Political Analysis (2011) 19:123–134
doi:10.1093/pan/mpq026

On the Fixed-Effects Vector Decomposition

Trevor Breusch, Michael B. Ward, Hoa Thi Minh Nguyen, and Tom Kompas
Crawford School of Economics and Government, The Australian National University, Canberra, ACT 0200, Australia
e-mail: trevor.breusch@anu.edu.au (corresponding author), michael.ward@anu.edu.au, hoa.nguyen@anu.edu.au, tom.kompas@anu.edu.au

This paper analyzes the properties of the fixed-effects vector decomposition estimator, an emerging and popular technique for estimating time-invariant variables in panel data models with group effects. This estimator was initially motivated on heuristic grounds, and advocated on the strength of favorable Monte Carlo results, but with no formal analysis. We show that the three-stage procedure of this decomposition is equivalent to a standard instrumental variables approach, for a specific set of instruments. The instrumental variables representation facilitates the present formal analysis that finds: (1) The estimator reproduces exactly classical fixed-effects estimates for time-varying variables. (2) The standard errors recommended for this estimator are too small for both time-varying and time-invariant variables. (3) The estimator is inconsistent when the time-invariant variables are endogenous. (4) The reported sampling properties in the original Monte Carlo evidence do not account for presence of a group effect. (5) The decomposition estimator has higher risk than existing shrinkage approaches, unless the endogeneity problem is known to be small or no relevant instruments exist.

1 Introduction

We analyze the properties of a recently introduced methodology for panel data, known as fixed-effects vector decomposition (FEVD), which Plümper and Troeger (2007a) developed to produce improved estimates in cases where traditional panel data techniques have difficulty. Researchers in many fields seek to exploit the advantages of such panel data. Having repeated observations across time for each group in a panel allows one, under suitable assumptions, to control for unobserved heterogeneity across the groups that might otherwise bias the estimates. Mundlak (1978) demonstrated that a generalized least squares approach to unobserved group effects, which treats them as random and potentially correlated with the regressor, gives rise to the traditional fixed-effects (FE) estimator. However, FE is a blunt instrument for controlling for correlation between observed and unobserved characteristics because it ignores any systematic average differences between groups. Thus, any potential explanatory factors that are constant longitudinally (time-invariant) will be ignored by the FE estimator. Likewise, any explanatory variables that have little within variation (i.e., slowly changing over the longitudinal dimension) will have little explanatory power and will result in imprecise coefficient estimates that have large standard errors (SEs).

Hausman and Taylor (1981) had previously shown that a better estimator than FE is available if some of the explanatory variables are known to be uncorrelated with the unobserved group effect, thus described as exogenous explanatory variables. The Hausman-Taylor (HT) estimator is an instrumental variables (IV) procedure that combines aspects of both FE and random-effects estimation. Given a sufficient number of exogenous regressors, the HT procedure allows time-invariant variables to be kept in the model. It also provides more efficient estimates than FE for the coefficients of the exogenous time-varying variables. The downside of the HT estimator resides in specifying the exogeneity status for each of the time-varying and time-invariant variables in the model. In many practical applications such detailed specification is onerous.

Plümper and Troeger introduced FEVD as an alternative that seemed to be superior to HT because it requires fewer explicit assumptions yet seemed to always have more desirable sampling properties. Like the FE estimator, and unlike HT, FEVD does not require specifying the exogeneity status of the explanatory variables. Like the HT estimator, and unlike FE, the FEVD procedure gives coefficient estimates for

Authors' note: Supplementary materials for this article are available on the *Political Analysis* Web site.

© 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/mpq026 Published online by Cambridge University Press

time-invariant (and slowly-changing) variables as well as the time-varying variables. Plümper and Troeger motivated the FEVD procedure on heuristic grounds, and advocated it on the strength of favorable results in a Monte Carlo simulation study. In particular, the simulation indicated that FEVD has superior sampling properties for time-invariant explanatory variables.

Although the FEVD procedure comes out of the empirical political science literature, it is rapidly finding application in many other areas including social research and economics. At last count, there were well over 200 references in Google Scholar to this emerging estimation methodology. Several empirical studies report SEs for FEVD-based estimates that are strikingly smaller than estimates based on traditional methods. There is, however, little formal analysis of the FEVD procedure in this literature.

The present paper is a remedy to the lack of formal analysis. We demonstrate that the FEVD coefficient estimator can be equivalently written as an IV estimator, which serves to demystify the nature of the three-stage FEVD procedure and its relationship with other estimators. As one immediate benefit, the IV representation allows us to draw on a standard toolkit of results.

First, using the IV variance formula, we show that the FEVD SEs for coefficients of both the time-varying and time-invariant variables are uniformly *too small*. In the case of the latter variables, the discrepancy in the FEVD SEs is *unbounded* and grows with the length of the panel and with the variance of the group effects.

Second, using the moment-condition representation, we prove that the coefficients of the time-varying variables in FEVD are exactly the same as in FE. This result is apparent in many of the practical studies that list FE estimates beside FEVD estimates, but it is hardly mentioned in the existing analytical material. An immediate implication is that FEVD estimates, like FE, are *inefficient* if any of the time-varying variables are exogenous.

Third, FEVD usually produces lower variance estimates of time-varying coefficients than HT in small samples. However, it does so by including invalid instruments that produce *inconsistent* estimates. So, even with massive quantities of data those FEVD estimates will deviate from the truth.

Further developments can also be made to the estimator, to exploit the ideas in FEVD while avoiding the problems of that procedure. The advantage of FEVD will be found in smaller samples where the large sample concept of consistency does not dominate. The Monte Carlo simulation studies by Plümper and Troeger (2007a) and Mitze (2009) show a trade-off between bias and efficiency in which FEVD often appears to be better than either FE or HT under quadratic loss. We present Monte Carlo evidence that a standard shrinkage approach combines the desirable small sample properties of FEVD with the desirable large sample properties of the HT estimator, so that it has superior risk to both FEVD and HT over a wide region of the parameter space.

