How do I interpret a regression model when some variables are log transformed?

Authored by
stats writer

June 30, 2024

A regression model is a statistical tool used to analyze the relationship between a dependent variable and one or more independent variables. In some cases, it may be necessary to log transform certain variables in order to better understand and interpret the results of the model. This is often done when the data is skewed or when there is a non-linear relationship between the variables.

Interpreting a regression model with log transformed variables requires a different approach than a model with non-transformed variables. Instead of directly interpreting the coefficients, one must first transform them back to their original scale. This can be done by taking the inverse of the log transformation, typically using the exponential function.

Additionally, the interpretation of the coefficients may also change when log transformations are applied. For example, in a non-transformed model, a coefficient of 0.5 would indicate that a one-unit increase in the independent variable leads to a half-unit increase in the dependent variable. However, in a log transformed model, this would be interpreted as a one-unit increase in the independent variable leading to a 50% increase in the dependent variable.

In summary, when interpreting a regression model with log transformed variables, it is crucial to transform the coefficients back to their original scale and to consider the new interpretation of the coefficients. This can help provide a more accurate understanding of the relationship between the variables.

FAQ How do I interpret a regression model when some variables are log transformed?

Introduction

In this page, we will discuss how to interpret a regression model when some
variables in the model have been log transformed.
The example data can be downloaded here (the file is in .csv format). The
variables in the data set are writing, reading, and math scores ((textbf{write}), (textbf{read}) and

**(textbf{math})),
the log transformed writing (lgwrite) and log transformed math scores (lgmath)
and (textbf{female}). For these examples, we have taken the natural log (ln). All the examples are done in Stata, but they can be easily
generated in any statistical package. In the examples below, the variable (textbf{write}) or its log
transformed version will be used as the outcome variable. The examples are used for
illustrative purposes and are not intended to make substantive sense. Here is a
table of different types of means for variable (textbf{write}).**

Variable | Type Obs Mean
-------------+--
write | Arithmetic 200 52.775 51.45332 54.09668
| Geometric 200 51.8496 50.46854 53.26845
| Harmonic 200 50.84403 49.40262 52.37208
--

Outcome variable is log transformed

Very often, a linear relationship is hypothesized

**between a log transformed
outcome variable and a group of predictor variables. Written mathematically, the
relationship follows the equation**

**begin{equation}
log(y_i) = beta_0 + beta_1 x_{1i} + cdots + beta_k x_{ki} + e_i ,
end{equation}**

**where (y) is the outcome variable and (x_1, cdots, x_k) are the predictor variables.
In other words, we assume that (log(y) - mathbf{x}^T boldsymbolbeta)
is normally distributed, (or (y) is log-normal conditional on all the covariates). Since this is just
an ordinary least squares regression, we can easily interpret a
regression coefficient, say (beta_1), as the expected change in log of (y) with respect to a one-unit
increase in (x_1) holding all other variables at any fixed value, assuming that (x_1)
enters the model only as a main effect. But what if we want to know**

what happens to the outcome variable (y) itself for a one-unit increase in (x_1)?
The natural way to do this is to interpret the exponentiated regression
coefficients, (exp(beta)), since exponentiation is
the inverse of logarithm function.

Let's start with the intercept-only model.

```
------------------------------------------------------------------------------
lgwrite | Coef. Std. Err. t P>|t|
-------------+----------------------------------------------------------------
intercept | 3.948347 .0136905 288.40 0.000 3.92135 3.975344
------------------------------------------------------------------------------
```

(log(textbf{write}) = beta_0 = 3.95).

We can say that (3.95) is the unconditional expected mean of log of (textbf{write}).
Therefore the exponentiated value is (exp(3.948347) = 51.85). This is the geometric mean of
(textbf{write}). The emphasis here is that it is the geometric mean
instead of the arithmetic mean. OLS regression of the

original variable (y) is used to
to estimate the expected arithmetic mean and OLS regression of the log
transformed outcome variable is to estimated the expected geometric mean of the
original variable.

Now let's move on to a model with a single binary predictor variable.

