

Quantitative and Formal Methods in IR

Class 3: Multiple Linear Regression

Dave Armstrong
University of Oxford
Department of Politics and International Relations
Center for Research Methods in the Social Sciences
Nuffield College

February 27, 2009

For today's class, we will be using the dataset we've been using `class3.dta`, which you can get from my website, <http://www.quantoid.net/IRStats.php>.

1 Linear Regression as Difference of Means

I suggested in class that the OLS model is nothing more than a powerful difference of means tester. Let's see how this works. First, let's do a t-test of the null hypothesis that there is no difference in the mean of `gdp/capita` for free states versus not-free states. To do this, you'll first have to make a binary variable out of the `freedomstat` variable.

```
. recode freedomstat (1/2=0 "Not/Partly Free") (3=1 "Free"), into(free)
(176 differences between freedomstat and free)

. ttest gdppc, by(free) unpaired unequal

Two-sample t test with unequal variances
-----+-----
      Group |      Obs      Mean   Std. Err.   Std. Dev.   [95% Conf. Interval]
-----+-----
Not/Part |      101   5952.734   721.1529   7247.497   4521.987   7383.48
   Free |       75  17324.88   1267.405   10976.05  14799.52  19850.24
-----+-----
combined |      176  10798.82   800.4426   10619.07   9219.055  12378.58
-----+-----
      diff |           -11372.15    1458.21           -14259.22   -8485.073
-----+-----
      diff = mean(Not/Part) - mean(Free)                t =  -7.7987
Ho: diff = 0                Satterthwaite's degrees of freedom = 120.339

      Ha: diff < 0                Ha: diff != 0                Ha: diff > 0
Pr(T < t) = 0.0000            Pr(|T| > |t|) = 0.0000            Pr(T > t) = 1.0000
```

Now, let's do the regression where `gdp/capita` is the dependent variable and `free` is the independent variable.

Table 1: Differences in the Mean of GDP/Capita of Each Civilization vs. the “Other” Category

| Civilization | Difference |
|----------------|------------|
| African | |
| Islamic | |
| Latin American | |
| Orthodox | |
| Western | |

| Source | SS | df | MS | Number of obs = | 176 |
|----------|------------|-----|------------|-----------------|--------|
| Model | 8.8743e+09 | 5 | 1.7749e+09 | F(5, 170) = | 27.78 |
| Residual | 1.0860e+10 | 170 | 63879574.9 | Prob > F = | 0.0000 |
| | | | | R-squared = | 0.4497 |
| | | | | Adj R-squared = | 0.4335 |
| Total | 1.9734e+10 | 175 | 112764681 | Root MSE = | 7992.5 |

| gdppc | Coef. | Std. Err. | t | P> t | [95% Conf. Interval] |
|----------|-----------|-----------|-------|-------|----------------------|
| african | -4100.438 | 1969.366 | -2.08 | 0.039 | -7988 -212.8768 |
| islamic | 453.1272 | 1798.59 | 0.25 | 0.801 | -3097.32 4003.575 |
| latinam | -467.0987 | 2188.828 | -0.21 | 0.831 | -4787.882 3853.684 |
| orthodox | 2225.985 | 2551.627 | 0.87 | 0.384 | -2810.97 7262.94 |
| western | 16542.14 | 1836.143 | 9.01 | 0.000 | 12917.57 20166.72 |
| _cons | 7855.791 | 1263.72 | 6.22 | 0.000 | 5361.186 10350.4 |

So, remember from lecture that when we have a variable that has k categories, we include $k - 1$ binary variables and leave one of them out as the reference category. The coefficient for `african`, for instance, can be interpreted as the difference between African civilizations and the others is -4100. That is, African civilizations, on average, have GDP/capita \$4100 lower than the countries coded “other”. Western countries have GDP/capita of around \$16542 more than countries coded “other”. This matches nicely with the results from above. Regression, then, tells us the mean of the dependent variable, the best guess we could make, given the information we have available (in this case, the civilization to which the country belongs).

2 Linear Regression with Continuous Independent Variable(s)

We need to talk some about regression diagnostics - these are things that tell us how our model fits the data. We need to look at more than whether the coefficient fits or not.

