Published in Graham J. Hooley and Michael K. Hussey (Eds.),
Quantitative Methods in Marketing, Second Edition.
London: International Thompson Business Press, 1999, pp. 92-119.

Forecasting for Marketing

J. Scott Armstrong
The Wharton School, University of Pennsylvania

Roderick J. Brodie
Department of Marketing, University of Auckland

Research on forecasting is extensive and includes many studies that have tested alternative methods in order to determine which ones are most effective. We review this evidence in order to provide guidelines for forecasting for marketing. The coverage includes intentions, Delphi, role playing, conjoint analysis, judgmental bootstrapping, analogies, extrapolation, rule-based forecasting, expert systems, and econometric methods. We discuss research about which methods are most appropriate to forecast market size, actions of decision makers, market share, sales, and financial outcomes. In general, there is a need for statistical methods that incorporate the manager's domain knowledge. This includes rule-based forecasting, expert systems, and econometric methods. We describe how to choose a forecasting method and provide guidelines for the effective use of forecasts including such procedures as scenarios.

INTRODUCTION

Forecasting has long been important to marketing practitioners. For example, Dalrymple (1987), in his survey of 134 U.S. companies, found that 99 percent prepared formal forecasts when they used formal marketing plans. In Dalrymple (1975), 93 percent of the companies sampled indicated that sales forecasting was one of the most critical' aspects, or a 'very important' aspect of their company's success. Jobber, Hooley and Sanderson (1985), in a survey of 353 marketing directors from British textile firms, found that sales forecasting was the most common of nine activities on which they reported.

Managers' forecasting needs vary considerably. They may need to forecast the size and growth of a market or product category. When strategic issues are being considered, they need to forecast the actions and reactions of key decision makers such as competitors, suppliers, distributors, governments, their own actions, and complementors' (organizations with whom they cooperate) actions. These actions can help to forecast market share. The resulting forecasts allow one to calculate a sales forecast. If strategic issues are not important, one can extrapolate sales directly. Finally, by forecasting costs and using the sales forecast, one can forecast profits and other financial outcomes. These forecasting needs and their relationships are illustrated in Figure 6.1.

The purpose of this chapter is to provide guidance to managers about the use of formal forecasting methods in marketing. In developing the guidelines, it is recognized that managers may have negative attitudes towards the usefulness of formal forecasting. This may have occurred because they have used poor forecasts in the past, because they had unrealistic expectations about forecasting accuracy, or because they do not like it when forecasts conflict with their beliefs about the future.

FIG 6.1. Needs for marketing forecasts

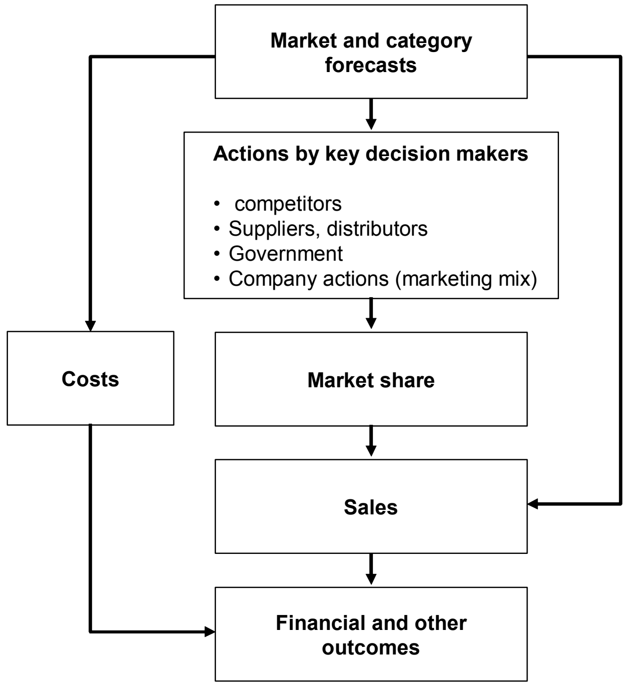

The guidelines draw on the evidence collected in the Forecasting Principles Project, which is described on the website forecastingprinciples.com. The evidence is provided in *Principles of Forecasting: A Handbook far Researchers and Practitioners* (2001) edited by Scott Armstrong.

FORECASTING METHODS

Forecasting involves methods that derive from judgmental sources and from statistical sources. These methods and the relationships between them are shown in the flowchart in Figure 6.2. Going down the flowchart, there is an increasing amount of integration between judgmental and statistical data and procedures. This integration, which has been studied by researchers in the last decade, can improve forecast accuracy (Armstrong and Collopy, 1998).

We provide only a brief description of the methods and their application. More detailed descriptions are provided in forecasting textbooks such as Makridakis, Wheelwright and Hyndman (1998).

FIG 6.2. Characteristics of forecasting methods and their relationships
(dotted lines represent possible relationships)

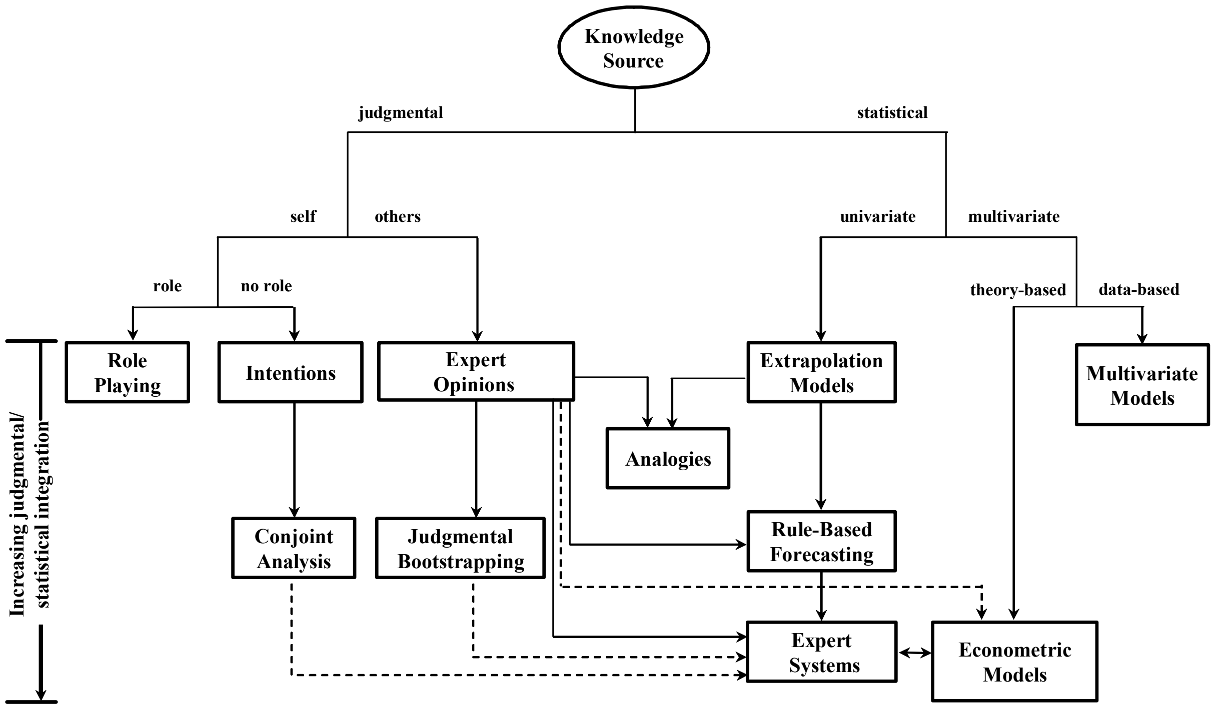

Methods Based on Judgment

Intentions

With intentions surveys, people are asked to predict how they would behave in various situations. Intentions surveys are widely used when sales data are not available, such as for new product forecasts. There is much empirical research about the best way to assess intentions and Morwitz (2001) draws upon this to develop principles for using intentions in forecasting.

