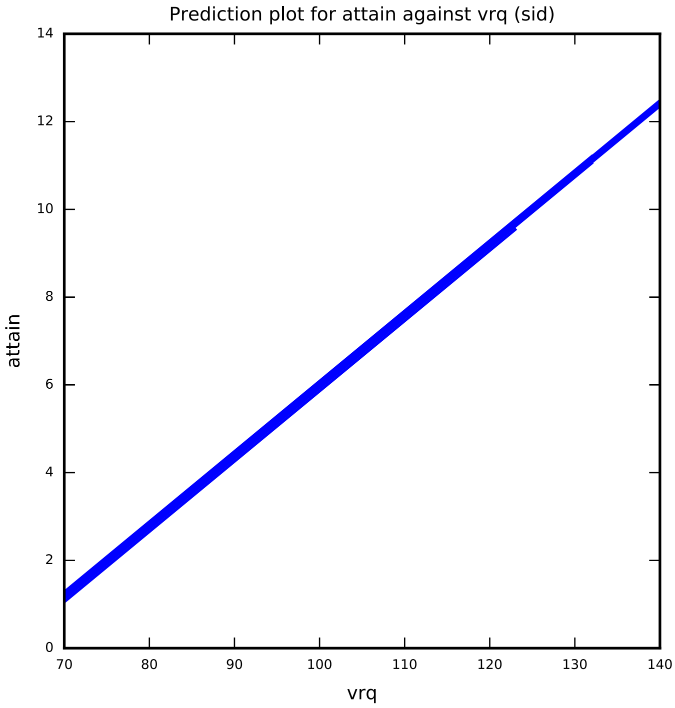

Prediction plot for attain against vrq (sid)
attain
vrq
0
2
4
6
8
10
12
14
70
80
90
100
110
120
130
140

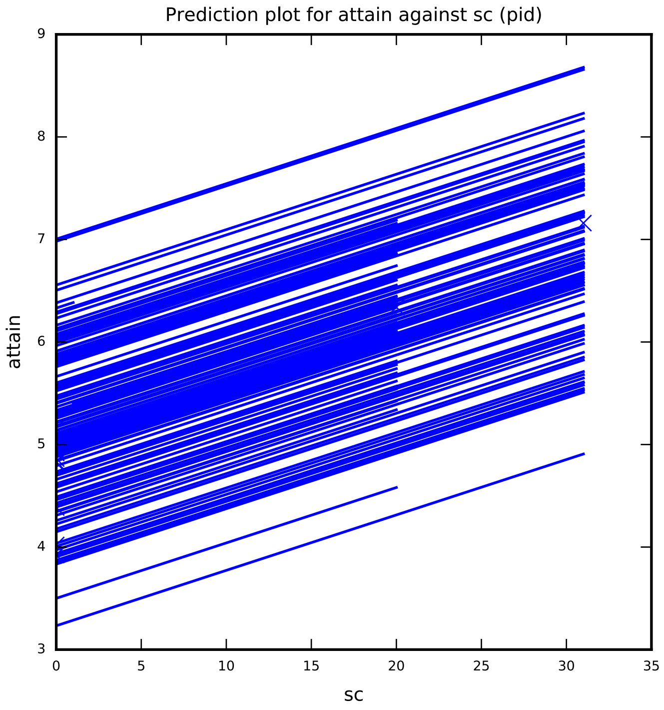
Prediction plot for attain against sc (pid)
attain
sc
3
4
5
6
7
8
9
0
5
10
15
20
25
30
35

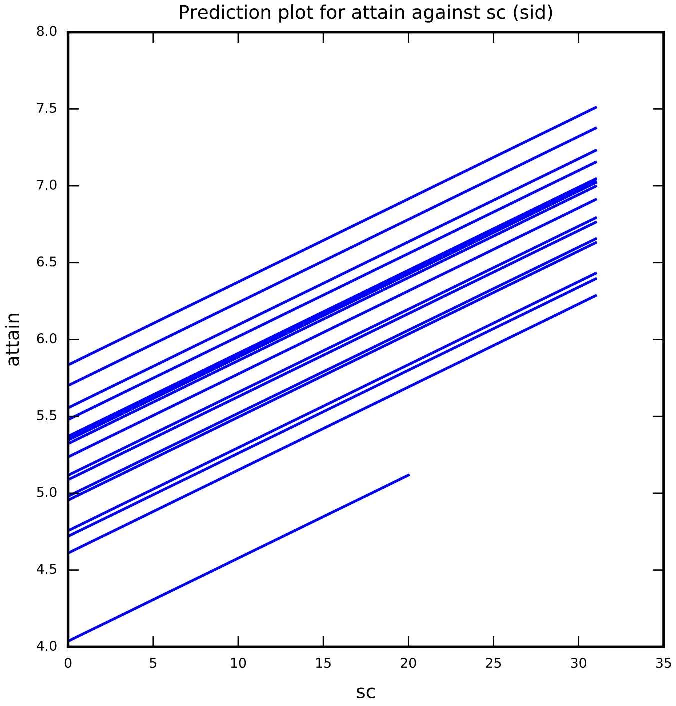
Prediction plot for attain against sc (sid)
attain
sc
4.0
4.5
5.0
5.5
6.0
6.5
7.0
7.5
8.0
0
5
10
15
20
25
30
35

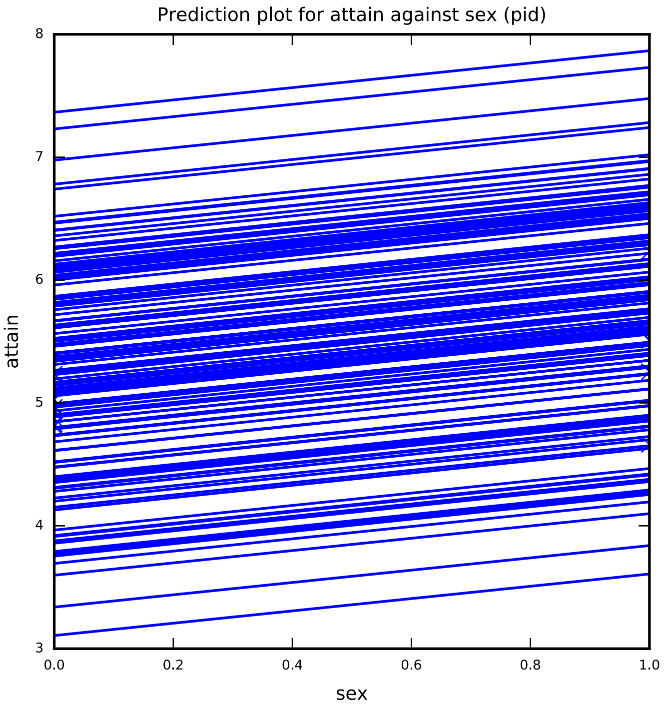
Prediction plot for attain against sex (pid)
attain
sex
8
7
6
5
4
3
0.0
0.2
0.4
0.6
0.8
1.0

