Warsaw School of Economics
Institute of Econometrics
Department of Applied Econometrics

Department of Applied Econometrics Working Papers

Working Paper No. 2-10

Oversampling of stochastic processes

D.S.G. Pollock
University of Leicester

OVERSAMPLING OF STOCHASTIC PROCESSES

By D.S.G. POLLOCK

University of Leicester

Email: stephen_pollock@sigmapi.u-net.com

Discrete-time ARMA processes can be placed in a one-to-one correspondence with a set of continuous-time processes that are bounded in frequency by the Nyquist value of π radians per sample period. It is well known that, if data are sampled from a continuous process of which the maximum frequency exceeds the Nyquist value, then there will be a problem of aliasing. However, if the sampling is too rapid, then other problems will arise that will cause the ARMA estimates to be severely biased. The paper reveals the nature of these problems and it shows how they may be overcome.

Keywords: Stochastic Differential Equations, Band-Limited Stochastic Processes, Oversampling.

1. Introduction

Econometricians are liable to regard discrete-time autoregressive moving-average (ARMA) models as approximations to continuous-time stochastic differential equations. (For a recent exposition of this point of view, see Sims 2010.) In common with other statisticians, they tend to assume that the differential equations are powered by the increments of a Wiener process. A Wiener process is formed from the accumulation of a continuous stream of infinitesimal impulses. This generates a trajectory that is everywhere continuous but nowhere differentiable.

The impulses that underlie a Wiener process correspond to Dirac delta functions. These generalised functions are zero-valued everywhere except at a single point. The Fourier transform of a Dirac function indicates that it comprises sinusoidal elements of uniform but infinitesimal amplitudes that cover an unbounded range of frequencies.

A Wiener process is a non-anticipatory process that has the martingale property, whereby its future evolution cannot be predicted—and the best estimate of its future values is its present value. This property is exploited extensively in financial mathematics, which is one of the reasons for the prevalence of Wiener processes.

In many cases of financial modelling, the concomitant property of the unbounded frequencies, which must seem unreasonable, can be overlooked. However, in the context of macroeconomics, where the trajectories of the variables are dominated by trends and by motions of low frequency, it seems inappropriate to propose that the forcing functions should be Wiener processes.

The notion of unbounded frequency is also at odds with the modern technology of digital signal processing and communications engineering. Here, the assumption is made that signals that convey information are bounded in frequency. In that case, the Shannon–Nyquist sampling theorem indicates that, given a sufficient rate of sampling, the information of the signal can be conveyed by a sequence of values sampled at regular intervals.

A question arises of whether or not the Shannon–Nyquist theorem can be applied profitably to econometric time series. If a discretely sampled econometric data sequence were to convey all of the information within the underlying continuous process, then a model formulated in discrete time might provide an appropriate and an accurate description of the underlying process. In order to capture the information, the rate of sampling must be such that at least two observations occur in the time that it takes the sinusoidal element of highest frequency within the process to complete its cycle.

In principle, is should be straightforward to recognise cases where the rate sampling exceeds the necessary minimum value. The periodogram of the data, which shows the squares of the estimated amplitudes of the sinusoidal elements, should reveal values that are close to zero for the frequencies that exceed the maximum value within the underlying process. The squared amplitudes denote the power or the variance of the sinusoidal elements.

In practice, matters are often complicated by the presence of trends within the data. The effect of a trend will be to create a disjunction in the periodic extension of the data at the points where the end of one replication of the sample joins the beginning of the next replication. Such a disjunction will create a slew of spectral power that extends across the entire frequency range. The result is what Granger (1966) described as the typical spectral shape of an economic data sequence.

The problem can be avoided only by a careful detrending of the data that ensures that there are smooth transitions where the beginning and the end of the data are joined and that any structural breaks within the run of the data are absorbed by the trend function. It seems that it is partly for want of adequate detrending that econometricians have failed to recognise the prevalence of the frequency limits to their data.

The typical band-limited frequency structure can be illustrated by a data sequence that requires only a simple method detrending. Figure 1 shows the logarithms of U.S. quarterly gross domestic product (GDP) for the period 1968–2007 through which a trend line has been interpolated. The trend has been calculated using the filter of Hodrick and Prescott (1980, 1997) with the smoothing parameter set to the value of 1600. (The filter in question is also attributable to Leser 1961.)

Figure 2 shows the periodogram of the deviations from the trend. It will be seen that the spectral structure extends no further than the frequency of $\pi/4$ radians per quarter. The remainder of the periodogram comprises what may be described as a dead space that contains only the traces of minor elements of noise.

One might hope to characterise the essential dynamic properties of the US GDP process by fitting an ARMA model of relatively low orders. Thus, a second-order autoregressive AR(2) model with complex conjugate roots within the autoregressive polynomial should serve to represent the dynamics. The argument of the complex roots should characterise the angular velocity the process, or equivalently, the length of its cycles, and their modulus should characterise its damping properties.

