12

Short-Term Analysis of Macroeconomic Time Series

Agustín Maravall

12.1. Introduction

In this paper I shall be discussing the statistical treatment of short-term macroeconomic data. In particular, I shall focus on monthly time series of standard macroeconomic aggregates, such as monetary aggregates, consumer price indices, industrial production indices, export and import series, and employment series. The statistical treatment considered is that aimed at helping economic policy-makers in short-term control, and at facilitating the monitoring and interpretation of the economy by analysts in general. The purpose of the statistical treatment is to answer two basic questions, well summarized by P. G. Wodehouse when answering a question concerning his physical shape:

The day before yesterday, for instance, [the weighing machine in my bathroom] informed me—and I don't mind telling you, J. P., that it gave me something of a shock—that I weighed seventeen stone nine. I went without lunch and dined on a small biscuit and a stick of celery, and next day I was down to eleven stone one. This was most satisfactory and I was very pleased about it, but this morning I was up again to nineteen stone six, so I really don't know where I am or what the future holds. (Wodehouse 1981: 577)

The two separate—though related—questions are:

1. Where are we?
2. Where are we heading?

Of course, forecasting provides the answer to question 2. The answer to question 1 usually consists of an estimation of the present situation, free of seasonal variation; on occasion, variation judged transitory is also removed. Thus, seasonal adjustment and trend estimation are used to answer question 1. For monthly macroeconomic series, it is often the case that seasonal variation dominates the short-run variability of the series.

In the following sections I will briefly review the recent evolution of the statistical methodology used in this context, and will provide a (justified) forecast of how I would expect it to evolve over the next ten years. The discussion will address the evolution in terms of research and in terms of practical applications (such as, for example, official seasonal adjustment). While the former typically contains a lot of noise, practical applications lag research by several (sometimes many) years.

Whenever confronted with forecasting an event, one is bound to look first at the present and past history. I shall do that, and from the (critical) look my forecast will emerge in a straightforward and unexciting manner. The discussion centres on tools used for short-term analysis; at the end I present an example that illustrates the unreliability of these tools when our horizon is a long-term one.

12.2. Short-Term Forecasting

I shall start with a very brief mention of short-term forecasting in economics. Leaving aside judgemental (or 'expert') forecasting, in the remote past some deterministic models, for example models with linear trends and seasonal dummies, were used for short-term forecasting. This practice gave way to the use of *ad hoc* filters, which became popular in the 1950s and 1960s; examples are the exponentially weighted moving average method (Winters 1960; Cox 1961), and the discounted least squares method of Brown (1962). In the 1970s, the work of Box and Jenkins (1970) provoked a revolution in the field of applied forecasting; this revolution was further enhanced by another factor: the discovery by economic forecasters of the Kalman filter (see, e.g. Harrison and Stevens 1976). The Box–Jenkins approach offered a powerful, easy-to-learn and easy-to-apply methodology for short-term forecasting; the Kalman filter provided a rather convenient tool to apply and extend the methodology.

The outcome was a massive spread of statistical models (simple parametric stochastic processes) and in particular of the so-called autoregressive integrated moving average (ARIMA) model, and of its several extensions, such as intervention analysis and transfer function models (see Box and Tiao 1975; and Box and Jenkins 1970), and of the closely related structural time series models (see Harvey 1989). One can safely say that ARIMA models are used every day by thousands of practitioners. Some use is also made of multivariate versions of these models (see e.g. Litterman 1986), although multivariate extensions have often been frustrating. A good review of the economic applications of time-series models is contained, for example, in Mills (1990).

The point I wish to make at present is that, overall, stochastic-model-based forecasting has become a standard procedure. At present, two important directions of research are:

1. *multivariate extensions*, where the research on cointegration and common factors may lead to an important break (through a reduction in dimensionality and an improved model specification); see, for example, Gouriéroux (1992) and Hamilton (1994);
2. *nonlinear extensions*, such as the use of bilinear, ARCH, GARCH, stochastic parameter, Markov chain models, and so on (see, for example, Johansen 1996, and Tiao and Tsay 1989).

For the type of series considered, I would expect direction 1 eventually to play a very important role in applied short-term forecasting over the next decade. As for the future impact of direction 2, I doubt that in the next few years stochastic nonlinear models will become a standard tool for the average practitioner, working with monthly or quarterly series.

12.3. Unobserved Components Estimation and Seasonal Adjustment

Going back to question 1 of Section 12.1, we proceed to the problem of estimating the relevant underlying evolution of an economic variable, that is, to seasonal adjustment; we shall also consider some trend estimation issues. Two good references that describe the state of the art concerning seasonal adjustment over the last 20 years are Den Butter and Fase (1991) and Hylleberg (1992). As with forecasting, deterministic models were used in the distant past. At present, however, except for some isolated cases, official agencies producing monthly seasonally adjusted data do not remove seasonality with deterministic models.

It is widely accepted by practitioners that, typically, seasonality in macroeconomic series is of the moving type, for which the use of filters is appropriate. If x_t denotes the series of interest (perhaps in logs), n_t the seasonally adjusted series, and s_t the seasonal component, a standard procedure is to assume

$$x_t = n_t + s_t,$$

and to estimate s_t by

$$\hat{s}_t = C(B)x_t, \tag{1}$$

where B is the lag operator such that $B^k x_t = x_{t-k}$ (k integer), and $C(B)$ is the linear and symmetric filter,

$$C(B) = c_0 + c_1(B + F) + \ldots + c_r(B^r + F^r), \tag{2}$$

with $F = B^{-1}$. Of course, the filter $C(B)$ is designed to capture variability of the series in a (small) interval around each seasonal frequency. The seasonally adjusted series is, in turn, estimated by

$$\hat{n}_t = [1 - C(B)]x_t = A(B)x_t, \tag{3}$$

and $A(B)$ is also a centred, linear, and symmetric filter.

The symmetric and complete filters (1) and (3) cannot be used to estimate s_t or n_t when t is close to either end of the series. Specifically, at time T, when the available series is $(x_1, \ldots, x_T)$, estimation of s_t with (1) for $t < r$ requires unavailable starting values of x; analogously, estimation of s_t for $t > T - r$ requires future observations not yet available. Therefore, the centred and symmetric filter characterizes 'historical' estimates. For recent enough periods asymmetric filters have to be used, which yield preliminary estimators. As time passes and new

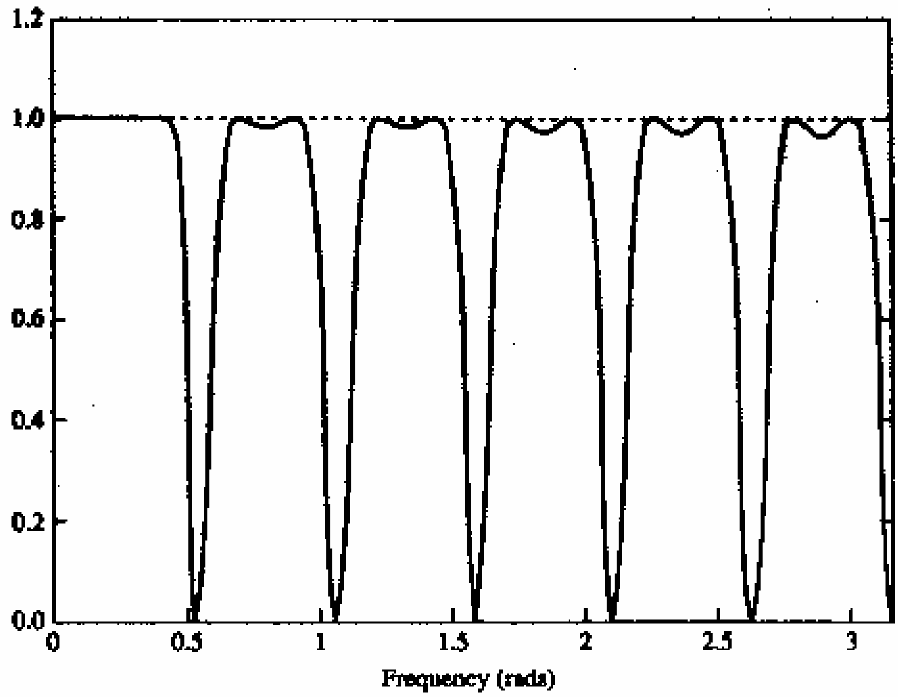

Fig. 12.1. X11: Transfer Function

observations become available, those preliminary estimators will be revised until the historical or final estimator is eventually obtained. We shall come back to this issue later.

In the same way that 1970 marks an important date for applied short-term forecasting (i.e., the year of publication of the book by Box and Jenkins), 1967 marks a crucial event in the area of seasonal adjustment. That event was the appearance of the program X11, developed by the US Bureau of the Census (see Shiskin *et al.* 1967). Except for some outlier treatment (which we shall ignore), X11 can be seen as a sequence of linear filters, and hence as a linear filter itself (see e.g. Hylleberg, 1986). For the discussion that follows, we use the parsimonious approximation to X11 (historical filter) of Burridge and Wallis (1984). Fig. 12.1 plots the transfer function (in the frequency domain) of the X11 monthly filter $A(B)$. This function represents, for each frequency, the proportion of the variability of x_t that is used to estimate the seasonally adjusted series. It is seen how the X11 filter passes the variation associated with all frequencies, except for some small intervals around the seasonal ones.

Over the next decade, X11 spread at an amazing speed, and many thousands of series came to be routinely adjusted with X11. It was an efficient and easy-to-use procedure that seemed to provide good results for many series (although the meaning of 'good' for seasonal adjustment is somewhat unclear). Yet, towards the end of the 1970s some awareness of X11 limitations started to develop. Those limitations were mostly associated with the rigidity of the X11 filter, i.e. with its *ad hoc*, relatively fixed structure ('fixed' includes the case in which a few fixed options may be available). It is to these limitations that I turn next.

12.4. Limitations of *Ad Hoc* Filtering

I shall provide simple illustrations of some major limitations of fixed *ad hoc* filters. For illustration we use X11 run by default

1. The danger of spurious adjustment is illustrated in Fig. 12.2. In a white-noise series, with spectrum that of part (*a*) in the figure, X11 will extract a seasonal component, with spectrum that of part (*b*). This spectrum is certaintly that of a seasonal component, but the series had no seasonality to start with.

2. In the previous white-noise series, trivially, the filter $A(B)$ used to seasonally adjust the series should simply be 1; at the other extreme, if the observed series has a spectrum as in part (*b*) of the figure, the filter to adjust the series seasonally should obviously be 0, since the series contains only seasonal variation. The filter should thus depend on the characteristics of the series.

