# POL 681 Lecture Notes: Statistical Interactions

## 1 Preliminaries

To this point, the linear models we have considered have all been interpreted in terms of "additive" relationships. That is, the relationship between two covariates, $X_1$ and $X_2$ in the context of the model

$$\hat{Y} = \hat{\beta}_0 + \hat{\beta}_1 X_1 + \hat{\beta}_2 X_2,$$

is additive. This implies that the partial effect of each independent variable is the same, regardless of the specific value at which the other covariate is held constant.

To fix some ideas, suppose we estimated a model and obtained the following coefficients:

$$\hat{Y} = 88.81 - .766(X_1),$$

where $X_1$ is a quantitative covariate. The predicted regression line would look like that shown in Figure 1. The interpretation of the model is standard: for a 1 unit increase in $X_1$, the expected value of $Y$ decreases by about .766.

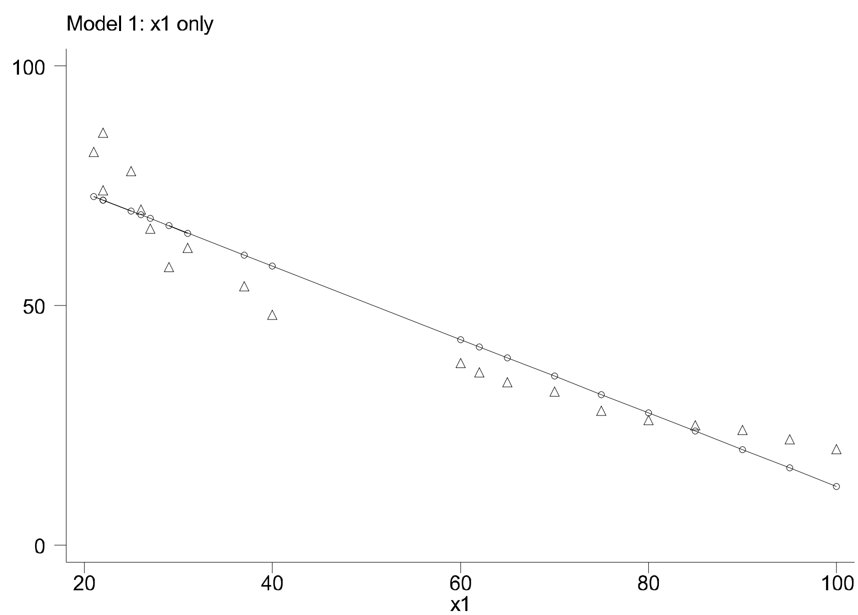


Figure 1: *Predicted Values of $Y$ given $X_1$.*

Now, suppose that we believed there were two distinct subpopulations that were theoretically related to $Y$. To account for these two groups, we create a dummy variable called $D_1$ such that $D_1 = 1$ denotes the first group and $D_1 = 0$ denotes the second group. We estimate a model and obtain,

$$\hat{Y} = 86.43 + -.665(X_1) - 5.91(D_1),$$

where the coefficient for $D_1$ is the partial regression coefficient for the dummy variable. This is interpreted as giving us the expected change in $Y$ when $D_1 = 1$; that is, for group 1, the expected value of $Y$ decreases by about -5.91. Graphically, this model would give us the results shown in Figure 2.

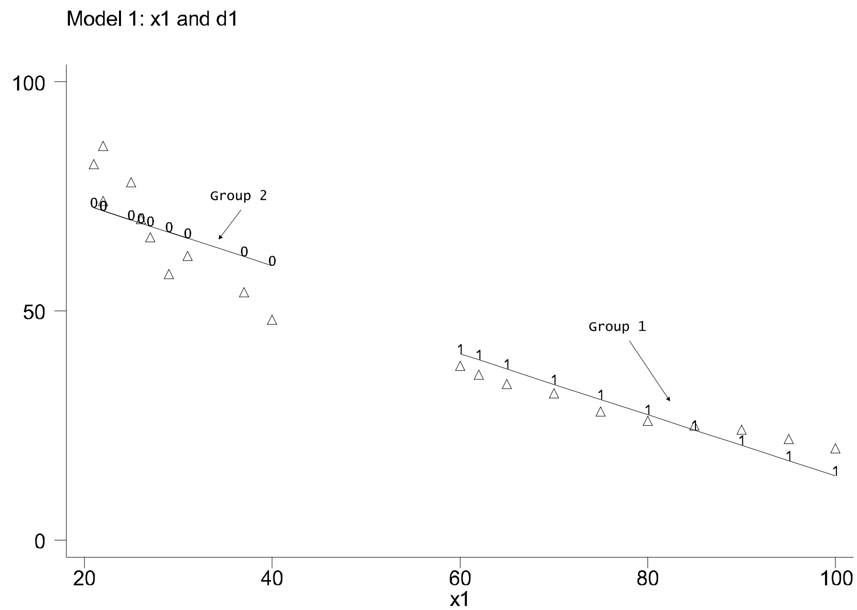


Figure 2: *Predicted Values of $Y$ given $X_1$ and $D_1$. Note the parallel slopes that occur by computing the separate regression lines for each group.*

The idea of "additivity" is important here: our interpretation of $X_1$ does not change as the value of $D_1$ changes. This is what gives rise to the parallel slopes that you see in Figure 2. Similarly, our interpretation of $D_1$ does not change as a function of $X_1$. The relationship between $D_1$ and $Y$ is accounted for by the offset in the two parallel slopes (which is equal to -5.91). The "effects" are additive because to understand the expected value of $Y$ for both variables, we simply "add" one coefficient to the other.

