

Four Steps to Forecast Total Market Demand

William Barnett (HBR July 1988 Issue)

Recent history is filled with stories of companies and sometimes even entire industries that have made grave strategic errors because of inaccurate industry wide demand forecasts. For example: In 1974, U.S. electric utilities made plans to double generating capacity by the mid-1980s based on forecasts of a 7% annual growth in demand. Such forecasts are crucial since companies must begin building new generating plants five to ten years before they are to come on line. But during the 1975–1985 period, load actually grew at only a 2% rate. Despite the postponement or cancellation of many projects, the excess generating capacity has hurt the industry financial situation and led to higher customer rates.

The petroleum industry invested \$500 billion worldwide in 1980 and 1981 because it expected oil prices to rise 50% by 1985. The estimate was based on forecasts that the market would grow from 52 million barrels of oil a day in 1979 to 60 million barrels in 1985. Instead, demand had fallen to 46 million barrels by 1985. Prices collapsed, creating huge losses in drilling, production, refining, and shipping investments.

In 1983 and 1984, 67 new types of business personal computers were introduced to the U.S. market, and most companies were expecting explosive growth. One industry forecasting service projected an installed base of 27 million units by 1988; another predicted 28 million units by 1987. In fact, only 15 million units had been shipped by 1986. By then, many manufacturers had abandoned the PC market or gone out of business altogether.

The inaccurate suppositions did not stem from a lack of forecasting techniques; regression analysis, historical trend smoothing, and others were available to all the players. Instead, they shared a mistaken fundamental assumption: that relationships driving demand in the past would continue unaltered. The companies didn't foresee changes in end-user behavior or understand their market's saturation point. None realized that history can be an unreliable guide as domestic economies become more international, new technologies emerge, and industries evolve.

As a result of changes like these, many managers have come to distrust traditional techniques. Some even throw up their hands and assume that business planning must proceed without good demand forecasts. I disagree. It is possible to develop valuable insights into future market conditions and demand levels based on a deep understanding of the forces behind total-market demand. These insights can sometimes make the difference between a winning strategy and one that flounders.

A forecast of total-market demand won't guarantee a successful strategy. But without it, decisions on investment, marketing support, and other resource allocations will be based on hidden, unconscious assumptions about industrywide requirements, and they'll often be wrong. By gauging total-market demand explicitly, you have a better chance of controlling your company's destiny. Merely going through the process has merit for a management team. Instead of just coming out with pat answers, numbers, and targets, the team is forced to rethink the competitive environment.

Total-market forecasting is only the first stage in creating a strategy. When you've finished your forecast, you're not done with the planning process by any means. There are four steps in any total-market forecast:

1. Define the market.
2. Divide total industry demand into its main components.
3. Forecast the drivers of demand in each segment and project how they are likely to change.
4. Conduct sensitivity analyses to understand the most critical assumptions and to gauge risks to the baseline forecast.

Defining the Market

At the outset, it's best to be overly inclusive in defining the total market. Define it broadly enough to include all potential end users so that you can both identify the appropriate drivers of demand and reduce the risk of surprise product substitutions.

The factors that drive forecasts of total-market size differ markedly from those that determine a particular product's market share or product-category share. For example, total-market demand for office telecommunications products nationally depends in part on the number of people in offices and their needs and habits, while total demand for PBX systems depends on how they compare on price and benefits with substitute products like the local telephone company's central office switching service. Beyond this, demand for a particular PBX is a function of price and benefit comparisons with other PBXs.

In defining the market, an understanding of product substitution is critical. Customers might behave differently if the price or performance of potential substitute products changes. One company studying total demand for industrial paper tubes had to consider closely related uses of metal and plastic tubes to prevent customer switching among tubes from biasing the results.

Understand, too, that a completely new product could displace one that hitherto had comprised the entire market—like the electronic calculator, which eliminated the slide rule. For a while after AT&T’s divestiture, the Bell telephone companies continued to forecast volume of long-distance calls by using historical trend lines of their revenues—as if they were still part of a monopoly. Naturally, these forecasts grew more inaccurate with time as end users were presented with new choices. The companies are now broadening their market definitions to take account of heightened competition from other long-distance carriers.

There are several ways you can make sure you include all important substitute products (both current and potential). From interviews with industrial customers you can learn about substitutes they are studying or about product usage patterns that imply future switching opportunities. Moreover, market research can lead to insights about consumer products. Speaking with experts in the relevant technologies or reviewing technological literature can help you identify potential developments that could threaten your industry.

Finally, careful quantification of the economic value of alternative products to different customers can yield deep insights into potential switching behavior—for example, how oil price movements would affect plastics prices, which in turn would affect plastic products’ ability to substitute for metal or paper.

Analyses like these can lead to the construction of industry demand curves—graphs representing the relationship between price and volume. With an appropriate definition, the total-industry demand curves will often be steeper than demand curves for individual products in the industry. Consumers, for example, are far more likely to switch from Maxwell House to Folgers coffee if Maxwell House’s prices increase than they are to stop buying coffee if all coffee prices rise.

In some cases, managers can make quick judgments about market definition. In other cases, they’ll have to give their market considerable thought and analysis. A total-market forecast may not be critical to business strategy if market definition is very difficult or the products under study have small market shares. Instead, your principal challenge may be to understand product substitution and competitiveness. One company analyzed the potential market for new consumer food cans, and it concluded that growth trends in food product markets were not critical to the strategy question. What was critical was knowing the value positions of the new packages relative to metal cans, glass jars, and composite cans. So the company spent time on that subject.

Dividing Demand into Component Parts

The second step in forecasting is to divide total demand into its main components for separate analysis. There are two criteria to keep in mind when choosing segments: make each category small and homogeneous enough so that the drivers of demand will apply consistently across its various elements; make each large enough so that the analysis will be worth the effort. Of course, this is a matter of judgment.

You may find it useful in making this judgment to imagine alternative segmentations (based on end-use customer groups, for example, or type of purchase). Then hypothesize their key drivers of demand (discussed later) and decide how much detail is required to capture the true situation. As the assessment continues, managers can return to this stage and reexamine whether the initial decisions still stand up.

Managers may wish to use a “tree” diagram like the accompanying one constructed by a management team in 1985 to study demand for paper. In this disguised example, industry data permitted the division of demand into 12 end-use categories. Some categories, like business forms and reprographic paper, were big contributors to total consumption; others, such as labels, were not. One (other converting) was fairly large but too diverse for deep analysis. The team focused on the four segments that accounted for 80% of 1985 demand. It then developed secondary branches of the tree to further dissect these categories and to determine their drivers of demand. It analyzed the remaining segments less completely (that is, via a regression against broad macroeconomic trends).

Components of Uncoated White Paper Making Up Total Demand (thousands of tons)

	End-Use Category	Percent of Total 1995 Demand
Total Demand	Business Forms	25%
	Commercial Printing	25
	Reprographics	20
	Envelopes	10
	Other Converting	5
	Stationery and Tablet	5
	Books	5
	Directories	1 or less
	Catalogs	
	Magazines	
	Inserts	
	Labels	

Reviewed in Depth

Other companies have used similar methods to segment total demand. One company divided demand for maritime satellite terminals by type of ship (e.g., seismic ships, bulk/cargo/container ships). Another divided demand for long-distance telephone service into business and residential customers and then subdivided it by usage level. And a third segmented consumer appliances into three purchase types—appliances used in new home construction, replacement appliance sales in existing homes, and appliance penetration in existing homes.

In thinking about market divisions, managers need to decide whether to use existing data on segment sizes or to commission research to get an independent estimate. Reliable public information on historical demand levels by segment is available for many big U.S. industries (like steel, automobiles, and natural gas) from industry associations, the federal government, off-the-shelf studies by industry experts, or ongoing market data services. For some foreign markets and less well-researched industries in the United States, like the labels industry, you may have to get independent estimates. Even with good data sources, however, the readily available information may not be divided into the best categories to support an insightful analysis. In these cases, managers must decide whether to develop their forecasts based on the available historical data or to undertake their own market research programs, which can be time-consuming and expensive.


