OLS with Heterogeneous Coefficients

Nikolas Mittag*
CERGE-EI

October 10, 2024

Abstract: Regressors often have heterogeneous effects in the social sciences, which are usually modeled as unit-specific slopes. OLS is frequently applied to these correlated coefficient models. I first show that without restrictions on the relation between slopes and regressors, OLS estimates can take any value including being negative when all individual slopes are positive. I derive a simple formula for the bias in the OLS estimates, which depends on the covariance of the slopes with the squared regressor. While instrumental variable methods still allow estimation of (local) average effects under the additional assumptions that the instrument is independent of the coefficients in the first stage and reduced form equations, the results here imply complicated biases when these assumptions fail.

Keywords: correlated coefficient model, heterogeneous effects
JEL Classification Numbers: C21, C26

*I would like to thank Stanislav Anatolyev, Jon Davis, Štěpán Jurajda, Derek Wu and Paolo Zaccia for comments that improved this note. Address: CERGE-EI, Politických vězňů 7, Prague 1, 111 21, Czech Republic, nikolasmittag@posteo.de

A defining characteristic of the social sciences is that the subjects under study, such as individuals, firms or countries, act in their own way, i.e. that the effects of their characteristics differ between units of observations. Despite the likely presence of such individual-level heterogeneity, the vast majority of empirical studies estimates regression models with constant coefficients. Even though it is unknown what OLS estimates unless the effects are independent of the regressors, the estimated OLS coefficients are often interpreted as averages of the individual effects. In this note, I show that applying OLS in the presence of heterogeneous slopes can be severely misleading. In fact, OLS estimates can be negative even when all individual slopes are positive. More generally, they can take any value depending on the relationship between slopes and regressors. I derive the bias in the OLS estimates, which depends on the covariance of the slopes with the squared regressor. The bias formula implies that OLS estimates the average effect only when slopes are uncorrelated with the square of the regressor. This condition is likely to be violated in many applications, e.g. because optimizing agents select characteristics based on their effects or because there are decreasing returns to scale. Instrumental variable methods can still be used to estimate (local) average effects, but require a similar restriction. Standard IV requires both the first stage and the reduced form coefficients to be uncorrelated with the squared coefficients.[1] The results here imply complicated biases when these assumptions fail, suggesting that assessing their validity should receive more attention in practice.

The standard way to analyze heterogeneity in, among others, the evaluation of social programs is the correlated coefficient model. See e.g. Heckman and Vytlacil (2007) for discussion and further references. In its simplest, univariate form, the correlated coefficient model relates an outcome y_i to a regressor x_i linearly at the individual level, i.e.

$$y_i = x_i \beta_i + \varepsilon_i \tag{1}$$

[1] Many other methods require or invoke the stronger assumption that the instrument and the coefficients are independent.

ε_i is a standard error term that is assumed to be uncorrelated with regressors and coefficients. For simplicity, I will assume that all variables are demeaned and thus omit the intercept.[2]

