Studies in Nonlinear Dynamics and Econometrics

Quarterly Journal
April 1998, Volume 3, Number 1
The MIT Press

Studies in Nonlinear Dynamics and Econometrics (ISSN 1081-1826) is a quarterly journal published electronically on the Internet by The MIT Press, Cambridge, Massachusetts, 02142. Subscriptions and address changes should be addressed to MIT Press Journals, Five Cambridge Center, Cambridge, MA 02142; tel.: (617) 253-2889; fax: (617) 577-1545; e-mail: journals-orders@mit.edu. Subscription rates are: Individuals $40.00, Institutions $130.00. Canadians add 7% GST. Prices subject to change without notice.

The Decomposition of Economic Relationships by Time Scale Using Wavelets: Expenditure and Income

James B. Ramsey
Camille Lampart
Department of Economics
New York University
New York, USA

Abstract. *Economists have long known that time scale matters, in that the structure of decisions as to the relevant time horizon, degree of time aggregation, strength of relationship, and even the relevant variables differ by time scale. Unfortunately, until recently it was difficult to decompose economic time series into orthogonal time-scale components except for the short and long run, in which the former is dominated by noise. This paper uses wavelets to produce an orthogonal decomposition of some economic variables by time scale over six different time scales. The relationship of interest is the permanent income hypothesis. We confirm that time-scale decomposition is very important for analyzing economic relationships and that a number of anomalies previously noted in the literature are explained by these means. The analysis indicates the importance of recognizing variations in phase between variables when investigating the economic relationships.*

Keywords. wavelets, time scale, permanent-income hypothesis

Acknowledgments. The financial and institutional support of the C. V. Starr Center for Applied Economics is gratefully acknowledged. We would also like to thank the reviewers whose comments improved the exposition substantially.

1 Introduction

The concept of sequences of time scales delineating different modes of behavior, or different relationships between variables, is one that is a commonplace in the physical and biological sciences. In the physical sciences, the notion of time scale captures the notion of the distinction between "slow" and "fast" variables; the former are variables whose values are changing slowly relative to those of the latter. For example, in the analysis of turbulence, there is a cascade of time scales each of which determines a particular mode of behavior. In the biological sciences, time scales determine the relevant parameters and extent of evolutionary changes. And in economics the idea of time scale is recognized in terms of the concept of *time period*.

For decades, the idea of "time period" in economic analysis was enshrined in textbooks as the "short run" and the "long run." Classic text authors, including Marshall, Edgeworth, Schumpeter, Hicks, and others, realized that there were more time periods involved in economic decision making, but the pedagogical advantages of just two periods dominated the relevance of many periods. In any event, it was recognized early in the profession that the time period of analysis, or as we would now term the matter, the *time scale of the analysis*, is very important for determining those aspects of decision making that are relatively more important and those that are relatively less important. Similarly in economics, one can envisage a cascade of time scales within which different levels of decisions are being made, and different modes of behavior exist. Some decisions are taken with long horizons, others are taken with short horizons. The choice of time scale determines not only the length of the period over which one requires forecasts of future events, but the very choice of variables that are to be the focus of the decision maker's processing of information.

For example, to use the illustrations of the classical economists, a plant manager planning an extension of the plant, or ordering fuel, or adjusting the plant output rate, are all decisions that can be formulated in terms of different time scales of decision making, from decades to hours. In each case, the manager focuses on a particular time horizon for gathering information. Further, the manager will most likely take into consideration quite different variables in the different cases. In the plant-planning process, the key issues are the anticipated long-run conditions of demand for the firm's product, which depend in turn on the growth of the industry and the firm's relative position in the industry. For ordering fuel, the concern is only with current inventory stocks and the price of fuel over the next few months. And in the recovery decision, the manager is probably only concerned with overtime labor rates and the probability of breakdown from running machinery too intensely.

In this paper we investigate some of these issues empirically by using wavelet analysis that enables one to separate different time scales of variation in the data. We investigate the role of time scale in economic relationships in terms of the permanent-income hypothesis. Recently, the concept has been under empirical investigation, and there is some controversy about the nature of the relationship and the extent to which the relevant theory, or theories, are supported by the empirical evidence. We attempt to shed light on these issues by separating the empirical analysis of the relationship between variables into that between the variables after separation into time-scale components.

1.1 Consumption and Income

The analysis of the consumption function during the post-WWII period has been extensive. Friedman's original and seminal conception of the permanent-income hypothesis (Friedman 1957) was an attempt at that time to reconcile the Keynesian, linear, short-run income hypothesis with the obvious empirical finding that long-run (say averaged over decades) consumption was proportional to income. Friedman's solution essentially was to propose an "errors in the variables" model to reconcile the short- and long-run empirical results. Later, Friedman (1963) complained that his concept of the "horizon" had been misunderstood; we believe that the approach in this paper using time scales through the medium of wavelets returns to the essence of his original idea. Further work recognized the role of the interest rate in determining *permanent income*, and made distinctions between the consumption of different types of goods.

The next major wave of research interest grew out of the discovery that economic variables seemed to be dominated by trend and random-walk components, and the importance of the rational expectations hypothesis. In addition, a more careful distinction between consumption and expenditure was made in that it was recognized explicitly that what entered the utility function was the flow of services that the goods provided. As a consequence, there developed two parallel literatures, one on consumption that restricted its attention to nondurable goods and one that dealt with the timing of durable goods purchases; see for example Grossman and Laroque (1989). An early key article on the "consumption function" was that by Hall (1978), which stimulated a substantial literature by "confirming" the revised permanent-income hypothesis, which stated that the marginal utility of consumption and therefore consumption itself evolves according to a random walk with trend. A major implication was that only lagged consumption had any explanatory role to play for current consumption. In subsequent work, Flavin (1981) rejected the permanent-income hypothesis as postulated by Hall, but then Nelson (1987) was able to retrieve the hypothesis in part; the difference in results depended on how one detrended the data.

Deaton (1987) added to the debate by pointing out that while the permanent-income hypothesis in its new rational expectations guise with income as a random walk implied that nondurable consumption should be relatively volatile, the evidence indicated that nondurable consumption was relatively smooth. This is now termed the *Deaton paradox*. Quah (1990) rescued the hypothesis by returning to the original Friedman observation that one should split income into permanent and transitory components.

Campbell and Mankiw (1990) generalized the permanent-income hypothesis by generating a model that related expected changes in consumption to expected changes in income. They came to the conclusion that the ratio was as high as 50%. Further, they noted the empirical regularity that changes in the real interest rate did not explain changes in consumption. Christiano, Eichenbaum, and Marshall (1991) generated a model in continuous time in which consumption persistence was due to time aggregation or the effects of exogenous technology shocks. Using quarterly data, they found it difficult to distinguish between the two models, and weakly accepted the permanent-income hypothesis.

Molana (1991) concentrated on the role of wealth and interest rates. Molana modeled the contemporaneous variation in the data in terms of an error-corrections model. While the model provided additional reasons for the empirical failure of the simple random-walk model and reintroduced the role of the interest rate, the empirical results were implausible (for example, that the long-run wealth elasticity for consumption was two). In examining the empirical evidence, Campbell and Mankiw (1989) argued for a "rule of thumb model of consumer behavior" instead of the permanent-income hypothesis, and expanded that there was no observed role for interest rates.

In terms of long-run trends, Viard's analysis (1993) was pragmatic, and examined the impact of the post-1973 productivity slowdown, which implied a long-run decline in real income that in turn should have implied an increase in the saving rate. The empirical evidence was the opposite: the saving rate declined.

An alternative view of the decision-making process was that the relationships between variables was in fact between the frequency components of the variables. In Engle (1974), the idea was expressed that there might be several ranges of frequencies for consumption and income such that different coefficient values might relate consumption and income at different frequency ranges. Engle used band-spectrum estimation techniques to evaluate the relationship between consumption and income. He allowed for different relationships at low frequencies and high frequencies, but was able to continuously adjust the break between low and high frequencies. Engle's results were remarkably consistent in that the two regression results were virtually identical, and adjusting the break between low and high frequencies did not alter the results.

Recently, Corbae, Ouliaris, and Phillips (1991) returned to this question with tools developed since the 1974 Engle article. This article extended the Engle results in two important ways. First, the authors allowed for the trend in both consumption and income by recasting the analysis into a cointegration framework. Second, they derived a theoretical model that indicated that as the frequency of consumption and income rose, the marginal propensity to consume fell from its maximum of one at zero frequency (the [very] long-run marginal propensity to consume (MPC) is one because there is no bequest motive in the model and the representative consumer exhausts his wealth in the limit as time goes to infinity). Corbae, Ouliaris, and Phillips (1991) used seasonally adjusted quarterly per-capita data from 1948–1990. The results were that the zero-frequency MPC was about 0.73, less than one with very high probability, and that the constancy of MPC across higher frequencies was also rejected. However, the direct estimates yielded the result that the high-frequency MPC was larger than that for zero frequency. Redoing the experiment using first-differenced data yielded the more theoretically satisfactory result of lower MPCs at higher frequencies. In interpreting these results, the reader should recall that the model used by Corbae, Ouliaris, and Phillips (1991) was a simple representative agent model incorporating a quadratic utility function, which clearly could not be a reasonable representation except in a small region around the "bliss point."