Note that while such segmentation is sufficient for forecasting total demand, it may not create categories useful for developing a marketing strategy. A single product may be driven by entirely different factors. One study of industrial components found that consumer industry categories provided a good basis for projecting total-market demand but gave only limited help in formulating a strategy based on customer preferences: distinguishing those who buy on price from those who buy on service, product quality, or other benefits. Such buying-factor categories generally do not correlate with the customer industry categories used for forecasting. A strong sales force, however, can identify customer preferences and develop appropriate account tactics for each one.

Forecasting the Drivers of Demand

The third step is to understand and forecast the drivers of demand in each category. Here you can make good use of regressions and other statistical techniques to find some causes for changes in historical demand. But this is only a start. The tougher challenge is to look beyond the data on which regressions can easily be based to other factors where data are much harder to find. Then you need to develop a point of view on how those other factors may themselves change in the future.

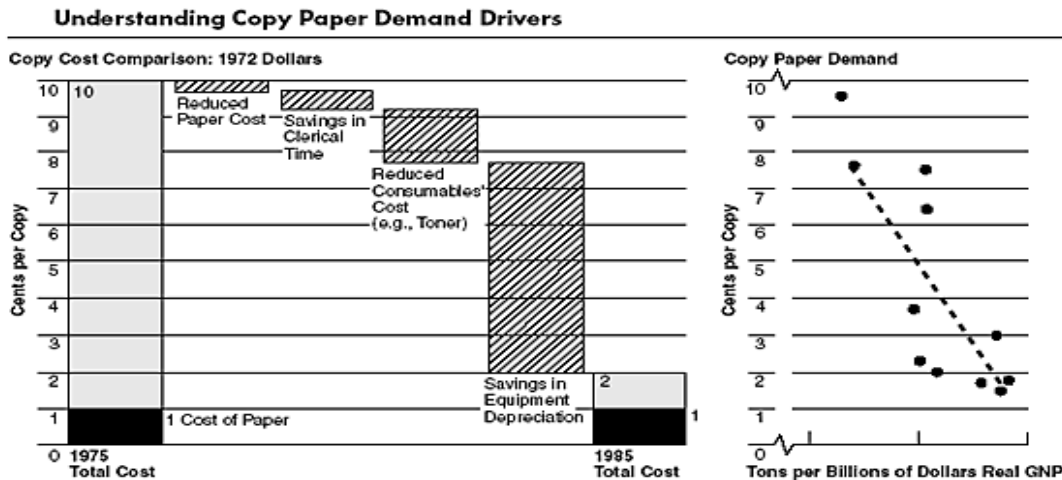
An end-use analysis from the commodity paper example, reprographic paper, is shown in the accompanying chart. The management team, using available data, divided reprographic paper into two categories: plain-paper copier paper and nonimpact page printer paper. Without this important differentiation, the drivers of demand would have been masked, making it hard to forecast effectively.

Drivers of Demand for Reprographic Paper



In most cases, managers can safely assume that demand is affected both by macroeconomic variables and by industry-specific developments. In looking at plain-paper copier paper, the team used simple and multiple regression analyses to test relationships with macroeconomic factors like white-collar workers, population, and economic performance. Most of the factors had a significant effect on demand. Intuitively, it also made sense to the team that the level of business activity would relate to paper consumption levels. (Economists sometimes refer to growth in demand due to factors like these as an “outward shift” in the demand curve—toward a greater quantity demanded at a given price.)

Demand growth for copy paper, however, had exceeded the real rate of economic growth and the challenge was to find what other factors had been causing this. The team hypothesized that declining copy costs had caused this increased usage. The relationship was proved by estimating the substantial cost reductions that had occurred, combining those with numbers of tons produced over time, and then fashioning an indicative demand curve for copy paper. (See the chart “Understanding Copy Paper Demand Drivers.”) The clear relationship between cost and volume meant that cost reductions had been an important cause of past demand growth. (Economists sometimes describe this as a downward-shifting supply curve leading to movement down the demand curve.)



Further major declines in cost per copy seemed unlikely because paper costs were expected to remain flat, and the data indicated little increase in price elasticity, even if cost per copy fell further. So the team concluded that usage growth (per level of economic performance) was likely to continue the flattening trend begun in 1983: growth in copy paper consumption would be largely a function of economic growth, not cost declines as in the past. The team then reviewed several econometric services forecasts to develop a base case economic forecast.

Similar studies have been performed in other industries. A simple one was the industrial components analysis mentioned before, a case where the total forecast was used as background but was not critical to the company's strategy decision. Here the team divided demand into its consuming industries and then asked experts in each industry for production forecasts. Total demand for components was projected on the assumption that it would move parallel to a weight-averaged forecast of these customer industries. Actual demand three years later was 2% above the team's prediction, probably because the industry experts underestimated the impact of the economic recovery of 1984 and 1985.

In another example, a team forecasting demand for maritime satellite terminals extrapolated past penetration curves for each of five categories of ships. These curves were then adjusted for major changes in the shipping industry (e.g., adding the depressing effect of the growing oil glut, taking out of these historical trends the unnatural demand growth that had been caused by the Falklands war). The actual figure three years later was within 1% of the forecast.

Knowing the drivers of demand is crucial to the success of any total-market demand forecast. In 1974, as I mentioned earlier, most electric utilities used an incomplete total-demand forecast to predict robust demand growth. In the early 1980s, one company's management team, however, decided to study potential changes in end-user demand as well. The team divided electricity demand into the three traditional categories: residential, commercial, and industrial. It then profiled differences in residential demand because of more efficiency in home appliances and changes in home size and the ratio of multi-unit to single-family dwellings. Industrial demand was analyzed by evaluating the future of several key consuming industries, paying special attention to changes in their total production and electricity use. This end-use approach sharply reduced the utility's initial forecasts and led to cancellation of two \$700 million generating plants then in the planning stage.

In 1983, forecasters in the U.S. personal computer industry were saying that demand would continue to rise at a rapid rate because there were 50 million white-collar workers and only 8 million installed PCs. One company, however, did a more detailed demand forecast that showed that growth would soon flatten out. It found that more than two-thirds of white-collar workers either did not require PCs in their jobs—actors and elevator operators, for instance—or were supported mostly by inexpensive terminals linked to large computers, as in the case of many clerical workers. The potential market was not big enough to support the growth rate. Indeed, the market began to flatten the next year.

Forecasting total demand became crucial for another company that was thinking about acquiring a maker of video games. Many thought that low overall market penetration (10% of U.S. households) signified a lot of room for growth before the market became saturated, when about 50% of the households would have games. Using available data, however, the management team created categories based on family income and children's ages. The analysis made clear that the main target market, upper-income families with children, was already well penetrated. Families with incomes exceeding \$50,000 and children between the ages of 6 and 15 already were 75% penetrated. This finding convinced management that demand would fall and that the proposed acquisition did not make sense. The dramatic decline in video game sales shortly thereafter confirmed the wisdom of this judgment.

Conducting Sensitivity Analyses

Managers who rely on single-point demand forecasts run dangerous risks. Some of the macroeconomic variables behind the forecasts could be wrong. Despite the best analysis, moreover, the assumptions behind the other demand drivers could also be wrong, especially if discontinuities loom on the horizon. Imaginative marketers who ask questions like "What things could cause this forecast to change dramatically?" produce the best estimates. They are more likely to identify potential risks and discontinuities—developments in competing technologies, in customer industry competitiveness, in supplier cost structures—than those who do not. So once a baseline forecast is complete, the challenge is to determine how far it could be off target.

At one level, such a sensitivity analysis can be done by simply varying assumptions and quantifying their impact on demand. But a more targeted approach usually provides better insight.

