

[Description](#)[Remarks and examples](#)[References](#)[Also see](#)

Description

After you fit a model using one of the ERM commands, you can generally interpret the coefficients in the usual way. You can also use `margins` to produce counterfactuals, derivatives, contrasts, potential outcomes, treatment effects, and effects for any type of change in the covariates.

In this entry, we discuss how to interpret coefficients, how to use `margins`, and how to use `predict`. We demonstrate how this works for a simple linear model, and we discuss how the same `margins` and `predict` commands work for nonlinear and random-effects models.

Remarks and examples

Remarks are presented under the following headings:

[*Use margins*](#)[*Endogenous covariates*](#)[*How to interpret coefficients*](#)[*How to use and interpret margins*](#)[*How to use margins in models without endogenous covariates*](#)[*How to use margins with endogenous covariates*](#)[*margins with predict\(asf\)*](#)[*margins with predict\(fixedasf\)*](#)[*When to use which*](#)[*Using margins with nonlinear and random-effects models*](#)[*Advanced options: Using margins predict\(base\(\)\) and predict\(fix\(\)\)*](#)

Use margins

In many models, including many that can be fit by the ERM commands, the coefficients have a structural interpretation. An example of a structural interpretation is, “What would we expect to happen to our dependent variable if we increased the value of a covariate by one unit for everyone in the population?” So long as there are no endogenous covariates in the main equation and your model is correctly specified, the coefficients from all models fit by ERM commands have a structural interpretation. For linear models you fit with `eregress`, `eintreg`, `xteregress`, and `xteintreg`, the coefficients almost always still have a structural interpretation, even if the model has endogenous covariates.

What can we do if our main equation does have endogenous covariates and we are not using a linear estimator? Use `margins`, which automatically produces estimates of margins, derivatives, and effects that fully account for any endogenous covariates. Even if you have a linear model where the coefficients have a structural interpretation, we recommend you still use `margins`. The default results from `margins` will always have a structural interpretation, and you do not need to worry whether your model meets the criteria that allow the coefficients to be interpreted.

Here is a simple example. If you fit a model,

```
. eregress y x1 x2
```

and then type

```
. margins, dydx(x1)
```

`margins` will return an estimate of the expected change in y given an instantaneous unit change in x_1 . And, because the model is linear, this estimate, its standard error, and confidence intervals will all match those from the coefficient on x_1 in the results of `eregress`.

With linear models, this result holds even if we believe x_1 is endogenous. If we type

```
. eregress y x1 x2, endogenous(x1 = z1, nomain)
```

followed by

```
. margins, dydx(x1)
```

the results from `margins` will again match those for the coefficient x_1 from `eregress`.

Things change if we are interested in a result that is not linear in the coefficient estimates. Consider the following:

```
. eprobit y x1 x2, endogenous(x1 = z1, nomain)
```

If we are interested in the average change in the probability of observing $y = 1$ in the population for an instantaneous unit change in x_1 , and we use only the coefficients on x_1 and x_2 , we would not account for the contribution $e_i \cdot y$. Prior to the instantaneous change, the probability is

$$\Pr(y_i = 1) = \Pr(\beta_0 + \beta_1 x_{1i} + \beta_2 x_{2i} + e_i \cdot y > 0)$$

Say we computed the probability as $\Phi(\beta_0 + \beta_1 x_{1i} + \beta_2 x_{2i})$. This will not be at the proper starting point on the nonlinear normal cumulative density curve, because we do not account for the correlation of the unobservables and the covariates. You do not need to worry about this when you use ERM postestimation commands. If you type

```
. margins, dydx(x1)
```

probability computations account for $e_i \cdot y$ and its correlation with the covariates. The effects produced have a structural interpretation. That is to say, they have an interpretation as though we could change values in the population and observe their effect. So the bottom line is use `margins`. Assuming your model is correctly specified, the default results from `margins` will always have a structural interpretation. See [Blundell and Powell \(2003\)](#), [Imbens and Newey \(2009\)](#), and [Wooldridge \(2010\)](#) for a detailed description of structural functions with linear and nonlinear endogenous models.