To illustrate the point, we use the well-known 'Airline model' of Box and Jenkins (1970: ch. 9). It is a model appropriate for monthly series displaying trend and seasonality. For series in logs, the model implies that the annual difference of the monthly rate of growth is a stationary process. The Airline model is, on the one hand, a model often encountered in practice; on the other hand, it provides an excellent reference example. The model is given by the equation

$$\nabla\nabla_{12}x_t = (1 - \theta_1 B)(1 - \theta_{12}B^{12})\, a_t + \mu, \tag{4}$$

where μ is a constant, a_t is a white-noise innovation (with variance V_a), $\nabla = 1 - B$, $\nabla_{12} = 1 - B^{12}$, and $-1 < \theta_1 < 1$, $0 < \theta_{12} < 1$. The series x_t, generated by (4), accepts a rather sensible decomposition into trend, seasonal, and irregular component (see Hillmer and Tiao 1982). As θ_1 approaches 1, model (4) tends towards the model

$$\nabla_{12}x_t = (1 - \theta_{12}B^{12})\, a_t + \mu_0 + \mu t,$$

with deterministic trend. Similarly, when θ_{12} becomes 1, the seasonal component becomes deterministic. Thus, the parameter $\theta_1(\theta_{12})$ may be interpreted as a measure of how close to deterministic the trend (seasonal) component is.

In the frequency domain, this 'closer to deterministic' behaviour of a component is associated with the width of the spectral peaks. Thus, for example, Fig. 12.3 displays the spectra of two series both following models of the type (4). The one with the continuous line contains more stochastic seasonal variation, in accordance with the wider spectral peaks for the seasonal frequencies. The seasonal component in the series with spectrum given by the dotted line will be more stable, and hence closer to deterministic. Since the X11 seasonal adjustment filter displays holes of fixed width for the seasonal frequencies, it follows that X11 will underadjust when the width of the seasonal peak in the series

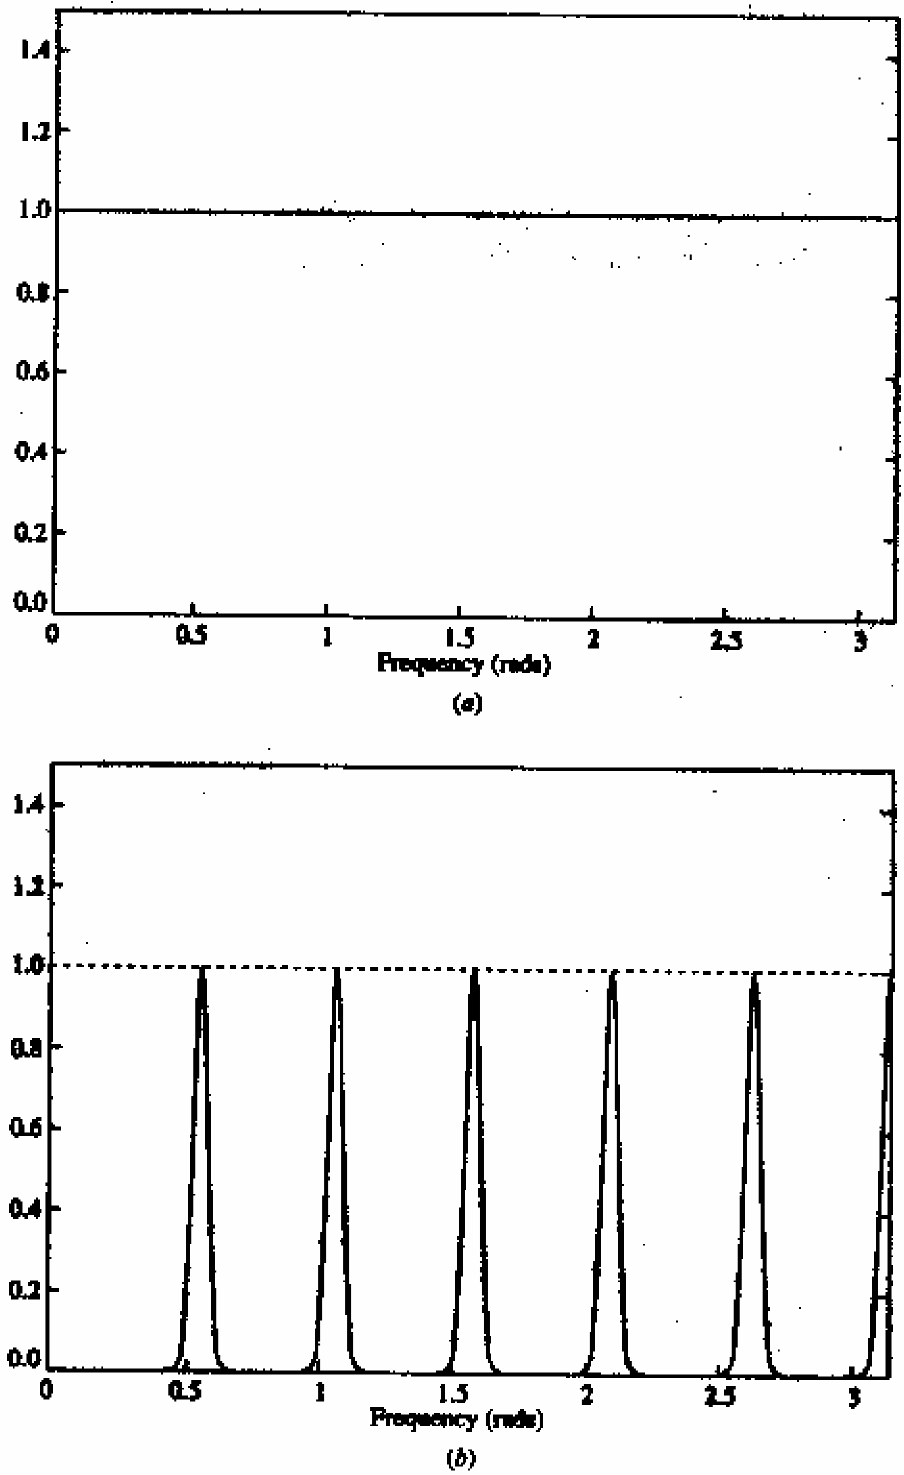

Fig. 12.2. (*a*) Spectrum: White-Noise Variable
(*b*) Spectrum: Seasonal Component in the White Noise (X11)

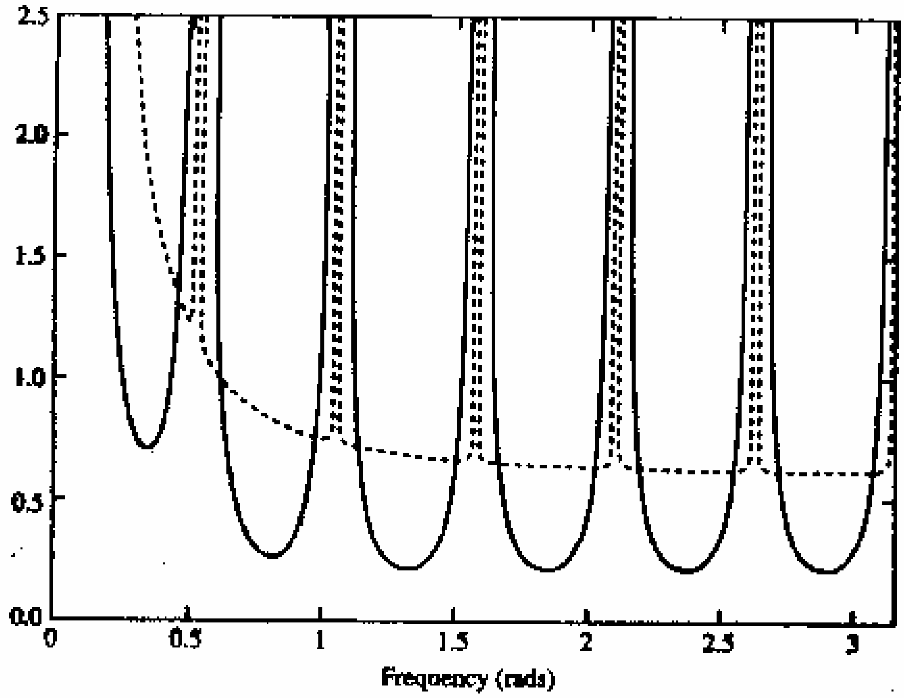

Fig. 12.3. Spectra: Stable and Unstable Seasonality

spectrum is larger than that captured by the X11 filter. Fig. 12.4(*a*) illustrates this situation, and Fig. 12.4(*b*) displays the spectrum of the estimated seasonally adjusted series obtained in this case. The underadjustment is reflected in the two peaks that remain in the neighbourhood of each seasonal frequency: obviously, X11 has not removed all seasonal variation from the series.

On the other hand, X11 will overadjust (i.e. will remove too much variation from the series) when the width of the seasonal spectral peaks are narrower than those captured by X11. This effect is evidenced in Fig. 12.5: the holes in the seasonally adjusted series spectrum (part (*b*) of the figure) are now too wide.

3. Another limitation of X11 is the lack of a proper framework for detecting the cases in which its application is inappropriate. On the one hand, diagnostics are few and difficult to interpret. Moreover, when found inappropriate, there is no systematic procedure to overcome the inadequacies.

4. Even when appropriate, X11 does not contain the basis for proper inference. For example, what are the standard errors associated with the estimated seasonal factors? This limitation has important policy implications (see Bach *et al.* 1976; Moore *et al.* 1981). In short-term monetary control, if the monthly target for the rate of growth of M_1 (seasonally adjusted) is 10 per cent and actual growth for that month turns out to be 13 per cent, can we conclude that growth has been excessive and, as a consequence, raise short-term interest rates? Can

