

Applied Regression Analysis: Final Report

Due on May 9, 2017

with Joshua French, Ph.D.

MATH 4387.001

by

Faith Blair

Luke Mattingly

Wyatt “Miller” Miller

INTRODUCTION

Why should someone be interested in this problem? (Why is it interesting to analyze this data set?)

As students prepare for future semesters they typically want the most information they can get before choosing their courses. For some students having a good instructor is very important to them. Resources like *RateMyProfessors.com* offer students a way to “rate” their previous instructors, in order to provide useful information for future students looking to take a course with that teacher. While it could be argued that the opinions of these students are subjective, and that the website could be augmented by non-students (outside the population of interest), the success of *RateMyProfessors.com* would imply a certain validity in their model of perceived quality.

Background information people need to know to understand the problem and/or data set.

This data is collected by presumed students going to *RateMyProfessors.com* and rating a professor. We will be looking at a data set that consists of 366 observations of professor ratings on a single Midwestern college campus. While this sample is far too constrained to formulate conclusions about the entire population of students in the country, it will at least give us a better understanding of perceived quality on this campus and will provide conjecture material for a larger study.

This information was collected between the years of 1999 and 2009, and asks the student to provide observational data, such as gender, number of years teaching, perceived easiness, clarity, etc. Each instructor in this data set has received at least 10 student ratings.

Is there any prior research on your topic that might be helpful for the audience?

Student Consensus on RateMyProfessors.com by Bleske-Rechek, et al (2011)¹, used this data set for an in-depth analysis on a similar topic.

From where did the data come? Is this an experiment or observational study? Who collected the data? Why was the data collected? (if you werent the one doing the collecting)

The original data is readily available and public at *RateMyProfessors.com*. This data is observational, as the input comes from online reviews collected on the website. It should be noted that the data set we’ll be working with has further calculated variables (such as standard deviation) that are not normally available on the *RateMyProfessors.com* website - presumably calculated for the study mentioned above. The data was collected by Bleske-Rechek, et al to, “investigate reliability in student ratings by analyzing variance in the ratings of instructor quality and easiness.”

What are the questions of interest that you hope to answer?

1. Are Helpfulness, Clarity, Easiness, and Rater interest all statistically significant in the model?
2. Does Gender have an effect on perception of Overall Quality, given the model of best fit?
3. Does the Chili Pepper have an effect on perception of Overall Quality, given the model of best fit?

1. DATA SUMMARY The raw project data includes the summaries of the ratings of 366 instructors at one large campus in the Midwest from Bleske-Rechek and Fritsch (2011). Each instructor included in the data had at least 10 ratings over a several year period. **Students provided ratings on 5 point scales.** The data file provides the averages ratings and additional characteristics of the instructors.

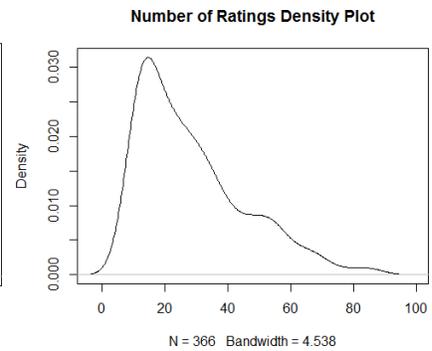
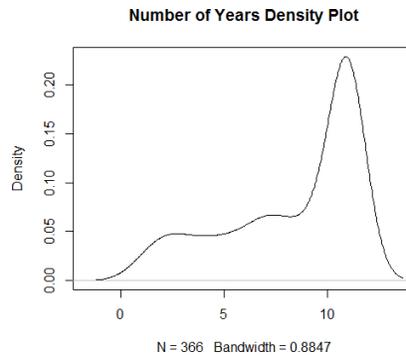
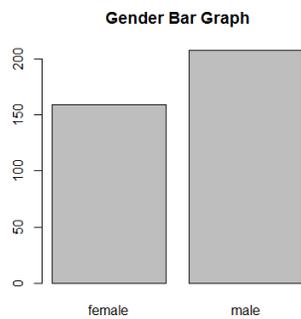
Field	Description (directly from R)
<code>gender</code>	Instructor gender, a factor with levels female male
<code>numYears</code>	A numeric vector, number of years in which this instructor had ratings between 1999 and 2009
<code>numRaters</code>	Number of ratings
<code>numCourses</code>	Number of different course titles included in the rating for this instructor
<code>pepper</code>	A factor with levels no and yes. In addition to rating for quality, instructors are rated as attractive or not. A value of yes means that the consensus is that the instructor is attractive
<code>discipline</code>	A factor with levels Hum for humanities, SocSci for social sciences, STEM for science, technology, engineering and mathematics and Pre-prof for professional training
<code>dept</code>	A qualitative variable with department names
<code>quality</code>	Average overall quality rating, between 1, worst, to 5, best (calculated)
<code>helpfulness</code>	Average helpfulness rating, between 1, worst, to 5, best (surveyed)
<code>clarity</code>	Average clarity rating, between 1, worst, to 5, best (surveyed)
<code>easiness</code>	Average easiness rating, between 1, worst, to 5, best (surveyed)
<code>raterInterest</code>	Average rater interest, between 1, lowest, to 5, highest (surveyed)
<code>sdQuality</code>	Standard Deviation of quality rating
<code>sdHelpfulness</code>	Standard Deviation of helpfulness rating
<code>sdClarity</code>	Standard Deviation of clarity rating
<code>sdEasiness</code>	Standard Deviation of easiness rating
<code>sdRaterInterest</code>	Standard Deviation of the Rater's interest

Next, we see numerical and graphical summaries of the project data set. The data values output in a preliminary summary were consistent with the data, i.e., the min, max, mean etc were within the bounds of the data with no obvious anomalies. There were no missing values or extreme data values we could detect.

There is some natural skewness in some of the data - number of raters, years of experience etc, which is to be expected given the type of data collected.

Variable	Min	1 st Q	Median	Mean	3 rd Q	Max	[Other]	[Other]
gender							159 (female)	207 (male)
numYears	1.00	6.00	10.00	8.35	11.00	11.00		
numRaters	10.00	15.00	24.00	28.58	37.00	86.00		
numCourses	1.00	3.00	4.00	4.25	5.00	12.00		
pepper							46 (yes)	320 (no)
quality	1.41	2.936	3.61	3.58	4.25	4.98		
helpfulness	1.36	3.07	3.66	3.63	4.35	5.00		
clarity	1.33	2.87	3.60	3.53	4.21	5.00		
easiness	1.39	2.55	3.15	3.14	3.69	4.90		
raterInterest	1.10	2.93	3.31	3.31	3.69	4.91		

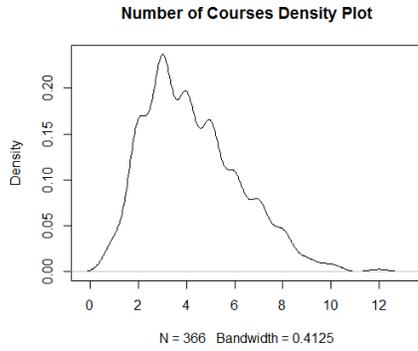
Next we'll inspect the graphical summaries:



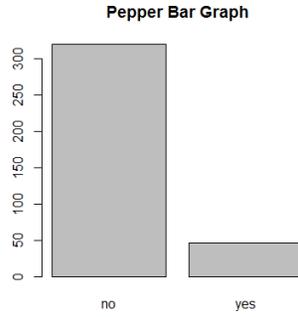
From the bar graph we learn that there are more male instructors than female in the data.

The number of years density plot is skewed left, so we can conclude that the majority of professors have been teaching for at least 8 years between 1999 and 2009. We can see that the most frequent number of years taught between 1999 and 2009 is 11 years.

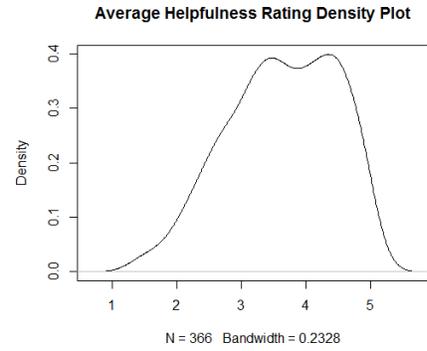
The number of ratings density plot is skewed right, so we can conclude that the majority of the professors have less than 40 total ratings, with the most frequent number of ratings falling between 10 and 20.



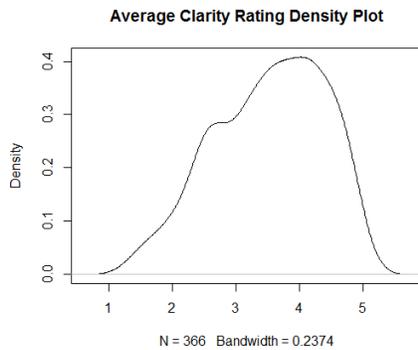
The number of courses density plot is skewed right, so we can conclude that the majority of professors have less than 6 courses that are rated. From the graph we can see the most frequent number of courses is 3. We can also see that there are a few unusually large values.



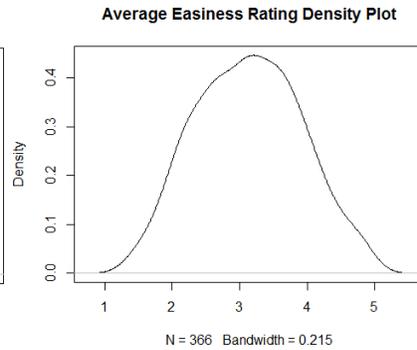
From the bar graph we learn that more instructors were rated as not attractive than those that were related as attractive.



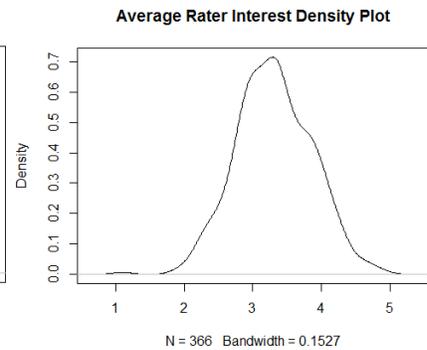
The average helpfulness rating density plot is bimodal, with the two most frequent average helpfulness ratings falling a little higher than 3, and at about 4.5.



The average clarity rating density plot is skewed left so we conclude the majority of instructors have high average clarity ratings, with the most frequent average rating falling around 4.



