Testing for Unit Roots: What Should Students Be Taught?

John Elder and Peter E. Kennedy

Abstract: Unit-root testing strategies are unnecessarily complicated because they do not exploit prior knowledge of the growth status of the time series, they worry about unrealistic outcomes, and they double- or triple-test for unit roots. The authors provide a testing strategy that cuts through these complications and so facilitates teaching this dimension of the unit-root phenomenon. *F* tests are used as a vehicle for understanding, but *t* tests are recommended in the end, consistent with common practice.

Key words: teaching econometrics, unit roots

JEL codes: A220, A230, C220

Graphing GDP (gross domestic product) over time suggests that it is a time series growing at a constant rate, with several bumps along the way. Consider the following two ways of modeling this time series, where y_t represents GDP at time t, θ is the growth rate, and ε_t is an error term with mean one:

Model A: $y_t = y_0 e^{\theta t}\varepsilon_t$.

Model B: $y_t = y_{t-1} e^{\theta}\varepsilon_t$.

In model A, we represent the time series by specifying that at time 0, GDP began at y_0 and then from that base grew at a compound rate of 100*θ percent, with an error playing a role in each year. In model B, we represent the time series by specifying that GDP grows by 100*θ percent from the year before, with an error playing a role in each year. These two specifications both appear reasonable, and *a priori* there seems little to choose between them—without doing some empirical tests, the choice between them is subjective. Which of these two specifications would you choose?

When we ask students to choose between these two specifications on subjective grounds, they typically split their votes. Indeed, most believe that these specifications are inherently equivalent and so which is chosen is not of much import. They are surprised to learn that it is very important which specification represents the real world. To see why this is the case, take logs of both models to get

Model A′: $\ln y_t = \ln y_0 + \theta t + \ln\varepsilon_t$.

John Elder *(e-mail: John_Elder@ndsu.nodak.edu) is an assistant professor of finance at North Dakota State University.* ***Peter E. Kennedy*** *is a professor of economics at Simon Fraser University.* *This article is a revised version of a presentation at a conference at Middlebury College in November 1999. The authors thank Dorian Owen and Simon Xu for thoughtful comments.*

Model B′: $\ln y_t = \ln y_{t-1} + \theta + \ln\varepsilon_t$.

Now substitute repeatedly for the lagged ln*y* value in model B′ to get it in a form comparable to that of model A′.

$$\text{Model B}'': \ln y_t = \ln y_0 + \theta t + \sum_{i=1}^{t} \ln\varepsilon_i ,$$

which reveals that the error term in model B is doing something dramatically different than in model A. In model A, an error term affects what is happening in the current time period but has no effect on what happens in succeeding time periods. In contrast, in model B, an error term affects what happens in the current time period and also in every succeeding time period. In model B, a shock to GDP persists, but in model A, it disappears after the current time period. This has profound implications for macroeconomic theory—are shocks, policy or otherwise, permanent or transitory?

It also has profound implications for econometric estimation and testing. If $\ln y_t$, is evolving according to model B, it embodies a growing number of error components, so its variance is increasing without limit. This causes big trouble for test statistics when variables with this character are involved in regressions.

Such variables are said to have unit roots, a terminology arising from the fact that in model B′ the coefficient on $\ln y_{t-1}$ is unity. If this coefficient were less than unity, a shock to GDP would not persist—it would die out over time, a result that can easily be seen by calculating what the composite error term in model B″ would look like under this circumstance. If two variables with unit roots are regressed on each other, spurious results are obtained—*t* statistics are misleadingly high, as are *R* squares, and Durbin-Watson statistics are very low. To guard against being misled in this respect, it is important for researchers to test for whether the variables they are using in regressions contain unit roots; if unit roots are found to be present, an alternative estimation procedure should be undertaken.

The appropriate alternative estimation procedure typically involves checking for co-integration and specifying an error correction model, both of which lie beyond the scope of this article; most modern econometrics texts explain these techniques. Our purpose in this study is to exposit what has become the mandatory initial step in modern time-series analysis—testing for unit roots.