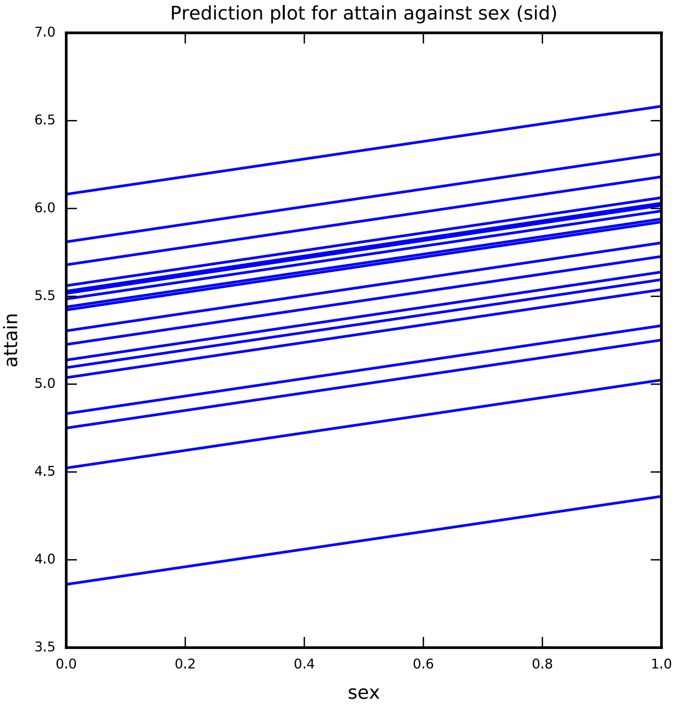
Prediction plot for attain against sex (sid)
attain
sex
7.0
6.5
6.0
5.5
5.0
4.5
4.0
3.5
0.0
0.2
0.4
0.6
0.8
1.0

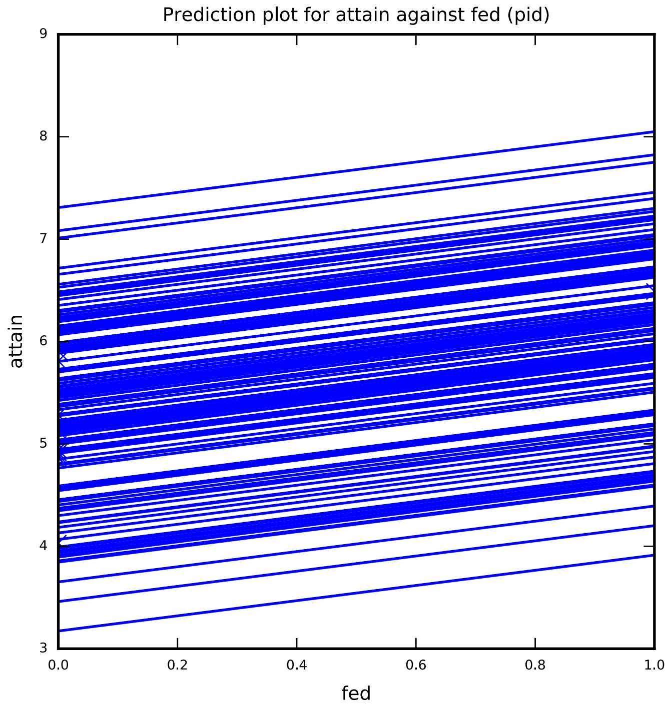
Prediction plot for attain against fed (pid)
attain
fed
9
8
7
6
5
4
3
0.0
0.2
0.4
0.6
0.8
1.0

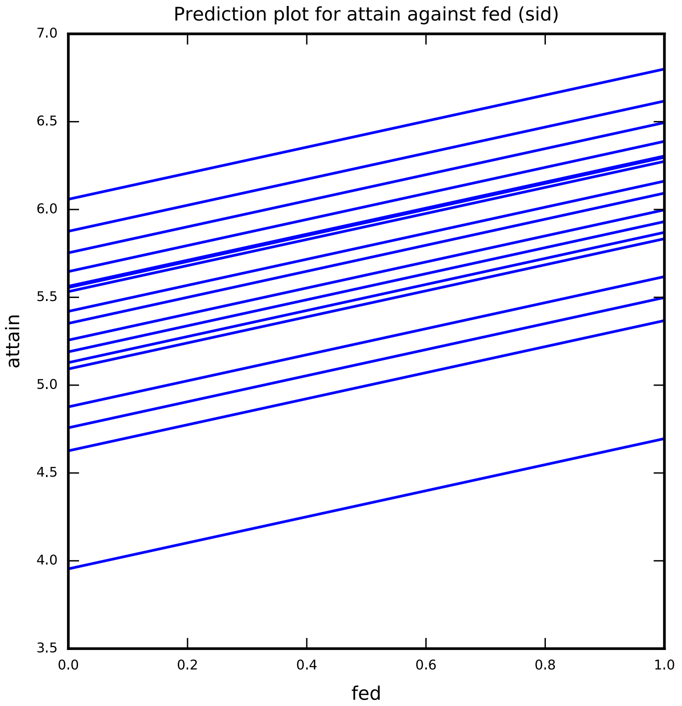
Prediction plot for attain against fed (sid)
attain
fed
7.0
6.5
6.0
5.5
5.0
4.5
4.0
3.5
0.0
0.2
0.4
0.6
0.8
1.0

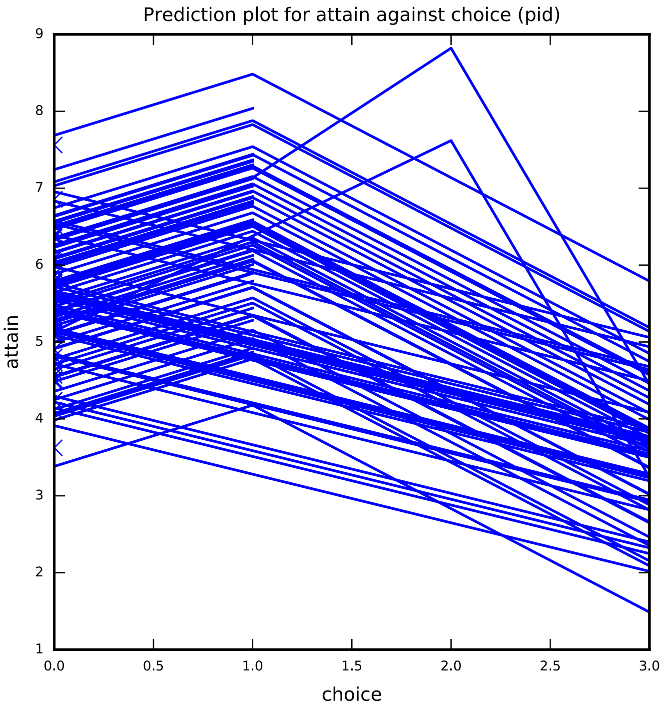
Prediction plot for attain against choice (pid)
attain
choice
9
8
7
6
5
4
3
2
1
0.0
0.5
1.0
1.5
2.0
2.5
3.0

