December 1997

Can the Use of a DSS Improve Decision Making?

Michael Lawrence
University of New South

Marcus O'Connor
University of New South

Can the Use of a DSS Improve Decision Making?

Michael Lawrence & Marcus O'Connor
School of Information Systems
University of New South Wales, Sydney 2052
Phone: (02)9385.4413; Fax: (02)9662.4061
Email: M.Lawrence@unsw.edu.au

Executive Summary
With a computer on virtually every manager's desk it must seem that increasingly important business decisions will be made using sophisticated DSS's to provide high quality advice to the decision maker. With an ever increasing supply of data available, it is never doubted by management scientists that mathematical models can provide this quality of advice. And it is certainly never doubted that a decision maker provided with good advice will be able to take advantage of it. This paper examines these two issues - can a DSS provide useful advice and when provided with good advice is the decision maker able to take advantage of it and add value to it? A number of studies are drawn on to reflect on these questions including three studies recently carried out by the authors. Most of these studies concern sales forecasting: a task requiring the analysis of (structured) time series information and its integration with contextual and domain information (unstructured). A field study of sales forecasting practice in 13 manufacturing organisations who, like the majority of organisations, develop their forecasts judgementally, reveals that DSS advice could have been useful in a majority of these companies. On the other hand a number of studies are discussed which indicate that managers make poor use of good advice provided to them. This may arise from a difficulty in integration of the advice with the subjective opinion. To investigate this a study was carried out in which the user was given help to integrate the DSS advice with subjectively formed opinions. However this did not improve the use of the information. Thus although we concluded ex post that DSS advice could have been useful in a majority of the companies surveyed, it is not certain that if a DSS had been implemented, the sales forecasts would have been any better.

1. Introduction
Two assumptions underlie the Decision Support Systems (DSS) field:

1. In decision making where there exists both structured and unstructured data, a DSS has the potential to provide sound advice.

2. An executive provided with sound advice from a DSS will:
   (a) generally make a better decision than he could without the DSS and
   (b) generally modify the DSS advice to improve the decision.

That is, 2(a) says that the DSS provides value to the decision maker and 2(b) says that the decision maker provides value over and above that of the DSS.

Even though these assumptions are basic to the DSS field, they do not appear to have been examined. Perhaps they appear to be so obvious that they do not do need to be stated or investigated. This paper addresses these two assumptions reformulated as Research Questions 1 and 2. As we are concerned with both the value and quality of DSS advice we need to limit our attention to those DSS which produce advice. In the classification hierarchy of Alter (Alter 1980), only his highest level 'Suggestion Systems" produce advice. (In his lowest level systems, which he calls a 'File Drawer System" the DSS is little more than a source of raw data.) This paper discusses evidence from both field surveys and from laboratory experiments and shows that while assumption 1 appears true, assumption 2(a) appears to be weakly supported while there is little evidence supporting 2(b). The implications of these findings are discussed.

2. Background

Alter has proposed a hierarchy of DSS which has as the lowest level "File Drawer Systems" and as the highest level "Suggestion Systems" (Alter 1980). File Drawer Systems, as their name suggests, are essentially data access and retrieval systems. On the other hand Suggestion Systems perform calculations which generate advice in the form of a suggested decision for the decision maker. Typically, the decision maker has extra information available to him which is not in the DSS - the unstructured information which by definition cannot be included in the model which incorporates only the structured information. On this basis he is likely to modify the DSS advice to improve the decision.

There is little information on the extent of practical use of DSS's but what exists suggests that most of the DSS in use are of the File Drawer variety or at the lower end of the Alter hierarchy. There is evidence for the use of systems at the higher end (e.g. Smith et al 1992, Stone, 1992) but the extent of use appears to be much more limited than the opportunities available. We outline one such task domain below and use it to investigate the research questions.

