Journal of Econometrics 27 (1985) 371–381. North-Holland

DURBIN–WATSON TESTS FOR SERIAL CORRELATION IN REGRESSIONS WITH MISSING OBSERVATIONS*

Jean-Marie DUFOUR and Marcel G. DAGENAIS

Université de Montréal, Montréal, Québec, Canada H3C 3J7

Received October 1983, final version received October 1984

We study two Durbin–Watson type tests for serial correlation of errors in regression models when observations are missing. We derive them by applying standard methods used in time series and linear models to deal with missing observations. The first test may be viewed as a regular Durbin–Watson test in the context of an extended model. We discuss appropriate adjustments that allow one to use all available bounds tables. We show that the test is locally most powerful invariant against the same alternative error distribution as the Durbin–Watson test. The second test is based on a modified Durbin–Watson statistic suggested by King (1981a) and is locally most powerful invariant against a first-order autoregressive process.

1. Introduction

The most common test against the autocorrelation of errors in regression models is the bounds test of Durbin and Watson (1950, 1951, 1971). This test is easy to compute, exact in small samples (under standard assumptions) and possesses optimal power properties against first-order serial dependence [see Durbin and Watson (1950, pp. 423–425; 1971, pp. 13–15), King (1980, 1981a)]. When observations are either missing or excluded from the regression (e.g. because they are viewed as outliers), there is obviously a difficulty in performing the Durbin–Watson (DW) test. The main paper on testing for autocorrelation when observations are missing is by Savin and White (1978). They discuss two Durbin–Watson type tests: the d_* and d' tests. The d_* test is a popular solution which consists of dropping from the numerator of the DW statistic all the differences between residuals at the ends of each gap. Though this has intuitive appeal, it also modifies the structure of the test

*This work was supported by the Natural Sciences and Engineering Research Council of Canada, the Social Sciences and Humanities Research Council of Canada and the Fonds F.C.A.C. du Québec. The authors thank Maxwell King, Jan Kmenta, James MacKinnon, Lonnie Magee, Roch Roy and two anonymous referees for several useful comments.

statistic and usual tables are not applicable.[1] The d' test uses the DW statistic that one obtains by simply ignoring the presence of gaps in the data. The d' statistic yields an exact test for which some of the existing tables can be applied [e.g., Durbin and Watson (1951), Savin and White (1977)]. It is the test most strongly recommended by Savin and White (1978). Further work by Richardson and White (1979) on the power of the d' and d_* tests indicates that none uniformly dominates the other and, in any case, differences are small.

On the other hand, the d' test has three less attractive characteristics. First, the d' statistic contains comparisons between residuals separated by lags greater than one: clearly, this is not a natural feature in a test against first-order serial correlation. Second, closing the gaps modifies the structure of the regressor matrix; consequently, we cannot generally apply specialized tables, like those published by King (1981b, 1983) for regressions with a trend or seasonal dummy variables. Third, if we consider the serial dependence schemes against which the DW test was traditionally gauged, we can obtain other tests for which we can prove optimal power properties (as we will see below).

In this paper, we study two alternative Durbin–Watson type tests that retain the advantages of the d' test and avoid its shortcomings. The first one is obtained by applying standard techniques used in linear models and time series analysis to deal with missing observations. The test statistic is based on least-squares residuals and is very easy to compute from widely available programs. One can obtain valid significance bounds from all available DW bounds tables, provided appropriate adjustments are made on the numbers of observations and coefficients: in particular, we show that this adjustment depends on the number of gaps in the data, not the number of missing observations. We prove also that the test is locally most powerful invariant (LMPI) against precisely the same alternative error distribution as the one for which the standard DW test is locally most powerful invariant. Besides, one method of computing the test allows the investigator to work as if the data set were complete. If a bounds test is inconclusive, one can compute exact probability tails by directly applying computer routines developed for the standard DW statistic (in the complete sample case). The second test is obtained by applying the same techniques to the modified DW statistic suggested by King (1981a): while remaining computationally convenient, this test has the neat property of being locally most powerful invariant against a first-order autoregressive process on the errors.

In section 2, we derive and define the test statistics. In section 3, we discuss their distribution under the null hypothesis. Finally, in section 4, we show that the tests suggested have optimal local power properties.

