International Journal of Forecasting 8 (1992) 545–557
North-Holland

Business planning under uncertainty

Will we attain our goal?

Lilian Shiao-Yen Wu and J.R.M. Hosking
IBM Research Division, T.J. Watson Research Center, Yorktown Heights, NY 10598, USA

Jeanne M. Doll
IBM US Marketing and Services, 1133 Westchester Avenue, White Plains, NY 10604, USA

Abstract: In business, planning often starts with an annual financial target which is then apportioned into a set of monthly targets, called the track. As the year progresses planners must assess, by comparing the actual measurement with the track, whether the measurement is on track to reach the target or whether action needs to be taken in order to reach the target. Wu (*The Statistician*, 1988, 37, 141–152) has briefly described three planning charts, called WINEGLASS, SHIPWRECK and OUTLOOK, that enable this assessment to be made objectively. Here we give a detailed description of Wu's approach. We give a complete description of the charts and of the calculations needed to construct them, and we present examples of these calculations. We also describe a way of setting the current year's monthly targets on the basis of previous years' data. The calculations are based on a statistical model, a modification of Wu's Track Uncertainty Model, of the variability of the actual measurement about the track.

Keywords: Planning, Statistical modeling, Tracking, Wineglass chart.

1. Introduction

In business, planning often starts with financial targets from which goals for financial controls and sales of individual products are generated. There are two focuses in this planning process. One is producing the annual goals. The other is assessing, as the year progresses, whether the goals set are being attained or whether action needs to be taken to achieve the goals. To make these assessments, the annual target is further apportioned into monthly targets. The monthly targets are called the track. Often the track is constructed by a planner and is based on the planner's judgement of such matters as manufacturing capacity, product announcements, price changes and seasonal variations in demand. When the track is based on a planner's judgement, it will be called a *planner track*. Sometimes it is desired to make assessments based on history alone, excluding the planner's subjective judgement. When a track is based on history alone, it will be called a *historical track*.

This business planning framework is somewhat different from a typical forecasting problem. The distinction between planning (setting targets and monitoring whether they can be attained) and forecasting (extrapolation of a time series) is rarely made in the forecasting literature, but has recently been described by Bretschneider (1991).

Wu (1988) described how to model and quan-

Correspondence to: L.S.-Y. Wu, IBM Research Division, T.J. Watson Research Center, Yorktown Heights, NY 10598, USA

tify the uncertainties in the planning process based on a planner track, and mentioned three planning charts, called WINEGLASS, SHIPWRECK and OUTLOOK, that can be used to assess the attainability of goals and to make annual business risk outlooks. In this paper we describe how the three planning charts are constructed, and how the analyses that they involve are interrelated. We also show how to quantify the uncertainty and construct the three planning charts when using a historical track rather than a planner track.

For clarity of explanation we will for the most part suppose that the measurement being assessed is the sales of a product. For sales, the objective is to attain a high sales total at the end of the year, so we call sales a 'high objective' or 'the larger the better' measurement. However, our approach applies equally well to measurements, such as expense, the nature of which is 'low objective' or 'the smaller the better'.

2. The planning framework

The planning framework is as follows. For the current year there is an annual target $\boldsymbol{T}$. At month I in the current year, monthly sales to date are $Y_1, \ldots, Y_I$. Given these data, we seek answers to the following questions.

(1) Are sales on track to make the annual target?

(2) Given an optimistic view of the future, will we make up the deficit (or, in the case of expense, will we eliminate the excess) by the end of the year?

(3) What are reasonable outlooks, high, medium and low, for the year's total sales?

The WINEGLASS, SHIPWRECK and OUTLOOK charts provide answers to these questions. Two components are required to construct the charts: the *track*, a set of monthly targets $T_1, \ldots, T_{12}$, that sum to the annual target; and an estimate of the accuracy of the track.

Two kinds of track can be used. A *planner track* is the planner's judgement of how the target can be reached. A *historical track* is based solely on history and excludes the planner's judgement: its construction is described in Subsection 4.2. The accuracy of the track is estimated from past data. The estimate is based on how close the seasonality in the track was to the actual seasonality in the past. By 'seasonality' we mean the way that the annual total is divided among the months (e.g. 5% in January, 12% in February, etc.).

3. A statistical model

Our mathematical framework is as follows. Let Y_i denote sales in month i of the current year; let T_i be the track for month i and $\boldsymbol{T}$ the annual target. We use the model

$$Y_i = gT_i + \tau_i \,, \tag{1}$$

where the τ_i are independent random variables with mean 0 and variance $\omega^2 gT_i g\boldsymbol{T}$. The amount by which g differs from 1 represents the error in setting the annual target: if $g = 1$, then the mean value of the annual sales total $Y_1 + \cdots + Y_{12}$ is exactly $\boldsymbol{T}$. The τ_i represent the error in apportioning the track.

The structure of the variance term $\omega^2 gT_i g\boldsymbol{T}$ arises from a mixture of theory and practice. For low-volume products, simple models of customer behaviour imply that month-to-month variation in sales can be modeled by a Poisson distribution [Wu (1988, Section 4.2)]. For high-volume products a compound Poisson model is appropriate [Wu (1988, Section 4.3)]. The gT_i term in the variance makes the variance of Y_i proportional to the mean of Y_i, in accordance with these Poisson models. That the variance is proportional to T_i also means that the year-to-date, actual-to-track ratio $(Y_1 + \cdots + Y_I)/(T_1 + \cdots + T_I)$ is the minimum-variance linear unbiased estimator of g based on sales data available at month I. Because this ratio is a natural measure for comparing year-to-date sales with the track, and is frequently used as such by planners, it is logically consistent for it also to be the statistically optimal estimator of g. The presence of the $g\boldsymbol{T}$ term in the variance means that the remaining constant term ω^2, the *WINEGLASS uncertainty*, is dimensionless and is unaffected by changing the scale of the observations. The precise form of the $g\boldsymbol{T}$ term – for example, that it should be $g\boldsymbol{T}$ rather than just $\boldsymbol{T}$ – is chosen to ensure that outlooks, described in subsection

5.3 below, remain the same if the annual target is changed.

Because the variance of τ_i is positive only when $T_i > 0$, our procedures are appropriate only for strictly positive data such as sales and expense, and not for measurements such as earnings (sales minus expense) when these could plausibly take negative values. In practice our procedures are often not affected by the presence of occasional negative numbers in the data, but the circumstances in which this is true depend on the details of how individual numbers affect the computations involved in the estimation of g and ω^2.

