HIERARCHICAL LINEAR MODELS FOR ELECTORAL RESEARCH:
A Worked Example in Stata

Marco R. Steenbergen
Department of Political Science
University of Zurich

2012

Abstract

These notes discuss a worked example of a hierarchical linear model with extensive discussion about modeling decisions and interpretation. Note that variable names and Stata commands appear in a `typewriter` font.

Setup

Imagine we are interested in evaluations of Barack Obama in 2008. We distinguish between two levels of analysis: the individual (level-1) and the `state` in which he/she resides (level-2). Evaluations of Barack Obama are measured on a feeling thermometer (`obamafeel`), which we treat as a continuous outcome. Our primary level-1 predictor is the race of the respondent (1 = `black`; 0 = white). In addition, we consider partisanship (`pid`) and the individual's ideological distance to Obama (`iddist`) as predictors. Partisanship runs from 0 (= strong Democrat) to 6 (= strong Republican), while ideological distance runs from 0 (the respondent and Obama are placed at the same ideological position) to 10 (the respondent and Obama are at opposite ends of the ideological spectrum). The analysis is based on the 2008 American National Election Studies.[1]

[1] The data are downloadable from `www.electionstudies.org`.

Random Effects ANOVA

As a starting point for the analysis, it is useful to run a random effects ANOVA. This allows us to determine what portion of the variance in candidate evaluations is due to cross-state differences as compated to individual differences. The model may be written as

$$\begin{aligned} obamafeel_{ij} &= \beta_{0j} + \epsilon_{ij} \\ \beta_{0j} &= \gamma_{00} + \delta_{0j} \end{aligned}$$

for individual i in state j. Here, β_{0j} can be interpreted as the mean evaluation of Obama in a state, whereas γ_{00} is the grand mean (i.e., the mean across all individuals and states). The level-1 error term (ϵ_{ij}) shows how an individual's evaluation deviates from the mean evaluation in the state in which he/she resides. The level-2 error term (δ_{0j}) shows how the mean evaluation in a particular state deviates from the grand mean. In addition to estimating γ_{00}, we are interested in estimating $Var(\epsilon_{ij}) = \sigma^2$ and $Var(\delta_{0j}) = \tau_{00}$.

We estimate the model using Stata's `xtmixed` command, using the following syntax:

```
xtmixed obamafeel if flagmis==0, || state:, var
```

where `flagmis=0` selects those cases that are not missing on the level-1 covariates or level-2 covariate (to be discussed later). The `var` option causes Stata to report variances rather than standard deviations, which is the default in `xtmixed`. The output is shown in Figure 1.

The output is divided into four sections. At the top one finds descriptive information about the analysis. This is followed by a listing of the fixed effects estimates. The third section displays the variance component estimates. Finally, at the bottom of the output, a likelihood ratio test is shown.

Focusing on the descriptive part, we see that there are 1673 level-1 units in the analysis (`number of obs`), and 33 level-2 units (`number of groups`). The cluster size varies from 11 to 258, with an average cluster size of just shy of 51 observations. The log-likelihood is -7973.552.[2]

In the fixed effects part, we see an estimate associated with `_cons`. This estimate of 64.54 is the estimate of the grand mean γ_{00}. Thus, averaging acros respondents

[2]The parts about the Wald test are left empty because there are no level-1 covariates in the model whose effects would be tested.

```
Mixed-effects ML regression                     Number of obs      =      1673
Group variable: state                           Number of groups   =        33

                                                Obs per group: min =        11
                                                               avg =      50.7
                                                               max =       258


                                                Wald chi2(0)       =         .
Log likelihood =  -7973.552                     Prob > chi2        =         .

------------------------------------------------------------------------------
   obamafeel |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   64.54253   1.793141    35.99   0.000     61.02804    68.05702
------------------------------------------------------------------------------

------------------------------------------------------------------------------
  Random-effects Parameters  |   Estimate   Std. Err.     [95% Conf. Interval]
-----------------------------+------------------------------------------------
state: Identity              |
                   var(_cons) |   81.25812   28.19668      41.16221    160.4113
-----------------------------+------------------------------------------------
                var(Residual) |   782.5182   27.37138       730.669    838.0468
------------------------------------------------------------------------------
LR test vs. linear regression: chibar2(01) =    52.57 Prob >= chibar2 = 0.0000
```

