Political Science Research and Methods Vol 3, No. 2, 399–408 May 2015
© The European Political Science Association, 2014
doi:10.1017/psrm.2014.32

Should I Use Fixed or Random Effects?

TOM S. CLARK AND DREW A. LINZER*

Empirical analyses in social science frequently confront quantitative data that are clustered or grouped. To account for group-level variation and improve model fit, researchers will commonly specify either a fixed- or random-effects model. But current advice on which approach should be preferred, and under what conditions, remains vague and sometimes contradictory. This study performs a series of Monte Carlo simulations to evaluate the total error due to bias and variance in the inferences of each model, for typical sizes and types of datasets encountered in applied research. The results offer a typology of dataset characteristics to help researchers choose a preferred model.

In social science research, it is common to confront data that are clustered or grouped into higher-level units. One of the most frequently encountered challenges when modeling these data arises when the dependent variable exhibits group-level variation beyond what can be explained by the independent variables alone. In these cases, fitting a standard linear regression or generalized linear model without accounting for the grouped nature of the observations can lead to poorly fitting models and misleading estimates of both the effect of independent variables of interest and of the precision of those estimates (Beck and Katz 1995; Greene 2012).

The two dominant approaches to remedy this problem are the use of so-called *fixed-effects* or *random-effects* models.[1] Although much has been written on the theoretical properties of both approaches (for example, Kreft and DeLeeuw 1998; Robinson 1998; Kennedy 2003; Frees 2004; Gelman 2005; Wilson and Butler 2007; Arceneaux and Nickerson 2009; Wooldridge 2010; Greene 2012), recommendations for applied researchers are often confusing—or even contradictory (Gelman and Hill 2007, 245). Often they are made with reference to idealized datasets with very large sample sizes, or using divergent standards for assessing model quality. There remains little consistent guidance for researchers trying to decide how best to model the data they have on hand. They are left to wonder: "Should I use fixed or random effects?"

In this article, we offer practical guidance for researchers choosing between fixed- and random-effects models. As we describe below, both models entail a series of assumptions that might be violated in any given dataset. Under certain conditions, random effects models can introduce bias, but reduce the variance of estimates of coefficients of interest. Fixed-effects estimates will be unbiased, but may be subject to high sample dependence. We argue that researchers ought not to place undue weight on minimizing either bias or variance, but rather consider the trade-off between the two in either model. While it is true that under a

* Tom Clark is Asa Griggs Candler Professor of Political Science, Emory University, 1555 Dickey Drive, Atlanta, GA 30030 USA (email: tom.clark@emory.edu). Drew Linzer is Assistant Professor, Department of Political Science, Emory University (email: drew@votamatic.org). We thank Kyle Beardsley, Justin Esarey, Andrew Gelman, Kosuke Imai, Benjamin Lauderdale, Jeffrey Lax and Jamie Monogan for helpful discussions and feedback. Nigel Lo provided valuable research assistance.

[1] We note that there is considerable confusion in the literature concerning the meanings of these terms (for a discussion, see Gelman 2005, 20). We employ them here as we believe most applied researchers use them: as shorthand for two modeling approaches. As we describe below, fixed effects refer to a series of dummy variables for the units from which grouped data arise, while random effects refer to an estimator that assumes unit effects are drawn from an underlying, modeled distribution. There are certainly many other modeling options one might consider.

https://doi.org/10.1017/psrm.2014.32 Published online by Cambridge University Press

random-effects specification there may be bias in the coefficient estimates if the covariates are correlated with the unit effects, it does not follow that *any* correlation between the covariates and the unit effects implies that fixed effects should be preferred. What should be judged instead is how *much* bias is created, and how much variance would be introduced by using fixed effects instead. After all, except in exceptional circumstances, there will always be some level of correlation between the covariates and the unit effects, and thus at least minimal bias (for example, Angrist and Pischke 2009, 223n2). The question is, how much is too much?

We compare the performance of the fixed- and random-effects models using a series of Monte Carlo experiments that vary the sample size, effect size of the independent variable, correlation between the independent variable and unit effects, and, crucially, whether the majority of variation in the dependent variable is within or between units. We then calculate how consistently each model recovers the true coefficient of interest. Drawing upon our simulation results, we derive a set of guidelines that applied researchers can use to determine how best to specify models for grouped dependent variables.

