Long-Run Covariability*

Ulrich K. Müller and Mark W. Watson
Princeton University
Department of Economics
Princeton, NJ, 08544

August 2016
Preliminary and Incomplete

Abstract

We develop inference methods about the degree of long-run comovement of two time series. The parameters of interest are defined in terms of the population second-moment properties of low-frequency trends computed from the data. These trends are similar to low-pass filtered data and are designed to extract variability that corresponds to periods longer than the span of the sample divided by $q/2$, where q is a small number, such as 12. We numerically determine confidence sets that control coverage over a wide range of potential bivariate persistence patterns, which include arbitrary linear combinations of I(0), I(1), near unit roots and fractionally integrated models. In an application to U.S. economic data, we quantify the long-run covariability of a variety of series, such as those giving rise to the "great ratios", nominal exchange rates and relative nominal prices, unemployment rate and inflation, earnings and stock prices, etc.

JEL classification: C22, C53, E17

Keywords: band-pass regression, cointegration, fractional integration, great ratios

*We thank participants at several seminars for their comments as this work has evolved. Support was provided by the National Science Foundation through grant SES-1627660.

1 Introduction

Economic theories often have stark predictions about the covariability of variables over long-horizons: consumption and income move proportionally (permanent income/life cycle model of consumption) as do nominal exchange rates and relative nominal prices (long-run PPP), the unemployment rate is unaffected by the rate of price inflation (vertical long-run Phillips curve), and so forth. But there is a limited set of statistical tools to investigate the validity of these long-run propositions. This paper expands this set of tools.

Two fundamental problems plague statistical inference about long-run phenomena. First, inference depends critically on the data's long-run persistence. Random walks yield statistics with different probability distributions than i.i.d. data, for example, and observations from persistent autoregressions or fractionally integrated processes yield statistics with their own unique probability distributions. Second, there are few "long-run" observations in the samples typically used in empirical analyses of long-run relations, so sample information is limited. Taken together these two problems conspire to make long-run inference particularly difficult: proper inference depends critically on the exact form of long-run persistence, but there is limited sample information available to empirically determine this persistence.

The most well-known example of faulty inference due to a mistaken assumption about persistence is Granger and Newbold's (1974) "spurious regression", where standard OLS inference leads to grossly misleading conclusions when applied to independent $I(1)$ variables. The last 40 years has seen important progress developing inference for specific classes of stochastic processes (most notably for $I(0)$ and integrated/cointegrated processes), but several aspects of the resulting inference remains fragile. For example, while HAC standard errors lead to reliable inference in $I(0)$ settings with limited serial correlation, the resulting hypothesis tests exhibit substantial size distortions for stationary series with high serial correlation (e.g., den Haan and Levin (1997), Kiefer, Vogelsang, and Bunzel (2000), and Müller (2014)). Inference in cointegrated models is well-developed (e.g., Engle and Granger (1987), Johansen (1988), Phillips (1991), Stock and Watson (1993)), but these models have knife-edge implications about long-run covariability (cointegrated variables have unit long-run correlations) and efficient inference methods are not robust to small departures from the model's assumed exact unit autoregressive roots (Elliott (1998)). Variables that are highly but not perfectly correlated in the long-run, or are highly persistent, but perhaps without exact unit roots, fall outside the standard cointegration framework.

This paper develops methods designed to provide reliable inference about long-run covariability for a wide range of persistence patterns (encompassing $I(0)$, $I(1)$, and many other forms of long-run persistence) and that are applicable regardless of the degree of long-run correlation. The methods rely on low-frequency averages of the data to measure the data's long-run variability and covariability. These long-run data summaries have proven useful for constructing long-run covariance matrices and associated test statistics in $I(0)$ settings (e.g., Müller (2004, 2007), Phillips (2005), Sun (2013), and Lazarus, Lewis, and Stock (2016)), but also for conducting inference about more general patterns of long-run persistence and measuring uncertainty about long-run predictions (Müller and Watson (2008, forthcoming)). A key simplification offered by these averages is that they are normally distributed in large samples even though the stochastic process generating the data may exhibit substantial persistence (Müller and Watson (2016)). Large-sample inference about covariability parameters is thus transformed into finite-sample problems involving a handful of normal random variables, and, while these problems are "non-standard," they can be solved using previously developed statistical methods paired with modern computing power.

The paper's goal is to provide empirical researchers with an easy-to-use method for constructing confidence intervals for long-run correlation coefficients, linear regression coefficients, and standard deviations of regression errors. These confidence intervals are both valid over a wide range of persistence patterns (for example, the 90% confidence interval for the long-run correlation coefficient includes the true value of the coefficient with probability of at least 90%) and nearly optimal (in the sense of having close to shortest expected length; see Section 4 for details). As discussed in Section 3, the procedures allow for $I(0)$, $I(1)$, near unit roots, fractionally integrated models, and linear combinations of variables with these forms of persistence. Using a set of pre-computed "approximate least favorable distributions", the confidence intervals readily follow from the formulae discussed in Section 4.[1]

The outline of the paper is as follows. The next section defines the notion of long-run variability and covariability that is used in throughout the paper. It is defined in terms of the population properties of long-run projections, which are usefully thought of as low-pass filers (e.g., Baxter and King (1999)), Hodrick and Prescott (1997)). The discussion

[1]The replication files contain a matlab program for computing these confidence intervals, available at www.princeton.edu\~mwatson. This program uses the approximate least favorable distributions discussed in Section 4 and the appendix, which are also available in the replication files.

is carried out in the context of two leading empirical examples, the long-run relationship between consumption and GDP and between short- and long-term nominal interest rates. In the long-run projections we employ, long-run variability and covariability is equivalently captured by the covariability of a small number of trigonometrically weighted averages of the data. Section 3 derives the large-sample normality of these averages and introduces a flexible parameterization of the joint long-run persistence properties of the underlying stochastic process. The large-sample framework developed in Section 3 reduces the problem of inference about long-run covariability parameters into the problem of inference about the covariance matrix of a low dimensional multivariate normal random vector. Section 4 reviews relevant methods for solving this finite sample problem. Section 5 uses the resulting inference methods to empirically study several familiar long-run relations involving balanced growth (GDP, consumption, investment, labor income, and productivity), the term structure of interest rates, the Fisher correlation (inflation and interest rates), the Phillips correlation (inflation and unemployment), PPP (exchange rates and price ratios), and the long-run relationship between stock prices, dividends and earnings. Section 6 examines the robustness of Section 5's empirical conclusions to changes in the periodicities defining the "long-run", and to alternative choices for the information set used for inference.

2 Long-run projections and covariability

2.1 Two empirical examples of long-run covariability

We begin by examining the long-run covariability of GDP and consumption and of short- and long-term nominal interest rates. These data will motivate and illustrate the methods developed in this paper.

Consumption and income: One of the most celebrated and studied long-run relationship in economics concerns income and consumption. The long-run stability of consumption/income ratio is one of economics' "Great Ratios" (Klein and Kosobud (1961)); the dynamic implications of this stability inspired early work on error-correction models (e.g., Sargan (1964) and Davidson, Hendry, Srba, and Yeo (1978)), and these in turn motivated Granger's formulation of cointegration (Granger (1981)). While early analysis provided empirical support for the cointegration of consumption and income (e.g., Campbell (1987), King, Plosser, Stock, and Watson (1991), Cochrane (1994)), more recent work has come to the opposite conclusion (see Lettau and Ludvigson (2013) for discussion and references). Whether or not consumption and income are cointegrated (i.e., have an exact unit autoregressive root and exact unit long-run correlation), even a casual glance at the data suggests the two variables move together closely in the long run.

Consider, for example, the evolution of U.S. real per-capita GDP and consumption over the post-WWII period. In the 17 years from 1948 through 1964, GDP increased by 62% and consumption increased by 52%. Over the next 17 years (1965-1981) both GDP and consumption grew more slowly, by only 30%. Growth rebounded during 1982 to 1998, when GDP grew by 43% and consumption increased 55%, but slowed again over 1999-2015 when GDP grew by only 17% and consumption increased by only 23%. Over these 17-year periods, there was substantial variability in the average annual rate of growth of GDP (2.9%, 1.4%, 2.1%, and 0.9% per year, respectively over the sub-samples), and these changes were roughly matched by consumption (annual average growth rates of 2.5%, 1.5%, 2.6%, and 1.2%). In this sense, GDP and consumption exhibited substantial long-run variability and covariability over the post-WWII period.[2]

There are two distinct notions of "long-run" implicit in this calculation. The most obvious is that each period makes up 17 years, approximately twice the length of the typical business cycle. But another is that each period encompasses a large fraction (1/4) of the full 1948-2015 sample period. Our statistical framework defines long-run in this latter way: long-run statistical analysis involves inference about characteristics of stochastic processes that govern the evolution of averages of the data over periods that are large relative to the available sample.

With this is mind, the first two panels of Figure 1 plot the average growth rates of GDP and consumption over six non-overlapping sub-samples in 1948-2015. Figure 1.a plots the averages growth rates against time, and Figure 1.b is a scatterplot of the six average growth rates for consumption against corresponding values for GDP. Each of the six sub-samples contains 11.3 year (45 quarters), spans of history longer than the typical business cycle, and arguably capture "long-run" variability in GDP and consumption. And, each represents

[2] Consumption is personal consumption expenditures (including durables) from the NIPA; Section 5 shows results for non-durables, services, and durables separately. Both GDP and consumption are deflated by the PCE deflator, so that output is measured in terms of consumption goods, and expressed in per-capita terms using the civilian non-institutionalized population over the age of 16. The data appendix contains data sources and descriptions for all data used in this paper.

Figure 1: Long-run average growth rates of GDP and consumption

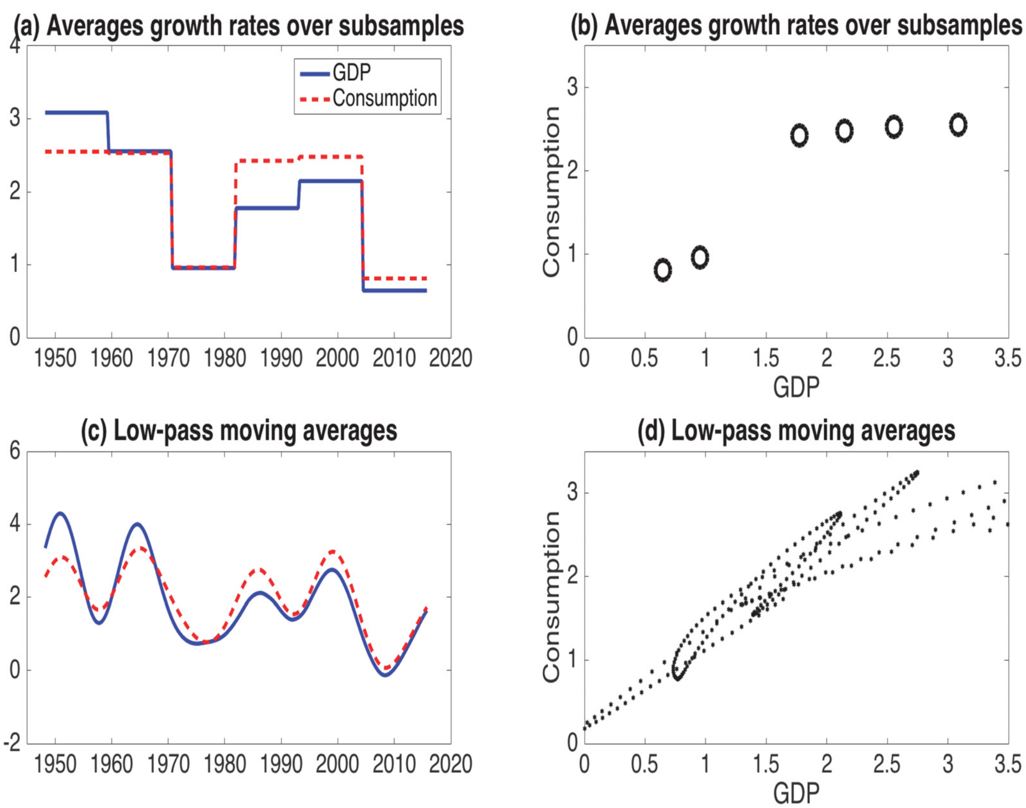

Notes: Panel (a) shows sample average of the variables over the period shown. Panel (c) shows low-pass moving averages using a frequency cutoff corresponding to *T*/6 periods (approximately 11 years). Panels (b) and (d) are scatterplots of the variables in (a) and (c).

a substantial fraction (1/6) of the sample and is a long-run observation in a statistical sense. Average GDP and consumption growth over these subsamples exhibited substantial variability and (from the scatter plot) roughly one-for-one covariability.

