

Tidbits on Exploring Interaction Effects in Logistic Regression Analyses

UCSF Center for Aging in Diverse Communities, Analysis Core

Tor Neilands and Jennifer Toman

Investigators conducting aging-related health disparities research frequently need to specify statistical models containing one or more interaction terms. Such terms are one method for exploring whether an association between two variables differs depending on the level of a third variable. For example, does the association between race/ethnicity and dementia (or Alzheimer's) differ by level of SES? What follows are some observations regarding interactions that we have noted in the years we've provided statistical and methodological consultations and run statistical analyses for health disparities researchers who are conducting aging-related health disparities research.

The examples below are based on Hosmer and Lemeshow's low birth weight dataset featured in the Stata statistical analysis program's logistic regression command's help file. The dataset can be accessed within Stata by issuing the command `webuse lbw` from the Stata command line. We chose this dataset because it is readily available and contains both continuous and categorical predictors. We used Stata for the examples that follow because Stata is one of the most frequently used statistical computing programs at UCSF; however, the approaches covered in the examples can be implemented in other general purpose statistical programs. For instance, we supply links to online resources for conducting some of the methods described below using SAS.

We centered this document on logistic regression because we've noticed that logistic regression is one of the most popular modeling approaches used in the medical sciences and because interactions in logistic regression models are generally more challenging to unpack than are interactions in linear models. Although the demonstration dataset contains independent observations, the approaches described below can be readily extended to data obtained via complex survey sampling methods and to multilevel datasets – the estimation command syntax will change to meet the requirements of those types of datasets, but the model specification and post-estimation commands illustrated below should still apply regardless of the independence of observations.

Part 1: Two Categorical Predictors

Suppose an investigator is interested in fitting a statistical model to examine whether babies having low birthweights (`low`) is associated with the mother having smoked during pregnancy (`smoke`) and with whether the mother has hypertension (`ht`). Furthermore, the investigator is interested in whether there is an interaction between mother's hypertension status and smoking status. To address this question, the investigator envisions logistic model of the following form:

$$\text{logit}(\text{low}) = b * \text{constant} + b1 * \text{smoke} + b2 * \text{ht} + b3 * \text{smoke} * \text{ht}$$

Here `low` is the outcome, coded 1 if the baby had a low birthweight and 0 otherwise. Low birthweight is defined as the baby's weight being below < 2500 grams. The *logit* is the log of the odds of a mother having a low birthweight baby. Hypertension is coded as 0 if the mother did not have hypertension and 1 if the mother had hypertension. Smoking status is coded as 0 if the mother was a nonsmoker and 1 if the mother was a smoker (this variable's categories were already labeled in the Stata version of the dataset such that nonsmokers were labeled "nonsmoker" and smokers were labeled "smoker."). When this model is fitted in Stata, here are the model results using the default modeling syntax. The `i.` prefix instructs Stata to treat the predictor as a categorical variable. The double hash (`##`) syntax instructs Stata to include in the model smoking status and hypertension as well as their interaction (Stata refers to this way to specify models as *factor variable* notation). The `-base(line)-` option, which appears after the comma below, tells Stata to report base levels for factor variables and interactions.

```
. logistic low i.smoke##i.ht, base
```

Logistic regression

Number of obs = 189

LR chi2(3) = 9.14

Prob > chi2 = 0.0274

Pseudo R2 = 0.0390

Log likelihood = -112.76355

	low	Odds ratio	Std. err.	z	P> z	[95% conf. interval]	

smoke							
Nonsmoker		1	(base)				
Smoker		2.134286	.717116	2.26	0.024	1.104715	4.123393
ht							
0		1	(base)				
1		4.426666	3.528532	1.87	0.062	.928062	21.11429
smoke#ht							
Smoker#1		.5271082	.6519033	-0.52	0.605	.0466845	5.951507
_cons		.3012048	.0687171	-5.26	0.000	.1926053	.4710377

Note: _cons estimates baseline odds.

Simple Main Effects: We sometimes hear investigators refer to the constituent effects of an interaction term as “main effects.” By “constituent effects” we mean the separate predictor variable components that are used to create an interaction term. In our example, the constituent effects would be the 0/1-coded variables `smoke` and `ht`. We contend that the use of the term “main effect” to describe these effects can be confusing because “main effect” often has another meaning, which we will mention below when we discuss estimating the effects of one predictor averaged over levels of another predictor. More important than which term is being used to describe a given effect is an understanding of exactly what the term is estimating in the model.

