Macroeconomic Dynamics, **3**, 1999, 69–83. Printed in the United States of America.

DETECTION AND MODELING OF REGRESSION PARAMETER VARIATION ACROSS FREQUENCIES

WITH AN APPLICATION TO TESTING THE PERMANENT INCOME HYPOTHESIS

HUI BOON TAN
Universiti Putra Malaysia

RICHARD ASHLEY
Virginia Polytechnic Institute and State University

A simple technique for directly testing the parameters of a time-series regression model for instability across frequencies is presented. The method can be implemented easily in the time domain, so that parameter instability across frequency bands can be conveniently detected and modeled in conjunction with other econometric features of the problem at hand, such as simultaneity, cointegration, missing observations, and cross-equation restrictions. The usefulness of the new technique is illustrated with an application to a cointegrated consumption-income regression model, yielding a straightforward test of the permanent income hypothesis.

Keywords: Permanent Income, Parameter Variation, Frequency Domain

1. INTRODUCTION

The notion that relationships between macroeconomic time series vary across frequencies has a distinguished history. Early authors expressed and analyzed this variation in the time domain, distinguishing between the short period versus the long period (Marshall, 1920), between the short run and the long run (Keynes, 1936), or between transitory income and permanent income (Friedman, 1957). Later workers explicitly utilized the frequency domain as the field of discourse, e.g., Engle (1974, 1978), Lucas (1980), Geweke (1982, 1986), Mills (1982), Summers (1983, 1986), Cochrane (1989), Phillips (1991), Thoma (1992, 1994), Corbae et al. (1994), and Lee (1994).

In view of all of this activity, it is not surprising that several frameworks already exist for analyzing the frequency dependence of time-series relationships.

Address correspondence to: Richard Ashley, Department of Economics (0316), Virginia Polytechnic Institute and State University, Blacksburg, VA 24061, USA; e-mail: ashleyr@vt.edu.

 1365-1005/99 $9.50

Geweke (1982), for example, provides a measure of how the strength of a relationship varies with frequency in a linear model. However, Tan and Ashley (1997) argue on fundamental grounds that no linear model can capture frequency-dependent relationships of the sort at issue here; they find that Geweke's measure actually only quantifies the degree to which an innovation in one series yields low- or high-frequency variation in the other series. The other approaches cited above generally trace their roots to either the band spectral regression model introduced by Engle (1974), and or to the system spectral regression model of Phillips (1991).

These spectral regression methods are asymptotically valid, conditional on an a priori choice of which frequency bands to consider, but their actual utility is rather limited. As ordinarily formulated, the band spectral regression approach is a single-equation technique and only allows a test of whether a regression coefficient is different over a single band of frequencies. The Phillips (1991) approach explicitly considers a multidimensional cointegrated system, but it is so sophisticated that it is difficult to imagine it coming into widespread use, particularly where adaptations to specific data problems might require modifications to the technique or where the credibility of the results hinges on the reader fully understanding the methodology. In any case, the detection of parameter variation across frequencies implies that models not allowing for this variation are misspecified, but this realization is of little value unless the parameter variation can be incorporated in ordinary modeling efforts.

We propose a new procedure for detecting and modeling regression-parameter variation across frequency bands. This procedure has a number of advantages over currently available techniques:

1. It does not require large samples, unlike Phillips' approach.
2. It does not require that the frequency bands be specified a priori.
3. It does not require specialized estimation software.
4. It is implemented in the time domain and is fully compatible with whatever econometric techniques and/or software are already in use.
5. It is so simple that it is potentially accessible to a variety of applied economists.

The band-spectrum regression approach of Engle (1974) is reviewed briefly in Section 2, leading to a description of the proposed technique in Section 3. This new technique is used to detect and model frequency dependence in the income coefficient of a simple dynamic consumption-income relation in Section 4. Section 5 provides concluding remarks.

2. BAND-SPECTRUM REGRESSION MODEL OF ENGLE (1974)

Suppose that the k explanatory variables in a multiple regression model have been numbered in such a way that the issue at hand is whether, and in what way, the relationship between the dependent variable Y_t and the kth explanatory variable

https://doi.org/10.1017/S1365100599010032 Published online by Cambridge University Press

X_{tk} depends on frequency. For simplicity of exposition, consider the ordinary multiple regression model,

$$\boldsymbol{Y} = X\boldsymbol{\beta} + \boldsymbol{\epsilon}, \qquad \boldsymbol{\epsilon} \sim N[0, \sigma^2 I], \qquad r(X) = k. \tag{1}$$

Clearly, β_k is the coefficient of interest. The approach proposed below applies equally well to more complex situations in which X is stochastic, where ϵ is non-Gaussian, where $\text{var}(\epsilon) = \sigma^2\Omega \neq \sigma^2 I$, where this is the mth equation in a system of simultaneous equations, and so forth; this simple model is considered here solely for expositional clarity.