However, when such a model is fitted to the data, it transpires, almost invariably, that the roots of the autoregressive polynomial are real-valued. (For a testimony to this, see Pagan 1997.) Figure 2 shows such an outcome, which can be discerned from the shape of the parametric spectral density function, or "spec-
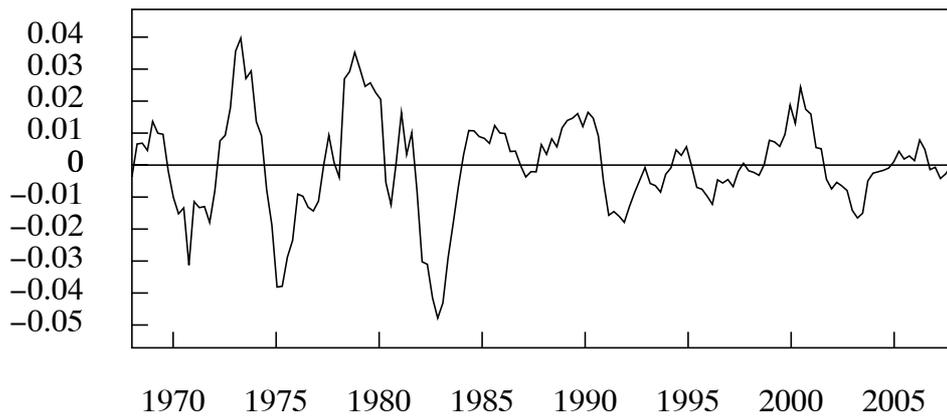
Figure 1. The deviations of the logarithmic quarterly index of Real US GDP from an interpolated trend. The observations are from 1968 to 2007. The trend is determined by a Hodrick–Prescott (Leser) filter with a smoothing parameter of 1600.
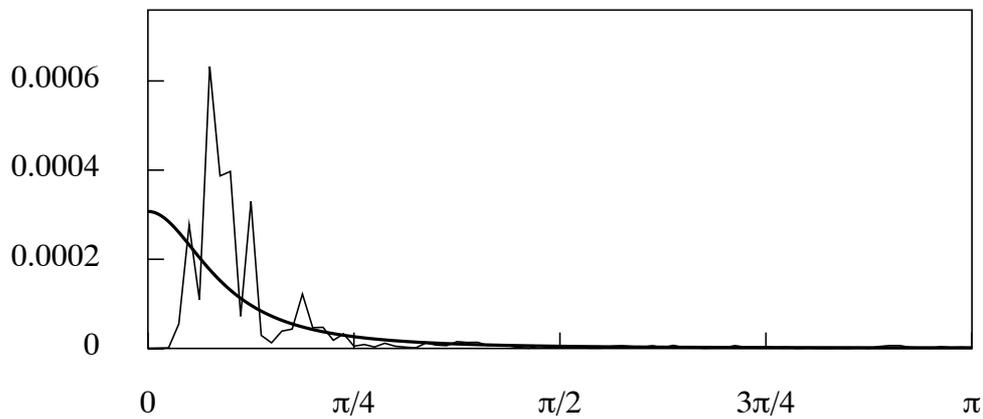
Figure 2. The periodogram of the data points of Figure 1 overlaid by the parametric spectral density function of an estimated regular AR(2) model.
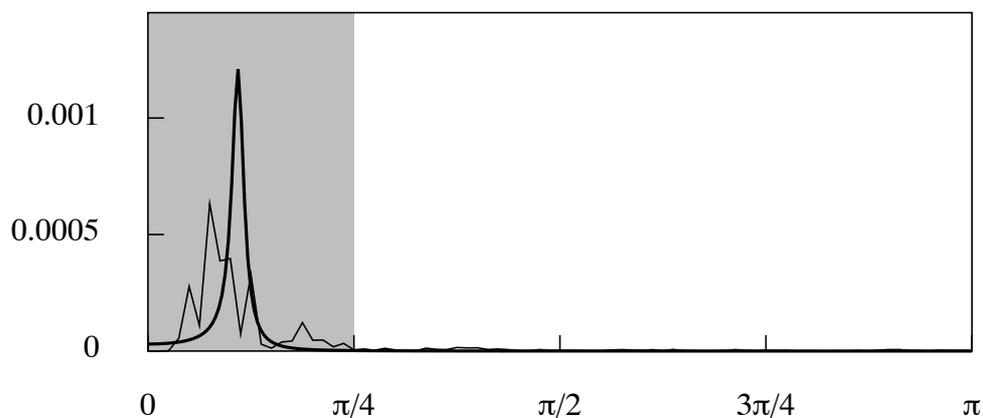
Figure 3. The periodogram of the data points of Figure 1 overlaid by the spectral density function of an AR(2) model estimated from band-limited data.

trum", of the estimated model.

A reasonable recourse in the face of such difficulties is to employ an estimator in which greater weight is placed on the low-frequency elements of the data than on the high-frequency elements. If the interest is confined to the business cycles, then it is reasonable to attribute weighs of unity to the elements within the range of the business cycle frequencies and to attribute weights of zero to those elements that lie outside the range. This is tantamount to applying an unweighted estimator to data that is band-limited to a low-frequency range, which may have been created by applying a perfect lowpass filter to the original data. (For a description of such a filter, see Pollock, 2009)

The consequence of applying an unrestricted estimator to data that are strictly band-limited will be to create an estimated autoregressive polynomial in which the complex roots approach the perimeter of the unit circle. The effect will be to misrepresent the damping properties of the process. Figure 3 illustrates this outcome.

The only difference between the circumstances depicted in Figures 2 and 3 concerns an almost imperceptible noise contamination that is found throughout the interval $(\pi/4, \pi]$. Whereas this noise is affecting the estimates that give rise to the parametric spectrum of Figure 2, is has been eliminated from estimation of the spectrum of Figure 3.

In this paper, we investigate the effects of applying an unrestricted estimator of an ARMA model to a band-limited data sequence that has been sampled at an excessive rate from a continuous process that is limited in its frequency content. This investigation is preceded by an analysis of the effects of a continuous band-limited process. We show that, if the data are sampled at a rate that is precisely attuned to the highest frequency that is present, then there will be a one-to-one correspondence between the continuous-time processes and their discrete-time representations.

In the case of data that have been sampled over-rapidly, it is possible to reconstitute the underlying continuous process from the Fourier ordinates of the sample. Thereafter, this process can be resampled at the appropriate rate, which is lower than the original rate. Then, an ARMA model can be estimated from the resampled data without incurring the problems that have been outlined above.

2. The Sampling Process

The sampling theorem of Shannon and Nyquist establishes that, if the maximum frequency within a continuous square-integrable function is ω_c radians per unit period, then all of the information within the function can be conveyed via a sequence of values sampled regularly from the function at intervals of π/ω_c units of time. (See Shannon, 1949.)

The theorem implies that, if samples are taken at unit intervals, then the maximum frequency that can be detected is at π radians per interval, which is the so-called Nyquist frequency. Moreover, if the underlying function contains frequencies in excess of π, then, via the process of aliasing, these will be confounded with frequencies that lie within the Nyquist Interval $[0, \pi]$.

The interval $[0, \pi]$ is appropriate to the case where the function in question is expressed in terms of sines and cosines. When it is expressed equivalently in terms
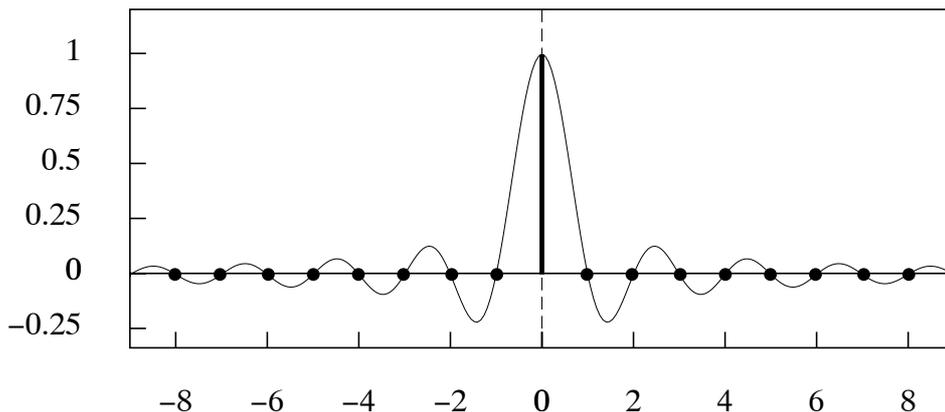
Figure 4. The sinc function wave-packet $\phi(t) = \sin(\pi t)/\pi t$ comprising frequencies in the interval $[0, \pi]$.

of complex exponential functions, the Nyquist frequency range becomes $[-\pi, \pi]$. In the case of a function with a maximum frequency of π radians per sample interval, it will be possible, in theory, to reconstitute the continuous function from its sampled values by associating to each of them a so-called sinc function kernel scaled by that value.