The assumption of independence in the τ_i implies that any apparent serial correlation in the observed Y_i is adequately explained by the month-to-month variation of the track. This assumption is plausible if the track has been competently set; alternatively it is possible to modify the model so as to include serial correlation among the τ_i. However, the nature and magnitude of serial dependence are difficult to estimate accurately from the small data sets with which we typically deal. When assessing sales performance within a year we are therefore content to ignore the possibility of serial dependence. More complicated time-series models can, however, be appropriate for setting a planner track at the start of a year [Wu, Ravishanker and Hosking (1991)].

Some calculations concerned with WINEGLASS, SHIPWRECK and OUTLOOK charts require the specification of the distribution of the τ_i. We generally assume this distribution to be Normal. This assumption is made mainly for convenience, but seems to be plausible for typical sales data.

4. Tracks and their accuracy

4.1. Planner tracks: Estimating their accuracy

The accuracy of the track is related to the ω^2 parameter of model (1). To get a rough interpretation of ω, note that the actual-to-track ratio for month i, Y_i/T_i, has mean g and variance $\omega^2 g^2 \boldsymbol{T}/T_i$. Usually g is close to 1 and for a typical month $\boldsymbol{T}/T_i \approx 12$. Therefore $\omega\sqrt{12}$ may roughly be though of as the standard deviation of a typical month's actual-to-track ratio.

To estimate ω^2 we make use of data from previous years. For example, if there are two years of history, then the accuracy of the track is computed from a weighted average of the errors in using the planner seasonalities to forecast the actual seasonalities for the past two years.

In the general case, suppose that H years of history are to be used and index the years in reverse order as J (the current year), $J-1, \ldots, J-H$. For each year j there is an annual target $\boldsymbol{T}_{/j}$, a track $T_{1/j}, \ldots, T_{12/j}$, and actual monthly sales $Y_{1/j}, \ldots, Y_{12/j}$. Model (1) is extended to cover different years by using the subscript i/j to indicate month i of year j. We assume that

$$Y_{i/j} = g_{/j} T_{i/j} + \tau_{i/j}, \tag{2}$$

where $\tau_{i/j}$ has mean 0 and variance $\omega^2 g_{/j} T_{i/j} g_{/j} \boldsymbol{T}_{/j}$. The g parameter is permitted to vary from year to year and is therefore written $g_{/j}$. The ω^2 parameter is assumed to remain constant from year to year and can be estimated separately from each year's data. From the data for year j the best linear unbiased estimator of $g_{/j}$ is

$$\tilde{g}_j = \frac{Y_{1/j} + \cdots + Y_{12/j}}{T_{1/j} + \cdots + T_{12/j}}. \tag{3}$$

Because the variance of $\tau_{i/j}$ is proportional to $T_{i/j}$, it is natural to base an estimator of ω^2 on the sum of squares of weighted residuals $\Sigma\, \hat{\tau}_{i/j}^2 / T_{i/j}$, where $\hat{\tau}_{i/j} = Y_{i/j} - \tilde{g}_j T_{i/j}$. It is straightforward to show that

$$E\left(\sum_{i=1}^{12} \hat{\tau}_{i/j}^2 / T_{i/j}\right) = 11 \omega^2 g_{/j}^2 \boldsymbol{T}_{/j}. \tag{4}$$

Therefore a reasonable estimator of ω^2 using year j's data is

$$\tilde{\omega}_j^2 = \frac{1}{11} \sum_{i=1}^{12} \frac{(Y_{i/j} - \tilde{g}_j T_{i/j})^2}{\tilde{g}_j^2 T_{i/j} \boldsymbol{T}_{/j}}. \tag{5}$$

This estimator is not unbiased, but no simple unbiased estimator exists. Note from (3) and (5) that the estimators $\tilde{\omega}_j^2$ are dimensionless: they measure the difference between the seasonality in the track and the seasonality in the actual data, not the error in setting the annual target.

An overall estimator of ω^2 is obtained by forming an average of the estimators $\tilde{\omega}_j^2$, $j = J-H, \ldots, J-1$. One reasonable choice is the arithmetic mean, an unweighted average. However, we usually prefer to use the weighted average

$$\tilde{\omega}^2 = \sum_{j=1}^{H} w_j^{(H)} \tilde{\omega}_{J-j}^2 , \tag{6}$$

the weights $w_j^{(H)}$ being defined in eqn. (13) below: the sum in eqn. (6) then assigns relatively more weight to more recent years' data, thereby providing some protection against evolutionary changes of ω^2 over time.

The number of years of planner tracks to be used may require careful consideration. Our recommendation is to use only those years for which the accuracy of the planner track is likely to be a reasonable estimate of the current year's planner track accuracy. For example, suppose that two years ago the distribution of a product was changed from selling directly to customers to selling through wholesalers. Since the seasonality in how customers buy and how wholesalers stock inventory may differ a great deal, the accuracy of a planner track may also change when the method of distribution is changed. It would therefore be sensible in this case to use only two years of past data to estimate ω.

4.2. *Historical tracks, and estimating their accuracy*

There may not always be a planner track available, and even when there is it may be preferred to base the track on the pattern of seasonality in past years' data rather than on a planner's subjective judgement. The 'historical track' described in this subsection is based on past years' data and can be constructed at the start of the year, before any data have been obtained for the current year. Suppose as before that H years of history $J-H, \ldots, J-1$, are to be used. The seasonality of the data in year j is obtained by rescaling the data so that the monthly values sum to 1:

$$X_{i/j} = \frac{Y_{i/j}}{Y_{1/j} + \cdots + Y_{12/j}} . \tag{7}$$

Then given an annual target $\boldsymbol{T}$, the historical track for the current year J is defined to be the annual target multiplied by a weighted average of the seasonalities in the H available years of history:

$$T_i = \boldsymbol{T} \sum_{j=1}^{H} w_j^{(H)} X_{i/J-j} . \tag{8}$$

To ensure that $T_1 + \cdots + T_{12} = \boldsymbol{T}$, the weights must sum to 1:

$$\sum_{j=1}^{H} w_j^{(H)} = 1 . \tag{9}$$

The weighting scheme is based on exponential smoothing. Given a quantity X_j observed in successive years, exponential smoothing predicts X_J given all previous values, $X_{J-1}, X_{J-2}, X_{J-3}, \ldots$, by

$$\begin{aligned} \hat{X}_J &= \theta\{X_{J-1} + (1-\theta)X_{J-2} + (1-\theta)^2 X_{J-3} + \cdots\} \\ &= \sum_{j=1}^{\infty} \theta(1-\theta)^{j-1} X_{J-j} . \end{aligned} \tag{10}$$