----
lgwrite | Coef. Std. Err. t P>|t|
-------------+--
female | .1032614 .0265669 3.89 0.000 .050871 .1556518
intercept | 3.89207 .0196128 198.45 0.000 3.853393 3.930747

(begin{split}
log(textbf{write}) & = beta_0 + beta_1 times textbf{female} \
& = 3.89 + .10 times textbf{female} .
end{split})

Before diving into the interpretation of these

parameters, let's get the
means of our dependent variable, (textbf{write}), by gender.

males

Variable | Type Obs Mean
-------------+--
write | Arithmetic 91 50.12088 47.97473 52.26703
| Geometric 91 49.01222 46.8497 51.27457
| Harmonic 91 47.85388 45.6903 50.23255
----

females

Variable | Type Obs Mean
-------------+--
write | Arithmetic 109 54.99083 53.44658 56.53507
| Geometric 109 54.34383 52.73513 56.0016
| Harmonic 109 53.64236 51.96389 55.43289

Now we can map the parameter estimates to the geometric means for the two
groups. The intercept of (3.89) is the log of geometric

mean of (textbf{write}) when (textbf{female} = 0), i.e., for males. Therefore, the exponentiated value of it is the geometric mean for the male group: (exp(3.892) = 49.01). What can we say about the coefficient for (textbf{female})? In the log scale, it is the difference in the expected arithmetic means of the log of (textbf{write}) between the female students and male students. In the original scale of the variable (textbf{write}), it is the ratio of the geometric mean of (textbf{write}) for female students over the geometric mean of (textbf{write}) for male students, (exp(.1032614) = 54.34383 / 49.01222 = 1.11). In terms of percent change, we can say that switching from male students to female students, we expect to see about (11%) increase in the geometric mean of writing scores.

Last, let's look at a model with multiple predictor variables.

----

lgwrite | Coef. Std. Err. t P>|t|

-------------+--

female | .114718 .0195341 5.87 0.000 .076194 .153242

read | .0066305 .0012689 5.23 0.000 .0041281 .0091329

math | .0076792 .0013873 5.54 0.000 .0049432 .0104152

intercept | 3.135243 .0598109 52.42 0.000 3.017287 3.253198

(begin{split}
log(textbf{write}) & = beta_0 + beta_1 times textbf{female} + beta_2 times textbf{read} + beta_3 times textbf{math} \
& = 3.135 + .115 times textbf{female} + .0066 times textbf{read} + .0077 times textbf{math} .
end{split})

The exponentiated coefficient (exp(beta_1))
for (textbf{female}) is the ratio of the expected geometric mean for the female
students group over the expected geometric mean for the male students group, when
(textbf{read}) and (textbf{math}) are held at some fixed

value. Of course, the expected geometric means for the male and female students group will be different for different values of (textbf{read}) and (textbf{math}). However, their ratio is a constant: (exp(beta_1)). In our example, (exp(beta_1) = exp(.114718) approx 1.12) . We can say that writing scores will be (12%) higher for the female students than for the male students. For the variable (textbf{read}), we can say that for a one-unit increase in (textbf{read}), we expect to see about a (0.7%) increase in writing score, since (exp(.0066305) = 1.006653 approx 1.007). For a ten-unit increase in (textbf{read}), we expect to see about a (6.9%) increase in writing score, since (exp(.0066305 times 10) = 1.0685526 approx 1.069).

The intercept becomes less interesting when the predictor variables are not centered and are continuous. In this particular model, the intercept is the expected mean for (log(textbf{write})) for male ((textbf{female} =

0)) when (textbf{read}) and (textbf{math}) are equal to zero.

In summary, when the outcome variable is log transformed, it is natural to interpret the exponentiated regression coefficients. These values correspond to changes in the ratio of the expected geometric means of the original outcome variable.

