Identifying Non-linearities In Fixed Effects Models

Craig T. McIntosh[♣] and Wolfram Schlenker[♠]

October 2006

Abstract

We discuss the use of quadratic terms in models which include fixed effects or dummy variables. Contrary to the standard interpretation wherein fixed effects models are identified solely by deviations from the group mean, a quadratic explanatory variable causes group means to re-enter the identification. We illustrate the conceptual difference between a global quadratic functional form across units and a quadratic functional form within groups, and suggest applications where there are important economic distinctions between them. The relationship between these forms of non-linearity is derived, and we use Monte Carlo simulations to illustrate the results of mis-specification. A hybrid estimator is introduced which allows for consistent estimation of non-linearities when both forms of convexity are present.

♣ Corresponding author, School of International Relations and Pacific Studies, University of California at San Diego, 9500 Gilman Drive, MC 0519, La Jolla, CA 92093.Email: ctmcintosh@ucsd.edu. Phone: (858) 822-1125.

♠ Department of Economics and School of International and Public Affairs, Columbia University.

Identifying Non-linearities In Fixed Effects Models

Abstract

We discuss the use of quadratic terms in models which include fixed effects or dummy variables. Contrary to the standard interpretation wherein fixed effects models are identified solely by deviations from the group mean, a quadratic explanatory variable causes group means to re-enter the identification. We illustrate the conceptual difference between a global quadratic functional form across units and a quadratic functional form within groups, and suggest applications where there are important economic distinctions between them. The relationship between these forms of non-linearity is derived, and we use Monte Carlo simulations to illustrate the results of mis-specification. A hybrid estimator is introduced which allows for consistent estimation of non-linearities when both forms of convexity are present.

Fixed effects regressions are a straightforward way of estimating marginal effects when we suspect that unexplained, time-invariant components are common to members of a given group. Measuring each variable in deviations from group means, sometimes called the *within transformation*[1], provides consistent results when additive unobserved fixed effects are present, even if these fixed effects are correlated with other observed explanatory variables. The intuitive meaning of the coefficients measured by the within estimator is 'how much does a one-unit change in X, relative to the group mean, push the outcome Y away from the group mean'? In a model where the DGP is linear, redefining the variables in this way is innocuous and allows us to estimate consistent marginal parameters under unobserved heterogeneity.

What, however, happens when we wish to ask non-linear questions of data in which fixed effects are present? The purpose of this paper is to illustrate that a basic tension exists between the standard interpretation of 'within-group variation' in fixed effects regression and the concept of estimating non-linear effects. A strict definition of within variation would imply that the marginal parameter on a certain deviation from the group mean is invariant to the value of the group mean itself. If we are using our data to look for a global non-linear relationship over the distribution of a covariate, however, this definition cannot be satisfied. The marginal effects of a relationship which is globally quadratic in X must depend on the non-demeaned value of X, and so inherently cannot be identified in deviations from a group mean. For this reason, pure within variation could not be used to identify a DGP in which Y is a globally non-linear function of the value of X.

Does this imply that fixed effects models are therefore unable to measure global non-linearities? The answer to this question is no, and the somewhat surprising reason is that quadratic fixed effects models in fact include a component of between variation in their identification. This arises because the standard application includes X^2 as an explanatory variable in a fixed effects model, meaning that the variable is first squared and then demeaned (see Behrman & Deolalikar (1993), Hummels & Levinsohn (1995), Lafontaine & Shaw (1999), and Crossley et al. (2001) for examples from the applied literature). We illustrate that the demeaned squared variable is itself a function of the group mean, and so a standard quadratic fixed effects model uses a source of variation which is not strictly 'within'. In this way quadratic fixed effects models successfully recover a global non-linear relationship precisely because the group mean is indirectly included.

A different way of thinking about non-linear relationships using within variation is to ask whether the data display a quadratic relationship in deviations from the group mean.

[1]See for example (Wooldridge 2002, p.267).