[1] See Wallis (1972), Dhrymes (1978, pp. 174–175), Savin and White (1978). For a long time, this solution was suggested in the widely used TSP program [see Hall and Hall (1974, 1980)].

2. Durbin–Watson statistics with missing observations

We consider the following regression model:

$$y_t = x_t'\beta + u_t, \qquad t = 1, \ldots, n, \tag{1}$$

where y_t is the dependent variable, x_t is a $k \times 1$ vector of non-stochastic explanatory variables ($k < n$), β is $k \times 1$ vector of coefficients, and u_t is a random disturbance. The regressor matrix has full column rank and the u_t's are normally distributed with mean zero and the same variance. We wish to test the null hypothesis (H_0) that the errors u_t, $t = 1, \ldots, n$, are independent against an alternative of the form

$$\begin{aligned} &u_t = \rho u_{t-1} + \varepsilon_t, \\ &\varepsilon_t \overset{\text{ind}}{\sim} \mathrm{N}(0, \sigma^2), \quad |\rho| < 1, \quad t = 0, \pm 1, \ldots, \end{aligned} \tag{2}$$

with

$$\mathbf{H}_1^+ : \rho > 0 \quad \text{or} \quad \mathbf{H}_1^- : \rho < 0.$$

Let $\hat{\beta}$ be the ordinary least-squares estimate of β and $\hat{u}_t = y_t - x_t\hat{\beta}$, $t = 1, \ldots, n$. When no observation is missing, the DW statistic is given by

$$\begin{aligned} d &= \sum_{t=1}^{n-1} (\hat{u}_t - \hat{u}_{t+1})^2 \Big/ \sum_{t=1}^{n} \hat{u}_t^2 \\ &= 2(1 - r_1) - \left[\left(\hat{u}_1^2 + \hat{u}_n^2 \right) \Big/ \sum_{t=1}^{n} \hat{u}_t^2 \right], \end{aligned} \tag{3}$$

where

$$r_1 = \sum_{t=1}^{n-1} \hat{u}_t \hat{u}_{t+1} \Big/ \sum_{t=1}^{n} \hat{u}_t^2.$$

Suppose now that a subset M of the observations is missing. Let $S = \{1, \ldots, n\} \setminus M$ be the set of the available observations, m the number of elements in M ($0 < m < n - k$), and $n_1 = n - m$ the number of elements in S. We assume also that S contains the observations at the ends of the sample ($t = 1$ and n) and that the regressor matrix of the sample S has full column rank. If we define

$$y_* = [y_t]_{t \in S}, \qquad X_* = [x_t']_{t \in S},$$

where the rows are put in chronological order (from top to bottom), the residuals obtained by running the regression on the incomplete sample are given by

$$\hat{v}_* = y_* - X_*\hat{\beta}_* \equiv [\hat{v}_t]_{t\in S}, \tag{4}$$

where

$$\hat{\beta}_* = (X_*'X_*)^{-1}X_*'y_*.$$

The d' test of Savin and White (1978) is based on the standard DW statistic obtained from this regression, ignoring the fact that observations are missing. One sees easily that this yields an exact test.

A 'gap' in the data is a run (i.e., an uninterrupted set) of missing observations. Define the function $h(t)$ as the length of the gap immediately following an observation t. Clearly, $h(t) = 0$ whenever the observation following t is not missing; the next non-missing observation after t is $t+1+h(t)$. Let g be the number of gaps in the sample ($1 \le g \le m$). If we define $T = S\setminus\{n\}$ and set $\hat{v}_t = 0$ when observation t is missing, we can write the d' statistic as

$$d' = 2(1 - r_1') - \left[(\hat{v}_1^2 + \hat{v}_n^2)\Big/\sum_{t=1}^{n}\hat{v}_t^2\right], \tag{5}$$

where

$$r_1' = \sum_{t\in T}\hat{v}_t\hat{v}_{t+1+h(t)}\Big/\sum_{t=1}^{n}\hat{v}_t^2.$$

From the latter expression we see clearly that the estimate of the first autocorrelation coefficient on which d' is implicitly based contains comparisons of residuals separated by lags greater than 1. Clearly, autocorrelations with opposite signs (at different lags) may tend to offset each other.