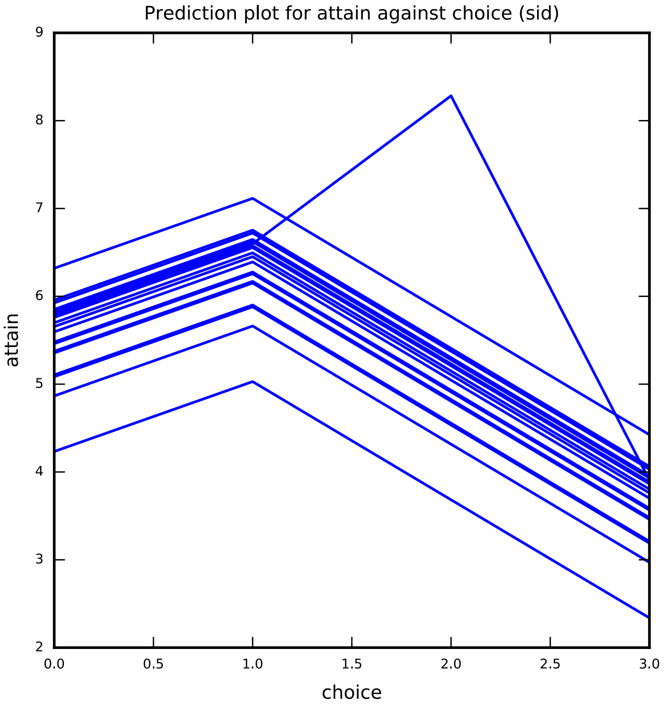
Prediction plot for attain against choice (sid)
attain
choice
9
8
7
6
5
4
3
2
0.0
0.5
1.0
1.5
2.0
2.5
3.0

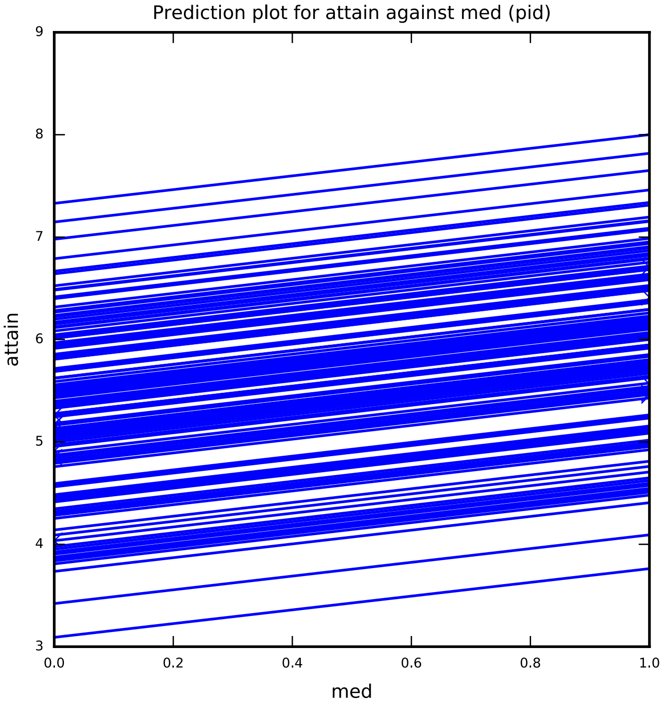
Prediction plot for attain against med (pid)
attain
med
9
8
7
6
5
4
3
0.0
0.2
0.4
0.6
0.8
1.0

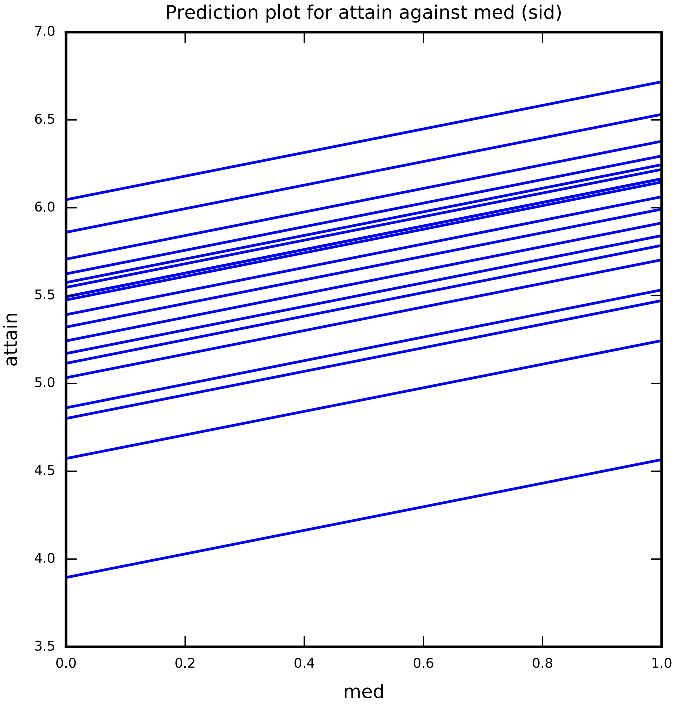
Prediction plot for attain against med (sid)
attain
med
7.0
6.5
6.0
5.5
5.0
4.5
4.0
3.5
0.0
0.2
0.4
0.6
0.8
1.0

Looking at correlations between predictors

Our next step is to check that none of the correlations between the predictor variables are too great as this could cause estimation problems when we add the predictors to the model together. To do this we look at all correlations between the predictor variables that have been identified as significant univariably and are thus candidates to be added to the model.

The correlations are as follows:

Variables	Correlation
(sc, vrq)	0.189
(sex_1, vrq)	0.096
(sex_1, sc)	-0.029
(fed_1, vrq)	0.102
(fed_1, sc)	0.095
(fed_1, sex_1)	0.055
(choice_2, vrq)	0.037
(choice_2, sc)	0.024
(choice_2, sex_1)	-0.002
(choice_2, fed_1)	0.015
(choice_3, vrq)	0.018
(choice_3, sc)	0.0
(choice_3, sex_1)	0.03
(choice_3, fed_1)	-0.018
(choice_3, choice_2)	-0.008
(choice_4, vrq)	-0.113
(choice_4, sc)	-0.04
(choice_4, sex_1)	-0.03
(choice_4, fed_1)	-0.046
(choice_4, choice_2)	-0.055
(choice_4, choice_3)	-0.006
(med_1, vrq)	0.096

Variables	Correlation
(med_1, sc)	0.048
(med_1, sex_1)	0.056
(med_1, fed_1)	0.513
(med_1, choice_2)	0.045
(med_1, choice_3)	0.02
(med_1, choice_4)	-0.041

Correlations greater than 0.8 (in magnitude) are worth looking at as they may result in model fitting problems when both predictors are included.

Performing multivariable model selection - random intercept models

In this next stage we will look at the best random intercepts model using only main effects for the variables to be considered. You have chosen to perform forward pass which is a quicker method than full forward selection. It may therefore not explore as many possible models. The predictor variables are considered in turn based on their significance in the univariable analysis and each is added to the current model. If the resulting model is a significant improvement then the predictor is kept in the model otherwise it is removed. Attention then moves on to the next predictor until all predictors are considered.

You have chosen to use Wald tests to compare models. These work by looked at estimates and standard error matrices for each predictor to assess significance and run quicker than the alternative methods as they do not need to run submodels. It should be noted that the Wald test is an unusual choice for MCMC estimation even though we offer it here.

The most significant predictor in the univariable analysis was sc so our starting point in multivariable modelling is the model:

$$\mathrm{attain}_i = \beta_0 \mathrm{sc}_i + \beta_1 \mathrm{intercept}_i + u^{(2)}_{0,pid_i} + u^{(3)}_{0,sid_i} + e_i$$

Variable	Coefficient	SD	ESS	p value	Significance
sc	0.054	0.00462	3374	< 0.001	***
Intercept	5.166	0.169	240		
Between sid Variance	0.312	0.173	783		
Between pid Variance	0.909	0.182	769		
Level 1 Variance	7.861	0.192	3124		

Adding variable sc was a significant improvement and so we retain it in the model.

