Chapter 2

INTERPRETING LOGISTIC REGRESSION COEFFICIENTS

As is true for nonlinear transformations more generally, the effects of the independent variables in logistic regression have multiple interpretations. Effects exist for probabilities, odds, and logged odds, and the interpretation of each effect varies.

To preview, the effects of the independent variables on the logged odds are linear and additive—each *X* variable has the same effect on the logged odds regardless of its level or the level of other *X* variables—but the units of the dependent variable, logged odds, have little intuitive meaning. The effects of the independent variables on the probabilities have intuitive meaning but are nonlinear and nonadditive—each *X* variable has a different effect on the probability depending on its level and the level of the other independent variables. Despite the interpretable units, the effects on probabilities are less easily summarized in the form of a single coefficient. The interpretation of the effects of the independent variables on the odds offers a popular alternative. The odds have more intuitive appeal than the logged odds and can still express effects in single coefficients, but the effects on odds are multiplicative rather than additive.

This chapter examines the multiple ways to interpret effects in logistic regression results. It gives particular attention to interpretations of probability effects, the most informative but also the most complex way to understand logistic regression results.

Logged Odds

One interpretation directly uses the coefficients obtained from the estimates of a logistic regression model. The logistic regression coefficients show the change in the predicted logged odds of experiencing an event or having a characteristic for a one-unit increase in the independent variables, holding other independent variables constant. The coefficients are similar to linear regression coefficients in that a single linear and additive coefficient summarizes the relationship. The difference is that the dependent variable takes the form of logged odds.

Consider an example. Returning to the 2017 National Health Interview Survey (NHIS) data and the binary outcome measure of currently smokes, a simple model includes continuous measures of age (26–85+) and years of education (0–18) plus categorical measures of gender (a dummy variable with males coded 1), race (four dummy variables with whites as the referent), and Hispanic ethnicity (with Hispanics coded 1). The sample size is 23,786. Selected output from a logistic regression in Stata produces the results in Table 2.1.

For the continuous variables, the predicted logged odds of smoking on average decrease by .183 with a 1-year increase in education and by .024 with a 1-year increase in age, controlling for other predictors. For the categorical variables, a change of one unit implicitly compares the indicator group to the reference or omitted group. The coefficient of .254 for gender indicates that the predicted logged odds of smoking are higher by .254 for men than women. The coefficients for race show that, compared to whites, the log odds of smoking are lower by .084 for African Americans, higher by .288 for Native Americans, lower by .810 for Asian Americans, and higher by .436 for multi-race respondents. The gap relative to whites is largest for Asian Americans with the controls, but the gap between Asian Americans and multi-race respondents is still larger. An additional measure shows that the logged odds of smoking are lower by 1.040 for Hispanics than non-Hispanics.

Table 2.1 Stata Output: Logistic Regression Model of Current Smoking, NHIS 2017

Smoker	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
Education	-.1830747	.0067791	-27.01	0.000	-.1963616	-.1697878
Age	-.023966	.0011756	-20.39	0.000	-.0262701	-.0216619
Gender						
Male	.2535793	.0369909	6.86	0.000	.1810785	.3260801
Race						
African American	-.0835037	.0574981	-1.45	0.146	-.1961979	.0291905
Native American	.2882139	.1523181	1.89	0.058	-.0103242	.586752
Asian American	-.8100691	.1092474	-7.41	0.000	-1.02419	-.595948
Multiple Race	.4363179	.1180186	3.70	0.000	.2050056	.6676301
Ethnicity						
Hispanic	-1.039563	.0700205	-14.85	0.000	-1.176801	-.9023256
_cons	2.053864	.127611	16.09	0.000	1.803751	2.303976

The coefficients represent the relationship, as in ordinary regression, with a single coefficient. Regardless of the value of an independent variable—small, medium, or large—or the values of the other independent variables, a one-unit change has the same effect on the dependent variable. According to the model, the difference in the logged odds of smoking between white women and men is the same as the difference in the logged odds between Asian-American women and men. Similarly, the effect of education in the model does not differ between men and women or between any of the race-ethnic groups. Indeed, logistic regression aims to simplify the nonlinear and nonadditive relationships inherent in treating probabilities as dependent variables.

Despite the simplicity of their interpretation, the logistic regression coefficients, as mentioned, lack a meaningful metric and offer little substantive information other than the sign. Statements about the effects of variables on changes in logged odds reveal little about the relationships and do little to help explain the substantive results. Interpreting the substantive meaning or importance of the coefficients requires something more than reporting the expected changes in logged odds.

Tests of Significance

Tests of significance often receive much attention, perhaps too much attention, in logistic regression. If the coefficients have little intuitive meaning in terms of substantive importance, it is easy to note the statistically significant and nonsignificant coefficients. Then, the signs of the significant coefficients offer a crude but quick summary of the results.

As in regression, the size of a coefficient relative to its standard error provides the basis for tests of significance in logistic regression. The logistic regression procedures in Stata and R present the coefficient divided by its standard error, which can be evaluated with the z distribution. The significance of the coefficient—the likelihood that the coefficient in the sample could have occurred by chance alone when the population parameter equals 0—is then interpreted as usual. However, since we know little about the small sample properties of logistic regression coefficients, tests of significance for samples less than 100 prove risky (Long, 1997, p. 54).

Table 2.1 lists, along with the coefficients, the standard errors of the coefficients, the z values, the probabilities of the z values under the null hypothesis, and 95% confidence intervals around the coefficients. With a sample size of 23,786, all but two of the coefficients reach statistical significance at the .001 level. The coefficients for Native Americans and African Americans are not significant at the usual .05 level. Although the logistic regression results treat race as four separate dummy variables, it is of course a single categorical measure. It is important to test for the significance of all the racial categories together using a procedure discussed in the next chapter.

The logistic regression procedure in SPSS calculates the Wald statistic for a (two-tailed) test of a single coefficient. Table 2.2 shows SPSS output from the same logistic regression model as presented above. The coefficients and standard errors are identical. The Wald test appears different but in fact is the same as the *z* value squared (Hosmer, Lemeshow, & Sturdivant, 2013).[7] The Wald statistic in SPSS has a chi-square distribution with one degree of freedom. As before, all coefficients except two are significant at .001.

Statistical significance has obvious importance but depends strongly on sample size. The *p* values provide little information on the strength or substantive meaning of the relationship. Large samples, in particular, can produce significant *p* values for otherwise small and trivial effects. Despite the common reliance of studies on statistical significance (and the sign of the coefficient) in interpreting logistic regression coefficients, *p* values best serve only as an initial hurdle to overcome before interpreting the coefficient in other ways.

Table 2.2 SPSS Output: Logistic Regression Model of Current Smoking, NHIS 2017

		B	S.E.	Wald	df	Sig.	Exp(B)
Step 1[a]	Education	−.183	.007	729.302	1	.000	.833
	Age	−.024	.001	415.612	1	.000	.976
	Gender (1)	.254	.037	46.994	1	.000	1.289
	Race			76.138	4	.000	
	Race (1)	−.084	.057	2.109	1	.146	.920
	Race (2)	.288	.152	3.580	1	.058	1.334
	Race (3)	−.810	.109	54.982	1	.000	.445
	Race (4)	.436	.118	13.668	1	.000	1.547
	Ethnicity (1)	−1.040	.070	220.420	1	.000	.354
	Constant	2.054	.128	259.040	1	.000	7.798

[a]Variable(s) entered on step 1: Education, Age, Gender, Race, and Ethnicity.

Odds

The second interpretation comes from transforming the logistic regression coefficients so that the independent variables affect the odds rather than the logged odds of the dependent variable. Recall that the odds equal the probability of a binary outcome divided by one minus the probability, or $P/(1 - P)$. To find the effects on the odds, take the exponent or antilogarithm of the logistic regression coefficients. Exponentiating both sides of the logistic regression equation eliminates the log of the odds and shows the influences of the variables on the odds. The transformation from logged odds to odds for logistic regression with multiple predictors is as follows:

$$\ln\left(P/1-P\right) = b_0 + b_1X_1 + b_2X_2,$$
$$e^{\ln\left(P/1-P\right)} = e^{b_0} + e^{b_1X_1 + b_2X_2},$$
$$P/1-P = e^{b_0} \times e^{b_1X_1} \times e^{b_2X_2}.$$

With the odds rather than the logged odds as the outcome, the right-hand side of the equation becomes multiplicative rather than additive.