Such curvature, which we refer to as 'within non-linearity', would be invariant to the value of the mean. A natural way to measure within non-linearity is to use a quadratic in the within variation, which entails first demeaning the covariate and then squaring it, rather than squaring then demeaning. This within-unit non-linearity would have a centering point for each fixed effect, whereas global non-linearity has only a single centering point across the whole distribution of the explanatory variable. Conceptually, the distinction is the following: are the non-linearities in the data found in deviations from the mean of the fixed-effect group to which an observation belongs, or in deviations from the mean of the sample as a whole? While the difference between these two terms arises from a somewhat arcane property of quadratic functions, it is a relatively simple matter to include explanatory variables which allow us simultaneously to test for both effects.

Figure 1 illustrates the difference between these two different forms of non-linearity. Global non-linearity implies that the rate of change in the outcome Y varies as one moves along some independent variable X, while within non-linearity implies that it depends only on the movement of X away from the within-group mean of X.

In what follows we illustrate a number of applied examples in which the two forms of non-linearity are conceptually distinct and so the ability to identify them separately is of empirical interest.

Planning Horizons

Farmers make planting decisions to maximize yields given their local climate, which is known prior to planting. Specific realizations of weather in a given planting cycle, however, are stochastic and *ex post* to planting decisions (Schlenker 2006). We could use panel data with weather realizations on the right-hand side to explain yields. In this case the 'global' non-linearity explains how optimized yields change as the predictable, average weather (climate) changes. Within non-linearities, on the other hand, are caused by unexpected variations around the local mean temperature, and so represent weather shocks. In this case global concavity measures the curvature in optimal crop yields as we move across climatic zones, whereas within concavity reveals the yield response of crops already in the ground to unanticipated yearly weather shocks. In a simulation of the impact of global warming where we expect both the mean and the variance of temperature to increase, these effects are distinct and both are important.

Collective Decisionmaking

Imagine analyzing an individual-level household survey where fixed effects are implemented at the household level, using a quadratic in income to explain the savings rate. The quadratic tests whether the marginal propensity to save changes with income. If there is global concavity in this relationship, then individuals' marginal savings rates will decrease as the absolute level of income increases, regardless of whether income is above or below the household mean. To the extent that individual savings decisions are related to relative position within the household income distribution, however, we expect to see a non-linearity which is centered on the household mean rather than the population mean. In this case, a low-income individual in a household displays different savings behavior than a high-income individual in the same household, and so savings rates are likely to display nonlinearities within groups as well, potentially, as between units. When we include income linearly it is irrelevant whether the comparison point is the household or population mean, but for the quadratic term this distinction becomes important. In this case, then, within convexity measures how the marginal propensity to save changes in relation to household mean income, and global convexity picks up non-linearities across the whole income distribution.

Relativistic Preferences

Take the use of race dummies in a standard Mincerian wage regression. If a quadratic on education is included, this implies that all racial groups have a single centering point for any non-linearities. However, if members of a group are judged by the discriminatory labor market in relative terms, this is not where we should look for a non-linearity. In practice, if a black individual has a level of education which is above the black mean but below the white mean, how is their relative educational attainment judged? If their wage reflects their education relative to the population mean, we would observe any non-linearities in a global form. If, however, members of a group are rewarded by the market relative to the mean within their respective groups, then non-linearities would be present within units. Again the distinction is only relevant for a non-linear model. Here, within convexity measures whether there are increasing returns to being more educated than the average member of your group, and global convexity measures increasing returns to education for any individual at any level of education.

Panel Estimation of Kuznets Relationships

It has been argued that the recent development of detailed panel datasets on inequality (Anand & Kanbur 1993, Ram 1997) and on industrial pollutants (Antweiler et al. 1998) allow for the Kuznets Hypothesis (KH) to be tested as the author intended it, namely in a dynamic context. Such tests naturally consist of using fixed effects to eliminate the unobserved heterogeneity which might bias the cross-section, together with a quadratic in income per capita or capital/labor ratios. Since this hypothesis posits a single, global inverted-U shape it is correct to use a global quadratic as the test. It is entirely possible, however, that within non-linearity exists in the data as well. Short-term rigidities or capacity constraints in the economy may mean that the outcome is non-linear in growth observed in a manner that is related strictly to deviations from the country-level mean. In this case inequality (for the standard KH) or emissions (for the environmental KH) will display non-linearities which are not related to the level of GDP at which the growth occurs. Interestingly, we can think of a story (for example, if rapid expansion causes countries to resort to dirtier sources of energy) wherein the sign of this within non-linearity (convexity) is opposite to the global concavity suggested by the KH. In what follows, we show that tests for global concavity will be biased if we fail to account for this other form of non-linearity in the data.