Hence, the expected value of $Y$ given $D_1 = 1$ and $X_1 = 53.1$ (its mean) is $\hat{Y} = 86.43 + -.665(53.1) - 5.91(1)$ which is 45.21. Note that to obtain this number, we simply added the two coefficients together (after multiplying them by 53.1 and 1 respectively). This is illustrative of the additivity property of the linear regression model. This property ensures parallel slopes when one of the regressors is dummy variable (the slope will be offset by the value of the dummy variable, as illustrated in Figure 2).

However, what if the additivity property does not hold? This would imply that the effect of some covariate on the dependent variable would be influenced by changes in the value of another covariate; that is, the interpretation of, say, $X_1$ on $Y$ would be *conditional* on some other covariate, say $D_1$.

To begin to see this, suppose we reestimated our regression model, but this time, estimated a separate model for observations in group 1 and observations in group 2. For group

1, our model gives us

$$\hat{Y} = 62.11 - .43(X_1),$$

and for group 2, our model gives us

$$\hat{Y} = 118.35 - 1.81(X_1).$$

Note the differences here. For group 1, the slope coefficient tells us that for a unit change in $X_1$, the expected change in $Y$ is -.43; for group 2, that expected change is -1.81 (a little over 4 times that of group 1).

Note that if the additivity property held, the relationship between $X_1$ and $Y$ would be nearly the same for the two groups. To see this, return to our definition of additivity: the partial effect of each independent variable is the same, regardless of the specific value at which the other covariate is held constant. In the "two models" approach, we are essentially holding $D_1$ constant and we can easily see that our interpretation of $X_1$ drastically changes, depending on which value $D_1$ assumes. This is suggestive that the additivity property may not hold. If we graphed the results from the two models above, we would get the results shown in Figure 3.

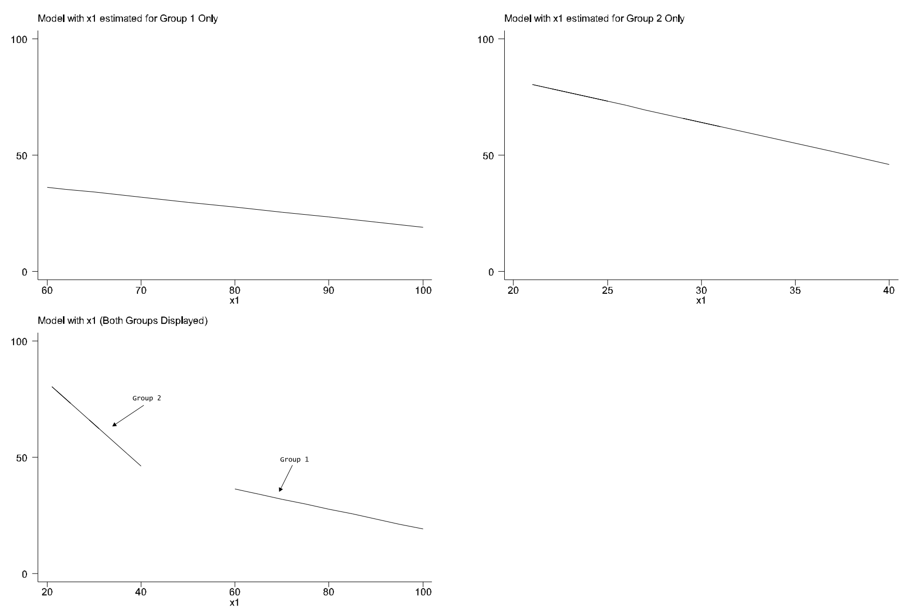


Figure 3: *Predicted Values of Y given $X_1$. The first panel is the predicted regression function based on the observations for group 1; the second panel is the regression function based on the observations for group 2; panel 3 gives both functions. Note the non-parallel slopes that occur by computing the regression lines based on separate models.*

The point to note is how the slopes change depending on which group we are looking at. The slope for group 1 is considerably flatter than that for group 2 (which is consistent with our separate regression models). In combining the two graphs, this point is made even clearer: the slopes are not parallel and the interpretation of $X_1$ *is not the same* for each group.

This suggests that the relationship between $X_1$ and $Y$ may be conditional on $D_1$; that is, $D_1$ may *moderate* the relationship.

To see this, suppose we estimated a model treating $X_1$ as a function of $D_1$. In doing this, we obtain

$$\hat{X}_1 = 28 + 50.2(D_1),$$

where the coefficients are interpreted as follows: the expected value of $X$ for group 1 is 78.2; the expected value of $X$ for group 2 is 28. Here we see that the average level of $X_1$ is much higher for group 1 than group 2. Hence, it seems that $X_1$ is related to $D_1$. In terms of our additivity assumption, it will *not* be the case that we can correctly interpret $X_1$ by holding $D_1$ constant: the partial effect changes, depending on which group we're looking at.

This suggests that the relationship between $D_1$, $X_1$ and $Y$ is *multiplicative.* That is, the effect of $X_1$ on $Y$ may be multiplicatively higher in one group than when compared to the second group. This leads to consideration of an *interactive model*, which has the form

$$\hat{Y} = \hat{\beta}_0 + \hat{\beta}_1 D_1 + \hat{\beta}_2 X_1 + \hat{\beta}_3(D_1 X_1),$$

where the last term denotes the interaction term. It is easy to see where the model gets its name: we are interacting, multiplicatively, $X_1$ with $D_1$. Note what the interaction term is giving you. This term is 0 for group 2 (why?) and nonzero for group 1; hence, you're getting an estimate of the slope of $X_1$ for $D_1$. For group 2, the relationship between $X_1$ is simply given by the coefficient for $X_1$. Thus

$$\hat{Y}_{G1} = \hat{\beta}_0 + \hat{\beta}_1 D_1 + (\hat{\beta}_2 + \hat{\beta}_3)X_1,$$

gives us the regression function for Group 1 and

$$\hat{Y}_{G2} = \hat{\beta}_0 + \hat{\beta}_2 X_1,$$

gives us the regression function for Group 2.

