Variance Shifts, Structural Breaks and Stationarity Tests

Fabio Busetti
Research Department
Bank of Italy

A.M.Robert Taylor
Department of Economics
University of Birmingham

February 2002

Abstract

This paper considers the problem of testing the null hypothesis of stochastic stationarity in time series which are characterized by variance shifts at some (known or unknown) point in the sample. It is shown that existing stationarity tests can be severely biased in the presence of such shifts, either oversized or undersized with associated spurious power gains or losses respectively, depending on the values of the breakpoint parameter and on the ratio of the pre- to post-break variance. Under the assumption of a serially independent Gaussian error term with known breakdate and known variance ratio, a locally best invariant (LBI) test of the null hypothesis of stationarity in the presence of variance shifts is then derived: both the test statistic and its asymptotic null distribution depend on the breakpoint parameter and, in general, also on the variance ratio. Modifications of the LBI test are proposed for which the limiting distribution is independent of such nuisance parameters and belongs to the family of Cramér-von Mises distributions. One such modification is particularly appealing in that it is simultaneously exact invariant to variance shifts and to structural breaks in the slope and/or level of the series. Monte Carlo simulations demonstrate that the power loss from using our modified statistics in place of the LBI statistic is not big, even in the neighbourhood of the null hypothesis, and particularly for series with shifts in the slope and/or level. The tests are extended to cover the cases of weakly dependent error processes and unknown breakpoints. The implementation of the tests are illustrated using output, inflation and exchange rate data series.

Address for Correspondence: Robert Taylor, Department of Economics, University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom. Tel: +44 121 414 6657. Email: R.Taylor@bham.ac.uk

1 Introduction

Applied economists and econometricians have recently focused attention on the question of whether or not the variability in macro-economic time-series has changed over time. For example, Kim and Nelson (1999), McConnell and Perez Quiros (2000) and Koop and Potter (2000) all find empirical evidence which strongly suggests that the volatility of US GDP has declined over the past twenty years. Moreover, McConnell et al. (1999) report a decline in volatility in each of the major components of GDP, while Chauvet and Potter (2001) also report decreased volatility in both aggregate consumption and income and in aggregate employment, the last finding also confirmed by results in Warnock and Warnock (2000). Watson (1999) also finds evidence of a decrease in the variability of short-term US interest rates since 1985. Most recently, Sensier and van Dijk (2001) report that around 90 % of the 215 monthly series constituting the macroeconomic database of Stock and Watson (1999) show a significant break (at the 5 % level) in volatility over their sample period (1959 to 1996), using conventional tests for a single structural change at an unknown point in the innovation variance to an autoregressive model; see Sections 3 and 4 of Sensier and van Dijk (2001) for full details of their testing procedure. Sensier and van Dijk (2001) find strong evidence that most real variables and price variables have undergone a reduction in variability since the 1980s.

Structural breaks in the level and/or slope (trend) are also known to be a very important empirical feature of many macroeconomic data sets. Indeed, Sensier and van Dijk (2001) explicitly control for the possibility of such effects in their analysis. A striking example of a series which undergoes a simultaneous break in both level and variance is provided by the nominal exchange rate series of Thailand's Baht to the U.S. Dollar. The Bank of Thailand took the decision to float the Thai Baht, from a previously fixed exchange rate, on July 2nd 1997. One would therefore anticipate a large structural break across the two regimes. To that end, consider Figure 1 which graphs the monthly exchange rate series for the period 1991.M2-1999.M12. The data correspond to the level of the exchange rate on the first day of each month and were obtained from Datastream.

Figure 1 about here

It is quite apparent from a visual inspection of Figure 1 that the series does indeed undergo a large level shift between the observed values for July 1997 and August 1997, as

expected. However, in addition to the shift in the level of the series, and exactly as one would expect for a switch from a fixed to floating exchange rate regime, it is evident from Figure 1 that the variability in the series is also much larger under the floating regime. Indeed, the standard deviation in the floating regime is estimated to be some 12.6 times greater than that under the fixed regime.

The problem of structural breaks in the level/slope of a series has proved to be of considerable interest in the unit root testing literature. Perron (1989) demonstrated that the conventional augmented Dickey-Fuller [ADF] tests are not consistent against processes which are stochastically stationary about a broken trend. Although the ADF tests are consistent against stationary processes about a broken level, Perron (1989) shows that their power is vastly reduced relative to the case where no break occurs. Moreover, Leybourne et al. (1998) have shown that the standard ADF test is somewhat over-sized against unit root series with level shifts, in cases where the break occurs near the start of the series; that is, it shows a tendency to spuriously reject the unit root null. Perron (1989) proposed modifications of the ADF test, constructed so as to be invariant to level/slope breaks, for the case of a single breakpoint occurring at a known point in the sample. Extensions of Perron's work to allow for an unknown breakpoint have been considered by Zivot and Andrews (1992) and Banerjee et al. (1992), and for multiple breaks by, inter alia, Bai and Perron (1998).

Two recent papers have analysed the effects of a single break in the innovation variance of an integrated process on the ADF test. Hamori and Tokihisa (1997) report a small tendency for ADF tests from a regression including no deterministic terms to over-reject the unit root hypothesis when there is an increase in the innovation variance, while Kim et al. (2000) report severe over-rejections where either a constant or a linear trend is included in the ADF regression and the innovation variance decreases at a fixed point. Kim et al. (2000) then suggest modifications of the ADF tests which have pivotal limiting null distributions, regardless of the structural break in the variance. The breakpoint need not be known and these tests are also invariant to slope/level breaks (assuming a constant/trend is included in the test regressions), provided they occur at the same point as the variance shift.

Within the so-called stationarity testing literature, Busetti and Harvey (2001) [BH] have recently proposed generalisations of the tests of Nyblom and Mäkeläinen (1983) [NM] and Kwiatkowski et al. (1992) [KPSS] to allow for structural breaks in the slope and/or level of the process at a given point in time: in these tests the null hypothesis is stationarity (around

a broken deterministic level/slope) against the alternative of a unit root. While the NM and KPSS statistics asymptotically diverge for processes that are stochastically stationary about deterministic trends with structural breaks, the tests developed by BH are exact invariant to these breaks. These tests are, however, not invariant to breaks in variance. The aim of this paper is to generalise the BH approach, in the same way that Kim et al. (2000) have generalised the level/slope-break ADF tests, and develop tests which are invariant to breaks in slope/level and variance. We develop tests both for known and unknown breakpoints. Although our discussion is exemplified through the case of a single break, generalisations to multiple breaks are also discussed.

The plan for this paper is as follows. In Section 2 we outline a structural time series model which accommodates structural breaks in both level/slope and variance, generalising the model considered in, inter alia, BH, NM and KPSS. In Section 3 we review the class of locally best invariant (LBI) tests against a non-stationary level in this model for the case where the variance is fixed throughout the sample. We then analyse the limiting behaviour of these tests in cases where the variance displays a single structural break at some point in the sample. We demonstrate that this causes a shift in the limiting null distribution of the tests whose precise form depends both on the location of the break and on the ratio of the post- to pre-break variance. In Section 4 we demonstrate that where both the break date and the aforementioned variance ratio are known, then the LBI test against a unit root may be derived. A modification of this statistic, allowing for the case where the variance ratio is unknown, is shown to have a limiting null distribution which belongs to the Cramér-von Mises family. We also derive the distribution of this statistic under local alternatives. In Section 5, we further modify this test to derive a statistic which has a pivotal limiting distribution in the case where the process may display a break in any or all of its level, slope and variance at any given point in time. In practice it may well be hard to identify the source of a break in a time series and hence a test which is simultaneously robust to shifts in level, slope and variance would appear highly desirable for the practitioner. In Section 6, we consider further modifications of the above statistics for cases where the breakdate is unknown, and discuss consistent estimation of the break date.

In Section 7 we use Monte Carlo methods to compare the finite sample size and power properties of the stationarity tests of Sections 2-5 in cases where there is a variance shift and/or a structural break in the deterministic kernel. In Section 8 we present generalisations

of the model proposed in Section 2 which allow for cases where the irregular error process is weakly dependent. Moreover, the nature of the weak dependence in the irregular error process may also change at the breakpoint. Applications of the tests discussed in the paper to the output and price series from the database of Stock and Watson (1999), where any possibly breakdates cannot be assumed known, and to the Thai exchange rate series of Figure 1, where the breakdate is known, are considered in Section 9. Section 10 concludes the paper. Appendix A contains proofs of the stated Propositions, while Appendix B contains material relating to the output and prices series of Section 9.1.

2 The Model

In order to accommodate structural breaks in the slope/level of the process and/or structural breaks in the variance we consider the following structural time series model for the scalar process y_t, $t = 1, \ldots, T$,

$$y_t = d_t + \sigma_t(\mu_t + \epsilon_t), \quad \epsilon_t \sim IN(0, \sigma^2), \quad \sigma_t > 0, \tag{2.1}$$

$$\mu_t = \mu_{t-1} + \eta_t, \quad \eta_t \sim IN(0, \sigma^2\sigma_\eta^2). \tag{2.2}$$

The initial value μ_0 is assumed to be zero with no loss of generality. The deterministic kernel $d_t = \mathbf{x}_t'\beta$ where $\mathbf{x}_t$ is a $(p \times 1)$, $p < T$, fixed sequence, whose first element is fixed at unity throughout (so that (2.1) always contains an intercept term), with associated parameter vector β. We assume that ϵ_t and η_t are mutually independent. From (2.1), observe that the variance of the irregular component is allowed to display time-varying behaviour. The variance of the error driving the level also varies with the irregular variance, being a fixed multiple thereof. Generalisations of the data generating process [DGP] (2.1)-(2.2) to allow for weak dependence in the error process $\{\epsilon_t\}$ are discussed in Section 8.

For the purposes of this paper we focus on the case where σ_t^2 of (2.1)-(2.2) displays a single structural break at time $[T\tau]$, $\tau \in (0, 1)$; that is,

$$\sigma_t^2 = h_t(\tau) + k^2\overline{h}_t(\tau), \quad k > 0, \tag{2.3}$$

where $h_t(\tau) \equiv 1(t \leq [T\tau])$, $1(\cdot)$ the indicator function, and $\overline{h}_t(\tau) \equiv 1 - h_t(\tau)$. Without loss of generality, we have set $\sigma_t^2 = 1$ for all time periods prior to the break. This allows us to interpret the parameter k^2 as the ratio of the post- to pre-break variance.