The remainder of the paper is organized as follows. First, we demonstrate algebraically that a quadratic term in the fixed effects estimator captures non-linearity between units. We suggest a natural way of estimating non-linearity within units, and a hybrid estimator that can measure both forms of non-linearity simultaneously. We conclude with Monte Carlo simulations to illustrate under what circumstances a misspecified model will still give consistent estimates.

1 Model

To state the problem more formally, let x_{it} in group i in period t be a covariate drawn from a distribution with group mean μ_i. Let us suppose that the true data-generating process can be written as follows:

$$y_{it} = \beta_1 x_{it} + \beta_2 x_{it}^2 + \beta_3 (x_{it} - \mu_i)^2 + c_i + \epsilon_{it} \quad (1)$$

This equation allows for three separate mechanisms through which y_{it} can exhibit a non-linear relationship with x_{it}.

The first is that the fixed effects (c_i) themselves may be non-linear in the covariate, or that the conditional mean outcome across units is convex (or concave) in the distribution of x. This could be verified by running the dummy-variable fixed effect regression and then plotting the coefficients on the dummies against the mean value of x for each group. Given, however, that the purpose of using the within transformation is precisely to remove c_i from outcomes, we assume that this convexity is not of interest to researchers using fixed effects and so we do not discuss it further.

The two forms of non-linearity that are relevant with the use of fixed effects are the global non-linearity that comes from a quadratic in x_{it} and the within non-linearity that comes from a quadratic in $(x_{it} - \mu_i)$. If $\beta_3 = 0$ in (1) than the model reverts to the case of a quadratic functional form with fixed effects that has been estimated by previous authors in the literature and which we label the 'global' model. On the other hand, if $\beta_2 = 0$, only deviations from group mean enter the specification and we label the model as 'within' model accordingly. Finally, the "hybrid" model allows both β_2 and β_3 to be nonzero.

Consider the standard estimate of the 'global' model:

$$y_{it} = b_1 x_{it} + b_2 x_{it}^2 + c_i + e_{it} \tag{2}$$

The key insight into the difference between global and within non-linearity comes from the fact that equation (2) first squares the covariate and then demeans it. Recall that $\mu_i = \mathbb{E}[x_{it}|i]$, which can be estimated by $\bar{x}_i = \frac{1}{T}\sum_{t=1}^{T} x_{it}$. To fix notation, let $\ddot{x}_{it} = x_{it} - \bar{x}_i$ and $\overline{x_i^2} = \frac{1}{T}\sum_{t=1}^{T} x_{it}^2$. The quantity which we arrive at by squaring and then demeaning, which we denote by $\dot{x}_{it}^2$, can be written as $x_{it}^2 - \overline{x_i^2}$. From here we can add and subtract the group mean to get: $\dot{x}_{it}^2 = [x_{it} - \bar{x}_i + \bar{x}_i]^2 - \overline{x_i^2} = \ddot{x}_{it}^2 + 2\ddot{x}_{it}\bar{x}_i + [\bar{x}_i]^2 - \overline{x_i^2}$. In other words, the canonical use of non-linear fixed effects given by (2) does not measure non-linearity within units, as the terminology used to describe the estimator might suggest, because $\dot{x}_{it}^2$ is in general not equal to $\ddot{x}_{it}^2$. Further, we see that by squaring the covariate before demeaning it, we re-introduce a function of the mean of the covariate into the identification.

