

Low-Frequency Robust Cointegration Testing*

Ulrich K. Müller

Mark W. Watson

Princeton University

Princeton University

Department of Economics

Department of Economics

Princeton, NJ, 08544

and Woodrow Wilson School

Princeton, NJ, 08544

June 2007

(This version: May 2012)

Abstract

Standard inference in cointegrating models is fragile because it relies on an assumption of an $I(1)$ model for the common stochastic trends, which may not accurately describe the data's persistence. This paper discusses efficient low-frequency inference about cointegrating vectors that is robust to this potential misspecification. A simple test motivated by the analysis in Wright (2000) is developed and shown to be approximately optimal in the case of a single cointegrating vector.

JEL classification: C32, C12

Keywords: stochastic trends, persistence, size distortion, interest rates, term spread

*We thank participants of the Cowles Econometrics Conference, the NBER Summer Institute and the Greater New York Metropolitan Area Econometrics Colloquium, and of seminars at Chicago, Cornell, Northwestern, NYU, Rutgers, and UCSD for helpful discussions. Support was provided by the National Science Foundation through grants SES-0518036 and SES-0617811.

1 Introduction

The fundamental insight of cointegration is that while economic time series may be individually highly persistent, some linear combinations are much less persistent. Accordingly, a suite of practical methods have been developed for conducting inference about cointegrating vectors, the coefficients that lead to this reduction in persistence. In their standard form, these methods assume that the persistence is the result of common $I(1)$ stochastic trends,¹ and their statistical properties crucially depend on particular characteristics of $I(1)$ processes. But in many applications there is uncertainty about the correct model for the persistence which cannot be resolved by examination of the data, rendering standard inference potentially fragile. This paper studies efficient inference methods for cointegrating vectors that is robust to this fragility.

We do this using a transformation of the data that focuses on low-frequency variability and covariability. This transformation has two distinct advantages. First, as we have argued elsewhere (Müller and Watson (2008)), persistence (“trending behavior”) and lack of persistence (“non-trending, $I(0)$ behavior”) are low-frequency characteristics, and attempts to utilize high-frequency variability to learn about low-frequency variability are fraught with their own fragilities.² Low-frequency transformations eliminate these fragilities by focusing attention on the features of the data that are of direct interest for questions relating to persistence. In particular, as in Müller and Watson (2008), we suggest focussing on below business cycle frequencies, so that the implied *definition* of cointegration is that error correction terms have a flat spectrum below business cycle frequencies. The second advantage is an important by-product of discarding high frequency variability. The major technical challenge when conducting robust inference about cointegrating vectors is to control size over the range of plausible processes characterizing the model’s stochastic common trends. Restricting attention to low frequencies greatly reduces the dimensionality of this challenge.

The potential impact of non- $I(1)$ stochastic trends on standard cointegration inference has long been recognized. Elliott (1998) provides a dramatic demonstration of the fragility of standard cointegration methods by showing that they fail to control size when the common stochastic trends are not $I(1)$, but rather are “local-to-unity” in the sense of Bobkoski

¹See, for instance, Johansen (1988), Phillips and Hansen (1990), Saikkonen (1991), Park (1992) and Stock and Watson (1993).

²Perhaps the most well-known example of this fragility involves estimation of HAC standard errors, see Newey and West (1987), Andrews (1991), den Haan and Levin (1997), Kiefer, Vogelsang, and Bunzel (2000), Kiefer and Vogelsang (2005), Müller (2007) and Sun, Phillips, and Jin (2008).

(1983), Cavanagh (1985), Chan and Wei (1987) and Phillips (1987).³ In a bivariate model, Cavanagh, Elliott, and Stock (1995) propose several procedures to adjust critical values from standard tests to control size over a range of values of the local-to-unity parameter, and their general approach has been used by several other researchers; Campbell and Yogo (2006) provides a recent example. Stock and Watson (1996), Jansson and Moreira (2006) and Elliott, Müller, and Watson (2012) go further and develop inference procedures with specific optimality properties in the local-to-unity model.

An alternative generalization of the $I(0)$ and $I(1)$ dichotomy is based on the fractionally integrated model $I(d)$, where d is not restricted to take on integer values (see, for instance, Baillie (1996) or Robinson (2003) for introductions). Fractional cointegration is then defined by the existence of a linear combination that leads to a reduction of the fractional parameter. A well-developed literature has studied inference in this framework: see, for instance, Velasco (2003), Robinson and Marinucci (2001, 2003), Robinson and Hualde (2003) and Chen and Hurvich (2003a, 2003b, 2006). As in the local-to-unity embedding, however, the low-frequency variability of the common stochastic trends is still governed by a single parameter, since (suitably scaled) partial sums of a fractionally integrated series converge to fractional Brownian motions, which are only indexed d . In contrast to the local-to-unity framework, this decisive parameter can be consistently estimated, so that the uncertainty about the exact nature of the stochastic trend vanishes in this fractional framework, at least under the usual asymptotics.

Yet, Müller and Watson (2008) demonstrate that relying on below business cycle variation, it is a hopeless endeavor to try to consistently discriminate between, say, local-to-unity and fractionally integrated stochastic data spanning 50 years. Similarly, Clive Granger discusses a wide range of possible data generating processes beyond the $I(1)$ model in his Frank Paish Lecture (1993) and argues, sensibly in our opinion, that it is fruitless to attempt to identify the exact nature of the persistence using the limited information in typical macro time series. While local-to-unity and fractional processes generalize the assumption of $I(1)$ trends, they do so in a very specific way, leading to worries about the potential fragility of these methods to alternative specifications of the stochastic trend.

As demonstrated by Wright (2000), it is nevertheless possible to conduct inference about a cointegrating vector without knowledge about the precise nature of the common stochastic trends. Wright's idea is to use the $I(0)$ property of the error correction term as the identifying property of the true cointegrating vector, so that a stationarity test of the model's putative

³Also see Elliott and Stock (1994) and Jeganathan (1997).

error correction term is used to conduct inference about the value of the cointegrating vectors. Because the common stochastic trends drop out under the null hypothesis, Wright's procedure is robust in the sense that it controls size under any model for the common stochastic trend. But the procedure ignores the data beyond the putative error correction term, and is thus potentially quite inefficient.

Section 2 of this paper provides a formulation of the cointegrated model in which the common stochastic trends follow a flexible limiting Gaussian process that includes the $I(1)$, local-to-unity, and fractional/long-memory models as special cases. Section 3 discusses the low-frequency transformation of the cointegrated model. Throughout the paper, inference procedures are studied in the context of this general formulation of the cointegrated model. The price to pay for this generality is that it introduces a potentially large number of nuisance parameters that characterize the properties of the stochastic trends and the relationship between the stochastic trends and the model's $I(0)$ components. In our framework, none of these nuisance parameters can be estimated consistently. The main challenge of this paper is thus to study efficient tests in the presence of nuisance parameters under the null hypothesis, and Sections 4–6 address this issue.

Using this framework, the paper then makes six contributions. The first is derive lower bounds on size distortions associated with trend specifications that are more general than those maintained under a test's null hypothesis. For example, for tests constructed under a maintained hypothesis that the stochastic trends follow an $I(1)$ process, we construct lower bounds on the test's size when the stochastic trends follow a local-to-unity or more general stochastic process. Importantly, these bounds are computed not for a specific test, but rather for *any* test with a pre-specified power. The paper's second contribution is an upper bound on the power for any test that satisfies a pre-specified rejection frequency under a null that may be characterized by a vector of nuisance parameters (here the parameters that characterize the stochastic trend process). The third contribution is implementation of a computational algorithm that allows us to compute an approximation to the lowest upper power bound and, when the number of nuisance parameters is small, a feasible test that approximately achieves the power bound.⁴ Taken together these results allow us to quantify both the power gains associated with exploiting restrictions associated with specific stochastic trend processes (for example, the power gains associated with the specializing the local-to-unity process to the $I(1)$ process), and the size distortions associated with these power gains when the

⁴The second and third contributions are applications of general results from a companion paper, Elliott, Müller, and Watson (2012), applied to the cointegration testing problem of this paper.

stochastic trend restrictions do not hold. Said differently, these results allow us to quantify the benefits (in terms of power) and costs (in terms of potential size distortions) associated with restrictions on the stochastic process characterizing the stochastic trend. Section 4 derives these size and power bounds in a general framework, and section 5 computes them for our cointegration testing problem.

The fourth contribution of the paper takes up Wright’s insight and develops efficient tests based only on the putative error-correction terms. We show that these tests have a particularly simple form when the alternative hypothesis restricts the model’s stochastic trends to be $I(1)$. The fifth contribution of the paper is to quantify the power loss associated with restricting tests to those that use only the error-correction terms rather than all of the data. This analysis shows that, in the case of single cointegration vector, a simple-to-compute test based only on the error-correction terms essentially achieves the full-data power bound for a general stochastic trend process, and is thus the efficient test. These results are developed in section 6.

The paper’s sixth contribution is empirical. We study the post-WWII behavior of long-term and short-term interest rates in the United States. While the levels of the interest rates are highly persistent, a suitably chosen linear combination of them is not, and we ask whether this linear combination corresponds to the term spread, the simple difference between long and short rates. More specifically we test whether with the cointegrating coefficient linking long rates and short rates is equal to unity. This value cannot be rejected using a standard efficient $I(1)$ test (Wald or LR versions of Johansen’s (1991) test), and we show that this result continues to hold under a general trend process. Of course, other values of the cointegrating coefficient are possible both in theory and in the data, and we construct a confidence set for the value of the cointegrating coefficient allowing for a general process and compare it the confidence set constructed used standard $I(1)$ methods. These results are presented in section 7.

2 Model

Let p_t , $t = 1, \dots, T$ denote the $n \times 1$ vector of variables under study. This section outlines a time domain representation of the cointegrated model for p_t in terms of canonical variables representing a set of common trends and $I(0)$ error correction terms. The common trends are allowed to follow a flexible process that includes $I(1)$, local-to-unity, and fractional models as special cases, but aside from this generalization, the cointegrated model for p_t is standard.

To begin, p_t is transformed into two components, where one component is $I(0)$ under the null hypothesis and the other component contains elements that are not cointegrated. Let β denote an $n \times r$ matrix whose linearly independent columns are the cointegrating vectors, let β_0 denote the value of β under the null, and $y_t = \beta_0' p_t$. The elements in y_t are the model's error correction terms under the null hypothesis. Let $x_t = \delta' p_t$ where δ is $n \times k$ with $k = n - r$, and where the linearly independent columns of δ are linearly independent of the columns of β_0 , so that the elements of x_t are not cointegrated under the null. Because the cointegrated model only determines the column space of the matrix of cointegrating vectors, the variables y_t and x_t are determined up to transformations $(y_t, x_t) \rightarrow (A_{yy}y_t, A_{xx}x_t + A_{xy}y_t)$, where A_{yy} and A_{xx} are non-singular. Most extant inference procedures are invariant (or asymptotically invariant) to these transformations, and, as discussed in detail below, our analysis will also focus on invariant tests.

2.1 Canonical Variable Representation of y_t and x_t

We will represent y_t and x_t in terms of a common stochastic trend vector v_t and an $I(0)$ vector z_t

$$\begin{aligned} y_t &= \Gamma_{yz}z_t + \Gamma_{yv}v_t \\ x_t &= \Gamma_{xz}z_t + \Gamma_{xv}v_t, \end{aligned} \tag{1}$$

where z_t is $r \times 1$, v_t is $k \times 1$, and Γ_{yz} and Γ_{xv} have full rank. In this representation, the restriction that y_t is $I(0)$ corresponds to the restriction $\Gamma_{yv} = 0$. All of the test statistics discussed in this paper are invariant to adding constants to the observations, so that constant terms are suppressed in (1). As a technical matter, we think of $\{z_t, v_t\}_{t=1}^T$ (and thus also $\{x_t, y_t\}_{t=1}^T$) as being generated from a triangular array; we omit the additional dependence on T to ease notation. Also, we write $[x]$ for the integer part of $x \in \mathbb{R}$, $\|A\| = \sqrt{\text{tr } A'A}$ for any real matrix A , $x \vee y$ for the maximum of $x, y \in \mathbb{R}$, ' \otimes ' for the usual Kronecker product and ' \Rightarrow ' to indicate weak convergence.

Let $W(\cdot)$ denote a $n \times 1$ standard Wiener process. The vector z_t is a canonical $I(0)$ vector in the sense that its partial sums converge to a $r \times 1$ Wiener process

$$T^{-1/2} \sum_{t=1}^{\lfloor sT \rfloor} z_t \Rightarrow S_z W(s) = W_z(s), \text{ where } S_z S_z' = I_r. \tag{2}$$

The vector v_t is a common trend in the sense that scaled versions of its level converge to a stochastic integral with respect to $W(\cdot)$. For example, in the standard $I(1)$ model,

$T^{-1/2}v_{[sT]} \Rightarrow \int_0^s HdW(t)$, where H is a $k \times n$ matrix and (H', S'_z) has full rank. More general trend processes, such as the local-to-unity formulation, allow the matrix H to depend on s and t . The general representation for the common trends used in this paper is

$$T^{-1/2}v_{[sT]} \Rightarrow \int_{-\infty}^s H(s, t)dW(t) \quad (3)$$

where $H(s, t)$ is sufficiently well behaved to ensure that there exists a cadlag version of the process $\int_{-\infty}^s H(s, t)dW(t)$.⁵

2.2 Invariance and Reparameterization

As discussed above, because cointegration only identifies the column space of β , attention is restricted to tests that are invariant to the group of transformations

$$(y_t, x_t) \rightarrow (A_{yy}y_t, A_{xx}x_t + A_{xy}y_t) \quad (4)$$

where A_{yy} and A_{xx} are non-singular, but (A_{yy}, A_{xx}, A_{xy}) are otherwise unrestricted real matrices.

The restriction to invariant tests allows a simplification of notation: because the test statistics are invariant to the transformations in (4), there is no loss of generality setting $\Gamma_{yz} = I_r$, $\Gamma_{xv} = I_k$, and $\Gamma_{xz} = 0$. With these values, the model is

$$\begin{aligned} y_t &= z_t + \Gamma_{yv}v_t \\ x_t &= v_t. \end{aligned} \quad (5)$$

2.3 Restricted Versions of the Trend Model

We will refer to the general trend specification in (3) as the “unrestricted” stochastic trend model throughout the remainder of the paper. The existing literature on efficient tests relies on restricted forms of the trend process (3) such as $I(1)$ or local-to-unity processes, and we compute the potential power gains associated with these and other *a priori* restrictions on $H(s, t)$ below. Here we describe five restricted versions of the stochastic trend.

⁵The common scale $T^{-1/2}$ for the $k \times 1$ vector v_t in (3) is assumed for convenience; with an appropriate definition of local alternatives, the invariance (4) ensures that one would obtain the same results for any scaling of v_t . For example, for an $I(2)$ stochastic trend scaled by $T^{-3/2}$, set $H(s, t) = \mathbf{1}[t \geq 0](s - t)H$, with the $k \times n$ matrix H as in the $I(1)$ case.