Here θ is a parameter taking a value in the range $0 < \theta \leq 1$. In practice it is likely that only a short stretch of past data is available, so the infinite sum in (10) must be modified. We use an approach similar to the 'backcasting' method mentioned by Abraham and Ledolter (1983, p. 88). The predictor (10) is optimal, in the sense of having a minimum mean-square error among all linear predictors, when the X_j are generated by the integrated first-order moving-average process.

$$X_j - X_{j-1} = a_j - (1-\theta)a_{j-1} , \tag{11}$$

in which the sequence $\{a_j\}$ consists of uncorrelated identically distributed random variables with finite variance. Given a finite stretch of the past, $X_{J-H}, \ldots, X_{J-1}$, the optimal linear predictor of X_J is the linear combination

$$\hat{X}_J^{(H)} = \sum_{j=1}^{H} w_j^{(H)} X_{J-j} , \tag{12}$$

in which the weights $w_j^{(H)}$ are chosen to minimize $E(X_J - \hat{X}_J^{(H)})^2$. For model (11) it can be shown, using the methods of Brockwell and Davis (1987, Section 9.5), that the weights are given by

$$w_j^{(H)} = \frac{\theta\{(1-\theta)^{j-1} + (1-\theta)^{2H-j}\}}{1-(1-\theta)^{2H}}, \quad j = 1, \ldots, H. \tag{13}$$

The $w_j^{(H)}$ are therefore a natural and logically consistent set of weights to use for exponential smoothing of a finite stretch of history, and so we use them in eqn. (8) to construct the historical track for the current year. For $0 < \theta < 1$, the weights $w_j^{(H)}$ are all positive and decrease as j increases: thus relatively more weight is given to more recent years' data. As $\theta \to 0$, $w_j^{(H)} \to 1/H$ for each j, assigning equal weight to each year's data; as $\theta \to 1$, $w_1^{(H)} \to 1$ and $w_j^{(H)} \to 0$ for $j \neq 1$, assigning all of the weight to the most recent year.

Often the most natural way to describe a weighting scheme is by the weight applied to the most recent year. Equation (13) can be parameterized by $w_1^{(H)}$, which takes values between $1/H$ and 1 and from which θ and thence the other $w_j^{(H)}$ can be determined.

To estimate ω^2 for a historical track we use the same approach as with planner tracks, based on eqns. (3)–(6). This requires the comparison of actual data with historical tracks for past years. Historical tracks can be computed for every year except the first, so the number of individual years' estimates of ω^2 which contribute to the overall estimate is one fewer than the number of years of history. For example, if there are 3 years of history, then the accuracy of the historical track is computed from a weighted average of: (i) the error in using the seasonality of year 1 to forecast year 2; and (ii) the error in using a weighted average of the seasonalities of years 1 and 2 to forecast year 3.

In general, when H years of history are used, historical tracks can be computed for years $J-1$, $J-2, \ldots, J-H+1$. The historical track for year $J-j$ is based on data for the $H-j$ years $J-k$, $k = j+1, \ldots, H$, and, by the same argument as that leading to eqn. (8), is given by

$$T_{i/J-j} = \boldsymbol{T}_{/J-j} \sum_{k=j+1}^{H} w_{k-j}^{(H-j)} X_{i/J-k}; \tag{14}$$

again the weights $w_j^{(\cdot)}$ are as defined in eqn. (13). Estimators of ω^2 are obtained, via eqns. (3) and (5), for years $J-H+1, \ldots, J-1$. Corresponding to eqn. (6), an overall estimator of ω^2 is

$$\tilde{\omega}^2 = \sum_{j=1}^{H-1} w_j^{(H-1)} \tilde{\omega}_{J-j}^2. \tag{15}$$

Equation (14) involves the annual targets for past years, $\boldsymbol{T}_{/J-j}$, but it is not important that these be known. Equations (3) and (5) show that the magnitude of the target for year j does not affect $\tilde{\omega}_j^2$ and therefore does not affect $\tilde{\omega}^2$. This should not be surprising since, as noted earlier, ω^2 is dimensionless and measures the accuracy of the seasonality in the track and not the accuracy of the annual target.

As with planner tracks, the number of years of history that should be used when constructing a historical track requires careful consideration. Our recommendation is to use only those years for which the seasonality is likely to be similar to that of the current year. If the pattern of seasonality in the data has changed markedly in the past, then data before the change should be discarded.

Some data sets contain only weak seasonality, or none at all. Our analyses and planning charts still apply in these cases. However, it may be more sensible to use a flat planner track, $T_i = \boldsymbol{T}/12$ for all i, rather than to estimate a historical track from past years' seasonalities. To choose between the two, it is reasonable to compute the overall estimator $\tilde{\omega}^2$ using both candidate tracks and to use the track that yields the smaller value of $\tilde{\omega}^2$.

A further consideration is the choice of the parameter θ, or equivalently $w_i^{(H)}$, in the weighting scheme used in (8) and (14). If the pattern of seasonality in the data remains fairly static over the years, it is reasonable to choose $\theta = 0$ or $w_1^{(H)} = 1/H$, giving equal weight to each year. If the pattern of seasonality is evolving, a larger value of θ or $w_1^{(H)}$ is appropriate. In our experience, the choice $w_1^{(H)} = \frac{1}{2}$ often agrees with planners' intuitive judgement of an appropriate weighting.

When using planner tracks, estimation of ω^2 using eqn. (6) requires one year of history for both the measurement and the track. When using historical tracks and eqn. (15), two years of history are required. Sometimes the amount of history available is insufficient: in these cases it is often possible to obtain an estimate of ω^2 from data for another measurement. For example, for a new product where there is no history, a suitable estimate of ω^2 would be an estimate $\tilde{\omega}^2$ derived from data for a similar or predecessor product. Another example arises when planner tracks are used and history is available for the actual sales data but not for the track. Historical tracks can be calculated for past years, and the resulting estimate of ω^2 can be applied to the current year's planner track. This estimate of ω^2 may be further modified subjectively: for example, if it is thought that planner tracks tend to be more accurate than historical tracks, then the estimate of ω^2 obtained from historical tracks can be reduced before being used with the planner track.