THE PROBLEM

We consider the linear model for observations $i = 1 \ldots N$ grouped into units $j = 1 \ldots J$,

$$y_i = \alpha_{j[i]} + \beta x_i + \varepsilon_i; \qquad \varepsilon_i \sim N(0, \sigma_y^2). \tag{1}$$

The effect of x on y, denoted β, is the primary quantity of interest. We assume that β is the same within each unit.[2] However, even after accounting for the effect of x, there may still remain additional variation in the overall level of y across units. The *unit effect* α_j captures the amount by which predictions of y in unit j must be adjusted upward or downward, given knowledge only of x. The notation $j[i]$ indicates the unit j of observation i.

One interpretation of the unit effects is that they represent ignorance about all of the other systematic factors that predict y, other than x. If these factors were known, they could ostensibly be included as additional covariates in the model, thus "explaining" the extra variation in y and eliminating variation in α_j across units. Since these variables are not included in the model, we capture their effects with α_j instead. The variation in α_j might also be partially or completely nonsystematic, due simply to stochastic noise.

With few exceptions, failing to allow for the possibility that α_j varies by unit will lead to biased estimates of β. If we assume that the unit effects are all equivalent—that is, $\alpha_j = \alpha_k$ for all j and k —then Equation 1 reduces to the *pooled* model

$$y_i = \alpha + \beta x_i + \varepsilon_i; \qquad \varepsilon_i \sim N(0, \sigma_y^2). \tag{2}$$

The pooled regression model is appropriate if α_j does not vary once x is included as an independent variable. The pooled model will also not produce bias in estimates of β if the unit effects differ but are uncorrelated with x. In most applications, however, the unit effects are associated to some degree with x, so variation in α_j must be modeled in order to avoid faulty inferences about β.

[2] Although there are many instances in which a researcher may wish to allow β to vary by unit, Equation 1 represents the most commonly encountered modeling scenario. Our notation follows that of Gelman and Hill (2007, 256–7).

https://doi.org/10.1017/psrm.2014.32 Published online by Cambridge University Press

TWO SOLUTIONS: FIXED AND RANDOM EFFECTS

There are two standard approaches for modeling variation in α_j: *fixed effects* and *random effects*. The fixed-effects model is a linear regression of y on x, which adds to the specification a series of indicator variables z_j for each unit, such that $z_{j[i]} = 1$ if observation i is in unit j, and $z_{j[i]} = 0$ otherwise:

$$y_i = \sum_{j=1}^{J} \alpha_j z_{j[i]} + \beta x_i + \varepsilon_i; \qquad \varepsilon_i \sim N(0, \sigma_y^2). \tag{3}$$

The coefficients $\hat{\alpha}_j$ that are computed for each z_j are taken as estimates of α_j.

In the random-effects model, the α_j are instead assumed to follow a probability distribution, with parameters estimated from the data. This distribution is typically normal, with average unit effect μ_α and variance σ_α^2, which describes by how much the other unit effects vary around the mean.

$$y_i = \alpha_{j[i]} + \beta x_i + \varepsilon_i; \qquad \alpha_j \sim N(\mu_\alpha, \sigma_\alpha^2); \qquad \varepsilon_i \sim N(0, \sigma_y^2). \tag{4}$$

As Gelman and Hill (2007, 258) note, the random-effects estimator is equivalent to the fixed-effects estimator when we assume that $\alpha_j \sim N(\mu_\alpha, \infty)$ rather than $\alpha_j \sim N(\mu_\alpha, \sigma_\alpha^2)$.