If σ_t is fixed for all t, then (2.1)-(2.2) belongs to the class of models considered by Nabeya and Tanaka (1988). Within that class, NM consider $\mathbf{x}_t = 1$, Nyblom (1986) considers $\mathbf{x}_t = (1, t)'$, while BH consider, among others, the following two scenarios: $\mathbf{x}_t = (1, \overline{h}_t(\tau))'$, and $\mathbf{x}_t = (1, t, \overline{h}_t(\tau), t\overline{h}_t(\tau))'$. In the first of these scenarios the process displays a one-off break in its deterministic level at time $t = [\tau T]$, while in the second there is a structural break in both the level and the slope of the mean of the process. In both cases the break point $\tau \in (0, 1)$ is assumed known. Finally, KPSS consider both $\mathbf{x}_t = 1$ and $\mathbf{x}_t = (1, t)'$, allowing for weak dependence in $\{\epsilon_t\}$.

3 Existing Tests against a Stochastic Trend

Consider (2.1)-(2.2) with σ_t fixed. With no loss of generality, we set $\sigma_t = 1$, $t = 1, \ldots, T$. Nabeya and Tanaka (1988) demonstrate that the LBI test for the null hypothesis of stationarity against the unit root (stochastic trend) alternative, viz,

$$\mathsf{H}_0 : \sigma_\eta^2 = 0 \text{ versus} \tag{3.1}$$

$$\mathsf{H}_1 : \sigma_\eta^2 > 0. \tag{3.2}$$

is given by,

$$\mathcal{NM} = T^{-2}\hat{\sigma}^{-2} \sum_{t=1}^{T} \left(\sum_{j=1}^{t} e_j \right)^2 > \ell \tag{3.3}$$

where: e_t, $t = 1, \ldots, T$, are the Ordinary Least Squares (OLS) residuals obtained from the regression of y_t on $\mathbf{x}_t$, $t = 1, ..., T$; $\hat{\sigma}^2 = T^{-1} \sum_{t=1}^{T} e_t^2$, and ℓ is a suitably chosen constant (see Ferguson, 1967, p.235). The special cases of $d_t = \beta_0$ and $d_t = \beta_0 + \beta_1 t$ are considered in NM and Nyblom (1986) respectively, while BH consider, among others, $d_t = \beta_{0,1} h_t(\tau) + \beta_{0,2} \overline{h}_t(\tau)$ and $d_t = (\beta_{0,1} + \beta_{1,1} t) h_t(\tau) + (\beta_{0,2} + \beta_{1,2} t) \overline{h}_t(\tau)$, the broken level and broken level and slope models, respectively.

The limiting distribution of the statistic $\mathcal{NM}$ of (3.3) under H_0 of (3.1) is non-standard and depends on d_t. Although Gaussianity has been assumed in (2.1)-(2.2) thus far, all of the limiting results which follow also hold under considerably weaker, martingale difference, conditions on $\{\epsilon_t, \eta_t\}$; see Stock (1994,pp.2745,2794-2799) for details. Assuming that the vector $\mathbf{x}_t$ satisfies the following conditions (see Phillips and Xiao, 1998): there exists a scaling

matrix $\boldsymbol{\delta}_T$ and a bounded piecewise continuous function $\mathbf{x}(r)$ such that (a) $\boldsymbol{\delta}_T \mathbf{x}_{\lfloor Tr \rfloor} \to \mathbf{x}(r)$ as $T \to \infty$ uniformly in $r \in [0,1]$, and (b) $\int_0^1 \mathbf{x}(r)\mathbf{x}(r)'dr$ is positive definite, then

$$\mathcal{NM} \Rightarrow \int_0^1 \mathsf{B}_{\mathsf{x}}(r)^2 dr \tag{3.4}$$

where $\Rightarrow$ denotes weak convergence and

$$\mathsf{B}_{\mathsf{x}}(r) = \mathsf{W}_0(r) - \int_0^1 \mathbf{x}(r)' d\mathsf{W}_0(r) \left(\int_0^1 \mathbf{x}(r)\mathbf{x}(r)'dr \right)^{-1} \int_0^r \mathbf{x}(s)ds, \quad r \in [0,1], \tag{3.5}$$

$\mathsf{W}_0(r)$ a standard Brownian motion process. In the case where $\mathbf{x}_t = 1$, as in NM, $\mathsf{B}_{\mathsf{x}}(r) \equiv \mathsf{B}_0(r) = \mathsf{W}_0(r) - r\mathsf{W}_0(1)$, $r \in [0,1]$, a standard Brownian bridge process, and the right member of (3.4) is a first level Cramér-von Mises distribution with one degree of freedom, denoted $CvM_1(1)$; see Harvey (2001). Where $\mathbf{x}_t = (1,t)'$, as in Nyblom (1986), $\mathsf{B}_{\mathsf{x}}(r) \equiv \mathsf{B}_0(r) - 6r(1-r)\int_0^1 \mathsf{B}_0(s)ds$, a standard second level Brownian bridge process. Where $\mathbf{x}_t = (1, t, ..., t^{p-1})'$, $1 \leq p < \infty$, $\mathsf{B}_{\mathsf{x}}(r)$ takes the form of a standard p-th level Brownian bridge processes (see MacNeill, 1978), and the right member of (3.4) is a pth level generalised Cramér-von Mises distribution with one degree of freedom, denoted $CvM_p(1)$.

Under the fixed alternative H_1 of (3.2), the statistic $\mathcal{NM}$ of (3.3) is $O_p(T)$, regardless of the specific form of $\mathbf{x}_t$. This result is demonstrated in, inter alia, Leybourne and McCabe (1994) and KPSS. However, a more useful approximation to the finite-sample power of the LBI statistic is often obtained by considering the limiting distribution of $\mathcal{NM}$ under the local alternative

$$\mathsf{H}_c : \sigma_\eta^2 = c^2/T^2; \tag{3.6}$$

cf. Tanaka (1996, p.368). Notice that, for $c = 0$, H_c coincides with H_0 of (3.1).

Under H_c of (3.6) the limiting distribution of $\mathcal{NM}$ weakly converges to a weighted sum of separate functionals of two independent standard Brownian motion processes, the relative weightings on the two functionals determined by the non-centrality parameter, c. Provided $\mathbf{x}_t$ again satisfies the regularity conditions given above (3.4), it is easy to demonstrate that

$$\mathcal{NM} \Rightarrow \int_0^1 \left[\mathsf{B}_{\mathsf{x}}(r) + c \int_0^r \mathsf{W}_{c,\mathsf{x}}(s)ds \right]^2 dr \equiv \int_0^1 \mathsf{V}_{\mathsf{x}}(r)^2 dr, \tag{3.7}$$

where $\mathsf{B}_{\mathsf{x}}(r)$ is as defined in (3.5) and

$$\mathsf{W}_{c,\mathsf{x}}(r) \equiv \mathsf{W}_c(r) - \int_0^1 \mathsf{W}_c(r)\mathbf{x}(r)'dr \left(\int_0^1 \mathbf{x}(r)\mathbf{x}(r)'dr \right)^{-1} \mathbf{x}(r), \tag{3.8}$$

with $\mathsf{W}_c(r)$ a standard Brownian motion process independent of $\mathsf{W}_0(r)$. Where $\mathbf{x}_t = 1$, $\mathsf{V}_\mathsf{x}(r) \equiv \mathsf{V}_0(r) = \mathsf{B}_0(r) + c(\int_0^r \mathsf{W}_c(s)ds - r\int_0^1 \mathsf{W}_c(\tau)d\tau)$. Similarly, if $\mathbf{x}_t = (1,t)'$, $\mathsf{V}_\mathsf{x}(r) \equiv \mathsf{V}_0(r) - 6r(1-r)\int_0^1 \mathsf{V}_0(s)ds$. The limiting power function of $\mathcal{NM}$ under H_c of (3.6) for these two cases are graphed in Tanaka (1996, Figure 10.4, p.389).

We now consider the behaviour of the statistic $\mathcal{NM}$ of (3.3) when applied to a process y_t generated by (2.1)-(2.2) with $d_t = \mathbf{x}_t'\beta$, as before, but where σ_t^2 displays a single structural break at time $[T\tau]$, $\tau \in (0,1)$, as in (2.3). The limiting null distribution of $\mathcal{NM}$ when applied in this case is now given in Proposition 3.1.

Proposition 3.1 Let y_t be generated by (2.1)-(2.2) with σ_t^2 satisfying (2.3). Then under H_0 of (3.1)

$$\mathcal{NM} \Rightarrow \frac{\int_0^\tau \mathsf{G}_{1,\mathsf{x}}(r)^2 dr + \int_\tau^1 \mathsf{G}_{2,\mathsf{x}}(r)^2 dr}{\tau + (1-\tau)k^2}, \tag{3.9}$$

where

$$\begin{aligned}
\mathsf{G}_{1,\mathsf{x}}(r) &= \mathsf{B}_\mathsf{x}(r) + (1-k)\int_\tau^1 \mathbf{x}(r)' d\mathsf{W}_0(r) \left(\int_0^1 \mathbf{x}(r)\mathbf{x}(r)' dr\right)^{-1} \int_0^r \mathbf{x}(s)ds, \\
\mathsf{G}_{2,\mathsf{x}}(r) &= k\mathsf{B}_\mathsf{x}(r) + (1-k)\left\{\mathsf{W}_0(\tau) - \int_0^\tau \mathbf{x}(r)' d\mathsf{W}_0(r) \left(\int_0^1 \mathbf{x}(r)\mathbf{x}(r)' dr\right)^{-1} \int_0^r \mathbf{x}(s)ds\right\},
\end{aligned}$$

where $\mathsf{B}_\mathsf{x}(r)$ is as defined in (3.5), constructed from the standard Brownian motion $\mathsf{W}_0(r)$.

Remark 3.1: As an example, for the case $\mathbf{x}_t = 1$ the formula in (3.9) holds with

$$\begin{aligned}
\mathsf{G}_{1,\mathsf{x}}(r) &\equiv \mathsf{B}_0(r) + (1-k)r\left\{\mathsf{W}_0(1) - \mathsf{W}_0(\tau)\right\}, \quad r \in [0,\tau] \\
\mathsf{G}_{2,\mathsf{x}}(r) &\equiv k\mathsf{B}_0(r) + (1-k)(1-r)\mathsf{W}_0(\tau), \quad r \in (\tau,1],
\end{aligned}$$

where $\mathsf{B}_0(r) = \mathsf{W}_0(r) - r\mathsf{W}_0(1)$ is a standard Brownian bridge.

Remark 3.2: Notice that if $k^2 = 1$, so that no break actually occurs, (3.9) reduces to the right member of (3.4), appropriate to $\mathbf{x}_t = 1$. Otherwise, we see that the limiting null distribution of $\mathcal{NM}$ of (3.3) is not pivotal and depends both on the ratio of the post- to pre-break variances, k^2, and on the break fraction τ. The finite-sample dependence of the distribution of $\mathcal{NM}$ on k and τ is explored numerically in Section 7.