There is a simple alternative way of treating missing observations when estimating serial correlation coefficients: it just consists of dropping the cross-products where a missing observation would normally appear. This method was used by several authors in the time series literature, mainly for estimation purposes [see Jones (1962), Parzen (1963), Marshall (1980)]. A statistic suggested by this method is the following one:

$$d_0 = 2(1-\bar{r}_1) - \left[\left(\hat{v}_1^2 + \hat{v}_n^2 \right) \Big/ \sum_{t=1}^{n} \hat{v}_t^2 \right]$$

$$= \sum_{t=1}^{n-1} (\hat{v}_t - \hat{v}_{t+1})^2 \Big/ \sum_{t=1}^{n} \hat{v}_t^2, \tag{6}$$

where

$$\bar{r}_1 = \sum_{t=1}^{n-1} \hat{v}_t \hat{v}_{t+1} \Big/ \sum_{t=1}^{n} \hat{v}_t^2;$$

since $\hat{v}_t = 0$ whenever t is missing, all cross-products involving a missing $\hat{v}_t$ in $\bar{r}_1$ cancel. We will see below that tests based on d_0 may be performed in the same way as standard DW tests, except for the critical values.

King (1981a) showed that a slightly modified DW statistic yields a LMPI test of H_0 against autoregressive processes of order 1 (as opposed to an approximately LMPI test). The procedure followed above to obtain d_0 can be applied to this statistic. We then get

$$d_0' = d_0 + \left[\left(\hat{v}_1^2 + \hat{v}_n^2 \right) \Big/ \sum_{t=1}^{n} \hat{v}_t^2 \right] = 2(1-\bar{r}_1), \tag{7}$$

where $\hat{v}_t = 0$, when t is missing.

We see immediately that d_0 is a simple function of the least-squares residuals. To obtain critical values, we will use an alternative method of deriving the statistic d_0 (section 3). From it, we will see that we can get valid bounds for d_0 by using standard Durbin–Watson tables with $n_1 + g$ observations and $k + g$ regression coefficients (where n_1 is the effective number of observations and g the number of gaps). Further, in many situations, we will find that this alternative form of the test is computationally convenient to obtain exact probability tails. The same remarks will apply to d_0', provided one uses an appropriate table.

3. Null distribution

Consider the regression

$$y_t = x_t'\beta + \sum_{i=1}^{g} \gamma_i \delta(t, k_i) + u_t, \qquad t \in S \cup K, \tag{8}$$

where $K = \{k_1, \ldots, k_g\}$, k_i is the index of the first observation in the ith gap ($k_i \neq 1, n$), $\delta(t, s) = 1$ if $t = s$, and $\delta(t, s) = 0$ otherwise.

Then, $S \cup K$ is the set of all available observations plus the first observation in each gap. Assume (provisionally) that y_t and $\boldsymbol{x}_t'$ for $t \in K$ are also available. If we estimate (8) by least-squares (with the observations in chronological order), the estimate of $\boldsymbol{\beta}$ obtained is $\hat{\boldsymbol{\beta}}_*$, i.e., the estimate of $\boldsymbol{\beta}$ based on the available sample S, and the residuals are $\hat{v}_t = y_t - \boldsymbol{x}_t'\hat{\boldsymbol{\beta}}_*$, for $t \in S$, and $\hat{v}_t = 0$, for $t \in K$. This is due to the inclusion of a dummy variable for each observation in K [see Salkever (1976), Dufour (1980)]. Further, one can check easily that the Durbin–Watson statistic from this regression is identical with d_0 in (6) and does not require knowing any value of y_t or $\boldsymbol{x}_t$ for $t \in K$ (so that we do not need the assumption that these are available). Since the model (8) has $n_1 + g$ observations and $k + g$ coefficients and since all the conditions for the DW statistic to have its standard distribution are satisfied, we can find valid bounds for the critical values of this statistic by using a standard DW table with $n_1 + g$ observations and $k + g$ coefficients. It is straightforward to see that these bound corrections also apply to d_0' provided one uses a table appropriate for King's (1981a) modification of the DW statistic.

We may note here that the technique of using observation-specific dummy variables has been fruitfully applied in the past to deal with various problems: missing observations [Bartlett (1937), Wilkinson (1960)], outlier detection [Gentleman and Wilk (1975), John and Draper (1978)], prediction [Salkever (1976), Fuller (1980)] and structural change analysis [Dufour (1980, 1981)]. The tests derived in section 2 may be viewed as an application of this technique.