In Engle's approach, equation (1) is premultiplied by the complex-valued orthogonal matrix W whose (j, t)th element is

$$w_{jt} = \exp[i2\pi(j-1)(t-1)/T] \tag{2}$$

to obtain

$$\begin{aligned} W\boldsymbol{Y} &= WX\boldsymbol{\beta} + W\boldsymbol{\epsilon}, \qquad W\boldsymbol{\epsilon} \sim N[0, \sigma^2 I], \\ \boldsymbol{Y}^* &= X^*\boldsymbol{\beta} + \boldsymbol{\epsilon}^*, \qquad \boldsymbol{\epsilon}^* \sim N[0, \sigma^2 I], \end{aligned} \tag{3}$$

where $\boldsymbol{Y}^* = W\boldsymbol{Y}$, etc. By construction, the jth element of $\boldsymbol{Y}^*$, of $\boldsymbol{\epsilon}^*$, and of each column of X^* is the finite Fourier transform of the analogous column vector of time-domain observations, evaluated at frequency $2\pi(j-1)/T$. Note that ϵ^* is not identical to ϵ, but it has the same distribution because W is an orthogonal matrix; in contrast, the coefficient vector $\boldsymbol{\beta}$ is unaffected by the transformation. Most important, each of the T elements of $\boldsymbol{Y}^*$ and of the columns of X^* is a weighted sum of the data from all T time periods—by construction, these T observations now correspond to the frequencies $0, 2\pi(1/T), 2\pi(2/T), \ldots, 2\pi([T-1]/T)$.

Least-squares estimation of equation (3) requires specialized software because Y^* and X^* are complex-valued and, in any case, yields the same results as estimating the time-domain model, equation (1). Equation (3), however, makes it possible to test whether the $Y_t - X_{tk}$ relationship is different over one band of frequencies versus the rest. Engle (1974) provides a Chow-type test of such a hypothesis that can be performed using ordinary (real-valued) regression software.

Engle's approach has three serious drawbacks, however. The first is only expositional, but illustrates the kind of arcana that limits the acceptance of most spectral techniques: The test requires that frequency $2\pi([T-j]/T)$ be included in the frequency band to be tested whenever frequency $2\pi j/T$ is included. Thus, to test whether β_k for observations 2 and 3 [corresponding to the two lowest nontrivial frequencies, $2\pi(1/T)$ and $2\pi(2/T)$] is the same as it is for the other frequencies, the band of frequencies examined would have to include not only $2\pi(1/T)$ and $2\pi(2/T)$, but also the two *highest* frequencies: $2\pi([T-2]/T)$ and $2\pi([T-1]/T)$. This is correct, but the explanation involves enough spectral theory that the uninitiated observer is likely to simply lose interest in the results at that point.

Second, Engle's test only addresses the constancy of β_k across partitions of the sample into two frequency bands. Because the technique does not provide any convenient means for visualizing the variation in β_k across frequencies, this partitioning is somewhat arbitrary, which limits the usefulness of the resulting test.[1]

Finally, the band-spectrum regression test is closed-ended—it can indicate that a regression parameter varies across frequencies and even, through a sequence of tests, say something about the manner in which it varies, but it says little about how to improve the specification of the original regression model.

3. BLOCKWISE TIME-DOMAIN SPECTRAL REGRESSION

3.1. Overview

The procedure proposed here consists of three steps:

1. Transform the original time-domain regression model, equation (1), into a *real-valued* frequency-domain regression model using a transformation matrix based on the finite sine and cosine transformations. This amounts to premultiplying equation (1) by a real-valued orthogonal matrix, A, defined below.
2. Allow for variation in β_k across m frequency bands—i.e., across m groups of observations in the frequency-domain regression model—using dummy variables.
3. Back-transform the resulting regression model to the time domain to estimate $\beta_1 \cdots \beta_{k-1}$ and the m dummy variable coefficients. This back-transformation merely involves premultiplying the regression equation, augmented by the dummy variable terms, by the transpose of A.

The appropriate number of frequency bands, m, is determined using standard model selection tools—e.g., minimizing the Schwarz criterion. Parameter constancy across the frequency bands corresponds to the null hypothesis that all m dummy variable coefficients are equal; this can be tested readily using standard methods.

Actually, the back-tranformation of step 3 is not always necessary; the estimation and parameter stability testing can be done just as easily in the frequency-domain regression model. However, if β_k is found to vary significantly across the m frequency bands, then it is convenient in subsequent modeling to replace X_{tk} with the m back-transformed dummy variables; indeed, that is precisely what is done in the regression model of step 3.

3.2. Real-Valued Frequency Domain Regression Model

Engle (1974) premultiplies the original time-domain regression model, equation 1, by a complex-valued matrix W whose T rows are the finite Fourier transform coefficients for frequencies $2\pi(j-1)/T$, $j = 1 \ldots T$. Here, equation 1 instead is premultiplied by the real-valued transformation matrix A, first suggested by

Harvey (1978), with (j, t)th element:

$$a_{j,t} = \begin{cases} \left(\frac{1}{T}\right)^{\frac{1}{2}} & \text{for } j = 1; \\ \left(\frac{2}{T}\right)^{\frac{1}{2}} \cos\left[\frac{\pi j(t-1)}{T}\right] & \text{for } j = 2, 4, 6, \ldots, \\ & \quad (T-2) \text{ or } (T-1); \\ \left(\frac{2}{T}\right)^{\frac{1}{2}} \sin\left[\frac{\pi (j-1)(t-1)}{T}\right] & \text{for } j = 3, 5, 7, \ldots, \\ & \quad (T-1) \text{ or } T; \\ \left(\frac{1}{T}\right)^{\frac{1}{2}} (-1)^{t+1} & \text{for } j = T \quad \text{and} \quad T \text{ is even}, \\ & \quad t = 1, \ldots, T. \end{cases} \tag{4}$$