The odds are a function of the exponentiated constant e^{b_0} multiplied by the exponentiated product of the coefficient and $X_1\,(e^{b_1X_1})$ and the exponentiated product of the coefficient and $X_2\,(e^{b_2X_2})$. The effect of each variable on the odds (rather than the logged odds) thus comes from taking the antilog of the coefficients. If not already presented in the output, the exponentiated coefficients can be obtained using any calculator by typing the coefficient and then the e^x function. The exponentiated coefficients of −.183, −.024, and .254 from Tables 2.1 and 2.2 equal, respectively, .833, .976, and 1.289. These are conveniently listed in the last column of the SPSS output and can be easily obtained with options in Stata.

The fact that the equation determining the odds is multiplicative rather than additive shifts the interpretation of the exponentiated coefficients. In an additive equation, a variable has no effect when its coefficient equals 0. The predicted value of the dependent variable sums the values of the variables times the coefficients; when adding 0, the predicted value does not change. In a multiplicative equation, the predicted value of the dependent variable does not change when multiplied by a coefficient of 1. Therefore, 0 in the additive equation corresponds to 1 in the multiplicative equation. Furthermore, the exponential of a positive number exceeds 1 and the exponential of a negative number falls below 1 but above 0 (as the exponential of any number is always greater than 0).

For the exponentiated coefficients, then, a coefficient of 1 leaves the odds unchanged, a coefficient greater than 1 increases the odds, and a coefficient smaller than 1 decreases the odds. Moreover, the more distant the coefficient from 1 in either direction, the greater the effect in changing the odds. Recall as well that the odds are not symmetric around 1. They vary between 0 and 1 on one end, but from 1 to positive infinity on the other.

Interpretation

To illustrate the interpretations of the exponential coefficients, or the effects on odds, Table 2.3 presents logistic regression output from R, again using the model of current smoking. The commands required to obtain the exponentiated logistic regression coefficients and the format of the output differs from Stata and SPSS. But the results are the same.

Table 2.3 R Output: Logistic Regression Model of Current Smoking, NHIS 2017

```
Coefficients:
                                  Std.
                  Estimate       Error.     z value   Pr(>|z|)
 (Intercept)       2.053864   0.127611     16.095    < 2e-16 ***
 Education        -0.183075   0.006779    -27.006    < 2e-16 ***
 Age              -0.023966   0.001176    -20.387    < 2e-16 ***
 Gender            0.253579   0.036991      6.855   7.12e-12 ***
 Race.f2          -0.083504   0.057498     -1.452   0.146422
 Race.f3           0.288214   0.152318      1.892   0.058466  .
 Race.f4          -0.810069   0.109247     -7.415   1.22e-13 ***
 Race.f5           0.436318   0.118019      3.697   0.000218 ***
 Ethnicity.f1     -1.039563   0.070020    -14.847    < 2e-16 ***

                       Odds Ratio and Confidence Interval
 (Intercept)         7.7979707       6.0750055     10.0186332
 Education           0.8327060       0.8216888      0.8438194
 Age                 0.9763189       0.9740669      0.9785661
 Gender              1.2886296       1.1985124      1.3855425
 Race.f2             0.9198877       0.8210606      1.0286780
 Race.f3             1.3340426       0.9826226      1.7869569
 Race.f4             0.4448273       0.3570509      0.5481475
 Race.f5             1.5470004       1.2225119      1.9425137
 Ethnicity.f1        0.3536091       0.3077680      0.4049986
```

For example, exponentiating the coefficient of −.183 indicates that the predicted odds of smoking are reduced by a multiplicative factor of .833 with a 1-year increase in years of education, holding the other predictors constant. If, hypothetically, the odds of smoking for someone with 12 years of education equal 0.300, then the predicted odds of smoking for someone with 13 years of education equal .300 × .833 or .250. The exponentiated coefficient for age of .976 indicates that the odds of smoking are reduced by a multiplicative factor of .976 with 1-year increase in age. If the predicted odds at age 25 are .400, then the predicted odds at age 26 would fall to .390 (or .400 × .976).

The same relationships can be restated in terms of odds ratios. The ratio of the predicted odds of smoking for someone with 13 years of education to someone with 12 years of education, or for someone with 18 years of education to someone with 17 years of education, equals the exponentiated logistic regression coefficient of .833. The ratio of predicted odds for someone aged 26 years (or age 56) to someone aged 25 years (or age 55) equals .976. Thus, the exponentiated coefficient shows the ratio of odds for those one-unit higher to those one-unit lower on the independent variable.

For categorical predictors in the form of dummy variables, a similar interpretation follows. The exponentiated coefficient for men of 1.289 indicates that their odds of smoking are higher than those for women by a factor of 1.289. Here, a one-unit increase defines the comparison of men to the reference group of women. If the predicted odds of smoking equal .200 for women, they equal .200 × 1.289 or .258 for men. Equivalently, the ratio of the odds of smoking for men to women is 1.289. The exponentiated coefficient of 1.547 shows higher odds of smoking for multi-race respondents compared to whites. The odds for Hispanics are lower by a factor of .354 than non-Hispanics, and the ratio of odds for Hispanics to non-Hispanics is .354.

Since the distance of an exponentiated coefficient from 1 indicates the size of the effect, a simple calculation can further aid in interpretation. The difference of a coefficient from 1 exhibits the increase or decrease in the odds for a unit change in the independent variable. In terms of a formula, the exponentiated coefficient minus 1 and times 100 gives the percentage increase or decrease due to a unit change in the independent variable:

$$\%\Delta = \left(e^{b} - 1\right) \times 100.$$

For education, the exponentiated coefficient says that the odds of smoking decline by 16.7% or are 16.7% lower with an increase of 1 year in education. This appears more meaningful than to say the logged odds decline by .183. The size of the effect on the odds also depends on the units of measurement of the independent variables—the change in odds for variables measured in different units do not warrant direct comparison. Still, the interpretation of percentage change in the odds has intuitive appeal.[8]

For men, the exponentiated logistic regression coefficient of 1.289 means that the odds of smoking are 28.9% higher than for women. The exponentiated coefficient for Hispanics of .354 indicates that their odds of participating are 64.6% lower than for non-Hispanics.

In interpreting the exponentiated coefficients, remember that they refer to multiplicative changes in the odds rather than probabilities. It is incorrect to say that an additional year of education makes smoking 16.7% less probable or likely, which implies probabilities rather than odds. More precisely, the odds of smoking are .833 times smaller or 16.7% smaller with an additional year of education.

Probabilities

The third strategy of interpreting the logistic regression coefficients involves translating the effects on logged odds or odds into the effects on probabilities. Since the relationships between the independent variables and probabilities are nonlinear and nonadditive, they cannot be fully represented by a single coefficient. The effect on the probabilities has to be identified at a particular value or set of values. The choice of values to use in evaluating the effect on the probabilities depends on the concerns of the researcher and the nature of the data, but an initial strategy has the advantage of simplicity: examine the effect on the probability for a typical case.

Before interpreting probability effects from logistic regression, it helps to introduce two related concepts of predicted probabilities and marginal effects.