We can obtain the same DW statistic by introducing a dummy variable for each missing observation (instead of only one per gap) and by estimating the model

$$y_t = \boldsymbol{x}_t'\boldsymbol{\beta} + \sum_{i=1}^{m} \bar{\gamma}_i \delta(t, m_i) + u_t, \qquad t = 1, \ldots, n, \tag{9}$$

where m_i is the ith missing observation. This suggests using $n_1 + m$ and $k + m$ (where $m \geq g$) for reading the table. However, the inconclusive region is then generally wider than the one based on $(n_1 + g, k + g)$, in the sense that the latter will be *contained* in the former (as can be verified by looking at any table). When one of the gaps contains several observations, the corresponding region can actually be much wider, to the point of making the bounds test practically useless. Thus the bounds obtained from $(n_1 + g, k + g)$ are preferable.[2]

[2] Durbin–Watson tests were used previously in the context of models that contain observation-specific dummy variables [e.g. Dufour (1981), Honohan and McCarthy (1982)]. However, none of these authors used the appropriate method of performing the bounds test.

When the test based on d_0 is inconclusive, it is possible to approximate the critical values. Several approximations are discussed by Durbin and Watson (1971). All these can be applied directly to d_0, provided we consider the extended models (8) or (9). A theoretically more attractive procedure is to compute the tail probability associated with $\hat{d}_0$, the observed value of d_0: $\alpha_0 = \mathrm{P}[d_0 \leq \hat{d}_0]$ for a test against H_1^+. We can do this in principle by using the procedures described by Imhof (1961) or Pan Jie-Jian (1968). These methods avoid inconclusive tests but may involve a heavy computational cost (in programming or computer time).

Note here that, if we have an algorithm that computes tail probabilities of the standard DW statistic for an arbitrary regression model, we can obtain the tail areas of d_0 by simply considering one of the extended models (8) or (9).[3] For the practitioner, this may be especially convenient.

Similar methods can be applied to obtain approximate critical values or exact tail probabilities of the statistic d_0'. But algorithms must be adjusted accordingly, since the statistic has a slightly different structure.

4. Power

We will now examine some theoretical power properties of the tests suggested above. From section 3, we know that the statistic d_0 can be viewed as the standard DW statistic computed after estimating the extended model (9). Using this fact, we will first see that d_0 has optimal power properties in the context of this extended model. Then, we will show that a similar result remains valid if we assume that the data are generated by the original model (1) and restrict attention to test functions which do not depend on the missing observations.

Following Durbin and Watson (1971, sec. 3), consider the alternative error density

$$f(\boldsymbol{u}) = K_2 \exp\left(- \frac{1}{2\sigma^2} \boldsymbol{u}' \boldsymbol{B} \boldsymbol{u} \right), \tag{10}$$

where $\boldsymbol{u} = (u_1, u_2, \ldots, u_n)'$, K_2 is the appropriate normalizing constant and B is a matrix such that

$$\boldsymbol{u}'\boldsymbol{B}\boldsymbol{u} = (1+\rho^2) \sum_{t=1}^{n} u_t^2 - \rho\left(u_1^2 + u_n^2\right) - 2\rho \sum_{t=2}^{n} u_t u_{t-1}. \tag{11}$$

[3] Such algorithms are not widely available in statistical or econometric packages. One package that does compute the marginal significance levels for the DW statistic is the SHAZAM package [White (1978)].

The density (10) is a close approximation of the error density when $\boldsymbol{u}$ follows a stationary first-order autoregressive process [as described in (2)]. If we assume that the data are generated by the extended model (9)–(10), we can conclude from Durbin and Watson (1971) that the test with critical region $d_0 < d_0(\alpha)$ is a locally most powerful invariant test for testing H_0: $\rho = 0$ against H_1^+: $\rho > 0$; $d_0(\alpha)$ is the critical value for a test of level α. The group of transformations G under which the invariance property holds is the set of all transformations of the form

$$\tilde{y}_t = c_0 y_t + \boldsymbol{x}_t'\boldsymbol{b} + \boldsymbol{D}_t'\boldsymbol{d}, \qquad t = 1, \dots, n, \tag{12}$$