The first model, which we will refer to as the G -model, restricts $H(s, t)$ to satisfy

$$H(s, t) = G(s, t)S_v, \quad (6)$$

where $G(s, t)$ is $k \times k$ and S_v is $k \times n$ with $S_v S_v' = I_k$ and (S_z', S_v') nonsingular. In this model, the common trend depends on $W(\cdot)$ only through the $k \times 1$ standard Wiener process $W_v(\cdot) = S_v W(\cdot)$, and this restricts the way that v_t and z_t interact. In this model

$$T^{-1/2}v_{\lfloor sT \rfloor} \Rightarrow \int_{-\infty}^s G(s, t) dW_v(t), \quad (7)$$

and the covariance between the Wiener process characterizing the partial sums of z_t , W_z , and W_v is equal to the $r \times k$ matrix $R = S_z S_v'$. Standard $I(1)$ and local-to-unity formulations of cointegration satisfy this restriction and impose additional parametric restrictions on $G(s, t)$.

The second model further restricts (7) so that $G(s, t)$ is diagonal:

$$G(s, t) = \text{diag}(g_1(s, t), \dots, g_k(s, t)). \quad (8)$$

An interpretation of this model is that the k common trends evolve independently of one another (recall that W_v has identity covariance matrix), where each trend is allowed to follow a different process characterized by the functions $g_i(s, t)$.

The third model further restricts the diagonal- G model so that the k stochastic trends converge weakly to a stationary continuous time process. We thus impose

$$g_i(s, t) = g_i^S(s - t), \quad i = 1, \dots, k. \quad (9)$$

The stationary local-to-unity model (with an initial condition drawn from the unconditional distribution), for instance, satisfies this restriction.

Finally, we consider two parametric restrictions of G :

$$G(s, t) = \mathbf{1}[t > 0]I_k \quad (10)$$

which is the $I(1)$ model, and

$$G(s, t) = \mathbf{1}[t > 0]e^{C(s-t)} \quad (11)$$

which is the multivariate local-to-unity model, where C is the $k \times k$ diffusion matrix of the limiting Ornstein-Uhlenbeck process (with zero initial condition).⁶

⁶The $I(1)$ specification in (10) is the same as the $I(1)$ specification given below (2) because the invariance in (4) implies that the trend models are unaffected by premultiplication of $H(s, t)$ (or $G(s, t)$) by an arbitrary non-singular $k \times k$ matrix.

2.4 Testing Problem and Local Alternatives

The goal of the paper is to study asymptotically efficient tests for the value of the cointegrating vectors with controlled rejection probability under the null hypothesis for a range of stochastic trend specifications. The different orders of magnitude of z_t and v_t in (2) and (3) suggest a local embedding of this null hypothesis against alternatives where $\Gamma_{yv} = T^{-1}B$ for B a constant $r \times k$ matrix, so that in model (5),

$$T^{-1/2} \sum_{t=1}^{\lfloor sT \rfloor} y_t \Rightarrow S_z W(s) + B \int_0^s \int_{-\infty}^u H(u, t) dW(t) du.$$

In this parametrization, the null hypothesis becomes

$$H_0 : B = 0, H(s, t) \in \mathcal{H}_0 \tag{12}$$

where $H(s, t)$ is restricted to a set of functions \mathcal{H}_0 , that, in the unrestricted trend model includes functions sufficiently well behaved to ensure that there exists a cadlag version of the process $\int_{-\infty}^s H(s, t) dW(t)$, or more restricted versions of $H(s, t)$ as in (6), (8), (9), (10), or (11).

Since our goal is to consider efficient tests of the null hypothesis (12), we also need to specify the alternative hypothesis. Our results below are general enough to allow for the derivation of efficient tests against any particular alternative with specified $B = B_1$ and stochastic trend process $H(s, t) = H_1(s, t)$,

$$H_1 : B = B_1, H(s, t) = H_1(s, t) \tag{13}$$

or, more generally, for tests that are efficient in the sense of maximizing weighted average power against a set of values for B_1 and stochastic trend models $H_1(s, t)$.

Our numerical results, however, focus on alternatives in which the stochastic trend v_t is $I(1)$, so that $H_1(s, t)$ satisfies (6) and (10). This is partly out of practical considerations: while there is a wide range of potentially interesting trend specification, the computations for any particular specification are involved, and these computational complications limit the number of alternatives we can usefully consider. At the same time, one might consider the classical $I(1)$ model as an important benchmark against which it is useful to maximize power—not necessarily because this is the only plausible model under the alternative, but because a test that performs well against this alternative presumably has reasonable power properties for a range of empirically relevant models. We stress that despite this focus on

the $I(1)$ stochastic trend model for the alternative hypothesis (13), we restrict attention to tests that control size for a range of models under the null hypothesis (12). The idea is to control the frequency of rejections under the null hypothesis for any stochastic trend model in \mathcal{H}_0 , so that the rejection of a set of cointegrating vectors cannot simply be explained by the stochastic trends not being exactly $I(1)$. In this sense, our approach is one of “robust” cointegration testing, with the degree of robustness governed by the size of the set \mathcal{H}_0 .

2.5 Summary

To summarize, this section has introduced the time domain representation of the cointegrated model with a focus on the problem of inference about the space of cointegrating vectors. In all respects except one, the representation is the standard one: the data are expressed as a linear function of a canonical vector or common trends and a vector of $I(0)$ components. Under the null, certain linear combinations of the data do not involve the common trends. Because the null only restricts the column space of the matrix of cointegrating vectors, attention is restricted to invariant tests. The goal is to construct tests with best power for an alternative value for the matrix of cointegrating vectors under a particular model for the trend (or best weighted average power for a collection of B_1 and $H_1(s, t)$). The formulation differs from the standard one only in that it allows the model for the trend process under the null to be less restrictive than the standard formulation. Said differently, because of potential uncertainty about the specific form of the trend process, the formulation restricts attention to tests that control size for a range of different trend processes. This generalization complicates the problem of constructing efficient tests by introducing a potentially large number of nuisance parameters (associated with the trend process) under the null hypothesis.

3 Low-Frequency Representation of the Model

Cointegration is a restriction on the low-frequency behavior of time series, and as discussed in the introduction, we therefore focus on the low-frequency behavior of (y_t, x_t) . This low-frequency variability is summarized by a small number, q , of weighted averages of the data. In this section we discuss these weighted averages and derive their limiting behavior under the null and alternative hypotheses.

3.1 Low-Frequency Weighted Averages

We use weights associated with the cosine transform, where the j 'th weight is given by $\Psi_j(s) = \sqrt{2} \cos(j\pi s)$. For any sequence $\{a_t\}_{t=1}^T$, the j 'th weighted average will be denoted by

$$A_T(j) = \int_0^1 \Psi_j(s) a_{\lfloor sT \rfloor + 1} ds = \iota_{jT} T^{-1} \sum_{t=1}^T \Psi_j\left(\frac{t-1/2}{T}\right) a_t \quad (14)$$

where $\iota_{jT} = (2T/j\pi) \sin(j\pi/2T) \rightarrow 1$ for all fixed j . As demonstrated by Müller and Watson (2008), the weighted averages $A_T(j)$, $j = 0, \dots, q$, essentially capture the variability in the sequence corresponding to frequencies below $q\pi/T$.

We use the following notation: with a_t a $h \times 1$ vector time series, let $\Psi(s) = (\Psi_1(s), \Psi_2(s), \dots, \Psi_q(s))'$ denote the $q \times 1$ vector of weighting functions, and $A_T = \int_0^1 \Psi(s) a'_{\lfloor sT \rfloor + 1} ds$ the $q \times h$ matrix of weighted averages of the elements of a_t , where $\Psi_0(s)$ is excluded to make the results invariant to adding constants to the data. Using this notation, the $q \times r$ matrix Y_T and the $q \times k$ matrix X_T summarize the variability in the data corresponding to frequencies lower than $q\pi/T$. With $q = 12$, (Y_T, X_T) capture variability lower than the business cycle (periodicities greater than 8 years) for time series that span 50 years (postwar data) regardless of the sampling frequency (months, quarters, weeks, etc.). This motivates us to consider the behavior of these matrices as $T \rightarrow \infty$, but with q held fixed.

The large-sample behavior of X_T and Y_T follows from the behavior of Z_T and V_T . Using the assumed limits (2) and (3), the continuous mapping theorem, and integration by parts for the terms involving Z_T , one obtains

$$\begin{bmatrix} T^{1/2} Z_T \\ T^{-1/2} V_T \end{bmatrix} \Rightarrow \begin{bmatrix} Z \\ V \end{bmatrix} \quad (15)$$

where

$$\begin{bmatrix} \text{vec } Z \\ \text{vec } V \end{bmatrix} \sim \mathcal{N} \left(0, \begin{bmatrix} I_{rq} & \Sigma_{ZV} \\ \Sigma_{VZ} & \Sigma_{VV} \end{bmatrix} \right) \quad (16)$$

with

$$\begin{aligned} \Sigma_{VZ} &= \int_0^1 \left(\int_{t \vee 0}^1 [H(s, t) \otimes \Psi(s)] ds \right) [S_z \otimes \Psi(t)]' dt \\ \Sigma_{VV} &= \int_{-\infty}^1 \left(\int_{t \vee 0}^1 [H(s, t) \otimes \Psi(s)] ds \right) \left(\int_{t \vee 0}^1 [H(s, t) \otimes \Psi(s)] ds \right)' dt. \end{aligned} \quad (17)$$

The relative scarcity of low-frequency information is thus formally captured by considering the weak limits (2) and (3) as pertinent only for the subspace spanned by the weight function

$\Psi(\cdot)$, yielding (15) as a complete characterization of the relevant properties of the error correction term z_t and the common stochastic trend v_t .

Using $\Gamma_{yv} = T^{-1}B$, equation (5) implies that $Y_T = Z_T + T^{-1}V_T B'$ and $X_T = V_T$. Thus,

$$\begin{bmatrix} T^{1/2}Y_T \\ T^{-1/2}X_T \end{bmatrix} \Rightarrow \begin{bmatrix} Y \\ X \end{bmatrix} = \begin{bmatrix} Z + VB' \\ V \end{bmatrix} \quad (18)$$

where

$$\begin{bmatrix} \text{vec } Y \\ \text{vec } X \end{bmatrix} \sim \mathcal{N}(0, \Sigma_{(Y,X)}) \quad (19)$$

with

$$\Sigma_{(Y,X)} = \begin{bmatrix} I_r \otimes I_q & B \otimes I_q \\ 0 & I_k \otimes I_q \end{bmatrix} \begin{bmatrix} I_r \otimes I_q & \Sigma_{ZV} \\ \Sigma_{VZ} & \Sigma_{VV} \end{bmatrix} \begin{bmatrix} I_r \otimes I_q & 0 \\ B' \otimes I_q & I_k \otimes I_q \end{bmatrix}. \quad (20)$$

3.2 “Best” Low-Frequency Hypothesis Tests

We consider invariant tests of H_0 against H_1 given in (12) and (13) based on the data $\{y_t, x_t\}_{t=1}^T$. Because we are concerned with the model’s implications for the low-frequency variability of the data, we restrict attention to tests that control asymptotic size for all models that satisfy (18)-(20). Our goal is to find an invariant test that maximizes power subject to this restriction, and for brevity we will refer to such a test as a “best” test. Müller (2011) considers the general problem of constructing asymptotically most powerful tests subject to asymptotic size control over a class of models such as ours. In our context, his results imply that asymptotically best tests correspond to the most powerful invariant tests associated with the limiting distribution (19).

Thus, the relevant testing problem has a simple form: $\text{vec}(Y, X)$ has a normal distribution with mean zero and covariance matrix that depends on B . Under the null $B = 0$, while under the alternative $B \neq 0$. Tests are restricted to be invariant to the group of transformations

$$(Y, X) \rightarrow (YA'_{yy}, XA'_{xx} + YA'_{xy}) \quad (21)$$

where A_{yy} and A_{xx} are nonsingular, and A_{yy} , A_{xx} , and A_{xy} are otherwise unrestricted. Thus, the hypothesis testing problem becomes the problem of using an invariant procedure to test a restriction on the covariance matrix of a multivariate normal vector.

4 Bounds on Power and Size

The general version of the hypothesis testing problem we are facing is a familiar one: Let U denote a single observation of dimension $m \times 1$. (In our problem, U corresponds to the maximal invariant for (Y, X)). Under the null hypothesis U has probability density $f_\theta(u)$ with respect to some measure μ , where $\theta \in \Theta$ is a vector of nuisance parameters. (In our problem, the vector θ describes the stochastic trend process under the null hypothesis and determines $\Sigma_{(Y,X)}$ via (17) and (20)). Under the alternative, U has known density $h(u)$. (Choices for $h(u)$ for our problem will be discussed in Subsection 5.2.1.) Thus, the null and alternative hypothesis are

$$\begin{aligned} H_0 &: \text{The density of } U \text{ is } f_\theta(u), \theta \in \Theta \\ H_1 &: \text{The density of } U \text{ is } h(u), \end{aligned} \tag{22}$$

and possibly randomized tests are (measurable) functions $\varphi : \mathbb{R}^m \mapsto [0, 1]$, where $\varphi(u)$ is the probability of rejecting the null hypothesis when observing $U = u$, so that size and power are given by $\sup_{\theta \in \Theta} \int \varphi f_\theta d\mu$ and $\int \varphi h d\mu$, respectively.

This section presents two probability bounds, one on power and one on size, in this general problem. The power bound is taken from Elliott, Müller, and Watson (2012) (EMW), and provides an upper bound on the power of any valid test of H_0 versus H_1 . As discussed in EMW, this power bound is useful for two reasons. First, when the dimension of θ is small, numerical methods developed in EMW can be used to construct a test that essentially attains the power bound, so the resulting test is the approximately efficient test. For example, in our problem, we will use the EMW algorithm to compute the approximately efficient test when x_t follows a local-to-unity process. The second reason the bound is useful is that it can be computed even when the dimension of θ is large, where it is infeasible to directly construct the efficient test using numerical methods. The resulting power bound makes it possible to evaluate the potential power shortfalls of existing tests. For example, when x_t follows the general stochastic trend process (3) and there is a single cointegrating vector ($r = 1$), the power of a low-frequency of a test originally proposed by Wright (2000) essentially coincides the power bound and this allows us to conclude that the test is approximately efficient.

The second probability bound focuses on size, and is motivated in our context by the following concern. Suppose that H_0 specifies that the stochastic trend follows an $I(1)$ process, and consider a test that exploits features of the $I(1)$ process to increase power. Uncertainty about the trend process means that it is useful to know something about the rejection frequency of tests under null hypotheses that allow for more general trends, such as the

unrestricted trend model (3) or other less restricted versions described above. We provide a lower bound on this rejection frequency, where a large value of this lower bound highlights the fragility of tests that exploit a particular H_0 to obtain more powerful inference.

4.1 An Upper Bound on Power

A standard device for problems such as (22) is to consider a Neyman-Pearson test for a related problem in which the null hypothesis is replaced with a mixture

$$H_\Lambda : \text{The density of } U \text{ is } \int f_\theta d\Lambda(\theta)$$

where Λ is a probability distribution for θ with support in Θ . The following lemma (taken from EMW) shows that the power of the Neyman-Pearson test of H_Λ versus H_1 provides an upper power bound for tests of H_0 versus H_1 .

