

Improving the Interpretation of Random Effects Regression Results*

Soren Jordan[†]
Andrew Q. Philips[‡]

December 10, 2021
Forthcoming at *Political Studies Review*

Abstract

Mummolo and Peterson (2018) improve the use and interpretation of fixed effects models by pointing out that unit intercepts fundamentally reduce the amount of variation of variables in fixed effects models. Along a similar vein, we make two claims in the context of random effects models. First, we show that potentially large reductions in variation, in this case caused by quasi-demeaning, also occur in models using random effects. Second, in many instances, what authors claim to be a random effects model is actually a pooled model after the quasi-demeaning process, affecting how we should interpret the model. A literature review of random effects models in top journals suggests both points are currently not well understood. To better help users interested in improving their interpretation of random effects models, we provide Stata and R programs to easily obtain post-estimation quasi-demeaned variables.

Keywords: methodology, random effects, demeaning, interpretation

Word Count: (abstract) 140; (manuscript) about 4200 (plus two figures and three tables)

*We would like to thank the graduate students in the University of Colorado's 2019 Summer Workshop in Political Methodology for the inspiration to write this paper, the editor, two anonymous reviewers and participants of Texas A&M's "Dynamic Pie" Group, especially Allyson Benton, Natália de Paula Moreira and Guy Whitten. Brian Ezeonu provided excellent research assistance. However, any errors and omissions are our own.

[†]sorenjordanpols@gmail.com. Associate Professor, Department of Political Science, Auburn University, Haley 8024, Auburn, AL 36849.

[‡]andrew.philips@colorado.edu. Assistant Professor, Department of Political Science, University of Colorado Boulder, UCB 333, Boulder, CO 80309-0333.

Mummolo and Peterson (2018) discuss how fixed effects (FE) models in panel data—including unit-specific dummy variables to allow the intercept to shift up or down for each unit—estimate only within-unit variance. Thus, coefficients describe only the effect of a one-unit increase occurring *within* a unit, not between units. FE models have several advantages, namely that they exclude the possibility of endogeneity caused by a correlation between the included covariates and any omitted time-invariant unit effects. However, FE complicate interpretation, since variation over time is generally appreciably smaller than overall variation (i.e., what we would interpret with a global summary). To avoid problematic interpretation, especially overstating the substantive effects of the model, Mummolo and Peterson suggest accounting for this within-transformation, since it will almost always result in smaller levels of variance for the variable of interest to consider when constructing counterfactuals.

In the spirit of Mummolo and Peterson’s contribution, we note that a related panel data model also has the potential to complicate substantive conclusions about the magnitude of effects due to the nature of the estimation process: models incorporating random unit intercepts, commonly referred to as “random” or “mixed” effects (RE). RE offer a data-driven compromise between a fully pooled model and FE model by partially (“quasi”) demeaning observations by their unit averages. Moreover, this estimator is popular; recent applications range from economic voting (Valdini and Lewis-Beck 2018) to the spread of economic, social, and cultural rights (Bjornskov and Mchangama 2019), to public opinion about government expenditures (Busemeyer et al. 2019).

In this note, we make two novel claims for scholars interested in using RE models. First, like FE, RE models have the potential to appreciably reduce variation in a variable. Thus, if users are interested in exploring the substantive significance of their results they run the risk of making “extreme counterfactuals” far away from the transformed data space the model was estimated on (King and Zeng 2006). Since no popular software currently informs users of the range of these quasi-demeaned variables, we show that it is often the case that users are making interpretations—and thus theoretical conclusions—based on unrealistic counterfactual shifts in their variables of interest.

Second, due to characteristics in the data, sometimes what authors claim is a RE model is in fact no different from a fully pooled model. While this does not alter the reported estimates, it *does* mean that any reported advantages of using RE over a pooled model are not applicable, and might leave users ill-informed about the nature of their data (for instance, by thinking that unit heterogeneity is resulting in demeaning, while in reality units are so similar that they can be fully pooled).

These are claims about RE models, *not* criticisms. Instead, like Mummolo and Peterson, we seek to help applied users better understand what is going on “under the hood” and interpret their RE results: specifically by (1) adjusting counterfactual interpretations for the range of quasi-demeaned data, and (2) always reporting the variance components of RE models (and noting when they result in fully pooled models). We proceed as follows. First, since the notion that RE models lie somewhere between a fully pooled and FE model might be unfamiliar to most practitioners—for instance, Clark and Linzer (2015) devote only a few sentences to this—we discuss how quasi-demeaning occurs. Second, we survey articles using RE in the literature to demonstrate that our claims are currently not well understood. Third, we use two examples to illustrate why our claims matter substantively. We conclude by offering suggestions on reporting these quantities of interest with software we provide as well as describing alternate modeling approaches.

