A Portmanteau Test for Serially Correlated Errors in Fixed Effects Models

Atsushi Inoue
North Carolina State University

Gary Solon
University of Michigan

May 2004

Abstract We propose a portmanteau test for serial correlation of the error term in a fixed effects model. The test is derived as a conditional Lagrange multiplier test, but it also has a straightforward Wald test interpretation. In Monte Carlo experiments, the test displays good size and power properties.

The authors are grateful for comments from David Drukker, Christian Hansen, Serena Ng, and Jeffrey Wooldridge.

A Portmanteau Test for Serially Correlated Errors in Fixed Effects Models

I. Introduction

Empirical researchers frequently use longitudinal data to estimate fixed effects models of the form

$$y_{it} = c_i + \beta' x_{it} + \varepsilon_{it} \quad (1)$$

where $i = 1, 2, ..., N$ indexes cross-sectional units (such as individuals, firms, states in a country, or countries) and $t = 1, 2, ..., T$ indexes time periods. In analyses of longitudinal microdata, T typically is fairly small. The K explanatory variables in the x_{it} vector are commonly assumed to be strictly exogenous, while the "fixed effect" c_i is a time-invariant unit-specific effect that may be correlated with elements of x_{it}, but not with the error term ε_{it}. If ε_{it} is i.i.d. $N(0, \sigma^2)$, the efficient estimator of β is the "fixed effects estimator" that applies ordinary least squares (OLS) to the mean-differenced regression of $y_{it} - \bar{y}_i$ on $x_{it} - \bar{x}_i$ where $\bar{y}_i = \sum_{t=1}^{T} y_{it}/T$ and $\bar{x}_i = \sum_{t=1}^{T} x_{it}/T$. An alternative way of computing the same estimator is to apply OLS to the regression of y_{it} on x_{it} and a vector of unit-specific dummy variables.

Often, however, the error term is not i.i.d., but instead is serially correlated. This occurs in longitudinal data for the same reasons it frequently occurs in single time series - mainly because of left-out variables that evolve gradually over time. Quite strangely, researchers who learned in introductory econometrics always to check for serial correlation when estimating time series regressions completely forget this lesson when estimating fixed effects regressions with multiple time series. When Kezdi (2002) scoured three recent years' issues of the *American Economic Review*, *Journal of Political Economy*, and *Quarterly Journal of Economics*, he found that, of the 42 articles that estimated fixed effects models, 36 paid no attention whatsoever to the serial correlation issue. Similarly, Bertrand, Duflo, and Mullainathan (2004), who focused on the "differences in differences" special case in which the explanatory variable of main interest is a binary policy variable, located 65 articles that appeared in the same journals plus three applied field journals over the 1990-2000 period, and they found that 60 of those studies totally ignored serial correlation. The trouble with this state of affairs is that ignoring serial correlation in the fixed effects context has the same poor consequences that it has with a single time series: it leads to inconsistent estimation of standard errors and hence to inappropriate hypothesis tests, and it also leads to inefficient estimation of the regression coefficients.

We conjecture that practitioners' inattention to serial correlation in fixed effects models is due to a lack of simple diagnostics. Therefore, in the next section, we present a straightforward portmanteau statistic for testing the null hypothesis of no serial correlation against a general alternative that at least some of the autocorrelations are nonzero.[1] Our test can be applied in the fixed effects context much as the Box-Pierce statistic is with a single time series. When our test rejects the null hypothesis, as it often will, practitioners should proceed in the same three ways that they do in the time series context. First, they should consider whether the error term's serial correlation is a symptom of model misspecification.[2] Second, at a minimum, they should use a robust covariance matrix estimator to correct their estimated standard errors (Arellano, 1987; Kezdi, 2002). Third, they should consider attempting more efficient coefficient estimation through a feasible generalized least squares procedure (Kiefer, 1980; Nickell, 1980; Bhargava, Franzini, and Narendranathan, 1982; Solon, 1984b; Hansen, 2003).