Lemma 1 *Let φ_Λ be the best level α test of H_Λ against H_1 . Then for any level α test φ of H_0 against H_1 , $\int \varphi_\Lambda h d\mu \geq \int \varphi h d\mu$.*

Proof. *Since φ is a level α test of H_0 , $\int \varphi f_\theta d\mu \leq \alpha$ for all $\theta \in \Theta$. Therefore, $\int \int \varphi f_\theta d\mu d\Lambda(\theta) = \int \int \varphi f_\theta d\Lambda(\theta) d\mu \leq \alpha$ (where the change in the order of integration is allowed by Fubini's Theorem), so that φ is also a level α test of H_Λ against H_1 . The result follows by the definition of a best test. ■*

This result is closely related to Theorem 3.8.1 of Lehmann and Romano (2005) which provides conditions under which a least upper bound on the power for tests H_0 versus H_1 is associated with a “least favorable distribution” for θ , and that using this distribution for Λ produces the least upper power bound. The least favorable distribution Λ^* has the characteristic that the resulting φ_{Λ^*} is a level α test for testing H_0 versus H_1 . Said differently, if φ_{Λ^*} is the best level α test of H_{Λ^*} against H_1 and is also a level α test for testing H_0 versus H_1 , then $\varphi^* = \varphi_{\Lambda^*}$, that is φ_{Λ^*} is the most powerful level α test of H_0 versus H_1 . Unfortunately, while the test associated with the least favorable distribution solves the testing problem (22), there is no general and constructive method for finding the least favorable distribution Λ^* (and it does not always exist).

With this in mind, Lemma 1 is stated so that Λ is not necessarily the least favorable distribution. That is, the bound in Lemma 1 holds for any probability distribution Λ . The goal of the numerical analysis carried out below is to choose Λ to approximate the least upper bound. Importantly, even if one cannot identify the least favorable distribution, Lemma 1

shows that the power of φ_Λ provides a valid bound for the power of any test of H_0 versus H_1 , for any Λ .

4.2 A Lower Bound on Size under an Auxiliary Null

Now consider the "Larger" auxiliary null hypothesis

$$H_L : \text{The density of } U \text{ is } f_\theta(u), \theta \in \Theta_L$$

with an associated mixture

$$H_{\Lambda_L} : \text{The density of } U \text{ is } \int f_\theta d\Lambda_L(\theta)$$

where Λ_L has support in Θ_L . (In our problem H_L will be a null hypothesis that allows for a less restricted trend process than under H_0 . Thus if H_0 allows only for an $I(1)$ trend, H_L might allow for a local-to-unity trend or one of the more general trends processes discussed in Section 2.)

Consider any test φ of level α_0 under H_0 with power of at least β . The following lemma provides a lower bound on the rejection frequency under the auxiliary null H_L .

Lemma 2 (a) *The problem*

$$\begin{aligned} & \min_{\varphi} \int \varphi \int f_\theta d\Lambda_L(\theta) d\mu \\ \text{s.t. } & \int \varphi \int f_\theta d\Lambda_0(\theta) d\mu \leq \alpha_0 \quad \text{and} \quad \int \varphi h d\mu \geq \beta. \end{aligned}$$

is solved by $\varphi^* = \mathbf{1}[h \geq \lambda_1 \int f_\theta d\Lambda_0(\theta) + \lambda_2 \int f_\theta d\Lambda_L(\theta)]$, where λ_1 and λ_2 are non-negative constants associated with the constraints.

(b) Let $\alpha_L = \int \varphi^* \int f_\theta d\Lambda_L(\theta) d\mu$ denote the minimized value of the objective function in (a). Let φ be a level α_0 test under H_0 and of power of at least β . Then $\sup_{\theta \in \Theta_L} \int \varphi f_\theta d\mu \geq \alpha_L$.

Proof. (a) is a variant of the generalized Neyman-Pearson Lemma (Theorem 3.6.1 in Lehmann and Romano (2005)). (b) Note that φ satisfies the two constraints in the problem given in (a), so that $\int \varphi \int f_\theta d\Lambda_L(\theta) d\mu \geq \alpha_L$, and the result follows from $\sup_{\theta \in \Theta_L} \int \varphi f_\theta d\mu \geq \int \varphi \int f_\theta d\Lambda_L(\theta) d\mu$. ■

This lemma is particularly useful in conjunction with Lemma 1: Suppose application of Lemma 1 implies that no 5% level test of a relatively restricted H_0 can have power of more

than, say, 70%. This suggests that there could indeed exist a 5% level test φ with power, say, 67%, and one might want to learn about the size properties of such tests under the more general null hypothesis H_L . Lemma 2 provides a way of computing a lower bound on this size that is valid for *any* test with power of at least 67%. So if this size distortion is large, then without having to determine the class of 5% level tests of H_0 with power of at least 67%, one can already conclude that all such tests will be fragile. In the numerical section below, we discuss how to determine suitable Λ_0 and Λ_L to obtain a large lower bound α_L (λ_1 and λ_2 are determined through the two constraints on φ^*).

5 Computing Bounds

In this section we compute the power and size bounds from the last section. The analysis proceeds in four steps. First, we derive the density of the maximal invariant of (Y, X) ; this density forms the basis of the likelihood ratio. Second, since the density of the maximal invariant depends on the covariance matrix of (Y, X) , we discuss the parameterization of $\Sigma_{(Y,X)}$ under the null and alternative hypotheses. In the third step we describe how the mixing distributions Λ , Λ_0 and Λ_L are chosen to yield tight bounds. Finally, we present numerical values for the bounds.

5.1 Density of a Maximal Invariant

Recall that we are considering tests that are invariant to the group of transformations $(Y, X) \rightarrow (YA'_{yy}, XA'_{xx} + YA'_{xy})$ where A_{yy} and A_{xx} are nonsingular, and A_{yy} , A_{xx} , and A_{xy} are otherwise unrestricted. Any invariant test can be written as a function of a maximal invariant (Theorem 6.2.1 in Lehmann and Romano (2005)), so that by the Neyman-Pearson lemma, the most powerful invariant test rejects for large values of the likelihood ratio statistic of a maximal invariant. The remaining challenge is the computation of the density of a maximal invariant, and this is addressed in the following theorem.

Theorem 1 *If $\text{vec}(Y, X) \sim \mathcal{N}(0, \Sigma_{(Y,X)})$ with positive definite $\Sigma_{(Y,X)}$ and $q > r + k$, the density of a maximal invariant of (21) has the form*

$$c(\det \Sigma_{(Y,X)})^{-1/2} (\det V'_{0Y} \Sigma_{(Y,X)}^{-1} V_{0Y})^{-1/2} (\det \Omega)^{-1/2} E_{\omega} [|\det(\omega_Y)|^{q-r} |\det(\omega_X)|^{q-r-k}]$$

where c does not depend on $\Sigma_{(Y,X)}$, ω_Y and ω_X are random $r \times r$ and $k \times k$ matrices,

respectively, with $(\text{vec } \omega'_y, \text{vec } \omega'_x)' \sim \mathcal{N}(0, \Omega^{-1})$,

$$\Omega = D'_{YX} \Sigma_{(Y,X)}^{-1} D_{YX} - D'_{YX} \Sigma_{(Y,X)}^{-1} V_{0Y} (V'_{0Y} \Sigma_{(Y,X)}^{-1} V_{0Y})^{-1} V'_{0Y} \Sigma_{(Y,X)}^{-1} D_{YX},$$

$D_{YX} = \text{diag}(I_r \otimes Y, I_k \otimes X)$, $V_{0Y} = (0'_{r \times r}, I_k \otimes Y)'$, and E_ω denotes integration with respect to ω_Y and ω_X , conditional on (Y, X) .

Theorem 1 shows that the density of a maximal invariant can be expressed in terms of absolute moments of determinants of jointly normally distributed random matrices, whose covariance matrix depends on (Y, X) . We do not know of a useful and general closed-form solution for this expectation; for $r = k = 1$, however, Nabeya's (1951) results for the absolute moments of a bivariate normal yields an expression in terms of elementary functions, which we omit for brevity. When $r + k > 2$, the moments can be computed via Monte Carlo integration. However, computing accurate approximations is difficult when r and k are large, and the numerical analysis reported below is therefore limited to small values of r and k .

5.2 Parameterization of $\Sigma_{(Y,X)}$

Since the density of the maximal invariant of Theorem 1 depends on $\Sigma_{(Y,X)}$, the derivation of efficient invariant tests requires specification of $\Sigma_{(Y,X)}$ under the alternative and null hypothesis. We discuss each of these in turn.

5.2.1 Specification of $\Sigma_{(Y,X)}$ under the Alternative Hypothesis

As discussed above, we focus on the alternative where the stochastic trends follow an $I(1)$ process, so that $H(s, t)$ satisfies (6) and (10). There remains the issue of the value of B (the coefficients that show how the trends affect Y) and R (the correlation of the Wiener processes describing the $I(0)$ variables, z_t , and the common trends, v_t). For these parameters, we consider point-valued alternatives with $B = B_1$ and $R = R_1$; the power bounds derived below then serve as bounds on the asymptotic power envelope over these values of B and R . Invariance reduces the effective dimension of B and R somewhat, and this will be discussed in the context of the numerical results presented below.

5.2.2 Parameterization of $\Sigma_{(Y,X)}$ under the Null Hypothesis

From (20), under the null hypothesis with $B = 0$, the covariance matrix $\Sigma_{(Y,X)}$ satisfies

$$\Sigma_{(Y,X)} = \begin{bmatrix} I_{rq} & \Sigma_{ZV} \\ \Sigma_{VZ} & \Sigma_{VV} \end{bmatrix}.$$

The null model's specification of the stochastic trend determines the $rq \times kq$ matrix Σ_{ZV} and the $kq \times kq$ matrix Σ_{VV} by the formulae given in (17). Since these matrices contain a finite number of elements, it is clear that even for nonparametric specifications of $H(s, t)$, the effective parameter space for low-frequency tests based on (Y, X) is finite dimensional. We collect these nuisance parameters in a vector $\theta \in \Theta$.

Section 2 discussed several trend processes, beginning with the general process given in (3) with an unrestricted version of $H(s, t)$, and then five restricted models: (i) the “ G ” model in (6), (ii) the “Diagonal” model (8), (iii) the “Stationary” model (9), (iv) the local-to-unity model (11), and (v) the $I(1)$ model (10). The appendix discusses convenient parameterization for $\Sigma_{(Y,X)}$ for these five restricted models, and the following lemma provides the basis for parameterizing $\Sigma_{(Y,X)}$ when $H(s, t)$ is unrestricted.

Lemma 3 (a) *For any $(r+k)q \times (r+k)q$ positive definite matrix Σ^* with upper left $rq \times rq$ block equal to I_{rq} , there exists an unrestricted trend model with $H(s, t) = 0$ for $t < 0$ such that $\Sigma^* = E[\text{vec}(Z, V)(\text{vec}(Z, V))']$.*

(b) *If $r \leq k$, this $H(s, t)$ can be chosen of the form $H(s, t) = G(s, t)S_v$, where (S'_z, S'_v) has full rank.*

The lemma shows that when $H(s, t)$ is unrestricted (or $r \leq k$ and $H(s, t) = G(s, t)S_v$ with G unrestricted) the only restriction that the null hypothesis imposes on $\Sigma_{(Y,X)}$ is that $\Sigma_{YY} = I_{rq}$.⁷ In other words, since Σ_{ZV} and Σ_{VV} have $rkq^2 + kq(kq + 1)/2$ distinct elements, an appropriately chosen θ of that dimension determines $\Sigma_{(Y,X)}$ under the null hypothesis in the unrestricted model, and in the model where $H(s, t) = G(s, t)S_v$ for $r \leq k$.

⁷Without the invariance restriction (21), this observation would lead to an analytic least favorable distribution result: Factor the density of (Y, X) under the alternative into the product of the density of Y , and the density of X given Y . By choosing Σ_{VZ} and Σ_{VV} under the null hypothesis appropriately, the latter term cancels, and the Neyman-Pearson test is a function of Y only. In section 6 below we consider the efficient Y -only invariant test and compare its power to the approximate (Y, X) power bound.

5.3 Approximating the Least Upper Power and Greatest Lower Size Bound

We discuss two methods to approximate the power bound associated with the least favorable distribution from Section 4, and use the second method also to determine a large lower size bound. First, we discuss a generic algorithm developed in Elliott, Mueller and Watson (2012) that simultaneously determines a low upper bound on power, and a level α test whose power is close to that bound. The computational complexity is such, however, that it can only be applied when θ is low-dimensional; as such, it is useful for our problem only in the $I(1)$ and local-to-unity stochastic trend model for $r = k = 1$. Second, when the dimension of θ is large we choose Λ (and Λ_0 and Λ_L for Lemma 2) so the null and alternative distributions are close in a numerically convenient metric. Two numerical results suggest that this second method produces a reasonably accurate estimate of the lower bound: the method produces power bounds only marginally higher than the first method (when the first method is feasible), and when $r = 1$ we find that the method produces a power bound that can be achieved by a feasible test that we present in the next section.

We discuss the two methods in turn.

5.3.1 Low Dimensional Nuisance Parameter

Suppose that $LR_\Lambda = h(U) / \int f_\theta(U) d\Lambda(\theta)$ is a continuous random variable for any Λ , so that by the Neyman-Pearson Lemma, φ_Λ is of the form $\varphi_\Lambda = \mathbf{1}[LR_\Lambda > cv_\Lambda]$, where the critical value cv_Λ is chosen to satisfy the size constraint $\int \int \varphi_\Lambda f_\theta d\mu d\Lambda(\theta) = \alpha$. Then by Lemma 1, the power of φ_Λ , $\beta_\Lambda = \int \varphi_\Lambda h d\mu$, is an upper bound on the power of any test that is level α under H_0 . If Λ is not the least favorable distribution, then φ_Λ is not of size α under H_0 , i.e. $\sup_{\theta \in \Theta} \int \varphi_\Lambda f_\theta d\mu > \alpha$. Now consider a version of φ_Λ with a size-corrected critical value $cv_\Lambda^c > cv_\Lambda$, that is $\varphi_\Lambda^c = \mathbf{1}[LR_\Lambda > cv_\Lambda^c]$ with cv_Λ^c chosen to satisfy the size constraint $\sup_{\theta \in \Theta} \int \varphi_\Lambda^c f_\theta d\mu = \alpha$. Because the size adjusted test φ_Λ^c is of level α under H_0 , the least upper power bound must be between β_Λ^c and β_Λ . Thus, if β_Λ^c is close to β_Λ , then β_Λ^c serves as a good approximation to the least upper bound.

The challenge is to find an appropriate Λ . This is difficult because, in general, no closed form solutions are available for the size and power of tests, so that these must be approximated by Monte Carlo integration. Brute force searches for an appropriate Λ are not computationally feasible. Elliott, Müller, and Watson (2012) develop an algorithm that works well (in the sense that it produces a test with power within ε , where ε is a small

pre-specified value) in several problems when the dimension of θ is small, and we implement their algorithm here.

5.3.2 High Dimensional Nuisance Parameter

The dimension of θ can be very large in our problem: even when $r = k = 1$, the model with unrestricted stochastic trend leads to θ of dimension $q^2 + q(q + 1)/2$ so that θ contains 222 elements when $q = 12$. Approximating the least upper power bound directly then becomes numerically intractable. This motivates a computationally practical method for computing a *low* (as oppose to *least*) upper power bound.

The method restricts Λ so that it is degenerate with all mass on a single point, say θ^* , which is chosen so that the null distribution of the maximal invariant of Theorem 1 is close to its distribution under the alternative. Intuitively, this should make it difficult to distinguish the null from the alternative hypothesis, and thus lead to a low power bound. Also, this choice of θ^* ensures that draws from the null model look empirically reasonable, as they are nontrivial to distinguish from draws of the alternative with an $I(1)$ stochastic trend.