Figure 1.c sharpens the analysis by plotting "low-pass" moving averages of the series designed to isolate variation in the series with periods longer than 11 years.[3] Sample variation in these moving averages is much like the variation in the subsample averages of Figure 1.a, but Figure 1.c captures the smooth transition of the series from high-growth to low-growth periods. The scatterplot of these moving averages is plotted in Figure 1.d. Like Figure 1.b, it shows the close relationship between long-run movements in consumption and GDP, but it also shows the high degree of serial correlation in the moving averages.

A convenient device for handling this serial correlation is to use projections on low-frequency periodic functions in place of the low-pass moving averages. To be specific, let x_t, $t = 1, ..., T$ denote a time series (e.g., growth rates of GDP or consumption). We use cosine functions for the periodic functions; let $\Psi_j(s) = \sqrt{2}\cos(js\pi)$ denote the function with period $2/j$ (where the factor $\sqrt{2}$ simplifies a calculation below), $\Psi(s) = [\Psi_1(s), \Psi_2(s), \ldots, \Psi_q(s)]'$ denote a vector of these functions with periods 2 through $2/q$, and Ψ_T denote the $T \times q$ matrix with t'th row given by $\Psi((t-1/2)/T)'$, so the j'th column of Ψ_T has period $2T/j$. Our empirical analysis uses $q = 12$ which captures periodicities longer than $T/6$, and this defines the long-run variation in the data the analysis is designed to capture. The projection of x_t onto $\Psi((t-1/2)/T)$ for $t = 1, ..., T$ yields the fitted values

$$\widehat{x}_t = X_T'\Psi((t-1/2)/T) \tag{1}$$

where X_T are the projection (linear regression) coefficients, $X_T = (\Psi_T'\Psi_T)^{-1}\Psi_T' x_{1:T}$, where $x_{1:T}$ is the $T \times 1$ vector with t'th element given by x_t. The matrix Ψ_T has two properties that simplify calculations and interpretation. First, $\Psi_T' l_T = 0$ where l_T is a vector of ones, so that $\widehat{x}_t$ also corresponds to the projection of $x_t - \overline{x}_{1:T}$ onto $\Psi((t-1/2)/T)$, where $\overline{x}_{1:T}$ is the sample mean. Second, $T^{-1}\Psi_T'\Psi_T = I_q$, so X_T corresponds to simple cosine-weighted averages of the data (i.e., are the "cosine transforms" of $\{x_t\}$)

$$X_T = T^{-1}\sum_{t=1}^{T}\Psi((t-1/2)/T)x_t. \tag{2}$$

[3] These were computed using an ideal low-pass filter for periods longer than $T/6$ truncated after $T/2$ terms. The series were padded with pre- and post-sample backcasts and forecasts constructed from an AR(4) model.

Figure 2: Long-run projections of GDP and consumption growth rates

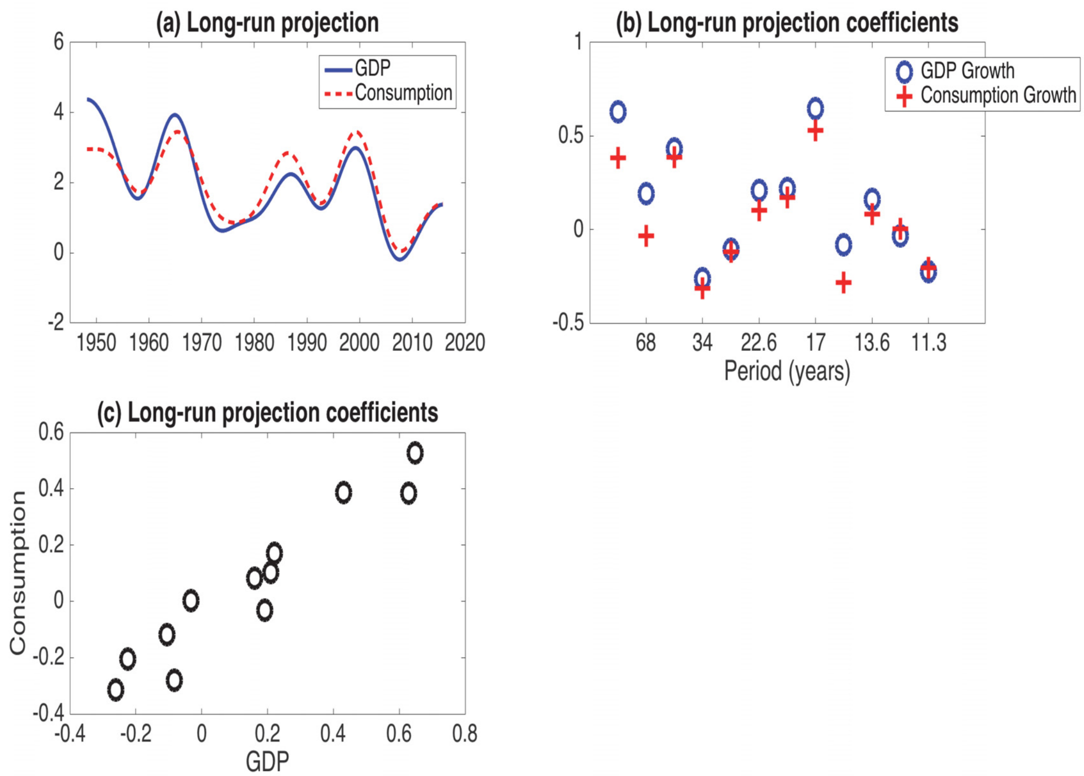

Notes: Panel (a) plots the projections of the data onto the low-frequency cosine terms discussed in the text, where sample means have been added to projections so they are consistent with the low-pass moving averages plotted in Figure 1(c). Panel (b) plots the projection coefficients (X_{jT}, Y_{jT}) against period $2T/j$ (in years). Panel (c) is a scatterplot of the variables from (b).

Letting (x_t, y_t) denote the growth rates of GDP and consumption, the long-run projections $(\widehat{x}_t, \widehat{y}_t)$ are plotted in Figure 2.a. Except for minor differences near the endpoints, these long-run projections essentially coincide with the low-pass moving average plotted in Figure 1.c, so both capture the same long-run sample variability in the data. An advantage of the long-run projections is that they are fully summarized by the projection coefficients (X_T, Y_T), a relatively small number of cosine-weighted averages of the sample data. Figure 2.b plots the projection coefficients, (X_{jT}, Y_{jT}) against the period of the corresponding cosine term, $2T/j$. Evidently, there is substantial variation and covariation in the projection coefficients. Indeed, the scatterplot of (X_{jT}, Y_{jT}) shown in Figure 2.c suggests a roughly one-to-one relationship between the cosine transforms.

The orthogonality of the cosine regressors Ψ_T leads to a tight connection between the variability and covariability in the long-run projections $(\widehat{x}_t, \widehat{y}_t)$ plotted in Figure 2.a and the cosine transforms (X_{jT}, Y_{jT}) plotted in Figure 2.b and 2.c:

$$T^{-1}\sum_{t=1}^{T}\begin{pmatrix}\widehat{x}_t\\ \widehat{y}_t\end{pmatrix}\begin{pmatrix}\widehat{x}_t & \widehat{y}_t\end{pmatrix} = T^{-1}\begin{pmatrix}X_T'\\ Y_T'\end{pmatrix}\Psi_T'\Psi_T\begin{pmatrix}X_T & Y_T\end{pmatrix} = \begin{bmatrix}X_T'X_T & X_T'Y_T\\ Y_T'X_T & Y_T'Y_T\end{bmatrix}. \quad (3)$$

Thus, sample covariability in the time series projections coincides with sample covariability in the cosine transforms.

Short-term and Long-term interest rates. The second empirical example involves short- and long-term nominal interest rates, as measured by the rate on 3-month U.S. Treasury bills, x_t, and the rate on 10-year U.S. Treasury bonds, y_t, from 1953 through 2015. The levels of these interest rates are highly serially correlated, but the term spread, $y_t - x_t$, far less so. Early cointegration work (e.g., Campbell and Shiller (1987)) modeled the level of interest rates as $I(1)$, and short- and long-rates as cointegrated. Later empirical analysis of the term structure (e.g., Dai and Singleton (2000), Diebold and Li (2006)) model the levels of interest rates as a function of small number of dynamic common factors that lead to common, but less than unit-root, long-run persistence.

Figure 3 plots the levels of short- and long-term interest rates, (x_t, y_t), along with their long-run projections, $(\widehat{x}_t, \widehat{y}_t)$, and cosine transforms, (X_T, Y_T). The long-run projections capture the rise in interest rates from the beginning of the sample through the early 1980s and then their subsequent decline (Figure 3.a). These long-swings in the level of interest rates lead to relatively larger values in the long-period cosine transforms (Figure 3.b). The projections for long-term interest rates closely track the projections for short-term rates and,

Figure 3: Short- and long-term interest rates

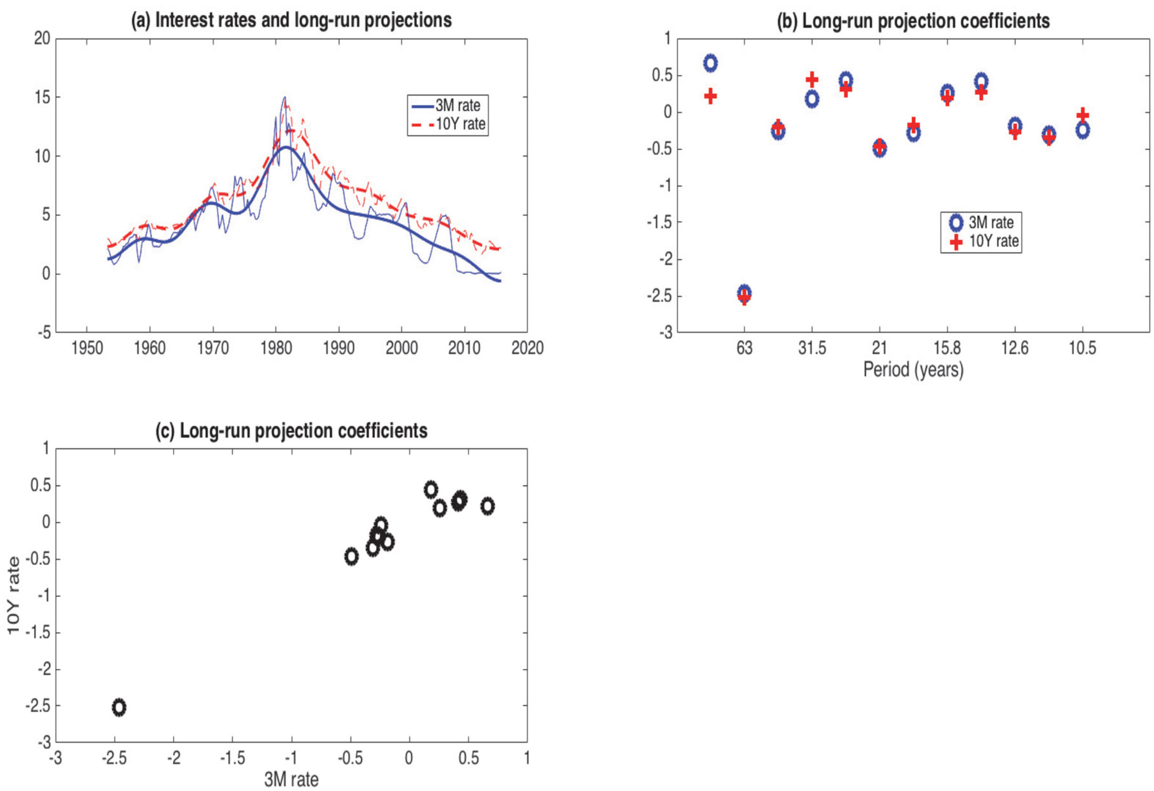

Notes: Panel (a) plots the data and projections of the data onto the low-frequency cosine terms discussed in the text, where sample means have been added to projections). Panel (b) plots the projection coefficients (X_{jT}, Y_{jT}) against period $2T/j$ (in years). Panel (c) is a scatterplot of the variables from (b).

given the connection between between the projections and cosine transforms, X_{jT} and Y_{jT} are highly correlated (Figure 3.c).

