

ARMA Models and the Box–Jenkins Methodology

SPYROS MAKRIDAKIS* AND MICHÈLE HIBON

INSEAD, France

ABSTRACT

The purpose of this paper is to apply the Box–Jenkins methodology to ARIMA models and determine the reasons why in empirical tests it is found that the post-sample forecasting the accuracy of such models is generally worse than much simpler time series methods. The paper concludes that the major problem is the way of making the series stationary in its mean (i.e. the method of differencing) that has been proposed by Box and Jenkins. If alternative approaches are utilized to remove and extrapolate the trend in the data, ARMA models outperform the models selected through Box–Jenkins methodology. In addition, it is shown that using ARMA models to seasonally adjusted data slightly improves post-sample accuracies while simplifying the use of ARMA models. It is also confirmed that transformations slightly improve post-sample forecasting accuracy, particularly for long forecasting horizons. Finally, it is demonstrated that AR(1), AR(2) and ARMA(1,1) models can produce more accurate post-sample forecasts than those found through the application of Box–Jenkins methodology. © 1997 by John Wiley & Sons, Ltd.

J. forecast. **16:** 147–163, 1997

No. of Figures: 12. No. of Tables: 0. No. of References: 44.

KEYWORDS time-series forecasting; ARMA models; Box–Jenkins, empirical studies; M-Competition

AutoRegressive (AR) models were first introduced by Yule in 1926. They were subsequently supplemented by Slutsky who in 1937 presented Moving Average (MA) schemes. It was Wold (1938), however, who combined both AR and MA schemes and showed that ARMA processes can be used to model a large class of stationary time series as long as the appropriate order of p , the number of AR terms, and q , the number of MA terms, was appropriately specified. This means that a general series x_t can be modelled as a combination of past x_t values and/or past e_t errors, or

$$x_t = \phi_1 x_{t-1} + \phi_2 x_{t-2} + \cdots + \phi_p x_{t-p} + e_t - \theta_1 e_{t-1} - \theta_2 e_{t-2} - \cdots - \theta_q e_{t-q} \quad (1)$$

Using equation (1) for modelling real-life time series requires four steps. First the original series, x_t , must be transformed to become stationary around its mean and its variance. Second, the

* Correspondence to: Spyros Makridakis, INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France

appropriate order of p and q must be specified. Third, the value of the parameters $\phi_1, \phi_2, \dots, \phi_p$ and/or $\theta_1, \theta_2, \dots, \theta_q$ must be estimated using some non-linear optimization procedure that minimizes the sum of square errors or some other appropriate loss function. Finally, practical ways of modelling seasonal series must be envisioned and the appropriate order of such models specified.

The utilization of the theoretical results suggested by Wold, expressed by equation (1), to model real-life series did not become possible until the mid-1960s when computers, capable of performing the required calculations to optimize the parameters of equation (1), became available and economical. Box and Jenkins (1976, original edition 1970) popularized the use of ARMA models through the following: (1) providing guidelines for making the series stationary in both its mean and variance, (2) suggesting the use of autocorrelations and partial autocorrelation coefficients for determining appropriate values of p and q (and their seasonal equivalent P and Q when the series exhibited seasonality), (3) providing a set of computer programs to help users identify appropriate values for p and q , as well as P and Q , and estimate the parameters involved and (4) once the parameters of the model were estimated, a diagnostic check was proposed to determine whether or not the residuals, e_t , were white noise, in which case the order of the model was considered final (otherwise another model was determined using (2) and steps (3) and (4) were repeated). If the diagnostic check showed random residuals then the model developed was used for forecasting or control purposes assuming, of course, constancy, that is, the order of the model and its non-stationary behaviour, if any, would remain the same during the forecasting, or control, phase.