Since the density of the maximal invariant is quite involved, θ^* is usefully approximated by a choice that makes the multivariate normal distribution of $\text{vec}(Y, X)$ under the null close to its distribution under the alternative, as measured by a convenient metric. We choose θ^* to minimize the Kullback-Leibler divergence (KLIC) between the null and alternative distributions. Since the bounds from Lemmas 1 and 2 are valid for any mixture, numerical errors in the KLIC minimization do not invalidate the resulting bound. Details are provided in the appendix.

5.4 Numerical Bounds

Table 1 shows numerical results for power and size bounds for 5% level tests with $q = 12$. Results are shown for $r = k = 1$ (panel A), $r = 1$ and $k \geq 2$ (panel B), and $r = 2$ and $k = 1$ (Panel C).⁸ Numerical results for larger values of $n = r + k$ are not reported because of the large number of calculations required to evaluate the density in large models.

Power depends on the values of B and R under the alternative, and results are presented for various values of these parameters. Because of invariance, when $r = 1$ (as in panels A and B), or $k = 1$ (as in panel C), the distribution of the maximal invariant depends on B and R

⁸The results shown in panel B were computed using the KLIC minimized value of θ for the model with $r = 1$ and $k = 2$. The appendix shows the resulting bounds are valid for $k \geq 2$.

only through $\|B\|$, $\|R\|$, and, if $\|R\| > 0$, on $\text{tr}(R'B)/(\|B\| \cdot \|R\|)$. Thus, in panel A, where $r = k = 1$, results are shown for two values of $|B|$, three values of $|R|$ and for $R \cdot B < 0$ and $R \cdot B > 0$, while panels B and C show results for three values of $\omega = \text{tr}(R'B)/(\|B\| \cdot \|R\|)$ when $\|R\| > 0$. All of the results in Table 1 use the KLIC minimized values of θ as described in the last subsection. Table 2 compares this KLIC-based bounds to to the numerical least upper power bounds when the parameter space is sufficiently small to allow calculation of the numerical least upper bounds.

To understand the formatting of Table 1, look at panel A. The panel contains italicized and non-italicized numerical entries. The non-italicized numbers are power bounds, and the italicized numbers are size bounds. The first column in the table shows the trend specification allowed under H_0 . The first entry, labelled “*unr*” corresponds to the unrestricted trend specification in (3) and the other entries correspond to the restricted trend processes discussed in Section 2. Because $r = k = 1$, there are no restrictions imposed by the assumption that $H(s, t) = G(s, t)S_v$ or that G is diagonal, so these models are not listed in panel A. Stationarity ($G(s, t) = G(s - t)$) is a restriction, and this is the second entry in the first column. The final two entries correspond to the local-to-unity (“*LTU*”) and $I(1)$ restrictions. The numerical entries shown in the rows corresponding to these trend models are the power bounds. For example, the non-italicized entries in the first numerical column show power bounds for $|R| = 0$ and $|B| = 7$, which are 0.36 for the unrestricted null, 0.41 when the trend is restricted to be stationary, 0.50 when the trend is restricted to follow a local-to-unity process, and 0.50 when the trend is further restricted to follow an $I(1)$ process.

The second column of panel A shows the auxiliary null hypotheses H_L , corresponding to the null hypothesis H_0 , shown in the first column. The entries under H_L represent less restrictive models than H_0 . For example, when H_0 restricts the null to be stationary, an unrestricted trend process (“*unr*”) is shown for H_L , while when H_0 restricts the trend to be $I(1)$, the less restrictive local-to-unity, stationary, and unrestricted nulls are listed under H_L . The numerical entries for these rows (shown in italics in the table) are the lower size bounds for H_L for 5% level tests under H_0 and with power that is 3 percentage points less than the corresponding power bound shown in the table. For example, from the first numerical column of panel A, the power bound for the $I(1)$ version of H_0 is 0.50. For any test with size no larger than 5% under this null and with power of at least 0.47 ($= 0.50 - 0.03$), the size under a null that allows an unrestricted trend (“*unr*” under H_L) is at least 12%, the size under a null that restricts the trend to be stationary is at least 8%, and the size under a null that restricts the trend to follow a local-to-unity process is at least 4%.

Looking at the entries in Panel A, two results stand out. First, and not surprisingly, restricting tests so that they control size for the unrestricted trend process leads to a non-negligible reduction in power. For example, when $|B| = 7$, and $R = 0$, the power bound is 0.36, for tests that control size for unrestricted trends, the bound increases to 0.41 for tests that control size for stationary trends, and increases to 0.50 for tests that control size for local-to-unity or $I(1)$ trend processes. Second, whenever there is a substantial increase in power associated with restricting the trend process, there are large size distortions under the null hypothesis without this restriction. For example, Elliott's (1998) observation that efficient tests under the $I(1)$ trend have large size distortions under a local-to-unity process is evident in the table. From the table, when $|B| = 7$, $|R| = 0.9$, and $R \cdot B > 0$, the power bound for the null with an $I(1)$ trend is 0.95, but any test that controls size for this null and has power of at least 0.92 will have size that is greater than 0.50 when the trend is allowed to follow a local-to-unity process. However, addressing Elliott's (1998) concern by controlling for size in the local-to-unity model, as in the analysis of Stock and Watson (1996) or Jansson and Moreira (2006) does not eliminate the fragility of the test. For example, with the same values of B and R , the power bound for the null that allows for a local-to-unity trend is 0.67, but any test that controls size for this null and has a size of at least 0.64 will have a size greater than 0.32 when the trend is unrestricted.

Panels B ($r = 1$ and $k = 2$) and C ($r = 2$ and $k = 1$) show qualitatively similar results. Indeed these panels show even more fragility of tests that do not allow for general trends. For example, the lower size bound for the unrestricted trend null exceeds 0.50 in several cases for tests that restrict trends to be $I(1)$, local-to-unity, or stationary.

When $r = k = 1$, it is feasible to approximate the least upper power bound for the $I(1)$ and local-to-unity trend restrictions using the method described in subsection 5.3.1. As described in the appendix, the approximate least upper bounds (LUB) in Table 2 are no more than 2.5 percentage points above the actual least upper bound, apart from Monte Carlo error. The differences with the KLIC minimized power bounds are small, suggesting that the bounds in Table 1 are reasonably tight.

6 Efficient Y -Only Tests

The primary obstacle for constructing efficient tests of the null hypothesis that $B = 0$ is the large number of nuisance parameters associated with the stochastic trend (the parameters that determine $H(s, t)$). These parameters govern the values of Σ_{ZV} and Σ_{VV} , which in turn

determine Σ_{YX} and Σ_{XX} . Any valid test must control size over all values of these nuisance parameters. Wright (2000) notes that this obstacle can be avoided by ignoring the x_t data and basing inference only on y_t , since under the null hypothesis, $y_t = z_t$. This section takes up Wright’s suggestion and discusses efficient low-frequency “Y-only” tests.⁹

We have two related goals. The first is to study the power properties of these tests relative to the power bounds computed in the last section. As it turns out, when $r = 1$ (so there is only a single cointegrating vector), this Y-only test essentially achieves the power bound, so the test efficiently uses all of the information in Y and X . Given the efficiency property of the Y-only test, the second goal is to develop simple formulae for implementing the test. We discuss these in reverse order, first deriving a convenient formulae for the test statistic and then studying the power of the resulting test.

6.1 Efficient Tests against General Alternatives

The distribution of $\text{vec } Y \sim \mathcal{N}(0, \Sigma_{YY})$ follows from the derivations in Section 3: Under the null hypothesis, $\Sigma_{YY} = I_{rq}$, and under the alternative, Σ_{YY} depends on the local alternative B , the properties of the stochastic trend and its relationship with the error correction term Z . For a particular choice of alternative, the testing problem thus becomes $H_0 : \Sigma_{YY} = I_{rq}$ against $H_1 : \Sigma_{YY} = \Sigma_{YY1}$, and the invariance requirement (21) becomes

$$Y \rightarrow YA'_{yy} \text{ for arbitrary nonsingular } r \times r \text{ matrices } A_{yy}. \quad (23)$$

The density of the maximal invariant is given in the following theorem.

Theorem 2 (a) *If $\text{vec } Y \sim \mathcal{N}(0, \Sigma_{YY})$ with positive definite Σ_{YY} and $q > r$, the density of a maximal invariant to (23) has the form*

$$c_1(\det \Sigma_{YY})^{-1/2}(\det \Omega_Y)^{-1/2}E_{\omega_Y}[|\det(\omega_Y)|^{q-r}]$$

where c_1 does not depend on Σ_{YY} , ω_Y is an $r \times r$ random matrix with $\text{vec } \omega_Y \sim \mathcal{N}(0, \Omega_Y^{-1})$, $\Omega_Y = (I_r \otimes Y)' \Sigma_{YY}^{-1} (I_r \otimes Y)$, and E_{ω_Y} denotes integration with respect to the distribution of ω_Y (conditional on Y).

⁹Wright (2000) implements this idea using a “stationarity” test of the $I(0)$ null proposed by Saikkonen and Luukkonen (1993), using a robust covariance matrix as in Kwiatkowski, Phillips, Schmidt, and Shin (1992) for the test proposed in Nyblom (1989). This test relies on a consistent estimator of the spectral density matrix of z_t at frequency zero. But consistent estimation requires a lot of pertinent low frequency information, and lack thereof leads to well-known size control problems (see for example, Kwiatkowski, Phillips, Schmidt, and Shin (1992), Caner and Kilian (2001), and Müller (2005)). These problems are avoided by using the low-frequency components of y_t only; see Müller and Watson (2008) for further discussion.

(b) If in addition, $\Sigma_{YY} = \tilde{V}_{YY} \otimes \tilde{\Sigma}_{YY}$ where \tilde{V}_{YY} is $r \times r$ and $\tilde{\Sigma}_{YY}$ is $q \times q$, then the density simplifies to

$$c_2(\det \tilde{\Sigma}_{YY})^{-r/2} \det(Y' \tilde{\Sigma}_{YY}^{-1} Y)^{-q/2}$$

where c_2 does not depend on Σ_{YY} .

As in Theorem 1, part (a) of this theorem provides a formula for the density of a maximal invariant in terms of absolute moments of the determinant of a multivariate normal matrix with a covariance matrix that depends on the data. Part (b) provides an explicit and simple formula when the covariance matrix is of a specific Kronecker form. This form arises under the null hypothesis with $\Sigma_{YY} = I_{rq}$, and under alternatives where each of the r putative error correction terms in y_t have the same low-frequency covariance matrix. For a simple alternative hypothesis with $\Sigma_{YY1} = \tilde{V}_{YY1} \otimes \tilde{\Sigma}_{YY1}$, the best test then rejects for large values of $\det(Y'Y)/\det(Y' \tilde{\Sigma}_{YY1}^{-1} Y)$. The form of weighted average power maximizing tests over a set of alternative covariance matrices Σ_{YY1} are also easily deduced from Theorem 2 parts (a) and (b).

6.2 Efficient Tests against $I(1)$ Alternatives

As discussed above, the numerical results in this paper focus on the benchmark alternative where the stochastic trend follows an $I(1)$. Under this alternative, y_t follows a multivariate “local level model” (cf. Harvey (1989)), which is the alternative underlying well-known “stationarity” tests such as Nyblom and Mäkeläinen (1983), Nyblom (1989), Kwiatkowski, Phillips, Schmidt, and Shin (1992), Nyblom and Harvey (2000), Jansson (2004), and others. Thus, suppose that the stochastic trend satisfies (6) and (10), so that

$$T^{-1/2} \sum_{t=1}^{\lfloor sT \rfloor} y_t \Rightarrow W_z(s) + B \int_0^s W_v(t) dt. \quad (24)$$

The optimal test therefore depends on the value of B and the correlation matrix $R = S_z S'_v = E[W_z(1)W_v(1)']$ under the alternative. The invariance constraint (23), as well as the properties of Σ_{YY} arising under (24) lead to certain cancellations and simplifications, that we discuss in detail in the appendix. The end-result is the following Corollary to Theorem 2.

Corollary 1 *Let*

$$LFST(b) = \det(Y'Y)/\det(Y'(I_q + b^2 D)^{-1} Y) \quad (25)$$

for some scalar $b > 0$, where $D = \text{diag}(\pi^{-2}, (2\pi)^{-2}, \dots, (q\pi)^{-2})$.

(a) When $r = 1$, the test that rejects for large values of $\text{LFST}(b)$ is uniformly most powerful against all alternatives (24) with $\|B\| = b$ and any value of R .

(b) When $1 < r \leq k$, the test that rejects for large values of $\text{LFST}(b)$ is uniformly most powerful against all alternatives (24) with B and R such that the r eigenvalues of BB' are equal to b^2 , and there exists an orthogonal $k \times k$ matrix Q_k for which BQ_kR' is a symmetric $r \times r$ matrix.

The statistic (25) was developed as a "low-frequency stationarity test" and denoted "LFST" in Müller and Watson (2008) and we continue to use that label here. The corollary suggests that the LFST is a powerful Y -only test for inference about the cointegrating vector in a variety of cases. When $r > k$, the point optimal test statistic under similar conditions as described in part (b) of the Corollary is given by $\xi(b)$ derived in the appendix. In contrast to $\text{LFST}(b)$, the statistic $\xi(b)$ does not have a simple closed-form expression, as it involves the form of the density of the maximal invariant derived part (a) of Theorem 2.

As a practical matter, it is much more convenient to continue to rely on $\text{LFST}(b)$ with $Y = Y_T$ even if $r > k$, although at the cost of somewhat reduced power, as discussed in the next subsection. Table 3 presents 10%, 5%, and 1% critical values for the point-optimal $\text{LFST}(10/\sqrt{r})$ for various values of r and q , where the alternative $b = 10/\sqrt{r}$ is chosen so that 5% test has approximately 50% power (cf. King (1988)). Since the low-frequency transformation of the putative error correction terms Y_T are linear functions of the hypothesized cointegrating vector(s), the boundaries of the confidence set obtained by inverting the point-optimal LFST (that is, the values of β_0 that do not lead to a rejection) simply solve a quadratic equation. The finding of an empty confidence set may be interpreted as evidence against the notion that the system satisfies $r \geq 1$ cointegrating relationships.

6.3 Power of Efficient Y -only Tests

Table 4 shows the power of the point-optimal LFST and the corresponding power envelope of the Y -only test for $r = 1$ in panel A, and $r = 2$ and $k = 1$ in panel B. In panel A, the power envelope is given by the $\text{LFST}(b)$ evaluated at the value of $b = B$ under the alternative, while the point-optimal LFST is evaluated at $b = 10$. The power of the point-optimal test is very close to the Y -only power envelope. A more interesting comparison involves the power of the point-optimal LFST with the (Y, X) -power bound computed in the last section. Because the Y -only tests control size for any trend process, the relevant comparison is the unrestricted

$H(s, t)$ bound shown in panels A and B of Table 1. The power of the point-optimal LFST differs from the Table 1 power bound by no more than 0.01 when $|B| = 7$ and no more than 0.03 when $|B| = 14$. Thus, for all practical purposes, the point-optimal LFST corresponds to an efficient test in models with a single cointegrating vector ($r = 1$).

