Partial Correlations: MLR Coefficients and Endogeneity

- ***Interpreting MLR Coefficients: What's New?, What's Left?***
- ***… and Partial Correlations***
- ***Example I: Bodyfat***
- ***Returning to Endogeneity and Partial Correlations***
- ***Example II: Bodyfat, again***

Interpreting MLR Coefficients: *What's New?, What's Left?*

Consider the standard MLR model in which *y* is being reressed on RHS variables *x*, *w*, *z*, etc. using OLS. You saw previously that estimated OLS/MLR coefficient for, say, *x*, could be generated by estimating the SLR model in which *What'sLeft* of the dependent variable *y* (the part of *y* unexplained by the other RHS variables) is regressed on *What'sNew* about *x* (the part of *x* unexplained by the other RHS variables). Or put differently:

- $x = WhatsNew_x + \hat{x}$, and
- $y = WhatsLeft_y + \hat{y}$,

where $\hat{x}$ and $\hat{y}$ are the predicteds when x and y are, respectively, regressed on the remaining RHS variables.

And so in this way, estimated MLR coefficients can be generated by a well chosen SLR model. Further, since SLR coefficients essentially reflect correlations (subject to a standard deviations adjustment),[1] it's useful to interpret MLR coefficients as capturing the correlation between *WhatsLeft* of the dependent variable, say *y*, and *WhatsNew* about an explanatory variable, say *x*. Or to be more precise, the sign of the estimated MLR coefficient for *x* will agree with the sign of the correlation between *WhatsLeft* of the dependent variable *y* and *WhatsNew* about *x*.

... and Partial Correlations

We call the correlation of between *WhatsLeft* of the dependent variable *y* and *WhatsNew* about *x* the ***partial correlation*** of *y* and *x*. Here's why:

> The *partial correlation* of y and x (given a bunch of other explanatory variables in the MLR model), is the correlation between the (unexplained) residuals after separately regressing *y*

---

[1] Recall that in SLR models, $\beta_x = \rho_{xy}\dfrac{S_y}{S_x}$. And since $S_y > 0$ and $S_x > 0$, $sign(\beta_x) = sign(\rho_{xy})$.

and *x* on those other RHS variables. So, in a sense, in generating *WhatsLeft* of the dependent variable *y* and *WhatsNew* about *x*, we have *partialed out* the effects of the other RHS variables (on *y* and on *x*), and are evaluating the correlation of what remains.

Finally, and just to say it again: The sign of every MLR coefficient reflects the ***partial correlation*** between the particular RHS variable and the dependent variable. And since partial correlations are not the same as regular correlations, it is a mistake to assume that the sign of every MLR coefficient reflects the simple ***correlation*** between the particular RHS variable and the dependent variable. You need to *partial out* the effects of the other RHS variables.

This last insight will be helpful in attempts to sign the impact of endogeneity. We'll turn to that shortly… but first, here's an example, working with the bodyfat dataset, and two explanatory variables *abd* (waist size) and *wgt_kg* (weight in kilograms).

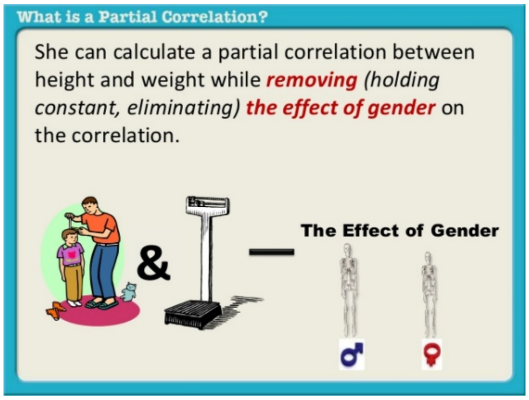

Example I: Bodyfat and wgt_kg

The simple correlation between *brozek* and *wgt_kg* is positive:

```
. corr brozek wgt_kg

             |   brozek   wgt_kg
-------------+------------------
      brozek |   1.0000
      wgt_kg |   0.6132   1.0000
```

To generate the partial correlation between *brozek* and *wgt_kg* given *abd*, we need to generate *WhatsLeft* of the dependent variable *brozek* (given *abd*) and *WhatsNew* about *wgt_kg*, (again, given *abd*), and then compute their correlation:

