The Stata Journal (2004)
4, Number 2, pp. 154–167

Computing interaction effects and standard errors in logit and probit models

Edward C. Norton
Department of Health Policy and Administration
University of North Carolina at Chapel Hill
Hua Wang
Department of Health Policy and Administration
University of North Carolina at Chapel Hill
Chunrong Ai
Department of Economics
University of Florida and Tsinghua University, China

Abstract. This paper explains why computing the marginal effect of a change in two variables is more complicated in nonlinear models than in linear models. The command `inteff` computes the correct marginal effect of a change in two interacted variables for a logit or probit model, as well as the correct standard errors. The `inteff` command graphs the interaction effect and saves the results to allow further investigation.

Keywords: st0063, inteff, interaction terms, logit, probit, nonlinear models

1 Introduction

Applied researchers often estimate interaction terms to infer how the effect of one independent variable on the dependent variable depends on the magnitude of another independent variable. For example, is the effect of car weight on gas mileage the same for both domestic and foreign cars? To answer this question, we can run a regression to predict gas mileage as a function of weight, a dummy variable for foreign, and the interaction between the two. If the coefficient on the interaction term is statistically significant, there is a difference between domestic and foreign cars in how additional weight affects mileage.

Interaction terms are also used extensively in nonlinear models, such as logit and probit models. Unfortunately, the intuition from linear regression models does not extend to nonlinear models. The marginal effect of a change in both interacted variables is not equal to the marginal effect of changing just the interaction term. More surprisingly, the sign may be different for different observations. The statistical significance cannot be determined from the z-statistic reported in the regression output. The odds-ratio interpretation of logit coefficients cannot be used for interaction terms.

Despite the common use of interaction terms, most applied researchers misinterpret the coefficient of the interaction term in nonlinear models. A review of 13 economics journals listed on JSTOR (*www.jstor.org*) found 72 articles published between 1980 and

 st0063

2000 that used interaction terms in nonlinear models (Ai and Norton 2003). None of the studies interpreted the coefficient on the interaction term correctly. A recent article by DeLeire (2000) is a welcome exception.

The Stata command `inteff` computes the correct marginal effect of a change in two interacted variables for a logit or probit model. It also computes the correct standard errors. The `inteff` command will work if the interacted variables are both continuous variables, if both are dummy variables, or if there is one of each. In addition, it will graph the interaction effect and save the results to allow further investigation.

2 Estimation of interaction effects

2.1 Linear models

In linear models, the interpretation of the coefficient of the interaction between two variables is straightforward. Let the continuous dependent variable y depend on two independent variables x_1 and x_2, their interaction, and a vector of additional independent variables X, including the constant term. The expected value of the dependent variable, conditional on the independent variables, is

$$E\left[y|x_1,x_2,X\right] = \beta_1 x_1 + \beta_2 x_2 + \beta_{12} x_1 x_2 + X\beta$$

where the βs are unknown parameters. If X is independent of x_1 and x_2, then the interaction effect of the independent variables x_1 and x_2 is β_{12} for both continuous and discrete interacted variables. The statistical significance of the interaction effect can be tested with a single t test on the coefficient β_{12}.

2.2 Nonlinear models

The intuition from linear models, however, does not extend to nonlinear models. To illustrate, consider a probit model similar to the previous example, except that the dependent variable y is a dummy variable. The conditional mean of the dependent variable is

$$\begin{aligned} E\left[y|x_1,x_2,X\right] &= \Phi\left(\beta_1 x_1 + \beta_2 x_2 + \beta_{12} x_1 x_2 + X\beta\right) \\ &= \Phi\left(u\right) \end{aligned}$$

where Φ is the standard normal cumulative distribution and u denotes the index $\beta_1 x_1 + \beta_2 x_2 + \beta_{12} x_1 x_2 + X\beta$. Suppose that x_1 and x_2 are continuous. The marginal effect of just the interaction term $x_1 x_2$ is

$$\frac{\partial \Phi\left(u\right)}{\partial\left(x_1 x_2\right)} = \beta_{12}\Phi'\left(u\right)$$

Most applied researchers interpret this as the interaction effect. However, the full interaction effect is the cross-partial derivative of the expected value of y

$$\frac{\partial^2 \Phi\left(u\right)}{\partial x_1 \partial x_2} = \beta_{12}\Phi'\left(u\right) + \left(\beta_1 + \beta_{12} x_2\right)\left(\beta_2 + \beta_{12} x_1\right)\Phi''\left(u\right)$$

This equation shows clearly that the interaction effect is not equal to $\beta_{12}\Phi'(u)$.