Begin such an analysis by thinking through and quantifying the areas of greatest strategic risk. One company's strategy decision may be affected only if demand is well below the baseline forecast; in another case, big risks may result from small forecasting errors.

Next, gauge the likelihood of such a development. In the white paper example, the baseline forecast called for continued market growth, though below historical levels. In any particular year, demand could fluctuate with the economy, but the

critical question was whether demand would at some point begin a long decline. If so, the companion supply-curve analysis indicated that prices would probably fall dramatically.

The team created two scenarios of a gradual decline, one based largely on changes in the economy and the other on changes in assumed end-use trends. These scenarios showed what would make demand fall (e.g., different rates of decline in copier prices) and thereby provided a basis for evaluating the likelihood of a downturn.

Determining an Appropriate Effort

The forecasting framework outlined above can work for both comprehensive and simple assessments, but there are different ways to carry out these analyses. A big challenge in demand forecasting (just as with other types of market analysis) is to gauge the appropriate effort for the project's purpose. It's useful to ask: "How much do I need to know to make the decision at hand?"

Managers can invest a lot of time in such analyses—the paper example took about 8 man-weeks and the large-scale electricity forecast about 14 man-weeks. Some companies have forecasting departments who work year-round on these subjects. The more thorough, though time-consuming, approach generates greater confidence, and the effort will be appropriate where the demand projection can significantly influence corporate strategy (whether to make a several hundred million dollar capital investment, for example), or where there is great uncertainty about total demand.

Often, however, the issues are not complicated, time is limited, or the total demand forecast is not important enough to merit that commitment (for example, the company is looking to add a couple of points to its small market share). In such cases, managers should proceed quickly and inexpensively. They can, for example, rely on experts' judgment or unsophisticated regressions to forecast drivers of demand. Even the limited approaches can yield insights. Furthermore, beginning the demand analysis process can help managers determine whether important demand issues exist that should be analyzed in greater depth.

Total-demand forecasting can be important to strategy decisions. Developing independent forecasts through the four-step framework I've outlined will not only lead to better recommendations but also help build conviction and consensus for action by creating understanding of the drivers of demand and the risks in forecasts.

Even when the work is sound, though, uncertainties will remain: discontinuities will still be difficult to predict, especially if they are rooted in momentous political, macroeconomic, or technological changes. But managers who push their thinking through the steps in this framework will have a better chance of finding these discontinuities than those who do not. And those who base their business strategies on a solid knowledge of demand will stand a much greater chance of making wise investments and competing effectively.

How to Choose the Right Forecasting Technique

John C. Chambers Satinder K. Mullick Donald D. Smith (HBR July 1971 Issue)

In virtually every decision they make, executives today consider some kind of forecast. Sound predictions of demands and trends are no longer luxury items, but a necessity, if managers are to cope with seasonality, sudden changes in demand levels, price-cutting maneuvers of the competition, strikes, and large swings of the economy. Forecasting can help them deal with these troubles; but it can help them more, the more they know about the general principles of forecasting, what it can and cannot do for them currently, and which techniques are suited to their needs of the moment. Here the authors try to explain the potential of forecasting to managers, focusing special attention on sales forecasting for products of Corning Glass Works as these have matured through the product life cycle. Also included is a rundown of forecasting techniques.

To handle the increasing variety and complexity of managerial forecasting problems, many forecasting techniques have been developed in recent years. Each has its special use, and care must be taken to select the correct technique for a particular application. The manager as well as the forecaster has a role to play in technique selection; and the better they understand the range of forecasting possibilities, the more likely it is that a company's forecasting efforts will bear fruit.

The selection of a method depends on many factors—the context of the forecast, the relevance and availability of historical data, the degree of accuracy desirable, the time period to be forecast, the cost/ benefit (or value) of the forecast to the company, and the time available for making the analysis.

These factors must be weighed constantly, and on a variety of levels. In general, for example, the forecaster should choose a technique that makes the best use of available data. If the forecaster can readily apply one technique of acceptable accuracy, he or she should not try to “gold plate” by using a more advanced technique that offers potentially greater accuracy but that requires nonexistent information or information that is costly to obtain. This kind of trade-off is relatively easy to make, but others, as we shall see, require considerably more thought.

Furthermore, where a company wishes to forecast with reference to a particular product, it must consider the stage of the product’s life cycle for which it is making the forecast. The availability of data and the possibility of establishing relationships between the factors depend directly on the maturity of a product, and hence the life-cycle stage is a prime determinant of the forecasting method to be used.

Our purpose here is to present an overview of this field by discussing the way a company ought to approach a forecasting problem, describing the methods available, and explaining how to match method to problem. We shall illustrate the use of the various techniques from our experience with them at Corning, and then close with our own forecast for the future of forecasting.

Although we believe forecasting is still an art, we think that some of the principles which we have learned through experience may be helpful to others.

Manager, Forecaster & Choice of Methods

A manager generally assumes that when asking a forecaster to prepare a specific projection, the request itself provides sufficient information for the forecaster to go to work and do the job. This is almost never true.

Successful forecasting begins with a collaboration between the manager and the forecaster, in which they work out answers to the following questions.

1. What is the purpose of the forecast—how is it to be used? This determines the accuracy and power required of the techniques, and hence governs selection. Deciding whether to enter a business may require only a rather gross estimate of the size of the market, whereas a forecast made for budgeting purposes should be quite accurate. The appropriate techniques differ accordingly.

Again, if the forecast is to set a “standard” against which to evaluate performance, the forecasting method should not take into account special actions, such as promotions and other marketing devices, since these are meant to change historical patterns and relationships and hence form part of the “performance” to be evaluated.

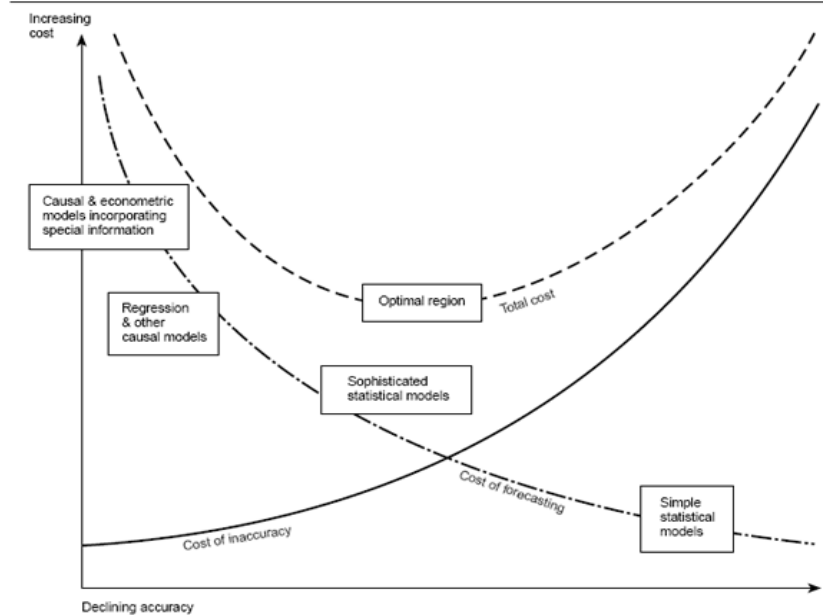
Forecasts that simply sketch what the future will be like if a company makes no significant changes in tactics and strategy are usually not good enough for planning purposes. On the other hand, if management wants a forecast of the effect that a certain marketing strategy under debate will have on sales growth, then the technique must be sophisticated enough to take explicit account of the special actions and events the strategy entails.

Techniques vary in their costs, as well as in scope and accuracy. The manager must fix the level of inaccuracy he or she can tolerate—in other words, decide how his or her decision will vary, depending on the range of accuracy of the forecast. This allows the forecaster to trade off cost against the value of accuracy in choosing a technique.