The approach proposed by Box and Jenkins came to be known as the Box–Jenkins methodology to ARIMA models, where the letter ‘I’, between AR and MA, stood for the ‘Integrated’ and reflected the need for differencing to make the series stationary. ARIMA models and the Box–Jenkins methodology became highly popular in the 1970s among academics, in particular when it was shown through empirical studies (Cooper, 1972; Nelson, 1972; Elliot, 1973; Narasimham *et al.*, 1974; McWhorter, 1975; for a survey see Armstrong, 1978) that they could outperform the large and complex econometric models, popular at that time, in a variety of situations.

EMPIRICAL EVIDENCE

The popularity of the Box–Jenkins methodology to ARIMA models was shaken when empirical studies (Groff, 1973; Geurts and Ibrahim, 1975; Makridakis and Hibon, 1979; Makridakis *et al.*, 1982, 1993; Huss, 1985; Fildes *et al.*, 1997), using real data, showed that simple methods were equally or more accurate than Box–Jenkins when post-sample comparisons were made. Today, after many debates, it is accepted by a large number of researchers that in empirical tests Box–Jenkins is not an accurate method for post-sample time-series forecasting, at least in the domains of business and economic applications where the level of randomness is high and where constancy of pattern, or relationships, cannot be assured.

The purpose of this paper is to examine the post-sample forecasting accuracy of ARIMA models in order to determine the contribution to such accuracy of each of its elements (i.e. seasonality, stationarity, the order of the ARMA model, and the necessity that the residuals of the ARMA model must be random). It is concluded that the major problem with the Box–Jenkins methodology is the way that the series are made stationary in their mean. When




Figure 1. Schematic representation of the Box–Jenkins methodology. (Source: Makridakis *et al.*, 1983)

alternative ways of dealing with and extrapolating the trend are provided, ARMA models are slightly more accurate than the corresponding time-series methods that extrapolate the trend in the time series.

THE STEPS (ELEMENTS) OF THE BOX–JENKINS METHODOLOGY

Figure 1 presents the four steps of the Box–Jenkins methodology. This section examines each of these steps and discusses its possible contribution to post-sample forecasting accuracy.

Stationarity

Before equation (1) can be used the series should be stationary in its mean and variance. The Box–Jenkins methodology suggests short and seasonal (long) differencing to achieve stationarity in the mean, and logarithmic or power transformation to achieve stationarity in the variance. The value of both differencing and transformations have been questioned. Pierce (1977) argued that differencing is not an appropriate way of making the data stationary and instead he proposed linear detrending. Nelson and Plosser (1982) argued that some series could be better made stationary through differencing and others through linear detrending. Others (Parzen, 1982; Newton and Parzen, 1984; Meese and Geweke, 1984) have used a pre-filter consisting of a long-memory AR model to capture possible non-stationarity in the series before using a regular ARMA model.

Box and Jenkins suggest logarithmic or power transformations to achieve stationarity in the variance. The value of such transformations to improve post-sample forecasting accuracy has also been debated and no agreement has been reached as to whether or not transformations are helpful (Chatfield and Prothero, 1973). At the same time, it is clear that transformations require personal judgement and the possibility of making errors, even when utilized by high-level academic experts

(see comments on the paper by Chatfield and Prothero, 1973). At the empirical level there is also no evidence that logarithmic or power transformations improve post-sample forecasting accuracy (Granger and Nelson, 1978; Makridakis and Hibon, 1979; Meese and Geweke, 1984).

Seasonality

In case the series are seasonal, the Box–Jenkins methodology proposes multiplicative seasonal models coupled with long-term differencing, if necessary, to achieve stationarity in the mean. The difficulty with such an approach is that there is practically never enough data available to determine the appropriate level of the seasonal ARMA model with any reasonable degree of confidence. Users therefore proceed through trial and error in both identifying an appropriate seasonal model and in selecting the correct long-term (seasonal) differencing. In addition, seasonality requires more data to estimate the appropriate model parameters. There has apparently been no empirical work to test whether or not deseasonalizing the data first, using a decomposition procedure (a suggestion made by Durbin, 1979), and subsequently using the Box–Jenkins method on the seasonally adjusted data improves post-sample forecasting accuracy.