There are four important implications of this equation for nonlinear models. First, the interaction effect could be nonzero, even if $\beta_{12} = 0$. For example, for a probit model with $\beta_{12} = 0$, the interaction effect is

$$\left.\frac{\partial^2 \Phi(u)}{\partial x_1 \partial x_2}\right|_{\beta_{12}=0} = \beta_1 \beta_2 \Phi''(u)$$

Second, the statistical significance of the interaction effect cannot be tested with a simple t test on the coefficient of the interaction term β_{12}. Instead, the statistical significance of the entire cross derivative must be calculated. Third, the interaction effect is conditional on the independent variables, unlike the interaction effect in linear models. (It is well known that the marginal effect of a single, uninteracted variable in a nonlinear model is conditional on the independent variables.) Fourth, because there are two additive terms, each of which can be positive or negative, the interaction effect may have different signs for different values of covariates. Therefore, the sign of β_{12} does not necessarily indicate the sign of the interaction effect.

In summary, for nonlinear models to compute the magnitude of the interaction effect, one must compute the cross derivative of the expected value of the dependent variable. The test for the statistical significance of the interaction effect must be based on the estimated cross-partial derivative, not on the coefficient of the interaction term. The main objective of this paper is to introduce a Stata command that will calculate the correct interaction effect and standard errors for logit and probit models.

Stata's `mfx` and `dprobit` commands are useful for estimating the marginal effect of a single variable, given specific values of the independent variables. However, these commands should never be used when a variable is interacted with another or has higher order terms. In those cases, `mfx` and `dprobit` will estimate the wrong marginal effect. The `mfx` and `dprobit` commands do not know if a variable is interacted with another or has higher order terms, so they cannot take the full derivative with respect to that variable. Also, the results of the `mfx` and `dprobit` commands are misleading because the marginal effects for any nonlinear model differ for each observation, yet `mfx` and `dprobit` only report one marginal effect per variable. The examples in this paper will drive home the point that there is a distribution of magnitudes of marginal effects, often with opposite signs.

The Stata command `predictnl` can be used to derive all the results found with `inteff`. However, because `predictnl` is so general, allowing for nonlinear predictions after any Stata estimation command, the user must be able to write correct formulas of the marginal effects in vector notation. We believe that because `inteff` is easier to use than `predictnl`, it will lead to fewer user errors and encourage more researchers to calculate correct interaction effects.

2.3 General formulas

In a nonlinear model, the dependent variable is a nonlinear function $F(u)$ of the index of independent variables. For example, in the logit and probit models, the dependent variable of interest, F, is the probability that $y = 1$. For logit and probit models, define the *interaction effect* to be the change in the predicted probability that $y = 1$ for a change in both x_1 and x_2.

When the interacted variables are both continuous, the interaction effect is the double derivative with respect to x_1 and x_2:

$$\begin{aligned}\frac{\partial^2 F(u)}{\partial x_1 \partial x_2} &= \frac{\partial \left\{(\beta_1 + \beta_{12} x_2) f(u)\right\}}{\partial x_2} \\ &= \beta_{12} f(u) + (\beta_1 + \beta_{12} x_2)(\beta_2 + \beta_{12} x_1) f'(u)\end{aligned}$$

where $f(u) = F'(u)$ and $f'(u) = F''(u)$.