From this model, we will not only get differences in intercepts (due to the offset between the two groups), but we will *also* get differences in the slope. For Group 1, the slope is $\hat{\beta}_2 + \hat{\beta}_3$ and for Group 2, the slope is $\hat{\beta}_2$.

Using the data we've been working with, I estimate the interactive model and obtain

$$\hat{Y} = 118.35 - 56.24D_1 - 1.81X_1 + 1.38D_1X_1.$$

This model is *not* interpretable as an additive model (though it is a linear model). The reason is that the impact of $X_1$ on $Y$ is conditional on which group you are looking at. This raises an important point. Sometimes analysts will refer to the coefficient for $X_1$ as a "main effect." That is, sometimes the interpretation is given that the coefficient for $X_1$ represents the constant effect of $X_1$. It does not.

The terms "interaction" and "main effects" were adopted from the ANOVA method. In the context of ANOVA, an "interaction" refers to the effect of a factor averaged over another factor, and the main effect represents the average effect of a single variable.

In the case of multiple regression, the "main effect" terminology is not justified although it is commonly assumed that the coefficients for $D_1$ and $X_1$ are "main effects." In the presence of an interaction, these coefficients in no instance represent a constant effect of the independent variable on the dependent variable. The use of the term main effect implies that the coefficients are somehow interpretable alone, when they actually represent a portion of the effect of the corresponding variable on the dependent variable.

To see this, let's back out the predicted regression function for each of the two groups. For Group 1, the regression function is given by

$$\hat{Y} = 118.35 - 56.24D_1 + (-1.81 + 1.38)X_1,$$

and for Group 2, the regression function is given by

$$\hat{Y} = 118.35 - 1.81X_1.$$

If we can compute these functions, we can graph them and I do so in Figure 4. Here we see nonparallel slopes. Further, it should be clear that the results we obtain using the interaction term *are identical to the results obtained using the separate regressions approach* (refer back to Figure 3). This is not coincidental: in the interactive model, we are conditioning the relationship of $X_1$ on $D_1$. This, in effect, gives us separate models within a single model. This is the intuition of an interactive model.

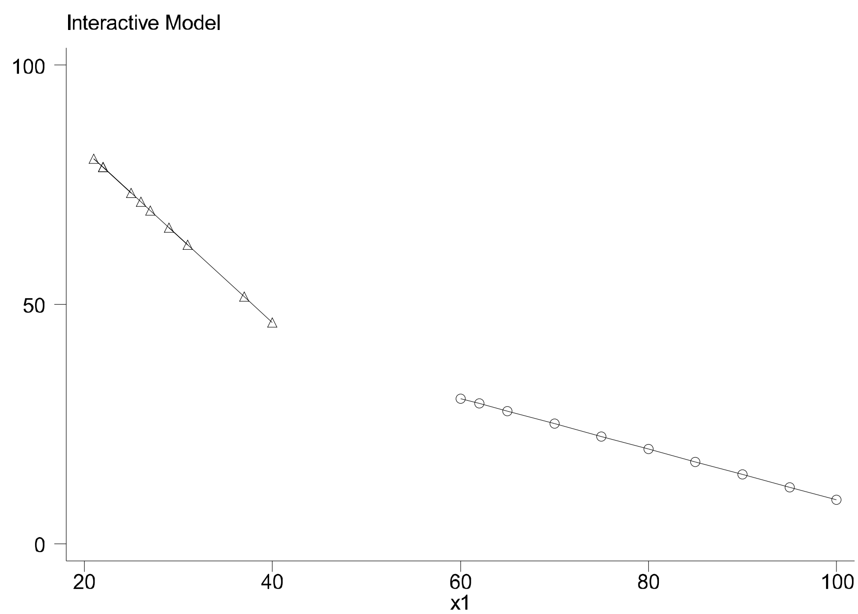


Figure 4: *Predicted Regression Functions for Groups 1 and 2 from Interactive Model.*

Let me point out one last thing. Recall in the model where we treated $X_1$ as a function of $D_1$. There, we got the expected value of $Y$ for each group. Let me redraw Figure 4, but this time inserting a vertical reference line at the mean of $X_1$ for each group. This is shown in Figure 5.

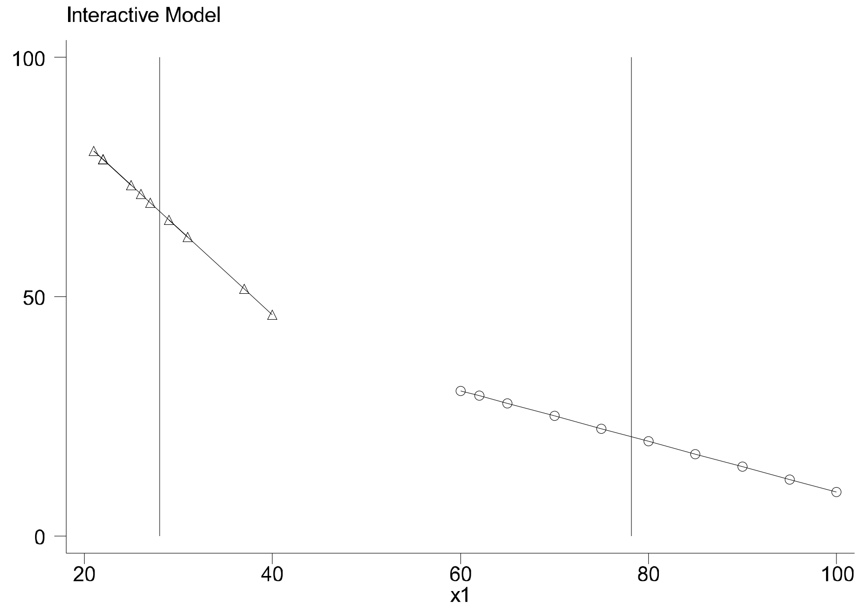


Figure 5: *Predicted Regression Functions for Groups 1 and 2 from Interactive Model with Vertical Reference Line Drawn through $X_1$ means for each group.*

Note that the vertical reference line goes through the points 28, which corresponds to the mean of $X_1$ for Group 2 and 78.2, which corresponds to the mean of $X_1$ for Group 1. This figure illustrates the point made in the regression of $X_1$ on $D_1$: the separate slopes pass through the group means, which are equivalent to the regression estimates from this submodel.