So what does the standard non-linear fixed effects estimator measure? A simple way to see this is to use our DGP to write out the mean of the outcome for each fixed effect unit. Taking averages over group i we get

$$\bar{y}_i = \beta_1 \bar{x}_i + \beta_2 \overline{x_{it}^2} + \beta_3 \overline{(x_{it} - \mu_i)^2} + c_i + \bar{\epsilon}_i \tag{3}$$

First noting that $(x_{it} - \mu_i)^2 - \overline{(x_{it} - \mu_i)^2} = x_{it}^2 - 2x_{it}\mu_i + \mu_i^2 - \frac{1}{T}\sum_{t=1}^{T}[x_{it}^2 - 2x_{it}\mu_i + \mu_i^2] =$

$x_{it}^2 - 2x_{it}\mu_i + \mu_i^2 - \overline{x_i^2} + 2\bar{x}_i\mu_i - \mu_i^2 = \dot{x}_{it}^2 - 2\ddot{x}_{it}\mu_i$ we can subtract (3) from (1) to give an explicit representation of both the global and the within variation in the conditional mean as:[2]

$$\begin{aligned}\mathbb{E}[\ddot{y}_{it}|x_{it}] &= \beta_1\ddot{x}_{it} + \beta_2\dot{x}_{it}^2 + \beta_3\left[\ddot{x}_{it}^2 - \overline{\ddot{x}_i^2}\right] &(4)\\ &= [\beta_1 - 2\beta_3\mu_i]\,\ddot{x}_{it} + [\beta_2+\beta_3]\,\dot{x}_{it}^2 &(5)\\ &= [\beta_1 + 2\beta_2\mu_i]\,\ddot{x}_{it} + [\beta_2+\beta_3]\,\ddot{x}_{it}^2 + [\beta_2+\beta_3]\left[[\bar{x}_i]^2 - \overline{x_i^2}\right] &(6)\end{aligned}$$

Several important features of the applied use of non-linear functional forms with fixed effects are apparent from these expressions. Equation (4) gives the correctly specified model to measure both within and global non-linearities simultaneously. When these two forms of non-linearity take the additively separable form given in (1), one has to include two squared terms: the demeaned squared variable to measure 'global' non-linearities and the squared demeaned variable to measure 'within-group' non-linearities. We label this the the 'hybrid estimator'.

Equations (5) and (6) illustrate the consequences of using a misspecified model, wherein only one form of the non-linearity is included as an explanatory variable when in fact both are present in the DGP. Equation (5) shows that if we attempt to estimate equation (2) in the presence of within-group non-linearities, bias can enter not only the estimate of the quadratic term but the linear term as well. Squaring-then-demeaning is the correct way to recover global non-linearity only if there is no within-group non-linearity present in the data, in which case $\beta_3 = 0$, and so we recover $b_1 = \beta_1$ and $b_2 = \beta_2$.

Equation (6) illustrates the reverse problem, and shows that an estimating equation which includes only within-group non-linearity will be misspecified unless $\beta_2 = 0$, i.e., there is no global non-linearity in the DGP. If $\beta_2 \neq 0$, both linear and quadratic terms are potentially biased.

2 Monte-Carlo Simulations

The previous section defined the traditional 'global' form of non-linearity, a 'within' model that only relies on deviations from the group mean, and a 'hybrid' model that allows for both sort of quadratic relationship. If the true Data Generating Process (DGP) is from either

[2]We present the asymptotic case where the number of observations per group goes towards infinity and hence $\bar{x}_i = \mu_i$. In finite samples $\bar{x}_i \neq \mu_i$ which will induce attenuation bias on the within term, but this bias decreases in the number of observations per group.

model and a different model besides the hybrid is estimated, the regression is misspecified and the coefficient estimates might hence be biased. The goal of this section is to show how these coefficients will be biased and outline a special case where some of the coefficients are still unbiased even if an inaccurate model is estimated.