When the interacted variables are both dummy variables, the interaction effect is the discrete double difference:

$$\begin{aligned}\frac{\Delta^2 F(u)}{\Delta x_1 \Delta x_2} &= \frac{\Delta \left\{F(\beta_1 + \beta_2 x_2 + \beta_{12} x_2 + X\beta) - F(\beta_2 x_2 + X\beta)\right\}}{\Delta x_2} \\ &= F(\beta_1 + \beta_2 + \beta_{12} + X\beta) \\ &\quad - F(\beta_1 + X\beta) - F(\beta_2 + X\beta) + F(X\beta)\end{aligned}$$

When one continuous variable and one dummy variable are interacted, the interaction effect is the discrete difference (with respect to x_2) of the single derivative (with respect to x_1):

$$\begin{aligned}\frac{\Delta \frac{\partial F(u)}{\partial x_1}}{\Delta x_2} &= \frac{\Delta \left\{(\beta_1 + \beta_{12} x_2) f(u)\right\}}{\Delta x_2} \\ &= (\beta_1 + \beta_{12}) f\left\{(\beta_1 + \beta_{12}) x_1 + \beta_2 + X\beta\right\} - \beta_1 f(\beta_1 x_1 + X\beta)\end{aligned}$$

Ai and Norton (2003) derive the standard errors for the interaction effect in logit and probit models, applying the Delta method. For the case of two dummy variables, the asymptotic variance of the estimated interaction effect is estimated consistently by

$$\frac{\partial}{\partial \beta'} \left\{\frac{\Delta^2 F(u)}{\Delta x_1 \Delta x_2}\right\} \widehat{\Omega}_\beta \frac{\partial}{\partial \beta} \left\{\frac{\Delta^2 F(u)}{\Delta x_1 \Delta x_2}\right\}$$

where $\widehat{\Omega}_\beta$ is a consistent covariance estimator of $\widehat{\beta}$. For continuous variables, we replace the discrete difference operator Δ with the partial derivative operator.

This paper focuses on the most common type of interactions, those between two variables. The correct interpretation for a model with three interacted variables requires taking three derivatives (or three discrete differences), following on the logic of the previous section.

The interpretation is also complicated if, in addition to being interacted, a variable has higher order terms—for example, if age squared is included in addition to age and age interacted with marital status.

For all these more complicated models, the principle is the same: take derivatives or discrete differences.

2.4 Logit formulas

For the logit model, $F(u)$ is the familiar logit cumulative distribution function:

$$F(u) = \frac{1}{1 + e^{-(\beta_1 x_1 + \beta_2 x_2 + \beta_{12} x_1 x_2 + X\beta)}}$$

When the interacted variables are both continuous, the interaction effect is the cross derivative with respect to x_1 and x_2:

$$\begin{aligned}\frac{\partial^2 F(u)}{\partial x_1 \partial x_2} &= \beta_{12} \{F(u)(1 - F(u))\} \\ &\quad + (\beta_1 + \beta_{12} x_2)(\beta_2 + \beta_{12} x_1) \Big[F(u) \big\{1 - F(u)\big\}\big\{1 - 2F(u)\big\}\Big]\end{aligned}$$

When the interacted variables are both dummy variables, the interaction effect is the discrete double difference:

$$\begin{aligned}\frac{\Delta^2 F(u)}{\Delta x_1 \Delta x_2} &= \frac{1}{1 + e^{-(\beta_1 + \beta_2 + \beta_{12} + X\beta)}} \\ &\quad - \frac{1}{1 + e^{-(\beta_1 + X\beta)}} - \frac{1}{1 + e^{-(\beta_2 + X\beta)}} + \frac{1}{1 + e^{-X\beta}}\end{aligned}$$

When one continuous variable and one dummy variable are interacted, the interaction effect is the discrete difference (with respect to x_2) of the single derivative (with respect to x_1):

$$\begin{aligned}\frac{\Delta \frac{\partial F(u)}{\partial x_1}}{\Delta x_2} &= (\beta_1 + \beta_{12}) \left(\begin{array}{c} F\big\{(\beta_1 + \beta_{12})x_1 + \beta_2 + X\beta\big\} \\ \times \big(1 - F\big\{(\beta_1 + \beta_{12})x_1 + \beta_2 + X\beta\big)\big\} \end{array} \right) \\ &\quad - \beta_1 \Big[F(\beta_1 x_1 + X\beta) \big\{1 - F(\beta_1 x_1 + X\beta)\big\}\Big]\end{aligned}$$