If you want to understand more about why and how, keep reading.

Endogenous covariates

Sometimes, care must be taken interpreting the coefficients from models when your main equation contains endogenous covariates. Endogenous covariates in the main equation cause problems, which means that if your model has no endogenous covariates in the main equation, you have no problems. The following models have no endogenous covariates in the main equation:

```
. eregress y x1 x2
. eregress y x1 x2 c.x1#c.x2
. eregress y x1 x2, select(selected = x1 z1 z2)    ///
    endogenous(z2 = z3 z4, nomain)

. xtprobit y x1 x2
. xtprobit y x1 x2 c.x1#c.x2
. xtprobit y x1 x2, select(selected = x1 z1 z2)    ///
    endogenous(z2 = z3 z4, nomain)
```

We showed examples with `eregress` and `xtprobit`. We could just as well have shown examples with any of the other ERM commands. Note that the last model for each command we showed has an endogenous covariate, but it is *not* in the main equation.

In any case, if you have no endogenous covariates in the main equation, you interpret coefficients and use `margins` and `predict` just as you usually would. If you have endogenous covariates in `eprobit` or `eoprobit`, do not attempt to interpret your coefficients. Rather, use `margins`, and interpret the results as you normally would.

In the rest of the manual entry, when we write about models with or without endogenous covariates, we mean models with or without endogenous covariates in the main equation.

Models with endogenous covariates in the main equation require care in interpretation, even if you fit a model as simple as

```
. eregress y x1, endogenous(x1 = z1, nomain)
```

There are four ways endogenous covariates can end up in the main equation:

1. You specify `endogenous(x1 = ...)` to add variable `x1` to the main equation.
2. You specify `endogenous(x1 = ..., nomain)`, *and* you include `x1` in the main equation.
3. You specify `entreat(treated = ...)` to handle endogenous treatment effects. `entreat()` itself adds endogenous covariate `treated` to the main equation.
4. You specify `select(selected = ...)` to handle endogenous selection, *and* you include `selected` in the main equation. `select()` makes variable `selected` endogenous, but it does not automatically add it to the main equation.

In what follows, we will show examples of endogenous covariates added to the main equation by option `endogenous()`, but we could have added them in any of the above ways.

In this manual entry, we depart from our usual practice of naming exogenous covariates `x1`, `x2`, ... and naming endogenous covariates `w1`, `w2`, ... We depart from this practice because we will introduce a situation and then say, “if `x1` is exogenous, do this; if it is endogenous, do something else”.

How to interpret coefficients

For `eregress`, `eintreg`, `xteregress`, and `xteintreg`, you can almost always interpret your coefficients in the usual way. This is true even if your model has endogenous covariates in the main equation. What do we mean by “the usual way”?

Say you are interested in the effect of covariate `x1`. Whether you have typed

```
. eregress y1 x1 x2
```

or

```
. eregress y1 x1 x2, endogenous(x1 = z1, nomain)
```

or even

```
. eintreg y1 y2 x1 x2, endogenous(x1 = x2 z1, nomain) ///
                        endogenous(z1 = x2 z2, nomain) ///
                        select(selected = x2 z3 z4)
```

you will have fit a model where

$$y1_i = \dots + \beta_1 x1_i + \dots$$

You interpret the fitted coefficient β_1 as the change in `y1` for a one-unit change in `x1`. That is true whether `x1` is an exogenous or an endogenous covariate. That interpretation sounds obvious, but we will see cases later where we must be more specific about the questions we ask regarding changes to `x1`.

Even if `x1` is interacted with another covariate, you still interpret the coefficients in the usual way. Say you have the model

```
. eregress y1 x1 c.x1#c.x2 x2
```

you will have fit a model where

$$y1_i = \dots + \beta_1 x1_i + \beta_2 x1_i \times x2_i + \dots$$

So a one-unit change in `x1` leads to a $\beta_1 + \beta_2 x2$ change in `y1`. Again, this is true whether `x1` is exogenous or endogenous.