to obtain the real-valued frequency-domain regression equation,

$$\begin{aligned} A\boldsymbol{Y} &= AX\boldsymbol{\beta} + A\boldsymbol{\epsilon}, \qquad A\boldsymbol{\epsilon} \sim N[0, \sigma^2 I], \\ \boldsymbol{Y}^{**} &= X^{**}\boldsymbol{\beta} + \boldsymbol{\epsilon}^{**}, \qquad \boldsymbol{\epsilon}^{**} \sim N[0, \sigma^2 I], \end{aligned} \tag{5}$$

where $\boldsymbol{Y}^{**} = A\boldsymbol{Y}$, etc. A is known to be orthogonal,[2] and so, $\boldsymbol{\epsilon}^{**}$ is not identical to $\boldsymbol{\epsilon}$ of equation (1), but it has the same distribution; $\boldsymbol{\beta}$ itself is unchanged. Equation (5) can be estimated directly with ordinary regression software because the elements of $\boldsymbol{Y}^{**}$ and X^{**} are real-valued.

The relationship between the rows of A [the "observations" in equation (5)] and the corresponding frequencies is summarized in Table 1. The second and third

TABLE 1. Frequency corresponding to each observation in equation (5)

Observation No.	Frequency
1	0
2	$\pi\{1/T\}$
3	$\pi\{1/T\}$
4	$\pi\{2/T\}$
5	$\pi\{2/T\}$
6	$\pi\{3/T\}$
7	$\pi\{3/T\}$
…	
$T-2$	$\pi\{(T/2)-1\}/T$
$T-1$	$\pi\{(T/2)-1\}/T$
T	π

https://doi.org/10.1017/S1365100599010032 Published online by Cambridge University Press

rows of A are the coefficients in the finite sine and cosine transforms at the lowest nonzero frequency, $\pi(1/T)$; the fourth and fifth rows are the coefficients for the transforms at the next frequency, $\pi(2/T)$, etc. The first row of A corresponds to zero frequency; if T is even, there is also a single row corresponding to the highest frequency, π.

Some authors express these frequencies divided by a factor of 2π. This makes no difference; the essential feature is that the low and high frequencies associated with the rows of A correspond in a straightforward way to what is meant by phrases such as "the relationship at low frequencies" or "the high-frequency component."

Consider how the various rows of A operate on a column vector of observations on a time series. The first row averages all T observations together with equal weights; any fluctuation in the time series that largely averages out to zero over the entire period will yield a small component at frequency zero. The third row averages all T observations also, but its weights make one complete sine oscillation during the course of the sample. Thus, any fluctuation in the time series that takes place sufficiently quickly as to largely average out to zero during either the first half or the second half of the sample will contribute little to the third observation in the frequency domain. The second row is similar, but here the weights make one complete cosine oscillation during the sample, so that any fluctuation in the time series that takes place sufficiently quickly as to largely average out to zero during either the first quarter, the middle half, or the fourth quarter of the sample will contribute little to the second observation in the frequency domain. Similarly, the weights for rows 4 and 5 make two complete cosine and sine oscillations, respectively, during the course of the sample, and so, fluctuations in the time series must complete themselves (average out to zero) something like twice as quickly in order to contribute little to observations 4 and 5 in the frequency domain. Clearly, the transformation corresponding to the low-frequency rows of the A matrix is ignoring the quickly fluctuating parts of the data vector and thereby extracting the most smoothly and slowly varying components. Finally, suppose that T is even and consider the highest-frequency row of A. This row simply averages $T/2$ changes in the data; clearly, it is ignoring any slowly varying components of the data vector and extracting the most quickly varying component.[3]

3.3. Testing for Parameter Stability in the Frequency Domain

Because the real-valued spectral regression model, equation (5), is an ordinary regression equation—the only difference being that its T observations correspond to the frequencies given in Table 1 rather than to time periods—the stability of the $\boldsymbol{Y}_t - X_{tk}$ relationship across frequencies can be tested using established methods for assessing the stability of regression coefficients. A number of such methods exist, including those of Chow (1960), Farley et al., (1975), Brown et al., (1975), Garbade (1977), LaMotte and McWhorter (1978), Ashley (1984), and Watson and Engle (1985).

The stabilogram test given by Ashley (1984) is used here. In the present context this test partitions the sample of T observations in the frequency domain into m more-or-less equal subsamples or frequency bands. These bands are generally somewhat unequal in length, partly because m does not always divide T evenly and partly because there are two observations at each nonzero frequency if T is odd. (If T is even, there is also a single observation at frequency π.)