The results in panel B of Table 4 are somewhat different. Here, because $r > k = 1$, the Y -only power envelope is computed using the $\xi(\|B\|)$ statistic derived in the appendix, and not by the LFST. The numerical results show that the relative power of the LFST depends on both B and R , and the loss can be large when B and R are large and orthogonal ($\omega = 0$). Comparing results in panel B of Table 4 to the corresponding results in panel C of Table 1, shows that in this case there are potentially additional power gains associated with using data on both Y and X . For example, when $\|B\| = 20$, $\|R\| = 0.9$, and $\omega = 0$, the Y -only power envelope is 0.86, while the (Y, X) power bound is 0.94.

7 Interest Rates and the Term Spread

Figure 1 plots monthly values of the 1 year and 10 year U.S. Treasury bond rates ($R^{1\text{Year}}$ and $R^{10\text{Year}}$) from 1953:4-2011:12. Evidently, both series are persistent, but over long periods they closely move together suggesting a common source of long-run variability. Figure 2 plots the monthly values of the term spread, the difference between $R^{10\text{Year}}$ and $R^{1\text{Year}}$. This series appears to be less persistent than the levels of the interest rates in figure 1. Taken together, the figures suggest $R^{10\text{Year}}$ and $R^{1\text{Year}}$ are cointegrated with a cointegrating coefficient that is close to unity.

The first two rows of Table 5 show statistics that summarize the persistence in the levels of $R^{1\text{Year}}$ and $R^{10\text{Year}}$. The DFGLS test does not reject the null hypothesis of a unit root in the series, and the 95% confidence interval for the largest AR root constructed by inverting the Dickey-Fuller statistic (Stock (1991)), yields a range values between 0.98 and 1.00 for both interest rates. The usual Nyblom/KPSS (cf. Nyblom (1989) and Kwiatkowski, Phillips, Schmidt, and Shin (1992)) statistic rejects the $I(0)$ null with a p -value of less than 1%. The estimates of the long-memory fractional integration parameter d constructed from Geweke and Porter-Hudak (1983) regressions (implemented as described in Robinson (2003)) suggest values of d that differ from $d = 0$ (from levels versions of the regression), but also differ from $d = 1$ (from the differences version of the regression). Finally, the LFST with q chosen to capture below business cycle frequency (periods greater than 8 years) rejects the $I(0)$ null for the levels of the two series. Taken together these results suggest that the levels of the

interest rates are highly persistent, but with a stochastic trend process that might well differ from the exact $I(1)$ model.

Figure 2 suggests a value of the cointegrating coefficient that is equal to unity, but other values are equally plausible. To take one example, consider a simple version of the expectations theory of the term structure in which $R_t^{10\text{Year}} = \frac{1}{10} \sum_{i=0}^9 R_{t+i|t}^{1\text{Year}} + e_t$, where $R_{t+i|t}^{1\text{Year}}$ is the time t expectation of $R_{t+i}^{1\text{Year}}$ and e_t is an $I(0)$ term premium. Suppose that $R_{t+i}^{1\text{Year}}$ follows an $AR(1)$ model with coefficient ρ . In this case $R_t^{10\text{Year}} = \beta(\rho)R_t^{1\text{Year}} + e_t$ with $\beta(\rho) = 1$ when $\rho = 1$ and $\beta(\rho) = \frac{1}{10} \frac{1-\rho^{10}}{1-\rho}$ for $\rho \neq 1$. When ρ is close to unity, the interest levels are highly persistent, but $R_t^{10\text{Year}} - \beta(\rho)R_t^{1\text{Year}}$ is not, so that $\beta(\rho)$ is the cointegrating coefficient. Note that $\beta(\rho)$ can sharply differ from $\beta = 1$ even when ρ is close unity. For example, $\beta(0.99) = 0.96$ and $\beta(0.98) = 0.91$.¹⁰

Estimating β using using Johansen's ML estimator in VECM with 12 lags produces $\hat{\beta} = 0.96$ with a standard error of 0.06. Thus, this $I(1)$ estimator does not reject the null that $\beta = 1$. Given the uncertainty about the exact nature of the trend process, it is useful to ask whether this result is robust to alternative models. This is investigated in the final row of Table 5, which summarizes the persistence properties of the term spread. The DFGLS test rejects the unit root null, although the 95% confidence interval for the largest root still suggests considerable persistence. This persistence is also evident in the GPH regression estimate of d , and the small p -value of the Nyblom/KPSS statistic. However, the LFST does not reject the $I(0)$ null and so fails to reject the null hypothesis that $\beta = 1$. The low-frequency behavior of the term spread is thus consistent with the $I(0)$ model, that is the data does not reject the notion that all persistence and other non-i.i.d. dynamics of the term spread are associated with variations at business cycle or higher frequencies. Non-trivial business cycle dynamics are, of course, highly plausible for the term spread, so that the low-frequency focus of the LFST becomes a potentially attractive notion of cointegration in this setting. In contrast, the more traditional measures in Table 5 that indicate some degree of deviation from the $I(0)$ model all depend on a relatively wider range of frequencies.¹¹

Finally, while a value of $\beta = 1$ is not rejected using standard $I(1)$ inference or low-frequency inference that is robust to the $I(1)$ trend specification, there are a range of other values that are not rejected. The 95% confidence interval for β constructed using the standard $I(1)$ estimates ranges from 0.83 to 1.08. The coverage of this confidence interval is

¹⁰Valkanov (2003) uses this observation in a local-to-unity framework to estimate the local-to-unity parameter by inverting an estimate of the cointegrating coefficient $\beta(\rho)$.

¹¹For instance, the GPH regressions leading to a significant estimate of d rely on a reasonably large number of periodogram ordinates ($[T^{0.65}] = 71$), which contain information about periodicities as low as 10 months.

suspect, though, because of its reliance on the $I(1)$ trend specification. Inverting the LFST statistic produces a confidence interval that ranges from 0.85 to 1.31. Thus the robust procedure yields a similar lower bound for the confidence interval, but a higher upper bound.

A Appendix

A.1 Proof of Theorem 1

Write $Y = (Y'_1, Y'_2, Y'_3)'$ and $X = (X'_1, X'_2, X'_3)'$, where Y_1 and X_1 have r rows, and Y_2 and X_2 have k rows. Consider the one-to-one mapping $h : \mathbb{R}^{q \times n} \mapsto \mathbb{R}^{q \times n}$ with

$$h(Y, X) = Q = \begin{pmatrix} Q_{Y1} & Q_{X1} \\ Q_{Y2} & Q_{X2} \\ Q_{Y3} & Q_{X3} \end{pmatrix} = \begin{pmatrix} Y_1 & Y_1^{-1}X_1 \\ Y_2(Y_1)^{-1} & X_2 - Y_2Y_1^{-1}X_1 \\ Y_3(Y_1)^{-1} & (X_3 - Y_3Y_1^{-1}X_1)(X_2 - Y_2Y_1^{-1}X_1)^{-1} \end{pmatrix}.$$

A straightforward calculation shows that $(\text{vec } Q'_{Y2}, \text{vec } Q'_{Y3}, \text{vec } Q'_{X3})$ is a maximal invariant to (21). The inverse of h is given by

$$h^{-1}(Q) = \begin{pmatrix} Q_{Y1} & Q_{Y1}Q_{X1} \\ Q_{Y2}Q_{Y1} & Q_{X2} + Q_{Y2}Q_{Y1}Q_{X1} \\ Q_{Y3}Q_{Y1} & Q_{X3}Q_{X2} + Q_{Y3}Q_{Y1}Q_{X1} \end{pmatrix}.$$

Using matrix differentials (cf. Chapter 9 of Magnus and Neudecker (1988.)), a calculation shows that the Jacobian determinant of h^{-1} is equal to $(\det Q_{Y1})^{q-r+k}(\det Q_{X2})^{q-k-r}$. The density of Q is thus given by

$$(2\pi)^{-qn/2}(\det \Sigma_{(Y,X)})^{-1/2} |\det Q_{Y1}|^{q-r+k} |\det Q_{X2}|^{q-k-r} \exp[-\frac{1}{2}(\text{vec } h^{-1}(Q))' \Sigma_{(Y,X)}^{-1} (\text{vec } h^{-1}(Q))]$$

and we are left to integrate out Q_{Y1} , Q_{X1} and Q_{X2} to determine the density of the maximal invariant.

Now consider the change of variables from Q_{Y1} , Q_{X1} , Q_{X2} to ω_Y , ω_X and ω_{YX}

$$\begin{aligned} Q_{Y1} &= Y_1 \omega_Y \\ Q_{X1} &= \omega_Y^{-1} Y_1^{-1} X_1 \omega_X - \omega_Y^{-1} \omega_{YX} \\ Q_{X2} &= (X_2 - Y_2 Y_1^{-1} X_1) \omega_X \end{aligned}$$

with Jacobian determinant $(\det Y_1)^r (\det(X_2 - Y_2 Y_1^{-1} X_1))^k \det(-\omega_Y)^{-k}$. Noting that with this change, $h^{-1}(Q) = (Y \omega_Y, X \omega_X - Y \omega_{YX})$, we find that the density of the maximal invariant is equal to

$$\begin{aligned} &\int (2\pi)^{-qn/2} (\det \Sigma_{(Y,X)})^{-1/2} |\det Y_1|^{q+k} |\det(X_2 - Y_2 Y_1^{-1} X_1)|^{q-r} |\det \omega_Y|^{q-r} |\det \omega_X|^{q-k-r} \\ &\cdot \exp[-\frac{1}{2} \text{vec}(Y \omega_Y, X \omega_X - Y \omega_{YX})' \Sigma_{(Y,X)}^{-1} \text{vec}(Y \omega_Y, X \omega_X - Y \omega_{YX})] d(\text{vec } \omega'_Y, \text{vec } \omega'_X, \text{vec } \omega'_{YX})'. \end{aligned}$$

Since $\text{vec}(Y\omega_Y, X\omega_X - Y\omega_{YX}) = D_{YX} \text{vec}(\omega_Y, \omega_X) - V_{0Y} \text{vec}(\omega_{YX})$, we have

$$\begin{aligned} & \text{vec}(Y\omega_Y, X\omega_X - Y\omega_{YX})' \Sigma_{(Y,X)}^{-1} \text{vec}(Y\omega_Y, X\omega_X - Y\omega_{YX}) \\ &= \text{vec}(\omega_Y, \omega_X)' D_{YX}' \Sigma_{(Y,X)}^{-1} D_{YX} \text{vec}(\omega_Y, \omega_X) \\ & \quad - 2 \text{vec}(\omega_Y, \omega_X) D_{YX}' \Sigma_{(Y,X)}^{-1} V_{0Y} \text{vec}(\omega_{YX}) + \text{vec}(\omega_{YX})' V_{0Y}' \Sigma_{(Y,X)}^{-1} V_{0Y} \text{vec}(\omega_{YX}). \end{aligned}$$

The result now follows from integrating out ω_{YX} by 'completing the square'.

A.2 Proof of Theorem 2

The proof to part (a) mimics the proof to Theorem 1 and is omitted. To prove part (b), note that because of invariance, we can set $\tilde{V}_{YY} = I_r$ without loss of generality, so that $\det \Sigma_{YY} = (\det \tilde{\Sigma}_{YY})^r$, $\Omega_Y = (I_r \otimes Y' \tilde{\Sigma}_{YY}^{-1} Y)$ and $(\det \Omega_Y)^{-1/2} = \det(Y' \tilde{\Sigma}_{YY}^{-1} Y)^{-r/2}$. Since $(\text{vec } \omega_Y)' \Omega_Y (\text{vec } \omega_Y) = \text{tr}(\omega_Y' Y' \tilde{\Sigma}_{YY}^{-1} Y \omega_Y)$, the density in part (a) of the Theorem becomes proportional to

$$(\det \tilde{\Sigma}_{YY})^{-r/2} \det(Y' \tilde{\Sigma}_{YY}^{-1} Y)^{-r/2} \int |\det \omega_Y|^{q-r} \exp[-\frac{1}{2} \text{tr}(\omega_Y' Y' \tilde{\Sigma}_{YY}^{-1} Y \omega_Y)] d(\text{vec } \omega_Y).$$

Let $\tilde{\omega}_Y = (Y' \tilde{\Sigma}_{YY}^{-1} Y)^{1/2} \omega_Y$, so that $|\det \omega_Y|^{q-r} = \det(Y' \tilde{\Sigma}_{YY}^{-1} Y)^{-(q-r)/2} |\det \tilde{\omega}_Y|^{q-r}$ and $\text{vec } \omega_Y = (I_r \otimes (Y' \tilde{\Sigma}_{YY}^{-1} Y)^{-1/2}) \text{vec } \tilde{\omega}_Y$, and the Jacobian determinant of the transformation from ω_Y to $\tilde{\omega}_Y$ is $\det(I_r \otimes (Y' \tilde{\Sigma}_{YY}^{-1} Y)^{-1/2}) = \det(Y' \tilde{\Sigma}_{YY}^{-1} Y)^{-r/2}$. Thus, the density is proportional to

$$(\det \tilde{\Sigma}_{YY})^{-r/2} \det(Y' \tilde{\Sigma}_{YY}^{-1} Y)^{-q/2} \int |\det \tilde{\omega}_Y|^{q-r} \exp[-\frac{1}{2} \text{tr}(\tilde{\omega}_Y' \tilde{\omega}_Y)] d(\text{vec } \tilde{\omega}_Y),$$

and the result follows.

A.3 Proof of Lemma 3

We first establish a preliminary result.

Lemma 4 *For any $t > 0$ and integer κ , the functions $\Psi_l : [0, t] \mapsto \mathbb{R}$ with $\Psi_l(s) = \sqrt{2} \cos(\pi l s)$, $l = 1, \dots, \kappa$ are linearly independent.*

Proof. Choose any real constants c_j , $j = 1, \dots, \kappa$, so that $\sum_{j=1}^{\kappa} c_j \Psi_j(s) = 0$ for all $s \in [0, t]$. Then also $\sum_{j=1}^{\kappa} c_j \Psi_j^{(i)}(0) = 0$ for all $i > 0$, where $\Psi_j^{(i)}(0)$ is the i th (right) derivative of Ψ_j at $s = 0$. A direct calculation shows $\Psi_j^{(i)}(0) = (-1)^{i/2} \sqrt{2} (\pi j)^i$ for even i . It is not hard to see that the $\kappa \times \kappa$ matrix with j, i th element $(-1)^{i/2} (\pi j)^i$ is nonsingular, so that $\sum_{j=1}^{\kappa} c_j \Psi_j^{(i)}(0) = 0$ for $i = 2, 4, \dots, 2\kappa$ can only hold for $c_j = 0$, $j = 1, \dots, \kappa$. ■

For the proof Lemma 3 we construct $H(s, t)$ such $\text{vec } Z = \int_0^1 (I_r \otimes \Psi(t)) S_z dW(t)$ and $\text{vec } V = \int_0^1 \int_t^1 (H(s, t) \otimes \Psi(s)) ds dW(t)$ have the specified covariance matrix. The proof of the slightly more difficult part (b), where $H(s, t) = G(s, t) S_v$, is based on the following observations:

- (i) Ignoring the restriction on the form of $\text{vec } V$, it is straightforward to construct an appropriate multivariate normal vector $\text{vec } V$ from a linear combination of $\text{vec } Z$ and ζ , where $\zeta \sim \mathcal{N}(0, I_{kq \times kq})$ independent of Z .
- (ii) Suppose that $R = S$ was allowed, where $S = (I_r, 0_{r \times (k-r)})$. Then $S_z = S S_v$, $\text{vec } Z \sim \int_0^1 F_Z(t) S_v dW(t)$ for $F_Z(t) = S \otimes \Psi(t)$, and one can also easily construct ζ as in (i) via $\zeta = \int_0^1 F_\zeta(t) S_v dW(t)$ by an appropriate choice of F_ζ . Since Ito-Integrals are linear, one could thus write $\text{vec } V = \int_0^1 F(t) S_v dW(t)$ with F a linear combination of F_Z and F_ζ , using observation (i).
- (iii) For any matrix function $F : [0, 1] \mapsto \mathbb{R}^{kq \times k}$ that is equal to zero on the interval $(1 - \varepsilon, 1]$ for some $\varepsilon > 0$, one can set $G(s, t) = (I_k \otimes \Psi(s)' J(t)^{-1}) F(t)$, where $J(t) = \int_t^1 \Psi(s) \Psi(s)' ds$ and obtain $\int_0^1 \int_t^1 (G(s, t) \otimes \Psi(s)) ds S_v dW(t) = \int_0^1 F(t) S_v dW(t)$, since for any matrix A with k rows and vector v , $A \otimes v = (I_k \otimes v) A$.