1 The link between several panel data models

Consider a model with one regressor that does not take into account any potential heterogeneity across units i observed across time t :

$$y_{it} = \alpha + \beta x_{it} + \varepsilon_{it} \tag{1}$$

This is often called a fully pooled model since a single global intercept, α , is estimated. No unit-specific differences are captured in this intercept; thus, any unit-specific (i.e.,

time-invariant) differences in y_{it} will end up in the error term. Moreover, standard errors will likely be incorrect because the estimator assumes observations are independent and identically distributed.

A popular alternative to this adds intercepts for each unit:

$$y_{it} = \alpha_i + \beta x_{it} + \varepsilon_{it} \quad (2)$$

and is known as a FE model.¹ Here, any time-invariant unit-specific differences—in other words, any variation between units—are soaked up by the α_i terms. The remaining variation explained by β is only within-unit.²

A third modeling choice employs a random intercept:

$$y_{it} = \alpha + \beta x_{it} + v_{it} \quad (3)$$

where the stochastic term is now a composite of fundamental error, ε_{it} (which we assume is well-behaved, as in any model), as well as error specific to a particular unit (after partialing out the “global” intercept α); thus, $v_{it} = u_i + \varepsilon_{it}$. RE models account for unit-specific differences by allowing for variation around a global intercept, α . To do this, both the variance around the stochastic term ε_{it} , σ_ε^2 , and the variance around α , which is given by σ_α^2 , are estimated.³

¹An equivalent model is often estimated by subtracting unit means from each variable (and the constant), and is known as the within-transformation: $(y_{it} - \bar{y}_i) = \beta(x_{it} - \bar{x}_i) + (\varepsilon_{it} - \bar{\varepsilon}_i)$.

² Within variance for, say, x_{it} , is given as $\frac{\sum_{i=1}^N \sum_{t=1}^T (x_{it} - \bar{x}_i)^2}{N \times T - 1}$, while between variance is $\frac{\sum_{i=1}^N (\bar{x}_i - \bar{\bar{x}})^2}{N - 1}$, where N and T are the total number of units and time points (respectively), \bar{x}_i are unit-means, and $\bar{\bar{x}}$ is the grand mean.

³i.e., we assume $\varepsilon_{it} \sim N(0, \sigma_\varepsilon^2)$ and $\alpha_i \sim N(\alpha, \sigma_\alpha^2)$. Although we use panel data nota-

The tradeoffs between these three approaches are not strictly of interest here, although there are well-known differences between pooled, FE and RE (Clark and Linzer 2015). Generally speaking, a fully pooled model (Eq. 1) is never advisable in the panel data context (Clark and Linzer 2015), since any correlation between the observables, x_{it} , and any time invariant unit effects (α_i) will lead to biased estimates of β ; at best, it is still inefficient, since it does not take into account the structure of the data (i.e., observations across time are nested within units), which will likely lead to unit-specific heteroskedasticity. Like the fully pooled model, RE (Eq. 3) also suffers from potential endogeneity caused by any correlation between the included covariates in the model and the (unobserved) time-invariant unit-specific effects in the error term, although efficiency gains can sometimes outweigh such bias (Clark and Linzer 2015). FE models in contrast, are unable to include time-invariant regressors (since they will be perfectly correlated with α_i), and are biased in dynamic models with short time points, or under certain forms of dynamic misspecification (Plümper and Troeger 2019).

This choices might appear somewhat limited; FE models estimate only within effects, while pooled and RE models combine both between and within effects. An emerging alternative is to estimate these effects separately—variants of which are the Mundlak, “hybrid”, or random effects within-between (REWB) models (Bell and Jones 2015; Bell, Fairbrother and Jones 2019):

$$y_{it} = \alpha + \beta_W(x_{it} - \bar{x}_i) + \beta_B\bar{x}_i + v_{it} \quad (4)$$