Some (not all) predictor variables are log transformed

Occasionally, we also have some predictor variables being log transformed. In this section, we will take a look at an example where some predictor variables are log-transformed, but the outcome variable is in its original scale.

```
------------------------------------------------------------------------------
write | Coef. Std. Err. t P>|t|
-------------+----------------------------------------------------------------
female | 5.388777 .9307948 5.79 0.000 3.553118 7.224436
lgmath | 20.94097 3.430907 6.10 0.000 14.17473 27.7072
lgread | 16.85218 3.063376 5.50 0.000 10.81076 22.89359
```

intercept | -99.16397 10.80406 -9.18 0.000 -120.4711 -77.85685

Written in equation, we have

(begin{split}
textbf{write} & = beta_0 + beta_1 times textbf{female} + beta_2 times log(textbf{math}) + beta_3 times log(textbf{read}) \
& = -99.164 + 5.389 times textbf{female} + 20.941 times log(textbf{math}) + 16.852 times log(textbf{read}).
end{split})

Since this is an OLS regression, the interpretation of the regression
coefficients for the non-transformed variables are unchanged from an OLS
regression without any transformed variables. For example, the
expected mean difference in writing scores between the female and male students is about (5.4) points, holding the other predictor variables constant. On the other hand,

due to the log transformation, the estimated effects of (textbf{math}) and (textbf{read}) are no longer linear, even though the effect of (log(textbf{math})) and (log(textbf{read})) are linear. The plot below shows the curve of predicted values against the reading scores for the female students group holding math score constant.

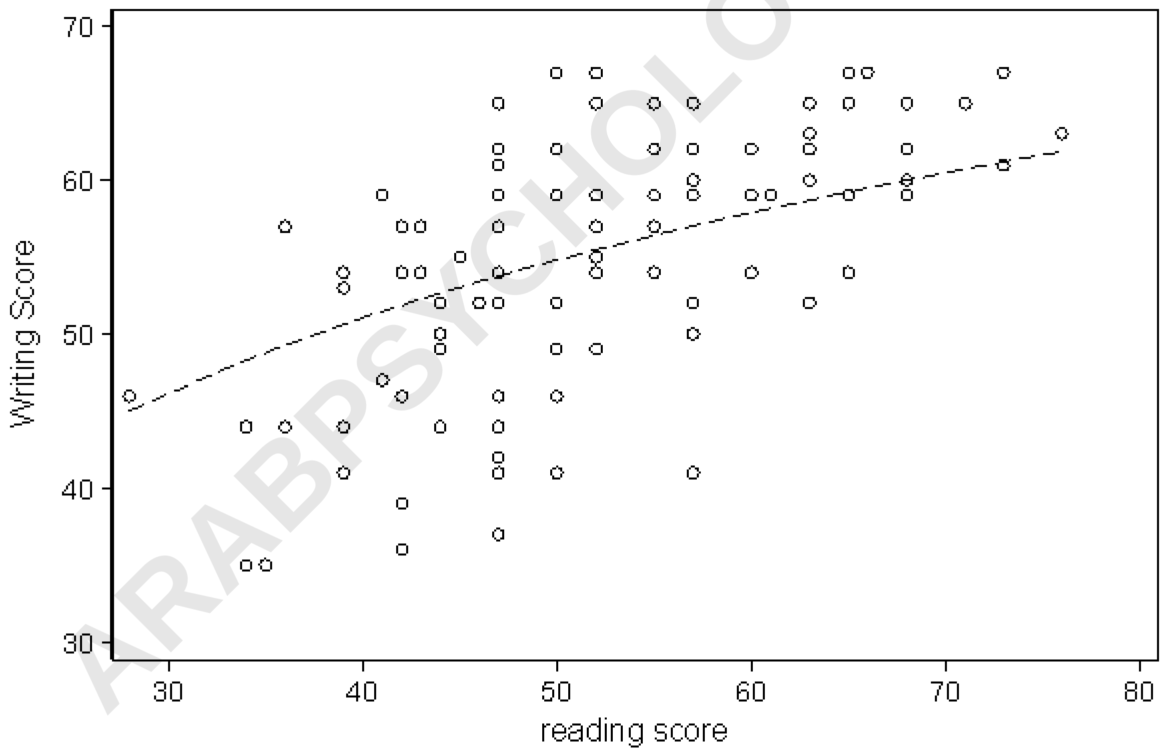

How do we interpret the coefficient of (16.852) for the variable of log of reading score? Let's take two values of reading score, (

r_1) and (r_2). The expected mean difference in writing score at (r_1) and (r_2), holding the other predictor variables constant, is (textbf{write}(r_2) - textbf{write}(r_1) = beta_3 times = beta_3 times).