The estimators of the coefficient β under the two models are mathematically related. Let $\mathbf{x}_j$ and $\mathbf{y}_j$ represent the n observations in the jth unit. Following Greene (2012, 373), the generalized least squares estimator of β for the random-effects model is:

$$\hat{\beta}_{RE} = \left(\sum_{j=1}^{J} \mathbf{x}'_j \Sigma^{-1} \mathbf{x}_j \right)^{-1} \left(\sum_{j=1}^{J} \mathbf{x}'_j \Sigma^{-1} \mathbf{y}_j \right), \tag{5}$$

where, letting $\mathbf{I}$ be the $n \times n$ identity matrix and $\mathbf{i}$ be an $n \times 1$ column of ones,

$$\Sigma^{-1/2} = \frac{1}{\sigma_y} \left[\mathbf{I} - \frac{\theta}{n} \mathbf{i}\mathbf{i}' \right] \tag{6}$$

with

$$\theta = 1 - \frac{\sigma_y}{\sqrt{\sigma_y^2 + n\sigma_\alpha^2}}. \tag{7}$$

As $\sigma_\alpha^2 \to \infty$ or $n \to \infty$, $\theta \to 1$, making the random-effects estimator $\hat{\beta}_{RE}$ reduce to the fixed-effects estimator $\hat{\beta}_{FE}$ (Greene 2012, 361).

How to Choose?

In some cases, choosing between a fixed- and random-effects specification will follow directly from a researcher's theoretical model. Or a researcher may have application-specific concerns about the appropriateness of the assumptions underlying either model. Unfortunately, however, too often one's theoretical model does not dictate a particular specification, and the theoretical assumptions about which one might be concerned do not lend themselves to empirical evaluation. In these situations, neither specification emerges as an obvious choice. In this context, we investigate empirically which model offers better inferences about the quantities of interest. Both models have potential advantages—as well as disadvantages—to consider when selecting an approach.

https://doi.org/10.1017/psrm.2014.32 Published online by Cambridge University Press

The problem of high variance

The estimate of β in the fixed-effects model may, under certain conditions, produce estimates that are highly sample dependent—that is, overly sensitive to the random error in a given dataset. Suppose that there are few observations per unit, or that x varies little within each unit, relative to the amount of variation in y. Estimates of the *within-unit* effects of x on y may then diverge considerably from the true effect due to chance alone.[3] If there is also a small number of units, then many of the within-unit effects may diverge from the true effects in the same direction, leading the estimate of β to be quite different from the true β. This lack of robustness to potentially anomalous samples is what is meant by the fixed-effects model having high *variance*.

A related drawback of fixed-effects models is that they require the estimation of a parameter for each unit—the coefficient on the unit dummy variable. This can substantially reduce the model's power and increase the standard errors of the coefficient estimates. The problem is worsened when the within-unit sample size is very small, as the unit effects alone may account for most of the variation in the dependent variable.

Random-effects models enable estimation of β with lower sample-to-sample variability by *partially pooling* information across units (Gelman and Hill 2007, 258). By estimating the variance parameter σ_α^2 in Equation 4, the random-effects estimator forms a compromise between the fixed-effects and pooled models. Groups with outlying unit effects will have their respective α_j *shrunk* back toward the mean, μ_α. This brings estimates of β away from the less stable fixed-effects estimate and closer to the more stable (albeit potentially biased) pooled estimate. The effects of shrinkage will be greatest for units containing fewer observations, especially when estimates of σ_α^2 are close to zero.

The problem of bias

The most serious drawback of the random-effects approach is the problem of bias that partial pooling can introduce in estimates of β. To eliminate this bias, the random-effects estimator requires there to be no correlation between the covariate of interest, x, and the unit effects, α_j. To illustrate, suppose that there is a variable z that predicts y but is not included as a covariate in the random-effects model. As a result of omitting z from the model specification, the higher or lower levels of y in unit j due to z are instead accounted for by the unit effects α_j. For there to be no bias in estimates of the coefficient on x, there must be no correlation between x and z—and, hence, no correlation between x and α_j, implying no confounding due to the omitted z. Since the random-effects model does not estimate separate unit effects, any correlation between x and α_j can imply an omitted variable z that produces bias in estimates of β. The greater the magnitude of the correlation between x and α_j, the greater the bias in estimates of β.[4]

Scholars are sometimes advised to use a Hausman (1978) specification test to detect violations of the random-effects modeling assumption that the explanatory variables are orthogonal to the unit effects. A "significant" test result is taken as evidence of a correlation between x and α_j, implying that the random-effects model should be rejected in favor of the fixed-effects model. However in most applications, the true correlation between the covariates and unit effects is not exactly zero. Therefore, if the Hausman test fails to reject the null hypothesis of orthogonality, it is most likely not because the true correlation is zero—and, hence, that the random-effects estimator is unbiased.