II. A Portmanteau Test

We can rewrite the model in equation (1) in matrix notation as

$$y_i \;=\; c_i \ell_T + X_i\beta + \varepsilon_i \tag{2}$$

where ℓ_T is the T-dimensional column vector of ones, $y_i = [y_{i1}\ y_{i2}\ \cdots\ y_{iT}]'$, $X_i = [x_{i1}\ x_{i2}\ \cdots\ x_{iT}]'$, and $\varepsilon_i = [\varepsilon_{i1}\ \varepsilon_{i2}\ \cdots\ \varepsilon_{iT}]'$. Letting $\Sigma = E(\varepsilon_i \varepsilon_i')$, we wish to test the null hypothesis $\Sigma = \sigma^2 I_T$ against the alternative that at least some off-diagonal elements of Σ are nonzero. To devise a powerful test, we will start with a Lagrange multiplier (LM) approach under the assumption that ϵ_i is normally distributed. It will turn out, though, that the resulting test has a straightforward Wald test interpretation even when ε_i is nonnormal.

[1]Existing tests for serial correlation in fixed effects models are discussed below in Section III.

[2]For example, Solon (1984a), upon finding large positive autocorrelations at low orders and large negative ones at high orders, recognized that he needed to add state-specific time trends to his model. Note that an advantage of our test relative to existing alternatives is that it gives attention to higher-order autocorrelations, which sometimes may help with identifying specification problems.

Because of the incidental parameters c_i, one cannot construct an LM test based on the likelihood function. Thus we construct a conditional likelihood function based on a sufficient statistic for the individual specific effect c_i (see Chamberlain, 1980, for this approach to the logit model for panel data). When $\Sigma = \sigma^2 I_T$, the sufficient statistic is $\bar{y}_i$. When $\Sigma \neq \sigma^2 I_T$, however, the sufficient statistic is $(\ell_T' \Sigma^{-1} y_i)/(\ell_T' \Sigma^{-1} \ell_T)$. Then the conditional log-likelihood function is given by

$$\begin{aligned} \ln L &= -\frac{N}{2}\ln|\Sigma| - \frac{N}{2}\ln(\ell_T' \Sigma^{-1} \ell_T) \\ &\quad -\frac{1}{2}\mathrm{tr}\left[\left(\Sigma^{-1} - \frac{\Sigma^{-1}\ell_T \ell_T' \Sigma^{-1}}{\ell_T' \Sigma^{-1} \ell_T}\right)\sum_{i=1}^{N} \varepsilon_i \varepsilon_i'\right]. \end{aligned} \tag{3}$$

Because of the mean-differencing transformation, the estimated covariance matrix is singular and is of rank $T-1$. Thus we will test $(T-1)(T-2)/2$ zero-restrictions on the off-diagonal elements of the $(T-1)\times(T-1)$ matrix obtained from deleting the kth column and row of the estimated covariance matrix. Let $D_{k,T}$ denote the $T^2 \times (T-1)(T-2)/2$ Jacobian matrix of $\mathrm{vec}(\Sigma)$ with respect to the lower-diagonal elements of the $(T-1)\times(T-1)$ matrix obtained from deleting the kth column and row of the covariance matrix, where vec is the vec operator that transforms a matrix into a column vector by stacking the columns of the matrix. When $k=2$ and $T=3$, for example,

$$D_{k,T} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}.$$

The gradient of the conditional log-likelihood function with respect to the lower-diagonal elements of the $(T-1)\times(T-1)$ matrix is