Dummy variables $(D_j^1 \cdots D_j^m)$ are created, one for each of these subsamples, such that $D_j^s = X_{jk}^{**}$ if observation j is in frequency band s and $D_j^s = 0$ otherwise. The regression model,

$$Y^{**} = X_{\{k\}}^{**}\beta_{\{k\}} + D\gamma + \nu^*, \qquad \nu^* \sim N[0, \sigma^2 I] \tag{6}$$

then is estimated, where $X_{\{k\}}^{**}$ is X^{**} with its kth column removed, $\boldsymbol{\beta}_{\{k\}}$ is the β vector with its kth component removed, and D is the $T \times m$ matrix $[\boldsymbol{D}^1 \cdots \boldsymbol{D}^m]$. The number of frequency bands, m, is chosen by minimizing a corrected goodness-of-fit measure, such as the Schwarz criterion.[4] Finally, the null hypothesis that all m components of γ are the same—which corresponds to the null hypothesis that β_k, the coefficient quantifying the strength of the $Y_t - X_{tk}$ relationship, is constant across frequencies—is tested using standard methods.[5]

The stabilogram test is a natural choice in the present context for several reasons. It is straightforward to implement using standard regression software, yet simulations reported by Ashley (1984) show that its power is similar to that of the alternative tests in samples of modest size. Moreover, the stabilogram itself, a plot of the estimated 95% confidence intervals for $\hat{\gamma}_1 \cdots \hat{\gamma}_m$, provides a convenient way to visualize the variation in the strength and/or sign of the $Y_t - X_{tk}$ relationship across frequencies. Finally, the stabilogram dummy variables, $\boldsymbol{D}^1 \cdots \boldsymbol{D}^m$, can be back-transformed easily to yield a time-domain version of the test.

3.4. Testing for Frequency Dependence in the Time Domain

Because the transformation matrix A is orthogonal, its transpose is its inverse. Consequently, premultiplying equation (6) by A^t yields an equivalent time-domain regression model,

$$\begin{aligned} A^t\boldsymbol{Y}^{**} &= A^t X_{\{k\}}^{**}\boldsymbol{\beta}_{\{k\}} + A^t D\gamma + A^t\boldsymbol{\nu}^*, \qquad A^t\nu^* \sim N[0, \sigma^2 I], \\ \boldsymbol{Y} &= X_{\{k\}}\boldsymbol{\beta}_{\{k\}} + D^*\gamma + \boldsymbol{\nu}, \qquad \nu \sim N[0, \sigma^2 I], \end{aligned} \tag{7}$$

where $\boldsymbol{Y}$ is the (time domain) vector of dependent variable data in equation (1), and $X_{\{k\}}$ is a $T \times (k-1)$ matrix consisting of the first $k-1$ columns of X, the matrix of observations on the explanatory variables in equation (1).

The m columns of D^* are essentially filtered versions of the kth column of X. Figure 1 illustrates the relationship between the D and D^* matrices for $m = 3$ where X_{tk} is PPI_t, the deviation of the growth rate in the U.S. Producer Price Index from its sample mean, over a sample of 72 observations from June 1982 to

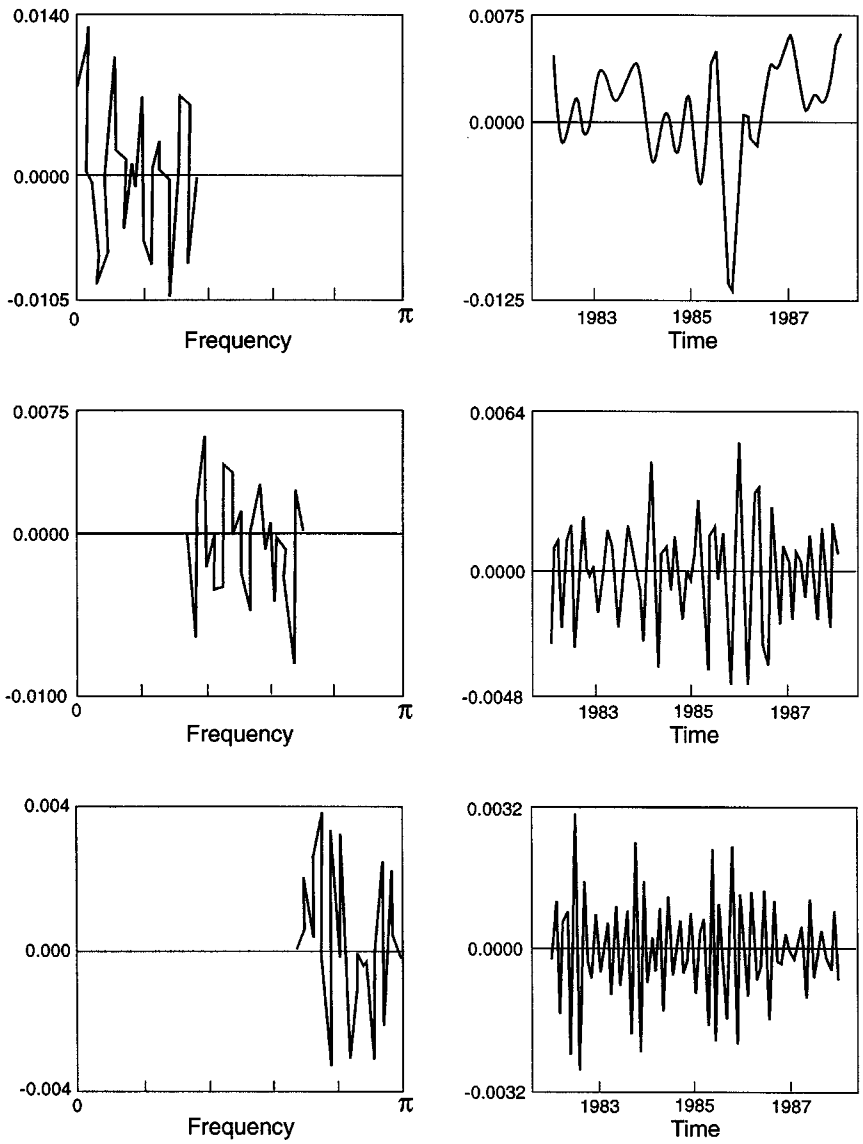

FIGURE 1. Frequency-domain dummy variables for PPI_t and time-domain equivalents.