In each of the following graphs of Figure 2 we use a DGP that is a convex combination of the 'global' and the 'within' quadratic term:

$$y_{it} = 3x_{it} - \alpha x_{it}^2 + 2(1-\alpha)(x_{it} - \mu_i)^2 + c_i + \epsilon_{it} \tag{7}$$

where $\alpha \in [0,1]$ is plotted on the x-axis. The extreme cases are $\alpha = 1$ where the DGP has only a 'global' quadratic component and $\alpha = 0$ where the DGP has only a 'within' quadratic component. For all intermediate values, the true DGP is a linear combination of both. We purposefully pick a low variance σ_ϵ so we can detect bias, and a larger number of observations per group (1000) so the group means $\bar{x}_i$ are estimated with high precision.[3] We use ten fixed effect groups in all of our simulations.[4]

While we vary the DGP in each plot along the x-axis, the model specification in the estimation is held constant for each plot. The left column uses a traditional 'global' quadratic functional form $y_{it} = b_1 x_{it} + b_2 x_{it}^2 + c_i + e_{it}$ and the lines in each graph display the bias (in percent) for the coefficient estimates of b_1 and b_2. Accordingly, if $\alpha = 1$, the model only includes a 'between' quadratic term in the DGP and hence the model is correctly specified and there is no bias in the estimation. Similarly, the second column uses a 'within' model $y_{it} = b_1 x_{it} + b_2(x_{it} - \bar{x}_i)^2 + c_i + e_{it}$ in the estimation, which is correctly specified for the special case where $\alpha = 0$. The third column uses a hybrid estimator $y_{it} = b_1 x_{it} + b_2 x_{it}^2 + b_3(x_{it} - \bar{x}_i)^2 + c_i + e_{it}$ that correctly identifies both parameters.

The rows distinguish whether the variance over all group members is constant across groups, e.g., if the support/range of variations around a group mean is constant for all groups and the individual x_{it} are randomly drawn. Intuitively, if the variance among all group members increases systematically in the group mean $\bar{x}_i$ in the presence of 'within-group' non-linear effects, part of that 'within-group' variation gets picked up as global non-linearity.

The first row shows that if the variance of the x_{it} in a group is constant across all groups,

[3] The smaller the number of observations per group, the larger the potential estimation error of the group mean $\bar{x}_i$, which will induce attenuation bias for the 'within' term.

[4] The 10 groups $i = 1...10$ are uniform random draws from the interval $[20(i-1), 20(i-1)+10]$ in case the variance of the x_{it} is constant across groups. In case the variance of x_{it} increases in $\bar{x}_i$, the x_{it} are uniform random draws from the interval $[20(i-1), 20(i-1)+20i]$. The standard deviation of the errors is $\sigma_\epsilon = 0.01$.

then the 'global' model in the first column will still give an unbiased estimate of the 'global' nonlinearity β_2 even if within-group non-linearities are present, while the linear term β_1 is biased. All studies that we are aware of that previously used fixed effects with quadratic functional forms used a 'global' model: Behrman & Deolalikar (1993) use quadratics in schooling and education to explain log wages and hours worked, Lafontaine & Shaw (1999) use quadratics in number of outlets and years franchising, with FE for each franchise to explain royalty rates and franchise fees, and Crossley et al. (2001) use quadratics in years since migration to explain likelihood that immigrants receive public benefits with dummies for the arrival/year cohorts. Given any 'within-group' non-linearities, linear and quadratic estimation of these effects would in general be biased.

A quadratic term is often added as a 'robustness check' on the linear term. Yet, the linear term is itself biased in the presence of within non-linearities in a traditional model of global non-linearity. If the variance is not constant across groups, not only the linear term but also the global quadratic terms are biased by the mis-specification of estimating a global model only.

A pure 'within-group' model, on the other hand, will give biased estimates of both the linear and quadratic term even if the variance over x_{it} within a group is constant across all groups.

The hybrid estimator is consistent for all data-generating processes. Since we purposefully pick a small variance σ_ϵ, the estimated coefficients are hardly different from the true values and the computed bias is indistinguishable from zero. The difference between estimated coefficients and true parameter values of the DGP in the first two columns hence represent mis-specification bias and are not due to noise in the estimation.

3 Conclusions

Even though fixed effects or dummy variables are often combined with quadratic terms, we show in this paper that the resulting coefficients have to be interpreted with some care. It is common to equate the use of fixed effects with the notion that all identification comes from 'within' variation, i.e., the variation of a variable around its mean. However, the standard use of quadratics with fixed effects uses group means to identify a global non-linear (quadratic) relationship. This is not the same as non-linearities in deviations from group means. While the use of fixed effects are equivalent to a joint demeaning of the dependent and all independent variables, the problem arises from the fact the demeaned squared variable is *not* the same as the square of the demeaned variable. We therefore generalize previous models by including both a demeaned squared term (labelled 'global' quadratic term) and the square of the demeaned variable (labelled 'within-group' quadratic term).