[3] In the bivariate linear regression model, $Var(\hat{\beta})$ increases with smaller values of $Var(x)$, and with larger values of σ_y^2, the conditional variance of y given x (Greene 2012, 48).

[4] We set aside the issue of causal inference with observational data and focus explicitly on the bias and variance of these two estimators under varying data-generating processes.

https://doi.org/10.1017/psrm.2014.32 Published online by Cambridge University Press

Rather, it is likely that the test has insufficient statistical power to reliably distinguish a small correlation from zero correlation. When using the random-effects model, there will still be bias (if perhaps negligible) in estimates of β, even if the Hausman test does not find a significant result. Of course, in many cases, a biased (random-effects) estimator can be preferable to an unbiased (fixed-effects) estimator if the former provides sufficient variance reduction over the latter. The Hausman test does not help evaluate this trade-off.

Practical considerations

In addition to these theoretical considerations, there are practical and technical issues that researchers might take into account when deciding between a fixed- and random-effects estimator. For example, it is very common for a researcher to want to include in the specification an important covariate of interest that does not vary within units. In this case, the unit-invariant predictor will be perfectly collinear with the set of unit dummy variables, making it impossible to estimate the unique effects of that variable. Alternatively, the independent variable may exhibit extremely minimal variation within each unit. In time-series cross-sectional data, independent variables that change very gradually over time are frequently referred to as *slow moving* or *sluggish*. If the correlation between the sluggish covariate and the unit fixed effects is high enough, this can destabilize estimates of the effect of the independent variable. Plümper and Troeger (2007, 2011) propose alternative modeling strategies for data that exhibit these characteristics, which rely on estimating fixed effects and then decomposing those estimates to assess the effect of the sluggish variables.

What if a researcher is interested in making predictions about units that are not in the dataset? When employing a fixed-effects estimator, making out-of-sample predictions is not possible because the unit effects for unobserved units are unknown. In the random-effects specification, out-of-sample predictions are feasible using the model's estimate of the underlying distribution of unit effects in the population.

EVALUATING THE BIAS-VARIANCE TRADE-OFF

We perform a series of Monte Carlo experiments to determine the conditions under which a fixed- or random-effects model provides better estimates of β. Our study investigates variation in the number of units, the number of observations within each unit, the strength of correlation between x and the unit effects, the strength of association between x and y, and the amount of variation in x within units.

To simulate a data-generating process in which observations are clustered by units, we first generate a series of J within-unit means $\bar{x}_j$, and corresponding unit effects α_j, by sampling from a bivariate normal distribution centered at zero:

$$\begin{bmatrix} \alpha_j \\ \bar{x}_j \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} \right). \tag{8}$$

The variances of both α_j and $\bar{x}_j$ are fixed at 1. The off-diagonal covariances ρ control the amount of correlation between the independent variable and the unit effects. We then draw n observations of x_i within each unit $j = 1 \ldots J$ from a normal distribution with mean $\bar{x}_j$ and standard deviation σ_x. The total sample size is $J \times n$. Finally, we apply Equation 1 to produce y_i as a linear function of x_i, with slope β, unit-level constant terms α_j and within-unit error variance $\sigma_y^2 = 1$.

We choose hypothetical values of J, n, ρ, σ_x and β to mimic typical features of quantitative social science datasets (Table 1). In the terminology of longitudinal data analysis, our

https://doi.org/10.1017/psrm.2014.32 Published online by Cambridge University Press

TABLE 1 *Parameters Manipulated in the Monte Carlo Experiments and their Assumed Values*

Parameter	Description	Values
J	Number of units	10, 40, 100
n	Observations per unit	5, 20, 50
ρ	Correlation between independent variable and unit effects	0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.95
σ_x	Within-unit standard deviation of x	0.2 (sluggish), 1 (standard)
β	Within-unit effect of x on y	0, 0.5, 1, 2

simulations assess both *short panels*, in which $J>n$, as well as *long panels*, in which $n>J$.[5] To control the variation in x, we set $\sigma_x = 0.2$ to represent a sluggish independent variable, and $\sigma_x = 1$ for what we refer to as the *standard* case in which individuals differ more greatly within units, making the units more similar on average. Finally, we allow the correlation between the unit effects α_j and the means $\bar{x}_j$ to vary from $\rho = 0$ to $\rho = 0.95$.[6]

For each simulated dataset, we estimate the fixed-effects model (Equation 3), the random-effects model (Equation 4) and the pooled model (Equation 2), and record the estimates of $\hat{\beta}$ produced by each.[7] We then repeat this process for 2,000 simulated datasets for each combination of values in Table 1.