Judging by textbook expositions and comments such as those of Ayat and Burridge (2000, 74), the augmented Dickey-Fuller (ADF) (1981) test has become the most popular of many competing tests in the literature. Thus, it is natural for instructors in undergraduate econometrics courses to focus their exposition of unit-root testing on the ADF test. The essence of this test is to run the regression indicated by model B′, and test if the coefficient of $\ln y_{t-1}$ is unity using a *t* test.

A crucial ingredient of this test, not well recognized in textbooks, is that a testing *strategy* is required, as opposed to mere calculation of a single test statistic. This strategy is necessary to determine if an intercept, an intercept plus a time trend, or neither an intercept nor a time trend, should be included in the regression run to conduct the unit-root test. Including too many of these deterministic regressors results in lost power, whereas not including enough of them biases the test in favor of the unit-root null. As discussed below, inclusion of an intercept, or an intercept and a time trend, is necessary to allow representation of the alternative hypothesis competing against the null of a unit root. In the example presented earlier, it appears from looking at model B′ that an intercept should be included; but the alternative, model A′, contains both an intercept and a trend, suggesting that in testing for a unit root (a coefficient of unity on $\ln y_{t-1}$) both an intercept and a time trend should be included in the model B′ regression. Because the data-generating structure (presence of an intercept or trend, for example) is unknown, testing for a unit root involves simultaneously determining whether an intercept and/or a time trend belong; this, in turn, requires a testing strategy.

Because of these complications, many students, and probably some instructors (if our own experience is any guide), find the subject of testing for unit roots somewhat overwhelming. Our purpose in this article is to propose a unit-root testing strategy that, in the end, corresponds to what is taught in textbooks and is popular in empirical work but which is more readily understood by students. Further, for the benefit of students, we enhance this exposition by using Monte Carlo results to illustrate why special critical values are required for its implementation. To sharpen the focus, we confine ourselves to the testing strategy and so deliberately do not discuss many other facets of unit-root testing that an instructor would normally include in his or her exposition, a caveat to which we return in our conclusion.

UNIT-ROOT TESTING STRATEGIES

It is widely recognized in the literature that a testing strategy is needed when testing for a unit root. Perron (1988), Dolado, Jenkinson, and Sosvilla-Rivero (1990), Holden and Perman (1994), Enders (1995), and Ayat and Burridge (2000) propose such strategies, all of which are more complicated than we believe is necessary, particularly for students. They begin by estimating the equation

$$y_t = \rho\, y_{t-1} + \alpha + \beta t + \varepsilon_t \text{ in the form} \qquad (1)$$

$$\Delta y_t = (\rho - 1)y_{t-1} + \alpha + \beta t + \varepsilon_t$$

and then test for $\rho = 1$ (the unit root) while worrying about whether α and/or β are zero or nonzero.[1]

These strategies are complicated for three reasons.

1. They do not exploit prior knowledge of the growth status of the variable under test, forcing their strategies to cover all possibilities. For example, unemployment clearly does not have a long-run growth trend, and so for this variable, unit-root testing should begin by setting the trend coefficient $\beta = 0$, unless a plot of the data reveals an apparent trend during the time period covered by the data. (Graphing variables against time should always be the first step in any time-series analysis.)

2. They worry about outcomes that are not realistic. For example, consider the case of $\rho = 1$ and $\beta \neq 0$, which implies the simultaneous existence of a unit root and a trend. This is thought to be unrealistic,[2] as noted for example by Perron (1988, 304) and by Holden and Perman (1994, 63). Omitting such unreasonable

cases from consideration should simplify these strategies. Furthermore, such a simpler approach can avoid examples such as that in Enders (1995, 259), in which his strategy “works itself into an uncomfortable corner” by concluding that real GNP (gross national product) follows an implausible process. Such results are probably Type II errors resulting from the poor power of unit root tests.

3. They double- and triple-test for $\rho = 1$ as their testing strategies proceed,[3] presumably to alleviate the low power of unit-root tests. Although this may have some merit, it flirts with pretest bias, and, of more importance, complicates the strategy for students as well as practitioners. For students, this hinders understanding, and for practitioners, it may well cause them to use no strategy.