Order of ARMA model

The order of the ARMA model is found by examining the autocorrelations and partial autocorrelations of the stationary series. Box and Jenkins (1976) provided both a theoretical framework and practical rules for determining appropriate values for p and q as well as their seasonal counterparts P and Q . The possible difficulty is that often more than one model could be considered, requiring the user to choose one of them without any knowledge of the implications of his or her choice on post-sample forecasting accuracy since, according to the Box–Jenkins methodology, any model which results in random residuals is an appropriate one. Box and Jenkins do recommend the principle of parsimony meaning that a simpler (having fewer parameters) model should be selected in case more than one model is possible, but there has been very little systematic work to determine if this suggestion results in improvements in post-sample forecasting accuracy.

Estimating the model's parameters

This part of the Box–Jenkins methodology is the most straightforward one. The non-linear optimization procedure, based on the method of steepest descent (Marquardt, 1963), is used to estimate the parameter values of p and/or q (and their seasonal equivalent P and/or Q). Apart from occasional problems when there is no convergence (in which case another model is used) the estimation provides no special difficulties except for its inability to guarantee a global optimum (a common problem of all non-linear algorithms). The estimation is completely automated requiring no judgemental inputs, and therefore testing, as all computer programs use the same algorithm in applying the Marquardt optimization procedure.

Diagnostic checks

Once an appropriate model had been chosen and its parameters estimated, the Box–Jenkins methodology required examining the residuals of the actual values minus those estimated through the model. If such residuals are random, it is assumed that the model is appropriate. If not, another model is considered, its parameters estimated, and its residuals checked for randomness. In practically all instances a model could be found to result in random residuals. Several tests (e.g. Box and Pierce, 1970) have been suggested to help users determine if overall the

residuals are indeed random. Although it is a standard statistical procedure not to use models whose residuals are not random, it might be interesting to test the consequences of lack of residual randomness on post-sample forecasting accuracy.

POST-SAMPLE FORECASTING ACCURACY: PERSONALIZED VERSUS AUTOMATIC BOX-JENKINS

The Makridakis and Hibon (1979) study, the M-Competition (Makridakis *et al.*, 1982), the M2-Competition (1993) as well as many other empirical studies (Schnaars, 1986; Koehler and Murphree, 1988; Geurts and Kelly, 1986; Watson *et al.*, 1987; Collopy and Armstrong, 1992) have demonstrated that simple methods such as exponential smoothing outperform, on average, the Box-Jenkins methodology to ARMA models. Figures 2(a), (b) and (c) show the MAPE (Mean Absolute Percentage Errors), for various forecasting horizons, of Naive 2 (a deseasonalized random walk model), Single exponential smoothing (after the data have been deseasonalized, if necessary, and the forecasts subsequently reseasonalized) as well as those of the Box-Jenkins methodology found in the M2-Competition (Makridakis *et al.*, 1993), the M-Competition (Makridakis *et al.*, 1982), and the Makridakis and Hibon (1979) study. The results indicate that single smoothing outperformed 'Box-Jenkins' overall and in most forecasting horizons, while Naive 2 also does better than 'Box-Jenkins', although by a lesser amount. The results of Figure 2 are surprising since it has been demonstrated that single exponential smoothing is a special case of ARMA models (Cogger, 1974; Gardner and McKenzie, 1985). Moreover, it makes little sense that Naive 2, which simply uses the latest available value, taking seasonality into account, does so well in comparison to the Box-Jenkins methodology.

In the M-Competition (Makridakis *et al.*, 1982) the Box-Jenkins method was run on a subset of 111 (one out of every nine) series from the total of the 1001 series utilized. The reason was that the method required personal judgement, making it impractical to use all 1001 series as the expert analyst had to model each series individually, following the various steps described in the previous section, and spending, on average, about one hour before a model could be confirmed as appropriate for forecasting purposes (see Andersen and Weiss, 1984).