In the next section, under the assumption that both the break date and the ratio of the post- to pre-break variance are known, we derive the exact LBI test of H_0 of (3.1) against H_1 of (3.2). We subsequently show that we may drop the assumption that the ratio of the post-

to pre-break variance is known, by a suitable modification of the LBI statistic. The resulting (feasible) statistic is shown to have a familiar limiting distribution. We first consider the case where τ is known, as is the case with the data on the exchange rate in Thailand considered in the Introduction and in Section 9.2, but subsequently extend this to the case where τ might not be known, as will be the case for the price and output series analysed in Section 9.1.

4 The LBI statistic under a Known Variance Shift

From (2.1)-(2.2) with σ_t^2 following (2.3) it is straightforwardly seen that

$$\mathbf{y} \equiv (y_1, \ldots, y_T)' \sim N_T(\mathbf{X}\beta, \sigma^2 \mathbf{\Omega}^{1/2} \mathbf{D}(\sigma_\eta^2) \mathbf{\Omega}^{1/2}), \tag{4.1}$$

where $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_T]'$, $\mathbf{D}(\sigma_\eta^2) = \left[\sigma_\eta^2 \mathbf{V} + \mathbf{I}_T\right]$, $\mathbf{I}_T$ the $(T \times T)$ identity matrix, $\mathbf{V}$ a $T \times T$ matrix whose (i,j)th element is the minimum of i and j, $i, j = 1, \ldots, T$, and $\mathbf{\Omega}^{1/2}$ is a $(T \times T)$ diagonal matrix whose first $[T\tau]$ diagonal elements are unity and whose remaining diagonal elements are k. Pre-multiplying (4.1) by $\mathbf{P}$, the inverse of $\mathbf{\Omega}^{1/2}$, yields the transformed model

$$\mathbf{y}^* \sim N_T(\mathbf{X}^*\beta, \sigma^2 \mathbf{D}(\sigma_\eta^2)), \tag{4.2}$$

where $\mathbf{y}^* \equiv \mathbf{P}\mathbf{y}$ and $\mathbf{X}^* \equiv \mathbf{P}\mathbf{X}$. Also notice that $\mathbf{D}(0) = \mathbf{I}_T$; cf. King and Hillier (1985,p.99).

From King and Hillier (1985,p.99), it follows immediately from (4.2) that the LBI test of H_0 of (3.1) against H_1 of (3.2) is defined by the critical region:

Proposition 4.1 The LBI test of H_0 of (3.1) against H_1 of (3.1) is defined by the critical region,

$$\mathcal{S} = T^{-2}\hat{\sigma}_*^{-2} \sum_{t=1}^{T} \left(\sum_{j=t}^{T} e_j^* \right)^2 > \ell, \tag{4.3}$$

where $\mathbf{e}_* \equiv (e_1^*, \ldots, e_T^*)'$, are the OLS residuals from the regression of $\mathbf{y}^*$ on $\mathbf{X}^*$, and $\hat{\sigma}_*^2 = T^{-1}\mathbf{e}_*'\mathbf{e}_*$, and ℓ is a suitably chosen constant.

Remark 4.1: Although we constrained $\mathbf{i} \equiv (1, ..., 1)'$ to lie in the range space of $\mathbf{X}$, $\mathbf{i}$ will not, in general, lie in the range space of $\mathbf{X}^*$. Obvious exceptions occur where $k^2 = 1$, or where either $\mathbf{x}_t = (1, \overline{h}_t(\tau))'$ or $\mathbf{x}_t = (1, t, \overline{h}_t(\tau), t\overline{h}_t(\tau))'$. Consequently, while $\sum_{t=1}^{T} e_t \equiv 0$, $\sum_{t=1}^{T} e_t^*$ is not necessarily identically zero. Therefore, and in contrast to $\mathcal{NM}$ of (3.3), we

may not, in general, re-write (4.3) equivalently with the inner summation as $\sum_{j=1}^{t} e_j^*$; cf. KPSS, Appendix.

Remark 4.2: Observe that if there is no variance break, i.e. $k^2 = 1$, $\mathcal{S}$ of (4.3) and $\mathcal{NM}$ of (3.3) coincide.

The limiting distribution of the LBI statistic $\mathcal{S}$ of (4.3) under H_c of (3.6) may be established straightforwardly, given the results in Section 2, and is stated in the following Proposition. Notice that the limiting distribution under H_0 obtains on setting $c = 0$.

Proposition 4.2 Under H_c of (3.6) the LBI statistic $\mathcal{S}$ of (4.3) weakly converges to

$$\int_0^1 \left[\mathsf{B}_{\mathsf{X}^*}(1) - \mathsf{B}_{\mathsf{X}^*}(r) + c \int_r^1 \mathsf{W}_{c,\mathsf{X}^*}(s) ds \right]^2 dr, \tag{4.4}$$

where, appropriate to $\mathbf{X}^*$, $\mathsf{B}_{\mathsf{X}^*}(r)$ and $\mathsf{W}_{c,\mathsf{X}^*}(r)$ are as given in (3.5) and (3.8) respectively.

Remark 4.3: When $\mathbf{x}_t = 1$ the formula above holds with

$$\begin{aligned} \mathsf{B}_{\mathsf{X}^*}(r) &= \mathsf{W}_0(r) - \alpha_0(r;\tau,k) \left(k^{-1}\mathsf{W}_0(1) + (1-k^{-1})\mathsf{W}_0(\tau) \right), \\ \mathsf{W}_{c,\mathsf{X}^*}(r) &= \mathsf{W}_c(r) - \alpha_c(r;\tau,k) \int_0^1 \left(1 + k^{-1}\overline{h}(r;\tau) \right) \mathsf{W}_c(r) dr, \end{aligned}$$

where

$$\alpha_0(r;\tau,k) = \left(r + \tau(1-k^{-1})\overline{h}(r;\tau) \right) d(k,\tau)^{-1},$$

$$\alpha_c(r;\tau,k) = \left(1 - \left(1-k^{-1}\right)\overline{h}(r;\tau) \right) d(k,\tau)^{-1},$$

$$\overline{h}(r;\tau) = 1(r > \tau), \qquad d(k,\tau) = \tau + (1-\tau)k^{-1}.$$

Remark 4.4: For the cases $\mathbf{x}_t = (1, \overline{h}_t(\tau))'$ and $\mathbf{x}_t = (1, \overline{h}_t(\tau), t, t\overline{h}_t(\tau))'$ the asymptotic null distribution of $\mathcal{S}$ of (4.3) coincides with that given in Cases 1 and 2 of Proposition 3.1 of BH respectively, and hence is invariant to the variance ratio parameter k. Critical values from these limiting distributions are provided in BH, Tables 1(a)-1(b), p.134.

A large drawback of the statistic $\mathcal{S}$ of (4.3) is that it is only LBI where the ratio of the post- to pre-break variances, k^2, is known. Where this ratio is not known there is no LBI test. Under both H_0 of (3.1) and H_c of (3.6) k^2 may be consistently estimated, for example using $\hat{k}^2 = \tilde{\kappa}_2^2 \tilde{\kappa}_1^{-2}$, where $\tilde{\kappa}_1^2$ and $\tilde{\kappa}_2^2$ are as defined above (5.1). Hence, replacing k^2 by $\hat{k}^2$ in the construction of $\mathcal{S}$ of (4.3) would maintain the limiting distribution given by the right member of (4.4). However, with the exception of the cases discussed in Remark 4.4,

this limiting distribution depends on k^2, via the dependence on k of $\mathsf{B}_{\mathsf{X}^*}(r)$ and $\mathsf{W}_{c,\mathsf{X}^*}(r)$. Consequently, a test modified in this way would be of little practical use: to obtain critical points from the statistic's limiting null distribution one would need to know k^2. One could use critical values pertaining to the estimated value of k^2, but this seems rather unappealing.

With the above in mind, we now suggest a modified (feasible) version of $\mathcal{S}$ of (4.3), again maintaining the assumption that the breakdate, τ, is known. This modification yields a statistic whose limiting distribution is invariant to k^2 under both H_0 of (3.1) and H_c of (3.6).

Proposition 4.3 Let y_t be generated by (2.1)-(2.2) with σ_t^2 satisfying (2.3) and τ known. Define $\hat{\mathbf{e}} \equiv (\hat{e}_1, \ldots \hat{e}_T)'$ as the vector of OLS residuals from the regression of $\hat{\mathbf{y}} \equiv (y_1^*, \ldots, y_T^*)'$, where $y_t^* \equiv y_t/\widehat{\sigma}_t$, with $\widehat{\sigma}_t^2 = h_t(\tau)\sum_{t=1}^{[T\tau]} e_t^2/[T\tau] + \overline{h}_t(\tau)\sum_{t=[T\tau]+1}^{T} e_t^2/(T-[T\tau])$, on $\bar{\mathbf{X}} \equiv [\mathbf{X}_1 \mid \mathbf{X}_2]$, where $\mathbf{X}_j$ spans the range space of $\mathbf{X} \odot \mathbf{H}_j$, $j = 1, 2$, with $\mathbf{H}_1 = (\mathbf{h}, ..., \mathbf{h})$, $\mathbf{h} = (h_1(\tau), \ldots, h_T(\tau))'$and $\mathbf{H}_2 = \mathbf{i}_T\mathbf{i}_k' - \mathbf{H}_1$, $\mathbf{i}_j$ denoting a j-vector of ones, and $\odot$ the Hadamard (elementwise) product. Then under H_c of (3.6) the modified statistic

$$\begin{aligned}
\mathcal{S}^* &= T^{-2}\sum_{t=1}^{T}\left(\sum_{j=1}^{t}\hat{e}_j\right)^2 && (4.5)\\
&= T^{-2}\sum_{t=1}^{[T\tau]}\left(\sum_{j=1}^{t}\hat{e}_j\right)^2 + T^{-2}\sum_{t=[T\tau]+1}^{T}\left(\sum_{j=[T\tau]+1}^{t}\hat{e}_j\right)^2 \\
&\Rightarrow \tau^2\int_0^1 \mathsf{V}_{\mathsf{X}_1}(r)^2 dr + (1-\tau)^2\int_0^1 \mathsf{V}_{\mathsf{X}_2}(r)^2 dr, && (4.6)
\end{aligned}$$

where, appropriate to $\mathbf{X}_j$, $\mathsf{V}_{\mathsf{X}_j}(r)$, $j = 1, 2$, are independent processes each of the form given in (3.7) constructed from the mutually independent standard Brownian motions $\mathsf{W}_j(r)$ and $\mathsf{W}_{c,j}(r)$, $j = 1, 2$, respectively.

