Econometric Theory, **24**, 2008, 72–87. Printed in the United States of America.
DOI: 10.1017/S0266466608080055

TESTING FOR TREND

FABIO BUSETTI
Bank of Italy

ANDREW HARVEY
Cambridge University

The paper examines various tests for assessing whether a time series model requires a slope component. We first consider the simple *t*-test on the mean of first differences and show that it achieves high power against the alternative hypothesis of a stochastic nonstationary slope and also against a purely deterministic slope. The test may be modified, parametrically or nonparametrically, to deal with serial correlation. Using both local limiting power arguments and finite-sample Monte Carlo results, we compare the *t*-test with the nonparametric tests of Vogelsang (1998, *Econometrica* 66, 123–148) and with a modified stationarity test. Overall the *t*-test seems a good choice, particularly if it is implemented by fitting a parametric model to the data. When standardized by the square root of the sample size, the simple *t*-statistic, with no correction for serial correlation, has a limiting distribution if the slope is stochastic. We investigate whether it is a viable test for the null hypothesis of a stochastic slope and conclude that its value may be limited by an inability to reject a small deterministic slope.

1. INTRODUCTION

The question of whether a time series exhibits a trend is an important one. In other words, does the series show a persistent upward or downward movement over time that can be extrapolated into the future? If no trend is present, first differences of the series have zero mean, and a time series model will have an eventual forecast function for the level that is constant.

The main issues are best understood by considering a model that specifically contains a stochastic trend, namely,

The second author thanks the Economic and Social Research Council (ESRC) for support as part of a project on Dynamic Common Factor Models for Regional Time Series, grant L138 25 1008. Support from the Bank of Italy is also gratefully acknowledged. Earlier versions of this paper were presented at the meeting on Frontiers in Time Series held in Olbia, Italy, in June 2005 and at the NSF/NBER Time Series conference in Heidelberg, Germany, in September 2005; we are grateful to several participants for helpful comments. We also thank Peter Phillips, Robert Taylor, Jesus Gonzalo, and a number of other participants at the Unit Root and Co-integration Testing meeting in Faro, Portugal, for helpful comments. We are grateful to Paulo Rodrigues and two referees for their comments. Address correspondence to Andrew Harvey, Faculty of Economics, Cambridge University, Sidgwick Avenue, Cambridge CB3 9DD, United Kingdom; e-mail: Andrew.Harvey@econ.cam.ac.uk.

 © 2008 Cambridge University Press 0266-4666/08 $15.00

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

$$y_t = \mu_t + \upsilon_t, \qquad t = 1, \ldots, T, \tag{1}$$

$$\mu_t = \mu_{t-1} + \beta_{t-1} + \eta_t, \qquad \eta_t \sim NID(0, \sigma_\eta^2), \tag{2}$$

$$\beta_t = \beta_{t-1} + \zeta_t, \qquad \zeta_t \sim NID(0, \sigma_\zeta^2), \tag{3}$$

where μ_t is the level of the trend, β_t is the slope, υ_t is a zero mean Gaussian stationary component, and the notation $NID(0, \sigma^2)$ denotes normally and independently distributed with mean zero and variance σ^2. If both variances σ_η^2 and σ_ζ^2 are zero, the trend is deterministic. When only σ_ζ^2 is zero, the slope is fixed, and the trend reduces to a random walk with drift. Allowing σ_ζ^2 to be positive but setting σ_η^2 to zero gives an *integrated random walk* (IRW) trend, which when estimated tends to be relatively smooth. For all variants of the model, the eventual forecast function is a linear trend with slope b_T, where b_T is the estimator of β_T.

Most of the literature has focused on models with a deterministic slope, that is, $\sigma_\zeta^2 = 0$, and has investigated tests of the null hypothesis that this slope, β, is zero. If the notion of a stochastic slope is entertained, the question is whether or not to include this component, β_t, in the model. More formally the joint hypothesis to be tested is $H_0: \beta_0 = 0$ and $\sigma_\zeta^2 = 0$.

Testing may be carried out nonparametrically or by fitting a model. When an unobserved components or structural time series model (STM) is used, υ_t and the level and slope disturbances, η_t and ζ_t, respectively, are usually assumed to be mutually independent. If the model is estimated with a fixed slope, the null hypothesis that the slope is zero may be tested using its t-statistic. We will call this test $t_\beta(STM)$. If we allow for the possibility of the slope being stochastic, a t-test may be carried out by estimating a model and then setting σ_ζ^2 to zero while fixing the other parameters at their estimated values. (Because one of the variance parameters may be concentrated out of the likelihood function and reestimated, the other variances are fixed relative to it.) We will call this test $t_\beta^*(STM)$. Either way the alternative hypothesis is the general one of *permanent* slope, that is, a nonstationary component that is deterministic in a limiting case. In neither case is it necessary to assume that σ_η^2 is positive.

The aforementioned t-tests may also be carried out within an autoregressive integrated moving average (ARIMA) framework. The tests would be essentially as in Kim, Pfaffenzeller, Rayner, and Newbold (2003), except that these authors do not consider the possibility that they might be used against a stochastic trend.

Testing for trend can also be approached nonparametrically. There are a number of ways in which this may be done, but the most straightforward possibility is to base a test on the average change—growth rate if in logarithms—in the series, that is,

$$\hat{\beta} = T^{-1} \sum_{t=1}^{T} \Delta y_t = (y_T - y_0)/T, \tag{4}$$

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

where—to simplify notation—it is assumed that y_0 is also observed. The statistic can be regarded as a *contrast* between the first and last observations. The (asymptotic) distribution of $\hat{\beta}$ depends on whether the trend is deterministic, integrated of order one, denoted I(1), or I(2). If the trend is I(1), $T^{1/2}\hat{\beta}$ converges to a limiting normal distribution with variance σ_L^2, where σ_L^2 is the long-run variance of Δy_t. This suggests a nonparametric test t-statistic in which $T^{1/2}\hat{\beta}$ is divided by a consistent estimator of σ_L. Unfortunately, when the trend is deterministic, there is a problem because $\sigma_L^2 = 0$; the nonparametric tests proposed by Vogelsang (1998) do not suffer from this problem, nor does the local-to-unit root approach of Canjels and Watson (1997) to constructing a confidence interval for β. On the other hand, if the trend is I(2), $T^{-1/2}\hat{\beta}$ has a limiting normal distribution. In addition to providing part of the proof that tests such as the one based on $T^{1/2}\hat{\beta}$ are consistent against the alternative of a stochastic slope, this result also suggests the possibility of a test in which the slope is stochastic under the null.