First, let’s run the model and see what we get.

```
. reg new_polrts gdppc
```

| Source | SS | df | MS | Number of obs = | 176 |
|--------|----|----|----|-----------------|-----|
|--------|----|----|----|-----------------|-----|

| -----+----- | | | | F(1, 174) = 35.36 | |
|-------------|------------|-----|------------|------------------------|----------|
| Model | 26.6915535 | 1 | 26.6915535 | Prob > F | = 0.0000 |
| Residual | 131.348502 | 174 | .754876446 | R-squared | = 0.1689 |
| -----+----- | | | | Adj R-squared = 0.1641 | |
| Total | 158.040055 | 175 | .90308603 | Root MSE | = .86884 |

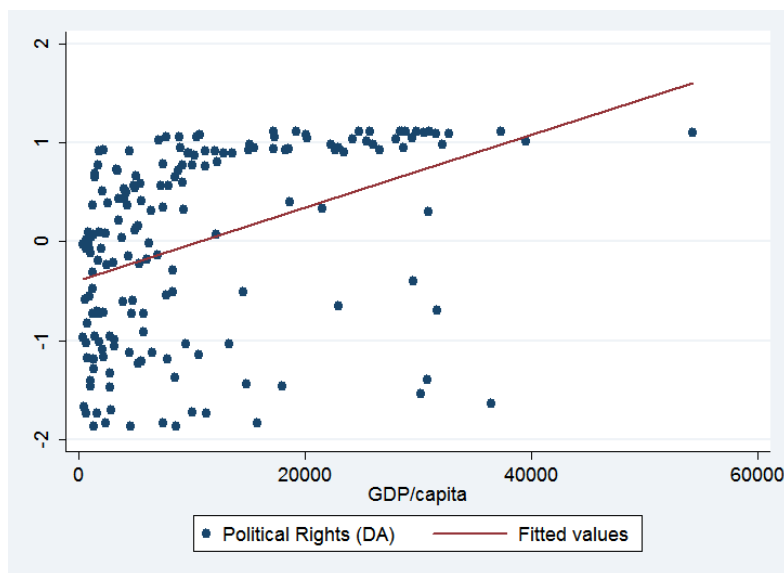
| -----+----- | | | | | | |
|-------------|--|-----------|-----------|-------|-------|----------------------|
| new_polrts | | Coef. | Std. Err. | t | P> t | [95% Conf. Interval] |
| gdppc | | .0000368 | 6.18e-06 | 5.95 | 0.000 | .0000246 .000049 |
| _cons | | -.3930351 | .093541 | -4.20 | 0.000 | -.5776562 -.2084141 |

How can we interpret this? The mean of political rights when GDP/capita equals zero is -0.393. Is this interesting? Not really because GDP/capita never equals zero in our dataset. For every dollar increase in GDP/capita, political rights increase by 0.000037. How much do political rights increase on average for every \$1000 increase in GDP/capita? Well, it's $0.000037 \times 1000 = 0.037$. The R^2 is around 0.17, which is not all that bad for one independent variable. So, on its face, this model looks alright. Let's start thinking about the assumptions we have to make:

2.1 Assumption 1: Linearity


How can we assess linearity? Well, for simple models with one independent variable we can simply look at the graph between the two variables that includes the fitted line:

```
twoway (scatter new_polrts gdppc) || (lfit new_polrts gdppc)
```




Does this relationship look linear? Not especially, it looks pretty curvy. So, are we simply done - we just pack up this research question and start on a new project? Not exactly. Often times, we can transform our data so that the relationship is linear. I know, this usually raises some red flags for people - they get uncomfortable when we transform the data to fit a particular functional form (i.e., a linear model). It is perfectly acceptable to do this as the information remains the same (i.e., there will be a one-to-one mapping of the transformed data onto the old data) and we will be able to transform the results back to talk about the data in their original metric.

How do we choose a transformation? Well, usually by eye. One way we can do this is with the “bulge rule” invented by Mosteller and Tukey. Let’s look at a picture and then we will discuss.



The bulging rule tells us, when the bulge in the data are pointing in a certain way, what we should do to either x and/or y . Generally, good advice is to start by transforming x . Which direction is the bulge here? It looks like it points toward the upper-left. This means we should either decrease the power on x from 1 to either $\frac{1}{2}$ (the square root), the log, etc... Often in International Relations, models that include either money or population will have these quantities logged to deal with this problem. Let’s try logging gdp/capita to see what happens:

```
gen log_gdppc = log(gdppc)
twoway (scatter new_polrts log_gdppc) || (lfit new_polrts log_gdppc)
```



