Standardized Coefficients in Logistic Regression

Richard Williams, University of Notre Dame, https://www3.nd.edu/~rwilliam/
Last revised October 10, 2024

NOTE: Long and Freese's `spost13` programs are used in this handout; specifically, the `listcoef` command, which is part of `spost13`, is used. Long's 1997 Regression Models for Categorical and Limited Dependent Variables provides a brief substantive discussion on pp. 69-71.

Overview. Long and Freese discuss alternative ways of standardizing variables that may help with interpretation. They primarily talk about these techniques with regards to logistic, multinomial logistic, and ordinal regression models, but they may be useful for OLS regression as well. Their `listcoef` command illustrates these different alternatives. I'll first present some preliminary results that will make it easier to understand what `listcoef` is doing.

```
. use https://www3.nd.edu/~rwilliam/statafiles/glm-logit.dta, clear
. sum

    Variable |        Obs        Mean    Std. Dev.        Min         Max
-------------+--------------------------------------------------------------
       grade |         32      .34375    .4825587          0           1
         gpa |         32    3.117188    .4667128       2.06           4
        tuce |         32     21.9375    3.901509         12          29
         psi |         32       .4375    .5040161          0           1

. logit  grade gpa tuce i.psi, nolog

Logistic regression                             Number of obs   =         32
                                                LR chi2(3)      =      15.40
                                                Prob > chi2     =     0.0015
Log likelihood = -12.889633                     Pseudo R2       =     0.3740


-----------------------------------------------------------------------------
       grade |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+---------------------------------------------------------------
         gpa |   2.826113   1.262941     2.24   0.025     .3507938    5.301432
        tuce |   .0951577   .1415542     0.67   0.501    -.1822835    .3725988
       1.psi |   2.378688   1.064564     2.23   0.025       .29218    4.465195
       _cons |  -13.02135   4.931325    -2.64   0.008    -22.68657    -3.35613
-----------------------------------------------------------------------------

. fitstat

[A lot of output omitted]

------------------------+-------------
Variance of             |
                    e   |       3.290
                 y-star |       7.210
```

```
.  listcoef, std help

logit (N=32): Unstandardized and standardized estimates

  Observed SD:  0.4826
    Latent SD:  2.6851


-------------------------------------------------------------------------------
            |          b         z     P>|z|    bStdX     bStdY    bStdXY    SDofX
------------+------------------------------------------------------------------
        gpa |     2.8261     2.238    0.025     1.319     1.053     0.491    0.467
       tuce |     0.0952     0.672    0.501     0.371     0.035     0.138    3.902
      1.psi |     2.3787     2.234    0.025     1.199     0.886     0.447    0.504
   constant |   -13.0213    -2.641    0.008         .         .         .        .
-------------------------------------------------------------------------------
        b = raw coefficient
        z = z-score for test of b=0
    P>|z| = p-value for z-test
    bStdX = x-standardized coefficient
    bStdY = y-standardized coefficient
   bStdXY = fully standardized coefficient
    SDofX = standard deviation of X
```

In the `listcoef` output, the column labeled `b` (which the `logit` command labels as `Coef.`) gives the unstandardized (metric) coefficients. The columns labeled `z` and `P>|z|` are also the same as in the `logit` output. The other columns (which were presented because I used the `std` option) give information that is relevant to different types of standardization. The `help` option added the descriptions of what each part of the output means.

Full Standardization. With full standardization, both the X and the Y* variables are standardized to have a mean of 0 and a standard deviation of 1. It is similar to standardization in OLS regression (with the important difference that Y* is a latent variable and not observed; we'll see why this is important later). In the `listcoef` output, the fully standardized coefficients are in the column labeled `bStdXY`. [NOTE: As `fitstat` shows, the variance of Y* is 7.21, which means its standard deviation is 2.685 – the same as what `listcoef` reports.]

The results show you that a 1 standard deviation increase in gpa results, on average, in almost half a standard deviation increase (.4912) in the log odds of getting an A.

If you know the metric coefficients and the standard deviations of the the x's and y*, you can compute the standardized coefficients the same way you do in OLS:

$$b'_k = b_k * \frac{s_{x_k}}{s_{y*}}$$

So, for example, to get the fully standardized effect of gpa,

$$b'_{gpa} = b_{gpa} * \frac{s_{gpa}}{s_{y*}} = 2.82611 * \frac{.4667}{2.685} = .4912$$

X-Standardization. An intermediate approach is to standardize only the X variables. In the `listcoef` output, in the column labeled `bStdX`, the Xs are standardized but Y* is not. Hence, by standardizing the Xs only, you can see the relative importance of the Xs. We see that a 1 standard deviation increase in gpa produces, on average, a 1.319 increase in the log odds of getting an A. (To get the X-Standardized coefficient, just multiply b_k by the standard deviation of x_k, e.g. for gpa 2.82611 * .4667 = 1.319.) Note that, if your goal is to compare the effects of Xs measured in different metrics, X-Standardization alone is sufficient.