The average easiness rating density plot is normally distributed, and centered around an average rating of about 3.



Average rater interest appears to be normally distributed, with a few unusually low ratings.

Please note that we are immediately omitting all standard deviation calculations, as these were calculated specifically for use in the Bleske-Rechek et al. paper [1]. While these are used to produce reliable and interesting results in their research, it is beyond the scope of our paper and also not part of the original data set. We are also omitting departmental qualitative variables and the response, `quality`.

2. COLLINEARITY ANALYSIS

MODEL SUMMARY *Please note: we analyze the summary for completeness' sake. It is not used as a primary decision tool in our analysis. Also, `gender` and `pepper` variables are omitted due to binary nature.* Little is gleaned from the summary in terms of correlation (Appendix, Figure 2.1). Of note are the p -values of `helpfulness` and `clarity`, which match precisely - this is improbable behavior. Furthermore, the only other regressor that maintains a strong p -value is `numRaters`.

PAIRWISE CORRELATION *Please note: we analyze the covariance matrix for completeness' sake. It is not used as a primary decision tool in our analysis. Also, `gender` and `pepper` variables are omitted due to binary nature.* This analysis (Appendix, Figure 2.2) offers a more telling story: `clarity` and `helpfulness` appear to maintain a strong correlation from both the plot and the pairwise correlation matrix (0.92), while moderate correlations appear to exist between `helpfulness` and `easiness`, and `clarity` and `easiness`.

VARIANCE INFLATION FACTOR (VIF) This analysis (Appendix, Figure 2.3) begins to clarify some of the issues in the original model: `helpfulness` and `clarity` appear to be correlated. It is likely that a recommendation to omit these two variables will be made.

Furthermore, it was discovered that `helpfulness` and `clarity` are not only correlated, but that they directly produce `quality` by means of a linear transformation (averaging). While this offers an explanation of collinearity, it also provides motive to omit both variables since leaving either one in would provide systemic issues in our analysis.

CONDITION INDICES Our final collinearity analysis (Appendix, Figure 2.4) begins with the original model. Note that the binary, qualitative variables of `gender` and `pepper` have been omitted from this model; we wish to reserve them for later hypothesis testing.

We believe there is sufficient, systemic reason to omit both `helpfulness` and `clarity` immediately from the model. After removal, we note there still exists a relatively high condition number. Between `easiness` and `raterInterest`, the removal of `raterInterest` results in an acceptable condition number ($27.831 < 30$).

MODEL STATEMENT Thus our current model, based solely off of correlation (and incidental systemic

removal) is:

$$\hat{y}_{quality} = \beta_0 + \beta_1 x_{numYears} + \beta_2 x_{numRaters} + \beta_3 x_{numCourses} + \beta_4 x_{easiness} + \beta_5 x_{gender} + \beta_6 x_{pepper}.$$

3. MODEL SELECTION

Variable Selection allows us to find the “best” subset of predictors. We want the simplest model that adequately explains the data. Unnecessary regressors will add noise to all model estimates.

Taking our colinearity analysis into consideration, we are given the model of:

$$\hat{y}_{quality} = \beta_0 + \beta_1 x_{numYears} + \beta_2 x_{numRaters} + \beta_3 x_{numCourses} + \beta_4 x_{easiness} + \beta_5 x_{gender} + \beta_6 x_{pepper}.$$

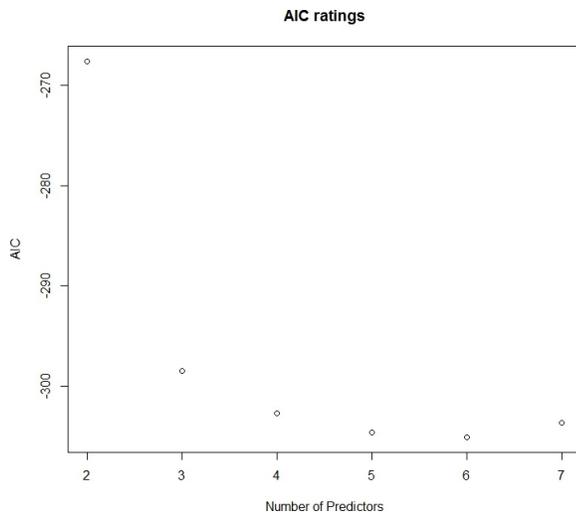
When performing Model Selection we need to fulfill two aspects 1. Search Strategy 2. Criterion for model comparisons.

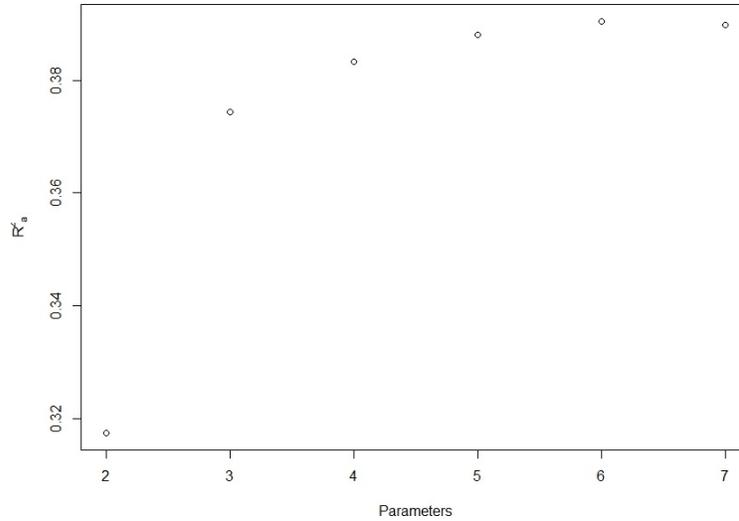
We choose to begin with an exhaustive search, which looks at all possible models using all available regressors, this is the best subset searching strategy due to having few variables.

	numYears	numRaters	numCourses	easiness	genderMale	pepperYes
1	FALSE	FALSE	FALSE	TRUE	FALSE	FALSE
2	FALSE	FALSE	FALSE	TRUE	FALSE	TRUE
3	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE
4	TRUE	FALSE	FALSE	TRUE	TRUE	TRUE
5	TRUE	FALSE	TRUE	TRUE	TRUE	TRUE
6	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE

Next let’s choose which model is best for us. We’ll use Criterion such as AIC, Adjusted R^2 , and Cp Statistics Our goal is to choose a model which minimizes AIC, the first term in the AIC is based on RSS which is made by improving the fit.

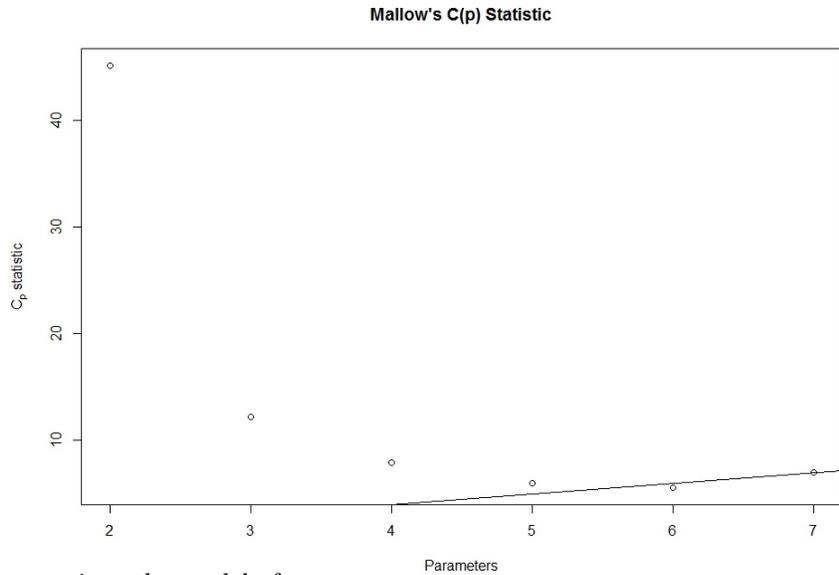
In plot to the right, we choose a number of predictors based on minimizing AIC. We observe that when $P = 6$ it suggests the five predictor model from above as being the best.





Next we compare the Adjusted R^2 to Parameters. We're looking for the largest Parameter value. We can see in the plot to the left that $P=6$ yet again is suggesting the same five predictor model as our AIC criteria.

Now we look at Mallows' Cp Statistic (right). For the Cp stat we want the smallest Cp value that is closest to the line $x=y$. Yet again we see it indicates 6 Parameters, meaning the five predictor model is ideal.



As a result of Model Selection, we are given the model of:

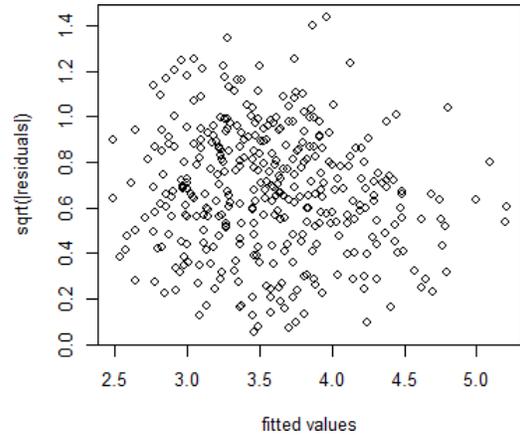
$$\hat{y}_{quality} = \beta_0 + \beta_1 x_{numYears} + \beta_2 x_{numCourses} + \beta_3 x_{easiness} + \beta_4 x_{gender} + \beta_5 x_{pepper}.$$

Since variable selection can be affected by both Outliers/Influential observations, as well as the transformation of variables, we will next iterate over our model by performing diagnostics.