where $0 < c_0 < +\infty$, $-\infty < b_j < +\infty$, $j = 1, \dots, k$, $-\infty < d_i < +\infty$, $i = 1, \dots, m$, $\boldsymbol{D}_t = [\delta(t, m_i)\colon m_i \in M]'$, $\boldsymbol{b} = (b_1, \dots, b_k)'$ and $\boldsymbol{d} = (d_1, \dots, d_m)'$; M is the set of the m missing observations as defined in section 2. Further, when the columns of the $n \times (k+m)$ regressor matrix $Z = [\boldsymbol{Z}_1, \dots, \boldsymbol{Z}_n]'$, where $\boldsymbol{Z}_t = [\boldsymbol{x}_t', \boldsymbol{D}_t']'$, are linear combinations of $k+m$ eigenvectors of B, the same test is uniformly most powerful for testing H_0 against H_1^+ [from Anderson (1948)].

Let L be the set of possible probability distributions of $\boldsymbol{y} = (y_1, \dots, y_n)'$ when (9) and (10) hold, and let V be the class of test functions $\phi(\boldsymbol{y})$ invariant under all transformations in G. Distributions in L may be indexed by the parameter vector $\boldsymbol{\theta} = (\boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma, \rho)$, where $\boldsymbol{\beta} = (\beta_1, \dots, \beta_k)'$ and $\boldsymbol{\gamma} = (\bar{\gamma}_1, \dots, \bar{\gamma}_m)'$. The group G induces a group $\bar{G}$ on the parameter vector $\boldsymbol{\theta}$: when the transformation (12) is applied to $\boldsymbol{y}$, $\boldsymbol{\theta}$ becomes $\tilde{\boldsymbol{\theta}} = (c_0\boldsymbol{\beta} + \boldsymbol{b}, c_0\boldsymbol{\gamma} + \boldsymbol{d}, c_0\sigma, \rho)$. We can see easily that $S(\boldsymbol{\theta}) \equiv \rho$ is a maximal invariant under $\bar{G}$ and, thus, any test $\phi(\boldsymbol{y})$ invariant under G has a distribution which depends on the single parameter ρ [see Lehmann (1959, p. 220)]. The level and the power function of any test in V do not depend on the values of $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$ and σ.

Since we are mainly interested in samples generated by the original model (1) and where the observations in M are missing, we will now assume that the distribution of $\boldsymbol{y}$ belongs to the class L_1 described by eqs. (1) and (10) and show that the test based on d_0 remains locally most powerful among invariant tests which do not depend on the missing observations. In this case, we consider the set V_1 of tests of the form $\phi(\boldsymbol{y}_*)$ which are invariant under the group G_1 of transformations

$$\tilde{y}_t = c_0 y_t + \boldsymbol{x}_t'\boldsymbol{b}, \qquad t = 1, \dots, n, \tag{13}$$

where $0 < c_0 < +\infty$, $-\infty < b_i < +\infty$, $i = 1, \dots, k$; $\boldsymbol{y}_* = [y_t]_{t \in S}$ is the vector of available observations. A test of the form $\phi(\boldsymbol{y}_*)$ can be viewed as a test $\phi(\boldsymbol{y})$ where the missing observations are not used.

Since $\boldsymbol{D}_t = \boldsymbol{0}$ for $t \in S$, any transformation in the group G can be written as

$$\begin{aligned} \tilde{y}_t &= c_0 y_t + \boldsymbol{x}_t'\boldsymbol{b} && \text{if} \quad t \in S, \\ &= c_0 y_t + \boldsymbol{x}_t'\boldsymbol{b} + d_i && \text{if} \quad t = m_i \in M, \end{aligned} \tag{14}$$

where d_i is an arbitrary real number and $c_0 > 0$. It is immediate from (14) that any test of the form $\phi(\boldsymbol{y}_*)$ which is invariant under G_1 is also invariant under G: hence $V_1 \subseteq V$. Conversely, the fact that d_i, $i = 1, \ldots, m$, can take arbitrary values implies that any test $\phi(\boldsymbol{y})$ invariant under the group G must be invariant with respect to all possible changes of the missing observations: tests in V do not depend on the missing observations $\{y_t\colon t \in M\}$ and thus have the form $\phi(\boldsymbol{y}_*)$. Further, if we set $\boldsymbol{d} = \boldsymbol{0}$, we see that any test invariant under G is also invariant under G_1. Hence $V \subseteq V_1$, and since $V_1 \subseteq V$, we have $V_1 = V$.