FIGURE 1
EXAMPLE OF NEGATIVE OLS ESTIMATES WHEN ALL INDIVIDUAL SLOPES ARE POSITIVE

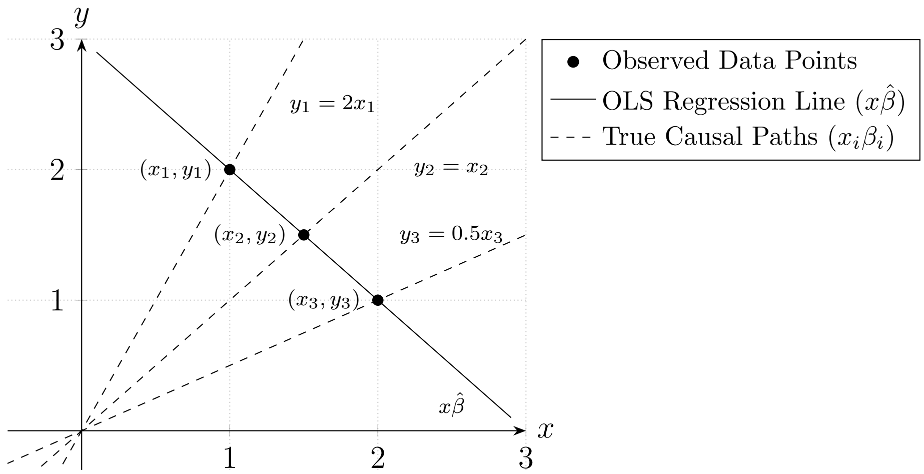

To see how heterogeneous slopes affect the observed data and hence estimates, first consider the expectation of the outcome conditional on the regressor. In the standard model without coefficient heterogeneity, i.e. when $\beta_i = \beta \; \forall i$, the conditional expectation of y given x is $x_i\beta$, which is linear in x. In the presence of heterogeneity, β_i is a random variable as well, so that the conditional expectation becomes $\mathbb{E}[y_i|x_i] = x_i\mathbb{E}[\beta_i|x_i]$. Notice that this function is non-linear in x_i unless $\mathbb{E}[\beta_i|x_i]$ is constant in x. Unless $\mathbb{E}[\beta_i|x_i]$ is known, x_i and β_i cannot be separated without further information such as exogenous variation in x. To see this, note that $\beta_i = \frac{f(x_i)}{x_i}$ results in $\mathbb{E}[y|x] = f(x)$ for any function $f()$. Thus, without any restrictions on the relation between x_i and β_i, coefficient heterogeneity can produce any aggregate functional form from a model that is linear at the unit level. For example, if $\beta_i = ax_i$, then $\mathbb{E}[y|x] = ax^2$, turning the linear relationship into a quadratic function. Figure 1 shows that the observed conditional expectation (solid line) can be severely misleading: Even though all individual slopes are positive, the slope of the conditional expectation and hence the OLS coefficient are negative. It is obvious from Figure 1 that by moving the

[2]The intercept will typically be biased and individual specific intercepts may cause further bias. This bias is well known and simple to analyze using matrix notation.

observed points along the true causal paths (dashed lines), the OLS estimates can take any value for a given distribution of individual slopes.

The example above raises the questions what OLS estimates when slopes are heterogeneous and under which conditions it still estimates an average effect. As Heckman, Schmierer and Urzua (2010) point out, applying OLS to (1) yields the average effect of x_i on y_i, $\bar{\beta} = \mathbb{E}[\beta_i]$, when x_i and β_i are independent, which implies that $\mathbb{E}[\beta_i|x_i]$ is a constant. To understand what OLS estimates more generally, consider the formula for univariate OLS coefficients without an intercept:

$$\mathbb{E}[\hat{\beta}] = \frac{\mathbb{E}[x_i y_i]}{Var(x_i)} = \frac{\mathbb{E}[x_i^2 \beta_i]}{Var(x_i)} \tag{2}$$

Where the second equality follows from the standard OLS assumption that the error is uncorrelated with the regressors. By the definition of the covariance, $\mathbb{E}[AB] = \mathbb{E}[A]\mathbb{E}[B] + Cov(A, B)$, so

$$= \frac{\mathbb{E}[x_i^2]\mathbb{E}[\beta_i]}{Var(x_i)} + \frac{Cov(x_i^2, \beta_i)}{Var(x_i)} \tag{3}$$

Since $\mathbb{E}[x_i] = 0$ by construction, $\mathbb{E}[x_i^2] = Var(x_i)$, yielding the final result that

$$\mathbb{E}[\hat{\beta}] = \bar{\beta} + \frac{Cov(x_i^2, \beta_i)}{Var(x_i)} \tag{4}$$

This result clearly shows that a necessary and sufficient condition for OLS to estimate $\bar{\beta}$ is that x_i^2 does not predict β_i linearly. Thus, the requirement for OLS to yield the average effect is substantially weaker than independence of x_i and β_i, which is sufficient, but not necessary. OLS still estimates the average effect when coefficients and regressors are correlated (or when there is dependence of higher order moments) as long as $Cov(x_i^2, \beta_i) = 0$. This condition will hold, among others, if $\mathbb{E}[\beta_i|x_i]$ is symmetric around a point on the y-axis, for example when it is linear or more generally a polynomial of odd order that is not shifted horizontally. However, if the coefficients are correlated with x_i^2, e.g. because (conditional average) effects are larger or smaller when x_i is large in absolute value, the OLS estimate will differ from the average effect.[3]

For the multivariate case with a demeaned regressor matrix X and a $K \times 1$ vector of average coefficients $\bar{\beta}$, it is straightforward to show that the k^{th} element of $X'y$ is $\sum_{l=1}^{K} Cov(x_{ki}, x_{li}\beta_{li})$. Using the same arguments as above, the expectation of the OLS coefficient is $\mathbb{E}[\hat{\beta}] = \bar{\beta} + (X'X)^{-1}C$ where the k^{th} element of C is $\sum_{l=1}^{K} Cov(x_{ki}x_{li}, \beta_{li}) = \sum_{l=1}^{K} E[x_{ki}x_{li}\beta_{li}]$.[4]

If an instrumental variable is available, several methods allow for estimation of $\bar{\beta}$ or other features of the distribution of effects even when x is endogenous. See Mogstad and Torgovitsky (2024) for a review. These results obviously still apply to the case at hand where x is independent of or uncorrelated with ε. However, in addition to the standard IV assumptions, these estimation strategies assume that the instrument is independent of β_i and the coefficients in the first stage equation, γ_i. While the standard IV assumptions are typically discussed or tested, little attention is paid to these additional restrictions in practice. Yet the results for OLS imply that heterogeneous coefficients lead to bias when the IV is related to the first or second stage coefficients.