Figure 1: Stata `xtmixed` Output for a Random Effects ANOVA

and states, the expected feeling thermometer rating of Obama was just shy of 65. The standard error is 1.79, resulting in a z-test statistic of 35.99, and a p-value of .000. Clearly, the grand mean is significantly different from zero, which should not come as too great of a surprise considering that Obama won the election. More telling is the 95% confidence interval, which runs roughly from 61.03 to 68.06. On feeling thermometers, a value of 50 is sometimes interpreted as a neutral candidate evaluation. Since this value is not included in the confidence interval, one has clear evidence of a positive attitude toward Obama on the average.

In the third section of the output, we find the estimates of the variance components. The estimate of τ_{00} is given as `var(_cons)` and is 81.26, with a 95% confidence interval running from 41.16 to 160.41. The estimate of σ^2 is given as `var(Residual)` and is 782.52, with a 96% confidence interval from 730.67 to 838.05.

The final portion of the output provides the likelihood (LR) ratio test statistic for the null hypothesis that $\tau_{00} = 0$, i.e., that there is no cross-state variation in evaluations of Obama. The LR test compares to fit of the random effects ANOVA to that of a an ordinary regression model with a constant only. The latter model allows only for individual variation. The test statistic has one degree of freedom—we are testing just one variance component—and comes out to 52.57. The reported

p-value is .000 and should actually be halved to obtain a less conservative test.[3] In this case, halving does not affect the conclusion: the null hypothesis has to be rejected—there is evidence of cross-state variation in the evaluations of Obama.

Intra-class Correlation Based on the estimates in the third part of the output, we can compute the intra-class correlation (ICC). Using the definition of the ICC, its estimator is

$$\hat{\rho} = \frac{\hat{\tau}_{00}}{\hat{\tau}_{00} + \hat{\sigma}^2}$$

Substituting 81.26 for $\hat{\tau}_{00}$ and 782.52 for $\hat{\sigma}^2$, the ICC is estimated as $\hat{\rho} = .094$. This means that around 9 percent of the variance in Obama evaluations is due to differences across states, with the remaining 91 percent attributable to individual differences.

Random Intercept Model

We now add level-1 covariates to the model, for now assuming fixed effects. The intercept, however, is allowed to vary across states, in order to accommodate cross-state differences in the baseline evaluation of Obama. The mixed model that is to be estimated is thus

$$obamafeel_{ij} = \gamma_{00} + \gamma_{10} pid_{ij} + \gamma_{20} iddist_{ij} + \gamma_{30} black_{ij} + \delta_{0j} + \epsilon_{ij}$$

This can be estimated in Stata using

```
xtmixed obamafeel pid iddist black if flagmis==0, || state:,
var
```

The estimates are shown in Figure 2.

All three covariates have statistically significant effects. The Wald test shows, not surprisingly, that they are also jointly significant ($W = 1380.88$, $p = .000$). As one moves up one unit on the partisanship scale, ratings of Obama are expected to decrease by around 7 points. Compared to whites, the expected ratings of blacks are some 12 points higher. An anomalous result is that for ideological distance, as the coefficient suggests that an increased distance to Obama on the left-right scale