For example, in production and inventory control, increased accuracy is likely to lead to lower safety stocks. Here the manager and forecaster must weigh the cost of a more sophisticated and more expensive technique against potential savings in inventory costs.

Exhibit I shows how cost and accuracy increase with sophistication and charts this against the corresponding cost of forecasting errors, given some general assumptions. The most sophisticated technique that can be economically justified is one that falls in the region where the sum of the two costs is minimal.

Exhibit I Cost of Forecasting Versus Cost of Inaccuracy For a Medium-Range Forecast, Given Data Availability



Once the manager has defined the purpose of the forecast, the forecaster can advise the manager on how often it could usefully be produced. From a strategic point of view, they should discuss whether the decision to be made on the basis of the forecast can be changed later, if they find the forecast was inaccurate. If it can be changed, they should then discuss the usefulness of installing a system to track the accuracy of the forecast and the kind of tracking system that is appropriate.

2. What are the dynamics and components of the system for which the forecast will be made? This clarifies the relationships of interacting variables. Generally, the manager and the forecaster must review a flow chart that shows the relative positions of the different elements of the distribution system, sales system, production system, or whatever is being studied.

Exhibit II displays these elements for the system through which CGW's major component for color TV sets—the bulb—flows to the consumer. Note the points where inventories are required or maintained in this manufacturing and distribution system—these are the pipeline elements, which exert important effects throughout the flow system and hence are of critical interest to the forecaster.

All the elements in dark gray directly affect forecasting procedure to some extent, and the color key suggests the nature of CGW's data at each point, again a prime determinant of technique selection since different techniques require different kinds of inputs. Where data are unavailable or costly to obtain, the range of forecasting choices is limited.

The flow chart should also show which parts of the system are under the control of the company doing the forecasting. In Exhibit II, this is merely the volume of glass panels and funnels supplied by Corning to the tube manufacturers.


In the part of the system where the company has total control, management tends to be tuned in to the various cause-and-effect relationships, and hence can frequently use forecasting techniques that take causal factors explicitly into account.

The flow chart has special value for the forecaster where causal prediction methods are called for because it enables him or her to conjecture about the possible variations in sales levels caused by inventories and the like, and to determine which factors must be considered by the technique to provide the executive with a forecast of acceptable accuracy.

Once these factors and their relationships have been clarified, the forecaster can build a causal model of the system which captures both the facts and the logic of the situation—which is, after all, the basis of sophisticated forecasting.

3. How important is the past in estimating the future? Significant changes in the system—new products, new competitive strategies, and so forth—diminish the similarity of past and future. Over the short term, recent changes are unlikely to cause overall patterns to alter, but over the long term their effects are likely to increase. The executive and the forecaster must discuss these fully.

Exhibit II Flow Chart of TV Distribution System



Three General Types

Once the manager and the forecaster have formulated their problem, the forecaster will be in a position to choose a method.

There are three basic types—qualitative techniques, time series analysis and projection, and causal models. The first uses qualitative data (expert opinion, for example) and information about special events of the kind already mentioned, and may or may not take the past into consideration.

The second, on the other hand, focuses entirely on patterns and pattern changes, and thus relies entirely on historical data.

The third uses highly refined and specific information about relationships between system elements, and is powerful enough to take special events formally into account. As with time series analysis and projection techniques, the past is important to causal models.

These differences imply (quite correctly) that the same type of forecasting technique is not appropriate to forecast sales, say, at all stages of the life cycle of a product—for example, a technique that relies on historical data would not be useful in forecasting the future of a totally new product that has no history.

The major part of the balance of this article will be concerned with the problem of suiting the technique to the life-cycle stages. We hope to give the executive insight into the potential of forecasting by showing how this problem is to be approached. But before we discuss the life cycle, we need to sketch the general functions of the three basic types of techniques in a bit more detail.

Qualitative techniques

Primarily, these are used when data are scarce—for example, when a product is first introduced into a market. They use human judgment and rating schemes to turn qualitative information into quantitative estimates.

The objective here is to bring together in a logical, unbiased, and systematic way all information and judgments which relate to the factors being estimated. Such techniques are frequently used in new-technology areas, where development of a product idea may require several “inventions,” so that R&D demands are difficult to estimate, and where market acceptance and penetration rates are highly uncertain.

The multi-page chart “Basic Forecasting Techniques” presents several examples of this type (see the first section), including market research and the now-familiar Delphi technique.¹ In this chart we have tried to provide a body of basic information about the main kinds of forecasting techniques. Some of the techniques listed are not in reality a single method or model, but a whole family. Thus our statements may not accurately describe all the variations of a technique and should rather be interpreted as descriptive of the basic concept of each.

BASIC FORECASTING TECHNIQUES

Technique	A. Qualitative Methods		
	1. Delphi Method	2. Market Research	3. Panel Consensus
Description	A panel of experts is interrogated by a sequence of questionnaires in which the responses to one questionnaire are used to produce the next questionnaire. Any set of information available to some experts and not others is thus passed on to the others, enabling all the experts to have access to all the information for forecasting. This technique eliminates the bandwagon effect of majority opinion.	The systematic, formal, and conscious procedure for evolving and testing hypotheses about real markets.	This technique is based on the assumption that several experts can arrive at a better forecast than one person. There is no secrecy, and communication is encouraged. The forecasts are sometimes influenced by social factors, and may not reflect a true consensus.
Accuracy Short term (0–3 months) Medium term (3 months–2 years) Long term (2 years & up)	Fair to very good Fair to very good Fair to very good	Excellent Good Fair to good	Poor to fair Poor to fair Poor
Identification of turning points	Fair to good	Fair to very good	Poor to fair
Typical applications	Forecasts of long-range and new-product sales, forecasts of margins.	Forecasts of long-range and new-product sales, forecasts of margins.	Forecasts of long-range and new-product sales, forecasts of margins.
Data required	A coordinator issues the sequence of questionnaires, editing and consolidating the responses.	As a minimum, two sets of reports over time. One needs a considerable collection of market data from questionnaires, surveys, and time series analyses of market variables.	Information from a panel of experts is presented openly in group meetings to arrive at a consensus forecast. Again, a minimum is two sets of reports over time.
Cost of forecasting* With a computer Is calculation possible without a computer?	\$2,000+ Yes	\$5,000+ Yes	\$1,000+ Yes
Time required to develop an application & make a forecast	2 months +	3 months +	2 weeks +
References	North & Pyke, “Probes’ of the Technological Future,” <i>HBR</i> May–June 1969, p. 68.	Bass, King & Pessemeier, <i>Applications of the Sciences in Marketing Management</i> (New York, John Wiley & Sons, Inc., 1968).	—————

A. Qualitative Methods (continued)		B. Time Series Analysis & Projection	
4. Visionary Forecast	5. Historical Analogy	1. Moving Average	2. Exponential Smoothing
A prophecy that uses personal insights, judgment, and, when possible, facts about different scenarios of the future. It is characterized by subjective guesswork and imagination; in general, the methods used are non-scientific.	This is a comparative analysis of the introduction and growth of similar new products that bases the forecast on similarity patterns.	Each point of a moving average of a time series is the arithmetic or weighted average of a number of consecutive points of the series, where the number of data points is chosen so that the effects of seasonals or irregularity or both are eliminated.	This technique is similar to the moving average, except that more recent data points are given more weight. Descriptively, the new forecast is equal to the old one plus some proportion of the past forecasting error. Adaptive forecasting is somewhat the same except that seasonals are also computed. There are many variations of exponential smoothing: some are more versatile than others, some are computationally more complex, some require more computer time.
Poor Poor	Poor Good to fair	Poor to good Poor	Fair to very good Poor to good
Poor	Good to fair	Very poor	Very poor
Poor	Poor to fair	Poor	Poor
Forecasts of long-range and new-product sales, forecasts of margins.	Forecasts of long-range and new-product sales, forecasts of margins.	Inventory control for low-volume items.	Production and inventory control, forecasts of margins and other financial data.
A set of possible scenarios about the future prepared by a few experts in light of past events.	Several years' history of one or more products.	A minimum of two years of sales history, if seasonals are present. Otherwise, less data. (Of course, the more history the better.) The moving average must be specified.	The same as for a moving average.
\$100+ Yes	\$1,000+ Yes	\$.005 Yes	\$.005 Yes
1 week+	1 month+	1 day-	1 day-
—————	Spencer, Clark & Hoguet, <i>Business & Economic Forecasting</i> (Homewood, Illinois, Richard D. Irwin, Inc., 1961).	Hadley, <i>Introduction to Business Statistics</i> (San Francisco, Holden-Day, Inc., 1968).	Brown, "Less Risk in Inventory Estimates," HBR July-August 1959, p. 104.