In the results presented above, the effect of `smoke` (OR = 2.13) represents the ratio of the odds of a baby having a low birth weight if the mother is a smoker relative to the odds of a baby having a low birth weight if the mother is not a smoker **when the mother is not hypertensive** (i.e., `ht` = 0). Similarly, the effect of `ht` (OR = 4.43) represents the ratio of the odds of a baby having a low birth weight if the mother is hypertensive relative to the odds of a baby having a low birth weight if the mother is not hypertensive **when the mother is not a smoker** (i.e., `smoke` = 0). These effects are sometimes referred to in the statistical literature as *simple main effects* or *simple effects*. These effects are typically estimated following a statistically significant interaction. The interaction between `smoke` and `ht` in our model is not statistically significant at the conventional 5% level with a p-value of 0.605, so for practical purposes we would not include this interaction in a final analytic model. However, we will still estimate the simple main effects for demonstration purposes.

What about the corresponding simple effects of `smoke` when `ht=1` and `ht` when `smoke=1`? Those simple effects do not appear on the above output, but fortunately it is relatively straightforward to obtain them from the Stata `-contrast- post-estimation` command:

```
. contrast smoke@ht, effects nowald or
```

```
Contrasts of marginal linear predictions
```

```
Margins: asbalanced
```

	Odds ratio	Std. err.	z	P> z	[95% conf. interval]	
-----+-----						
smoke@ht						
(Smoker vs base) 0	2.134286	.717116	2.26	0.024	1.104715	4.123393
(Smoker vs base) 1	1.124999	1.339017	0.10	0.921	.109149	11.59537

```
. contrast ht@smoke, effects nowald or
```

```
Contrasts of marginal linear predictions
```

```
Margins: asbalanced
```

	Odds ratio	Std. err.	z	P> z	[95% conf. interval]	
-----+-----						
ht@smoke						
(1 vs base) Nonsmoker	4.426666	3.528532	1.87	0.062	.928062	21.11429
(1 vs base) Smoker	2.333332	2.206424	0.90	0.370	.3656547	14.88956

The `@` operator in the first `contrast` command requests estimation of the simple main effects of smoker vs. non-smoker at each of the two levels of hypertension. The `@` operator in the second `contrast` command requests estimation of the simple main effects of hypertension at each of the two levels of smoking status. The `effects` option requests that the estimated coefficient (that is, the log-odds) be printed; the `or` option requests that the exponentiation of the log-odds regression coefficient estimate b_j : $\exp(b)$, be printed and labeled as an odds ratio. The `nowald` option omits printing of multiple degrees-of-freedom Wald tests. Since each predictor involved in the interaction in this example has two levels, each has only $2-1 = 1$ degree-of-freedom, so the Wald test results would be redundant with the effects output. In applications involving interaction terms where one or more constituent variables have 3 or more levels, the Wald test output will provide non-redundant information so the `nowald` option should be dropped from the command statement in such instances.

Main Effects without the “Simple”: It is important to distinguish between simple main effects of the type presented above and another type of main effect frequently estimated and tested in models with interactions: averaged effects. Averaged effects are often referred to as “main effects” (without the “simple” qualifier). These main effects are the effect of one predictor that is a constituent component of an interaction term averaged over the levels of the other variable(s) used to form the interaction term. Here’s an illustration based on our low-birth-weight babies model:

```
. contrast smoke, effects or nowald
```

```
Contrasts of marginal linear predictions
```

```
Margins: asbalanced
```

	Odds ratio	Std. err.	Z	P> z	[95% conf. interval]
smoke (Smoker vs base)	1.549539	.9581993	0.71	0.479	.4611464 5.206742

```
. contrast ht, effects or nowald
```

```
Contrasts of marginal linear predictions
```

```
Margins: asbalanced
```

	Odds ratio	Std. err.	Z	P> z	[95% conf. interval]
ht (1 vs base)	3.213857	1.987376	1.89	0.059	.9564516 10.79917

As above when estimating simple main effects, we used the Stata contrast command. However, this time the syntax is different. We no longer list the two predictors forming the interaction separated by the @ sign. Instead, we specify a single predictor. As before, the `effects`, `or`, and `nowald` options estimate odds ratios and associated statistics and omit Wald tests. Here we see that the odds ratio for smokers vs. non-smokers averaged over hypertensive and non-hypertensive participants is estimated to be 1.55 whereas the odds ratio for hypertensive vs. non-hypertensive participants averaged over smokers and non-smokers is estimated to be 3.21.

We like using the term “main effects” to refer to the averaged effects estimated immediately above and the terms “simple main effects” or “simple effects” to refer to the estimation of an effect for a predictor within levels a second predictor when both predictors are constituents of an interaction term. As mentioned above, the specific terms used are less important than understanding the differences in what is being estimated using each approach and choosing the approach best suited to address one’s research questions. Confusion between these two types of effects has been discussed in the statistical literature (e.g., Jaccard, 2001, p. 20).