[3] Remember that we are testing a variance component, for which the alternative hypothesis is of necessity one-sided, namely $\tau_{00} > 0$ (negative variances, which would be allowed under a two-sided test, do not make sense).

```
Mixed-effects ML regression                     Number of obs      =      1673
Group variable: state                           Number of groups   =        33

                                                Obs per group: min =        11
                                                               avg =      50.7
                                                               max =       258


                                                Wald chi2(3)       =   1380.88
Log likelihood = -7476.7372                     Prob > chi2        =    0.0000

------------------------------------------------------------------------------
   obamafeel |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         pid |   -6.99571   .2956337   -23.66   0.000    -7.575142   -6.416279
      iddist |   1.332214   .2142992     6.22   0.000     .9121954    1.752233
       black |   12.27659   1.405161     8.74   0.000     9.522521    15.03065
       _cons |   72.56996   1.507307    48.15   0.000      69.6157    75.52423
------------------------------------------------------------------------------

------------------------------------------------------------------------------
  Random-effects Parameters  |   Estimate   Std. Err.     [95% Conf. Interval]
-----------------------------+------------------------------------------------
state: Identity              |
                   var(_cons) |   11.18923   5.948333      3.947216    31.71827
-----------------------------+------------------------------------------------
                var(Residual) |   439.6678   15.36251      410.5657    470.8327
------------------------------------------------------------------------------
LR test vs. linear regression: chibar2(01) =    11.20 Prob >= chibar2 = 0.0004
```

Figure 2: Stata `xtmixed` Output for a Random Intercept Model

actually improves his ratings. This result may possibly be a consequence of model misspecification, i.e., ideological distance picking up some other effect.

There is evidence of variation in the intercepts. The variance of the δ_{0j}s is estimated as 11.19, which is sizable. Moreover, comparing the fit of the random intercept model to that of a regression model yields $LR = 11.20$ with a halved p-value of .000. Thus, we can clearly reject the null hypothesis that the intercept is the same across all of the states, as the regression model assumes.

Explained Variance How much of the variance do the three level-1 covariates explain? To answer this question, we compare the variance in the outcome in the random intercept model and that in the random effects ANOVA. Specifically, we evaluate

$$R_1^2 \quad = \quad 1 - \frac{Var_{New}(Y)}{Var_{Old}(Y)}$$

where $Var_{New}(Y)$ is the variance in Obama feelings in the random intercept model and $Var_{Old}(Y)$ is the variance in the random effects ANOVA. From the estimates in Figure 1, $Var_{Old} = 81.26 + 782.52 = 863.78$. From Figure 2, we get $Var_{New} = 11.19 + 439.67 = 450.86$. Consequently, $R_1^2 = .48$ so that the three covariates

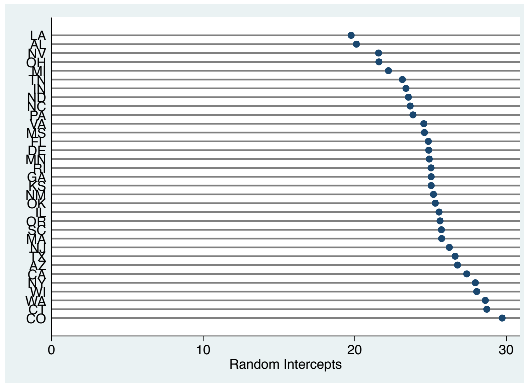

Figure 3: Empirical Bayes Estimates of the Random Intercepts Across States

account for about 48 percent of the variation in the outcome.[4]

Empirical Bayes Estimates It is useful to get a sense of how the intercepts vary across the states. To that effect, we compute

$$\hat{\beta}_{0j} \quad = \quad \hat{\gamma}_{00} + \hat{\delta}_{0j}$$

where $\hat{\delta}_{0j}$ is the empirical Bayes residual. We can then display the estimates in a dot plot. The following syntax accomplishes this:

```
predict ebres*, reffects level(state)
gen beta0=obamafeel[_cons]+ebres1
graph dot (mean) beta0, over(V081201a, gap(V081201a) sort(beta0))
cw linetype(line) ytitle(Random Intercepts)
```

where `V081201a` is a variable containing the abbreviation of the state name. The resulting estimates are shown in Figure 3. We see that the smallest estimated intercept (19.77) occurs for Louisiana, and the highest one (29.73) occurs for Colorado. The spread in intercept values is considerable, further evidence that cross-state variation in the evaluations of Obama was considerable in 2008.

