

Model Coefficients, Adjusted Predictions, & Marginal Effects: A Summary of How All Three are Related

Richard Williams and Junrong Sheng, University of Notre Dame, <https://academicweb.nd.edu/~rwilliam/>
Last revised September 10, 2024

Critical Point: Adjusted Predictions and Marginal Effects are *functions* of a model's coefficients. While a model's coefficients may have limited intuitive appeal, the APs and MEs can make their substantive and practical significance much clearer and more tangible.

- In the Logit02 handout, we examined how being in an experimental teaching program affected the likelihood of getting an A in the course. The logistic regression coefficient for the variable `psi` was 2.379 and was statistically significant. This clearly implied that being in PSI made you more likely to get an A – but it didn't tell you how much more likely.
 - That same handout showed how you could plug in different values for the X variables and compute the predicted probability of an A for someone with the specified characteristics. For example, we saw that a B student with a score of 20 on TUCE in a traditional classroom would have less than a 7% chance of getting an A. But, if that same person were in PSI instead, their probability of getting an A would jump to over 43%. Telling somebody that they are 37 percentage points more likely to get an A if they are in PSI probably has a lot more intuitive appeal to them than saying their log odds of success increase by 2.379 if they are in `psi`.
 - It may be worth reviewing the Logit02 handout if you don't remember how we were able to go from the model coefficients to the predicted probabilities. *Coefficients and predicted probabilities are not the same, but predicted probabilities for a given set of X values are a function of (i.e. can be computed from) the model coefficients.*
 - Similar calculations showed that the straight A student in a traditional classroom would have a 66% chance of getting an A, but if they were instead in PSI they would have a 95% chance of success.
 - Conversely, the student with the lowest GPA in the class (2.06) only had a 0.6% chance of getting an A in a traditional classroom. Put that person in PSI, and their chance of getting an A jumped to 6.13%. Which is still pretty low, but better than it would be by not being in PSI.
 - In short, the logistic regression coefficient of 2.379 for PSI told us that being in PSI was a good thing. But, by plugging in different values for the Xs and calculating the predicted probabilities given those X values, we got a far better feel for how beneficial being in PSI was, and we also saw that the benefits of PSI were dependent on the values of other variables in the model
- Similarly, the logistic regression coefficient of .059 for age in the Margins01 handout demonstrated that getting older is bad for your health. This is not surprising! But the handout also showed that by plugging in different values for age and the other X variables, the effects of age can be made much more tangible. In the examples given, a 20-year-old was expected to have less than a 1% chance of having diabetes, while for a 70-year-old the predicted probability was about 11%. For me at least, that provides a much clearer feel for the effect of age than did the logistic regression coefficient.

- Margins01 and other Margins handouts further elaborated on the types of computations you might do. (The [Appendix](#) reviews how and why the following calculations are done.)
 - The Stata 17 Base Reference Manual (p. 1414) says “A margin is a statistic based on a fitted model calculated over a dataset in which some of or all the covariates are fixed at values different from what they really are. For instance, after a linear regression fit on males and females, the marginal mean (margin of mean) for males is the predicted mean of the dependent variable, where every observation is treated as if it represents a male; thus, those observations that in fact do represent males are included, as well as those observations that represent females. The marginal mean for females would be similarly obtained by treating all observations as if they represented females.”
 - The Margins handouts examined various special cases of margins. We showed how you might compare two otherwise “average” individuals who differ on one trait. (We also saw there were different ways of defining “average”. In this handout I am going to deal with cases where other X values are held at their observed values, but you could also hold them at their mean values.)
 - For example, in Margins01, we computed [Average Adjusted Predictions](#) for race. We did this by first treating every respondent as though they were a Black person (regardless of what their race actually was), then treating every respondent as though they were a White person (again, regardless of what their race really was), and then computing their predicted probability of having diabetes given their substituted race and their actual values on the other X variables in the model.
 - So, in the diabetes example, when using Average Adjusted Predictions, the average Black person had an 8.4% predicted probability of having diabetes, while the average White person had only a 4.4% predicted probability.
 - [The difference between those two numbers is the Average Marginal Effect of race](#), i.e. $8.4\% - 4.4\% = 4\%$. That is, on average Black individuals are 4 percentage points more likely to have diabetes than are White individuals. The AME gives you a single number for comparing groups, which can be quite handy, but it also obscures some detail (e.g. the AME would have also been 4% if the AAPs had been 25% and 21%, or 64% and 60%).
 - To put it more formally, Marginal Effects indicate how change in the value of X changes the value of the outcome. Or, as Stephanie Glen puts it (<https://www.statisticshowto.com/marginal-effects/>), “Marginal effects tells us how a dependent variable (outcome) changes when a specific independent variable (explanatory variable) changes. Other covariates are assumed to be held constant.” In this case, the AME for race shows you how the probability of having diabetes changes when the value of race changes, holding all other X variables at their observed values.
 - Note that, in this case, [race is a discrete/categorical variable, not a continuous variable. AMEs for continuous independent variables are computed differently and have much less intuitive appeal](#). Margins02 explains this. I greatly prefer the approach taken by the `mcp` command, which computes adjusted predictions for several different possible values of continuous independent variables and then displays the results graphically.