2.5 Probit formulas

For the probit model, $F(\cdot)$ is the familiar normal, cumulative distribution function

$$F(u) = \Phi(\beta_1 x_1 + \beta_2 x_2 + \beta_{12} x_1 x_2 + X\beta)$$

When the interacted variables are both continuous, the interaction effect is the double derivative with respect to x_1 and x_2:

$$\frac{\partial^2 F(u)}{\partial x_1 \partial x_2} = \big\{\beta_{12} - (\beta_1 + \beta_{12} x_2)(\beta_2 + \beta_{12} x_1) u\big\} \phi(u)$$

When the interacted variables are both dummy variables, the interaction effect is the discrete double difference:

$$\begin{aligned}\frac{\Delta^2 F(u)}{\Delta x_1 \Delta x_2} &= \Phi(\beta_1+\beta_2+\beta_{12}+X\beta)\\ &\quad -\Phi(\beta_1+X\beta)-\Phi(\beta_2+X\beta)+\Phi(X\beta)\end{aligned}$$

When one continuous variable and one dummy variable are interacted, the interaction effect is the discrete difference (with respect to x_2) of the single derivative (with respect to x_1):

$$\begin{aligned}\frac{\Delta\frac{\partial F(u)}{\partial x_1}}{\Delta x_2} &= (\beta_1+\beta_{12})\,\phi\big\{(\beta_1+\beta_{12})x_1+\beta_2+X\beta\big\}\\ &\quad -\beta_1\phi(\beta_1 x_1+X\beta)\end{aligned}$$

2.6 Odds ratio

Many researchers, especially epidemiologists, prefer to fit logit models than probit models because of the odds-ratio interpretation of the logit coefficients. Before explaining why this interpretation does not work for a model with interaction terms, we review the derivation of odds ratios. The *odds* are the ratio of a probability p to one minus the probability:

$$\begin{aligned}p &= \frac{1}{1+e^{-X\beta}}\\ \text{odds} &= \frac{p}{1-p}=\frac{1}{e^{-X\beta}}=e^{X\beta}\in[0,\infty)\end{aligned}$$

The *odds ratio* is the ratio of odds for two different observations that differ only in the value of one explanatory variable. This is easiest to understand for a dummy variable. Consider the probability of smoking, which depends on whether the person is female, as well as on many explanatory variables (X). The odds ratio for gender (holding all other variables constant) is the odds for female (female = 1) divided by the odds for male (female = 0):

$$\begin{aligned}\text{odds for female} &= \frac{\Pr(\text{smoke}|\text{female})}{1-\Pr(\text{smoke}|\text{female})}=e^{\beta_f\,\text{female}+X\beta}\\ \text{odds for male} &= \frac{\Pr(\text{smoke}|\text{male})}{1-\Pr(\text{smoke}|\text{male})}=e^{X\beta}\\ \text{odds ratio} &= \frac{\text{odds for female}}{\text{odds for male}}=e^{\beta_f}\end{aligned}$$

Even though the odds ratio is difficult to understand conceptually (it is the ratio of ratios, and honestly, who understands that?), it is widely used for two reasons. First, it is easy to calculate, requiring only the exponentiation of the estimated coefficient, and can be reported directly by Stata. Second, when p is small, the odds ratio is a good approximation to the risk ratio, which is easy to understand conceptually. The *risk ratio* is the ratio of two probabilities. For example, the risk ratio for the smoking example is the probability of smoking for women divided by the probability of smoking for men:

$$\text{risk ratio} = \frac{\Pr(\text{smoke}|\text{female})}{\Pr(\text{smoke}|\text{male})}$$

If the risk ratio equals 1.5, for example, women are fifty percent more likely to smoke than men, holding all other variables constant. Unfortunately, researchers often interpret odds ratios as if they were risk ratios, even when p is not close to zero and the approximation is not close. For an example of how odds ratios can be misreported by researchers and the media, see commentary by Schwartz, Woloshin, and Welch (1999).