Accounting for Variation in the Intercepts

We now seek to account for the variation in the intercepts and do so by including the percentage of blacks in each state (`pctblack`). Thus, we estimate the following

[4]For comparison, the OLS regression R^2 is .46, which is comparable.

```
Mixed-effects ML regression                     Number of obs      =      1673
Group variable: state                           Number of groups   =        33

                                                Obs per group: min =        11
                                                               avg =      50.7
                                                               max =       258


                                                Wald chi2(4)       =   1405.09
Log likelihood = -7473.6068                     Prob > chi2        =    0.0000

------------------------------------------------------------------------------
   obamafeel |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         pid |   -6.95944   .2957349   -23.53   0.000    -7.53907    -6.37981
      iddist |   1.329306   .2139513     6.21   0.000     .9099689    1.748643
       black |   13.02433   1.434679     9.08   0.000    10.21242    15.83625
    pctblack |  -.2181398   .0833902    -2.62   0.009   -.3815817   -.054698
       _cons |   75.28452   1.766732    42.61   0.000    71.82179    78.74726
------------------------------------------------------------------------------

------------------------------------------------------------------------------
  Random-effects Parameters  |   Estimate   Std. Err.     [95% Conf. Interval]
-----------------------------+------------------------------------------------
state: Identity              |
                   var(_cons) |   7.271424   4.766288      2.012205    26.27646
-----------------------------+------------------------------------------------
                var(Residual) |   439.6309   15.35853      410.5362    470.7876
------------------------------------------------------------------------------
LR test vs. linear regression: chibar2(01) =     5.37 Prob >= chibar2 = 0.0102
```

Figure 4: Stata `xtmixed` Output for a Hierarchical Linear Model with Random Intercept and Level-2 Covariates

mixed model:

$$\begin{aligned} obamafeel_{ij} &= \gamma_{00} + \gamma_{01} pctblack_j + \gamma_{10} pid_{ij} + \gamma_{20} iddist_{ij} + \gamma_{30} black_{ij} + \\ &\quad \delta_{0j} + \epsilon_{ij} \end{aligned}$$

We can also think of this as a model in which the state-level intercepts are given by

$$\beta_{0j} = \gamma_{00} + \gamma_{01} pctblack_j + \delta_{0j}$$

Estimation of the model is straightforward: one simply adds `pctblack` to the list of covariates in the `xtmixed` command. This produces the estimates shown in Figure 4.

Focusing on the fixed effects, we observe that the level-1 covariates remain statistically significant. The newly added level-2 covariate is also statistically significant at conventional levels. We see that Obama evaluations tend to be more negative in states with higher shares of black people. Specifically, for every percentage point increase in the share of blacks, ratings of Obama are expected to drop by 2/10 of a point.

Explained Variance Addition of the level-2 covariate has reduced the size of the level-2 variance component, which is now 7.27. The variance component remains statistically significant: the LR test statistic is 5.37, with a halved p-value of .005. It is thus clear that the percentage of blacks in a state does not fully account for all of the variation in the intercepts but, due to the reduced variance component, we can say that it accounts for some.

To obtain a better sense of the variance that is explained by the percentage of blacks we compute a level-2 R^2 measure, which is given by

$$R_2^2 = 1 - \frac{\tau_{00}^{New}}{\tau_{00}^{Old}}$$

Here, τ_{00}^{Old} is the level-2 variance component from the random intercept model, which is estimated as 11.19 (see Figure 2). Further, τ_{00}^{New} is the level-2 variance component from the model with the level-2 covariate and random intercept; this is estimated as 7.27 (see Figure 4). Thus, $R_2^2 = 1 - (7.27/11.19) = .35$. This means that the percentage of blacks in a state accounts for around 35 percent of the variation in the intercepts.