Remark 4.5: The limiting null representation in (4.6) depends on the breakpoint parameter τ. Again this is inconvenient, since it means that critical values from the limiting null distribution ($c = 0$) will vary with τ. However, a minor modification to $\mathcal{S}^*$ of (4.5) yields a statistic with a well-known limiting distribution under H_c of (3.6) which is invariant to τ:

$$\begin{aligned}
\mathcal{S}^{**} &= \tau^{-2}\left[T^{-2}\sum_{t=1}^{[T\tau]}\left(\sum_{j=1}^{t}\hat{e}_j\right)^2\right] + (1-\tau)^{-2}\left[T^{-2}\sum_{t=[T\tau]+1}^{T}\left(\sum_{j=[T\tau]+1}^{t}\hat{e}_j\right)^2\right] && (4.7)\\
&\Rightarrow \int_0^1 \left[\mathsf{V}_{\mathsf{X}_1}(r)^2 + \mathsf{V}_{\mathsf{X}_2}(r)^2\right] dr. && (4.8)
\end{aligned}$$

Remark 4.6: Given the definitions of $\mathbf{H}_j$, $j = 1, 2$, and $\widehat{\sigma}_t$ in Proposition 4.3, it is clear how $\mathcal{S}^*$ of (4.5) and $\mathcal{S}^{**}$ of (4.7) easily generalise to accommodate the case of multiple variance breaks. In such a case, the limiting distribution in (4.8) would be of the form $\int_0^1 \left[\sum_{j=1}^{g} \mathsf{V}_{\mathsf{x}_j}(r)^2\right] dr$, $g-1$ being the number of breaks and the $\mathsf{V}_{\mathsf{x}_j}(r)$, $j = 1, \ldots, g$, g independent processes each of the form given in (3.7).

Remark 4.7: If $\mathbf{x}_t = (1, t, ..., t^{p-1})'$, or indeed if $\mathbf{x}_t = (1, t, ..., t^{p-1}, \overline{h}_t(\tau), t\overline{h}_t(\tau), \ldots, t^{p-1}\overline{h}_t(\tau))'$, then the right member of (4.8) for $c = 0$ is a $CvM_p(2)$ distribution, regardless of τ. This following from the additive property of independent $CvM_p(1)$ distributions; again, see Harvey (2001). Critical points from the $CvM_1(2)$ and $CvM_2(2)$ distributions are tabulated in Nyblom and Harvey (2000).

Remark 4.8: Under H_1 of (3.2), it is straightforward to demonstrate that $\mathcal{S}^*$ of (4.5), $\mathcal{S}^{**}$ of (4.7) and $\mathcal{NM}$ of (3.3) are all consistent at rate $O_p(T)$.

5 Testing against a stochastic level when there is a break in the mean and/or variance at a known point

As discussed above, BH consider the problem of testing H_0 of (3.1) against H_1 of (3.2) in the context of (2.1)-(2.2) with $k^2 = 1$, in, among others, the following two scenarios: $\mathbf{x}_t = (1, \overline{h}_t(\tau))'$, and $\mathbf{x}_t = (1, t, \overline{h}_t(\tau), t\overline{h}_t(\tau))'$, which correspond to a structural break in the level and in both level and slope respectively, maintaining the assumption that $\tau \in (0, 1)$ is a known breakpoint.

The test formed from the statistic $\mathcal{S}^{**}$ of (4.7) is biased (both exactly and asymptotically) in cases where there is a structural break in the deterministic kernel, d_t, not accounted for by appropriate change dummies in $\mathbf{x}_t$. A simple example of this is where a change occurs in the underlying level of the process, but $\mathbf{x}_t$ contains only a constant term. This bias is attributable to the resulting inconsistent estimation of the sub-sample variances, these being formed from the residuals obtained from regressing $\mathbf{y}_t$ on $\mathbf{x}_t$.

As a variance break is likely to be accompanied by a structural change in the level and/or slope of the series, it is therefore important to have a test which is exact invariant to such

changes. This is obtained by using the following statistic

$$\mathcal{S}^{***} = \tau^{-2}\tilde{\kappa}_1^{-2}\left[T^{-2}\sum_{t=1}^{[T\tau]}\left(\sum_{j=1}^{t}\tilde{e}_j\right)^2\right] + (1-\tau)^{-2}\tilde{\kappa}_2^{-2}\left[T^{-2}\sum_{t=[T\tau]+1}^{T}\left(\sum_{j=[T\tau]+1}^{t}\tilde{e}_j\right)^2\right], \quad (5.1)$$

where $\tilde{\mathbf{e}} \equiv (\tilde{e}_1, \ldots, \tilde{e}_T)'$ are the OLS residuals from the regression of $\mathbf{y}$ on $\bar{\mathbf{X}}$, $\tilde{\kappa}_1^2 = [T\tau]^{-1}\sum_{t=1}^{[T\tau]}\tilde{e}_t^2$ and $\tilde{\kappa}_2^2 = (T-[T\tau])^{-1}\sum_{t=[T\tau]+1}^{T}\tilde{e}_t^2$. In the case of $\mathcal{S}^{***}$, the partial sums of residuals in each sub-sample are now scaled by a sub-sample variance estimator which is consistent, irrespective of whether a structural change has occurred in the deterministic kernel or not. Notice that if the range spaces of $\mathbf{X}$ and $\overline{\mathbf{X}}$ are identical, then $\mathcal{S}^{**}$ and $\mathcal{S}^{***}$ are numerically identical. It is transparently seen that the limiting distribution of $\mathcal{S}^{***}$ is as given by the right member of (4.8), and that this holds irrespective of changes in the deterministic kernel, with Remarks 4.7 and 4.8 also remaining apposite. As in Remark 4.6, $\mathcal{S}^{**}$ is easily generalised to accommodate the case of multiple breaks. Consequently, the case where, for example, a variance shift occurs at one point in the sample and a level/slope change occurs at a different point is also permitted.

If one was unsure as to whether the structural change in the series was due to any or all of a variance shift, a shift in level, or a shift in trend, then $\mathcal{S}^{***}$ of (5.1) would clearly be preferred. If structural change occurs in the variance only, one would anticipate the finite-sample power of the $\mathcal{S}^{***}$ and $\mathcal{S}^{**}$ of (4.7) statistics to be very similar, given that they differ only in the variance estimators used. Finally, we note that $\mathcal{S}^{***}$ is also exact invariant to k, since, unlike in $\mathcal{S}^{**}$, the variance estimators, $\tilde{\kappa}_1^2$ and $\tilde{\kappa}_2^2$, are always constructed from the residuals $\{\tilde{e}_t\}_{t=1}^{T}$. These issues are explored numerically in Section 7.

6 Modifications to allow for an Unknown Breakpoint

A simple way to deal with the case of unknown break-date is to use a two-stage procedure. In the first stage, the breakpoint is estimated superconsistently, in the sense that the estimator converges to the true value faster than the usual rate $T^{1/2}$. This estimate is then used in computing the statistics of the previous sections, as if it were the true (unknown) breakpoint. By virtue of the superconsistency of the breakpoint estimator, the limiting null distribution of the resultant statistics is then unaltered. Moreover, since the statistics diverge under the alternative, irrespective of the breakdate used, the two-stage tests are asymptotically valid.

If the structural break occurs in the variance only, then a superconsistent estimator of the breakpoint is given by

$$\widehat{\tau}_V = \underset{\tau}{argmin}\, T^{-1} \sum_{t=1}^{T} \widehat{\widehat{e}}_t(\tau)^2, \tag{6.1}$$

where $\left(\widehat{\widehat{e}}_1(\tau), ..., \widehat{\widehat{e}}_T(\tau)\right)$ are the OLS residuals from regressing e_t^2 on $(h_t(\tau), \overline{h}_t(\tau))$ and, in turn, $(e_1, ..., e_T)$ are the OLS residuals from regressing $\mathbf{y}$ on $\mathbf{X}$. Under H_0 of (3.1), the squared residuals e_t^2 are consistent estimators of the squared unobserved errors $\sigma_t^2 \varepsilon_t^2$, $t = 1, ..., T$. Consequently, under H_0, the objective function of (6.1) is asymptotically equivalent to the function minimized by the superconsistent estimator of Bai (1994, p. 454).

If the structural break occurs in the level of the series, irrespective of whether or not there exists a corresponding break in the variance, then the appropriate breakpoint estimator is given by

$$\widehat{\tau}_L = \underset{\tau}{argmin}\, T^{-1} \sum_{t=1}^{T} \widetilde{e}_t(\tau)^2, \tag{6.2}$$

where $(\widetilde{e}_1(\tau), ..., \widetilde{e}_T(\tau))$ are the OLS residuals from regressing on y_t on $(h_t(\tau), \overline{h}_t(\tau), \overline{\mathbf{x}}_t')'$, where $\overline{\mathbf{x}}_t$ contains the elements of $\mathbf{x}_t$ except for the constant. In this case, the estimator (6.2) is exactly that proposed in Bai (1994, 1997). Finally, if there is a shift in either or both of the level and the slope of the series, then the appropriate breakpoint estimator is given by (6.2) but with the residuals now obtained from the regression of y_t on $(h_t(\tau), \overline{h}_t(\tau), th_t(\tau), t\overline{h}_t(\tau), \overline{\overline{\mathbf{x}}}_t')'$ where $\overline{\overline{\mathbf{x}}}_t$ omits both the constant and linear trend terms from $\mathbf{x}_t$. Given that the statistic $\mathcal{S}^{**}$ of (4.7) is not invariant to structural change in the level and/or slope of the process, the breakpoint estimator $\widehat{\tau}_L$ is clearly only useful in the context of $\mathcal{S}^{***}$ of (5.1).

7 Numerical Results

In this Section we use Monte Carlo simulation methods to investigate the finite sample size and power properties of the stationarity tests considered in this paper. We generate data from the DGP (2.1)-(2.2), setting $\sigma^2 = 1$ with no loss of generality. Using the notation of previous sections, we investigate the impacts of varying: the ratio of the post- to pre-break standard deviation, k, among $k = 4, 2, 4/3, 1, 3/4, 1/2, 1/4$; the signal-to-noise ratio, σ_η, among $\sigma_\eta =$

$0, 0.025, 0.05, 0.10$, and the breakpoint location, τ, among $\tau = 0.3, 0.5, 0.7$.[1] All experiments were programmed using the random number generator of the matrix programming language Ox 2.20 of Doornik (1998), over $N = 10,000$ Monte Carlo replications. All reported results relate to tests run at the nominal 5% asymptotic level for a sample of size $T = 200$. Results for other choices of T and the nominal level are not qualitatively different from those reported but can be obtained from the authors on request.

In Table 1 we report the empirical rejection frequencies, size under H_0 of (3.1) and power under H_1 of (3.2), of the $\mathcal{NM}$ statistic of (3.3). We consider the following cases for the deterministic kernel d_t: (i) $\mathbf{x}_t = 1$, (ii) $\mathbf{x}_t = (1, t)'$, (iii) $\mathbf{x}_t = (1, \overline{h}_t(\tau))'$, (iv) $\mathbf{x}_t = (1, t, \overline{h}_t(\tau), t\overline{h}_t(\tau))'$. We label the resulting $\mathcal{NM}$ statistics $\mathcal{NM}_1$, $\mathcal{NM}_2$, $\mathcal{BH}_1$ and $\mathcal{BH}_2$, respectively.[2]. Case (i) is the statistic considered in NM, Case (ii) that of Nyblom (1986), while Cases (iii) and (iv) are proposed in BH.