Since the M-Competition was completed, several empirical studies have shown that automatic Box-Jenkins approaches (Hill and Fildes, 1984; Libert, 1983; Texter and Ord, 1989) performed about the same or better in terms of post-sample accuracy as the personalized approach followed by Andersen and Weiss (1984). Figure 3 shows the results of a specific automatic Box-Jenkins program (Stellwagen and Goodrich, 1991) together with that of the personalized approach utilized by Andersen and Weiss in the M-Competition. Figure 3 illustrates that for the 111 series used in the comparison the post-sample accuracies of the automatic and personalized approaches are about the same. We could, therefore, use an automatic Box-Jenkins version to perform our comparisons using all the 1001 series of the M-Competition.

ATTRIBUTING THE DIFFERENCES IN POST-SAMPLE FORECASTING ACCURACIES

Since we found no substantive differences between the personalized and automatic versions of Box-Jenkins (see Figure 3), we have utilized *all* the 1001 series of the M-Competition using an




Figure 2. The post-sample forecasting accuracy of Box–Jenkins, Naive 2, and single exponential smoothing. (a) M2-Competition; (b) M-Competition; (c) Makridakis and Hibon study

automatic Box–Jenkins procedure (Stellwagen and Goodrich, 1991). Such a large sample of 1001 series allow us to attribute more reliably the differences in post-sample forecasting accuracy to the various elements (steps) of the Box–Jenkins methodology.




Figure 3. Box–Jenkins: mean absolute percentage error (MAPE): automatic versus personalized (M-Competition: 111 series)

DESEASONALIZING THE SERIES

In the discussion of the Makridakis and Hibon (1979) paper it was suggested (Durbin, 1979) that the Box–Jenkins methodology should also be applied to the seasonally adjusted data to determine the effect of seasonality. This suggestion is being tested in this study.

X_t , the original data, can be deseasonalized by dividing it by the seasonal index, S_j , computed through the classical decomposition method (Makridakis *et al.*, 1983), or

$$X'_t = X_t / S_j \quad (2)$$

where S_j is the seasonal index corresponding to the j th month, if the data are monthly, or the j th season if quarterly. If the data are not seasonal, all indices are set to equal 1.

Once the forecasts have been computed using the automatic Box–Jenkins program, they can be reseasonalized by multiplying them by the corresponding seasonal index, or

$$\hat{X}_t = \hat{X}'_t S_j$$

Figure 4 shows the MAPE of the original and deseasonalized versions of the automatic Box–Jenkins using all 1001 series of the M-Competition. Using ARIMA models on the deseasonalized data results in more accurate post-sample forecasts *consistently*, although the differences between the two approaches are *small* and *not statistically significant*. As it is easier and much simpler to apply ARIMA models to the deseasonalized series, this study suggests that at least for the 1001 series of the M-Competition it is preferable to use ARIMA models to seasonally adjusted data.




Figure 4. MAPE: original versus seasonally adjusted data

LOG OR POWER TRANSFORMATIONS

Figure 5 shows the MAPEs when log or power transformations were employed, when necessary, to achieve stationarity in the variance of the original data and the seasonally adjusted data. There is a very small improvement when logarithmic or power transformations are applied to the raw




Figure 5. MAPE: original, deseasonalized and transformed data

data, but the differences are not statistically significant except for horizon 18. However, the differences are consistent and increase as the forecasting horizon becomes longer. This finding is not in agreement with previous ones which have concluded that power or log transformations do *not* improve at all post-sample forecasting accuracy. As transformations improve forecasting accuracy it must be determined whether the extra work required to make these transformations justifies the small improvements found, and whether or not such statistically insignificant improvements (except for horizon 18) will be also found with other series than those of the M-Competition.

TRANSFORMATIONS FOR ACHIEVING STATIONARITY IN THE MEAN

To the approach of differencing suggested by Box and Jenkins (1976) for achieving stationarity in the mean there are several alternatives employing various ways to remove the trend in the data. The trend, T_t , can be modelled as:

$$T_t = f(t) \quad (3)$$

where $t = 1, 2, 3, \dots, n$. In the case of a linear trend equation (3) becomes

$$T_t = a + bt \quad (4)$$

where a and b are sample estimates of the linear regression coefficients α and β in

$$T_t = \alpha + \beta t + u_t$$

where u_t is an independent, normally distributed error term with zero mean and constant variance. Alternatively, other types of trends can be assumed, or various pre-filters can be applied for removing the trend.