- But, however you define “average,” averages can conceal important details. It would NOT be correct to say that, if you had two otherwise identical people, one a Black person, one a White person, the Black person would be four percentage points more likely to have diabetes. No; while the **average** difference between White and Black individuals may be 4 percentage points, the actual predicted difference in any given case depends on the values of other variables in the model.
 - For example, when we looked at **Adjusted Predictions at Representative Values** (APRs) we saw that a 20-year-old White Person had a 0.6% probability of having diabetes, while an otherwise identical 20-year-old Black person had a 1.2% probability of having diabetes. Yes, the 20-year-old Black person is more likely to have diabetes, but the difference is far less than the average of 4 percentage points. This reflects the fact that, at least back in 1980, 20-year-olds were unlikely to have diabetes no matter what their race was.
 - Conversely, when the same calculations are done for 70 years, we found that a 70-year-old White person had a 10.34% chance of having diabetes while for the 70-year-old Black person the estimate was 19.12%. This is more than double the average difference of 4 percentage points we found earlier.

Summary

- Way back in the Logit02 handout, we showed how to go from model coefficients to Adjusted Predictions and predicted probabilities. You may want to review that handout. It may not have been clear at the time, but we were laying the groundwork for what was to come. For pedagogical purposes, I showed you how to do things by hand so you would get an understanding of what was being done, but ever since then we’ve let Stata do most of the calculations for us.
- Programs like `margins`, `mcp`, `mtable`, `mchange`, and other routines are basically applying the principles learned early on. They are making life easy for you, doing lots of calculations for you, and giving you easy-to-read (or at least easier-to-read) output. But they are basically just automating things demonstrated earlier in the course.
- With discrete/categorical independent variables, the Marginal Effect is the difference between 2 Adjusted Predictions, e.g. the Marginal Effect for a variable like Married would be the difference between the Adjusted Prediction for married people and the Adjusted Prediction for unmarried people.
- For continuous independent variables, the computation of Marginal Effects is more complicated and MEs are harder to interpret. At least they are for me, but economists are fonder of them than I am. I find that approaches taken by programs like `mcp`, which instead compute several adjusted predictions using several possible values of the X continuous variable, produce easier to interpret results.
- So again, the most critical thing to understand: Adjusted Predictions and Marginal Effects are **functions** of a model’s coefficients. All three are related but they are not the same things. While a model’s coefficients may have little intuitive appeal, the APs and MEs that can be computed from the coefficients can make their substantive and practical significance much clearer and more tangible.

Appendix: Computing Average Adjusted Predictions and Average Marginal Effects (Adapted from Margins01)

The following example shows, both intuitively and with Stata code, what AAPs and AMEs are and how they are computed.

Intuitive Explanation. Intuitively, in this example, the AAPs and AME for the variable black are computed as follows:

- Go to the first case. Treat that case as though they were a White person, regardless of what the person's race actually is. Leave all other independent variable values for the case as is. Compute the probability this person (if he or she were a White person) would have diabetes.
- Now do the same thing, this time treating the case as though they were a Black person.
- The difference in the two probabilities just computed (i.e. the Adjusted Predictions for the case) is the Marginal Effect for that case
- Repeat the process for every case in the sample
- Compute the average of all the Adjusted Predictions and Marginal Effects you have computed. This gives you the AAPs and the AME for the variable black.

Corresponding Stata Code

```
. webuse nhanes2f, clear
. logit diabetes i.black i.female age, nolog
```

```
Logistic regression                                Number of obs = 10,335
                                                    LR chi2(3)      = 374.17
                                                    Prob > chi2     = 0.0000
Log likelihood = -1811.9828                        Pseudo R2      = 0.0936
```

diabetes		Coefficient	Std. err.	z	P> z	[95% conf. interval]	
black							
Black		.7179046	.1268061	5.66	0.000	.4693691	.96644
female							
Female		.1545569	.0942982	1.64	0.101	-.0302642	.3393779
age		.0594654	.0037333	15.93	0.000	.0521484	.0667825
_cons		-6.405437	.2372224	-27.00	0.000	-6.870384	-5.94049

```
. * Average Adjusted Predictions (AAPs)
. margins black
```

Predictive margins
Model VCE: OIM

Number of obs = 10,335

Expression: Pr(diabetes), predict()

		Delta-method				
	Margin	std. err.	z	P> z	[95% conf. interval]	
black						
Not Black	.0443248	.0020991	21.12	0.000	.0402107	.0484389
Black	.084417	.0084484	9.99	0.000	.0678585	.1009756

```
. * Average Marginal Effects (AMEs)
. margins, dydx(black)
```

Average marginal effects
Model VCE: OIM

Number of obs = 10,335

Expression: Pr(diabetes), predict()
dy/dx wrt: 1.black

		Delta-method				
	dy/dx	std. err.	z	P> z	[95% conf. interval]	
black						
Black	.0400922	.0087055	4.61	0.000	.0230297	.0571547

Note: dy/dx for factor levels is the discrete change from the base level.

```
. * Replicate AAPs and AME for black without using margins
. clonevar xblack = black
. quietly logit diabetes i.xblack i.female age, nolog
. replace xblack = 0
(1,086 real changes made)
```

```
. predict adjpredwhite if e(sample)
(option pr assumed; Pr(diabetes))
. replace xblack = 1
(10,337 real changes made)
. predict adjpredblack if e(sample)
(option pr assumed; Pr(diabetes))
. gen meblack = adjpredblack - adjpredwhite
. sum adjpredwhite adjpredblack meblack
```

Variable	Obs	Mean	Std. dev.	Min	Max
adjpredwhite	10,335	.0443248	.0362422	.005399	.1358214
adjpredblack	10,335	.084417	.0663927	.0110063	.2436938
meblack	10,335	.0400922	.0301892	.0056073	.1078724