One class of potential, but largely unrealised DSS application has been of particular interest to the authors for some years. This is the task of sales forecasting. Forecasting combines both structured and unstructured elements and has been represented as being a prime candidate system for a DSS (van Dissel, Borgman and Beulens, 1990; Lim and O'Connor, 1996). The structured information consists of the time series data and perhaps correlated data such as macro economic time series or promotional activity, although these latter series are most frequently subjectively included into the forecasts as unstructured information. The unstructured information consists of the special knowledge available to the forecaster such as that relating to the product, the market, the competition and likely economic developments. We use the forecasting task as a typical DSS application area in order to examine the two Research Questions above. First we give some background on this task and the extent to which DSS technology is used to support it.

Surveys of forecasting practice in Australia, USA and the UK indicate the continuing high use of judgemental and opinion based methods in preference to quantitative methods (e.g. Dalrymple, 1987; Tarranto, 1989; Sanders and Manrodt, 1994). These studies of business (mostly sales) forecasting practice reveal that only around 10% of the firms surveyed use quantitatively based forecasting techniques and that the number of firms who have tried and subsequently abandoned these techniques is about double the number currently using them. This situation cannot be blamed on lack of supply of forecasting technology - there has been an enthusiastic response by system developers and there is a wide range of inexpensive and user-friendly PC windows and main-frame based DSS software to choose from. These systems benefit from the great volume of forecasting research conducted over the last 25 years. As might be expected, this lack of use has been a cause of concern to researchers and management educators who believe that valuable technology is remaining under-utilised (Makridakis, 1988; Armstrong, 1994).

Improvements in forecasting could have a significant impact on business. Most businesses are involved in regularly developing forecasts and losses through inaccurate forecasts can be high (Makridakis, 1988). For example, organisations involved in warehousing and distribution incur holding and obsolescence costs on their inventory which can be around 25% per annum. In addition lost business from stockouts may be even more damaging. But the preference revealed in these studies for judgemental methods cannot be dismissed as the action of a business community unconcerned by forecast errors. The surveys cited show that business does keep records of forecast accuracy and that they are unlikely to continue to use a method which does not perform (Dalrymple, 1988). Nor can the preference be attributed to lack of knowledge of quantitative forecasting technology. Armstrong (1994) comments that an increasing exposure to forecasting technology has not increased its use.

With this background of little use being made of DSS technology in forecasting it is appropriate to examine, in this task domain, the first research question:

1. In decision making where there exists both structured and unstructured data, does a DSS have the potential to provide sound advice?

We would anticipate that a DSS would be able to provide worthwhile advice to a decision maker by removing inconsistency (Bowman, 1963), by rapidly adjusting to change (O'Connor et al, 1993) and by removing human biasses (Tversky and Kahneman, 1974) as well as by its modelling power to incorporate much more of the available information (Blattberg and Hoch, 1990). On the other hand we do know that human judgement is able to generate forecasts as good as the best available techniques (Lawrence, 1983; Lawrence, Edmundson and O'Connor, 1985; Lawrence and O'Connor, 1992), and that a combination of judgement and model can outperform either (Lawrence, Edmundson and O'Connor, 1986; Edmundson, Lawrence and O'Connor, 1988)

3. Is Forecasting Support Needed?

To answer this question we conducted a field survey of companies developing their sales forecasts judgementally - that is with no DSS or technology supplied advice. The sample group of companies used for the study comprised 13 large Australian national and international manufacturing based organisations selling branded consumer, frequently purchased goods and infrequently purchased durable items. Representatives of sales forecasting management in each company were interviewed in person and the objectives of the study explained. Data security and confidentiality were an important consideration for most of the organisations so the participating companies are identified only by number. Monthly actual sales for a selected range of products (identified by the company as important products for achieving good forecasting accuracy) covering generally a 12 month period were obtained from the participating companies. In total there were around 4,500 actual sales values and 24,000 forecasts in the data base.

