

Notes 8: The linear model and interpretation

ECO 231W - Undergraduate Econometrics

Prof. Carolina Caetano

1 Interpretation of the linear regression model

We presented the linear regression model, which said that

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_k x_k + u,$$

and $\mathbb{E}[u|x_1, x_2, \dots, x_k] = 0$. This model implied that

$$\mathbb{E}[y|x_1, \dots, x_k] = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_k x_k, \quad (1)$$

and thus the model has two implications that greatly restrict reality:

1. The expected outcome is indeed linearly related to the regressors, as evidenced in equation (1). No parabola, no log. Linear.
2. The things we don't observe are expected to be the same (zero) for all values of the regressors, as evidenced in $\mathbb{E}[u|x_1, x_2, \dots, x_k] = 0$.

We need to interpret the two implications above in order to understand what restrictions we are imposing in the world. Today we will focus on the first: linearity. Next class we will tackle the second condition.

1.1 Linearity

The expected outcome is linear. This means that

$$\mathbb{E}[y|x_1, x_2, \dots, x_k] = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_k x_k,$$

What does this mean in English? It means that for a given value of x_1, \dots, x_k , the expected outcome is a line. That's not really English, is it? Let's try with one of our examples. It means that:

If we know how many classes (x_1), office hours (x_2) and sections (x_3) the student attended, and we know how many hours he or she studied (x_4), then we expect his or her final grade (y) to be a linear function of the classes, office hours, sections, and hours studying.

I don't know about you, but I don't think that's English yet. One thing that is not quite clear is the term: "linear function." What do I mean by that? Sure, you know the equation of a line. Do you know what characterizes a linear function? Let's try plotting a line:

A line is characterized by the following principle: _____

We can say exactly the same for x_2, \dots, x_k as well. Now let's go back to the example.

If we know how many classes, office hours and sections the student attended, and we know how many hours he or she studied, then if we vary the number of classes he attends, we expect his final grade to vary exactly the same, no matter what was the original number of classes that he attended. The same can be said if we vary the office hours, sections, or hours studying.

1.2 Interpretation of the slope coefficients

What we were just discussing is closely related to the meaning of the β 's in this model. Let's begin with β_1 . Let's do a bit of math. Remember that

$$\mathbb{E}[y|x_1, x_2, \dots, x_k] = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_k x_k.$$

Suppose that we begin with $x_1 = 0$. Leave x_2, \dots, x_k alone, they could be anything.

Now, suppose that we increase one unit of x_1 . Now $x_1 = 1$, but (here is the most important detail) we leave everything else constant. Yes, *ceteris paribus*, remember? It means that x_2, \dots, x_k are still the same. Then,

If we subtract the first equation from the second, we get: Intuitively, β_1 is how much

we expect y to vary when we increase x_1 from zero to 1, and leave everything else constant.

But wasn't the whole point of a line that the outcome will vary just as much no matter what was the original point? Let's imagine that the original $x_1 = 15$. Then Now we

increase one unit of x_1 , so $x_1 = 16$, but we leave everything else constant. Then

Again, subtract the first from the second: So, β_1 is also how much we expect y to vary

when we increase x_1 from 15 to 16, and leave everything else constant.

We could do this forever with any number we wish. As long as we keep increasing one

unit, the change in the outcome will be exactly β_1 . So:

What is β_1 ? _____

Going back to the example:

β_1 is how much we expect the final grade to vary when we force the student to take one more class, but leave everything else (office hours, sections and hours studying) constant.

Is β_1 the causal effect of x_1 ? Yes. But this interpretation only works if both condition holds. Don't forget that we also required that $\mathbb{E}[u|x_1, \dots, x_k] = 0$. So, we are always indirectly assuming that everything else unobservable is also expected to be constant (in fact, it's expected to be zero!). So yes, if this is indeed the case, then we pinned down the comparability requirement for causal effects. Yes, if everything else is constant, and we vary x_1 , then the outcome will vary β_1 .

Can you think of what will happen if we vary x_1 in 2 units? What if we vary 7 units? I recommend that you give this a thought. I will ask something of this nature in the exam.

Observe also that nothing we did here is special about β_1 . We could do everything just the same for any of the other β . Everything but the causal interpretation. Remember, we write models for the variable of interest. Reserve the causal interpretations for the variable of interest, leave the controls alone. We will discuss more about this later in the course.

1.3 Interpretation of the intercept

Now we move to the interpretation of the intercept. Again:

$$\mathbb{E}[y|x_1, x_2, \dots, x_k] = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_k x_k.$$

Now think that $x_1 = 0, x_2 = 0, \dots, x_k = 0$. Then

So, β_0 is the expected outcome when $x_1 = \dots = x_k = 0$. In English in an example:

β_0 is how much we expect the final grade to be for a student that did not go to class, office hours, or sections, and did not study at all.

In this example, β_0 is very intuitive. Although it is not common, we can all conceive of such a student. However, sometimes β_0 has a very silly interpretation. This will be easy to see when we talk about qualitative data (dummies) in the next class.