Role playing

A person's role may be a dominant factor in some situations, such as in predicting how someone in a firm would behave in negotiations. Role playing is useful for making forecasts of the behavior of individuals who are interacting with others, and especially when the situation involves conflict. The key principle here is to provide a realistic simulation of the interactions. It is a method that has considerable potential for forecasting although, currently, it is seldom used (Armstrong, 2001b).

Expert opinions

Expert opinion studies differ substantially from intentions surveys. When an expert is asked to predict the behavior of a market, there is no need to claim that this is a representative expert. Quite the contrary, the expert may be exceptional.

One principle is to combine independent forecasts from a group of experts, typically 5 to 20 (Ashton and Ashton, 1985). The required level of expertise is surprisingly low (Armstrong 1985). The preferred procedure is to weight each expert's forecast equally.

The accuracy of expert forecasts can be improved through the use of structured methods, such as the Delphi procedure. Delphi is an iterative survey. procedure in which experts make forecasts for a problem, receive anonymous summary feedback on the forecasts made by other experts, and then make a further forecast. For a summary of the evidence on the accuracy of Delphi versus unstructured judgment, see Rowe and Wright (2001). One principle is that experts' forecasts should generally be independent of one another. Focus groups always violate this principle; as a result, they should not be used in forecasting.

Conjoint analysis

Intentions can be explained by relating consumers' intentions to various factors that describe the situation. For example, by asking consumers to state their intentions to purchase for a variety of different product offerings, it is possible to infer how the factors relate to intended sales. This can be done by regressing intentions against the factors, a procedure which is known as "conjoint analysis." It is a method that is used in many areas of marketing, but particularly for new product decisions. It is based on sound principles, such as using experimental design to create situations and soliciting independent intentions from a sample of potential customers (Wittink and Bergesteun, 2001).

Judgmental bootstrapping

As with conjoint analysis, one can develop a model of the expert. This approach, known as judgmental bootstrapping, converts subjective judgments into objective procedures. Experts are first asked to make predictions for a series of conditions. For example, they could make forecasts for next year's sales in various geographical regions. This process is then converted to a set of rules by regressing the forecasts against the information used by the forecaster. Once developed, judgmental bootstrapping models offer a low-cost procedure for making forecasts. They almost always provide an improvement in accuracy in comparison to judgmental forecasts, although these improvements are typically modest (Armstrong, 2001a).

Methods Based on Statistical Sources

Extrapolation

Extrapolation methods use historical data on the series of interest. Exponential smoothing is the most popular and cost effective of the extrapolation methods. It implements the principle that more recent data should be weighted more heavily and also seeks to "smooth" out seasonal and/or cyclical fluctuations to predict the direction in which the trend is moving.

Alternatively, one may simply make judge mental extrapolations of historical data. Judgmental extrapolations are preferable to quantitative extrapolations when there have been large recent changes in the sales level and where there is relevant knowledge about the item to be forecast (Armstrong and Collopy, 1998).

An important principle for extrapolation is to use long time series when developing a forecasting model. Yet Focus Forecasting, one of the most widely used time series forecasting software packages, does not do this; as a result, its forecasts are less accurate than alternative procedures (Gardner anti Anderson, 1997).

Another principle for extrapolation is to use reliable data. 'the existence of retail scanner data means that reliable data can be obtained for existing products. Scanner data is detailed, accurate, timely, and inexpensive. As a result, forecast accuracy should improve, especially because of the reduction in the error of assessing the current status. Not knowing where you are starting from has often been a major source of error in predicting future values. Scanner data may also be used for early identification of trends.

Empirical studies have led to the conclusion that relatively simple extrapolation methods perform as well as more complex methods. For example, the Box Jenkins procedure, one of the more complex approaches, has produced no measurable gains in forecast accuracy relative to simpler procedures (Makridakis, et al., 1984;

Armstrong, 1985). More recently forecasters have extensively examined another complex procedure, neural networks. Neural networks are computer intensive methods that use decision processes analogous to that of the human brain. Like the brain, they have the capability of learning and updating their parameter estimates as experience is acquired. Neural networks have not produced more accurate forecasts than other methods (Chatfield, 1993). Some promising work has been done in the study of market share forecasting (Agrawal and Schorling, 1996) and in predicting consumer's choice (West et al., 1997), but our advice at present is to ignore neural nets.

Rule-based forecasting

Quantitative extrapolation methods make no use of managers' knowledge of the time series. The basic assumption is that the causal forces that have affected a historical series will continue over the forecast horizon. This assumption is sometimes false. Rule based forecasting is a type of expert system that allows one to integrate managers' knowledge about the domain with time series data in a structured and inexpensive way. For example, in many cases a useful guideline is that trends should be extrapolated only when they agree with managers' prior expectations. When the causal forces are contrary to the trend in the historical series, forecast errors tend to be large (Armstrong and Collopy, 1993). While such problems may occur only in a small percentage of cases, their effects can be disastrous.

Analogies

Experts can identify analogous situations. Extrapolation of results from these situations can be used to predict the situation of interest (Duncan, Gore and Szczypula, 2001). For example, to assess the loss in sales when the patent protection for a drug is removed, one might examine results for previous cases, especially if the drugs are similar.

Expert systems

As the name implies, expert systems use the rules of experts. These rules are typically created from protocols, whereby the forecaster talks about what he is doing while making forecasts. The real promise, however, is for expert systems to draw upon empirical results of relationships that come from econometric studies. In fact, this is a common way to construct expert systems. Expert opinion, conjoint analysis and bootstrapping can also aid in the development of expert systems.

Multivariate time series methods

Despite much research effort, there is little evidence that multivariate time series provide benefits for forecasting. As a result, these methods are not discussed here.

Econometric methods

Econometric methods use prior knowledge (theory) to construct a model. This involves selecting causal variables, identifying the expected directions of the relationships, imposing constraints on the relationships to ensure that they are sensible, and selecting functional forms. In most marketing problems, one can also make reasonable prior estimates for the magnitude of the relationships, such as for price or advertising elasticities. Data from the situation can then be used to update the estimates, especially if one has sufficient amounts of relevant and reliable data.

Econometric models have the advantage that they can relate directly to planning and decision making. They can provide a framework to examine the effects of marketing activity as well as key aspects of the market and the environment, thus providing information for contingency planning. While some causal variables can be forecast with a reasonable level of accuracy (e.g. demographic changes), others are more difficult to forecast (e.g. changes in fashion and competitors' actions).

Econometric methods are most useful when:

- strong causal relationships with sales or other entities are expected;
- these causal relationships are known or they can be estimated;
- large changes are expected to occur in the causal variables over the forecast horizon; and
- these changes in the causal variables can be accurately forecast or controlled, especially with respect to their direction.

If any of these conditions does not hold (which is typical for short-range forecasts), econometric methods are less likely to contribute to accuracy.

FORECASTING MARKET SIZE

Market size is influenced by environmental factors such as economic conditions, population, ability to purchase, social trends, technological change, or government legislation. For example, demographic factors such as the size and age distribution of the population, distribution of disposable income, culture, and religious factors influence the market for alcoholic beverages.