First, logistic regression produces a predicted value for each observation in the data, but the predicted value can take the form of logged odds or probabilities. The predicted logits or logged odds are calculated for each observation by substituting that observation's values on the independent variables, multiplying by the estimated logit coefficients, and summing the products. The predicted probabilities can then be obtained by using the formula transforming logits to probabilities. As presented in Chapter 1, the formula shows that the probabilities are a function of the logits, L_i:

$$P_i = (e^{L_i})/(1 + e^{L_i}).$$

Of course, the predicted values can be obtained directly from statistical packages. For example, the logistic regression model of smoking in Tables 2.1

to 2.3 can generate predicted logits and probabilities for each of the 23,786 observations. The summary statistics from Stata in Table 2.4 list the following:

- The outcome (Smoker) has values of only 0 or 1, with a mean of .154 (i.e., 15.4% of the sample currently smokes).
- The predicted logits from the model vary between −4.649 and 1.636. Persons with negative values show relatively low predicted smoking and those with positive values show relatively high predicted smoking. The predicted logits are skewed toward nonsmoking.
- The predicted probabilities, which have limits of 0 and 1, vary from .009 to .837. Persons with probabilities below .5 are less likely to smoke and persons with probabilities above .5 are more likely to smoke. The mean predicted probability is the same as the mean of the dependent variable and shows that most of the sample does not smoke.

The table also presents the comparison of the predicted probabilities from a linear regression, which have no limits of 0 and 1. They vary between −.168 and .595. These values illustrate the point that, unlike linear regression, logistic regression keeps predicted probabilities within the limits.

Second, marginal effects refer to the influence of independent variables on a dependent variable. A marginal effect is defined in general terms as the change in the expected value of a dependent variable associated with a change in an independent variable, holding other independent variables constant at specified values. In linear regression, the marginal effect is simply the slope coefficient for an independent variable. In logistic regression, however, the marginal effect on probabilities varies. It is not fully represented by a single coefficient.

Table 2.4 Stata Output: Summary Statistics for Observed Values of Current Smoking and Predicted Values From Logistic Regression and Regression Models of Current Smoking, NHIS 2017

Variable	Obs	Mean	Std. Dev.	Min	Max
Smoker	23,786	.1539141	.3608739	0	1
Logit_ Smoker	23,786	-1.837902	.6406989	-4.648903	1.636395
Prob_Smoker	23,786	.1539141	.0829396	.0094813	.8370438
Reg_Smoker	23,786	.1539141	.0797695	-.1680726	.5951874

There are two varieties of marginal effects. One involves marginal change in continuous independent variables and the other involves discrete change in categorical independent variables. The two types of marginal effects associated with each type of variable involve different calculations and strategies for interpreting the results.

Continuous Independent Variables

One way to understand the marginal effect of a continuous independent variable on probabilities involves calculating the linear slope of the tangent line of the nonlinear curve at a single point. The slope of the tangent line is defined by the partial derivative of the nonlinear equation relating the independent variables to the probabilities (Agresti, 2013, p. 164). The partial derivative shows the change in the outcome for an infinitely small or marginal change in the predictor. More intuitively, it represents a straight line that meets the logistic curve at a single point. Figure 2.1 depicts the tangent line where the logistic curve intersects $Y = P = .76$. The tangent line identifies the slope only at that particular point, but it allows for easy interpretation. Its slope shows the linear change in the probability at a single point on the logistic curve.

The change in probability or the linear slope of the tangent line comes from a simple equation for the partial derivative. The partial derivative equals

$$\partial P/\partial X_{\mathrm{k}} = b_k \times P \times (1 - P).$$

Simply multiply the logistic regression coefficient by the selected probability P and 1 minus the probability.

The formula for the partial derivative nicely reveals the nonlinear effects of an independent variable on probabilities. The effect of b (in terms of logged odds) translates into a different effect on the probabilities depending on the level of P. The effect will be at its maximum when P equals .5 since $.5 \times .5 = .25$, $.6 \times .4 = .24$, $.7 \times .3 = .21$ and so on. The closer P comes to the ceiling or floor, the smaller the value $P(1 - P)$, and the smaller the effect a unit change in X has on the probability.

Multiplying the coefficient times $.5 \times .5$ shows the maximum effect on the probabilities, but may overstate the influence for a sample in which the split on the dependent variable is not so even. Substituting the mean of the dependent variable, P, in the formula gives a more typical effect. For smoking, the logistic regression coefficient for years of education equals −.183, and the mean of the dependent variable or the probability of smoking equals .154. The marginal change at the mean equals $-.183 \times .154 \times .846$

Figure 2.1 Tangent line of logistic curve at $Y = P = .76$.

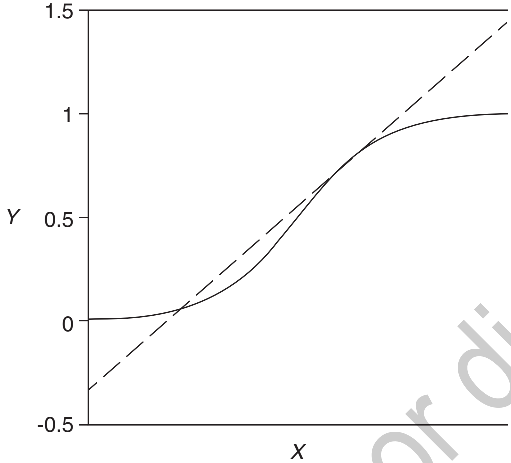

or −.024. A marginal or instantaneous change in education reduces the probability of smoking by .024. The effect reaches its maximum of −.046 when $P = .5$.

While this example illustrates the logic underlying marginal effects for continuous variables in logistic regression, it oversimplifies things. A question remains: At what values should the marginal effect be calculated? We want to calculate a marginal effect in a way that best represents the relationships for the sample. To do that, three types of marginal effects are commonly recommended (Breen, Karlson, & Holm, 2018; Long & Freese, 2014; Williams, 2012). There are marginal effects at the means, marginal effects at representative values, and average marginal effects (AME). Consider each in turn.

Three Types of Marginal Effects

First, the marginal effect at the means is calculated when all independent variables in the model take their mean value. The predicted probability is obtained from multiplying each logistic regression coefficient times the mean of the corresponding independent variable, summing the products and the intercept, and transforming the predicted logit into a predicted probability. Then, this predicted probability can be used with the formula for the partial derivative to calculate marginal effects for each continuous independent variable.

Table 2.5 Stata Output: Marginal Effects at Means From Logistic Regression Model of Current Smoking, NHIS 2017

```
Expression     : Pr(Smoker), predict()
dy/dx w.r.t.   : Education Age 1.Gender 2.Race 3.Race 4.Race
                 5.Race 1.Ethnicity
at             : Education       =     13.90932 (mean)
                 Age             =       54.294 (mean)
                 0.Gender        =     .5514588 (mean)
                 1.Gender        =     .4485412 (mean)
                 1.Race          =     .8100143 (mean)
                 2.Race          =     .1101909 (mean)
                 3.Race          =     .0113092 (mean)
                 4.Race          =     .0501976 (mean)
                 5.Race          =     .0182881 (mean)
                 0.Ethnicity     =     .8853107 (mean)
                 1.Ethnicity     =     .1146893 (mean)
```

	dy/dx	Delta-method Std. Err.	z	P>|z|	[95% Conf.	Interval]
Education	-.0216849	.0007684	-28.22	0.000	-.023191	-.0201788
Age	-.0028387	.0001353	-20.98	0.000	-.0031039	-.0025736
Gender						
Male	.030344	.0044637	6.80	0.000	.0215954	.0390926
Race						
African American	-.0098693	.0066399	-1.49	0.137	-.0228833	.0031447
Native American	.0388422	.0224991	1.73	0.084	-.0052551	.0829396
Asian American	-.0733998	.0072834	-10.08	0.000	-.0876749	-.0591247
Multiple Race	.0618207	.0190291	3.25	0.001	.0245243	.0991171
Ethnicity						
Hispanic	-.0924188	.0044527	-20.76	0.000	-.1011459	-.0836917

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

Marginal effects at the means are best done with program commands. In Stata, the margins command generates marginal effects at the means. Table 2.5 displays the output from a margins command following the logistic regression ("margins, dydx(*) atmeans"). For the moment, we can focus on the two continuous measures of education and age. The results show a marginal effect of education, or the effect for an infinitely small change in education is −.022 when all independent variables are at their mean values or for a hypothetical person who is average on all characteristics. This indicates that the expected probability of smoking decreases by .022 with a marginal change in education. Alternatively, one can view −.022 as the slope of the tangent line at the predicted probability when the independent variables are at their means. Note that the top of the table lists the means for each independent variable used in the calculations.