5. Three charts for making assessments

5.1. The WINEGLASS chart: 'Are sales on track?'

Exhibit 1 shows an example of a WINEGLASS chart, which derives its name from its shape. The chart is used to assess whether year-to-date sales are on track to make the annual target, or whether action is needed to achieve the target. 'On track' means that the actual performance is statistically consistent with paths that at the end of the year exactly reach the annual target. The chart shows, for each month: (i) year-to-date cumulative sales as a percentage of the year-to-date cumulative track, and (ii) upper and lower bounds that each month contain a prechosen proportion q, the *on-track level*, of the paths that at the end of the year exactly reach the annual target. We commonly choose the on-track level to be 0.80. If the year-to-date cumulative sales total is inside the bounds, then sales are on track to make the annual target. On the other hand, if the year-to-date cumulative sales total is outside the bounds, then sales are not on track to make the annual target and action may be required to bring sales back on track.

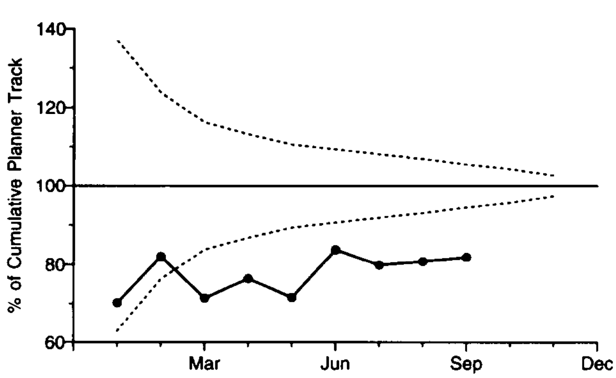

Exhibit 1. Example of a WINEGLASS chart.

The interpretation of 'on track' used with the WINEGLASS chart gives rise to an unusual situation for statistical inference. A planner might consider the annual track to be correct if $g = 1$, so that each Y_i has expectation T_i. But for assessing business performance, what matters is whether at the end of the year the target is met, and the WINEGLASS chart is designed to assess whether sales are in this sense on track. The chart tests whether the year-to-date sales performance, as measured by the year-to-date ratio of cumulative sales to the track,

$$\hat{g}_l = \frac{Y_1 + \cdots + Y_l}{T_1 + \cdots + T_l}, \tag{16}$$

is typical of a year that would meet the plan by having $Y_1 + \cdots + Y_{12} = T_1 + \cdots + T_{12}$, i.e. $\hat{g}_{12} = 1$. The WINEGLASS bounds are therefore based on the distribution of $\hat{g}_l$ conditional on $\hat{g}_{12} = 1$. Under the assumptions of model (1) it is easily seen that the random vector $[\hat{g}_l \quad \hat{g}_{12}]^{\mathrm{T}}$ has mean vector $[g \quad g]^{\mathrm{T}}$ and covariance matrix

$$\omega^2 g^2 \begin{bmatrix} \dfrac{T_1 + \cdots + T_{12}}{T_1 + \cdots + T_l} & 1 \\ 1 & 1 \end{bmatrix}. \tag{17}$$

Assuming Normality of the Y_i, it follows from standard theory [e.g. Rao (1973), p. 522)] that the conditional distribution of $\hat{g}_l$ is Normal with $E(\hat{g}_l \mid \hat{g}_{12} = 1) = 1$ and

$$\operatorname{var}(\hat{g}_l \mid \hat{g}_{12} = 1) = \omega^2 g^2 \frac{T_{l+1} + \cdots + T_{12}}{T_1 + \cdots + T_l}. \tag{18}$$

This variance is naturally estimated by setting $g = 1$, because 1 is the best linear unbiased estimator of g when $\hat{g}_{12} = 1$, and substituting $\tilde{\omega}^2$ for ω^2. This yields the 'WINEGLASS variance'

$$VW_I = \tilde{\omega}^2 \frac{T_{I+1} + \cdots + T_{12}}{T_1 + \cdots + T_I} . \qquad (19)$$

Thus the WINEGLASS bounds are constructed to contain the central $100q\%$ of a Normal distribution with mean 1 and variance given by eqn. (19). Letting ξ_t denote the $100t$ percentile of the standard Normal distribution, the WINEGLASS bounds for month I are $1 \pm \xi_{(1+q)/2} VW_I^{1/2}$. Note that these bounds can be calculated even before any data are measured in the current year.

An alternative method of constructing the WINEGLASS bounds is to use the on-track level q to set the simultaneous probability that an entire year's sequence $\hat{g}_1, \ldots, \hat{g}_{11}$, will all lie within the bounds, rather than the marginal probability that each month's $\hat{g}_I$ will lie within its own bounds. WINEGLASS bounds based on simultaneous probability are discussed in Wu, Hosking and Doll (1990). We prefer to use marginal probabilities, because at month I we want to base our assessments only on the year-to-date cumulative sales, not on the entire sequence of sales in the preceding months. In particular, if year-to-date sales at month I are within the WINEGLASS bounds, it seems natural to conclude that sales are on track even if in some previous month the year-to-date sales were outside the bounds; however, a WINEGLASS chart constructed using simultaneous probability would under these circumstances compel us to conclude that sales were not on track.

5.2. The SHIPWRECK chart: 'Will we recover from our deficit by year's end?'

Exhibit 2 shows an example of a SHIPWRECK chart, which derives its name from its shape. The chart shows the year-to-date cumulative deviation of sales from the track, together with a lower bound which marks the largest deficit from which we can reasonably hope to recover: specifically, if the actual deviation falls below the lower bound, then, even assuming that sales for the rest of the year conform to the plan in the sense that $g = 1$ in model (1), there is less than a prechosen probability p, the *recovery level*, that the full-year sales will reach the annual target. We commonly choose the recovery level to be 0.10. If the deviation from the track lies above the bound, then there is probability at least p that the annual target will be attained, provided that sales for the rest of the year are distributed about a mean level equal to the track. If the deviation lies below the bound, then the probability that the annual target will be attained is less than p. This is another indication that action may be needed to bring sales back on track.

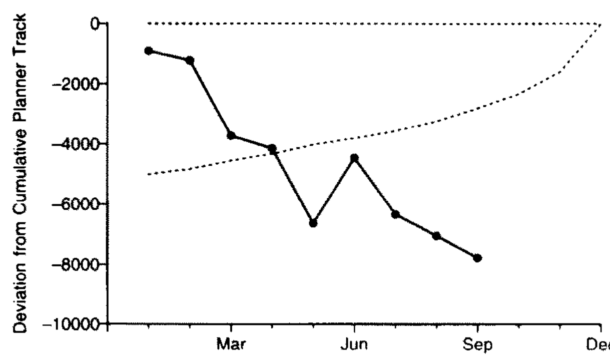

Exhibit 2. Example of a SHIPWRECK chart.