When the trend is deterministic, $\mathrm{Var}(\hat{\beta})$ is of $O(1/T^2)$. However, least squares regression on a time trend yields an estimator of β that converges at a faster rate: it needs to be multiplied by $T^{3/2}$ to yield a limiting distribution. Hence a t-test on the slope from this regression—which may be carried out parametrically or nonparametrically depending on whether a model for v_t is specified—will be more powerful than a test based on $\hat{\beta}$.

Although the nonparametric t-test based on $T^{1/2}\hat{\beta}$ has power against a changing slope, the IRW trend model suggests an alternative approach based on signal extraction whereby the average change in the series is measured by subtracting the estimated level of the trend at the beginning from the level at the end. We can then test the significance of the change by reference to the root mean square error (RMSE) of the difference; see Visser and Molenaar (1995) and Harvey (2001). When σ_ζ^2 is zero, the deterministic trend model is obtained, and the test is the parametric t-test referred to at the end of the previous paragraph. In addition to providing a test as to whether a model should include a slope, contrasting the beginning and end of a particular series is also important for measuring how much it has actually gone up or down over the period in question.

The IRW model offers a different perspective on testing for trend by assuming that an extracted trend should be slowly changing and smooth. Although this view of the world contrasts with the fixed slope I(1) paradigm, it should be noted that a simple IRW model provides the statistical rationale for the Hodrick–Prescott filter in macroeconomics and is closely related to the cubic spline in statistics. The contrast between different ways of modeling trend can be illustrated by an example concerned with global warming. Figure 1 shows the global annual surface land and marine air anomalies with respect to the 1950–1979 average; see Parker, Folland, and Jackson (1995). A deterministic trend[1] with an AR(1) plus white noise stationary part has a slope of 0.0054, whereas a simple random walk plus drift with an additive irregular compo-

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

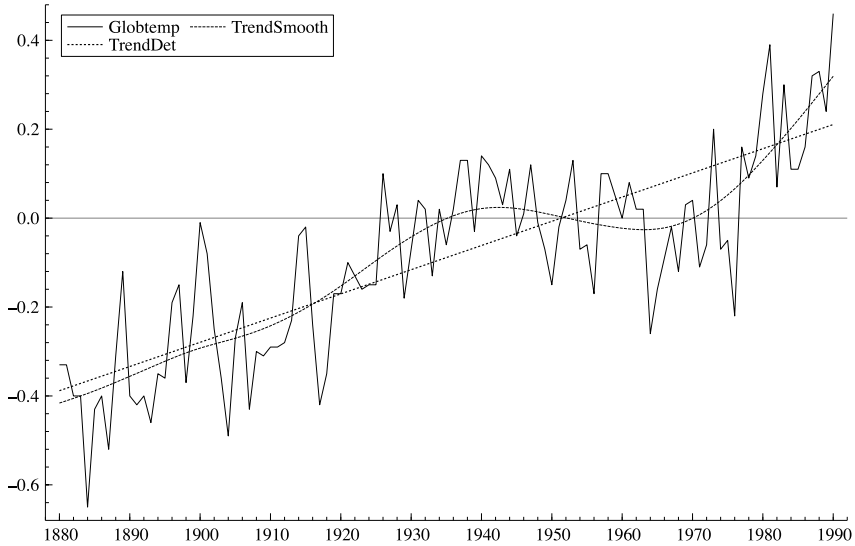


FIGURE 1. Global temperature anomaly series with integrated random walk (smooth) and deterministic trends plotted against year.

nent gives an estimated slope of 0.0065. An IRW trend fits almost as well as a random walk plus drift trend and does even better if the irregular component is replaced by an AR(1). The slope at the end of the series is 0.0215, and this contrasts sharply with the estimates for the fixed slopes in the random walk plus drift and the deterministic trend models. The implications for forecasting are clear from plots of the deterministic and IRW trends. Diagnostic tests and goodness of fit statistics offer no clear guidance as to the choice of trend: the deterministic trend model gives the best fit, but only just.

The plan of the paper is as follows. Section 2 considers the null hypothesis of no slope against the alternative of a permanent slope. We note that this hypothesis may be tested by a variant of the stationarity test where it is assumed that there is no slope under the null hypothesis. The test may be carried out parametrically using the residuals from a fitted model or nonparametrically using the approach of Kwiatkowski, Phillips, Schmidt, and Shin (1992). As such it offers an alternative to the *t*-test on the slope. Section 3 provides the limiting representation of these tests and shows their consistency. Their performance is compared with that of a number of other tests, including the $t\text{-}PS^1$ test of Vogelsang (1998). Section 4 investigates how the *t*-test fares when the trend is deterministic. Section 5 picks up on the result that when the trend is I(2), $T^{-1/2}\hat{\beta}$ has a limiting distribution. It is shown that the standardized slope, $\hat{\beta}/\hat{\sigma}$, where $\hat{\sigma}^2$ is the sample variance of the first differences, has a

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

known (nonstandard) limiting distribution and so can be used to construct a nonparametric test of the null hypothesis of a stochastic slope. Section 6 concludes.

2. TESTING AGAINST STOCHASTIC AND DETERMINISTIC SLOPE

Taking first differences in (1), and assuming for simplicity of notation that y_0 is also observed, gives

$$\Delta y_t = \beta_{t-1} + \omega_t, \qquad t = 1, 2, \ldots, T, \tag{5}$$

where $\omega_t = \eta_t + \Delta \upsilon_t$. More generally, assuming that ω_t is a zero mean weakly dependent process with long-run variance $\sigma_L^2 > 0$ allows us to obtain asymptotic representations for nonparametrically modified tests of the null hypothesis that there is no slope, that is, $\beta_t = 0$ for all t.

Consider the model given by (5) and (3). If the slope is assumed to be deterministic, that is, $\sigma_\zeta^2 = 0$, a nonparametric test of the hypothesis that it is zero can be set up in first differences using the test statistic

$$t_\beta(m) = T^{1/2}\hat{\beta}/\hat{\sigma}_L(m), \tag{6}$$

where the denominator is the square root of an estimator of the long-run variance based on m lags. Various options for the kernel and guidelines for choosing m may be found in Andrews (1991). Here we use a Bartlett window.

The test can be carried out by fitting a parametric model and testing the significance of the estimate of the fixed slope β. Either an autoregressive moving average (ARMA) model can be fitted to first differences, or a structural time series model may be estimated. The latter may have some attraction when the series is such that it is natural to include components like cycles and seasonals in a model for the levels. As noted in the introduction there is a variant in which the model is fitted with allowance made for a stochastic slope. These parametric tests are still valid if $\sigma_\eta^2 = 0$ in (2), and they have the same asymptotic representation as $t_\beta(m)$, given in Proposition 1 in Section 3.