COMPARING RMSE

We compute the root mean square error (RMSE) of $\hat{\beta}$ from the fixed-effects, random-effects and pooled models. When there is a standard amount of variation in the independent variable, there is no appreciable difference in the quality of inferences between the fixed- and random-effects models, except at very small numbers of observations per unit, and extremely high correlation between the independent variable and unit effects (Figure 1). The conventional understanding that any correlation between regressors and unit effects necessarily results in unallowable levels of error in the random-effects estimator is therefore unfounded. Researchers should feel secure using either fixed- or random-effects models under standard conditions, as dictated by the practical and theoretical aspects of a given application. Either way, both approaches are strictly preferable to the pooled model.

In the case of a sluggish independent variable, there is a range of scenarios in which the RMSE of $\hat{\beta}$ is lower with random effects than with fixed effects, even when the random-effects estimate is biased (Figure 2). The smaller the dataset, the greater the potential support for a random-effects specification if the correlation between x and the unit effects is sufficiently low. When the number of units and/or observations per unit is small—fewer than 200 total observations, as a rule of thumb—the RMSE of the fixed-effects estimator can be quite high. This is due to the high variance of the fixed-effects estimator in small samples, despite its unbiasedness. The random-effects model, in contrast, can have much lower variance in small datasets. As long

[5] Applications with unbalanced panels (i.e., where n varies by unit) can be mapped to our results according to the average number of observations per unit.

[6] Negative correlation will lead to the same results, but in the opposite direction. In any given dataset, a researcher may obtain an approximate estimate of ρ by fitting the fixed-effects model, and then computing the correlation between the estimated unit effects and the within-unit means of the independent variable.

[7] The simulation is performed in R (R Development Core Team 2014). We estimate the random-effects model using the function lmer in the lme4 package (Bates, Maechler and Bolker 2011).

https://doi.org/10.1017/psrm.2014.32 Published online by Cambridge University Press