Whatever the approach being followed, T_t can be computed and subsequently used to achieve stationarity assuming an additive

$$x_t = X_t - \hat{T}_t \quad (5)$$

or multiplicative trend

$$x_t = \frac{X_t}{\hat{T}_t} \quad (6)$$

Figure 6 shows the forecasts of the data made stationary through differencing (the approach suggested by Box–Jenkins) and that through linear detrending using expression (4). Figure 6 shows that the linear trend is slightly worse, in terms of post-sample forecasting accuracy, for short forecasting horizons and a little better for longer ones than the method of the first differences. The results of the linear trend improve for forecasting horizons 15 to 18 (see Figure 7), although the differences are small and for most horizons non-statistically significant. This finding suggests that the two approaches produce equivalent results with an improvement of differencing for short forecasting horizons and the opposite holding true for long ones (see Figure 7). This makes sense as differencing better captures short-term trends and linear regression long-term ones.




Figure 6. MAPE: achieving stationarity; linear trend versus differencing




Figure 7. Percentage improvement of linear trend versus differencing

DAMPENING THE TREND

In reality few trends increase or decrease consistently making differencing and linear extrapolation not the most accurate ways of predicting their continuation. For this reason the forecasting literature recommends dampening the extrapolation of trends as a function of their

randomness (see Gardner and McKenzie, 1985). In this study this dampening is achieved in the following four ways:

(1) *Damped exponential trend:*

$$T'_{t+l} = S_t \sum_{i=1}^l \phi^i T'_t$$

$$l = 1, 2, 3, \dots, m$$

where $S_t = \alpha X_t + (1 - \alpha)(S_{t-1} + T'_{t-1})\phi$ and $T'_t = \beta(S_t - S_{t-1}) + (1 - \beta)T'_{t-1}\phi$ and where α , β and ϕ are smoothing parameters found by minimizing the sum of the square errors between the actual values and those predicted by the model forecasts. This method of damped trend has been proposed by Gardner and McKenzie (1985).

(2) *Horizontal extrapolation of the trend through single exponential smoothing:*

$$T'_{t-l} = \alpha X_t + (1 - \alpha)T'_t$$

$$l = 1, 2, 3, \dots, m$$

where α is a smoothing parameter found by minimizing the sum of square errors. The method of single exponential smoothing has been proposed by Brown (1959) and is widely used by business firms and the military.

(3) *The AR(1) extrapolation of the linear trend:* Instead of extrapolations the linear trend of expression (4) as

$$\hat{T}_{t+l} = a + b(n + l) \quad (7)$$

where $l = 1, 2, 3, \dots, m$ we can instead use a pre-filter of the form,

$$\hat{T}_{t+l} = a + b(n + l)\phi^l \quad (8)$$

where ϕ is the AR(1) parameter calculated from the available data. As the value of ϕ is smaller than 1, the trend in expression (8) is dampened depending upon the value of the autoregressive coefficient ϕ .

(4) *The AR(1) extrapolation of the trend found through differencing:* The trend of the latest differencing can be damped by multiplying it by ϕ^l where ϕ is the AR(1) parameter calculated from the available data. In such a case the trend is damped by the exponent l since ϕ is smaller than 1 in absolute value:

$$\hat{T}'_{t+l} = \phi^l \hat{T}_{t+l-1} \quad (9)$$

Figure 8 shows the results of the four methods of dampening the trend versus the method of differencing advocated by Box and Jenkins and the linear trend suggested by Pierce (1977). All




Figure 8. MAPE (a) ARMA with trend of damped smoothing, linear extrapolation and differencing. (b) ARMA with trend of single smoothing, linear trend and differencing. (c) ARMA with AR(1) damped trend, trend of linear extrapolation and differencing. (d) ARMA with trend of AR(1), linear trend extrapolation and differencing

four ways of damped trend outperform the method of differencing consistently and in all but a few exceptions in Figure 8(d) that of the linear trend too. This finding suggests that the approach of achieving stationarity in the mean is crucial and that neither the method of differencing nor that of linear trend are the most accurate ways for doing so. Instead more effective ways of extrapolating the trend must be found through either dampening it or other alternatives using pre-filters.