In the next section, we introduce the notation to be used and describe the three-stage FEVD estimator. We summarize the connections between these stages in a theorem, which we prove by comparing the various moment conditions. This approach demonstrates naturally the description of the FEVD estimator as IV. Section 3 compares the correct IV variance formula with the formula implicit in the SEs of the three-stage FEVD approach. The main results are summarized in two further theorems. We also provide an empirical example to illustrate these results and some from the previous section. In Section 4, we examine the relationships between estimators in more detail, allowing the possibility of a trade-off between bias and variance to produce an estimator with lower mean-squared error (MSE). Section 5 reports some Monte Carlo evidence in the spirit of Plümper and Troeger that demonstrates the superiority of a standard shrinkage estimator. Section 6 has some overall conclusions.

2 The Model

The data are ordered so that there are N groups each of T observations. The model for a single scalar observation is

$$y_{it} = X_{it}\beta + Z_i\gamma + u_i + \mathrm{e}_{it}, \quad \text{for} \quad i = 1, \ldots, N \quad \text{and} \quad t = 1, \ldots, T. \tag{1}$$

Here, X_{it} is a $k \times 1$ vector of time-varying explanatory variables, and Z_i is a $p \times 1$ vector of time-invariant explanatory variables.[1] The parameters β, γ, the group effect u_i, and the error term e_{it} are all unobserved.

[1]The setup here describes a balanced panel with observations on every t for each i, but the ideas extend to unbalanced panels with more complicated notation. A constant can be represented in this model by including a vector of ones as part of the time-invariant elements, Z.

https://doi.org/10.1093/pan/mpq026 Published online by Cambridge University Press

Some elements of X_{it} or Z_i are correlated with the group effect u_i, in which case we call those variables *endogenous*. Otherwise, we call those variables *exogenous*. With endogenous explanatory variables standard linear regression techniques may produce estimates of the unknown parameters that are inconsistent in the sense that they do not converge to the true parameter values as the sample size grows large. One standard approach to this endogeneity problem is to use the IV technique developed by Hausman and Taylor.

2.1 *Notation*

The presentation is considerably simplified by introducing some projection matrix notation. Let

$$D = I_N \otimes \iota_T, \tag{2}$$

where I_N is an $N \times N$ identity matrix and ι_T is a $T \times 1$ vector of ones. That is, D is a matrix of dummy variables indicating group membership. For any matrix M, we use $P_M = M(M'M)^{-1}M'$ to indicate the projection matrix for M, and we use $Q_M = I - P_M$ to indicate the projection matrix for the null space of M. For example,

$$P_D = D(D'D)^{-1}D' = \frac{1}{T}\left(I_N \otimes \iota_T \iota_T^{'}\right) \tag{3}$$

is the matrix that projects a vector onto D. This particular projection produces a vector of group means. That is, $P_D y = \{\bar{y}_i\} \otimes \iota_T$,where $\bar{y}_i = \frac{1}{T}\sum_{t=1}^{T} y_{it}$. Also,

$$Q_D = I_{NT} - P_D \tag{4}$$

is the matrix that produces the within-group variation. That is, $Q_D y = \{y_{it} - \bar{y}_i\}$ is the $NT \times 1$ vector of within-group differences.

2.2 *The FEVD Estimator*

The FEVD estimator proceeds in three stages, which we detail below. To sharpen the analysis, we assume that the elements of Z are exactly time invariant (not just slowly changing), so that $P_D Z = Z$. An explicit analysis of the slowly changing case yields qualitatively similar insights.

Stage 1: Perform an FE regression of y on the time-varying X. The moment condition corresponding to an FE regression is

$$(y - Xb)' Q_D X = 0. \tag{5}$$

The unexplained component after this first step is $y - Xb$. The group average of the unexplained component is $P_D(y - Xb)$.

Stage 2: Regress the group average of the unexplained component from the first step on the time-invariant Z. The moment condition is $(P_D(y - Xb) - Zg)' Z = 0$. Using the fact that $P_D Z = Z$, this moment condition can be equivalently written as

$$(y - Xb - Zg)' Z = 0. \tag{6}$$

The group-average residuals from this regression are

$$h = P_D(y - Xb - Zg). \tag{7}$$

Stage 3: Regress y on X, Z, and h. The coefficients from this step are the final FEVD estimates. The moment conditions are

$$(y - X\beta - Z\gamma - h\delta)' [X, Z, h] = 0. \tag{8}$$

Theorem 1. The solution for β is b from Stage 1; the solution for γ is g from Stage 2; and the solution for δ is one.

https://doi.org/10.1093/pan/mpq026 Published online by Cambridge University Press

Proof. We need to verify that the moment conditions (8) are satisfied at $\beta = b$, $\gamma = g$, and $\delta = 1$. This requires that

$$(y - Xb - Zg - h)'[X, Z, h] = 0. \tag{9}$$

Substituting in the definition of h from equation (7) and gathering terms, this simplifies to

$$(y - Xb - Zg)'Q_D[X, Z, h] = 0. \tag{10}$$

Using the fact that $Q_D Z = 0$, this further simplifies to

$$(y - Xb)'Q_D[X, Z, h] = 0. \tag{11}$$

The first set of equalities in equation (11) must be satisfied since it is identical to the moment condition (5) that defines b. The second set of equalities must be satisfied since $Q_D Z = 0$. Similarly, the third set of equalities must be satisfied since $Q_D h = 0$, which follows from the definition of h in equation (7) and the fact that $Q_D P_D = 0$. □

2.2 *IV Representation*

Using Theorem 1, we can show that the FEVD estimator can also be expressed as an IV estimator for a particular set of instruments. The major benefit of using the IV representation is that one can draw on a standard toolkit of results. Theorem 1 shows that the FEVD estimates of β are identical to the standard FE estimator b from Stage 1. This estimator is defined by the moment condition (5). Theorem 1 also shows that the FEVD estimates of γ are equivalent to the estimator of g from Stage 2. This estimator is defined by the moment condition (6). Combining both moment conditions, and using the fact that $Q_D Z = 0$, the full moment conditions for the FEVD estimator are

$$(y - X\beta - Z\gamma)'[Q_D X, Z] = 0. \tag{12}$$

In other words, the FEVD estimator is equivalent to an IV estimator using the instruments $Q_D X$ and Z.