These marginal effects make for a more intuitive interpretation than logged odds or odds. The coefficients of −.022 for education and −.003 for age represent changes on a probability scale ranging from 0 to 1. Although interpretations still depend on the measurement units and variation of the independent variables, they rely on more familiar units for the dependent variable. Always remember, however, that these effects are specific to a predicted probability determined by the means of the independent variables.

Second, marginal effects at representative values use a typical case on the independent variables rather than the means of the independent variables. Note that the means do not typically refer to the characteristics of an actual person. In the NHIS sample, the mean education of 13.91 and the mean age of 54.29 are not observed values for the measures. The concern is more obvious for the categorical variables. The mean for gender is .449 (44.9% male) and the mean for African American is .110. Obviously, the NHIS categorical measures do not treat gender or race as a proportion. To represent a typical case, we might select values for someone with 12 years of education, 45 years old, male (Gender = 1), white (Race = 1), and non-Hispanic (Ethnicity = 0).

Margins in Stata can again be used to obtain these marginal effects ("margins, dydx(*) at (Education = 12 Age = 45 Gender = 1 Race = 1 Ethnicity = 0)"). Table 2.6 lists the values used to obtain the predicted probability and the marginal effects at those values. The predicted probability for a person with these representative values is .275, which is higher than the predicted probability of .137 when the independent variables are at their means. The expected change in the probability of smoking for an infinitely small change in education at these representative values is −.037 (vs. −.022 previously). Other models at different representative values could give substantially different results, however.

Third, the average marginal effect is obtained differently. It first calculates the marginal effect for each observation by using the actual values

Table 2.6 Stata Output: Marginal Effects at Representative Values From Logistic Regression Model of Current Smoking, NHIS 2017

```
Expression     :  Pr(Smoker), predict()
dy/dx w.r.t.   :  Education Age 1.Gender 2.Race 3.Race 4.Race
                  5.Race 1.Ethnicity
at             :  Education     =       12
                  Age           =       45
                  Gender        =        1
                  Race          =        1
                  Ethnicity     =        0
```

	dy/dx	Delta-method Std. Err.	z	P>|z|	[95% Conf.	Interval]
Education	-.0365248	.0016434	-22.23	0.000	-.0397457	-.0333038
Age	-.0047814	.0002657	-18.00	0.000	-.0053021	-.0042607
Gender						
Male	.0476225	.0069547	6.85	0.000	.0339915	.0612535
Race						
African American	-.0163434	.0111053	-1.47	0.141	-.0381095	.0054226
Native American	.0610315	.0339325	1.80	0.072	-.0054749	.127538
Asian American	-.1307434	.014079	-9.29	0.000	-.1583378	-.103149
Multiple Race	.0948479	.0273988	3.46	0.001	.0411472	.1485486
Ethnicity						
Hispanic	-.1568755	.0088443	-17.74	0.000	-.17421	-.139541

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

of each observation rather than the means or representative values. For an independent variable, there are as many marginal effects as observations in the analysis. It then calculates the average of those marginal effects. Note the difference in strategy. The marginal effects at the means and representative values define a single marginal effect for each independent variable. The average marginal effect defines a distribution of marginal effects for the sample and then computes the mean of the distribution for each independent variable.

Table 2.7 presents the AME obtained from the margins command in Stata ("margins, dydx(*)"). The AME are similar but not identical to those obtained at the means or representative values. Holding other covariates constant, the average change in smoking for an infinitely small change in education is −.023 (vs. −.022 and −.037 previously).

Each type of marginal effect has strengths and weaknesses (Muller & MacLehose, 2014). The marginal effects at the means represent central tendency and in that sense are typical of the sample. However, the means on all independent variables define a hypothetical example rather than a real person, group, organization, or other unit of analysis. The marginal effects at representative values are based on observed values of the independent variables but may not be typical of the sample. Although researchers will select key groups or characteristics that are common in the sample, the representative values miss those in other groups or with other

Table 2.7 Stata Output: Average Marginal Effects From Logistic Regression Model of Current Smoking, NHIS 2017

```
Expression    : Pr(Smoker), predict()
dy/dx w.r.t. : Education Age 1.Gender 2.Race 3.Race 4.Race
               5.Race 1.Ethnicity
```

	dy/dx	Delta-method Std. Err.	z	P>|z|	[95% Conf.	Interval]
Education	-.0225815	.0008187	-27.58	0.000	-.024186	-.020977
Age	-.0029561	.0001436	-20.59	0.000	-.0032375	-.0026747
Gender						
Male	.0315158	.0046228	6.82	0.000	.0224553	.0405763
Race						
African American	-.0102466	.0069154	-1.48	0.138	-.0238005	.0033073
Native American	.0396333	.0226697	1.75	0.080	-.0047986	.0840652
Asian American	-.0783194	.0079958	-9.80	0.000	-.0939909	-.0626479
Multiple Race	.0626203	.0189212	3.31	0.001	.0255354	.0997052
Ethnicity						
Hispanic	-.0988868	.0049608	-19.93	0.000	-.1086097	-.0891638

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

characteristics. The average marginal effect can be viewed as the effect for a case picked at random from the sample (Breen et al., 2018). It has the advantages of using all observed values in the sample and thereby representing everyone. Long and Freese (2014, Chapter 6) offer a qualified recommendation for the average marginal effect or AME: "Broadly speaking, we believe that the AME is the best summary of the effect of a variable."

Note that, even after selecting one of the marginal effects, any single coefficient showing the change in probability is potentially misleading. The coefficient will not fully reflect the nonlinear and nonadditive relationship of the independent variables with the probabilities. To get more information but also add more complexity, a researcher might compute marginal effects for a range of values on the independent variables and present the marginal effects for the extremes as well as the middle of the distribution. Ways to present a more complete summary of the range of influences of a variable on probabilities are discussed in more detail below.

Categorical Independent Variables

The partial derivative works best with continuous variables for which small changes in the independent variables are meaningful. For dummy variables, the relevant change occurs from 0 to 1, and the tangent line for infinitely small changes in *X* makes little sense. Instead, the marginal effect for categorical variables is best shown by the discrete change from one category to another. It is possible to compute predicted probabilities for two categories and then measure the difference in the probabilities. This marginal effect refers to a discrete change in the independent variable rather than a marginal change. The two may approximate one another, but calculating the predicted probabilities for categorical independent variables based on the discrete change makes more sense. Remember, however, that the marginal effect based on differences in predicted probabilities, like the partial derivative, varies across points on the logistic curve and specific predicted probabilities.

The three strategies for estimating marginal effects on probabilities for continuous independent variables apply to categorical independent variables. Each strategy finds a predicted probability for the omitted group, finds a predicted probability for the dummy variable group, and subtracts the former probability from the latter. The differences in strategies come from setting the values assigned to the other independent variables. The three most common strategies can be reviewed and adapted to categorical independent variables using the same model and examples in Tables 2.5 to 2.7.

First, the marginal effect at the means uses the predicted values for categorical variables when one value is 0, the other value is 1, and all other independent variables in the model take their mean values. The discrete change in the predicted values from 0 and 1 for the categorical comparison defines the marginal effect.

In Table 2.5, the Stata output notes at the bottom that the marginal effect (symbolized by dy/dx) for factor levels (values of a categorical independent variable) is the discrete change from the base level. For gender, the base level or reference category refers to females, and the marginal effect for males of .030 shows that the expected probability of smoking is higher for males than females by .030, when the other independent variables are held constant at their means. With education, age, time, and the marital status categories taking their mean values, additional runs show that the predicted probability of smoking is .1244 for females and .1547 for males. The difference of .0303 is the same as the gender coefficient for males in Table 2.5.

For race, the base level or reference category is white. The coefficients of −.010 for African Americans, .039 for Native Americans, −.073 for Asian Americans, and .062 for multi-race respondents show considerable diversity in smoking across the groups. The marginal effect is largest for Asians, with the difference in expected probabilities of −.073 compared to whites, when other independent variables are held constant at their means.