Now consider the odds ratio when there is an interaction between two dummy variables, x_1 and x_2. The common interpretation is that the odds ratio for the interaction term equals $\exp(\beta_{12})$. This is not true. The expression $\exp(\beta_{12})$ is the ratio of odds ratios:

$$\begin{aligned}
\text{odds ratio for } x_1|_{x_2=1} &= \frac{\frac{\Pr(y=1|x_1=1;x_2=1)}{1-\Pr(y=1|x_1=1;x_2=1)}}{\frac{\Pr(y=1|x_1=0;x_2=1)}{1-\Pr(y=1|x_1=0;x_2=1)}} \\
&= \frac{e^{\beta_1+\beta_2+\beta_{12}+X\beta}}{e^{\beta_2+X\beta}} \\
\text{odds ratio for } x_1|_{x_2=0} &= \frac{\frac{\Pr(y=1|x_1=1;x_2=0)}{1-\Pr(y=1|x_1=1;x_2=0)}}{\frac{\Pr(y=1|x_1=0;x_2=0)}{1-\Pr(y=1|x_1=0;x_2=0)}} \\
&= \frac{e^{\beta_1+X\beta}}{e^{X\beta}} \\
\text{ratio of odds ratios for } x_1 \text{ and } x_2 &= e^{\beta_{12}}
\end{aligned}$$

Not only is $\exp(\beta_{12})$ not a risk ratio, it is not even an odds ratio.

3 Syntax

The new command `inteff` calculates the interaction effect, standard error, and z-statistic for each observation for either logit or probit when two variables have been interacted. The interacted variables cannot have higher order terms, such as squared terms. The command is designed to be run immediately after fitting a logit or probit model.

`inteff` *varlist* [`if` *exp*] [`in` *range*] [`, savedata(`*filename*[`, replace`]`)` `savegraph1(`*filename*[`, replace`]`)` `savegraph2(`*filename*[`, replace`]`)`] where *varlist* must be the same as the fitted logit or probit model and must include at least four variables. The order of these first four variables must be dependent variable, independent variable 1, independent variable 2, and the interaction between independent variables 1 and 2. Other independent variables can be added after the interaction term, i.e., starting from the fifth position.

If the interaction term (at the fourth position) is a product of a continuous variable and a dummy variable, the first independent variable x_1 has to be the continuous variable, and the second independent variable x_2 has to be the dummy variable. The order of the second and third variables does not matter if both are continuous or both are dummy variables.

4 Options

`savedata(`*filename*[`, replace`]`)` specifies the path and filename of computed data to be saved. This gives the researcher the option of further investigation. Saved data include five variables, in the following order:

1. predicted probability
2. interaction effect (calculated by conventional linear method)
3. interaction effect (calculated by the method suggested in this paper)
4. standard error of the interaction effect
5. z-statistic of the interaction effect

The variables all have meaningful names. For example, after we run a logit model, the five variables would be `_logit_phat`, `_logit_linear`, `_logit_ie`, `_logit_se`, and `_logit_z`. The prefix for probit models is `_probit`.

`savegraph1(`*filename*[`, replace`]`)` and `savegraph2(`*filename*[`, replace`]`)` save the graphs with the name and path designated by the user. The `inteff` command generates two scatter graphs. Both plot predicted probabilities on the x-axis. The first graph plots two interaction effects (one is calculated by the method suggested in this paper, and the other one is calculated by the conventional linear method) against predicted probabilities. The second graph plots z-statistics of the interaction effect against predicted probabilities.

5 Examples

5.1 Data

We illustrate the use of the `inteff` command with two examples. Both examples analyze data from the 2000 Medical Expenditure Panel Survey (MEPS), which can be used to compute nationally representative estimates of health care use and expenditures. These examples are intended to be illustrative. These data are available to the public from the Agency for Healthcare Research and Quality web site (*www.meps.ahrq.gov*).

The sample includes all adults age 21–64 who have complete information on all variables. Out of the original sample of 25,096 people, 12,365 meet these criteria, with most exclusions due to the age restriction.

The dependent variable is whether the person had an office-based physician visit during the calendar year 2000. In this adult sample, two-thirds did have an office-based physician visit, and one-third did not.

The mean age is 42, 47 percent are male, 64 percent are married, 81 percent are white, 14 percent are black, and 4 percent are neither white nor black. The average number of years of education is just over 12, and we also control for household income categories defined relative to the poverty level. The vast majority have some form of private health insurance, with 8 percent covered by public insurance and 18 percent uninsured. We control for severe health-status problems, defined as problems with activities of daily living, instrumental activities of daily living, vision, or hearing. To control for broad geographic differences in access to health care, we include variables for the four census regions of the country and whether the person lives in a metropolitan statistical area (78 percent).