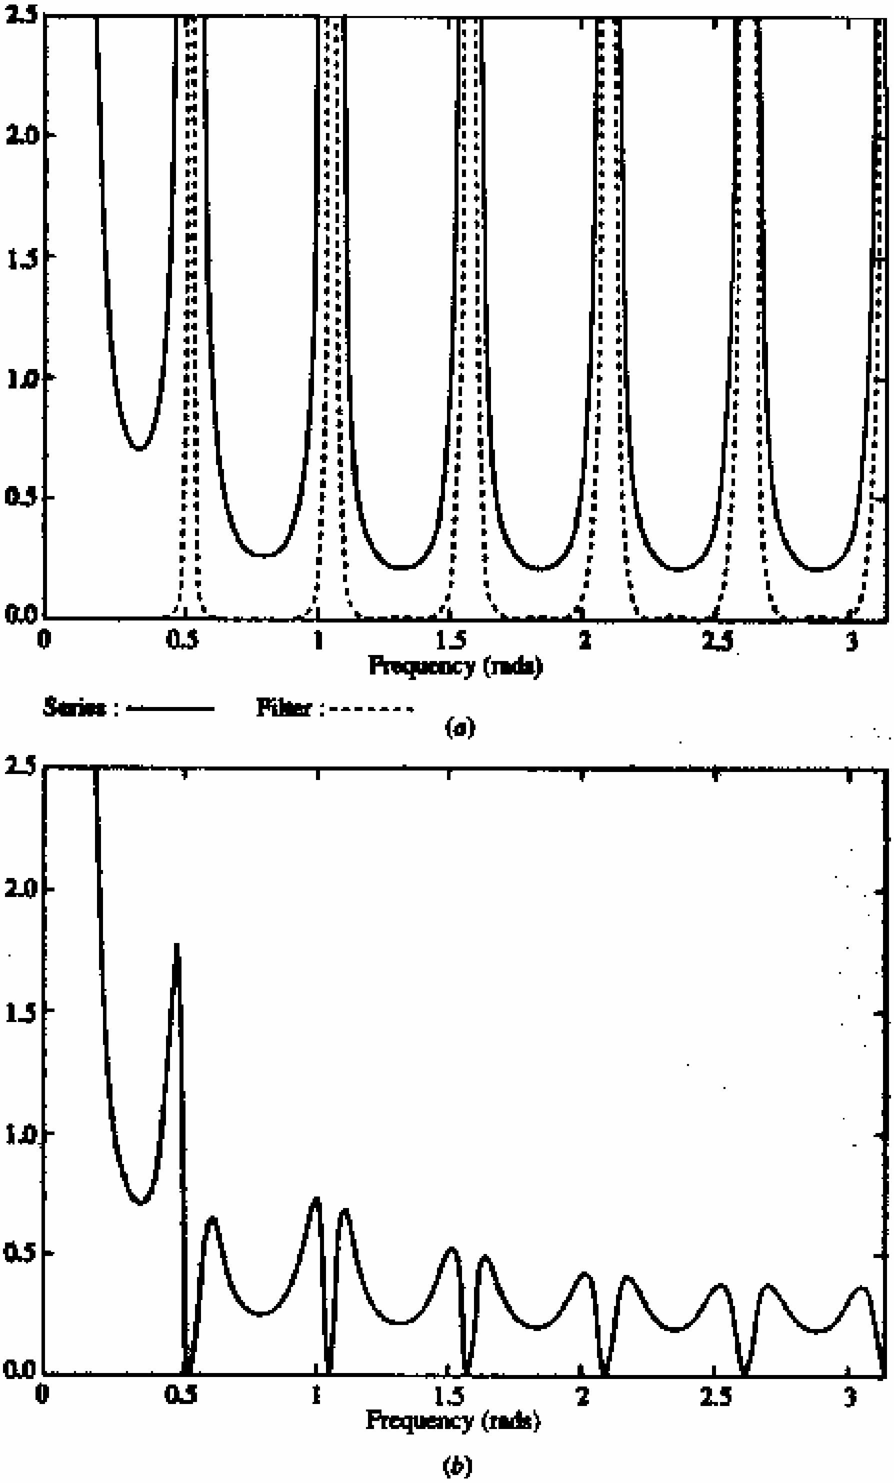

Fig. 12.4. (*a*) X11 Underadjustment: Series and Seasonal Filter
(*b*) X11 Underadjustment: Seasonally Adjusted Series

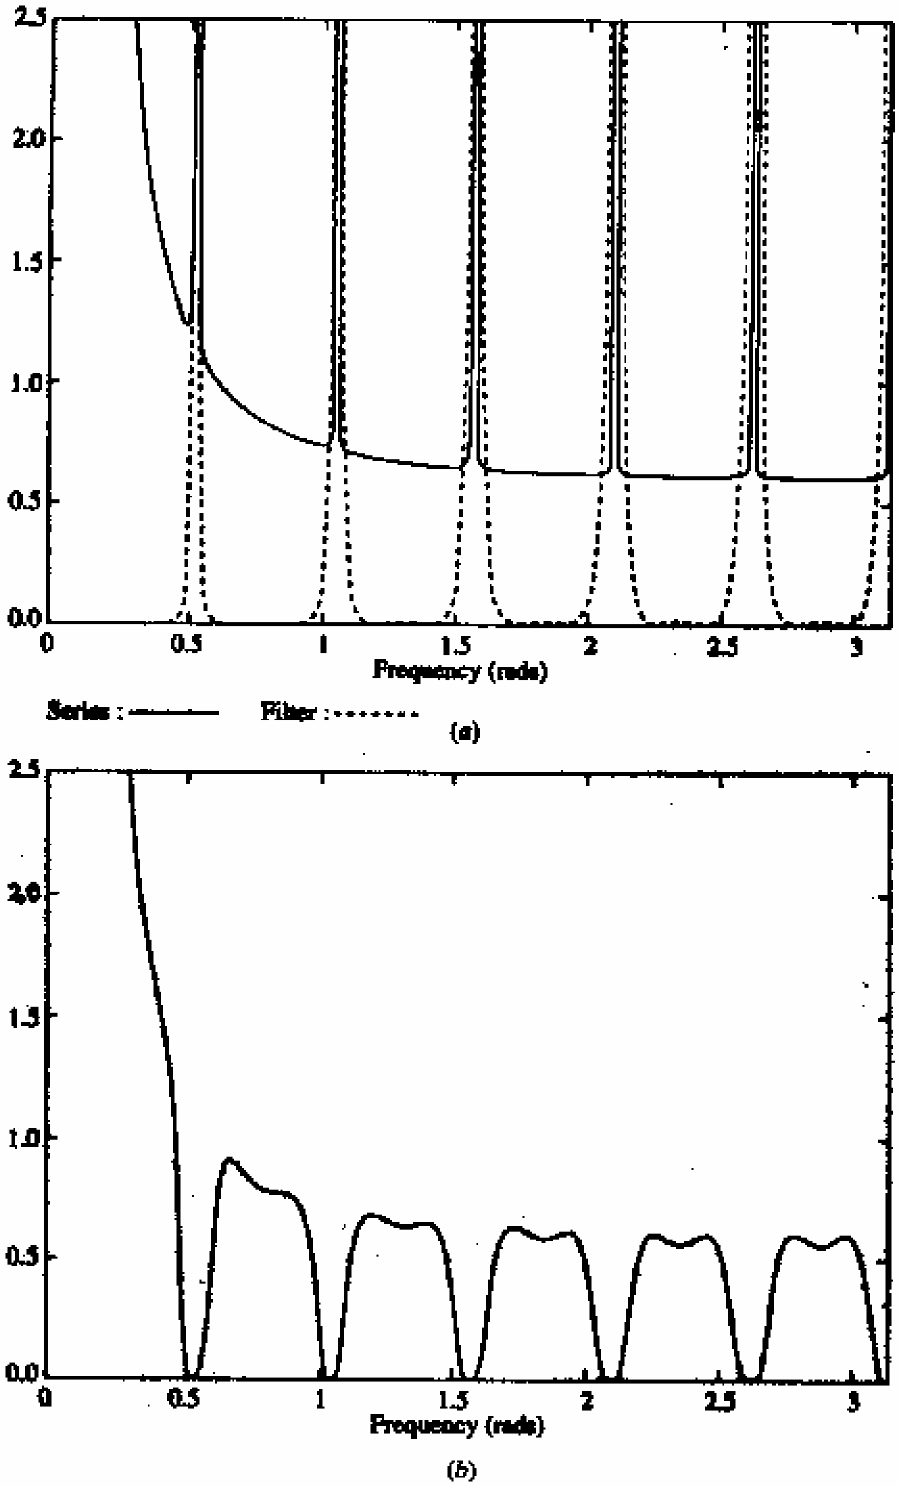

Fig. 12.5. (*a*) X11 Overadjustment: Series and Seasonal Filter
(*b*) X11 Overadjustment: Seasonally Adjusted Series

the 3 percentage point (p.p.) difference be attributed to the error implied by the estimation of the seasonally adjusted series? Similarly, when assessing the evolution of unemployment, if the series of total employment grows by 90,000 persons in a quarter, and the seasonal effect for that quarter is estimated as an increase of 50,000 persons, can we assume that the increase has been more than a pure seasonal effect?

5. In the same way that X11 does not provide answers to these questions, it does not allow us to compute optimal forecasts of the components. (Seasonal factors for the year ahead are simply computed by adding to this year factors one-half of the difference between them and the factors one year before. Of course, there is no measure of the uncertainty associated with these forecasts.)

6. Although X11 computes separate estimates of the trend, seasonal, and irregular components, their statistical properties are not known. Therefore, it is not possible to answer questions such as, for example, whether the trend or the seasonally adjusted series provide a more adequate signal of the relevant underlying evolution of the series (see Kenny and Durbin 1982; Moore *et al.* 1981; and Maravall and Pierce 1986).

To overcome some of those limitations, throughout the years X11 has been subject to modifications. In particular, the program X11 ARIMA, developed by Statistics Canada (see Dagum 1980), improved upon X11 in several ways. First, it incorporated several new elements for diagnosis. Perhaps more relevantly, it provided better estimators of the components at the end of the series. This was achieved by replacing the *ad hoc* X11 filters for the preliminary estimators with a procedure in which the series is extended with ARIMA forecasts, so that the filter $A(B)$ can be applied to the extended series (and hence to more recent periods). In fact, X11 ARIMA has replaced X11 in many standard applications.

At present, the US Bureau of the Census has just completed a new program for seasonal adjustment: X12 ARIMA (see Findley *et al.* 1998). The program follows the direction of X11 ARIMA, and incorporates some new sets of diagnostics and some new model-based features, having to do with the treatment of outliers and the estimation of special effects.

Be that as it may, practical applications (such as 'official' seasonal adjustment by agencies) lag with respect to research. So, let us now turn to the evolution of research during the last twenty years.

12.5. The Model-Based Approach

Towards the end of the 1970s and beginning of the 1980s, a new approach to the problem of estimating unobserved components in time series, and in particular to seasonal adjustment, was developed. The approach combined two elements: the use of simple parametric time-series models (mostly of the ARIMA type); and the use of signal extraction techniques. Although there were earlier attempts at using signal extraction on time-series models (see Nerlove *et al.* 1979), these attempts were of limited interest because they were restricted to stationary series, while economic series are typically non-stationary.

The model-based approach has taken two general directions. One is the so-called ARIMA-model-based methodology, and some relevant references are Burman (1980), Hillmer and Tiao (1982), Bell and Hillmer (1984), and Maravall and Pierce (1987). The second direction follows the so-called structural time-series methodology, and some important references are Engle (1978), Harvey and Todd (1983), and Gersch and Kitagawa (1983). We shall refer to them as the AMB and STS approaches, respectively. Both are closely related, and share the following basic structure (for a more complete description, see Maravall 1995).

The observed series $[x_t] = [x_1, \ldots, x_t]$ can be expressed as the sum of several orthogonal components,

$$x_t = \sum_i x_{it},$$

where each component x_{it} may be expressed as an ARIMA process (with Gaussian innovations). Thus, for example, the model for the trend, p_t, may be of the type

$$\nabla^2 p_t = \theta_p(B)\, a_{pt},$$

where $\theta_p(B)$ is a polynomial in B, and the model for the seasonal component, s_t, is often of the form (for monthly series):

$$(1 + B + \ldots + B^{11}) s_t = \theta_s(B)\, a_{st},$$

specifying that the sum of 12 consecutive seasonal components is a zero-mean stationary process (with 'small' variance). While the trend and seasonal components are typically non-stationary, the irregular component is a zero-mean stationary process, often simply white noise. Since the sum of ARIMA models yields an ARIMA model, the observed series x_t also follows an ARIMA model, say

$$\phi(B)\, x_t = \theta(B)\, a_t, \tag{5}$$

where $\phi(B)$ contains the stationary and non-stationary autoregressive roots.