Second, marginal effects at representative values use a typical case on the independent variables rather than the means of the independent variables. To reexamine the discrete change for gender, the predicted value for both males and females can be obtained when the other independent variables are set, for example, at 12 years of education, 45 years old, white (Race = 1), and non-Hispanic (Ethnicity = 0). The results are shown in Table 2.6. The marginal effect for gender equals .048, which is larger than the marginal effect at the means. The marginal effect for Hispanics relative to non-Hispanics is −.157 at the selected representative values. As always, using different representative values can give substantially different results.

Third, the AME for the categorical variables are listed in Table 2.7. The average marginal effect calculates the predicted probability for each observation when gender equals 0 and the other independent variables equal their observed values. The same is done when gender equals 1 and the difference is obtained for each observation. The average of the differences equals the average marginal effect. Table 2.7, which presents the average marginal effect for all independent variables, shows coefficients of .032 for gender and −.099 for Hispanic ethnicity.

The same strengths and weaknesses in the three summary marginal effects of continuous independent variables apply to categorical independent variables.

The interpretation shifts from the change in probabilities due to an infinitely small change in a continuous independent variable to the discrete change from a base level to a group level for a categorical independent variable. Otherwise, the issues faced in interpreting probability effects still apply. Along with selecting the type of marginal effect, one should consider examining the marginal effects at varied levels of the independent variables. An easy way to do this is discussed below.

Graphing Marginal Effects

The three ways to summarize marginal effects—at the means, at representative values, and the average marginal effect—do not represent the variation of the effects around the average. A single coefficient is quite useful but incomplete. One way to capture this variation is to select a set of values at which to calculate marginal effects. For example, the marginal effect for a continuous variable can be calculated when the variable takes values –2, –1, 0, 1, and 2 standard deviations from its mean and the other independent variables are at their means, at representative values, or actual values in AME. The AME of education are −.038, −.030, −.022, −.015, and −.010, respectively, at –2, –1, 0, 1, and 2 standard deviations from its mean. The marginal effects get weaker at higher levels of education, as the probability of smoking gets lower. Alternatively, marginal effects might be calculated when a continuous independent variable takes its maximum, mean, and minimum values. The marginal effects of education are −.039 at 0, the education minimum, −.022 at the education mean, and −.013 at 18, the education maximum.

An easier procedure involves graphing the marginal effects. In Stata, the average marginal effect for each value of education again comes from the margins command ("margins, dydx(education) at(education = (0(1)18))") followed by "marginsplot." The results of this command are shown in Figure 2.2 The graph plots the average marginal effect on the *Y* axis for each observed value of education on the *X* axis. The bars represent the 95% confidence interval around the marginal effects. The scale for the marginal effects in the *Y* axis varies from only −.05 to −.01 and the marginal effects range from −.013 to −.043.

Note the nonlinear relationship of education with the probability of smoking. The graph reveals the strongest marginal effects for those with few years of completed schools, who also tend to smoke more than others. The peak negative effect of −.043 occurs with four years of education. The marginal effect moves toward 0, reaching −.013 at 18 years of education. The pattern of the marginal effects reaffirms the point made early that marginal effects are strongest when predicted probabilities are near .5, a level at which those with little education come closest. As smoking falls to levels farther from .5 and closer to 0 for those with advanced education, the negative marginal effect becomes weaker. Similar results could be obtained for marginal effects at the means or marginal effects at representative values.

The pattern of marginal effects in Figure 2.2 demonstrates the nonlinearity in probability effects. That is, the average marginal effect of education varies in a nonlinear pattern with the level of education. Graphs can also demonstrate nonadditivity, or how the marginal effect of one independent variable on probabilities varies with the level of another variable. Consider gender and age. Figure 2.3 shows the marginal effect from the discrete change in gender for each age. The marginal effect of being male is larger at younger ages when current smoking is higher. The marginal effect decreases at older ages, as many former smokers have quit or died. The marginal effect falls from .045 at age 26 to .018 at age 85 and over. The change is not large but nonetheless illustrates how the difference between men and women varies with age.

Figure 2.2 Average marginal effects of education at values of education from logistic regression model of current smoking, NHIS 2017.

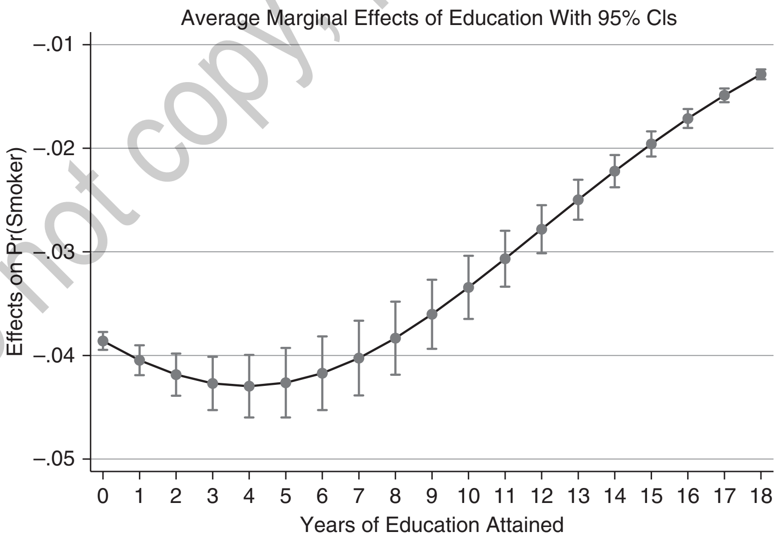

Figure 2.3 Average marginal effects of gender (1 = males) at values of age from logistic regression model of current smoking, NHIS 2017.

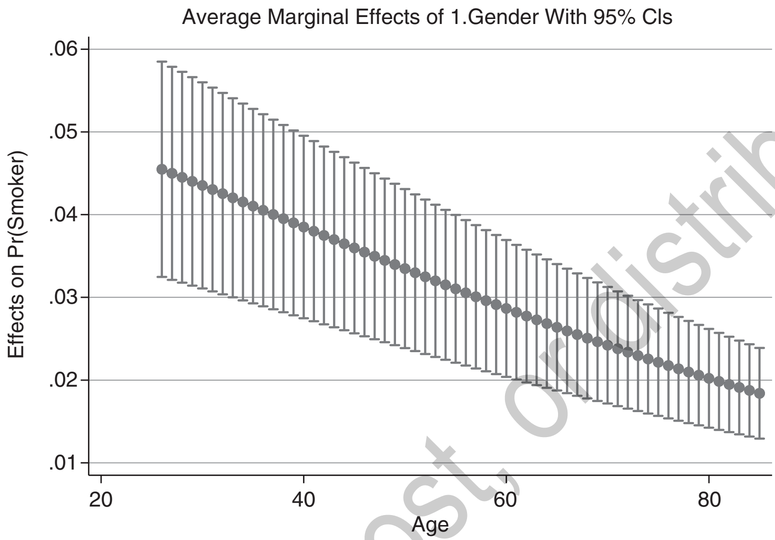

Figure 2.4 Average marginal effects of education at values of age from logistic regression model of current smoking, NHIS 2017.