Relationships between the simple main effects and the interaction term: The simple main effects are also related to the odds ratio estimate for the interaction term. The interaction term is a ratio of two odds ratios. The two odds ratios are the simple main effects odds ratios for a predictor when its value is 1 versus when its value is 0 (Jaccard, 2001, pp. 20-21). Returning to the simple main effects output to estimate the odds of having a low-birth-weight baby for hypertensive vs. non-hypertensive mothers at each level of smoking status, we have:

```
. contrast ht@smoke, effects nowald or
Contrasts of marginal linear predictions
Margins: asbalanced
```

	Odds ratio	Std. err.	Z	P> z	[95% conf. interval]
-----+-----					
ht@smoke					
(1 vs base) Nonsmoker	4.426666	3.528532	1.87	0.062	.928062 21.11429
(1 vs base) Smoker	2.333332	2.206424	0.90	0.370	.3656547 14.88956

Dividing 2.333332 by 4.426666 yields .5271082, which is the odds ratio estimate for the interaction term shown in the initial logistic regression model results table. Within rounding error, the same estimate can be obtained from the other simple main effects output we examined comparing smoking status within each level of hypertension. Interpreting the ratio of two odds ratios can be difficult, so decomposing the interaction odds ratio into its component simple main effects can be helpful to facilitate a clear interpretation of the interaction term.

Part 2: One Categorical Predictor and One Continuous Predictor

We can apply the same principles from the discussion above to a scenario where there is one categorical predictor and one continuous predictor. Let’s revisit the Hosmer and Lemeshow dataset. This time let’s use logistic regression to model low birth weights as a function of mother’s race/ethnicity and mother’s age. Race/ethnicity is a three-level categorical variable coded 1 if the mother is White-identified, 2 if the mother is Black-identified, and 3 if the mother is Other-identified. In contrast, age is a continuous variable ranging from 14 years to 45 years. A *c* is placed in front of *age* in the model statement to designate age as continuous.

```
. logistic low c.age##i.race, base
```

```
Logistic regression                                Number of obs =    189
                                                    LR chi2(5)      =    8.09
                                                    Prob > chi2    = 0.1512
Log likelihood = -113.28973                        Pseudo R2     = 0.0345
```

low	Odds ratio	Std. err.	Z	P> z	[95% conf. interval]
-----+-----					
age	.9457412	.0435732	-1.21	0.226	.8640819 1.035118
race					
White	1 (base)				
Black	.2060127	.4310946	-0.75	0.450	.0034096 12.44763
Other	1.908774	3.258261	0.38	0.705	.0672605 54.16874
race#c.age					
Black	1.112336	.1026986	1.15	0.249	.9282117 1.332984
Other	.9955155	.0736227	-0.06	0.952	.861188 1.150795
_cons	1.190902	1.315057	0.16	0.874	.1367522 10.37093

Note: _cons estimates baseline odds.

Since the race predictor has three levels, it is useful to produce a two degree-of-freedom Wald test of the joint hypothesis that (a) Black differs from White and (b) Other differs from White:

```
. testparm race#c.age

( 1)  [low]2.race#c.age = 0
( 2)  [low]3.race#c.age = 0

      chi2( 2) =      1.52
      Prob > chi2 =      0.4667
```

This omnibus test does not demonstrate a statistically significant interaction between race and age at the conventional 5% level of significance. Nonetheless, for illustration purposes we will move forward with decomposing the interaction.

There are different ways this interaction could be unpacked. For example, one approach that is well described is to generate predicted probabilities and plot them. That approach is demonstrated nicely in detail here: <https://stats.oarc.ucla.edu/stata/faq/how-can-i-understand-a-categorical-by-continuous-interaction-in-logistic-regression-stata-12/>, so we do not discuss it further here. Instead, we present an approach that builds on the logic of the simple main effects concept presented in the previous section.

Nested Fixed Effects Model: If we visualize a two-way line plot with the log-odds (logit) of low birthweight on the Y-axis and values of the continuous variable `age` on the X-axis, one possible approach we could take is to estimate separate slopes for the regression of (the log-odds of) `low` onto `age` for each race group. A way to do that is to return to the original model we fitted but instead fit an *equivalent model* that contains different parameters. One such model is a *nested fixed effects* model. The nested fixed effects model replaces the interaction effect we estimated in the original model with a set of race group-specific slopes. Here is how to specify the nested fixed effects model in Stata:

```
. logistic low i.race c.age#i.race

Logistic regression                                Number of obs =      189
                                                    LR chi2(5)         =      8.09
                                                    Prob > chi2        = 0.1512
Log likelihood = -113.28973                       Pseudo R2         = 0.0345
```

	low	Odds ratio	Std. err.	Z	P> z	[95% conf. interval]	
-----+-----							
race							
Black		.2060127	.4310946	-0.75	0.450	.0034096	12.44763
Other		1.908774	3.258261	0.38	0.705	.0672605	54.16874
race#c.age							
White		.9457412	.0435732	-1.21	0.226	.8640819	1.035118
Black		1.051982	.0841687	0.63	0.526	.8992985	1.230588
Other		.9415	.0544649	-1.04	0.297	.84058	1.054536
_cons		1.190902	1.315057	0.16	0.874	.1367522	10.37093

Note: `_cons` estimates baseline odds.