Our next step is to consider adding variable choice to the current model.

$$\mathrm{attain}_i = \beta_0 \mathrm{sc}_i + \beta_1 \mathrm{choice\_2}_i + \beta_2 \mathrm{choice\_3}_i + \beta_3 \mathrm{choice\_4}_i + \beta_4 \mathrm{intercept}_i + u^{(2)}_{0,pid_i} + u^{(3)}_{0,sid_i} + e_i$$

Variable	Coefficient	SD	ESS	p value	Significance
sc	0.052	0.00455	2604	< 0.001	***
choice_2	0.718	0.209	2316	< 0.001	***
choice_3	2.491	1.596	2913		
choice_4	-1.78	0.235	3023		
Intercept	5.197	0.167	203		
Between sid Variance	0.285	0.164	430		
Between pid Variance	0.898	0.168	744		
Level 1 Variance	7.703	0.19	2540		

Adding variable choice was a significant improvement and so we retain it in the model.

Our next step is to consider adding variable vrq to the current model.

$$\text{attain}_i = \beta_0\text{sc}_i + \beta_1\text{choice\_2}_i + \beta_2\text{choice\_3}_i + \beta_3\text{choice\_4}_i + \beta_4\text{vrq}_i + \beta_5\text{intercept}_i + u^{(2)}_{0,pid_i} + u^{(3)}_{0,sid_i} + e_i$$

Variable	Coefficient	SD	ESS	p value	Significance
sc	0.0274	0.00342	3236	< 0.001	***
choice_2	0.458	0.153	3191	< 0.001	***
choice_3	1.386	1.194	3488		
choice_4	-0.634	0.177	4720		
vrq	0.155	0.00283	3240	< 0.001	***
Intercept	-9.739	0.281	2836		
Between sid Variance	0.0126	0.0154	201		
Between pid Variance	0.223	0.0545	479		
Level 1 Variance	4.171	0.102	2968		

Adding variable vrq was a significant improvement and so we retain it in the model.

Our next step is to consider adding variable fed to the current model.

$$\begin{aligned}\text{attain}_i =& \beta_0\text{sc}_i + \beta_1\text{choice\_2}_i + \beta_2\text{choice\_3}_i + \beta_3\text{choice\_4}_i + \beta_4\text{vrq}_i \\ &+ \beta_5\text{fed\_1}_i + \beta_6\text{intercept}_i + u^{(2)}_{0,pid_i} + u^{(3)}_{0,sid_i} + e_i\end{aligned}$$

Variable	Coefficient	SD	ESS	p value	Significance
sc	0.0266	0.00332	3316	< 0.001	***
choice_2	0.451	0.149	3168	< 0.001	***
choice_3	1.517	1.207	4405		
choice_4	-0.608	0.175	3446		
vrq	0.154	0.00282	2985	< 0.001	***
fed_1	0.331	0.0798	3310	< 0.001	***
Intercept	-9.746	0.281	2790		
Between sid Variance	0.0121	0.0139	207		
Between pid Variance	0.213	0.0521	574		
Level 1 Variance	4.155	0.102	3808		

Adding variable fed was a significant improvement and so we retain it in the model.

Our next step is to consider adding variable med to the current model.

$$\mathrm{attain}_i = \beta_0 \mathrm{sc}_i + \beta_1 \mathrm{choice\_2}_i + \beta_2 \mathrm{choice\_3}_i + \beta_3 \mathrm{choice\_4}_i + \beta_4 \mathrm{vrq}_i$$
$$+ \beta_5 \mathrm{fed\_1}_i + \beta_6 \mathrm{med\_1}_i + \beta_7 \mathrm{intercept}_i + u^{(2)}_{0,pid_i} + u^{(3)}_{0,sid_i} + e_i$$

Variable	Coefficient	SD	ESS	p value	Significance
sc	0.0266	0.00335	3122	< 0.001	***
choice_2	0.44	0.154	3374	< 0.001	***
choice_3	1.411	1.198	3840		
choice_4	-0.605	0.175	3659		
vrq	0.154	0.00281	3263	< 0.001	***
fed_1	0.219	0.0926	3746	0.018	*
med_1	0.204	0.0867	3927	0.019	*
Intercept	-9.764	0.277	2904		
Between sid Variance	0.00937	0.0125	305		
Between pid Variance	0.216	0.0517	604		
Level 1 Variance	4.148	0.101	3451		

Adding variable med was a significant improvement and so we retain it in the model.

Our next step is to consider adding variable sex to the current model.

$$\begin{aligned}\text{attain}_i =& \beta_0\text{sc}_i + \beta_1\text{choice\_2}_i + \beta_2\text{choice\_3}_i + \beta_3\text{choice\_4}_i + \beta_4\text{vrq}_i \\ &+ \beta_5\text{fed\_1}_i + \beta_6\text{med\_1}_i + \beta_7\text{sex\_1}_i + \beta_8\text{intercept}_i + u^{(2)}_{0,pid_i} + u^{(3)}_{0,sid_i} + e_i\end{aligned}$$

Variable	Coefficient	SD	ESS	p value	Significance
sc	0.0269	0.00338	3319	< 0.001	***
choice_2	0.442	0.152	3514	< 0.001	***
choice_3	1.354	1.202	3806		
choice_4	-0.598	0.177	4046		
vrq	0.154	0.00279	3001	< 0.001	***
fed_1	0.215	0.0928	3729	0.021	*
med_1	0.199	0.0879	3852	0.023	*
sex_1	0.12	0.0703	4044	0.089	N/S
Intercept	-9.781	0.278	2518		
Between sid Variance	0.0109	0.0139	206		
Between pid Variance	0.216	0.0535	602		
Level 1 Variance	4.145	0.101	3144		

Adding variable sex did not significantly improve the model, so we remove it from the model.

We have considered all variables so now run our final random intercepts model.

$$\begin{aligned}\mathrm{attain}_i = {} & \beta_0 \mathrm{sc}_i + \beta_1 \mathrm{choice\_2}_i + \beta_2 \mathrm{choice\_3}_i + \beta_3 \mathrm{choice\_4}_i + \beta_4 \mathrm{vrq}_i \\ & + \beta_5 \mathrm{fed\_1}_i + \beta_6 \mathrm{med\_1}_i + \beta_7 \mathrm{intercept}_i + u^{(2)}_{0,pid_i} + u^{(3)}_{0,sid_i} + e_i\end{aligned}$$

Variable	Coefficient	SD	ESS	p value	Significance
sc	0.0266	0.00335	3122	< 0.001	***
choice_2	0.44	0.154	3374	< 0.001	***
choice_3	1.411	1.198	3840		
choice_4	-0.605	0.175	3659		
vrq	0.154	0.00281	3263	< 0.001	***
fed_1	0.219	0.0926	3746	0.018	*
med_1	0.204	0.0867	3927	0.019	*
Intercept	-9.764	0.277	2904		
Between sid Variance	0.00937	0.0125	305		
Between pid Variance	0.216	0.0517	604		
Level 1 Variance	4.148	0.101	3451		

This is our final model.

Choosing interactions

You have chosen not to investigate interactions and so this page is empty.

Adding random slopes

You have chosen not to look at random slopes and so this page is blank.

Analysing the residuals

Here we look at the residuals from the model and plot them in various ways.