Therefore, the test with critical region $d_0 < d_0(\alpha)$ is locally most powerful for testing H_0 against H_1^+ among (level-α) tests which are invariant under G_1 and do not depend on the missing observations. This holds irrespective of the values of $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$ and σ. Further, since L_1 is a subset of L (obtained by setting $\boldsymbol{\gamma} = \boldsymbol{0}$), this property still holds if we assume that the distribution of $\boldsymbol{y}$ belongs to L_1.

It is important to note here that the inclusion relationship $L_1 \subseteq L$ may suggest that a more powerful test, that would exploit the information $\boldsymbol{\gamma} = \boldsymbol{0}$, could be obtained.[4] As noted above, if the distribution of $\boldsymbol{y}$ belongs to L, the family of invariant tests to consider is V; on the other hand, if the distribution of $\boldsymbol{y}$ belongs to L_1 and observations are missing, it is natural to take V_1 as the appropriate family of invariant tests. However, it turns out here that the sets V and V_1 coincide. In other words, information about $\boldsymbol{\gamma}$ only affects missing observations which do not enter into the test statistics under consideration. Then, since the distribution of any test in V does not depend on the parameters $(\boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma)$ for any distribution in L (and thus in L_1), the optimality of d_0 in V under L implies its optimality in V_1 under L_1. The proof is thus complete.

It is easy to verify that the test based on d' as well as all the other tests considered by Savin and White (1978) are invariant under the transformation groups described above. Therefore, the test based on d_0 is at least as powerful and thus usually more powerful than each of these tests, at least in the neighborhood of the null hypothesis. This holds when the alternative error distribution is simply the same as the one postulated by Durbin and Watson (1971, sec. 3).

Consider now the statistic d_0'. The latter is King's modification of the DW statistic obtained after estimating the extended model (9). Like d_0, it does not

[4] For an illustration of the fact that exploiting restrictions on regressions coefficients can increase the power of the resulting autocorrelation test, see King (1981c).

depend on the missing observations and is invariant under the groups G and G_1. In the context of the extended model and using theorem 3 of King (1981a), we get that a critical region of the form $d_0' < d_0'(\alpha)$ is a LMPI test of H_0: $\rho = 0$ against H_1^+: $\rho > 0$, when the alternative error distribution corresponds to a first-order autoregressive process as given in (2). Then, by an argument analogous to the one given above, we see that the test based on d_0' is most powerful, at least in the neighborhood of the null hypothesis, among all invariant tests based on the available observations; the transformation group is defined by (13). This test has the neat property of being LMPI against the standard model of stationary first-order autoregressive process. From this point of view, it has a theoretical advantage over the tests considered by Savin and White (1978) and the test based on d_0. A similar result holds for the problem of testing H_0 against H_1^-: $\rho < 0$; in this case, the critical region has the form $d_0' > d_0''(\alpha)$. On the other hand, we must keep in mind that the test is not necessarily most powerful for all values of ρ. King (1981a, sec. 4) gives examples showing that the standard DW test can become more powerful than the modified test as the value of $|\rho|$ increases.

In theory, tests based on d_0' are not more difficult to perform than those based on d_0. However, tables applicable to d_0' are not as extensive as those applicable to d_0; further, canned programs routinely compute the standard DW statistic (and possibly, its tail probabilities) but do not produce the modified statistic. Because of its convenience, we thus recommend the test based on d_0 to practitioners.

Applications of the tests discussed in this article as well as further details on their implementation are discussed in a working paper [Dagenais and Dufour (1984)].

References

Anderson, T.W., 1948, On the theory of testing serial correlation, Skandinavisk Aktuarietidskrift 31, 88–116.

Bartlett, M.S., 1937, Some examples of statistical methods of research in agriculture and biology, Journal of the Royal Statistical Society, Suppl. 4, 137–170.