The sinc function $\phi(t)$ is the Fourier transform of a frequency-domain rectangle supported on the interval $[-\pi, \pi]$; and it is just a sine function to which a hyperbolic taper has been applied:

$$\phi(t) = \frac{\sin(\pi t)}{\pi t} = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{i\omega t} d\omega. \tag{1}$$

This is illustrated in Figure 4.

The following expression for the continuous function is derived by applying a sinc function kernel to each element of the sequence $\{x_t; t = 0, \pm 1, \pm 2, \ldots\}$:

$$x(t) = \sum_{k=-\infty}^{\infty} x_k \frac{\sin\{\pi(t-k)\}}{\pi(t-k)} = \sum_{k=-\infty}^{\infty} x_k \phi(t-k). \tag{2}$$

Here, $t \in \mathcal{R}$ is to be regarded as a continuous index of time. The formula indicates that there is a one-to-one correspondence between continuous functions limited in frequency to the Nyquist interval and their sampled ordinates taken at unit intervals in time.

Although the Shannon–Nyquist theorem is proved in relation to a square integrable function, it is clear that the construction above can be extended to encompass stationary stochastic processes. Thus, it might be assumed that the doubly-infinite sequence $\{x_t; t = 0, \pm 1, \pm 2, \ldots\}$ is generated by an ordinary autoregressive moving-average process.

The sequence of sinc functions $\phi(t-k); k \in \mathcal{Z} = \{0, \pm 1, \pm 2, \ldots\}$ constitutes an orthogonal basis for the set of all continuous analytic functions that lie within the Nyquist frequency interval. To show this, let $\phi(\omega)$ be the Fourier transform of

5

$\phi(t)$ and consider the following autoconvolution:

$$
\begin{aligned}
\int_t \phi(t)\phi(\tau - t)dt &= \int_t \phi(t)\left\{\frac{1}{2\pi}\int_\omega \phi(\omega)e^{i\omega(\tau-t)}d\omega\right\}dt \\
&= \frac{1}{2\pi}\int_\omega \phi(\omega)\left\{\int_t \phi(t)e^{-i\omega t}dt\right\}e^{i\omega\tau}d\omega \\
&= \frac{1}{2\pi}\int_\omega \phi(\omega)\phi(\omega)e^{i\omega\tau}d\omega.
\end{aligned}
\tag{3}
$$

The symmetry of $\phi(t)$ allows us to write $\phi(\tau - t) = \phi(t - \tau)$, whereas the idempotency of $\phi(\omega)$ gives $\phi^2(\omega) = \phi(\omega)$. Together, these two conditions indicate that $\phi(t)$ is its own autocorrelation function. Therefore, the condition

$$
\phi(t) = 0 \quad \text{for} \quad t \in \{\pm 1, \pm 2, \ldots\},
\tag{4}
$$

which is manifest in the formula of (1), indicates that sinc functions separated by integer distances are mutually orthogonal.

When the set of sinc functions $\{\phi(t - k); k \in \mathcal{Z}\}$ at unit displacements are sampled at the integer values of t, the result is nothing but the set of unit impulses at the integer points. This constitutes a basis for the set of all sequences defined over the set of integers.

A continuous function $y(t)$ that is limited by the frequency value $\omega_c < \pi$ can be expressed in the manner of (2) as

$$
y(t) = \sum_{j=-\infty}^{\infty} y_j \frac{\sin\{\omega_c(t - j)\}}{\omega_c(t - j)} = \sum_{j=-\infty}^{\infty} y_j \phi_c(t - j),
\tag{5}
$$

where j in an index that demarcates time intervals of $\pi/\omega_c > 1$ units. The set of sinc functions $\{\phi_c(t - j); k \in \mathcal{Z}\}$ constitutes an orthogonal basis for functions limited to the interval $[-\omega_c, \omega_c]$. Let the sequence $\{h_t; t = 0, \pm 1, \pm 2, \ldots\}$ be the ordinates sampled from the function $\phi_c(t)$ a unit intervals of time. Then, according to the Shannon–Nyquist theorem, there is

$$
\phi_c(t) = \sum_{k=-\infty}^{\infty} h_k \phi(t - k),
\tag{6}
$$

which is an instance of a wavelet dilation equation. (See, for example, Burrus *et al.*, 1998). Putting this expression into (5) gives the expression for $y(t)$ in terms of the orthogonal basis corresponding to the Nyquist interval:

$$
y(t) = \sum_{j=-\infty}^{\infty} y_j \left\{\sum_{k=-\infty}^{\infty} h_k \phi(t - j - k)\right\}.
\tag{7}
$$

The comparison of this expression with the more parsimonious expression of (5) shows the advantage of employing a set of basis functions that cover the same frequency interval as the function $y(t)$.

The reconstruction or interpolation of a function in the manner suggested by the sampling theorem is not possible in practice, because it requires summing an infinite number of sinc functions, each of which is supported on the entire real line. Nevertheless, a continuous band-limited periodic function, defined on a finite interval, can be reconstituted from a finite number of wrapped or periodic sinc functions, which are Dirichlet kernels by another name. The Dirichlet kernel is obtained by sampling the frequency-domain rectangle of the sinc function.