In most of the papers that have been reviewed above and many others that we have not cited, the consumption/income relationship is regarded at most as a twofold distinction between permanent and transitory income. However, it is plausible that consumers have different time horizons for different consumption decisions. Consequently, it is productive to consider that the relationship between consumption and income may depend on the range of time scales involved, that is, different relationships, certainly different coefficient values, may apply over different time scales. In one sense, this idea is nothing more than a generalization of Friedman's original concept. Further, when examining macro data, we postulate that this idea of time scales affecting the qualitative nature of the relationship applies to both consumption functions and durable goods purchases.

In contrast, the standard representative agent model, say in the statement of the maximization of consumption of services over time, assumes that any unit of consumption that is being considered over any time horizon is the same: there is no difference in decision making for purchases in buying paper clips from buying a house; in short, there are no time scales, and there is only one mode of behavior. The distinction in this example is not on the relative size of the purchase, nor on the infrequency of purchase, but the time intervals over which the agent is operating, and the focus of the agent when making the decision. In the former case, the decisions are over a time scale of weeks, and decade changes in income are irrelevant. The latter is over a time scale of years to decades, and temporary fluctuations in income are ignored. While our analysis is empirical at the macro level, the finding that time scales do matter raises the interesting question as to the form of the individual's utility function and the coordination of decisions at various time scales. We believe that this provides an interesting and potentially fruitful approach to individual decision making.

However, there are some lessons that have been learned from these exercises. First, it is reasonably clear that the permanent-income hypothesis is correct in the idea that there is some difference between the MPC in the long run and the MPC at other time scales. Second, not enough is known about the mechanisms generating the data to formulate reliably a specific model such that inference can be reduced to the relatively simple process of estimation within the confines of a known model. Third, over the past 30–50 years, evidence has accumulated for nonstationarity in the data series that goes beyond the simple presumption of a trend, or of an integrated process. Fourth, one might well suspect that economic relationships may differ over different time scales, in that different time horizons are involved. At least one should be prepared to consider that relationships over decades may involve different levels of evaluation than decisions over the next 10 hours.

1.2 Major Implications for Research

Several strands of difficulty appear in this literature. The first has to do with nonstationarity of the data beyond that which might be incorporated in a unit root, or a trend. At both the theoretical and empirical levels, there are reasons to suspect that the consumption/income relationship might have evolved over the decades since the war. Long-run changes in the relative proportions of income spent on various consumption categories may affect the consumption-income relationship, as would the obvious effects caused by broad large-scale changes in demographics.

The next strand of problems involves the partially recognized need for analyzing relationships between variables conditioned on a given time scale of analysis. We have made the case that the choice of time scale will determine not only the time horizon involved in decision making, but also the degree of time aggregation, or averaging, and the choice of variables that enter the relationships.

While there is a core of agreement over some basics in the relationship, there is considerable disagreement over everything else. These disagreements have waxed and waned as different researchers have tried various and ingenious methods for examining the basic relationships. The conjunction of different theories, different data sets and different econometric technologies has probably inhibited our ability to learn more rapidly from the observed data. The only matter that is certain is that simple, noncomplex, stable relationships between consumption and income do not exist; otherwise, they would by now have been revealed and tested many times over.

We conclude this review of the literature by stating that the main requirements of an exploratory analysis involve the recognition that nonstationarity in addition to the so-called unit-root problem are relevant issues; that an allowance for effects that are localized in time is needed; that relationships may vary by time scale of analysis; and that, given the lack of knowledge about the underlying functional forms involved, nonparametric procedures will be required. This combination of conditions indicates that as an exploratory tool, an analysis using wavelets should prove useful. To this topic we now turn.

2 Wavelet Analysis: An Introduction

The discussion that follows is meant to give only an intuitive feel for wavelet analysis, and to stress the points at the heart of the examination of the statistical properties of the data; a more elaborate statement is contained in Ramsey and Lampart (1997) and references therein (see also Ramsey and Lampart [1998]). For a thorough review of the basics of wavelet analysis, Chui (1992) is an excellent reference, and Daubechies (1992) contains a detailed analysis of the mathematical properties of wavelets; Strang (1989) provides a succinct review of the basics. A useful recent article that helps to link wavelet analysis to more conventional time-series analysis is Priestley (1996), and a useful nontechnical review article is Rioul and Vetterli (1991). An important article that explores the use of the wavelet approach to the estimation of complex signals contaminated by noise is Donoho et al. (1995). Further, Brillinger (1994) develops some of the relevant distribution theory.

While there are many useful properties of wavelets, in this paper we are interested in only three major facets of wavelet analysis: the ability to handle nonstationary data, localization in time, and the resolution of the signal in terms of the time scale of analysis. The nonstationarity that we are concerned about is a broader notion than the existence of a mere unit-root process, in that some allowance should be made for variation in the processes over time as well as for time-local effects. The evaluation of the effects of time scale on the relationships between economic variables is central to our purpose.

Wavelet analysis has points of comparison and contrast to Fourier analysis. Recognizing both is important for understanding what wavelet analysis can bring to the examination of a data series. Both procedures involve the projection of a signal onto an orthonormal set of components, trigonometric in the case of Fourier series representations, "wavelets" (that are to be defined below) in the case of wavelet analysis. Fourier projections are most naturally defined for functions restricted to $L^2(0, 2\pi)$, because Fourier series have infinite energy, but finite power, when extended to being defined over the entire real line. Intuitively, this is the source of the difficulty, in that a single disturbance to a signal affects the analysis at all frequencies, and that a single disturbance in time is interpreted by Fourier analysis as an event of period T, where T is the length of the observed series. While Fourier analysis allows for the complex superposition of individual harmonics, or *waves*, the maintained hypothesis is that over any subsegment of the observed time series the precise same frequencies hold at the same amplitudes; the signal is *homogeneous over time*.

In contrast, the functions that are to be represented by wavelets are naturally defined within $L^2(R)$, and are not necessarily homogeneous over time. This means that the basis functions to be used must have narrow compact support so that the basis functions rapidly converge to zero as the index t approaches $\pm$ infinity; such basis functions are known as *wavelets*. This is in contrast to the trigonometric functions that have been traditionally associated with "waves." However, the wavelet approach brings its own difficulty in that functions that do not have narrow support have to be approximated by a sequence of functions that do have support. It is this difficulty that leads to the necessity for defining wavelets with respect to specific locations, and then considering a sequence of such functions, each indexed by a particular location. By stringing together a sequence of such wavelet functions, each localized to a particular position on the time axis, complex functions can be approximated.

There is another aspect of wavelets that produces much of their appeal, and that aspect derives from the "rescaling" capability of wavelets. Instead of considering a single sequence of basis functions $g(t, k)$, where k denotes a sequence of positions about which the function $g(\cdot)$ is centered, we consider a double sequence of functions:

$$g(t) = \frac{1}{\sqrt{s}} g\left(\frac{t-k}{s}\right), \tag{1}$$

where s is a sequence of scales. The term $\frac{1}{\sqrt{s}}$ maintains the norm of $g(\cdot)$ at one. The basis function $g(\cdot)$ is centered at k with a scale of s; that is, the energy of $g(\cdot)$ is concentrated in a neighborhood of k the size of which is proportional to s. Essentially, as s is increased, the length of the support of $g(\cdot)$ in terms of t increases. For example, if the support of $g(\cdot)$ for $s = 1$ is $[-d, d]$ when $k = 0$, then the effect of s is to broaden the support to $[-ds, ds]$. In effect, the rescaling characteristic of wavelets in the time domain is similar to the rescaling of frequencies in the frequency domain for Fourier analysis. The process is also known as *integral dilation*. In Fourier analysis, all frequencies are generated by rescaling, or dilating, a single fundamental frequency. If ω_0 is the fundamental frequency, then the function to be approximated is projected onto a sequence of expressions of the form:

$$\{e^{-is\omega_0 t}\}. \tag{2}$$

In short, 2π-periodic functions are representable by integral dilations of the single function $e^{-i\omega_0 t}$.

Dilation in the time domain has advantages in addition to those stemming from frequency dilations. In the latter case, one is considering projecting the entire signal onto ever lower frequencies. In the former case with wavelets, a localized component of the signal is projected onto an ever broader base. The two important aspects here are that the projection is for a *local* component of the signal, and that each projection is onto a wavelet whose narrow compact support is a function of s. In a sense, the choice of a dilation (or scale) level indicates the size of the "packets" used to represent the signal. An important aspect of the wavelet approach is that the sequence of basis functions is orthogonal across scales, so that any signal can be easily decomposed in terms of its constituent scales. The view using wavelets is that the total variation of the data centered at any given point in time is obtained by adding together the components extracted at each of the admissible scales. A broad-support wavelet yields information on signal variations on a large scale; a small-support wavelet yields information on signal variations on a small scale. One way to visualize the process is in terms of maps, in that a large-scale map gives the broad picture without details and a small-scale map fills in the details.

An alternative view is that obtained by aggregating the data to various levels using a windowing technique, which can in turn be interpreted as a projection onto the vector $\{1/N, 1/N, \ldots, 1/N\}$, where there are N observations selected to be in each projection, $N = \{2, 3, 4, 5 \ldots\}$. However, by taking this view, one soon recognizes the nonoptimality of such a sequence of projections, in that at the very least orthogonality of the projections has been lost, not to mention that one does not have a basis. Consequently, the wavelet approach potentially adds to our understanding of the relationships between variables by enabling us to separate by orthogonal decomposition the relevant time scales that may be involved.