3 Variance Formulae

Using standard results for IV estimators, the asymptotically correct sampling variance of the FEVD procedure is

$$V_{\mathrm{IV}}(\beta, \gamma) = (H'W)^{-1} H' \Omega H (W'H)^{-1}, \quad \text{for} \quad H = [Q_D X, Z] \quad \text{and} \quad W = [X, Z]. \tag{13}$$

Here, H is the matrix of instruments and W is the matrix of explanatory variables. Ω is the covariance of the residual, $u_i + \mathrm{e}_{it}$, which can be expressed as

$$\Omega = \sigma_{\mathrm{e}}^2 I_{NT} + \sigma_u^2 I_N \otimes \iota_T \iota_T' = \sigma_{\mathrm{e}}^2 Q_D + \left(\sigma_{\mathrm{e}}^2 + T\sigma_u^2\right) P_D. \tag{14}$$

Using straightforward algebraic manipulation of equation (13), we will later separately expand out the variances of β and of γ for more detailed inspection.

We now compare the correct IV variance formula with the FEVD variance formula. Plümper and Troeger state that the sampling variance of the FEVD estimator can be obtained by applying the standard ordinary least squares (OLS) formula to the Stage 3 regression. Therefore,

$$V_{\mathrm{FEVD}}(\beta, \gamma, \delta) = s^2([X, Z, h]'[X, Z, h])^{-1} = s^2 \begin{pmatrix} X'X & X'Z & X'h \\ Z'X & Z'Z & Z'h \\ h'X & h'Z & h'h \end{pmatrix}^{-1}. \tag{15}$$

https://doi.org/10.1093/pan/mpq026 Published online by Cambridge University Press

Here, $s^2 = \|y - X\beta - Z\gamma - h\|^2 \Big/ dof$, where dof is the degrees of freedom. By application of equation (7), the expression for s^2 can be simplified to

$$s^2 = \|Q_D(y - X\beta)\|^2 \Big/ dof, \tag{16}$$

which we note is the standard textbook FE estimator for σ_e^2 when $dof = NT - N - k$ (see, e.g., Wooldridge 2002, p. 271).[2]

Now consider the variance of β. The FEVD variance formula for β is the top-left block of the overall FEVD variance formula in equation (15); using the partitioned-inverse formula this submatrix can be written as

$$V_{\text{FEVD}}(\beta) = s^2 \left(X' Q_{[Z,h]} X\right)^{-1}. \tag{17}$$

By expanding out equation (13), the correct variance for β can be written as

$$V_{\text{IV}}(\beta) = \sigma_e^2 (X' Q_D X)^{-1}. \tag{18}$$

Note that this is exactly the textbook FE variance formula.

Now we note from equation (16) that s^2 is a consistent estimator of σ_e^2. However, the matrices in the FEVD formula (equation 17) and the correct formula (equation 18) differ. The FEVD variance formula for β must therefore be incorrect, and we can show the direction of the error.

Theorem 2. The FEVD variance formula for coefficients on time-varying variables is too small.

Proof. Now $P_D[Z, h] = [Z, h]$, so that $P_D P_{[Z,h]} = P_{[Z,h]}$. Such a relationship between projection matrices implies that $P_D - P_{[Z,h]}$ is positive semi-definite (in matrix shorthand, $P_D \geqslant P_{[Z,h]}$). So, $Q_D \leqslant Q_{[Z,h]}$. That $(X' Q_{[Z,h]} X)^{-1} \leqslant (X' Q_D X)^{-1}$ follows immediately. This inequality will almost always be strict because the $p+1$ variables $[Z, h]$ cannot span the whole of the N-dimensional space of group operator D, and the X's have arbitrary within-group variation. □

The FEVD formula for the variance of β is *biased* in that it systematically *understates* the true sampling variance of the estimator. The essential inequality does not disappear as N gets larger, so the formula is also *inconsistent*. The usual reported SEs will be *too small*.

Now, consider the variance of γ. The FEVD variance formula for γ is the middle block of the overall FEVD variance formula in equation (15). Using an alternative representation of the partitioned inverse, this submatrix can be written as

$$V_{\text{FEVD}}(\gamma) = s^2 (Z'Z)^{-1} \left(I + Z' [X, h] \left([X, h]' Q_Z [X, h] \right)^{-1} [X, h]' Z (Z'Z)^{-1} \right). \tag{19}$$

Note that $Z'h = 0$, so that in the partitioned central matrix of the second term only the submatrix corresponding to X will be selected. Then, we have the simplification of equation (19),

$$V_{\text{FEVD}}(\gamma) = s^2 (Z'Z)^{-1} + s^2 (Z'Z)^{-1} Z' X (X' Q_Z X)^{-1} X' Z (Z'Z)^{-1}. \tag{20}$$

In contrast, by expanding out equation (13), the correct variance for γ can be written as

$$V_{\text{IV}}(\gamma) = \sigma_e^2 (Z'Z)^{-1} + T\sigma_u^2 (Z'Z)^{-1} + \sigma_e^2 (Z'Z)^{-1} Z' X (X' Q_D X)^{-1} X' Z (Z'Z)^{-1}. \tag{21}$$

[2] The usual OLS formula for the SEs from the Stage 3 regression would calculate the scale term using $dof = NT - k - p - 1$, where p is the number of Z variables including the constant and the final minus one allows for the additional regressor h. This divisor would clearly produce an inconsistent estimator of σ_e^2 for large N and small T. Plümper and Troeger (2007a, 129) mention briefly an adjustment to the degrees of freedom and, although they do not give an explicit formula, their software employs the divisor $dof = NT - N - k - p + 1$ (Plümper and Troeger 2007b). This adjustment would yield a consistent estimate of σ_e^2, but it is nonstandard and slightly biased. To sharpen the subsequent analysis, we use the standard unbiased estimator of σ_e^2, in which $dof = NT - N - k$.

https://doi.org/10.1093/pan/mpq026 Published online by Cambridge University Press

Again, s^2 is a consistent estimator of σ_e^2, so the first term in equations (20) and in (21) is essentially the same. However, the expressions are otherwise different, so the FEVD variance formula for γ must also be incorrect. Again, we can show the direction of the error.