We use several examples to motivate the economic intuition behind these distinct forms of non-linearity. While the exact interpretation of each form of non-linearity will depend on the context, both non-linearities are likely to be present jointly in many cases. We propose a hybrid estimator with two quadratic terms that allows us to make the distinction an empirical one, because it consistently estimates both forms of non-linearities.

Finally we show with the help of Monte Carlo simulations that the hybrid estimator gives consistent estimates for either form of non-linearity, but only for certain case will a traditional 'global' model give consistent estimates in case there are within-group non-linearities.

References

Anand, S. & Kanbur, S. M. R. (1993), 'The kuznets process and the inequality-development relationship', *Journal of Development Economics* **40**, 25–52.

Antweiler, W., Copeland, B. & Taylor, S. (1998), 'Is free trade good for the environment?'.

Behrman, J. R. & Deolalikar, A. B. (1993), 'Unobserved household and community heterogeneity and the labor market impact of schooling: A case study for indonesia', *Economic Development and Cultural Change* **41**(3), 461–488.

Crossley, T. F., McDonald, J. T. & Worswick, C. (2001), 'Immigrant benefit receipt revisited: Sensitivity to the choice of survey years and model specification', *Journal of Human Resources* **36**(2), 379–397.

Hummels, D. & Levinsohn, J. (1995), 'Monopolistic competition and international trade: Reconsidering the evidence', *Quarterly Journal of Economics* **110**(3), 799–836.

Lafontaine, F. & Shaw, K. L. (1999), 'The dynamics of franchise contracting: Evidence from panel data', *Journal of Political Economy* **107**(5), 1041–1080.

Ram, R. (1997), 'Level of economic development and income inequality: Evidence from the postwar developed world', *Southern Economic Journal* **64**(2), 576–583.

Schlenker, W. (2006), 'Inter-annual weather variation and crop yields', *Mimeo. Columbia University* .