Specifically, consider the case of the just-identified IV estimator, which is the ratio of the OLS estimates from the reduced form equation ($y_i = z_i\beta_i\gamma_i + \eta$) and the first stage equation ($x_i = z_i\gamma_i + \nu$). Using the formulas from above yields

$$\text{plim } \hat{\beta}^{IV} = \frac{\mathbb{E}[\widehat{\beta\gamma}^{OLS}]}{\mathbb{E}[\hat{\gamma}^{OLS}]} = \frac{\mathbb{E}[\beta_i\gamma_i] + \frac{Cov(z_i^2,\beta_i\gamma_i)}{Var(z_i)}}{\mathbb{E}[\gamma_i] + \frac{Cov(z_i^2,\gamma_i)}{Var(z_i)}} \tag{5}$$

$$= \frac{\mathbb{E}[\beta_i\gamma_i]}{\mathbb{E}[\gamma_i]} \cdot \frac{\mathbb{E}[\gamma_i]}{\mathbb{E}[\gamma_i] + \frac{Cov(z_i^2,\gamma_i)}{Var(z_i)}} + \frac{Cov(z_i^2, \beta_i\gamma_i)}{\mathbb{E}[\gamma_i]Var(z_i) + Cov(z_i^2, \gamma_i)} \tag{6}$$

$$= \beta^{IV} \cdot \frac{\bar{\gamma}}{\bar{\gamma} + \frac{Cov(z_i^2,\gamma_i)}{Var(z_i)}} + \frac{Cov(z_i^2, \beta_i\gamma_i)}{\mathbb{E}[\gamma_i z_i^2]} \tag{7}$$

[3]Observe that for binary x, $\frac{Cov(x_i^2,\beta_i)}{Var(x_i)} = \frac{Cov(x_i,\beta_i)}{Var(x_i)}$ is the coefficient from regressing β_i on an intercept and x, which is $\mathbb{E}[y_i|x=1] - \mathbb{E}[y_i] = \mathbb{E}[\beta_i|x=1] - \mathbb{E}[\beta_i]$. Thus $\mathbb{E}[\hat{\beta}] = \mathbb{E}[\beta_i|x=1]$, showing the well-known result that OLS still estimates the average effect of treatment on the treated in the presence of heterogeneity.

[4]Using $Cov(x_{ki}x_{li}, \beta_{li}) = \mathbb{E}[x_{ki}x_{li}\beta_{li}] - \mathbb{E}[x_{ki}]\mathbb{E}[x_{li}\beta_{li}]$ and $\mathbb{E}[x_{ki}] = 0$.

The first two lines follow from the definitions and rearranging terms, the last line makes use of the fact that $Var(z_i) = \mathbb{E}[z_i^2]$ and $Cov(z_i^2, \gamma_i) = \mathbb{E}[\gamma_i z_i^2] - \mathbb{E}[\gamma_i]\mathbb{E}[z_i^2]$. The IV estimator thus suffers from a multiplicative bias that is the ratio of the average first stage effect to the OLS estimate of the first stage coefficient. This bias arises from the effect of heterogeneity on the first stage estimate. In addition, there is an additive bias term that arises from the bias due to heterogeneity in the reduced form equation (and is also scaled by the same multiplicative term). The bias thus depends on unknown quantities that enter the formulas in a complex way, making it difficult to assess in practice. Thus, it is crucial to take the validity of the additional assumption that the IV neither predicts first stage nor reduced form equation coefficients as serious as the standard IV assumptions.

In summary, this note shows that OLS estimates can be misleading about the effects of a regressor when effects are heterogeneous. OLS estimates can be negative even when all individual slopes are positive. OLS estimates the average effect if and only if the slopes are uncorrelated with the squared regressor. This condition is weaker than the requirement of independence that is typically invoked, but likely to be violated in many studies. For example, decreasing returns to scale likely induce a negative correlation, while one would expect optimizing agents to choose x_i such that it is (positively) correlated with β_i. If this condition fails, the OLS estimate differs from the average slope by the covariance of β_i and x_i^2 divided by the variance of x_i. Since the variance of x_i is known and positive, this simple result allows researchers to assess whether there is bias and determine its sign if they have information on how slopes change with the square of the regressor. Instrumental variables still identify (local) average effects under the additional assumption that the instrument is independent of the coefficients in the first stage and the reduced form equation, but estimates suffer from complicated biases when these assumptions fail. This result underlines that these assumptions should be scrutinized more in practice, because violations may often be as problematic as failure of the standard IV validity assumption. More generally, the results that heterogeneity leads to biased OLS coefficients suggest that researchers should probe more for the presence of heterogeneity and the conditions under which their estimates remain unbiased in the presence of heterogeneous effects.

References

Heckman, James J., and Edward J. Vytlacil. 2007. "Econometric Evaluation of Social Programs, Part I: Causal Models, Structural Models and Econometric Policy Evaluation." In *Handbook of Econometrics*. Vol. 6b, , ed. James J. Heckman and Edward Leamer, Chapter 70, 4779–4874. Amsterdam:Elsevier.

Heckman, James J., Daniel Schmierer, and Sergio Urzua. 2010. "Testing the correlated random coefficient model." *Journal of Econometrics*, 158(2): 177–203.

Mogstad, Magne, and Alexander Torgovitsky. 2024. "Instrumental Variables with Unobserved Heterogeneity in Treatment Effects." NBER Working Paper w30162.