If $\omega_t \sim NID(0, \sigma_\omega^2)$ and $\zeta_t \sim NID(0, \sigma_\zeta^2)$, the locally best invariant (LBI) test of the null hypothesis of a deterministic slope against the alternative of a nonstationary stochastic slope, that is, $H_0: \sigma_\zeta^2 = 0$ against $H_1: \sigma_\zeta^2 > 0$, is to reject for large values of

$$\zeta^1 = T^{-2}\hat{\sigma}^{-2}\sum_{t=1}^{T}\left(\sum_{i=1}^{t}(\Delta y_i - \hat{\beta})\right)^2 = T^{-2}\hat{\sigma}^{-2}\sum_{t=1}^{T}(y_t - y_0 - t\hat{\beta})^2, \tag{7}$$

where $\hat{\sigma}^2$ is the sample variance of the first differences. This is the test of Nyblom and Mäkeläinen (1983) applied to first differences; Bailey and

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

Taylor (2002) show that it is still optimal when ω_t and ζ_t are correlated. In deriving the test from the LBI principle, one initially obtains the summations running in reverse, that is, from $t = T$ to i, but as a consequence of fitting the slope it can be shown that the two statistics are identical. Asymptotically, ζ^1 has the Cramér–von Mises distribution, denoted CvM_1, under the null hypothesis.

If the preceding test statistic is formed (with reverse partial sums) without subtracting $\hat{\beta}$ it will be LBI against $H_1: \sigma_\zeta^2 > 0$ for zero initial conditions, that is, $\beta_0 = 0$. However, it is also a consistent test against a deterministic slope. Its asymptotic distribution under the null is a different member of the Cramér–von Mises family, denoted CvM_0; see Nyblom (1989). The statistics constructed with forward and reverse partial sums, ζ_F^0 and ζ_R^0, respectively, are no longer identical, but they have the same asymptotic distribution under the null hypothesis. Although ζ_R^0 is the LBI test, assuming that β_0 is zero is somewhat arbitrary, which is why ζ_F^0 is considered also.

To be strictly LBI, the statistics ζ_F^0 and ζ_R^0 should have the sample mean square rather than the sample variance in the denominator. However, the limiting distribution under the null hypothesis is the same. The local asymptotic distribution is also the same, but because dividing by the sample variance makes the test statistic bigger, the power is higher in small samples and there appears to be no adverse effect on size.

Parametric and nonparametric forms of the preceding LBI tests may be constructed. In the nonparametric case, $\hat{\sigma}_L^2(m)$ replaces $\hat{\sigma}^2$ just as in the t-test. A parametric statistic can be constructed by fitting an ARIMA or STM.

3. LIMITING REPRESENTATIONS AND LOCAL ASYMPTOTIC POWER

For a model given by (5) and (3) in which ω_t is a weakly dependent process with positive long-run variance, σ_L^2, and ζ_t need not be Gaussian, we consider the limiting behavior of the nonparametrically modified statistics $t_\beta(m)$, $\zeta_F^0(m)$, $\zeta_R^0(m)$, and $\zeta^1(m)$, where $\zeta^1(m)$ is as in (7) with $\hat{\sigma}$ replaced by $\hat{\sigma}_L(m)$, and similarly for $\zeta_F^0(m)$ and $\zeta_R^0(m)$, and $m \to \infty$ with $m^2/T \to 0$. The following proposition provides the limiting distribution of the statistics under the local alternative hypothesis

$$H_{1,T}: \beta_0 = c_d \sigma_L / \sqrt{T}, \qquad \sigma_\zeta^2 = c_s^2 \sigma_L^2 / T^2, \tag{8}$$

where c_d, c_s are fixed constants; the proof is contained in the Appendix.

PROPOSITION 1. *Consider the model (5)–(3) with ω_t being a weakly dependent process as in Stock (1994, p. 2745), with $\sigma_L^2 > 0$. Let $W_0(r)$ and $W_1(r)$ be independent standard Wiener processes for $r \in [0,1]$. Then under $H_{1,T}$,*

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

$$t_\beta(m) \xrightarrow{d} V(1;c_d,c_s), \tag{9}$$

$$\zeta_F^0(m) \xrightarrow{d} \int_0^1 V(r;c_d,c_s)^2\,dr, \tag{10}$$

$$\zeta_R^0(m) \xrightarrow{d} \int_0^1 (V(1;c_d,c_s) - V(r;c_d,c_s))^2\,dr, \tag{11}$$

$$\zeta^1(m) \xrightarrow{d} \int_0^1 V^*(r;c_s)^2\,dr, \tag{12}$$

where

$$V(r;c_d,c_s) = W_0(r) + c_d\,r + c_s\int_0^r W_1(s)\,ds,$$

$$V^*(r;c_s) = W_0(r) - rW_0(1) + c_s\int_0^r \left(W_1(s) - \int_0^1 W_1(u)\,du\right) ds.$$

Remark 1. For $c_s = 0$, the limiting distribution of $t_\beta^2(m)$ is a noncentral chi-square with one degree of freedom and noncentrality parameter equal to c_d^2, whereas that of $\zeta^1(m)$ is a standard Cramér–von Mises, CvM_1, distribution. For $c_s = c_d = 0$, $t_\beta(m)$ is asymptotically standard normal, and the limiting distributions of $\zeta_F^0(m)$ and $\zeta_R^0(m)$ are CvM_0.

Consistency of the tests follows after showing that the statistics diverge under the fixed alternative hypotheses of deterministic and stochastic slope. This is done in the following proposition; the proof is in the Appendix. Note that the test based on $\zeta^1(m)$ is not consistent against a purely deterministic slope.