1. Generate *WhatsLeft* of *y:*

```
. reg brozek abd
. predict whatsleft, resid
```

2. Generate *WhatsNew* about *x:*

```
. reg wgt_kg abd
. predict whatsnew, resid
```

3. Generate the correlation between *WhatsLeft* of *y* and *WhatsNew* about *x*:

```
. corr whatsleft whatsnew

             | whatsl~t whatsnew
-------------+------------------
   whatsleft |   1.0000
    whatsnew |  -0.4093   1.0000
```

Surprise! While the ***simple*** correlation between *brozek* and *wgt_kg* is positive, once you have controlled for the effects of *abd*, the ***partial*** correlation between *brozek* and *wgt_kg* is negative. And so it should be no surprise that the *wgt_kg* coefficient is positive in the SLR model in which *brozek* is regressed on *wgt_kg*, and negative in the MLR model when *abd* has been added in as an explanatory variable:

```
. reg brozek wgt_kg
. reg brozek wgt_kg abd

--------------------------------------------
                      (1)             (2)
                   brozek          brozek
--------------------------------------------
wgt_kg              0.357***       -0.301***
                  (12.27)         (-7.08)

abd                                 0.915***
                                  (17.42)

_cons              -9.995***       -41.35***
                  (-4.18)        (-17.14)
--------------------------------------------
N                     252             252
--------------------------------------------
t statistics in parentheses
* p<0.05, ** p<0.01, *** p<0.001
```

Returning to Endogeneity … and Partial Correlations

So now it should be clear that we were a bit fast and loose earlier in saying that OVB (omitted variable bias) was effectively driven by the product of two ***simple*** correlations. It is in fact driven by two ***partial*** correlations… the correlations are between *WhatsLeft* of a LHS variable, and *WhatNew* about some (additional) RHS variable.

To illustrate, let's return to the case in which the Full Model has three RHS variables, *x*, *z* and *w*, and we are interested in evaluating the impact on the estimated coefficients for the surviving variables, *x* and *z*, when *w* is dropped from the model.

From before, we know that the two MLR SRFs of interest in calculating the OVB impacts are:

- Full Model: SRF_y: $\hat{y} = \hat{\beta}_0 + \hat{\beta}_x x + \hat{\beta}_z z + \hat{\beta}_w w$
- Collinearity Regression Model: SRF_w: $\hat{w} = \hat{\alpha}_0 + \hat{\alpha}_x x + \hat{\alpha}_z z$

In this case, the OVB for the *x* coefficient will be defined by

- $OVB_x = \hat{\alpha}_x \hat{\beta}_w$ (the product of the SRF_w *x* coeff and the SRF_y *w* coeff)

Since the two components in this calculation, $\hat{\alpha}_x$ *and* $\hat{\beta}_w$, are MLR coefficients, they will also be coefficients in SLR models with *WhatsNew* as the explanatory variable:

- $\hat{\beta}_w$: Start with the Full Model, and regress *WhatsLeft* of y on *WhatsNew* about *w* in that model. $\hat{\beta}_w$ will be the estimated slope coefficient in that SLR model, and effectively reflect the ***partial*** correlation between *y* and *w* (given the rest of the Full Model).

- $\hat{\alpha}_x$: Now start with the Collinearity Regression Model, and regress *WhatsLeft* of *w* on *WhatsNew* about x in that model. $\hat{\alpha}_x$ will be the estimated slope coefficient in that SLR model, and effectively reflect the ***partial*** correlation between *w* and *x* (given the rest of the Collinearity Regression Model, variable *z*).

So the effects of endogeneity are driven by *partial* correlations… and not by *simple* correlations. And if you can sign the *partial* correlations, you can sign the omitted variable impact.

Sounds simple… but that can be a real challenge. While we often have a good intuitive sense of the signs of simple correlations, who can intuit a partial correlation (especially when there are lots and lots of RHS variables in the model)? But if you want to get the sign of omitted variable bias right, that's exactly the challenge! ***Good luck!***

Here's an update to the earlier summary chart. As before, the Full MLR Model has dependent variable *y*, RHS variables *x* and *z* and possibly additional RHS variables. *z* is dropped from the Full Model, and we are interested in the impact on the estimated *x* coefficient:

Omitted Variable Bias (dependent variable *y*; drop z)**:** impact on the *x* coeff. = $\hat{\beta}_z\hat{\alpha}_x$

		partial correlation between *y* and omitted variable *z* (Full MLR Model)		
partial correlation between *x* and omitted variable *z* (Collinearity Regression)		positive	zero	negative
positive		positive	0	negative
zero		0	0	0
negative		negative	0	positive