Including both the unit-demeaned and unit-means of x_{it} separately provides β_W , the within effect (i.e., same as β in the FE model), and β_B , the between effect. Eq. 4 is

tion with time and unit notation, our claims extend to any model that estimates random intercepts for a hierarchical, upper, “level 2” structure in which observations are nested (level i in our notation (c.f., Bell and Jones 2015; Bell, Fairbrother and Jones 2019)).

estimated using RE, as the error components are assumed to be drawn from a normal distribution. REWB has clear advantages. In the context of the “standard” RE model of Eq. 3, REWB can “test whether the assumption of equal within and between effects is true” by seeing if $\beta_W = \beta_B$ (Bell, Fairbrother and Jones 2019, p. 1057). More generally, REWB helps avoid the RE vs. FE false dichotomy often discussed by scholars. Practitioners should consider estimating a REWB model to determine if there are separate within and between effects to uncover. Despite REWB’s advantages, we believe there is still a need to offer interpretative advice to users of the “standard” RE shown in Eq. 3. For one, when certain conditions are met, the RE model remains an appropriate choice as it is a more efficient estimator than REWB.⁴ Moreover, since FE and RE are still the predominant models in the field,⁵ we stress that there is a need for interpreting these models correctly when reviewers or readers expect to see them.

The pooled, FE, RE and REWB models discussed above offer different approaches to handling unit heterogeneity. In fact, we can explicitly formulate the RE model as a compromise between the completely pooled and completely within-variance (FE) approaches. This formulation better allows us to understand how we might arrive at interpretation challenges due to quasi-demeaning when using RE. To better see this, first write the unit-specific average value of each variable (i.e., averaging across time points t) for each unit i as $\bar{y}_i = \bar{\alpha}_i + \beta\bar{x}_i + \bar{\epsilon}_i$. The RE approach operates by conceptualizing and estimating a demeaning parameter, θ , that accounts for unit-specific averages

⁴ When the within and between effects of Equation 4 are equivalent (Bell, Fairbrother and Jones 2019, p. 1057), which practitioners can test by estimating the REWB model.

⁵ In Bell and Jones (2015, p. 134), for instance, less than 10% of random effects models in articles also mentioned “Mundlak.” Even leading proponents of REWB like Bell, Fairbrother and Jones (2019, p. 1056), refer to contemporary examples of this approach as “rare.”

(Wooldridge 2010, 326-328):

$$(y_{it} - \theta \bar{y}_i) = (\alpha - \theta \bar{\alpha}_i) + \beta(x_{it} - \theta \bar{x}_i) + (\varepsilon_{it} - \theta \bar{\varepsilon}_i) \quad (5)$$

If we let $\theta = 0$, then Eq. 5 becomes the pooled model; that is to say, there is no demeaning by the unit means for each variable. Another way of picturing this is that the pooled model fully—and equally—incorporates variation within units as well as between units. In contrast, if we let $\theta = 1$, then Eq. 5 becomes the FE model, since unit means are subtracted out, leaving only within variance.

In contrast to the fully pooled and FE model, the RE model allows θ , the proportion of demeaning from the unit averages, to vary, such that $0 \leq \theta \leq 1$. Therefore, RE can be thought of as a compromise between the fully pooled model (where within and between variance are equally weighted) and the fixed effects model (where within variance is completely removed since $\theta = 1$). As θ gets closer to one, estimation results are based on more within variation relative to between, and thus results should approximate a fixed effects model. In practice, we do not need to make a theoretical determination about the value of θ , as software programs automatically perform this quasi-demeaning. This requires the total number of T time points in the estimated sample for each unit, as well as the two variance components from Eq. 3; “fundamental” error variance, σ_ε^2 , and error variance attributable to unit-specific differences, σ_α^2 (Clark and Linzer 2015). With these three values, θ is estimated as:⁶

$$\hat{\theta} = 1 - \sqrt{\frac{\hat{\sigma}_\varepsilon^2}{T \hat{\sigma}_\alpha^2 + \hat{\sigma}_\varepsilon^2}} \quad (6)$$

Estimation of $\hat{\sigma}_\varepsilon^2$ and $\hat{\sigma}_\alpha^2$ is typically performed via feasible generalized least squares;

⁶Eq. 6 assumes a balanced dataset; in other words, the total number of time points is the same for all units. If this is not the case, we end up indexing T_i , meaning θ_i can differ by unit: $\hat{\theta}_i = 1 - \sqrt{\frac{\hat{\sigma}_\varepsilon^2}{T_i \hat{\sigma}_\alpha^2 + \hat{\sigma}_\varepsilon^2}}$.

for instance, by using the `xtreg, re` command in Stata or in the `plm` package in R (Croissant and Millo 2008). The estimation of the unit variance and error variance differs across various algorithms, typically, by how they handle unbalanced panels and estimation in small samples (c.f., Amemiya 1971; Swamy and Arora 1972). RE estimation can also be performed using maximum likelihood, which involves simultaneously estimating σ_ε^2 and σ_α^2 and the coefficients in a single likelihood function. However, the same demeaning occurs.