PROPOSITION 2. *Under the fixed alternative hypothesis of deterministic slope, $\beta_0 \neq 0$ and $\sigma_\zeta^2 = 0$, $t_\beta(m)$ is of $O_p(T^{1/2})$, $\zeta_F^0(m)$ and $\zeta_R^0(m)$ are of $O_p(T)$, and $\zeta^1(m) \xrightarrow{d} \int_0^1 V^*(r;0)^2\,dr$ as in Proposition 1. Under the fixed alternative hypothesis of stochastic slope, $\beta_0 = 0$ and $\sigma_\zeta^2 > 0$, $t_\beta(m)$ is of $O_p((T/m)^{1/2})$, and $\zeta_F^0(m)$, $\zeta_R^0(m)$, $\zeta^1(m)$ are of $O_p(T/m)$.*

The asymptotic representations given in Proposition 1 can be used to compare the power of the tests against local deviations from the null hypothesis in the direction of deterministic and/or stochastic slope. Some results are reported in Table 1 in terms of the percentage of rejections. Specifically, we have generated 50,000 replications of the limiting random variables defined in (9)–(12) by replacing the continuous time Wiener processes W_0 and W_1 by their discrete counterparts (dividing the unit interval into 1,000 parts)

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

TABLE 1. Simulated local asymptotic power of the tests ($\times$ 100)

	c_d	0	0.5	1.0	2.0	3.0	4.0
$c_s = 0$	$t_\beta(m)$	4.9	7.7	16.7	51.8	85.2	97.9
	$\zeta_F^0(m)$	5.0	7.3	14.5	44.1	78.0	95.4
	$\zeta_R^0(m)$	4.9	7.1	14.6	44.3	77.8	95.3
	$\zeta^1(m)$	4.9	4.9	4.9	4.9	4.9	4.9
	$t\text{-}PS^1$	4.8	5.5	6.9	12.2	18.2	22.8
	$T^{-1}W$	5.3	7.3	13.7	39.2	68.9	89.4
	$t\text{-}PS^1_{FD}$	5.7	8.0	14.8	40.5	69.5	88.4
$c_s = 2.50$	$t_\beta(m)$	26.4	28.5	34.1	52.4	72.6	87.7
	$\zeta_F^0(m)$	19.2	21.3	27.4	47.5	70.2	87.2
	$\zeta_R^0(m)$	30.1	31.5	35.5	49.9	66.7	81.8
	$\zeta^1(m)$	12.7	12.7	12.7	12.7	12.7	12.7
	$t\text{-}PS^1$	6.5	6.7	7.4	10.0	13.3	17.0
	$T^{-1}W$	18.9	20.1	24.4	39.5	58.6	75.9
	$t\text{-}PS^1_{FD}$	14.9	16.7	21.3	37.4	57.8	76.5
$c_s = 5$	$t_\beta(m)$	52.4	53.1	54.8	60.5	68.6	77.3
	$\zeta_F^0(m)$	43.6	44.7	46.9	55.5	66.5	77.9
	$\zeta_R^0(m)$	56.9	57.4	58.4	62.4	68.0	74.8
	$\zeta^1(m)$	31.0	31.0	31.0	31.0	31.0	31.0
	$t\text{-}PS^1$	5.6	5.8	6.1	6.9	8.6	10.4
	$T^{-1}W$	35.0	35.5	37.2	43.6	52.7	62.9
	$t\text{-}PS^1_{FD}$	24.8	25.6	27.9	36.7	48.5	61.6
$c_s = 10$	$t_\beta(m)$	74.3	74.2	74.5	75.6	76.8	78.9
	$\zeta_F^0(m)$	70.5	70.7	71.1	72.9	75.6	79.3
	$\zeta_R^0(m)$	78.7	78.9	79.0	79.6	80.6	82.0
	$\zeta^1(m)$	61.3	61.3	61.3	61.3	61.3	61.3
	$t\text{-}PS^1$	2.7	2.6	2.8	3.0	3.5	3.9
	$T^{-1}W$	48.0	48.3	48.6	50.4	52.8	56.4
	$t\text{-}PS^1_{FD}$	33.2	33.3	34.2	37.3	42.2	48.0
$c_s = 50$	$t_\beta(m)$	94.8	94.8	94.9	94.8	94.8	94.8
	$\zeta_F^0(m)$	98.8	98.8	98.7	98.8	98.8	98.8
	$\zeta_R^0(m)$	99.1	99.1	99.1	99.2	99.1	99.1
	$\zeta^1(m)$	99.1	99.1	99.1	99.1	99.1	99.1
	$t\text{-}PS^1$	0.1	0.1	0.1	0.1	0.1	0.1
	$T^{-1}W$	56.1	56.2	56.1	56.1	56.1	56.1
	$t\text{-}PS^1_{FD}$	36.7	36.7	36.7	36.8	36.9	37.2

and computing the rejection probabilities for tests run at the 5% level of significance.

In addition we report the rejection frequencies for the tests $t\text{-}PS^1$ and $T^{-1}W$ of Vogelsang (1998) and for the *t-PS* test of Zambrano and Vogelsang (2000),

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

here denoted as $t\text{-}PS^1_{FD}$, which is similar to $t\text{-}PS^1$ but computed on the constant term of first differences. The power of these tests is computed by direct simulation of the data generating process (5)–(3) with $\omega_t \sim NID(0,1)$, for $T = 1{,}000$ under the (local) alternative hypothesis $\beta_0 = c_d/\sqrt{1{,}000}$, $\sigma_\eta = c_s/1{,}000$; given that T is large and the alternative hypothesis is scaled by the sample size these rejection frequencies can be seen as approximation of the local limiting power. Recall that the $t\text{-}PS^1$ test is asymptotically valid whether or not y_t contains a unit root, whereas $T^{-1}W$ requires a unit root in y_t. Note that although the nonparametric versions of our tests $t_\beta(m), \zeta^0_F(m), \zeta^0_R(m), \zeta^1(m)$ require a unit root in y_t, this is not the case for the parametric tests $t_\beta(STM)$ and $t^*_\beta(STM)$ based on estimating an unobserved component model or for their analogues based on ARIMA models. The limiting power of $t_\beta(STM)$, $t^*_\beta(STM)$, and their ARIMA analogues is the same as that of $t_\beta(m)$.

As expected the $t_\beta(m)$ test is most powerful against a deterministic slope. For example, for $c_d = 2$ (and $c_s = 0$), its local asymptotic power is 0.518, as opposed to 0.441 and 0.443 for the tests based on $\zeta^0_F(m)$ and $\zeta^0_R(m)$ respectively. Note that the asymptotic power of the $\zeta^1(m)$ test against a deterministic slope is always equal to its size. The $\zeta^0_R(m)$ test achieves the highest power against a stochastic slope starting at zero, that is, $\beta_0 = c_d = 0$; indeed it corresponds to the LBI test for this case. Thus with $c_s = 5$, the power of the $\zeta^0_R(m)$ test is 0.569, whereas that of $t_\beta(m)$ is only 0.524. However, the $t_\beta(m)$ test dominates both the $\zeta^0_F(m)$ and $\zeta^1(m)$ tests, for which the powers are 0.436 and 0.310 respectively. On the other hand, the power of $\zeta^0_F(m)$ is slightly greater than that of $\zeta^0_R(m)$ when c_d is high and c_s is not too large. Of course, if β_{T+1} rather than β_0 had been assumed to be zero, the powers of $\zeta^0_F(m)$ and $\zeta^0_R(m)$ would have been interchanged. The $\zeta^1(m)$ test is invariant to c_d, and it is dominated by all the other tests except when $c_s = 50$. However, this may be useful insofar as a nonrejection by $\zeta^1(m)$ and rejection by the other tests is an indication of deterministic slope. Overall, it seems that the $t_\beta(m)$ test is the best compromise. Even when $c_d = 0$ and $c_s = 50$ its power is only a little below those of the $\zeta^0_F(m)$ and $\zeta^0_R(m)$ tests. Furthermore there may be occasions when a one-sided alternative is plausible, in which case the $t_\beta(m)$ would become even more powerful.