May 1988. The three diagrams on the left side of Figure 1 plot the 72 elements of each of the three columns of the D matrix against frequency; the three diagrams on the right plot the 72 elements of each of the analogous columns of the D^* matrix against time. The sum of these three columns of D^* precisely reproduces the sample data on PPI_t.[6]

By partitioning $X_{tk} = \text{PPI}_t$ into these three components and allowing equation (7) to fit potentially different coefficients to each, this procedure allows for the possibility that fluctuations in the low-frequency component of PPI_t—i.e., changes in its smooth, relatively predictable, local trend—might affect Y_t differently than the erratic, relatively unpredictable, high-frequency component of PPI_t. As with equation (3), testing for frequency dependence in the $Y_t - \text{PPI}_t$ relationship involves nothing more than testing the null hypothesis that $\gamma_1 = \gamma_2 = \gamma_3$.

The frequency-domain test [using equation (6)] and the time-domain test [using equation (7)] are both straightforward parameter-restriction hypothesis tests on an ordinary regression equation; because they are equivalent, they give identical results. However, the time-domain approach has a major advantage in that the disaggregation of X_{tk} across frequencies provided by the m columns of D^* can be used in *any* time-domain econometric modeling effort to test and/or allow for frequency dependence in relationships involving X_{tk}, using whatever econometric technique (cointegration, 2SLS, probit, etc.) is appropriate for that context.

4. ILLUSTRATIVE APPLICATION: A TEST OF THE PERMANENT INCOME HYPOTHESIS

Numerous macroeconomic relationships are thought to vary with frequency: money-income, money-inflation, output-inflation, inflation-interest rate. Yet, surely the canonical example in which theory suggests that such variation is important is the consumption-income relation. Therefore, the technique proposed above is applied here to test for and model frequency dependence in the coefficient on income in a dynamic model for U.S. consumption expenditures.

Monthly data on real consumption expenditures and real personal income (GMCQ_t and GMPY82_t, respectively) are obtained from the CITIBASE Data Bank for the period February 1959 to October 1991—the full interval over which GMPY82_t is available. The sample behavior of each of these time series clearly is dominated by a unit root, and so, the analysis is done using the logarithmic growth rate of each series, denoted c_t and y_t, below.

Plotted, c_t and y_t both appear to be covariance stationary over the entire time period. Consequently, the sample initially was partitioned into two 196-month subsamples, to provide for model cross validation and/or postsample forecasting. However, the c_t–y_t relationship becomes unstable during the second subsample, and so, results are reported here only for the first period, which runs from March 1959 through June 1975.[7]

There is ample reason to expect GMCQ_t and GMPY82_t to be cointegrated; consequently, the Engle and Granger (1987) two-step estimation method is used, yielding the estimated cointegrating equation

$$\log(\text{GMCQ}_t) = 0.5303 + 0.9221 \log(\text{GMPY82}_t) + \hat{\nu}_t \qquad \textbf{(8)}$$

https://doi.org/10.1017/S1365100599010032 Published online by Cambridge University Press

and the estimated error-correction equation

$$c_t = \underset{(4.86)}{2.591} - \underset{(3.13)}{1.892}\hat{\nu}_{t-1} + \underset{(3.48)}{0.274}y_t + \hat{\epsilon}_t, \qquad \overline{R}^2 = 0.065, \quad DW = 2.03. \tag{9}$$

The fitting errors from the cointegrating equation ($\hat{\nu}_t$) are correlated serially, but a time plot indicates that they are $I(0)$ and covariance stationary. The significance of the estimated coefficient on $\hat{\nu}_{t-1}$ in the error-correction equation confirms that cointegration is present. Additional lags in c_t and y_t entered with insignificant coefficients.[8]

Before examining the coefficient on y_t for instability across frequencies, it was tested first for instability of the ordinary kind, across the sample observations. It appears to be quite stable over this time period: Partitioning the sample period into six approximately equal subperiods and estimating a separate coefficient on y_t for each subperiod yields the stabilogram (a plot of the six dummy variable coefficient estimates and their associated 95% confidence intervals versus time) given in Figure 2. Although the coefficient is a bit larger and more precisely estimated in the latter part of the sample period, the null hypothesis that all six coefficients are identical cannot be rejected. Analogous tests on stabilogram regressions with 2 to 10 subperiods yield similar results: The coefficients differ from one another only at significance levels ranging from 12% to 54%.