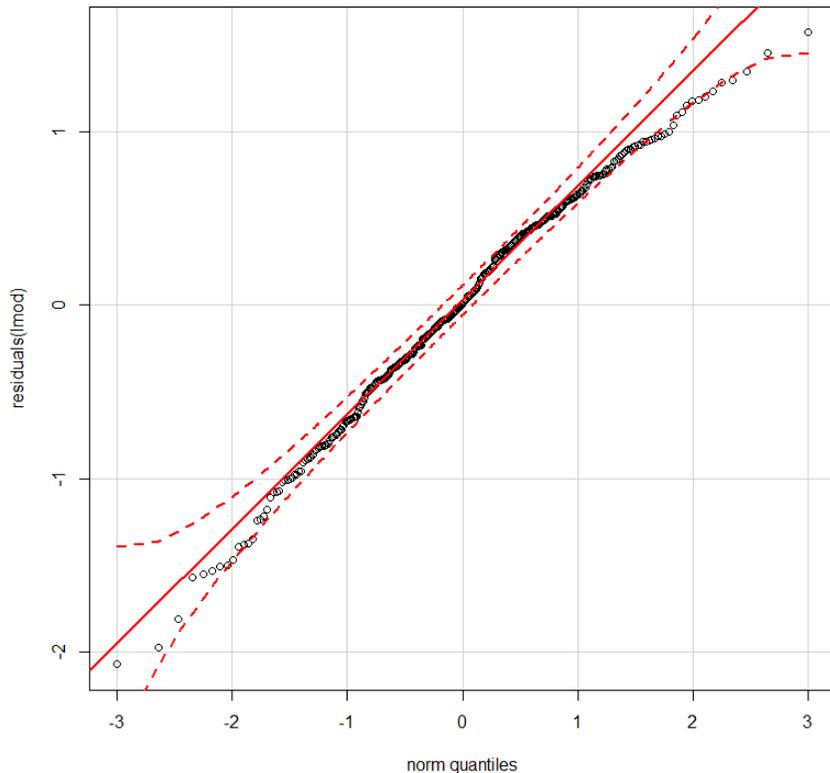
4. DIAGNOSTICS

PRELIMINARY ANALYSIS While not used formally in this analysis (due to lack of explanation), an examination of R's `gv1ma()` function (Appendix, Figure 4.1) reveals that the assumptions of kurtosis and heteroscedasticity are satisfied. However, the remaining assumptions are not met. We will attempt to explain these results in the following section.

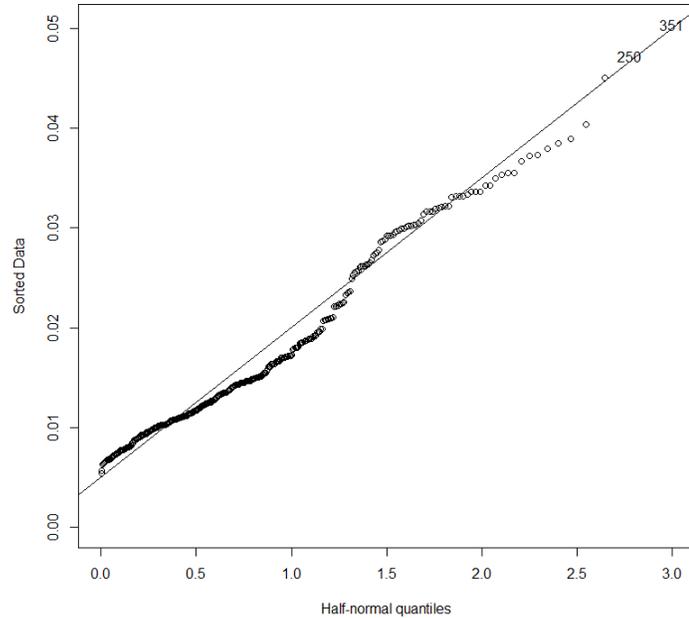
CONSTANT VARIANCE We note there is an apparent bound of residuals as the fitted values increase, possibly the result of a structural problem which we will address later in this section. While square-rooting the residuals does appear to diminish this effect, it does not eliminate it (R code in Appendix, Figure 4.2; relevant plot to the right).



NORMALITY Observe the Q-Q plot with confidence bands to the right (R code in Appendix, Figure 4.3; plot left). Note that the normality assumption is indeed slightly breached, both left and right. However, the majority of the breach occurs to the right, indicating right-skew.

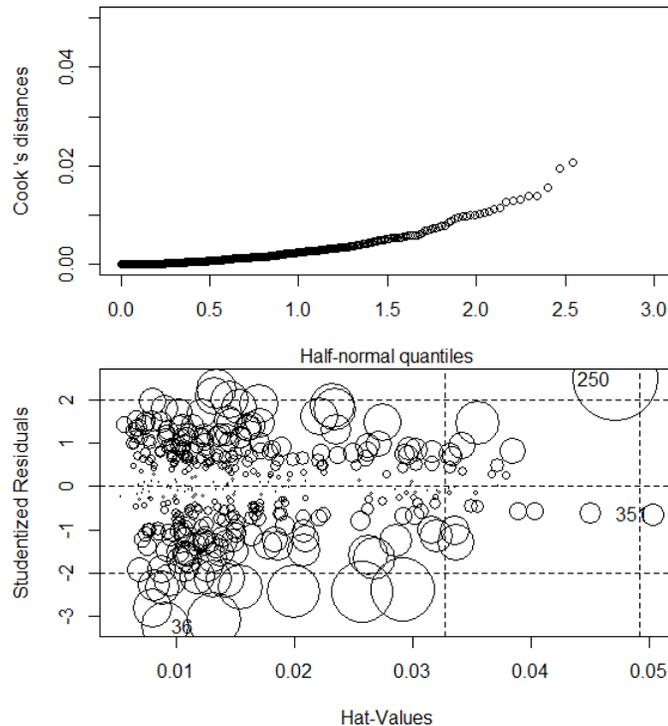


LEVERAGE POINTS Despite this model’s breach of the normality assumption, there do not appear to be any leverage points that significantly affect the model. Observe the plot to the right (R code in Appendix, Figure 4.4): there appears to be a (roughly) linear relation between the sorted data and the half-normal quantiles. The lack of a significant off-shooting data point indicates a lack of leverage points.

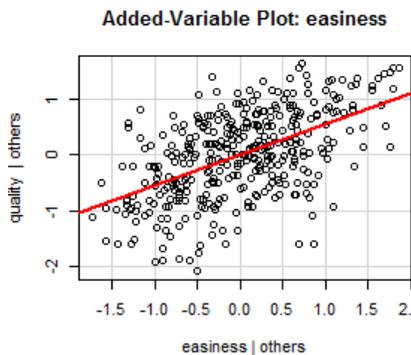
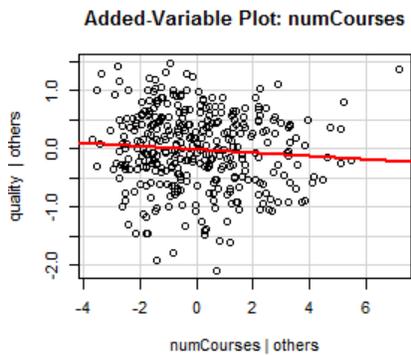
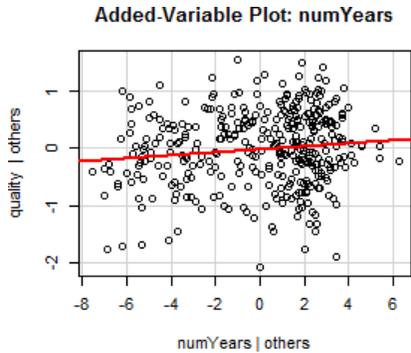


OUTLIERS By the Bonferonni p -value (R code in Appendix, Figure 4.5), we detect observation 36 as an outlier. In the next section we will examine the possibility for influential data.

INFLUENTIAL POINTS Observe on the plots to the right (R code in Appendix, Figure 4.6) that there are two potential influence points indicated on the plot of studentized residuals versus hat-values (observations 250 and 351), and none indicated on the plot of Cook’s distances. Since neither of these observations are indicated as outliers or leverage points, it is unlikely that they are candidates for removal. However, for completeness’ sake we will consider their removal. In Appendix Figure 4.7, we observe the `summary()` of the original linear model, followed by a sequential removal of these two points (complete with abbreviated summaries). We note that there is no significant change between summaries, thus we are not inclined to remove any observation as an influential point.

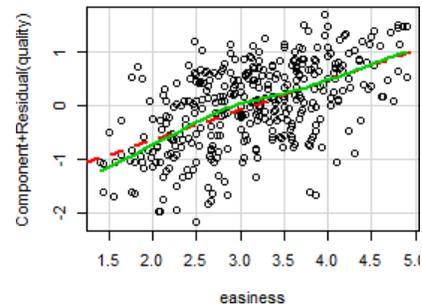
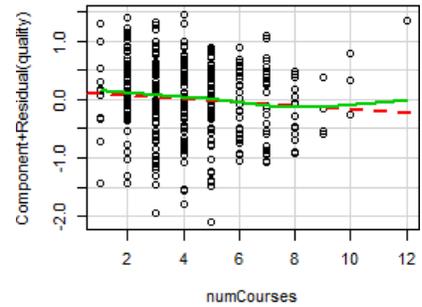
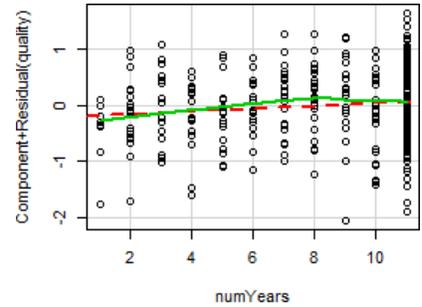


STRUCTURE As noted at the beginning of this section, there appears to be a distinct structural problem that cannot be readily transformed. We will further this examination by considering the added-variable and component-plus-residual plots (below) of the quantitative data in our linear model.



By the added-variable plots (left), the slope of the red lines somewhat indicate a linear relation between quality vs. numYears, and quality vs. numCourses. There appears to be a stronger linear relation between quality vs. easiness.

By the component-plus-residual plots (right), there is a slight indication that transformations may be called for on numYears and numCourses, however further investigation reveals that basic transformations do not improve our model significantly.



In Appendix Figure 4.8 (and the accompanying Appendix plot), we attempt four basic transformations: numYears^2 , numCourses^2 , $\log(\text{numYears})$, and $\log(\text{numCourses})$. While some of these transformations improve the linearity of the cpr plot (e.g. numYears^2 , there were no significant changes in summary output). Furthermore, none of these models improved our `gvlma()` output. It was thus concluded to not transform any of the quantitative variables in this model.

5. HYPOTHESIS In this section we have identified three questions we wish to answer:

1. Are Helpfulness, Clarity, Easiness, and Rater interest all statistically significant in the model?
2. Does Gender have an effect on perception of Overall Quality, given the model of best fit?
3. Does the Chili Pepper have an effect on perception of Overall Quality, given the model of best fit?