We said you can “almost always interpret your coefficients in the usual way”. When can you not? You cannot interpret them in the usual way when all the following are true:

1. The covariate you are trying to interpret is endogenous or is an endogenous treatment.
2. If the covariate is endogenous, it is either binary or ordinal and is so declared in the `endogenous()` option using suboption `probit` or `oprobit`.
3. That covariate is in the main equation.
4. There is a second endogenous covariate in the main equation.
5. You have designated that each level (category) of the covariate you are interpreting has a different outcome error variance. Or you have designated that the correlation of the outcome error with the other endogenous errors varies by the levels of the covariate you are interpreting. You specify these cases by adding suboption `povariance` or suboption `pocorrelation` to the equation for the endogenous covariate of interest.

Whew! We did say that you could “almost always interpret your coefficients in the usual way”.

Here is one way to specify such a model,

```
. eregress y1 y2 x1 x2, endogenous(x1 = x2 z1, probit povariance nomain) ///
    endogenous(x2 = z2, nomain)
```

The coefficient on x_2 can be interpreted in the usual way. The coefficient on x_1 cannot. Why not? The conditional-on- x_2 expectation for y_1 depends on the conditional-on- x_2 expectation of the error for y_1 . Because there is a different error variance when $x_1 = 0$ and when $x_1 = 1$, their expectation no longer cancels out when we take the expected value of the effect. That's the "intuitive" answer. Were we conditioning on the observed value of x_1 in the effect (evaluating the treatment effect on the treated), we would have the same situation. The expectation of the errors would not cancel out. See [Treatment](#) in [eregress](#) for the full mathematical explanation.

For all other models, the best approach is to use `margins`. That should give you comfort, not concern—`margins` is a clear and safe way to form inferences and to measure and test effects. In fact, feel free to use `margins` rather than the coefficients even in regressions where you can "interpret your coefficients in the usual way". `margins` will give you exactly the same answers that you will get by looking at the coefficients. `margins` also makes it easy to ask what happens if you increase x_1 by 100, rather than by 1. Or to ask what happens if you give each person an additional 100 units of x_1 beyond his or her current endowment. In models with interactions or models with treatments, such questions can be tedious to answer from the coefficients.

To be completely honest, the coefficients from `eprobit` and `eoprobit` models without endogenous covariates can be interpreted in the same way as the coefficients from `probit` and `oprobit` models. The coefficients are in standard-deviation-of-the-latent-dependent-variable units. If you understood that, great, go ahead. If you did not, use `margins` for all post hoc inferences after `probit`, `oprobit`, `eprobit`, `xteprobit`, `eoprobit`, and `xteoprobit` models. With `margins`, you can easily make and test statements about how your covariates determine the levels of the probability of a positive outcome and how changes in your covariate change that probability.

How to use and interpret margins

You can always interpret the results from `margins` as being structural results. That means we can interpret them as though we were able to manipulate values in the population and compute the effect of those manipulations on the dependent variable. Obviously, this assumes the model is correctly specified.

When you have endogenous covariates, `margins` fully accounts for the correlation that gave rise to the endogeneity. For an in-depth discussion of interpretation of results of models with endogenous covariates, see [Blundell and Powell \(2003\)](#), [Imbens and Newey \(2009\)](#), and [Wooldridge \(2010\)](#).

How to use margins in models without endogenous covariates

If your models include no endogenous covariates in the main equation, you can use `margins` in the ordinary way. Here is how you would ordinarily use `margins`. The following model has no endogenous covariates:

```
. use https://www.stata-press.com/data/r19/ermexample
    (Artificial ERM example data)
. eregress y x1 x2 c.x1#c.x2
    (output omitted)
```

The model fit is

$$y_i = \beta_0 + \beta_1 x_{1i} + \beta_2 x_{2i} + \beta_3 x_{1i} x_{2i} + e_i \cdot y$$

Assume that our interest is in the effect of x_1 . One way to interpret the effect is to interpret the coefficients: a one-unit increase in x_1 increases y by $\beta_1 + \beta_3 x_2$. Another way to interpret the effect is by using counterfactuals. In these data, what would be the average change in y if x_1 were increased by 1? margins will tell us if we type

```
. margins, at(x1=generate(x1)) at(x1=generate(x1+1)) contrast(at(r) nowald)
Contrasts of predictive margins                                     Number of obs = 200
Model VCE: OIM
Expression: Average structural function mean, predict()
1._at: x1 = x1
2._at: x1 = x1+1
```

	Delta-method			
	Contrast	std. err.	[95% conf. interval]	
_at (2 vs 1)	1.109641	.1750625	.7665246	1.452757

You can learn about margins, its features, and its syntax in [\[R\] margins](#). We will tell you enough, however, so that everything we say will make sense.