Wooldridge, J. (2002), *Econometric Analysis of Cross-section and Panel Data.*

Figure 1: "Global" and "Within" Quadratic Forms

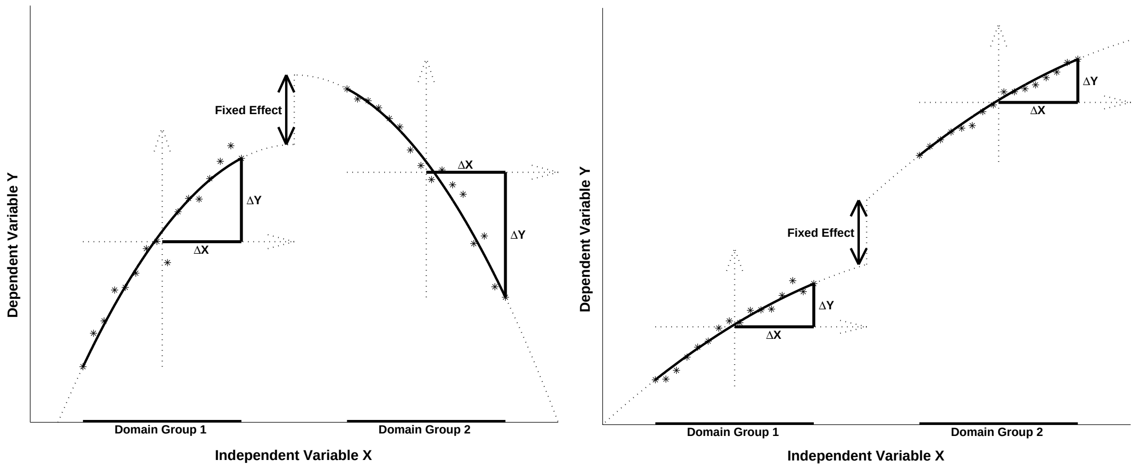

Notes: Both graphs display the case of fixed effects with two disjoint groups. The left graph assumes a global quadratic form, while the right graph assumes a constant quadratic functional form around the mean of each group ('within'). The use of fixed effects amounts to a joint demeaning of the dependent and independent variable, as indicated by the dotted axes that pass through the mean of both variables for each group. The stars indicate individual observations. In case of the 'within' quadratic functional form, a given deviation Δx from the mean of the independent variable results in the same impact Δy, i.e., only the 'within' variation in a group identifies the coefficient. However, note that for the case of the 'global' quadratic functional form, a given deviation Δx from the mean of the independent variable results in different impacts Δy that dependent on the absolute value of x. Hence both the variation within *and between* groups are responsible for the identification.

Figure 2: Bias for Convex Combinations of "Global" and "Within" DGP

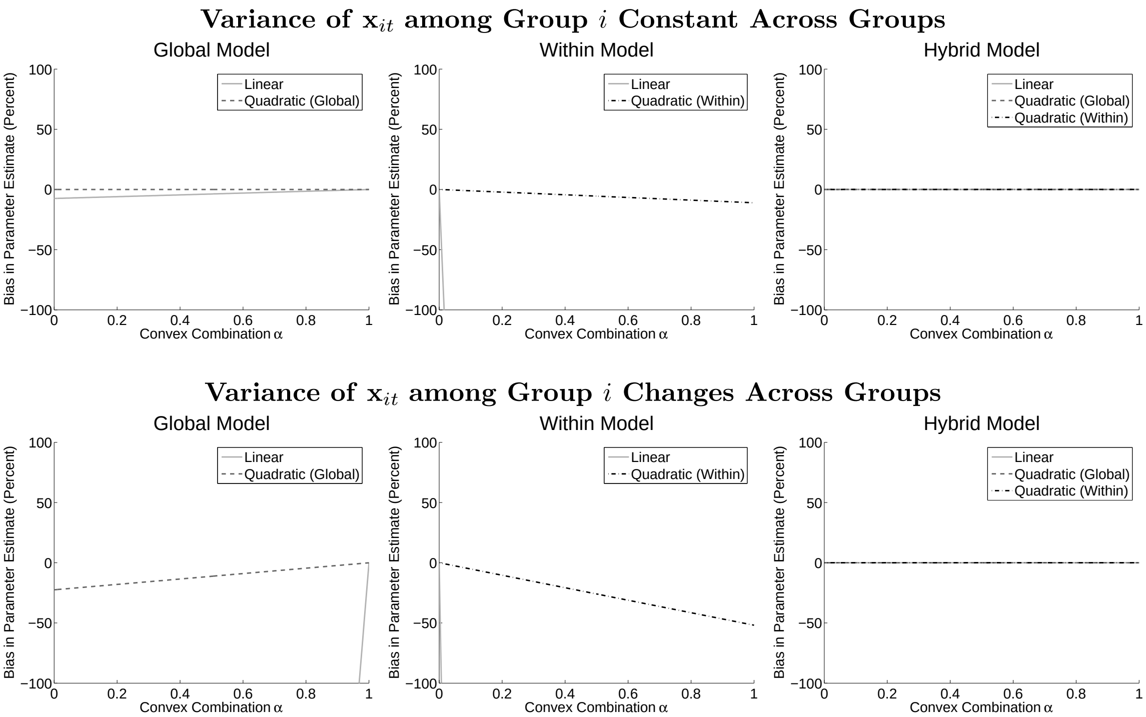

Notes: All plots use a spectrum of data generating processes $y_{it} = 3x_{it} - \alpha x_{it}^2 + 2(1-\alpha)(x_{it} - \bar{x}_i)^2 + c_i + \epsilon_{it}$ that are convex combinations of 'global' and 'within-group' quadratic functional forms, where $\alpha \in (0, 1)$ varies along the x-axis. The first column uses the standard 'global' quadratic estimator, the second column a 'within-group' quadratic estimator, and the third column a "hybrid" estimator. The first column assumes that the variance of the x_{it} for each group i is constant across al groups, while the second row assumes that the variance of the x_{it} within each group is increasing in $\bar{x}_i$.