The following proof follows this outline, but three complications are addressed: $R = S$ is not allowed; the matrix function F needs to be zero on the interval $(1 - \varepsilon, 1]$, which does not happen automatically in the construction in (ii); one must verify that the process $\int_0^s G(s, t) S_v dW(t)$ admits a cadlag version.

Set S_z to be the first r rows of I_n . Since $\Psi_l(1 - s) = (-1)^l \Psi_l(s)$ for all $l \geq 1$, Lemma 4 implies that $J(t) = \int_t^1 \Psi(s) \Psi(s)' ds$ and $I_q - J(t)$ are nonsingular for any $t \in (0, 1)$. The $rq \times 1$ random vector $\text{vec } Z_\varepsilon = \int_0^{1-\varepsilon} (S_z \otimes \Psi(s)) dW(s)$ thus has nonsingular covariance matrix $I_r \otimes \Sigma_q^\varepsilon$, where $\Sigma_q^\varepsilon = I_q - J(1 - \varepsilon)$. Also, since

$$\Sigma^* = \begin{pmatrix} I_r \otimes I_q & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}$$

is positive definite, so is $I_{rq} - \Sigma_{12} \Sigma_{22}^{-1} \Sigma_{21}$, so that we can choose $0 < \varepsilon < 1$ such that $I_r \otimes \Sigma_q^\varepsilon - \Sigma_{12} \Sigma_{22}^{-1} \Sigma_{21}$ is positive definite. With that choice of ε , also $\Sigma_{22} - \Sigma_{21} (I_r \otimes \Sigma_q^\varepsilon)^{-1} \Sigma_{12}$ is positive definite.

For part (a) of the lemma, define the $[0, 1] \mapsto \mathbb{R}^{kq \times n}$ function $F_a(t) = A_a (I_n \otimes \Psi(t))$, where $A_a = (A_{a1}, A_{a2})$ with $A_{a1} = \Sigma_{21} (I_r \otimes \Sigma_q^\varepsilon)^{-1}$ and $A_{a2} = (\Sigma_{22} - \Sigma_{21} (I_r \otimes \Sigma_q^\varepsilon)^{-1} \Sigma_{12})^{1/2} (I_k \otimes (\Sigma_q^\varepsilon)^{-1/2})$.

For part (b) of the lemma, choose $0 < \rho < 1$ so that $\Sigma_{22} - \rho^{-2} \Sigma_{21} (I_r \otimes \Sigma_q^\varepsilon)^{-1} \Sigma_{12}$ is positive definite. Choose S_v to be the first k rows of I_n multiplied by ρ , so that $R = S_z S_v' = \rho S$. Let $\tilde{\Psi}_1(s)$ be scaled residuals of a continuous time projection of $\mathbf{1}[s \leq 1 - \varepsilon] \Psi_{q+1}(s)$ on $\{\mathbf{1}[s \leq 1 - \varepsilon] \Psi_l(s)\}_{l=1}^q$ on the unit interval, and let $\tilde{\Psi}_j(s)$, $j = 2, \dots, kq$ be the scaled residuals of continuous time projection

of $\mathbf{1}[s \leq 1 - \varepsilon]\Psi_{q+j}(s)$ on $\{\mathbf{1}[s \leq 1 - \varepsilon]\Psi_l(s)\}_{l=1}^q$ and $\{\mathbf{1}[s \leq 1 - \varepsilon]\tilde{\Psi}_l(s)\}_{l=1}^{j-1}$. By Lemma 4, $\tilde{\Psi}_j(s)$, $j = 1, \dots, kq$, are not identically zero, and we can choose their scale to make them orthonormal. Let $\tilde{\Psi}(s) = (\tilde{\Psi}_1(s), \dots, \tilde{\Psi}_{kq}(s))'$, the $k \times 1$ vector $\iota_k = (1, 0, \dots, 0)'$, and $A_b = (A_{b1}, A_{b2})$ with $A_{b1} = \rho^{-1}\Sigma_{21}(I_r \otimes \Sigma_q^\varepsilon)^{-1}$ and $A_{b2} = (\Sigma_{22} - \rho^{-2}\Sigma_{21}(I_r \otimes \Sigma_q^\varepsilon)^{-1}\Sigma_{12})^{1/2}$. Now define the $[0, 1] \mapsto \mathbb{R}^{kq \times n}$ function

$$F_b(t) = A_b \begin{pmatrix} S \otimes \Psi(t) \\ \iota'_k \otimes \tilde{\Psi}(t) \end{pmatrix} S_v.$$

For both parts, that is for $i \in \{a, b\}$, set

$$H_i(s, t) = (I_k \otimes \Psi(s)'J(t)^{-1})F_i(t) \text{ for } t \in [0, 1 - \varepsilon]$$

and $H_i(s, t) = 0$ otherwise. With this choice

$$\begin{aligned} \text{vec } V_i &= \int_0^1 \int_t^1 (H_i(s, t) \otimes \Psi(s)) ds dW(t) \\ &= \int_0^{1-\varepsilon} \int_t^1 ((I_k \otimes \Psi(s)\Psi(s)'J(t)^{-1})F_i(t)) ds dW(t) \\ &= \int_0^{1-\varepsilon} F_i(t) dW(t). \end{aligned}$$

Thus

$$\begin{aligned} E[(\text{vec } V_i)(\text{vec } V_i)'] &= \int_0^{1-\varepsilon} F_i(t)F_i(t)' dt \\ E[(\text{vec } V_i)(\text{vec } Z)'] &= \int_0^{1-\varepsilon} F_i(t)(S_z \otimes \Psi(t))' dt \end{aligned}$$

since $\text{vec}(Z - Z_\varepsilon) = \int_{1-\varepsilon}^1 (I_r \otimes \Psi(t))S_z dW(t)$ is independent of $\text{vec } V_i$. A direct calculation now shows that $\int_0^{1-\varepsilon} F_a(t)F_a(t)' dt = A_a(I_n \otimes \Sigma_q^\varepsilon)A_a'$, $\int_0^{1-\varepsilon} F_a(t)(S_z \otimes \Psi(t))' dt = A_a(S_z' \otimes \Sigma_q^\varepsilon)$, $\int_0^{1-\varepsilon} F_b(t)F_b(t)' dt = A_b \text{diag}(I_r \otimes \Sigma_q^\varepsilon, I_{kq})A_b'$ and $\int_0^{1-\varepsilon} F_b(t)(S_z \otimes \Psi(t))' dt = \rho A_b(S_z' \otimes \Sigma_q^\varepsilon)$, so that from the definitions of A_i , $E[(\text{vec } V_i)(\text{vec } V_i)'] = \Sigma_{22}$ and $E[(\text{vec } V_i)(\text{vec } Z)'] = \Sigma_{21}$, as required.

It thus remains to show that the processes $\int_0^s H_i(s, t) dW(t)$, $i \in \{a, b\}$, admit a cadlag version.

Recall that $\|A\|$ is the Frobenius norm of the real matrix A , $\|A\| = \sqrt{\text{tr } A'A}$, which is submultiplicative. If $v \sim \mathcal{N}(0, \Sigma)$, then $E[\|v\|^4] = E[(v'v)^2] = 2 \text{tr}(\Sigma^2) + (\text{tr } \Sigma)^2 \leq 3(\text{tr } \Sigma)^2$, so that with $\int_s^t H_i(u, \lambda) dW(\lambda) \sim \mathcal{N}(0, \int_s^t H_i(u, \lambda)H_i(u, \lambda)' d\lambda)$, we find

$$\begin{aligned} E[\|\int_s^t H_i(u, \lambda) dW(\lambda)\|^4] &\leq 3(\text{tr} \int_s^t H_i(u, \lambda)H_i(u, \lambda)' d\lambda)^2 \\ &\leq 3(\int_s^t \|H_i(u, \lambda)\|^2 d\lambda)^2. \end{aligned}$$

Thus, for $0 \leq s < t \leq 1$, we have with $\psi(s) = d\Psi(s)/ds$

$$\begin{aligned}
& E\left[\left\|\int_0^t H_i(t, \lambda)dW(\lambda) - \int_0^s H_i(s, \lambda)dW(\lambda)\right\|^4\right] \\
&= E\left[\left\|\int_0^s (H_i(t, \lambda) - H_i(s, \lambda))dW(\lambda) + \int_s^t H_i(t, \lambda)dW(\lambda)\right\|^4\right] \\
&\leq 3\left[\int_0^s \|H_i(t, \lambda) - H_i(s, \lambda)\|^2 d\lambda + \int_s^t \|H_i(t, \lambda)\|^2 d\lambda\right]^2 \\
&\leq 3[k^2(\sup_{0 \leq \lambda \leq 1-\varepsilon} \|J(\lambda)^{-1}\|^2 \|F_i(\lambda)\|^2)(\|\Psi(s) - \Psi(t)\|^2 + (t-s) \sup_{0 \leq \lambda \leq 1} \|\Psi(\lambda)\|^2)]^2 \\
&\leq 3k^4(\sup_{0 \leq \lambda \leq 1-\varepsilon} \|J(\lambda)^{-1}\|^4 \|F_i(\lambda)\|^4)(\sup_{0 \leq \lambda \leq 1} \|\psi(\lambda)\|^2 + \sup_{0 \leq \lambda \leq 1} \|\Psi(\lambda)\|^2)^2 (t-s)^2
\end{aligned}$$

where the last inequality follows from $\Psi(t) - \Psi(s) = (t-s) \int_0^1 \psi(s + \lambda(t-s))d\lambda$, so that by Kolmogorov's continuity theorem (p. 14 of Oksendal (2000)), there exist continuous (and thus cadlag) versions of the stochastic processes $\int_0^s H_i(s, t)dW(t)$, $i \in \{a, b\}$.

A.4 Parameterization of $\Sigma_{(Y,X)}$ under H_0 in the restricted trend models

G -model with $r > k$: Because of invariance, it is without loss of generality to assume that the first $r-k$ rows of R are equal to zero, so that the first $r-k$ columns of Z are independent of V . The joint distribution of V and the last k rows of Z are then just as in the model with $r = k$, so that Lemma 3 implies that in the G -model with $r > k$, $\Sigma_{(Y,X)}$ is of the form $\Sigma_{(Y,X)} = \text{diag}(I_{r-k} \otimes I_q, \Sigma_k^*)$ under the null hypothesis, where Σ_k^* is any positive definite $k^2q \times k^2q$ matrix with upper left $kq \times kq$ block equal to the identity matrix. The nuisance parameter θ is thus of dimension $k^2q^2 + kq(kq+1)/2$

Diagonal G -model: Let Z_V and ζ be $q \times k$ random matrices with $\text{vec}(Z, Z_V, \zeta) \sim \mathcal{N}(0, \Sigma_{(Z, Z_V, \zeta)})$, where

$$\Sigma_{(Z, Z_V, \zeta)} = \text{diag}\left(\left[\begin{array}{cc} I_r & R \\ R' & I_k \end{array}\right], I_k\right) \otimes I_q.$$

A construction as in the proof of Lemma 3 implies that the j 'th column of V can be chosen as an arbitrary linear combination of the j 'th column of Z_V and the j 'th column of ζ , $j = 1, \dots, k$ (subject to the constraint that the resulting matrix is positive definite). Thus, $\Sigma_{(Y,X)}$ may be parametrized as $\Sigma_{(Y,X)} = A_{(Z, Z_V, \zeta)} \Sigma_{(Z, Z_V, \zeta)} A'_{(Z, Z_V, \zeta)}$, where

$$A_{(Z, Z_V, \zeta)} = \begin{pmatrix} I_{rq} & 0 & 0 \\ 0 & \text{diag}(A_{V,1}, A_{V,2}, \dots, A_{V,k}) & \text{diag}(L_{\zeta,1}, L_{\zeta,2}, \dots, L_{\zeta,k}) \end{pmatrix},$$

$A_{V,j}$ are arbitrary $q \times q$ matrices and $L_{\zeta,j}$ are arbitrary lower diagonal $q \times q$ matrices. Including R , θ is thus of dimension $rk + kq^2 + kq(q+1)/2$.

In the stationary diagonal model where $G(s, t) = \text{diag}(g_1^S(s-t), \dots, g_k^S(s-t))$, we set g_j^S to be step functions

$$g_j^S(x) = \sum_{i=1}^{n_g} \xi_{j,i} \mathbf{1} \left[\frac{i-1}{n_g+1} \leq \frac{x}{1+x} < \frac{i}{n_g+1} \right] \quad (26)$$

for $n_g = 40$ and some scalar parameters $\xi_{j,i}$, $j = 1, \dots, k$, $i = 1, \dots, n_g$. The steps occur at the points $i/(n_g + 1 - i)$, so that more flexibility is allowed for small values of x (half of the points are associated with values of x less than 1, for example). The values of Σ_{ZV} and Σ_{VV} then follow from (17). In this specification θ contains the kn_g coefficients $\xi_{j,i}$ and the rk coefficients in the correlation matrix R . While the specification (26) only captures a subset of all possible covariance matrices $\Sigma_{(Y,X)}$ in the (nonparametric) stationary model, any test that controls size for all functions $H(s, t)$ of the form $H(s, t) = \text{diag}(g_1^S(s-t), \dots, g_k^S(s-t))S_v$ *a fortiori* has to control size for the specification (26). The upper bounds on power of tests that control size for all values of θ under (26) are therefore also upper bounds for tests that control size for the generic stationary stochastic trend model.

In the local-to-unity model, a straightforward (but tedious) calculation determines the value of $\Sigma_{(Y,X)}$ as function of the $k \times k$ matrix C and the $r \times k$ correlation matrix R , so that θ is of dimension $k^2 + rk$. Finally, the $I(1)$ model is a special case of the local-to-unity model with $C = 0$.