3.1 Analysis Methodology

Analysis of the research question requires the calculation of forecasts (such as might comprise the advice of a DSS) and a comparison of their accuracy with that of the company forecasts. If the DSS forecast is equally accurate or more accurate than the company forecast we can say that research question 1 is true - a DSS can provide sound advice. However, if the DSS forecast is less accurate, we can make no definite judgement on the research question. It may have provided sound advice up to the limit of the accuracy of the structured data and saved the decision maker valuable time. But the unstructured data known to the decision maker may have made such a difference that the judgemental forecast exceeds the accuracy of the DSS forecast. On balance though, if the DSS is significantly less accurate than management judgement, we suggest that the DSS would be seen to be of little value.

Two quantitative forecasts were estimated: a naive (the last observation projected forward) and an exponential smoothing forecast. Both these techniques performed well in the Makridakis forecasting competition (Makridakis et al, 1982; Makridakis et al, 1993). The exponential smoothing forecast was optimised on the first half of each data series and then the errors calculated over the full length of the series. Even though this overstates its accuracy, its performance was not overall greatly different from the naive forecast, and in a number of companies, its error was much worse than naive. Because of these difficulties it was decided to only use the naive method for comparison with the company forecast results. However we note that a custom developed DSS carefully tuned for the requirements of each organisation must exceed in accuracy the forecasts used in this comparison.

The forecast accuracy was measured using mean absolute percentage error (MAPE) and symmetric mean absolute percentage error (SMAPE) (Makridakis, 1993), defined as

MAPE is the mean of $\text{APE} = |(\text{actual} - \text{forecast})|/\text{forecast}$

SMAPE is the mean of $\text{SAPE} = 2|(\text{actual} - \text{forecast})|/(\text{actual} + \text{forecast})$

The error for each time period is averaged, for each company, over the products to give an average accuracy for a company. These accuracies are reported in the results. Their statistical significance is calculated using paired t-tests where the pairing is over the same product in the same period.

3.2 Results

Table 1 presents the results for the comparison of the one month ahead forecast accuracies. This table contains the MAPE and SMAPE scores for the naive and the company judgemental forecasts. For the MAPE and the SMAPE measures, a lower value indicates a more accurate forecast. The SMAPE measure avoids the asymmetry and extreme value problems of the MAPE measure which are particularly evident when the actual drops suddenly to near zero, driving the MAPE to a very large number. Despite the differences in the measures, the two sets of results are in agreement as to whether the company or naive forecast is more accurate. In addition, the table indicates for which companies the naive forecast accuracy (using the SMAPE measure in a paired t-test) is statistically significantly different from the company forecast accuracy (at the 0.001 level). The table shows that for 4 out of the 13 companies, the one month ahead judgemental company forecasts were significantly more accurate than the naive. However for 2 companies, the naive forecasts were significantly more accurate, and for 7 of the comparisons, the difference in accuracy was insignificant. Thus, for only about a third of the forecast comparisons was the management judgemental forecast significantly more accurate. It would appear, on balance, that the research question is supported - DSS forecasts appear to be able to provide generally good advice. Had seasonal adjustment of the naive forecast been performed, it would most likely have demonstrated much improved accuracy over the unseasonally adjusted naive forecast. We now examine another aspect of the quality of the company forecasts - whether they are unbiassed and efficient. If they do not possess these desirable qualities then it is likely that providing an anchor point for the forecasters in the form of DSS advice may be able to improve the accuracy of the final forecast.

		MAPE		SMAPE			
Co. No	cases	Company	Naive	Company	Naive	T-test	Best
1	104	22	18	17	18	n.s.	
2	169	1087	33	71	30	**	naive
3	153	31	46	19	31	**	company
4	239	104	172	57	77	**	company
5	294	23	24	20	23	n.s.	
6	376	71	45	36	36	n.s.	
7	284	43	51	51	32	**	naive
8	252	42	58	31	48	**	company
9	95	21	35	18	30	**	company
10	755	76	50	31	31	n.s.	
11	1552	114	84	37	37	n.s.	
12	149	27	23	23	23	n.s.	
13	165	28	21	23	20	n.s.	