Assume that the data comprise three subgroups in which we have a special interest. For instance, we want to know how an increase in x_1 would affect each subgroup. margins can tell us that too.

```
. margins, at(x1=generate(x1)) at(x1=generate(x1+1)) contrast(at(r) nowald)
> over(group)
Contrasts of predictive margins                                     Number of obs = 200
Model VCE: OIM
Expression: Average structural function mean, predict()
Over:      group
1._at: 0.group
        x1 = x1
    1.group
        x1 = x1
    2.group
        x1 = x1
2._at: 0.group
        x1 = x1+1
    1.group
        x1 = x1+1
    2.group
        x1 = x1+1
```

	Delta-method			
	Contrast	std. err.	[95% conf. interval]	
_at@group				
(2 vs 1) 0	.5561469	.1960937	.1718102	.9404835
(2 vs 1) 1	1.123401	.1754062	.7796108	1.46719
(2 vs 1) 2	1.641114	.2153742	1.218988	2.063239

margins helps us to understand changes that are different in each observation. If we had the simple model `eregress y x1 x2`, we know the effect of incrementing x_1 is to increase y by $\hat{\beta}_1$, which might be 3. The change would be 3 in every observation. In the model we have, however, the effect of incrementing x_1 is to increase y by $\beta_1 + \beta_3 x_2$. The average effect depends on the distribution of x_2 .

`margins` helps us to understand how a change affects the average in our data and subgroups of our data. We are using our sample as a proxy for the population and subpopulations, but that is what we usually do in statistics. We assume that our sample is representative. The issues are the same as we discussed in [ERM] Intro 5.

If our sample is representative but we want `margins` to report population-based standard errors, we need to specify `vce(robust)` when we fit the model:

```
. eregress y x1 x2 c.x1#c.x2, vce(robust)
```

If our sample is not representative, we can weight it with the inverse probability that its observations were sampled from the underlying population. If we want `margins` to report population-based standard errors, we can type

```
. eregress y x1 x2 c.x1#c.x2 [pw = weight], vce(robust)
```

or type

```
. eregress y x1 x2 c.x1#c.x2 [pw = weight]
```

We can type either because specifying `[pw=weight]` implies `vce(robust)`.

Even when we do specify or imply `vce(robust)`, `margins` will report sample standard errors by default. To obtain population-based standard errors, we must specify or imply `vce(robust)` when we fit the model, and when we use `margins`, we must specify its `vce(unconditional)` option:

```
. margins, at(x1=generate(x1)) at(x1=generate(x1+1)) contrast(at(r)) ///
    vce(unconditional)
```

In the linear regression example we have been discussing, we included an interaction in the model and used `margins` to report averages. We used `margins` because the interaction caused changes to vary observation by observation. Probit and ordered probit models produce predictions that vary observation by observation even in models with no interactions. Consider the following probit model, which is almost the simplest one possible:

```
. eprobit y_p x1
```

The model is

$$\Pr(\text{positive outcome}) = \Pr(\beta_0 + \beta_1 x1_i + e_i \cdot y_p > 0) = \text{normal}(\beta_0 + \beta_1 x1)$$

Assume that our interest is in `x1` just as it was previously. The effect of a one-unit increase in `x1` is to increase the normal index by $\hat{\beta}_1$. Simple, right? No, it is not. The effect in probabilities of that change varies observation by observation. Here is how the results vary if $\hat{\beta}_1$ were 0.5 and we incremented `x1` by 1. The effect depends on each subject's initial probability of a positive outcome:

Subject's original Pr(pos. outcome)	Increment by	Subject's new Pr(pos. outcome)	Difference
0.01	0.5 s.d.	0.03	0.02
0.10	0.5 s.d.	0.22	0.12
0.20	0.5 s.d.	0.37	0.17
0.40	0.5 s.d.	0.60	0.20
0.50	0.5 s.d.	0.69	0.19
0.60	0.5 s.d.	0.77	0.17
0.90	0.5 s.d.	0.96	0.06
0.99	0.5 s.d.	1.00	0.01

A subject whose original probability was 0.40 experiences an increase of 0.20 when x1 is incremented by 1. Meanwhile, a subject whose probability was 0.90 experiences a mere 0.06 increase.

Using margins, we can obtain the average changes in probabilities in the data due to incrementing x1 by 1. We type

```
. margins, at(x1=generate(x1)) at(x1=generate(x1+1)) contrast(at(r) nowald)
Contrasts of adjusted predictions                                Number of obs = 200
Model VCE: OIM
Expression: Average structural function probability, predict()
1._at: x1 =    x1
2._at: x1 = x1+1
```

	Delta-method			
	Contrast	std. err.	[95% conf. interval]	
_at (2 vs 1)	.2961685	.0287644	.2397912	.3525458

We can obtain the changes for each of the three subgroups too:

```
. margins, at(x1=generate(x1)) at(x1=generate(x1+1)) contrast(at(r) nowald)
> over(group)
Contrasts of adjusted predictions                                Number of obs = 200
Model VCE: OIM
Expression: Average structural function probability, predict()
Over:      group
1._at: 0.group
       x1 =    x1
    1.group
       x1 =    x1
    2.group
       x1 =    x1
2._at: 0.group
       x1 = x1+1
    1.group
       x1 = x1+1
    2.group
       x1 = x1+1
```

	Delta-method			
	Contrast	std. err.	[95% conf. interval]	
_at@group (2 vs 1) 0	.3857775	.051078	.2856664	.4858885
(2 vs 1) 1	.2944176	.0294406	.2367152	.3521201
(2 vs 1) 2	.2096478	.0202614	.1699363	.2493594

Counterfactuals are useful in complicated linear models—we had an interaction in ours—and in non-linear models whether simple or complicated.

How to use margins with endogenous covariates

Let's start with a reasonably simple model, namely,

```
. eregress y x1 x2, endogenous(x1 = z1, nomain)
```

The model is

$$\begin{aligned}y_i &= \beta_0 + \beta_1 x1_i + \beta_2 x2_i + e_i.y \\ x1_i &= \gamma_0 + \gamma_1 z1_i + e_i.x1\end{aligned}$$

where $\rho = \text{corr}(e.x1, e.y)$ and is nonzero.

Let's imagine that y is a health outcome and $x1$ is a 0/1 variable indicating whether a treatment was administered that is expected to improve the outcome. Observations are people, and people choose for themselves whether to have the treatment. Given the story, we *should* fit the model by typing

```
. eregress y i.x1 x2, endogenous(x1 = z1, probit nomain)
```

Nonetheless, we are going to fit the model without the `probit` specification and factor-variable notation for endogenous covariate $x1$:

```
. eregress y x1 x2, endogenous(x1 = z1, nomain)
```

We omit `probit` only because it will be easier for us to explain. We need to show you some math, and the math will be simpler in the linear model case.

What is important is that ρ is likely to be nonzero, no matter how the model is fit. ρ is the correlation between $e.y$ and $e.x1$. $e.y$ includes all the unobserved things that affect how well the treatment works. $e.x1$ includes all the unobserved things that affect whether individuals choose the treatment. ρ is likely to be nonzero and positive because people who believe that they are more likely to benefit from the treatment ($e.y > 0$) should be more likely to choose the treatment ($e.x1 > 0$).