There is an understandable temptation to regard time scale and frequency as similar concepts. However, as indicated by Priestley (1996), there is only an intuitive and indirect connection between frequency and the scale of the analysis. It is only true in the simplest of cases that large-scale wavelets are associated with low frequencies, in that the detection of low frequencies requires components with wide support in the time domain. Correspondingly, the idea that the analysis of high frequencies requires a high sampling rate which is provided by components that have narrow support is also a naive interpretation. At each scale the strength of the projection determines the magnitude of the corresponding coefficient. In essence, one is extracting at the chosen scale those components of the data that resonate most strongly at that scale. The link between scale and frequency is only interpretable when the data are stationary; the wavelet projections are still meaningful even if there are no oscillations in the data.

Consider for example a given scale, say 2^3 (or 8) months, and that the signal contains components that are at that scale. It might still be true that the power of the projections depends in a cyclic manner on the value of k, the index of position; for example, the power of the projections indexed by k might oscillate with a period of 10 years, or indeed any period. More importantly, this intuitive idea of "packets of information" that arises out of the dilation procedure can be used and interpreted, even when a signal cannot usefully be regarded as a superposition of trigonometric components. Priestley (1996) illustrates this idea in the context of a signal with a fixed frequency but a time-varying amplitude; Fourier analysis, of course, incorrectly detects power at all frequencies.

We can consider two ways of examining a given signal. At a fixed point in time, say k, we can examine how the localized signal projects onto wavelets of varying length of support. In effect this procedure examines the signal at the given location for variations in the strength of projections onto wavelets of various scales, or dilations. The other view is to examine how the signal projections vary by position as indicated by the value of k, but at a given fixed scale, or dilation, s.

The literature on wavelets is growing rapidly, but so far in economics the development has been sparse. Goffe (1994) illustrated the application of wavelets to nonstationary data, and Gilbert (1995) attempted, using quarterly data, to examine macro relationships for rapid regime shifts. In the two papers by Ramsey and Zhang (1995, 1996), an effort was made to use waveform dictionaries to analyze financial data; waveform dictionaries provide a generalization of both wavelets and Fourier analysis. The emphasis in these two papers was on examining the time-frequency (Wigner) distributions. Ramsey, Usikov, and Zaslavsky (1995) pursued an approach that was common in the earlier wavelet literature by querying the statistical self-similarity of financial data. The detection of discontinuities and the occurrence of sharp cusps were explored by Truong and Patil (1991) and by Wang (1995). Further work on financial data and the role of fractional differencing was explored by Greenblatt (1996), Hog (1996), and Jensen (1996).

One source of confusion in understanding the role of wavelets is that the wavelet approach can be used in many different ways; see for example the discussion in Rioul and Vetterli (1991). Some of these procedures, or aspects of wavelet analysis, might be useful in economic and financial analyses, but not necessarily all. For example, the early emphasis in using wavelets to explore ideas of self-similarity was not very productive, nor did that analysis lead to useful insights into economic and financial mechanisms; see for example Ramsey, Usikov, and Zaslavsky (1995). Priestley (1996) indicated some of the difficulties involved in shifting from models that are essentially deterministic with low levels of noise to the stochastic processes that are more common in the economics and finance literature; see also the discussion by Donoho et al. (1995).

2.1 Technical Details

There are a variety of functions that have been developed for use as the fundamental wavelet that is to be dilated by s and translated by k; for example, the Haar, the Daublet, the Symmlet, etc. Depending on normalization rules, there are two types of wavelets within a given function/family, *father* and *mother wavelets*.

$$\begin{aligned}\Phi_{j,k} &= 2^{-\frac{j}{2}}\Phi\left(\frac{t-2^jk}{2^j}\right), \\ \Psi_{j,k} &= 2^{-\frac{j}{2}}\Psi\left(\frac{t-2^jk}{2^j}\right),\end{aligned} \tag{3}$$

$$\textit{Father wavelets:} \quad \int \Phi(t)dt = 1, \text{ and}$$
$$\textit{Mother wavelets:} \quad \int \Psi(t)dt = 0.$$

Father wavelets are used for the "lowest frequency" smooth components; those requiring wavelets with the widest support and mother wavelets are used for the "higher frequency" detail components. Father wavelets are used for the "trend components" and mother wavelets are used for all the deviations from trend. While a sequence of mother wavelets is used to represent a function, only one father wavelet is used.

Any function $f(t)$ to be represented by a wavelet analysis can be built up as a sequence of projections onto father and mother wavelets indexed by both $\{k\}$, $k = \{0, 1, 2, \ldots\}$ and by $\{s\} = 2^j$, $\{j = 1, 2, 3 \ldots\}$. In actual data analysis using discretely sampled data, it is necessary to create a lattice over which the calculations will be made. Mathematically, it is convenient to use a dyadic expansion as illustrated in Equation 3.

The coefficients in the expansion are given by the projections:

$$\begin{aligned} s_{J,k} &= \int f(t)\Phi_{J,k}(t)dt, \\ d_{j,k} &= \int f(t)\Psi_{j,k}(t)dt, \text{ and} \\ j &= 1, 2, \ldots, J, \end{aligned} \tag{4}$$

where J is the maximum scale sustainable by the number of data points. The representation of the signal $f(t)$ can now be given by:

$$\begin{aligned} f(t) &= \sum_k s_{J,k}\Phi_{J,k}(t) + \sum_k d_{J,k}\Psi_{J,k}(t) \\ & \sum_k d_{J-1,k}\Psi_{J-1,k}(t) + \cdots \\ & \cdots + \sum_k d_{1,k}\Psi_{1,k}(t). \end{aligned} \tag{5}$$

We can represent the approximation in a more revealing manner for our purposes:

$$\begin{aligned} f(t) &= S_J + D_J + D_{J-1} + \cdots D_j + \cdots D_1 \\ \text{where } S_J &= \sum_k s_{J,k}\Phi_{J,k}(t), \\ \text{and } D_j &= \sum_k d_{j,k}\Psi_{j.k}(t), \ j = 1, 2, \ldots, J. \end{aligned} \tag{6}$$

When n, the number of observations, is divisible by 2^J, then the number of coefficients of each type is given by the following:

- at the finest scale, 2^1: $\frac{n}{2}$ coefficients $d_{1,k}$;
- at the next scale, 2^2: $\frac{n}{2^2}$ coefficients $d_{2,k}$;
- at the coarsest scale, 2^J: $\frac{n}{2^J}$ coefficients $d_{J,k}$;
- at the coarsest scale, 2^J: $\frac{n}{2^J}$ coefficients $s_{J,k}$;

so that:

$$n = \frac{n}{2} + \frac{n}{4} + \cdots + \frac{n}{2^{J-1}} + \frac{n}{2^J} + \frac{n}{2^J}. \tag{7}$$

For the discrete wavelet transform we can summarize the string of coefficients by:

$$w = \begin{pmatrix} s_J \\ d_J \\ d_{J-1} \\ \vdots \\ d_1 \end{pmatrix} = Wf, \tag{8}$$

where W is an orthogonal matrix that summarizes the transform and f is the n-dimensional vector of discrete observations on the function. However, most of the coefficients of w are zero, or very close to zero. The matrix w is sparse, because the wavelet transform is "information compressing." The product of the wavelet associated with each coefficient and its coefficient is termed an *atom*. The vector defined by the product of the elements of each row of w with the corresponding wavelets creates a *crystal*. A crystal is made up of atoms along a lattice, either by location for a given scale (as in our case), or by scale for a given location. A group of atoms that is not distributed along a lattice is termed a *molecule*; for example, the atoms that are collected in the top 100 atoms by size of modulus of coefficient would create a molecule.

We can define the multiresolution decomposition of a signal by specifying:

$$S_J \text{ coarsest scale, and}$$

$$S_{J-1} = S_J + D_J.$$

In general,

$$S_{j-1} = S_j + D_j,$$

where $\{S_J, S_{J-1}, \ldots, S_1\}$ is a sequence of multiresolution *approximations* of the function $f(t)$, at ever increasing levels of refinement. The corresponding multiresolution *decomposition* of $f(t)$ is given by:

$$\{S_J, D_J, D_{J-1}, \ldots D_j, \ldots, D_1\}.$$

The sequence of terms: $S_J, D_J, D_{J-1}, \ldots, D_1$ represents a set of *orthogonal* signal components that provide representations of the signal at resolutions 1 to J; each D_{J-k} provides the orthogonal increment to the representation of the function $f(t)$ at the scale, or resolution 2^{J-k}.

With each set of macroeconomic data, the idea is to examine in some detail the relationship between the variables when the variation in each variable has been restricted to a specific scale. Instead of looking at the relationship between expenditure and income "averaged" over all time scales, we examine the relationship at each time scale separately. If $C[S_J]_t$ and $Y[S_J]_t$, and $C[D_j]_t$ and $Y[D_j]_t$ represent consumption and income at each scale as determined by the wavelet decomposition, recall from Equation 6 that we can estimate a sequence of regressions using:

$$\begin{aligned} C[S_J]_t &= \alpha + \beta Y[S_J]_t + \varepsilon_t, \text{ or} \\ C[D_j]_t &= \alpha + \beta Y[D_j]_t + \varepsilon_t, \; j = 1, 2, \ldots, J-1. \end{aligned}$$

3 Data Used in the Analysis

The details of the data used and their sources are listed in Table 1. Monthly data, seasonally unadjusted where available, were used to have a sufficiently high sampling rate to carry out the wavelet analysis, to introduce as little preprocessing of the data as possible, and to allow for the effects of short-term dynamics on the results. In all cases, of course, we observed expenditure, not consumption per se. The real interest rate was calculated by subtracting an annual inflation rate from the U.S. Treasury bill rate for one-year bonds. The expected inflation rate at time t was proxied by the actual inflation rate from time t to time $t+1$. The inflation rate was based on the implicit deflator for personal consumption expenditures on durables.