How does that look, better? Well, it does look some better, but it's still not perfect. We're probably not going to get a more linear relationship through transformation again. Here's the regression output:

```
. reg new_polrts log_gdppc
```

| Source | SS | df | MS | Number of obs = | 176 |
|----------|------------|-----|------------|-----------------|--------|
| Model | 31.0122955 | 1 | 31.0122955 | F(1, 174) = | 42.48 |
| Residual | 127.02776 | 174 | .730044596 | Prob > F = | 0.0000 |
| Total | 158.040055 | 175 | .90308603 | R-squared = | 0.1962 |
| | | | | Adj R-squared = | 0.1916 |
| | | | | Root MSE = | .85443 |

| new_polrts | Coef. | Std. Err. | t | P> t | [95% Conf. Interval] |
|------------|-----------|-----------|-------|-------|----------------------|
| log_gdppc | .3593182 | .0551299 | 6.52 | 0.000 | .2505089 .4681276 |
| _cons | -3.128798 | .484976 | -6.45 | 0.000 | -4.085991 -2.171605 |

Now that we have created the log of GDP/capita, how do we interpret this in terms of actual GDP/capita figures? One way we can interpret this is:

As GDP/capita (not the log) goes up by one percent, we would expect political rights to change by $0.01 \times 0.36 = 0.0036$ units. Notice that the percentage change in the independent variable is what we multiply the coefficient by. A 10% change in GDP/capita should lead to a $0.1 \times 0.36 = 0.036$ change in expected political rights. We can also consider what an actual change in GDP/capita would do to the predictions. Below, we consider three one-thousand unit changes i) from 1000 to 2000, ii) from 9000-10000 and iii) from 50000 to 51000.

```
. sum gdppc log_gdppc
```

| Variable | Obs | Mean | Std. Dev. | Min | Max |
|-----------|-----|----------|-----------|----------|----------|
| gdppc | 176 | 10798.82 | 10619.07 | 400 | 54285.46 |
| log_gdppc | 176 | 8.719056 | 1.171571 | 5.991465 | 10.90201 |

What we can do, is see what an interesting change does to the predicted change in political rights. The somewhat problematic thing is that the *marginal effect* of GDP/capita is not constant because we performed a non-linear transformation. What this really means is that the effect of one-unit change in GDP/capita on political rights will be different depending on where you start. We have the following model:

$$\widehat{\text{new_polrts}} = -3.13 + 0.36 \log(\text{GDP/capita})$$

We can see what a one-thousand dollar change in GDP/capita will mean when we change from 1000 to 2000.

$$\begin{aligned}
\widehat{\text{new_polrts}}_0 &= -3.13 + 0.36\log(\text{GDP/capita}) \\
&= -3.13 + 0.36 * \log(1000) \\
&= -3.13 + 0.36 \times 6.9 \\
&= -0.64 \\
\widehat{\text{new_polrts}}_1 &= -3.13 + 0.36 * \log(2000) \\
&= -3.13 + 0.36 \times 7.6 \\
&= -0.39 \\
\widehat{\text{new_polrts}}_1 - \widehat{\text{new_polrts}}_0 &= -0.39 - (-0.64) \\
&= 0.25
\end{aligned}$$

Now, let's look at the difference in predictions for a one-thousand dollar change in GDP/capita from 9000 to 10000.

$$\begin{aligned}
\widehat{\text{new_polrts}}_0 &= -3.13 + 0.36\log(\text{GDP/capita}) \\
&= -3.13 + 0.36 * \log(9000) \\
&= -3.13 + 0.36 \times 9.1 \\
&= 0.146 \\
\widehat{\text{new_polrts}}_1 &= -3.13 + 0.36 * \log(10000) \\
&= -3.13 + 0.36 \times 9.2 \\
&= 0.182 \\
\widehat{\text{new_polrts}}_1 - \widehat{\text{new_polrts}}_0 &= 0.182 - 0.146 \\
&= 0.036
\end{aligned}$$

This predicted change in political rights moving from 9000 to 10000 dollars is much smaller than the change in predicted political rights for the same size change from 1000 to 2000 dollars. Let's do one more at the high end of the variable from 50000 to 51000:

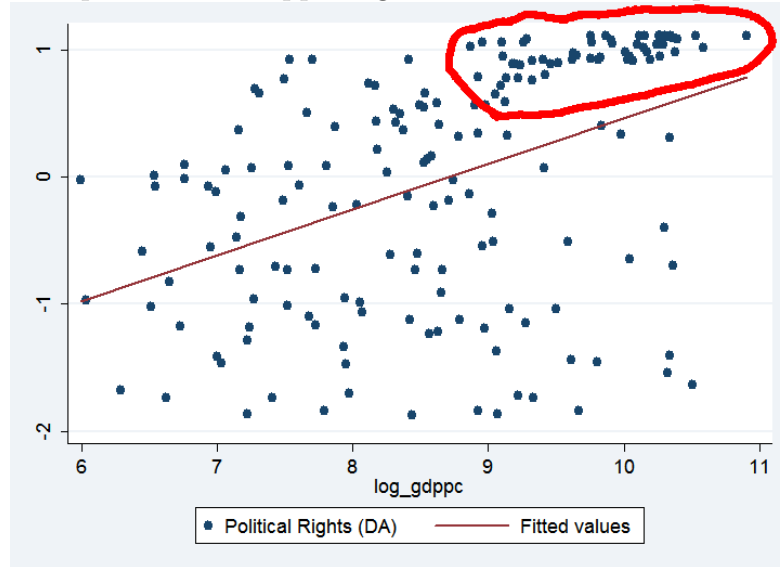
$$\begin{aligned}
\widehat{\text{new_polrts}}_0 &= -3.13 + 0.36\log(\text{GDP/capita}) \\
&= -3.13 + 0.36 * \log(50000) \\
&= -3.13 + 0.36 \times 10.82 \\
&= 0.765 \\
\widehat{\text{new_polrts}}_1 &= -3.13 + 0.36 * \log(51000) \\
&= -3.13 + 0.36 \times 10.84 \\
&= 0.772 \\
\widehat{\text{new_polrts}}_1 - \widehat{\text{new_polrts}}_0 &= 0.772 - 0.765 \\
&= 0.007
\end{aligned}$$

At the very high end of the data, the predicted change in political rights (0.007) for a one-thousand dollar change in GDP/capita (50000 to 51000) is about 3% of the size of

the change in predicted political rights for a one-thousand dollar change at the low end (from 1000 to 2000). When we log a variable it is then said to have a *decreasing marginal effect* - this means that as we move up the scale of the independent variable, a one-unit change has a smaller effect on the dependent variable.

2.2 Outliers

Remember what we said about outliers - that they pull the mean in the direction of the outlier. Look at the points in the upper right-hand corner of the plot below:

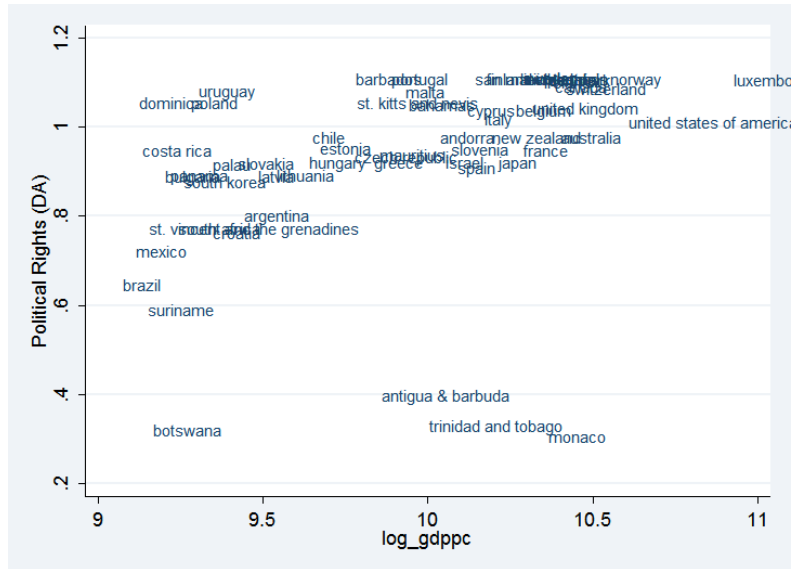


Remember, the points on the line tell us the mean of the dependent variable given the value of the independent variable. Given what we know about the mean and outliers, what effect do you think the points circled in the upper-left have on regression line?