Once the models are specified, the unobserved components are estimated as the minimum mean squared error (MMSE) estimator

$$\hat{x}_{it} = E(x_{it} \mid [x_t]), \tag{6}$$

and this conditional expectation is computed with signal extraction techniques.[1] The MMSE estimator (6) obtained in the model-based approach is also a linear filter, symmetric, centred, and convergent in the directions of both the past and

[1] This technique is a fairly general procedure that can be applied to a variety of statistical problems besides unobserved components estimation. In particular, forecasting can be seen as the particular case when $\hat{x}_{it}$ is the estimator of a future observation; another well-known application is interpolation of missing values.

the future. Thus, as was the case with filter (2), the filter applies to historical estimates, and the problem of preliminary estimation and revisions again reappears. The model-based approach offers an optimal solution: the observed series are extended with forecasts (and backcasts) as needed, and the symmetric and centred filter can then be applied to the extended series.[2]

The decomposition of x_t into unobserved components presents a basic identification problem. In general, the AMB and STS methods use somewhat different assumptions in order to reach identification; these different assumptions, of course, lead to differences in the specification of the component models. It is the case, however, that the STS trend and seasonal components can be expressed as the ones obtained from an AMB approach with superimposed orthogonal white noise (see Maravall 1985). Ultimately, the crucial assumption for identification of the components concerns the amount of variance assigned to the irregular; the AMB approach, in order to maximize the stability of the trend and seasonal components, maximizes the irregular component variance.

Besides these differences in the specification of the component models, there are some additional ones between the two approaches. The AMB method starts by specifying the model for the observed series, following standard Box–Jenkins techniques. From this aggregate model, the component models are then derived, and the conditional expectation (6) is obtained with the Wiener–Kolmogorov filter (see Whittle 1963; and Bell 1984). On the contrary, the STS method starts by directly specifying the models for the components, and uses the Kalman filter to compute the conditional expectation (6); see Harvey (1989).

A simple example can illustrate the basic differences between the two model-based approaches. Assume a non-seasonal series with possibly a unit root. The STS method would likely estimate the model

$$x_t = p_t + u_t, \tag{7a}$$

where p_t, the trend, follows the random walk model

$$\nabla p_t = b_t, \tag{7b}$$

and b_t is white noise, orthogonal to the white-noise irregular u_t. The parameters that have to be estimated are two, namely the variances $V(b_t)$ and $V(u_t)$.

That basic model implies that the observed series x_t follows an IMA(1, 1) model, say

$$\nabla x_t = (1 - \theta B)\, a_t. \tag{8}$$

A potential problem of the STS is that, since it does not include a prior identification stage, the model specified may be inappropriate. On the contrary, the AMB approach first identifies the model with standard ARIMA-identification tools. (There are indeed many available.) If, in the AMB approach, the model

[2] In terms of the observed values, the filter will be, of course, asymmetric; see Cleveland and Tiao (1976).

identified for the observed series turns out to be of the type (8), then the decomposition becomes

$$x_t = p_t^* + u_t^*, \tag{9a}$$

$$\nabla p_t^* = (1 + B)\, b_t^*, \tag{9b}$$

and b_t^* is white noise, orthogonal to the white-noise irregular u_t^*. The factor $(1+B)$ in the MA part of (9b) implies that the spectrum of p_t^* is monotonically decreasing in the range $(0, \pi)$, with a zero at frequency π. It will be true that $V(u_t^*) > V(u_t)$, and, in fact, p_t can be expressed as

$$p_t = p_t^* + c_t,$$

where c_t is white noise, orthogonal to b_t^*, and with variance $[V(u_t^*) - V(u_t)]$.

The example illustrates some additional differences. It is straightforward to find, for example, that model (7), the STS specification, implies the constraint $\theta \geq 0$. (Otherwise the irregular has negative spectrum.) This constraint disappears in the AMB approach.

The parameters that have to be estimated in the STS approach are the variance of the innovations in the components, $V(u_t)$ and $V(b_t)$; in the AMB approach, the parameters to estimate are those of a standard ARIMA model. Of course, the STS approach has a parsimonious representation in terms of the component's models, and is likely to produce unparsimonious ARIMA expressions for the observed series. On the contrary, the AMB approach estimates a parsimonious ARIMA for the observed series, and the derived models for the components may well be unparsimonious. In both cases, estimation of the model is made by maximum likelihood. The component's estimators in the STS approach are obtained with the Kalman filter-smoother, while in the AMB approach the Wiener–Kolmogorov filter is used. If the former filter offers more programming flexibility, the Wiener–Kolmogorov is more informative for analytical purposes.

Be that as it may, despite the differences, both methods share the same basic structure of ARIMA components–ARIMA aggregate, where the components estimators are the expectations conditional on the available observations (their least-squares projections). Ultimately, both represent valid approaches. My (probably biased) view is that the AMB method, by using the data to identify the model, is less prone to misspecification. I find it reasonable, moreover, in the absence of additional information, to provide trend and seasonal components as smooth as possible, within the limits of the overall stochastic behaviour of the observed series. Further, the AMB method typically implies direct estimation of fewer parameters, and provides results that are quite robust and numerically stable. On the other hand, the state space–Kalman filter format in the STS methodology offers the advantage of its programming and computational simplicity and flexibility. In any case, both methods provide interesting and relatively powerful tools for unobserved components estimation in linear stochastic processes. It is worth noticing that many *ad hoc* procedures can be given a minimum

MSE-model-based interpretation for particular ARIMA models (see e.g. Cleveland and Tiao 1976; Burridge and Wallis 1984; and Watson 1986).

12.6. The Virtues of a Model-Based Method

The major advantage of a model-based method is that it provides a convenient framework for straightforward statistical analysis. To illustrate the point, I return to the six examples used when illustrating the limitations of *ad hoc* filtering is Section 12.4. As the model-based method, we use the program SEATS ('Signal Extraction in ARIMA Time Series; see Gomez and Maravall 1996), which enforces the AMB approach as in Burman (1980).

1. The danger of spurious adjustment is certainly attenuated: if a series is white-noise, it would be detected at the identification stage, and no seasonal adjustment would be performed.

2. The dangers of underadjustment (Fig. 12.4) and overadjustment (Fig. 12.5) are also greatly reduced. The parameters of the ARIMA model will adapt themselves to the width of the spectral peaks present in the series. Fig. 12.6 illustrates the seasonal adjustment (with the AMB method) of the series in Fig. 12.4: part (*a*) illustrates how the filter adapts itself to the seasonal spectral peaks, and part (*b*) shows how the spectrum of the estimated seasonally adjusted series shows no evidence now of underadjustment. AMB seasonal adjustment of the series with a very stable seasonal (the series of Fig. 12.5) is displayed in Fig. 12.7. The filter now captures a very narrow band, and the spectrum of the adjusted series estimator does not provide evidence of overadjustment.

3. To illustrate how the model-based approach can provide elements of diagnostics, we use an example from Maravall (1987). The example also illustrates how, when the diagnostic is negative, one can proceed in order to improve upon the results.

When adjusting with X11 the Spanish monthly series of insurance operations (a small component of the money supply), the program indicated that there was too much autocorrelation in the irregular estimator, $\hat{u}_t$. In fact, the lag – 1 autocorrelation of $\hat{u}_t$ was 0.42. This seems large, but what would be the correct value for X11? There is no proper answer to this question.

For a model-based method with a white-noise irregular component, u_t, the MMSE estimator $\hat{u}_t$ has the autocorrelation funcion (ACF) of the 'inverse' model of (5), that is of the model obtained by interchanging the AR and the MA polynomials. Hence, given the model for the observed series, the theoretical value of the ACF for $\hat{u}_t$ is easily obtained. For the model-based interpretation of X11 (for which u_t is white-noise), one finds $\rho_1(\hat{u}_t) = -0.2$ with a standard error of 0.1. Thus, a 95 per cent confidence interval for ρ_1 would be, approximately, (–0.40, 0). Since the value obtained, 0.42, is far from the interval, in the model-based approach it is clear that there is indeed too much autocorrelation in the irregular.

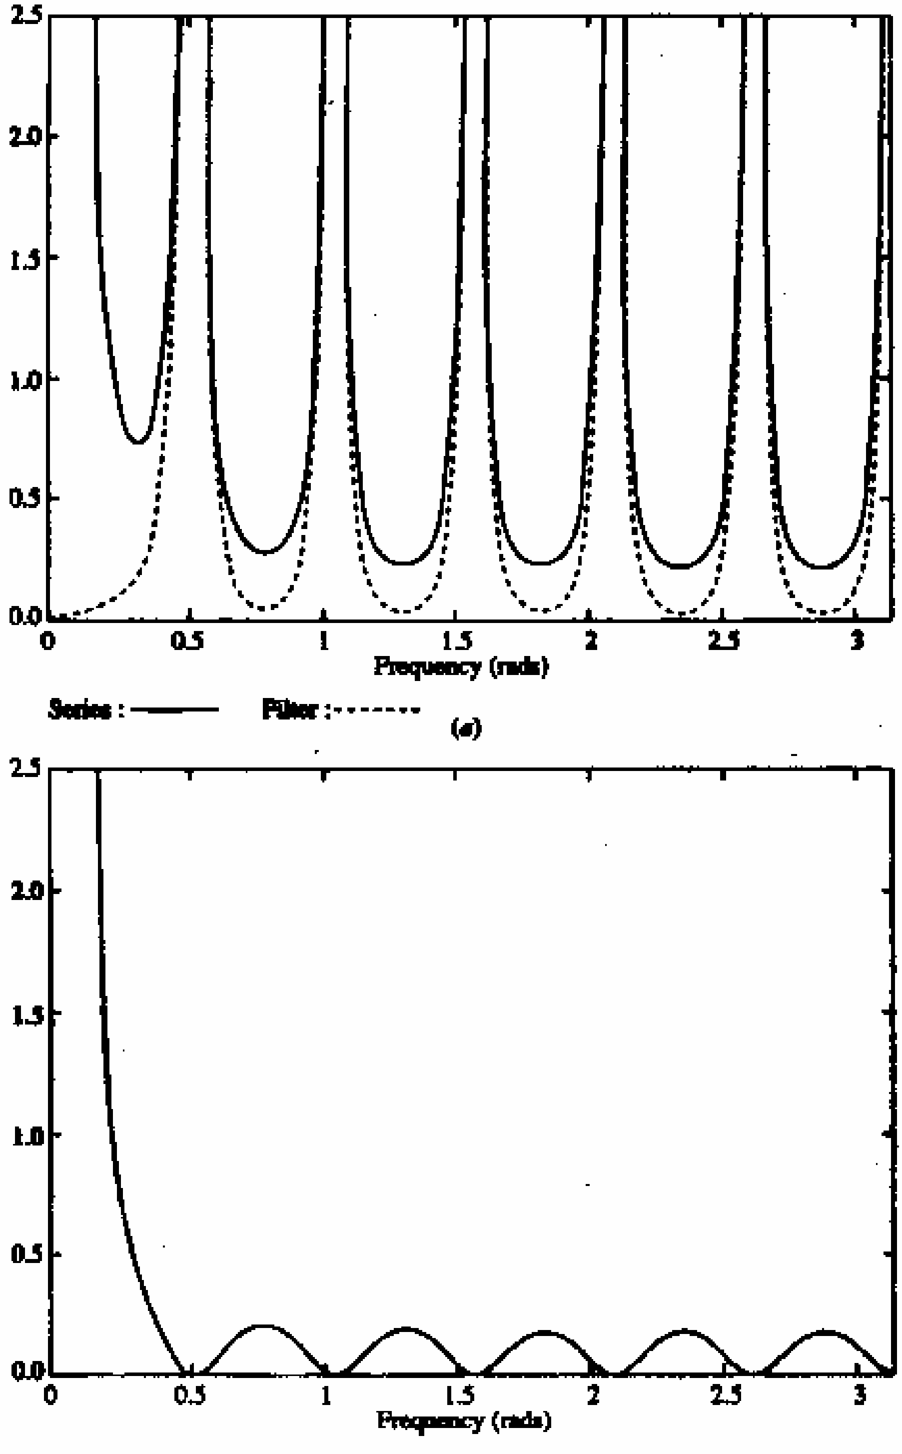

Fig. 12.6. (*a*) Unstable Seasonal: Series and AMB Seasonal Filter
(*b*) Unstable Seasonal: Seasonally Adjusted Series