BASIC FORECASTING TECHNIQUES (continued)

Technique	C. Causal Methods (continued)			
	5. Economic Input-Output Model	6. Diffusion Index	7. Leading Indicator	8. Life-Cycle Analysis
Description	Econometric models and input-output models are sometimes combined for forecasting. The input-output model is used to provide long-term trends for the econometric model; it also stabilizes the econometric model.	The percentage of a group of economic indicators that are going up or down, this percentage then becoming the index.	A time series of an economic activity whose movement in a given direction precedes the movement of some other time series in the same direction is a leading indicator.	This is an analysis and forecasting of new-product growth rates based on S-curves. The phases of product acceptance by the various groups such as innovators, early adapters, early majority, late majority, and laggards are central to the analysis.
Accuracy	Short term (0–3 months) Medium term (3 months–2 years) Long term (2 years & up)	Not applicable Good to very good	Poor to good Poor to good	Poor to good Poor to good
Identification of turning points	Good to excellent	Very poor	Very poor	Poor to good
Typical applications	Company sales for industrial sectors and subsectors.	Forecasts of sales by product class.	Forecasts of sales by product class.	Forecasts of new-product sales.
Data required	The same as for a moving average and X-11.	The same as an intention-to-buy survey.	The same as an intention-to-buy survey + 5 to 10 years' history.	As a minimum, the annual sales of the product being considered or of a similar product. It is often necessary to do market surveys.
*Cost of forecasting With a computer Is calculation possible without a computer?	\$100,000 No	\$1,000 Yes	\$1,000 Yes	\$1,500 Yes
Time required to develop an application & make a forecast	6 months+	1 month+	1 month+	1 month+
References	Evans & Preston, "Discussion Paper #138," Wharton School of Finance & Commerce, The University of Pennsylvania.	Evans, <i>Macro-economic Activity: Theory, Forecasting & Control</i> (New York, Harper & Row Publishers, Inc., 1969).	Evans, <i>Macro-economic Activity: Theory, Forecasting & Control</i> (New York, Harper & Row Publishers, Inc., 1969).	Bass, "A New Product Growth Model for Consumer Durables," <i>Management Science</i> , January 1969.

A disclaimer about estimates in the chart is also in order. Estimates of costs are approximate, as are computation times, accuracy ratings, and ratings for turning-point identification. The costs of some procedures depend on whether they are being used routinely or are set up for a single forecast; also, if weightings or seasonals have to be determined anew each time a forecast is made, costs increase significantly. Still, the figures we present may serve as general guidelines. The reader may find frequent reference to this gate-fold helpful for the remainder of the article.

Time series analysis

These are statistical techniques used when several years' data for a product or product line are available and when relationships and trends are both clear and relatively stable.

One of the basic principles of statistical forecasting—indeed, of all forecasting when historical data are available—is that the forecaster should use the data on past performance to get a “speedometer reading” of the current rate (of sales, say) and of how fast this rate is increasing or decreasing. The current rate and changes in the rate—“acceleration” and “deceleration”—constitute the basis of forecasting. Once they are known, various mathematical techniques can develop projections from them.

The matter is not so simple as it sounds, however. It is usually difficult to make projections from raw data since the rates and trends are not immediately obvious; they are mixed up with seasonal variations, for example, and perhaps distorted by such factors as the effects of a large sales promotion campaign. The raw data must be massaged before they are usable, and this is frequently done by time series analysis.

Now, a time series is a set of chronologically ordered points of raw data—for example, a division's sales of a given product, by month, for several years. Time series analysis helps to identify and explain:

- Any regularity or systematic variation in the series of data which is due to seasonality—the “seasonals.”
- Cyclical patterns that repeat any two or three years or more.
- Trends in the data.
- Growth rates of these trends.

(Unfortunately, most existing methods identify only the seasonals, the combined effect of trends and cycles, and the irregular, or chance, component. That is, they do not separate trends from cycles. We shall return to this point when we discuss time series analysis in the final stages of product maturity.)

Once the analysis is complete, the work of projecting future sales (or whatever) can begin.

We should note that while we have separated analysis from projection here for purposes of explanation, most statistical forecasting techniques actually combine both functions in a single operation.

A future like the past: It is obvious from this description that all statistical techniques are based on the assumption that existing patterns will continue into the future. This assumption is more likely to be correct over the short term than it is over the long term, and for this reason these techniques provide us with reasonably accurate forecasts for the immediate future but do quite poorly further into the future (unless the data patterns are extraordinarily stable).

For this same reason, these techniques ordinarily cannot predict when the rate of growth in a trend will change significantly—for example, when a period of slow growth in sales will suddenly change to a period of rapid decay.

Such points are called turning points. They are naturally of the greatest consequence to the manager, and, as we shall see, the forecaster must use different tools from pure statistical techniques to predict when they will occur.

Causal models:

When historical data are available and enough analysis has been performed to spell out explicitly the relationships between the factor to be forecast and other factors (such as related businesses, economic forces, and socioeconomic factors), the forecaster often constructs a causal model.

A causal model is the most sophisticated kind of forecasting tool. It expresses mathematically the relevant causal relationships, and may include pipeline considerations (i.e., inventories) and market survey information. It may also directly incorporate the results of a time series analysis.

The causal model takes into account everything known of the dynamics of the flow system and utilizes predictions of related events such as competitive actions, strikes, and promotions. If the data are available, the model generally includes factors for each location in the flow chart (as illustrated in Exhibit II) and connects these by equations to describe overall product flow.

If certain kinds of data are lacking, initially it may be necessary to make assumptions about some of the relationships and then track what is happening to determine if the assumptions are true. Typically, a causal model is continually revised as more knowledge about the system becomes available.

Again, see the gatefold for a rundown on the most common types of causal techniques. As the chart shows, causal models are by far the best for predicting turning points and preparing long-range forecasts.

Methods, Products & the Life Cycle

At each stage of the life of a product, from conception to steady-state sales, the decisions that management must make are characteristically quite different, and they require different kinds of information as a base. The forecasting techniques that provide these sets of information differ analogously. Exhibit III summarizes the life stages of a product, the typical decisions made at each, and the main forecasting techniques suitable at each.

Exhibit III Types of Decisions Made Over a Product's Life Cycle, with Related Forecasting Techniques

Stage of life cycle	Product development	Market testing & early introduction	Rapid growth	Steady state
Typical decisions	Amount of development effort Product design Business strategies	Optimum facility size Marketing strategies, including distribution & pricing	Facilities expansion Marketing strategies Production planning Sales	Promotions, specials Pricing Production planning Inventories
Forecasting techniques	Delphi method Historical analysis of comparable products Priority pattern analysis Input-output analysis Panel consensus	Consumer surveys Tracking & warning systems Market tests Experimental designs	Statistical techniques for identifying turning points Tracking & warning systems Market surveys Intention-to-buy surveys	Time series analysis & projection Causal & econometric models Market surveys for tracking & warning Life-cycle analysis

Equally, different products may require different kinds of forecasting. Two CGW products that have been handled quite differently are the major glass components for color TV tubes, of which Corning is a prime supplier, and Corning Ware cookware, a proprietary consumer product line. We shall trace the forecasting methods used at each of the four different stages of maturity of these products to give some firsthand insight into the choice and application of some of the major techniques available today.