# 2 Extensions

Let's reset our ideas about interactions and further illustrate why the interpretation in terms of "main effects" is not appropriate.

Suppose we posit the following model

$$\hat{Y} = \hat{\beta}_0 + \hat{\beta}_1 D_1 + \hat{\beta}_2 X_1 + \hat{\beta}_3 D_1 X_1,$$

where the last term is an interaction term between the covariates $X_1$ and $D_1$. Suppose $X_1$ is a quantitative variable and $D_1$ is a dummy variable. As noted above, the relationship between $X_1$ and $Y$ is conditional, in the interactive model, on $D_1$. This gives rise to two submodels: one for group 1 and one for group 2. For group 1, the regression function is given by

$$\begin{aligned}\hat{Y} &= \hat{\beta}_0 + \hat{\beta}_1 1 + \hat{\beta}_2 X_1 + \hat{\beta}_3 1 X_1 \\ &= \hat{\beta}_0 + \hat{\beta}_1 + (\hat{\beta}_2 + \hat{\beta}_3) X_1.\end{aligned}$$

And for group 2, the regression function is given by

$$\begin{aligned}\hat{Y} &= \hat{\beta}_0 + \hat{\beta}_1 0 + \hat{\beta}_2 X_1 + \hat{\beta}_3 0 X_1 \\ &= \hat{\beta}_0 + \hat{\beta}_2 X_1.\end{aligned}$$

Hence, for group 1, the slope is given by $(\hat{\beta}_2 + \hat{\beta}_3)$ and the intercept is given by $(\hat{\beta}_0 + \hat{\beta}_1)$, whereas for group 2, the slope is given by $\hat{\beta}_2$ and the intercept is given by $\hat{\beta}_0$. In no case is the "effect" between $X_1$ and $Y$ constant across groups; hence there is no main effect.

This is all easy to see in the case of a dummy variable (and even easier to see if you interact two dummy variables together); however, what about the case with two quantitative variables?

Suppose we are interested in the following model,

$$\hat{Y} = \hat{\beta}_0 + \hat{\beta}_1 X_1 + \hat{\beta}_2 X_2 + \hat{\beta}_3 X_1 X_2,$$

where the two covariates are quantitative variables and the last term denotes the interaction between them. In estimating this kind of model, our intuition is that the slope of $X_1$ on $Y$ is conditional on $X_2$ *and* the slope of $X_2$ on $Y$ is conditional on $X_1$. Since there are no main effects, the relationship of one independent variable on the dependent variable is conditional on the other independent variable. Thus, we believe that the relationship is nonadditive.

The question arises, naturally, as to how to interpret this model? At first glance, it seems to pose no special problems, and in general, it doesn't. In order to derive the conditional regression function for $Y$ and $X_1$ conditional on some specific value of $X_2$ we obtain

$$\hat{Y} = \hat{\beta}_0 + \hat{\beta}_2 X_2 + (\hat{\beta}_1 + \hat{\beta}_3 X_2) X_1;$$

and to obtain the conditional regression function for $Y$ and $X_2$ conditional on some specific value of $X_1$, we obtain

$$\hat{Y} = \hat{\beta}_0 + \hat{\beta}_1 X_1 + (\hat{\beta}_2 + \hat{\beta}_3 X_1) X_2.$$

In these models, the intercept gives us the expected value of $Y$ when all the covariates are equal to 0. In the first model, the coefficient for $\hat{\beta}_3$ gives us the change in the slope of $Y$ on $X_1$ associated with a unit change in $X_2$; in the second model, $\hat{\beta}_3$ gives us the change in the slope of $Y$ on $X_2$ associated with a unit change in $X_1$.

The key thing to see here is that the slope between one covariate is governed by the other covariate. To evaluate the impact one covariate has on the other, it is useful to note that the coefficients $\hat{\beta}_1$ and $\hat{\beta}_2$ denote the baselines around which the slopes vary. This is easy to see.

Suppose that $X_2 = 0$, then the first model above reduces to

$$\begin{aligned}\hat{Y} &= \hat{\beta}_0 + \hat{\beta}_2 0 + (\hat{\beta}_1 + \hat{\beta}_3 0) X_1 \\ &= \hat{\beta}_0 + \hat{\beta}_1 X_1.\end{aligned}$$

Suppose that $X_1 = 0$, then the second model above reduces to

$$\begin{aligned}\hat{Y} &= \hat{\beta}_0 + \hat{\beta}_1 0 + (\hat{\beta}_2 + \hat{\beta}_3 0) X_2 \\ &= \hat{\beta}_0 + \hat{\beta}_2 X_2.\end{aligned}$$

In both cases, we can see that the models represent the "baseline" case from which the conditional slope coefficients can be evaluated. In the first model, the coefficients $\hat{\beta}_0$ and $\hat{\beta}_1$ give the intercept and slope for the regression of $Y$ on $X_1$ *when* $X_2$ *is equal to 0.* In the second model, the coefficients $\hat{\beta}_0$ and $\hat{\beta}_2$ give the intercept and slope for the regression of $Y$ on $X_2$ *when* $X_1$ *is equal to 0.*

Through this, we again can see the conditional nature of the regression model when interactions are applied. In the standard least squares model without interactions, the coefficients for the covariates tell us something much different than they do in the interactive case. In the interactive case, the coefficients are conditional on the value of some covariate; hence, the effect of one covariate, say $X_1$, on $Y$ is not constant across the full range of values taken by the other covariate, $X_2$.