$$\begin{aligned} \nabla \ln L &= \frac{1}{2}\sum_{i=1}^{N} D_{k,T}' \mathrm{vec}\left[-\Sigma^{-1} + \Sigma^{-1}\varepsilon_i\varepsilon_i'\Sigma^{-1} - \frac{\Sigma^{-1}\ell_T\ell_T'\Sigma^{-1}\varepsilon_i\varepsilon_i'\Sigma^{-1}}{\ell_T'\Sigma^{-1}\ell_T} - \frac{\Sigma^{-1}\varepsilon_i\varepsilon_i'\Sigma^{-1}\ell_T\ell_T'\Sigma^{-1}}{\ell_T'\Sigma^{-1}\ell_T}\right] \\ &\quad + \frac{1}{2}\sum_{i=1}^{N} D_{k,T}' \left[\frac{\Sigma^{-1}\ell_T \otimes \Sigma^{-1}\ell_T}{\ell_T'\Sigma^{-1}\ell_T} + \mathrm{tr}(\Sigma^{-1}\ell_T\ell_T'\Sigma^{-1}\varepsilon_i\varepsilon_i')\frac{\Sigma^{-1}\ell_T \otimes \Sigma^{-1}\ell_T}{(\ell_T'\Sigma^{-1}\ell_T)^2}\right] \end{aligned}$$

When it is evaluated under the null hypothesis, it can be written as

$$\begin{aligned}\nabla \ln L_0 &= \frac{1}{2\tilde{\sigma}^2}\sum_{i=1}^{N} D'_{k,T}\text{vec}(e_i e_i' - \tilde{\sigma}_2 M) \\ &= \frac{1}{2\tilde{\sigma}^2}\sum_{i=1}^{N} D'_{k,T}\text{vec}(e_i e_i' - \text{trace}(e_i' e_i) M) \\ &= \frac{1}{2\tilde{\sigma}^2}\sum_{i=1}^{N} D'_{k,T} M_2 \text{vec}(\varepsilon_i \varepsilon_i' - \sigma^2 I_T) \qquad (4)\end{aligned}$$

where $e_i = M\varepsilon_i$, $M = I_T - (1/T)\ell_T \ell_T'$, $\tilde{\sigma}^2 = \sum_{i=1}^{N} e_i' e_i/(N(T-1))$ and $M_2 = M \otimes M' - (\text{vec}M)(\text{vec}M)'/(T-1)$. Let

$$V = E\left[\text{vec}(\varepsilon_i \varepsilon_i' - \sigma^2 M)\text{vec}(\varepsilon_i \varepsilon_i' - \sigma^2 M)'\right].$$

Then under the null hypothesis, the distribution of the infeasible LM statistic

$$N^{-\frac{1}{2}}\sum_{i=1}^{N} \text{vec}(e_i e_i' - \tilde{\sigma}^2 M)' D_{k,T}[D'_{k,T} M_2 V M_2 D_{k,T}]^{-1} N^{-\frac{1}{2}}\sum_{i=1}^{N} D'_{k,T}\text{vec}(e_i e_i' - \tilde{\sigma}^2 M)$$

converges to a χ^2 distribution with degrees of freedom $(T-1)(T-2)/2$ as $N \to \infty$. Thus, our test will be applicable in the typical panel data setting in which T may be small, but N is large.

To make the test operational we will proceed as follows: First, write the fixed effects estimator $\hat{\beta}$ as

$$\hat{\beta} = (\sum_{i=1}^{N} X_i M X_i')^{-1}\sum_{i=1}^{N} X_i M y_i,$$

and let

$$\begin{aligned}\hat{e}_{it} &= y_{it} - \bar{y}_i - \hat{\beta}'(x_{it} - \bar{x}_i), \\ \hat{\sigma}^2 &= \frac{1}{N(T-1)}\sum_{i=1}^{N}\sum_{t=1}^{T}\hat{e}_{it}^2.\end{aligned}$$

Second, compute the feasible LM statistic:

$$LM = N^{-\frac{1}{2}}\sum_{i=1}^{N} \text{vec}(\hat{e}_i \hat{e}_i' - \hat{\sigma}_2 M)' D_{k,T}[D'_{k,T} M_2 \hat{V} M_2 D_{k,T}]^{-1} N^{-\frac{1}{2}}\sum_{i=1}^{N} D'_{k,T}\text{vec}(\hat{e}_i \hat{e}_i' - \hat{\sigma}_2 M)$$

where

$$\hat{V} = \frac{1}{N}\sum_{i=1}^{N} \text{vec}(\hat{e}_i\hat{e}_i' - \hat{\sigma}^2 M)\text{vec}(\hat{e}_i\hat{e}_i' - \hat{\sigma}^2 M)'.$$

Theorem 1. Suppose that

(a) X_i, ε_i are iid and have finite fourth moments.

