How to interpret coefficients and calculate marginal effects in Discrete Choice Models

Xiang Ao

March 26, 2008

1 Binary Response Models

1.1 Models

In a linear model

$$y_i = X_i'\beta + \epsilon,$$

$$E(y_i) = X_i'\beta$$

When y_i is binary (meaning it takes two values, 1 or 0), We can also do an OLS regression on it, but a more popular way is to use a non-linear model. The most popular choices for modeling binary response are logit model and probit model. Both of them use the same idea: use a link function to map the binary variable into a continuous variable which is a linear function of the predictors.

Suppose p_t is the probability that $y_i = 1$. Then

$$p_i = E(y_i) = g^{-1}(X_i'\beta),$$

where $g^{-1}()$ is a function that maps from the linear predictor to y_i, which is often called index function or inverse link function.

This says that the expectation of y_i (equivalently, p_i) is not a linear function of x_i's, but by using the link function, we are still able to model a transformed p_i as a linear function of x_i's:

$$g(p_i) = X_i'\beta.$$

$g^{-1}()$ can be a normal CDF then we have a probit model:

$$g^{-1}(z) = \Phi(z) = \int_{-\inf}^{z} \frac{1}{\sqrt{2\pi}\exp[-(t^2/2)]} dt,$$

or it can be a logit:

$$g^{-1}(z) = \Lambda(z) = e^z/(1 + e^z).$$

1.2 Marginal Effects

Often times we are interested in the marginal effects of the predictors. In a linear model, it's easy: the coefficient on x_1 means when x_1 increase by one unit, how much will y increase.

In a binary response model,

$$\frac{\partial p}{\partial x} = \frac{\partial g^{-1}(X_i'\beta)}{\partial x}$$

In the case of probit:

$$\frac{\partial p}{\partial x} = \frac{\partial \Phi(X_i'\beta)}{\partial x} = \phi(X_i'\beta)\beta,$$

where $\phi()$ is the density function of normal distribution.

In the case of logit:

$$\frac{\partial p}{\partial x} = \frac{\partial \Lambda(X_i'\beta)}{\partial x} = \gamma(X_i'\beta)\beta,$$

where $\gamma()$ is the density function for logistic distribution:

$$\gamma(X_i'\beta) = \Lambda(X_i'\beta)(1 - \Lambda(X_i'\beta)) = p(1-p).$$

Therefore for logit, it's easier to calculate the marginal effect from coefficient estimate β, all you have to do is to plug in sample proportion mean $\bar{p}$ and coefficient estimate $\hat{\beta}$:

$$\frac{\hat{\partial p}}{\partial x} = \bar{p}(1-\bar{p})\hat{\beta}$$

The logit model can also be formulated as

$$\log(\frac{p}{1-p}) = \mathbf{X_i'}\beta \tag{1}$$

which says the logarithm of the odds (the ratio of the two probabilities) is equal to $\mathbf{X_i'}\beta$. Therefore,

$$p = \frac{\exp(X_i'\beta)}{1+\exp(\mathbf{X}_t\beta)} = \Lambda(X_i'\beta) \tag{2}$$

Often times we have discrete valued predictors, such as gender, or other dummy variables. In that case, it makes more sense to ask what's the effect of gender being 1 vs. 0 (female vs. male). What this question is can be formulated as:

$$E(p|x_1 = 1) - E(p|x_1 = 0).$$

In empirical analysis, it's often calculated by calculating the predicted probability by setting x_1 to 1 and by setting x_1 to 0. Then calculate the difference.

2 Count Data Models

In a count data model such as Poisson regression model or Negative Binomial model, we are modeling log of the expected counts:

$$\log(\mathrm{E}(y)) = \mathbf{X_i'}\beta \tag{3}$$

Therefore the marginal effect of x_1, for example, on $\mathrm{E}(y)$ is

$$\frac{\partial \log \mathrm{E}(y)}{\partial x} = \hat{\beta},$$

which means β is the marginal effect of x on log of the expected counts. Suppose the expected counts at x is μ_x, then

$$\frac{(\partial \mu_x)/\mu_x}{\partial x} = \hat{\beta},$$

which says that β represents the marginal effect of percentage change in μ_x with respect to x.

Another way to formulate this is to use the Incidence Rate Ratio Interpretation. Suppose the expected counts at x is μ_x, and expected counts at $x+1$ is μ_{x+1}.

$$\log(\mu_x) = x'\beta.$$

Therefore,

$$\mu_x = \exp(x'\beta), \tag{4}$$

and

$$\frac{\mu_{x+1}}{\mu_x} = \exp(\beta). \tag{5}$$

This says that the exponentiated β is the incidence rate ratio of the expected counts, if x increases by 1.

To calculate the marginal effect of x on y (or $\mathrm{E}(y)$), we need to calculate

$$\frac{\partial \mu_x}{\partial x} = \hat{\beta}\mu_x = \hat{\beta}\exp(X_i'\hat{\beta}),$$

if we evaluate μ_x at predicted value.

3 How to calculate marginal effects in Stata

Stata's built-in command for marginal effect is mfx, executed after the regression command (such as logit or poisson, for example). Scott Long's spost is a set of commands also used to calculate marginal effects, as well as other functionalities.

There are different options for mfx to calculate different marginal effects. For example, to calculate marginal effects for panel data negative binomial models in Stata, we should use the command

mfx, predict(nu0)

after the xtnbreg command, instead of mfx, which will calculate the linear prediction marginal effects.