Theorem 3. The FEVD variance formula for time-invariant variables is too small.

Proof. As shown in the proof of Theorem 2, $(X'Q_DX)^{-1} \geqslant \left(X'Q_{[Z,h]}X\right)^{-1}$ with almost certain strict inequality, so the last term in the FEVD variance formula (equation 20) understates the corresponding term in the correct variance expression (equation 21). The only exception would be the unlikely event that X and Z are exactly orthogonal, causing those terms to vanish. But even then, the FEVD variance formula will be an understatement because it omits the term $T\sigma_u^2(Z'Z)^{-1}$, which must be positive definite whenever there are random group effects. □

In general, the FEVD variance formula for γ is systematically *biased* and *inconsistent*. The usual reported SEs will be *too small*. The extent of the downward bias is unbounded. The correct variance expression includes a term that is directly proportional to the number of observations per group T and to the variance of the group effects σ_u^2. In contrast, the FEVD variance formula, and hence the SEs, are unaffected by these parameters. By increasing either or both these parameters, with everything else held constant, the extent of the downward bias in the FEVD variance formula becomes arbitrarily large.

3.1 *Empirical Example*

Reported results from the applied empirical literature align with these theoretical results. Table 1 presents our replication of Table 1 in Belke and Spies (2008) and shows results for pooled OLS (POLS), FE, FEVD, and HT. We add a column for the results from Stage 2 of FEVD and a row for the coefficient δ that arises in Stage 3 to further illustrate our theoretical results.[3] The first six variables only are shown for brevity. They include logged nominal gross domestic product (GDP) of the importing country *Lngdpim*, logged nominal GDP of the exporting country *Lngdpex*, and logged bilateral real exchange rate *Lrer*, as time-varying variables. The time-invariant variables shown are logged great circle distance in km *Ldist*, border length in km *Border*, and dummy for one or both countries being landlocked *Ll*. Results are estimated from a panel sample of $N = 420$ trading pairs for $T = 14$ years giving 5262 observations.

The coefficients for the first three (time varying) variables are the same for FE and FEVD, as shown by Theorem 1. To illustrate the second aspect of Theorem 1, the coefficients for the next three (time invariant) variables are exactly equal in Stage 2 and FEVD, and the solution for δ is one. Theorem 2 is illustrated by

Table 1 Partial replication of Belke and Spies (2008)

	(1) *POLS*	*(2)* *FE*	*(3)* *Stage 2*	*(4)* *FEVD*	*(5)* *HT*
Lngdpim	0.88	0.68		0.68	0.68
	(0.04)	(0.11)		(0.01)	(0.11)
Lngdpex	0.89	0.71		0.71	0.71
	(0.03)	(0.07)		(0.01)	(0.07)
Lrer	−0.01	0.13		0.13	0.13
	(0.01)	(0.06)		(0.01)	(0.06)
Ldist	−1.27		−1.41	−1.41	−1.75
	(0.11)		(0.04)	(0.00)	(0.16)
Border	−0.00		0.00	0.00	−0.00
	(0.00)		(0.00)	(0.00)	(0.00)
Ll	−0.16		−0.23	−0.23	−0.16
	(0.10)		(0.04)	(0.00)	(0.14)
δ				1.00	
				(0.00)	

Note. Robust SEs are in parentheses.

[3]We are grateful to those authors for supplying their data. We found some occasional small differences in reported SEs, probably due to use of ''robust'' SEs in the published results. The full table of our replication results is provided in the supplementary materials on the *Political Analysis* Web site.

https://doi.org/10.1093/pan/mpq026 Published online by Cambridge University Press

the way the first three FEVD reported SEs are systematically smaller than the FE ones, in an order of 0.01, 0.01, and 0.01, against 0.11, 0.07, and 0.06, even though the coefficients themselves are identical and the SE formula for FE is well established as being correct under the assumptions of the model.

It is a little harder to illustrate Theorem 3, which says that the FEVD SEs on the time-invariant variables are similarly understated. However, the HT estimator is just indentified in this case, which is the reason the HT coefficients and SEs for time-varying variables are exactly the same as FE. It is no surprise, then, that the coefficient estimates of three time-invariant variables (which are all exogenous) are generally similar for POLS, FEVD, and HT. As expected, the HT SEs are slightly larger but very close to those for POLS, in an order of 0.16, 0.00, and 0.14 against 0.11, 0.00, and 0.10. However, the FEVD SEs are very small, at 0.00 in every case for the precision that is shown. This is most implausible because one would not expect POLS to be generally less efficient, given the structure of this example.

Belke and Spies (2008) is the only paper to our knowledge that reports results for all methods including POLS, FE, FEVD, and HT. Several other applications report both FE and FEVD results (e.g., Krogstrup and Wälti 2008; Caporale et al. 2009; Mitze 2009). In the studies we examined, the FE t statistics were consistently smaller than those reported for FEVD time-varying variables—and often much smaller—except for few cases affected by robust SE formulae. Again, this is despite the fact that the coefficient estimators were actually identical by construction.

4 Comparison to Alternative Estimators

The FEVD estimator was introduced as an alternative to the HT estimator. By also expressing FEVD in its IV representation, we are able to develop insights into their comparative properties. Hausman and Taylor showed that the standard FE estimator is equivalent to an IV estimator with instrument set $Q_D X$. To that, they add any exogenous elements of X or of Z as further instruments.[4]

To see the relationship more clearly, decompose X and Z into exogenous and potentially endogenous sets: $X = [X_1, X_2]$ and $Z = [Z_1, Z_2]$, where the subscript 1 indicates exogenous variables and the subscript 2 indicates endogenous variables. The HT procedure is then an IV estimator that uses the instrument set $[Q_D X, X_1, Z_1]$. In contrast, the FEVD procedure is an IV estimator that uses the instrument set $[Q_D X, Z_1, Z_2]$.