Dagenais, M.G. and J.-M. Dufour, 1984, Durbin–Watson tests with missing observations: Applications and comparisons, Proceedings of the Business and Economic Statistics Section, American Statistical Association, forthcoming.

Dhrymes, P.J., 1978, Introductory econometrics (Springer-Verlag, New York).

Dufour, J.-M., 1980, Dummy variables and predictive tests for structural change, Economics Letters 6, 241–247.

Dufour, J.-M., 1981, Variables binaires et tests prédictifs contre les changements structurels: Une application à l'équation de St. Louis, L'Actualité Economique 57, 377–386.

Durbin, J. and G.S. Watson, 1950, Testing for serial correlation in least-squares regression, I, Biometrika 37, 409–428.

Durbin, J. and G.S. Watson, 1951, Testing for serial correlation in least-squares regression, II, Biometrika 38, 159–178.

Durbin, J. and G.S. Watson, 1971, Testing for serial correlation in least-squares regression, III, Biometrika 58, 1–19.

Fuller, W.A., 1980, The use of indicator variables in computing predictions, Journal of Econometrics 12, 231–243.

Gentleman, J.F. and M.B. Wilk, 1975, Detecting outliers, II: Supplementing the direct analysis of residuals, Biometrics 31, 387–410.

Hall, R.E. and B.H. Hall, 1974, TSP, Time series processor (Version 2.4), Technical paper no. 12 (Harvard Institute of Economic Research, Harvard University, Cambridge, MA). [Version (3.4), 1980].

Honohan, P. and C. McCarthy, 1982, On the use of Durbin–Watson type statistics where there are missing observations, The Statistician 31, 149–152.

Imhof, J.P., 1961, Computing the distribution of quadratic forms of normal variables, Biometrika 48, 419–426.

John, J.A. and N.R. Draper, 1978, On testing for two outliers or one outlier in two-way tables, Technometrics 20, 69–78.

Jones, R.H., 1962, Spectral analysis with regularly missed observations, Annals of Mathematical Statistics 2, 455–461.

King, M.L., 1980, Robust tests for spherical symmetry and their application to least-squares regression, Annals of Statistics 8, 1265–1271.

King, M.L., 1981a, The alternative Durbin–Watson test: An assessment of Durbin and Watson's choice of statistic, Journal of Econometrics 17, 51–66.

King, M.L., 1981b, The Durbin–Watson test for serial correlation: Bounds for regression with trend and/or seasonal dummy variables, Econometrica 49, 1571–1581.

King, M.L., 1981c, The Durbin–Watson bounds test and regressions without an intercept, Australian Economic Papers 20, 161–170.

King, M.L., 1983, The Durbin–Watson test for serial correlation: Bounds for regressions using monthly data, Journal of Econometrics 21, 357–366.

Lehmann, E.L., 1959, Testing statistical hypotheses (Wiley, New York).

Marshall, R.J., 1980, Autocorrelation estimation of time series with randomly missing observations, Biometrika 67, 567–570.

Pan Pie-Jian, 1968, Distribution of noncircular serial correlation coefficients, Selected Translations in Mathematical Statistics and Probability (printed for the Institute of Mathematical Statistics by the American Mathematical Society) 7, 281–291.

Parzen, E., 1963, On spectral analysis with missing observations and amplitude modulation, Sankhya A 25, 383–392.

Richardson, S.M. and K.J. White, 1979, The power of tests for autocorrelation with missing observations, Econometrica 47, 785–788.

Salkever, D.S., 1976, The use of dummy variables to compute predictions, prediction errors, and confidence intervals, Journal of Econometrics 4, 393–397.

Savin, N.E. and K.J. White, 1977, The Durbin–Watson test for serial correlation with extreme sample sizes or many regressors, Econometrica 45, 1989–1996.

Savin, N.E. and K.J. White, 1978, Testing for autocorrelation with missing observations, Econometrica 46, 59–67.

Wallis, K.F., 1972, Testing for fourth order autocorrelation in quarterly regression equations, Econometrica 40, 617–636.

White, K.J., 1978, A general computer program for econometric methods – SHAZAM, Econometrica 46, 239–240. [Version 4.4 (June 1982), University of British Columbia].

Wilkinson, G.N., 1960, Comparison of missing value procedures, Australian Journal of Statistics 2, 53–65.