The coefficient on y_t then was examined for stability across different frequencies using the procedure proposed above, based on equation (7). For $m = 2$, the D and D^* matrices are both 196×2. The first 99 elements of the first column of D are

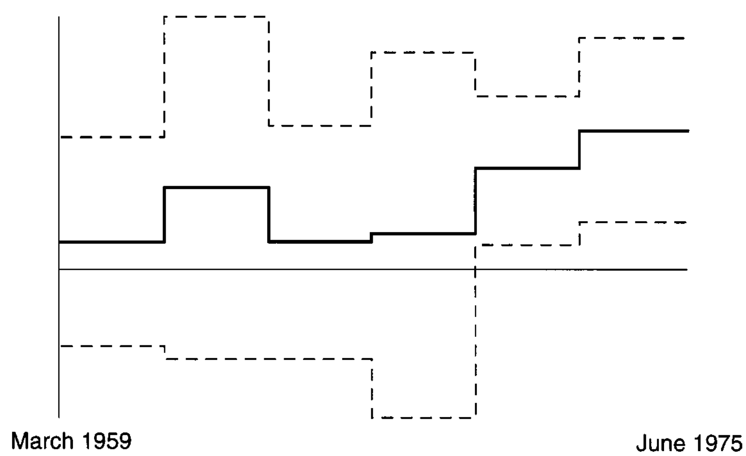

FIGURE 2. Time-domain stabilogram for y_t coefficient in equation (9). [Using the procedure suggested by Ashley (1984) the coefficient is tested for stability over time by partitioning the sample into six subperiods and estimating the coefficient separately over each using dummy variables. Here, the resulting six-parameter estimates and associated 95% confidence intervals are plotted against time. Other partitionings give similar results.]

the first 99 elements of $A\boldsymbol{y}$; the remaining elements are zero. The first 99 elements of the second column of D are zero; the remaining 97 elements are the last 97 elements of $A\boldsymbol{y}$. So, the D^* matrix is

$$D^* = A^t D = A^t \begin{bmatrix} \boldsymbol{e}_1 & 0 \\ \boldsymbol{e}_{99} & 0 \\ 0 & \boldsymbol{e}_{100} \\ 0 & \boldsymbol{e}_{196} \end{bmatrix} A\boldsymbol{y},$$

where $\boldsymbol{e}_i$ denotes the ith unit row vector of whatever length the context requires. The two bands are uneven in length because the first band contains the single observation corresponding to frequency zero and the 98 elements of $A\boldsymbol{y}$ corresponding to the first 49 nonzero frequencies, whereas the last band contains the 96 elements of $A\boldsymbol{y}$ corresponding to the next 48 nonzero frequencies plus the last element of $A\boldsymbol{y}$, which corresponds to frequency π.

Similarly, for $m = 3$, the D and D^* matrices have three columns. The first column of D has, as its first 65 elements, the corresponding elements of $A\boldsymbol{y}$; these consist of the frequency zero observation plus the observations for the first 32 nonzero frequencies. The second column of D has the next 66 elements equal to the corresponding elements of $A\boldsymbol{y}$; these are the observations for the next 33 frequencies. The third column of D contains the remaining 65 elements of $A\boldsymbol{y}$; these consist of 64 observations corresponding to the next 32 frequencies plus the single observation at frequency π. Thus, for $m = 3$, the D^* matrix is

$$D^* = A^t D = A^t D = A^t \begin{bmatrix} \boldsymbol{e}_1 & 0 & 0 \\ \boldsymbol{e}_{65} & 0 & 0 \\ 0 & \boldsymbol{e}_{66} & 0 \\ 0 & \boldsymbol{e}_{131} & 0 \\ 0 & 0 & \boldsymbol{e}_{132} \\ 0 & 0 & \boldsymbol{e}_{196} \end{bmatrix} A\boldsymbol{y},$$

and so forth.[9]

Equation (7) was estimated for $m = 2, 3, \ldots 10$; the optimal value for m is 4 using the Schwarz criterion. The resulting stabilogram for this model is plotted in Figure 3; the estimated model is

$$\begin{aligned} c_t = {} & \underset{(4.20)}{2.590} - \underset{(4.02)}{2.368}\, \hat{\nu}_{t-1} + \underset{(5.72)}{0.758}\, D_t^{(*1)} + \underset{(0.13)}{0.019}\, D_t^{(*2)} \\ & + \underset{(0.71)}{0.094}\, D_t^{(*3)} + \underset{(0.98)}{0.159}\, D_t^{(*4)} + \hat{\eta}_t, \end{aligned} \tag{10}$$

where $D_t^{(*j)}$ is the period-t observation from the jth column of D^*. The $F(3, 190)$ statistic for testing the null hypothesis that all four $D_t^{(*j)}$ coefficients are equal is 6.760, and so, this null hypothesis can be rejected at the 0.02% level. Evidently, it is either primarily or entirely the low-frequency component of y_t that affects c_t.

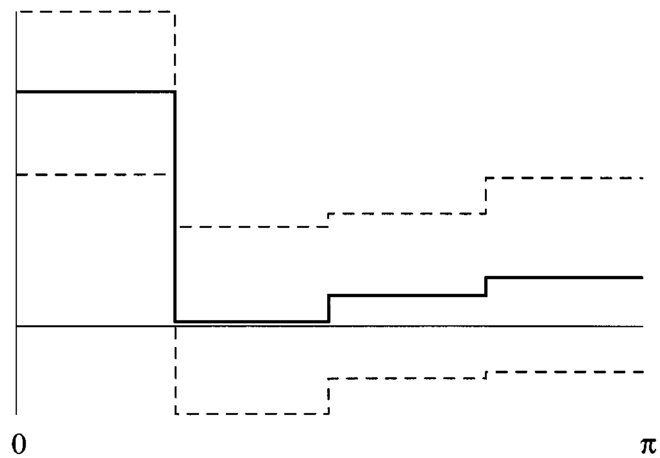

Figure 3. Four-band frequency-domain stabilogram for y_t coefficient [equation (10)]. [Here, the parameter estimates and associated 95% confidence intervals for the coefficients on $D^{*1} \cdots D^{*4}$ in equation (10) are plotted against frequency.]