To illustrate how θ depends on T , between-unit ($\hat{\sigma}_\alpha^2$) and fundamental ($\hat{\sigma}_\varepsilon^2$) error variance, Figure 1(a) plots values of θ from Eq. 6 across various combinations of σ_α^2 and σ_ε^2 when $T = 10$. Regardless of σ_ε^2 —the variance in the stochastic term—greater variance between units (higher values of σ_α^2) leads to larger values of θ . This means that more quasi-demeaning is occurring, and thus the RE model more closely approximates FE, which only uses within-variation. θ tends to be highest when the between-unit variance is large relative to the stochastic error variance. As the stochastic error variance increases, the proportion of quasi-demeaning decreases; quasi-demeaning tends to be lowest when the between-unit variance is small relative to the stochastic error variance. Figure 1(b) shows values of θ when $T = 100$. Now, θ very quickly increases as the between-unit variance increases, meaning that more within-unit variation is being analyzed, relative to between. This relationship between θ , $\hat{\sigma}_\varepsilon^2$, and $\hat{\sigma}_\alpha^2$ echoes other measures of the division of variance between levels, like the Variance Partition Coefficient (VPC) or the Intra-Class Correlation (ICC).⁷ When choosing a measure to report, though, θ explicitly accounts for unbalanced panels, has a consistent direction (as it is always expressed with $\hat{\sigma}_\varepsilon^2$ in the numerator), and plays a direct role in the underlying

⁷ In the case of a model with a continuous dependent variable and random intercepts, the VPC and ICC are both calculated as the same proportion: $\frac{\hat{\sigma}_x^2}{\hat{\sigma}_\alpha^2 + \hat{\sigma}_\varepsilon^2}$, where x can represent either ε or α (Leckie et al. 2020). Hence, $\hat{\theta}$ resembles the VPC weighted by the length of the panel.

quasi-demeaning happening in RE models.




Figure 1: The Proportion of Quasi-Demeaning, θ , Varies by σ_{α}^2 , σ_{ϵ}^2 and T

There are a few other important features surrounding quasi-demeaning in RE models. For one, the quasi-demeaning transformation applies to *all* variables in a model. In addition, the quasi-demeaning process sometimes estimates a higher level variance, σ_{α}^2 , as zero. While this may not be an econometric problem,⁸ it presents a poorly understood issue for interpretation: namely that the RE results will be identical to the pooled model. In this world, special care must be taken to *interpret* it as a fully pooled model, lacking any particular benefits of the RE approach. Taken together, quasi-demeaning and the RE model suggest special attention needs to be paid to the counterfactuals used when generating inferences (as the quasi-demeaning process explicitly changes the estimated variance of the variables), as well as whether RE results in pooled modeling.

⁸While Greene (2018, p. 409) suggests this is a model specification issue, Baltagi (2008, p. 24) argues it is “not a bad sin” since σ_{α}^2 is likely very small.

2 New insights or old hat?

A skeptic may claim that the above discussion and its implications are already well understood by authors using RE. To evaluate this, we reviewed the use of RE modeling in top political science journals (*American Political Science Review*, *American Journal of Political Science*, *Journal of Politics*, and *British Journal of Political Science*) from 2012 to 2020. A discipline that understands the link between FE, RE, and the fully pooled model should consistently be reporting the variance at each level ($\hat{\sigma}_\varepsilon^2$ and $\hat{\sigma}_\alpha^2$), the quasi-demeaning parameter $\hat{\theta}$ (or similar quantities like ICC or VPC), and should be accounting for this quasi-demeaning through the counterfactuals provided in the text. The findings are summarized in Table 1.