2.2.1 Identifying the Outliers

We can use a graph to identify the outliers. We want to look at the graph where $\log_gdppc > 9$ and where $new_polrts > .25$

```
twoway (scatter new_polrts log_gdppc if log_gdppc > 10 & new_polrts > .8, mlabel(cname) symbol(none))
```



It might still not be obvious which countries are there, so we could look, for instance, at the civilization scores of the countries up in that area of the graph:

```
. tab civ_small if log_gdppc > 9 & new_polrts > .25
```

| RECODE of civ (Civilizations (Henderson and Tucker)) | Freq. | Percent | Cum. |
|---|-------|---------|--------|
| Other | 12 | 20.34 | 20.34 |
| African | 3 | 5.08 | 25.42 |
| Latin American | 7 | 11.86 | 37.29 |
| Orthodox | 3 | 5.08 | 42.37 |
| Western | 34 | 57.63 | 100.00 |
| Total | 59 | 100.00 | |

It looks, as though many of those countries are in the Western civilization. What should happen if we add the civilizations into the model? Well, it will make a different intercept for each civilization. Let's do that

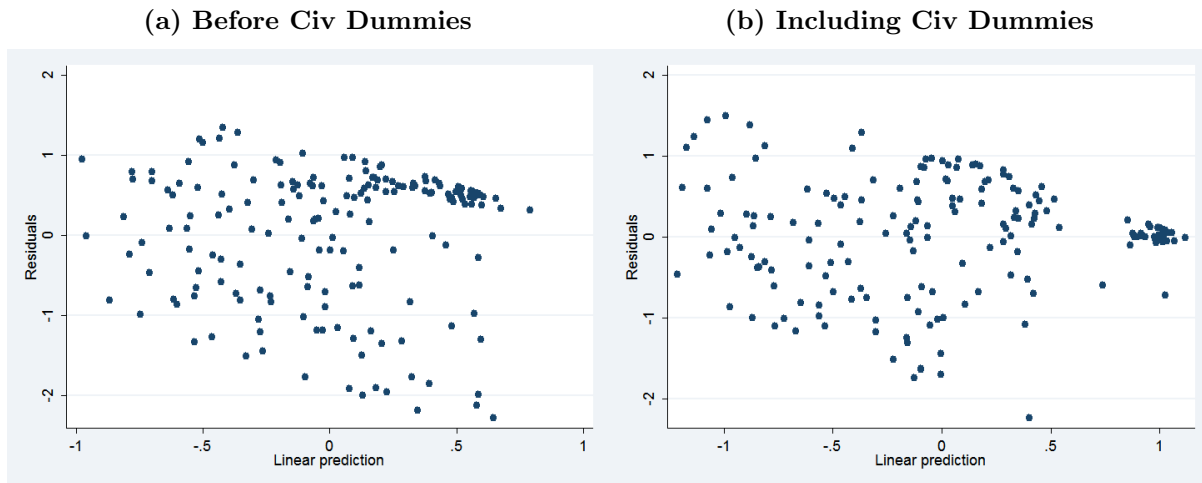
```
. reg new_polrts log_gdppc african-western
```

| Source | SS | df | MS | Number of obs = | 176 |
|----------|------------|-----|------------|-----------------|----------|
| Model | 73.4498214 | 6 | 12.2416369 | F(6, 169) = | 24.46 |
| Residual | 84.5902338 | 169 | .500533928 | Prob > F | = 0.0000 |
| Total | 158.040055 | 175 | .90308603 | R-squared | = 0.4648 |
| | | | | Adj R-squared | = 0.4458 |
| | | | | Root MSE | = .70748 |

| new_polrts | Coef. | Std. Err. | t | P> t | [95% Conf. Interval] |
|------------|-----------|-----------|-------|-------|----------------------|
| log_gdppc | .161227 | .0588314 | 2.74 | 0.007 | .045088 .277366 |
| african | -.292755 | .1810181 | -1.62 | 0.108 | -.650103 .064593 |
| islamic | -.9382975 | .1592205 | -5.89 | 0.000 | -1.252615 -.6239803 |
| latinam | .2534227 | .1948103 | 1.30 | 0.195 | -.1311524 .6379978 |
| orthodox | -.2616364 | .2284643 | -1.15 | 0.254 | -.7126478 .1893751 |
| western | .6491538 | .1847909 | 3.51 | 0.001 | .284358 1.01395 |
| _cons | -1.289391 | .5107095 | -2.52 | 0.013 | -2.297583 -.2811991 |

Now, how do we check to see if what we did makes sense? We can look at the residuals plotted against the log_gdppc variable.

Figure 1: Residual Versus Fitted Plots



The plots show the before the inclusion of the civilization dummies, there were some problems here at the high levels of `log_gdppc`, there were lots of observations above the line and not many below it. Remember, we assumed that the mean of the residuals was zero everywhere. When this assumption is violated, it often can suggest that we have forgotten to include an important variable. We included the civilization dummy variables. The second plot shows that after the inclusion of the civilization dummies, we have taken care of that problem. It seems that the mean of the residuals is probably everywhere zero. However, we have induced another problem - heteroskedasticity.