Before we begin, let us note how the situations differ for the two kinds of products:

For a consumer product like the cookware, the manufacturer's control of the distribution pipeline extends at least through the distributor level. Thus the manufacturer can effect or control consumer sales quite directly, as well as directly control some of the pipeline elements.

Many of the changes in shipment rates and in overall profitability are therefore due to actions taken by manufacturers themselves. Tactical decisions on promotions, specials, and pricing are usually at their discretion as well. The technique selected by the forecaster for projecting sales therefore should permit incorporation of such "special information." One may have to start with simple techniques and work up to more sophisticated ones that embrace such possibilities, but the final goal is there.

Where the manager's company supplies a component to an OEM, as Corning does for tube manufacturers, the company does not have such direct influence or control over either the pipeline elements or final consumer sales. It may be

impossible for the company to obtain good information about what is taking place at points further along the flow system (as in the upper segment of Exhibit II), and, in consequence, the forecaster will necessarily be using a different genre of forecasting from what is used for a consumer product.

Between these two examples, our discussion will embrace nearly the whole range of forecasting techniques. As necessary, however, we shall touch on other products and other forecasting methods.

1. Product Development

In the early stages of product development, the manager wants answers to questions such as these:

- What are the alternative growth opportunities to pursuing product X?
- How have established products similar to X fared?
- Should we enter this business; and if so, in what segments?
- How should we allocate R&D efforts and funds?
- How successful will different product concepts be?
- How will product X fit into the markets five or ten years from now?

Forecasts that help to answer these long-range questions must necessarily have long horizons themselves.

A common objection to much long-range forecasting is that it is virtually impossible to predict with accuracy what will happen several years into the future. We agree that uncertainty increases when a forecast is made for a period more than two years out. However, at the very least, the forecast and a measure of its accuracy enable the manager to know the risks in pursuing a selected strategy and in this knowledge to choose an appropriate strategy from those available.

Systematic market research is, of course, a mainstay in this area. For example, priority pattern analysis can describe consumers' preferences and the likelihood they will buy a product, and thus is of great value in forecasting (and updating) penetration levels and rates. But there are other tools as well, depending on the state of the market and the product concept.

For a defined market

While there can be no direct data about a product that is still a gleam in the eye, information about its likely performance can be gathered in a number of ways, provided the market in which it is to be sold is a known entity.

First, one can compare a proposed product with competitors' present and planned products, ranking it on quantitative scales for different factors. We call this product differences measurement.²

If this approach is to be successful, it is essential that the (in-house) experts who provide the basic data come from different disciplines—marketing, R&D, manufacturing, legal, and so on—and that their opinions be unbiased.

Second, and more formalistically, one can construct disaggregate market models by separating off different segments of a complex market for individual study and consideration. Specifically, it is often useful to project the S-shaped growth curves for the levels of income of different geographical regions.

When color TV bulbs were proposed as a product, CGW was able to identify the factors that would influence sales growth. Then, by disaggregating consumer demand and making certain assumptions about these factors, it was possible to develop an S-curve for rate of penetration of the household market that proved most useful to us.

Third, one can compare a projected product with an "ancestor" that has similar characteristics. In 1965, we disaggregated the market for color television by income levels and geographical regions and compared these submarkets with the historical pattern of black-and-white TV market growth. We justified this procedure by arguing that color TV represented an advance over black-and-white analogous to (although less intense than) the advance that black-and-white TV represented over radio. The analyses of black-and-white TV market growth also enabled us to estimate the variability to be expected—that is, the degree to which our projections would differ from actual as the result of economic and other factors.

The prices of black-and-white TV and other major household appliances in 1949, consumer disposable income in 1949, the prices of color TV and other appliances in 1965, and consumer disposable income for 1965 were all profitably considered in

developing our long-range forecast for color-TV penetration on a national basis. The success patterns of black-and-white TV, then, provided insight into the likelihood of success and sales potential of color TV.

Our predictions of consumer acceptance of Corning Ware cookware, on the other hand, were derived primarily from one expert source, a manager who thoroughly understood consumer preferences and the housewares market. These predictions have been well borne out. This reinforces our belief that sales forecasts for a new product that will compete in an existing market are bound to be incomplete and uncertain unless one culls the best judgments of fully experienced personnel.

For an undefined market

Frequently, however, the market for a new product is weakly defined or few data are available, the product concept is still fluid, and history seems irrelevant. This is the case for gas turbines, electric and steam automobiles, modular housing, pollution measurement devices, and time-shared computer terminals.

Many organizations have applied the Delphi method of soliciting and consolidating experts' opinions under these circumstances. At CGW, in several instances, we have used it to estimate demand for such new products, with success.

Input-output analysis, combined with other techniques, can be extremely useful in projecting the future course of broad technologies and broad changes in the economy. The basic tools here are the input-output tables of U.S. industry for 1947, 1958, and 1963, and various updatings of the 1963 tables prepared by a number of groups who wished to extrapolate the 1963 figures or to make forecasts for later years.

Since a business or product line may represent only a small sector of an industry, it may be difficult to use the tables directly. However, a number of companies are disaggregating industries to evaluate their sales potential and to forecast changes in product mixes—the phasing out of old lines and introduction of others. For example, Quantum-Science Corporation (MAPTEK) has developed techniques that make input-output analyses more directly useful to people in the electronics business today. (Other techniques, such as panel consensus and visionary forecasting, seem less effective to us, and we cannot evaluate them from our own experience.)

2. Testing & Introduction

Before a product can enter its (hopefully) rapid penetration stage, the market potential must be tested out and the product must be introduced—and then more market testing may be advisable. At this stage, management needs answers to these questions:

- What shall our marketing plan be—which markets should we enter and with what production quantities?
- How much manufacturing capacity will the early production stages require?
- As demand grows, where should we build this capacity?
- How shall we allocate our R&D resources over time?

Significant profits depend on finding the right answers, and it is therefore economically feasible to expend relatively large amounts of effort and money on obtaining good forecasts, short-, medium-, and long-range.

A sales forecast at this stage should provide three points of information: the date when rapid sales will begin, the rate of market penetration during the rapid-sales stage, and the ultimate level of penetration, or sales rate, during the steady-state stage.

Using early data

The date when a product will enter the rapid-growth stage is hard to predict three or four years in advance (the usual horizon). A company's only recourse is to use statistical tracking methods to check on how successfully the product is being introduced, along with routine market studies to determine when there has been a significant increase in the sales rate.

Furthermore, the greatest care should be taken in analyzing the early sales data that start to accumulate once the product has been introduced into the market. For example, it is important to distinguish between sales to innovators, who will try anything new, and sales to imitators, who will buy a product only after it has been accepted by innovators, for it is the latter group that provides demand stability. Many new products have initially appeared successful because of purchases by innovators, only to fail later in the stretch.

Tracking the two groups means market research, possibly via opinion panels. A panel ought to contain both innovators and imitators, since innovators can teach one a lot about how to improve a product while imitators provide insight into the desires and expectations of the whole market.

The color TV set, for example, was introduced in 1954, but did not gain acceptance from the majority of consumers until late 1964. To be sure, the color TV set could not leave the introduction stage and enter the rapid-growth stage until the networks had substantially increased their color programming. However, special flag signals like “substantially increased network color programming” are likely to come after the fact, from the planning viewpoint; and in general, we find, scientifically designed consumer surveys conducted on a regular basis provide the earliest means of detecting turning points in the demand for a product.