To calculate the SHIPWRECK bounds for month I, assume that the remaining months' sales follow model (1) with $g = 1$. The amount of deficit that is recovered in the remainder of the year is $(Y_{I+1} + \cdots + Y_{12}) - (T_{I+1} + \cdots + T_{12})$, which under the assumptions has mean 0 and variance $\omega^2 \boldsymbol{T}(T_{I+1} + \cdots + T_{12})$. The natural estimator of this variance is the 'SHIPWRECK variance'

$$VS_I = \tilde{\omega}^2 \boldsymbol{T}(T_{I+1} + \cdots + T_{12}) . \qquad (20)$$

Again assuming Normality of the Y_i, the SHIPWRECK bounds are constructed to contain all but the lower $100p\%$ of a Normal distribution with mean 0 and variance given by eqn. (20). The bound for month I is therefore $\xi_p VS_I^{1/2}$. Like the WINEGLASS bounds, it can be calculated before any data are obtained for the current year.

To compare the SHIPWRECK and WINEGLASS bounds, note first that a SHIPWRECK bound with recovery level p corresponds to a WINEGLASS lower bound with on-track level $q = 1 - 2p$: each bound is exceeded with probability p when the charts' respective assumptions are satisfied. Note also that the cumulative deviation of sales from the track may be written as $(\hat{g}_I - 1)(T_1 + \cdots + T_I)$. Thus the SHIPWRECK chart tests the deviation of $\hat{g}_I$ from 1, using the variance $VS_I/(T_1 + \cdots + T_I)^2$. This variance exceeds the WINEGLASS variance:

$$\frac{VS_I}{(T_1 + \cdots + T_I)^2} = \tilde{\omega}^2 \boldsymbol{T} \frac{T_{I+1} + \cdots + T_{12}}{(T_1 + \cdots + T_I)^2}$$
$$> \tilde{\omega}^2 \frac{T_{I+1} + \cdots + T_{12}}{T_1 + \cdots + T_I} = VW_I \,. \tag{21}$$

Thus a data point that exceeds the SHIPWRECK bound for recovery level p must also exceed the lower WINEGLASS bound for on-track level $1-2p$. Or, to put it more succinctly, it is easier to break a wineglass than to wreck a ship.

When assessing a 'low objective' measurement, such as expense, the natural question is not 'Will we recover from our deficit?' but 'Will we eliminate the excess of expenditure over the track?'. This question can be answered with an inverted version of the SHIPWRECK chart. Plot the cumulative deviation of expense from the track, together with upper bounds $\xi_{1-p} VS_i^{1/2}$, $i = 1, \ldots, 12$. From its shape, we call this chart the ROOFTOP chart. Its function may perhaps be described as to give warning when expenditure is going through the roof.

5.3. The OUTLOOK chart: 'What is our outlook range for the full year?'

Exhibit 3 shows an example of an OUTLOOK chart, which derives its name from its purpose. The chart shows, for each month thus far in the current year, three outlooks of the full-year sales total $Y_1 + \cdots + Y_{12}$. Outlooks made at month I are based on $Y_1, \ldots, Y_I$. Specifically, at month I a distribution of forecasts of the full-year sales total is computed, and outlooks are defined to be values that are attained by prechosen proportions of the distribution of forecasts. The OUTLOOK chart shows low, medium and high outlooks, corresponding to proportions p_L, p_M and p_H of the distribution of forecasts. We commonly use $p_L = 0.90$, $p_M = 0.50$ and $p_H = 0.10$. The purpose of the chart is to show the outlooks made at the current month, and to show how the outlooks have changed through the year. If the outlook range starts out wide in January and converges as more data become available, then we would have confidence in our outlooks. Exhibit 3 is an example of converging outlooks. On the other hand, if the outlooks fluctuate drastically from month to month with ranges that hardly overlap, this is an indication that the data are not following model (1) and therefore that outlooks made from the model are unreliable.

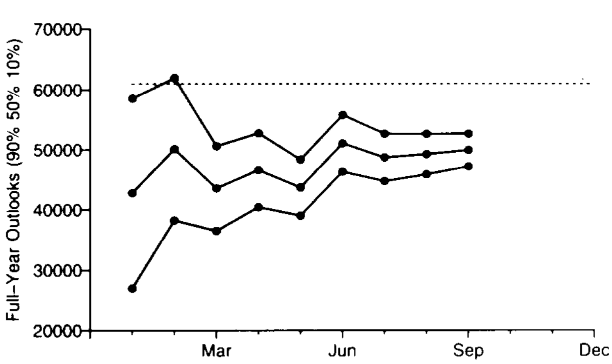

Exhibit 3. Example of an OUTLOOK chart.

At month I, the best linear unbiased estimator of the current year's future sales $Y_{I+1} + \cdots + Y_{12}$ is $\hat{g}_I(T_{I+1} + \cdots + T_{12})$, and so the best estimator of the year's total sales is

$$Y_1 + \cdots + Y_I + \hat{g}_I(T_{I+1} + \cdots + T_{12}) = \hat{g}_I \boldsymbol{T} \,. \tag{22}$$

The full-year sales may be written as $\hat{g}_{12}\boldsymbol{T}$, so the mean-square error of the estimator can be obtained from the joint second-order moments of $\hat{g}_I$ and $\hat{g}_{12}$ given in Subsection 5.1: it is

$$\boldsymbol{T}^2 E(\hat{g}_I - \hat{g}_{12})^2 = \omega^2 g^2 \boldsymbol{T}^2 \frac{T_{I+1} + \cdots + T_{12}}{T_1 + \cdots + T_I} \,. \tag{23}$$

This quantity is naturally estimated by the 'OUTLOOK variance'

$$VO_I = \tilde{\omega}^2 \hat{g}_I^2 \boldsymbol{T}^2 \frac{T_{I+1} + \cdots + T_{12}}{T_1 + \cdots + T_I} \,. \tag{24}$$

Thus at month I the best forecast of the full-year sales total is that it is distributed with mean $\hat{g}_I \boldsymbol{T}$ and variance VO_I, and outlooks are percentiles of this distribution. The proportion-p outlook is the value that is greater than or equal to proportion p of this forecast distribution; assuming Normality of the Y_i, it is given by $\hat{g}_I \boldsymbol{T} + \xi_{1-p} VO_I^{1/2}$. Low, medium and high outlooks are obtained by choosing suitable values $p_L > p_M > p_H$ of P.

