


Exercise #3 - Solution

Linear Regression – Diagnostics

This lab is designed to help you recall how to check the overall and case-by-case fit of the model to the data. Remember we do this to validate our model but also to ensure that we learn about the associations, and what influences them, as we go.

Begin by running the model with `-intvl_ln-` as the outcome and `-p_rct-`, `-hdsiz-`, `-pyear_ct-` and `-pyear_ct_sq-`.

1. Evaluate the assumption of homoscedasticity both graphically and statistically. What do you conclude? Explain the pattern of observations.



The standardized residuals show an apparent “fan” (or “cone”) pattern with larger variance at higher predicted values than at low predicted values. In reality, this is not easy to assess in a plot with so many points. The sheer volume of points with linear predictor values around 7 will naturally lead to a larger spread in the residuals in that range. One option to explore this further is to stratify the points into suitable intervals based on the fitted values, and then compute standard deviations of the standardized residuals in each of those intervals; we omit the details. The presence of heteroscedasticity is confirmed by different statistical tests:


```
. estat hettest
Breusch-Pagan / Cook-Weisberg test for heteroskedasticity
    Ho: Constant variance
    Variables: fitted values of intvl_ln

    chi2(1)      =    52.46
    Prob > chi2  =    0.0000
```

```
. estat imtest
Cameron & Trivedi's decomposition of IM-test
```

Source	chi2	df	p
Heteroskedasticity	153.62	13	0.0000
Skewness	194.35	4	0.0000
Kurtosis	145.04	1	0.0000

2. Evaluate the assumption of normality both graphically (using both a histogram and a normal probability plot) and statistically. What is your interpretation?



Both the histogram and the normal probability plot show departures from normality for the residuals (e.g., their distribution is left-skewed). This is confirmed by the Shapiro-Wilk test for normality.

```
. swilk stdres
```

Shapiro-Wilk W test for normal data

Variable	Obs	W	V	z	Prob>z
stdres	2987	0.96157	65.570	10.793	0.00000

3. Correcting problems of overall fit: It appears that we have problems with the overall fit. Often if we can correct one of the problems, that will correct both problems. So let's look and see what we might do to obtain somewhat more normally-distributed residuals.

- (a) First, start with the original outcome variable `-intvl-` and carry out a Box-Cox analysis using the `-boxcox-` command to see if there is a better transformation than “ln”. Note that we have already transformed Y to lnY, so we should go back and rerun the model with `-intvl-` as the outcome for this question. What does this approach suggest?

```
. boxcox intvl p_rct hdsiz e pye ar_ ct pye ar_ sq
...output omitted
```

```
Log likelihood = -24913.103
```

```
Number of obs   =      2987
LR chi2(4)      =      225.25
Prob > chi2     =      0.000
```

intvl	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
/theta	.1645652	.0177373	9.28	0.000	.1298008	.1993296

The Box Cox procedure suggests that the transformation which will improve the compliance with model assumptions (normality and homoscedasticity of errors) is $y^{0.165}$. However, applying this transformation does not produce a model which is much more satisfactory in terms of either normality or homoscedasticity.

```
. estat hettest
Breusch-Pagan / Cook-Weisberg test for heteroskedasticity
Ho: Constant variance
Variables: fitted values of intvl_bc

chi2(1)      =      83.17
Prob > chi2  =      0.0000
```

```
. estat imtest
Cameron & Trivedi's decomposition of IM-test
```

Source	chi2	df	p
Heteroskedasticity	245.49	13	0.0000
Skewness	66.38	4	0.0000
Kurtosis	248.40	1	0.0000
Total	560.27	18	0.0000

(b) Second, try applying the Box Cox method to the variable `intvl_ln`. In this instance we should obtain a $\theta = 1.886$ so in approximate terms we should model $-\text{intvl_ln}^2$. Does this help?

In this case, we are attempting a double transformation. First the data are log transformed and then squared. Unfortunately, it didn't work either (results not shown). It would in any case have been difficult to interpret the resulting transformation.

4. Subject by subject analysis of fit

- (a) We first look for poor fitting subjects (potential outliers). These will have standardized residual values >3 or <-3 . Are there any?

The residuals with the largest absolute values all have values >2.7 , as shown in the listing below.

```
. list id intvl_ln p_rct hdsiz p_year fit stdres delres if abs(stdres)>2.7 & stdres~.,
noobs
```

id	intvl_ln	p_rct	hdsiz	p_year	fit	stdres	delres
2489	3.526361	0	86	2000	6.6292	-2.963959	-2.967837
2357	4.110874	3	51	1995	7.042443	-2.799768	-2.802985
723	4.143135	1	24	1992	7.058354	-2.783901	-2.787059
2477	4.158883	1	44	1992	7.040209	-2.751371	-2.754408
668	4.158883	0	45	1993	7.03561	-2.74717	-2.750192
144	4.158883	3	84	1992	7.030485	-2.742186	-2.745189
445	4.127134	1	91	1995	6.979589	-2.724562	-2.727502
689	4.143135	1	78	1991	6.98419	-2.71317	-2.716069

The outlier test based on the deletion residuals indicates that deletion residuals numerically greater than 4.31 could be considered as significant outliers for this dataset. Based on this test, there are no significant outliers. Note however, that the test is based on the model assumptions being met, and that is clearly not the case here.