To answer the first question, we must fit our model, regress on `quality`, and then observe the p -values.

EASINESS is statistically significant at the $\alpha = 0.001$ level, with a p -value of $< 2 \cdot 10^{-16}$. Hence there is very strong evidence `easiness` is useful in explaining the average `quality` rating, when `gender`, `pepper`, `numYears` and `numCourses` are already in the model.

GENDER is statistically significant at the $\alpha = 0.05$ level, with a p -value of 0.0208. Hence there is very strong evidence `gender` is useful in explaining the average `quality` rating, when `easiness`, `pepper`, `numYears` and `numCourses` are already in the model.

PEPPER is statistically significant at the $\alpha = 0.001$ level, with a p -value of $8.11 \cdot 10^{-11}$. Hence there is very strong evidence `pepper` is useful in explaining the average `quality` rating, when `gender`, `easiness`, `numYears` and `numCourses` are already in the model.

NUMYEARS is statistically significant at the $\alpha = 0.05$ level, with a p -value of 0.0225. Hence there is moderate evidence `numYears` is useful in explaining the average `quality` rating, when `gender`, `pepper`, `easiness` and `numCourses` are already in the model.

NUMCOURSES is not statistically significant with a p -value of 0.1200. Hence there is no evidence `numCourses` is useful in explaining the average `quality` rating, when `gender`, `pepper`, `numYears` and `easiness` are already in the model.

MODEL AND INTERPRETATION

Here we show our model with the appropriate β values input.

$$\begin{aligned} E(\text{quality} | \text{numYears}, \text{numCourses}, \text{easiness}, \text{genderMale}, \text{pepperYes}) \\ = 1.56261 + 0.02622(\text{numYears}) - 0.02829(\text{numCourses}) + 0.55364(\text{easiness}) + 0.16240(\text{genderMale}) \\ + 0.68553(\text{pepperYes}) \end{aligned}$$

INTERPRETATION (`numYears`)

If two professors are identical to one another except that one professor has been teaching one more year than the other, then we predict the average quality rating for the professor with an extra year of teaching will be

0.02622 points higher than the other.

INTERPRETATION (`numCourses`)

If two professors are identical to one another except that one professor has one more course title included in their ratings, then we predict the average quality rating for the professor with more course titles will be 0.028929 less than the other professor.

INTERPRETATION (`easiness`)

If two professors are identical to one another except that one professor has an average easiness score that is one point higher than the other, then we predict the average quality rating will be about 0.55364 points higher than the other professor.

INTERPRETATION (`genderMale`)

If two professors are identical to one another except that one is male and the other is female, then we predict the average quality rating will be 0.16240 points higher for the female than the male.

INTERPRETATION (`pepperYes`)

If two professors are identical to one another except that one is rated as attractive and the other is not, then we predict the average quality rating for the attractive professor will be 0.68553 points higher than the other professor.

HYPOTHESIS TESTING We wanted to test whether the `gender` and `pepper` regressor variables should be simultaneously dropped from the model that already includes `easiness`, `numYears`, and `numCourses`.

$$H_0 : \beta_{gender} = 0 \text{ and } \beta_{pepper} = 0 \mid \beta_{easiness} \neq 0, \beta_{numYears} \neq 0, \beta_{numCourses} \neq 0$$

$$H_\alpha : \beta_{Gender} \neq 0 \text{ or } \beta_{Pepper} \neq 0 \mid \beta_{easiness} \neq 0, \beta_{numYears} \neq 0, \beta_{numCourses} \neq 0$$

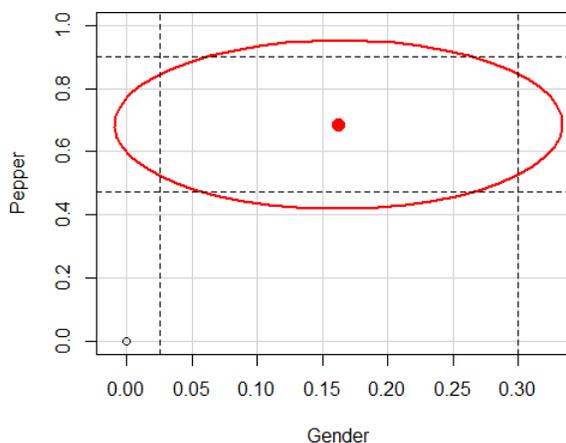
F-STATISTIC:

$$\begin{aligned} F &= [(RSS_\omega - RSS_\Omega)/(df_\omega - df_\Omega)] / (RSS_\Omega/df_\Omega) \\ &= [(173.16 - 153.91)/(362 - 360)] / (153.91/360) = 9.625/0.427527 \\ &= 22.5 \end{aligned}$$

P-VALUE: $6.075 \cdot 10^{-10}$

CONCLUSION IN CONTEXT There is very strong evidence that `gender` or `pepper` should be included in the model given that `easiness`, `numYears`, and `numCourses` are already in the model.

We can also see from the Ellipse Confidence Interval Graph Plot Below between **gender** and **pepper** that $\beta_{gender} = \beta_{pepper} = 0$ is not plausible because $(0, 0)$ falls into the rejection region.



We do not believe prediction is appropriate for this model. Not only are our assumptions violated (see conclusion below), but there are systemic issues with predicting students' perception: this data set is based on a single school. We would need a much larger, randomized pool of students to formulate useful predictions on the entire student population.

6. CONCLUSION

SUMMARIZING From our data exploration and analysis, we learned that the best linear model for our data set when regressing on Quality is the model that includes Gender, Pepper, Easiness, Number of Years, and Number of Courses. After running diagnostics on our data we found that our normality assumption failed (indicated by skewness). In addition, our linear model assumptions failed the global stat, and link function tests. We also found when testing our structure that there appears to be clear structure in the residuals versus fitted values plot, on the back end of the data, coming to a point. However, at this time we are unequipped to make adjustments to the data in order to correct this issue. Due to these issues it is important we note that our conclusions are unreliable. We are providing the tests and the interpretations to supply evidence of our ability to perform and interpret such tests, if our structure and normality assumptions were held. Finally, we must also note that these observations are based on students' *perceptions* of instructor quality, adding another layer of unreliability.

By finding appropriate test statistics and p -values associated with the variables in our model we are able to answer our questions of interest. The first question we asked was, "Does Gender have an effect on perception of Overall Quality given the model of best fit?" When comparing our full model of best fit which included **easiness**, **gender**, **pepper**, **numYears**, and **numCourses**, to the model with just **easiness**, **Pepper**, **numYears**, and **numCourses**; we found moderate evidence that **gender** was useful in explaining overall quality perception. The second question we looked into being, "Does the Chili Pepper have an effect on perception

of overall **quality** given the model of best fit?” We found when comparing our full model of best fit, which included **easiness**, **gender**, **pepper**, **numYears**, and **numCourses**, to the model with just **easiness**, **gender**, **numYears**, and **numCourses**, there was very strong evidence that the Chili Pepper was useful in explaining overall **quality** perception. The third question we proposed was “Are Helpfulness, Clarity, Easiness, and Rater Interest all statistically significant in the model of best fit?” As we researched more about how the data was collected from the paper “Student Consensus on *RateMyProfessors.com*” we learned that **helpfulness** and **clarity** are directly used to calculate **quality** perception [1]. Due to this we removed **helpfulness** and **clarity** from the model from the beginning. When we did diagnostics and checked for collinearity we found that **raterInterest** should be removed from the model. After both diagnostics and model selection we did test if **easiness** was statistically significant and found that it was significant with a p -value of less than $2 \cdot 10^{-16}$. So we can conclude that there is very strong evidence **easiness** is useful in explaining **quality** perception when **gender**, **pepper**, **numYears**, and **numCourses** are already in the model.

LEARNING As a team we learned a lot about model selection and diagnostics. The model selection and diagnostics sections were more difficult for us as we found a clear structure in our residuals versus fitted values that we did not know how to correct for. It was also difficult dealing with diagnostics as we learned our normality assumptions failed and therefore made our interpretations and hypotheses unreliable.

IMPROVING The major change we would make to improve the design of our study would be to find a way to transform the data that would fit our normality assumptions and a way to adjust for the evident structural issues. Due to the fact that these lead our conclusions to be unreliable this would be the most important improvement we could make so that conclusions could be drawn. To do this it may include using more of the variables from the original study of the data like the standard deviations of the rating variables.

FUTURE WORK To extend the research we would expand the data collection beyond just the large single campus in the Midwest. Due to the data being only from one specific campus, all of the interpretations and conclusions would be held within that campus and cannot be extended beyond that campus. By adding other campuses we could extend interpretations to other campuses. However, not only should we extend the sample size and parameters (more schools), but we would also recommend randomly selecting professors across those campuses to extend our interpretations to all the professors within the population.

Appendix

FIGURE 1.1

```

1 > summary(Rateprof)
2 Output:
3 gender          numYears          numRaters          numCourses          pepper          discipline
4 female:159      Min.      : 1.000      Min.      :10.00      Min.      : 1.000      no :320      Hum
   :134
5 male  :207      1st Qu.: 6.000      1st Qu.:15.00      1st Qu.: 3.000      yes: 46      SocSci :
   66
6           Median :10.000      Median :24.00      Median : 4.000
   :103
7           Mean   : 8.347      Mean   :28.58      Mean   : 4.251
   63
8           3rd Qu.:11.000      3rd Qu.:37.00      3rd Qu.: 5.000
9           Max.   :11.000      Max.   :86.00      Max.   :12.000
10
11          dept          quality          helpfulness          clarity          easiness
12 English   : 49      Min.      :1.409      Min.      :1.364      Min.      :1.333      Min.      :1.391
13 Math      : 34      1st Qu.:2.936      1st Qu.:3.069      1st Qu.:2.871      1st Qu.:2.548
14 Biology   : 20      Median :3.612      Median :3.662      Median :3.600      Median :3.148
15 Chemistry : 20      Mean   :3.575      Mean   :3.631      Mean   :3.525      Mean   :3.135
16 Psychology: 20      3rd Qu.:4.250      3rd Qu.:4.351      3rd Qu.:4.214      3rd Qu.:3.692
17 Spanish   : 20      Max.   :4.981      Max.   :5.000      Max.   :5.000      Max.   :4.900
18 (Other)   :203
19 raterInterest          sdQuality          sdHelpfulness          sdClarity          sdEasiness
20 Min.      :1.098      Min.      :0.09623      Min.      :0.0000      Min.      :0.0000      Min.      :0
   .3162
21 1st Qu.:2.934      1st Qu.:0.87508      1st Qu.:0.9902      1st Qu.:0.9085      1st Qu.:0
   .9045
22 Median :3.305      Median :1.15037      Median :1.2860      Median :1.1712      Median :1
   .0247
23 Mean   :3.310      Mean   :1.05610      Mean   :1.1719      Mean   :1.0970      Mean   :1
   .0196
24 3rd Qu.:3.692      3rd Qu.:1.28730      3rd Qu.:1.4365      3rd Qu.:1.3328      3rd Qu.:1
   .1485
25 Max.   :4.909      Max.   :1.67739      Max.   :1.8091      Max.   :1.8091      Max.   :1
   .6293
26
27 sdRaterInterest
28 Min.      :0.3015
29 1st Qu.:1.0848
30 Median :1.2167
31 Mean   :1.1965
32 3rd Qu.:1.3326
33 Max.   :1.7246