We then investigate the finite sample size and power properties of the LBI test $\mathcal{S}$ of (4.3), the statistics $\mathcal{S}^{**}$ of (4.7) and $\mathcal{S}^{***}$ of (5.1), which assume knowledge of the breakpoint τ, and the proposed modifications[3] to allow for an unknown breakpoint, denoted as $\mathcal{S}^{***}(\widehat{\tau}_L)$ and $\mathcal{S}^{***}(\widehat{\tau}_V)$, where the breakpoint is estimated using either $\widehat{\tau}_L$ of (6.2) or $\widehat{\tau}_V$ of (6.1); see Section 6. In Tables 2-5 we investigate these statistics for Cases (i)-(iv) of d_t respectively. In Table 3 the DGP contains a level shift, which we set equal to twice the standard deviation of the noise, 2σ, while in Table 5 we also introduce a slope shift set equal to 0.01σ. The asymptotic critical values for the LBI test $\mathcal{S}$ for cases (i) and (iii), that depend on the variance ratio k, were obtained by simulating empirical approximations to Brownian motions with samples of size 1000 over 10000 replications, and computing the expression (4.4) for $c = 0$. The true value of k was used in these computations.[4] This is, of course, infeasible in practice and hence the results reported for $\mathcal{S}$ in cases (i) and (iii) should be viewed as nothing more than benchmarking figures. Critical values appropriate for cases (ii) and (iv) were taken from BH; cf. Remark 4.4.

[1] Other values of these design parameters were also considered but yielded the same qualitative conclusions and hence are not reported.

[2] Asymptotic 5% critical values for each breakpoint parameter for $\mathcal{BH}_1$ and $\mathcal{BH}_2$ are taken from BH, p.134.

[3] We do not report results for the corresponding modifications of $\mathcal{S}^{**}$ because of the poor performance observed for this statistic.

[4] These critical values may be obtained from the authors on request.

Tables 1 – 5 about here

Consider first the properties of the $\mathcal{NM}_1$ statistic from Table 1. Here we observe an interesting symmetry under H_0 of (3.1). The test is under-sized when either $k > 1$ and $\tau < 0.5$ or $k < 1$ and $\tau > 0.5$ but over-sized when either $k > 1$ and $\tau > 0.5$ or $k < 1$ and $\tau < 0.5$. This symmetry breaks down under H_1 of (3.2): in general, for each breakpoint τ, the smaller is the variance in the first sub-sample (the higher is k) the higher is the power. The $\mathcal{NM}_2$ statistic appears rather less affected by k and τ than is $\mathcal{NM}_1$ but is still slightly over-sized for $k \neq 1$, relative to $k = 1$. Correspondingly, power when $\sigma_\eta > 0$ is generally higher for $k \neq 1$ than for $k = 1$, excepting the case where $\tau = 0.3$ and $k > 1$.

More dramatic consequences arise for the statistics which allow for level and level and slope shifts, $\mathcal{BH}_1$ and $\mathcal{BH}_2$ respectively. Under H_0 of (3.1), both tests are severely under-sized when either $k > 1$ and $\tau > 0.5$ or $k < 1$ and $\tau < 0.5$, and over-sized where either $k < 1$ and $\tau > 0.5$ or $k > 1$ and $\tau < 0.5$, exactly the opposite pattern as was noted for the $\mathcal{NM}_1$ statistic above. Unlike the $\mathcal{NM}_1$ statistic, these symmetries in the size properties of the $\mathcal{BH}_1$ and $\mathcal{BH}_2$ statistics also translate into their power properties when $\sigma_\eta > 0$. Where the tests are under-sized, the power losses, relative to the LBI test $\mathcal{S}$ of (4.3), are considerable; cf. Tables 3 and 5. The most dramatic examples occur where $k = 1/4$, $\tau = 0.3$ and $k = 4$, $\tau = 0.7$. Overall, the $\mathcal{BH}_2$ statistic appears slightly more sensitive to k and τ than $\mathcal{BH}_1$, again the antithesis of what was observed for the $\mathcal{NM}_1$ and $\mathcal{NM}_2$ statistics.

Consider now Table 2, which reports empirical rejection frequencies for the $\mathcal{S}$, $\mathcal{S}^{**}$, $\mathcal{S}^{***}$, $\mathcal{S}^{***}(\widehat{\tau}_L)$ and $\mathcal{S}^{***}(\widehat{\tau}_V)$ statistics for case (i); that is, where $\mathbf{x}_t = 1$. The $\mathcal{S}^{**}$ and $\mathcal{S}^{***}$ statistics, which assume knowledge of τ, behave very similarly under H_0 of (3.1), both displaying little deviation from the nominal 5% level. For a given value of σ_η under H_1 of (3.2), $\mathcal{S}^{***}$ generally displays higher power than $\mathcal{S}^{**}$. The power differences are greatest when the variance in the first sub-sample is smaller than that in the second. Under H_1 the variance estimators used in constructing both $\mathcal{S}^{**}$ and $\mathcal{S}^{***}$ all diverge; cf. KPSS. The power differences then reflect the differing finite sample behaviour of these variance estimators under H_1.

Compared with $\mathcal{S}$, the theoretical LBI test[5], both $\mathcal{S}^{**}$ and $\mathcal{S}^{***}$ suffer from some power losses near the null hypothesis, and, for $\mathcal{S}^{**}$, where $k \geq 1$. However, in this case the LBI test

[5] Recall that for k = 1, $\mathcal{S}$ is precisely the $\mathcal{NM}_1$ test.

is not a practical option, since, as discussed below Remark 4.4, in order to simulate critical values from its limiting distribution we must know both the breakpoint parameter τ and the variance ratio k^2. Interestingly, for given τ and σ_η, the power of $\mathcal{S}^{***}$ is very similar to that seen for the $\mathcal{BH}_1$ statistic when $k = 1$ in Table 1. It would seem therefore that the cost of buying invariance to k, in addition to changes in the level, is negligible. Finally notice that the exact invariance of $\mathcal{S}^{***}$ to k is clearly demonstrated in Table 2.

Where the breakpoint is unknown, the $\mathcal{S}^{***}(\widehat{\tau}_V)$ statistics, which uses the appropriate breakpoint estimator $\widehat{\tau}_V$ (6.1), displays quite robust size properties, but can be slightly over-sized in the presence of large variance shifts. As would be expected, the test displays some power loss, relative to the tests which assume knowledge of τ, although this is never dramatic. Notice finally that $\mathcal{S}^{***}(\widehat{\tau}_L)$, even though constructed with the inappropriate breakpoint estimator $\widehat{\tau}_L$ of (6.2), does not perform too badly. Its power loss, relative to $\mathcal{S}^{***}(\widehat{\tau}_V)$, is sizeable for cases where k is close to unity, but more acceptable when the variance shift is larger. In the latter case the large change in the variance is more likely to be identified as a level change.

Table 3 presents the results for the case of level shift in the DGP. As anticipated in section 5, $\mathcal{S}^{**}$ is severely affected by the level shift, with very severe under-sizing seen under H_0 of (3.1) and a corresponding massive power loss under H_1 of (3.2). The power losses are most pronounced for large k. For a fixed value of τ, power is higher the smaller is k. In contrast the exact invariance of $\mathcal{S}^{***}$ to both level shifts (and to k) in this case, as demonstrated in Section 5, is confirmed comparing the results in Table 3 with those in Table 2. More importantly, $\mathcal{S}^{***}$ displays virtually identical power properties to the LBI test, even for alternatives close to the null hypothesis. Recall from Remark 4.4 that in this case the critical values of the LBI test do not depend on k, so a feasible version of the LBI test could be formed in the manner outlined below Remark 4.4. However, there seems little point, given the observed properties of $\mathcal{S}^{***}$ relative to the LBI test.

Where level shifts occur, the breakpoint estimator $\widehat{\tau}_V$ of (6.1) is inappropriate. This is reflected in the properties of the $\mathcal{S}^{***}(\widehat{\tau}_V)$ statistic which is completely unreliable. The modification of $\mathcal{S}^{***}$ which uses the appropriate breakpoint estimator, $\mathcal{S}^{***}(\widehat{\tau}_L)$, appears to work well in practice. It tends towards slight over-sizing as k increases, although this is really only of note for $k = 4$, and is a purely finite sample effect. Generally speaking $\mathcal{S}^{***}(\widehat{\tau}_L)$ does not lose too much power, relative to $\mathcal{S}^{***}$, with the disparity the smallest for large k, as

expected given its slight over-sizing in such cases.

Turning to the results in Table 4, which correspond to the case where $\mathbf{x}_t = (1, t)'$, we see that these are qualitatively very much in line with the corresponding results for $\mathbf{x}_t = 1$ in Table 2. Specifically, $\mathcal{S}^{**}$ and $\mathcal{S}^{***}$ again perform similarly to one another. However, their performance is not as good as in Table 2: size distortions, where present, tend to be marginally greater than for $\mathbf{x}_t = 1$, while power is lower. However, the latter is to be expected, given that the same result holds for the theoretical LBI test $\mathcal{S}$. The exact invariance of $\mathcal{S}^{***}$ to k is again clearly demonstrated and, for given τ and σ_η, the power of $\mathcal{S}^{***}$ is also very similar to that seen for $\mathcal{BH}_2$ when $k = 1$ in Table 1. Analogous conclusions are also drawn for the modified statistics $S^{***}(\widehat{\tau}_L)$ and $S^{***}(\widehat{\tau}_V)$, relative to Table 2.

The results reported in Table 5, where the deterministic kernel contains a break in both the level and slope, are also largely comparable to those seen in Table 3. Again, $\mathcal{S}^{***}$ is the appropriate statistic, behaving exactly as in Table 4, while $\mathcal{S}^{**}$ suffers from severe under-sizing problems with associated power losses. If the breakpoint is unknown, $\mathcal{S}^{***}(\widehat{\tau}_V)$ is again clearly unusable in practice, while $\mathcal{S}^{***}(\widehat{\tau}_L)$ again works well in practice, provided the variance shift is not too large.