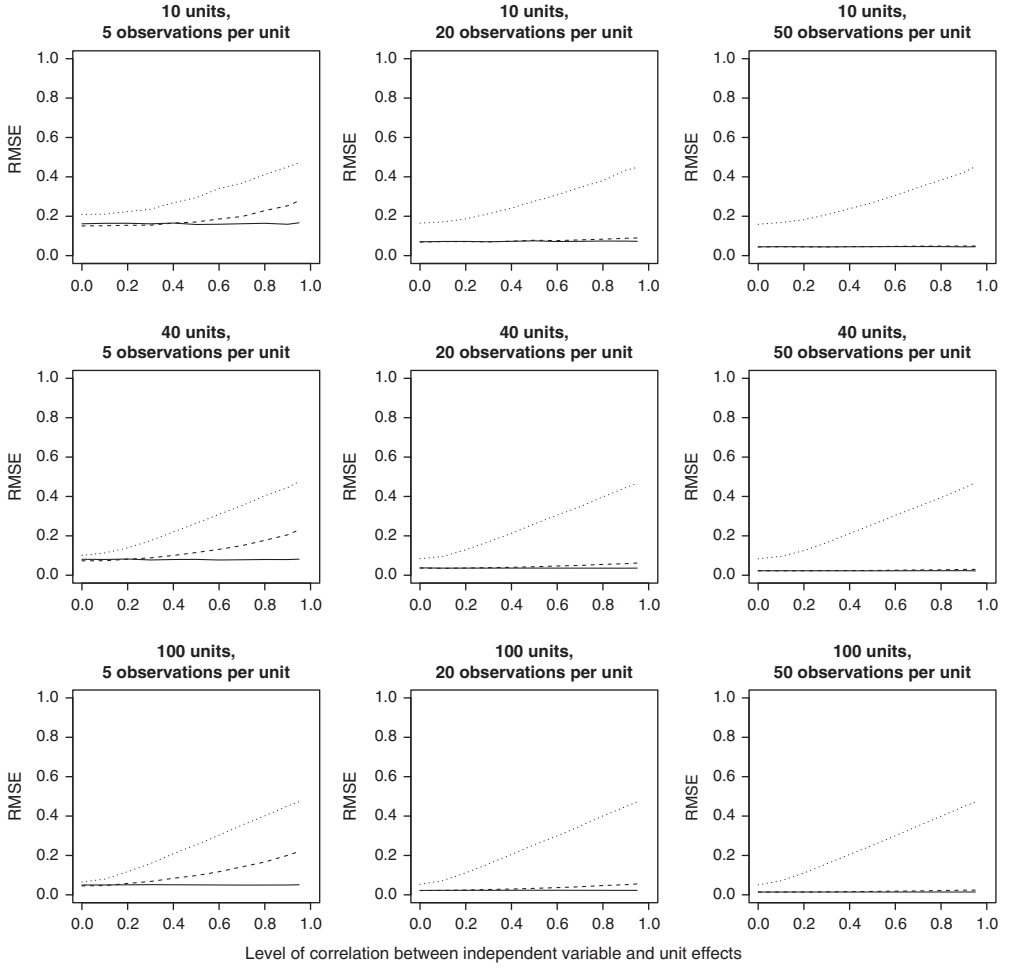


Fig. 1. Root mean squared error of slope parameter estimates—standard case
Note: lines represent the average RMSE of estimates $\hat{\beta}$ across multiple simulated datasets: fixed effects (solid line), random effects (dashed) and pooled (dotted). The horizontal axis is the true amount of correlation between $\bar{x}_j$ and unit effects α_j. An RMSE of zero indicates that estimates are both unbiased and subject to negligible estimation uncertainty. Each panel shows a particular combination of the number of units and number of observations per unit.

as the correlation between the covariate and unit effects is not too high, the bias in the random-effects estimator will be small enough that the lower variance of the estimator will produce a RMSE below that of the fixed-effects model. In the very smallest datasets, the random-effects estimator outperforms the fixed-effects estimator even when there are extreme violations of the assumption of zero correlation.

As the size of the dataset increases, and the variance of the fixed-effects estimates falls, there is less support for the random-effects model. In larger datasets, correlation of greater than 0.2 to 0.3 between the independent variable and unit effects leads to a preference for the fixed-effects model. Increasing the sample size decreases the RMSE of the random-effects estimator, but its bias at high levels of correlation is no longer sufficiently offset by efficiency gains relative to the

https://doi.org/10.1017/psrm.2014.32 Published online by Cambridge University Press