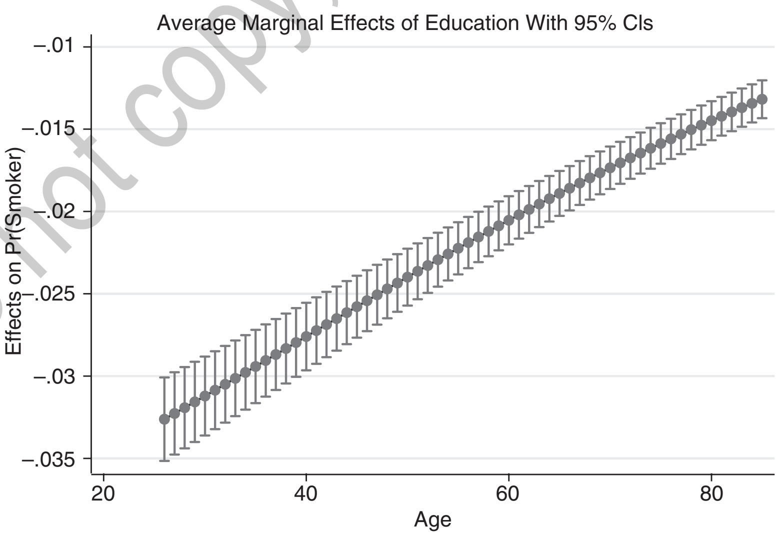

Nonadditivity also holds for continuous independent variables. Figure 2.4 plots the average marginal effect of education for selected ages. The negative marginal effect of education is strongest at the youngest ages, when current smoking is highest. It becomes weaker and closer to 0 at older ages when current smoking is lowest. Although the effect of education on the logged odds is the same across all ages, the effects on probabilities are nonadditive.

Graphing Predicted Probabilities

Marginal effects examine the change in predicted probabilities for an infinitely small change or discrete change in the independent variables. Some insight into the nature of marginal effects and, more generally, into the relationships of the independent variables with the dependent variable can come from examining the predicted probabilities themselves. Graphing of predicted probabilities, like graphing of marginal effects, proves helpful.

Figure 2.5 graphs the average predicted probabilities of men and women at selected ages (when the other independent variables take their observed values for each case). The predicted probabilities on the *Y* axis range from just above .05 to just below .3. The probability scale differs from the

Figure 2.5 Predicted probabilities by gender and age from logistic regression model of current smoking, NHIS 2017.

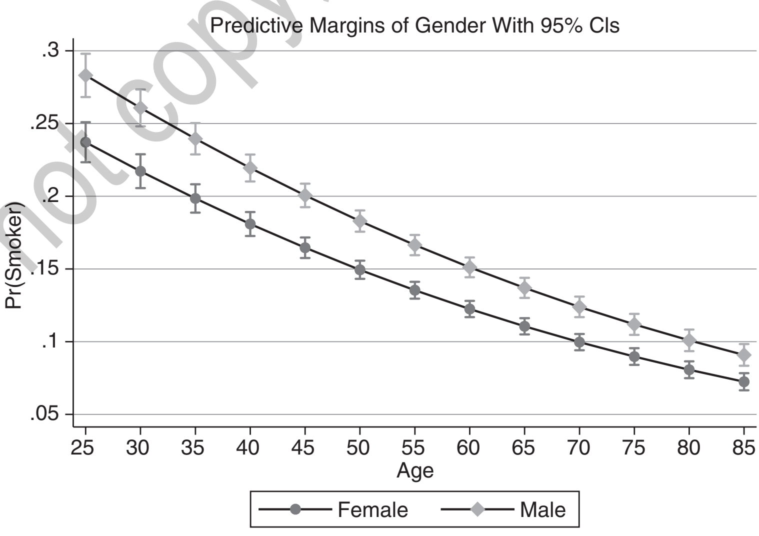

smaller scale for marginal effects. According to the graph, women have a lower probability of smoking than men. Furthermore, the gap between men and women, which reflects the average marginal effect of gender, varies with age. Although the change is not large, the gap between women and men appears narrower at older than younger ages. This affirms findings for the marginal effects of gender: Figure 2.3 shows these effects get smaller at older ages. The lines below again demonstrate the nonlinear and nonadditive relationships of gender and age with the predicted probabilities of smoking.

As one more example, Figure 2.6 presents the average predicted probabilities for non-Hispanics and Hispanics by education level. The line for non-Hispanics is higher than for Hispanics, but it declines more quickly with education. The gap thus narrows with education.

It is worth comparing the predicted probabilities from the logistic regression to the predicted probabilities from linear regression. Using the same independent variables as the logistic regression, the linear regression plus the margins command produces Figure 2.7. The predicted probabilities form a straight line, and the gap between non-Hispanics and Hispanics across years of education is constant. In other words, the marginal effects of Hispanic ethnicity and education are both linear and additive. The two

Figure 2.6 Predicted probabilities by Hispanic ethnicity and education from logistic regression model of current smoking, NHIS 2017.

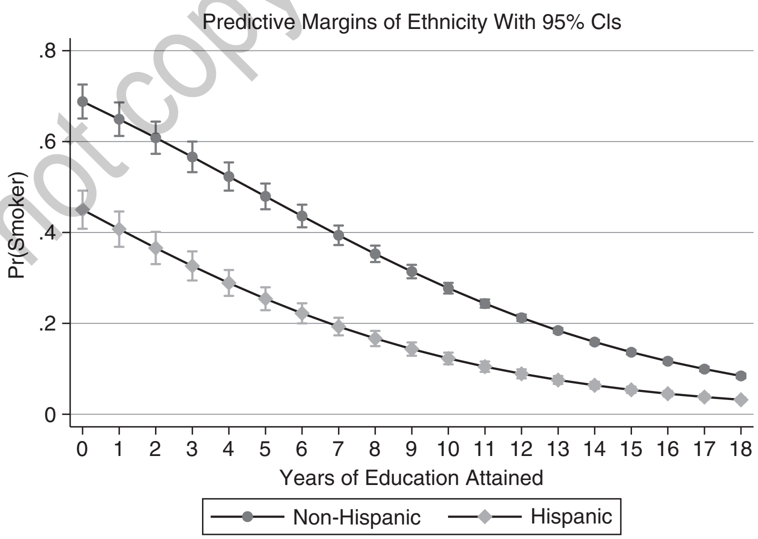

Figure 2.7 Predicted probabilities by Hispanic ethnicity and education from linear regression model of current smoking, NHIS 2017.

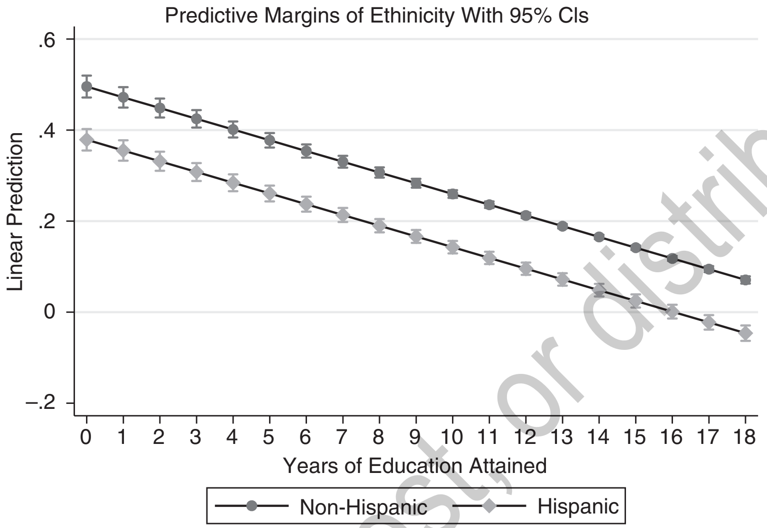

straight and parallel lines from the linear regression in Figure 2.7 contrast with the curved and nonparallel lines from the logistic regression in Figure 2.6. The linear regression shows the same gap at all levels of education, while the logistic regression shows a larger gap at low education than at high education. Reflecting this difference, the logistic regression further shows higher predicted smoking for non-Hispanics at low levels of education than the linear regression. Although the differences between the graphs are not huge, the one for the logistic regression appears more accurate.

Standardized Coefficients

Regression programs ordinarily present standardized coefficients along with unstandardized coefficients. Many find standardized coefficients to be helpful in interpreting regression results. Unstandardized coefficients show relationships between variables measured in their original metric. If the measurement units differ, as is typically the case, the unstandardized coefficients are not directly comparable. For the model of current smoking, a unit change in gender (from females to males) obviously differs from a 1-year change in completed education. Even comparisons of education and age are risky, despite both being measured in terms of years. For the sample, age varies from 26 to 85 and has a standard deviation of 16.6. In contrast, education varies from 0 to 18, with a standard deviation of 2.85. A 1-year change has a different meaning for these two variables.