Table 1: RE models and interpretation in top journals

Category	Number of Articles
Total that use RE	160
Report $\hat{\sigma}_\varepsilon^2$ and $\hat{\sigma}_\alpha^2$	5
Report numeric “variance of ...”	57
Report Intra-Class Correlation	3
Report Variance Partition Coefficient	1
Report $\hat{\theta}$	0
Report a variance estimate of 0	20 of 57 (35%)
Interpret substantively beyond a one-unit effect	102 of 130 (78%)
Considers within-unit range through interpretation	1 of 130 (1%)

Includes all tables of “random effects” or “mixed” models in the *American Political Science Review*, *American Journal of Political Science*, *Journal of Politics*, and *British Journal of Political Science* from 2012 to 2020.

We found 160 papers use RE models in the main text (instead of a robustness check in an appendix). Of these, well under half report $\hat{\sigma}_\varepsilon^2$ or $\hat{\sigma}_\alpha^2$, or any measure of variance, precluding an inference about the extent of quasi-demeaning in the data. Only four report a quantity describing the relative variance, and none of the articles explicitly reported θ , the parameter that captures the quasi-demeaning. Moreover, 20 of the articles reported an estimated variance of 0, suggesting that the authors were actually estimating a fully pooled model *despite* no discussion as such in the text. In fact, it was the opposite; authors would often vaunt the appropriateness of RE because of its between-unit variance properties, even though these were not being reflected in the

model.

With regards to substantive interpretation, 130 articles drew inferences beyond basic sign and significance. Only 28 reported the effect of a one-unit change. The remaining 102 reported counterfactuals relative to the original (not quasi-demeaned) range of the variables, sometimes as large as the full minimum to maximum: a comparison which is *never* observed within any unit after quasi-demeaning. Therefore, we think our claims, despite perhaps being transparent to experts, are still mishandled among practitioners, even in top journals. We further elaborate through applied examples below.

3 Substantive examples from the literature

Example I: Quasi-demeaning and avoiding extreme counterfactuals

To see the extent to which variation may be reduced beyond a counterfactual given in a paper, we replicate Williams (2017), who examines why some development projects like schools are started but never finished. Williams reports using the RE GLS estimator (Table 2, Model 1 in his article) to analyze Ghanaian local development projects in 327 districts from 2011 to 2013. The probability a project is unfinished is a function of whether the project is government funded, the vote share for the National Democratic Congress (NDC) party, as well as controlling for project type, construction type, project duration, and year. RE are included for districts, which are the subject of our investigation here. The model is strictly replicated in Table 2.

Our principal focus is the range of the transformed data once quasi-demeaning has occurred. Williams' key predictor is NDC vote share: the proportion of the vote for the National Democratic Congress, which can range from 0 (no votes) to 1 (all of the votes). Quasi-demeaning implies, however, that the proportion in any given district can *never* range from 0 to 1, since it will be adjusted downward by the unit-specific mean, according to Eq. 3. Following the formula in Eq. 6, we calculate the quasi-demeaned data

Table 2: Replication of Williams (2017), Table 2, Model 1

	RE Model	
Government-funded	-0.109***	(0.015)
NDC Vote Share	0.109	(0.082)
Constant	1.017***	(0.157)
σ_{ε}^2	0.412	
σ_{α}^2	0.226	

Note: Dependent variable is probability of project completion. Not shown: FE for 22 project types, FE for construction type, FE for years, and number of years since project start. Model appears exactly as in original article. $N = 4563$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

relative to the original data for the key predictor, NDC vote share. Compared to the original data, the range of quasi-demeaned data is dramatically reduced. Figure 2 illustrates exactly this; Williams interprets NDC across the range of 0.1 to 0.9, implying an effect of $[0.9 - 0.1] \times 0.109 = 0.8 \times 0.109 \approx 0.09$ unit increase in the probability a project is completed. However, once the data are quasi-demeaned, instead of moving NDC vote share from 0.1 to 0.9, the new \pm one standard deviation moves NDC vote share from 0.1 to just 0.29, implying an effect of only a $[0.29 - 0.10] \times 0.109 = 0.19 \times 0.109 \approx 0.02$ unit increase in the probability of a project being completed. Interpreting the original counterfactual leads to an estimated effect that is quadruple the size of a reasonable effect calculated using the demeaned data.

Williams goes on to estimate an interactive model between whether a project is government-funded and the NDC vote share, then shows a marginal effects plot displaying the effect of government funding across a range of NDC vote share from 0.1 to 0.9. Accounting for the quasi-demeaning, as shown in Figure 2, the 99.5 percentile of cases on the demeaned NDC vote share variable is 0.499. In other words, over half of the range of the interpreted interaction (all values over 0.50) *has no cases*, since less than 0.5% of units have a quasi-demeaned value of NDC vote share of over 0.5. Without special attention to this quasi-demeaning post-estimation, we are at significant risk of drawing implausible (at best) or impossible (at worst) counterfactuals from our RE models.