Random Slope and Intercept Model

We now add a random slope for the effect of black on evaluations of Obama. We can represent the resulting model through the following system of equations:

$$\begin{aligned} obamafeel_{ij} &= \beta_{0j} + \beta_{1j} pid_{ij} + \beta_{2j} iddist_{ij} + \beta_{3j} black_{ij} + \epsilon_{ij} \\ \beta_{0j} &= \gamma_{00} + \gamma_{01} pctblack_j + \delta_{0j} \\ \beta_{1j} &= \gamma_{10} \\ \beta_{2j} &= \gamma_{20} \\ \beta_{3j} &= \gamma_{30} + \delta_{3j} \end{aligned}$$

The model can be estimated using the following syntax:

```
xtmixed obamafeel pid iddist black pctblack if flagmis==0,
|| state:  black, covariance(indep) var
```

The request for an independent covariance matrix (`covariance(indep)`) suppresses the estimation of a covariance between the random intercept and slopes.[5]

[5] In this particular example, the covariance looks reasonable when we estimate it by specifying the option `covariance(unstruc)` in Stata, specifically, the estimate is -17.47, which translates

```
Mixed-effects ML regression                     Number of obs      =      1673
Group variable: state                           Number of groups   =        33

                                                Obs per group: min =        11
                                                               avg =      50.7
                                                               max =       258


                                                Wald chi2(4)       =   1230.33
Log likelihood = -7472.5017                     Prob > chi2        =    0.0000

------------------------------------------------------------------------------
   obamafeel |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         pid |  -6.929157   .2956243   -23.44   0.000     -7.50857   -6.349744
      iddist |   1.332853   .2144737     6.21   0.000      .912492    1.753213
       black |   13.48399   1.697982     7.94   0.000     10.15601    16.81198
    pctblack |  -.2920022   .0880059    -3.32   0.001    -.4644907   -.1195138
       _cons |   75.97266   1.791557    42.41   0.000     72.46128    79.48405
------------------------------------------------------------------------------

------------------------------------------------------------------------------
  Random-effects Parameters  |   Estimate   Std. Err.     [95% Conf. Interval]
-----------------------------+------------------------------------------------
state: Independent           |
                  var(black) |   15.52347   14.07504      2.625485    91.78424
                  var(_cons) |   7.156471   5.036224      1.801702    28.42595
-----------------------------+------------------------------------------------
               var(Residual) |   436.6301   15.36588      407.5287    467.8096
------------------------------------------------------------------------------
LR test vs. linear regression:       chi2(2) =     7.58   Prob > chi2 = 0.0226
```

Figure 5: Stata `xtmixed` **Output for a Hierarchical Linear Model with a Random Intercept and Slope**

The estimation results are shown in Figure 5.

Testing the Additional Variance and Covariance Components Looking at the Figure, we see a sizable variance component on race, suggesting that its effect varies across states. To see if we really need the variance component, we can perform a LR test comparing the new model to that from Figure 4. The latter model, which had a random intercept only, yielded a log-likelihood of -7473.607; this translates into a deviance of 14947.214.[6] The new model has a log-likelihood of -7472.502, corresponding to a deviance of 14945.003. The reduction in the deviance is thus 2.211, which is the LR test statistic. To obtain the p-value, we issue the command

```
di chi2tail(1,2,211)
```

which yields $p = .137$. When halved, this produces .069, which means that the variance component on the slope of race is only marginally significant. It is a

into a correlation of -.86. The reason we decided not to include the covariance is that its estimate turns into a corner solution (-1) once we include the cross-level interaction. To keep consistency across models, we therefore decided to omit the covariance. This is also justified by the fact that it is not statistically significant.

[6] Recall that the deviance is -2 times the log-likelihood.

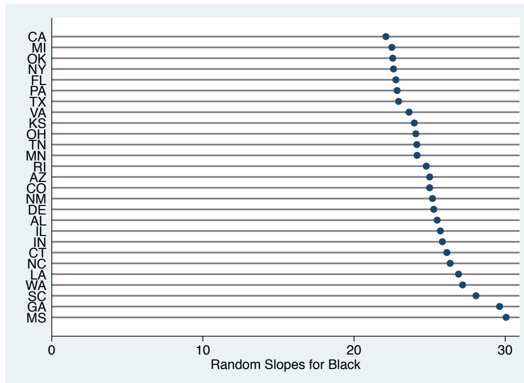

Figure 6: Empirical Bayes Estimates of the Random Slopes Across States

judgment call whether to conclude if this p is sufficiently small to conclude that the slope of race varies. One reason to assume it is that with $J = 33$ level-2 units, we severely lack in statistical power and might be willing to settle for a slightly larger type-I error rate.