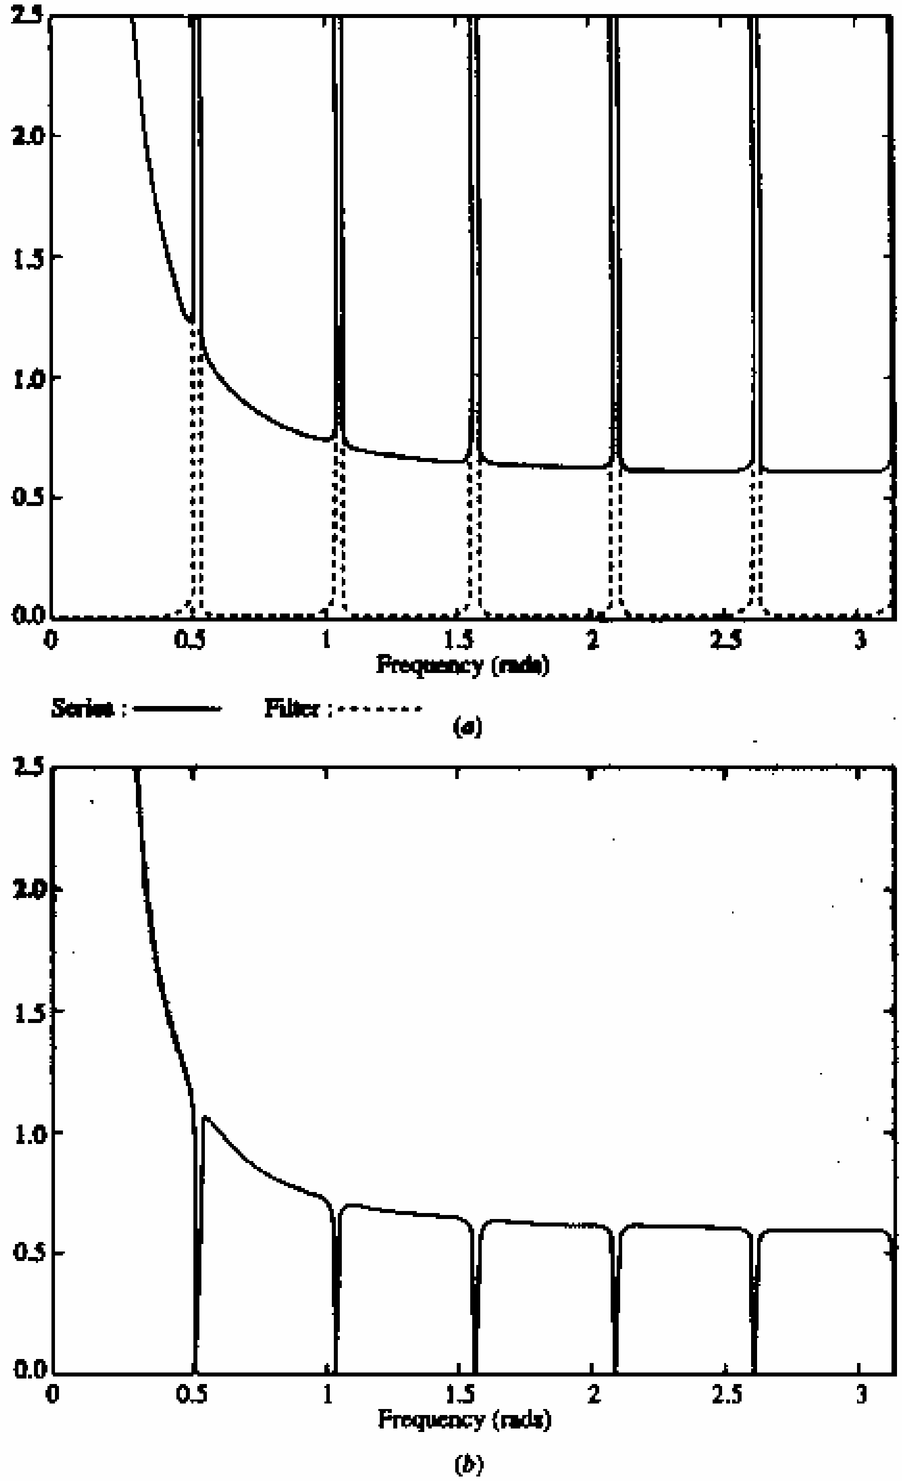

Fig. 12.7. (*a*) Stable Seasonal: Series and Seasonal Filter
(*b*) Stable Seasonal: Seasonally Adjusted Series

The relative large, positive autocorrelation in $\hat{a}_t$ seems to indicate underestimation of the trend, which, to some degree, contaminates the irregular. The ARIMA model for which X11 provides an MMSE filter contains the stationary transformation $\nabla\nabla_{12}$. Having had a negative diagnostic, back to the identification stage it was found that a model with the transformation $\nabla^2\nabla_{12}$, which allowed for a more stochastic trend, provided a better fit of the series. For this model, the theoretical value of $\rho_1(\hat{a}_t)$ was -0.83; AMB decomposition of the series with the new model specification yielded an irregular with $\hat{\rho}_1 = -0.82$, with the standard error (s.e.) = 0.05., perfectly in agreement with what should have been obtained.

Thus, the model-based approach offers a natural setup for carrying out diagnostics, and permits us to improve upon the results by applying the standard iterations (identification/estimation/diagnosis, and back to identification) of a model-building procedure.

4. As for the possibility of drawing inference, I mentioned the importance of measuring the errors associated with the estimated components.

Given the model, the estimator (6) contains two types of error. First, as mentioned in Section 12.3, when t is not far from the two ends of the series, a preliminary estimator will be obtained. Consider the case of concurrent estimation, that is the estimation of x_{it} when the last observation is x_t. As new observations become available, eventually the preliminary estimator will become the final one (i.e. the one that yields historical estimators). The difference between the preliminary and final estimator is the 'revision error'. The second type of error is the one contained in the final estimator, implied by the stochastic nature of the component. The revision error and the error in the final estimator are orthogonal (Pierce 1980).

With the model-based approach, it is straightforward to compute variances and autocorrelations of each type of error. Thus, in the examples used in point 4 of Section 12.4, the AMB method of SEATS yields the following answers. The standard error in the estimator of the monthly rate of growth of the seasonally adjusted Spanish monetary aggregate series is 1.95 percentage points of annualized growth. Thus, with a 95 per cent size we cannot reject the hypothesis that the measurement of 13 per cent growth is in agreement with the 10 per cent target. (If the size is reduced to 70 per cent, then the mea-sured growth becomes significantly different from the target.)

As for the quarterly series of Spanish employment, the standard error of the seasonal component estimator is equal to 19,000 persons. Thus, the 90,000 increase could be (barely) accepted as significantly more than the seasonal effect of 50,000.

5. The model-based approach provides MMSE forecasts of the components, as well as their associated standard errors. For example, for the Spanish monthly series of imports, the standard errors of the one- and twelve-periods-ahead forecasts for the original series, the seasonally adjusted series, and the trend are as follows (in percentages of the level):

No. of Periods ahead	Series	Seasonally adjusted series	Trend
1	11.6	11.0	5.4
12	14.9	14.6	11.0

The trend, thus, appears to be a considerably more precise forecasting tool.

The standard errors of the components provide answers to many problems of applied interest. For example, it is clear that optimal updating of preliminary seasonally adjusted data implies re-estimation whenever a new observation becomes available. This 'concurrent' adjustment implies a very large amount of work; in particular, it requires agencies producing data to change every month many series for many years. So, the overwhelming practice is to adjust once a year, and it is of interest to know how much precision is lost by this sub-optimal procedure. This can be easily computed, and, for the import series, moving from a once-a-year seasonal adjustment to a concurrent one decreases the root mean squared error (on average) by 10 per cent. Given real-life limitations, it would seem to me a case in which the improvement hardly justifies the effort.

6. In the previous point we compared the forecast errors of the trend and seasonally adjusted series. Since the two can be taken as alternative signals of the relevant underlying evolution of the series, it is of interest to look at a more complete comparison of their relative performances. Consider now the Spanish monthly series of exports. An Airline-type model fits the series well, although the series has a large forecast error variance. In terms of the components, this is associated with a large irregular component.

Starting with concurrent estimation (the case of most applied interest), the variances of the different types of error, expressed as fractions of the variance of a_t (the residuals of the ARIMA model), are as follows:

	Seasonally adjusted series	Trend
Revision error	0.073	0.084
Final estimation error	0.085	0.068
Total error	0.158	0.152

Therefore, the error contained in the concurrent estimator of the two signals is roughly equal. The error in the final estimator is smaller for the trend; in turn, the seasonally adjusted series is subject to smaller revisions.

In addition to the size of the full revision in the concurrent estimator, it is of interest to know how fast the estimator converges to the final value. After one year of additional data, for the trend component, 92 per cent of the revision standard deviation has been completed. The percentage drops to 28.8 per cent for the seasonally adjusted series. Thus, the trend estimator converges much faster to its final value.

Often, policy-makers or analysts are more interested in looking at rates of growth than at levels. Three of the most popular ones are the monthly rate of growth of the monthly series, r_1, the monthly rate of growth of a three-month moving average, r_3, and the annual rate of growth centred in the present month, r_{12} (that is, the growth over the last six months plus the forecasted growth over the next six months). For the export series, both r_1 and r_3 are annualized and the three rates are expressed in percentage points. The standard errors of the concurrent estimators of the three rates of growth are found to be:

	Series	Seasonally adjusted series	Trend
r_1	—	85.3	21.4
r_3	—	47.4	16.1
r_{12}	14.3	13.9	8.8

Thus, an attempt to follow the evolution of exports by looking at the monthly rate of growth of the seasonally adjusted series would be likely to induce a manic-depressive behaviour in policy-makers and analysts (similar to the one reported by Wodehouse).