** significant $p<0.001$

Table 1 One period ahead forecast error - comparison of company and naive forecasts

3.4 Use of Past Information

A number of studies have analysed analysts' forecasts of quarterly and annual company earnings and found evidence that the forecasts did not utilise available (past) information efficiently (Ali, Klein and Rosenfield, 1992; Mendenhall, 1991; Abarbanell and Bernard, 1991). Specifically they found evidence of bias and serial correlation in forecast errors. The question arises as to whether inefficiency and bias characterises monthly company sales forecasts, and whether they might be a cause of the poor accuracy of the final company forecasts.

To investigate the issue of the efficiency and bias of the company forecasts, (stepwise) regression analysis was used on the following equation:

$$\text{error}_t = a.\ \text{forecast}_t + b.\ \text{Error}_{t-1}$$

$$(\text{where, } \text{error}_t = \text{actual}_t - \text{forecast}_t)$$

If the forecasts are efficient, we expect that there will be no serial correlation in the errors from period to period - i.e. that b will be zero. In other words, people learn to accommodate the lessons from the past error in setting their new forecast. If there is no bias in the forecasts, we expect that a will be zero - i.e. there is no relationship between the error and the forecast. If there is a consistent pattern of under-forecasting (over-forecasting) a should be positive (negative). Table 2 presents the results of the regression analysis, together with (from Table 1) comparisons between company and naive forecast accuracy.

COMPANY	a	b	Best
1			
2	-0.38**	0.55**	naive
3	-0.096**	-	company
4	-0.197**	-	company
5	-0.137**	0.20**	
6			
7	-	0.679**	naive
8			company
9	0.064*	-	company
10	-0.124**	0.166**	
11	-0.065**	0.480**	
12	0.031*	0.47**	
13	-0.156**	-0.407**	

** significant $p<0.001$ * significant $p<0.01$

Table 2 Regression coefficients for bias and inefficiency in one period ahead forecasts. ($error_t = a\ fcst_t + b\ error_{t-1}$)

Table 2 indicates that for 9 out of the 13 companies a was significant, indicating that there was bias in the forecasts. For 7 of the companies there was an inefficiency in the forecasts (with b significantly different from zero). Interestingly, the 4 companies which were found to have better accuracy than the naive (from Table 1) all have efficient forecasts ($b = 0$). Companies 1 and 6 also had a and $b = 0$, although the forecasts for these companies were not better than the naive. All seven company forecasts where $b \neq 0$ (i.e., inefficient forecasts) were either worse than the naive, or equal to the naive. As mentioned, most company forecasts contained bias ($a \neq 0$). Even in the case of the 4 companies which were more accurate than the naive, there was a significant bias for 3 of them.

These results suggest that, if the bias and inefficiency is recognised and the company forecasts accordingly adjusted (through advice from a DSS), it could lead to greater accuracy. To test this proposition, a revised forecast was computed to remove the bias and inefficiency. The model was developed on all data available for the companies where there was a significant regression fit as detailed in Table 2. These results indicate that in 6 out of the 7 companies where the company forecasts were not more accurate than the naive and there was a significant regression, the procedure produced significantly better forecasts than the original company forecasts. Thus, when bias and inefficiency is removed, the revised forecasts were more accurate. Thus we conclude that there is convincing evidence that Research Question 1 is answered in the affirmative.

We now examine the second research question examining whether executives are able to benefit from advice given by a DSS and to add value to it.