Standardized coefficients have the advantage of showing relationships when the independent and dependent variables have a common scale. They can be understood as regression coefficients when all variables are measured as standard scores with means of 0 and standard deviations of 1. They then show the expected change in standard units of the dependent variable for a standard unit change in the independent variables, controlling for other independent variables. Given the comparable units of the variables, standardized coefficients help in comparing the relative strength of the relationships.

Unlike multiple regression programs, logistic regression programs do not routinely compute standardized coefficients. The problem with standardized coefficients in logistic regression stems partly from ambiguity in the meaning of standard scores or standard units for a binary dependent variable. Standardizing a binary variable merely translates values of 0 and 1 into two other values. If the mean of the dependent variable *Y* equals the probability *P*, the variance equals $P(1 - P)$ (Agresti, 2013, p. 117). Then, the standard score *z* has only two values:

$$Y \text{ values of 1 have } z \text{ values equal to } (1-P)/\sqrt{P(1-P)},$$

and

$$Y \text{ values of 0 have } z \text{ values equal to } (0-P)/\sqrt{P(1-P)}.$$

With only two values, a standardized binary dependent variable does not represent variation in the underlying probability of the outcome.

Standardizing the Independent Variables

One way around the problem involves semistandardizing coefficients in logistic regression (also called *X*-standardizing). Standardizing only the independent variables does not require a standard deviation for the binary dependent variable, but it still allows for some useful comparisons. The semistandardized coefficients show the expected change in the logged odds of the outcome associated with a standard deviation change in each of the independent variables. With a comparable metric for the independent variables, semistandardized coefficients reflect the relative importance of variables within a model.

Semistandardized coefficients can be obtained in one of two ways. First, they come from multiplying by hand the logistic regression coefficient for independent variables in their original metric by the standard deviation of the variables. The formula is simple and has some intuitive value in understanding how standardizing the independent variables works:

$$b_{yx\ semi} = b_{yx\ unstand} \times sd_x.$$

If the unstandardized coefficient shows the change in the logged odds of a unit change in *X*, then multiplying the coefficient by the standard deviation shows the change in the logged odds for a standard deviation change in *X*.

For example, the logistic regression coefficient for education in its original units is −.183 and the standard deviation of education is 2.85. The semistandardized coefficient equals:

$$-.183 \times 2.85 = -.522.$$

The logistic regression coefficient for age is −.024, but the standard deviation of 16.6 is larger. The semistandardized coefficient is

$$-.024 \times 16.6 = -.398.$$

It indicates that age (−.398) is less strongly associated with smoking than education (−.522).

Second, to obtain semistandardized coefficients directly, the independent variables can be standardized before being included in the logistic regression model. The resulting logistic regression coefficients will show the effects on the logged odds of a standard deviation change in each of the independent variables. Table 2.8 lists the results using this procedure. First note that, after standardizing the independent variables using the means and standard deviations of the variables for the sample used in the logistic regression analysis, all independent variables will have a mean of 0 and a standard deviation of 1. Next note that the logistic regression coefficients shown below for education and age match (within rounding error) the calculations done by hand. However, the list of all semistandardized coefficients together makes for easy comparison of the size of the coefficients. It can be seen that education has the strongest effect, followed by age and Hispanic ethnicity.

The interpretation of the outcome is still in logged odds, but taking the exponent of the coefficients will show the change in the odds for a standard deviation in the independent variables.

Standardizing the Independent Variables and Dependent Variable

Semistandardized coefficients, while simple to understand and calculate, have a limitation. They identify the effects of different independent variables for the same dependent variable and within the same model. However,

Table 2.8 Stata Output: Semi-Standardized Coefficients From Logistic Regression Model of Current Smoking With Standardized Independent Variables, NHIS 2017

Smoker	Coef.	Std. Err.	z	p>|z|	[95% Conf. Interval]	
zEducation	-.5212059	.0192999	-27.01	0.000	-.559033	-.4833787
zAge	-.3980747	.0195263	-20.39	0.000	-.4363455	-.3598038
zGender	.126119	.0183976	6.86	0.000	.0900603	.1621777
zAfricanAmer	-.0261479	.0180046	-1.45	0.146	-.0614363	.0091406
zNativeAmer	.0304768	.0161067	1.89	0.058	-.0010917	.0620454
zAsianAmer	-.1768843	.0238549	-7.41	0.000	-.2236391	-.1301295
zMultiRace	.0584639	.0158138	3.70	0.000	.0274695	.0894584
zEthnicity	-.3312603	.0223123	-14.85	0.000	-.3749915	-.287529
_cons	-1.837902	.0200325	-91.75	0.000	-1.877165	-1.798639

the scale for the dependent variable remains as logged odds, and comparing the effects of independent variables across models with different dependent variables can be misleading. Broader comparisons are possible, as in linear regression, with fully standardized coefficients. A fully standardized coefficient (also called *XY*-standardized coefficient) adjusts for the standard deviations of both *X* and *Y* as in the following formula:

$$b_{yx\,full} = b_{yx\,unstand} \times (sd_x / sd_y).$$

However, the problem of how to obtain the standard deviation of *Y* remains.

To obtain a meaningful measure of the standard deviation of a binary dependent variable, Long (1997, pp. 70–71) recommends using the predicted logits from the model. This approach is based on the idea that the observed binary values of an outcome in logistic regression are manifestations of an underlying latent continuous variable. This latent continuous variable is assumed to have a variance but one that is unobserved. However, the predicted logged odds from logistic regression have an observed variance that will reflect the underlying unobserved variance. In addition, the error term in the logistic regression equation has a variance, arbitrarily defined in the logistic distribution as $\pi^2/3$. Together, the variance of the predicted logits plus the variance of the error term offers an estimate of the variance of the unobserved continuous dependent variable. Taking the square root of the variance provides a measure of the standard deviation of the continuous latent variable. Using this standard deviation for the latent variable in the formula for the standardized coefficient will show that the standard deviation change in the logged odds for a one standard deviation unit change in the independent variables.

In practice, the variance of predicted logged odds can be obtained by saving the predicted logged odds and requesting descriptive statistics. For the model of smoking, the variance of the predicted logged odds is .410, and $\pi^2/3$ is 3.290. The sum is 3.700 and the square root of 3.700 is 1.924. To get the fully standardized coefficient for education, substitute values into the formula above:

$$-.183 \times (2.85/1.924) = -.271.$$

Although tedious and prone to error when done by hand, the same steps can be used to obtain fully standardized coefficients for the other independent variables. For Stata users, an easier ways to get the coefficients is available.

A Note on SPOST

Scott Long and Jeremy Freese (2014) have created a suite of programs called SPOST that can be used with Stata to interpret the coefficients obtained from nonlinear models. SPOST is free to download by Stata users. The authors' book presents a full introduction to the SPOST commands and their uses, and one of the many commands calculates semistandardized and fully standardized coefficients.

SPOST uses a simple command ("listcoef, std") after a logistic regression command. Table 2.9 presents the output from the SPOST command. It lists the coefficients, *z* values, and probabilities of *z* from the usual logistic regression output. The column labeled bStdX lists the semistandardized or *X*-standardized coefficients. These coefficients show the expected change in the logged odds for a one-standard deviation change in the independent variables. They match the coefficients created with the independent variables as standard scores but are easier to obtain. The standard deviations used to obtain these *X*-standardized coefficients are listed in the last column.