The first essential difference between these estimators is that the FEVD instrument set *excludes* the exogenous time-varying variables X_1. Of course, X_1 may have no members. In that case, the HT estimator for endogenous Z is not identified, so no useful comparisons can be made.[5] However, if X_1 has known members, then a more efficient estimator than FEVD could be created by augmenting the instrument set with X_1. The second essential difference is that the FEVD instrument set *includes* the potentially endogenous time-invariant variables Z_2. If these variables are in fact correlated with the group effect, then the FEVD estimator is inconsistent.

The FEVD and HT estimators coincide exactly when there are no exogenous elements of X and no endogenous elements of Z.[6] The FEVD procedure is thus primarily of interest when some Z may in fact be endogenous. The essential question raised by Plümper and Troeger is then whether it is better to use a biased and inconsistent but lower variance estimator, or a consistent but higher variance estimator. The question of whether a weak-instruments cure is worse than the disease is a sound one, which has been considered in other contexts by a variety of authors; see for example Bound, Jaeger, and Baker (1995).

Under an MSE loss function, neither the FEVD procedure nor the HT procedure will uniformly dominate the other. MSE can be expressed as variance plus bias squared. Thus, a consistent estimator such as HT will be preferable to the FEVD for sufficiently large sample size.[7] In contrast, for a small sample with a small endogeneity problem, it might be preferable to include the time-invariant endogenous variables Z_2 as instruments, as FEVD does. A more efficient estimator of this type than FEVD would be the IV estimator that augments the set of all valid instruments with Z_2, forming the instrument set $[Q_D X, X_1, Z]$.

[4] Hausman and Taylor describe $P_D X$ as the additional instrument, but this interpretation follows Breusch et al. (1989).

[5] Ideally, one would have theoretical grounds for identifying which elements of X are exogenous. As a practical matter, one could also use an over-identification test to confirm this assumption since the FE estimator of β is consistent.

[6] More precisely, the two estimators are identical when all elements of X are treated as if endogenous and all elements of Z are treated as if exogenous, regardless of the actual endogeneity status.

[7] Of course, consistency does require that valid instruments correlated with Z_2 exist.

https://doi.org/10.1093/pan/mpq026 Published online by Cambridge University Press

One conventional approach to finding a balance would be to select between the competing estimators based on a specification test (Baltagi, Bresson, and Pirotte 2003). If the test rejects the null hypothesis of no difference between estimators, then HT would be selected. Otherwise, the efficient estimator would be selected because the evidence of endogeneity is too weak. Selection of a final estimator based on the results of a preliminary test is known as a *pretest* procedure. Inference based on the SEs of the final selected estimator alone may be misleading; however, bootstrap techniques that include the model selection step can circumvent this problem (Wong 1997).

Since the work of James and Stein (1961), statisticians have understood that shrinking (biasing) an estimator toward a low-variance target can lower the MSE. An extensive literature suggests shrinkage approaches based on using a weighted average of two estimators when one estimator is efficient and the other is consistent; see for example Sawa (1973), Feldstein (1974), Mundlak (1978), Green and Strawderman (1991), Judge and Mittelhammer (2004), or Mittelhammer and Judge (2005). We consider a shrinkage estimator that combines the consistent but inefficient HT estimator and the efficient but possible inconsistent IV estimator. For purposes of illustration, we choose a particularly simple shrinkage approach, but the literature contains many variations on the basic theme, which will have different strengths and weaknesses. If the bias, variance, and covariance of two estimators are known, it is algebraically straightforward to find the weight that minimizes the MSE of a combined estimator. In particular, suppose one estimator ϕ is unbiased. The other estimator χ is biased but has lower variance. The shrinkage estimator then has the form $\chi+w(\phi-\chi)$, where w is the weight placed on the consistent estimator. Straightforward calculus shows that optimal weight that minimizes MSE is

$$w=\frac{\mu_\chi^2+\sigma_\chi^2-\sigma_{\chi\phi}}{\mu_\chi^2+\sigma_\chi^2+\sigma_\phi^2-2\sigma_{\chi\phi}}, \tag{22}$$

where bias is indicated by μ and where variance is indicated by σ.

Of course, the exact bias and variances will usually not be known; however, practical estimates of these terms are readily available for IV estimators. Mittelhammer and Judge (2005) show that plugging in such empirical estimates produces a practical weighted-average estimator. They choose a single w to minimize the sum of MSE over all coefficients. Since we are primarily interested in the MSE of a single coefficient in this analysis, we apply the solution for w, as presented in equation (22) which is the single covariate case of equation 3.5 of Mittelhammer and Judge. We use standard empirical estimates of the variance and covariance terms from application of the basic IV formula (equation 13). The difference between the two estimators provides our estimate of the bias of the efficient estimator since HT is asymptotically unbiased. Mittelhammer and Judge provide detailed discussion on calculating bootstrap percentiles and SEs, through application of Efron's bias–corrected and accelerated bootstrap (Efron 1987). The only change needed for the present context is to account for the panel structure, which is most simply done by resampling at the group level rather than resampling single observations independently.

5 Monte Carlo Evidence

In this section, we compare the practical performance under a range of conditions of various estimators for an endogenous time-invariant Z. In addition to the FEVD and HT estimators, we consider a pretest estimator and a shrinkage estimator. The pretest estimator selects between HT and the IV estimator based on the instrument set $[Q_D X, X_1, Z]$, which treats all Z as exogenous (as FEVD does) in addition to using the HT instruments. The pretest selection is based on the 95% critical value of the Durbin-Wu-Hausman specification test for exogeneity of Z (see, e.g., Davidson and MacKinnon 1993, 237). The shrinkage estimator assigns weights for the two estimators according to a first-stage empirical estimate of the formula in equation 22.