Figure 2: Quasi-demeaned versus untransformed vote share

Example II: Quasi-demeaning estimates no variance; RE are a pooled model

For the second applied example, we show how in certain instances, estimation of the quasi-demeaning parameter, $\hat{\theta}$, results in a model that is no different from the fully pooled model. We replicate Valdini and Lewis-Beck (2018), who examine how institutional rules affect the vote share of incumbent parties. Table 3, Model 1 shows the same results as the authors, using the RE GLS estimator as they do: except we report the estimated variances σ_{ε}^2 and σ_{α}^2 .⁹ Note that σ_{α}^2 , the partitioned error variance attributable to unit-specific differences, is estimated at 0. This means that $\hat{\theta} = 0$, thus no quasi-demeaning is taking place, and the estimator is equivalent to the fully-pooled model. As evidence of this, Model 2 in Table 3 shows the results using the fully pooled model

⁹Table 4, Model 1 in their article.

Table 3: RE *are* a fully pooled model when quasi-demeaning estimates no variance

	(1)		(2)	
	RE-GLS		Pooled Model	
Previous Incumbent Vote	0.530*	(0.284)	0.530*	(0.284)
GDP Growth Rate _{<i>t</i>-1}	2.076***	(0.638)	2.076***	(0.638)
Electoral Stability	0.309***	(0.068)	0.309***	(0.068)
Trade Openness	0.134**	(0.058)	0.134**	(0.058)
Concurrent Elections	2.756	(2.539)	2.756	(2.539)
GDP Growth Rate _{<i>t</i>-1} × Concurrent Elections	-0.894**	(0.447)	-0.894*	(0.447)
Constant	-19.383	(18.795)	-19.383	(18.795)
σ_{ε}^2	261.79		230.81	
σ_{α}^2	0		NA	

Note: Dependent variable is incumbent party vote share. Coefficients shown with robust standard errors clustered on 18 countries in parentheses for both models. $N = 92$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

estimator, which has exactly the same results as in Model 1. In their text, though, the authors suggest that a “full set of corrections for statistical efficiency” are employed, even though they are “limited by sample size” and the amount of variation in their data (Valdini and Lewis-Beck 2018, 418-419). However, the statistical “correction” did not include any quasi-demeaning: negating any supposed advantage of the RE model.

By no means are we criticizing the modeling strategy of Valdini and Lewis-Beck (2018). Moreover, they are in good company; recall from Table 1 that 35% of the articles that reported variance estimates actually estimated 0 variance: a RE model where quasi-demeaning resulted in a fully-pooled model without the researchers noting it as such. As discussed above, while econometricians differ about whether this implies fundamental model misspecification or something more benign (Baltagi 2008; Greene 2018), we observe that these instances of $\sigma_{\alpha}^2 = 0$ are never discussed in the main results as a fully pooled model. In other words: scholars are introducing, reporting, and interpreting fully-pooled models as if they are RE models. This misrepresentation indicates the need for a broad prescription on what to report when estimating RE models.

Conclusions and recommendations

We emphasize two claims. First—similar to FE models (Mummolo and Peterson 2018)—scholars need to account for the quasi-demeaned nature of their independent variables when interpreting RE models. Second, it is possible, maybe even common, for a model that a scholar believes to be a RE model to actually just be a fully pooled model due to a lack of higher-level variance to be estimated. This may not be a statistical problem, but when it happens, we should interpret the model more faithfully to what’s actually being estimated.

As such, we offer the following recommendations:

1. Scholars should discuss the substantive and statistical significance of their effects with plausible counterfactual changes that actually occur in the (transformed) sample space of data the model is estimated on. In order to assist users, we have written `qemean`, functions for both Stata and R that automatically calculate the proportion of demeaning that is occurring in their variables when a RE model is estimated.¹⁰ These programs generate these quasi-demeaned variables, from which users can create more realistic counterfactual scenarios (i.e., avoiding the extreme counterfactuals shown in Figure 2).
2. Scholars should report the estimated variance components of the RE model (σ_{ϵ}^2 and σ_{α}^2): currently reported by under half of RE models (Table 1). Additionally, it’s useful to report $\hat{\theta}$ (the estimated proportion of quasi-demeaning), which summarizes the partitioning of the variance at different levels in a single statistic. Values of $\hat{\theta}$ close to 1 indicate that the RE model is utilizing mostly within-variation, while small values of $\hat{\theta}$ close to 0 indicate that a substantial proportion of between variation is being used in estimation, which may lead to differences between the random and FE models. Using $\hat{\theta}$ as a diagnostic would help better identify cases where RE devolves to the fully pooled model. In these cases, scholars should