Consider a continuous function $x(t)$ defined the interval $[0, T)$. Such a function can be regarded as a single cycle of a periodic or circular function such that $x(t) = x(t + T)$; and, therefore, it has a Fourier series expansion in terms of the complex exponential functions $\exp(i\omega_j)$, where $\omega_j = 2\pi j/T$ is the jth Fourier frequency. If the function $x(t)$ is bounded by the Nyquist frequency, then the relevant expression for the series expansion of $x(t)$, together with the inverse transformation, which provides the Fourier coefficients ξ_j, is

$$x(t) = \sum_{j=0}^{T-1} \xi_j e^{i\omega_j t} \longleftrightarrow \xi_j = \frac{1}{T} \sum_{t=0}^{T-1} x(t) e^{-i\omega_j t}. \tag{8}$$

Here, $t \in \mathcal{R}$ is considered to be a continuous index of time. When $t = 0, 1, \ldots, T - 1$ is an integer, the equation represents of the discrete Fourier transform of the sampled sequence $\{x_t; t = 0, 1, \ldots, T - 1\}$ together with the inverse transform.

Putting the expressions for the Fourier ordinates at the sample points into the finite Fourier series expansion of the time function and commuting the summation signs gives

$$x(t) = \sum_{j=0}^{T-1} \left\{ \frac{1}{T} \sum_{k=0}^{T-1} x_k e^{-i\omega_j k} \right\} e^{i\omega_j t} = \frac{1}{T} \sum_{k=0}^{T-1} x_k \left\{ \sum_{j=0}^{T-1} e^{i\omega_j(t-k)} \right\}. \tag{9}$$

The inner summation of the final expression gives rise to the Dirichlet Kernel:

$$\phi_n^\circ(t) = \sum_{t=0}^{T-1} e^{i\omega_j t} = \frac{\sin(\{(T-1)/2\}\omega_1 t)}{\sin(\omega_1 t/2)}. \tag{10}$$

Thus, the Fourier expansion can be expressed in terms of the Dirichlet kernel, which is a circularly wrapped sinc function:

$$x(t) = \frac{1}{T} \sum_{t=0}^{T-1} x_k \phi_n^\circ(t - k). \tag{11}$$

The functions $\{\phi^\circ(t - k); k = 0, 1, \ldots, T - 1\}$ are appropriate for reconstituting a continuous periodic function $x(t)$ defined on the interval $[0, T)$ from its sampled ordinates $x_0, x_1, \ldots, x_{T-1}$. However, the periodic function can also be reconstituted by an ordinary Fourier interpolation

$$x(t) = \sum_{j=0}^{T-1} \xi_j e^{i\omega_j t} = \sum_{j=0}^{[T/2]} \left\{ \alpha_j \cos(\omega_j t) + \beta_j \sin(\omega_j t) \right\}, \tag{12}$$

where $[T/2]$ denotes the integral part of $T/2$ and where $\alpha_j = \xi_j - \xi_{-j}$ and $\beta_j = i(\xi_j + \xi_{-j})$ are the coefficients from the regression of the data on the sampled ordinates of the cosine and sine functions at the various Fourier frequencies.

3. Band-Limited Stochastic Processes

A continuous-time autoregressive moving-average process that is supported on the Nyquist interval $[-\pi, \pi]$ may be derived from an ordinary discrete-time process with unit time intervals by associating sinc functions to each of its ordinates. By virtue of the Shannon–Nyquist theorem, there will be a one-to-one correspondence between the discrete-time processes and the band-limited continuous-time processes.

The continuous-time ARMA process can be derived directly by applying an appropriate filter to a continuous band-limited white-noise process supported on the Nyquist interval. The continuous white-noise process is obtained from a train of sinc-function wave packets arriving regularly at unit intervals of time and having amplitudes that are distributed independently and identically in the manner of the ordinates of a discrete-time white-noise process.

Let $\{y_t; t = 0, \pm 1, \pm 2, \ldots\}$ be the sampled ordinates of the ARMA process and let $\{\varepsilon_t; t = 0, \pm 1, \pm 2, \ldots\}$ be the ordinates of the white-noise forcing function. The corresponding continuous-time functions are

$$y(t) = \sum_{k=-\infty}^{\infty} y_k \phi(t-k) \quad \text{and} \quad \varepsilon(t) = \sum_{k=-\infty}^{\infty} \varepsilon_k \phi(t-k), \tag{13}$$

respectively, where $t \in \mathcal{R}$ and $k \in \mathcal{Z}$ and where $\phi(t)$ is the sinc function kernel. The equation of the continuous-time ARMA process is

$$\sum_{j=0}^{p} \alpha_j y(t-j) = \sum_{j=0}^{q} \mu_j \varepsilon(t-j), \tag{14}$$

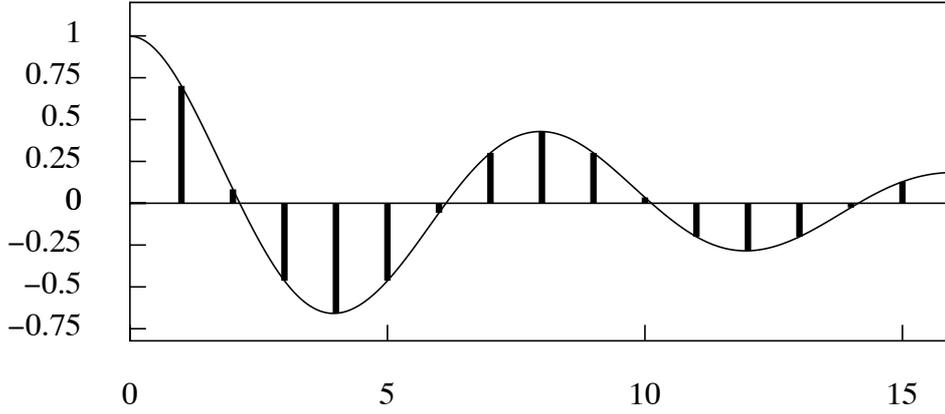where $\alpha_0 = 1$. This has a moving-average representation in the form of

$$y(t) = \sum_{j=0}^{\infty} \psi_j \varepsilon(t-j), \tag{15}$$

where the coefficients are from the series expansion of the rational function $\mu(z)/\alpha(z) = \psi(z)$. The autocovariance function of the band-limited white-noise process is

$$\gamma_\varepsilon(\tau) = \sigma_\varepsilon^2 \phi(\tau) = \sigma_\varepsilon^2 \frac{\sin(\pi\tau)}{\pi\tau}, \tag{16}$$

where $\tau \in \mathcal{R}$ and σ_ε^2 is the variance parameter. The autocovariance function of the continuous ARMA process is given by

$$\gamma(\tau) = E\left\{ \left[\sum_{i=0}^{\infty} \psi_i \varepsilon(t-\tau-i) \right] \left[\sum_{j=0}^{\infty} \psi_j \varepsilon(t-j) \right] \right\}$$
$$= \sum_{i=-\infty}^{\infty} \gamma_i \phi(\tau-i), \tag{17}$$

Figure 5. A continuous autocovariance function of an AR(2) process, obtained via the inverse Fourier transform of the spectral density function, together with the corresponding discrete-time autocovariances, calculated from the AR parameters.

where $\gamma_i = \sigma_\varepsilon^2 \sum_j \psi_j \psi_{j+i}$ is the ith autocovariance of the discrete-time process. It can be seen immediately that $\gamma(\tau) = \gamma_\tau$, when τ takes an integer value, and that the continuous-time autocovariance function is obtained from the discrete-time function by sinc-function interpolation.