Overall our simulation results show that for series with a structural break in the variance and/or the deterministic kernel at a known point in time, the statistic $\mathcal{S}^{***}$ seems to be preferred. Its performance is, in general, inferior to the theoretical LBI test only for cases (i) and (iii) of the deterministic kernel; in those cases, however, running the LBI test is not a feasible option, as its critical values depend on the variance ratio parameter k^2. Moreover, $\mathcal{S}^{***}$ has very similar finite sample power properties to the statistics proposed in BH which, although invariant to a shift in slope/level, are not invariant to k^2. Consequently, the additional benefit of (exact) invariance to k^2 seems to come at very little cost. If the breakpoint is unknown, some care must be taken in choosing its estimator. The $\mathcal{S}^{***}(\widehat{\tau}_V)$ appears to be the best option in the case of a shift in the variance only, while $\mathcal{S}^{***}(\widehat{\tau}_L)$ must definitely be chosen when there is a also shift in the slope and/or level; however, $\mathcal{S}^{***}(\widehat{\tau}_L)$ still performs reasonably well even in the former case.

8 Serial Correlation

In this section we discuss how the statistics $\mathcal{S}^{**}$ of (4.7) and $\mathcal{S}^{***}$ of (5.1), and their counterparts for the case of unknown breakpoint, may be adapted, via a non-parametric long-run variance correction, to be asymptotically pivotal when the observation error process $\{\epsilon_t\}$ is allowed to display weak dependence. Moreover, the precise nature of the dependence in the error process may also change at the breakpoint without affecting the limiting distribution of the non-parametrically modified statistics.

In the fixed variance case, $k^2 = 1$, KPSS and BH generalise (2.1)-(2.2) to the case where the observation error process $\{\epsilon_t\}$ satisfies the α -mixing conditions of Phillips and Perron (1988,p.336), with long run variance

$$\sigma_L^2 = \lim_{T\to\infty} T^{-1} E\left(\sum_{t=1}^{T} \epsilon_t\right)^2 .$$

These conditions permit $\{\epsilon_t\}$ to follow any finite-order, stable and invertible $ARMA$ process of unknown order and also allow for a degree of heterogeneity. Under the above assumptions, and H_0 of (3.1), it is straightforward to demonstrate that $(T\sigma_L)^{-2}\mathbf{e}'\mathbf{V}\mathbf{e} \Rightarrow \int_0^1 B_{\mathsf{x}}(r)^2 dr$. Consequently, $\mathcal{NM}$ of (3.3) does not have a pivotal null distribution. In order to obtain an asymptotically pivotal statistic, the unknown long run variance σ_L^2 must be replaced by a consistent estimator. KPSS therefore suggest replacing the OLS variance estimator $\hat{\sigma}^2$ used in constructing $\mathcal{NM}$ of (3.3) by the non-parametric estimator

$$\hat{\sigma}_L^2 = T^{-1}\sum_{t=1}^{T} e_t^2 + 2T^{-1}\sum_{i=1}^{m} w(i,m) \sum_{t=i+1}^{T} e_t e_{t-i}, \tag{8.1}$$

where $w(i,m) = 1 - i/(m+1)$, $i = 1,\ldots,m$, m the lag-truncation parameter. For $\mathbf{x}_t = 1$ and $\mathbf{x}_t = (1,t)'$, we denote this statistic as η_μ and η_τ, respectively; see KPSS, p.165. Similar modifications to the BH statistics are discussed in BH and denoted $\xi_1(\tau)$ and $\xi_2(\tau)$ for $\mathbf{x}_t = (1,\overline{h}_t(\tau))'$, and $\mathbf{x}_t = (1,t,\overline{h}_t(\tau),t\overline{h}_t(\tau))'$, respectively. As discussed in Stock (1994,p.2797), the rate conditions $m \to \infty$ and $m = o(T^{1/2})$ as $T \to \infty$ are sufficient to ensure that $\hat{\sigma}_L^2 \to^p \sigma_L^2$ under both H_0 of (3.1) and the local alternative H_c that the long run variance of η_t is $\sigma_L^2 c^2/T^2$.

The statistics $\mathcal{S}^{**}$ of (4.7) and $\mathcal{S}^{***}$ of (5.1) can be modified to allow for weak dependence in ϵ_t along exactly the same lines as the KPSS correction to $\mathcal{NM}$ of (3.3). This is achieved

simply by replacing the (OLS) estimators of the variances in the two sub-sample periods with the corresponding non-parametric long run variance estimators appropriate to each sub-sample. However, and unlike the KPSS modification of $\mathcal{NM}$ of (3.3), such modifications to the $\mathcal{S}^{**}$ and $\mathcal{S}^{***}$ statistics retain pivotal limiting null distributions even where the exact nature of the weak dependence in the observation error process is non-constant across the two sub-samples. The dynamics of the process $\{y_t\}$ can therefore also be allowed to change quite dramatically at the breakpoint while still retaining a pivotal Cramér-von Mises limiting null distribution.

For $\mathcal{S}^{**}$ of (4.7) this procedure involves replacing the OLS estimator $\widehat{\sigma}_t$ in (4.7) by the corresponding non-parametric estimator

$$\widehat{\sigma}^2_{L,t} = h_t(\tau)\hat{\sigma}^2_{1,L} + \overline{h}_t(\tau)\hat{\sigma}^2_{2,L}, \tag{8.2}$$

where $\hat{\sigma}^2_{1,L}$ and $\hat{\sigma}^2_{2,L}$ are obtained by applying the formula (8.1) to the OLS residuals e_t for the first and second sub-samples respectively. The statistic $\mathcal{S}^{***}$ of (5.1) may be similarly modified simply by replacing $\tilde{\kappa}^2_1$ and $\tilde{\kappa}^2_2$ in (5.1) by the corresponding long run variance estimators computed using the residuals $(\widetilde{e}_1, ..., \widetilde{e}_T)'$ of Section 5.

If the breakpoint, τ, is unknown, this can be estimated without modification of the strategy suggested in Section 6. This follows from Bai (1997), who demonstrates the superconsistency of the breakpoint estimator even under weak dependence assumptions of the type made above.

We repeated the simulation experiments of Section 7 for cases where the errors display weak dependence and where the non-parametric corrections discussed above are employed. We do not report these results here because qualitatively they neither add to nor contradict what was reported in Section 7. They merely demonstrate the well-known trade-off between size and power with respect to the lag truncation parameter, m, which has been extensively documented for this class of tests in, inter alia, KPSS (pp.169-173) and Leybourne and McCabe (1994,pp.160-161). These results are available from the authors on request.

9 Empirical Examples

9.1 Output and Inflation Series

According to the results reported in Sensier and van Dijk (2001), in the case of the set of ouput and inflation series from the monthly US database of macroeconomic series of Stock and Watson (1999), the hypothesis of fixed variability against a single structural break in the variance (at an unknown point) is rejected (at the 5 % level) for all series, and in most cases the statistical evidence points to a decrease in variability; see Sensier and van Dijk (2001, Section 4) for full details.

We now apply our battery of tests to the output and inflation series from Stock and Watson (1999). Specifically, we consider 16 series of industrial production and 17 series of inflation rates, the latter measured as the first difference of the logarithm of producer and consumer price indices, over the period 1967-1996. We identify the data by the same reference codes as given in Appendix A.2 of Stock and Watson (1999, pp.35-41), while graphs of the series are provided in Figures 3 and 4. For further details on the data see Stock and Watson (1999).

Tables 6 – 7 and Figures 3 – 4 about here

Tables 6 and 7 contain a summary of the results for industrial production and inflation rates respectively; more detailed results are contained in Tables B1-B2 of Appendix B. The outcomes of four statistics are reported: the standard stationarity test statistics η_τ and η_μ (for industrial production and inflation rates respectively) of KPSS which allow for neither level/slope nor variance breaks, the statistics $\mathcal{S}^{**}$ and $\mathcal{S}^{***}$ developed in this paper, and the statistics $\xi_1(\tau)$ and $\xi_2(\tau)$ (for industrial production and inflation rates respectively) of BH, which allow for level/slope shifts but not for the possibility of a variance change. In all cases the breakpoint was assumed unknown. In the case of the $\mathcal{S}^{***}$ and BH statistics the breakpoint was estimated using $\widehat{\tau}_L$ while for $\mathcal{S}^{**}$, $\widehat{\tau}_V$ was used. The statistics are computed across three values of the lag truncation parameter $m = m(4)$,$m(8)$ and $m(12)$, where, following KPSS, $m(x)$ is given by the formula

$$m(x) = integer\left(x\,(n/100)^{1/4}\right), \tag{9.1}$$

n being the sample size used to compute the statistic. The choice of lag truncation parameter reflects the usual trade-off between size and power; i.e., in general higher m corresponds to better size properties but lower power for the resulting test. Simulation experiments in KPSS provide some evidence for using $m(4)$ for a moderate sample size and $m(12)$ for larger samples.

Consider the results for the series of industrial production in Tables 6 and B1 first. Each of the statistics was computed for $\mathbf{x}_t = (1, t)'$, excepting the $\xi_2(\tau)$ statistic of BH where $\mathbf{x}_t = (1, t, \overline{h}_t(\tau), t\overline{h}_t(\tau))'$. For $m(12)$, the standard KPSS test rejects the null hypothesis of stationarity at the 5 % level for almost half of the series analysed. However, this inference is likely to be inappropriate since most of the series appear to have experienced some structural break during the sample period; cf. Figure 3. Estimating the breakpoints, we find some concentration in the years 1979-1981: 11 out of 16 series if $\widehat{\tau}_L$ is used and 8 series when using $\widehat{\tau}_V$. For the industrial production series, the former case seems the more appropriate, as can be seen from Figure 3.

The results in Table 6 show that, overall, the tests based on the $\mathcal{S}^{**}$ and $\mathcal{S}^{***}$ statistics provide much less evidence for a unit root in the industrial production series than do the standard KPSS tests: for $m(12)$, the null hypothesis of stationarity is rejected at 5% only for 3 series when $\mathcal{S}^{***}$ is used. Since trend-breaks seem to have occurred in many of these series, the arguments presented in BH tell us that the KPSS tests should be expected to be considerably over-sized. Moreover, in such cases $\mathcal{S}^{***}$would be the preferred statistic upon which to base our inference; cf. Section 7. The BH test, based on the same estimated breakpoint estimator as $\mathcal{S}^{***}$, but without allowing for a variance change, points towards still fewer rejections overall.

Table 6 also provides a strong body of evidence suggestive of a change from greater to lesser variability in the data; for most of the series both the square root of the estimated variance ratio, $\hat{k}$, and the estimated square root of the ratio of the long run variances, $\hat{k}^*$, the latter estimated using $m(12)$, are less than 0.75. For all but one series $\hat{k}$ is estimated to be less than unity. The one clear exception is the utilities series (IPUT) for which the breakpoint has been estimated in 1973. Since this corresponds to the first oil-price shock, one would anticipate relatively higher volatility in the series after the break. This feature, together with a trend-break at the same point, is clearly seen in Figure 3.

Tables 7 and B2 contain the corresponding results for the 17 series of inflation rates.