Thus, the best prediction of y that we can make for people like person 1 in our data—people who have the same value of $x1$, $x2$, and $z1$ —includes the effect of $\hat{\rho}$, albeit indirectly. The best prediction of y we can make for people like person 1 is that their expected value of y will be

$$\hat{y}_1 = \hat{\beta}_0 + \hat{\beta}_1 x1_1 + \hat{\beta}_2 x2_1 + \hat{e}_1.y$$

$\hat{e}_1.y$ is our estimate of the expected value of $e.y$ in the first observation. Expected values of errors are often 0, but not in this case. This one depends on ρ . Given that we know the values $x1_1$ and $z1_1$, we have an estimate of $e_1.x1$, namely,

$$\hat{e}_1.x1 = x1_1 - \hat{\gamma}_0 - \hat{\gamma}_1 z1_1$$

Because $e.x1$ and $e.y$ are correlated, we can produce an estimate of $e_1.y$ given $\hat{e}_1.x1$ and $\hat{\rho}$. It is a detail, but the formula is

$$\hat{e}_1.y = \frac{\rho \times \text{s.d.}(e.y)}{\text{s.d.}(e.x1)} \times \hat{e}_1.x1$$

The value of $\hat{e}_1.y$ can be calculated, and the best prediction we can make for people like person 1 includes it and is

$$\hat{y}_1 = \hat{\beta}_0 + \hat{\beta}_1 x1_1 + \hat{\beta}_2 x2_1 + \hat{e}_1.y$$

margins with predict(asf)

Let us temporarily consider x_1 to be continuous. We do this to consider what happens if we add 1 to x_1 .

What is the best prediction we can make for people like person 1 if x_1 was incremented by 1? It is

$$\hat{y}_1 = \hat{\beta}_0 + \hat{\beta}_1(x_{1_1} + 1) + \hat{\beta}_2 x_{2_1} + \hat{e}_{1.y}$$

The above is how `margins` with the default method, `predict(asf)`, makes the calculation for each observation in the data. Observation by observation, it calculates

$$\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1(x_{1_i} + 1) + \hat{\beta}_2 x_{2_i} + \hat{e}_{i.y} \quad (1)$$

`margins` includes $\hat{e}_{i.y}$ in the calculations. This is the best prediction for people like the people in our population conditioned on everything we know about them. Using the jargon in the literature, we say that it computes a structural function that averaged over your population gives you an average structural function.

But there is something more happening. You may have noticed that $\hat{e}_{i.y}$ is a function of the endogenous covariates, in this case x_1 . But we did not add 1 to x_1 in $\hat{e}_{i.y}$. We are asking `margins` to compute derivatives and effects taking the level of endogeneity as given by the data. Effects that condition on the level of endogeneity have a causal interpretation. To be precise, they have a structural function interpretation. For an in-depth discussion of structural functions, see [Blundell and Powell \(2003\)](#), [Imbens and Newey \(2009\)](#), and [Wooldridge \(2010\)](#).

Now, we return to considering x_1 to be binary.

margins with predict(fixedasf)

`predict(asf)` uses (1) and makes its predictions given how the world currently operates. People choose their values of x_1 , and the choice they make is correlated with the outcomes they expect.

`predict(fixedasf)` makes predictions for a world that operates differently. In the alternative world, x_1 is fixed at a value such as 1. This means that the population of people like person 1 is expanded from being all people like person 1 who made the same treatment choice to being all people like person 1 regardless of the treatment choice they made. In the expanded definition of people like person 1, the correlation between $e.y$ and $e.x_1$ is broken. The correlation is now 0, and the best prediction for people like person 1 sets $\hat{e}_{1.y}$ to 0:

$$\hat{y}_1 = \hat{\beta}_0 + \hat{\beta}_1 x_{1_1} + \hat{\beta}_2 x_{2_1} \quad (2)$$

In the jargon of statistics, x_1 is no longer endogenous—it is fixed, and the entire equation for x_1 becomes irrelevant.

When you specify `margins` with the `predict(fixedasf)` option, it makes the calculation for each person by using the approach used for person 1 in (2). It uses

$$\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1 x_{1_i} + \hat{\beta}_2 x_{2_i}$$

These observation-by-observation predictions are called potential outcomes when applied to treatment models. The averages based on them that `margins` reports are called potential-outcome means (POMs). These averages correspond to what would be observed in a world in which x_1 is fixed at a particular value.