The autocovariances of a discrete-time ARMA process are the coefficients of the series expansion of the generating function

$$\gamma(z) = \sigma_\varepsilon^2 \frac{\mu(z)\mu(z^{-1})}{\alpha(z)\alpha(z^{-1})}, \qquad (18)$$

where $\alpha(z) = \alpha_0 + \alpha_1 z + \cdots + \alpha_p z^p$ and $\mu(z) = \mu_0 + \mu_1 z + \cdots + \mu_q z^q$ are the autoregressive and moving-average polynomials respectively.

The spectrum of the process, denoted by $f(\omega)$, is the Fourier transform of the sequence of autocovariances. It is generated by the function $\gamma(z)$ when z travels around the perimeter of the unit circle in the complex plane. When $z = \exp\{-i\omega\}$ within $\gamma(z)$, the spectrum is generated when ω runs from $-\pi$ to π.

The spectrum of a continuous-time ARMA process that is limited in frequency to the Nyquist interval $[-\pi, \pi]$ has the same values over the interval as the periodic spectrum of the corresponding discrete-time process. The autocovariance function of the continuous-time process is given by the inverse Fourier integral transform of this spectrum. Thus

$$\gamma(\tau) = \int_{-\pi}^{\pi} e^{i\omega\tau} f(\omega) d\omega = \int_0^{\pi} 2\cos(\omega\tau) f(\omega) d\omega, \qquad (19)$$

where the second equality follows in consequence of the symmetry of $f(\omega) = f(-\omega)$.

Since no tractable analytic expression is available for evaluating this integral exactly, it must be approximated via a discrete cosine Fourier transform:

$$\gamma(\tau) \simeq \gamma_N^\circ(\tau) = \frac{2\pi}{N} \sum_{j=0}^{[N/2]} \cos(\omega_j \tau) f(\omega_j), \quad \omega_j = \frac{2\pi j}{N}. \qquad (20)$$

9

Here, N is the number of points sampled from the function $f(\omega)$ over the interval $[-\pi, \pi]$, whereas $\gamma_N^\circ(\tau)$ is the circular autocovariance for samples of size N. It will be found is $\gamma_N^\circ(\tau) \to \gamma(\tau)$ as $N \to \infty$.

An example of such an approximation is in Figure 5, where a sequence of $N = 256$ elements sampled from the function $f(\omega)$ have been transformed to create a continuous piecewise linear rendition of $\gamma(\tau)$ over the range $\tau \in [0, 16]$. The ordinates of the discrete-time autocovariance function are also plotted on the diagram. These have been obtained directly from the autoregressive parameters. Their coincidence with the ordinates of the continuous function testifies to the accuracy of the approximation.

4. Fitting ARMA Models to Oversampled Data

From the autocovariance function of a continuous-time process supported on the Nyquist interval $[-\pi, \pi]$, it will be possible to derive the autocovariances of a discrete-time process that would be obtained by over-rapid sampling.

If the sampled values are separated by intervals of $\omega_c/\pi < 1$ units of time, then the spectrum of the resulting discrete-time process will be supported on an interval $[-\omega_c, \omega_c]$, which is a subset of the Nyquist interval. Its spectrum will be zero-valued over the complementary interval within Nyquist interval, which has been described as a dead space.

Our purpose is to determine the limiting values of the estimates that would be obtained from the oversampled data of a continuous band-limited process. Having sampled the autocovariances from the continuous function, a method of moments can be used to infer the corresponding parameters.

In the case of a pure AR process, this is a matter of solving the Yule–Walker equations. In the case of an ARMA model, an iterative procedure is also required for finding the moving-average parameters via a Cramér–Wold factorisation. An effective procedure that exploits the Newton–Raphson algorithm has been expounded by Tunicliffe–Wilson (G.T. Wilson) (1969). The Yule–Walker procedure and the Cramér–Wold factorisation have been combined in a procedure, coded in both C and Pascal, that has been provided by Pollock (1999).

This procedure has been in incorporated in a program, OVERSAMPLE.PAS, that is available at the following web address:

<div align="center">

`http://www.le.ac.uk/users/dsgp1/`

</div>

Apart from determining the effects of oversampling, the program can be used to investigate the effects of a white-noise contamination of the data generated by the specified process.

Here, the program will be used to reaffirm what has already been revealed using an empirical data sequence. For the experiments, an ARMA(2, 2) process is specified, which is regarded as the truth. This is represented by the equation

$$y(t) + \alpha_1 y(t-1) + \alpha_2 y(t-2) = \varepsilon(t) + \mu_1 \varepsilon(t-1) + \mu_2 \varepsilon(t-2). \qquad (21)$$

The autoregressive parameters are set to $\alpha_1 = -1.0607$ and $\alpha_2 = 0.5625$. This is equivalent to specifying a pair of conjugate complex roots $\rho \exp\{\pm i\theta\}$ with a modulus $\rho = 0.75$ and with arguments of $\pm\theta = \pm 45°$. The moving-average parameters
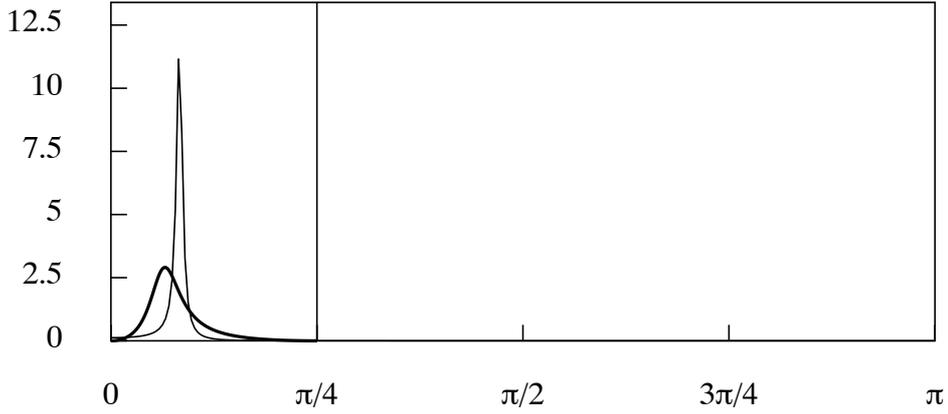
Figure 6. The parametric spectrum of an ARMA(2, 2) process, limited in frequency to π radians per period and oversampled at the rate of 4 observations per period, represented by a heavy line, together with the parametric spectrum of an AR(2) model fitted to the sampled autocovariances.

are $\mu_1 = 0$ and $\mu_2 = 1$. Their effect is to place zeros on the unit circle at zero frequency and at the Nyquist frequency of π radians per unit sample period.