Plümper and Troeger argue for the superiority of the FEVD procedure over the HT approach based on Monte Carlo evidence. Although our simulation design stays close to the original design where appropriate, our design differ from theirs in two fundamental respects.[8] The first difference is that in the Plümper and Troeger Monte Carlo study, the HT estimator was not actually consistent. This is because their data

[8]The authors graciously provided the original simulation code upon request.

https://doi.org/10.1093/pan/mpq026 Published online by Cambridge University Press

generating process had no correlation between X and Z. The fact that the available instruments had, by construction, zero explanatory power for the endogenous variable contrasts sharply with their characterization of the Monte Carlo results (130): "the advantages of the FEVD estimator over the Hausman-Taylor cannot be explained by the poor quality of the instruments." Plümper and Troeger note[11] that the advantage of FEVD persists in their experiments regardless of sample size. However, the asymptotic bias of an IV estimator is the same as the bias of OLS when the instruments are uncorrelated with the endogenous variable and thus irrelevant (Han and Schmidt 2001). In contrast, with a valid and relevant instrument, the bias of the IV estimator will approach zero asymptotically. We therefore consider scenarios in our simulation where the HT estimator is consistent, that is at where at least one instrument for the endogenous Z is valid and relevant.

The second difference is that our simulations account for random variation in the group effect, whereas the Plümper and Troeger code holds the effect (u) fixed across all replications. Mundlak (1978) shows there is no loss of generality in assuming the effect is random because the FE estimator and its related procedures can be described as inference conditional on the realizations of the effect in the sample. Further, the effect needs to be at least potentially random if the relationship between the effect and the regressors is to be described as *correlation*. As Mundlak shows, if the random effect is correlated with the group averages of regressors in unknown ways, then the optimal linear estimator in the random-effects *model* is in fact the FE *estimator*.

The code used by Plümper and Troeger does not simply fix the replicated effects at some sample realization, rather it uses the *Stata* command "corr2data" to fix the sample moments of the variables and the group effects exactly in every replication. The vector of effects is thereby "fixed" by making it exactly orthogonal to the exogenous variables, effectively excluding any practical influence of the group effect in the simulated data. That process does not simulate an FE model but rather one in which there is no group effect at all! By contrast, our random-effects simulation represents the situation where the analyst is uncertain of the magnitudes of the group effects.

We run a series of experiments that vary the degree of endogeneity and strength of instrument. The data generating process for our simulation is

$$y_{it} = 1 + 0.5x_1 + 2x_2 - 1.5x_3 - 2.5z_1 + 1.8z_2 + 3z_3 + u_i + \mathrm{e}_{it}. \tag{23}$$

Here, $[x_1, x_2, x_3]$ is a time-varying mean-zero orthonormal design matrix, fixed across all experiments. $[z_1, z_2]$ is a time-invariant mean-zero orthonormal design matrix, fixed across all experiments. z_3 is fixed for all replications in each experiment. z_3 has sample mean zero and variance 1 and is orthogonal to all other variables except x_1. The sample covariance of the group mean of x_1 with z_3 is set exactly to an experiment-specific level, which allows us to vary the strength of the instrument across experiments.[9] The idiosyncratic error term e is standard normal. The random effect u is drawn from a normal distribution in each replication. The expectation of u conditional on z_3 is ρz_3, where ρ works out to be the value of $\mathrm{cov}(z_3, u)$ set in the experimental design. All other variables are uncorrelated with u, and the variance of u conditional on all variables is 1.[10] The level of endogeneity is varied across experiments by changing the value of cov (z_3, u). Each experiment has 1000 replications, which vary the random components u and e. There are 30 groups (N) and 20 periods (T), as reported in Plümper and Troeger (2007a). In implementing the estimators, $[x_1, x_2, z_1, z_2]$ are treated as known exogenous, whereas $[x_3, z_3]$ are treated as potentially endogenous.

Figure 1 illustrates the simulation results for varying instrument strengths and endogeneity levels. The vertical axis in each panel is the square root of MSE of various estimators for the endogenous time-invariant variable z_3. The horizontal axis of each panel is the covariance between the random effect u and z_3. Each panel illustrates different instrument strength, as indicated by stronger instruments having higher correlation between the group-means of x_1 and the endogenous variable z_3. The four panels display the experiments for $\mathrm{corr}(\bar{x}_1, z_3) = 0.15, 0.30, 0.45,$ and 0.60, respectively.[11] Note that, within each panel,

[9] Conditional on a nonzero sample correlation of the endogenous variable and the instrument, the moments of the IV estimator exist, so the Monte Carlo MSE is well defined.

[10] The specified pattern of covariance is implemented through a Choleski decomposition approach.

[11] Because variances of $\bar{x}_1$ and z_3 are both 1, the covariance of these variables equals their correlation.

https://doi.org/10.1093/pan/mpq026 Published online by Cambridge University Press

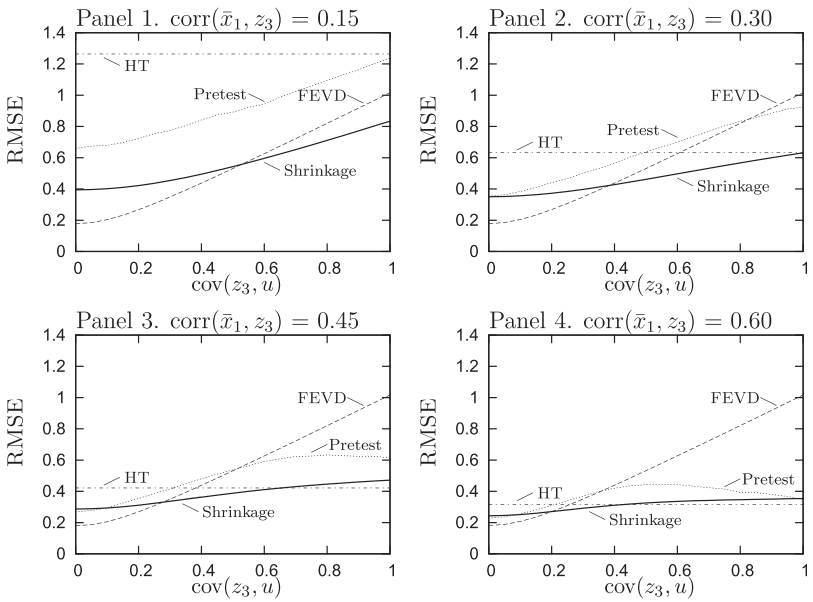


Fig. 1 Performance of the four estimators for varying instrument strengths.

the HT results are unchanging as a consequence of the experimental design. Also, across panels, the FEVD results are unchanging by design.