Y-Standardization. You can also standardize Y* only. The `listcoef` column labeled `bStdY` gives you the coefficients from when Y* is standardized but X is not. A 1 unit increase in gpa produces, on average, a 1.0525 standard deviation increase in Y*. To get the Y-standardized coefficient, just divide b_k by the standard deviation of Y*, e.g. for gpa 2.82611/2.685 = 1.0525.

If you don't include the `std` parameter, after a logistic regression `listcoef` does a variation of X-standardization, showing you the odds ratios (i.e. the factor change in the odds as X increases):

```
. listcoef

logit (N=32): Factor change in odds

  Odds of: 1 vs 0

-------------------------------------------------------------------------
            |         b         z      P>|z|        e^b   e^bStdX     SDofX
------------+------------------------------------------------------------
        gpa |    2.8261     2.238     0.025     16.880     3.740     0.467
       tuce |    0.0952     0.672     0.501      1.100     1.450     3.902
      1.psi |    2.3787     2.234     0.025     10.791     3.316     0.504
   constant |  -13.0213    -2.641     0.008          .         .         .
-------------------------------------------------------------------------
```

This tells you that a 1 unit increase in gpa multiplies the odds of success by 16.880. A 1 standard deviation increase in gpa multiplies the odds by 3.740. (Recall that the X-standardized coefficient is 1.3190; exp(1.3190) = 3.74.) See the help for `listcoef` for other options that may be useful.

Discussion. The usual argument for using standardized coefficients is that they provide a means for comparing the effects of variables measured in different metrics. This is true here as well. So, for example, you can see that a 1 standard deviation (SD) increase in gpa produces more change in the log odds of getting an A than does a 1 SD increase in tuce. Nevertheless, standardized effects tend to be looked down upon. It makes no sense to think about a one SD increase in a dummy variable like gender. Even for continuous variables, standardized coefficients are not very intuitive, e.g. how many of us think in terms of standard deviations? Worse, they can be very misleading. For example, if the standard deviations of variables differ across groups, the standardization of variables will also differ, causing coefficients to not be comparable across groups (e.g. in one group X might get divided by 10 while in another it gets divided by 7.)

There are, however, some unique concerns when using logistic regression and other GLMs. Unlike Y in OLS regression, the variance of Y* is not fixed; it will change as you add more

variables to the model. This can create problems in logistic regression that you do not have with OLS regression. Some authors (e.g. Winship & Mare, ASR 1984) therefore recommend Y-Standardization or Full-Standardization. We discuss this further in a later handout.

Appendix: Standardized Coefficients in OLS Regression

If you run `listcoef` after the `regress` command, the fully standardized coefficients are the same as the regression standardized coefficients, e.g.

```
. webuse nhanes2f, clear
. reg weight height age female black, beta
```

Source	SS	df	MS		Number of obs =	10337
					F(4, 10332) =	881.52
Model	620082.606	4	155020.652		Prob > F =	0.0000
Residual	1816944.64	10332	175.856044		R-squared =	0.2544
					Adj R-squared =	0.2542
Total	2437027.25	10336	235.7805		Root MSE =	13.261

weight	Coef.	Std. Err.	t	P>|t|	Beta
height	.7485279	.01966	38.07	0.000	.4709032
age	.1237255	.0078948	15.67	0.000	.1387257
1.female	-1.540187	.3721392	-4.14	0.000	-.0500913
1.black	3.679295	.4256284	8.64	0.000	.0734762
_cons	-59.05337	3.563342	-16.57	0.000	.

```
. listcoef, std help
```

regress (N=10337): Unstandardized and standardized estimates

  Observed SD: 15.3551
  SD of error: 13.2611

	b	t	P>|t|	bStdX	bStdY	bStdXY	SDofX
height	0.7485	38.074	0.000	**7.231**	0.049	**0.471**	9.660
age	0.1237	15.672	0.000	**2.130**	0.008	**0.139**	17.217
1.female	-1.5402	-4.139	0.000	**-0.769**	-0.100	**-0.050**	0.499
1.black	3.6793	8.644	0.000	**1.128**	0.240	**0.073**	0.307
constant	-59.0534	-16.572	0.000	.	.	.	.

```
      b = raw coefficient
      t = t-score for test of b=0
   P>|t| = p-value for t-test
   bStdX = x-standardized coefficient
   bStdY = y-standardized coefficient
  bStdXY = fully standardized coefficient
   SDofX = standard deviation of X
```

Note that, in OLS, while full standardization is frequently done, X-Standardization alone is enough to achieve the goal of comparing the effects of Xs measured in different metrics, and may be easier to interpret since Y is left in its original metric. So, for example, we can see that a 1 standard deviation in height results, in average, on a 7.23 kilogram increase in weight, whereas a 1 standard deviation increase in age results in an average increase of 2.13 kilograms.