Example II: Bodyfat and wgt_kg

Continuing with the previous example, let's add *hgt_m* to the model. Here are the results of the Full model and the model after *hgt_m* has been dropped:

```
--------------------------------------------
                      (1)             (2)
                   brozek          brozek
--------------------------------------------
wgt_kg             -0.265***       -0.301***
                  (-5.41)         (-7.08)

abd                 0.880***        0.915***
                  (15.19)         (17.42)

hgt_m              -4.652
                  (-1.43)

_cons              -32.66***       -41.35***
                  (-5.01)        (-17.14)
--------------------------------------------
N                     252             252
--------------------------------------------
t statistics in parentheses
* p<0.05, ** p<0.01, *** p<0.001
```

And here are the simple correlations:

```
. corr brozek wgt_kg abd hgt_m
(obs=252)

             |   brozek   wgt_kg      abd    hgt_m
-------------+------------------------------------
      brozek |   1.0000
      wgt_kg |   0.6132   1.0000
         abd |   0.8137   0.8880   1.0000
       hgt_m |  -0.0891   0.3083   0.0878   1.0000
```

If you worked with just the simple correlations, you's sign the OVB impact on both *wgt_kg* and *abd* as negative since the simple correlation of *brozek* and *hgt_m* is negative, and the simple correlations of *hgt_m* the two surviving variables, *wgt_kg* and *abd*, are both positive. And while you woud be correct in the case of *wgt_kg*, where the OVB impact is in fact negative, you'd be all wrong about the impact on the *abd* coefficient, which increases!

So be careful using *simple* correlations in signing OVB. But *partial* correlations will always give you the right answer.

To sign the OVB impacts of dropping *hgt_m* from the full model, we are interested in signing OLS coefficients in two MLR models:

- Full model: Sign the *hgt_m* coefficient in the full MLR model (*reg brozek wgt_kg abd hgt_m*)

  Step I: sign the partial correlation between *brozek* and *hgt_m* (given *wgt_kg* and *abd*)

- Collinearity regression model: Sign the *wgt_kg* and *abd* coefficients in the collinearity regression model (*reg hgt_m wgt_k abd*)

  Steps II.a and II.b: separately sign the partial correlations between *hgt_m*, and *wgt_kg* and *abd*

Step I: Sign the partial correlation between *brozek* and *hgt_m* (given *wgt_kg* and *abd*).

```
. reg brozek  wgt_kg abd
. predict whatsnewy, resid

. reg hgt_m wgt_kg abd
. predict whatsleftx, resid

. corr whatsnewy whatsleftx
(obs=252)

             | whatsn~y whatsl~x
-------------+------------------
   whatsnewy |   1.0000
  whatsleftx |  -0.0907   1.0000
```

The partial correlation between *brozek* and *hgt_m* (given *wgt_kg* and *abd*) is negative (-.0907).

Step II.a: Sign the partial correlation between *hgt_m*, and *wgt_kg* (given, or controlling for, *abd*).

```
. drop whatsnew* whatsleft*

.
. reg hgt_m abd
. predict whatsnewy, resid

. reg wgt_kg abd
. predict whatsleftx, resid

. corr whatsnewy whatsleftx
(obs=252)

             | whatsn~y whatsl~x
-------------+------------------
   whatsnewy |   1.0000
  whatsleftx |   0.5028   1.0000
```

The partial correlation between *hgt_m* and *wgt_kg* (given *abd*) is positive (0.5028). This implies that the OVB impact on *wgt_kg* will be ***negative*** (the product of a negative and a positive partial correlation)…. *Which it is!* (see above)

Step II.b: Sign the partial correlation between *hgt_m*, and *abd* (given, or controlling for *wgt_kg*).

```
. reg hgt_m wgt_kg
. predict whatsnewy, resid

. reg abd wgt_kg
. predict whatsleftx, resid

. corr whatsnewy whatsleftx
(obs=252)

             | whatsn~y whatsl~x
-------------+------------------
   whatsnewy |   1.0000
  whatsleftx |  -0.4250   1.0000
```

The partial correlation between *hgt_m* and *abd* (given *wgt_kg*) is negative (-0.4250). This implies that the OVB impact on *abd* will be ***positive*** (the product of two negative partial correlations)…. *Which it is!* (see above)