Textbook expositions of unit-root testing vary widely in character, but from our reading, most recommend using an ADF test of $\rho = 1$ by testing $\rho - 1 = 0$ in equation (1). The test statistic is the familiar t statistic but with special critical values employed to reflect its nonnormal (even asymptotically) distribution under the null of a unit root.[4] This test is one-sided because the alternative $\rho > 1$ is ruled out as implying unreasonable explosive behavior.

In contrast, the expositional strategy we suggest rests on F rather than t statistics in the ADF framework, rules out unreasonable cases *a priori*, resists the temptation to double- and triple test $\rho = 1$, and, when possible, exploits known information about the growth character of the variable under investigation. As should become evident in the next section, use of F tests rather than t tests serves to enhance student understanding of the complete story of what is going on with the unit-root phenomenon. Once this exposition has been achieved, however, we argue on power grounds that the F statistics be replaced in practice with t statistics, matching textbook recommendations.

A STUDENT-FRIENDLY TESTING STRATEGY

In many cases, particularly with macroeconomic data, it is reasonable to conclude, on the basis of theoretical considerations and by looking at a plot of data against time, that a variable is or is not growing. Such growth could occur via a deterministic time trend, or it could occur because the annual change in the variable is equal to a constant. In this latter case, the variable is equal to its lagged value plus an intercept and is referred to as having a unit root with drift; the term *drift* refers to this intercept. GDP, consumption, and investment, for example, are clearly growing over time, whereas rates, such as interest, inflation, and unemployment rates, do not grow in the long run. Consequently, it makes sense to begin one’s testing strategy, if possible, by exploiting such information.[5] To reflect this, we divide our testing strategy into three cases—it is known that in the long run y_t is growing (or shrinking), it is known that y_t is not growing, and there is no knowledge about y_t’s growth status.

Case 1: y_t Is Growing

In this case, in equation (1) either

(1) there is a unit root ($\rho = 1$), no time trend ($\beta = 0$), and a nonzero intercept

providing a drift term to create growth; or (2) there is no unit root ($\rho < 1$), but there is a time trend ($\beta \neq 0$), so that y_t is stationary around a deterministic time trend.

Strategy. Conduct an F test to test the joint null that $\rho = 1$ and $\beta = 0$. If this null is not rejected, we conclude that y_t has a unit root with drift. If this null is rejected, we have three possibilities: (1) $\rho \neq 1$ and $\beta = 0$; (2) $\rho \neq 1$ and $\beta \neq 0$; or (3) $\rho = 1$ and $\beta \neq 0$. The third of these possibilities, as noted earlier, is unreasonable and so is ruled out. The first of these possibilities is not consistent with our supposition that y_t is growing, and so it also is ruled out. We conclude that y_t is stationary around a deterministic time trend.

Could this test have been conducted, as most textbooks suggest, by ignoring the issue of $\beta = 0$ and just testing, via a t statistic, the null that $\rho = 1$? Yes, it could—if $\rho = 1$ is not rejected, $\beta = 0$ would have to follow if an unreasonable result is to be avoided, and if $\rho = 1$ is rejected, $\beta \neq 0$ would have to follow if our supposition that y_t is growing is correct. The choice depends on the relative power of the two tests. On the one hand, because when the null is false both ρ and β depart from their joint null values, it is likely that in general the F test will have more power. But on the other hand, F tests cannot exploit the one-sided nature of the alternative, as can t tests, suggesting that the t test will have more power. Practitioners and textbooks appear to have sided with the latter argument.[6]

A Monte Carlo experiment, using the structure reported below but with $\alpha = 0$, $\rho = 0.9$, and $\beta = 0.1$, was conducted to illustrate this. The estimated powers of the t and F tests were 0.30 and 0.35, respectively. Repeating this experiment, changing β to 0.05 results in estimated powers of the t and F tests of 0.22 and 0.20, respectively. A good student exercise is to explain step by step how such a Monte Carlo study would be conducted.