(b) $E(\varepsilon_i | X_i, c_i) = 0.$

(c) $\text{rank}[E(X_i' M X_i)] = \dim(x_{it}).$

(d) $T \geq 3.$

Then under the null hypothesis that $\Sigma = \sigma^2 I_T$, the LM test statistic is asymptotically distributed (as $N \to \infty$) as $\chi^2((T-1)(T-2)/2)$.

Proof: Since the regressors X_i are strictly exogenous and $\hat{\beta}$ is $\sqrt{N}$ consistent, it follows that

$$\begin{aligned}
& N^{-\frac{1}{2}}\sum_{i=1}^{N} D_{k,T}'\text{vec}(\hat{e}_i\hat{e}_i' - \hat{\sigma}^2 M) \\
= \; & N^{-\frac{1}{2}}\sum_{i=1}^{N} D_{k,T}'\text{vec}\left(e_i e_i' - \frac{\text{tr}(e_i e_i')}{T-1} M\right) + o_p(1) \\
= \; & N^{-\frac{1}{2}}\sum_{i=1}^{N} D_{k,T}'\left(M \otimes M - \frac{(\text{vec}M)(\text{vec}M)'}{T-1}\right)\text{vec}(e_i e_i' - \sigma^2 M) \\
& + o_p(1) \qquad (5)
\end{aligned}$$

where the second equality follows from equation 4 and Theorem 2 of Magnus and Neudecker(1999, p.30). Since $\hat{\beta}$ and $\hat{\sigma}^2$ are consistent, it follows that

$$D_{k,T}' M_2 \hat{V} M_2 D_{k,T} = D_{k,T}' M_2 V M_2 D_{k,T} + o_p(1). \qquad (6)$$

Combining (5) and (6) we obtain the desired result.

Although we have motivated our test as an LM test, inspection of the test statistic reveals that it also is a Wald statistic that checks whether the sample autocovariances of the fixed effects residuals $\hat{e}_{it}$ are significantly different from their population counterparts under the null hypothesis. As explained in Wooldridge (2002, pp. 270 and 274–275), autocovariances of the fixed effects residuals consistently estimate those of $e_{it} = \varepsilon_{it} - \bar{\varepsilon}_i$, not ε_{it}. As a result, under the null hypothesis that ε_{it} is serially uncorrelated, the sample autocovariances of $\hat{e}_{it}$ converge not to zero, but rather to $-\sigma^2/(T-1)$. Our test ascertains whether the discrepancies between the sample autocovariances and $-\hat{\sigma}^2/(T-1)$ are statistically significant. Equivalently, it tests whether the sample autocorrelations are significantly different from $-1/(T-1)$.

Our test can be usefully modified in two ways. First, in certain circumstances, especially when T is relatively large, it may be desirable to focus the test statistic on only the lower-order autocovariances. Including all the autocovariances may lead to a loss of power analogous to that from including too many orders of autocorrelation when the Box-Pierce test is used with a single time series.

Second, the test is readily adapted to the case of unbalanced panel data in which some observations are randomly missing. Let $s_i = [s_{i1}, ..., s_{iT}]'$ denote the T-dimensional column vector of selection indicators: $s_{it} = 1$ if x_{it} and y_{it} are observed and $s_{it} = 0$ otherwise. Then with some abuse of notation the modified LM statistic can be written as

$$LM = N^{-\frac{1}{2}} \sum_{i=1}^{N} \text{vec}(\hat{e}_i \hat{e}_i' - \hat{\sigma}_2 M_i)' D_{k,T} [D_{k,T}' \hat{W} D_{k,T}]^{-1} N^{-\frac{1}{2}} \sum_{i=1}^{N} D_{k,T}' \text{vec}(\hat{e}_i \hat{e}_i' - \hat{\sigma}_2 M_i)$$