We refer to them as fixed average structural functions. They too have a structural interpretation when we average over the whole population and when the covariates can be considered to be fixed exogenously, by fiat.

Note that in the case described in the prior paragraph, the results from `fixedasf` and `asf` are asymptotically equivalent. So you can continue to use the default predictions for `margins` if you wish. `fixedasf` simply provides an alternative computation.

When to use which

`margins` can produce counterfactuals in two ways.

With the default `predict(asf)` method, `margins` uses

$$\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1 \mathbf{x1}_i + \hat{\beta}_2 \mathbf{x2}_i + \hat{e}_i \cdot \mathbf{y}$$

for the values of `x1` and `x2` specified. The predictions are a function of `x1` and `x2` and the covariates appearing in the `x1` equation. Those covariates along with $\hat{\rho}$ go into the calculation of $\hat{e}_i \cdot \mathbf{y}$. These predictions correspond to how the current world operates.

When you specify `predict(fixedasf)`, `margins` uses

$$\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1 \mathbf{x1}_i + \hat{\beta}_2 \mathbf{x2}_i$$

where `x1` is fixed at the value specified. These predictions are based on the exogenous covariates in the main equation (`x2` in this case) and the value to which the fixed variable (`x1`) is set. These predictions correspond to a different world in which `x1` is no longer endogenous but is fixed to a particular value.

Using margins with nonlinear and random-effects models

Above, we showed you results for one-level (cross-sectional) linear models that are fit with `eregress`. That discussion extends naturally when fitting any of the other ERM models.

The formulas are more complicated when models are nonlinear, but the assumptions and their implications are the same.

What if we fit a random-effects model for panel data or grouped data? If we type

```
. xtregress y x1 x2, endogenous(x1 = z1, nomain)
```

the model is

$$\begin{aligned} y_{ij} &= \beta_0 + \beta_1 \mathbf{x1}_{ij} + \beta_2 \mathbf{x2}_{ij} + u_i \cdot \mathbf{y} + v_{ij} \cdot \mathbf{y} \\ \mathbf{x1}_{ij} &= \gamma_0 + \gamma_1 \mathbf{z1}_{ij} + u_i \cdot \mathbf{x1} + v_{ij} \cdot \mathbf{x1} \end{aligned}$$

We can rewrite this in terms of the combined errors $e_{ij} \cdot \mathbf{y} = u_i \cdot \mathbf{y} + v_{ij} \cdot \mathbf{y}$ and $e_{ij} \cdot \mathbf{x1} = u_i \cdot \mathbf{x1} + v_{ij} \cdot \mathbf{x1}$. Then, we have

$$\begin{aligned} y_{ij} &= \beta_0 + \beta_1 \mathbf{x1}_{ij} + \beta_2 \mathbf{x2}_{ij} + e_{ij} \cdot \mathbf{y} \\ \mathbf{x1}_{ij} &= \gamma_0 + \gamma_1 \mathbf{z1}_{ij} + e_{ij} \cdot \mathbf{x1} \end{aligned}$$

This produces an estimate of $e_{ij} \cdot \mathbf{y}$ that depends on estimates of $e_{ij} \cdot \mathbf{x1}$ and $\rho = \text{corr}(e_{ij} \cdot \mathbf{x1}, e_{ij} \cdot \mathbf{y})$.

Everything we said above about using default predictions (`predict(asf)`) with margins is true when we fit a random-effects model. To see this, we just replace $\hat{e}_{i \cdot y}$ with $\hat{e}_{ij \cdot y}$ in each of the formulas in the previous sections.

Advanced options: Using margins `predict(base())` and `predict(fix())`

We have simplified our lives by creating predictions that margins can consume and that provide a structural function (causal) interpretation (see [Blundell and Powell \[2003\]](#), [Imbens and Newey \[2009\]](#), and [Wooldridge \[2010\]](#)), namely, with `predict(asf)` and `predict(fixedasf)`. Behind the scenes, to compute these structural predictions, we use two advanced prediction options, `predict(base())` and `predict(fix())`. In most cases, the effects we want are given by margins with the default `predict(asf)` option. But perhaps, you may find exceptions to the rule. If this is the case, this section will be helpful; otherwise, you may skip this section without any harm.