Econometric methods have been used for environmental forecasting. Econometric researchers have devoted much effort to short-term forecasting, an area that has yielded unimpressive results. Econometric methods would be expected to be more useful for long-range forecasting because the changes in the causal variables are not swamped by random variations), as in the short run. Armstrong (1985, Chapter 15) reported seven empirical comparisons of methods used in long-range forecasting. In all comparisons, econometric methods were more accurate than extrapolations. Fildes (1985) located 20 studies on long-range forecasting; he found 15 where econometric methods were more accurate than other methods, 3 ties, and 2 showing econometric forecasts to be less accurate.

Improved environmental forecasts should lead to more accurate market forecasts. Surprisingly, research in this area indicates that forecasting errors are not particularly sensitive to the accuracy of environmental forecasts. Measurement error in the causal variables (e.g. the environmental inputs to a market forecasting model) had little impact on the accuracy of an econometric model in the few studies done on this topic (Armstrong, 1985). Moreover, conditional econometric forecasts (those made with actual data on the causal variables) have generally been found to be no more accurate than unconditional forecasts (where the causal variables themselves must be forecasted). Of 18 studies found, only 3 have shown conditional forecasts to be more accurate, 5 showed no difference, and 10 showed them to be less accurate (Armstrong, 1985; Rosenstone, 1983; and four studies from Fildes, 1985). A possible explanation for these strange findings is that the unconditional forecasts may have included subjective revisions that might have reduced the error in estimating starting values (current levels).

Methods based on judgment

Market forecasts are often based on judgment, especially; it seems, for relatively new or rapidly changing markets. This carries some risk, as research since the 1960s has identified biases that occur in judgmental forecasting. Among these biases are optimism, conservatism, anchoring, and availability.

The Delphi technique offers a useful way to implement many of the basic principles for expert forecasting. It uses (1) more than one expert, (2) unbiased experts, (3) structured questions, and (4) equal weights for each expert's forecast. It could be used to answer questions about market size such as: "By what percentage will the New Zealand wine market grow over the next 10 years?" But it is especially appropriate when one has scant relevant prior data. Thus, one might ask: "What proportion of U.S. households will subscribe to movies on demand over telephone or cable lines?"

Surprisingly, research indicates that high expertise in the subject area is not important for judgmental forecasts of change (Armstrong, 1985). The conclusion, then, is not to spend heavily to obtain the best experts to forecast change. On the other hand, one should avoid people who clearly have no expertise. Also, experts are helpful for assessing current levels.

Methods based on statistical sources

When considering forecasts of market size, one can use either time series extrapolation methods or econometric methods. Time series extrapolation is inexpensive. Econometric methods, while more expensive, are expected to be more accurate than judgmental methods and extrapolation.

Organizations should use systematic procedures for scanning the environment to be sure that they do not overlook variables that may have a large impact on their market. Periodic brainstorming with a group of experts should be sufficient to identify which variables to track, especially if the experts represent diverse areas. The key is to identify important variables and the direction of their effects. Once identified, crude estimates of the coefficients of these environmental variables are often sufficient in order to obtain useful forecasts.

When forecasting environmental factors related to market prices, it is important to remember what might be called Adam Smith's rule for forecasters: 'Forecasters cannot beat the market.' (Some people refer to this as the rule of efficient markets.) In other words, when an active market of buyers and sellers is at work (such as in stocks, bonds, money, commodities and land), forecasters have not had much success at finding methods that can improve upon the market's forecast. This rule assumes that the forecaster lacks inside information, so the market price is a reflection of available information. Judging from research since the 1930s, the market can forecast prices as effectively as can be done by any alternative forecasting procedure.

DECISION MAKERS' ACTIONS

The development of a successful marketing strategy sometimes depends upon having a good forecast of the actions and reactions of competitors. A variety of judgmental and statistical methods can be used to forecast competitive actions. These include:

- expert opinion (using experts who know about this and similar markets);
- intentions (ask the competitors how they would respond in a variety of situations);
- role playing (formal acting out of the interactions among decision makers for the firm and its competitors); and
- experimentation (trying the strategy on a small scale and monitoring the results). For example, if you lower your price, will competitors follow?

After forecasting their actions, the next step is to forecast their impact upon market size and market share.

It may also be important to forecast the actions of suppliers, distributors, complementors, government, and people within one's firm in order to develop a successful marketing strategy. Sometimes one may need to forecast the actions of other interest groups, such as 'concerned minorities.' For example, how would an environmental group react to the introduction of plastic packaging by a large fast food restaurant chain? A range of techniques similar to those for forecasting competitors' actions appears useful.

Role playing is well suited to predicting how decision makers will act. It provides substantially more accurate forecasts than can be obtained from expert opinion (Armstrong, 2001). In one study, role playing was used to forecast interaction between suppliers and distributors. Philco (called the Ace Company in the role play), a producer of home appliances, was trying to improve its share of a depressed market. They had developed a plan to sell appliances in supermarkets using a cash register tape discount plan. Secrecy was important because Philco wanted to be first to use this strategy. Implementation of such a plan depended upon the supermarket managers. Would the plan be acceptable to them? In this case, a simple role playing procedure produced substantially more accurate forecasts of the supermarket managers' decisions (8 of 10 groups were correct) than did unaided opinions (1 of 34 groups was correct). In the actual situation, the supermarket managers did accept the plan proposed by Philco. The superior accuracy of role playing relative to opinions seems to be due to its ability to provide a more realistic portrayal of the interactions.

Company plans typically require the cooperation of many people. For example, if the organization decides to implement a given marketing strategy, will it be able to carry out the plan? Sometimes an organization fails to do what it intends to do because of a lack of resources, misunderstanding, opposition by key stakeholders, or a lack of commitment to the plan by key people. The need to forecast organizational behavior is sometimes overlooked. Better forecasting here might lead to more realistic plans and to plans that are easier to implement.

Surveys of key decision makers in an organization may help to assess the likelihood that a given strategy can be implemented. Because those who are not committed to a plan may be reluctant to admit it, projective questions may be useful when asking about intentions to implement plans.

It is also important to predict the effects of the various actions. For example, in the Philco situation, the change in distribution channels led to substantial losses for all involved. One can make such forecasts by using expert judgment, judgmental bootstrapping, or econometric methods.

FORECASTING MARKET SHARE

The primary approaches to forecasting market share are:

- expert opinion;
- judgmental bootstrapping;
- extrapolation (statistical analysis of the market or of analogous markets);
- econometric methods (using relative prices, advertising, and product features).

In many cases, markets have reached a rough state of equilibrium. That is, the future is expected to consist of an extension of the same causal forces and the same types of actions. Under such conditions, a simple extrapolation of market share, such as the naive 'no change' model, is usually sufficient.

When large changes are expected, one should draw upon methods that incorporate causal reasoning. If tire changes are unusual, judgmental methods such as Delphi would be appropriate. If the changes are expected to be large, the causes are well understood, and if one lacks historical data, then judgmental bootstrapping is relevant.

In some cases, econometric models are relevant. They should be based on theory. Bell, Keeney and Little (1975) provide a start with their market share theorem, which states that market share changes depend upon the marketing effort for the brand (price, advertising, product, etc.) divided by the sum of marketing effort for all the brands in the market. There are many ways to formulate this model (Brodie et al., 2001)) and much prior research exists to help specify these models. For example, a meta-analysis by Tellis (1988) of 367 brand price elasticities, estimated using econometric models, reports a mean value of -2.5. Hamilton et al.'s (1997) analysis of 406 price brand elasticities also reported a value of -2.5. Generalizations can also be made about other measures of market activity, such as advertising elasticity.

