Political Analysis (2011) 19:119–122
doi:10.1093/pan/mpr010

Of Fixed-Effects and Time-Invariant Variables

Nathaniel Beck
Department of Politics, New York University, New York, NY 10012
e-mail: nathaniel.beck@nyu.edu

What follows is a longish controversy (two critiques, a reply and two rejoinders) over the quality of the estimates and associated SEs provided by Plümper and Troeger's (2007) ''fixed-effect vector decomposition'' (FEVD) procedure; Plümper and Troeger (PT) will refer to that article and not any persons. My role is to lay out some issues that separate the authors rather than to adjudicate between them. As with many controversies, a bit of heat is generated along with some light. Readers care a bit less than the authors about what was said when, but they do care a lot about what appropriate method to use when a panel data model has both unit-specific intercepts[1] and variables that are invariant over a unit. Thus, I also take it upon myself to discuss some things that I gleaned from this controversy; this discussion has a bit less heat than what follows, but of course readers should judge the evidence for themselves.

Why should anyone care about this controversy? Panel data always suffer from some potential omitted variable bias, bias due to a lack of measured data on important attributes of the individual or country that do not change over time. Thus, in labor economics, we usually cannot measure ability, but we know it is an important determinant of wages and it is likely to be correlated with other important variables (education), thus leading to omitted variable bias. Similar problems occur in comparative political economy, where institutional features of a country may be important but difficult to measure. The standard solution is to adjoint unit-specific intercepts to the specification.

At the same time, there are variables we care about that do not vary over time (time-invariant variables [TIV's]). In labor economics, these would be individual attributes such as gender, race, and education; in comparative political economy, these are often institutional variables related to the electoral system or issues of central bank independence.[2]

Scholars well beyond the authors in this controversy care a lot about the issues raised. The first econometric solution to the problem, instrumental variables, was proposed three decades ago by Hausman and Taylor (1981). HT, as I shall denote it, has almost 1500 Google Scholar citations; PT has about 250 such citations. Many scholars using panel data not only care about the effects of TIV's but also must rely on unit-specific intercepts to deal with unobserved heterogeneity.

Users of FEVD want to know whether it performs appropriately (both in terms of coefficient estimates and SEs) when applied to data that appear to fit its underlying assumptions. They also want to know if there is something new and different about the method. Both these issues are discussed analytically by Greene. As we shall see, the short answers are that FEVD performs well when its assumptions are met and that there is little new in this method. Researchers also want to know whether the assumptions behind FEVD are plausible, and how FEVD compares to other methods that make different assumptions. This is the basis of the very lucid Breusch, Ward, Nguyen, and Kompas contribution, and to my mind, this is the most interesting issue in the controversy.

Contributors to this controversy have similar issues but are in a somewhat different situation. Users care about the current version of *xtfevd*, the *Stata* implementation of FEVD (as of this writing 4.0 beta available at http://www.polsci.org/pluemper/ssc.html) while the contributors worked off whatever version of *xtfevd*

[1]These are usually called fixed effects. I want to use a different term since terms such as fixed effects and random effects are also used as ways of estimating panel data models.

[2]In comparative political economy, some institutional variables change only slightly and in only a few units. Although such variables do not technically pose a challenge to standard methods, they do cause serious problems in practice (see Green, Kim, and Yoon [2001] and the comments on it for a discussion of these issues). One nice feature of FEVD is that it is a *potential* solution to this problem, though one that involves trading off bias for lower variance. There is some interesting discussion of these tradeoffs in Breusch, Ward, Nguyen, and Kompas. Since my own view is that this is actually the interesting issue raised by PT, I hope that more attention will be paid to this issue in the future. It is clearly not enough to say that ordinary least squares (OLS) retains its optimal properties with the inclusion of almost time-invariant variables.

© The Author 2011. Published by Oxford University Press on behalf of the Society for Political Methodology.
All rights reserved. For Permissions, please email: journals.permissions@oup.com

https://doi.org/10.1093/pan/mpr010 Published online by Cambridge University Press

Table 1 Comparison of xtfevd with standard estimators (see the discussion surrounding Greene's Table 2 for definitions of the variables and other details)

	OLS/FE		*OLS FE's on TIV's*		*xtfevd*	
Variable	*Estimate*	*SE*	*Estimate*	*SE*	*Estimate*	*SE*
Constant	4.75	0.04	−1.84	0.19	2.91	0.200
EXP	0.0966	0.0012			0.0966	0.0012
WKS	0.0011	0.0006			0.0011	0.0006
OCC	−0.0249	0.0139			−0.0249	0.0139
IND	0.0208	0.0156			0.0208	0.0156
SOUTH	−0.0032	0.0346			−0.0032	0.0346
SMSA	−0.0438	0.0196			−0.0438	0.0196
MS	−0.0303	0.0190			−0.0303	0.0190
UNION	0.0342	0.0150			0.0342	0.0150
FEM			−0.126	0.128	−0.126	0.129
ED			0.146	0.014	0.146	0.014
BLK			−0.279	0.158	−0.279	0.158
Observations	4165		595		4165	