These two datasets differ markedly in their persistence: GDP and consumption growth rates are often modeled as low-order MA models, while nominal interest rates are highly serially correlated. Yet, the variables in both data sets exhibit substantial long-run variation and covariation which is readily evident in the long-run projections $(\widehat{x}_t, \widehat{y}_t)$ or equivalently (from (3)) the projection coefficients (X_T, Y_T). This suggests that the covariance/variance properties of (X_T, Y_T) are a useful starting point for defining the long-run covariability properties of stochastic processes exhibiting a wide range of persistent patterns.

2.2 A measure of long-run covariability using long-run projections

A straightforward definition of long-run covariability properties is based on the population analogue of the sample second moment matrices in (3). Let Σ_T denote the covariance matrix of $(X_T' \ Y_T')'$, partitioned as $\Sigma_{XX,T}$, $\Sigma_{XY,T}$, etc., and define

$$\Omega_T = T^{-1}\sum_{t=1}^{T} E\left[\begin{pmatrix} \widehat{x}_t \\ \widehat{y}_t \end{pmatrix}\begin{pmatrix} \widehat{x}_t & \widehat{y}_t \end{pmatrix}\right] = E\begin{bmatrix} X_T'X_T & X_T'Y_T \\ Y_T'X_T & Y_T'Y_T \end{bmatrix} = \begin{pmatrix} \mathrm{tr}(\Sigma_{XX,T}) & \mathrm{tr}(\Sigma_{XY,T}) \\ \mathrm{tr}(\Sigma_{YX,T}) & \mathrm{tr}(\Sigma_{YY,T}) \end{pmatrix} \tag{4}$$

where the equalities directly follow from (3).

The 2×2 matrix Ω_T is the average covariance matrix of the long-run projections $(\widehat{x}_t, \widehat{y}_t)$ in a sample of length T, and provides a summary of the variability and covariability of the long-run projections over repeated samples. Corresponding long-run correlation and linear regression parameters follow from the usual formulae

$$\begin{aligned} \rho_{xy,T} &= \Omega_{xy,T}/\sqrt{\Omega_{xx,T}\Omega_{yy,T}} \\ \beta_T &= \Omega_{xy,T}/\Omega_{xx,T} \\ \sigma^2_{y|x,T} &= \Omega_{yy,T} - (\Omega_{xy,T})^2/\Omega_{xx,T} \end{aligned} \tag{5}$$

where $(\Omega_{xy,T}, \Omega_{xy,T}, \Omega_{yy,T})$ are the elements of Ω_T. The linear regression coefficient β_T solves the population least-squares problem

$$\beta_T = \arg\min_b E\left[T^{-1}\sum_{t=1}^{T}(\widehat{y}_t - b\widehat{x}_t)^2\right],$$

so that β_T is the coefficient in the population best linear prediction of the long-run projection

$\hat{y}_t$ by the long-run projection $\hat{x}_t$,[4] $\sigma^2_{y|x,T}$ is the average variance of the prediction error, and $\rho^2_{xy,T}$ is the corresponding population R^2. These parameters thus measure the population comovement of the long-run variation of (x_t, y_t). Equivalently, by the second equality in (4), β_T also solves

$$\beta_T = \arg\min_b E\left[\sum_{j=1}^{q}(Y_{jT} - bX_{jT})^2\right]$$

with a corresponding interpretation $\sigma^2_{y|x,T}$ and $\rho^2_{xy,T}$. Thus, these parameters equivalently measure the (population) linear dependence in the scatter plots in Figures 2.c and 3.c.

The objective of the remaining analysis to develop inference about the parameters $(\rho_{xy,T}, \beta_T, \sigma^2_{y|x,T})$.

3 Asymptotic approximations and parameterizing long-run persistence and covariability

The long-run correlation and regression parameters are functions of Σ_T, the covariance matrix of (X_T, Y_T). This section takes up two related issues. The first is the asymptotic normality of the cosine-weighted averages (X_T, Y_T), which serves as the basis for the inference methods developed in Section 4 and provides large-sample approximation for the matrices Σ_T and Ω_T, and thus of the parameters of interest $\rho_{xy,T}$, β_T, and $\sigma^2_{y|x,T}$. The second issue involves parameterizing the form of long-run persistence and comovement, which determines the large sample value of Σ_T and Ω_T.

3.1 Large-sample properties of long-run sample averages

Because (X_T, Y_T) are smooth averages of (x_t, y_t), a central limit theorem effect suggests that these averages are approximately Gaussian under a range of primitive conditions about (x_t, y_t). The set assumptions under which asymptotic normality holds turns out to be reasonably broad, and encompasses many forms of potential persistence. Specifically, let $z_t = (x_t, y_t)'$ and suppose that Δz_t has moving average representation $\Delta z_t = C_T(L)\varepsilon_t$, where ε_t is a martingale difference sequence with non-singular covariance matrix, the coefficients

[4] The parameter β_T is closely related to a linear band-spectrum regression coefficient (Engle (1974)), corresponding to periods longer than $2T/q$.

in $C_T(L)$ die out sufficiently fast that Δz_t has a spectral density $F_{\Delta z,T}$, and ε_t and $C_T(L)$ satisfy other moment and decay restrictions given in Müller and Watson (2016, Theorem 1).[5] If the spectral density converges for all frequencies close to zero

$$T^{3-2\kappa}F_{\Delta z,T}(\omega/T) \to S_{\Delta z}(\omega)$$

in a suitable sense, then

$$T^{1-\kappa}\begin{pmatrix} X_T \\ Y_T \end{pmatrix} \Rightarrow \begin{pmatrix} X \\ Y \end{pmatrix} \sim \mathcal{N}(0,\Sigma), \tag{6}$$

and the finite-sample second moment matrix correspondingly converges to its large-sample counterpart (Müller and Watson (2016, Lemma 2))

$$T^{2-2\kappa}\,\mathrm{Var}\begin{pmatrix} X_T \\ Y_T \end{pmatrix} = T^{2-2\kappa}\Sigma_T \to \Sigma. \tag{7}$$

The limiting covariance matrix Σ in (6) and (7) is a function of the "local-to-zero" spectrum $S_{\Delta z}$ and the cosine weights $\Psi_j(s)$ that determine (X_T, Y_T); see Müller and Watson (2016) for additional details and an explicit formula. We make three comments about these large-sample results. First, they hold when the *first-difference* of z_t has a spectral density (and therefore has limited persistence); the *level* of z_t may (but doesn't have to) be more persistent. This is possible because the cosine averages sum to zero ($\Psi_T' l_T = 0$), so they do not extract zero-frequency variation in the data. If the *level* of z_t has a spectral density, then this restriction on the weights is not required and, for example, the sample mean of z_t also has a large-sample normal limit. Second, in common parameterizations of persistence models, the scale factor $T^{-\kappa}$ depends on the form of persistence; for example, the factor is $T^{-1/2}$ for $I(0)$ persistence and $T^{-3/2}$ for $I(1)$ persistence. However, we focus on inference procedures that do not depend on the scale of z_t (due to invariance or equivariance), so $T^{-\kappa}$ does not need to be known. Third, because $T^{2-2\kappa}\Sigma_T \to \Sigma$, then $T^{2-2\kappa}\Omega_T \to \Omega$ where Ω is defined as in the last expression of (4) with Σ in place of Σ_T. Correspondingly, $(\rho_{xy,T}, \beta_T, T^{2-2\kappa}\sigma^2_{y|x,T}) \to (\rho_{xy}, \beta, \sigma^2_{y|x})$ with the limits defined by (5) with Ω in place of Ω_T. Thus, a solution to the small-sample problem of inference about $(\rho_{xy}, \beta, \sigma^2_{y|x})$ from observing (X, Y) readily translates into a large-sample solution to inference about $(\rho_{xy,T}, \beta_T, \sigma^2_{y|x,T})$.

[5] The dependence of C_T and $F_{\Delta z,T}$ on the sample size T accommodates many forms of persistence that require double arrays as data generating process, such as autoregressive roots of the order $1 - c/T$, for fixed c. We omit the corresponding dependence of $z_t = (x_t, y_t)$ on T to ease notation.

3.2 Parameterizing long-run persistence and covariability

The limiting covariance matrix of the long-run projections, Ω, is a function of the covariance matrix of the cosine projections, Σ, which in turn is a function of the local-to-zero spectrum for the first-difference of z, $S_{\Delta z}$. The corresponding local-to-zero (pseudo-) spectrum for the level of z_t is $S_z(\omega) = \omega^{-2} S_{\Delta z}(\omega)$. In this section we discuss parameterizations of S_z, Σ, and Ω.

It is constructive to consider two leading examples. In the first, z_t is $I(0)$ with long-run covariance matrix Λ. In this case $S_z(\omega) \propto \Lambda$, and straightforward calculations show that $\Sigma = \Lambda \otimes I_q$ and $\Omega \propto \Lambda$, so the covariance matrix associated with the long-run projections corresponds to the usual long-run $I(0)$ covariance matrix. In this model, the cosine transforms (X_{jT}, Y_{jT}) plotted in Figures 2 and 3 are, in large samples and up to some deterministic scale, i.i.d. draws from a $\mathcal{N}(0, \Lambda)$ distribution. Inference about $\Omega = \Lambda$ and $(\rho_{xy}, \beta, \sigma^2_{y|x})$ thus follows from well-known small sample inference procedures for Gaussian data (see Müller and Watson (2016)). In the second example, z_t is $I(1)$ with Λ the long-run covariance matrix for Δz_t. In this case $S_z(\omega) \propto \omega^{-2}\Lambda$, and a calculation shows that $\Sigma = \Lambda \otimes D$, where D is a $q \times q$ diagonal matrix with j'th diagonal element $D_{jj} = (j\pi)^{-2}$. In this model, the cosine transforms (X_{jT}, Y_{jT}) plotted in Figures 2 and 3 are, in large samples and up to some deterministic scale, independent but heteroskedastic draws from $\mathcal{N}(0, (j\pi)^{-2}\Lambda)$ distributions. Thus $\Omega \propto \Lambda$, so the covariance matrix for long-run projections for z_t corresponds to the long-run covariance matrix for its first differences, Δz_t. By weighted least squares logic, inference for $I(1)$ processes follows after reweighting the elements of (X_{jT}, Y_{jT}) by the square roots of the inverse of the diagonal elements of D and then using the same results as in the $I(0)$ model.

GDP, consumption, short-, and long-term interest rates: Table 1 presents estimates and confidence sets for $(\rho_T, \beta_T, \sigma^2_{y|x.T})$ using (X_T, Y_T) with $q = 12$ for GDP and consumption (panel a) and short- and long-term interest rates (panel b). Results are presented for $I(0)$ and $I(1)$ models, and for a more general model of persistence introduced below. For now, focus on the $I(0)$ and $I(1)$ results. The point estimates shown in the table are MLEs, and confidence intervals for $(\beta_T, \sigma^2_{y|x,T})$ are computed using standard finite-sample normal linear regression formulae (after appropriate weighting in $I(1)$ model), and confidence sets for ρ_T are constructed as in Anderson (1984, section 4.2.2).