Heteroskedasticity can be fixed, or rather it can be accounted for. We don't "fix" it per se, but we can generate standard error estimates that are robust with respect to heteroskedasticity. That is to say, we can still make appropriate inferences about the coefficients in the presence of heteroskedasticity. It's hard to talk about what exactly these do without getting into some math, but basically they allow the residuals to have a different variance for every value of \hat{y} . Stata can produce these for you by putting the option `robust` after the comma in your regression command:

```
. reg new_polrts log_gdppc african-western, robust
```

```
Linear regression                Number of obs =    176
                                F( 6, 169) =    55.57
                                Prob > F      =    0.0000
                                R-squared      =    0.4648
                                Root MSE   =    .70748
```

| | Coef. | Robust Std. Err. | t | P> t | [95% Conf. Interval] | |
|------------------------|-----------|------------------|-------|-------|----------------------|-----------|
| <code>log_gdppc</code> | .161227 | .0677961 | 2.38 | 0.019 | .0273908 | .2950632 |
| <code>african</code> | -.292755 | .2214909 | -1.32 | 0.188 | -.7300002 | .1444902 |
| <code>islamic</code> | -.9382975 | .1861353 | -5.04 | 0.000 | -1.305747 | -.5708477 |
| <code>latinam</code> | .2534227 | .1947079 | 1.30 | 0.195 | -.1309503 | .6377957 |
| <code>orthodox</code> | -.2616364 | .2585704 | -1.01 | 0.313 | -.7720803 | .2488075 |
| <code>western</code> | .6491538 | .1545776 | 4.20 | 0.000 | .3440021 | .9543056 |
| <code>_cons</code> | -1.289391 | .6197626 | -2.08 | 0.039 | -2.512865 | -.0659174 |

Notice, that this does not change the coefficient estimates at all. This is a *post hoc* fix to the standard errors. That is, the model is estimated with the normal standard errors and then some fix is applied without re-estimating the model. Notice that some of the p-values have changed here. This is to be expected, though will not necessarily happen. Like I said, this didn't "fix" the heteroskedasticity - that is if we look at the graph of the residuals versus fitted, they will still look messed up. We are just able to make the right inferences in the presence of the messed up residuals.

Looking at the plot, this is certainly not perfect. However, it is better than it could be. There is one more step we can take with these data. We can not only allow each civilization to have its own intercept, we can allow each civilization to have its own slope of \log_gdppc . We are then getting estimates of:

$$\begin{aligned}
 PR_O &= a_O + b_O \log(GDP/capita)_O \\
 PR_A &= a_A + b_A \log(GDP/capita)_A \\
 PR_I &= a_I + b_I \log(GDP/capita)_I \\
 PR_L &= a_L + b_L \log(GDP/capita)_L \\
 PR_{Or} &= a_{Or} + b_{Or} \log(GDP/capita)_{Or} \\
 PR_W &= a_W + b_W \log(GDP/capita)_W
 \end{aligned}$$

```
. reg new_polrts african islamic latinam orthodox western log_gdppc
    afrxgdp islxgdp latxgdp orthxgdp westxgdp
```

| Source | SS | df | MS | Number of obs = | 176 |
|----------|------------|-----|------------|-----------------|----------|
| Model | 86.3882947 | 11 | 7.85348134 | F(11, 164) = | 17.98 |
| Residual | 71.6517605 | 164 | .436900979 | Prob > F | = 0.0000 |
| | | | | R-squared | = 0.5466 |
| | | | | Adj R-squared | = 0.5162 |
| Total | 158.040055 | 175 | .90308603 | Root MSE | = .66098 |

| new_polrts | Coef. | Std. Err. | t | P> t | [95% Conf. Interval] |
|------------|-----------|-----------|-------|-------|----------------------|
| african | 2.993281 | 1.262939 | 2.37 | 0.019 | .4995646 5.486997 |
| islamic | 5.377771 | 1.184767 | 4.54 | 0.000 | 3.038407 7.717134 |
| latinam | 1.166195 | 3.055931 | 0.38 | 0.703 | -4.867847 7.200237 |
| orthodox | 2.637315 | 3.009947 | 0.88 | 0.382 | -3.30593 8.580559 |
| western | 3.663915 | 1.941193 | 1.89 | 0.061 | -.1690377 7.496868 |
| log_gdppc | .545989 | .1018422 | 5.36 | 0.000 | .344898 .7470801 |
| afrxgdp | -.3883034 | .1564956 | -2.48 | 0.014 | -.6973093 -.0792974 |
| islxgdp | -.747072 | .139015 | -5.37 | 0.000 | -1.021562 -.472582 |
| latxgdp | -.1185928 | .3472236 | -0.34 | 0.733 | -.8041978 .5670121 |
| orthxgdp | -.3449979 | .3335723 | -1.03 | 0.303 | -1.003648 .3136523 |
| westxgdp | -.3602539 | .2014846 | -1.79 | 0.076 | -.7580922 .0375844 |
| _cons | -4.548367 | .8689226 | -5.23 | 0.000 | -6.264084 -2.832649 |