that was available when they did their analyses (given publication delays, a minimum of a year ago). As editor of the special issue in which published PT, I note with regret that I allowed PT to be published without explicit discussion of the estimation of the variance-covariance matrix. As a consequence, analysts, including the contributors here, were forced to figure out how FEVD estimated SEs by looking at whatever version of *xtfevd* was extant at the time. Code such at this is often written by one or two people, and it is often the case (as here) that busy political scientists who are not professional coders are not as careful about Revision Control Systems as are professional coders. As a consequence, it is very hard to know exactly which version of *xtfevd* the contributors relied on or how the code has changed in the last year (it has clearly changed from the 2.0 beta version which was available for a long time). In short, the contributors cannot be faulted for examining the extant version of a commonly used program, but a number of the issues raised in all the contributions are mostly irrelevant for the applied researcher.

So, and with no injustice meant to any contributor, I ran the current (4.0 beta) version of *xtfevd* (using the latest *Stata* 11.1) on the original data set used in Greene's contribution (see his Table 2 and the surrounding text for definitions of the variables and such). In Table 1, I report the results of a standard unit-specific intercept (fixed effects) estimation (using *Stata's xtreg*, *fe*), a regression of the estimated unit-specific estimates on the TIV's and *xtfevd* on the original data (so without the additional simulation step used by Greene which is irrelevant for my purposes).

The results are very clear; the coefficients and SEs of the time-varying coefficients are identical (not close, but identical) to those from the fixed-effects regression; the coefficients on the TIV's are identical to those obtained in the regression of the unit-specific intercepts on the TIV's and the SEs for those latter coefficients are similar to those obtained in the regression of the fixed effects on the TIV's. The only difference between the results reported here and the Greene's results relate to the SEs of all the coefficients. Given that *xtfevd* now reports SEs for the coefficients of the time-varying variables that are identical to those from a standard fixed-effects regression, obviously the current version of *xtfevd* has corrected any previous problems related to this part of the estimation.

The only issue left is the accuracy of the SEs for the coefficients of the TIV's. I ran Monte Carlo simulations of *xtfevd* with the data simulated exactly as in PT.[3] The average SE of the coefficient of the TIV was within a few percent of the standard deviation of the thousand estimated coefficients. These simulations (data not shown here) indicate that the SEs of the coefficients of the TIV estimated by *xtfevd*

[3]Thus, one time-varying covariate, x, one TIV, z, and a separate unit effect were drawn (making varying assumptions about their correlation), the sum of the TIV and unit effects were then regressed on the TIV with a residual, r, calculated, and then the dependent variable was generated as the sum of the x, TIV, r, and an SE term uncorrelated with anything else. One thousand runs were analyzed.

https://doi.org/10.1093/pan/mpr010 Published online by Cambridge University Press

are correct, that is, they accurately reflect the variation of those coefficients over the runs of the simulations *when the assumptions of PT coincide with the data generating process*.

From the applied researcher's perspective, *xtfevd* does what it claims to do and does it correctly if the data meet the assumptions implied by the method. However, as Greene clearly shows, it produces nothing different from the fixed-effects regression and the regression of those effects on the TIV's, both in terms of coefficient estimates and SEs.

Readers may be surprised that the *xtfevd* SEs for the TIV's are almost the same as those from the regression of the fixed effects on the TIV's, even though one is based on *NT* observations and the other on only *N* observations. Greene's algebra shows why this must be the case, but it is perhaps easier to understand what is going on by comparing a regression of $y_{i,t}$ on TIV's using *NT* observations and the ''between'' regression that regresses the unit means of the $y_{i,t}$ ($\bar{y}_i$) on the TIV's using only *N* observations. The key is to note that the error processes in the two regressions differ. The first regression estimates $\boldsymbol{\gamma}$ in

$$y_{i,t} = z_i\gamma + \varepsilon_{i,t}, \tag{1}$$

where the $\varepsilon_{i,t}$ are independent and identically distributed (''iid'') errors with variance σ^2. The between regression estimates $\boldsymbol{\gamma}$ in

$$\bar{y}_i = z_i\gamma + \bar{\varepsilon}_i. \tag{2}$$

Although obviously the $Z'Z$ matrix from the first regression is T times larger than the corresponding matrix from the second (between regression), the variance of the errors in the first regression is T times larger than in the second (since the latter error term is an average of T iid errors). Thus, the SEs in the regression of the fixed effects on the $\mathbf{z}_i$ are essentially the same as those obtained by *xtfevd*, other than for a small degrees of freedom correction (the number of time-varying covariates). Thus, if the data were generated as in PT, the coefficient estimates and SEs obtained by *xtfevd* are correct, but there is no magic here, they are the same as we would obtain by extant methods. There is also obviously no efficiency gain here (and I take some blame for allowing ''efficient'' to be the lead word in the title of PT). To my mind, efficiency is a bit of a red herring in this controversy since efficiency is only useful for comparing alternative estimators. Since the alternative estimators studied rely on different assumptions, it is impossible to compare their efficiency in general.