For GDP and consumption, there are only minor differences between the $I(0)$ and $I(1)$

Table 1: Long-run covariability estimates and confidence intervals from the I(0), I(1), and (*A,B,c,d*) models

a. GDP and consumption

		ρ	β	$\sigma_{y\|x}$
I(0)	Estimate	0.93	0.76	0.35
	67% CI	0.87, 0.96	0.67, 0.85	0.30, 0.46
	90% CI	0.81, 0.97	0.60, 0.92	0.26, 0.55
I(1)	Estimate	0.93	0.84	0.35
	67% CI	0.88, 0.96	0.74, 0.94	0.29, 0.45
	90% CI	0.82, 0.97	0.67, 1.01	0.26, 0.54
(*A,B,c,d*)	Estimate	0.91	0.76	0.09
	67% CI	0.83, 0.96	0.66, 0.86	0.05, 0.12
	90% CI	0.71, 0.97	0.48, 0.95	0.05, 0.19
	67% Bayes CS	*0.83, 0.96*	*0.66, 0.86*	*0.05, 0.12*
	90% Bayes CS	*0.71, 0.97*	*0.58, 0.95*	*0.05, 0.19*

b. Short- and long-term interest rates

		ρ	β	$\sigma_{y\|x}$
I(0)	Estimate	0.97	0.96	0.63
	67% CI	0.95, 0.98	0.89, 1.03	0.53, 0.81
	90% CI	0.93, 0.99	0.84, 1.08	0.47, 0.97
I(1)	Estimate	0.94	0.85	0.48
	67% CI	0.88, 0.96	0.76, 0.95	0.40, 0.62
	90% CI	0.82, 0.97	0.68, 1.03	0.36, 0.74
(*A,B,c,d*)	Estimate	0.96	0.92	0.15
	67% CI	0.92, 0.98	0.83, 1.05	0.09, 0.22
	90% CI	0.89, 0.98	0.75, 1.14	0.09, 0.40
	67% Bayes CS	*0.92, 0.97*	*0.83, 1.00*	*0.09, 0.22*
	90% Bayes CS	*0.89, 0.98*	*0.75, 1.07*	*0.09, 0.40*

Notes: The rows labeled "Estimate" are the maximum likelihood estimates using the large-sample distribution of the cosine transforms for the I(0) and I(1) models, and are the posterior median based on the *I*(d) model for the (*A,B,c,d*) model. "CI" denotes confidence interval, which is calculated as described in the text. "Bayes CS" are Bayes equal-tailed credible sets based on the posterior from the *I*(d) model.

estimates and confidence sets. The estimated long-run correlation is greater than 0.9, the lower range of the 90% confidence interval exceeds 0.8 in both the $I(0)$ and $I(1)$ models. Thus, despite the limited long-run information in the sample (captured here by the 12 observations making up (X_T, Y_T)), the evidence points to a large long-run correlation between GDP and consumption. The long-run regression of consumption onto GDP yields a regression coefficient that is estimated to be 0.76 in the $I(0)$ model and 0.84 in $I(1)$ model. This estimate is sufficiently accurate that $\beta = 1$ is not included in the 90% $I(0)$confidence set. The results for long-term and short-term nominal interest rates are similarly informative — for example, there is strong evidence that the series are highly correlated over the long-run — although the $I(0)$ and $I(1)$ results differ more sharply than for GDP and consumption. To take just one example, the 90% confidence interval for β ranges from 0.68 to 1.03 in the $I(1)$ model but is narrower (0.84 to 1.08) in the $I(0)$ model.

These empirical results raise two related questions: which of the $I(0)$ or $I(1)$ models fit the data better better, and more generally, are either of these models adequate descriptions of the long-run properties of the series? Figure 4 provides some suggestive evidence. The figure is constructed under the assumption that the data are generated as linear combinations of two independent fractional processes, the first $I(d_1)$ and the second $I(d_2)$. The model thus nests the $I(0)$ model ($d_1 = d_2 = 0$), the $I(1)$ model ($d_1 = d_2 = 1$), but also allows $d_1 \neq d_2$ to take on values between -0.4 and 1.0. Figure 4.a plots the (d_1, d_2) log-likelihood contours for GDP and consumption, and Figure 4.b shows the corresponding contours for short- and long-term interest rates.[6] Both plots are normalized so the log-likelihood value takes on a value of 0 for the $I(0)$ model. The likelihood surfaces imply that a wide range of (d_1, d_2) values are consistent with the data. For GDP and consumption, the $I(0)$ model fits much better than the $I(1)$ model (the likelihood difference is 3.4 log-points), but an intermediate model with $d_1 \approx d_2 = 0.3$ fits best. More to the point, a range of values of d_1, d_2 between -0.2 and 0.5 yield likelihoods within one log-point of the maximum. For interest rates, the $I(1)$ model fits much better than the $I(0)$ model (the likelihoods differ by 6.4 log-points), but an intermediate model with $(d_1, d_2) = (0.9, 0.7)$ achieves the best fit, with a value 0.8 log-points higher than the $I(1)$ model. Moreover, models with one highly persistent component ($d_1 \approx 0.9$) and one less persistent component ($d_1 \approx 0.2$) also provide good fits. Thus, while the data are sufficiently informative to rule out some models of persistence (the $I(1)$ model

[6] We report the asymptotic profiled likelihood based on the normal approximation (6) to (X_T, Y_T), with the linear combination parameters smaximized out for each value of the pair $(d_1,\, d_2)$.

Figure 4: Log-likelihood values in the *I*(*d*) model

a. GDP and consumption

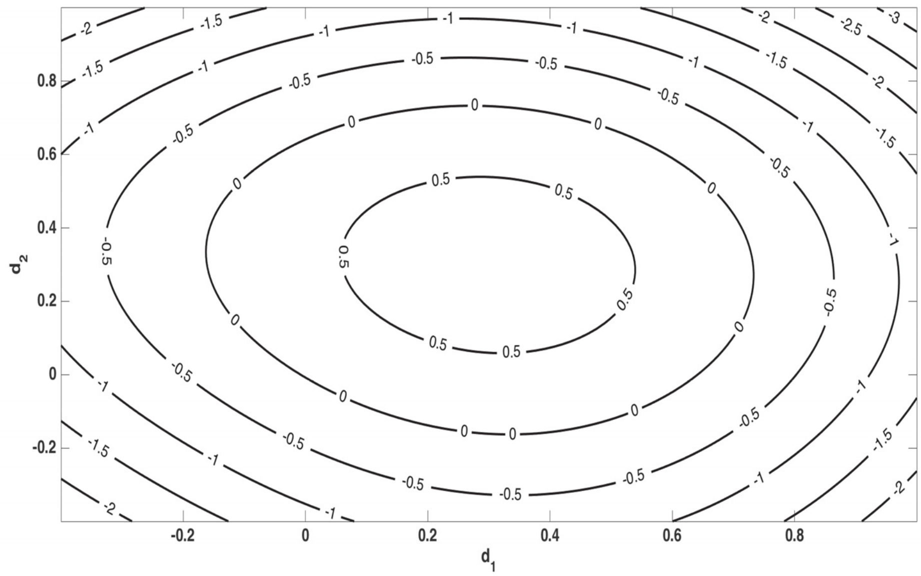

b. Short- and long-term interest rates

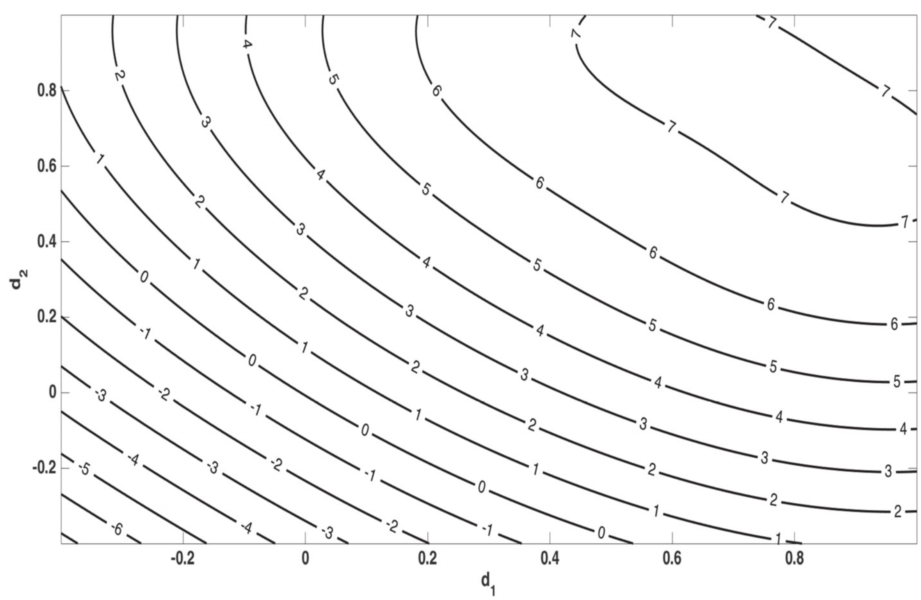

Notes: The figures show the log-likelihood contours for d_1 and d_2 for the bivariate *I*(*d*) model. The log-likelihood is normalized to be equal to zero for the *I*(0) model.

for GDP/consumption and the $I(0)$ model for interest rates), they are consistent with a range of persistence parameters.

3.2.1 (A, B, c, d) model

The shape of the local-to-zero spectrum determines the long-run persistence properties of the data, and misspecification of this persistence leads to faulty inference about long-run covariability. Thus, parameterizing S_z is a crucial issue for inference about long-run covariability. Addressing this issue faces a familiar trade-off: the parameterization needs to be sufficiently flexible to yield reliable inference about long-run covariability for a wide range of economically-relevant stochastic processes, and yet be sufficiently constrained to be tractable. $I(0)$ persistence generates a flat local-to-zero spectrum, and $I(1)$ persistence generates a local-to-zero spectrum proportional to ω^{-2}. Both of these models are tractable, but tightly constrain the spectrum. This limits their usefulness as general models for conducting inference about long-run covariability. The bivariate $I(d)$ model adds more flexibility at the cost of a few additional parameters, but it may be too tightly parameterized to provide reliable inference for the range of persistence patterns seen in economic time series.

With this trade-off in mind, we use a parameterization that nests and generalizes a range of models previously used to model persistence in economic time series. The parameterization is a bivariate extension of the univariate (b, c, d) model used in Müller and Watson (forthcoming) and yields a local-to-zero spectrum of the form

$$S_z(\omega) \propto A \begin{pmatrix} (\omega^2 + c_1^2)^{-d_1} & 0 \\ 0 & (\omega^2 + c_2^2)^{-d_2} \end{pmatrix} A' + BB' \tag{8}$$

where A is unrestricted and B is lower triangular.[7]

This model generates the standard spectral shapes: $A = 0$ yields the $I(0)$ model; $B = 0$, $c = 0$ and $d_1 = d_2 = 1$ yields the $I(1)$ model; $B = 0$, $d_1 = d_2 = 1$ yields a model with two AR roots local-to-unity; $B = 0$ and $c = 0$ yields the bivariate $I(d)$ introduced above. Other choices of (A, B, c, d) yield models that combine persistent and non-persistent components (as in cointegrated or "local-level" models) but go beyond the usual $I(0)/I(1)$ formulations.

[7] This is the spectrum of a bivariate Whittle-Matérn (c.f., Lindgren (2013)) process with time series representation $z_t = A\tau_t + e_t$, where τ_t is a bivariate process with independent elements, $(1 - \phi_{i,T}L)^{d_i}\tau_i = T^{-d_i/2}\varepsilon_{it}$, $\phi_{i,T} = (1 - c_i/T)$, $\varepsilon_t \sim I(0)$ with long-run variance equal to I_2, $e_t \sim I(0)$ with long-run variance equal to BB', and zero long-run covariance with ε_t.

The cost of the (A, B, c, d) model's flexibility is that it contains 11 parameters as opposed to just 3 in the $I(0)$ and $I(1)$ models or 6 in the bivariate $I(d)$ model. Yet, as we discuss in the next section, it is still possible conduct valid inference even in the loosely parameterized (A, B, c, d) model.

4 Constructing confidence intervals for ρ, β, and $\sigma_{y|x}$

4.1 An overview

There are several approaches one might take to construct confidence intervals for the parameters ρ, β , and $\sigma_{y|x}$. As a general matter, the goal is to compute confidence intervals that are as informative ("narrow") as possible, subject to the coverage constraint that they contain the true value of the parameter of interest with a pre-specified probability. We construct confidence intervals by explicitly solving a version of this problem.