The most notable feature of Fig. 1 is that neither HT nor FEVD uniformly dominates the other. If reasonably strong instruments are available to implement the HT procedure, and endogeneity is an issue, HT can greatly outperform FEVD as shown in Panel 4 because the higher variance of HT is compensated by lower bias.[12] For all cases when endogeneity is absent (or is mild), FEVD will be the most efficient estimator, as shown at the far left of all panels, because FEVD exploits the true (or approximately true) restriction that z_3 is uncorrelated with u. If the investigator has strong prior reason to believe that endogeneity is not an issue, it makes sense to use that information. Indeed, with informative priors over endogeneity, using a Bayesian procedure that minimizes risk against that prior would be the ideal approach. However, usually, the investigator will be using FE, HT, or FEVD precisely because of concern that endogeneity might be a significant problem.

Rather than relying solely on prior information about the degree of endogeneity, the investigator can rely on evidence from within the data set. Both the shrinkage and the pretest estimators are in this spirit. The shrinkage estimator in particular exhibits remarkably good risk characteristics across all ranges of all four panels, and it clearly dominates the pretest approach under MSE loss. Indeed, the shrinkage estimator often has an MSE lower than *both* the HT and the FEVD estimators, and never is much worse than the better of the two. The Monte Carlo evidence suggests that a shrinkage estimator would almost certainly be the best choice in the absence of prior information that the endogeneity problem is quite small.[13] More generally, if incomplete or uncertain prior information is available, alternatives that explicitly model that information, such as traditional Bayesian techniques or recent variants such as Bayesian model averaging (Hoeting et al. 1999), will likely be the best approach.

[12]The discussion here focuses on the small sample properties. When N is very large, HT will always outperform FEVD if there is endogeneity and valid and relevant instruments exist. For a modest example of relative estimator performance as N grows, see the Appendix, where the case of $N = 300$ and $T = 2$ is illustrated.

[13]Although our focus is on estimator performance, it is worth noting that the Monte Carlo results do confirm that the asymptotic variance formula in equation (13) provides unbiased estimates of the HT and FEVD sampling variance, when σ_e^2 and σ_u^2 are calculated with appropriate degrees of freedom corrections for small sample. Further, the bootstrap quantiles for the shrinkage estimator are reasonably accurate, confirming the results of Mittelhammer and Judge (2005).

https://doi.org/10.1093/pan/mpq026 Published online by Cambridge University Press

6 Conclusions

The FEVD estimator of Plümper and Troeger (2007a) offers the analyst of panel data a way to include time-invariant (and slowly changing) variables in the presence of group effects that are possibly correlated with the explanatory variables. Thus, it appears superior to the existing leading approaches of FE (which omits the time-invariant variables) and Hausman-Taylor (which requires specifying the exogeneity status of each explanatory variable). The motivation of Plümper and Troeger for the procedure was mostly heuristic and their evidence came from Monte Carlo experiments showing that FEVD often displays better MSE properties than both FE and HT. The procedure can be implemented in three easy stages or even more conveniently in the *Stata* package provided by Plümper and Troeger (2007b). This procedure has proved popular with panel data analysts.

Our analytical results and revised Monte Carlo experiments challenge the value of FEVD. Is it still a useful tool?

We find that the coefficients of all the time-varying variables after the three stages of FEVD are exactly the same as FE in the first stage. This fact is sometimes seen in the empirical applications but rarely commented upon with any clarity. Obviously, there is no gain in using FEVD over the simpler FE if these coefficients are the objects of interest. Further, if something is known about the exogeneity of explanatory variables, then these estimates are inefficient because they ignore the extra information. What is worse, unlike the simple first-stage FE, the SEs from FEVD are too small—sometimes very much too small, judging from our empirical example and other published applications. In this case, FEVD is a definite step backward.

The main attraction of FEVD is its ability to estimate coefficients of time-invariant explanatory variables. But, again the third stage is questionable. The same coefficient estimates are given in the second stage, which is a simple regression of the group-averaged residuals from FE on the time-invariant variables. The purported value of the third stage is to correct the SEs, but this reasoning is now known to be false. Indeed, there will be cases where the second-stage SEs—even though they are known to be wrong—will be more accurate than those from the third stage. The example we have provided in Section 3 shows this possibility.