- (b) Lets now look for leverage subjects; these could have big impact on the model. Look for cases with extreme values. Are there many? What are their characteristics?

The leverage threshold value is = 0.00335 (or 0.0050 if you consider the more conservative estimate). There are quite a few observations with leverage values above the threshold. In general, they are either quite large herds or ones with large values of p_rct.

```
. list id intvl_ln p_rct hdsiz p_year fit stdres lev if lev>.02 & lev~., clean
```

	id	intvl_ln	p_rct	hdsiz	p_year	fit	stdres	lev
42.	436	5.752573	86	169	1989	7.933545	-2.159655	.0705684
65.	2901	5.313206	56	301	1998	7.361374	-1.982958	.027724
105.	62	5.231109	66	401	2001	7.104988	-1.826807	.0410773
491.	220	6.426488	65	299	1995	7.640965	-1.181592	.037215
598.	221	5.47227	7	450	1999	6.491442	-.9855473	.0254031
862.	222	5.749393	2	435	2000	6.339141	-.5702394	.0252234
938.	2075	7.463937	74	105	1993	7.964081	-.4905893	.0528041
1028.	2843	7.056175	55	280	1997	7.435472	-.3669779	.0264391
1310.	2065	7.914252	102	472	1991	7.968274	-.0542788	.0972719

- (c) We now move to identifying subjects that actually have a large influence on the model (using either or both of Cook's D or Dfits). Are there many? Do any in particular stand out?

Leverage is only a measure of "potential influence". However, the same pattern holds for observations with large Cook's D or Dfits ... they are generally large herds or have high values of -p_rct-. The threshold values for both Cook's D and dfits are 0.00134 and 0.0818.

```
. list id intvl_ln p_rct hdsiz p_year stdres fit cook if cook>0.007 & cook~.,
noobs clean
```

id	intvl_ln	p_rct	hdsiz	p_year	stdres	fit	cook
1159	4.941642	27	46	1991	-2.315489	7.358566	.0076161
220	6.426488	65	299	1995	-1.181592	7.640965	.0107933
2901	5.313206	56	301	1998	-1.982958	7.361374	.0224246
62	5.231109	66	401	2001	-1.826807	7.104988	.0285913
436	5.752573	86	169	1989	-2.159655	7.933545	.0708258

- (d) Finally, since `p_rct` is our main exposure of interest, we might see if any of the observations have a particularly large influence on the estimate of that coefficient. Again, what do you find?

As expected, the observations which have the most influence on the coefficient for `-p_rct` are herds which had high values for this predictor. The five largest `dbetas` for this predictor had more than 30 reactors (although some herds with larger number of reactors had smaller `dbetas`).

```
. list id intvl_ln p_rct hdsiz p_year stdres fit db_rct if abs(db_rct)>0.1 & db_rct~=. ,
noobs clean
```

id	intvl_ln	p_rct	hdsiz	p_year	stdres	fit	db_rct
436	5.752573	86	169	1989	-2.159655	7.933545	-.5800523
62	5.231109	66	401	2001	-1.826807	7.104988	-.2870666
2901	5.313206	56	301	1998	-1.982958	7.361374	-.2746634
220	6.426488	65	299	1995	-1.181592	7.640965	-.2015166
1159	4.941642	27	46	1991	-2.315489	7.358566	-.1879207
64	5.31812	30	83	1993	-1.994868	7.39961	-.1696355
142	4.779123	24	107	1991	-2.377164	7.263377	-.1461776
390	5.777652	33	115	1993	-1.565816	7.410426	-.1410908
2395	5.003946	17	8	1991	-2.157934	7.260216	-.1155009
2075	7.463937	74	105	1993	-.4905893	7.964081	-.1148877
2725	5.198497	19	29	1991	-1.979187	7.267729	-.112664
2205	5.241747	26	158	1989	-1.824712	7.146574	-.1106877
1066	5.662961	26	104	1997	-1.481023	7.209951	-.1000493
4	8.699348	42	120	1990	1.202291	7.449911	.1448241

- (e) Try refitting the model but leave out either the observations which have the biggest impact on overall fit, or the ones which most influence the estimate of the coefficient for `-p_rct`. Does this make much difference?

Neither of these re-fits are very different from the original model.

Conclusion

At this point you might be quite despondent (probably verging on suicidal) about not being able to fix the problems of heteroscedasticity and non-normality. Three more advanced approaches that we could take to solving this problem are as follows.

- 1. Use robust standard errors which are less susceptible to violations of the major assumptions. (These will be introduced when we talk about clustered data.)*
- 2. The data are really “time-to-event” data so survival analysis methods may be more applicable. (These are discussed later in the semester.)*
- 3. We could categorize the outcome into 2 levels (short vs long) and fit a logisitic model or into >2 levels and fit some sort of multinomial model. The former will covered in this semester; the latter will not. However this approach entails the loss of a lot of potentially useful information.*
- 4. We could go back to Ireland and regenerate data that did not have the structural problem of the relationship between “year” and “interval” built in.*