We have concluded that a t test of $\rho = 1$ when estimating equation (1) is the appropriate strategy to employ in this case. This t statistic does not, under the null of a unit root, have a t distribution,[7] and so special critical values are required, as provided for example in Hamilton (1994, 763–64).[8] These critical values were obtained via Monte Carlo studies; an example of the results is provided in Figure 1. Here, 15,000 artificial sets of observations of sample size 100 from the unit-root process $y_t = 0.5 + y_{t-1} + \varepsilon_t$ were generated with ε_t drawn from a standard normal. The resulting distribution of t values is drawn as *t*100 in Figure 1 and is to be contrasted with the corresponding true t distribution, labeled *t*true.

Case 2: y_t Is Not Growing

In this case, in equation (1), $\beta = 0$ and so the estimating equation becomes

$$\Delta y_t = (\rho - 1)\, y_{t-1} + \alpha + \varepsilon_t, \tag{2}$$

where either there is (1) unit root ($\rho = 1$) and a zero intercept ($\alpha = 0$) so there is no drift term to create growth; or (2) no unit root ($\rho < 1$) and a nonzero intercept, so that y_t is stationary around mean $\alpha/(1 - \rho)$.

We are ruling out here the unlikely result that the mean of a stationary economic variable is 0—the interest, inflation, and unemployment rates are not likely to have 0 means, for example. In balder terms, Davidson and MacKinnon

FIGURE 1
Sampling Distribution of *t* Statistics for Case 1

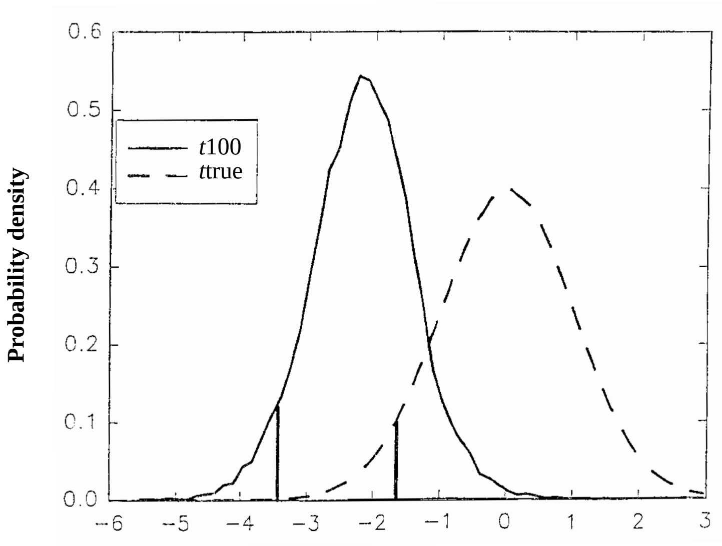

(1993, 702) note that "Testing with zero intercept is extremely restrictive, so much so that it is hard to imagine ever using it with economic time series."

Strategy. Conduct an *F* test to test the joint null that $\rho = 1$ and $\alpha = 0$. If this null is not rejected, we conclude that y_t has a unit root. If this null is rejected, we have three possibilities: (1) $\rho \neq 1$ and $\alpha = 0$; (2) $\rho = 1$ and $\alpha \neq 0$; or (3) $\rho \neq 1$ and $\alpha \neq 0$. The first of these possibilities, as noted earlier, is unreasonable and so is ruled out. The second possibility contradicts our prior knowledge that y_t is not growing, and so it also is ruled out. We conclude that y_t is stationary.

Could this test have been conducted, as most textbooks suggest, by ignoring the issue of $\alpha = 0$ and just testing, via a *t* statistic, the null that $\rho = 1$? Yes, it could—if $\rho = 1$ is not rejected, $\alpha = 0$ would have to follow if our supposition that y_t is not growing is correct, and if $\rho = 1$ is rejected, $\alpha \neq 0$ would have to follow if an unreasonable result is to be avoided. The choice depends on the relative power of the two tests. In this case, the *t* test has higher power. When the unit-root null is false, the value of the *F* test statistic, remarkably, is unaffected by the value of the intercept.[9] Consequently, because the *t* test is one-sided, it will have higher power.