Empirical Bayes Estimates To obtain a sense of the variation in the effect of race, we again compute the empirical Bayes estimates. The syntax for doing so is analogous to what we did for the random intercept model. The results are shown in Figure 6. As this figure reveals, the smallest effect occurs in California, where the gap between blacks and whites in terms of their evaluations of Obama is only about 22 points. The largest effect occurs in Mississippi, where the gap between blacks and whites is around 30 points. While perhaps not overwhelming, these differences are still noteworthy.

A Hierarchical Linear Model with a Cross-Level Interaction

We conclude our analysis of Obama feeling thermometer ratings by considering a possible explanation for the random slope on race. Our hypothesis is that the racial divide in the ratings depends on the percentage of blacks in the states. In states with a large black population, we expect the divide to be greater than in states were this is not the case. One reason could be that the perceived threat of a black presidential candidate to whites is greater when there is a sizable black population in

```
Mixed-effects ML regression                     Number of obs      =      1673
Group variable: state                           Number of groups   =        33

                                                Obs per group: min =        11
                                                               avg =      50.7
                                                               max =       258

                                                Wald chi2(5)       =   1433.90
Log likelihood = -7464.9839                     Prob > chi2        =    0.0000

------------------------------------------------------------------------------
   obamafeel |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         pid |  -6.876589   .2948488   -23.32   0.000    -7.454482   -6.298696
      iddist |    1.29745   .2129829     6.09   0.000     .8800115    1.714889
       black |   4.034204   2.592524     1.56   0.120    -1.047049    9.115457
    pctblack |  -.4300338   .0987236    -4.36   0.000    -.6235285   -.2365392
       inter |   .5742413   .1378223     4.17   0.000     .3041146    .8443679
       _cons |   77.59071   1.864426    41.62   0.000      73.9365    81.24491
------------------------------------------------------------------------------

------------------------------------------------------------------------------
  Random-effects Parameters  |   Estimate   Std. Err.     [95% Conf. Interval]
-----------------------------+------------------------------------------------
state: Independent           |
                  var(black) |   4.99e-15   3.89e-14      1.17e-21    2.13e-08
                  var(_cons) |   7.961062   5.045215      2.298971     27.5682
-----------------------------+------------------------------------------------
               var(Residual) |   434.7776    15.1956       405.992    465.6042
------------------------------------------------------------------------------
LR test vs. linear regression:       chi2(2) =     6.16   Prob > chi2 = 0.0459
```

Figure 7: Stata `xtmixed` Output for a Hierarchical Linear Model with a Cross-Level Interaction

the state. To bring this explanation into the model, we modify the random intercept and slope model by formulating the following level-2 model for the effect of race: $\beta_{3j} = \gamma_{30} + \gamma_{31} pctblack_j + \delta_{3j}$. This produces the following mixed model:

$$\begin{aligned} obamafeel_{ij} &= \gamma_{00} + \gamma_{01} pctblack_j + \gamma_{10} pid_{ij} + \gamma_{20} iddist_{ij} + \\ &\quad \gamma_{30} black_{ij} + \gamma_{31} pctblack_j \times black_{ij} + \\ &\quad \delta_{0j} + \delta_{1j} black_{ij} + \epsilon_{ij} \end{aligned}$$

Here $pctblack_j \times black_{ij}$ is the cross-level interaction term. In our analysis, this interaction is called `inter`. It can simply be added to the list of predictors in the Stata syntax for the random slope and intercept model. The model estimates are shown in Figure 7.

Explained Variance We see that the variance component on black has gone practically to zero. This means that percentage blacks in the state explains 100 percent of the variance in the slopes for race. This may seem like an unreasonable amount, but level-2 R^2 estimates like these are not unheard of in multilevel analysis.