Figure 6 shows, via the heavy line, the parametric spectrum of the specified process when it has be subject to sampling at the rate of 4 observations per unit period, which confines the plotted spectrum to the frequency band $[0, \pi/4]$. With this rate of sampling, the effective modulus of the autoregressive roots is $0.9306 = (0.75)^{1/4}$ and their effective arguments are $\pm 11.25° = \pm 45°/4$. The emptiness of the dead space in the interval $(\pi/4, \pi]$, which signifies the absence of any noise contamination, can be interpreted either as the result of the application to the data of a perfect low pass filter or as a consequence of a weighting scheme within the estimator that sets the spectral ordinates within this interval to zero.

The figure also shows the parametric spectrum of an AR(2) model fitted to the autocovariances sampled at this rate from their continuous function. It can be seen that the spectral spike of the model is excessively prominent and that its peak is at a higher frequency that the peak of the spectrum of the true process.

The estimated autoregressive parameters are $\alpha_1 = -1.9062$ and $\alpha_2 = 0.9735$, which correspond to conjugate complex roots with a modulus $\rho = 0.9866$ and with arguments of $\pm \theta = \pm 14.9873°$. The polar parameters imply cyclical fluctuations that are both more rapid and more persistent than the actual fluctuations.

Figure 7 shows the parametric spectrum of the oversampled process supported on the spectrum of a minimal white-noise contamination that extends of over the entire frequency range of $[0, \pi]$ and which is represented by a broken line that is barely raised above the horizontal axis. This is to be construed as a discrete-time process that contaminates the individual observations, as opposed to a continuous background noise. The variance of the noise is 0.05 of the variance of the ARMA(2, 2) process. The figure also shows the parametric spectrum of an AR(2) model that has been fitted to the autocovariances of the contaminated process.

The difference between the true spectrum and the spectrum of the fitted model is considerable. The fitted spectrum betrays the fact that the autoregressive poly-
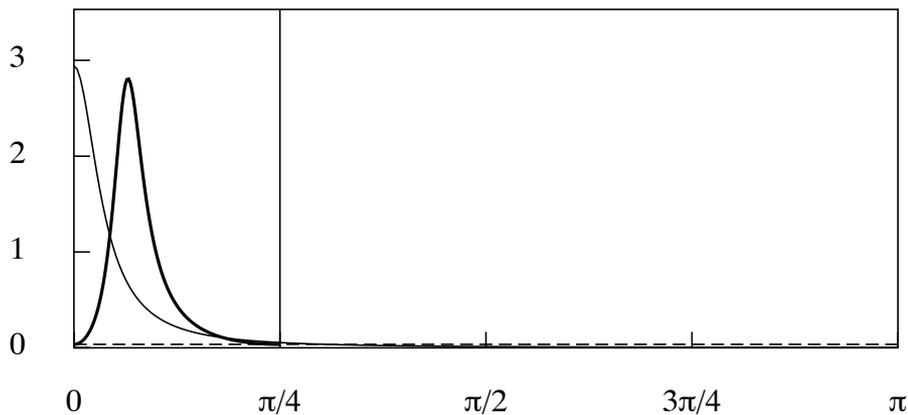
Figure 7. The parametric spectrum of the oversampled ARMA(2, 2) process, represented by a heavy line, supported on the spectrum of a white-noise contamination, together with the parametric spectrum of an AR(2) model fitted to the sampled autocovariances.

nomial of the model has a pair or real-valued roots in place of the conjugate complex roots of the process. The estimated autoregressive parameters are $\alpha_1 = -1.0386$ and $\alpha_2 = 0.1291$ and the corresponding real-valued roots of the autoregressive polynomial are $\lambda_1 = 0.1443$ and $\lambda_2 = 0.8943$.

The results of these experiments can be explained by reference to the auto-covariance function. When the rate of sampling is excessive, the autocovariances will be sampled at points that are too close to the origin, where the variance is to be found. Then, their values will decline at a diminished rate. This reduction in the rate of convergence is reflected in the modulus of the estimated complex roots, which indicates a rate of damping that understates the true value.

The opposite effect is experienced when there is a noise contamination. Then, the full variance of the noise will be added to the variance of the underlying process. Nothing will be added to the adjacent sampled ordinates of autocovariance function. Therefore, the sampled autocovariances will decline at an enhanced rate. If this rate of convergence exceeds the critical value, then there will be a transition from cyclical convergence to monotonic convergence. In that case, the estimated autoregressive roots will be real-valued, which belies the cyclical nature of the true process.

There are some simple strategies for alleviating the problems of noise contamination. These methods exploit the fact that the parameters of a discrete-time AR(p) process can be inferred from any sequence of $2p + 1$ consecutive elements of its autocovariance function—and from no more than $p + 1$ elements if one fully exploits the symmetry of the autocovariance function.

To avoid the noise contamination that might beset the leading autocovariances, one can infer the parameters from the sequence $\gamma_q, \gamma_{q+1}, \dots, \gamma_{q+2p}$, for some $q > 0$, instead inferring them from of the sequence $\gamma_0, \gamma_1, \dots, \gamma_p$. (This is neglecting the fact that, in a limited sample, the higher-order autocovariances loose some reliability, since they embody fewer sample points.)