A.5 Kullback-Leibler Divergence Problem of Section 5.3.2

Let Σ_1 denote the covariance matrix of $\text{vec}(Y, X)$ under a specific $I(1)$ alternative as described in Subsection 5.2.1 (that is, for specific values of $B = B_1$ and $R = R_1$), let $\Sigma_0(\theta)$ with $\theta \in \Theta$ be the covariance matrix of $\text{vec}(Y, X)$ under the null for the relevant restrictions on the stochastic trend, and define the $nq \times nq$ matrix (recall that $n = r + k$)

$$A(\gamma) = \begin{bmatrix} \gamma_{yz} \otimes I_q & 0 \\ \gamma_{xz} \otimes I_q & \gamma_{xv} \otimes I_q \end{bmatrix}$$

where γ_{yz} is $r \times r$, γ_{xz} is $k \times r$, and γ_{xv} is $k \times k$. This yields $A(\gamma) \text{vec}(Y, X) \sim \mathcal{N}(0, A(\gamma)\Sigma_0(\theta)A(\gamma)')$. Denote the Kullback-Leibler divergence between the $nq \times 1$ distributions $\mathcal{N}(0, \Sigma_1)$ and $\mathcal{N}(0, \Sigma_0)$ as $K(\Sigma_1, \Sigma_0) = \frac{1}{2} \ln(\det \Sigma_1 / \det \Sigma_0) + \frac{1}{2} \text{tr}(\Sigma_0^{-1} \Sigma_1) - nq$. The value of θ^* is chosen to numerically solve

$$\min_{\gamma \in \mathbb{R}^{r^2+k^2+kr}} K(\Sigma_1, A(\gamma)\Sigma_0(\theta^*)A(\gamma)') = \min_{\theta \in \Theta, \gamma \in \mathbb{R}^{r^2+k^2+kr}} K(\Sigma_1, A(\gamma)\Sigma_0(\theta)A(\gamma)'), \quad (27)$$

that is, θ^* numerically minimizes the Kullback-Leibler divergence (or KLIC) between the null and alternative densities of (Y, X) , allowing for transformations as described by $A(\gamma)$ under the null. While these transformations do not affect the implied distribution of the maximal invariant, they do in general lead to a different θ^* , which we found to yield a slightly lower upper bound. The

minimization problem is over a high dimensional parameter, but the objective function is quickly computed and well behaved, so that numerical minimization is feasible.

A.6 Validity of $r = 1, k = 2$ bounds for $r = 1, k > 2$

Here we show that the power bounds computed using $r = 1$ and $k = 2$ also serve as power bounds for models with $r = 1$ and all values of $k > 2$.

To see why, first consider the alternative $I(1)$ model as described in subsection 5.2.1, $Y = Z + VB'$ and $X = V$. Let P be a $k \times k$ orthonormal matrix whose last $k - 2$ rows are orthogonal to R and B , and whose second row is orthogonal to R . Partition $X = (X_{12}, X_{3k})$, where X_{12} contains the first two columns of X and X_{3k} contains the remaining $k - 2$ columns. By invariance, there is no loss in generality in setting $X = \tilde{X}P = (\tilde{X}_{12}, \tilde{X}_{3k})P$, so that $Y = Z + \tilde{X}PB' = Z + \tilde{X}_{12}B'_{12}$, where \tilde{X}_{12} and B'_{12} are the first two columns of \tilde{X} and B , respectively, and the last $k - 1$ columns of \tilde{X} (and thus \tilde{X}_{3k}) are independent of Z . The group of transformations

$$(Y, \tilde{X}_{12}, \tilde{X}_{3k}) \rightarrow (YA_{yy}, \tilde{X}_{12}\tilde{A}_{xx} + YA_{xy}, \tilde{X}_{3k}) \quad (28)$$

for nonsingular A_{yy} and \tilde{A}_{xx} is a subset of the transformations $(Y, \tilde{X}) \rightarrow (YA_{yy}, \tilde{X}\tilde{A}_{xx} + YA_{xy})$, so the best invariant test to (28) is at least as powerful as the best invariant test to (21). Let \tilde{Q}_{12} be a maximal invariant to $(Y, \tilde{X}_{12}) \rightarrow (YA_{yy}, \tilde{X}_{12}\tilde{A}_{xx} + YA_{xy})$, such that $\{\tilde{Q}_{12}, \tilde{X}_{3k}\}$ is a maximal invariant to (28). Since \tilde{X}_{3k} is independent of (Y, \tilde{X}_{12}) , the density of $\{\tilde{Q}_{12}, \tilde{X}_{3k}\}$ under the alternative factors as $f_{a, \tilde{Q}_{12}} \cdot f_{a, \tilde{X}_{3k}}$.

For all null models discussed in subsection 5.2.2, it is possible to choose $X = (X_{12}, X_{3k}) = V$ in a way such that X_{3k} is independent of X_{12} with marginal distribution $f_{0, X_{3k}} = f_{a, \tilde{X}_{3k}}$, (i.e. it corresponds to the $I(1)$ model) and the possibilities for X_{12} and its relationship with Z are the same as in the version of the model with $k = 2$. Thus, with this choice, the term $f_{a, \tilde{X}_{3k}}$ cancels in the likelihood ratio test of the maximal invariant to (28), and the testing problem corresponds precisely to the model with $k = 2$.¹² An upper bound for the model with $r = 1$ and $k = 2$ is therefore also an upper bound for the model with $r = 1$ and $k > 2$.

A.7 Derivation of LFST(b) and $\xi(b)$ of Section 6.2

A straightforward calculation shows that Σ_{YY} can be written as

$$\Sigma_{YY} = (I_r \otimes I_q) + (BB' \otimes D) + (BR' \otimes F) + (RB' \otimes F') \quad (29)$$

¹²This is not strictly true for the stationary G -model, which excludes $I(1)$ stochastic trends. But the low-frequency transformation of the suitably scaled stationary local-to-unity model converges in mean squared to the $I(1)$ model as the local-to-unity parameter approaches zero (cf. Section 2.4 of Müller and Watson (2008)), so that the additional discriminatory power from X_{3k} can be made arbitrarily small, and the conclusion continues to hold.

where F and D are $q \times q$ matrices, D is the diagonal matrix defined in Corollary 1, and $F = [f_{ij}]$, with $f_{ij} = 0$ if i and j are both even or odd, and $f_{ij} = 4/[\pi^2(i^2 - j^2)]$ otherwise. Thus, $F = -F'$. (The simple diagonal form of D is due to the particular choice of the weighting functions Ψ in (14); see Section 2.3 in Müller and Watson (2008)). Noting that VQ'_k has the same distribution as V , $Y = Z + VB'$ implies that Σ_{YY} is the same for $B = B_1$ and $B = B_1Q_k$. Thus, if there exists a suitable Q_k that makes BQ_kR' symmetric, the last two terms in (29) cancel. The Corollary now follows from the Neyman-Pearson Lemma and Theorem 2 part (b).

Furthermore, for $r > k$, assume that BB' has k eigenvalues equal to b^2 , and let $Q_r \text{diag}(b^2 I_k, 0I_{r-k})Q'_r$ be the spectral decomposition of BB' . The invariance $Y \rightarrow YA'$ implies that the density of the maximal invariant is the same for $B_1 = B$ and for $B_1 = AB$. Thus, with $A = Q'_r$, $Q'_r BB' Q_r = \text{diag}(b^2 I_k, 0I_{r-k})$, so that by Theorem 2 part (a) and the Neyman-Pearson Lemma, the best test rejects for large values of

$$\xi(b) = \det(Y'Y)^{(q+k-r)/2} \det(Y'(I_q + b^2 D)^{-1}Y)^{-k/2} E_{\omega_Y} [|\det(\omega_Y)|^{q-r}]$$

where $\text{vec } \omega_Y \sim \mathcal{N}(0, \Omega_Y^{-1})$ and $\Omega_Y = \text{diag}(I_k \otimes Y'(I_q + b^2 D)^{-1}Y, I_{r-k} \otimes Y'Y)$.

References

- ANDREWS, D. W. K. (1991): “Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation,” *Econometrica*, 59, 817–858.
- BAILLIE, R. T. (1996): “Long Memory Processes and Fractional Integration in Econometrics,” *Journal of Econometrics*, 73, 5–59.
- BOBKOSKI, M. J. (1983): “Hypothesis Testing in Nonstationary Time Series,” *unpublished Ph.D. thesis, Department of Statistics, University of Wisconsin*.
- CAMPBELL, J. Y., AND M. YOGO (2006): “Efficient Tests of Stock Return Predictability,” *Journal of Financial Economics*, 81, 27–60.
- CANER, M., AND L. KILIAN (2001): “Size Distortions of Tests of the Null Hypothesis of Stationarity: Evidence and Implications for the PPP Debate,” *Journal of International Money and Finance*, 20, 639–657.
- CAVANAGH, C. L. (1985): “Roots Local To Unity,” *Working Paper, Harvard University*.
- CAVANAGH, C. L., G. ELLIOTT, AND J. H. STOCK (1995): “Inference in Models with Nearly Integrated Regressors,” *Econometric Theory*, 11, 1131–1147.
- CHAN, N. H., AND C. Z. WEI (1987): “Asymptotic Inference for Nearly Nonstationary AR(1) Processes,” *The Annals of Statistics*, 15, 1050–1063.
- CHEN, W. W., AND C. M. HURVICH (2003a): “Estimating Fractional Cointegration in the Presence of Polynomial Trends,” *Journal of Econometrics*, 117, 95–121.
- (2003b): “Semiparametric Estimation of Multivariate Fractional Cointegration,” *Journal of the American Statistical Association*, 463, 629–642.
- (2006): “Semiparametric Estimation of Fractional Cointegration Subspaces,” *The Annals of Statistics*, 34, 2939–2979.
- DEN HAAN, W. J., AND A. T. LEVIN (1997): “A Practitioner’s Guide to Robust Covariance Matrix Estimation,” in *Handbook of Statistics 15*, ed. by G. S. Maddala, and C. R. Rao, pp. 309–327. Elsevier, Amsterdam.

- ELLIOTT, G. (1998): “The Robustness of Cointegration Methods When Regressors Almost Have Unit Roots,” *Econometrica*, 66, 149–158.
- ELLIOTT, G., U. K. MÜLLER, AND M. W. WATSON (2012): “Nearly Optimal Tests When a Nuisance Parameter is Present Under the Null Hypothesis,” *Working Paper, Princeton University*.
- ELLIOTT, G., AND J. H. STOCK (1994): “Inference in Time Series Regression When the Order of Integration of a Regressor is Unknown,” *Econometric Theory*, 10, 672–700.
- GEWEKE, J., AND S. PORTER-HUDAK (1983): “The Estimation and Application of Long Memory Time Series Models,” *Journal of Time Series Analysis*, 4, 221–238.
- GRANGER, C. W. J. (1993): “What Are We Learning About the Long-Run?,” *The Economic Journal*, 103, 307–317.
- HARVEY, A. C. (1989): *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- JANSSON, M. (2004): “Stationarity Testing with Covariates,” *Econometric Theory*, 20, 56–94.
- JANSSON, M., AND M. J. MOREIRA (2006): “Optimal Inference in Regression Models with Nearly Integrated Regressors,” *Econometrica*, 74, 681–714.
- JEGANATHAN, P. (1997): “On Asymptotic Inference in Linear Cointegrated Time Series Systems,” *Econometric Theory*, 13, 692–745.
- JOHANSEN, S. (1988): “Statistical Analysis of Cointegration Vectors,” *Journal of Economic Dynamics and Control*, 12, 231–254.
- (1991): “Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models,” *Econometrica*, 59, 1551–1580.
- KIEFER, N., AND T. J. VOGELSANG (2005): “A New Asymptotic Theory for Heteroskedasticity-Autocorrelation Robust Tests,” *Econometric Theory*, 21, 1130–1164.

- KIEFER, N. M., T. J. VOGELSANG, AND H. BUNZEL (2000): “Simple Robust Testing of Regression Hypotheses,” *Econometrica*, 68, 695–714.
- KING, M. L. (1988): “Towards a Theory of Point Optimal Testing,” *Econometric Reviews*, 6, 169–218.
- 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,” *Journal of Econometrics*, 54, 159–178.
- LEHMANN, E. L., AND J. P. ROMANO (2005): *Testing Statistical Hypotheses*. Springer, New York.
- MAGNUS, J. R., AND H. NEUDECKER (1988.): *Matrix Differential Calculus*. Wiley, New York.
- MÜLLER, U. K. (2005): “Size and Power of Tests of Stationarity in Highly Autocorrelated Time Series,” *Journal of Econometrics*, 128, 195–213.
- (2007): “A Theory of Robust Long-Run Variance Estimation,” *Journal of Econometrics*, 141, 1331–1352.
- (2011): “Efficient Tests under a Weak Convergence Assumption,” *Econometrica*, 79, 395–435.
- MÜLLER, U. K., AND M. W. WATSON (2008): “Testing Models of Low-Frequency Variability,” *Econometrica*, 76, 979–1016.
- NABEYA, S. (1951): “Absolute Moments in 2-Dimensional Normal Distribution,” *Annals of the Institute of Statistical Mathematics*, 3, 2–6.
- NEWBY, W. K., AND K. WEST (1987): “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.
- NYBLOM, J. (1989): “Testing for the Constancy of Parameters Over Time,” *Journal of the American Statistical Association*, 84, 223–230.

- NYBLOM, J., AND A. HARVEY (2000): “Tests of Common Stochastic Trends,” *Econometric Theory*, 16, 176–199.
- 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.
- OKSENDAL, B. (2000): *Stochastic Differential Equations: An Introduction with Applications*. Springer, Berlin, 5 edn.
- PARK, J. Y. (1992): “Canonical Cointegrating Regressions,” *Econometrica*, 60, 119–143.
- PHILLIPS, P. C. B. (1987): “Towards a Unified Asymptotic Theory for Autoregression,” *Biometrika*, 74, 535–547.
- PHILLIPS, P. C. B., AND B. E. HANSEN (1990): “Statistical Inference in Instrumental Variables Regression with I(1) Processes,” *Review of Economic Studies*, 57, 99–125.
- ROBINSON, P. M. (2003): “Long-Memory Time Series,” in *Time Series with Long Memory*, ed. by P. M. Robinson, pp. 4–32. Oxford University Press, Oxford.
- ROBINSON, P. M., AND J. HUALDE (2003): “Cointegration in Fractional Systems with Unknown Integration Orders,” *Econometrica*, 71, 1727–1766.
- ROBINSON, P. M., AND D. MARINUCCI (2001): “Narrow-Band Analysis of Nonstationary Processes,” *The Annals of Statistics*, 29, 947–986.
- (2003): “Semiparametric Frequency Domain Analysis of Fractional Cointegration,” in *Time Series with Long Memory*, ed. by P. M. Robinson, pp. 334–373. Oxford University Press, Oxford.
- SAIKKONEN, P. (1991): “Estimation and Testing of Cointegrating Regressions,” *Econometric Theory*, 7, 1–27.
- SAIKKONEN, P., AND R. LUUKONEN (1993): “Point Optimal Tests for Testing the Order of Differencing in ARIMA Models,” *Econometric Theory*, 9, 343–362.