4 Methods Used

Our main interests are in the reconstructions of the time series by crystals, that is, using a sequence of atoms distributed by location for a given choice of scale, and in the examination of the relations between crystals across variables. The process is exploratory. Our objective is to examine the extent to which an allowance for different effects by scale and for variations in the relationships over time lead to insight into the total variation of the signal over time. Our conclusion is that these topics merit further study and the development of procedures that are designed for the classes of statistical requirements indicated by our preliminary analysis.

We used the Wavelets package produced by StatSci of MathSoft that was written by A. Bruce and H. Gao (1994). We choose as the basic wavelet the Symmlet, which is as designated as S12. This wavelet is a

Table 1
The data used in the analysis.

Name	Dates	Units	Source	Code
Real disposable personal income	1960:5–1994:4	Billions 1987$	Citibase	GMYDQ
Real personal consumption expenditures (PCE)	1960:5–1994:4	Billions 1987$	Citibase	GMCQ
Real PCE—durables	1960:5–1994:4	Billions 1987$	Citibase	GMCDQ
Real PCE—(durable) furniture, etc.	1960:5–1994:4	Billions 1987$	Citibase	GMCFQ
Real PCE—(durable) motor vehicles, etc.	1960:5–1994:4	Billions 1987$	Citibase	GMCAQ
Real PCE—(durable) other durables	1960:5–1994:4	Billions 1987$	Citibase	GMCOQ
Real PCE—nondurables and services	1960:5–1994:4	Billions 1987$	Citibase	GMCSNQ
U.S. Treasury bill, maturity one year	1960:5–1993:12	% per annum	Citibase	FYGT1
PCE—implicit deflator (durables)	1960:5–1994:4	1987 = 100	Citibase	GMDCD

compromise between competing requirements. The S12 is nearly symmetric, is intermediate in support length with 11 months, has 5 vanishing moments, and is twice differential. This choice of wavelet is an intermediate choice in that it has reasonably narrow, compact support, is fairly smooth, is nearly symmetric, and has a moderate degree of flexibility. We experimented with alternative choices of scaling functions and wavelets, but the qualitative results were very robust to such changes and the initial choice of wavelet seemed to be the best on balance. We used the boundary condition "infinite" to reflect the fact that the data contain a "trend" component, and are best considered by extending either side of the observed range. Alternative boundary conditions allow for estimation of a function over fixed interval lengths that are periodic, or that can be extended by specifying an approximating polynomial.

The idea is to examine in some detail the relationship between consumption, or expenditure, and income, when the variation in each variable has been restricted to a specific scale. Instead of looking at the relationship between consumption and income "averaged" over all time scales, we examine the relationship at each time scale separately.

5 Consumption-Income Relationship

Tables 2, 3, and 4 show the coefficient estimates from running a sequence of least-squares regressions of consumption on disposable income using the data described in the previous section. Figures 1, 2, and 3 illustrate some of the results obtained.

Figure 1 portrays the time paths of durable and nondurable goods expenditure at the S5 level; that is, this level includes the long-term drift in the data with respect to the longest time scale available for these data, $2^6 = 64$ months, using the father wavelet coupled with the deviations induced at the scale of 2^6 months using a mother wavelet. Figures 2 and 3 illustrate the relationship between total expenditure and income at scales varying from $2^5 = 32$ months to 2 months using mother wavelets. Recognize that at each scale level one is seeing the isolated effect of the given scale; rather at each scale, the graphs indicate the relationship between total expenditure and income where the variation in both variables has been restricted to the indicated scale. The graphs indicate clearly that except for the highest scale, there is only a weak relationship between the variables; indeed, below D4 there seems to be little of significance. The regression results reveal somewhat more structure, but nevertheless confirm that below S5, the relationship is very noisy.

A review of Table 2 indicates that the degree of fit of the regression of total consumption on income falls as we move to shorter time scales, and the slope coefficient declines to zero as the time scale of analysis declines. There is an exception to this pattern in that at the very lowest level of time-scale analysis, the degree of fit is greater than the previous level and the slope coefficient, which is statistically significant at even high confidence levels, and is commensurate with that obtained at a level between D2 and D3. The intercept terms are very small and operationally insignificant, even where formally significant in a statistical sense.

While we examined the residuals from the regressions for approximate "i.i.d. Gaussianity" and obtained surprisingly good results, it is nevertheless true that the distribution theory for regressions within the context of our analysis requires considerable further investigation. Hence, we regard the formal statements given below concerning the size of standard errors of estimate as indicative and subject to further and more detailed analysis. As we stated earlier, the current results are purely exploratory, and are meant to stimulate interest in the general approach.

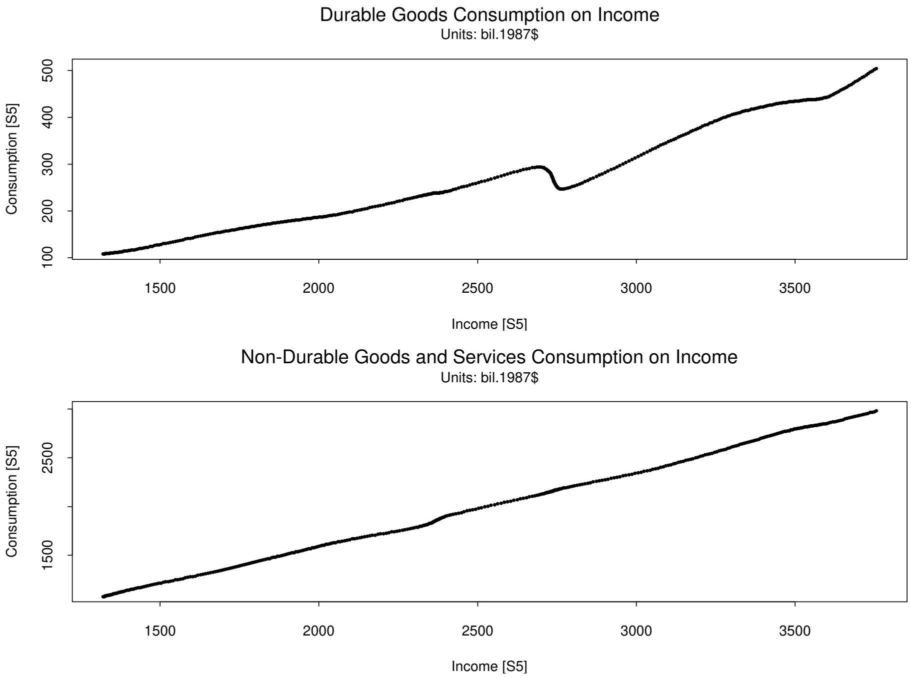

Figure 1
Plots of consumption on income using crystal (S5): 1960:6–1994:4.

Tables 3 and 4 portray these relationships by time scale in terms of a decomposition of total consumption into nondurable goods and services in one table and durable goods in the other table. The general conclusions that were obtained with respect to total consumption are confirmed with respect to the components. In each case the intercept term is insignificant in effect and is statistically insignificant for all regressions past the "trend fit." The slope coefficient for nondurable goods and services has a maximum of 0.78, but the pattern for degree of regression fit and the size of the slope coefficient across time scales reflects the same pattern observed for the total consumption variable.

The pattern of results for durable goods is similar to that for nondurable goods, but with some interesting differences. As expected, the slope coefficient for durable goods is much lower than that for nondurable goods at all time scales. However; there is the interesting difference that the slope coefficient at the highest time scale is less than that at the next level down. Or to express the matter another way, the marginal propensity to consume for durable goods has a maximum value at a time scale of $2^5 = 32$ months, which may well be plausible in that cars and household appliances dominate the durable goods index and a time horizon of 3 years seems in this context to be reasonable. Further; examining Figure 1 for durable goods, there is clear evidence of a shift in the relationship between consumption and disposable income. Although not directly observable from the figure, the shift occurred over the period of January 1979 through August 1981; the slope that prevailed before this extraordinary movement was regained by the end of the 1980s. The trend for nondurables and services does not exhibit similar behavior.