4. Does a DSS actually improve decisions?

The previous section has highlighted the value of decision support in the determination of forecasts. A conclusion that can be drawn is that the product forecasting teams in industry have considerable difficulty in integrating the contextual information (the unstructured information) with time series information (the structured information). Analysis of the discourse at the forecasting meetings revealed that the majority of the information under consideration was of a non time series nature. That is, information about the advertising programs coming up, labour market problems and supply issues tended to dominate the deliberations about the forecast for the next period. Nevertheless, despite the plethora of such rich and relevant information available to them, the participants don't seem to be able to adequately utilize it profitably. Systematic bias and an inability to learn from past errors seem to mitigate against any comparative advantage in the information set available to people over computer based models. There seems to be a clear need for decision support to improve decision quality. Most decisions are still in the province of human judgement - a DSS is designed to merely supplement judgement, not to replace it. However, even if such decision support is provided, the question remains whether people are able to use the DSS to improve the decision task. This is the essence of Research Question 2.

To empirically examine this question, laboratory based experiment have been conducted to see whether people are able to discern the value or contribution of the DSS in the forecasting task. A priori, the value of the decision support relates to the quality of the information provided to it. Past studies (Arkes et al, 1986; Ashton, 1990; Dawes et al, 1989; Peterson & Pitz, 1986) have often shown that people ignore high quality information provided to them. For example, Arkes et al (1986) found a consistent tendency for people to ignore highly accurate model information and that this pervasive tendency was not easily overcome unless people were explicitly told that 'people who try to do better than [the model]...actually do a lot worse. So just follow the model' (p97). Ashton (1990) speculated that such reluctance to rely on decision aids may be due to an over-estimation of one's own ability and the phenomenon of general over-confidence.

Lim and O'Connor (1995) have investigated the ability of people to recognise the value of a DSS producing quality forecasts in a situation where a monetary incentive was paid for accuracy. Their experimental design provided the subjects (graduate students in Commerce at the University of New South Wales) with a computer presented graph of a time series and asked for a forecast to be prepared. Another treatment group was provided with the graph together with Deseasonalised Exponential Smoothing (DSE) forecast of the time series and asked to produce a forecast. This group was told the DSE forecast was found through exhaustive comparison of forecasting methods to be one of the most accurate methods. A third group was given an artificially good forecast (QF)[1] and told it was prepared by a new and very superior forecasting method, and that they would be unlikely to be able to do better. Each experimental group produced forecasts over 30 successive trials with feedback given after each trial. The accuracies of these forecasts were calculated using the mean absolute percentage error (MAPE) measure. The results for this experiment are given below in Table 3 along with the accuracies of the forecasts given by both the DSE and the QF models.

Forecast Type	MAPE
Initial forecast	17.7
Revised forecast using DSE	16.2
Revised forecast using QF	8.9
Accuracy of DSE	12.7
Accuracy of QF	5.8

Table 3 Forecast Accuracy

Table 3 shows that the DSS advice improved the forecast accuracy from the unaided mean of 17.7% to 16.2% (for those receiving the DSE advice) and to 8.9% for those receiving the QF advice. However the improvements were still markedly below the mean accuracy for the forecasts provided to the subjects as advice. Thus despite a monetary incentive given for good performance and feedback provided on forecast accuracy, the subjects were unable to take full opportunity of the advice provided. This suggests that Research Question 2(a) is supported while 2(b) is not supported.

It is not enough, however, to provide people with highly reliable information. They need some mechanism to be able to use it in their decision making process. The previous section, as mentioned, has shown that people have problems in integrating the information in a profitable way. When faced with (say) two pieces of information, there are two tasks that need to be done to utilize them. First, one needs to specify the relative weights that should be placed on each piece of information. For example, one may decide to rely more heavily on one piece of information rather than the other. This is termed the cue weighting task. Although evidence is mixed as to the ability of

[1] This forecast was produced by averaging the DSE forecast with the actual value.

second task is to actually perform the combination or integration based on the weights determined in the first task. Numerous studies have shown that the task of combination is one in which people do not excel (for a review see Kleinmuntz, 1990). Thus, on the basis of past evidence, considerable advantages might accrue if some decision support is provided to aid in the combination or integration process. This study investigates the proposition that a DSS supporting the combination process is beneficial to the decision making process.