*Because of this conditional property, coefficient estimates from an interactive model will in general be much different than estimates derived from a model possessing the additivity property.*

The previous results assumed that the covariates have a meaningful 0 point. Having a 0 point was convenient because it gave the coefficients a natural interpretation: the baseline against which changes in the conditional slope could be evaluated. Of course the question naturally arises regarding the absence of a 0 on the covariates.

This is an important question and to understand why, focus on the submodels presented above. In the absence of a meaningful 0 term, the baseline "effects" are never observed in the data. Hence, the coefficient estimate (which is interpreted in terms of the covariate having a 0 point) and its standard error can be misleading, sometimes *very* misleading, because it is giving you the conditional relationship under a condition that doesn't actually hold in the data.

This is why careful interpretation of one's results is extremely important with interaction terms. The conditional slopes are still meaningful in the absence of a natural 0 point; however, the slopes are only interpretable or meaningful within the observed range of the data. Note that if you treat the model *as if* the covariates have a natural 0 point when they do not have a 0 point, then you are essentially extrapolating beyond the range of the data. This is generally a poor strategy to follow because you are drawing conclusions based on data points that actually don't exist (and may never exist in the real world).

So how do you interpret the interaction? Easy. You just follow the approach I outlined above, except you only compute the conditional slopes for observed values of the covariates.

Let's go through an example. Suppose I posit an interactive model with two quantitative covariates, $X_1$ and $X_2$. Using some simulated data, I get the following coefficient estimates:

$$\hat{Y} = 6.43 + .18X_1 + .99X_2 - .011X_1X_2.$$

To derive the conditional slope for $Y$ regressed on $X_1$ conditional on $X_2$, I rearrange the model and obtain

$$\hat{Y} = 6.43 + .99X_2 + (.18 - .011X_2)X_1.$$

To derive the conditional slope for $Y$ regressed on $X_2$ conditional on $X_1$, I rearrange the model and obtain

$$\hat{Y} = 6.43 + .18X_1 + (.99 - .011X_1)X_2.$$

The first two terms in each of the submodels, in theory, gives me the slope and intercept of the conditional slope for the case when the other covariate is 0; however, since neither covariate takes on the 0 value, I have to interpret the interaction in terms of the observed range of data.

For $X_1$, the range is 21 to 100; for $X_2$, the range is 40 to 100. One way to interpret the results is to compute the predicted regression function of the slope of $Y$ on $X_1$ for the maximum and minimum and mean values of $X_2$. I do this and show the results in Figure 6. The thing to note is the slope of the three lines, which represent hypothetical regression functions. These lines are created by substituting 40, 65, and 100 into the first model above and then allowing $X_1$ to freely vary. This gives us a sense of the conditional relationship and it tells us that as $X_2$ decreases, the slope between $Y$ and $X_1$ tends to decrease, that is, flatten out. To see this numerically, note that the estimated slope for $X_2 = 100$ is about -.92; for $X_2 = 40$, the estimated slope is about -.26. The dots in the graph represent the actual predicted values of $Y$ given the actual values of $X_1$ and $X_2$.

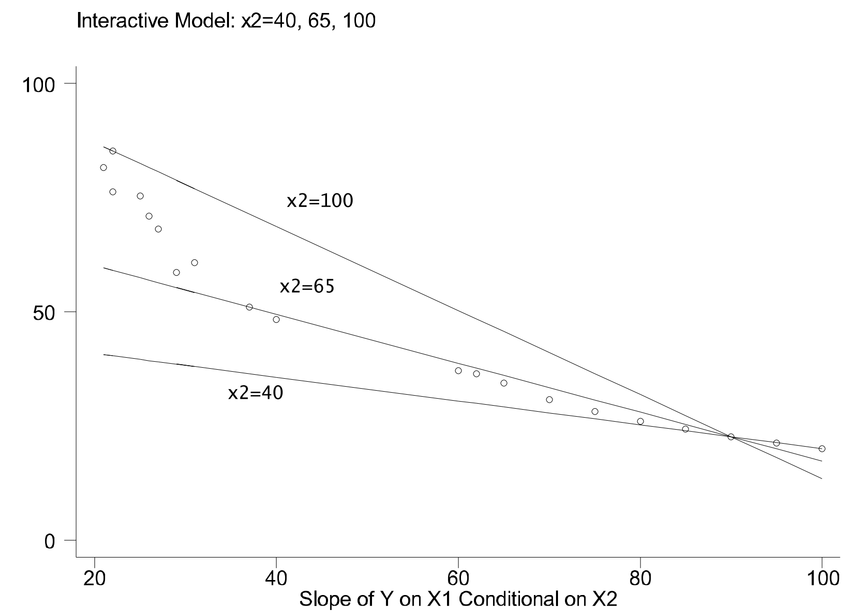


Figure 6: *Hypothetical Regression Function of $Y$ on $X_1$ conditional on $X_2$. The x-axis represents the variable $X_1$.*

I go through the same exercise for the slope of $Y$ on $X_2$ and present the results in Figure 7. The conclusion here is that as $X_1$ increases, the slope between $Y$ and $X_2$ tends to decrease (and indeed, it becomes negative for the upper bound on $X_1$. The range of the estimated slopes is -.11 (for $X_1 = 100$) to .76 (for $X_1 = 21$). Again, the dots represent the actual predicted values of $Y$ given the actual values of $X_1$ and $X_2$.