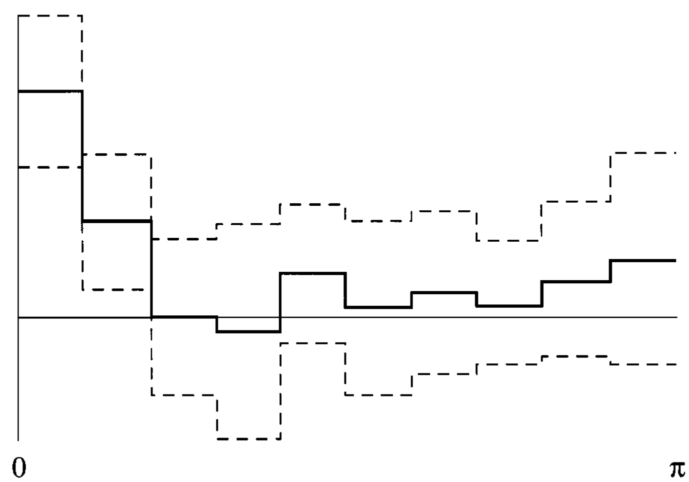

Figure 4. Ten-band frequency-domain stabilogram for y_t coefficient. [Here, the parameter estimates and associated 95% confidence intervals for coefficients on $D^{*1} \cdots D^{*10}$ are plotted against frequency in a model analogous to equation (10) only with 10 bands.]

Because the pattern of variation of the income coefficient with frequency is of interest in its own right, a more detailed stabilogram based on 10 frequency bands is plotted in Figure 4. This diagram indicates that the income coefficient is significantly positive only for the first two bands. The first band corresponds to frequencies zero through $\pi(9/196)$; the second band includes frequencies up

https://doi.org/10.1017/S1365100599010032 Published online by Cambridge University Press

through $\pi(20/196)$. The most quickly varying component in the first band completes nine full cycles during a sample of 196 months, for a period of 22 months; any fluctuation in y_t that substantially reverses itself within 11 months will have little impact on the dummy variable corresponding to this first band. Similarly, the most quickly varying component in the second band completes 19 full cycles in 196 months, for a period of 10 months; thus, a fluctuation in y_t that takes longer than 5 months but less than 11 months to complete itself will impact the dummy variable corresponding to this second band. Essentially, households seem to ignore income fluctuations that they expect to last for notably less than 6 months, give some weight to fluctuations that they expect to last for 6 to 12 months, and they base changes in their consumption spending decisions primarily on income fluctuations that they expect to last for a year or more.

5. CONCLUSIONS

It is shown in Section 3 that partitioning (filtering) a series y_t into m components, $D_t^{(*1)} \cdots D_t^{(*m)}$, corresponding to m frequency bands is actually quite straightforward. Once this partitioning is done, testing and allowing for frequency dependence in the relationship involves little more than replacing y_t with these m alternative regressors in the estimation equation. Similarly, the choice of how many frequency bands to consider does not require experience in spectral analysis—it is just another modeling decision of the usual form: choosing one set of regressors over another.

Frequency dependence is thought to characterize many of the most important relationships in macroeconomics, but detecting and modeling this aspect of these relationships has heretofore been considered a challenging and rather specialized endeavor. That is no longer the case. The results reported in Section 4 on the relationship between aggregate personal income and aggregate consumption expenditures illustrate how easy and effective it is to use the proposed technique to detect and model frequency variation in regression coefficients.

NOTES

1. The plot of MPC vs. frequency given by Engle (1974, Fig. 2) is not from the spectral regression at all—it is the gain of the estimated cross spectrum between the two series. This approach to visualizing the frequency dependence of β_k is only feasible for bivariate regression models and even then would be confounded by feedback relationships.

2. See Tan (1995). The rows of A are just linear combinations of the rows of W, reflecting the fact that the regression models of equations (1), (3), and (5) are all essentially equivalent.

3. In contrast, a low-frequency band in Engle's approach will contain observations at frequencies close to zero and close to 2π; and in Phillips' approach, the high-frequency band includes frequencies close to π as well as frequencies close to $-\pi$.

4. The Schwarz criterion [$\ln(\text{SSE}/T) + (k+m-1)\ln T/T$ here] is discussed by Judge et al. (1985, p. 245).

5. For equation (6), (RSS − URSS)($T - k - m + 1$)/[URSS($m - 1$)] is distributed $F(m-1, T-k-m+1)$ under this null hypothesis, where URSS is the sum of squared residuals and RSS is the sum of squared residuals from estimating equation (1) or (5).

6. The columns of D^* are reminiscent of the filtered series used by Cochrane (1989) and Lee (1994). Where band-pass filtering per se is the purpose, their filtering methods may be preferable. The frequency bands must be chosen arbitrarily in their approaches, however, and their filtered components do not aggregate back up to the original time-domain data. Also, their approaches consider the $Y_t - X_{tk}$ relationship at each frequency band in isolation, which limits the contact between their results and subsequent time-domain modeling.