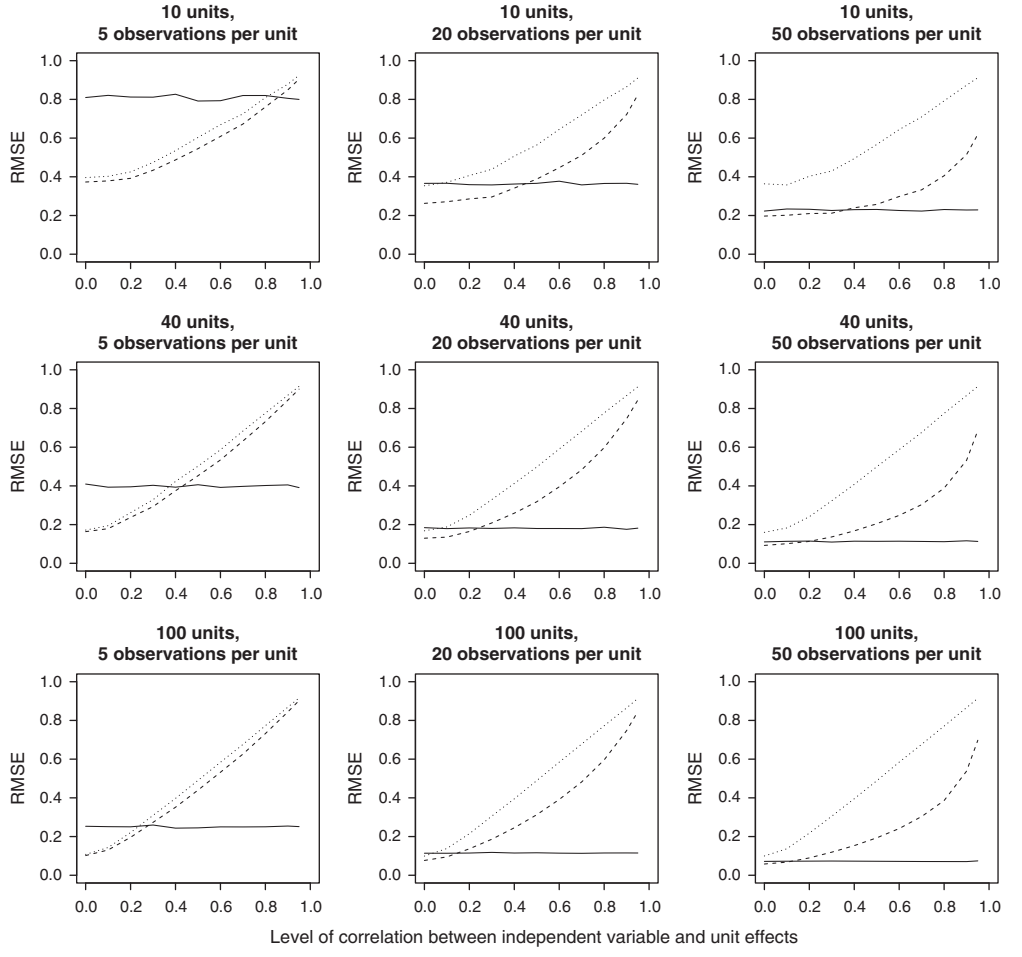


Fig. 2. Root mean squared error of slope parameter estimates—sluggish case
Note: lines represent the average RMSE of estimates $\hat{\beta}$ across multiple simulated datasets: fixed effects (solid line), random effects (dashed) and pooled (dotted). The horizontal axis is the true amount of correlation between $\bar{x}_j$ and unit effects αj. An RMSE of zero indicates that estimates are both unbiased and subject to negligible estimation uncertainty. Each panel shows a particular combination of the number of units and number of observations per unit.

fixed-effects model. At low (but still non-zero) levels of correlation, however, the random-effects model remains a superior choice.[8] The conventional wisdom that any violation of the random-effects model's assumption of zero correlation rules out its use is once again shown to be misguided. The presence of non-zero correlation between the independent variable and unit effects is neither a sufficient nor a necessary condition for choosing a fixed-effects model. Instead, the decision must be based on the amount of data in a study and the level of correlation between regressors and unit effects.

[8] The RMSE of each estimator is invariant to effect size. The preceding simulations set the within-unit effect of x on y at $\beta = 1$. Repeating the simulations with $\beta = 0$, $\beta = 0.5$ and $\beta = 2$, the results exactly match those shown in Figures 1 and 2.

https://doi.org/10.1017/psrm.2014.32 Published online by Cambridge University Press

Finally, we note that for a sluggish independent variable, under all conditions, the pooled estimator yields a higher RMSE than the random-effects estimator, though it can outperform the fixed-effects estimator when the data are sufficiently sparse. The superiority of the random-effects model over the pooled model increases with the number of observations per unit.

CONCLUSION

Scholars generally approach grouped data using either fixed-effects or random-effects models. Advice on which to choose typically emphasizes avoiding bias in estimates of the parameters of interest. Examining the RMSE of both estimators, however, we demonstrate that there is a range of conditions under which it may be worth accepting the bias in the random-effects model if it is associated with a sufficient gain in efficiency, leading to estimates that are closer, on average, to the true value in any particular sample. The most common objection to the use of random effects—the violation of a "critical" modeling assumption: that the regressor and the unit effects are uncorrelated—turns out to be an insufficient justification to prefer fixed over random effects. This condition will hold only under exceptional circumstances, and our simulations demonstrate that even in the presence of rather extreme violations of that assumption, the random-effects estimator can still be preferable to (or at least no worse than) the fixed-effects estimator.