¹⁰Available at <https://github.com/andyphilips/qdmean>.

simply use (and report) the results from a pooled model, or consider using a FE model, since none of the claimed advantages of RE apply.

3. Scholars might consider another approach entirely. While the points above are crucial given the popularity of RE in political science, alternatives exist. A fully-pooled model is one option; while it retains overall variation, it makes heroic assumptions about observation independence that are probably not met in most panel data applications. Despite this, under certain conditions it may still perform better than other models in the presence of dynamic misspecification (Plümper and Troeger 2019). As discussed above, an alternative approach is to parse the effect of covariates into specific within- and between-variation components (Bell and Jones 2015; Bell, Fairbrother and Jones 2019). This would allow users to examine counterfactual changes within a unit separately from those between units, as well as examine whether combining within and between effects (as the pooled does and the RE does to some extent) is valid.

Regardless of the modeling approach used, when working with panel data, users should always be aware of the assumptions behind each strategy, as well as understand any data transformations (or lack thereof) that have occurred.

References

- Amemiya, Takeshi. 1971. "The estimation of the variances in a variance-components model." *International Economic Review* pp. 1–13.
- Baltagi, Badi. 2008. *Econometric analysis of panel data*. John Wiley & Sons.
- Bell, Andrew and Kelvyn Jones. 2015. "Explaining fixed effects: Random effects modeling of time-series cross-sectional and panel data." *Political Science Research and Methods* 3(1):133–153.

- Bell, Andrew, Malcolm Fairbrother and Kelvyn Jones. 2019. "Fixed and Random effects models: making an informed choice." *Quality & Quantity* 53(2):1051–1074.
- Bjornskov, Christian and Jacob Mchangama. 2019. "Do social rights affect social outcomes?" *American Journal of Political Science* 63(2):452–466.
- Busemeyer, Marius R, Aurélien Abrassart, Spyridoula Nezi and Roula Nezi. 2019. "Beyond Positive and Negative: New Perspectives on Feedback Effects in Public Opinion on the Welfare State." *British Journal of Political Science* pp. 1–26.
- Clark, Tom S and Drew A Linzer. 2015. "Should I use fixed or random effects?" *Political Science Research and Methods* 3(2):399–408.
- Croissant, Yves and Giovanni Millo. 2008. "Panel data econometrics in R: The plm package." *Journal of Statistical Software* 27(2):1–43.
- Greene, William H. 2018. *Econometric analysis*. 8 ed. Pearson.
- King, Gary and Langche Zeng. 2006. "The dangers of extreme counterfactuals." *Political Analysis* 14(2):131–159.
- Leckie, George, William J. Browne, Harvey Goldstein, Juan Merlo and Peter C. Austin. 2020. "Partitioning variation in multilevel models for count data." *Psychological Methods* 25(6):787–801.
- Mummolo, Jonathan and Erik Peterson. 2018. "Improving the interpretation of fixed effects regression results." *Political Science Research and Methods* 6(4):829–835.
- Plümper, Thomas and Vera E Troeger. 2019. "Not so Harmless After All: The Fixed-Effects Model." *Political Analysis* 27(1):21–45.
- Schunck, Reinhard. 2013. "Within and between estimates in random-effects models: Advantages and drawbacks of correlated random effects and hybrid models." *The Stata Journal* 13(1):65–76.

- Swamy, PAVB and Swarnjit S Arora. 1972. "The exact finite sample properties of the estimators of coefficients in the error components regression models." *Econometrica: Journal of the Econometric Society* pp. 261–275.
- Valdini, Melody E and Michael S Lewis-Beck. 2018. "Economic Voting in Latin America: Rules and Responsibility." *American Journal of Political Science* 62(2):410–423.
- Williams, Martin J. 2017. "The political economy of unfinished development projects: Corruption, clientelism, or collective choice?" *American Political Science Review* 111(4):705–723.
- Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data*. 2 ed. MIT Press.