We start with level 1 residuals:

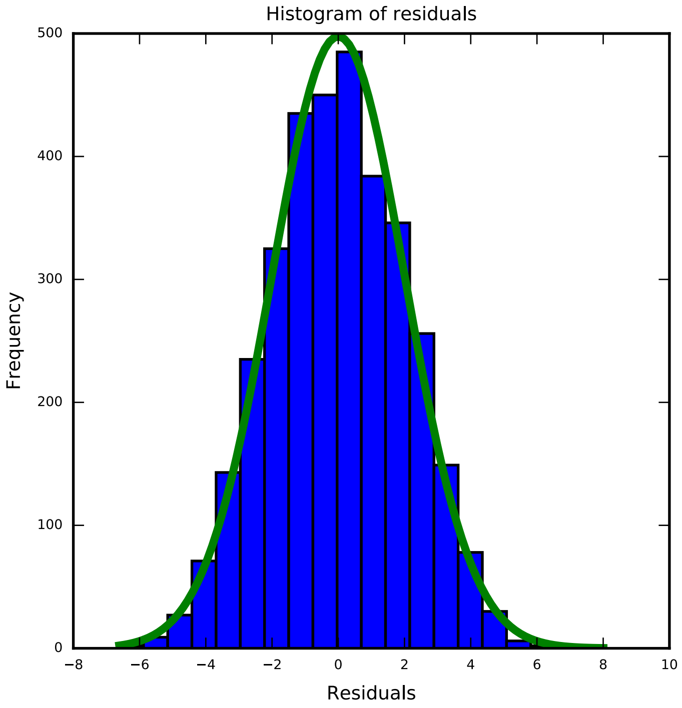

Here the distribution is reasonably symmetric with skewness value 0.016.

There are no obvious outliers in the residuals.

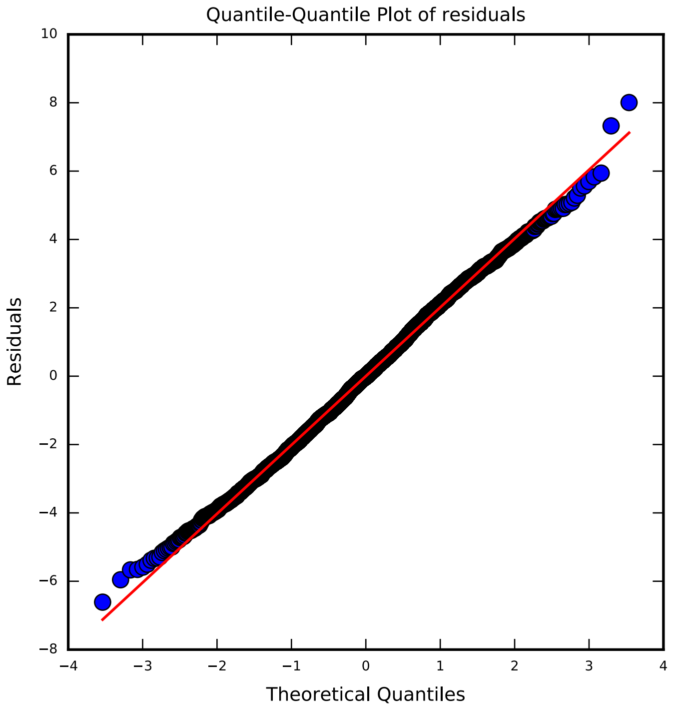

If the residuals are fairly normally distributed then the points in this graph should be close to the red line.

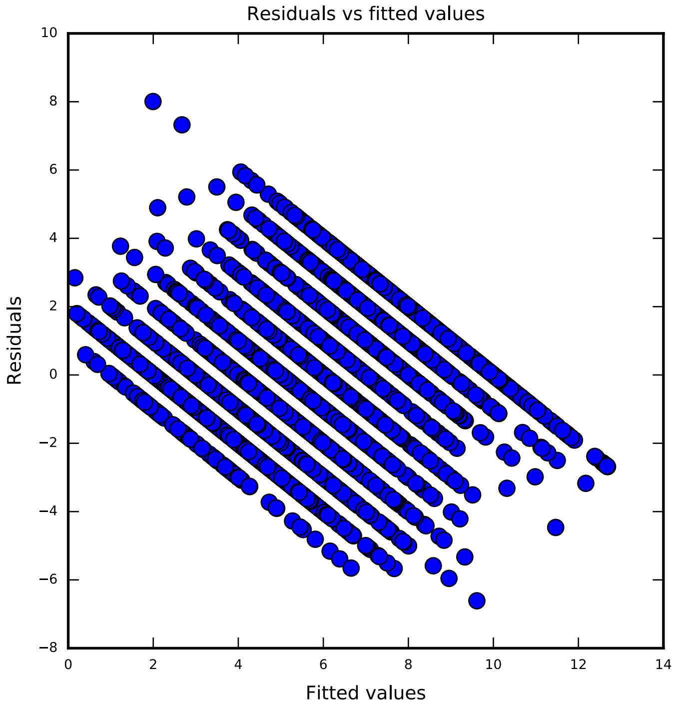

Here you should consider whether there are any patterns in this plot. Ideally we would like to see similar variability of the residuals across the range of fitted values.

Next the level 2 residuals for intercept:

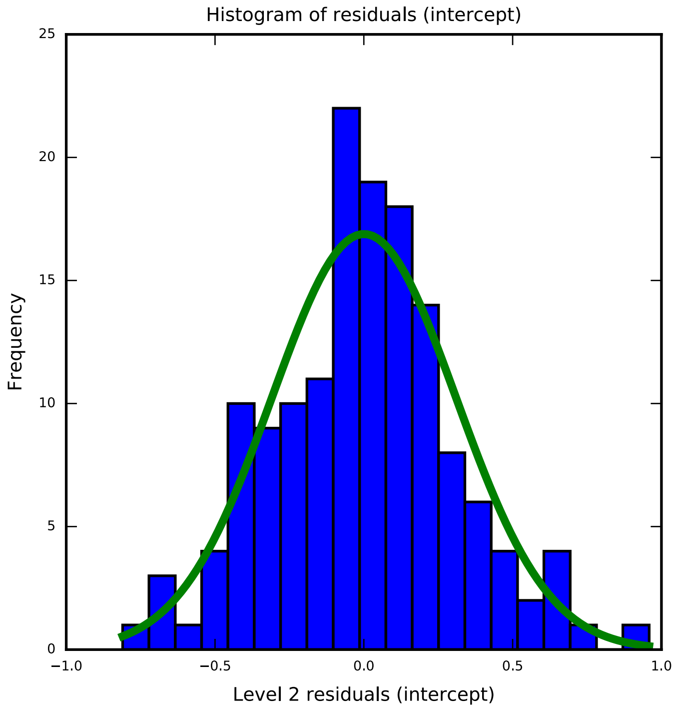

Here the distribution is reasonably symmetric with skewness value 0.118.

There are no obvious outliers in the residuals.