We offer a series of rules of thumb upon which researchers may rely when choosing between a fixed- or random-effects approach. When variation in the independent variable is primarily *within* units—that is, the units are relatively similar to one another on average—the choice of random versus fixed effects only matters at extremely high levels of correlation between the independent variable and the unit effects. Under these conditions, the appropriate model should be guided by the researcher's goals. When the independent variable exhibits only minimal within-unit variation, or is sluggish, there is a more nuanced set of considerations. In any particular dataset, the random-effects model will tend to produce superior estimates of β when there are few units or observations per unit, and when the correlation between the independent variable and unit effects is relatively low. Otherwise, the fixed-effects model may be preferable, as the random-effects model does not induce sufficiently high variance reduction to offset its increase in bias.

There are important limitations to the guidance offered here. We have not considered the problem of when the researcher hypothesizes that the effect of x on y varies across units, in which case one would need to employ interactive terms (fixed-effects approach) or a random coefficients model (random-effects approach). In addition, we have only considered the linear-regression model; we have not contemplated generalized linear models for dependent variables that may be dichotomous, ordinal or categorical. Finally, we have only compared two possible specifications—albeit the most commonly used ones in applied social science research. Other modeling choices abound. Nevertheless, the approach we have outlined to evaluate the performance of random- and fixed-effects models under possible violations of the model assumptions can be easily extended to test alternative models on a case-by-case basis.

REFERENCES

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics*. Princeton, NJ: Princeton University Press.

Arceneaux, Kevin, and David W. Nickerson. 2009. 'Modeling Certainty with Clustered Data: A Comparison of Methods'. *Political Analysis* 17(2):177–90.

Bates, Douglas, Martin Maechler, and Ben Bolker. 2011. *lme4: Linear Mixed-effects Models Wsing S4 Classes*. R package version 0.999375-38. Available from http://cran.R-project.org/package=lme4

https://doi.org/10.1017/psrm.2014.32 Published online by Cambridge University Press

Beck, Nathaniel, and Jonathan N. Katz. 1995. 'What To Do (and not To Do) with Time-Series Cross-Section Data'. *American Political Science Review* 89(3):634–47.

Frees, Edward W. 2004. *Longitudinal and Panel Data: Analysis and Applications in the Social Sciences*. New York: Cambridge University Press.

Gelman, Andrew. 2005. 'Analysis of Variance—Why It is More Important than Ever'. *The Annals of Statistics* 33(1):1–53.

Gelman, Andrew, and Jennifer Hill. 2007. *Data Analysis Using Regression and Multi-level/Hierarchical Models*. Cambridge: Cambridge University Press.

Greene, William H. 2012. *Econometric Analysis, Seventh Edition*. Upper Saddle River, NJ: Prentice Hall.

Hausman, Jerry A. 1978. 'Specification Tests in Econometrics'. *Econometrica* 46:1251–1271.

Kennedy, Peter E. 2003. *A Guide to Econometrics*. Cambridge, MA: MIT Press.

Kreft, Ita G.G., and Jan de DeLeeuw. 1998. *Introducing Multilevel Modeling*. London: Sage.

Plümper, Thomas, and Vera E. Troeger. 2007. 'Efficient Estimation of Time-Invariant and Rarely Changing Variables in Finite Sample Panel Analyses with Unit Fixed Effects'. *Political Analysis* 15(2):124–39.

——. 2011. 'Fixed-Effects Vector Decomposition: Properties, Reliability, and Instruments'. *Political Analysis* 19(2):147–64.

R Development Core Team. 2014. R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0. http://www.R-project.org

Robinson, G.K.. 1998. 'Variance Components'. In *Encyclopedia of Biostatistics*, edited by Peter Armitage and Theodore Colton, Vol. 6, 4713–19. Wiley.

Wilson, Sven E., and Daniel M. Butler. 2007. 'A Lot More to Do: The Sensitivity of Time-Series Cross-Section Analyses to Simple Alternative Specifications'. *Political Analysis* 15(2):101–23.

Wooldridge, Jeffrey M.. 2010. *Econometric Analysis of Cross Section and Panel Data*. 2nd edition. Cambridge, MA: MIT Press.

https://doi.org/10.1017/psrm.2014.32 Published online by Cambridge University Press