Equally, in applying the method of moments to an ARMA(p, q) model, which might be an appropriate model if the process is contaminated by noise, the sequence

12
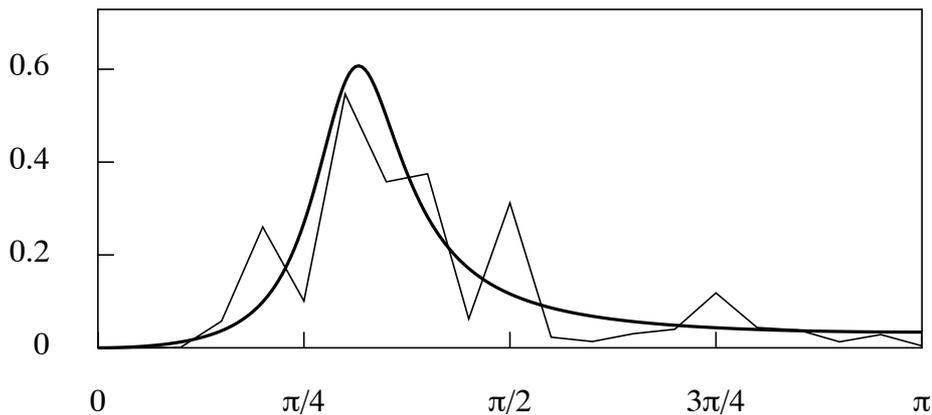
Figure 8. The periodogram of the data that have been filtered and subsampled at the rate of 1 observation in 4 overlaid by the parametric spectrum of an estimated ARMA(2, 1) model.

$\gamma_{q+1-p}, \gamma_{q+2-p}, \ldots, \gamma_{q+p}$ is used in calculating the autoregressive parameters, while the sequence $\gamma_0, \gamma_1, \ldots, \gamma_q$ is entailed in the calculation of the moving-average component. (See, for example, Pollock 1999, p. 545.) Evidence is available from the above-mentioned computer program to suggest that the problems besetting oversampled processes will be partially alleviated by increasing the autoregressive order and by including a moving-average component.

An alternative approach to improving the performance of an AR model relies upon a weighted Whittle estimator. This approach had been pursued by various authors including, notably, Haywood and Tunnicliffe–Wilson, (1997) and Proietti (2008). However, it is a delicate matter to find a weighting scheme that will allow one to navigate between the opposing hazards that are represented by the two experiments of this section. Nevertheless, it has been shown that an appropriate weighting scheme can serve to enhance greatly the forecasting performance of an ARMA model.

Another method, which may have a similar inspiration to that of Haywood and Tunnicliffe–Wilson (1997), is due to Morton and Tunnicliffe–Wilson (2004). They propose a model that incorporates a lowpass filter that depends upon a single estimable parameter. The filter serves to attribute appropriate weights to the Fourier elements in the low-frequency range.

The method that we propose for dealing with an oversampled frequency-limited process differs from all of the foregoing methods. We propose simply to reduce the rate of sampling. The first step is to reconstitute a continuous trajectory from the Fourier ordinates that lie within the frequency band in question.

The nonzero ordinates that lie outside the band are liable to represent elements of a contaminating noise, or else they may be the spectral signatures of a seasonal fluctuation that, in the circumstances, we would wish to remove from the data.

The second step is to sample the continuous trajectory at the rate that is precisely attuned to the highest frequency that it contains. Then, an ARMA model can be fitted in the usual way to the resampled data.

This strategy has been applied to the data represented in Figure 1. The

13

continuous trajectory that is constituted from the Fourier ordinates of the data that lie in the interval $[0, \pi/4]$ is represented in Figure 9 by the heavy line. An ARMA(2, 1) model has been fitted to values sampled from this trajectory at $1/4$ of the original sample rate, which makes this an annual rate.

In Figure 8, the parametric spectrum of the fitted model has been superimposed upon the periodogram of the subsampled data. It will be seen that this periodogram has exactly the profile of the spectral structure that is supported of the interval $[0, \pi/4]$ in Figures 2 and 3. The parametric spectrum fits the periodogram well.

In fitting the model, the noise in the interval $(\pi/4, \pi]$ has been removed, but no attempt has be made to remove whatever small quantity of noise there might be within the interval $[0, \pi/4]$ that supports the subsampled data. The first-order moving-average polynomial within the fitted model is the difference operator $1 - z$, which places a zero in the spectrum at zero frequency. (In terms of equation (21), there would be $\mu_1 = -1$ and $\mu_2 = 0$.)

To impose this feature on the model, the data can be accumulated via the summation operator, which is the inverse of the difference operator, wherafter the autoregressive parameters, with values of $\alpha_1 = -0.9487$ and $\alpha_2 = 0.6355$, can be estimated. Alternatively, when an unrestricted ARMA(2,1) is estimated from the subsampled data, a moving-average parameter of $\mu_1 = -0.\overline{9999}$ is delivered, together with autoregressive parameters that have virtually the values given above.

Although the periodogram of Figure 8 is zero-valued at the Nyquist frequency of π, there would be little advantage in including a zero at this frequency within the fitted model. The effect of such a zero would need to be confined to the neighbourhood of the Nyquist frequency; and for this to be achieved, it would be necessary to balance the zero with an autoregressive pole located nearby.

5. The Characteristics of a Band-Limited Process

A function that is limited in frequency is also an analytic function. For a proof of this, see, for example, Percival and Walden (1993). Such functions possess derivatives of all orders, albeit that the higher-order derivatives may be zero-valued. Knowing the values of an analytic function over a finite interval should allow one, in principle, to infer its values at all other points. The analytic nature of a stochastic band-limited function has some significant implications.

First, there is the differentiability of the function. In principle, this should enable one to determine its turning points by locating the points where its derivative is zero-valued. According to equation (12), a periodic function defined on the interval $[0, T)$ that is limited in frequency by the Nyquist value of π can be expressed as a weighted sum of T trigonometrical functions. For small values of T, it should be straightforward to express the derived function as the weighted sum of the derivatives of its trigonometrical elements.

In Figure 9, a smooth differentiable function has been interpolated through the data points previously displayed in Figure 1. The function has been synthesised from a limited number of the Fourier ordinates of the data that extend in frequency no further that the value of $\pi/4$. The derivative function is also plotted on the diagram; and the points where it cuts the horizontal axis mark the turning points of the smooth trajectory.
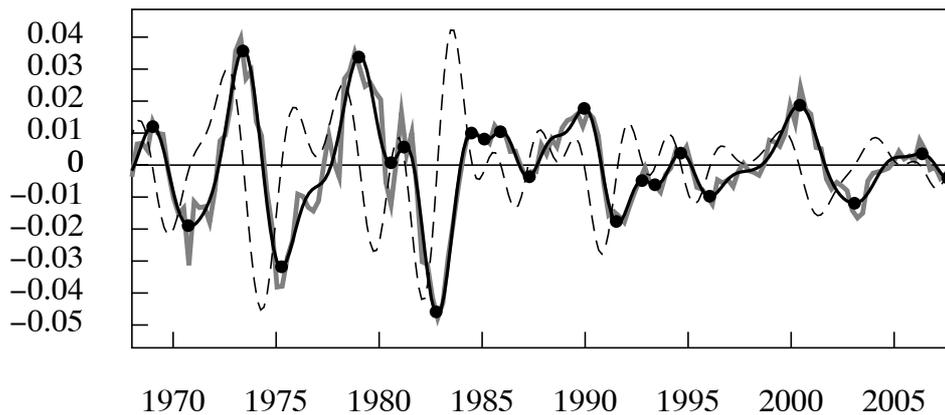
Figure 9. The turning points of the GDP index of Figure 1. The underlying trajectory is marked by the heavy line. Its turning points are marked by dots, which have been located by the zeros of the derivative function, which is represented by the discontinuous line. A plot of the original data is faintly visible beneath the curves.