The $t\text{-}PS^1$ test of Vogelsang (1998) has low power against the alternative of deterministic slope and has virtually no power against a stochastic slope. The $t\text{-}PS^1_{FD}$ test is more powerful than $t\text{-}PS^1$ but less powerful than $T^{-1}W$ (except for $c_s = 0$ when the two tests are essentially equivalent). However, $T^{-1}W$ is clearly dominated by the $t_\beta(m)$ test.

4. FINITE-SAMPLE BEHAVIOR FOR A DETERMINISTIC TREND WITH HIGHLY PERSISTENT DISTURBANCES

In this section we compare the $t_\beta(m)$ test with the $t\text{-}PS^1$, $T^{-1}W$ tests of Vogelsang (1998) and with the $t\text{-}PS^1_{FD}$ of Zambrano and Vogelsang (2000) when the

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

data generating process consists of a trend with a deterministic slope plus a highly persistent AR(1) process, that is,

$$y_t = \alpha + \beta_T t + v_t, \qquad t = 1, 2, \ldots, T, \tag{13}$$

$$v_t = \rho \varepsilon_{t-1} + \eta_t, \qquad \eta_t \sim NID(0,1), \tag{14}$$

$$\beta_T = c_d / \sqrt{T}, \tag{15}$$

where $c_d = 0, 0.5, 1, 1, 1.5, 2, 3$, $\rho = 1, 0.975, 0.95, 0.90$, and $T = 100$. When $\rho = 1$ the series Δy_t is I(0), and it is overdifferenced when $\rho < 1$. Overdifferencing implies that the long-run variance of Δy_t is equal to zero but $t_\beta(m)$ will still use the estimator $\hat{\sigma}_L^2(m)$: the question is whether the test might continue to work reasonably well in finite samples when ρ is close to one.

The preceding model is as in Table IV of Vogelsang (1998, p. 141), except that—to save space—we report results only for a pure AR(1) process instead of a more general ARMA(1,1). It is also the data generating process for the simulation results of Kim et al. (2003, Table II, p. 543). Here we essentially replicate the experiments of Vogelsang (1998) and Kim et al. (2003), adding the rejection frequencies for the $t_\beta(m)$ and the $t\text{-}PS_{FD}^1$ tests to the results already reported in those papers. Additionally, as a benchmark, we report the rejection frequencies of the infeasible generalized least squares (GLS) t-test (denoted t_{GLS}^*), obtained by computing the t-statistic on the trend coefficient after applying the Cochrane–Orcutt transformation to equation (14) taking ρ as known. Table 2 contains the simulated percentage rejections over 10,000 Monte Carlo replications. The random number generator of the matrix programming language Ox 2.20 was used.

For the case of $\rho = 1$ we obtain rejection frequencies similar to the case of $c_s = 0$ in Table 2, confirming that the local limiting power analysis of the previous section provides a good approximation in finite samples. Notice that in this case $m = 0$ would be the correct choice for the long-run variance estimator (and that $t_\beta(0)$ coincides with t_{GLS}^*); $m = 4, 8$ renders the t-test somewhat oversized.

When $\rho < 1$ all tests but the infeasible GLS are conservative, in the sense that the actual size is well below the nominal 5%. As c_d increases the power becomes nonnegligible. It is interesting to notice that the rejection frequencies of $t_\beta(m)$ increase with m, because adding extra lags tends to reduce the estimate of the long-run variance, a reflection of the overdifferenced nature of Δy_t. Despite being asymptotically valid, the $t\text{-}PS^1$ test still appears less preferable than $t_\beta(m)$ when there is high persistence in the data ($\rho = 0.975$ and 0.95); in these cases its power is very low for $c_d \geq 2$ (the $T^{-1}W$ and the $t\text{-}PS_{FD}^1$ tests display more power but not as high as $t_\beta(m)$).[2] The $t\text{-}PS^1$ test is more attractive when $\rho = 0.90$, because it has nonnegligible power for very small values of c_d, although it also displays a large probability of not rejecting the null hypothesis when $c_d > 2$. Figure 2 provides the graphs for the

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

TABLE 2. Simulated finite-sample power of tests against β_T ($\times$ 100)

ρ	c_d	GLS^*	$t_\beta(0)$	$t_\beta(4)$	$t_\beta(8)$	$t\text{-}PS^1$	$T^{-1}W$	$t\text{-}PS^1_{FD}$
1	0.0	5.3	5.3	6.4	8.0	5.3	5.6	5.4
	0.5	8.3	8.3	9.5	11.1	5.8	7.7	7.4
	1.0	16.8	16.8	18.6	21.0	8.0	13.9	14.0
	1.5	31.6	31.6	33.6	35.4	10.6	24.6	24.6
	2.0	50.5	50.5	52.3	54.5	13.7	38.3	38.0
	3.0	84.9	84.9	85.4	86.2	19.0	69.1	67.5
0.975	0.0	5.0	0.0	0.0	0.0	2.2	0.1	0.0
	0.5	6.5	0.1	0.3	0.7	3.8	0.7	0.6
	1.0	11.7	2.1	4.0	6.5	8.3	5.3	4.6
	1.5	20.0	16.1	22.1	27.3	13.2	18.1	17.5
	2.0	30.7	51.7	57.6	62.7	17.5	40.2	40.0
	3.0	57.5	98.1	98.1	98.0	24.0	81.1	81.1
0.95	0.0	5.0	0.0	0.0	0.0	1.7	0.0	0.0
	0.5	11.7	0.0	0.0	0.1	4.9	0.1	0.1
	1.0	30.7	0.3	1.4	3.8	11.9	3.5	3.2
	1.5	57.5	9.1	19.1	27.9	18.9	19.6	19.9
	2.0	81.4	51.1	63.1	70.9	24.6	48.2	52.4
	3.0	99.1	99.7	99.5	99.5	33.5	90.5	93.1
0.90	0.0	5.0	0.0	0.0	0.0	1.7	0.0	0.0
	0.5	30.7	0.0	0.0	0.0	10.1	0.0	0.0
	1.0	81.4	0.0	0.5	3.1	25.2	2.7	3.4
	1.5	99.1	3.7	18.8	36.3	37.6	27.0	35.0
	2.0	100.0	48.4	71.9	84.1	46.4	69.7	79.0
	3.0	100.0	100.0	100.0	99.9	57.3	99.0	99.7

empirical power functions of the tests $t_\beta(4)$, $T^{-1}W$, $t\text{-}PS^1$, and $t\text{-}PS^1_{FD}$ for $\rho =$ 1, 0.975, 0.95, 0.9.