At first glance, this model looks a lot like the first model. In fact, several things are the same in the two outputs. One is that the log-likelihoods of the two models are identical. The second is that the estimates and associated statistics for the constant term and the race term are also identical.

The key syntactical differences are that we have dropped the age term from the model and we have used only one hash (#) operator instead of the two (##) we used in the original model when specifying the interaction term. In the original model, `c.age##i.race` had Stata include terms for age and for race and as well as include their interaction. On the other hand, in the nested fixed effects model, the terms `i.race c.age#i.race` tell Stata to include only the race term and the interaction term for age with race; the age term is omitted. Therefore, the resulting interaction term subsection of the output incorporates age in a different way by listing the simple main effect slopes of age for White (OR=0.95), Black (OR=1.05), and Other (OR=0.94), respectively, as separate effects. The interpretation of these odds ratios is as follows: for every one-year increase in age, there is a 5% decrease in the odds of a White mother having a low birthweight baby. On the other hand, for every one-year increase in age, there is a 5% increase in the odds of a Black mother having a low birthweight baby. Finally, for every one-year increase in age, there is a 6% decrease in the odds of an Other-race-identified mother having a low birthweight baby.

Since we are discussing the results in terms of slopes, it could also be helpful to visualize the slopes. In our models, the logistic regression model's continuous predictors are assumed to be linearly related to the log-odds (logit) of the low birthweight outcome. We can plot the slopes of the log-odds/age relationship using the following Stata syntax:

```
. margins race, at(age=(14 24 34 44)) predict(xb)
```

```
Adjusted predictions
Model VCE: OIM
```

```
Number of obs = 189
```

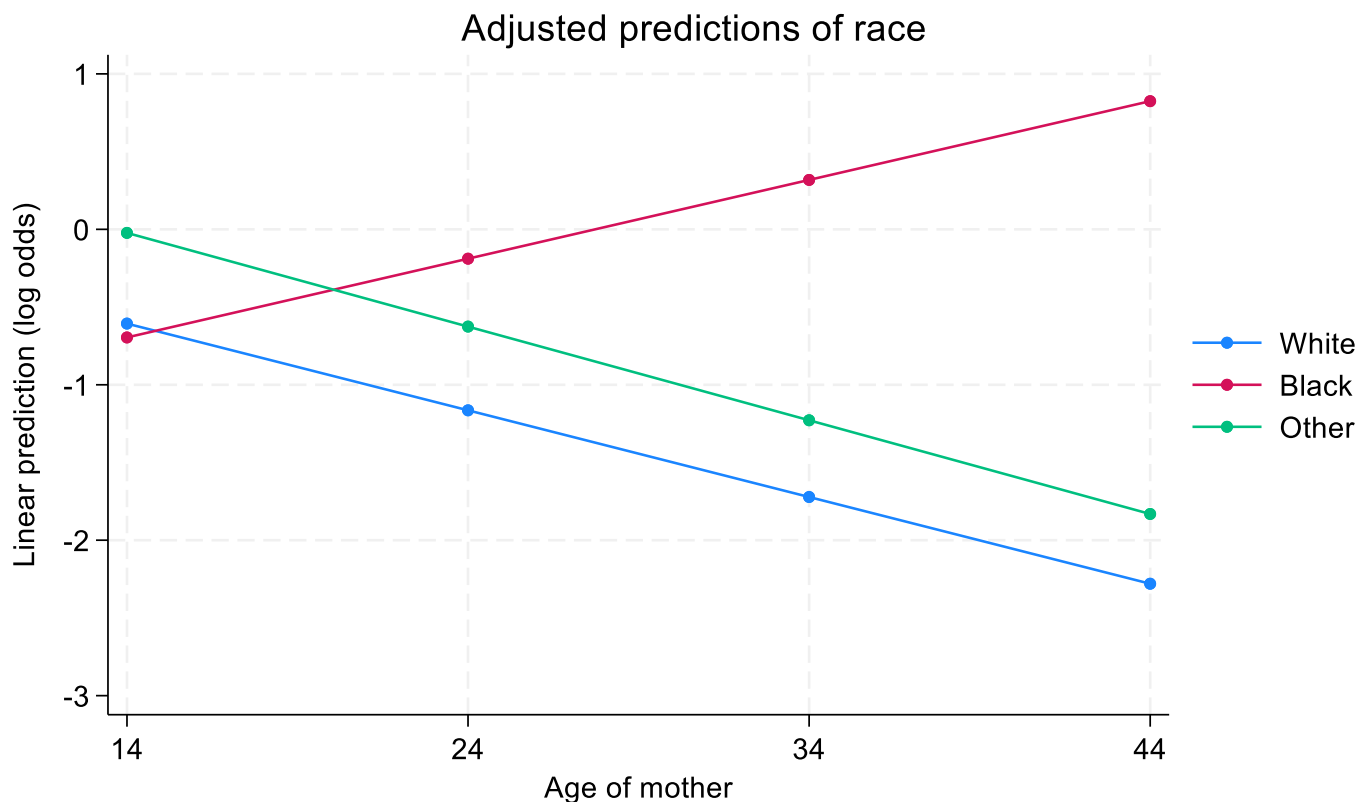
```
Expression: Linear prediction (log odds), predict(xb)
```

```
1._at: age = 14
2._at: age = 24
3._at: age = 34
4._at: age = 44
```