Standardizing both the dependent and independent variables gives the fully standardized or *XY*-standardized coefficients in the column labeled bStdXY. The fully standardized coefficient for education shows that a one standard deviation change in education is associated with an expected decrease in current smoking by −.271 standard deviations. Note that the standard deviation for the dependent variable refers to the underlying latent

Table 2.9 SPOST Output: Standardized Coefficients From Logistic Regression Model of Current Smoking, NHIS 2017

```
Observed SD : 0.3609
Latent SD   : 1.9236
```

	b	z	P>|z|	bStdX	bStdY	bStdXY	SDofX
Education	-0.1831	-27.006	0.000	-0.521	-0.095	-0.271	2.847
Age	-0.0240	-20.387	0.000	-0.398	-0.012	-0.207	16.610
Gender							
Male	0.2536	6.855	0.000	0.126	0.132	0.066	0.497
Race							
African American	-0.0835	-1.452	0.146	-0.026	-0.043	-0.014	0.313
Native American	0.2882	1.892	0.058	0.030	0.150	0.016	0.106
Asian American	-0.8101	-7.415	0.000	-0.177	-0.421	-0.092	0.218
Multiple Race	0.4363	3.697	0.000	0.058	0.227	0.030	0.134
Ethnicity							
Hispanic	-1.0396	-14.847	0.000	-0.331	-0.540	-0.172	0.319
constant	2.0539	16.095	0.000	.	.	.	.

variable. These coefficients again show the strongest effects from education, age, and Hispanic ethnicity.

The output lists other information that is helpful in understanding logistic regression. It is possible to standardize on *Y* but not *X*, just as it is possible to standardize on *X* but not *Y*. The table lists these *Y*-standardized coefficients. Also, the top of the table lists the observed standard deviation of the binary outcome of current smoking as .361, but for reasons mentioned above, that number is of limited value. The latent standard deviation of 1.924 is more meaningful. It refers to the underlying, unobserved distribution of smoking that produces the observed binary outcomes of 0 and 1.[9]

Caution is warranted in interpreting the meaning of a standard deviation change for a categorical independent variable, which is less straightforward than the meaning of a standard deviation change for a continuous independent variable. And comparison of standardized coefficient across groups is generally not recommended. But standardized coefficients are well suited for the interpretation of the relative strength of relationships within a single model and group.

Group and Model Comparisons of Logistic Regression Coefficients

Researchers often test theories and hypotheses by making comparisons across groups and models. The comparison across groups might come from estimating the same models for two or more groups, such as males and females or young, middle-aged, and older persons. The group comparisons are typically tested using interaction or moderation terms. The comparison across models might come from sequentially adding confounding or mediating independent variables to a model for the same group and sample and comparing the coefficients across the models. The effect of gender on an outcome might be first shown without controls and then examined with controls for education and work status. Given the use of the same sample, the first model is nested within the second, and the coefficients of an independent variable change from the first to the second model.

Although these strategies are appropriate with multiple regression, they are generally not appropriate with logistic regression coefficients. The editors of *American Sociological Review* (Mustillo, Lizardo, & McVeigh, 2018), in presenting guidelines for authors submitting articles, say the following: "don't use the coefficient of the interaction term to draw conclusions about statistical interaction in categorical models such as logit, probit, Poisson, and so on." The same problem noted for interaction terms applies to comparisons of coefficients across separate groups and across nested models. Comparing the size of logit coefficients or odds ratios or using tests of significance for differences in the size of the logit coefficients or odds ratios presents challenges for interpretation.

The sources of the problems are beyond the scope of this book, but several articles describe them in detail (Allison, 1999; Breen et al., 2018; Mood, 2010; Williams, 2009). Very briefly, comparisons must assume that the errors are the same across the multiple groups and models, but the errors for the logits are not known. Comparing coefficients across different groups and models, which may have different but unknown error variances, can be misleading. Breen et al. (2018) offer an overview of approaches that avoid or minimize the problem of comparing coefficients in logistic regression and related models. Possible solutions include *Y*-standardizing or fully *XY*-standardizing the coefficients, focusing on marginal probability effects, and using linear regression with robust standard errors. More complex solutions involve specifying additional model constraints when comparing coefficients across different groups (Allison, 1999; Williams, 2009) or residualizing the added control

Table 2.10 Stata Output: Logistic Regression Model of Current Smoking With the Interaction Between Hispanic Ethnicity and Education, NHIS 2017

Smoker	Coef.	Std. Err.	z	p>|z|	[95% Conf. Interval]	
Education	-.2233305	.0078442	-28.47	0.000	-.2387048	-.2079562
Ethnicity						
Hispanic	-3.115895	.2149463	-14.50	0.000	-3.537182	-2.694608
Ethnicity #c.Education		.				
Hispanic	.1802465	.0168571	10.69	0.000	.1472071	.2132858
Age	-.0243334	.0011835	-20.56	0.000	-.0266531	-.0220138
Gender						
Male	.25177	.0371321	6.78	0.000	.1789925	.3245476
Race						
African American	-.1187599	.0579308	-2.05	0.040	-.2323022	-.0052175
Native American	.2494087	.1524716	1.64	0.102	-.0494302	.5482476
Asian American	-.8317383	.1108915	-7.50	0.000	-1.049082	-.614395
Multiple Race	.4151453	.1182646	3.51	0.000	.183351	.6469396
_cons	2.622645	.139376	18.82	0.000	2.349473	2.895817

variables when comparing coefficients across nested models for the same sample (Breen et al., 2018).

To illustrate one approach, consider an example of group comparisons using interaction terms and interpretations based on the marginal effects on probabilities. A model allows education and Hispanic ethnicity to interact in the model of smoking, as presented in Table 2.10. The positive and significant interaction term indicates that the negative effect of Hispanic ethnicity is smaller at higher levels of education. The size of the Hispanic ethnicity logged odds coefficient should not be compared across education levels, but marginal effects on probabilities are appropriate for analysis (Long & Mustillo, 2018). Figure 2.8 shows the average marginal effect of Hispanic ethnicity as a discrete change across years of completed education as implied by the logistic regression model. As the graph shows, the negative marginal effect is smaller at higher levels of education.

Figure 2.8 Average marginal effects of Hispanic ethnicity at values of education from the logistic regression model of current smoking with the interaction between Hispanic ethnicity and education, NHIS 2017.

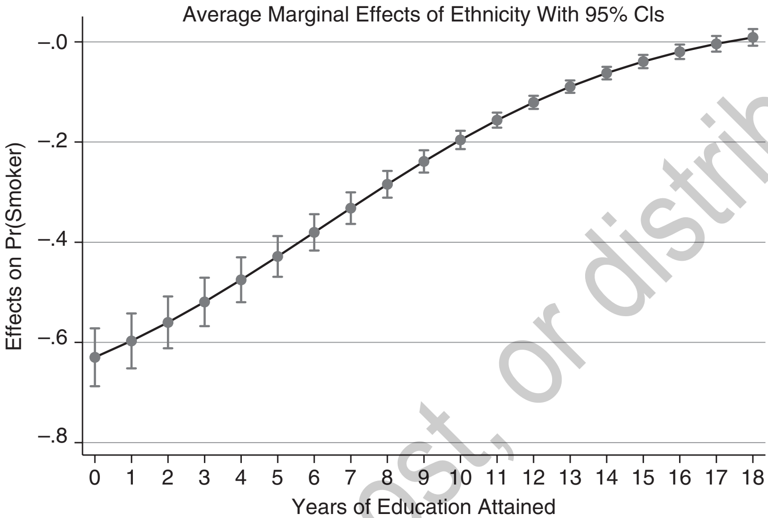

Despite its value, the single graph simplifies the complexities involved with interactions in logistic regression. Interpreting the probability effects in general requires care and thoroughness—a recommendation that is doubly important for models with interactions.

Summary

Logistic regression coefficients provide a simple linear and additive summary of the influence of a variable on the logged odds of having a characteristic or experiencing an event, but they lack an intuitively meaningful scale of interpretation of change in the dependent variable. Standard tests of significance offer another common way to interpret the results, but by themselves say little about the substantive meaning of the coefficients. Raising e to the coefficient b allows interpretation of the resulting coefficient in terms of multiplicative odds or percentage change in the odds. For still more intuitive coefficients, the marginal effects of independent variables on the probability of an outcome are helpful. However, effects on probabilities depend on the values of the independent variables at which the effect is calculated. Calculating standardized coefficients may help but warrant some caution given difficulties in standardizing the binary outcome. Making comparisons across groups or, equivalently, using statistical interaction terms in logistic regression models similarly warrants caution, as do comparisons across nested models. Care and thoroughness are needed to fully understand the relationships being modeled in logistic regression.