The figures for $T^{-1}W$, $t\text{-}PS^1$ reported in Table 2 are also very similar to those of Table IV of Vogelsang (1998). Vogelsang additionally reports the rejection frequencies for the feasible GLS (Bonferroni) test of Canjels and Watson (1997) based on a local-to-unit-root approximation. His results show that in nearly all cases the Bonferroni GLS test is dominated by either $T^{-1}W$ or $t\text{-}PS^1$. Vogelsang then graphs (cf. Figs. 2–5, pp. 138–139) the asymptotic power functions of $t\text{-}PS^1$ and $T^{-1}W$ based on the sequence of local alternatives (15) for a local-to-unit-root data generating mechanism. The power displayed in those graphs appears much greater than its finite-sample approximation reported in his Table IV and in our Table 2 and our Figure 2. We believe that his graphs refer to one-sided tests, as opposed to two-sided tests, and this would explain the apparent contradiction.[3]

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

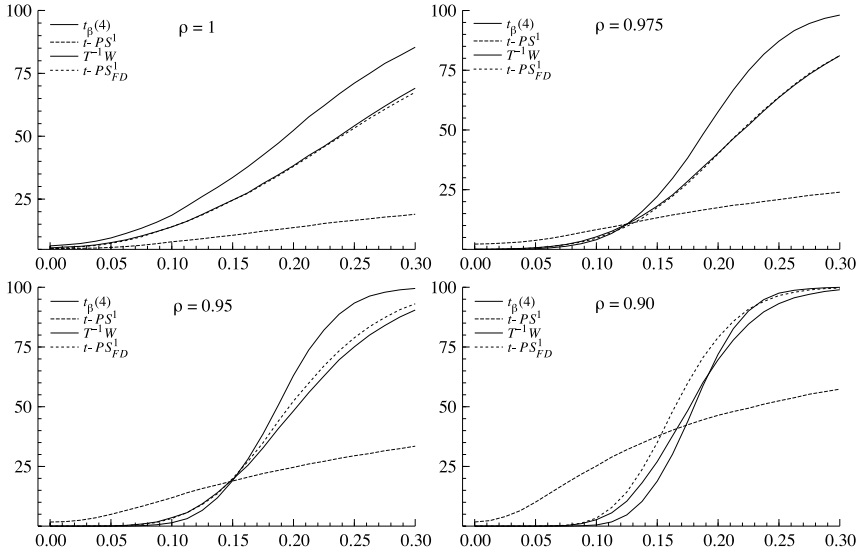


FIGURE 2. Empirical power functions of tests plotted against β_T. The data generating process is a deterministic linear trend model with $AR(1)$ errors (from $NID(0,1)$ innovations). The sample size is $T = 100$; β_T is the magnitude of the deterministic slope.

Table II of Kim et al. (2003) contains useful additional results for the data generating process (13)–(15). In particular the rejection frequencies from fitting an ARIMA(1,1,1) model, which in our notation corresponds to the parametric $t_\beta(ARIMA)$ test, are reported, showing that the parametric test clearly dominates. It is never oversized for $\rho < 1$, and it generally displays greater power. For example, for $\rho = 0.90$ and $c_d = 1$, the power of the parametric test equals 31% against 25.2% of $t\text{-}PS^1$. However further from the alternative, e.g., for $c_d = 2$, the power properties of $t_\beta(ARIMA)$ are largely comparable with those of $t_\beta(m)$ reported in our Table 2. The $t_\beta(ARIMA)$ test is oversized for $\rho = 1$, the empirical size being 11% for a sample of 100 observations, as a result of the fact that the ARIMA(1,1,1) structure is now overparametrized. Kim et al. (2003) also confirm previous results of Sun and Pantula (1999) on the degree of oversizing of the test obtained from estimating the correctly specified linear trend plus AR(1) model in finite samples: the empirical size of a nominal 5% test would be equal to 27% and 19% for ρ equal to 0.95 and 0.90, respectively.

In summary, the $t_\beta(m)$ test is conservative when the trend is deterministic, and this has implications for its power. Nevertheless it still has good power properties when the data are highly persistent, and in such circumstances it may be preferable to tests specifically designed to be robust to whether the series are I(0) or I(1).

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

5. PERMANENT SLOPE AS THE NULL

When the slope is stochastic, that is, σ_ζ^2 in (3) is positive, the standardized slope, $\hat{\beta}^* = \hat{\beta}/\hat{\sigma}_\omega$, where $\hat{\sigma}_\omega^2 = T^{-1}\sum_{t=1}^{T}(\Delta y_t - \hat{\beta})^2$, has an asymptotic distribution given by

$$\hat{\beta}^* = \hat{\beta}/\hat{\sigma}_\omega = t_\beta(0)/\sqrt{T} \xrightarrow{d} \left[\int_0^1 \underline{W}(r)^2\,dr\right]^{-1/2} \int_0^1 W(r)\,dr, \tag{16}$$

where $\underline{W}(r) \equiv W(r) - \int_0^1 W(r)\,dr$. No correction for serial correlation is needed. Note that the limiting distribution of $T^{-1/2}\hat{\beta}$ is $N(0, \frac{1}{3})$.

The preceding result suggests the setting up of a test of the null hypothesis of stochastic slope, that is, $\sigma_\zeta^2 > 0$, by rejecting for small values of $|\hat{\beta}^*|$. The 1%, 5%, 10% critical values, based on the asymptotic distribution, are 0.024, 0.118, and 0.239, respectively. The proof of (16) follows as a special case of Theorem 2 of Vogelsang (1998), though the motivation for his test, and the way it is used, is completely different. Our test derives from a proposal made by Bierens (2001) in the context of testing nonstationary cycles. Here the test can be regarded as a test (in differences) at frequency zero. However, we have made a slight modification in that Bierens constructs the denominator without subtracting the mean. In the present context this leads to the statistic $\hat{\beta}^\dagger = \hat{\beta}/(\Sigma(\Delta y_t)^2/T)^{1/2}$. The asymptotic distribution then has $W(r)$ in the denominator in addition to in the numerator. This makes virtually no difference to the 5% critical value, which is the same as before to three decimal places. Because the statistic is smaller than the one based on $\hat{\beta}^*$ it might be thought more likely to reject. However, as will be seen shortly, this appears to make virtually no difference in practice.