This means that as long as the percent increase in (textbf{read}) (the predictor variable) is fixed, we will see the same difference in writing score, regardless where the baseline reading score is. For example, we can say that for a (10%) increase in reading score, the difference in the expected mean writing scores will be always (beta_3 times log(1.10) = 16.85218 times log(1.1) approx 1.61).

Note:

Recalling the Taylor expansion of the function (f(x) = log(1 + x)) around (x_0 = 0), we have (log(1 + x) = x + mathcal{O}(x^2)).

Therefore, for a small change in the predictor variable we can approximate the difference in the expected mean of the dependent variable by multiplying the coefficient by the change in the

predictor variable. In our example we can say that for a (1%) increase in reading score,
the difference in the expected mean writing scores will be approximately (beta_3 times 0.01 = 16.85218 times 0.01 = .1685218).
If we use the log, the exact value will be (beta_3 times log(1.01) = 16.85218 times log(1.01) = .1676848).

Both the outcome variable and some predictor variables are log transformed

What happens when both the outcome variable and predictor variables are log
transformed? We can combine the two previously described situations into one.
Here is an example of such a model.

--

lgwrite | Coef. Std. Err. t P>|t|

-------------+--

female | .1142399 .0194712 5.87 0.000 .07584 .1526399
lgmath | .4085369 .0720791 5.67 0.000 .2663866 .5506872
read | .0066086 .0012561 5.26 0.000 .0041313 .0090859
intercept | 1.928101 .2469391 7.81 0.000 1.441102 2.415099

Written as an equation, we can describe the model:

(begin{split}
**log(textbf{write}) & = beta_0 + beta_1 times textbf{female} + beta_2 times log(textbf{math}) + beta_3 times textbf{read} **
& = 1.928101 + .1142399 times textbf{female} + .4085369 times log(textbf{math}) + .0066086 times textbf{read}.
end{split})

For variables that are not transformed, such as (textbf{female}), its exponentiated coefficient is the ratio of the geometric mean for the female to the geometric mean for the male students group. For example, in our example, we can say that the expected percent increase in geometric mean from male student group to female student group is about (12%) holding other variables constant, since (exp(.1142399) approx 1.12). For reading score, we can say that for a one-unit increase in reading score,

we expected to see
about (0.7%) of increase in the geometric mean of writing score, since (exp(.0066086)
= 1.007).

Now, let's focus on the effect of (textbf{math}). Take two values of (textbf{math}), (m_1)
and (m_2), and hold the other predictor variables at any fixed value. The equation above yields

(
begin{equation}
log(textbf{write}(m_2)) - log(textbf{write}(m_1)) = beta_2 times
end{equation})

It can be simplified to (log), leading to

(
begin{equation}
dfrac{textbf{write}(m_2)}{textbf{write}(m_1)} = left(dfrac{m_2}{m_1} right) ^ {beta_2}
end{equation})

This tells us that as long as the ratio of the

two math scores, (m_2/m_1) stays the same, the expected ratio of the outcome variable, (textbf{write}), stays the same. For example, we can say that for any (10%) increase in (textbf{math}) score, the expected ratio of the writing score will be ((1.10) ^ {beta_2} = (1.10) ^ .4085369 = 1.0397057). In other words, we expect about (4%) increase in writing score when math score increases by (10%).

Note:

Here also we can use an approximation method. Since, ((1 + x) ^ a approx 1 + ax) for a small value of (|a|x), therefore for a small change in the predictor variable we can approximate the expected ratio of the of the dependent variable by multiplying the coefficient by the ratio of the change in the predictor variable. For example, we can say that for any (1%) increase in (textbf{math}) score, the expected ratio of the writing score is approximately (1 + .01 times {beta_2} = 1 + .01 times .4085369 = 1.004085). The exact value will be ((1.01) ^ {beta_2} = (1.01)

^.4085369 = 1.004073).