This is just a bit more complicated version of what we did yesterday. However, just think of this as another plugging in numbers problem. For the "other" civilization, all of the dummy variables will be zero, which leaves:

$$-4.55 + 0.55 \log(GDP/capita)$$

For the African civilization, the `africa` variable gets a 1 and all other civilization dummies get a value of zero. This leaves:

$$-4.55 + 2.99 + 0.55 \log(\text{GDP/capita}) - 0.39 \log(\text{GDP/capita}) = -1.56 + (0.55 - 0.39) \log(\text{GDP/capita})$$

We could write out all of the other equations, too:

| | | |
|----------|--|--|
| Other | $-4.55 + 0.55 \log(\text{GDP/capita})$ | $-4.55 + 0.55 \log(\text{GDP/capita})$ |
| African | $-4.55 + 2.99 + 0.55 \log(\text{GDP/capita}) - 0.39 \log(\text{GDP/capita})$ | $-1.56 + 0.16 \log(\text{GDP/capita})$ |
| Islamic | $-4.55 + 5.38 + 0.55 \log(\text{GDP/capita}) - 0.75 \log(\text{GDP/capita})$ | $0.83 - 0.2 \log(\text{GDP/capita})$ |
| Latin Am | $-4.55 + 1.17 + 0.55 \log(\text{GDP/capita}) - 0.12 \log(\text{GDP/capita})$ | $-3.38 + 0.43 \log(\text{GDP/capita})$ |
| Orthodox | $-4.55 + 2.64 + 0.55 \log(\text{GDP/capita}) - 0.34 \log(\text{GDP/capita})$ | $-1.91 + 0.21 \log(\text{GDP/capita})$ |
| Western | $-4.55 + 3.66 + 0.55 \log(\text{GDP/capita}) - 0.36 \log(\text{GDP/capita})$ | $-0.89 + 0.19 \log(\text{GDP/capita})$ |

So, we basically have a different equation for each civilization. These are what we refer to as conditional coefficients and conditional intercepts. This just means that the intercept and slope of `log_gdppc` is conditional on some other variable - civilizations. What if we want to know whether the conditional coefficient is significant. For instance, let's look at the conditional coefficient for the African civilization.

```
lincom log_gdppc + afrxgdp
```

```
( 1) log_gdppc + afrxgdp = 0
```

| new_polrts | Coef. | Std. Err. | t | P> t | [95% Conf. Interval] |
|------------|----------|-----------|------|-------|----------------------|
| (1) | .1576857 | .1188235 | 1.33 | 0.186 | -.0769354 .3923068 |

This tells us that the conditional coefficient for `log_gdppc` for the African civilization is -.16. Since the p-value is bigger than 0.05, we cannot reject the null that this conditional coefficient is zero. That means that in Africa, political rights are not well explained by GDP/capita. We could do this for the other coefficients as well:

```
. lincom log_gdppc + islxgdp
```

```
( 1) log_gdppc + islxgdp = 0
```

| new_polrts | Coef. | Std. Err. | t | P> t | [95% Conf. Interval] |
|------------|----------|-----------|-------|-------|----------------------|
| (1) | -.201083 | .0946221 | -2.13 | 0.035 | -.3879176 -.0142484 |

```
. lincom log_gdppc + latxgdp
```

```
( 1) log_gdppc + latxgdp = 0
```

| new_polrts | Coef. | Std. Err. | t | P> t | [95% Conf. Interval] |
|------------|----------|-----------|------|-------|----------------------|
| (1) | .4273962 | .3319524 | 1.29 | 0.200 | -.2280552 1.082848 |

```
. lincom log_gdppc + orthxgdp
```


```
( 1) log_gdppc + orthxgdp = 0
```

| new_polrts | Coef. | Std. Err. | t | P> t | [95% Conf. Interval] |
|------------|----------|-----------|------|-------|----------------------|
| (1) | .2009912 | .3176455 | 0.63 | 0.528 | -.4262109 .8281932 |

```
. lincom log_gdppc + westxgdp
( 1) log_gdppc + westxgdp = 0
```

| new_polrts | Coef. | Std. Err. | t | P> t | [95% Conf. Interval] |
|------------|----------|-----------|------|-------|----------------------|
| (1) | .1857351 | .1738511 | 1.07 | 0.287 | -.1575399 .5290102 |