We have concluded that a *t* test of $\rho = 1$ when estimating equation (1) is the appropriate strategy to employ in this case. This *t* statistic does not, under the null of a unit root, have a *t* distribution,[10] and so special critical values are required, as provided, for example, in Hamilton (1994, 763–64).[11] These critical values were obtained via Monte Carlo studies, the results of an example of which are provided in Figure 2. Here, 15,000 artificial sets of observations of sample size 100 from the unit-root process $y_t = y_{t-1} + \varepsilon_t$ were generated with ε_t drawn from a

FIGURE 2
Sampling Distribution of *t* Statistics for Case 2

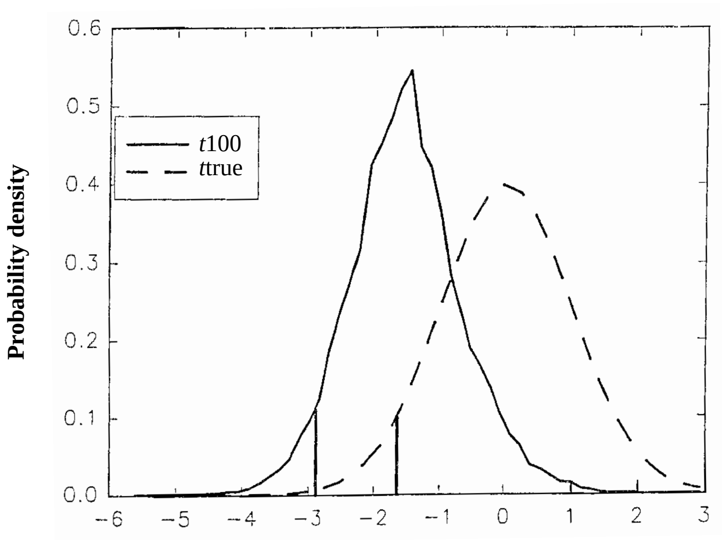

standard normal. The resulting distribution of t values is drawn as t100 in Figure 2 and is to be contrasted with the corresponding true t distribution, labeled $t$true.

Case 3: Growth Status of y_t Unknown

This case requires a sequence of tests and so is more complicated than the earlier cases that each involved only a single test. Indeed, one reason why the literature is so complicated on the issue of testing strategy is that existing strategies have been formulated to address this general case, which we believe is not typical of cases encountered in practice. In this general case, we cannot simplify by drawing on prior knowledge of y_t's growth status, but we can still simplify by ruling out unreasonable results.

The main problem here is that if the trend term is erroneously omitted, the tests are biased toward finding a unit root. This is easy to explain. If there is a trend but no trend term included in the regression, the only way the regression can capture the trend is by estimating a unit root and using the intercept (drift) to reflect the trend. On the other hand, including a trend term in the regression when it is inappropriate reduces the power of the unit root tests, for the same reason that including irrelevant explanatory variables increases variance. The strategy below is concerned more about the former problem than the latter.

Strategy. Conduct the t test for case 1 earlier.

1. Failure to reject will imply a unit root (with or without drift). Because we now “know” there is a unit root, we can test for the presence of the intercept (drift) by a simple t test of $\Delta y_t = 0$ from regressing Δy_t only on an intercept. Tra-

ditional critical values can be employed because Δy_t is stationary. The testing could stop here, as this second test has provided a structure for modeling purposes, but some researchers may wish to double-test for the unit root at this stage, to improve power.[12]

2. Rejection will imply no unit root (with or without a deterministic trend). Because unit-root tests are notoriously lacking in power (i.e., they very often tell us there is a unit root when there is no unit root), this is usually taken as firm evidence against a unit root, and so we can stop our test procedure. However, we may wish to know for modeling purposes if there is a trend. Because we now "know" there is no unit root, we can test for a deterministic trend by conducting a *t* test, using traditional critical values (because there is no unit root, traditional critical values are appropriate) of the null that $\beta = 0$ in equation (1). We do not test for $\alpha = 0$ because there is no reason to expect such an unreasonable result.