Generically, let θ denote the vector of parameters characterizing the probability distribution of (X, Y), and let Θ denote the parameter space. (In our context, θ denotes the parameters in (A, B, c, d).) Let $\gamma = g(\theta)$ denote the parameter of interest. ($\gamma = \rho$, β, or $\sigma_{y|x}$ for the problem we consider). Let $H(X, Y)$ denote a confidence interval for γ and $\text{vol}(H(X, Y))$ denote the length of the interval. The objective is to choose H so that it has small expected length, $E[\text{vol}(H(X, Y)]$, subject to coverage, $P(\gamma \in H(X, Y)) \geq 1 - \alpha$, where α is a pre-specified constant. Because the probability distribution of (X, Y) depends on θ, so will the expected length of $H(X, Y)$ and the coverage probability. By definition, the coverage constraint must be satisfied for all values of $\theta \in \Theta$, but one has freedom in choosing the value of θ over which expected length is to be minimized. As a general matter, let W denote a distribution that puts weight on different values of θ, so the problem becomes

$$\min_{H(.)} \int E_\theta(\text{vol}(H(X, Y))dW(\theta) \tag{9}$$

subject to

$$\sup_{\theta \in \Theta} P_\theta (\gamma \in H(X, Y)) \geq 1 - \alpha \tag{10}$$

where the objective function (9) emphasizes that the expected volume depends on the value of θ, with different values of θ weighted by W, and the coverage constraint (10) emphasizes that the constraint must hold for all values of θ in the parameter space Θ.

As noted by Pratt (1961), the expected length of confidence set for γ can be expressed in terms of the power of hypothesis tests of $H_0 : \gamma = \gamma_0$. The solution to (9)-(10) thus amounts to the determination of a family of most powerful hypothesis tests, indexed by γ_0. Elliott, Müller, and Watson (2015) suggest a numerical approach to compute corresponding approximate "least favorable distributions" for θ. We implement a version of those methods here; details are provided in the appendix. A key feature of the solution is that, conditional on the weighting function W and the least favorable distribution, the confidence sets have the familiar Neyman-Pearson form with a version of the likelihood ratio determining the values of γ included in the confidence interval.

While the resulting confidence intervals have (close to) smallest weighted expected length, they can have unreasonable properties for particular realizations of (X, Y). Indeed, for some values of (X, Y), the confidence intervals might be empty, with the uncomfortable implication that, conditional on observing these values of (X, Y), one is certain that the confidence interval excludes the true value. To avoid this, we follow Müller and Norets (2016) and restrict the confidence sets to be supersets of $1 - \alpha$ Bayes credible sets.

4.2 Some specifics

4.2.1 Invariance and equivariance

Correlations are invariant to the scale of the data. The linear regression of y onto x is the same as the regression of $y + bx$ onto x after subtracting b from the latter's regression coefficient. It is sensible to impose the same invariance/equivariance on the confidence intervals. Thus, letting H_ρ, H_β, and $H_{\sigma_{y|x}}$ denote confidence sets for ρ, β, and $\sigma_{y/x}$, we restrict these sets as follows:

$$\rho \in H_\rho(X, Y) \Leftrightarrow \rho \in H_\rho(b_x X, b_y Y) \text{ for } b_x b_y > 0 \tag{11}$$

$$\beta \in H_\beta(X, Y) \Leftrightarrow \frac{b_y \beta + b_{yx}}{b_x} \in H_\beta(b_x X, b_y Y + b_{yx} X) \text{ for } b_x, b_y \neq 0 \text{ and all values of } b_{yx} \tag{12}$$

$$\sigma_{y|x} \in H_{\sigma_{y|x}}(X, Y) \Leftrightarrow |b_y| \sigma_{y|x} \in H_{\sigma_{y|x}}(b_x X, b_y Y + b_{yx} X) \text{ for } b_x, b_y \neq 0 \text{ and all values of } b_{yx}. \tag{13}$$

These invariance/equivariance restrictions lead to two modifications to the solution to (9)-(10). First, they require the use of maximal invariants in place of the original (X, Y). The density of the maximal invariants for each of these transformations is derived in the appendix.

Second, because the objective function (9) is stated in terms of (X, Y), minimizing expected length by inverting tests based on the maximal invariant leads to a slightly different form of optimal test statistic. Müller and Norets (2016) develop these modifications in a general setting, and the appendix derives the resulting form of confidence sets for our problem.

4.2.2 Parameter space

The parameter space for $\theta = (A, B, c, d)$ is as follows: A and B are real, with B lower-triangular and (A, B) chosen so that Ω is non-singular, $0 \leq c_i \leq 400$, and $-0.4 \leq d_i \leq 1$, for $i = 1, 2$. Thus, the confidence intervals control coverage over a wide range of persistence patterns including processes less persistent than $I(0)$, as persistent as $I(1)$, local-to-unity autoregressions, and where different linear combinations of x and y may have markedly different persistence (as, for example, in a cointegrated model).

The confidence sets we construct require three distributions over θ: the weighting function W for computing the average length in the objective (9), the Bayes prior associated with the Bayes credible sets that serve as subsets for the confidence sets (Müller and Norets (2016)), and the least favorable distribution for θ that enforces the coverage constraint. The latter is endogenous to the program (9)-(10) and is approximated using numerical methods similar to those discussed in Elliott, Müller, and Watson (2015), with details provided in the appendix. We use the same distribution for W and the Bayes prior. Specifically, the distribution is based on the bivariate $I(d)$ model (so that $c_1 = c_2 = 0, B = 0$) with d_1 and d_2 independently distributed $U(-0.4, 1.0)$. Because of the invariance/equivariance restrictions, the scale of the matrix A is irrelevant and we set $A = R(\lambda_1)G(s)R(\lambda_2)$, where $R(\lambda)$ is a rotation matrix indexed by the angle λ, with λ_1 and λ_2 independently distributed $U[0, \pi]$. The relative eigenvalues of A are determined by the diagonal matrix $G(s)$, with $G_{11}/G_{22} = 15^s$ with s distributed $U[0, 1]$.

4.2.3 Empirical results for GDP, consumption, and interest rates

Table 1 in Section 3.2 above shows estimates for $(\rho_T, \beta_T, \sigma_{y|x,T})$ and confidence sets using the (A, B, c, d) model. The estimated value of $(\rho_T, \beta_T, \sigma_{y|x,T})$ is the median of the posterior using the $I(d)$-model prior, and the table also shows Bayes credible sets for this prior for comparison with the frequentist confidence intervals. For GDP and consumption, the (A, B, c, d) results look much like the results obtained for the $I(0)$ model. For most entries, the Bayes credible

Table 2: Coverage rates for efficient 90% confidence intervals with data generated by different stochastic processes

Efficient confidence set for	Data generated by:				
	I(0)	I(1)	I(0)+I(1)	I(*d*)	(*A*,*B*,*c*,*d*)
I(0)	0.90	0.01	0.01	0.01	0.01
I(1)	0.00	0.90	0.00	0.00	0.00
I(0)+I(1)	0.91	0.91	0.90	0.68	0.68
I(*d*)	0.90	0.90	0.87	0.90	0.87
(*A*,*B*,*c*,*d*)	0.91	0.90	0.90	0.90	0.90

sets are slightly larger than the $I(0)$ sets, presumably reflecting the possibility of persistence greater than $I(0)$, as was evident in Figure 4. The frequentist confidence intervals often coincide with Bayes intervals, but occasionally are somewhat wider. The results indicate that GDP and consumption are highly correlated in the long-run (the 90% confidence set is $0.71 \leq \rho \leq 0.97$) and the long-run regression coefficient of consumption onto GDP is large, but less than unity (the 90% confidence set is $0.48 \leq \beta \leq 0.95$). The results for interest rates are somewhat different. The confidence intervals (and Bayes credible sets) are roughly in-between the $I(0)$ and $I(1)$ intervals, a result consistent with the $I(d)$ likelihood values plotted in Figure 4. Substantively, the results indicate that long-run movements in short- and long-rates are highly correlated, and that a unit long-run response of long-rates to short-rates is consistent with these data.

4.3 Coverage properties of restricted versions of the (A, B, c, d) model

In this subsection we investigate the coverage distortions for confidence intervals constructed using misspecified models of persistence. Specifically we consider five models of persistence, and for each model we both generate data and construct confidence intervals for ρ. The data are generated using $\rho = 0$ and Table 2 shows the fraction of the confidence sets that include the true value $\rho = 0$.[8] The models considered are the $I(0)$ model ($S_z(\omega) \propto BB'$), the $I(1)$

[8]Results are shown for confidence sets that do not incorporate the Müller-Norets Bayes superset adjustment. Including this adjustment yields similar results.

model ($S_z(\omega) = \omega^{-2}AA'$), a bivariate "local-level" that includes $I(0)$ and $I(1)$ components ($S_z(\omega) \propto \omega^{-2}AA' + BB'$), the fractional $I(d)$ model ($S_z(\omega) \propto ADA'$, D diagonal with $D_{jj} = \omega^{-2d_j}$) and the general (A, B, c, d) model with $S_z(\omega)$ given by (8). Because data were generated and confidence intervals constructed by each of these five models, the table contains 25 entries. The columns indicate the model used to generate the data, the rows shows the model used to construct the confidence set, and the entries are fraction of confidence sets that contain the true value of ρ, minimized over the other parameters used to generate the data. The diagonal entries of the table are 0.90 indicating that each method has coverage 90% under its assumed data generating process. The off-diagonal differ from 0.90 and show the coverage distortions. For example, 90% $I(0)$ confidence sets have coverage of just 1% when the data are generated by the other four models. $I(1)$ confidence sets have similarly bad coverage when the data are not generated by the $I(1)$ model. The $I(0) + I(1)$ model encompasses both the $I(0)$ and $I(1)$ models, so the associated confidence intervals has good coverage for these models, but has coverage of only 68% in the $I(d)$ and (A, B, c, d) models. The $I(d)$ model encompasses the $I(0)$ and $I(1)$ models, and so has good coverage for these models. It does not encompass the the $I(0)+I(1)$ or (A, B, c, d) models, but exhibits only a small coverage distortion in these cases. Finally, the general (A, B, c, d) model encompasses all of the other models, and so controls coverage uniformly across these models.

Table 2 highlights the large coverage distortions associated with confidence intervals based on $I(0)$, $I(1)$, or $I(0) + I(1)$ models. These results echo results in the earlier literature on the fragility of $I(0)$ and $I(1)$ inference (e.g., den Haan and Levin (1997) for HAC inference in $I(0)$ models and Elliott (1998) for inference in cointegrated models). Table 2 suggests that inference based on the $I(d)$ model is much less fragile; indeed it offers near nominal coverage in Table 2. However, the $I(d)$ model does not fare as well in other contexts; for example Müller and Watson (forthcoming) show that $I(d)$ model yields long-run prediction sets with significant undercoverage when data are generated by a univariate analogue of the (A, B, c, d) model.

5 Empirical Analysis

The last section showed results for the long-run covariation between GDP and consumption and between short- and long-term nominal interest rates. In this section we use the same methods to investigate other important long-run correlations. We focus on two questions:

first, how much information does the sample contain about the long-run covariability, and second, what are the values of the long-run covariability parameters. A knee-jerk reaction to investigating long-run propositions in economics using, say, 68-year spans of data is that little can be learned, particularly so using analysis that is robust to a wide range of persistence patterns. In this case, even efficient methods for extracting relevant information from the data will yield confidence intervals that are so wide that they rule out few plausible parameter values. We find this to be true for some of the long-run relationships investigated below. But, as we have seen from the consumption-income and interest rate data, confidence intervals about long-run parameters can be narrow and informative, and this is true for several of the relationships that we now investigate.

5.1 Balanced growth correlations.

In the standard one-sector growth model, variations in per-capita GDP, consumption, investment, and in real wages arise from variations in total factor productivity (TFP). Balanced growth means that the consumption-to-income ratio, the investment-to-income ratio, and labor's share of total income are constant over the long run. This implies perfect pairwise long-run correlations between the logarithms of income, consumption, investment, labor compensation, and TFP. In this model, the long-run regression of the logarithm of consumption onto the logarithm of income has a unit coefficient, as do the same regressions with consumption replaced by investment or labor compensation. A long-run one-percentage point increase in TFP leads a long-run increase of $1/\alpha$ percentage points in the other variables, where α is labor's share of income. Of course, these implications involve the evolution of the variables over the untestable infinite long-run. That said, empirical analysis can determine how well these implications stand-up to the long-run variations that can be estimated using, for example, data spanning the post-WWII period. We use data for the U.S. and the methods discussed above to investigate these long-run balance growth propositions. The appendix contains a description of the data that are used.