	Margin	Delta-method std. err.	Z	P> z	[95% conf. interval]	
1#White	-.6062969	.4952164	-1.22	0.221	-1.576903	.3643095
1#Black	-.6956445	.7309639	-0.95	0.341	-2.128308	.7370184
1#Other	-.0227607	.5316463	-0.04	0.966	-1.064768	1.019247
2#White	-1.16416	.2426562	-4.80	0.000	-1.639757	-.6885624
2#Black	-.1888861	.4422463	-0.43	0.669	-1.055673	.6779007
2#Other	-.6255699	.2781805	-2.25	0.025	-1.170794	-.0803461
3#White	-1.722023	.5450407	-3.16	0.002	-2.790283	-.6537624
3#Black	.3178722	1.066379	0.30	0.766	-1.772193	2.407937
3#Other	-1.228379	.735816	-1.67	0.095	-2.670552	.2137938
4#White	-2.279885	.9796946	-2.33	0.020	-4.200051	-.3597194
4#Black	.8246306	1.832773	0.45	0.653	-2.767539	4.4168
4#Other	-1.831188	1.294129	-1.41	0.157	-4.367634	.7052583

```
Variables that uniquely identify margins: age race
```

```
. marginsplot, noci name(plot1)
```



The `predict(xb)` option in the margins command specification tells Stata to produce margins estimates on the log-odds scale. Log-odds may be less interpretable and less interesting than predicted probabilities. Predicted probabilities are the default when requesting margins, so we can omit the `predict(xb)` option to obtain predicted probabilities.

```
. margins race, at(age=(14 24 34 44))
```

```
Adjusted predictions
Model VCE: OIM
```

```
Number of obs = 189
```

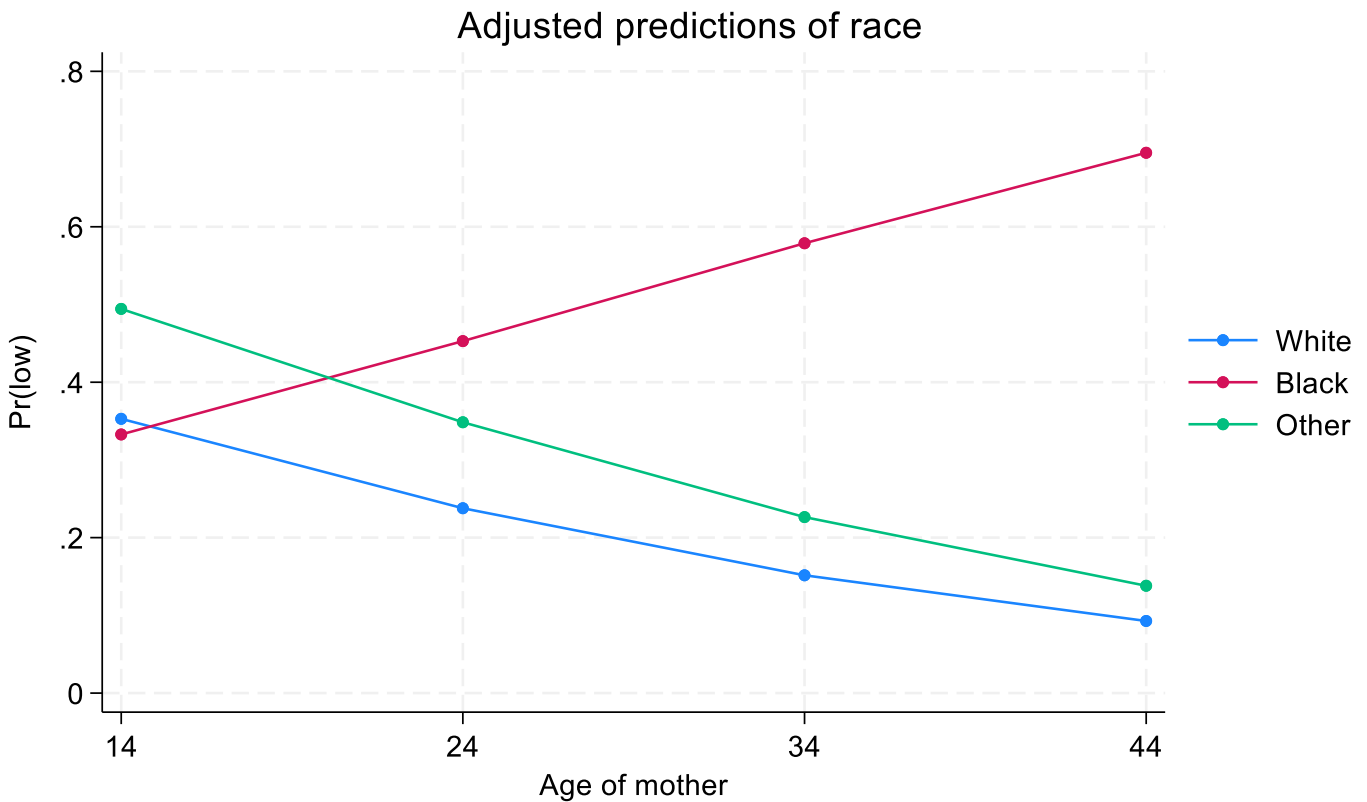
```
Expression: Pr(low), predict()
```

```
1. _at: age = 14
2. _at: age = 24
3. _at: age = 34
4. _at: age = 44
```