This strategy has two steps to it, which makes it more complicated than the strategies for cases 1 and 2. But it is much simpler than existing strategies in the literature, primarily because it does not double and triple test for the unit root. Advanced students may wish to investigate the logic of such multiple testing, but we believe this leads quickly to diminishing if not negative pedagogical returns and so lies outside the scope of this article.

CONCLUSION

We have provided a unit root testing strategy in a way students should find appealing. It cuts through the complications that surround testing strategies in the literature and clarifies what is going on with unit-root testing. Because we end up advocating the test procedure that most textbooks recommend (ADF *t* tests), we believe our exposition will be of value to instructors, as well as to practitioners.

Instructors should be aware, as we stressed in the introduction, that there is much more to teaching unit-root testing than the technical matter we have tried to clarify. A proper exposition of the unit-root phenomenon should include discussion of a variety of issues, such as the motivation for unit-root testing, implications of unit roots for estimation and hypothesis testing, explanation of technical details (such as the "augmentation" in ADF, selection of the ADF lag length, and why unit-root tests lack power), alternative tests for unit roots, problems caused by trend breaks, the concept of cointegration, the role of error correction models, and controversy surrounding the issue of testing for unit roots, as articulated for example in Cochrane (1991). The contribution of this article is to clarify one important dimension of the unit root phenomenon, allowing students more easily to develop the perspective on unit roots needed to ensure competent empirical work with time-series data.

NOTES

1. It should be clear that testing for a unit root ($\rho = 1$) corresponds to testing for a 0 coefficient on y_{t-1}. For expositional reasons, we have chosen to suppress the augmented part of the ADF test that requires that several lagged values of Δy_t be included as additional explanatory variables, to account for possible additional lagged values of y_t in the original specification, or to accommo-

date an autocorrelated ε_t. Critical values are the same as those used for the Dickey-Fuller (DF) test, which does not include lagged Δy_t values. Finally, we note that Ayat and Burridge (2000) also include a quadratic trend term; as is typical in the applied literature, we ignore this and other possible nonlinear trend specifications.

2. One reason why it is thought to be unrealistic for economic time series is that it implies an ever increasing (or decreasing) rate of change. An alternative way of viewing this issue is to interpret the unit-root test as testing for whether the detrended data have a unit root. Write

$$(y_t - \theta - \delta t) = \rho\,[y_{t-1} - \theta - \delta\,(t-1)] + \varepsilon_t\,,$$

where we want to test $\rho = 1$. This can be rewritten as

$$y_t = \rho\, y_{t-1} + \alpha + \beta t + \varepsilon_t\,,$$

where $\alpha = \theta\,(1-\rho) + \delta\rho$ and $\beta = \delta\,(1-\rho)$. From this, it is clear that when $\rho = 1$ (i.e., a unit root), $\beta = 0$, so there is no trend.

3. Suppose, for example, that the null $\rho = 1$ is accepted at an early stage of the testing strategy, with no constraints on α and β. At a later stage, it may be determined that $\beta = 0$, in which point the null $\rho = 1$ is tested again, constraining β to be 0. At a still later stage, it may be determined that $\alpha = 0$, prompting yet another test of $\rho = 1$, this time constraining both α and β to be 0.

4. As in note 2, if the null is true, the time trend disappears; including the irrelevant (under the null) time trend in the regression affects the distribution of the *t* statistic under the null, explaining the need for special critical values. The technical explanation of why the presence of a unit root gives rise to this phenomenon is beyond the scope of this article. Hamilton (1994, Table 17.3, 528–29) has an excellent summary of the different cases giving rise to nonstandard distributions.

5. This view is not novel. Enders (1995, 258), Hamilton (1994, 501), Stock (1994, 2829), and Wooldridge (2000, 584), among others, have clear statements to this effect.

6. Unfortunately, no textbook offers any rationale for its choice of *t* over *F* tests; this choice may have come about simply because practitioners are more comfortable with *t* than with *F* tests, or because Nelson and Plosser (1982) used *t* tests in their seminal application. Relative power depends on the extent to which β departs from 0, with larger β values increasing the *F* test power but not the *t* test power. A possible power-maximizing alternative here, implicit in some of the existing strategies, is to conduct both the *t* and the *F* tests outlined above, rejecting the null if either or both tests reject.