This finding suggests that the key to more accurate post-sample predictions is the 'I' of the Box–Jenkins methodology to ARIMA models. As statistical theory requires stationarity for applying ARMA models it cannot be blamed for the poor performance in terms of accuracy of models using data which are *not* stationary. Once the trend in ARMA models has been extrapolated in the same way as that of the more accurate of time-series methods, then their post-sample accuracy is superior to those methods, although by a small amount.

USING A RESTRICTED CLASS OF ARMA MODELS

By deseasonalizing the data first we can restrict the class of models being used to five major ones: AR(1), AR(2), MA(1), MA(2) and ARMA(1,1). An alternative to Box–Jenkins methodology is to run all five of them and select the one which minimizes the sum of square errors (SSE) for each series. Figure 9 shows the post-sample forecasting accuracies for AR models and suggests that




Figure 9. MAPE: model selected by the Box–Jenkins methodology and always using AR(1) or AR(2) models




Figure 10. MAPE: model selected by the Box–Jenkins methodology and always using MA(1) or MA(2) models

AR(1) and AR(2) are more accurate than those selected through the Box–Jenkins methodology. Figure 10 shows the post-sample accuracy of the MA(1) and MA(2). It suggests that the two MA models are worse than those selected by the Box–Jenkins methodology for shorter horizons and more accurate for long ones. Finally, Figure 11 shows the post-sample accuracy of always using an ARMA(1,1) model. Such a model produces post-sample accuracies which are superior to




Figure 11. MAPE: model selected by the Box–Jenkins methodology and always using an ARMA(1,1) model

those of the model selected through the Box–Jenkins methodology. This means that insistence on the Box–Jenkins methodology of achieving random residuals before a model is considered appropriate is not necessarily the only alternative to achieving more accurate post-sample forecasts through ARMA models. AR(1), AR(2) and ARMA(1,1) models applied to seasonally adjusted data provide at least as accurate post-sample results as those achieved through the Box–Jenkins methodology. On the other hand, MA models are not as accurate in comparison to AR or ARMA(1,1) models (see Figure 12). The extra advantage of AR(1), AR(2) or ARMA(1,1) models is that they are much easier to apply as they require less effort and computer time. It may be worth-while, therefore, to study the theoretical properties of AR and ARMA(1,1) models to determine why their post-sample accuracies match those of the wider class of ARMA ones. It may also be interesting to determine why the post-sample accuracy of strictly MA models is less accurate than those of AR at least for short- and medium-term forecasting horizons.

CONCLUSIONS

This paper has studied the various aspects of Box–Jenkins methodology applied to ARMA models. The major conclusion has been that the way that the data are made stationary in its mean is the most important factor determining post-sample forecasting accuracies. Most importantly, when the trend in the data is identified and extrapolated using the same procedure as in other methods that have been found to be more accurate in empirical studies than ARMA models perform consistently better than the models selected through the Box–Jenkins methodology. In addition, it was concluded that using seasonally adjusted data improves post-sample accuracies in a small but consistent manner, and that log and power transformations also contributed to small improvements in post-sample accuracies which become more pronounced for long forecasting




Figure 12. MAPE: model selected by the Box–Jenkins methodology and always using AR(1), MA(1), or ARMA(1,1) models

horizons. Finally, it was concluded that AR(1), AR (2) or ARMA(1,1) models produced as accurate post-sample predictions as those found by applying the automatic version of the Box–Jenkins methodology, suggesting that it is neither necessary, as far as post-sample accuracy is concerned, to study the autocorrelations and partial autocorrelations to determine the most appropriate ARMA model, nor to make sure that the residuals of such a model are necessarily random.