7. Also, the cointegrating relationship between GMCQ$_t$ and GMPY82$_t$ is notably different between the two subperiods.

8. Lagged personal income, y_{t-1}, remains insignificant when the contemporaneous term is eliminated, so it is not possible to analyze this relation in reduced form. The estimated coefficient on y_t is doubtless biased because of the joint endogeneity of c_t and y_t; elimination of this bias via instrumental variables is not undertaken here, however, because it would unduly complicate the example.

9. Software that computes D and D^* for given T, m, and y is available from the authors.

REFERENCES

Ashley, R.A. (1984) A simple test for regression parameter instability. *Economic Inquiry* 22, 253–267.

Brown, R.L. Durbin, J. & Evans, J.M. (1975) Techniques for testing the constancy of regression relationships over time, with comments. *Journal of the Royal Statistical Society (Series B)* 37, 149–192.

Chow, G. (1960) Tests of equality between subsets of coefficients in two linear regressions. *Econometrica* 28, 591–605.

Cochrane, J.H. (1989) The return of the liquidity effect: a study of the short-run relation between money growth and interest rates. *Journal of Business and Economic Statistics* 7, 75–83.

Corbae, D., Ouliaris, S. & Phillips, P.C.B. (1994) A reexamination of the consumption function using frequency domain regressions. *Empirical Economics* 19, 595–609.

Engle, R.F. (1974) Band spectrum regression. *International Economic Review* 15, 1–11.

Engle, R.F. (1978) Testing price equations for stability across spectral frequency bands. *Econometrica* 46, 869–881.

Engle, R.F. & Granger, C.W.J (1987) Co-integration and error correction: Representation, estimation and testing. *Econometrica* 55, 251–279.

Farley, J.U., Hinich, M. & McGuire, T.W. (1975) Some comparisons of tests for a shift in the slopes of a multivariate linear time series model. *Journal of Econometrics* 3, 297–318.

Friedman, M. (1957) *A Theory of the Consumption Function*, Princeton: Princeton University Press.

Garbade, K. (1977) Two methods for examining the stability of regression coefficients. *Journal of the American Statistical Association* 72, 54–63.

Geweke, J. (1982) Measurement of linear dependence and feedback between multiple time series. *Journal of the American Statistical Association* 77, 304–314.

Geweke, J. (1986) The superneutrality of money in the United States: An interpretation of the evidence. *Econometrica* 54, 1–21.

Harvey, A.C. (1978) Linear regression in the frequency domain. *International Economic Review* 19, 507–512.

Judge, G.G., Griffith, W.E., Hill, C.R., Lütkepohl, H. & Lee, T. (1985) *The Theory and Practice of Econometrics*, 2nd ed. New York: Wiley.

Keynes, J.M. (1936) *The General Theory of Employment, Interest and Money*. London: Macmillan.

LaMotte, L.R. & McWhorter, A. (1978) An exact test for the presence of random-walk coefficients in a linear regression model. *Journal of the American Statistical Society* 73, 816–820.

Lee, K. (1994) Causality of Nominal Money Growth at Various Frequency Bands: An Empirical Analysis Using Time Series Domain Chebyshev Filters. Working paper, University of Missouri.

Lucas, R.E. (1980) Two illustrations of the quantity theory of money. *American Economic Review* 70, 1005–1014.
Marshall, A. (1920) *Principles of Economics*, 8th ed., pp. 363–380. London: Macmillan.
Mills, T.C. (1982) Signal extraction and two illustrations of the quantity theory of money. *American Economic Review* 72, 1162–1168.
Phillips, P.C.B. (1991) Spectral regression for cointegrated time series. In W. Barnett, J. Powell, & G. Tauchen (eds.), *Nonparametric and Semiparametric Methods in Economics and Statistics*. Cambridge, England: Cambridge University Press.
Summers, L.H. (1983) The non-adjustment of interest rates: A study of the Fisher effect. In J. Tobin (ed.), *Macroeconomics, Prices and Quantities: Essays in Memory of Arthur M. Okun*. Washington, DC: Brookings Institution.
Summers, L.H. (1986) Estimating the long-run relationship between interest rates and inflation (a response to McCallum). *Journal of Monetary Economics* 18, 7–86.
Tan, H.B. (1995) Disentangling Low-Frequency Versus High-Frequency Economic Relationships Via Regression Parameter Stability Tests. Ph.D. Dissertation, Virginia Polytechnic Institute and State University.
Tan, H.B. & Ashley, R.A. (1998) On the inherent nonlinearity of frequency-dependent time series relationships. In P. Rothman (ed.), *Nonlinear Time Series Analysis of Macroeconomic and Financial Data*. Norwell: Kluwer Academic Press.
Thoma, M.A. (1992) The effects of inside and outside money on industrial product across spectral frequency bands. *The Review of Economics and Statistics* 74, 737–741.
Thoma, M.A. (1994) The effects of money growth on inflation rates across spectral frequency bands. *Journal of Money, Credit, and Banking* 26, 218–230.
Watson, M.W. & Engle, R.F. (1985) Testing for regression coefficient stability with a stationary AR(1) alternative. *The Review of Economics and Statistics* 67, 341–346.