Table 3 summarizes the results on the long-run correlations. The values above the main diagonal show point estimates constructed as the posterior median using the $I(d)$-model with prior discussed above, together with 67% confidence intervals using the general (A, B, c, d) model (shown in parentheses). The values below the main diagonal are the corresponding 90% confidence intervals using the (A, B, c, d) model. Table 4 reports results from selected

Table 3: Long-run correlations of GDP, consumption, investment, labor compensation, and TFP

	GDP	Cons.	Inv.	$w\times n$	TFP
GDP		0.91 (0.83, 0.96)	0.53 (0.29, 0.72)	0.98 (0.97, 0.99)	0.78 (0.64, 0.89)
Cons.	(0.71, 0.97)		0.53 (0.30, 0.72)	0.92 (0.86, 0.96)	0.70 (0.48, 0.82)
Inv.	(0.02, 0.81)	(0.03, 0.81)		0.51 (0.27, 0.71)	0.38 (0.05, 0.60)
$w\times n$	(0.95, 0.99)	(0.68, 0.97)	(0.02, 0.80)		0.72 (0.56, 0.85)
TFP	(0.45, 0.95)	(0.28, 0.91)	(-0.08, 0.71)	(0.38, 0.93)	

Notes. All variables are measured in growth rates. The entries above the diagonal show the median of the posterior distribution followed by the 67% confidence interval. The entries below the diagonal show the 90% confidence interval.

Table 4: Selected long-run regressions involving GDP, consumption, investment, labor compensation, and TFP

Y	*X*		β				$\hat{\sigma}_{y\|x}$
			$\hat{\beta}$	67% CI	90% CI		
Consumption	GDP		0.76	0.66, 0.86	0.48, 0.95		0.40
Investment	GDP		1.24	0.64, 1.78	0.21, 2.21		2.18
Labor comp. ($w\times n$)	GDP		1.29	1.22, 1.36	1.17, 1.41		0.27
GDP	TFP		1.21	0.91, 1.47	0.71, 1.71		0.74
Cons. (Nondurable)	GDP		0.35	0.13, 0.57	-0.06, 0.76		0.88
Cons. (Services)	GDP		0.83	0.66, 1.00	0.55, 1.21		0.59
Cons. (Durables)	GDP		1.86	1.47, 2.25	1.17, 2.55		1.49
Inv. (Nonresidential)	GDP		0.96	0.41, 1.46	-0.09, 1.89		2.18
Inv. (Residential)	GDP		2.17	0.87, 3.48	-0.27, 4.62		5.48
Inv. (Equipment)	GDP		0.81	0.12, 1.57	-0.41, 2.10		2.74

Notes. All variables are measured in growth rates. The entries were constructed from the long-run regression of the variable labeled *Y* onto the variable labeled *X*.

long-run regressions.

As reported in the previous section, the long-run correlation between GDP and consumption is large. Investment and GDP are less highly correlated; the *upper* bound of the 90% confidence interval is only 0.8 and the lower bound is close to zero. Labor income and GDP are highly correlated with a tightly concentrated 90% confidence interval of 0.95 to 0.99. The estimated long-run correlation of TFP and GDP is also high, although the correlation of TFP and the other variables appears to be somewhat lower.

Table 4 shows results from long-run regressions of the growth rates of consumption, investment, and labor income onto the growth rate of GDP, and the corresponding regression of GDP onto TFP. Labor compensation appears to vary more than one-for-one with GDP and (as reported above) consumption less than one-for-one. The long-run investment-GDP regression coefficient is imprecisely estimated. Disaggregating consumption into nondurables, durables, and services, suggests that durable consumption responds more to long-run variations in GDP than do services and non-durables. These long-run regression results are reminiscent of results using business cycle covariability, and in Section 6 we investigate the robustness of these results to the periodicities incorporated in the long-run analysis.

In summary, what has our 68-year sample been able to say about the balanced-growth implications of the simple growth model? First, that several of the variables are highly correlated over the long-run (labor compensation and GDP, consumption and GDP), and second that the long-run regression coefficient on GDP is different from unity for some variables (consumption and labor compensation). There is less information about the long-run covariability of investment with the other variables, although even here there are things to learn, such as the long-run correlation of investment and GDP is unlikely to much larger than 0.8.

5.2 Other long-run relations.

Table 5 shows long-run covariation measures for eight pairs of variables, using post-WWII U.S. data. (See the data appendix for description and sources.) We discuss each in turn.

CPI and PCE inflation. The first row of the table considers the long-run covariation of two widely-used measured of inflation, the first based on the consumer price index (CPI) and the second based on the price deflator for personal consumption expenditures (PCE). The Boskin Commission Report and related research (Boskin, Dulberger, Gordon, Griliches,

Table 5: Long-run covariation measures for selected variables

	Y	X	Trans.	ρ			β			$\hat{\sigma}_{y\vert x}$
				$\hat{\rho}$	67% CI	90% CI	$\hat{\beta}$	67% CI	90% CI	
1	CPI Infl.	PCE Infl.	L,L	0.98	0.96, 0.99	0.95, 0.99	1.13	1.07, 1.19	0.98, 1.24	0.10
2	3M rates	PCE Infl.	L,L	0.47	0.21, 0.82	-0.00, 0.91	0.73	0.34, 1.48	-0.09, 1.91	0.49
3	10Y rates	PCE Infl.	L,L	0.47	0.23, 0.82	-0.00, 0.91	0.66	0.30, 1.36	-0.03, 1.72	0.48
4	Un. Rate	PCE Infl.	L,L	0.25	-0.03, 0.58	-0.27, 0.82	0.21	-0.04, 0.45	-0.24, 0.78	0.43
5	Un. Rate	TFP	L,G	-0.65	-0.75, -0.35	-0.91, -0.15	-1.00	-1.39, -0.62	-1.64, -0.27	0.25
6	Stock Prices	Dividends	G,G	0.20	-0.05, 0.43	-0.30, 0.68	0.45	-0.17, 1.06	-0.60, 1.74	1.61
7	Stock Prices	Earnings	G,G	0.21	-0.04, 0.42	-0.27, 0.57	0.39	-0.15, 0.92	-0.53, 1.35	1.60
8	Exchange Rates	Rel. Price Ind.	G,G	0.41	0.13, 0.53	-0.08, 0.71	1.10	0.38, 1.82	-0.10, 2.46	1.49

Notes: The column labeled "Trans." indicates the transformation applied to the data with "L" denoting level (no transformation) and "G" denoting growth rate (in percentage points at annual rate using scaled first-differences of logarithms). Thus, the levels of inflation, interest rates, and the unemployment were used, and other variables were transformed to growth rates. The long-run regressions were computed from the regression of the variable labeled *Y* onto the variable labeled *X*.

and Jorgenson (1996), Gordon (2006)) highlights important methodological and quantitative differences in these two measures of inflation. For example, the CPI is a Laspeyres index, while the PCE deflator uses chain weighting, and this leads to greater substitution bias in the CPI. Differences in these inflation measures may change over time both because of the variance of relative prices (which affects substitution bias) and because measurement methods for both price indices evolved over the sample period.

The data are informative about the long-run covariability these two inflation measures. Table 5 suggests that PCE and CPI inflation are highly correlated in the long-run; the 90% confidence interval suggests that $\rho > 0.95$. The long-run regression of CPI inflation on PCE inflation yields an estimated slope coefficient that is 1.13 (90% confidence interval: $0.98 \leq \beta \leq 1.24$) suggesting a larger bias in the CPI during periods of high trend inflation.

Long-run Fisher correlation: The next two rows show the long-run covariation of inflation and short- and long-term nominal interest rates. The well-known Fisher relation (Fisher (1930)) decomposes nominal rates into an inflation and real interest rate component making it is interesting to gauge how much of the long-run variation in nominal rates can be explained by long-run variation in inflation. The long-run correlation of nominal interest rates and inflation is estimated to be approximately 0.5, although the confidence intervals indicate substantial uncertainty. A unit long-run regression coefficient of nominal rates onto inflation is consistent with data, but again the confidence intervals are very wide.[9] The data are therefore not very informative about the long-run Fisher correlation.

Long-run Phillips correlation: The next row of the table summarizes the long-run correlation between the unemployment and inflation. This estimated long-run Phillips correlation and slope coefficient are positive, but $\rho = \beta = 0$ is contained in the 67% confidence interval. That said, the confidence intervals are wide so that, like the Fisher correlation, the data are not very informative about the long-run Phillips correlation.

Unemployment and Productivity: The fifth row of the table investigates the long-run covariation of the unemployment rate and productivity growth. The data are informative about this covariability: there is a statistically significant negative long-run relationship between the variables, and a long-run one percentage point increase in the rate of growth of

[9] These estimates measure the long-run Fisher "correlation," not the long-run Fisher "effect". The long-run Fisher correlation considers variation from all sources, while the Fisher effect instead considers variation associated with exogenous long-run nominal shocks (e.g., Fisher and Seater (1993), King and Watson (1997)). A similar distinction holds for the Phillips correlation and the Phillips curve (King and Watson (1994)).

productivity is associated with an estimated one percentage point decline in the long-run unemployment rate. The large negative in-sample long-run correlation has been noted previously (e.g., Staiger, Stock, and Watson (2001)); the confidence intervals reported in Table 5 show that the correlation is unlikely to be spurious. We are unaware of an economically compelling theoretical explanation for the large negative correlation.

Stock Prices, Dividends, and Earnings: Present value models of stock prices show a tight link between long-run values of prices, dividends, and earnings (e.g., Campbell and Shiller (1987)). An implication of this long-run relation in a cointegration framework is that dividends, earnings, and stock prices share a common $I(1)$ trend, so that their growth rates are perfectly correlated in the long-run and the dividend-price or price-earning ratio is useful for predicting future stock returns. This latter implication has been widely investigated (see Campbell and Yogo (2006) for analysis and references). Table 5 shows the long-run correlation of stock prices with dividends and with earnings.[10] While there is considerable uncertainty about the value of the long-run correlation between prices and dividends or earnings, the data suggest that the correlation is not strong. For example, values above $\rho = 0.43$ are ruled out by the 67% confidence set and values above 0.68 are ruled out by the 90% sets.

Long-run PPP: The final row of the table shows results on the long-run correlation between nominal exchange rates (here the U.S. dollar/British pound exchange rate from 1971-2015) and the ratio of nominal prices (here the ratio of CPI indices for the two countries). Long-run PPP implies that the nominal exchange rate should move proportionally with the price ratio over long time spans, so the long-run growth rates of the nominal exchange rate and price ratios should be perfectly correlated. A large literature has tested this proposition in a unit-root and cointegration framework and obtained mixed conclusions. (See Rogoff (1996) and Taylor and Taylor (2004) for discussion and references). From the final row of Table 5, the growth rate of nominal exchange rates and relative nominal prices are positively correlated over the long-run, statistically significantly so at the 33% significance level, but the correlation is far from perfect ($\rho < 0.71$ based on the 90% confidence set.) We highlight two caveats. First, we use the post-Bretton Woods sample period, so the sample includes only 45 years, and using $q = 12$ cosine terms the long-run projections capture variability with periods of (approximately) 7 years or higher. This 7-year period may be sufficiently short

[10] The data are for the S&P, are updated updated versions of the data used in Shiller (2000), and were obtained from Robert Shiller's webpage. Table 5 is based on the data from 1947-2015.

that long-run adjustments have not occurred, something we investigate in the next section. Second, the price ratio uses relative CPIs, a large component of which includes non-traded goods which may be less tightly linked to exchange rates than prices of traded goods.

6 Alternative measures of long-run covariability

The empirical results in the last section relied on covariance measures associated with projections of the data onto $q = 12$ cosine functions capturing periodicities of $T/6$ or higher, where T is the length of the sample. Using data from 1948-2015 ($T = 68$ years) this analysis used periods longer than 11 years to define "long-run" variation and covariation. While 11 years is longer than typical business cycles, it does incorporates periods corresponding to what some researchers refer to as the "medium run" (Blanchard (1997), Comin and Gertler (2006)). In this section we consider measures of long-run covariability that focus on a subset of the q periods. This allows a comparison of, say, results from periods corresponding to the "medium-long run" and to those from the "longer-long run."