Let's think of the model given by

```
. eregress y1 x1 x2, endogenous(x1 = z1, nomain)
```

and described by

$$\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1(x1_i) + \hat{\beta}_2x2_i + \hat{e}_{i \cdot y}$$

margins with the default prediction will compute inferences for a given level of endogeneity, and it will assume $\hat{e}_{i \cdot y}$ is fixed. In other words, margins allows $x1$ to change in $\hat{\beta}_1(x1_i)$ but not in $\hat{e}_{i \cdot y}$. This is also true for $z1$. margins does this by creating copies of the variables $x1$ and $z1$ (let us call them $x1orig$ and $z1orig$) and then computing the following:

```
. margins ..., ... predict(base(x1=x1orig z1=z1orig))
```

`predict(base())` allows us to tell margins which endogenous equations and elements within them to fix at any value. The default `predict(asf)` method fixes the elements of the endogenous equation at its original values.

`predict(fixedasf)` proceeds similarly. Because we are fixing the correlation to be zero, we need to specify only the endogenous equations that we wish to be treated as fixed. In our example, we would type

```
. margins ..., ... predict(fix(x1))
```

You can appreciate how the bookkeeping becomes more involved as we increase the number of endogenous equations and exclude instruments such as $z1$.

But the utility of `predict(base())` and `predict(fix())` is not to manually compute what `predict(asf)` and `predict(fixedasf)` provide us. One case where this framework becomes handy is when we have an endogenous treatment equation and multiple endogenous equations. Say we fit

```
. eregress y1 x1, entreat(x1 = z1) endogenous(x2 = z1 z2)
```

We want to compute an average treatment effect for $x1$. We would type

```
. margins r.x1, predict(fix(x1))
```

When we type this, we ignore the correlation for the endogenous equation of $x1$ because the values of $x1$ here are exogenously fixed. Yet we are not ignoring the correlation of the endogenous equation for $x2$. This cannot be done with `predict(fixedasf)` because it will fix all endogenous equations. It cannot be done with `predict(asf)` either because it would incorporate the correlation of the endogenous treatment equation for $x1$.

This is a case where understanding `predict(fix())` and `predict(base())` matters. In fact, it is such an important case that we again have made it easy to obtain. Whenever you have an endogenous treatment equation, you can just type

```
. estat teffects
```

and get the average treatment effect.

We could have obtained a consistent estimate of the average treatment effect typing `margins r.x1` and using the default `predict(asf)` method. However, using `predict(fix())` is more efficient in this case.

References

- Blundell, R. W., and J. L. Powell. 2003. “Endogeneity in nonparametric and semiparametric regression models”. In *Theory and Applications: Eighth World Congress*. Advances in Economics and Econometrics, edited by M. Dewatripont, L. P. Hansen, and S. J. Turnovsky, vol. 2: 312–357. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511610257.011>.
- Imbens, G. W., and W. K. Newey. 2009. Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica* 77: 1481–1512. <https://doi.org/10.3982/ECTA7108>.
- Wooldridge, J. M. 2010. *Econometric Analysis of Cross Section and Panel Data*. 2nd ed. Cambridge, MA: MIT Press.

Also see

- [ERM] **Intro 9** — Conceptual introduction via worked example
- [ERM] **Example 1a** — Linear regression with continuous endogenous covariate

Stata, Stata Press, Mata, NetCourse, and NetCourseNow are registered trademarks of StataCorp LLC. Stata and Stata Press are registered trademarks with the World Intellectual Property Organization of the United Nations. StataNow is a trademark of StataCorp LLC. Other brand and product names are registered trademarks or trademarks of their respective companies. Copyright © 1985–2025 StataCorp LLC, College Station, TX, USA. All rights reserved.



For suggested citations, see the FAQ on [citing Stata documentation](#).