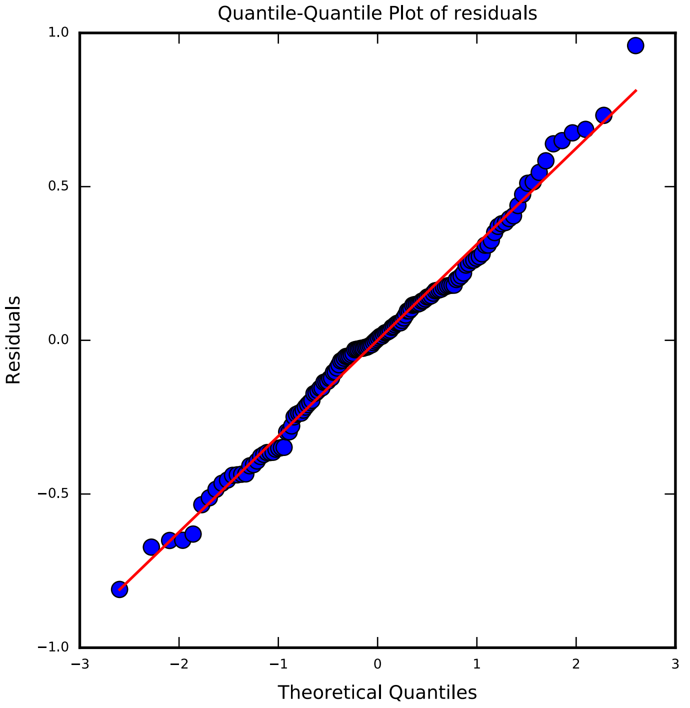

If the residuals are fairly normally distributed then the points in this graph should be close to the red line.

Next the level 3 residuals for intercept:

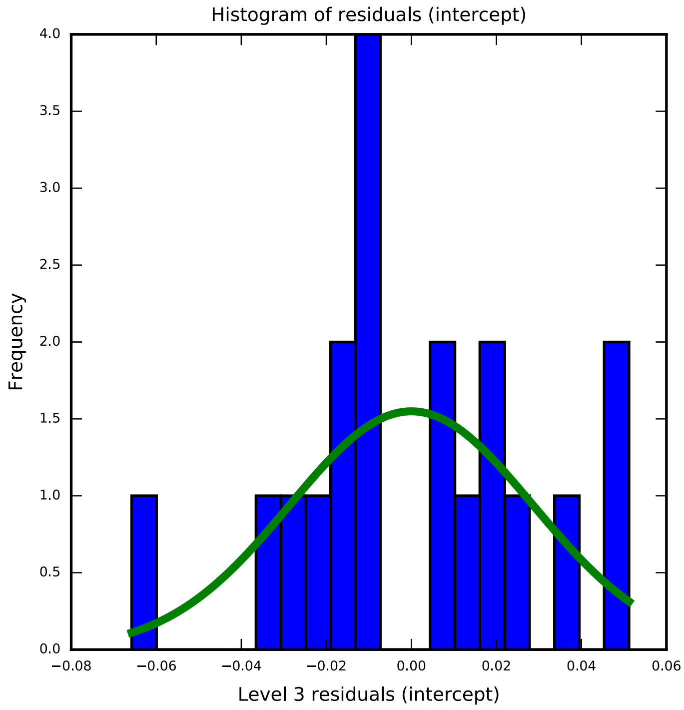

Here the distribution is reasonably symmetric with skewness value -0.068.

There are no obvious outliers in the residuals.

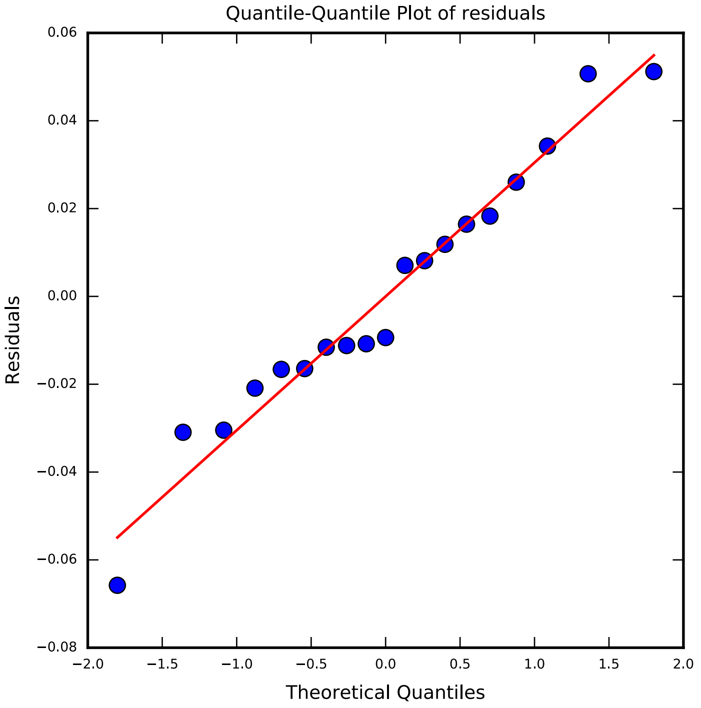

If the residuals are fairly normally distributed then the points in this graph should be close to the red line.

Looking at predictions

Having fitted a model with several predictors we might like to represent this model graphically. This is more difficult than when we have only one predictor and so for now we consider each predictor in turn and set all other predictors to their mean values.