Similar-product technique

Although statistical tracking is a useful tool during the early introduction stages, there are rarely sufficient data for statistical forecasting. Market research studies can naturally be useful, as we have indicated. But, more commonly, the forecaster tries to identify a similar, older product whose penetration pattern should be similar to that of the new product, since overall markets can and do exhibit consistent patterns.

Again, let’s consider color television and the forecasts we prepared in 1965.

For the year 1947–1968, Exhibit IV shows total consumer expenditures, appliance expenditures, expenditures for radios and TVs, and relevant percentages. Column 4 shows that total expenditures for appliances are relatively stable over periods of several years; hence, new appliances must compete with existing ones, especially during recessions (note the figures for 1948–1949, 1953–1954, 1957–1958, and 1960–1961).

Exhibit IV Expenditures on Appliances Versus All Consumer Goods
(In billions of dollars)

Year (1)	All consumer goods* (2)	Household appliances† (3)	Radio, TV & other† (4)	Totals of columns 3 & 4 (5)	Column 5 ÷ Column 2 (6)	Column 4 ÷ Column 2 (7)
1947	110.9	3.18	1.43	4.61	4.16%	1.29%
1948	118.9	3.47	1.48	4.95	4.16	1.23
1949	119.1	3.13	1.70	4.83	4.06	1.43
1950	128.6	3.94	2.46	6.40	4.98	1.91
1951	138.4	3.87	2.26	6.13	4.43	1.63
1952	143.3	3.82	2.37	6.19	4.32	1.65
1953	150.0	3.99	2.61	6.60	4.40	1.74
1954	151.1	4.02	2.74	6.77	4.48	1.81
1955	162.9	4.69	2.79	7.48	4.59	1.71
1956	168.2	4.89	2.87	7.76	4.61	1.71
1957	176.4	4.63	3.00	7.63	4.33	1.70
1958	178.1	4.44	3.07	7.51	4.22	1.72
1959	190.9	4.86	3.42	8.28	4.34	1.79
1960	196.6	4.74	3.62	8.36	4.25	1.84
1961	200.1	4.77	3.76	8.53	4.26	1.88
1962	212.1	5.01	3.94	8.95	4.22	1.86
1963	222.5	5.24	4.54	9.78	4.40	2.04
1964	237.9	5.74	5.41	11.15	4.69	2.27
1965	257.4	6.03	6.01	12.04	4.68	2.33
1966	277.7	6.77	6.91	13.68	4.93	2.49
1967	288.1	7.09	7.41	14.50	5.03	2.57
1968	313.9	7.80	7.85	15.65	4.99	2.50

*Data obtained from Survey of Current Business, Personal Consumption Expenditure Tables (U.S. Department of Commerce, July issues).
†Data obtained from the Survey of Current Business Statistics (U.S. Department of Commerce, 1969 Biennial Edition).

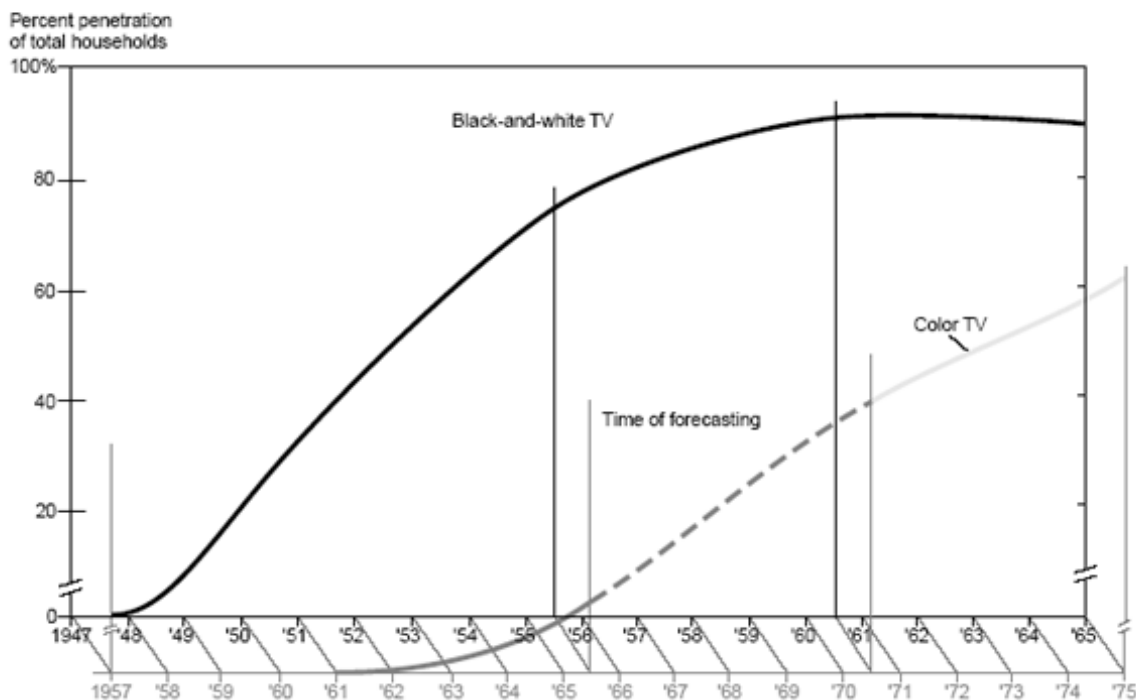
Certain special fluctuations in these figures are of special significance here. When black-and-white TV was introduced as a new product in 1948–1951, the ratio of expenditures on radio and TV sets to total expenditures for consumer goods (see column 7) increased about 33% (from 1.23% to 1.63%), as against a modest increase of only 13% (from 1.63% to 1.88%) in the ratio for the next decade. (A similar increase of 33% occurred in 1962–1966 as color TV made its major penetration.)

Probably the acceptance of black-and-white TV as a major appliance in 1950 caused the ratio of all major household appliances to total consumer goods (see column 5) to rise to 4.98%; in other words, the innovation of TV caused the consumer to start spending more money on major appliances around 1950. Our expectation in mid-1965 was that the introduction of

color TV would induce a similar increase. Thus, although this product comparison did not provide us with an accurate or detailed forecast, it did place an upper bound on the future total sales we could expect.

The next step was to look at the cumulative penetration curve for black-and-white TVs in U.S. households, shown in Exhibit V. We assumed color-TV penetration would have a similar S-curve, but that it would take longer for color sets to penetrate the whole market (that is, reach steady-state sales). Whereas it took black-and-white TV 10 years to reach steady state, qualitative expert-opinion studies indicated that it would take color twice that long—hence the more gradual slope of the color-TV curve.

Exhibit V Long-term Household Penetration Curves for Color and Black-and-White TV



At the same time, studies conducted in 1964 and 1965 showed significantly different penetration sales for color TV in various income groups, rates that were helpful to us in projecting the color-TV curve and tracking the accuracy of our projection.

With these data and assumptions, we forecast retail sales for the remainder of 1965 through mid-1970 (see the dotted section of the lower curve in Exhibit V). The forecasts were accurate through 1966 but too high in the following three years, primarily because of declining general economic conditions and changing pricing policies

We should note that when we developed these forecasts and techniques, we recognized that additional techniques would be necessary at later times to maintain the accuracy that would be needed in subsequent periods. These forecasts provided acceptable accuracy for the time they were made, however, since the major goal then was only to estimate the penetration rate and the ultimate, steady-state level of sales. Making refined estimates of how the manufacturing-distribution pipelines will behave is an activity that properly belongs to the next life-cycle stage.

Other approaches: When it is not possible to identify a similar product, as was the case with CGW’s self-cleaning oven and flat-top cooking range (Counterange), another approach must be used. For the purposes of initial introduction into the markets, it may only be necessary to determine the minimum sales rate required for a product venture to meet corporate objectives. Analyses like input-output, historical trend, and technological forecasting can be used to estimate this minimum. Also, the feasibility of not entering the market at all, or of continuing R&D right up to the rapid-growth stage, can best be determined by sensitivity analysis.