- STOCK, J. H. (1991): “Confidence Intervals for the Largest Autoregressive Root in U.S. Macroeconomic Time Series,” *Journal of Monetary Economics*, 28, 435–459.
- STOCK, J. H., AND M. W. WATSON (1993): “A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems,” *Econometrica*, 61, 783–820.
- (1996): “Confidence Sets in Regressions with Highly Serially Correlated Regressors,” *Working Paper, Harvard University*.
- SUN, Y., P. C. B. PHILLIPS, AND S. JIN (2008): “Optimal Bandwidth Selection in Heteroskedasticity-Autocorrelation Robust Testing,” *Econometrica*, 76, 175–794.
- VALKANOV, R. (2003): “Long-Horizon Regressions: Theoretical Results and Applications,” *Journal of Financial Economics*, 68, 201–232.
- VELASCO, C. (2003): “Gaussian Semiparametric Estimation of Fractional Cointegration,” *Journal of Time Series Analysis*, 24, 345–378.
- WRIGHT, J. H. (2000): “Confidence Sets for Cointegrating Coefficients Based on Stationarity Tests,” *Journal of Business and Economic Statistics*, 18, 211–222.

Table 1: Power and Size Bounds for 5% Tests ($q = 12$)

A. $r = k = 1$

H_0	H_B	$ B = 7$				$ B = 14$					
		$ R = 0$	$ R = 0.5$		$ R = 0.9$		$ R = 0$	$ R = 0.5$		$ R = 0.9$	
			$RB < 0$	$RB > 0$	$RB < 0$	$RB > 0$		$RB < 0$	$RB > 0$	$RB < 0$	$RB > 0$
<i>unr</i>		0.36	0.36	0.36	0.36	0.36	0.64	0.65	0.65	0.66	0.66
<i>stat</i>		0.41	0.52	0.40	0.89	0.44	0.70	0.80	0.68	0.98	0.72
	<i>unr</i>	0.06	0.14	0.05	0.65	0.07	0.07	0.17	0.05	0.47	0.06
<i>LTU</i>		0.50	0.65	0.59	0.95	0.67	0.81	0.92	0.81	1.00	0.87
	<i>unr</i>	0.13	0.28	0.21	0.79	0.32	0.15	0.42	0.17	0.49	0.20
	<i>stat</i>	0.08	0.09	0.15	0.07	0.20	0.11	0.13	0.13	0.03	0.14
<i>l(1)</i>		0.50	0.65	0.65	0.95	0.95	0.82	0.91	0.91	1.00	1.00
	<i>unr</i>	0.12	0.28	0.26	0.76	0.80	0.19	0.37	0.34	0.48	0.48
	<i>stat</i>	0.08	0.09	0.21	0.07	0.71	0.12	0.12	0.31	0.03	0.42
	<i>LTU</i>	0.04	0.04	0.06	0.04	0.50	0.04	0.03	0.16	0.00	0.25

B. $r = 1$ and $k \geq 2$

H_0	H_B	$ B = 7$						$ B = 14$							
		$\ R\ = 0$	$\ R\ = 0.5$			$\ R\ = 0.9$			$\ R\ = 0$	$\ R\ = 0.5$			$\ R\ = 0.9$		
			$\omega = -1$	$\omega = 0$	$\omega = 1$	$\omega = -1$	$\omega = 0$	$\omega = 1$		$\omega = -1$	$\omega = 0$	$\omega = 1$	$\omega = -1$	$\omega = 0$	$\omega = 1$
<i>unr</i>		0.36	0.37	0.36	0.36	0.36	0.36	0.36	0.63	0.64	0.64	0.65	0.65	0.65	0.66
<i>diag</i>		0.36	0.37	0.37	0.36	0.36	0.42	0.36	0.63	0.64	0.64	0.65	0.65	0.74	0.66
	<i>unr</i>	0.04	0.04	0.04	0.04	0.04	0.06	0.04	0.04	0.04	0.04	0.04	0.04	0.08	0.04
<i>stat</i>		0.40	0.48	0.42	0.39	0.83	0.59	0.44	0.68	0.77	0.71	0.67	0.97	0.84	0.71
	<i>unr</i>	0.05	0.10	0.07	0.05	0.55	0.19	0.07	0.06	0.12	0.08	0.05	0.45	0.20	0.06
	<i>diag</i>	0.05	0.10	0.06	0.05	0.55	0.15	0.07	0.06	0.12	0.07	0.05	0.45	0.14	0.06
<i>LTU</i>		0.46	0.57	0.49	0.53	0.92	0.70	0.64	0.74	0.84	0.77	0.76	1.00	0.85	0.85
	<i>unr</i>	0.09	0.18	0.10	0.15	0.71	0.34	0.27	0.09	0.22	0.12	0.12	0.50	0.20	0.18
	<i>diag</i>	0.09	0.18	0.09	0.15	0.71	0.31	0.27	0.09	0.22	0.13	0.12	0.50	0.16	0.18
	<i>stat</i>	0.07	0.07	0.06	0.12	0.08	0.10	0.15	0.06	0.08	0.06	0.10	0.05	0.05	0.12
<i>l(1)</i>		0.46	0.57	0.48	0.57	0.91	0.86	0.91	0.75	0.83	0.78	0.83	0.99	0.98	0.99
	<i>unr</i>	0.09	0.18	0.11	0.17	0.71	0.70	0.71	0.10	0.23	0.15	0.21	0.50	0.46	0.49
	<i>diag</i>	0.09	0.18	0.11	0.17	0.71	0.53	0.71	0.10	0.23	0.13	0.21	0.50	0.37	0.49
	<i>stat</i>	0.07	0.07	0.07	0.17	0.08	0.23	0.67	0.08	0.08	0.08	0.18	0.05	0.20	0.45
	<i>LTU</i>	0.04	0.04	0.04	0.05	0.04	0.22	0.46	0.04	0.03	0.07	0.10	0.01	0.22	0.28

Table 1 (Continued)

C. $r = 2$ and $k = 1$

H_0	H_B	$\ B\ = 10$						$\ B\ = 20$							
		$\ R\ =0$	$\ R\ = 0.5$			$\ R\ = 0.9$			$\ R\ =0$	$\ R\ = 0.5$			$\ R\ = 0.9$		
			$\omega=-1$	$\omega=0$	$\omega=1$	$\omega=-1$	$\omega=0$	$\omega=1$		$\omega=-1$	$\omega=0$	$\omega=1$	$\omega=-1$	$\omega=0$	$\omega=1$
	<i>unr</i>	0.45	0.47	0.50	0.48	0.50	0.69	0.51	0.69	0.73	0.75	0.73	0.79	0.94	0.79
	<i>G</i>	0.45	0.47	0.55	0.48	0.50	0.74	0.51	0.69	0.73	0.76	0.73	0.79	0.97	0.79
	<i>unr</i>	0.04	0.04	0.06	0.04	0.04	0.06	0.04	0.04	0.04	0.04	0.04	0.04	0.05	0.04
	<i>stat</i>	0.50	0.61	0.61	0.51	0.94	0.74	0.57	0.72	0.81	0.78	0.74	0.98	0.97	0.82
	<i>unr</i>	0.06	0.13	0.11	0.05	0.61	0.06	0.07	0.05	0.10	0.05	0.05	0.29	0.04	0.05
	<i>G</i>	0.06	0.13	0.07	0.05	0.61	0.04	0.07	0.05	0.10	0.05	0.05	0.29	0.02	0.05
	<i>LTU</i>	0.55	0.72	0.63	0.66	0.99	0.74	0.76	0.79	0.92	0.87	0.81	1.00	0.97	0.89
	<i>unr</i>	0.10	0.25	0.11	0.17	0.70	0.06	0.29	0.12	0.26	0.13	0.09	0.32	0.05	0.12
	<i>G</i>	0.10	0.25	0.08	0.17	0.70	0.04	0.29	0.12	0.26	0.13	0.09	0.32	0.02	0.12
	<i>stat</i>	0.06	0.09	0.05	0.15	0.07	0.04	0.19	0.07	0.13	0.10	0.08	0.03	0.03	0.09
	<i>I(1)</i>	0.55	0.70	0.63	0.70	0.97	0.74	0.97	0.84	0.91	0.88	0.91	1.00	0.97	1.00
	<i>unr</i>	0.10	0.22	0.11	0.24	0.69	0.06	0.69	0.16	0.26	0.15	0.29	0.32	0.05	0.32
	<i>G</i>	0.10	0.22	0.08	0.24	0.69	0.04	0.69	0.16	0.26	0.14	0.29	0.32	0.02	0.32
	<i>stat</i>	0.06	0.07	0.05	0.19	0.05	0.04	0.65	0.12	0.14	0.11	0.26	0.03	0.03	0.31
	<i>LTU</i>	0.04	0.04	0.04	0.06	0.02	0.04	0.45	0.10	0.03	0.04	0.18	0.00	0.02	0.21

Notes: Non-italicized numerical entries are upper bounds on power for 5% tests for the null hypothesis restricting the trend process as shown in the first column. The italicized entries are lower bounds on size for auxiliary trend processes given in columns for 5% level tests with power greater than the power bound less 3 percentage points. Abbreviations for the trend models in columns 1 and 2 are: (i) “*unr*” is the unrestricted model given in (3), (ii) “*G*” is the *G*-model in (6), (iii) “*diag*” is the diagonal *G*-model in (8), (iv) “*stat*” is the stationary *G* model given in (9), (v) “*LTU*” is the local-to-unity model given in (11), and (vi) “*I(1)*” is *I*(1) model given in (10). In panels B and C, $\omega = \text{tr}(R'B) / (\|R\| \|B\|)$. Results are based on 20,000 Monte Carlo replications.

Table 2: Comparison of KLIC Minimized and Approximate Least Upper Power Bounds
 ($r = k = 1, q = 12$)

H ₀	Bound	B = 7					B = 14				
		R = 0	R = 0.5		R = 0.9		R = 0	R = 0.5		R = 0.9	
			RB<0	RB>0	RB<0	RB>0		RB<0	RB>0	RB<0	RB>0
LTU	KLIC	0.51	0.66	0.59	0.95	0.66	0.81	0.93	0.80	1.00	0.86
	LUB	0.50	0.66	0.58	0.93	0.65	0.78	0.88	0.78	1.00	0.82
I(1)	KLIC	0.51	0.65	0.65	0.95	0.95	0.82	0.92	0.91	1.00	1.00
	LUB	0.50	0.65	0.65	0.94	0.94	0.81	0.90	0.90	1.00	1.00

Notes: The entries labeled “KLIC” are computed using the KLIC minimization discussed in Section 5.3.2. The entries labeled “LUB” are computed using the approximate least upper algorithm discussed in Section 5.3.1, and are by construction no more than 2.5 percentage points above the actual least upper bound. Results are based on 20,000 Monte Carlo replications.

Table 3: 1%, 5%, and 10% Critical Values for the LFST Statistic

q	r														
	1			2			3			4			5		
6	5.25	3.62	3.08	6.76	5.16	4.39	7.25	6.09	5.43	7.14	6.46	6.02	6.54	6.33	6.16
7	4.33	3.08	2.68	5.52	4.20	3.63	6.05	4.95	4.37	6.15	5.35	4.92	5.89	5.48	5.20
8	3.68	2.73	2.39	4.65	3.54	3.08	5.17	4.16	3.68	5.29	4.55	4.12	5.26	4.73	4.42
9	3.21	2.46	2.18	4.02	3.09	2.73	4.46	3.58	3.19	4.63	3.93	3.56	4.66	4.12	3.83
10	2.86	2.25	2.02	3.56	2.79	2.48	3.94	3.17	2.84	4.10	3.47	3.15	4.18	3.66	3.38
11	2.62	2.10	1.90	3.16	2.54	2.29	3.53	2.87	2.59	3.71	3.12	2.84	3.78	3.30	3.03
12	2.46	1.98	1.81	2.89	2.35	2.13	3.18	2.64	2.39	3.38	2.84	2.60	3.48	3.02	2.78
13	2.29	1.88	1.73	2.68	2.21	2.01	2.92	2.44	2.23	3.13	2.63	2.42	3.20	2.77	2.57
14	2.16	1.80	1.67	2.50	2.09	1.92	2.74	2.31	2.11	2.91	2.47	2.27	2.97	2.59	2.40
15	2.07	1.74	1.61	2.36	1.99	1.84	2.56	2.18	2.01	2.69	2.32	2.15	2.80	2.44	2.27
16	1.97	1.67	1.56	2.24	1.91	1.77	2.44	2.08	1.92	2.55	2.21	2.05	2.64	2.30	2.15
17	1.89	1.62	1.52	2.15	1.84	1.71	2.32	1.99	1.85	2.43	2.11	1.96	2.50	2.20	2.05
18	1.82	1.58	1.49	2.07	1.78	1.66	2.21	1.92	1.79	2.32	2.02	1.89	2.39	2.10	1.98

Notes: The table shows asymptotic critical for the LFST(b) statistic computed using $b = 10/\sqrt{r}$, where $LFST(b) = \det(Y'Y)/\det(Y'(I + b^2D)^{-1}Y)$, with $D = \text{diag}(d_1, \dots, d_k)$ and $d_i = (i\pi)^{-2}$. Results are based on 50,000 Monte Carlo replications.

Table 4: Power of Y -only 5% Tests ($q = 12$)

A. $r = 1$

	$ B = 7$	$ B = 14$
LFST(10)	0.36	0.63
Y-Only Power Envelope	0.36	0.64

B. $r = 2$ and $k = 1$

$\ B\ = 10$							$\ B\ = 20$						
Power of LFST($10/\sqrt{2}$)													
0.39							0.58						
Y-Only Power Envelope													
$\ R\ =0$	$\ R\ = 0.5$			$\ R\ = 0.9$			$\ R\ =0$	$\ R\ = 0.5$			$\ R\ = 0.9$		
	$\omega=-1$	$\omega=0$	$\omega=1$	$\omega=-1$	$\omega=0$	$\omega=1$		$\omega=-1$	$\omega=0$	$\omega=1$	$\omega=-1$	$\omega=0$	$\omega=1$
0.41	0.41	0.44	0.41	0.41	0.55	0.41	0.68	0.64	0.68	0.64	0.64	0.86	0.64

Notes: The power envelope is computed using the test statistic LFST($|B|$) in panel A and by $\xi(B)$ in panel B. In panel B, $\omega = \text{tr}(R'B) / (\|R\| \cdot \|B\|)$. Results are based on 20,000 Monte Carlo replications.

Table 5: Persistence Statistics for Interest Rates

	DFGLS p -value	Largest AR Root 95% Conf. Int.	Nyblom/KPSS p -value	d -GPH Level	d -GPH Difference	LFST p -value
$R^{1\text{Year}}$	0.12	0.98 – 1.00	< 0.01	0.31 (0.08)	-0.33 (0.08)	< 0.01
$R^{10\text{Year}}$	0.26	0.98 – 1.00	< 0.01	0.18 (0.08)	-0.20 (0.08)	< 0.01
$R^{10\text{Year}} - R^{1\text{Year}}$	< 0.01	0.95 – 0.99	< 0.01	0.21 (0.08)	-0.23 (0.08)	0.16

Notes: The DFGLS are computed using 12 lagged first differences of the series. The confidence interval for the largest AR root is computed by inverting the Dickey-Fuller τ' statistic (Stock (1991)) using an autoregression augmented with 12 lagged first differences. The Nyblom/KPSS test is implemented using a Newey-West long-run variance estimator with $[0.75T^{1/3}] + 1$ autocovariances. The GPH regressions are estimated using $[T^{0.65}]$ periodogram ordinates. The LFST statistics are computed using $q = 14$, which capture variability with periodicities greater than 8.3 years.

Figure 1: U.S. Treasury Bond Rates




Figure 2: Term Spread ($R^{10Year} - R^{1Year}$)