To motivate the new measures, look again at Figure 2.a which plots the projections of GDP and consumption growth rates onto $q = 12$ cosine regressors with periods that range from $T/6$ ($\approx$ 11 years) to $2T$ (136 years). Figures 5.b and 5.c show the corresponding projections onto the first $q_1 = 6$ of these cosine terms (with periods from $T/3 \approx 23$ years to $2T = 136$ years) and last $q_2 = 6$ cosine terms (with periods $T/12 \approx 11$ years to $2T/7 \approx 19$ years). The first of these captures the longer-long-run variation in the data, and the second captures the medium-long-run variability. Each can be studied separately. To differentiate these periodicities, we replace equation (4) with

$$\Omega_{i:j,T} = T^{-1}\sum_{t=1}^{T} E\left[\begin{pmatrix} \widehat{x}_{i:j,t} \\ \widehat{y}_{i:j,t} \end{pmatrix}\begin{pmatrix} \widehat{x}_{i:j,t} & \widehat{y}_{i:j,t} \end{pmatrix}\right] = E\begin{bmatrix} X'_{i:j,T}X_{i:j,T} & X'_{i:j,T}Y_{i:j,T} \\ Y'_{i:j,T}X_{i:j,T} & Y'_{i:j,T}Y_{i:j,T} \end{bmatrix} \tag{14}$$

where the subscript "$i:j$" notes that the projection is computed using the i through j cosine terms (i.e., the i through j columns of Ψ_T) corresponding to periods $2T/j$ through $2T/i$. Thus the longer-long-run periodicities shown in Figure 5.a correspond to the covariance matrix $\Omega_{1:6,T}$ (the first 6 cosine terms) and the medium-long-run periodicities in Figure 5.b correspond to $\Omega_{7:12,T}$ (the 7-12th cosine terms).

Throughout the paper we have used q to denote the number of low-frequency cosine terms that define the long-run periods of interest (perhaps divided further into longer-long

Figure 5: Long-run projections for GDP and consumption growth rates for different periodicities

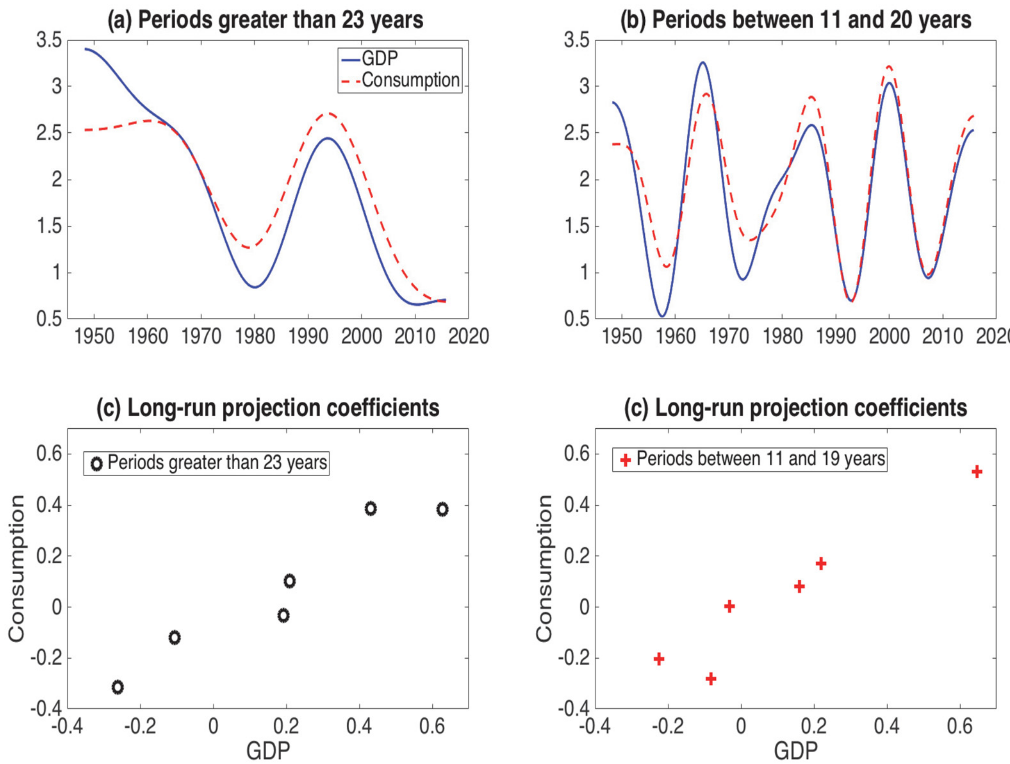

Notes: Panel (a) plots the projections of the data onto six cosine terms with periods 23-136 years. Panel (b) shows the projections onto six low-frequency terms with periods 11-19 years. Sample means have been added to both sets ofr projections. Panels (c) and (d) are scatterplots of the coefficients (cosine transforms) from panels (a) and (b).

and medium-long). But q plays another important role in the analysis. The value of Ω (or now $\Omega_{i:j}$) ultimately depends on the variability and persistence in the stochastic process as exhibited in the local-to-zero (pseudo-) spectrum S_z. This spectrum is parameterized by (A, B, c, d); see equation (8). We learn about the value of these parameters (and therefore the value of Ω) using the data $(X_{1:q,T}, Y_{1:q,T})$. Thus, q also denotes the sample variability in the data that is used to infer the value of the long-run covariance matrix Ω. So, while our interest might lie in the longer-long-run covariability captured in $\Omega_{1:6}$, the sample variability in $(X_{1:12,T}, Y_{1:12,T})$ might be used to learn about $\Omega_{1:6}$. While it is arguably most natural to match the variability in the data used for inference to the variability of interest, for example using $(X_{1:q,T}, Y_{1:q,T})$ to learn about $\Omega_{1:q}$, if the (A, B, c, d) model accurately characterizes the spectrum over a wider frequency band, then variability over this wider band can improve inference. But of course using a wider frequency band runs the risk of misspecification if the (A, B, c, d) model is a poor characterization of the spectrum over this wider range of frequencies. This is the standard trade-off of robustness and efficiency.

With these ideas in mind, Table 6 shows results for long-run correlation and regression parameters from $\Omega_{1:12}$, $\Omega_{1:6}$, and $\Omega_{7:12}$, corresponding the periods $T/6$ and higher, $T/3$ and higher, and $T/6$ through $2T/7$. Results are shown using inference based on the same $q = 12$ cosine transforms used in the sections above, but also using $q = 6$, so only lower frequency variability in the data is used to learn about (A, B, c, d), and with $q = 18$, so higher frequency variability is also used. Table 6.a shows results for long-run covariability of GDP, consumption, investment, labor compensation, and TFP. Table 6.b shows results for selected long-run relationships involving the other variables. (Results for all the pairs of variables shown in Table 3-5 are available in the appendix.)

The first block of results in Table 6.a are for consumption and GDP. The first row repeats earlier results using the $q = 12$ cosine terms to learn about $\Omega_{i:j}$ with $i = 1$ and $j = 12$. The other rows are for other values of q, i, and j. The results suggest remarkable stability across the different values of q, i, and j. Figure 5.c provides hints at this stability. It shows the scatter plot of $(X_{1:6,T}, Y_{1:6,T})$ and $(X_{7:12,T}, Y_{7:12,T})$ corresponding the projections plotted in panels 5.a and 5.b. The scatter plots corresponding to the different periodicities are quite similar, and this is reflected in the stability of the results shown in Table 6. This same stability across q, i, and j is evident for the other pairs of variables in Table 6.a. Looking closely at Table 6.a, there are subtle differences in the rows. For example, the confidence intervals for the parameters from $\Omega_{1:12}$ tend to be somewhat narrower using $q = 18$ than

Table 6: Long-run covariation measures for selected variables with $\Omega_{i:j}$
Inference based on q cosine transforms.

a. GDP, consumption, investment, labor compensation, and TFP

Y	X	q	$i{:}j$	$\rho_{i:j}$			$\beta_{i:j}$			$\hat{\sigma}_{i:j,y\vert x}$
				$\hat{\rho}$	67% CI	90% CI	$\hat{\beta}$	67% CI	90% CI	
Cons.	GDP	12	1:12	0.91	0.83, 0.96	0.71, 0.97	0.76	0.66, 0.86	0.48, 0.95	0.40
		12	1:6	0.92	0.82, 0.96	0.58, 0.98	0.74	0.55, 0.86	0.38, 0.96	0.30
		12	7:12	0.92	0.84, 0.96	0.74, 0.98	0.81	0.69, 0.92	0.60, 1.04	0.23
		18	1:12	0.91	0.84, 0.95	0.77, 0.97	0.75	0.66, 0.83	0.55, 0.91	0.40
		6	1:6	0.89	0.71, 0.96	0.53, 0.98	0.72	0.55, 0.88	0.37, 1.06	0.30
Inv.	GDP	12	1:12	0.53	0.29, 0.72	0.02, 0.81	1.24	0.64, 1.78	0.21, 2.21	2.18
		12	1:6	0.54	0.27, 0.76	-0.16, 0.91	1.18	0.58, 1.78	-0.27, 2.27	1.54
		12	7:12	0.53	0.25, 0.73	0.04, 0.82	1.30	0.70, 1.97	-0.03, 2.81	1.36
		18	1:12	0.63	0.42, 0.75	0.31, 0.81	1.60	1.05, 2.15	0.62, 2.58	2.41
		6	1:6	0.56	0.23, 0.80	-0.04, 0.91	1.08	0.53, 1.68	-0.12, 2.33	1.12
$w \times n$	GDP	12	1:12	0.98	0.97, 0.99	0.95, 0.99	1.29	1.22, 1.36	1.17, 1.41	0.27
		12	1:6	0.99	0.97, 0.99	0.95, 0.99	1.29	1.22, 1.37	1.17, 1.42	0.18
		12	7:12	0.98	0.95, 0.99	0.90, 0.99	1.29	1.21, 1.37	1.12, 1.45	0.19
		18	1:12	0.96	0.91, 0.97	0.87, 0.98	1.21	1.12, 1.31	1.04, 1.38	0.44
		6	1:6	0.98	0.95, 0.99	0.90, 0.99	1.30	1.21, 1.37	1.14, 1.44	0.16
GDP	TFP	12	1:12	0.78	0.64, 0.89	0.45, 0.95	1.21	0.91, 1.47	0.71, 1.71	0.74
		12	1:6	0.79	0.62, 0.93	0.43, 0.96	1.21	0.91, 1.53	0.68, 1.77	0.55
		12	7:12	0.73	0.56, 0.86	0.37, 0.91	1.15	0.83, 1.44	0.59, 1.71	0.44
		18	1:12	0.76	0.65, 0.87	0.49, 0.93	1.19	0.95, 1.46	0.74, 1.67	0.75
		6	1:6	0.71	0.41, 0.93	0.08, 0.96	1.09	0.71, 1.51	0.33, 1.93	0.57