REFERENCES

- Andersen, A. and Weiss, A., 'Forecasting: the Box–Jenkins approach', in Makridakis *et al.* (eds), *The Forecasting Accuracy of Major Time Series Methods*, Chichester: John Wiley, 1984.
- Armstrong, J. S., 'Forecasting with econometric methods: folklore versus fact with discussion', *Journal of Business*, **51** (1978) 549–600.
- Box, G. E. P. and Jenkins, G., *Time Series Analysis, Forecasting and Control*, San Francisco, CA: Holden-Day.
- Box, G. E. P. and Pierce, D. A., 'Distribution of the residual autocorrelations in autoregressive-integrated moving-average time series models', *Journal of the American Statistical Association*, **65** (1970), 1509–26.
- Brown, R. G., *Statistical Forecasting for Inventory Control*, New York: McGraw-Hill, 1954.
- Chatfield, C. and Prothero, D. L., 'Box–Jenkins seasonal forecasting: problems in a case study', *Journal of the Royal Statistical Society (A)*, **136** (1973), 295–336.
- Cogger, K. O., 'The optimality of general-order exponential smoothing', *Operations Research*, **22** (1974), 858–67.
- Collopy, F. and Armstrong, J. S., 'Rule-based forecasting', *Management Science*, **38** (1992), 1394–414.
- Cooper, R. L., 'The predictive performance of quarterly econometric models of the United States', in Hickman, B. G. (ed.), *Econometric Models of Cyclical Behavior*, New York: National Bureau of Economic Research, 1972.

- Durbin, J., 'Comment on: Makridakis, S. and Hibon, M., "Accuracy of forecasting: an empirical investigation"', *Journal of the Royal Statistical Society (A)*, **142** (1979), 133–4.
- Elliot, J. W., 'A direct comparison of short-run GNP forecasting models', *Journal of Business*, **46** (1973), 33–60.
- Fildes, R., Hibon, M., Makridakis, S. and Meade, N., 'Generalising about univariate forecasting methods: further empirical evidence', *International Journal of Forecasting*, forthcoming (1997).
- Gardner, E. S. Jr, and McKenzie, E., 'Forecasting trends in time series', *Management Science*, **31** (1985), 1237–46.
- Geurts, M. D. and Ibrahim, I. B., 'Comparing the Box–Jenkins approach with the exponentially smoothed forecasting model application to Hawaii tourists', *Journal of Marketing Research*, **12** (1975), 182–8.
- Geurts, M. D. and Kelly, J. P., 'Forecasting retail sales using alternative models', *International Journal of Forecasting*, **2** (1986), 261–72.
- Granger, C. W. J. and Nelson, H., 'Experience using the Box–Cox transformation when forecasting economic time series', Working Paper 78 (1), Department of Economics, UCSD, 1978.
- Groff, G. K., 'Empirical comparison of models for short-range forecasting', *Management Science*, **20**(1) (1973), 22–31.
- Hill, G. and Fildes, R., 'The accuracy of extrapolation methods: an automatic Box–Jenkins package SIFT', *Journal of Forecasting*, **3** (1984), 319–23.
- Huss, W. R., 'Comparative analysis of company forecasts and advanced time series techniques in the electric utility industry', *International Journal of Forecasting*, **1** (1985), 217–39.
- Jenkins, G. M., 'Some practical aspects of forecasting in organizations', *Journal of Forecasting*, **1** (1982), 3–21.
- Koehler, A. B. and Murphree, E. S., 'A comparison of results from state space forecasting with forecasts from the Makridakis competition', *International Journal of Forecasting*, **4** (1988), 45–55.
- Libert, G., 'The M-competition with a fully automatic Box–Jenkins procedure', *Journal of Forecasting*, **2** (1983), 325–8.
- Lusk, E. J. and Neves, J. O., 'A comparative ARIMA analysis of the 111 Series of the Makridakis competition', *Journal of Forecasting*, **3** (1984), 329–32.
- Makridakis, S. and Hibon, M., 'Accuracy of forecasting: an empirical investigation (with discussion)', *Journal of the Royal Statistical Society (A)*, **142** (1979), Part 2, 79–145 (lead article).
- Makridakis, S. *et al.*, 'The accuracy of extrapolative (time series methods): results of a forecasting competition', *Journal of Forecasting*, **1** (1982), No. 2, 111–53 (lead article).
- Makridakis, S., Wheelwright, S. and McGee, V., *Forecasting Methods and Applications* (2nd edition), New York: John Wiley, 1983.
- Makridakis, S. *et al.*, 'The M-2 Competition: A real-time judgmentally-based forecasting study', *International Journal of Forecasting*, **9** (1993), No. 1, 5–23.
- Marquardt, D. W., 'An algorithm for least squares estimation of non-linear parameters', *Soc. Indust. Appl. Math.*, **11** (1963), 431–41.
- McWhorter, A. Jr, 'Time series forecasting using the Kalman filter: an empirical study', *Proc. Amer. Statist. Ass., Business and Economics Section* (1975), 436–46.
- Meese, R. and Geweke, J., 'A comparison of autoregressive univariate forecasting procedures for macro-economic time series', *Journal of Business and Economic Statistics*, **2** (1984), 191–200.
- Narasimham, ? *et al.*, 'On the predictive performance of the BEA quarterly econometric model and a Box–Jenkins type ARIMA model', *Proceedings of the American Statistical Association: Business and Economics Section* (1974), 501–4.
- Nelson, C. R., 'The prediction performance of the FRB–MIT–PENN model of the US economy', *American Economic Review*, **5** (1972), 902–17.
- Nelson, C. R. and Plosser, C. I., 'Trends and random walks in macroeconomic time series: some evidence and implications', *Journal of Monetary Economics*, **10** (1982), 139–62.
- Newbold, P., 'The competition to end all competitions', in Armstrong, J. S. and Lusk, E. J. (eds), 'Commentary on the Makridakis time series competition (M-Competition)', *Journal of Forecasting*, **2** (1983), 276–79.
- Newton, H. J. and Parzen, E., 'Forecasting and time series model types of 111 economic time series', in Makridakis, S. *et al.* (eds), *The Forecasting Accuracy of Major Time Series Methods*, Chichester: John Wiley, 1984.