Each of the statistics was computed for $\mathbf{x}_t = 1$, excepting the $\xi_1(\tau)$ statistic of BH where $\mathbf{x}_t = (1, \overline{h}_t(\tau))'$. In no cases was a slope component included, since as is clear from Figure 4, none of the series displays linear trending behaviour. In this case both the KPSS and the BH tests provide strong evidence for nonstationarity, while our tests that allow for the presence of a variance shift tend to reject the null hypothesis much less frequently overall. For the series of inflation rates most of the estimated breakpoints are located in the years 1981-1982 and, as before, a decreased volatility in the series appears evident. Finally notice that our finding of decreased volatility, for both output and inflation, starting from the eighties is in line with the arguments on the efficiency of the Fed's monetary policy in the Volcker-Greenspan era, which appears to have contributed to a stabilization of the US economy; see, for example, Clarida et al. (2000).

9.2 Nominal Exchange Rate of the Thai Baht

Consider again the monthly nominal Thai Baht exchange rate series from Figure 1. Our question is whether the Baht-Dollar exchange rate is a stochastically stationary ($I(0)$) series or is a unit root ($I(1)$) series, controlling for the deterministic shift in both the level and variance of the series resulting from the known shift from a fixed to floating exchange rate on July 2nd 1997; $\tau = 0.73$. It is impossible to get much insight into this question from Figure 1. To that end, consider Figure 2. Here we have standardized the data in each sub-sample separately (by subtracting the sub-sample mean and dividing by the sub-sample standard deviation) and then joined the two standardised sub-series together at the time of the break. This gives a clearer picture of the underlying stochastic properties of the series. It is not clear from Figure 2 whether the series is $I(0)$or $I(1)$. Heuristically speaking, the $\mathcal{S}^{***}$statistic proposed in Section 5 tests the null of stationarity against the alternative of a unit root in the standardised series.

Figure 2 about here

Table 8 reports the application of the various stationarity tests discussed in this paper to the exchange rate data. The last three columns of Table 8 give critical values appropriate to each of the reported tests. Each of the tests were computed for $\mathbf{x}_t = 1$, excepting the $\xi_1(\tau)$ statistic of BH where $\mathbf{x}_t = (1, \overline{h}_t(\tau))'$. In no cases was a slope component included since, as is clear from Figure 2, the series does not display linear trending behaviour.

Table 8 about here

It is seen from Table 8 that the KPSS η_μ statistic, which ignores both the level and variance shift in the series, provides strong evidence of nonstationarity, its outcome emphatically rejecting the null of stationarity at the 1 % level, regardless of the lag truncation parameter m used in constructing the long run variance estimator,[6] $\hat{\sigma}_L^2$ of (8.1). This result is not unexpected, given that BH demonstrate that η_μ diverges under level shifts. In contrast, for all values of m, the outcome of BH's $\xi_1(\tau)$ statistic, which takes account of the level shift in the series, does not provide sufficient evidence to reject the null hypothesis of stationarity around a deterministic level shift, even at the 10% level. Again this is not unexpected, given the simulation results in Table 1 which show that, for the case of $\tau = 0.7$ and $k > 1$, this statistic is severely under-sized under the $I(0)$ null and, correspondingly, displays very low power under the $I(1)$ alternative.

The third and fourth rows of Table 8 report the statistics $\mathcal{S}^{**}$ and $\mathcal{S}^{***}$, in each case for $\mathbf{x}_t = 1$, modified via the nonparametric correction for serial correlation, as detailed in Section 8. Recall that for this case $\mathcal{S}^{***}$ is the appropriate test since the data contain a level shift occurring at a known point in time. In the case of the $\mathcal{S}^{***}$ statistic, the evidence is mixed: one can reject the $I(0)$ null hypothesis at the 1%, 5% and 10% levels for $m = 0$, $m = 1$ and $m = 2$ respectively, but cannot reject the null at the 10% level for any $m \geq 3$. Regardless of m, the outcome of the $\mathcal{S}^{**}$ statistic lies well below the 10% level critical value. However, in the latter case, recall from the simulation results reported in Table 3 that, where a level shift occurs, the $\mathcal{S}^{**}$ statistic is badly under-sized under the $I(0)$ null with an associated dramatic power loss under the $I(1)$ alternative.

Although, and unlike the output and price series considered in Section 9.1, the breakpoint is known in this example, it is also interesting to imagine that this was not the case and, hence, to consider $\mathcal{S}^{***}(\widehat{\tau}_V)$ and $\mathcal{S}^{***}(\widehat{\tau}_L)$, the modifications to allow for an unknown breakpoint, using the breakpoint estimators from Section 6. These statistics are reported in

[6] The formula (9.1) yields m(4) = 4 and m(8) = 8. However, the sample size is much smaller than for the series analysed in Section 9.1. Consequently, and according to the simulation results of KPSS, m(4) would seem the most appropriate choice for this sample size. We report outcomes for m = 0, . . . , 5, for completeness.

the final two rows of Table 8. Inference from the $\mathcal{S}^{***}(\widehat{\tau}_L)$ statistic coincides with that from $\mathcal{S}^{***}$ for $m \leq 1$. For $m = 2$ the outcome of $\mathcal{S}^{***}(\widehat{\tau}_L)$ allows rejection of the $I(1)$ null at the 5 % level while its outcome for $m = 3, 4, 5$ allows rejection at the 10% level. The estimated breakpoint from (6.2) sits between September 1997 and October 1997. The outcome of the $\mathcal{S}^{***}(\widehat{\tau}_V)$ statistic is similar to that of the $\hat{\eta}_\mu$ statistic, again as predicted by the simulation results in Table 3. In this case the (inconsistently) estimated breakpoint was found to be near the beginning of the sample.

10 Conclusions

This paper has focused on the problem of testing the null hypothesis that a series is stochastically stationary against a stochastic trend alternative, within the set-up of an unobserved components model, in cases where the variance of the series displays a structural shift. We have demonstrated that the limiting null distributions of conventional stationarity tests depend on both the breakpoint and on the ratio of the post- to pre-break variance. The extent of this dependence in finite samples was quantified via Monte Carlo experimentation. Under Gaussianity and assuming a known break and known post- to pre-break variance ratio, the exact LBI test against the stochastic trend alternative was derived. Feasible versions of the LBI test which relax these assumptions were subsequently suggested and shown to possess well-known pivotal limiting distributions. A numerical investigation into the size and power properties of the feasible statistics showed that they perform well in practice and, in one case, had the additional advantage of being exact invariant to simultaneous breaks in the level/slope of the process. Generalisations of the test statistics to accommodate weakly dependent non-Gaussian errors, whose pattern of dependence may change at the breakpoint, and the cases of an unknown and multiple breakpoints were provided. The application of the standard KPSS stationarity test, which does not take account of either the variance shift or the break in level occurring at the transition from a fixed to floating rate in the Thai exchange rate series gave a massive rejection of the null hypothesis of stationarity. Contrastingly, the stationarity tests of BH which allow for the break in level, but not the variance, gave no evidence against the null. These outcomes are consonant with the predictions from our Monte Carlo study. The tests developed in this paper, which are invariant to both a shift in the level and the variance of the process, were much more ambivalent with, on balance,

insufficient evidence to reject the null hypothesis. The tests were also applied to the ouput and price data sets of Stock and Watson (1999) where breakpoint locations could not be taken as known, a priori. We found a considerable body of evidence in favour of variance shifts in these series, in most cases with a decline in variability seen in the early 1980s, corroborating findings for the same series reported in Sensier and van Dijk (2001). As with the Thai series, the results provided by the tests proposed in this paper again provided a strong contrast with those of the conventional KPSS test.

References

[1] Bai, J., 1994, Least squares estimation of a shift in linear processes, Journal of Time Series Analysis 15, 453-472.

[2] Bai, J., 1997, Estimation of a change point in multiple regression models, Review of Economics and Statistics, 551-563.

[3] Bai, J. and P. Perron, 1998, Estimating and testing linear models with multiple structural changes, Econometrica 66, 47-78.

[4] Banerjee, A., R.L. Lumsdaine and J.H. Stock, 1992, Recursive and sequential tests of the unit-root and trend-break hypotheses: theory and international evidence, Journal of Business and Economic Statistics 10, 271-287.

[5] Busetti, F., 2002, Testing for (common) stochastic trends in the presence of structural breaks, Journal of Forecasting, forthcoming.

[6] Busetti, F. and A.C. Harvey, 2001, Testing for the presence of a random walk in series with structural breaks, Journal of Time Series Analysis 22, 127-50.

[7] Chauvet, M. and S.M. Potter, 2001, Recent changes in the US business cycle, The Manchester School, forthcoming.

[8] Clarida, R.H., J. Gali and M. Gertler 2000, Monetary policy rules and macroeconomic stability: evidence and some theory, Quarterly Journal of Economics 115, 147-80.

[9] Doornik J.A., 1998. Object-oriented matrix programming using Ox 2.0, London, Timberlake Consultants Press.

[10] Ferguson, T.S., 1967, Mathematical Statistics: A Decision Theoretic Approach, Academic Press: New York.

[11] Hamori, S. and A. Tokihisa, 1997, Testing for a unit root in the presence of a variance shift, Economics Letters 57, 245-253.

[12] Harvey, A.C., 2001, Testing in unobserved components models, Journal of Forecasting 20, 1-19.

[13] Kim, C.-J. and C.R. Nelson, 1999, Has the US economy become more stable? A Bayesian approach based on a Markov-switching model of the business cycle, Review of Economics and Statistics 81, 608-616.

[14] Kim, T.-H., S. Leybourne and P. Newbold, 2000, Unit root tests with a break in variance, Discussion Paper 00/12, Department of Economics, University of Nottingham.

[15] King, M.L. and G.H. Hillier, 1985, Locally best invariant tests of the error covariance matrix of the linear regression model, Journal of the Royal Statistical Society, (Series B) 47, 98-102.

[16] Koop, G. and S.M. Potter, 2000, Nonlinearity, structural breaks or outliers in economic time series?, in W.A. Barnett, D.F. Hendry, S. Hylleberg, T. Terasvirta, D. Tjostheim and A.H. Wurtz (eds.), Nonlinear Econometric Modelling in Time Series Analysis, Cambridge: CUP, pp. 61-78.

[17] Kwiatkowski, D., P.C.B. Phillips, P. Schmidt and 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 54, 159-178.

[18] Leybourne, S.J. and B.P.M. McCabe, 1994, A consistent test for a unit root, Journal of Business and Economic Statistics 12, 157-167.

[19] Leybourne, S.J., T.C. Mills and P. Newbold, 1998, Spurious rejections by Dickey-Fuller tests in the presence of a break under the null, Journal of Econometrics 87, 191-203.

[20] MacNeill, I.B., 1978, Properties of sequences of partial sums of polynomial regression residuals with applications to tests for change of regression at unknown times, The Annals of Statistics 6, 422-433.

[21] McConnell, M.M. and G. Perez Quiros, 2000, Output fluctuations in the United States: what has changed since the early 1980s?, American Economic Review 90, 1464-1476.