	Margin	Delta-method std. err.	Z	P> z	[95% conf. interval]	
_at#race						
1#White	.3529044	.113089	3.12	0.002	.1312539	.5745549
1#Black	.3327786	.162301	2.05	0.040	.0146744	.6508828
1#Other	.4943101	.1328944	3.72	0.000	.2338419	.7547782
2#White	.2379123	.043996	5.41	0.000	.1516817	.3241428
2#Black	.4529184	.1095813	4.13	0.000	.238143	.6676937
2#Other	.3485157	.0631616	5.52	0.000	.2247213	.4723102
3#White	.1516108	.0701059	2.16	0.031	.0142059	.2890158
3#Black	.5788056	.2599723	2.23	0.026	.0692693	1.088342
3#Other	.2264653	.1288993	1.76	0.079	-.0261728	.4791033
4#White	.0928026	.0824808	1.13	0.261	-.0688567	.2544619
4#Black	.6952184	.3883459	1.79	0.073	-.0659255	1.456362
4#Other	.1380968	.1540351	0.90	0.370	-.1638064	.44

. marginsplot, noci name(plot2)

Variables that uniquely identify margins: age race



The slight curves in the lines of predicted probabilities reflect the non-linear (exponential) nature of the logistic regression model.

Relationships between the simple main effects and the interaction term: As in the earlier example involving two categorical predictors, there is a relationship between the interaction terms from the original default specification of the logistic regression model and the simple main effects obtained from the nested fixed effects version of the model. Specifically, an interaction term in the default model represents the *difference* between a pair of slopes. To demonstrate this idea, we'll compare (a) Black-identified to White-identified participants and (b) Other-identified to White-identified participants using the Stata `lincom` post-estimation command:

```
. lincom 2.race#c.age - 1b.race#c.age
```

```
( 1) - [low]1b.race#c.age + [low]2.race#c.age = 0
```

low	Odds ratio	Std. err.	Z	P> z	[95% conf. interval]
(1)	1.112336	.1026986	1.15	0.249	.9282117 1.332984

```
. lincom 3.race#c.age - 1b.race#c.age
```

```
( 1) - [low]1b.race#co.age + [low]3.race#c.age = 0
```

low	Odds ratio	Std. err.	Z	P> z	[95% conf. interval]
(1)	.9955155	.0736227	-0.06	0.952	.861188 1.150795

The first `lincom` statement compares Black-identified to White-identified participants while the second `lincom` statement compares Other-identified participants to White-identified participants. The comparisons are performed by subtracting the White group's age slope from the Black group's and Other group's age slopes, respectively. Comparing the odds ratios obtained via the `lincom` statement from the nested fixed effects model to the interaction terms from the original default model reveals they are identical.

Part 3: Two Continuous Predictors

Interpretation of interactions among continuous-by-continuous predictors is especially difficult. While it is possible to interpret the interaction's odds ratio numerically as the amount by which one predictor's odds changes multiplicatively per unit increase in the second predictor (Jaccard, 2001, pp. 42-44), most investigators will probably find a visualization-based approach will yield more intuitive and informative results. One way to do that is to use Stata's `-margins-` command to produce a plot of predicted probabilities similar to the one shown in Part 2 above for the categorical-by-continuous predictor scenario, except that in the continuous-by-continuous predictor scenario, the analyst will estimate a 1-unit change in one predictor's impact on the predicted probabilities while the other predictor is held constant at different values. The resulting estimates are then plotted. This approach is demonstrated at <https://stats.oarc.ucla.edu/stata/faq/how-can-i-understand-a-continuous-by-continuous-interaction-in-logistic-regression-stata-12/>.