```

FIGURE 1.2

```

1 > plot(density(Rateprof$numYears, na.rm = TRUE), main = "Number of Years Density
   Plot")
2 > plot(density(Rateprof$numCourses, na.rm = TRUE), main = "Number of Courses
   Density Plot")
3 > plot(density(Rateprof$numRaters, na.rm = TRUE), main = "Number of Ratings
   Density Plot")
4 > plot(density(Rateprof$easiness, na.rm = TRUE), main = "Average Easiness Rating
   Density Plot")
5 > plot(density(Rateprof$helpfulness, na.rm = TRUE), main = "Average Helpfulness
   Rating Density Plot")
6 > plot(density(Rateprof$clarity, na.rm = TRUE), main = "Average Clarity Rating
   Density Plot")

```

```
7 > plot(density(Rateprof$raterInterest, na.rm = TRUE), main = "Average Rater
Interest Density Plot")
```

FIGURE 1.3

```
1 > gender1 <- Rateprof$gender
2 > gender.freq = table(gender1)
3 > barplot(gender.freq, main = "Gender Bar Graph")
4
5 > peppery <- Rateprof$pepper
6 > pepper.freq = table(peppery)
7 > barplot(pepper.freq, main = "Pepper Bar Graph")
```

FIGURE 2.1

```
1 > summary(lmod)
2
3 Call:
4 lm(formula = quality ~ gender + numYears + numRaters + numCourses +
5     pepper + helpfulness + clarity + easiness + raterInterest,
6     data = Rateprof)
7
8 Residuals:
9     Min       1Q   Median       3Q      Max
10 -0.74521 -0.00749  0.00243  0.01238  0.16384
11
12 Coefficients:
13             Estimate Std. Error t value Pr(>|t|)
14 (Intercept)  -0.0143292  0.0183639  -0.780  0.4357
15 gendermale    0.0026288  0.0048498   0.542  0.5881
16 numYears      0.0009020  0.0008273   1.090  0.2763
17 numRaters     -0.0003962  0.0001526  -2.597  0.0098 **
18 numCourses    -0.0013853  0.0012826  -1.080  0.2808
19 pepperyes     0.0013419  0.0078707   0.170  0.8647
20 helpfulness   0.5360712  0.0074231  72.216 <2e-16 ***
21 clarity       0.4637251  0.0071509  64.849 <2e-16 ***
22 easiness      0.0063883  0.0037449   1.706  0.0889 .
23 raterInterest -0.0010421  0.0049108  -0.212  0.8321
24 ---
25 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
26
27 Residual standard error: 0.04493 on 356 degrees of freedom
28 Multiple R-squared:  0.9972, Adjusted R-squared:  0.9971
29 F-statistic: 1.405e+04 on 9 and 356 DF, p-value: < 2.2e-16
```

FIGURE 2.2

```

1 > round(cor(Rateprof[,c ("numYears", "numRaters", "numCourses", "helpfulness", "
   clarity", "easiness", "raterInterest"))),4)
2
3 numYears numRaters numCourses helpfulness clarity easiness raterI...
4 numYears 1.0000 0.3420 0.2570 -0.0429 -0.0245 -0.1405 -0.0126
5 numRaters 0.3420 1.0000 0.2599 -0.0770 -0.0753 -0.1065 -0.1363
6 numCourses 0.2570 0.2599 1.0000 -0.0673 -0.0948 -0.0901 0.0233
7 helpfulness -0.0429 -0.0770 -0.0673 1.0000 0.9208 0.5635 0.4630
8 clarity -0.0245 -0.0753 -0.0948 0.9208 1.0000 0.5359 0.4611
9 easiness -0.1405 -0.1065 -0.0901 0.5635 0.5359 1.0000 0.2052
10 raterInterest -0.0126 -0.1363 0.0233 0.4630 0.4611 0.2052 1.0000
11 > plot ( Rateprof[,c ("numYears", "numRaters", "numCourses", "helpfulness", "
   clarity", "easiness", "raterInterest") ] )

```

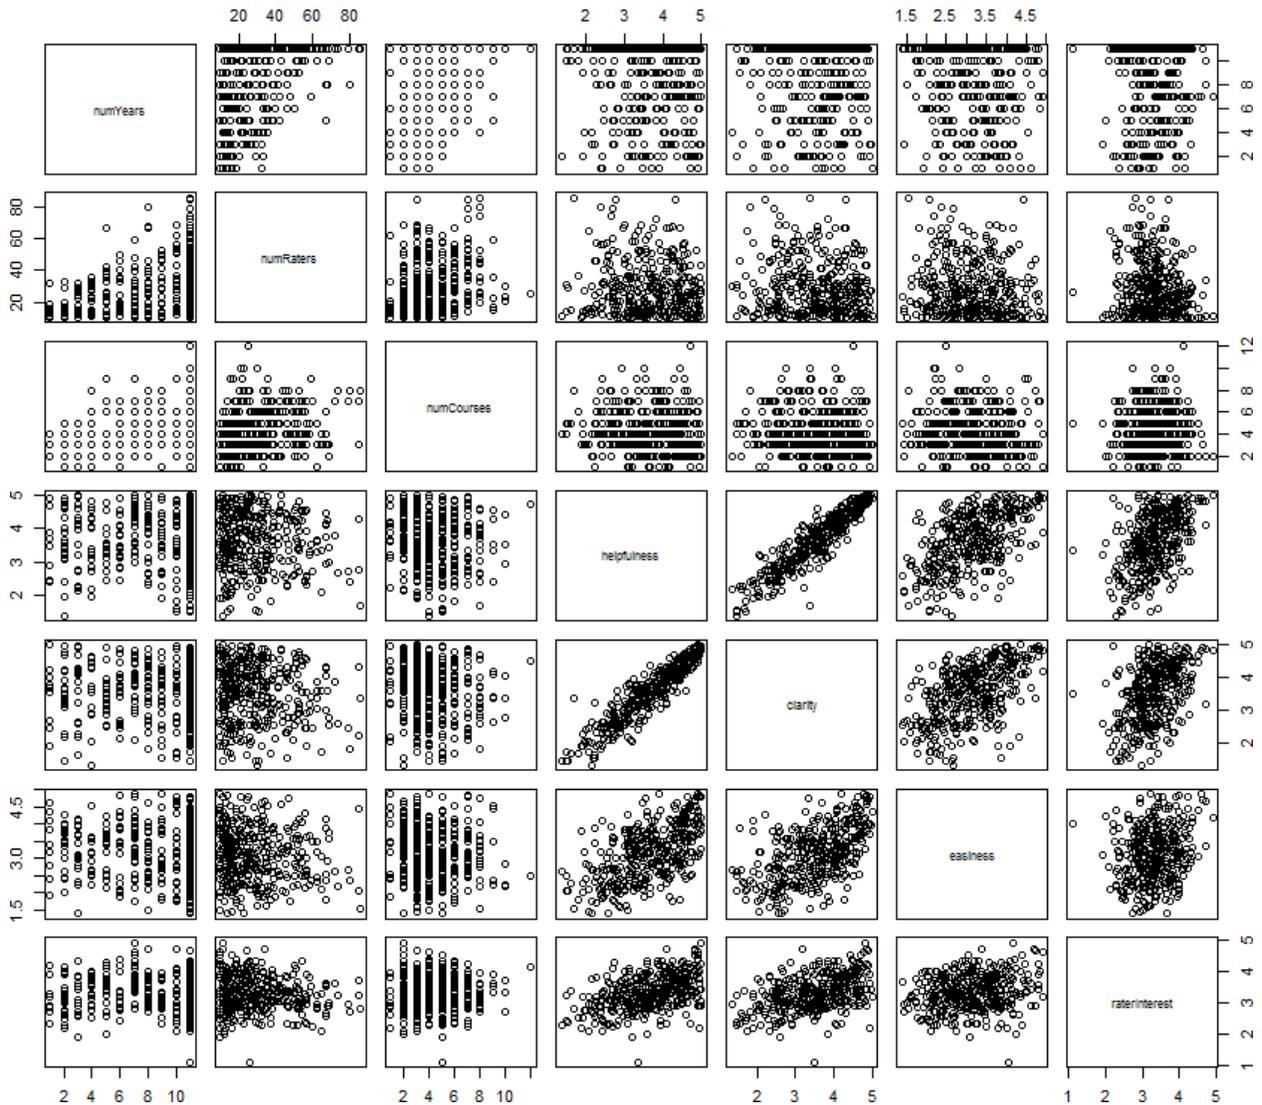


FIGURE 2.3

```

1 > round(vif(lmod),4)
2 gender numYears numRaters numCourses pepper helpfulness clarity easiness raterI...
3 1.0478 1.2678 1.203 1.1417 1.2342 7.0681 6.8206 1.5337 1.3312