In the discussion above, we mention that the plot for D5 is very different. The explanation is to be found in Figure 4. This figure shows the time-series plots of both consumption and income superimposed on each other. The interesting aspect is that from about 1976 on, variations in income and consumption at this time scale were in phase. Before this period, income and consumption were out of phase (although moving into phase). We also note that at this time scale, over the period during which the variables were in phase, consumption was less volatile than income. The reverse was true when the variables were out of phase. In Figure 5 we relate these events to the official NBER business-cycle dates for income and consumption; the NBER dates are shown in Table 9. It is interesting to note that at this time scale at least, it is consumption that is in phase with the business-cycle peaks and troughs, and income is in phase only after it has become in

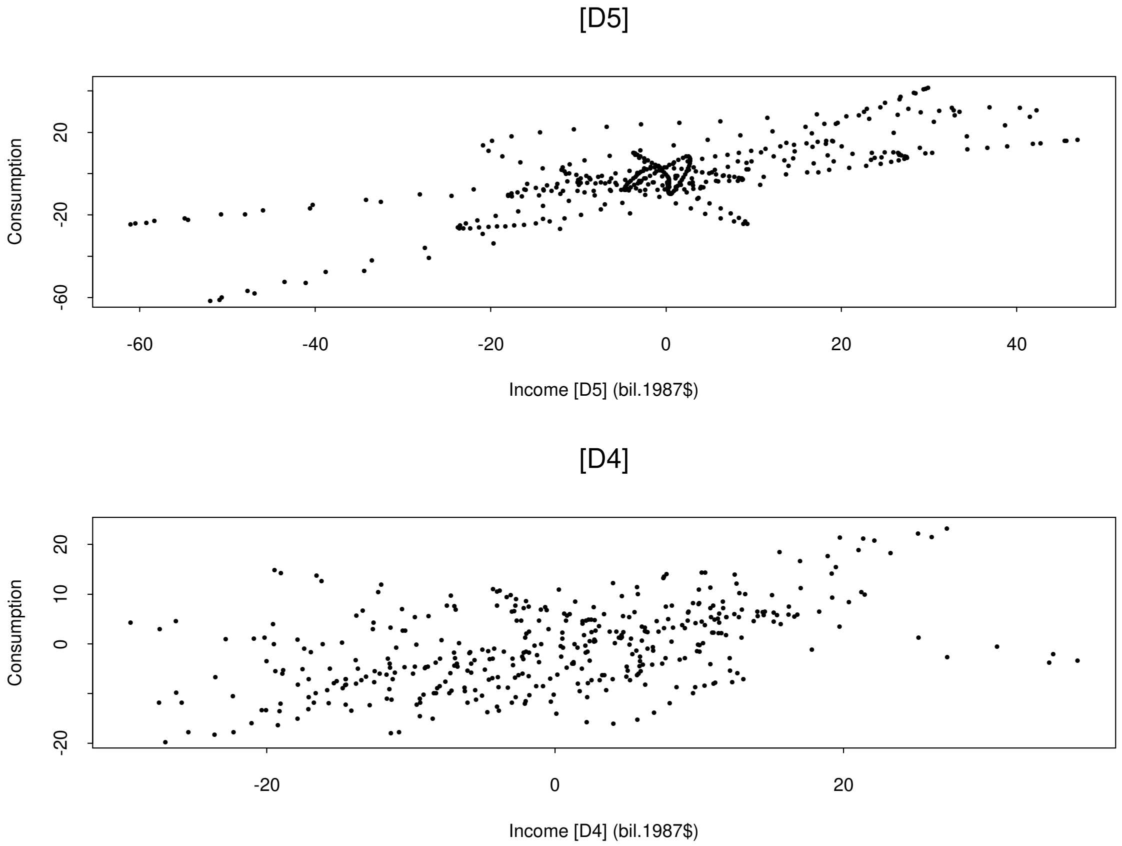

Figure 2
Plots of total consumption on income using crystals (D5) and (D4): 1960:5–1994:4.

phase with consumption. A closer examination of the figure indicates that there were substantial oscillations in income and consumption that were not associated with an NBER-declared cyclic peak or trough.

The results obtained above introduce an important generalization to modeling relationships; namely, the need to allow for variations in the timing of relationships. Thus, in this example, we might need to consider:

$$C_t = \beta Y_{t-d(X)} + \varepsilon_t,$$

where d(X) is a variable lag function that might well indicate a lead relationship for suitable values of X. The variable(s) represented by X could include institutional factors, variations in anticipated price changes, or interest-rate changes. The idea is that besides having to model the relationship between the levels of variables, we may also have to consider the timing relationship as an added component. For example, during the early period following the first oil-price shock in 1973, the then Federal Energy Office announced that in three weeks the price of gasoline would be allowed to rise; supplies fell dramatically and almost instantaneously, and demand rose as precipitously, so that excess demand was phenomenally high and the rate of change was swift. Both sides of the market were trying to rearrange the *timing* of their market actions; in this case both were also wishing to change their level of activity. Further, anticipated changes in interest rates or credit constraints can lead to a change in the timing of relationships.

5.1 Decomposition of Durable Goods Results by Category and the Role of Interest Rates

As noted above, the relationship between consumption and income shown in Figure 1 indicates a dramatic shift in the relationship beginning in about 1979. This observation is worth exploring further, and in addition it would be productive to examine durable goods expenditure as a function of the real interest rate, notwithstanding the negative results in the current literature. Tables 5, 6, and 7 summarize the regression

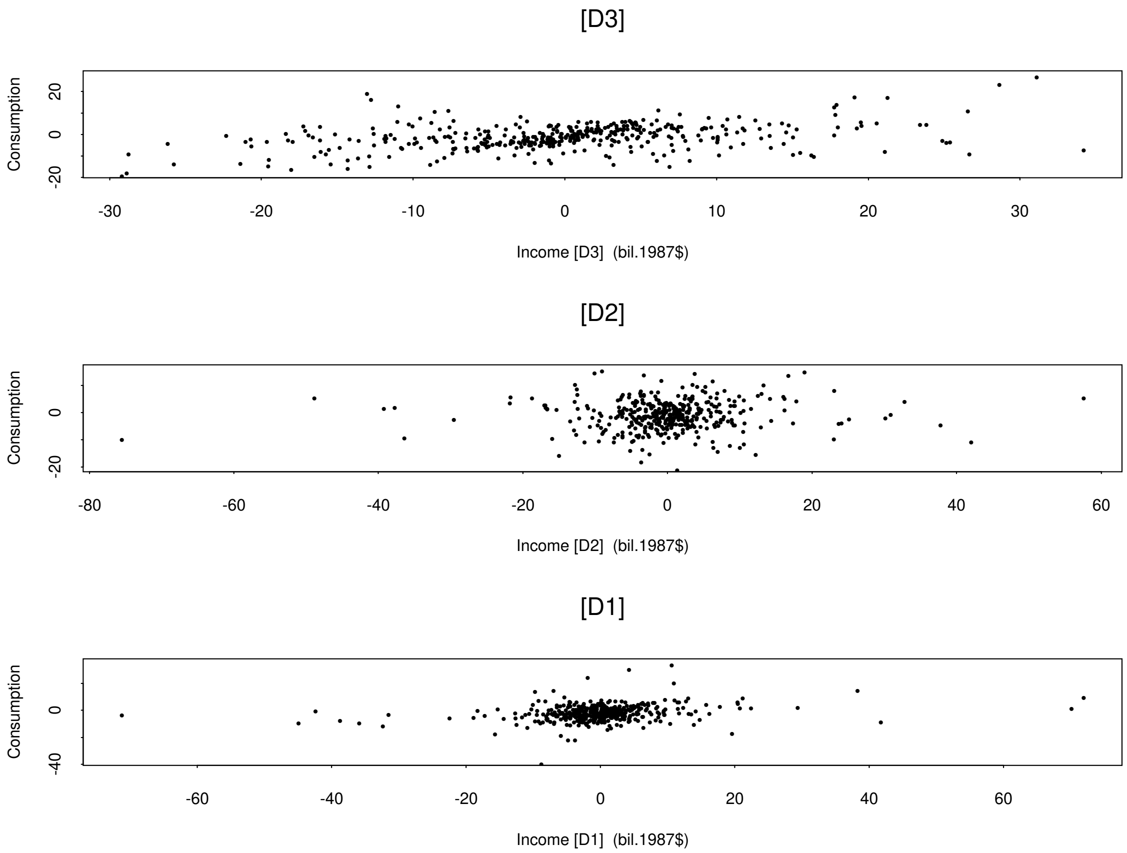

Figure 3
Plots of total consumptionon income using crystals(D3), (D2), and (D1): 1960:5–1994:4.

Table 2
Regressions of total consumptionon income for individualcrystals.

Crystal	Inter cept (standar d err or)	Slope Coefficient (standar d err or)	R^2
S5	−63.448 (6.206)	0.934 (0.002)	0.997
D5	0.217 (0.539)	0.669 (0.029)	0.573
D4	−0.032 (0.355)	0.338 (0.030)	0.239
D3	0.029 (0.272)	0.224 (0.028)	0.139
D2	0.0173 (0.270)	0.034 (0.027)	0.004
D1	0.001 (0.297)	0.143 (0.029)	0.057

results. The three subcategories of durable goods consumption are furniture and household items, automobiles and auto parts, and the category "other," which includes jewelry, boats, and sports equipment. Summary statistics on the subseries are shown in Table 8.

The main issue to be addressed is whether the shift in the relationship between income and consumption can be explained by the variation in the real interest rate. A corollary investigation is whether real interest rates play a more important and longer term role in the demand for durable goods as consumer theory would indicate.