Finally, the standard error of the one-period-ahead forecast of the series is 12.6 per cent of the level of the series. For the seasonally adjusted series, the corresponding forecast error becomes 11 per cent, and it drops to 4.5 per cent for the trend component.

From the previous results, it is clear that for the case of the export series (and a similar comment applies to the series of imports) the seasonally adjusted series provides a highly volatile and unstable signal, and the use of the trend in month-to-month monitoring seems certainly preferable.

12.7. The Next Ten Years

In the previous section I have tried to illustrate some of the advantages of a model-based approach in short-term analysis of macroeconomic data. In fact, the model-based approach can be a powerful tool, and it is gradually becoming available to the community of applied statisticians and economists. The speed of its diffusion, however, is damped by two basic problems. The first one is the inertia that characterizes beaurocratic institutions producing large amounts of economic data. (Old habits die hard!) The second is that, when dealing with many series, individual identification of the correct model for each series may seem, in practice, unfeasible.

This second limitation is, in my opinion, more apparent than real. There are already some automatic model identification procedures that can be enforced in an efficient and reliable manner (see e.g. Tsay and Tiao 1984, Beguin *et al.* 1980, and Gomez and Maravall 1998). They can also incorporate additional convenient features, such as automatic detection and correction for several types of outlier. In this way, they can be used routinely on large sets of series.[3]

My forecast for the next ten years will thus come as no surprise: model-based signal extraction with ARIMA-type models will increase its importance for practical applications, and will eventually replace the X11-based methodology (although this may take more than a decade). It is worth mentioning that the new Bureau of the Census program X12 (the successor to X11) contains a preadjustment program which is ARIMA-model-based, and hence represents a first move in the model-based direction. On the other side of the Atlantic, EUROSTAT is already using SEATS for routine adjustment of some of their series.

As for directions of new research, the extension of signal extraction to multivariate models seems to me a promising direction. I would expect to see multivariate models that incorporate the possibility that several series may share several components. This would permit a more efficient estimation of the components, and a more parsimonious multivariate model. Some preliminary steps in that direction can be found in Harvey *et al.* (1994), Fernandez-Macho *et al.* (1987) and Stock and Watson (1988). However, routine adjustment of hundreds of thousands of series will probably continue to be based on univariate filters for quite a few years.

Similarly to the case of forecasting, another obvious research direction is the extension of unobserved component models to nonlinear time series. It is the case, for example, that nonlinearity often affects seasonal frequencies, and hence should be taken into consideration when estimating seasonally adjusted series. Efforts in this direction are Kitagawa (1987), Harvey *et al.* (1992) and Fiorentini and Maravall (1996).

On a related, more practical and more important, issue, there is a forecast that I would really like to see realized: it concerns the practice of some data-producing agencies and companies of providing only seasonally adjusted data. We now know that seasonal adjustment, apart from the many problems pointed out by researchers (from Wallis 1974 to Ghysels and Perron 1993), seriously limits the usefulness of some of the most basic and important econometric tools. In particular, it forces us to work with non-invertible series, and hence autoregressive representations of the series, for example, are not appropriate (see Maravall 1995). By having to use the heavily distorted and distorting seasonally adjusted series, life for the economist is made unnecessarily difficult. The wish is thus that the damaging practice of making available only adjusted data will cease; the original, unadjusted data, should always be available.

[3] EUROSTAT (1996) presents the results of the automatic model identification procedure of the TRAMO program (see Gomez and Maravall 1996) for close to 15,000 series of economic indicators (all activities) in the 15 EU countries, USA, and Japan. They report that for 87% of the series the fit was good, for 11% it was reasonably acceptable, and for 2% only the fit was poor. They conclude that 'automatic modelization works much better than expected'.

To complete my statements about the future, I should add a last one, well known to any one who has been involved in actual forecasting: no matter what I may say, my forecast will most certainly be wrong.

12.8. Final Comment: Limitations of the Model-Based Approach

From the previous discussion, it would seem as if the use of a model-based approach is a panacea that will permit us to obtain proper answers to all questions. Yet this panacea is not a well-defined one. What do we really mean by a model? Ultimately the models we use are not properties of an objective reality that we manage to approximate, but figments of the researcher mind. In particular, the proper model to use can be defined only in terms of the problems one wishes to analyse. In this context, ARIMA models were devised for short-term analysis, yet they have been borrowed to deal with many other applications. We shall concentrate on one of these applications: the efforts by macroeconomists to measure the business cycle and analyse the behaviour of aggregate output. One of the directions of this research has been the attempt to measure the long-term effects of shocks to GNP and, in particular, to answer the question, does a unit innovation in GNP have a permanent effect on the level of GNP? This long-term effect of a unit innovation has been denoted 'persistence'. If x_t = log GNP follows the I(1) model,

$$\nabla x_t = \psi(B)\, a_t,$$

then the measure of persistence, m, can be defined as the effect of a unit innovation on the long-term forecast of x_t, or

$$m = \lim_{k\to\infty} (E_t x_{t+k} - E_{t-1} x_{t+k}) = \lim_{k\to\infty} \sum_{j=0}^{k} \psi_j = \psi(1), \tag{10}$$

since $a_t = 1$. Indeed, different values of m have been attributed to competing theories on the business cycle. Specifically, if $m > 1$, real factors, associated mostly with supply, would account for the business cycle; on the contrary, $m < 1$ would indicate that transitory, demand-type shocks play an important role in the generation of cycles. Whether the business cycle is driven by demand or by supply shocks has very different and important policy implications.

The standard procedure for estimating m has been to specify a model, then to fit it by a maximum likelihood (ML) or some least squares (LS) criteria, and then to use the parameter estimates for inference. We consider the quarterly series of US GNP.[4] For our purposes, a reasonable model is given by

$$\nabla x_t = (1 - \theta B)\, a_t + \mu, \tag{11}$$

[4] The series is the same as in Evans (1989): it consists of 144 observations.

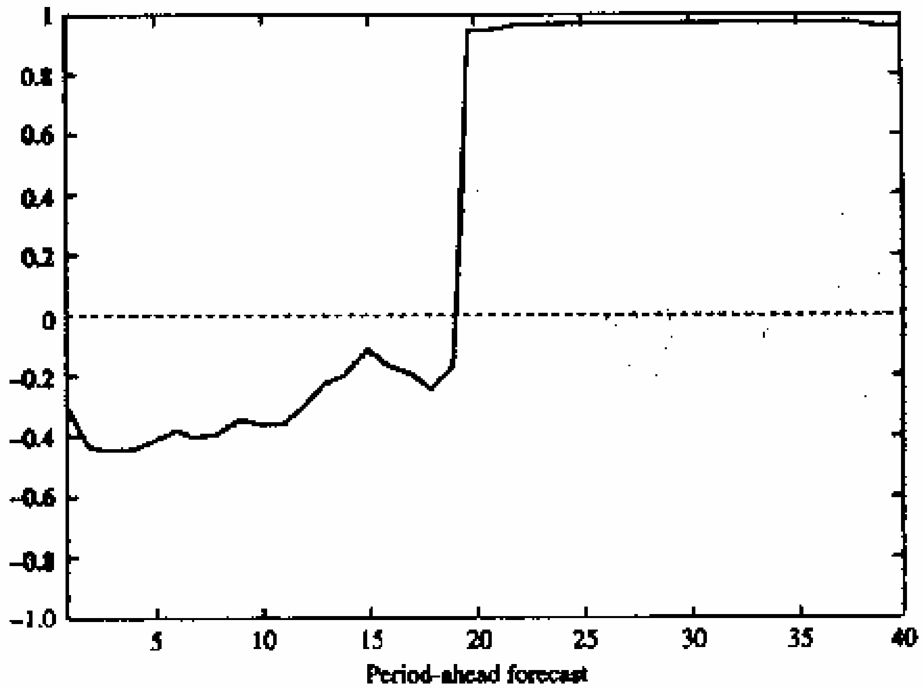

Fig. 12.8. Parameter Estimator as a Function of the Forecast Horizon

where μ is a constant. ML estimation yields $\hat{\theta} = -0.30$, and the residuals obtained seem to behave as white-noise. The measure of persistence becomes thus

$$\hat{m} = 1.30 \quad (\text{s.e.} = 0.08),$$

and hence it can be concluded that $m > 1$. This is in agreement with many univariate estimates of m found in the literature; see e.g. Campbell and Mankiw (1987).

Broadly, an ML or an LS criterion minimizes the sum of squares (SS) of the residuals a, or, approximately, the SS of the one-period-ahead (1 − p.a.) forecast error. Why select the 1 − p.a. forecast? If our interest is the long-run, and this is certainly the case when measuring persistence, why not minimize a long-term forecast error? Since models are always simplifications which imply some degree of misspecification, it is a well-known fact that minimizing the SS of the 1 − p.a. forecast error may yield parameter estimates that differ substantially from those that minimize the SS of the k − p.a. forecast (for k not close to 1). Some references are Gersch and Kitagawa (1983), Findley (1984), and Tiao and Xu (1992).

For our example, let $\hat{\theta}(k)$ denote the estimator that minimizes the SS of the (in-sample) k − p.a. forecast errors. Fig. 12.8 displays $\hat{\theta}(k)$ as a function of k. For $k < 20$ periods, the estimator fluctuates between −0.2 and −0.4; then it jumps quickly to 0.94, and for $k \geq 20$ it remains basically unchanged around that value. Curiously enough, the sample information seems to discriminate two values for θ:

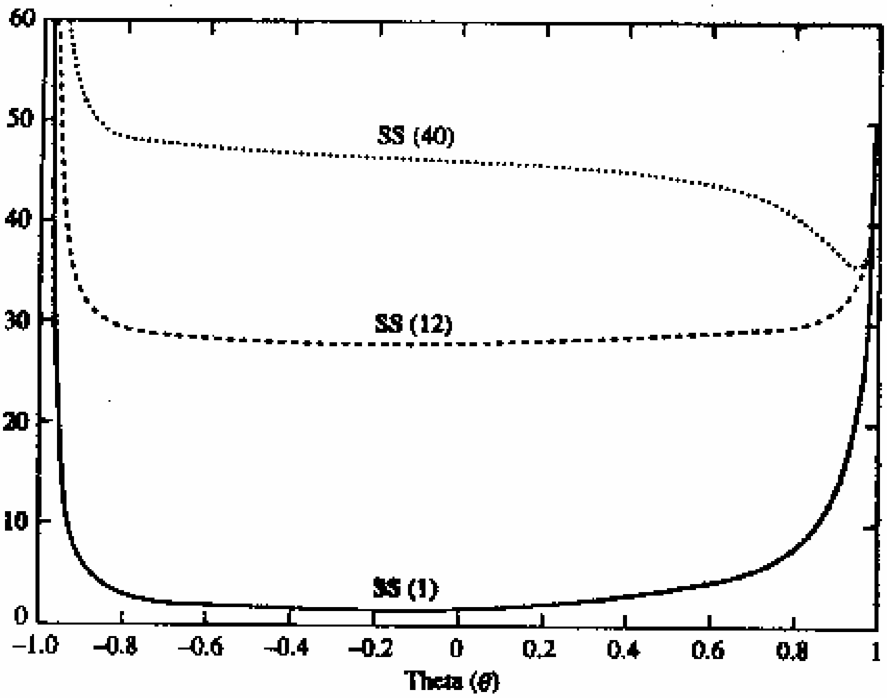

Fig. 12.9. Sum of Squares Function for Increasing the Forecasting Horizon

one for the short run (less than a five-year horizon) and one for the long run (more than a five-year horizon). A closer look at the behaviour of the SS function—Fig. 12.9—shows that for low values of k a minimum is found for a small, negative value of θ. Then, for the intermediate values $k = 10 - 20$, the SS function becomes very flat. As k becomes larger, the minimum for $\theta = 0.95$ becomes more and more pronounced. A similar behaviour of $\hat{\theta}(k)$ is obtained by Tiao and Xu (1992). This behaviour is, in fact, quite general. In Fabrizio (1995) the previous analysis is applied to the annual series of real GDP indices for 12 OECD countries in Maddison (1991). Again, they all display a roughly similar behaviour: $\hat{\theta}(k)$ is relatively low for k below some threshold, and $\hat{\theta}(k)$ becomes close to 1 for k above that threshold. Be that as it may, I find it intriguing that analysis of the data overwhelmingly produces this sharp and sudden distinction between short and long-term forecast.