Outlooks can also be made by specifying some level of full-year sales and finding the proportion of forecasts that attain this level. Often the specified level will be the annual target or some level related to it, such as the annual target plus 5%. These 'plan-based outlooks' are discussed in Wu, Hosking and Doll (1990).

An alternative method of constructing outlooks is possible. The OUTLOOK variance (24) is obtained by substituting $\hat{g}_I$ for g in eqn. (23), and does not take into account the variability of $\hat{g}_I$ itself. Outlooks can be constructed which take this variability into account. The computations are much more complicated but make very little difference to the final outlooks: details are given in Wu, Hosking and Doll (1990). We therefore consider the estimator (24) to be adequate.

To compare outlooks with the WINEGLASS bounds, note first that WINEGLASS bounds with on-track level q correspond to outlooks with proportions $p_L = (1+q)/2$ and $p_H = (1-q)/2$: each quantity then marks the upper or lower $100(1-q)/2$ percentile of the distribution from which it is calculated. Note also that the outlook at month I of full-year sales divided by the annual target is $\hat{g}_I$ and has mean-square error

$$VO_I/\boldsymbol{T}^2 = \hat{\omega}^2 \hat{g}_I^2 \frac{T_{I+1} + \cdots + T_{12}}{T_1 + \cdots + T_I} = \hat{g}_I^2 VW_I \,. \tag{25}$$

Thus the variance used to calculate the outlooks is less than the WINEGLASS variance when sales are below the track ($\hat{g}_I < 1$), but greater than the WINEGLASS variance when sales are above the track ($\hat{g}_I > 1$). This means that the low outlook becomes less than the annual target before the lower WINEGLASS bound is exceeded, whereas the upper WINEGLASS bound is exceeded before the high outlook becomes greater than the annual target.

When assessing a 'low objective' measurement, outlooks are still defined to be values that are attained by prechosen proportions of the distribution of forecasts of the full-year total. However, the natural interpretation of this definition is that the proportion-p outlook is the value that is *less than* or equal to the proportion p of the distribution of forecasts. Thus the proportion-p outlook at month I is $\hat{g}_I \boldsymbol{T} + \xi_p VO_I^{1/2}$.

6. Examples

6.1. Shipments

This example describes how the charts shown in Exhibits 1–3 were constructed. The data are for shipments of a product: this is a 'high objective' measurement. The analysis uses planner tracks and three years of history. The data are given in Exhibit 4 and illustrated in Exhibit 5.

First we estimate ω^2, using 1986–1988 shipments and planner tracks. We assume that previous years' data follow eqn. (2). For each year we calculate $\tilde{g}_j$ from eqn. (3) and $\tilde{\omega}_j^2$ from eqn. (5), obtaining $\tilde{\omega}_{86} = 5.19\%$, $\tilde{\omega}_{87} = 8.37\%$ and $\tilde{\omega}_{88} = 5.83\%$. The three estimates of ω^2 are combined using eqn. (6). We choose the weight applied to the most recent year of history, 1988, to be 0.5, i.e. $w_1^{(3)} = 0.5$. We solve eqn. (13) for θ with $H = 3$ and $j = 1$, obtaining $\theta = 0.47$. Using this value of θ in eqn. (13), we get $w_2^{(3)} = 0.29$ and

Exhibit 4
Planner tracks and shipments for a product for 1986 through 1989.

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Total
1986 track	2068	2948	4452	3282	4232	5060	3340	4200	4712	4792	5034	6780	50900
1986 ships	1061	2469	2986	2103	3697	3765	2321	2923	3857	3153	3315	3193	34843
1987 track	1917	2562	3263	2404	4279	4337	2064	2940	3663	2452	3418	2976	36275
1987 ships	1817	2473	3311	2517	3415	2699	2372	1370	4693	2921	4800	3564	35952
1988 track	2221	3914	4440	2617	4561	4822	2835	4014	4951	3581	4998	4246	47200
1988 ships	3272	3779	5935	4201	4693	5758	2329	4441	6588	4787	4684	3659	54126
1989 track	3012	3833	6155	4581	5866	4003	4051	5167	6219	5586	6672	5855	61000
1989 ships	2113	3509	3658	4156	3361	6193	2166	4457	5492				

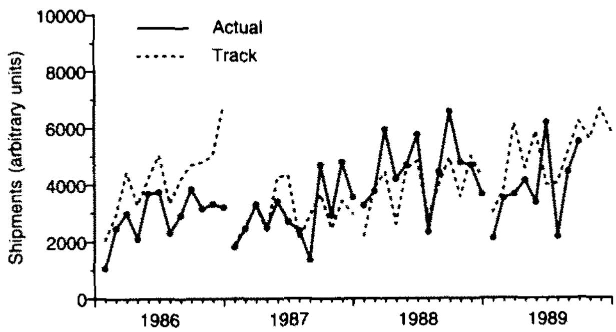

Exhibit 5. Planner tracks and shipments for a product for 1986 through 1989.

$w_3^{(3)} = 0.21$. Equation (6) now gives $\tilde{\omega} = 0.004305$, i.e. $\tilde{\omega} = 6.56\%$.

We choose the on-track level for the WINEGLASS chart to be 0.80. To compute the WINEGLASS bounds we need the percentile $\xi_{(1+q)/2}$ of the standard Normal distribution: it is $\xi_{0.9} = 1.282$. For January the WINEGLASS variance, found from eqn. (19), is $VW_1 = 0.08289$, and the numbers plotted on the WINEGLASS chart are: (i) the year-to-date shipments as a percentage of the year-to-date track,

$$100 \times \frac{2113}{3012} = 70.2 ;$$

(ii) the upper WINEGLASS bound,

$$100 \times (1 + \xi_{0.9} VW_1^{1/2}) = 136.9 ;$$

and (iii), the lower WINEGLASS bound,

$$100 \times (1 - \xi_{0.9} VW_1^{1/2}) = 63.1 .$$

Similar computations can be made for the other months. The resulting WINEGLASS chart is Exhibit 1. Since the year-to-date cumulative shipments total is below the lower bound, shipments are not on track to make the annual plan and action is likely to be required to bring them back on track.

From eqn. (20), the SHIPWRECK variance for January is $VS_1 = 1.5229 \times 10^7$. We choose the recovery level p to be 0.10, and we have $\xi_p = \xi_{0.1} = -1.282$. The SHIPWRECK bound for January is therefore $\xi_{0.1} VS_1^{1/2} = -5003$. The cumulative deviation of shipments from track for January is $2113 - 3012 = -899$. Corresponding values for other months can be computed similarly. The resulting SHIPWRECK chart is Exhibit 2.