[22] McConnell, M.M., P.C. Mosser and G. Perez Quiros, 1999, A decomposition of the increased stability of GDP growth, Federal Reserve Bank of New York - Current Issues in Economics and Finance 5(13).

[23] Nabeya, S. and K. Tanaka, 1988, Asymptotic theory of a test for the constancy of regression coefficients against the random walk alternative, Annals of Statistics 16, 218-235.

[24] Nyblom, J., 1986, Testing for deterministic linear trend in time series, Journal of the American Statistical Association 81, 545-549.

[25] Nyblom, J. and A.C. Harvey, 2000, Tests of common stochastic trends, Econometric Theory 16, 176-199.

[26] Nyblom, J. and T. Mäkeläinen, 1983, Comparisons of tests for the presence of random walk coefficients in a simple linear model, Journal of the American Statistical Association 78, 856-864.

[27] Perron, P., 1989, The great crash, the oil price shock, and the unit root hypothesis, Econometrica 57, 1361-1401.

[28] Phillips, P.C.B. and P. Perron, 1988, Testing for a unit root in time series regression, Biometrika 75, 335-346.

[29] Phillips, P.C.B. and Z. Xiao, 1999, A primer on unit root testing, Journal of Economic Surveys 12, 423-470.

[30] Sensier, M. and D. van Dijk, 2001, Short-term volatility versus long-term growth: evidence in US macroeconomic time series, Econometric Institute Report EI 2001-11.

[31] Stock, J.H., 1994, Unit roots, structural breaks and trends, in R.F. Engle and D.L. McFadden (eds.), Handbook of Econometrics 4, Elsevier Science, 2739-2840.

[32] Stock, J.H. and M.W. Watson, 1999, A comparison of linear and nonlinear univariate models for forecasting macroeconomic time series, in R.F. Engle and H. White (eds.), Cointegration, Causality and Forecasting: A Festschrift in Honour of Clive W.J. Granger, Oxford: OUP, pp. 1-44.

[33] Tanaka, K., 1996, Time series analysis: nonstationary and noninvertible distribution theory, Wiley: New York.

[34] Taylor, A.M.R., 2002, Locally optimal tests against seasonal unit roots, Journal of Time Series Analysis, forthcoming.

[35] Warnock, M.V.C. and F.E. Warnock, The declining volatility of US unemployment: was Arthur Burns right?, International Finance Discussion Paper 677, Board of Governors of the Federal Reserve System.

[36] Watson, M.W., 1999, Explaining the increased variability in long-term interest rates, Federal Reserve Bank of Richmond, Economic Quarterly 85, 71-96.

[37] Zivot, E. and D.W.K. Andrews, 1992, Further evidence on the great crash, the oil-price shock, and the unit-root hypothesis, Journal of Business and Economic Statistics 10, 251-270.

Appendix A

Proof of Proposition 3.1: Consider (2.1)-(2.2) under H_0 of (3.1) with σ_t^2 satisfying (2.3). We demonstrate the proof for $\mathbf{x}_t = 1$, $t = 1, ..., T$. The generalisation to the case of arbitrary $\mathbf{x}_t$, satisfying the regularity conditions of Section 2, is straightforward but tedious. Let $\overline{\epsilon}_t = \sigma_t \epsilon_t$. Then, for $r \leq \tau$,

$$T^{-1/2} \sum_{j=1}^{[Tr]} \overline{\epsilon}_j \Rightarrow \sigma \mathsf{W}_0(r), \tag{A.1}$$

where $\mathsf{W}_0(r)$ is a standard Brownian motion process, and for $r > \tau$

$$T^{-1/2} \sum_{j=1}^{[Tr]} \overline{\epsilon}_j \quad \equiv \quad T^{-1/2} \sum_{j=1}^{[Tr]} \epsilon_j + T^{-1/2}(k-1) \left(\sum_{j=[T\tau]+1}^{[Tr]} \epsilon_j \right)$$

$$\begin{aligned} &\Rightarrow \quad \sigma \mathsf{W}_0(r) + \sigma(k-1)\left[\mathsf{W}_0(r) - \mathsf{W}_0(\tau)\right] \\ &= \quad \sigma k \mathsf{W}_0(r) + \sigma(1-k)\mathsf{W}_0(\tau). \end{aligned} \tag{A.2}$$

Now, $e_t \equiv y_t - T^{-1}\sum_{j=1}^{T} y_j = \overline{\epsilon}_t - T^{-1}\sum_{j=1}^{T}\overline{\epsilon}_j$, and hence from (A.1), (A.2) and an application of the Continuous Mapping Theorem (CMT), we have that, for $r \leq \tau$,

$$\begin{aligned} \sigma^{-1}T^{-1/2}\sum_{j=1}^{[Tr]} e_j \quad &\Rightarrow \quad \mathsf{W}_0(r) - r\left\{k\mathsf{W}_0(1) + (1-k)\mathsf{W}_0(\tau)\right\} \\ &= \quad \mathsf{B}_0(r) + (1-k)r\left\{\mathsf{W}_0(1) - \mathsf{W}_0(\tau)\right\} \equiv \mathsf{G}_{1,\mathsf{x}}(r), \end{aligned} \tag{A.3}$$

where $\mathsf{B}_0(r) \equiv \mathsf{W}_0(r) - r\mathsf{W}_0(1)$, and for $r > \tau$,

$$\begin{aligned} \sigma^{-1}T^{-1/2}\sum_{j=1}^{[Tr]} e_j \quad &\Rightarrow \quad k\mathsf{W}_0(r) + (1-k)\mathsf{W}_0(\tau) - r\left\{k\mathsf{W}_0(1) + (1-k)\mathsf{W}_0(\tau)\right\} \\ &= \quad k\mathsf{B}_0(r) + (1-k)(1-r)\mathsf{W}_0(\tau) \equiv \mathsf{G}_{2,\mathsf{x}}(r). \end{aligned} \tag{A.4}$$

Next we observe that

$$\begin{aligned} \hat{\sigma}^2 \quad &= \quad T^{-1}\sum_{j=1}^{T} e_t^2 \\ &= \quad T^{-1}\sum_{j=1}^{[T\tau]} e_j^2 + T^{-1}\sum_{j=[T\tau]+1}^{T} e_j^2 \\ &\rightarrow^p \quad (\tau + (1-\tau)k^2)\sigma^2, \end{aligned} \tag{A.5}$$

by standard arguments. The stated result then follows from (A.3), (A.4), (A.5) and applications of the CMT.

Proof of Proposition 4.1: For a multivariate Normal model of the form (4.2) with general covariance matrix $\mathbf{D}(\sigma_\eta^2)$ satisfying $\mathbf{D}(0) = \mathbf{I}$, King and Hillier (1985,Equation 6,p.99) demonstrate that the LBI test rejects for large values of $(\mathbf{e}_*'\mathbf{A}\mathbf{e}_*)(\mathbf{e}_*'\mathbf{e}_*)^{-1}$, where $\mathbf{e}_* = (e_1^*, ..., e_T^*)'$ are the OLS residuals from regressing $\mathbf{y}^*$ on $\mathbf{X}^*$ and $\mathbf{A} = \frac{\partial}{\partial\sigma_\eta^2}\mathbf{D}(\sigma_\eta^2)\Big|_{\sigma_\eta^2=0} \equiv \mathbf{V}$. Consequently, $\mathbf{e}_*'\mathbf{A}\mathbf{e}_* = \sum_{t=1}^{T}\left(\sum_{j=t}^{T} e_j^*\right)^2$, and the stated proposition then follows immediately.

Proof of Proposition 4.2: Using results in Busetti (2002) and Taylor (2002), the partial sum process $\sigma^{-1}T^{-1/2}\sum_{j=1}^{[Tr]} e_j^*$ converges weakly to $\mathsf{V}_\mathsf{x}(r) \equiv \mathsf{B}_\mathsf{x}(r) + c\int_0^r \mathsf{W}_{c,\mathsf{x}}(s)ds$ under the local alternative $\mathsf{H}_c : \sigma_\eta^2 = c^2/T^2$; see also Section 2. Consequently, $\sigma^{-1}T^{-1/2}\sum_{j=[Tr]+1}^{T} e_j^* \Rightarrow$

$\mathsf{V}_{\mathsf{x}}(1) - \mathsf{V}_{\mathsf{x}}(r)$. The proposition then follows using the consistency of $\hat{\sigma}_*^2$ for σ^2 and applications of the CMT.

Proof of Proposition 4.3: Let $\widehat{\mathbf{e}} = (\widehat{e}_1, ..., \widehat{e}_T)'$ be the vector of OLS residuals from the regression

$$\widehat{\mathbf{y}} = \mathbf{X}_1\beta_1 + \mathbf{X}_2\beta_2 + error. \tag{A.6}$$

Suppose first, for simplicity of notation, that $\mathbf{X}\odot\mathbf{H}_i$, $i = 1, 2$, is of full column rank; that is, $\mathbf{X}_i = \mathbf{X}\odot\mathbf{H}_i$, $i = 1, 2$. Then the OLS residuals from (A.6) are

$$\widehat{e}_t = y_t/\widehat{\sigma}_t - \mathbf{x}_t' \left(h_t(\tau)\hat{\beta}_1 + \overline{h}_t(\tau)\hat{\beta}_2 \right),$$

$t = 1, ..., T$, where $\hat{\beta}_1$, $\hat{\beta}_2$ are the OLS estimators from (A.6) and $\widehat{\sigma}_t$ is as defined in the statement of the proposition. Using the DGP (2.1)-(2.2), for $t \leq [T\tau]$ we have

$$\widehat{e}_t = \frac{\sigma_t}{\widehat{\sigma}_t} (\mu_t + \epsilon_t) - \mathbf{x}_t' \left(\hat{\beta}_1 - \beta \right),$$

and thus, under the local alternative $\mathsf{H}_c : \sigma_\eta^2 = c^2/T^2$,

$$[T\tau]^{-\frac{1}{2}} \sum_{j=1}^{[T\tau s]} \widehat{e}_j \Rightarrow \mathsf{V}_{\mathsf{x},1}(s), \qquad s \in [0, 1],$$

by the same arguments as for the standard NM test and the fact that $\widehat{\sigma}_t$ is a consistent estimator for $\sigma\sigma_t$. Similarly, for $t > [T\tau]$,

$$[T(1-\tau)]^{-\frac{1}{2}} \sum_{j=[T\tau]+1}^{[T(\tau+r-\tau r)]} \widehat{e}_j \Rightarrow \mathsf{V}_{\mathsf{x},2}(r), \qquad r \in [0, 1],$$

under H_c. The proposition then follows using an application of the CMT. The same line of argument, only with a much more involved notation, can be straightforwardly used for the case when $\mathbf{X}\odot\mathbf{H}_i$, $i = 1, 2$, is not of full column rank.