In addition to being useful for policy issues, econometric models are sometimes more accurate forecasts than time-series extrapolations. Brodie et al.'s (2001) review of empirical evidence on market share forecasting concludes that econometric methods are most accurate when

1. there are strong causal relationships between the marketing mix variables and market share;
2. ample historical data exhibit sufficient variation to allow one to improve the estimates;
3. the causal variables can be forecast or controlled, especially with respect to their direction;
4. causal variables are expected to change substantially.

This implies that there is the need to be able to forecast large changes in competitors' actions.

SALES FORECASTS

Our assumption above was that one would prepare a market forecast and a market share forecast and then forecast sales by multiplying these components. Alternatively, sales can be forecasted directly. The direct approach seems most appropriate for short-range sales forecasting in situations where one is not concerned about assessing the effects of alternative strategies.

Methods based on judgment

One popular belief is that to improve forecasts, one should survey consumers about their desires, needs, plans, or expectations. The benefit of asking consumers depends on the situation. In particular, if sales (behavioral) data already exist, then it is obviously less important to seek information from consumers. However, intentions data can be used to improve the accuracy of sales extrapolation.

Focus groups are often used to forecast behavior. But this, procedure conflicts with certain principles. First, focus groups are seldom representative of the population. Second, the responses of each participant are influenced because they hear opinions stated by others. Third, the sample sizes are small. And fourth, questions for the participants are generally not well structured. As a result, focus groups should not be used in forecasting. In addition, no evidence exists to show that focus groups lead to improved forecasts.

Another approach is to solicit expert opinions. Expert opinion studies are widely used for forecasting of marketing problems. For example, forecasts are often obtained from the sales force. It is important to learn more about how to pose the questions, how to adjust for biases in these forecasts, and how to aggregate the responses.

Judgmental bootstrapping is likely to improve upon expert's sales forecasts. This is due largely to improved consistency.

Methods based on statistical sources

When a large number of sales forecasts are needed, the preferred method has been extrapolation. Here, relatively simple methods suffice. Sophistication beyond a modest level does not improve accuracy, but it does increase costs and reduce understanding. Marketers need a consistent set of forecasts for their products. If you are selling for computer parts, the forecasts for hard drives, disk drives and disks may be related to one another. On one forecast the aggregate (e.g. number of computer parts), then allocate percentages to the components ('top-down'). Alternatively; one could forecast for each part and then sum up the whole ('bottom-up'). Arguments can be made for each approach. The bottom-up seems best when one has reasonably good information on the parts. Empirical tests on the two approaches indicate that, in general, the bottom-up approach is more accurate (MacGregor, 2001)

Probably the most important principle when forecasting sales using data from intervals of less than a year (e.g. monthly data) is to adjust the data for seasonality. Dalrymple's (1987) survey results are consistent with this principle. Substantial improvements were also found in the large-scale study of time series by Makridakis et al. (1984). We believe that seasonal factors should be dampened, but no direct tests have yet been made.

Schnaars (1986) examined which extrapolation models are most accurate for annual sales forecasts for various consumer products. Two principles that helped were to dampen the trend and combine alternative forecasts. These principles improved accuracy in comparison with the rule “pick the model that provides the best fit to the historical sales.”

Some controversy exists as to whether mechanical extrapolations will do better than judgmental extrapolations. A study by Lawrence et al. (1985) concluded in favor of judgmental or 'eyeball' extrapolations, but Carbone and Gorr (1985) concluded the opposite. Of course, mechanical extrapolation methods are less expensive when many forecasts must be made, such as for inventory control.

New product sales

Sales forecasting for new products is a particularly important area, especially in view of the substantial investments and the likelihood of large forecasting errors. Forecasts are required at the different stages of product development to assist managers with the go/no-go decisions and then in planning the introduction of the new product.

Large errors are typical for new product forecasts. Tull (1967) estimated the mean absolute percentage error for new product sales to be about 65 percent. It is not surprising then, that pre-test market models have gained wide acceptance among business firms. Shocker and Hall (1986) provide an evaluation of some of these models Because of the lack of systematic and unbiased forecast validation studies, they conclude it is difficult to draw conclusions about which approach is best.

The choice of a forecasting model to estimate customer response depends on the stage of the product life cycle. As one moves from the concept phase to prototype, test market, introductory, growth, maturation, and declining stages, the relative value of the alternative forecasting methods changes. In general, the movement is from purely judgmental approaches to quantitative models. For example, intentions and expert opinions are vital to the concept and prototype stages. Later, expert judgment is useful as an input to quantitative models. Extrapolation methods may be useful in the early stages if it is possible to find analogous products (Claycamp and Liddy, 1969). In later stages, quantitative methods become more useful as less expensive sales and cost data become available.

When a new product is in the concept phase, a heavy reliance is usually placed on intentions surveys. Intentions to purchase new products are complicated because potential customers may not be sufficiently familiar with the proposed product and because the various features of the product affect one another (e.g., price, quality and distribution channel). This suggests the need to prepare a good description of the proposed product. The description may involve prototypes, visual aids, product clinics, or laboratory tests. However, brief descriptions are sometimes sufficient, as found in Armstrong and Overton's (1970) study of a new form of urban mass transportation.

In a typical intentions study, potential consumers are provided with a description of the product and the conditions of sale, and then are asked about their intentions to purchase. Rating scales of eleven points (0-10) are recommended. The scale should have verbal designations such as 0 = 'No chance, almost no chance (1 in 100)' to 10 = 'Certain, practically certain (99 in 100)'. It is best to state the question broadly about one's 'expectations' or 'probabilities' of purchase, rather than the narrower question of intentions. This distinction was proposed and tested by Juster (1966) and its importance has been shown to other empirical studies (Day et al., 1991).

Intentions surveys are useful when one of the following conditions holds: (1) the event is important; (2) the respondent (at least the intenders do); (3) the respondent can fulfill the plan; (4) events are unlikely to change the plan; (5) responses can be obtained; and (6) the respondent reports correctly. These conditions imply that intentions are more useful for short-term forecasts and business-to-business sales.

Intentions survey methodology has improved since the 1950s. Useful methods have been developed for selecting samples, compensating for nonresponse bias, and reducing response error. Dillman (2000) provides advice for designing intentions surveys. Improvements in this technology have been demonstrated by studies on voter intentions (Terry, 1979). Response error is probably the most important component of total error (Sudman and Birnbaum, 1982). Despite the improvements, the correspondence between intentions and sales is often not close, as shown in Morwitz (2001).

As an alternative to asking potential customers about their intentions to purchase, one can ask experts to predict how consumers will respond. For example, Wotruba and Thurlow (1976) discuss how opinions from members of the sales force can be used to forecast sales. One could also ask distributors or marketing executives to make forecasts. Experts may be able to make better forecasts if the problem is decomposed so that the parts are better known to them than the whole. Thus, if the task was to forecast the sales of high-definition television sets rather than making a direct forecast, one could break the problem into parts such as "How many households will there be in the U.S. in the forecast year?" "Of these households, what percentage will make more than $30,000 per year?" "Of these households, how many have not purchased a large screen TV in the past year?" and so on. The forecasts are obtained by multiplying the components. Decomposition is more accurate where there is much uncertainty about the direct or 'global' forecast (MacGregor, 2001). It turns out that much uncertainty is induced because people have difficulty in comprehending large numbers (operationalized as numbers over a million).