Another alternative which we like is to use contour plots. A contour plot displays regions of probability. Let's generate a contour plot for the interaction between the continuous predictors `age` and mother's last pre-pregnancy weight (`lwt`):

```
. logistic low c.age##c.lwt
```

Logistic regression

Number of obs =	189
LR chi2(3) =	7.53
Prob > chi2 =	0.0567
Pseudo R2 =	0.0321

Log likelihood = -113.56918

	low	Odds ratio	Std. err.	z	P> z	[95% conf. interval]	
age		.963454	.1503887	-0.24	0.811	.7095192	1.308271
lwt		.9877904	.0281114	-0.43	0.666	.9342014	1.044453
c.age#c.lwt		.9999798	.0011978	-0.02	0.987	.9976348	1.00233
_cons		5.404975	19.75623	0.46	0.644	.004183	6983.929

Note: `_cons` estimates baseline odds.

The range of `age` is from 14-45 and the range of `lwt` is from 80-250 lbs. To generate the contour plot, we will closely follow the example code provided by Chuck Huber of Stata Corp in his online demonstration of generating interaction contour plots using Stata at <https://www.stata.com/stata-news/news32-1/spotlight/>.

We first obtain the predicted probabilities of having a low-birth-weight baby across ages 14-45 (in increments of 1 year) and mother weights of 80-250 (in increments of 10 pounds), to yield a manageable number of plotting points).


```
. quietly margins, at(age=(14(1)45) lwt=(80(10)250)) saving(predictions, replace) nose
```

The `-quietly-` prefix suppresses output from `-margins-`. The `-nose-` option suppresses calculation of standard errors for the margins, which we will not use in the contour plot and which can take a long time to calculate. The `-saving-` options saves the resulting predicted probabilities in a separate dataset named "predictions." After loading the new dataset into memory, we rename the variables in the predictions dataset to have more meaningful names:

```
. use predictions, clear
. rename _at1 age
. rename _at2 weight
. rename _margin pr_low_b_weight
```

Now, we are ready to generate the plot using Stata's `-tway contour-` command:

```
. tway contour pr_low_b_weight weight age
```