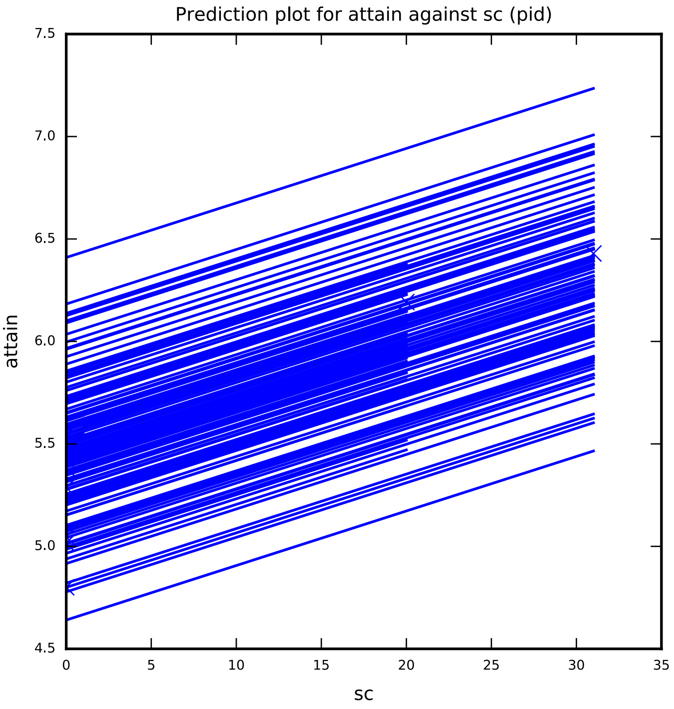

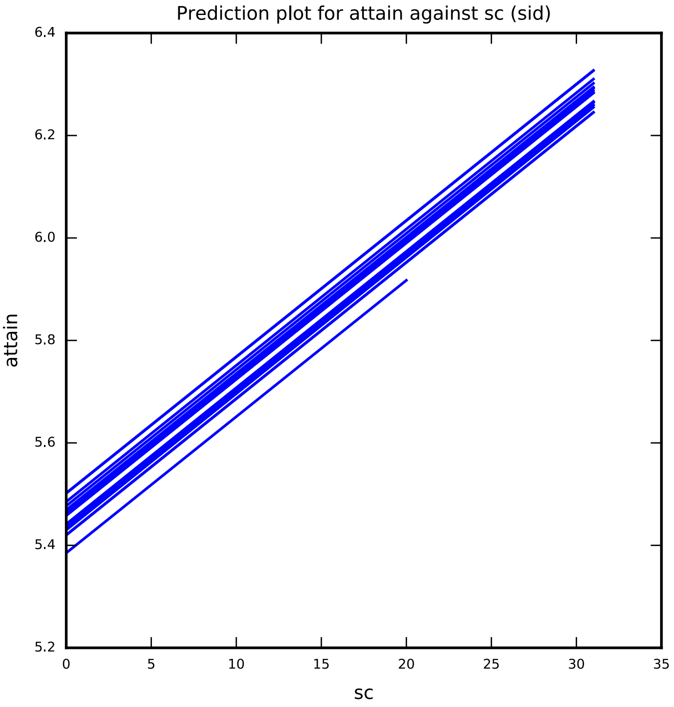

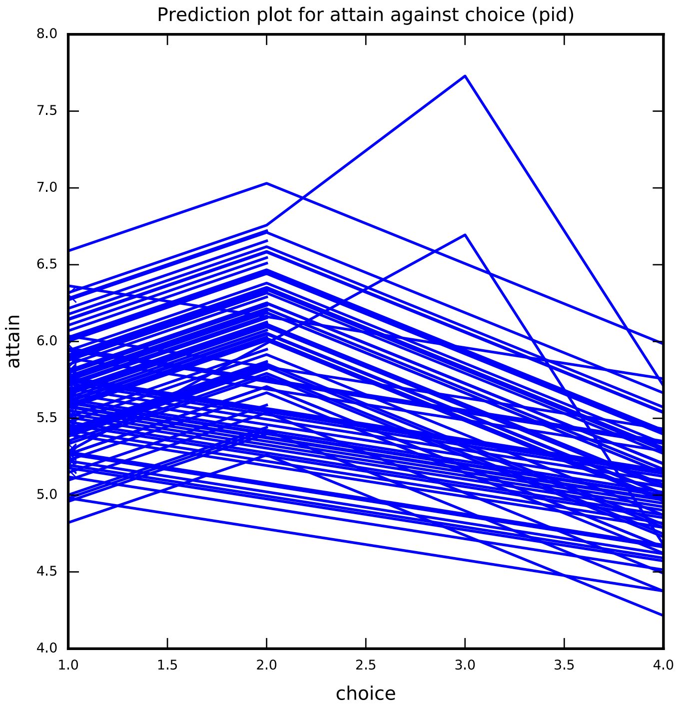
Prediction plot for attain against choice (pid)

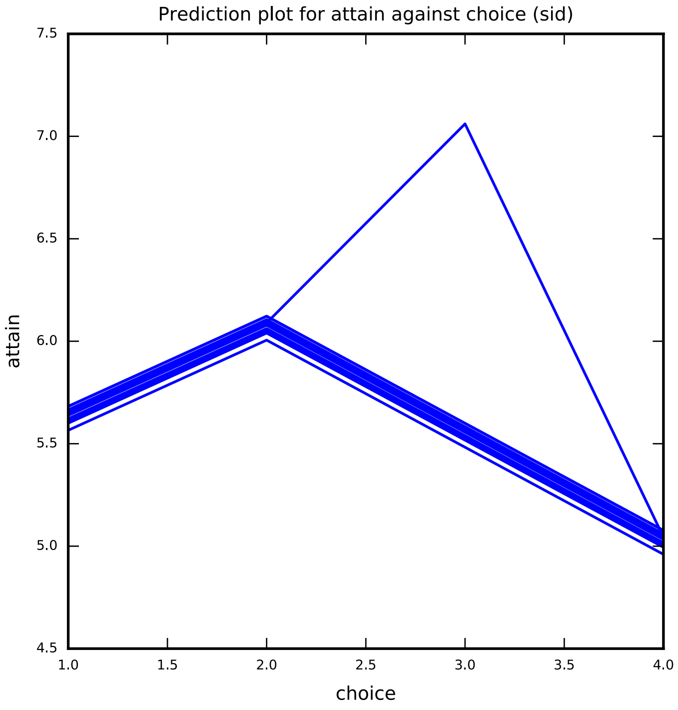
Prediction plot for attain against choice (sid)

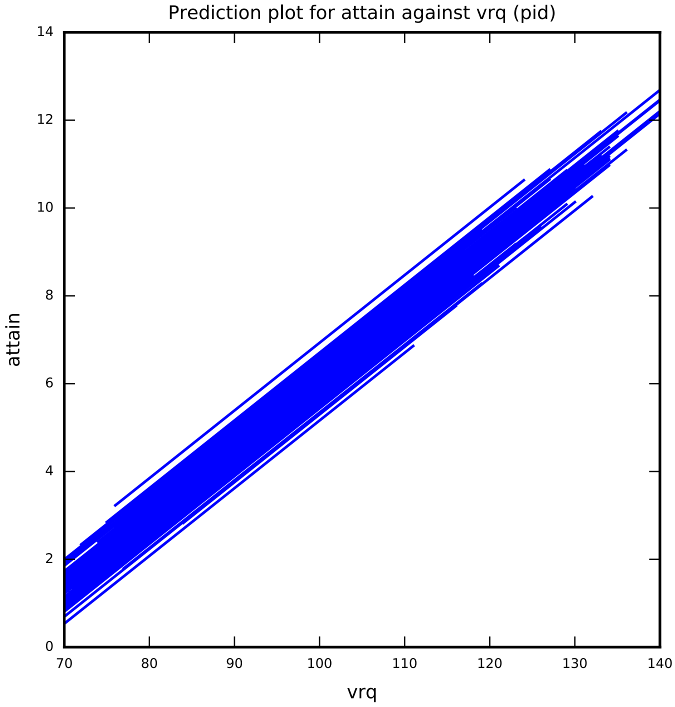

Prediction plot for attain against vrq (pid)

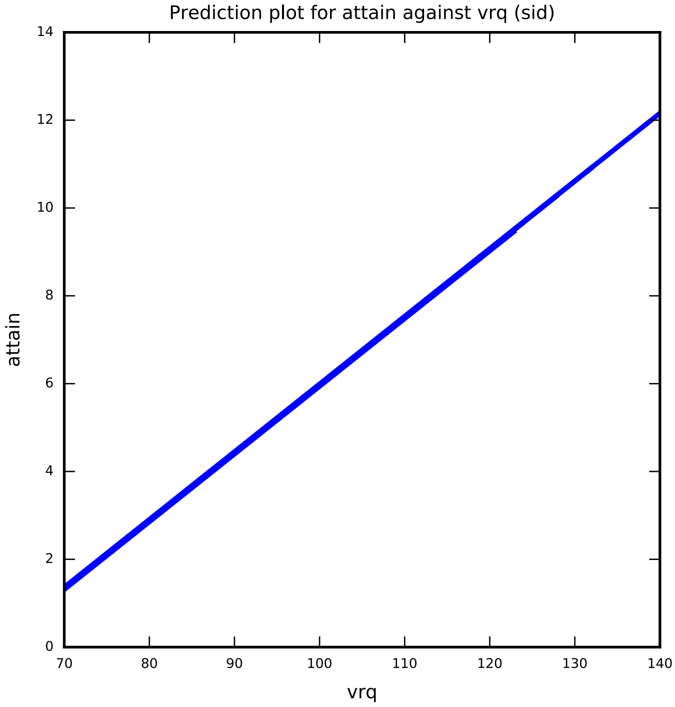

Prediction plot for attain against vrq (sid)