Unfortunately, experts are often subject in biases in new product forecasting (Tyebjee 1987). Sales people may try to forecast on the low side if the forecasts will be used to set quotas. Marketing executives may forecast high in their belief that this will gain approval for the project or motivate the sales force. If possible, avoid experts who would have obvious reasons to be biased. Another strategy is to use a heterogeneous group of experts in the hopes that their differing biases may cancel one another.

Producers often consider several alternative designs for a new product. In such cases, potential customers may be presented with a series of perhaps 20 or so alternative offerings. For example, various features of a personal computer, such as price, weight, battery life, screen clarity and memory might vary according to rules for experimental design (the basic ideas being that each feature should vary substantially and the variations among the features should not correlate with one another). The customer is forced to make trade-offs among various features. This is called 'conjoint analysis' because the consumers *consider* the product features *jointly.* This procedure is widely used by firms (Wittink and Bergestuen, 2001). An example of a successful application is the design of a new Marriott hotel chain (Wind et al., 1989). The use of conjoint analysis to forecast new product demand can be expensive because it requires large samples of potential buyers, the potential buyers may be difficult to locate, and the questionnaires are not easy to complete. Respondents must, of course, understand the concepts that they are being asked to evaluate. Although conjoint analysis rests on good theoretical foundations, little validation research exists in which its accuracy is compared with the accuracy of alternative techniques such as Delphi or judgmental forecasting procedures (Wittink and Bergestuen, 2001).

Expert judgments can be used in a manner analogous to conjoint analysis. That is, experts would make predictions about situations involving alternative product designs and alternative marketing plans. These predictions would then be related to the situations by regression analysis. (Following the philosophy for naming conjoint analysis, this could be called 'joint analysis.') It has advantages as compared to conjoint analysis in that few experts are needed (probably between five and ten). In addition, it can incorporate policy variables, such as advertising, that are difficult for consumers to assess.

Once a new product is on the market, it is possible to use extrapolation methods. Much attention has been given to the selection of the proper functional form to extrapolate early sales. The diffusion literature uses an S - shaped curve to predict new product sales. That is, growth builds up slowly at first, becomes rapid if word of mouth is good and if people see the product being used by others. Then it slows as it approaches a saturation level. A substantial literature exists on diffusion models. Despite this, the number of comparative validation studies is small and the benefits of choosing the best functional form are modest (Meade, 2001).

FORECASTING PROFITS AND OTHER OUTCOMES

Forecasts can be used to examine how each of the stakeholders will be affected. It may be useful to start the forecasting process with the analysis of stakeholders to ensure that the forecasts are relevant to decisions. For example, we might want to forecast whether a proposed plan will benefit consumers, how it might affect the local community, or how it might affect the long-term relationship with one of your complementors.

Forecasts of marketing costs can affect the marketing plan. Costs may be so high as to render a proposed plan unprofitable. Extrapolations are often used to forecast costs. Typically, unit costs decrease, but at a decreasing rate. Thus, a learning curve is often appropriate. This concept originated in educational psychology and was adopted by industrial engineering in the early 1900s. In simple terms, one estimates the percentage annual decrease in variable costs.

Sudden changes in costs can be forecasted by expert judgment, such as engineering estimates. Another approach is to use econometric models. Given the availability of relevant historical data, econometric models are especially useful for large changes in costs, such as those created by government actions. For example, costs of electricity vary substantially by geographic region due to the level of regulation. Econometric models might help to forecast prices given planned changes in regulation.

Assessing uncertainty

In addition to improving accuracy, forecasting is concerned with assessing uncertainty. This can help in managing risk.

Statisticians have given much attention to assessing uncertainty. They have relied heavily on tests of statistical significance. However, statistical significance is inappropriate for assessing uncertainty in forecasting. Furthermore, its use has been attacked as being misleading (e.g., see Cohen, 1994). It is difficult to find studies in marketing forecasting where statistical significance has made an important contribution.

Instead of statistical significance, the focus should be on prediction intervals. Chatfield (2000) summarizes research on prediction intervals. Unfortunately, prediction intervals are not widely used in practice. Rush and Page (1979) found a decreasing use of measures of uncertainty for metals forecasts from 22 percent of forecasts during the period 1910-1939 to only 8 percent during 1940-1964. Tull's (1967) survey noted that only 25 percent of 16 respondent companies said they provided confidence intervals with their forecasts. Dalrymple (1987) found that 48 percent did not use confidence intervals, and only 10 percent 'usually' used them.

The fit of a model to historical data is a poor way to estimate prediction intervals. It typically results in confidence intervals that are too narrow. It is best to simulate the actual forecasting procedure as closely as possible, and use the distribution of the resulting ex ante forecasts to assess uncertainty. For example, if you need to make forecasts for two years ahead, withhold enough data to be able to have a number of two-year ahead ex ante forecasts.

Methods based on judgment

Much work has been done on judgmental estimates of uncertainty. One of the key findings is that judges are typically overconfident (Arkes, 2001). This occurs even when subjects are warned in advance about this overconfidence phenomenon. Fortunately, there are procedures to improve the calibration of judges. Where possible, judges should be provided with frequent and well-summarized feedback of the outcomes of their predictions along with reasons why they were right or wrong. Arkes (2001) shows that when feedback is good, judges' estimates of confidence are well calibrated. For example, when weather forecasters say that there is about 60 percent chance of rain, it rains 60 percent of the time. This suggests that marketing forecasters should ensure that they receive feedback on the accuracy of their forecasts. The feedback should be frequent and it should summarize accuracy in a meaningful fashion.

In cases where good feedback is not possible, certain procedures can be used to improve estimates of confidence. One is to have experts write all the reasons why their forecasts might be wrong (Arkes, 2001). Another is the devil's advocate procedure, where someone is assigned for a short time to raise arguments about why the forecast might be wrong.

Still another way to assess uncertainty is to examine the agreement among judgmental forecasts. For example, Ashton (1985), in a study of forecasts of annual advertising sales for *Time* magazine, found that the agreement among the individual judgmental forecasts was a proxy for accuracy. Little evidence exists on this topic and it is not clear how to translate such information into prediction intervals. For example, in McNees'

(1992) examination of economic forecasts from 22 economists over 11 years, the actual values fell outside the range of their individual forecasts about 43 percent of the time.

Methods based on statistical data

Prediction intervals from quantitative forecasts tend to be too narrow even when based on ex ante n-ahead forecasts. Some empirical studies have shown that the percentage of actual values that fall outside the 95 percent prediction intervals is substantially greater than 5 percent, and sometimes greater than 50 percent (Makridakis et al., 1987). This occurs because the estimates ignore various sources of uncertainty. For example, discontinuities might occur over the forecast horizon. In addition, forecast errors in time series are often asymmetric, so this makes it difficult to estimate prediction intervals. This is likely to occur when the forecasting model uses an additive trend. The most sensible procedure is to transform the forecast and actual values to logs, then calculate the prediction intervals using logged differences. Interestingly, researchers and practitioners do not follow this advice (except where the original forecasting model has been formulated in logs). This procedure does not solve the situation where the trend extrapolation is contrary to the managers' expectations. Such errors are asymmetrical in logs. Evidence on the issue of asymmetrical errors is provided in Armstrong and Collopy (2001).

Loss functions can also be asymmetric. For example, the cost of a forecast that is too low by 50 units may differ from the cost if it is too high by 50 units. But this is a problem for the planner, not the forecaster.