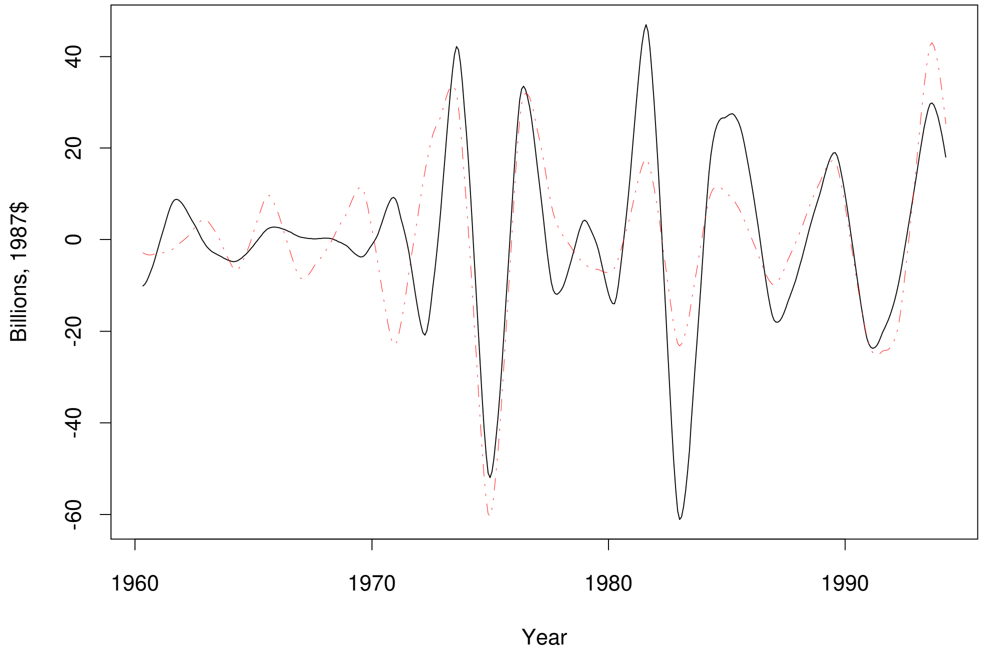

Figure 4
Time-series plots of total consumption and income using crystal (D5): 1960:5–1994:4. Income—solid line, consumption—dashed line.

Table 3
Regressions of nondurable goods and services consumption on income for individual crystals.

Crystal	Intercept (standard error)	Slope Coefficient (standard error)	R^2
S5	44.72 (3.253)	0.783 (0.001)	0.999
D5	0.195 (0.316)	0.392 (0.017)	0.572
D4	−0.032 (0.249)	0.203 (0.021)	0.188
D3	−0.00 (0.128)	0.128 (0.017)	0.120
D2	0.019 (0.142)	0.020 (0.014)	0.005
D1	−0.01 (0.170)	0.085 (0.016)	0.062

Before beginning the main investigation, we observe that the qualitative properties that were observed for all the categories combined are reflected in the three subcategories. That is, the degree of relationship as measured by R^2 declines with the reduction in scale. The strength of the effect of income on consumption as measured by the magnitude of the regression coefficient also declines with the reduction in scale; all intercept terms are statistically indistinguishable from zero, below the S5 scale. At the S5 scale the intercept terms are both statistically significant and substantial in size relative to the respective mean values. There is some indication that the relative importance of furniture and equipment is increasing relative to motor vehicles and parts.

Figure 6 shows the time-series plots at the S5 level of the three subcategories of durable consumption, furniture, autos, and other durables with straight-line approximations of the pre-1976 trend superimposed as a

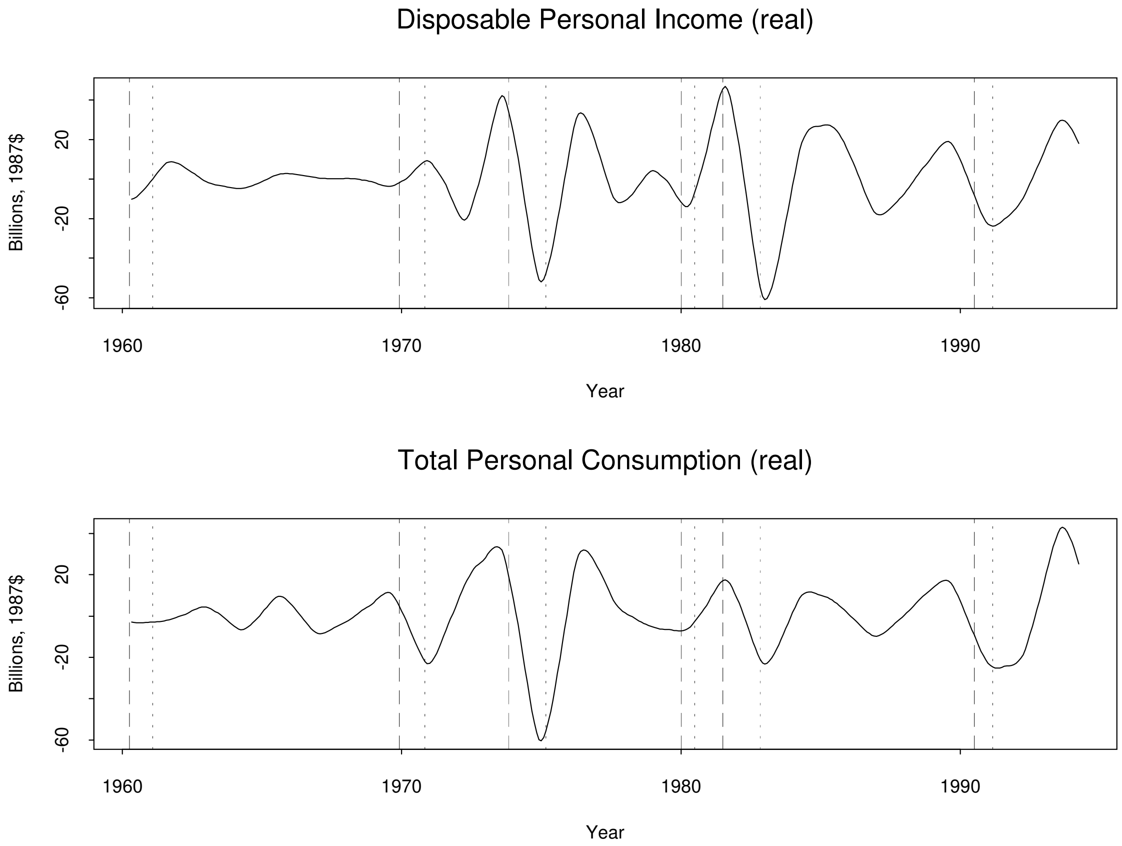

Figure 5
Real disposable income (D5), real personal consumption (D5), and NBER business cycle dates: 1960:5–1994:4. Troughs—dotted lines; peaks—dashed lines.

Table 4
Regressions of durable goods consumption on income for individual crystals.

Crystal	Intercept (standard error)	Slope Coefficient (standard error)	R^2
S5	−108.164 (3.754)	0.152 (0.001)	0.965
D5	0.021 (0.284)	0.277 (0.015)	0.453
D4	−0.001 (0.153)	0.135 (0.013)	0.211
D3	0.029 (0.190)	0.096 (0.019)	0.057
D2	−0.001 (0.213)	0.014 (0.021)	0.001
D1	0.002 (0.230)	0.058 (0.022)	0.016

Table 5
Regressions of furniture and household equipment on disposable income, the real interest rate, and an interaction term for individual crystals.

	Intercept (standard error)	Income (standard error)	Interest Rate (standard error)	Interaction (standard error)	R^2
S5	−31.3	0.054	−11.65	0.003	0.940
	5.15	0.002	1.76	0.001	
D5	−0.072	0.064	−0.149	−0.000	0.528
	0.056	0.003	0.071	0.002	
D4	−0.020	0.014	−0.585	0.011	0.166
	0.048	0.004	0.075	0.006	
D3	−0.003	0.015	0.275	0.001	0.143
	0.027	0.003	0.048	0.003	
D2	−0.003	0.009	0.171	0.006	0.058
	0.025	0.002	0.052	0.004	
D1	0.002	0.008	−0.071	0.001	0.027
	0.026	0.003	0.087	0.008	

Table 6
Regressions of motor vehicles and parts on disposable income, real interest rate, and an interaction term for individual crystals.

	Intercept (standard error)	Income (standard error)	Interest Rate (standard error)	Interaction (standard error)	R^2
S5	13.02	0.048	−16.52	0.005	0.954
	4.41	0.002	1.50	0.001	
D5	−0.188	0.180	0.152	0.030	0.289
	0.266	0.015	0.339	0.010	
D4	−0.001	0.097	−0.151	0.009	0.211
	0.111	0.010	0.175	0.013	
D3	−0.007	0.063	0.565	0.072	0.074
	0.177	0.018	0.314	0.020	
D2	−0.024	0.008	−0.925	0.032	0.016
	0.195	0.019	0.415	0.035	
D1	−0.014	0.036	−1.142	−0.007	0.015
	0.217	0.022	0.739	−0.087	

Table 7
Regressions of other durables on income, the real interest rate, and an interaction term for individual crystals.

	Intercept (standard error)	Income (standard error)	Interest Rate (standard error)	Interaction (standard error)	R^2
S5	−19.13	0.028	−2.82	0.001	0.993
	0.80	0.000	0.27	0.000	
D5	−0.018	0.026	0.242	0.001	0.438
	0.033	0.002	0.042	0.001	
D4	−0.005	0.017	−0.124	0.016	0.104
	0.036	0.003	0.057	0.004	
D3	−0.002	0.007	0.044	0.003	0.026
	0.025	0.003	0.044	0.003	
D2	−0.004	−0.004	0.104	0.003	0.011
	0.032	0.003	0.067	0.006	
D1	0.001	0.008	−0.462	0.025	0.068
	0.031	0.003	0.104	0.012	

Table 8
Summary statistics for durable goods consumption and the real interest rate.