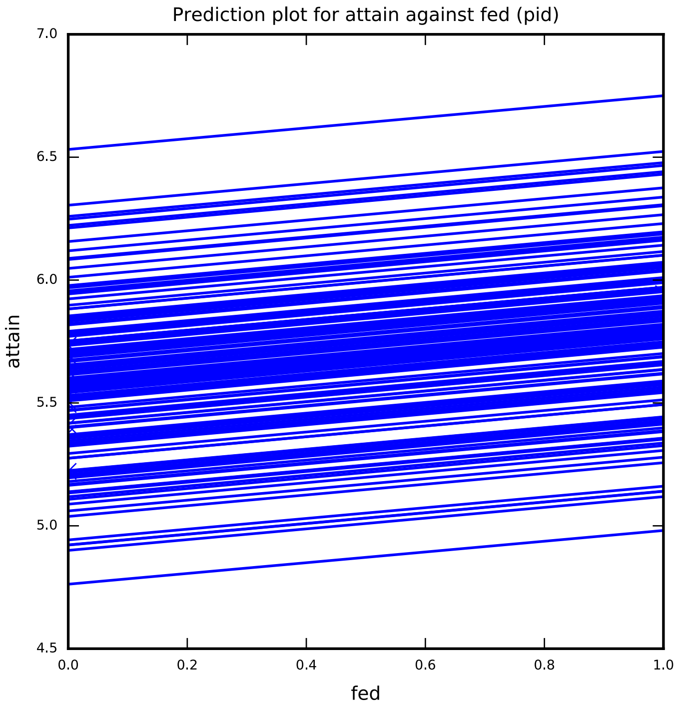
Prediction plot for attain against fed (pid)
attain
fed
7.0
6.5
6.0
5.5
5.0
4.5
0.0
0.2
0.4
0.6
0.8
1.0

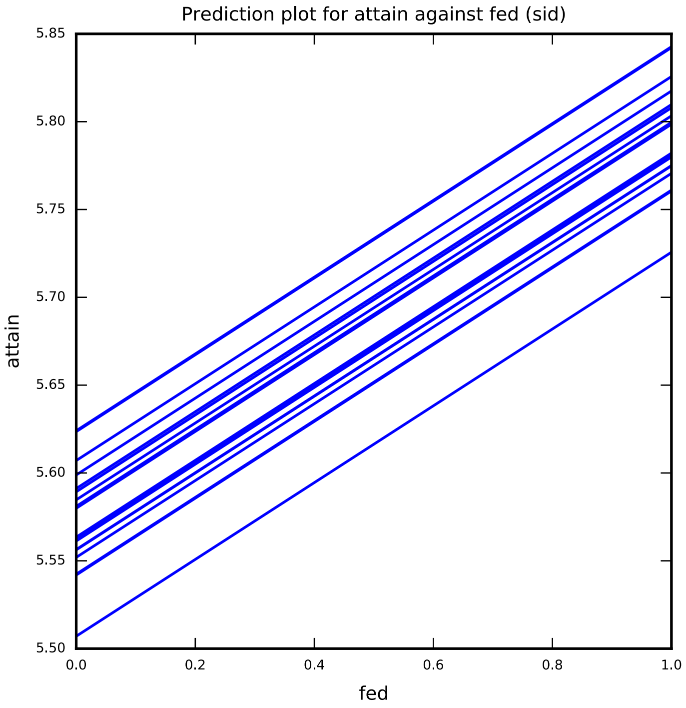
Prediction plot for attain against fed (sid)
attain
fed
5.85
5.80
5.75
5.70
5.65
5.60
5.55
5.50
0.0
0.2
0.4
0.6
0.8
1.0

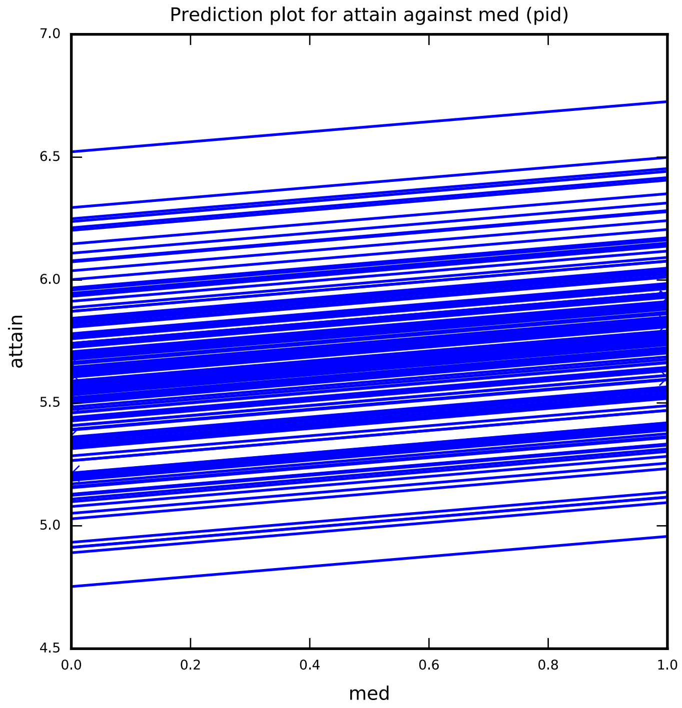
Prediction plot for attain against med (pid)
attain
med
7.0
6.5
6.0
5.5
5.0
4.5
0.0
0.2
0.4
0.6
0.8
1.0

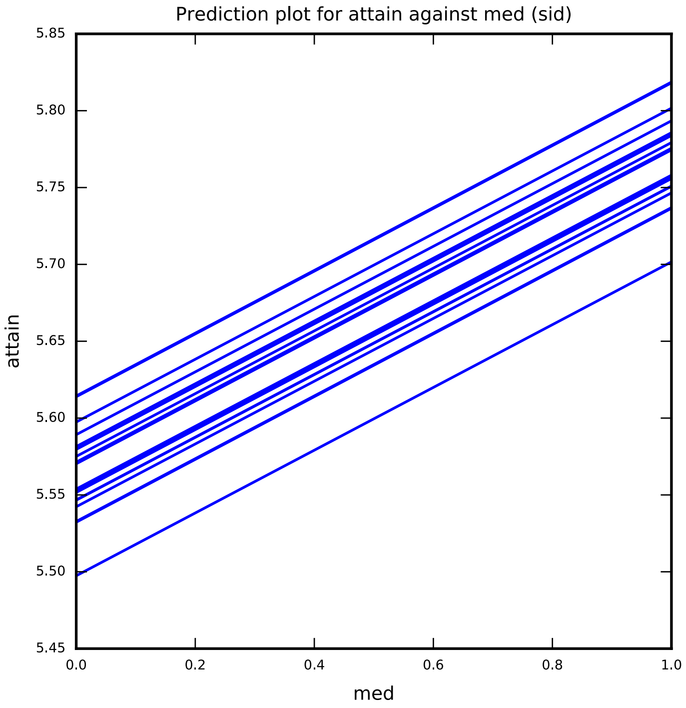
Prediction plot for attain against med (sid)
attain
med
5.85
5.80
5.75
5.70
5.65
5.60
5.55
5.50
5.45
0.0
0.2
0.4
0.6
0.8
1.0