where

$$\begin{aligned}
M_i &= I_T - s_i s_i'/(s_i' s_i), \\
M_{2,i} &= M_i \otimes M_i' - (\text{vec} M_i)(\text{vec} M_i)'/(s_i' s_i - 1), \\
\hat{\beta} &= (\sum_{i=1}^{N} X_i' M_i X_i)^{-1} \sum_{i=1}^{N} X_i' M_i y_i, \\
\hat{e}_i &= M_i (y_i - X_i \hat{\beta}), \\
\hat{\sigma}^2 &= \frac{1}{N} \sum_{i=1}^{N} \frac{1}{s_i' s_i - 1} \sum_{t=1}^{T} s_{it} \hat{e}_{it}^2, \\
\hat{W} &= \frac{1}{N} \sum_{i=1}^{N} M_{2,i} \text{vec}(\hat{e}_i \hat{e}_i' - \hat{\sigma}^2 M_i) \text{vec}(\hat{e}_i \hat{e}_i' - \hat{\sigma}^2 M_i)' M_{2,i}.
\end{aligned}$$

Theorem 2. Suppose that

(a) X_i, ε_i, s_i are iid and have finite fourth moments.

(b) $E(\varepsilon_i|X_i, c_i, s_i) \; = \; 0$.

(c) $\text{rank}[E(X_i' M_i X_i)] \; = \; \dim(x_{it})$.

(d) $s_i' s_i \geq 2$ for all i and $T \geq 3$.

(e) The rank of $E[M_{2,i}\text{vec}(\hat{e}_i\hat{e}_i' - \hat{\sigma}^2 M_i)\text{vec}(\hat{e}_i\hat{e}_i' - \hat{\sigma}^2 M_i)' M_{2,i}]$ is greater than or equal to $(T-1)(T-2)/2$.

Then under the null hypothesis that $E(\varepsilon_i\varepsilon_i'|s_i) = \sigma^2 I_T$, the modified LM test statistic is asymptotically distributed (as $N \to \infty$) as $\chi^2((T-1)(T-2)/2)$.

The proof of Theorem 2 is analogous to the proof of Theorem 1 and thus is omitted.

III. Monte Carlo Analyses

We have shown that, under the null hypothesis, our portmanteau test statistic converges to a $\chi^2((T-1)(T-2)/2)$ distribution as $N \to \infty$, but how applicable is that distribution when N is large but finite? To explore that question, we have conducted a Monte Carlo study in which the data generating process is equation (1) with $\beta = 0$, scalar $x_{it} \sim i.i.d.\, N(0,1)$, $\varepsilon_{it} \sim i.i.d.\, N(0,1)$. The sample sizes considered are $N = 50, 100, 250, 500$ and $T = 5, 8$. The number of Monte Carlo replications is set to 5000, and the deleted time period in the test statistic is $k = 1$.

Table 1 reports the actual rejection frequencies when the nominal size is 5%. In the experiments with $T = 5$, the empirical sizes come quite close to the nominal size. With $T = 8$, there are some mild size distortions at smaller N. These distortions mostly disappear by the time N reaches 500.

We also have conducted a series of Monte Carlo analyses to investigate the power of our test and compare it to the power of several other tests. Unlike our portmanteau test, most existing tests focus on the specific alternative of nonzero first-order autocorelation. For example, Bhargava, Franzini, and Narendranathan (1982), assuming normality of the error term, have developed a Durbin-Watson test against the alternative that the error term follows a first-order autoregression.[3] Wooldridge (2002, p. 275) has suggested applying OLS to the first-order autoregression of $\hat{e}_{it}$ and then performing a t-test of the hypothesis that the autoregressive coefficient equals

[3]Baltagi and Wu (1999) have proposed a related test. Hansen's (2003) analysis emphasizes feasible generalized least squares estimation, but his methods can be used to formulate a test against the alternative of a p^{th}-order autoregression.