If we compute the percentage increase in the MSE of the forecast from using the ML estimator instead of $\hat{\theta}(k)$, it is found that, for $k < 20$, that percentage is negligible (in line with the results in Weiss 1991, who considers values of $k \leq 4$). For $k = 24$, use of the ML estimator $\hat{\theta}(1)$ increases the MSE by 14 per cent; for $k = 32$, this percentage becomes 26 per cent, and for $k = 40$ it goes up to 31 per cent. Therefore, if our aim is the long-term forecast, it would seem quite inefficient to use as parameter $\hat{\theta}(1) = -0.30$: our MSE may deteriorate by more than 30 per cent.

It is easy to give an intuitive explanation for the behaviour of $\hat{\theta}(k)$. The good performance of ARIMA models is a result of their flexibility to adapt their

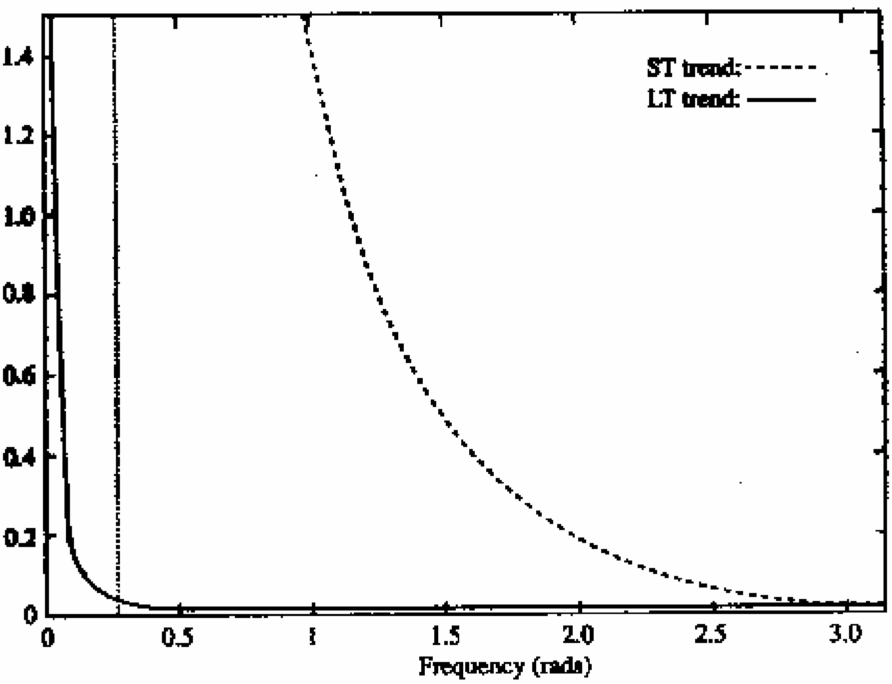

Fig. 12.10. Trend Spectrum: Short and Long-Term Models (five-year-period frequency)

forecast function to the short-run variability. Long-run extrapolation of this short-run flexibility will introduce too much noise in the long-term forecast.

Again, a look at the model components will prove helpful. The IMA(1, 1) model of (8) can be seen as the sum of a trend and an orthogonal white-noise component, where the variance of the noise component can take any value in the interval $(0, (1 + \theta)^2 V_a/4)$; see Box *et al.* (1978). In the trend-plus-noise decomposition of (8), the forecast of the series is the same as the forecast of the trend. The two models obtained by setting $\theta = -0.3$ and $\theta = 0.95$ will imply very different trend components. To compare them, for each of the two parameters, we select the decomposition of (8) that sets the variance of the noise equal to its maximum; this is the so-called canonical decomposition, and it maximizes the smoothness of the trend. For the two canonical decompositions corresponding to $\theta = -0.3$ and $\theta = 0.95$, the variances of the innovations in the trend component are 0.423 V_a and 0.001 V_a, respectively. When $\theta = 0.95$, therefore, the trend contains very little stochastic variability. The two spectra are compared in Fig. 12.10. This comparison shows that the trend implied by the model that is optimal for long-term forecasting is very stable, and picks up only very small frequencies. This is a sensible result: when interested in short-term analysis, we look at the month-to-month or quarter-to-quarter forecasts. Thus, for example, the variability of the series associated in the spectrum with the frequency corresponding to a period of five years should be a part of the forecast and of the trend. However, if we are forecasting twenty years ahead, the variance of the series corresponding to a five-year cycle should not be considered, and hence should not be a part of the trend: in twenty years, the (damped) five-year cycle has practically disappeared. This is precisely what Fig. 12.10 tells us. When looking at the long run, only movements in the series associated with very large periods, i.e. very small frequencies, are of interest. Fig. 12.11 compares the two (short-term and long-term) trend estimates and the two associated estimates of the noise component. The short-run noise reflects the estimator of a white-noise variable; the long-run noise instead allows for larger effects, since over a long span of time they approximately cancel out.

If the measure of persistence, which attempts to measure the effect of a shock on the very long-term forecast, is based on the model optimal for long-term forecasting, then

$$\hat{m} = 1 - 0.95 = 0.05,$$

quite different from the measure obtained before, and certainly below one. Yet the point is not to claim that this result points towards a business cycle dominated by demand shocks. The way I read it, the conclusion is that the trend model obtained in the standard ML estimation–ARIMA specification approach makes sense only for relatively short-term analysis. It is with this type of analysis that I have here been concerned, and it seems to me that the short-term tools we use may not be appropriate for long-term inference.

REFERENCES

Bach, G. L., Cagan, P. D., Friedman, M., Hildreth, C. G., Modigliani, F., and Okun, A. (1976), *Improving the Monetary Aggregates: Report of the Advisory Committee on Monetary Statistics*. Washington, DC: Board of Governors of the Federal Reserve System.

Beguin, J. M., Gourieroux, C., and Monfort, A. (1980), 'Identification of a Mixed Autoregressive-Moving Average Process: the Corner Method', in O. D. Anderson (ed.), *Time Series* (Proceedings of a March 1979 Conference). Amsterdam: North-Holland.

Bell, W. R. (1984), 'Signal Extraction for Nonstationary Time Series', *Annals of Statistics*, 12: 646–64.

—— and Hillmer, S. C. (1984), 'Issues Involved with the Seasonal Adjustment of Economic Time Series', *Journal of Business and Economic Statistics*, 2: 291–320.

Box, G. E. P., and Jenkins, G. M. (1970), *Time Series Analysis: Forecasting and Control*. San Francisco: Holden-Day.

—— and Tiao, G. C. (1975), 'Intervention Analysis with Applications to Economic and Environmental Problems', *Journal of the American Statistical Association*, 70: 71–9.

—— Hillmer, S. C., and Tiao, G. C. (1978), 'Analysis and Modeling of Seasonal Time Series', in A. Zellner (ed.), *Seasonal Analysis of Economic Time Series*. Washington, DC: US Department of Commerce, Bureau of the Census, 309–34.

Brown, R. G. (1962), *Smoothing, Forecasting and Prediction of Discrete Time Series*. Englewood Cliffs, NJ: Prentice-Hall.

Burman, J. P. (1980), 'Seasonal Adjustment by Signal Extraction', *Journal of the Royal Statistical Society* A, 143: 321–37.