```

FIGURE 2.4

```

1 > lmod <- lm(quality ~ numYears + numRaters + numCourses + helpfulness + clarity +
  easiness + raterInterest, Rateprof)
2 > X <- model.matrix(lmod)[-1]
3 > sqrt(eigen(crossprod(X,X))$val[1]/eigen(crossprod(X,X))$val)
4 [1] 1.000000 6.419884 15.145677 19.321868 57.575535 63.491312 146.986953
5 > colldiag(lmod, scale=FALSE, add.intercept=FALSE)
6 Condition
7 Index Variance Decomposition Proportions
8
9      numYears numRaters numCourses helpfulness clarity easiness raterInt...
10 1      1.000 0.000      0.184      0.000      0.000      0.000      0.000
11 2      6.420 0.153      0.789      0.011      0.000      0.000      0.001
12 3     15.146 0.783      0.002      0.039      0.003      0.003      0.010
13 4     19.322 0.022      0.023      0.880      0.002      0.002      0.005
14 5     57.576 0.013      0.001      0.000      0.016      0.026      0.961
15 6     63.491 0.026      0.001      0.061      0.037      0.051      0.009
16 7    146.987 0.002      0.000      0.008      0.943      0.918      0.014
17 > lmod <- lm(quality ~ numYears + numRaters + numCourses + easiness +
  raterInterest, Rateprof)
18 > X <- model.matrix(lmod)[-1]
19 > sqrt(eigen(crossprod(X,X))$val[1]/eigen(crossprod(X,X))$val)
20 [1] 1.000000 6.972014 17.109728 23.419540 59.091063
21 > colldiag(lmod, scale=FALSE, add.intercept=FALSE)
22 Condition
23 Index Variance Decomposition Proportions
24
25      numYears numRaters numCourses easiness raterInterest
26 1      1.000 0.000      0.191      0.000      0.000
27 2      6.972 0.247      0.792      0.017      0.002
28 3     17.110 0.603      0.014      0.503      0.014
29 4     23.420 0.110      0.002      0.455      0.139
30 5     59.091 0.040      0.002      0.025      0.845
31 > lmod_0 <- lm(quality ~ numYears + numRaters + numCourses + raterInterest,
  Rateprof)
32 > X <- model.matrix(lmod_0)[-1]
33 > sqrt(eigen(crossprod(X,X))$val[1]/eigen(crossprod(X,X))$val)
34 [1] 1.000000 7.181937 17.535442 31.334791
35 > colldiag(lmod_0, scale=FALSE, add.intercept=FALSE)
36 Condition
37 Index Variance Decomposition Proportions
38
39      numYears numRaters numCourses raterInterest
40 1      1.000 0.000      0.194      0.000
41 2      7.182 0.290      0.785      0.019
42 3     17.535 0.467      0.021      0.694
43 4     31.335 0.242      0.000      0.286
44 > lmod_1 <- lm(quality ~ numYears + numRaters + numCourses + easiness, Rateprof)
45 > X <- model.matrix(lmod_1)[-1]
46 > sqrt(eigen(crossprod(X,X))$val[1]/eigen(crossprod(X,X))$val)
47 [1] 1.000000 7.22862 17.52871 27.38139
48 > colldiag(lmod_1, scale=FALSE, add.intercept=FALSE)
49 Condition
50 Index Variance Decomposition Proportions
51
52      numYears numRaters numCourses easiness
53 1      1.000 0.000      0.193      0.000
54 2      7.229 0.336      0.783      0.021
55 3     17.529 0.517      0.022      0.735
56 4     27.381 0.147      0.001      0.243

```

FIGURE 3.1

```
1 > b <- regsubsets(quality ~ gender + pepper + easiness + numYears + numRaters +
  numCourses,data = Rateprof)
2 > rs <- summary(b) #summarize model that minimizes RSS for each p
3 > rs$which
4 (Intercept) gendermale pepperyes easiness numYears numRaters numCourses
5 1          TRUE      FALSE      FALSE      TRUE      FALSE      FALSE      FALSE
6 2          TRUE      FALSE      TRUE      TRUE      FALSE      FALSE      FALSE
7 3          TRUE      TRUE       TRUE      TRUE      FALSE      FALSE      FALSE
8 4          TRUE      TRUE       TRUE      TRUE      TRUE       FALSE      FALSE
9 5          TRUE      TRUE       TRUE      TRUE      TRUE       FALSE      TRUE
10 6         TRUE      TRUE       TRUE      TRUE      TRUE       TRUE       TRUE
```

FIGURE 3.2

```
1 > AIC <- 366*log(rs$rss/366) + (2:7)*2
2 > plot(AIC ~ I(2:7), ylab="AIC", xlab="Number of Predictors", main = "AIC ratings
  ")
```

FIGURE 3.3

```
1 > adjr = rs$adjr
2 > plot(adjr ~ p, xlab= "Parameters", ylab = expression({R^2}[a]), main= "Adjusted
  R Squared Ratings")
```

FIGURE 3.4

```
1 > cp = rs$cp
2 > plot(cp ~ p,xlab = "Parameters" ,ylab = expression(paste(C[p], " statistic")),
  main = "Mallow's C(p) Statistic")
```

FIGURE 4.1

```

1 > library(gvlma)
2 > lmod <- lm(quality ~ numYears + numCourses + easiness + gender + pepper,
3   Rateprof)
4 > gvlma(lmod)
5 Call:
6 lm(formula = quality ~ numYears + numCourses + easiness + gender +
7   pepper, data = Rateprof)
8
9 Coefficients:
10 (Intercept)      numYears      numCourses      easiness      gendermale      pepperyes
11      1.56261         0.02622        -0.02829         0.55364         0.16240         0.68553
12
13
14 ASSESSMENT OF THE LINEAR MODEL ASSUMPTIONS
15 USING THE GLOBAL TEST ON 4 DEGREES-OF-FREEDOM:
16 Level of Significance = 0.05
17
18 Call:
19 gvlma(x = lmod)
20
21
22 Value p-value Decision
23 Global Stat      13.18612 0.010401 Assumptions NOT satisfied!
24 Skewness         7.40233 0.006514 Assumptions NOT satisfied!
25 Kurtosis         0.09833 0.753837 Assumptions acceptable.
26 Link Function    5.36589 0.020534 Assumptions NOT satisfied!
27 Heteroscedasticity 0.31957 0.571869 Assumptions acceptable.

```

FIGURE 4.2

```

1 > plot(fitted(lmod), residuals(lmod), xlab = "fitted values", ylab = "residuals")
2 > plot(fitted(lmod), sqrt(abs(residuals(lmod))), xlab = "fitted values", ylab = "
3   sqrt(|residuals|)")

```

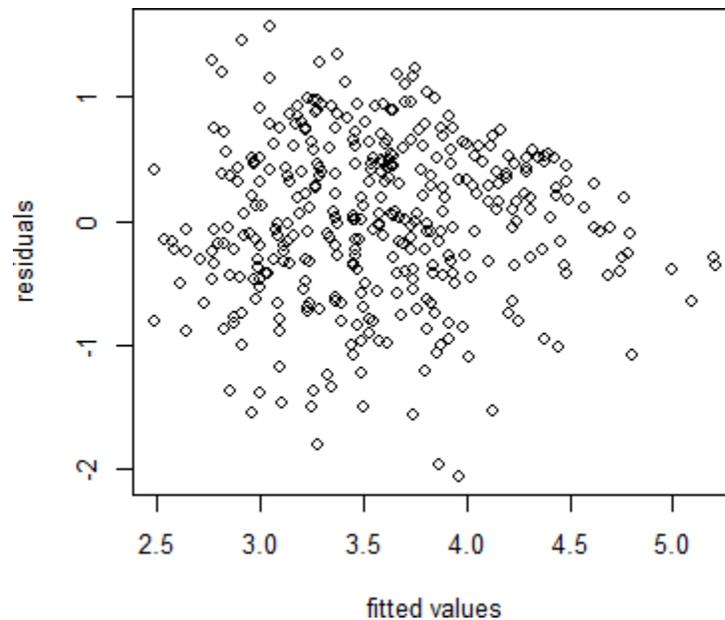


FIGURE 4.3

```
1 > library(car)
2 > qqPlot(residuals(lmod))
```

FIGURE 4.4

```
1 > halfnorm (hatvalues(lmod), nlab = 2)
```

FIGURE 4.5

```
1 > outlierTest(lmod)
2
3 No Studentized residuals with Bonferonni p < 0.05
4 Largest |rstudent|:
5   rstudent unadjusted p-value Bonferonni p
6 36 -3.222465          0.0013872          0.50772
```

FIGURE 4.6

```
1 > influencePlot(lmod)
2           StudRes           Hat           CookD
3 36  -3.2224651  0.009108198  0.015504419
4 250  2.4822783  0.047188045  0.050140857
5 351 -0.6386189  0.050363798  0.003610842
6 > halfnorm (cooks.distance(lmod), n=3, labs=row.names(teengamb), ylab = "Cook 's
   distances")
```

FIGURE 4.7 (Truncated output)

```

1 > lmod2 <- lm(quality ~ numYears + numCourses + easiness + gender + pepper,
  Rateprof, subset=(sdQuality != 0.72571800))
2 > lmod3 <- lm(quality ~ numYears + numCourses + easiness + gender + pepper,
  Rateprof, subset=(sdQuality != 0.80750000))
3 > summary(lmod)
4 # ...
5 Coefficients:
6           Estimate Std. Error t value Pr(>|t|)
7 (Intercept)  1.56261    0.19289   8.101 8.55e-15 ***
8 numYears     0.02622    0.01144   2.291  0.0225 *
9 numCourses  -0.02829    0.01815  -1.559  0.1200
10 easiness    0.55364    0.04580  12.088 < 2e-16 ***
11 gendermale  0.16240    0.06994   2.322  0.0208 *
12 pepperyes   0.68553    0.10861   6.312 8.11e-10 ***
13 ---
14 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
15
16 Residual standard error: 0.6538 on 360 degrees of freedom
17 Multiple R-squared:  0.3988, Adjusted R-squared:  0.3905
18 F-statistic: 47.77 on 5 and 360 DF, p-value: < 2.2e-16
19
20 > summary(lmod2)
21 # ...
22 Coefficients:
23           Estimate Std. Error t value Pr(>|t|)
24 (Intercept)  1.58675    0.19177   8.274 2.56e-15 ***
25 numYears     0.02673    0.01136   2.352  0.0192 *
26 numCourses  -0.03742    0.01839  -2.034  0.0427 *
27 easiness    0.55623    0.04549  12.228 < 2e-16 ***
28 gendermale  0.15785    0.06947   2.272  0.0237 *
29 pepperyes   0.68822    0.10784   6.382 5.41e-10 ***
30 ---
31 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
32
33 Residual standard error: 0.6492 on 359 degrees of freedom
34 Multiple R-squared:  0.4064, Adjusted R-squared:  0.3982
35 F-statistic: 49.16 on 5 and 359 DF, p-value: < 2.2e-16
36
37 > summary(lmod3)
38 # ...
39 Coefficients:
40           Estimate Std. Error t value Pr(>|t|)
41 (Intercept)  1.54834    0.19434   7.967 2.17e-14 ***
42 numYears     0.02663    0.01147   2.322  0.0208 *
43 numCourses  -0.02682    0.01831  -1.464  0.1440
44 easiness    0.55565    0.04595  12.093 < 2e-16 ***
45 gendermale  0.15934    0.07016   2.271  0.0237 *
46 pepperyes   0.69462    0.10963   6.336 7.06e-10 ***
47 ---
48 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
49
50 Residual standard error: 0.6544 on 359 degrees of freedom
51 Multiple R-squared:  0.3981, Adjusted R-squared:  0.3897
52 F-statistic: 47.49 on 5 and 359 DF, p-value: < 2.2e-16