What if the slope is purely deterministic? Then

$$p\lim \hat{\beta}^* = \beta/\sigma_\omega \quad \text{and} \quad p\lim \hat{\beta}^\dagger = \beta/\sqrt{\sigma_\omega^2 + \beta^2}. \tag{17}$$

In both cases the null is unlikely to be rejected unless the size of the deterministic slope is small relative to σ_ω^2. Specifically, at the 5% level of significance, the null is rejected by $\hat{\beta}^*$ with probability one as $T \to \infty$ only if $|\beta| < 0.118\sigma_\omega$.

To evaluate the properties of the $\hat{\beta}^*$ and $\hat{\beta}^\dagger$ tests, a series of Monte Carlo experiments were carried out for the model (5)–(3) with $\omega_t \sim NID(0, \sigma_\omega^2)$ and $\beta_0 = 0$. Table 3 shows probabilities of rejection at the 5% level of significance, estimated with 10,000 replications, over different values of $q^{1/2} = \sigma_\zeta/\sigma_\omega$ for $T = 100$. Results for the augmented Dickey–Fuller (ADF) test with m lags, denoted $ADF_0(m)$, are given so as to provide a benchmark; $ADF_1(m)$ indicates the inclusion of a constant. In practice, small values of q are most likely to arise, and so the case of $q = 0.01$ ($q^{1/2} = 0.1$) is of particular importance. When q is small, the size of the ADF test is well above the nominal 5%. The reason for this is well known—the reduced form of second differences contains a moving average root close to the unit circle, and hence the autoregressive approximation is poor. On the other hand the $\hat{\beta}^*$ and $\hat{\beta}^\dagger$ tests do rather well in that the rejection probability is 0.17 for $q = 0.01$, whereas when $q = 0$ it shoots up to 0.76.

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

Table 3. Percentage rejections for a random walk with stochastic drift data generating process, $T = 100$

$q^{1/2}$	0	0.1	0.25	0.5	1
$\hat{\beta}^{\dagger}$	76.1	16.8	8.0	5.8	5.0
$\hat{\beta}^{*}$	75.9	16.7	8.0	5.8	4.9
$ADF_0(5)$	100.0	71.6	24.7	8.8	5.3
$ADF_1(5)$	99.7	77.8	26.7	9.4	6.1

We also carried out simulations for $\beta_0 = 0.1$, 0.2, and 0.5. For nonzero q the rejection probabilities of the standardized slope tests and $ADF_0(m)$ are changed very little; $ADF_1(m)$ is unaffected anyway. (Indeed it can be shown that the local asymptotic distributions—as in Phillips and Perron, 1988—are independent of β_0.) When $q = 0$, there is a sharp change from $\beta_0 = 0.1$, where the probability of rejection is quite high, to $\beta_0 = 0.2$, where it is low; see Busetti and Harvey (2002) for details. This is exactly what one would expect given the probability limit in (17).

6. CONCLUSION

The t_β test is designed to test against a deterministic slope, but it is also consistent against the alternative hypothesis of a nonstationary stochastic slope. Overall, it seems to be the best option for testing the null hypothesis of no slope against the alternative of permanent slope. The parametric t-tests, which can be carried out in an unobserved components or ARIMA framework, do not need the series to be integrated of order one under the null hypothesis. The nonparametric test $t\text{-}PS^1$ proposed by Vogelsang (1998) also gets around this problem, but a price is paid in terms of power. Furthermore we find that when there is a highly persistent stationary component attached to a deterministic trend, the nonparametric t_β-test suffers little size distortion and in general is still more powerful than $t\text{-}PS^1$.

The standardized slope provides a simple nonparametric test of the null hypothesis of permanent slope. Unfortunately, a question mark hangs over the test because of its inability to reject a small deterministic slope.

Finally we note that most of the literature on testing for trend in econometrics has focused on deterministic trends and series that are stationary in first differences. Yet the assumption that underlying slopes, or growth rates, are constant is a strong one and is often implausible. This point is brought home by the global warming example. Using more recent data in Jones, New, Parker, Martin, and Rigor (1999), we find that an IRW trend, with an AR(1) disturbance, provides excellent unconditional forecasts over a postsample period from 1991 to 2004. On the other hand a deterministic trend model with AR(1) and white noise sta-

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

tionary components[4] grossly underpredicts. A model in which the AR(1) is replaced by a random walk also underpredicts but not by quite so much.[5]

NOTES

1. All the models were fitted using the STAMP package of Koopman, Harvey, Doornik, and Shephard (2000).

2. For $\rho = 0.975$ and $c_d \geq 2$ the $t_\beta(m)$ test rejects more frequently than the infeasible GLS test, but this is a finite-sample effect.

3. Bunzel and Vogelsang (2005) recently proposed a modification of the tests of Vogelsang (1998) based on the fixed-bandwidth asymptotic framework. The power improvement of the new test appears to be of the order of 0.05–0.1 for $\rho = 0.90$ (cf. their Fig. 7, p. 390); we thank a referee for pointing out this reference.

4. The AR(1) parameter is 0.89 in the deterministic trend model but only 0.33 in the IRW model. The inclusion of the AR(1) component in the IRW model gives a better fit, but it makes little difference to predictions.

5. The sum of absolute forecast errors from unconditional predictions for the IRW trend model is 7.11, for the deterministic trend it is 17.84, and for the random walk plus drift it is 12.99.

REFERENCES

Andrews, D.W.K. (1991) Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica* 59, 817–858.

Bailey, R.W. & A.M.R. Taylor (2002) An optimal test against a random walk component in a non-orthogonal unobserved components model. *Econometrics Journal* 5, 520–532.

Bierens, H.J. (2001) Complex unit roots and business cycles: Are they real? *Econometric Theory* 17, 962–983.

Bunzel, H. & T.J. Vogelsang (2005) Powerful trend function tests that are robust to strong serial correlation, with an application to the Prebisch-Singer hypothesis. *Journal of Business & Economic Statistics* 23, 381–394.

Busetti, F. & A.C. Harvey (2002) Testing for Drift in a Time Series. DAE Working paper 0237, University of Cambridge.