- Parzen, E., 'ARARMA models for time series analysis and forecasting', *Journal of Forecasting*, **1** (1982), 67–82.
- Pierce, D. A., 'Relationships—and the lack thereof—between economic time series, with special reference to money and interest rates', *Journal of the American Statistical Association*, **72** (1977), 11–26.
- Schnaars, S. P., 'A comparison of extrapolation models on yearly sales forecasts', *International Journal of Forecasting*, **2** (1986), 71–85.
- Slutsky, E., 'The summation of random causes as the source of cyclic processes', *Econometrica*, **5** (1937), 105–46.
- Stellwagen, E. A. and Goodrich, R. L., *Forecast Pro-Batch Edition*, New York: Business Forecast Systems, Inc., 1991.
- Texter, P. A. and Ord, J. K., 'Forecasting using automatic identification procedures: A comparative analysis', *International Journal of Forecasting*, **5** (1989), 209–15.
- Watson, M. W., Pastuszek, L. M., and Cody, E., 'Forecasting commercial electricity sales', *Journal of Forecasting*, **6** (1987), 117–36.
- Wold, H., *A Study in the Analysis of Stationary Time Series*, Stockholm: Almqvist & Wiksell, 1938.
- Yule, G. U., 'Why do we sometimes get nonsense-correlations between time series? A study in sampling and the nature of time series', *Journal of Royal Statistical Society*, **89** (1926), 1–64.

Authors' address:

Spyros Makridakis and **Michèle Hibon**, INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France.

Copyright of Journal of Forecasting is the property of John Wiley & Sons, Inc. / Business. The copyright in an individual article may be maintained by the author in certain cases. Content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.