The only significant conditional coefficient is the one for the Islamic world - there, as GDP/capita increases, political rights get worse. How would we do this for the intercepts? Now, let's go ahead and look at the residuals again:



This shows that again, the residuals look as though the mean is everywhere zero, perhaps more so than in the previous plot, however the fan-shape indicating heteroskedasticity is still a problem. Let's re-estimate the model above with robust standard errors and then re-calculate the conditional coefficients and their standard errors.

```
. reg new_polrts african islamic latinam orthodox western log_gdppc afrxgdp islxgdp latxgdp orthxgdp w
> estxgdp, robust
```

```
Linear regression                               Number of obs =    176
                                                F( 11, 164) =    46.31
                                                Prob > F      =    0.0000
                                                R-squared    =    0.5466
                                                Root MSE    =    .66098
```

| new_polrts | Coef. | Robust Std. Err. | t | P> t | [95% Conf. Interval] |
|------------|-----------|------------------|-------|-------|----------------------|
| african | 2.993281 | 1.543228 | 1.94 | 0.054 | -.0538772 6.040438 |
| islamic | 5.377771 | 1.387917 | 3.87 | 0.000 | 2.637281 8.118261 |
| latinam | 1.166195 | 1.710824 | 0.68 | 0.496 | -2.211887 4.544277 |
| orthodox | 2.637315 | 3.428734 | 0.77 | 0.443 | -4.13284 9.407469 |
| western | 3.663915 | 1.332659 | 2.75 | 0.007 | 1.032534 6.295296 |
| log_gdppc | .545989 | .1296443 | 4.21 | 0.000 | .290002 .8019761 |
| afrxgdp | -.3883034 | .1946084 | -2.00 | 0.048 | -.7725645 -.0040422 |
| islxgdp | -.747072 | .1598509 | -4.67 | 0.000 | -1.062703 -.431441 |
| latxgdp | -.1185928 | .2010911 | -0.59 | 0.556 | -.5156542 .2784685 |
| orthxgdp | -.3449979 | .3965785 | -0.87 | 0.386 | -1.128056 .4380601 |
| westxgdp | -.3602539 | .1478462 | -2.44 | 0.016 | -.6521813 -.0683265 |
| _cons | -4.548367 | 1.122128 | -4.05 | 0.000 | -6.764047 -2.332686 |

.
end of do-file

. lincom log_gdppc + afrxgdp

(1) log_gdppc + afrxgdp = 0

new_polrts | Coef. Std. Err. t P>|t| [95% Conf. Interval]
-----+-----
 (1) | .1576857 .1451372 1.09 0.279 -.1288928 .4442641

. lincom log_gdppc + islxgdp

(1) log_gdppc + islxgdp = 0

new_polrts | Coef. Std. Err. t P>|t| [95% Conf. Interval]
-----+-----
 (1) | -.201083 .0935129 -2.15 0.033 -.3857275 -.0164385

. lincom log_gdppc + latxgdp

(1) log_gdppc + latxgdp = 0

new_polrts | Coef. Std. Err. t P>|t| [95% Conf. Interval]
-----+-----
 (1) | .4273962 .1537205 2.78 0.006 .1238698 .7309226

. lincom log_gdppc + orthxgdp

(1) log_gdppc + orthxgdp = 0

new_polrts | Coef. Std. Err. t P>|t| [95% Conf. Interval]
-----+-----
 (1) | .2009912 .3747891 0.54 0.592 -.5390429 .9410252

. lincom log_gdppc + westxgdp

(1) log_gdppc + westxgdp = 0

new_polrts | Coef. Std. Err. t P>|t| [95% Conf. Interval]
-----+-----
 (1) | .1857351 .0710694 2.61 0.010 .0454062 .3260641

Now that we've dealt appropriately with the standard errors, you can see that three of the conditional coefficients are statistically significant - the Islamic, Latin American and Western coefficients are all statistically significantly different from zero.