$-1/(T-1)$. He emphasized that the standard error estimate in the denominator of the t-ratio must be robust to serial correlation. Another test, hinted at by Wooldridge (2002, pp. 282–283) and developed by Drukker (2003), is based on the residuals from OLS estimation of the first difference of equation (1). This test applies OLS to the first-order autoregression of the residuals and then performs a t-test of the hypothesis that the autoregressive coefficient equals $-1/2$. Again, the standard error estimate must be robust to serial correlation. Finally, Kezdi (2002) has proposed a White-type test that checks whether a covariance matrix estimate robust to serial correlation differs significantly from the conventional covariance matrix estimate that assumes no serial correlation.

The Monte Carlo analyses summarized in Table 2 compare the performance of all these tests in experiments with $T = 8$, $N = 500$, and 5000 replications. The four rows of the table correspond to experiments with four different data generating processes. The first (DGP1) is the same as in Table 1: the null hypothesis case of no serial correlation. In the second (DGP2), the error term ε_{it} follows a first-order autoregression with autoregressive parameter 0.4. In the third (DGP3), ε_{it} follows a second-order moving average process with first-order parameter 0.375 and second-order parameter 0.6. With these parameter values, the first- and second-order autocorrelations equal each other and are approximately 0.4. In each of these three data generating processes, $Var(\varepsilon_{it}) = 1$. In the fourth (DGP4), ε_{it} follows the nonstationary process

$$\varepsilon_{it} \;=\; v_{it} + \alpha_i t$$

where v_{it} and α_i are i.i.d. normal with zero mean, $\mathrm{Var}(v_{it}) = 0.5$, and $\mathrm{Var}(\alpha_i) = 0.02$. This experiment represents the situation in which misspecification of the fixed effects model (namely, the omission of the individual-specific linear time trends) may or may not be detected by serial correlation diagnostics.[4] The possibility of detection arises because the fixed effects residuals will be positively autocorrelated at low orders, negatively autocorrelated at high orders, and heteroskedastic.

Two general results from Table 2 are worth noting at the outset. First, as shown in the first row, all the tests display an empirical size reasonably close to the nominal size of 0.05. Second, as shown in the last column, the

[4]See Solon (1984a), Jacobson, LaLonde, and Sullivan (1993), Friedberg (1998), and Donohue and Levitt (2001) for examples of longitudinal analyses involving individual- or state-specific time trends.

Kezdi test has no power against any of the departures from the null hypothesis. This is by design: the regressor in our experiments is i.i.d. As a result, the conventional covariance matrix estimator remains consistent despite serial correlation (and, in DGP4, heteroskedasticity) of the error term. Kezdi's test, which compares conventional and robust covariance matrix estimates, discovers no problem. If consistency of standard error estimation were the only concern, this would be a good outcome. As noted in Section I, however, there are two other motives for serial correlation diagnostics: (1) to detect model misspecification and (2) to check whether more efficient estimation may be possible through feasible generalized least squares. Our experiments highlight the point that Kezdi's test sometimes lacks power for these purposes.

The other lessons from Table 2 are specific to the particular data generating processes. Under DGP2, the serial correlation of the error term is most pronounced at the first order. Because all of the tests other than Kezdi's are sensitive to this type of serial correlation, all of them show good power.

Under DGP3, the first- and second-order autocorrelations both are around 0.4. Most of the tests still are powerful, but not the Wooldridge-Drukker test based on the residuals from first-difference estimation. That test checks an implication of the null hypothesis that the first-differenced error term has a first-order autocorrelation of $-1/2$. The trouble is that the same implication applies to any process for ε_{it} in which the first- and second-order autocorrelations are the same (but not necessarily zero). By design, the $MA(2)$ process in DGP3 has that property, so the Wooldridge-Drukker test has no power in this case. The more general lesson is that the Wooldridge-Drukker test will lack power for detecting serial correlation of ε_{it} whenever the first- and second-order autocorrelations are similar.

Under DGP4, an important manifestation of the fixed effects model's misspecification is negative higher-order autocorrelations of the residuals. As a result, our portmanteau test tends to show much better power than tests focused on first-order autocorrelation. The reason that our own test specialized to first-order autocovariances is powerful is that it is sensitive to the heteroskedasticity that causes the first-order autocovariances from different time periods to differ from each other.