SELECTING, EVALUATING AND USING FORECASTING METHODS

At a minimum, the use of new forecasting methods depends upon knowledge about them. While the acceptance of the new forecasting methods has been slow in the past, this is changing. The traditional methods of gaining acceptance, such as attending courses, reading textbooks, and using consultants, are being augmented by the Internet. The latest methods can be fully disclosed on web sites and they can be incorporated into expert systems and software packages. For example, the complete set of rules for role-based forecasting is kept available and up to date and can be accessed through the forecasting principles site (www-marketing.wharton.upenn.edu/forecast). Expert systems and software packages have the ability, then, to incorporate the cumulative wisdom about forecasting software. Users would apply the latest findings unless they chose to override the system to avoid them.

Rationally, one would expect that it helps to base the selection of the proper method upon previous research. So we describe how generalizations from prior research can aid in the selection of a forecasting method. Typically, however, one also needs to evaluate the performance of methods in the given situation.

Choosing a method based on prior research

The choice of the best forecasting method for any particular situation is not a simple task and sometimes more than one method may be appropriate. To provide guidance, Armstrong (2001c) used the findings from forecastingprinciples.com to develop a flow chart to aid the selection of the most appropriate forecasting method for any particular situation.

The first issue that the analyst needs to address is whether many data points are available. If not, judgmental procedures are called for.

For judgmental procedures, the next issues are whether the situation involves interaction among decision makers and whether large changes are involved. For large changes, is policy analysis involved, and if it is, what is the best source of evidence?

If one has a large quantity of data, does this consist of time series data? The next issue is whether there is knowledge about the expected empirical relationships. For example, meta-analyses have been done so that, in most situations, excellent prior knowledge exists about price elasticities (Tellis, 1988). If empirical knowledge of relationships is available, use econometric models. In addition, one should consider domain knowledge, such as a manager's knowledge about the situation.

For time series situations where one lacks causal knowledge, extrapolation is appropriate. If there is no prior knowledge about relationship, but domain knowledge exists (such as if a manager knows that sales will increase), use rule-based forecasting.

In situations where one does not have times series data and also has no prior knowledge about relationships, analogies are appropriate if domain knowledge is lacking. But given domain knowledge, expert systems should be used.

Figure 6.3 summarizes the above guidelines. While these represent the major considerations, the list is not comprehensive. Furthermore, the conditions may not always be clear. In such cases, one should use different approaches to the problem. The forecasts from these approaches can then be combined. To illustrate the use of the flow chart, we provide some examples.

Figure 6.3. Selection tree for forecasting methods

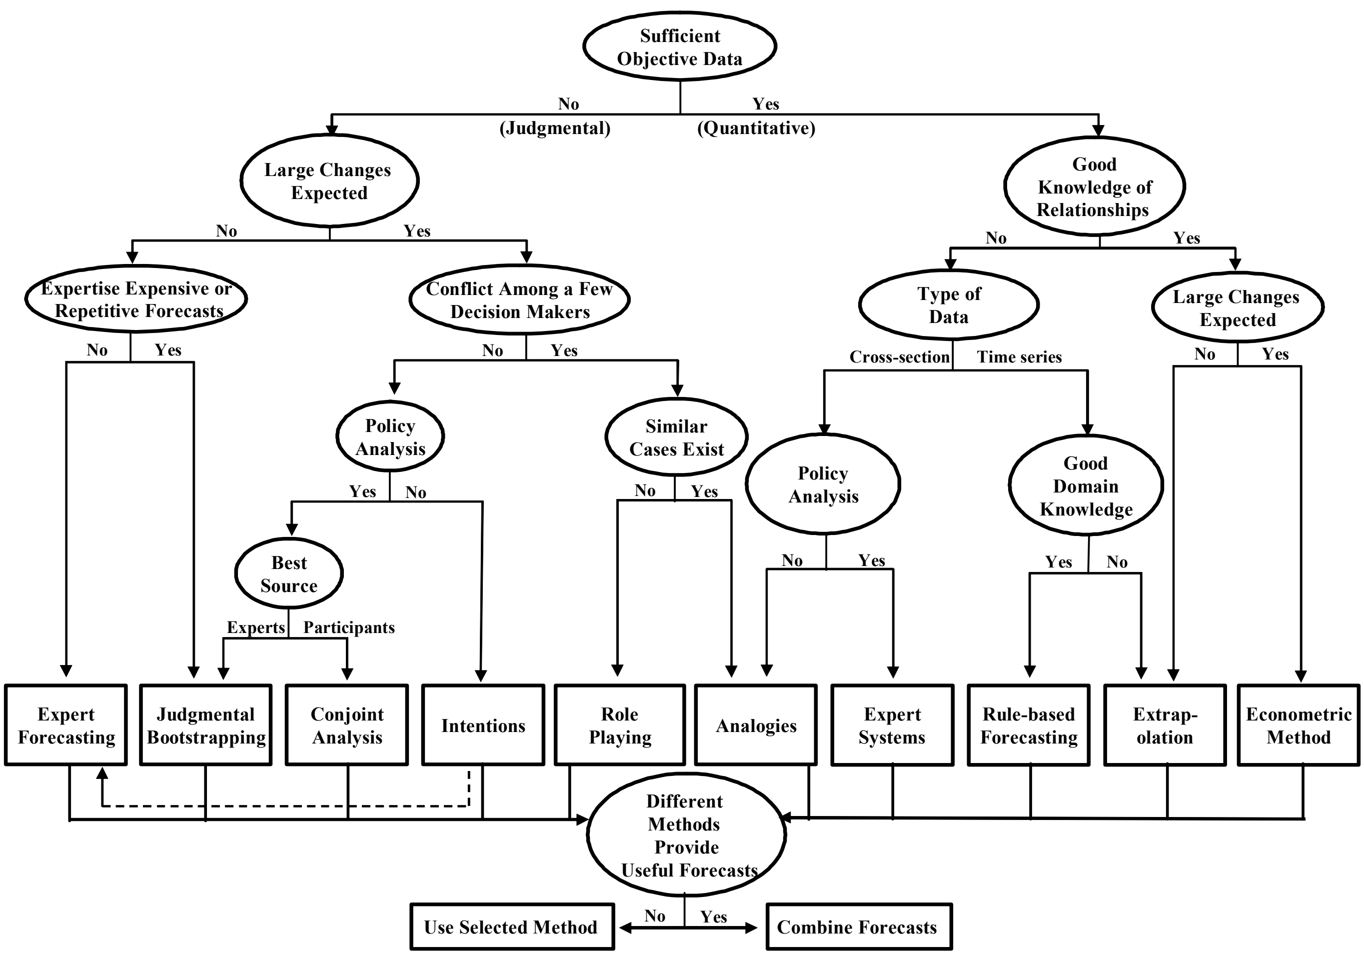

Market size

Consider a case where the New Zealand Wine Institute wishes to forecast the market size for white wine in New Zealand for the next five years. Ample time series data are available for wine consumption and the environmental factors influencing the demand. These include demographic and economic variables such as the size and distribution of the population, the changing age composition, and disposable income. Data are available about the price of white wine and the substitute products such as red wine and other alcoholic beverages. Also there is substantial empirical knowledge about the relationships from studies on the alcoholic beverage market conducted in New Zealand and other countries. These studies allow for generalizations to be made about price and income elasticities. Thus, an econometric modeling approach is recommended.

Decision makers' actions

Assume that the New Zealand Wine Institute is trying to assess likely effects of government protection and subsidies for local production in some of the newer export markets. In this case, little objective data are available, there is the possibility of interaction amongst the decision makers (i.e. New Zealand and foreign government trade officials), and the effects of the outcomes of these negotiations could be large. In this case, role playing or the use of experts is recommended.