This method of finding the turning points of an econometric data sequence differs markedly from the conventional method of Bry and Boschan (1971); and it is arguable that it is both simpler and more robust. It is notable that the method is practical only in so far as T has a limited value. The method is inconceivable in the context of a doubly-infinite sequence of an aperiodic nature, which is the usual subject of the basic theory of time-series analysis.

Next is the question of the predictability of a band-limited stochastic process. This topic has been pursued extensively in the literature of electrical engineering, where it is generally assumed that the data constitute a semi-infinite sequence stretching indefinitely into the past.

It has been observed that the only requirement for perfect predictability is that the process should be sampled at a rate that exceeds the Nyquist value, so as to allow two observations to be made in less time than it takes the data element of highest frequency to execute its cycle. (See Mugler 1990, for example.)

Similarly, it has been shown by Lyman *et al.* (2000) that, if the past values of the process are known over an interval of arbitrary positive length, then the mean-squared prediction error may be made arbitrarily small, regardless of how far in the future the prediction is made. It has been noted, however, that the result presupposes the absence of any rounding error or noise contamination in the observations.

The provenance of these theoretical results can be explained by resolving the analytic function into a sequence of sinc functions in the manner of equation (2). Given that the sinc functions have infinite supports in the time domain, it must be the case that, if the rate of sampling exceeds the critical Nyquist value, then every observation that is taken will comprise elements of all of the functions, be they ever so remote for the point of observation. By taking a denumerable infinity of observations, information would be gathered on every sinc function; and, by these means, the entire trajectory of the analytic function could be determined.

However, this result is best described an analytic fantasy. It depends crucially

on the infinite supports of the frequency-bounded sinc functions. This feature is an artefact of our analytic method. For, it is an inescapable fact that a function that is limited in frequency, as are the sinc functions, cannot also be limited in time. Conversely, if a function is limited in time, then its expression as a weighted combination of perpetual trigonometric functions entails an infinite number of such functions over an unbounded range of frequencies.

A dilemma of Fourier analysis is that one is required to choose between functions that are limited in frequency and functions that are limited in time. (This issue has been discussed incisively by Slepian, 1983.) In this paper, we have chosen to limit the frequencies. The dilemma could be avoided by adopting an alternative analytic framework. In place of a basis of trigonometric functions, one might choose to use a basis of wavelet functions on finite supports. Such functions would be limited in scale, which is analogous to a frequency limitation. However, the tractability of a Fourier analysis discourages us from following such an approach.

References

Bry, G., and C. Boschan, (1971), *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs,* National Bureau of Economic Research.

Burrus, C.S., R.A. Gopinath and H. Guo, (1998), *Introduction to Wavelets and Wavelet Transforms: A Primer,* Prentice Hall, New Jersey.

Granger, C.W.J., (1966), The Typical Spectral Shape of an Economic Variable, *Econometrica,* 34, 150–161.

Haywood, J., and G. Tunnicliffe–Wilson, (1997), Fitting Time Series Models by Minimising Multistep-Ahead Errors: A Frequency Domain Approach, *Journal of the Royal Statistical Society, Series B (Methodological),* 59, 237–254.

Hodrick, R.J., and E.C. Prescott, (1980), *Postwar U.S. Business Cycles: An Empirical Investigation,* Working Paper, Carnegie–Mellon University, Pittsburgh, Pennsylvania.

Hodrick R.J., and E.C. Prescott, (1997), Postwar U.S. Business Cycles: An Empirical Investigation, *Journal of Money, Credit and Banking,* 29, 1–16.

Leser, C.E.V., (1961), A Simple Method of Trend Construction, *Journal of the Royal Statistical Society, Series B,* 23, 91–107.

Lyman, R.J., W.W. Edmonson, S. McCullough, and M. Rao (2000), The Predictability of Continuous-Time, Bandlimited Processes, *IEEE Transaction of Signal Processing,* 48. 311–315.

Morton, A.S., and G. Tunnicliffe–Wilson, (2004), A Class of Modified High-Order Autoregressive Models with Improved Resolution of Low Frequency Cycles, *Journal of Time Series Analysis,* 25, 235–250.

Mugler, D.H., (1990), Computationally-effcient Linear Prediction of a Band- Limited Signal and its Derivative, *IEEE Transactions on Information Theory,* IT-36, 589–596.

Pagan, A., (1997), Towards an Understanding of Some Business Cycle Characteristics, *The Australian Economic Review,* 30, 1–15.

Percival, D.B., and A.T. Walden, (1993), *Spectral Analysis for Physical Applications: Multitaper and Conventional Techniques,* Cambridge University Press, Cambridge.

Pollock, D.S.G., (1999), *A Handbook of Time-Series Analysis, Signal Processing and Dynamics,* Academic Press, London.

Pollock, D.S.G., (2009), Realisations of Finite-sample Frequency-selective Filters, *Journal of Statistical Planning and Inference,* 139, 1541–1558.

Proietti, T., (2008), Band Spectral Estimation for Signal Extraction, *Economic Modelling,* 25, 54–69.

Shannon, C.E., (1949), Communication in the Presence of Noise, *Proceedings of the Institute of Radio Engineers,* 37, 10–21. Reprinted in 1998, *Proceedings of the IEEE,* 86, 447–457.

Sims, C.A., (2010), Continuous and Discrete Time Models, pps. 60–67 in S.N. Durlauf and L.E. Blume, (eds.), *Macroeconometrics and Time Series Analysis,* Palgrave Macmillan, Basingstoke.

Slepian, D., (1983), Some Comments on Fourier Analysis, Uncertainty and Modeling, *SIAM Review,* 25, 379–393.

Wilson, G.T., (1969), Factorisation of the Covariance Generating Function of a Pure Moving Average Process, *SIAM Journal of Numerical Analysis,* 6, 1–7.