Canjels, W. & M.W. Watson (1997) Estimating deterministic trends in the presence of serially correlated errors. *Review of Economics and Statistics* 79, 184–200.

Harvey, A.C. (2001) Trend analysis. In D.R. Brillinger (ed.), *Encyclopedia of Environmetrics,* vol. 4, pp. 2243–2257. Wiley.

Jones, P.D., M. New, D.E. Parker, S. Martin, & I.G. Rigor (1999) Surface air temperature and its changes over the past 150 years. *Reviews of Geophysics* 37, 173–199.

Kim, T., S. Pfaffenzeller, T. Rayner, & P. Newbold (2003) Testing for linear trend with application to relative primary commodity prices. *Journal of Time Series Analysis* 24, 539–552.

Koopman, S.J., A.C. Harvey, J.A. Doornik, & N. Shephard (2000) *STAMP 6.0, Structural Time Series Analyser, Modeller and Predictor.* Timberlake Consultants.

Kwiatkowski, D., P.C.B. Phillips, P. Schmidt, & Y. Shin (1992) Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics* 44, 159–178.

Nyblom, J. (1989) Testing for the constancy of parameters over time. *Journal of the American Statistical Association* 84, 223–230.

Nyblom, J. & T. Mäkeläinen (1983) Comparison of tests for the presence of random walk coefficients in a simple linear model. *Journal of the American Statistical Association* 78, 856–864.

Parker, D.E., C.K. Folland, & M. Jackson (1995) Marine surface temperature: Observed variations and data requirements. *Climatic Change* 31, 559–600.

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press

Phillips, P.C.B. & P. Perron (1988) Testing for a unit root in time series regression. *Biometrika* 75, 335–346.

Stock, J.H. (1994) Unit roots, structural breaks and trends. In R.F. Engle & D.L. McFadden (eds.), *Handbook of Econometrics*, vol. 4, pp. 2739–2840. North-Holland.

Sun, H. & S.G. Pantula (1999) Testing for trends in correlated data. *Statistics and Probability Letters* 41, 87–95.

Visser, H. & J. Molenaar (1995) Trend estimation and regression analysis in climatological time series: An application of structural time series models and the Kalman filter. *Journal of Climate* 8, 969–979.

Vogelsang, T.J. (1998) Trend function hypothesis testing in the presence of serial correlation. *Econometrica* 66, 123–148.

Zambrano, E. & T.J. Vogelsang (2000) A simple test of the law of demand for the United States. *Econometrica* 68, 1013–1022.

APPENDIX

Proof of Proposition 1. Define the independent partial sum processes $S_{0,[Tr]} = T^{-1/2}\sum_{j=1}^{[Tr]}\omega_j$, $r \in [0,1]$, where $[Tr]$ denotes the integer part of Tr, and $S_{1,[Tr]} = c_s^{-1}T^{1/2}\sum_{j=1}^{[Tr]}\zeta_j$. Under the local alternative hypothesis $H_{1,T}$ of (8), $S_{0,[Tr]}$ and $S_{1,[Tr]}$ satisfy a multivariate invariance principle such that

$$T^{-1/2}(S_{0,[Tr]}, S_{1,[Tr]}) \Rightarrow \sigma(W_0(r), W_1(r)), \qquad r \in [0,1], \tag{A.1}$$

jointly, where $W_0(r)$ and $W_1(r)$ are independent standard Wiener processes. Then

$$T^{-1/2}\sum_{t=1}^{[Tr]}\Delta y_t = T^{-1/2}\sum_{t=1}^{[Tr]}\omega_t + c_d\sigma_L[Tr]/T + T^{-1/2}\sum_{t=1}^{[Tr]}\sum_{j=1}^{t-1}\zeta_j \Rightarrow \sigma_L V(r;c_d,c_s),$$

where, for $r \in [0,1]$, $V(r;c_d,c_s) = W_0(r) + c_d r + c_s\int_0^r W_1(s)\,ds$ and $T^{1/2}\hat{\beta} \Rightarrow \sigma_L V(1;c_d,c_s)$. As $\hat{\sigma}_L^2(m) \xrightarrow{p} \sigma_L^2$ (see Stock, 1994, p. 2799), an application of the continuous mapping theorem (CMT) gives (9)–(10)–(11).

As $\Delta y_t - \hat{\beta} = (\omega_t + \beta_{t-1}) - T^{-1}\sum_{j=1}^{T}(\omega_j - \beta_{j-1})$, by (A.1) and the CMT we obtain that, under $H_{1,T}$, $T^{-1/2}\sum_{t=1}^{[Tr]}(\Delta y_t - \hat{\beta}) \Rightarrow \sigma_L V^*(r;c_s)$, where $V^*(r;c_s) = W_0(r) - rW_0(1) + c_s\int_0^r(W_1(s) - \int_0^1 W_1(u)\,du)\,ds$. Thus an application of the CMT delivers (12). Note that, because the statistic $\zeta^1(m)$ is constructed with demeaned first differences, its limiting distribution is not influenced by the presence of a (local or fixed) slope. ■

Proof of Proposition 2. Under $\beta_0 \neq 0$ and $\sigma_\zeta^2 = 0$, it is easy to see that $\hat{\beta}$ is square root consistent and that the partial sum process $\sum_{t=1}^{[Tr]}\Delta y_t$ is of $O_p(T)$. Because it still holds that $\hat{\sigma}_L^2(m) \xrightarrow{p} \sigma_L^2$, we immediately obtain that $t_\beta(m)$ is of $O_p(T^{1/2})$, whereas $\zeta_F^0(m)$ and $\zeta_R^0(m)$ are of $O_p(T)$. Then, because the limiting distribution of $\zeta^1(m)$ is unaffected by the presence of a slope, it readily follows that $\zeta^1(m) \xrightarrow{d} \int_0^1 V^*(r;0)^2\,dr$.

Under $\beta_0 = 0$ and $\sigma_\zeta^2 > 0$, the partial sum processes $\sum_{t=1}^{[Tr]}\Delta y_t$ and $\sum_{t=1}^{[Tr]}(\Delta y_t - \hat{\beta})$ are of $O_p(T^{3/2})$, whereas $\hat{\beta} = O_p(T^{1/2})$. Then, because $\hat{\sigma}_L^2(m) = O_p(Tm)$ (compare Kwiatkowski et al., 1992, p. 168), it follows that $t_\beta(m)$ is of $O_p((T/m)^{1/2})$, whereas $\zeta_F^0(m)$, $\zeta_F^0(m)$, $\zeta^1(m)$ are of $O_p(T/m)$. ■

https://doi.org/10.1017/S0266466608080055 Published online by Cambridge University Press