7. As in note 2, if the null is true, the time trend disappears; including the irrelevant (under the null) time trend in the regression affects the distribution of the *t* statistic under the null, explaining the need for special critical values. This result is also true for the distribution of the *F* statistic that the *t* statistic has replaced—special critical values are required.

8. Case 4 in Hamilton (1994, Table B6, 763) reports critical values for the *t* statistic. For large sample sizes, the 5 percent critical value is –3.41 versus –1.65 for the usual *t* test. Case 4 (Table B7 on p. 764) reports critical values for the *F* statistic. For large sample sizes, the 5 percent critical value is 6.25 versus 2.99 for the usual *F* test. Hamilton has taken these values from Fuller (1976, 373) and from Dickey and Fuller (1981, 1063). Technically, for the ADF test only the asymptotic critical values are identical to those of the DF test, but applications universally adopt the same critical values for finite sample sizes. Davidson and MacKinnon (1993, 708) argue that ". . . all finite-sample critical values for unit-root tests depend on one or another highly unrealistic assumptions about the error terms, usually that they are $N(0, \sigma^2)$. Asymptotic critical values are valid much more generally. . . ."

9. This unusual result happens because when calculating the restricted sum of squared errors the unit root eliminates any role for the intercept through first differencing. For an explanation, see Elder and Kennedy (2000).

10. If the null is true, the intercept disappears; including the irrelevant (under the null) intercept in the regression affects the distribution of the *t* statistic under the null, explaining the need for special critical values.

11. Case 2 in Hamilton (1994, Table B6, 763) reports critical values for the *t* statistic. For large sample sizes, the 5 percent critical value is –2.86 versus –1.65 for the usual *t* test. Case 2 in (Table B7, p. 764) reports critical values for the *F* statistic. For large sample sizes, the 5 percent critical value is 4.59 versus 2.99 for the usual *F* test. Hamilton has taken these values from Fuller (1976, 373) and from Dickey and Fuller (1981, 1063).

12. If this second test tells us there is no intercept, we can retest for a unit root, without including a trend in the ADF equation.

REFERENCES

Ayat, L., and P. Burridge. 2000. Unit root tests in the presence of uncertainty about the non-stochastic trend. *Journal of Econometrics* 95 (1): 71–96.

Cochrane, J. H. 1991. Comment. In *NBER macroeconomics annual,* O. J. Blanchard and S. Fischer, eds., 201–10. Cambridge, Mass.: MIT Press.

Davidson, R., and J. G. MacKinnon. 1993. *Estimation and inference in econometrics.* Oxford: Oxford University Press.

Dickey, D. A., and W. A. Fuller. 1981. Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica* 49 (4): 1057–72.

Dolado, J., T. Jenkinson, and S. Sosvilla-Rivero. 1990. Cointegration and unit roots. *Journal of Economic Surveys* 4 (3): 249–73.

Elder, J., and P. E. Kennedy. 2000. *F* versus *t* tests for unit roots. Working paper.

Enders, W. 1995. *Applied econometric time series.* New York: Wiley.

Fuller, W. A. 1976. *Introduction to statistical time series.* New York: Wiley.

Hamilton, J. D. 1994. *Time series analysis.* Princeton: Princeton University Press.

Holden, D., and R. Perman. 1994. Unit roots and cointegration for the economist. In B. B. Rao, ed., *Cointegration for the applied economist,* 47–112. New York: St. Martin's.

Nelson, C. R., and C. Plosser. 1982. Trends and random walks in macroeconomic time series. *Journal of Monetary Economics* 10 (2): 139–62.

Perron, P. 1988. Trends and random walks in macroeconomic time series. *Journal of Economic Dynamics and Control* 12 (12): 297–332.

Stock, J. M. 1994. Unit roots, structural breaks and trends. In R. F. Engle and D. L. McFadden, eds., *Handbook of econometrics,* vol. 4, 2739–841. Amsterdam: North Holland.

Wooldridge, J. 2000. *Introductory econometrics: A modern approach.* Cincinnati: South-Western.