Table 6: continued
b. Selected other variables

Y	X	q	i:j	$\rho_{i:j}$			$\beta_{i:j}$			$\hat{\sigma}_{i:j,y\|x}$
				$\hat{\rho}$	67% CI	90% CI	$\hat{\beta}$	67% CI	90% CI	
10Y rates	3M rates	12	1:12	0.96	0.92, 0.98	0.89, 0.98	0.92	0.83, 1.05	0.75, 1.14	0.70
		12	1:6	0.96	0.91, 0.98	0.88, 0.98	0.93	0.83, 1.09	0.75, 1.18	0.61
		12	7:12	0.94	0.87, 0.97	0.76, 0.98	0.87	0.67, 0.97	0.59, 1.05	0.28
		18	1:12	0.95	0.91, 0.97	0.83, 0.98	0.87	0.75, 1.04	0.67, 1.12	0.80
		6	1:6	0.96	0.92, 0.98	0.84, 0.99	0.97	0.86, 1.08	0.74, 1.20	0.61
3M rates	PCE Infl.	12	1:12	0.47	0.21, 0.82	-0.00, 0.91	0.73	0.34, 1.48	-0.09, 1.91	2.20
		12	1:6	0.52	0.23, 0.87	-0.02, 0.95	0.84	0.38, 1.63	0.06, 2.06	1.92
		12	7:12	0.23	-0.04, 0.56	-0.46, 0.70	0.36	-0.21, 0.72	-0.57, 1.07	0.88
		18	1:12	0.54	0.28, 0.82	0.09, 0.91	0.82	0.47, 1.50	0.03, 1.85	2.20
		6	1:6	0.65	0.30, 0.87	0.00, 0.95	1.02	0.55, 1.64	0.13, 2.16	1.98
Un. Rate	TFP	12	1:12	-0.65	-0.75, -0.35	-0.91, -0.15	-1.00	-1.39, -0.62	-1.64, -0.27	1.06
		12	1:6	-0.78	-0.91, -0.47	-0.96, -0.23	-1.21	-1.60, -0.79	-1.94, -0.44	0.75
		12	7:12	-0.31	-0.56, 0.00	-0.69, 0.32	-0.53	-1.00, 0.20	-1.34, 0.71	0.65
		18	1:12	-0.42	-0.76, -0.27	-0.91, -0.08	-0.91	-1.34, -0.49	-1.64, -0.19	1.17
		6	1:6	-0.89	-0.95, -0.72	-0.98, -0.57	-1.79	-2.34, -1.35	-2.75, -1.00	0.51
Stock Prices	Dividends	12	1:12	0.20	-0.05, 0.43	-0.30, 0.68	0.45	-0.17, 1.06	-0.60, 1.74	7.27
		12	1:6	0.36	-0.00, 0.68	-0.25, 0.91	0.76	-0.05, 1.56	-0.60, 3.16	5.26
		12	7:12	0.03	-0.46, 0.29	-0.68, 0.47	-0.11	-0.79, 0.69	-1.28, 1.31	4.03
		18	1:12	0.39	0.14, 0.55	-0.06, 0.68	0.75	0.31, 1.24	-0.08, 1.63	6.34
		6	1:6	0.42	0.02, 0.71	-0.20, 0.87	1.15	0.28, 1.94	-0.50, 2.72	5.56

Notes: Results are shown for a subset of the variables listed in Tables 3-5. Results are based on $\Omega_{i:j}$ (col. 4 lists i and j) and sample information in q cosine transforms (col. 3 shows q).

using $q = 12$, consistent with a modest amount of additional information using a larger value of q. The same result holds for results for $\Omega_{1:6}$ computed using $q = 6$ and $q = 12$.

The results summarized in Table 6.b show much of the same stability as Table 6.a, but there are some notable differences. For example, the point estimates suggest a somewhat larger Fisher correlation over longer periods (greater than 23 years) than over shorter periods (11 to 19 years), and the same holds for stock prices and dividends. In both cases however, the confidence intervals remain wide. And, the puzzling negative correlation between the unemployment rate and TFP appears to be stronger over the longer-long run than over the medium-long run.

7 Concluding remarks

This paper has focused on inference about long-run covariability of two time series. Just as with previous frameworks, such as cointegration analysis, it is natural to consider a generalization to a higher dimensional setting. For example, this would allow one to determine whether the significant long-run correlation between the unemployment rate and productivity is robust to including a control for, say, some measure of human capital accumulation.

Many elements of our analysis generalize to n time series in a straightforward manner: The analogous definition of Ω_T is equally natural as a second-moment summary of the covariability of n series, and gives rise to corresponding regression parameters, such as coefficients from a $n-1$ dimensional multiple regression, corresponding residual standard deviations and population R^2s.[11] Multivariate versions of Ω_T can also be used for long-run instrumental variable regressions. As shown in Müller and Watson (2016), the Central Limit Theorem that reduces the inference question to one about the covariance matrix of a multivariate normal holds for arbitrary fixed n. The (A, B, c, d) model of persistence naturally generalizes to a n dimensional system. And, confidence sets for multiple regression parameters satisfy natural invariance and equivariance constraints, which reduces the number of effective parameters.

Having said that, our numerical approach for constructing (approximate) minimal-length confidence sets faces daunting computational challenges in a higher order system: The quadratic forms that determine the likelihood require $O(n^2q^2)$ floating point operations. Worse still, even for n as small as $n = 3$, the number of parameters in the (A, B, c, d) model

[11] Müller and Watson (2016) provide the details of inference in the $I(0)$ model.

is equal to 21. So even after imposing invariance or equivariance, ensuring coverage requires an exhaustive search over a high dimensional nuisance parameter space.

At the same time, it would seem to be relatively straightforward to determine Bayes credible sets also for larger values of n: Under our asymptotic approximation, the (A, B, c, d) parameters enter the likelihood through the covariance matrix of a $nq \times 1$ multivariate normal, so with some care, modern posterior samplers should be able to reliably determine the posterior for any function of interest. Of course, such an approach does not guarantee frequentist coverage, and the empirical results will depend on the choice of prior in a non-trivial way. In this regard, our empirical results in the bivariate system show an interesting pattern: Especially at a lower nominal coverage level, for many realizations, there is no need to augment the Bayes credible set computed from the bivariate fractional model. This suggests that the frequentist coverage of the unaltered Bayes intervals is not too far below the nominal level, so these Bayes sets wouldn't be too misleading even from a frequentist perspective.[12] While this will be difficult to exhaustively check, this pattern might well generalize also to larger values of n.

[12] In fact, a calculation analogous to those in Table 2 shows that the 67% Bayes set contains the true value of $\rho = 0$ at least 64.1% of the time in the bivariate (A, B, c, d) model, and the 95% Bayes set has coverage of 82.8%.

8 Appendix

In preparation

References

BAXTER, M., AND R. G. KING (1999): "Measuring business cycles: approximate band-pass filters for economic time series," *Review of economics and statistics*, 81(4), 575–593.

BLANCHARD, O. J. (1997): "The Medium Run," *Brooking Panel on Economic Activity*, pp. 89–158.

BOSKIN, M. J., E. DULBERGER, R. GORDON, Z. GRILICHES, AND D. JORGENSON (1996): "Toward a More Accurate Measure of the Cost of Living," *Final Report to the Senate Finance Committee.*

CAMPBELL, J. Y. (1987): "Does Savings Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis," *Ecomometrica*, 55, 1249–1273.

CAMPBELL, J. Y., AND R. J. SHILLER (1987): "Cointegration and Tests of Present Value Models," *Journal of Political Economy*, 95, 1062–1088.

CAMPBELL, J. Y., AND M. YOGO (2006): "Efficient Tests of Stock Return Predictability," *Journal of Financial Economics*, 81, 27–60.

COCHRANE, J. (1994): "Permanent and Transitory Components of GNP and Stock Prices," *Quarterly Journal of Economics*, CIX, 241–266.

COMIN, D., AND M. GERTLER (2006): "Medium Term Business Cycles," *American Economic Review*, 96, 523–551.

DAI, Q., AND K. SINGLETON (2000): "Specification analysis of affine term structure models," *Journal of Finance*, 55, 1943–1978.

DAVIDSON, J. E., D. F. HENDRY, F. SRBA, AND S. YEO (1978): "Econometric Modelling of the Aggregate Time-Series Relationship Between Consumers' Expenditure and Income in the United Kingdom," *Economic Journal*, 88, 661–692.

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. 299–342. Elsevier, Amsterdam.

Diebold, F. X., and C. Li (2006): "Forecasting the Term Structure of Government Bond Yields," *Journal of Econometrics*, 130, 337–364.

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 (2015): "Nearly Optimal Tests When a Nuisance Parameter is Present Under the Null Hypothesis," *Econometrica*, 83, 771–811.

Engle, R. F. (1974): "Band spectrum regression," *International Economic Review*, pp. 1–11.

Engle, R. F., and C. W. J. Granger (1987): "Co-Integration and Error Correction: Representation, Estimation, and Testing," *Econometrica*, 55, 251–276.

Fisher, I. (1930): *The Theory of Interest*. MacMillan, New York.

Fisher, M. E., and J. D. Seater (1993): "Long-run Neutrality and Superneutrality in an ARIMA Framework," *American Economic Review*, 83, 402–415.

Gordon, R. J. (2006): "The Boskin Commission Report: A Retrospective One Decade Later," *NBER working paper 12311*.

Granger, C. W. J. (1981): "Some Properties of Time Series Data and their Use in Econometric Model Specification," *Journal of Econometrics*, 16, 121–130.

Granger, C. W. J., and P. Newbold (1974): "Spurious Regressions in Econometrics," *Journal of Econometrics*, 2, 111–120.

Hodrick, R. J., and E. C. Prescott (1997): "Postwar US business cycles: an empirical investigation," *Journal of Money, Credit and Banking*, pp. 1–16.

Johansen, S. (1988): "Statistical Analysis of Cointegration Vectors," *Journal of Economic Dynamics and Control*, 12, 231–254.

Kiefer, N. M., T. J. Vogelsang, and H. Bunzel (2000): "Simple Robust Testing of Regression Hypotheses," *Econometrica*, 68, 695–714.

King, R., C. I. Plosser, J. H. Stock, and M. W. Watson (1991): "Stochastic Trends and Economic Fluctuations," *American Economic Review*, 81, 819–840.

King, R. G., and M. W. Watson (1994): "The Post-War U.S. Phillips Curve: A Revisionist Econometric History," *Carnegie-Rochester Conference on Public Policy*, 41, 157–219.

——— (1997): "Testing long-run neutrality," *Federal Reserve Bank of Richmond Economic Quarterly*, 83, 69–101.

Klein, L. R., and R. F. Kosobud (1961): "Some Econometrics of Growth: Great Ratios of Economics," *Quarterly Journal of Economics*, LXXV, 173–198.

Lazarus, E., D. J. Lewis, and J. Stock (2016): "HAR Inference: Kernel Choice, Size Distortions, and Power Loss," *Working Paper, Harvard University.*

Lettau, M., and S. Ludvigson (2013): *Shocks and Crashes.* MIT Press, Cambridge MA.

Lindgren, G. (2013): *Stationary Stochastic Processes: Theory and Applications.* Chapman & Hall/CRC.

Müller, U. K. (2004): "A Theory of Robust Long-Run Variance Estimation," *Working paper, Princeton University.*

——— (2007): "A Theory of Robust Long-Run Variance Estimation," *Journal of Econometrics*, 141, 1331–1352.

——— (2014): "HAC Corrections for Strongly Autocorrelated Time Series," *Journal of Business and Economic Statistics*, 32, 311–322.

Müller, U. K., and A. Norets (2016): "Credibility of Confidence Sets in Nonstandard Econometric Problems," *forthcoming, Econometrica.*

Müller, U. K., and M. W. Watson (2008): "Testing Models of Low-Frequency Variability," *Econometrica*, 76, 979–1016.

——— (2016): "Low-Frequency Econometrics," *Working Paper, Princeton University.*

——— (forthcoming): "Measuring Uncertainty about Long-Run Predictions," *Review of Economic Studies.*

Phillips, P. (1991): "Optimal Inference in Cointegrated Systems," *Ecomometrica*, 59, 283–306.

Phillips, P. C. B. (2005): "HAC Estimation by Automated Regression," *Econometric Theory*, 21, 116–142.

Pratt, J. W. (1961): "Length of Confidence Intervals," *Journal of the American Statistical Association*, 56, 549–567.

Rogoff, K. (1996): "The Purchasing Power Parity Puzzle," *Journal of Economic Literature*, 34, 647–668.

Sargan, J. D. (1964): "Wages and Prices in the UK: A Study in Econometric Methodology," in *Econometric Analysis for National Planning*, ed. by P. E. Hart, G. Mills, and J. K. Whitaker, pp. 25–26, London. Butterworths.

Shiller, R. J. (2000): *Irrational Exuberance*. Princeton University Press, Princeton.

Staiger, D., J. H. Stock, and M. W. Watson (2001): "Prices, Wages, and the U.S. NAIRU in the 1990s," in *The Roaring Nineties*, ed. by A. B. Krueger, and R. Solow, pp. 3–60, New York. The Russell Sage Foundation.

Stock, J. H., and M. W. Watson (1993): "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems," *Econometrica*, 61, 783–820.

Sun, Y. (2013): "Heteroscedasticity and Autocorrelation Robust F Test Using Orthonormal Series Variance Estimator," *The Econometrics Journal*, 16, 1–26.

Taylor, A. M., and M. P. Taylor (2004): "The Purchasing Power Parity Debate," *Journal of Economic Perspectives*, 18, 135–158.