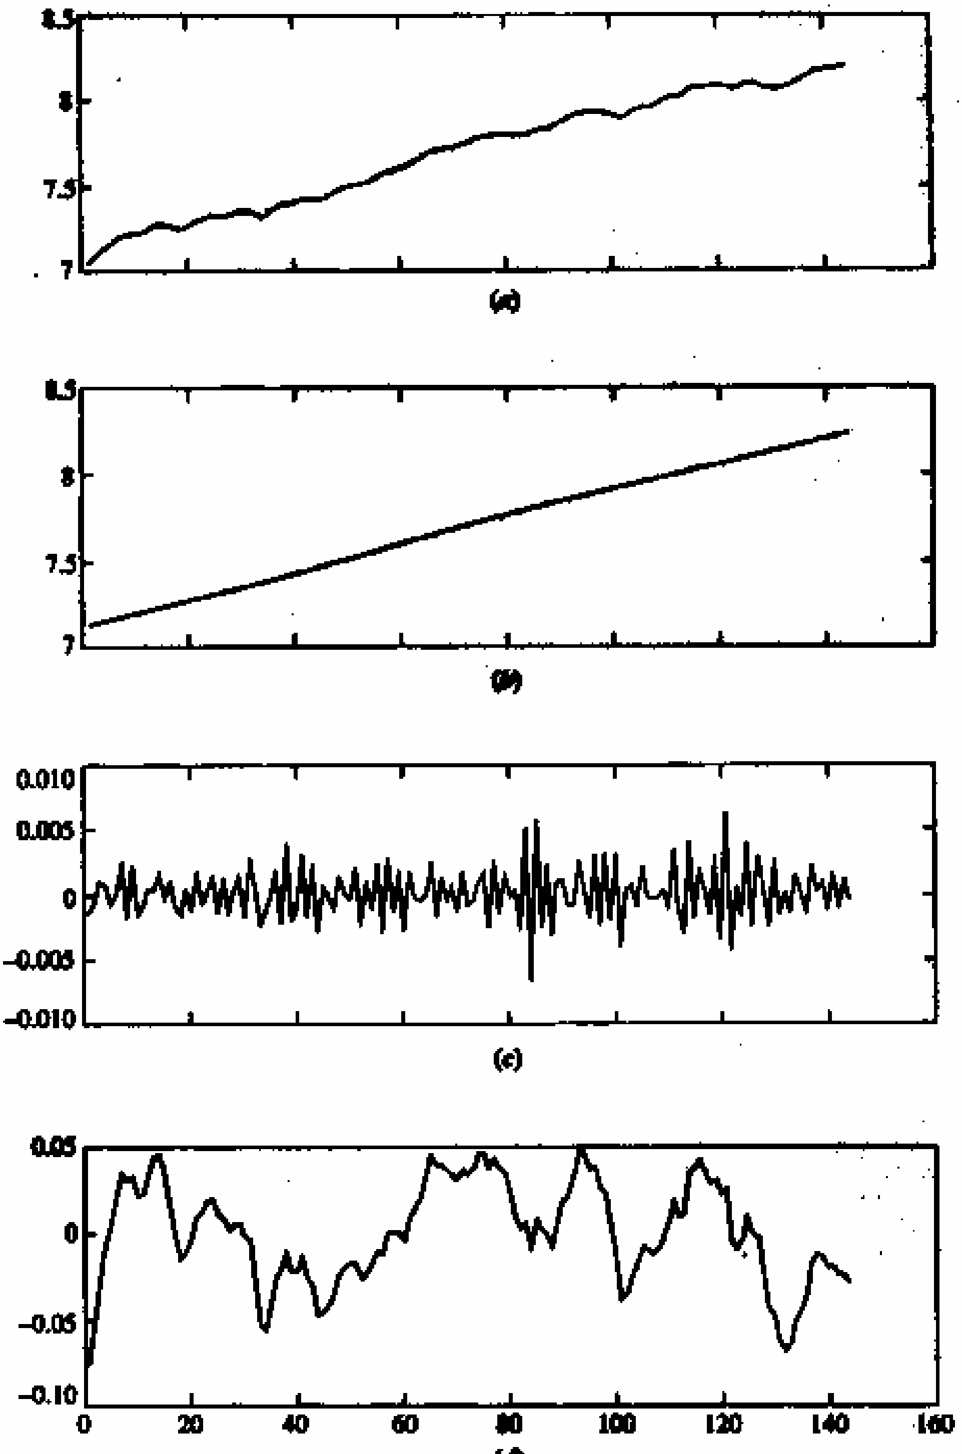

Fig. 12.11. (*a*) Short-run Trend
(*b*) Long-run Trend
(*c*) Short-run Noise
(*d*) Long-run Noise

Burridge, P., and Wallis, K. F. (1984), 'Unobserved Components Models for Seasonal Adjustment Filters', *Journal of Business and Economic Statistics*, 2: 350–9.

Campbell, J. Y., and Mankiw, N. G. (1987), 'Permanent and Transitory Components in Macro-economic Fluctuations', *American Economic Review Proceedings*, 1987: 111–17.

Cleveland, W. P., and Tiao, G. C. (1976), 'Decomposition of Seasonal Time Series: A Model for the X-11 Program', *Journal of the American Statistical Association*, 71: 581–7.

Cox, D. R. (1961), 'Prediction by Exponentially Weighted Moving Averages and Related Methods', *Journal of the Royal Statistical Society* B, 23: 414–22.

Dagum, E. B. (1980), 'The X11 ARIMA Seasonal Adjustment Method', Statistics Canada, Catalogue 12–564E.

Den Butter, F. A. G., and Fase, M. M. G. (1991), *Seasonal Adjustment as a Practical Problem*. Amsterdam: North-Holland.

Engle, R. F. (1978), 'Estimating Structural Models of Seasonality', in A. Zellner (ed.), *Seasonal Analysis of Economic Time Series*, Washington, DC: US Department of Commerce, Bureau of the Census, 281–97.

Evans, G. W. (1989), 'Output and Unemployment Dynamics in the United States: 1950–1985', *Journal of Applied Econometrics*, 4: 213–37.

Eurostat (1996), 'Is the ARIMA Model Adapted to Economic Time Series?' Presented at the Informal Group on Seasonal Adjustment, Eurostat/D3/SA/04.

Fabrizio, S. (1995), 'The Measurement of Persistence of Economic Shocks', Ph.D. Dissertation, Department of Economics, European University Institute.

Fernandez-Macho, F., Harvey, A. C., and Stock, J. H. (1987), 'Forecasting and Interpolation using Vector Autoregressions with Common Trends', *Annales d'Economie et de Statistique*, 6/7: 279–88.

Findley, D. F. (1984), 'On Some Ambiguities Associated with the Fitting of ARMA Models to Time Series', *Journal of Time Series Analysis*, 5: 217–27.

—— Monsell, B. C., Bell, W. R., Otto, M. C., and Chen, S. (1998), 'New Capabilities and Methods of the X12ARIMA Seasonal Adjustment Program', *Journal of Business and Economics Statistics*, 16, 127–52.

Fiorentini, G. and Maravall, A. (1996), 'Unobserved Components in Arch Models: An Application to Seasonal Adjustment', *Journal of Forecasting*, 15, 175–201.

Gersch, W., and Kitagawa, G. (1983), 'The Prediction of Time Series with Trends and Seasonalities', *Journal of Business and Economic Statistics*, 1: 253–64.

Ghysels, E., and Perron, P. (1983), 'The Effect of Seasonal Adjustment Filters on Tests for a Unit Root', *Journal of Econometrics*, 55: 57–98.

Gomez, V. and Maravall, A. (1998), 'Automatic Modeling Methods for Univariate Series', Working Paper 9808, Research Department, Bank of Spain.

—— and Maravall, A. (1996), 'Programs TRAMO (Time series Regression with Arima noise, Missing observations, and Outliers) and SEATS (Signal Extraction in Arima Time Series). Instructions for the User, Working Paper 9628, Servicio de Estudios, Banco de España.

Gouriéroux, C. (1992), *Modeles ARCH et Applications Financieres*, Paris: Económica.

Hamilton, J. D. (1994), *Time Series Analysis*, Princeton, NJ: Princeton University Press.

Harrison, P. J., and Stevens, C. F. (1976), 'Bayesian Forecasting', *Journal of the Royal Statistical Society* B, 38: 205–47.

Harvey, A. C. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.

—— and Todd, P. H. J. (1983), 'Forecasting Economic Time Series with Structural and Box–Jenkins Models: a Case Study', *Journal of Business and Economic Statistics*, 1: 299–306.

—— Ruiz, E., and Sentana, E. (1992), 'Unobserved Component Time Series Models with ARCH Disturbances', *Journal of Econometrics*, 52: 129–57.

——, ——, and Shephard, N. G. (1994), 'Multivariate Stochastic Variance Models', *Review of Economic Studies*, 61, 247–64.

Hillmer, S. C., and Tiao, G. C. (1982), 'An ARIMA-Model-Based Approach to Seasonal Adjustment', *Journal of the American Statistical Association*, 77: 63–70.

Hylleberg, S. (1986), *Seasonality in Regression*. New York: Academic Press.

—— (ed.) (1992), *Modeling Seasonality*. New York: Oxford University Press.

Johansen, S. (1996), *Likelihood Based Inference in Cointegrated Vector Autoregressive Models*, Oxford: Oxford University Press.

Kenny, P., and Durbin, J. (1982), 'Local Trend Estimation and Seasonal Adjustment of Economic and Social Time Series', *Journal of the Royal Statistical Society* A, 145: 1–28.

Kitagawa, G. (1987), 'Non-Gaussian State Space Modeling of Nonstationary Time Series', *Journal of the American Statistical Association*, 82: 1032–63.

Litterman, R. B. (1986), 'A Statistical Approach to Economic Forecasting', *Journal of Business and Economic Statistics*, 4: 1–5.

Maddison, A. (1991), *Dynamic Forces in Capitalist Development: a Long-Run Comparative View*. Oxford: Oxford University Press.

Maravall, A. (1985), 'On Structural Time Series Models and the Characterization of Compon-ents', *Journal of Business and Economic Statistics*, 3: 350–5.

—— (1987), 'On Minimum Mean Squared Error Estimation of the Noise in Unobserved Component Models', *Journal of Business and Economic Statistics*, 5: 115–20.

—— (1995), Unobserved Components in Economic Time Series, in H. Pesaran, P. Schmidt, and M. Wickens (eds.), *The Handbook of Applied Eco-nometrics*, i. Oxford: Basil Blackwell.

—— and Pierce, D. A. (1986), 'The Transmission of Data Noise into Policy Noise in US Monetary Control', *Econometrica*, 54: 961–79.

—— —— (1987), 'A Prototypical Seasonal Adjustment Model', *Journal of Time Series Analysis*, 8: 177–93.

Mills, T. C. (1990), *Time Series Techniques for Economists*. Cambridge: Cambridge University Press.

Moore, G. H., Box, G. E. P., Kaitz, H. B., Stephenson, J. A. and Zellner, A. (1981), *Seasonal Adjustment of the Monetary Aggregates: Report of the Committee of Experts on Seasonal Adjustment Techniques*. Washington, DC: Board of Governors of the Federal Reserve System.

Nerlove, M., Grether, D. M., and Carvalho, J. L. (1979), *Analysis of Economic Time Series: A Synthesis*. New York: Academic Press.

Pierce, D. A. (1980), 'Data Revisions in Moving Average Seasonal Adjustment Procedures', *Journal of Econometrics*, 14: 95–114.

Shiskin, J., Young, A. H., and Musgrave, J. C. (1967), 'The X11 Variant of the Census Method II Seasonal Adjustment Program', Washington, DC: Bureau of the Census, Technical Paper 15.

Stock, J. H., and Watson, M. W. (1988), 'Testing for Common Trends', *Journal of the American Statistical Association*, 83: 1097–1107.

Tiao, G. C. and Tsay, R. S. (1989), 'Model Specification in Multivariate Time Series', *Journal of the Royal Statistical Society B*, 51, 157–213.

—— and Xu, D. (1992), 'Robustness of MLE for Multi-Step Predictions: the Exponential Smoothing Case', Technical Report 117, Statistics Research Center, Graduate School of Business, University of Chicago.

Tsay, R. S., and Tiao, G. C. (1984), 'Consistent Estimates of Autoregressive Parameters and Extended Sample Autocorrelation Function for Stationary and Non-Stationary ARMA Models', *Journal of the American Statistical Association*, 79: 84–96.

Wallis, K. F. (1974), 'Seasonal Adjustment and Relations between Variables', *Journal of the American Statistical Association*, 69: 18–31.

Watson, M. W. (1986), 'Univariate Detrending Methods with Stochastic Trends', *Journal of Monetary Economics*, 18, 49–75.

Weiss, A. A. (1991), 'Multi-Step Estimation and Forecasting in Dynamic Models', *Journal of Econometrics*, 48: 135–49.

Whittle, P. (1963), *Prediction and Regulation by Linear Least-Square Methods*. London: English Universities Press.

Winters, P. R. (1960), 'Forecasting Sales by Exponentially Weighted Moving Averages', *Management Science*, 6: 324–42.

Wodehouse, P. G. (1981), *Wodehouse on Wodehouse*. London: Penguin Books.