Market share

What if a wine company wants to assess the effects that a new marketing strategy might have on its market? Here, scanner data are available, including market share data for the competing brands in the various product categories and data about the marketing activity for the competing brands. Also, there is substantial empirical knowledge about the relationships between the relevant variables from the response elasticities of the various marketing variables (e.g. price, point of sale promotion, advertising, and retail availability). Thus, econometric methods are recommended.

Sales

Assume that a wine company is considering planting a premium variety of grapes and the market is not familiar with this type of wine. Assume also that not many data are available, there is likely to be little interaction with decision makers, and the plantings require a large investment for the company. Hence a formal judgmental method would be recommended which could involve either consumers (conjoint analysis) or experts (Delphi or judgmental bootstrapping).

Profits and other outcomes

Now consider the case where a wine company is investigating the profitability of an investment in a new premium wine. This requires forecasts of costs and sales revenues. Considerable data about production and marketing costs are typically available, although they may not relate specifically to this new venture. Thus analogies are recommended. If knowledge exists about the factors that affect costs, an econometric analysis would be appropriate.

Statisticians have relied upon sophisticated procedures for analyzing how well models fit historical data. They then select the model with the best fit. Typically, this has been of little value for the selection of forecasting methods. Forecasters should not use measures of fit (such as R^2 or the standard error of the estimate of the model) because they have. little relationship to forecast accuracy. This conclusion is based on a series of .studies that go back at least to Ferber (1956). For a summary see Armstrong (2001c).

Ex ante forecasts from realistic simulations of the actual situation faced by the forecaster are likely to provide useful information about the expected accuracy of a forecasting model. By ex ante, we mean that the analyst uses only information that would be available at the time of an actual forecast.

Traditional error measures, such as the mean square error (MSE), do not provide a reliable basis for comparison of forecasting methods (Armstrong and Collopy, 1992). The median absolute percentage error (MdAPE) is more appropriate because it is invariant to scale and is not influenced by outliers. For comparisons using a small set of series, it is desirable to also control for degree of difficulty in forecasting. One measure that does this is the median relative absolute error (MdRAE), which compares the error for a given model against errors for the naive, no change forecast (Armstrong and Collopy, 1992).

In situations where uncertainty exists as to the best method, one should use two or more relevant methods, then combine the forecasts. Considerable research suggests that, lacking well structured domain knowledge, equally weighted averages are typically as accurate as other weighting schemes (Armstrong, 2001d). Equal weights combined forecasts produce consistent, though modest, improvements in accuracy and they reduce the likelihood of large errors. They are especially useful when the component methods differ substantially from one another. For example, Blattberg and Hoch (1990) obtained improved sales forecast by equally weighting managers' judgmental forecasts and forecasts from a quantitative model.

There is an exception to the equal weights principle. The selection and weighting of forecasting methods can be improved by using domain knowledge (about the item to be forecast), as shown in research on rule-based forecasting (Collopy and Armstrong, 1992). Domain knowledge can be structured, especially with respect to trend expectations. These expectations, along with a consideration of the features of the data (e.g., discontinuities), enable improvements in the weightings assigned to various extrapolations.

Rule based forecasting is just one of a number of procedures for integrating various methods. In particular, we have been stressing the need to integrate judgmental and statistical methods. Much research has been done on this topic in the past decade (Armstrong and Collopy, 1998). This integration is more effective because the judgments are collected in a systematic manner and then used as inputs to the quantitative models, rather than simply as adjustments the outputs. Unfortunately, the latter procedure is commonly used.

Using forecasts effectively

Forecasts that contradict management's expectations have much potential value. However, they are often ignored (Griffith and Wellman, 1979). One way to avoid this problem is to gain agreement on what forecasting procedures to use prior to presenting the forecasts. This may involve making adjustments to the forecasting method in order to develop forecasts that will be used.

Another way to gain acceptance of forecasts is to ask decision makers to decide in advance what decisions they will make given different possible forecasts. Are the decisions affected by the forecasts?

Prior agreements on process and on decisions can greatly enhance the value of forecasts, but they are difficult to achieve in many organizations. The use of scenarios offers an aid to this process. Scenarios involve writing detailed stories of how decision makers would handle alternative possibilities for the future. Decision makers project themselves into the situation and they write the stories about what they did in that situation. They should be written in the past tense. Detailed instructions for writing scenarios are summarized in Gregory and Duran (2001). Scenarios are effective in getting forecasters to accept the possibility that certain events might occur. They should not be used to make forecasts, however, because they distort subjective probability estimates.

CONCLUSIONS

Significant gains have been made in forecasting for marketing, especially since the 1960. Advances have occurred in the development of methods based on judgment, such as Delphi, role playing, intentions studies, opinions surveys, and bootstrapping. They have also occurred for methods based on statistical data, such as extrapolation, role based forecasting, and econometric methods. In the 1990s, gains have come from the integration of statistical and judgmental forecasts.

General principles

- Domain knowledge should be used.
- When making forecasts in highly uncertain situations, be conservative. For example, the trend should be dampened over the forecast horizon.
- When making forecasts in highly uncertain situations, use more than one method and combine the forecasts using equal weights.
- Complex methods have not proven to be more accurate than relatively simple methods. Given their added cost and the reduced understanding among users, highly complex procedures cannot presently be justified.
- When possible, forecasting methods should use data on actual behavior, rather than judgments or intentions, to predict behavior.
- Methods that integrate judgmental and statistical data and procedures (e.g., rule-based forecasting) can improve forecast accuracy in many situations.
- Overconfidence occurs with quantitative and judgmental methods

Methods based on judgment

- When using judgment, a heavy reliance should be placed on structured procedures such as Delphi, role playing, and conjoint analysis.
- Role playing is useful to predict the decisions or variations in practice involved in conflict situations, such as in negotiation.
- In addition to seeking good feedback, forecasters should explicitly list all the things that might be wrong about their forecast. This will produce better calibrated prediction intervals.

Methods based on statistical data

- With the increased availability of data, econometric models play an increasingly important role in forecasting market size, market share, and sales.
- Methods should be developed primarily on the basis of theory, not data.

Finally, efforts should be made to ensure forecasts are free of political considerations in a firm. To help with this, emphasis should be on gaining agreement about the forecasting methods, rather than the forecasts. Also, for important forecasts, decisions on their use should be made before the forecasts are provided. Scenarios are helpful in guiding this process.

REFERENCES

Agrawal, D. and Schorling, C. (1996). “Market share forecasting: An empirical comparison of artificial networks and multinomial logit model,” *Journal of Retailing* 72 (4), 383-407.

Arkes, M. R. (2001), “Overconfidence in judgmental forecasting,” in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 495-515.

Armstrong, J. S. (1985), *Long-term Forecasting: From Crystal Ball to Computer* (second ed.). New York: John Wiley.

Armstrong, J. S. (2001a), “Judgmental bootstrapping: Inferring experts' rules for forecasting,” in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 171-192.

Armstrong, J. S. (2001b), “Role playing: A method to forecast decisions,” in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 15-30.

Armstrong, J. S. (2001c), “Evaluating forecasting methods,” in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 365-382.

Armstrong, J. S. (2001d), “Combining forecasts,” in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 417-439.

Armstrong, J. S. and Collopy, F. (1992), “Error measures for generalizing about forecasting methods empirical comparisons,” *International Journal of Forecasting*, 8, 69-80.