```

FIGURE 4.8 (Truncated output)

```

1 > lmod1 <- lm(quality ~ numYears + I(numYears^2) + numCourses + easiness + gender
  + pepper, Rateprof)
2 > lmod2 <- lm(quality ~ numYears + numCourses + I(numCourses^2) + easiness +
  gender + pepper, Rateprof)
3 > lmod3 <- lm(quality ~ numYears + I(log(numYears)) + numCourses + easiness +
  gender + pepper, Rateprof)
4 > lmod4 <- lm(quality ~ numYears + numCourses + I(log(numCourses)) + easiness +
  gender + pepper, Rateprof)
5 > gvlma(lmod1)
6
7 Call:
8 lm(formula = quality ~ numYears + I(numYears^2) + numCourses +
9   easiness + gender + pepper, data = Rateprof)
10
11 Coefficients:
12 (Intercept) numYears I(numYears^2) numCourses easiness gendermale pepperyes
13   1.321938  0.127810      -0.007281  -0.030381  0.547801   0.161652  0.700914
14
15
16 ASSESSMENT OF THE LINEAR MODEL ASSUMPTIONS
17 USING THE GLOBAL TEST ON 4 DEGREES-OF-FREEDOM:
18 Level of Significance = 0.05
19
20 Call:
21 gvlma(x = lmod1)
22
23           Value  p-value           Decision
24 Global Stat    13.99849 0.007300 Assumptions NOT satisfied!
25 Skewness       7.13596 0.007555 Assumptions NOT satisfied!
26 Kurtosis       0.05934 0.807536 Assumptions acceptable.
27 Link Function   6.52272 0.010651 Assumptions NOT satisfied!
28 Heteroscedasticity 0.28047 0.596391 Assumptions acceptable.
29 > gvlma(lmod2)
30
31 Call:
32 lm(formula = quality ~ numYears + numCourses + I(numCourses^2) +
33   easiness + gender + pepper, data = Rateprof)
34
35 Coefficients:
36 (Intercept) numYears numCourses I(numCourses^2) easiness gendermale pepperyes
37   1.80680  0.02714   -0.14026      0.01127  0.54730   0.15627  0.69139
38
39
40 ASSESSMENT OF THE LINEAR MODEL ASSUMPTIONS
41 USING THE GLOBAL TEST ON 4 DEGREES-OF-FREEDOM:
42 Level of Significance = 0.05
43
44 Call:
45 gvlma(x = lmod2)
46
47           Value  p-value           Decision
48 Global Stat    15.9397 0.003101 Assumptions NOT satisfied!
49 Skewness       7.9075 0.004923 Assumptions NOT satisfied!
50 Kurtosis       0.4124 0.520744 Assumptions acceptable.
51 Link Function   7.1894 0.007334 Assumptions NOT satisfied!
52 Heteroscedasticity 0.4303 0.511819 Assumptions acceptable.
53 > gvlma(lmod3)

```

```

54
55 Call:
56 lm(formula = quality ~ numYears + I(log(numYears)) + numCourses +
57     easiness + gender + pepper, data = Rateprof)
58
59 Coefficients:
60 (Intercept) numYears I(log(numYears)) numCourses easiness gendermale pepperyes
61     1.38695 -0.03620          0.35732   -0.03004  0.54990    0.16348    0.70120
62
63
64 ASSESSMENT OF THE LINEAR MODEL ASSUMPTIONS
65 USING THE GLOBAL TEST ON 4 DEGREES-OF-FREEDOM:
66 Level of Significance = 0.05
67
68 Call:
69 gvlma(x = lmod3)
70
71                Value p-value                Decision
72 Global Stat      13.7759 0.008046 Assumptions NOT satisfied!
73 Skewness         7.2223 0.007200 Assumptions NOT satisfied!
74 Kurtosis         0.1099 0.740234 Assumptions acceptable.
75 Link Function    6.1064 0.013469 Assumptions NOT satisfied!
76 Heteroscedasticity 0.3372 0.561429 Assumptions acceptable.
77 > gvlma(lmod4)
78
79 Call:
80 lm(formula = quality ~ numYears + numCourses + I(log(numCourses)) +
81     easiness + gender + pepper, data = Rateprof)
82
83 Coefficients:
84 (Intercept) numYears numCourses I(log(numCourses)) easiness gendermale pepperyes
85 1.63770      0.02668      0.02053          -0.20416  0.54937    0.16054    0.68901
86
87
88 ASSESSMENT OF THE LINEAR MODEL ASSUMPTIONS
89 USING THE GLOBAL TEST ON 4 DEGREES-OF-FREEDOM:
90 Level of Significance = 0.05
91
92 Call:
93 gvlma(x = lmod4)
94
95                Value p-value                Decision
96 Global Stat      14.6813 0.005410 Assumptions NOT satisfied!
97 Skewness         7.5824 0.005894 Assumptions NOT satisfied!
98 Kurtosis         0.2389 0.624975 Assumptions acceptable.
99 Link Function    6.4615 0.011024 Assumptions NOT satisfied!
100 Heteroscedasticity 0.3986 0.527831 Assumptions acceptable.
101 > summary(lmod1)
102
103 Call:
104 lm(formula = quality ~ numYears + I(numYears^2) + numCourses +
105     easiness + gender + pepper, data = Rateprof)
106
107 Residuals:
108     Min       1Q   Median       3Q      Max
109 -2.12070 -0.40753  0.02513  0.46038  1.61747
110
111 Coefficients:
112                Estimate Std. Error t value Pr(>|t|)

```

```

113 (Intercept)      1.321938    0.238681    5.539 5.91e-08 ***
114 numYears        0.127810    0.060717    2.105  0.0360 *
115 I(numYears^2)  -0.007281    0.004274   -1.704  0.0893 .
116 numCourses     -0.030381    0.018145   -1.674  0.0949 .
117 easiness        0.547801    0.045810   11.958 < 2e-16 ***
118 gendermale      0.161652    0.069757    2.317  0.0210 *
119 pepperyes       0.700914    0.108699    6.448 3.66e-10 ***
120 ---
121 Signif. codes:  0    ***    0.001    **    0.01    *    0.05    .    0.1    1
122
123 Residual standard error: 0.6521 on 359 degrees of freedom
124 Multiple R-squared:  0.4036, Adjusted R-squared:  0.3937
125 F-statistic:  40.5 on 6 and 359 DF,  p-value: < 2.2e-16
126
127 > summary(lmod2)
128
129 Call:
130 lm(formula = quality ~ numYears + numCourses + I(numCourses^2) +
131     easiness + gender + pepper, data = Rateprof)
132
133 Residuals:
134     Min       1Q   Median       3Q      Max
135 -2.01419 -0.42344  0.00212  0.47517  1.48825
136
137 Coefficients:
138             Estimate Std. Error t value Pr(>|t|)
139 (Intercept)   1.806796   0.241812   7.472 6.08e-13 ***
140 numYears      0.027144   0.011429   2.375  0.0181 *
141 numCourses   -0.140265   0.069559  -2.016  0.0445 *
142 I(numCourses^2) 0.011266   0.006757   1.667  0.0963 .
143 easiness      0.547296   0.045847  11.937 < 2e-16 ***
144 gendermale    0.156271   0.069864   2.237  0.0259 *
145 pepperyes     0.691394   0.108398   6.378 5.52e-10 ***
146 ---
147 Signif. codes:  0    ***    0.001    **    0.01    *    0.05    .    0.1    1
148
149 Residual standard error: 0.6522 on 359 degrees of freedom
150 Multiple R-squared:  0.4034, Adjusted R-squared:  0.3935
151 F-statistic:  40.46 on 6 and 359 DF,  p-value: < 2.2e-16
152
153 > summary(lmod3)
154
155 Call:
156 lm(formula = quality ~ numYears + I(log(numYears)) + numCourses +
157     easiness + gender + pepper, data = Rateprof)
158
159 Residuals:
160     Min       1Q   Median       3Q      Max
161 -2.09605 -0.40564  0.05068  0.46388  1.60778
162
163 Coefficients:
164             Estimate Std. Error t value Pr(>|t|)
165 (Intercept)   1.38695   0.21914   6.329 7.36e-10 ***
166 numYears     -0.03620   0.03898  -0.929  0.3537
167 I(log(numYears)) 0.35732   0.21337   1.675  0.0949 .
168 numCourses   -0.03004   0.01814  -1.657  0.0985 .
169 easiness      0.54990   0.04574  12.022 < 2e-16 ***
170 gendermale    0.16348   0.06977   2.343  0.0197 *
171 pepperyes     0.70120   0.10874   6.448 3.65e-10 ***