5.2 Logit with two continuous variables interacted

The first example includes the interaction between age and number of years of education, both continuous variables. In this example, we fitted a logit model, although the results for probit would be virtually identical. The standard errors are adjusted for clustering on person id (`pid`). The full model also controls for race, marital status, income, health status, and geographic region (summarized by the global variable `$x`), but the results for these variables are not reported for brevity.

```
. logit $y age educ ageeduc male ins_pub ins_uni $x, nolog cluster(pid)
Logit estimates                                   Number of obs   =      12365
                                                  Wald chi2(23)   =    9745.78
                                                  Prob > chi2     =     0.0000
Log pseudo-likelihood = -6889.3644                Pseudo R2       =     0.1189
                              (standard errors adjusted for clustering on pid)
------------------------------------------------------------------------------
             |               Robust
    opvisits |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         age |   .0419025   .0070434     5.95   0.000     .0280977    .0557073
        educ |    .127117   .0236365     5.38   0.000     .0807903    .1734437
     ageeduc |  -.0013739   .0005168    -2.66   0.008    -.0023869   -.0003609
        male |  -.9765431   .0348741   -28.00   0.000    -1.044895    -.908191
     ins_pub |   .5829237   .1043102     5.59   0.000     .3784794     .787368
     ins_uni |  -.8781526   .0541354   -16.22   0.000     -.984256   -.7720491
  (output omitted)
       _cons |  -1.559739   .3379041    -4.62   0.000    -2.222019   -.8974595
------------------------------------------------------------------------------
```

After running this logit model, we invoke the `inteff` command with the same variable list, and save the data.

```
. inteff $y age educ ageeduc male ins_pub ins_uni $x,
> savedata(d:\data\logit_inteff,replace) savegraph1(d:\data\figure1, replace)
> savegraph2(d:\data\figure2, replace)
Logit with two continuous variables interacted
file d:\data\logit_inteff.dta saved
(file d:\data\figure1.gph saved)
(file d:\data\figure2.gph saved)
```

Variable	Obs	Mean	Std. Dev.	Min	Max
_logit_ie	12365	-.0003334	.0001145	-.0005798	.0001607
_logit_se	12365	.0001003	.0000311	4.81e-06	.000323
_logit_z	12365	-3.401374	1.245229	-6.228868	7.130231

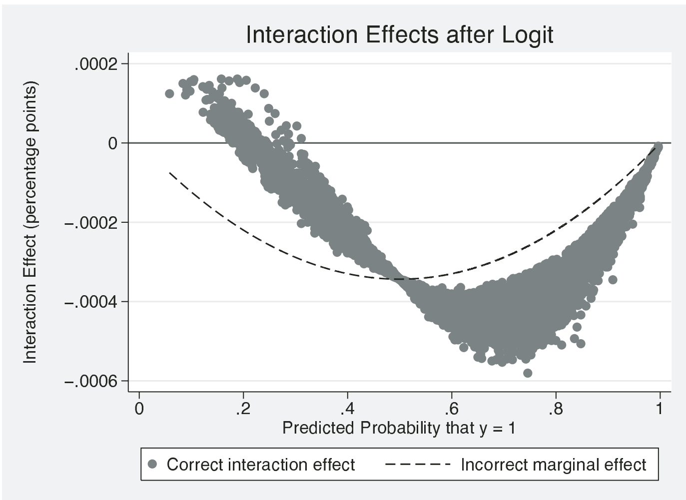

Figure 1.

(*Continued on next page*)

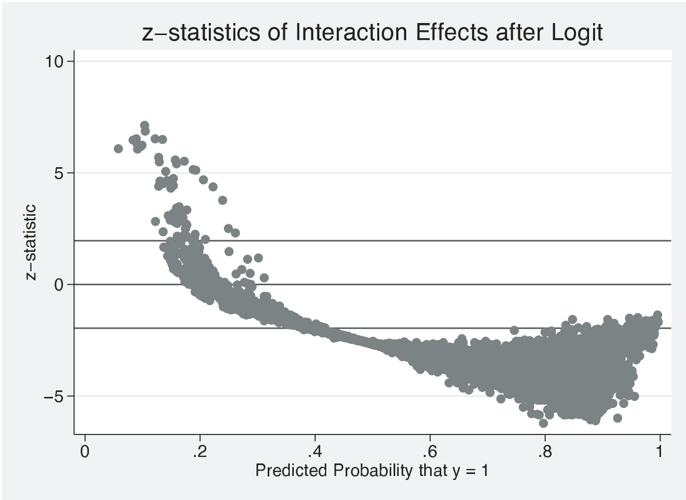

Figure 2.