Index	Mean	Median	Standard Deviation
Durable goods consumption: total	274.5	256.1	113.3
Durable goods consumption: furniture, etc.	99.6	90.45	40.6
Durable goods consumption: motor vehicles, etc.	126.0	118.3	45.7
Durable goods consumption: other durables	48.9	50.4	21.3
Real interest rate	2.86	3.67	3.11

Table 9
NBER business cycle dates.

Trough	Peak
	1960:4
1961:2	1969:12
1970:11	1973:11
1975:3	1980:1
1980:7	1981:7
1982:11	1990:7
1991:3	

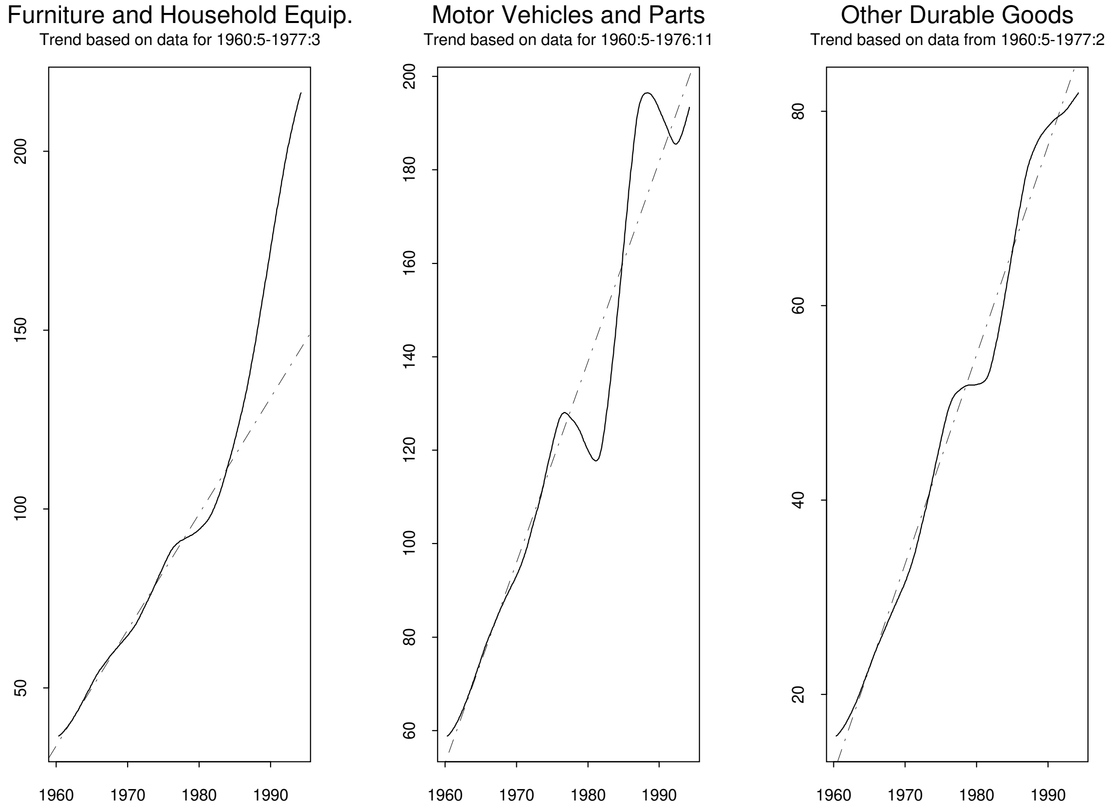

Figure 6
The (S5) crystals for components of personal consumption expenditure on durable goods with trend lines (dotted): 1960:5–1994:4.

visual reference for the shift in trend after 1979. One notices immediately that the three categories behave in different ways after 1979. After a minor slackening in consumption, furniture increases at a faster pace than before. Autos are the most affected, and account for much of the dip shown in the total durable goods index in Figure 1. The "other" category has a more moderate fluctuation.

One common reaction to observing this situation is to rationalize the result as a reaction to an oil price shock. We are unpersuaded that this is a reasonable first line of attack, because the timing of the change is apparently anticipatory of the actual oil price rise. Further, the far greater price rise in real terms that occurred in November 1973 that was even more of an unanticipated shock did not produce a similar reaction to automobile purchases. More convincing is the observation that net of our allowance for the interest rate effect to be discussed below, there is very little evidence for any special "shock."

The model that was used to investigate the potential relevance of interest rates at each scale level was a linear model that allowed for an interaction term between interest rates and income. If we define consumption by C_{it} for each of the three categories of consumption, Y_t as real disposable income, and r_t as the real interest rate, the regression model assuming an error term e_{it} that is approximately Gaussian in distribution is:

$$\begin{aligned} C_{it} &= \alpha_{0,i} + \alpha_{1,i} Y_t + \alpha_{2,i} r_t + \alpha_{3,i} Y_t r_t + e_{it}, \\ i &= 1, 2, 3. \end{aligned} \tag{9}$$

This model allows for an interaction term between income and the interest rate. By allowing for the effect of the interest rate to depend in part on the level of income, we were able to allow for the "interest rate effect" to be state dependent.

Analysis of the residuals indicated that the assumption of approximate Gaussianity without significant evidence for autocorrelation was justified.

Before examining the interest rate coefficients themselves, the pattern of coefficient values for income are interesting in the case of motor vehicles and parts. The size of the income coefficient for the scale of $2^5 = 32$ months is nearly four times greater than for the trend levels, and about twice for the next lower scale level. Furniture and household equipment exhibits this "business cycle" scale concentration to a very limited extent, and the "other" category, like nondurable goods, illustrates a monotonic decline in income coefficient values. Consequently, it appears that the anomalous coefficient behavior for durable goods as a whole is explained almost entirely by motor vehicles and parts. This would make some intuitive sense, in that one can easily believe that relative to the other components, automobile consumption is likely best analyzed on a three-year time scale. However, the age of the car fleet is now approximately seven years and has been at least six for some time. An interesting question is how these relationships change in response to variations in the age of the car fleet.

The interest rate effects, reported in Tables 5, 6, and 7, are revealing. First of all, the interest rate effect at the highest scale is strongest for automobiles, considerably less for furniture, and least for "other," even after allowing for the differences in units of measurement. Not only is the interest rate very highly significant statistically, its effects on the levels of consumption at the top scales are considerable. In all cases the interest rate effect falls off dramatically as the scale is reduced from 64 months, and in most cases the interest coefficient becomes insignificant; this is particularly true for motor vehicles and parts.

Furniture and household equipment provides an interesting contrast. The interest rate is statistically significant with a negative sign to the D4 level, or down to a scale of 16 months. However, for variations within a year, using time scales of four and eight months, the interest rate effect is seemingly positive and significant. One might speculate that these results represent timing effects in that with interest rates rising, customers are able to buy at the old interest rate when suppliers in anticipation of a demand downturn attempt to decrease inventory levels.

The only real puzzle is provided by the regressions for "other goods," in which there is a strong positive interest rate effect at business-cycle time scales. The remaining interest rate coefficients are statistically insignificant. This result could be interpreted as a substitution effect in which the demand for furniture and motor vehicles rises with a lowering of interest rates so that there is a corresponding decrease in the demand for "other goods."

If we examine the role of the interest rate at *all* time scales using S4, that is, we examine the relationships at all time scales up to 32 months, then for every category the interest rate effect is statistically insignificant; it is never above a t-ratio of 1.4. These results help to explain the nonsignificance of the real interest rate variable in regressions using the variables at all time scales: the lack of relationship that holds at the shorter scales dominates the theoretically expected results at the very highest scales.

A similar regression using nondurable goods also indicated that at the highest scale there was a strong interest rate effect of the correct sign with a t-ratio of 6.1, but that for all scales up to S4, there were marginally significant positive interest rate effects with a t-ratio of 2.5. One might hazard the idea that at business-cycle frequencies and below, high interest rates dampen the long-term durable goods consumption, freeing up resources for short-term nondurable goods consumption. However, the empirical support for this result is not clear, so that these remarks are purely speculative until confirmation can be obtained.

In any event, the overall results unambiguously show that the interest rate effect is a long (not a short) time-scale effect, and has its maximum impact on those components that intuitively one would expect to be most interest rate sensitive. With respect to income, it is true for automobiles and furniture that the maximum income effect is at level $2^5 = 32$ months, rather than at 64 months, whereas the "other" category more closely follows the monotonic decline found for nondurable goods.

6 Summary and Conclusions

A key element in any set of economic time series of any length is the potential presence of nonstationarity beyond the notion of a unit root. Further, the theoretically imprecise statement of the functional relationships between economic variables has led many researchers to consider nonparametric procedures. An important class of such procedures is that provided by wavelets. More importantly, wavelets have some characteristics that make them particularly suitable as a vehicle for analyzing economic data. Because of the translation and scale properties, nonstationarity in the data is not a problem when using wavelets, and prefiltering is not needed. Further, because of the flexibility in choice of basis function, that is, choice of wavelet function, and because of the property of "narrow" compact support, wavelets are particularly well suited to handle complex signals that involve cusps, discontinuities, and rapid changes in modeling regime.

For the objectives of this paper, an even more important property of wavelets involves the separation of time scales of variation into a sequence of scales that can be decomposed orthogonally. A key premise underlying the research in this paper is that the relationships between economic variables may be better expressed in terms of restrictions to given time scales. The idea of time scales in structural relationships is ubiquitous in other fields (including biology), and has in fact a long history in economics. Very simple versions of this notion are captured by concepts of the "short run" and the "long run" and by Friedman's formulation of the permanent-income hypothesis. For example, the relevant variables, decision horizons, length of window for time averages, and so on will vary considerably between the decision to buy a house, or a car, or dinner, or paper clips.