```

```

172 ---
173 Signif. codes:  0    ***    0.001    **    0.01    *    0.05    .    0.1    1
174
175 Residual standard error: 0.6522 on 359 degrees of freedom
176 Multiple R-squared:  0.4035, Adjusted R-squared:  0.3935
177 F-statistic: 40.47 on 6 and 359 DF,  p-value: < 2.2e-16
178
179 > summary(lmod4)
180
181 Call:
182 lm(formula = quality ~ numYears + numCourses + I(log(numCourses)) +
183     easiness + gender + pepper, data = Rateprof)
184
185 Residuals:
186      Min       1Q   Median       3Q      Max
187 -2.04694 -0.42327 -0.00303  0.47966  1.47451
188
189 Coefficients:
190             Estimate Std. Error t value Pr(>|t|)
191 (Intercept)    1.63770    0.21140   7.747 9.74e-14 ***
192 numYears        0.02668    0.01146   2.328  0.0205 *
193 numCourses      0.02053    0.05900   0.348  0.7281
194 I(log(numCourses)) -0.20416    0.23480  -0.869  0.3852
195 easiness        0.54937    0.04608  11.922 < 2e-16 ***
196 gendermale      0.16054    0.07000   2.294  0.0224 *
197 pepperyes       0.68901    0.10872   6.338 7.00e-10 ***
198 ---
199 Signif. codes:  0    ***    0.001    **    0.01    *    0.05    .    0.1    1
200
201 Residual standard error: 0.6541 on 359 degrees of freedom
202 Multiple R-squared:  0.4001, Adjusted R-squared:  0.3901
203 F-statistic: 39.9 on 6 and 359 DF,  p-value: < 2.2e-16
204
205 > crPlot(lmod1, variable = "numYears")
206 > crPlot(lmod1, variable = "I(numYears^2)")
207 > crPlot(lmod1, variable = "numCourses")
208 > crPlot(lmod1, variable = "easiness")
209 > crPlot(lmod2, variable = "numYears")
210 > crPlot(lmod2, variable = "numCourses")
211 > crPlot(lmod2, variable = "I(numCourses^2)")
212 > crPlot(lmod2, variable = "easiness")
213 > crPlot(lmod3, variable = "numYears")
214 > crPlot(lmod3, variable = "I(log(numYears))")
215 > crPlot(lmod3, variable = "numCourses")
216 > crPlot(lmod3, variable = "easiness")
217 > crPlot(lmod4, variable = "numYears")
218 > crPlot(lmod4, variable = "numCourses")
219 > crPlot(lmod4, variable = "I(log(numCourses))")
220 > crPlot(lmod4, variable = "easiness")

```

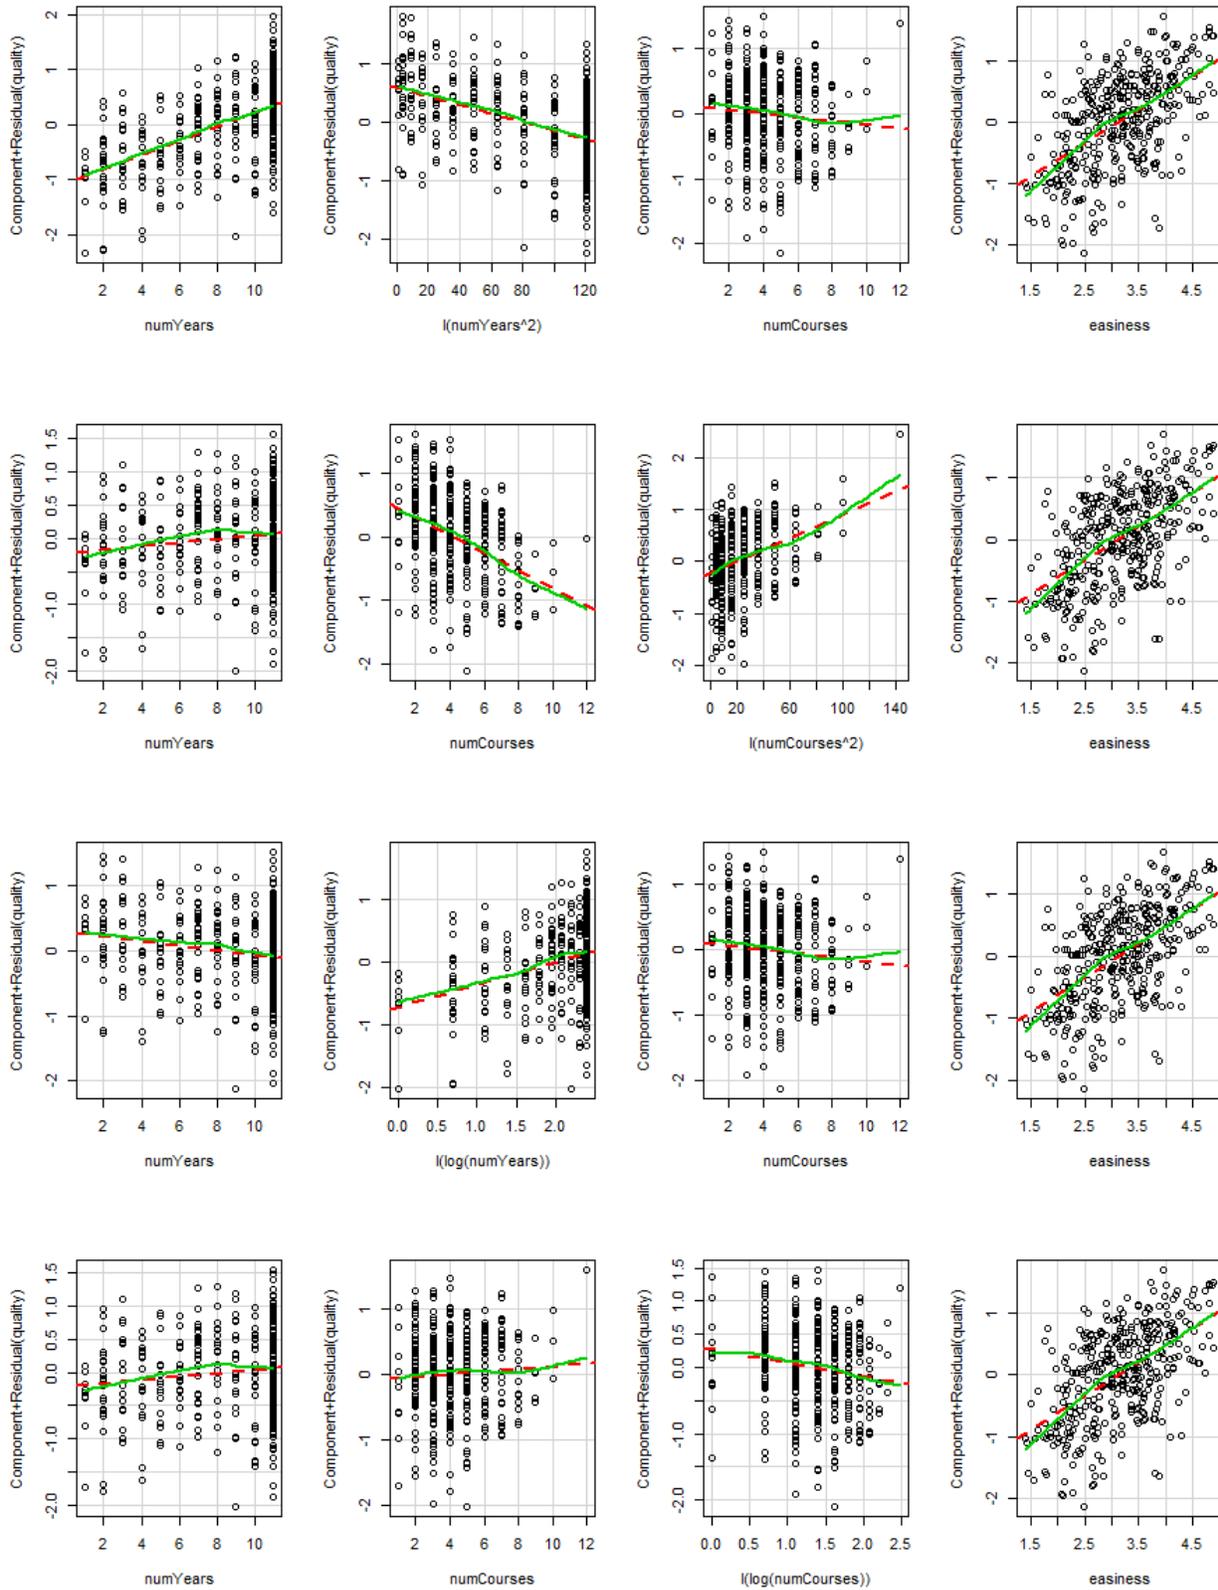


FIGURE 5.1

```

1 > data(Rateprof, package = "alr4")
2 > lmod <- lm(quality ~ easiness + gender + pepper + numYears + numCourses, data =
  Rateprof)
3 > summary(lmod)
4
5 Output:
6 Call:
7 lm(formula = quality ~ easiness + gender + pepper + numYears +
8     numCourses, data = Rateprof)
9
10 Residuals:
11      Min       1Q   Median       3Q      Max
12 -2.07057 -0.41488  0.01105  0.47686  1.57304
13
14 Coefficients:
15             Estimate Std. Error t value Pr(>|t|)
16 (Intercept)  1.56261    0.19289   8.101 8.55e-15 ***
17 easiness     0.55364    0.04580  12.088 < 2e-16 ***
18 gendermale   0.16240    0.06994   2.322  0.0208 *
19 pepperyes   0.68553    0.10861   6.312 8.11e-10 ***
20 numYears     0.02622    0.01144   2.291  0.0225 *
21 numCourses  -0.02829    0.01815  -1.559  0.1200
22 ---
23 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
24
25 Residual standard error: 0.6538 on 360 degrees of freedom
26 Multiple R-squared:  0.3988, Adjusted R-squared:  0.3905
27 F-statistic: 47.77 on 5 and 360 DF, p-value: < 2.2e-16

```

FIGURE 5.2

```

1 lmod2 <- lm(quality ~ easiness + numYears + numCourses, data = Rateprof)
2 anova(lmod2, lmod)
3
4 #95% CI Ellipse
5 library(car)
6 confidenceEllipse(lmod, which.coef = c(3, 4), levels = 0.95, ylim = c(0,1), ylab =
  "Pepper", xlab = "Gender")
7 # add origin to plot
8 points(0, 0)
9 abline(v = confint(lmod)[3,], lty = 2)
10 abline(h = confint(lmod)[4,], lty = 2)

```

Works Cited

- [1] Bleske-Rechek, A. and Fritsch, A. (2011). Student Consensus on RateMyProfessors.com. Practical Assessment, Research & Evaluation, 16(18), <http://pareonline.net/getvn.asp?v=16&n=18>

WORK BREAKDOWN

FAITH BLAIR Graphical summaries and interpretations; hypothesis testing; model interpretations; introduction; conclusion.

LARRY BREEDEN Graphical and numerical summaries.

LUKE MATTINGLY Model selection; preliminary diagnostic analysis; introduction; document editing and \LaTeX content injection.

WYATT MILLER Collinearity analysis; final diagnostic analysis; introduction; document formatting, layout, content injection, and \LaTeX encoding.