From eqn. (24) the OUTLOOK variance for January is $VO_1 = 1.518 \times 10^8$. The outlooks for January can now be computed as described at the end of Subsection 5.3. We choose the outlook proportions to be 0.9, 0.5 and 0.1. Because shipments is a 'high objective' measurement, outlooks are defined to be values that are greater than or equal to the proportions 0.9, 0.5 and 0.1 of the distribution of forecasts. The proportion-0.9 outlook is

$$\hat{g}_1 \boldsymbol{T} + \xi_{0.1} VO_1^{1/2} = 26\,999 .$$

Similarly the proportion-0.5 outlook is 42 793 and the proportion-0.1 outlook is 58 588. Similar computations are made for the other months. The resulting OUTLOOK chart is Exhibit 3.

6.2. Cost

This example is for the cost of producing a product: this is a 'low objective' measurement. The analysis uses historical tracks and two years of history. The data are given in Exhibit 6 and illustrated in Exhibit 7. Exhibit 7 also contains the historical tracks for 1988 and 1989, as computed below and given in Exhibit 8. The annual target for 1989, the current year, is 4500.

The historical track for 1989 is computed using eqns. (7) and (8). We choose the weight applied to the most recent year of history to be 0.5: i.e. $w_1^{(2)} = 0.5$. This corresponds to letting $\theta \to 0$ in eqn. (13) and implies that $w_2^{(2)} = 0.5$.

Exhibit 6
Cost of producing a product for 1987 through 1989.

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Total
1987 cost	236	272	343	248	315	442	286	300	370	304	290	477	3883
1988 cost	212	266	348	277	336	401	292	364	442	365	465	568	4336
1989 cost	250	215	385	240	485								

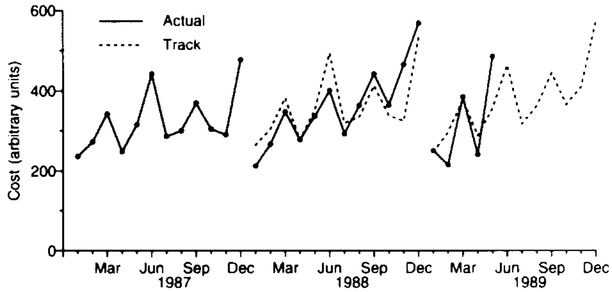

Exhibit 7. Historical tracks and actual cost of producing a product for 1987 through 1989.

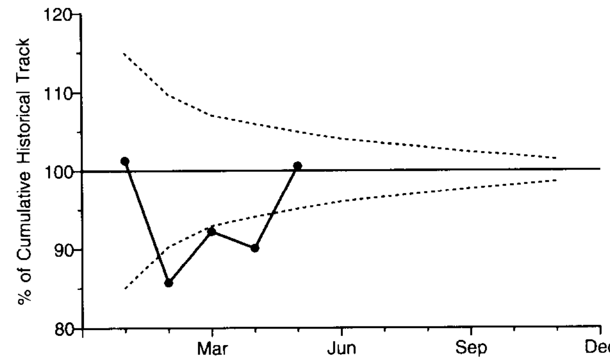

Exhibit 9. WINEGLASS chart for cost.

The historical track for 1989 is tabulated in the second row of Exhibit 8.

To estimate ω^2 we compute the accuracy of historical tracks for previous years. There are only two years of history, so the only 'previous year' for which a historical track can be computed is 1988. Using eqn. (14) with $j = 1$ we can compute the historical track for 1988 cost, based on 1987 seasonality and the 1988 annual target $\boldsymbol{T}_{/88}$. Since the magnitude of $\boldsymbol{T}_{/88}$ does not affect $\tilde{\omega}^2$, as noted in the discussion following eqn. (15), we choose $\boldsymbol{T}_{/88}$ to be 4336, the same as the actual total cost for 1988. This ensures that $\tilde{g}_{88} = 1$. The estimate of ω^2 from the 1988 data is computed from eqn. (5): since $\tilde{g}_{88} = 1$, we have

$$\tilde{\omega}^2_{88} = \frac{1}{11} \sum_{i=1}^{12} \frac{(Y_{i/88} - T_{i/88})^2}{T_{i/88}\boldsymbol{T}_{/88}} = 0.002279\,,$$

i.e. $\tilde{\omega}_{88} = 0.0477 = 4.77\%$. Because there are only two years of history, the weighted average in eqn. (15) reduces to $\tilde{\omega}^2 = \tilde{\omega}^2_{88}$.

Construction of the WINEGLASS chart proceeds exactly as in the previous example. A WINEGLASS chart with on-track level 0.50 is shown in Exhibit 9. Since the year-to-date cumulative total of cost is within the bounds, cost is on track to make the annual target.

Because cost is a 'low objective' measurement, we use a ROOFTOP chart rather than a SHIPWRECK chart. The computations are almost the same as for a SHIPWRECK chart: for example, from eqn. (20), the ROOFTOP variance for January is

$$VS_1 = \tilde{\omega}^2\boldsymbol{T}(T_2 + \cdots + T_{12}) = 43\,620\,.$$

We choose the recovery level p to be 0.05, and we have $\xi_{1-p} = \xi_{0.95} = 1.645$. The ROOFTOP bound for January is therefore $\xi_{0.95} VS_1^{1/2} = 344$. The cumulative deviation of cost from track for January is +3. Corresponding values for other months can be computed similarly. The resulting ROOFTOP chart is shown in Exhibit 10.

Outlooks can now be computed as described at the end of Subsection 5.3. We choose the

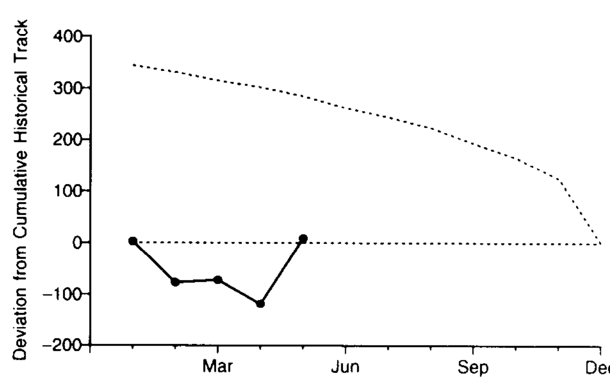

Exhibit 10. ROOFTOP chart for cost.

Exhibit 8
Historical tracks for cost for 1988 and 1989.