Not only are both age and number of years of education statistically significant at conventional levels, their interaction is also (z-statistic is −2.66). The main effects imply that persons who are older and have more years of education are more likely to have an office-based visit during the year. However, after running the `inteff` command, we learn that the mean interaction effect is negative (−.0003334) and varies widely. For some observations, the interaction effect is positive, and for others, it is negative (see figures 1 and 2).

The interaction effect depends on other covariates. In this example, for people whose predicted probability of having a physician visit is around 0.2 (toward the left end of figure 1), the interaction effect between age and education is positive for half of them and negative for the other half. If we look at the right side of figure 1, where people have a predicted probability of having a physician visit around 0.8, their interaction effects are all negative. In terms of the significance of the interaction effects, for the left group of people whose predicted probability is about 0.2, only a few have statistically significant interaction effects. On the other hand, for the right group of people whose predicted probability is around 0.8, the interaction effects are mostly significant.

5.3 Probit with two dummy variables interacted

The second example includes the interaction between gender and insurance status. The effect of gender on having an office-based physician visit may depend on having insurance. Many young men choose not to purchase health insurance, figuring that they are unlikely to need medical care, while many women of child-bearing age are uninsured but ineligible for Medicaid or other public insurance. In this example, we fit a probit model. As before, the standard errors are adjusted for clustering on person id (`pid`), and the full model also controls for additional variables, but these are not reported for brevity.

```
. probit $y male ins_uni male_uni age educ ins_pub $x, nolog cluster(pid)
Probit estimates                                  Number of obs   =      12365
                                                  Wald chi2(23)   =    9391.46
                                                  Prob > chi2     =     0.0000
Log pseudo-likelihood =  -6897.391                Pseudo R2       =     0.1179
                              (standard errors adjusted for clustering on pid)

                             Robust
    opvisits       Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]

        male   -.5922717   .0240826   -24.59   0.000    -.6394726   -.5450707
     ins_uni   -.5653776   .0482592   -11.72   0.000    -.6599639   -.4707912
    male_uni    .0539772   .0605362     0.89   0.373    -.0646716    .1726261
         age    .0146619   .0012311    11.91   0.000     .0122489    .0170749
        educ    .0407643   .0047575     8.57   0.000     .0314397    .0500888
     ins_pub    .3275272   .0608681     5.38   0.000     .2082279    .4468265
  (output omitted)
       _cons   -.4692864   .0889194    -5.28   0.000    -.6435652   -.2950076
```

After running this probit model, we invoke the `inteff` command with the same variable list, and save the data.

```
. inteff $y male ins_uni male_uni age educ ins_pub $x,
> savedata(d:\data\probit_inteff, replace)
> savegraph1(d:\data\figure3, replace) savegraph2(d:\data\figure4, replace)
Probit with two dummy variables interacted
file d:\data\probit_inteff.dta saved
(file d:\data\figure3.gph saved)
(file d:\data\figure4.gph saved)
    Variable         Obs        Mean    Std. Dev.        Min         Max

   _probit_ie      12365   -.0092839    .0294776  -.0578116    .0829161
   _probit_se      12365    .0218298    .0023465   .0046057    .0314373
    _probit_z      12365   -.5169928    1.522319  -5.561593    5.530833
```

(*Continued on next page*)

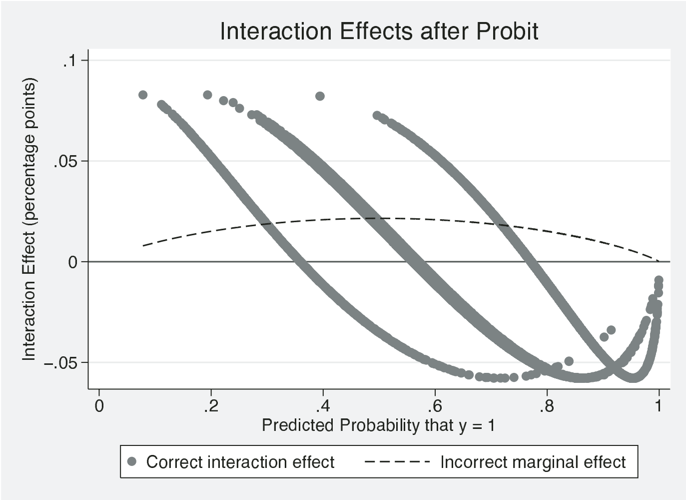

Figure 3.