Simple Slope Estimates Ignoring, i.e., averaging over, the level-1 and level-2 error components, the mean of the outcome variable is given by

$$\begin{aligned}\mu_{ij} &= \gamma_{00} + \gamma_{01} pctblack_j + \gamma_{10} pid_{ij} + \gamma_{20} iddist_{ij} + \\ &\quad \gamma_{30} black_{ij} + \gamma_{31} pctblack_j \times black_{ij}\end{aligned}$$

If we are interested in the partial effect of black on this mean, then we compute

$$\frac{\partial \mu_{ij}}{\partial black_{ij}} = \gamma_{30} + \gamma_{31} pctblack_j$$

This is the simple slope equation for race.

We can assess the impact of percentage black on the effect of race in a variety of ways. First, we can set some values for the percentage of blacks, say the minimum, median, and maximum. We can then use Stata's `lincom` command to assess the slope of race at these values. This is done in Figure 8. We see that the predicted race gap in Obama evaluation ranges from 4.7, which is not quite significant, when the percentage black is at the minimum (of 1.2) to 25.3 when percentage black is at its maximum (of 37).

We can also plot the simple slope over the entire range of percentage black and add a 95% confidence band. To compute this band, we need to compute the variance of the simple slope, which is

$$\begin{aligned}\widehat{Var}[\hat{\gamma}_{30} + \hat{\gamma}_{31} pctblack_j] &= \widehat{Var}[\hat{\gamma}_{30}] + 2 pctblack_j \widehat{Cov}[\hat{\gamma}_{30}, \hat{\gamma}_{31}] + \\ &\quad pctblack_j^2 \widehat{Var}[\hat{\gamma}_{31}]\end{aligned}$$

The variances and covariances of the estimators can be obtained by giving the

```
estat vce
```

command. This yields $\widehat{Var}[\hat{\gamma}_{30}] = 6.721$, $\widehat{Var}[\hat{\gamma}_{31}] = 0.019$, and $\widehat{Cov}[\hat{\gamma}_{30}, \hat{\gamma}_{31}] = -.298$. The 95% confidence interval can then be computed as

$$\hat{\gamma}_{30} + \hat{\gamma}_{31} pctblack_j \pm 1.96 \sqrt{\widehat{Var}[\hat{\gamma}_{30} + \hat{\gamma}_{31} pctblack_j]}$$

The resulting graph is shown in Figure 9.

```
. lincom black+1.2*inter

 ( 1)  [obamafeel]black + 1.2*[obamafeel]inter = 0
```

obamafeel	Coef.	Std. Err.	z	P>|z|	[95% Conf. Interval]	
(1)	4.723294	2.456248	1.92	0.054	-.0908633	9.537451

```
. lincom black+11.8*inter

 ( 1)  [obamafeel]black + 11.8*[obamafeel]inter = 0
```

obamafeel	Coef.	Std. Err.	z	P>|z|	[95% Conf. Interval]	
(1)	10.81025	1.52692	7.08	0.000	7.817544	13.80296

```
. lincom black+37*inter

 ( 1)  [obamafeel]black + 37*[obamafeel]inter = 0
```

obamafeel	Coef.	Std. Err.	z	P>|z|	[95% Conf. Interval]	
(1)	25.28113	3.266159	7.74	0.000	18.87958	31.68269

Figure 8: The Impact of Race on Obama Evaluations at Different Values of Percentage Black

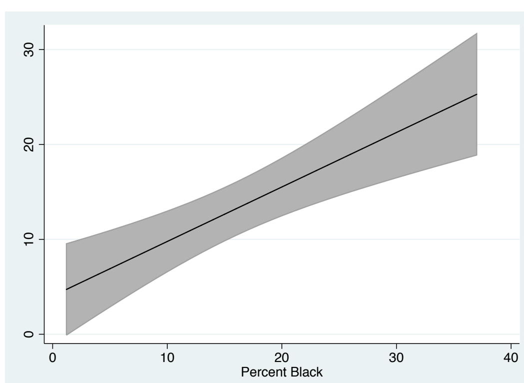

Figure 9: The Simple Slope for Black on Obama Evaluations