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Total
1988 track	264	304	383	277	352	494	319	335	413	339	324	533	4336
1989 track	247	296	379	287	357	464	317	363	444	366	409	571	4500

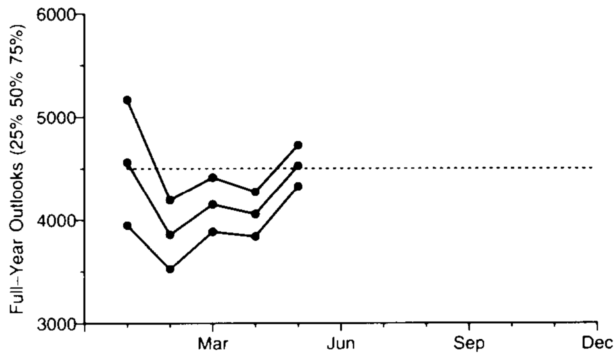

Exhibit 11. OUTLOOK chart for cost.

outlook proportions to be 0.25, 0.5 and 0.75. Because cost is a 'low objective' measurement, outlooks are defined to be values that are less than or equal to the proportions 0.25, 0.5 and 0.75 of the distribution of forecasts. For January the proportion-0.25 outlook is

$$\hat{g}_1 \boldsymbol{T} + \xi_{0.25} VO_1^{1/2} = 3950 ,$$

the proportion-0.5 outlook is 4559 and the proportion-0.75 outlook is 5168. Corresponding values for other months can be computed similarly. The resulting OUTLOOK chart is shown in Exhibit 11.

7. Concluding remarks

The WINEGLASS, SHIPWRECK and OUTLOOK charts were first described by Wu (1988). Model (1) is essentially the Total Error Equation of the Track Uncertainty Model of Wu (1988). There are some notational differences between the models, but the only significant difference is that the variance of Y_i in model (1) is $\omega^2 gT_i g\boldsymbol{T}$, whereas in the Track Uncertainty Model it was, in our current notation, $(K + C)gT_i$. As noted in Section 3, the new form of the variance makes it straightforward to combine variance estimates made from data with different numerical scales. Wu (1988) also discussed how to model the 'inherent variability' component of planning uncertainty. For some kinds of data (e.g. revenue at the corporate level, or total sales of a high-volume product with no data available for sales disaggregated across regions or business units), it is difficult to estimate the inherent variability separately from the systematic error in apportioning the track, the other component of τ_i. We have therefore preferred to work throughout with assessments and planning charts based on the total error τ_i. The WINEGLASS bounds in Wu (1988) are parameterized in terms of p, the probability that year-to-date sales will lie outside the bounds. The on-track level q used in Subsection 5.1 is of course $1 - p$.

The WINEGLASS, SHIPWRECK and OUTLOOK charts described in Section 5 involve calculations that assume that the Y_i have a Normal distribution. For some measurements, such as sales of a low-volume product, that are counts of small numbers of units, the monthly numbers are small and discrete and the assumption of Normality is clearly invalid. Such data are often well described by a Poisson distribution [Wu (1988)]. To assume a Normal distribution when the true distribution is Poisson causes small inaccuracies in the three charts. However, these inaccuracies are in most practical situations too small to matter: an example is given in Wu, Hosking and Doll (1990). We therefore consider it reasonable to construct WINEGLASS bounds, SHIPWRECK bounds and outlooks assuming Normality of the Y_i, even though the data may be small and discrete.

The charts in Section 5 are based on dividing a year into 12 months for planning purposes, but similar charts can be constructed for other planning frameworks. For example, if planning is done on a quarterly basis, then model (1) may still be used but with the subscript i denoting a quarter rather than a month. Calculations similar to those of Sections 4 and 5 can be made to obtain WINEGLASS bounds, SHIPWRECK bounds and outlooks.

Similar circumstances can arise when, as sometimes happens, a planner track is reset during the course of a year. Suppose, for example, that it is recognized early in the year that the annual target will not be met, and that at the end of June a new target is set of the second half of the year and a new planner track is constructed to reach the target. It then seems most reasonable to regard the months July through December as a 'year' in themselves, and to base assessments of whether the new target will be met on a WINEGLASS chart constructed for this 6-month 'year'.

Software

WINEGLASS, an IBM Experimental Software program for the IBM PC to calculate the Wineglass chart and outlooks, is available at no charge from the first author. The program requires DOS 2.1 or higher and 350K of memory. It is available on $3\frac{1}{2}$ in. or $5\frac{1}{4}$ in. diskette.

Acknowledgements

We are grateful for the many contributions and ideas of Byrd Ball, without whom this work would not have been possible. We also thank Eugene Bershatsky, Charles Bruce, Betty Flehinger, David Ing, and John Patrick, who have contributed to this work and related work in the area of planning uncertainty.

References

Abraham, B. and J. Ledolter, 1983, *Statistical Methods for Forecasting* (Wiley, New York).

Bretschneider, S., 1991, "Some important issues related to the gap between research of forecasting and the practice of forecasting", *International Institute of Forecasters Newsletter*, 4, no. 4.

Brockwell, P.J. and R.A. Davis, 1987, *Time Series: Theory and Methods* (Springer-Verlag, New York).

Rao, C.R., 1973, *Linear Statistical Inference and Its Applications*, 2nd edn. (Wiley, New York).

Wu, L.S.-Y., 1988, "Business planning under uncertainty: Quantifying variability", *The Statistician*, 37, 141–152.

Wu, L.S.-Y., J.R.M. Hosking and J.M. Doll, 1990, "Business planning under uncertainty: Will we attain our goal?", Research Report RC16120, IBM Research Division, Yorktown Heights, NY.

Wu, L.S.-Y., N. Ravishanker and J.R.M. Hosking, 1991, "Forecasting for business planning: A case study of IBM product sales", *Journal of Forecasting*, 10, 579–595.

Biographies: Lilian Shiao-Yen WU has been a Research Staff Member at the T.J. Watson Research Center of IBM Research Division since 1974. She received a Ph.D. in Applied Mathematics from Cornell University in 1974. Her interests include forecasting, business planning and data analysis.

J.R.M. HOSKING has been a Research Staff member at the T.J. Watson Research Center of IBM Research Division since 1986. He received a Ph.D. in Statistics from the University of Southampton in 1979. His interests include time-series analysis, distribution theory and flood frequency analysis.

Jeanne M. DOLL is a Senior Planner and Statistical Consultant in Finance and Planning for the IBM US Division. She received a Ph.D in Statistics from Rutgers University in 1990. Her interests include forecasting, time series analysis and empirical Bayes estimation.