With this done, we can turn to the more interesting issue of whether the PT assumptions are sensible in applied work. The specification for panel data with both time-varying and time-invariant variables, unit-specific intercepts, and an independent error process is

$$y_{i,t} = \mathbf{x}_{i,t}\beta + \mathbf{z}_i\gamma + I_i\kappa_i + \varepsilon_{i,t} \tag{3}$$

where I_i is an indicator variable for unit i. As the contributors note, γ and κ are not identified, which of course is the problem. Greene proposes that we drop any attempt to estimate $\boldsymbol{\gamma}$, which may or not be ideal depending on the interests of the researcher.

PT propose instead to estimate

$$y_{i,t} = \mathbf{x}_{i,t}\beta + \mathbf{z}_i\gamma + \eta\hat{h}_i + \varepsilon_{i,t}, \tag{4}$$

where $\hat{h}_i$ is the residual from the regression of the estimated unit-specific intercepts in a regression of $y_{i,t}$ on $\mathbf{x}_{i,t}$ and, as contributors note, $\eta = 1$. Of critical importance is that $\hat{h}_i$ is a residual, so it is by construction orthogonal to $\mathbf{z}_i$ and it is not a stochastic error. The coefficients and SEs of equation (4) are what is correctly estimated by *xtfevd*.

The critical assumption is that the measured TIV's ($\mathbf{z}_i$) are given all the shared explanatory power of both themselves and the unit-specific intercepts, as in any partial regression setup. This is a very strong assumption, and researchers need to think very hard about whether it applies to their research context. In the political economy context, the assumption may be plausible; here the TIV's are usually institutional variables that we might plausibly believe are the principal reason why countries differ (as compared to other unmeasured variables that make up the remaining country specific intercept). It *may* be reasonable to

https://doi.org/10.1093/pan/mpr010 Published online by Cambridge University Press

assume that the institutional variables account for the maximal portion of the country-specific intercepts. In labor economics, the assumption *may* be less plausible. Unmeasured idiosyncratic factors may well be more important than sociological characteristics in the determination of wages. As with any identification issue, the resolution is theoretical and not empirical.

Breusch, Ward, Nguyen, and Kompas discuss other assumptions that are plausible, based on the HT instrumental variable approach. HT assumes that the $x_{i,t}$ can be parsed into two sets, one of which is independent of the unit effects. It may be the case that the unit means of these latter variables are correlated with the TIV's that are correlated with the theoretical (unobserved) unit-specific intercepts. If there are such $x_{i,t}$, their unit means can serve as instruments for these TIV's. Again, no empirical evidence can be brought to bear on this issue, but researchers can of course think about the credibility of this assumption. It may well be in labor economics, we can find time-varying variables that are related to sociological variables but unrelated to underlying measures of ability. Thinking about this issue will push readers to agree more with either Plümper and Troeger or Breusch, Ward, Nguyen, and Kompas. Without making some additional assumptions, equation (3) is clearly unidentified. Whether one might want to impose additional assumptions to allow estimation of γ of course depends on the importance of estimating $x_{i,t}$ and whether any of the needed assumptions are sufficiently plausible. Thus, only applied researchers, not econometricians, can in the end decide which of the various prescriptions offered is best, and, sadly, this decision cannot be based on empirical evidence. (Once various assumptions are made, Breusch, Ward, Nguyen, and Kompas also show that FEVD can be combined with instrumental variables to improve estimation.)

Finally, my hope, at least for comparative political economy, is that the issues raised here can become irrelevant. The unit-specific intercepts are simply properties of the various political and economic systems that we currently cannot or have not measured. In the large N panels used in labor economics, it is likely the case that the unexplained person-to-person variation is idiosyncratic and likely both not worth and impossible to model. But in comparative political economy, we ought to be able to model why nations differ and to make the unmodeled unit heterogeneity small enough to be ignorable. To my mind, this should be our goal. But until that goal is reached, researchers will have to take seriously the arguments in this controversy. The commenters have clearly laid out their positions. Some issues can be settled by using the most recent version of *xtfevd*, but the more important issues can only be settled by theoretical argument.

References

Green, Donald, Soo Yeon Kim, and David Yoon. 2001. Dirty pool. *International Organizations* 55:441–68.

Hausman, Jerry A., and William E. Taylor. 1981. Panel data and unobservable individual effects. *Econometrica* 49:1377–98.

Plümper, Thomas, and Vera Troeger. 2007. Efficient estimation of time-invariant and rarely changing variables in finite sample panel analyses with unit fixed effects. *Political Analysis* 15:124–39.

https://doi.org/10.1093/pan/mpr010 Published online by Cambridge University Press