A separate concept that can be easily investigated in the context of a wavelet decomposition of economic variables is that of the timing of economic relationships and its variation over the business cycle. For example, the timing of the relationship between money and income may vary with Federal Reserve policy, even to the point where monetary changes sometimes lead and sometimes lag income changes. Another example is provided by the timing of consumption income variations, which may be affected by expected changes in prices or in interest rates. These variations in timing are recognized as phase drift in terms of the spectral analysis of the data, even after restricting attention to a specific time scale. That is, as in the consumption-income relationship, there may be time periods in which the timing of the normal relationship between variables may change; variables that were in phase may move out of phase, and variables that were out of phase with each other may move further apart in time, or come together. The importance of phase drift is that if it is ignored in an empirical study, one may easily conclude that the relationship is a complicated nonlinear function, when in fact the relationship can be quite simple, but subject to phase drift.

The case study used in this paper is that involving the relationship between consumption and income. We can conclude that an appropriate way to model consumption-income relationships during the post-war period is that at the time scale dominated by trend, a simple proportional relationship exists and the marginal propensity to consume is for all goods and services about 0.94. As one reduces the time scale of analysis, the degree of fit and the slope of the consumption-income relationship declines monotonically, except for the very lowest scale, which partially reverses the decline.

The results for the scale analysis of the expenditure/income relationship for durables are even more interesting. While the overall evidence of time scale variation is as important in durables as elsewhere, there are two other facts that are of interest. The first involves an unexpected drop in durable goods expenditure at the highest scale of analysis between early 1979 and mid-1981 (approximately). The second fact involves the maximum value for the MPC, which for durable goods is at the scale of 32 months. In fact, this result is almost entirely due to automobile purchases, and there is some evidence that three years is the relevant time horizon for automobile purchases, at least over the historical period studied.

The former fact, the dip in the relationship between durable goods consumption and income, seems to be explained by the strong dependence of the relationship on the real interest rate. Indeed, we discovered that at the longest time scales for both durable goods and for nondurable goods that the real interest rate had the theoretically expected signs, was very highly significant, and explained a substantial proportion of the total variation in consumption, or expenditure. However, at shorter time scales the interest rate effect declined dramatically for durable goods, and for nondurable goods there seemed to be a moderate positive effect. We speculated that this might represent the effect that when interest rates are high and have their maximum effect on durable goods purchases, resources are released for shorter run nondurable goods purchases.

In almost all cases the residuals from the regressions with interest rates included were close in distribution to Gaussian, and seemed to be both identically distributed and uncorrelated.

While further work is necessary to pursue these tentative findings, they are of interest for those wishing to model macro data successfully. The immediate requirements for further work involve first the confirmation, possibly using data from other countries, of the qualitative results presented in this paper. The second step is to explore the potential causes of the changes that were observed in the time series at the various scales. The empirical results lead us to rethink our theoretical development to include the role of time scale as a characterization of economic data and to recognize that sometimes the "loss of simple relationships between variables" merely indicates the occurrence of a phase shift in the relationship, that is, a change in the timing of the relationship between variables. Of course, the corresponding challenge is to discover why the phase shift occurs, and what factors bring the shift to closure.

References

Brillinger, David R. (1994). "Uses of cumulants in wavelet analysis." *Proceedings of SPIE Conference on Advanced Signal Processing*, 2296: 2–18.

Bruce, A., and H. Gao (1994). *S+WAVELETS User's Manual Version 1.0*. Seattle, WA: MathSoft, Inc., StatSci Division, Mktg@StatSci.com.

Campbell, J. Y., and N. G. Mankiw (1989). "Consumption, income, and interest rates: Reinterpreting the times series evidence." *NBER Macroeconomics Annual 1989*. Cambridge, MA: NBER.

Campbell, J. Y., and N. G. Mankiw (1990). "Permanent income, current income and consumption." *Journal of Business & Economic Statistics*, 8(3):265–279.

Christiano, L. J., M. Eichenbaum, and D. Marshall (1991). "The permanent income hypothesis revisited." *Econometrica*, 59(2):397–423.

Chui, C. K. (1992). *An Introduction to Wavelets*. San Diego, CA: Academic Press.

Corbae, D., S. Ouliaris, and P. C. Phillips (1991). "A reexamination of the consumption function using frequency domain regressions." Department of Economics Working Paper Series 91-25, University of Iowa.

Daubechies, I. (1992). *Ten Lectures on Wavelets, CBMS-NSF Regional Conference Series on Applied Mathematics*, vol. 61. Philadelphia, PA: Society for Industrial and Applied Mathematics.

Deaton, A. S. (1987). "Life-cycle models of consumption: Is the evidence consistent with the theory?" In T. F. Bewley (ed.) *Advances in Econometrics: Fifth World Congress*, vol. 2. New York: Cambridge University Press.

Donoho, D., I. Johnstone, G. Kerkyacharian, and D. Picard (1995). "Wavelet shrinkage: Asymptopia (with discussion)?" *Journal of the Royal Statistical Society Series B (Methodological)*, 57(2): 301–369.

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

Flavin, M. A. (1981). "The adjustment of consumption to changing expectations about future income." *Journal of Political Economy*, 89(5): 974–1009.

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

Friedman, M. (1963). "Windfalls, the 'horizon,' and related concepts in the permanent-income hypothesis." In Carl Christ, et al. (eds.) *Measurement in Economics: Studies in Mathematical Economics and Econometrics in Memory of Yehuda Grunfeld*. Stanford, CA: Stanford University Press, pp. 3–28.

Gilbert, S. (1995). "Structural change: Estimation and testing by wavelet regression." Mimeograph, University of California at San Diego.

Goffe, W. L. (1994). "Wavelets in macroeonomics: An introduction." In D. Belsley (ed.), *Computational Techniques for Econometrics and Economic Analysis*. The Netherlands: Kluwer Academic, pp. 137–149.

Greenblatt, Seth A. (1996). "Atomic decomposition of financial data." Working Paper, Center for Quantitative Economics and Computing, Department of Economics, University of Reading, UK.

Grossman, S. J., and G. Laroque (1989). "Asset pricing and optimal portfolio choice in the presence of illiquid consumption goods." *Econometrica*, 59:1221–1248.

Hall, R. E. (1978). "Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence." *Journal of Political Economy*, 86(61):971–987.

Hog, Esben (1996). "Fractional integration: A wavelet analysis approach." Working Paper, Department of Information Science, The Aarhus School of Business, Denmark.

Jensen, Mark J. (1996). "An alternative maximum likelihood estimator of long memory processes using compactly supported wavelets." Working Paper, Department of Economics, Southern Illinois University at Carbondale.

Molana, H. (1991). "The time series consumption function: Error correction, random walk and the steady state." *The Economic Journal*, 101:382–403.

Nelson, C. R. (1987). "A reappraisal of recent tests of the permanent income hypothesis." *Journal of Political Economy*, 95(3):641–646.

Priestley, M. (1996). "Wavelets and time-dependent spectral analysis." *Journal of Time Series Analysis*, 17(1):85–103.

Quah, Dennis (1990). "Permanent and transitory movements in labor income: An explanation for 'excess smoothness' in consumption." *Journal of Political Economy*, 98(3):449–475.

Ramsey, James B., and Camille Lampart (1997a). "The decomposition of economic relationships by time scale using wavelets." Working Paper 97-08, C. V. Starr Center, New York University.

Ramsey, James B., and Camille Lampart (1998). "The decomposition of economic relationships by time scale using wavelets: Money and income." *Macroeconomic Dynamics*, 2:49–71.

Ramsey, J. B., D. Usikov, and G. M. Zaslavsky (1995). "An analysis of U.S. stock price behavior using wavelets." *Fractals*, 3(2):377–389.

Ramsey, J. B., and Z. Zhang (1995). "The analysis of foreign exchange data using waveform dictionaries." *Conference on High Frequency Dynamics*. Zurich: Olsen and Associates.

Ramsey, J. B., and Z. Zhang (1996). "The application of wave form dictionaries to stock market index data." In J. Kadtke and Y. A. Kravtsov (eds.), *Predictability of Complex Dynamical Systems*. New York: Springer-Verlag, pp. 189–205.

Rioul, Olivier, and M. Vetterli (1991). "Wavelets and signal processing." *IEEE Signal Processing Magazine*, (October):14–38.

Strang, Gilbert (1989). "Wavelets and dilation equations." *Siam Review*, 31:613–627.

Truong, Young K., and Patil, Prakesh (1991). "On estimating possibly discontinuous regressions involving stationary time series." Working Paper, Department of Biostatistics, University of North Carolina at Chapel Hill.

Viard, A. D. (1993). "The productivity slowdown and the savings shortfall: A challenge to the permanent income hypothesis." *Economic Inquiry*, 31:549–563.

Wang, Yazhen (1995). "Jump and sharp cusp detection by wavelets." *Biometrika*, 82:385–397.

Editorial Policy

The *SNDE* is formed in recognition that advances in statistics and dynamical systems theory may increase our understanding of economic and financial markets. The journal will seek both theoretical and applied papers that characterize and motivate nonlinear phenomena. Researchers will be encouraged to assist replication of empirical results by providing copies of data and programs online. Algorithms and rapid communications will also be published.

ISSN 1081-1826