4.1 Study Design

As mentioned, there were two research questions at issue in the design of the laboratory experiment. First, were people able to perceive differences in the reliability or accuracy of the information provided to them? Second, does a DSS that aids in the combination process actually lead to improved performance?

To this end, a forecasting support system (a DSS for forecasting) was designed to aid in the production of final forecasts from time series data. Subjects were provided with a display of time series data on which they were required to produce an initial forecast - i.e. one that was based only on the past time series data. They were then provided with a computer produced statistical forecast of one of three levels of reliability - low, medium and high reliability. The three levels of reliability were constructed such that, at the low level, judgemental extrapolation was likely to be more accurate: and at the high reliability level it was most unlikely that people could have made a forecast that was as accurate.

A DSS was provided to half the subjects. The DSS provided people with the opportunity to specify weights they wished to attach to their initial forecast and the statistical forecast provided to them. But they were not required to make any calculation of a final forecast - they only had to specify the weights they wished to assign. Thus, the DSS was designed to test the proposition that people are relatively good at specifying their weightings for the information that is to be combined, but are bad at the actual task of integration. It was expected that people would, over time, come to assign a weight to the provided statistical forecast that was commensurate with its reliability. In order to see the effects of the DSS, half the subjects were provided with a DSS to aid in the process of information integration and half were required to do the task on their own. There were 30 repeats of the task, to examine any learning effect over time. Thus, the experiment was a 3 (information reliability) x 2 (DSS) design that was examined over the course of the 30 trials. There were 48 subjects in the experiment, all of who were mature post-graduate students in the Faculty of Commerce and Economics at the University of New South Wales.

The task procedure at each trial was for the subject to produce an initial judgemental forecast based solely on the past time series information (the initial forecast). The statistical forecast was then revealed to them. They were then required to produce a new forecast based on both their initial forecast and the statistical forecast. The contribution of the DSS to the task was evaluated by the improvement in accuracy (measured in APE) of the final forecast over their initial judgement (viz. $IMP = APE_{initial} - APE_{revised}$). Thus, a positive IMP indicates that the error for the revised forecast was lower than for the initial judgement. Table 4 provides the mean IMP across decision support and reliability conditions.

	INFORMATION RELIABILITY		
	LOW	MEDIUM	HIGH
NO DSS	-0.57	5.24	12.78
DSS	0.21	6.52	8.62

Table 4 - IMP means across decision support and reliability conditions.

Overall, there was no difference in IMP between the DSS and no-DSS conditions. However, where a highly reliable statistical forecast was provided, there was a significant difference in IMP between the decision support groups ($t(830)=3.21, p<.001$). However, as Table 4 reveals, the no-DSS group was more accurate than the DSS group. So, for this reliability condition, the DSS actually made the final forecasts worse, compared to the no-DSS condition.

To further investigate the weights people placed on the two pieces of information, a regression analysis was performed of the statistical forecast on the final forecast. This provided information on the weights people actually placed on the statistical forecast (and by definition the initial judgemental forecast). Table 5 contains the results of the regression analysis of the weights people placed on the (provided) statistical forecast for the three reliability conditions across the DSS and no-DSS groups, together with the optimal model.

	INFORMATION RELIABILITY		
	LOW	MEDIUM	HIGH
NO-DSS	0.103	0.545	0.831
DSS	0.279	0.472	0.722
OPTIMAL	0.095	0.835	1.069

Table 5 Regression weights of the statistical forecasts.

In all cases, the weights specified by the subjects in the DSS condition were further from the optimal than those used by the non-DSS subjects. These results clearly confirm those of Table 4, that the DSS which took away the necessity for tedious calculations, was detrimental to forecast accuracy. These results question whether the problem with information integration lies in the combination component. They suggest that there may also be a problem with the specification of the weights themselves.