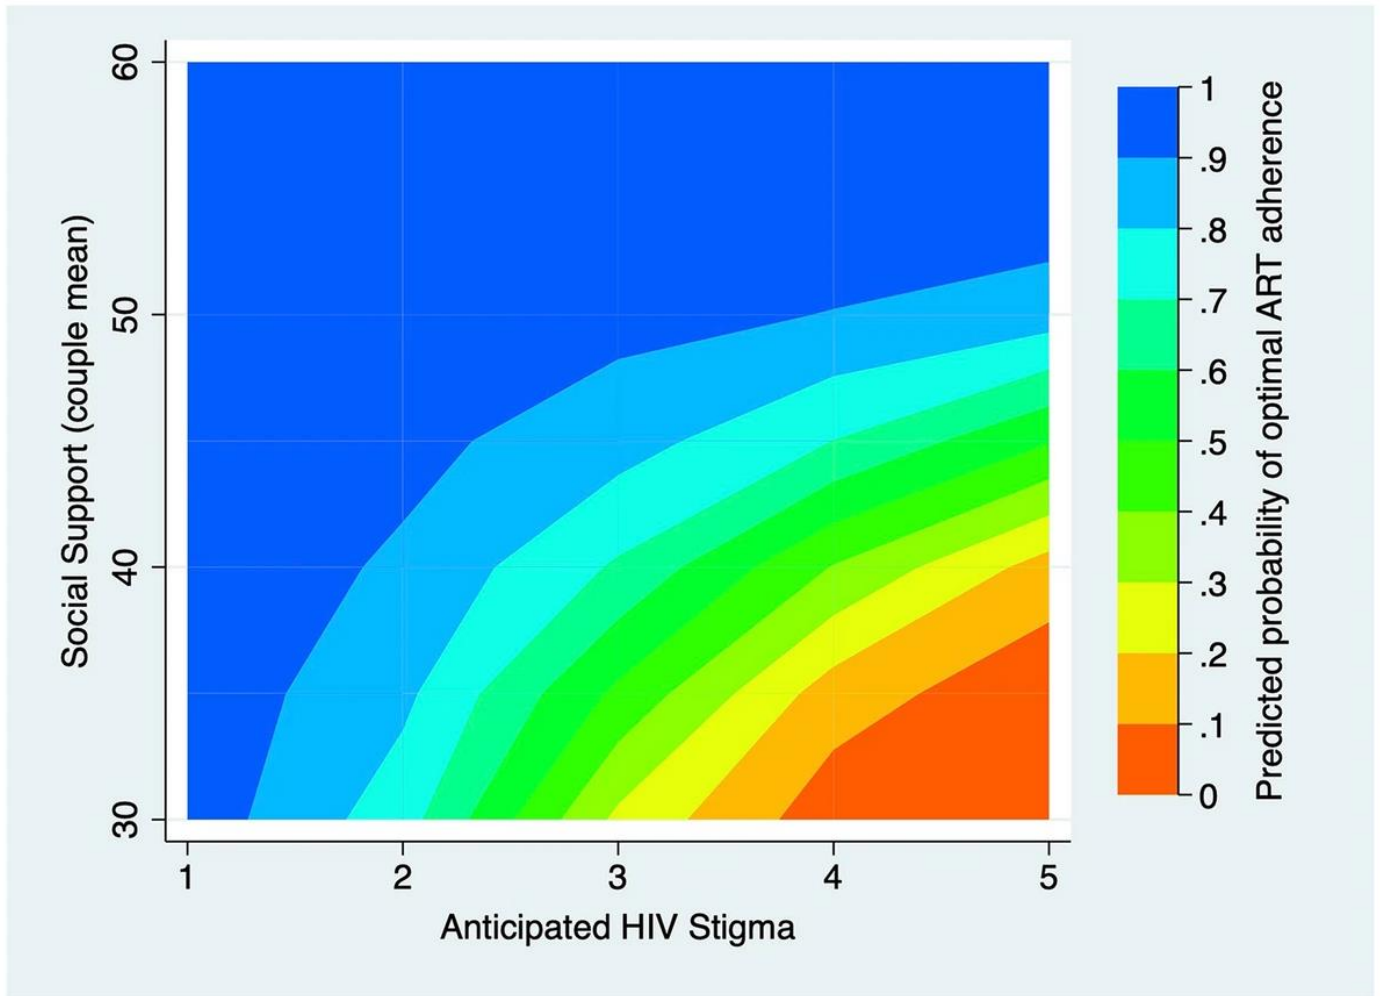


Since there is little evidence of interaction between mother's age and mother's weight in predicting the odds of a low-birth-weight baby, the plot shows straight contour lines demarcating three regions of probability. The odds of a mother having a low-birth-weight baby appear highest when the mother is young and has low weight, though since the interaction term is not statistically significant at the conventional 5% level, these findings should be interpreted with extreme caution, if at all.

As a contrasting – and more compelling – example, Gutin and colleagues investigated associations between continuous measures of anticipated HIV stigma and romantic partner social support, and the probability of optimal HIV antiretroviral therapy medication adherence, published in *BMC Public Health* (<https://bmcpublichealth.biomedcentral.com/articles/10.1186/s12889-023-16762-w>). They found a statistically significant interaction between HIV stigma and partner support (OR = 1.10, $p=.032$) in explaining antiretroviral therapy (ART) adherence. They depicted this interaction using a contour plot (<https://bmcpublichealth.biomedcentral.com/articles/10.1186/s12889-023-16762-w/figures/1>), reproduced below with permission from Dr. Gutin and the journal under the Creative Commons license agreement (<https://creativecommons.org/licenses/by/4.0/>):

Fig. 1

From: [Supportive couple relationships buffer against the harms of HIV stigma on HIV treatment adherence](#)



Reading the plot from left to right, we see that the predicted probability of ART adherence is especially poor (i.e., low, the red color in the figure) when anticipated HIV stigma is high and partner support is low. Alternatively, ART adherence is high (the dark blue color in the figure) when anticipated stigma is low and social support is high.

References

Jaccard J. *Interaction Effects in Logistic Regression*. Vol 07-135. Thousand Oaks, CA: Sage Publications, Inc.; 2001.

Relevant Resources for Stata Users:

- <https://stats.oarc.ucla.edu/stata/faq/everything-you-always-wanted-to-know-about-contrasts-but-were-afraid-to-ask/>
- <https://stats.oarc.ucla.edu/stata/seminars/interactions-stata/>
- <https://stats.oarc.ucla.edu/stata/faq/how-can-i-understand-a-categorical-by-continuous-interaction-in-logistic-regression-stata-12/>
- <https://stats.oarc.ucla.edu/stata/faq/how-can-i-understand-a-continuous-by-continuous-interaction-in-logistic-regression-stata-12/>
- Mitchell (2021). *Interpreting and Visualizing Regression Models Using Stata: Second Edition*. College Station, TX: Stata Press.

Relevant Resources for SAS Users:

- <https://blogs.sas.com/content/iml/2019/05/30/visualize-interaction-effects-regression.html>
- <https://blogs.sas.com/content/iml/2016/06/22/sas-effectplot-statement.html>

Acknowledgements

We appreciate review by Drs. Estie Hudes, Anita Stewart, Sarah Gutin and permission from Dr. Gutin to share the example contour plot from Dr. Gutin's publication.