The point of this exercise is to note that inclusion of, and interpretation of interaction terms, even in the face of quantitative variables without a natural 0 point is straightforward. There are caveats: you need to be careful in interpreting the results. Graphs like the ones

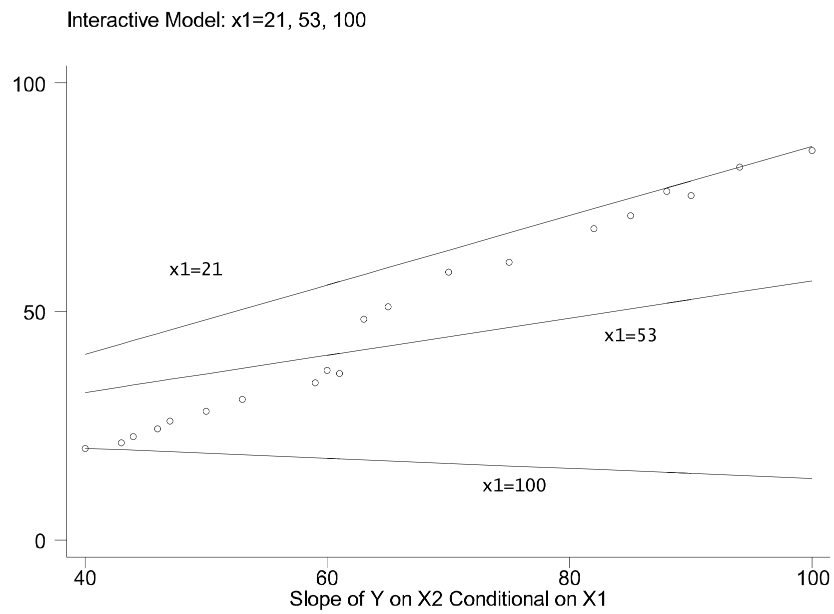


Figure 7: *Hypothetical Regression Function of $Y$ on $X_2$ conditional on $X_1$. The x-axis represents the variable $X_2$.*

shown above are problematic because many (most!) of the points on the predicted lines are not actual data points—they are hypothetical.

# 3 Uncertainty

If there is considerable confusion in the literature regarding the interpretation of interaction effects, there is even more confusion regarding the interpretation of the standard errors associated with interaction terms and their constituent terms.

It is not at all uncommon to estimate a statistical interaction and find that the standard errors, relative to the coefficient, are very large, and indeed, statistically insignificant. This insignificance can be misleading (although it won't always be misleading). The reason is most acute when one is using quantitative variables in interactive models.

Understand that the standard error reported to you on your `Stata` output—or any statistical software's output—is the standard error for the case when $X_1$ and $X_2$, respectively, are zero. (See Friedrich's article). That is, *the standard errors are only applicable to a particular range of the interactive effect.*

This is crucial to understand. Just as the relationship between $X_1$ and $Y$ in an interactive model is conditional on $X_2$, *so to are the standard errors.* Hence, the standard errors are **conditional** standard errors. This has the important implication that the interactive effect may be significant over a certain range of the data, and insignificant over another part of the range of the data.

This implies that you will need to compute the standard error for interactive terms, as

Stata will not do this for you. Hence, if you think interactions are important and want to make a big deal about them, then it is incumbent upon you to provide the estimates of uncertainty around the conditional slopes across ranges of your interaction term.

This is easily done, fortunately. The general formula to compute the standard errors of the conditional slopes is given by

$$s.e._{\hat{\beta}_1+\hat{\beta}_3} = \sqrt{\text{var}(\hat{\beta}_1) + X_2^2\text{var}(\hat{\beta}_3) + 2X_2\text{cov}(\hat{\beta}_1\hat{\beta}_3)},$$

for the slope of $Y$ on $X_1$ conditional on $X_2$; and

$$s.e._{\hat{\beta}_2+\hat{\beta}_3} = \sqrt{\text{var}(\hat{\beta}_2) + X_1^2\text{var}(\hat{\beta}_3) + 2X_1\text{cov}(\hat{\beta}_2\hat{\beta}_3)},$$

for the slope of $Y$ on $X_2$ conditional on $X_1$. (The proof of this result comes from the standard variance and covariance results for two random variables; see Friedrich, p. 810).

Note that this formula is not a constant; it is determined by the level of the covariate upon which the slope is conditioned. In the first case, this is $X_2$; in the second case, this is $X_1$.

Let's work through our example. Recall that we estimated the following model:

$$\hat{Y} = 6.43 + .18X_1 + .99X_2 - .011X_1X_2.$$

The conditional slope for $Y$ regressed on $X_1$ conditional on $X_2$ was given by

$$\hat{Y} = 6.43 + .99X_2 + (.18 - .011X_2)X_1,$$

and the conditional slope for $Y$ regressed on $X_2$ conditional on $X_1$ was given by

$$\hat{Y} = 6.43 + .18X_1 + (.99 - .011X_1)X_2.$$

The quantity of primary interest is the term in the last two equations in parentheses. This, as you recall, gives us the conditional slope (or what Friedrich calls the "metric effect."). What we want to do is compute the standard error for these conditional slopes, which will vary as $X_2$ (or $X_1$) varies.

Looking at the first submodel, suppose that $X_2$ is set at 100. The conditional slope is equal to -.92 (if you replicate this and use Stata estimates, this estimate will be about -.90 because of rounding differences). The question to ask is whether or not this conditional slope coefficient is different from 0. To answer this question, we need to compute the standard error. In order to compute the standard error, we need to back out of our coefficient estimates, the variances and covariances of the parameters. In Stata, this is easily obtained by typing one of two commands. The easiest is to type

```
vce
```

after the model is estimated. This will display the variance-covariance matrix for you; if you want the matrix saved for later reference, then type

```
matrix V =e(V)
```

and then type

```
matrix list V
```

to call the matrix. The latter approach gives more precise numbers.