Over the course of the 30 trials, people reduced their persistent faith in their own initial forecast. Table 6 presents the regression weights for the three reliability conditions over three time blocks (periods 1-10, 11-20, and 21-30).

	TIME BLOCKS (periods)			
	1-10	11-20	21-30	Average
HIGH	0.676	0.821	0.852	0.791
MEDIUM	0.524	0.529	0.464	0.498
LOW	0.222	0.205	0.117	0.168

Table 6 Regression weights for the statistical forecast over three time blocks.

Clearly, with the high reliability forecasts, people gradually increased their reliance on the provided data. For the low reliability condition, there was some late reduction in the statistical forecast. The average statistics show that, in general, people were able to differentiate and utilize the different reliability of the information provided to them.

In conclusion, we speculated that people have two problems to deal with in the task of integrating information for decision making. The first was the task of deciding which information to weight more than others (the perception problem) and the second was the task of actually combining them after the weights had been decided in the first task (the combination problem). Past studies have suggested that people have considerable difficulty with the latter. So a DSS to support it should improve judgement. This laboratory study has shown that the DSS actually made the forecasts worse. Thus, there are problems in both of the above phases/tasks - the perception and the combination tasks. Whilst this study does not suggest that people do not have a problem with the combination task, it does imply that the perception of the importance of information cues needs to be supported in any DSS, not just a simple aid to help combinations. Support of this type, however, is a much more complex task that presents a considerable challenge to the designer of DSS.

5. Discussion and Conclusions

Through the foregoing reviews of past and current research, we have demonstrated that while, ex post, we can demonstrate that good DSS advice is needed and generally available to improve decision making, it is by no means certain that it will be used to good advantage. From the evidence of laboratory trials, there appears to be a considerable loss of value in the step of the decision maker taking the DSS advice and converting it into a decision. Far from adding value to the decision, the decision maker appears to reduce value. Thus, it is not sufficient for a DSS to just provide good advice, even when feedback is provided to show how good that advice is. There is another step needed to guide the decision maker in the use of that advice. However, deciding how to do that is by no means simple. One such means was discussed which decomposed the decision maker's task into two stages - one stage of weighting the information cues and one stage of combining them.

However a system designed to help the decision maker with these two stages was not successful, indicating probably that the decomposition confused rather than helped.

References

Abarbanell, J.S., and Bernard, V.L., (1992), 'Tests of analysts' overreaction/underreaction to earnings information as an explanation for anomalous stock price behaviour", *Journal of Finance*, 47, 1181-1207.

Ali, A., Klein, A. and Rosenfield, J., (1992), 'Analysts use of information about permanent and transitory components in forecasting annual EPS", *The Accounting Review*, 67, 183-198.

Alter, S.L., *Decision Support Systems: Current practice and continuing challenges*. Addison Wesley, Reading, Mass (1980).

Arkes, H. R., Dawes, R. M. and Christensen, C., (1986), 'Factors influencing the use of a decision rule in a probabilistic task', *Organizational Behaviour and Human Decision Processes*, 37,93-110.

Armstrong, J.S., Review of Sanders and Manrodt, *International Journal of Forecasting*, 10 (1994) 471-472.

Ashton, R.H. 'Pressure and performance in accounting decision settings: paradoxical effects of incentives, feedback and justification", *Journal of Accounting Research* 28 (1990), 148-180.

Blattberg, R. C. and Hoch, S. J., (1990), 'Database models and managerial intuition - 50% model + 50% manager', *Management Science*, 36, 887-899.

Bowman, E.H. "Consistency and Optimality in Managerial Decision Making", *Management Science*, 9 (1963) 310-321.

Dalrymple, D.J. (1987), Sales forecasting practices: results of a United States survey. *International Journal of Forecasting*, 3, 379-392.