Although our own tests perform well in all of these experiments, it is obvious that our portmanteau test will be suboptimal in certain circumstances. For example, when serial correlation is most pronounced at the first order and T is sufficiently large, the portmanteau test that uses all autocovariances must surely be less powerful than a test focused on first-order autocorrelation. Our recommendation to practitioners is to use both a portmanteau test and a test for first-order autocorrelation. We are convinced that following this advise would be a major improvement over the typical current practice of ignoring serial correlation altogether.

References

Arellano, Manuel (1987), "Computing Robust Standard Errors for Within-Groups Estimators," *Oxford Bulletin of Economics and Statistics*, 49, 431–434.

Baltagi, Badi H. and Ping X. Wu (1999), "Unequally Spaced Panel Data Regressions with AR(1) Disturbances," *Econometric Theory,* 15, 814–823.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan (2004), "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics*, 119, 249–275.

Bhargava, A., L. Franzini, and W. Narendranathan (1982), "Serial Correlation and the Fixed Effects Model," *Review of Economic Studies*, 49, 533–549.

Chamberlain, Gary (1980), "Analysis of Covariance with Qualitative Data," *Review of Economic Studies,* 47, 225–238.

Donohue, John J., III, and Steven D. Levitt (2001), "The Impact of Legalized Abortion on Crime," *Quarterly Journal of Economics*, 116, 379–420.

Drukker, David M. (2003), "Testing for Serial Correlation in Linear Panel-Data Models," *Stata Journal*, 3, 168–177.

Friedberg, Leora (1998), "Did Unilateral Divorce Raise Divorce Rates?" *American Economic Review*, 88, 608–627.

Hansen, Christian (2003), "Generalized Least Squares Estimation in Differences-in-Differences and Other Panel Models," unpublished.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan (1993), "Earnings Losses of Displaced Workers," *American Economic Review,* 83, 685–709.

Kezdi, Gabor (2002), "Robust Standard Error Estimation in Fixed-Effects Panel Models," unpublished.

Kiefer, Nicholas M. (1980), "Estimation of Fixed Effect Models for Time Series of Cross-Sections with Arbitrary Intertemporal Covariance," *Journal of Econometrics*, 14, 195–202.

Magnus, J.R. and H. Neudecker (1999), *Matrix Differential Calculus with Applications in Statistics and Econometrics,* Revised Edition, Wiley: New York.

Nickell, Stephen (1980), "Correcting the Biases in Dynamic Models with Fixed Effects," Working Paper No. 133, Industrial Relations Section, Princeton University.

Solon, Gary (1984a), "The Effects of Unemployment Insurance Eligibility Rules on Job Quitting Behavior," *Journal of Human Resources*, 19, 118–126.

Solon, Gary (1984b), "Estimating Autocorrelations in Fixed-Effects Models," Technical Working Paper No. 32, National Bureau of Economic Research.

Wooldridge, Jeffrey M. (2002), *Econometric Analysis of Cross Section and Panel Data*, MIT Press: Cambridge, MA.

Table 1.

Empirical Size of Portmanteau Tests for Serial Correlation

N	$T=5$	$T=8$
50	0.048	0.030
100	0.046	0.064
250	0.054	0.067
500	0.053	0.054

Notes: The nominal size is 0.05.

Table 2.

Empirical Power of Tests for Serial Correlation

		Our portmanteau test	Our test specialized to first-order autocovariances	Bhargava et al.	Wooldridge fixed effects	Wooldridge-Drukker first differences	Kezdi
DGP1:	no serial correlation	0.054	0.046	0.049	0.045	0.048	0.041
DGP2:	AR(1)	1.000	1.000	1.000	1.000	1.000	0.039
DGP3:	MA(2)	1.000	1.000	1.000	1.000	0.054	0.041
DGP4:	individual-specific trend	1.000	1.000	0.247	0.192	0.823	0.047

Notes: The nominal size is 0.05, $T = 8$, and $N = 500$. We implement the test of Bhargava et al. with critical values based on an asymptotic normal approximation.