In doing this for my example, I obtain the following variances and covariances that can be substituted into the formula above. For the slope of $Y$ on $X_1$ conditional on $X_2 = 100$, I obtain

$$s.e._{\hat{\beta}_1+\hat{\beta}_3} = \sqrt{.0034 + 100^2(.000002042) + 2(100)(-.000067)},$$

which gives me a standard error estimate of about .102. Given the conditional slope estimate for this case was about -.92, I see that the $t$ ratio is about -9, which would be significant at any level. In Figures 8 and 9, I've plotted the conditional slope estimates as well as the upper and lower 95 percent confidence limits around the estimate (which were obtained by calculating the standard error and multiplying it by the critical $t$ value on 16 degrees of freedom which is 2.12 [$n = 20$ in this example]).

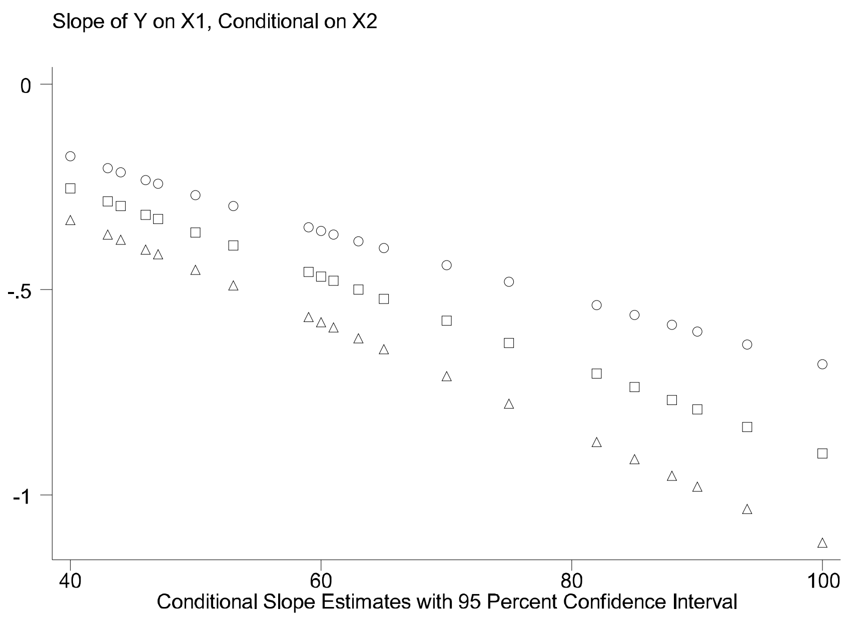


Figure 8: *Conditional Slopes with 95 percent confidence intervals. The x-axis represents the variable* $X_2$.

As can be seen, the slope of $Y$ on $X_1$ conditional on $X_2$ is in general more precisely estimated than the slope of $Y$ on $X_2$, conditional on $X_1$. The major point I want you to understand here is that if you are going to interpret these interactions, you need to be aware that the range of effects are not all going to be statistically different from 0.

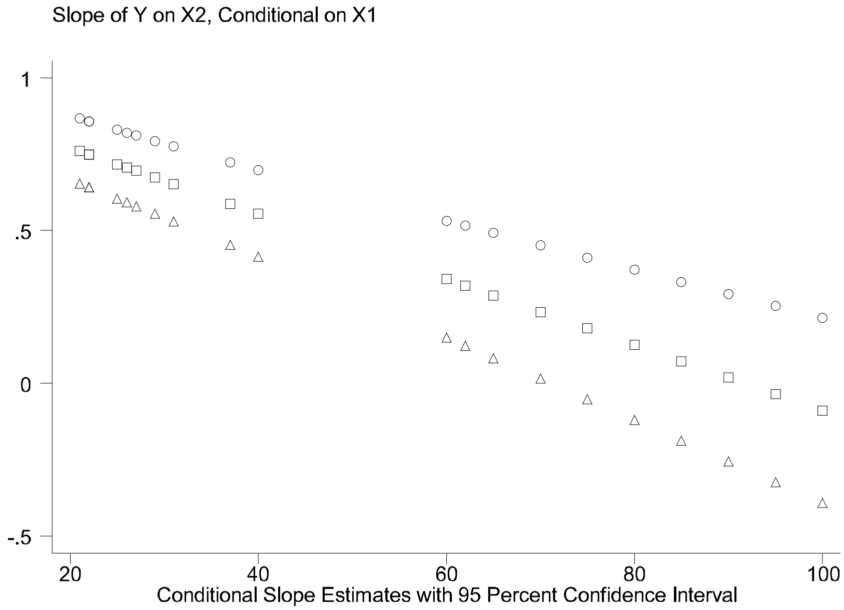


Figure 9: *Conditional Slopes with 95 percent confidence intervals. The x-axis represents the variable* $X_1$*.*

To illustrate these points further, we could compute the $t$ ratios for each of the predicted slope coefficients. This could be done easily be deriving the conditional slope and dividing it by its standard error. As these are $t$ ratios, any $t$ exceeding the critical $t$ would be statistically significant. For these data, I set $\alpha = .05$ and compute the two-tail critical $t$, which is 2.12. In Figures 10 and 11 I plot the absolute value of the $t$ ratios with respect to $X_2$ (in Figure 10) and $X_1$ (in Figure 11).

In Figure 10, I find that all of the estimated $t$ ratios exceed the critical $t$ of 2.12. This implies that the conditional slopes are statistically different from 0, and implies that across the full range of $X_2$, the interaction effect seems to be significant. This is in contrast to the results given in Figure 11. Here we see that once $X_1$ exceeds 75 (or so), the estimated $t$ ratios drop below the critical threshold (which is denoted by the horizontal reference line at 2.12). *For this range of data, the statistical interaction is not different from 0.* This is illustrative of the larger point that the statistical interaction need not be significant across the full range of the conditioning variable, in this case, $X_1$.