Dawes, R.M., Faust, D. and Meehl, P.E. "Clinical versus actuarial judgement" *Science*, 243 (1989), 1668-1773.

Edmundson, R., Lawrence, M., & O'Connor, M. (1988), 'The use of non time series information in sales forecasting: a case study", *Journal of Forecasting*, 7, 201-211.

Funaro. J.F., 'An empirical analysis of five descriptive models for cascaded inference", *Organisational Behaviour and Human Decision Processes*, 14 (1975), 186-206.

Kleinmuntz, B., (1990), 'Why We Still Use Our Heads Instead of Formulas: Toward an Integrative Approach", *Psychological Bulletin*. 107, 296-310.

Lawrence, M. (1983), 'An exploration of some practical issues in the use of quantitative forecasting models", *Journal of Forecasting*, 1, 169-179.

Lawrence, M. and O'Connor, M., (1992), "Exploring Judgemental Forecasting", *International Journal of Forecasting*, 8, 15-26.

Lawrence, M., Edmundson, R. and O'Connor, M. "An examination of the accuracy of judgmental extrapolation of time series", *International Journal of Forecasting*, 1 (1985) 14-25.

Lawrence, M., Edmundson, R. and O'Connor, M. "The accuracy of combining judgmental and statistical forecasts", *Management Science*, 32 (1986) 1521-1532.

Lim J. and M. O'Connor, 'Judgemental adjustment of initial forecasts: its effectiveness and biases", *Journal of Behavioural Decision Making*, 8 (1995) 149-168.

Lim J.S., and O'Connor M.J. 'Judgemental Forecasting with interactive forecasting support systems", *Decision Support Systems* 16 (1996) 339-357.

Makridakis, S. "Metaforecasting", *International Journal of Forecasting*, 4 (1988) 467-491.

Makridakis, S. (1993), 'Accuracy measures: theoretical and practical concerns", *International Journal of Forecasting*, 9, 527-529.

Makridakis, S., Anderson, A., Carbone, R., Fildes, R., Hibon, M., Lewandowski, R., Newton, J., Parzen, E and Winkler, R. "The accuracy of extrapolation (time series) methods: results of a forecasting competition", *Journal of Forecasting*, 1 (1982) 111-153.

Makridakis, S., Chatfield, C., Hibon, M., Lawrence, M., Mills, T., Ord, K., and Simmons, L, (1993), "The M2 competition: a real time judgemental based forecasting study", *International Journal of Forecasting*, 9, 5-22.

Mendenhall, R. R. (1991), 'Evidence on the possible underweighting of earnings-related information", *Journal of Accounting Research*, 29, 170-179.

O'Connor, M., Remus, W., & Griggs, K. (1993), 'Judgemental forecasting in times of change", *International Journal of Forecasting*, 9, 163-172.

Peterson, D.K. and Pitz, G.F. 'Effect of input from a mechanical model on clinical judgement" *Journal of Applied Psychology*, 71 1 (1986), 163-167.

Sanders, N. and Manrodt, K., (1994), 'Forecasting practices in US corporations: survey results", *Interfaces*, 24, 92-100.

Smith, B.C., Leimkuhler, J.F., and Darrow, R.M., 'Yield management at American Airlines", *Interfaces* 22, No 1 (1992) 8-31.

Stone, L.D. 'Search for the SS Central America: Mathematical treasure hunting" *Interfaces* 22, No 1 (1992) 32-54.

Tarranto, G.M. (1990), "Sales forecasting practices: results from an Australian survey", University of NSW.

Tversky, A. and Kahneman, D. "Judgement under uncertainty: heuristics and biases", *Science*, 185 (1974) 1124-1131.

van Dissel, H.G., Borgman, H.P. and Beulens, A.J. 'Task allocation between DSS and problem owner: The example of Box and Jenkins time series analysis". *Decision Support Systems* 6, (1990) 339-345.