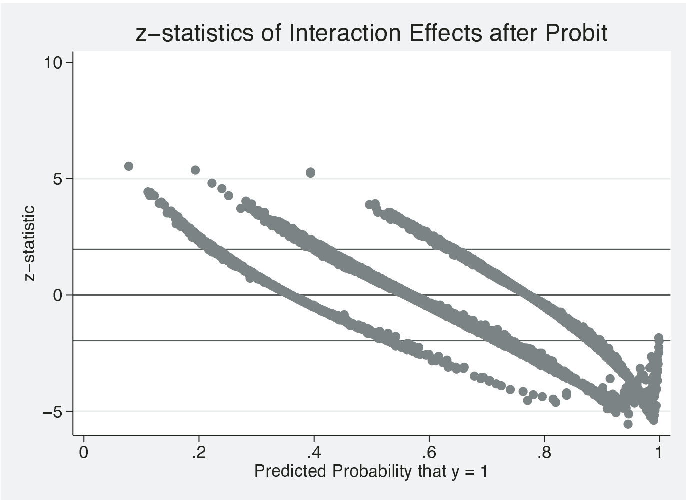

Figure 4.

In this example, both the main terms are highly statistically significant. However, unlike before, the interaction between male and uninsured is not statistically significant. In a linear model, we could conclude from such results that the interaction effect is essentially zero. However, in the nonlinear probit, we see that the magnitude and statistical significance ranges widely. Despite the lack of statistical significance of the coefficient on the interaction term, the full interaction effect is large and statistically significant for many observations (see figures 3 and 4). This shows that only looking at the table of results can be misleading.

6 Conclusion

Interaction effects are complicated to compute and interpret in nonlinear models. Because of their widespread use, however, having a command to compute them is important for applied researchers. The new command `inteff` allows users to compute the magnitude, sign, and statistical significance of interaction effects in logit and probit models.

The results of the two examples are typical of the patterns we have found after computing interaction effects for a wide range of problems. The interaction effect has a wave shape when plotted against predicted values. Some interaction effects are positive, and some are negative, no matter what the sign of the coefficient on the interaction term. For predicted values equal to .5, the interaction effect is $\beta_{12}\Phi'(u)$ for the probit case. There is wide variation in the statistical significance of the interaction effect.

There are two limitations to the `inteff` command. One is that the code will only work for logit and probit models, even though the issue applies to all nonlinear models, such as tobit and count models. In addition, the command will only work for the interaction between two variables that do not also have higher order terms. For example, the command would yield the wrong answer if, in the first example, age squared was also included as an independent variable. For other nonlinear models, interactions between more than two variables or interactions of variables with higher-order terms use the Stata command `predictnl` with great care.

7 References

Ai, C. R. and E. C. Norton. 2003. Interaction terms in logit and probit models. *Economics Letters* 80(1): 123–129.

DeLeire, T. 2000. The wage and employment effects of the Americans with Disabilities Act. *Journal of Human Resources* 35(4): 693–715.

Schwartz, L. M., S. Woloshin, and H. G. Welch. 1999. Misunderstandings about the effects of race and sex on physicians' referrals for cardiac catheterization. *Journal of the American Medical Association* 341(4): 279–283.

About the Authors

Edward C. Norton is an Associate Professor in the Department of Health Policy and Administration at the University of North Carolina at Chapel Hill. His research on health economics is in the areas of long-term care, econometrics, mental health, and managed care.

Hua Wang is a Ph.D. student in the Department of Health Policy and Administration. Her concentration is in economics.

Chunrong Ai is an Associate Professor in the Department of Economics at the University of Florida. He research is on econometrics, microeconomics, international economics, and industrial organization.