Of course had you not computed these standard errors, you would not have known this by simply looking at the `Stata` output. This tells you that to correctly employ an interactive model, you need to work harder than in the case of a simple additive model of interpreting your results.

## 3.1 A Brief Note on Centering

Recall my discussion about the problem of an absence of a natural 0 point. This can, in some cases, be overcome by mean-centering your covariates. This is easily done in `Stata`

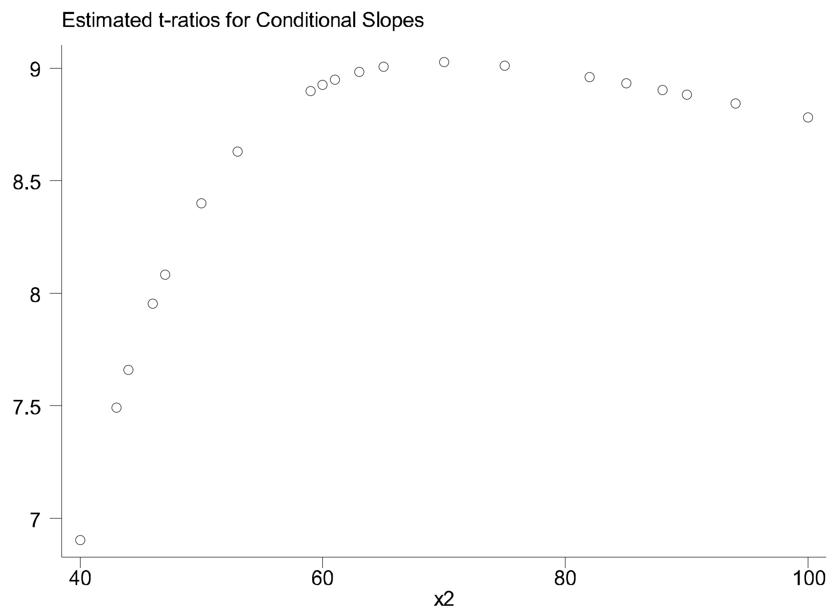


Figure 10: *Estimated $t$ ratios for the slope of $Y$ on $X_1$ conditional on $X_2$. The $x$ axis represents variable $X_2$. The critical $t$ is 2.12 (for $\alpha = .05$, two-tail). Note that all estimated ratios exceed the critical value.*

through the following command:

gen x1cent=x1-$\bar{X}_1$

where $\bar{X}$ denotes the mean of $X_1$. For values of $X$ that equal the mean, this new variable will assume the value of 0. Centering may or may not be helpful and centering *will not alter your substantive conclusions* because you are still working with the same quantities. The coefficients may change in value, and the intercept will always change, because the means of the variables are adjusted to 0.

If you want to verify this, go through the exercises just given to you, but first mean-center $X_1$ and $X_2$. Create an interaction of the new mean-centered variables and estimate the regression model. If you generate the quantities of interest, for example the quantities shown in the Figures, they will be identical to the results for the uncentered case. I'm not going to present these results, but you can easily generate them yourself.

# 4 Conclusion

If you can follow this presentation, you have a substantial advantage over workers who incorrectly interpret, or who don't interpret interactive relationships. The great failing is to not account for the uncertainty in parameter estimates and to fail to recognize the conditional nature of the "effects." Once you understand this, you'll see that interaction terms are not

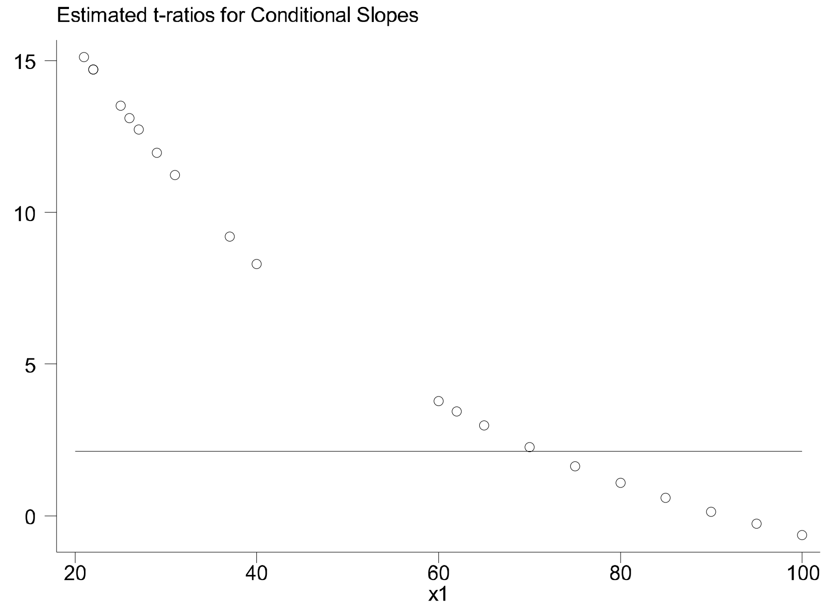


Figure 11: *Estimated $t$ ratios for the slope of $Y$ on $X_2$ conditional on $X_1$. The $x$ axis represents variable $X_1$. The critical $t$ is 2.12 (for $\alpha = .05$, two-tail). Note that not all estimated ratios exceed the critical value. When $X_1$ exceeds 75, the $t$ ratios drop below the critical threshold (given by the horizontal reference line). This implies that for this range of $X_1$, the interaction is not significantly different from 0.*

only easy to implement, but they reveal potentially interesting information. You just have to work harder to uncover that information.

Use of interactions can be extended to a wide range of applications. We will next turn to functional form and learn that interactions can be used profitably in using a linear model to obtain "nonlinear" estimates. We will thus consider segmented slope models and piecewise linear functions.