So if FEVD is the label to describe the three-stage procedure, it cannot be recommended for making inferences about any of the coefficients. The coefficient estimator, however, also represents a particular choice of instruments in standard IV. Dropping the three-stage methodology and reverting to an explicit IV approach would allow correct SEs to be obtained in the cases where the estimator is consistent. However, since all the time-invariant variables are used as instruments, the FEVD estimator will be inconsistent if

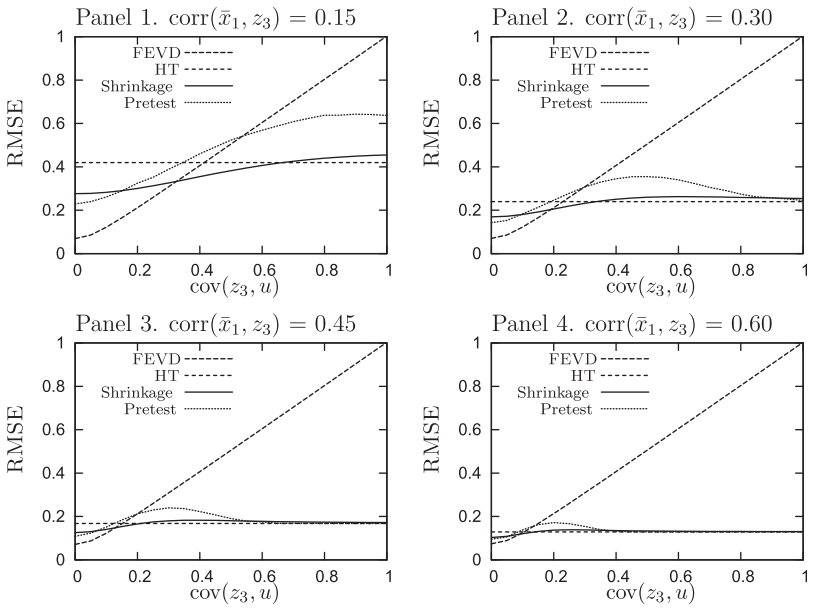


Fig. 2 Relative estimator performance when $N = 300$ and $T = 2$.

https://doi.org/10.1093/pan/mpq026 Published online by Cambridge University Press

any of these are endogenous. The value of this estimator relative to others then depends on the trade-off between inconsistency and inefficiency.

When the objective is reduced MSE, the literature is replete with other methods such as shrinkage estimators known to have good properties. We have provided one such estimator that clearly dominates the FEVD estimator over much of the parameter space and also limits the risk in regions where the FEVD risk is unbounded. We demonstrate the feasibility of such an estimator with SEs found empirically by bootstrapping. In undertaking these investigations, we have also uncovered an explanation for the misleading evidence favoring FEVD that was suggested in the previous Monte Carlo studies.

Appendix

Monte Carlo Results for Large N and Small T

In applications such as labor market studies, the number of groups can be quite large, often in the tens of thousands, since there may be a distinct group for each individual in the study. Figure 2 presents a modest example of the relative behavior of the four estimators as the number of groups grows larger. Each panel in Fig. 2 illustrates the same parameter settings as the corresponding panel in Fig. 1. The simulation code for the figures is identical, except for the N and T settings. Although the overall number of observations is the same in the two figures, the larger number of groups provides more information about the time-invariant variables. Panel 4 illustrates that the relative performance of FEVD can be quite poor for reasonable parameter settings and a modest number of observations.

References

Baltagi, Badi H., Georges Bresson, and Alain Pirotte. 2003. Fixed effects, random effects or Hausman–Taylor? A pretest estimator. *Economics Letters* 79:361–9.

Belke, Ansgar, and Julia Spies. 2008. Enlarging the EMU to the east: What effects on trade? *Empirica* 35:369–89.

Bound, John, David A. Jaeger, and Regina M. Baker. 1995. Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *Journal of the American Statistical Association* 90: 443–50.

Breusch, Trevor S., Grayham E. Mizon, and Peter Schmidt. 1989. Efficient estimation using panel data. *Econometrica* 57:695–700.

Caporale, Guglielmo M., Christophe Rault, Robert Sova, and Anamaria Sova. 2009. On the bilateral trade effects of free trade agreements between the EU-15 and the CEEC-4 countries. *Review of World Economics* 145:189–206.

Davidson, Russell, and James G. MacKinnon. 1993. *Estimation and inference in econometrics*. New York: Oxford University Press.

Efron, Bradley. 1987. Better bootstrap confidence intervals. *Journal of the American Statistical Association* 82:171–85.

Feldstein, Martin. 1974. Errors in variables: A consistent estimator with smaller MSE in finite samples. *Journal of the American Statistical Association* 69:990–6.

Green, Edwin J., and William E. Strawderman. 1991. A James-Stein type estimator for combining unbiased and possibly biased estimators. *Journal of the American Statistical Association* 86:1001–6.

Han, Chirok, and Peter Schmidt. 2001. The asymptotic distribution of the instrumental variable estimators when the instruments are not correlated with the regressors. *Economics Letters* 74:61–6.

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

Hoeting, Jennifer A., David Madigan, Adrian E. Raftery, and Chris T. Volinsky. 1999. Bayesian model averaging: A tutorial. *Statistical Science* 14:382–401.

James, W., and Charles Stein. 1961. Estimation with quadratic loss. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability.* Vol. 1, ed. J. Neyman, 361–79. Berkeley, CA; University of California Press.

Judge, George G., and Ron C. Mittelhammer. 2004. A semiparametric basis for combining estimation problems under quadratic loss. *Journal of the American Statistical Association* 99:479–87.

Krogstrup, Signe, and Sébastien Wälti. 2008. Do fiscal rules cause budgetary outcomes? *Public Choice* 136:123–38.

Mittelhammer, Ron C., and George G. Judge. 2005. Combining estimators to improve structural model estimation and inference under quadratic loss. *Journal of Econometrics* 128:1–29.

Mitze, Timo. 2009. *Endogeneity in panel data models with time-varying and time-fixed regressors: to IV or not IV?* Ruhr Economic Paper No. 83.

Mundlak, Yair. 1978. On the pooling of time series and cross section data. *Econometrica* 46:69–85.

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

———. 2007b. xtfevd.ado version 2.00 beta. http://www.polsci.org/pluemper/xtfevd.ado.

Sawa, Takamitsu. 1973. The mean square error of a combined estimator and numerical comparison with the TSLS estimator. *Journal of Econometrics* 1:115–32.

Wong, Ka-fu. 1997. Effects on inference of pretesting the exogeneity of a regressor. *Economics Letters* 56:267–71.

Wooldridge, Jeffrey M. 2002. *Econometric analysis of cross section and panel data*. Cambridge, MA: MIT Press.

https://doi.org/10.1093/pan/mpq026 Published online by Cambridge University Press