Armstrong, J. S. and Collopy, F. (1993), “Causal forces: Structuring knowledge for time series extrapolation,” *Journal of Forecasting*, 12, 103-115.

Armstrong, J. S. and Collopy, F. (1998), “Integration of statistical methods and judgment for time series forecasting: Principles from empirical research, “in G. Wright and P. Goodwin (Eds.), *Forecasting with Judgmen*t. Chicliester:John Wiley.

Armstrong, J. S. and Collopy, F. (2001), “Identification of asymmetric prediction intervals through causal forces,” *Journal of Forecasting*, 20, 273-283.

Armstrong, J. S. and Overton, T. (1970), “Brief vs. comprehensive descriptions in measuring intentions to purchase,” *Journal of Marketing Researh*, 8, 114-117.

Ashton, A. H. (1985), “Does consensus imply accuracy in accounting studies of decision making?” *Accounting Review,* 60, 173-185.

Ashton, A. H. and Ashton, R H. (1985), “Aggregating subjective forecasts: Some empirical results,” *Management Science*, 31, 1499-1508.

Bell, D. E., Keeney, R. L. and Little, J. D._C. (1975), “A market share theorem,” *Journal of Marketing Research*, 12, 136-141.

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

Brodie, R. J., Danaher, P,, Kumar, V. and Leeflang, R. (2000), “Econometric models for forecasting market share,” in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 597-611.

Carbone, R. and Gorr, W. (1985), “Accuracy of judgmental forecasting of time series,” *Decision Sciences*, 16, 153-160.

Chatfield, C. (1993), “Neural networks: forecasting breakthrough or passing fad?” *International Journal of Forecasting*, 9, 1-3.

Chatfield, C. (2001), “Prediction intervals for time series,” in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 475-494.

Claycamp, H. J. and Liddy, L. E. (1969), “Prediction of new product performance: An analytical approach,” *Journal of Marketing Research*, 6, 414-420.

Cohen, J. (1994), “The earth is round ($p < .05$),” *American Psychologist*, 49, 997-1003.

Collopy, F. and Armstrong, J. S. (1992), "Rule-based forecasting: Development and validation of an expert systems approach to combining time series extrapolations," *Management Science*, 38, 1394-1414.

Dalrymple, D. J. (1975), "Sales forecasting: Methods and accuracy," *Business Horizons*, 18, 69-73.

Dalrymple, D. J. (1987), "Sales forecasting practices: Results from a U.S. survey," *International Journal of Forecasting,* 3, 379-391.

Day, D. B., Gan, B., Gendall, P. and Esslemont, D. (1991), "Predicting purchase behavior," *Marketing Bulletin*, 2, 18-30 (Full text available at http://marketing-bulletin. massey.ac.nz/).

Dillman, D. A. (2000), *Mail and Internet Surveys.* New York: John Wiley.

Duncan, G., Gorr, W. and Szczypula, J. (2001), "Forecasting analogous time series," in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 195-213.

Ferber, R. (1956), "Are correlations any guide to predictive value?" *Applied Statistics*, 5 (June), 113-122.

Fildes, R. (1985), "Quantitative forecasting - the state of the art: Econometric models," Journal of the Operational Research Society, 36, 549-580.

Gregory, L. and Duran, A. (2001), "Scenarios and acceptance of forecasts," in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 519-541.

Griffith, J. R. and Wellman, B. T. (1979), "Forecasting bed needs and recommending facilities plans for community hospitals: A review of past performance," *Medical Care,* 17, 293-303.

Hamilton, W, East, R. and Kilafatis, S. (1997), "The measurement and utility of brand price elasticities," *Journal of Marketing Management,* 13, 285-298.

Jobber, D., Hooley G. and Sanderson, S. (1985), "Marketing in a hostile environment: The British textile industry," *Industrial Marketing Managemen*t, 14, 35-41.

Juster, T. (1966), "Consumer buying intentions and purchase probability: An experiment in survey design," *Journal of the American Statistical Association*, 61, 658-696.

Lawrence, M. J., Edmondson, R. H. and O'Connor, M. J. (1985), "An examination of the accuracy of judgmental extrapolation of time series," *International Journal of Forecasting*, 1, 25-35.

MacGregor, D. (2001), "Decomposition for judgmental forecasting and estimation," in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 107-123.

Makridakis, S., et al. (1984), *The Forecasting Accuracy of Major Times Series Methods.* New York: John Wiley.

Makridakis, S., Hibon, M., Lusk, F. and Belhadjali, M. (1987), "Confidence intervals: An empirical investigation of tine series in the M-competition," *International Journal of Forecasting,* 3, 489-508.

Makridakis, S., Wheelwright, S. C., and Hyndman, R. J. (1998), *Forecasting Methods for Management,* Third edition. New York: John Wiley.

McCloskey, D. N. and Ziliak, S. T. (1996), "The standard error of regressions," *Journal of Economic Literature*, 34, 97-114.

McNees, S. K. (1992), "The uses and abuses of 'consensus' forecasts," *Journal of Forecasting*, 11, 703-710.

Meade, N. and Islam, T. (2001), "Forecasting the diffusion of innovations: Implications for time series extrapolation," in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 4577-595.

Morwitz, V. (2001), "Methods for forecasting from intentions data," in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 33-56.

Perry, C. (1979), "Certain problems in election survey methodology," *Public Opinion Quarterly*, 43, 312-325.

Rosenstone, S. J. (1983), *Forecasting Presidential Elections*. New Haven, CN: Yale University Press.

Rowe, G. and Wright, G. (2001), "Expert opinions in forecasting role of the Delphi technique," in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell, MA: Kluwer Academic Publishers, pp. 125-144.

Rush, H. and Page, W. (1979), "Long-term metals forecasting: The track record," *Futures*, 11, 321-337.

Schnaars, S. P. (1986), "An evaluation of rules for selecting an extrapolation model on yearly sales forecasts," *Interfaces* 16 (Nov-Dec.), 100-107.

Shocker, A. D. and Hall, W. G. (1986), "Pretest market models: A critical evaluation," *Journal of Product Innovation Management*, 3, 86-107.

Sudman, S. and Bradburn, N. R. (1982), *Asking Questions: A Practical Guide to Questionnaire Design.* San Francisco: Jossey-Bass.

Tellis, G. J. (1988),"The price elasticity of selected demand," *Journal of Marketing Research*, 25, 331-341.

Tull, D. S. (1967), "The relationship of actual and predicted sales and profits in new product introductions," *Journal of Business*, 411, 233-250.

Tyebjee, T. T. (1987), "Behavioral biases in new product forecasting," *International Journal of Forecasting*, 3, 393-404.

West, P. M., Brockett, R. L., and Golden, L.L. (1997), "A comparative analysis of neural networks and statistical methods for predicting consumer choice," *Marketing Science*, 16, 370-391.

Wind, J., Green, P. E, Shifflet, D., and Scarbrough, M. (1989), "Courtyard by Marriott: Designing a hotel facility with consumer-based marketing," *Interfaces*, 19, 1, 25-47.

Wittink, D. R, and Bergestuen, T. (2001), "Forecasting with conjoint analysis," in J. S. Armstrong (Ed.) *Principles of Forecasting: Handbook for Researchers and Practitioners*. Norwell MA: Kluwer Academic Publishers, pp. 147-167.

Wotruba, T. R. and Thurlow, M.L. (1976), "Sales force participation in quota setting and sales forecasting," *Journal of Consumer Research*, 8, 162-171.