Predicting rapid growth

To estimate the date by which a product will enter the rapid-growth stage is another matter. As we have seen, this date is a function of many factors: the existence of a distribution system, customer acceptance of or familiarity with the product concept, the need met by the product, significant events (such as color network programming), and so on.

As well as by reviewing the behavior of similar products, the date may be estimated through Delphi exercises or through rating and ranking schemes, whereby the factors important to customer acceptance are estimated, each competitor product is rated on each factor, and an overall score is tallied for the competitor against a score for the new product.

As we have said, it is usually difficult to forecast precisely when the turning point will occur; and, in our experience, the best accuracy that can be expected is within three months to two years of the actual time.

It is occasionally true, of course, that one can be certain a new product will be enthusiastically accepted. Market tests and initial customer reaction made it clear there would be a large market for Corning Ware cookware. Since the distribution system was already in existence, the time required for the line to reach rapid growth depended primarily on our ability to manufacture it. Sometimes forecasting is merely a matter of calculating the company's capacity—but not ordinarily.

3. Rapid Growth

When a product enters this stage, the most important decisions relate to facilities expansion. These decisions generally involve the largest expenditures in the cycle (excepting major R&D decisions), and commensurate forecasting and tracking efforts are justified.

Forecasting and tracking must provide the executive with three kinds of data at this juncture:

- Firm verification of the rapid-growth rate forecast made previously.
- A hard date when sales will level to “normal,” steady-state growth.
- For component products, the deviation in the growth curve that may be caused by characteristic conditions along the pipeline—for example, inventory blockages.

Forecasting the growth rate

Medium- and long-range forecasting of the market growth rate and of the attainment of steady-state sales requires the same measures as does the product introduction stage—detailed marketing studies (especially intention-to-buy surveys) and product comparisons.

When a product has entered rapid growth, on the other hand, there are generally sufficient data available to construct statistical and possibly even causal growth models (although the latter will necessarily contain assumptions that must be verified later).

We estimated the growth rate and steady-state rate of color TV by a crude econometric-marketing model from data available at the beginning of this stage. We conducted frequent marketing studies as well.

The growth rate for Corning Ware Cookware, as we explained, was limited primarily by our production capabilities; and hence the basic information to be predicted in that case was the date of leveling growth. Because substantial inventories buffered information on consumer sales all along the line, good field data were lacking, which made this date difficult to estimate. Eventually we found it necessary to establish a better (more direct) field information system.

As well as merely buffering information, in the case of a component product, the pipeline exerts certain distorting effects on the manufacturer's demand; these effects, although highly important, are often illogically neglected in production or capacity planning.

Simulating the pipeline

While the ware-in-process demand in the pipeline has an S-curve like that of retail sales, it may lag or lead sales by several months, distorting the shape of the demand on the component supplier.

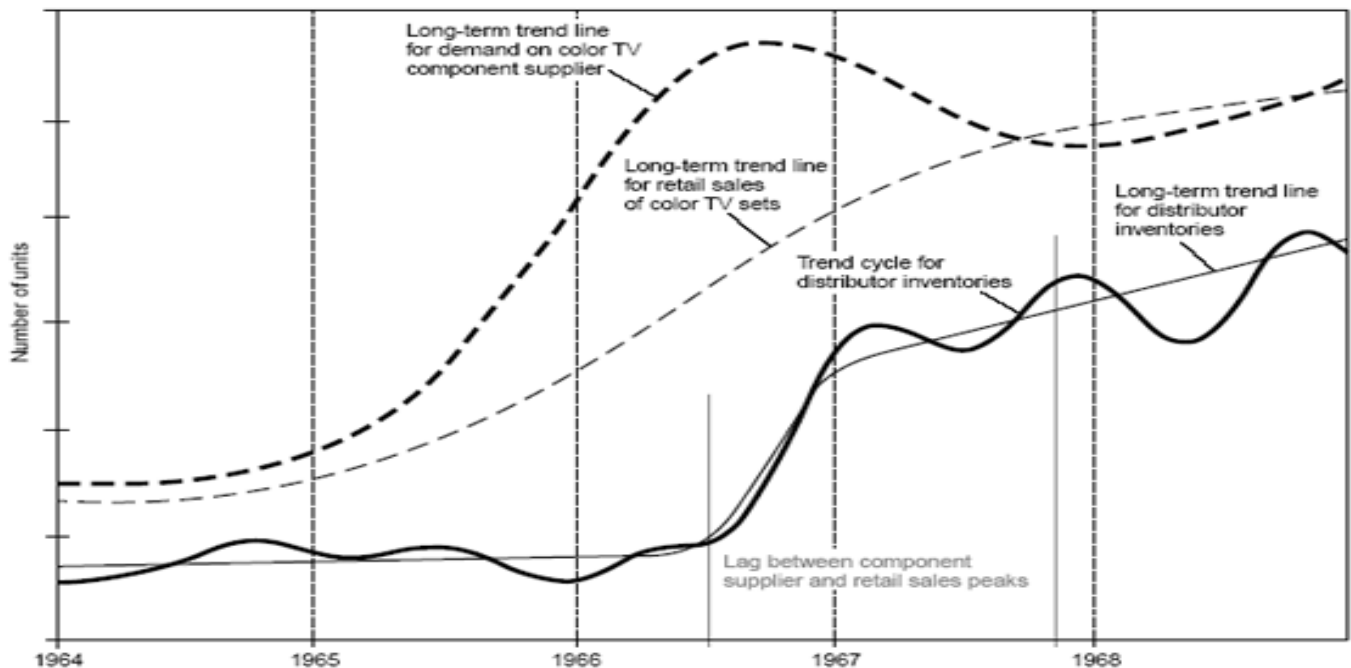
Exhibit VI shows the long-term trend of demand on a component supplier other than Corning as a function of distributor sales and distributor inventories. As one can see from this curve, supplier sales may grow relatively sharply for several months and peak before retail sales have leveled off. The implications of these curves for facilities planning and allocation are obvious.

Exhibit VI Patterns for Color-TV Distributor Sales, Distributor Inventories, and Component Sales Note: Scales are different for component sales, distributor inventories, and distributor sales, with the patterns put on the same graph for illustrative purposes.

Here we have used components for color TV sets for our illustration because we know from our own experience the importance of the long flow time for color TVs that results from the many sequential steps in manufacturing and distribution (recall Exhibit II). There are more spectacular examples; for instance, it is not uncommon for the flow time from component supplier to consumer to stretch out to two years in the case of truck engines.

To estimate total demand on CGW production, we used a retail demand model and a pipeline simulation. The model incorporated penetration rates, mortality curves, and the like. We combined the data generated by the model with market-share data, data on glass losses, and other information to make up the corpus of inputs for the pipeline simulation. The simulation output allowed us to apply projected curves like the ones shown in Exhibit VI to our own component-manufacturing planning.

Exhibit VI Patterns for Color-TV Distributor Sales, Distributor Inventories, and Component Sales



Simulation is an excellent tool for these circumstances because it is essentially simpler than the alternative—namely, building a more formal, more “mathematical” model. That is, simulation bypasses the need for analytical solution techniques and for mathematical duplication of a complex environment and allows experimentation. Simulation also informs us how the pipeline elements will behave and interact over time—knowledge that is very useful in forecasting, especially in constructing formal causal models at a later date.

Tracking & warning

This knowledge is not absolutely “hard,” of course, and pipeline dynamics must be carefully tracked to determine if the various estimates and assumptions made were indeed correct. Statistical methods provide a good short-term basis for estimating and checking the growth rate and signaling when turning points will occur.

In late 1965 it appeared to us that the ware-in-process demand was increasing, since there was a consistent positive difference between actual TV bulb sales and forecasted bulb sales. Conversations with product managers and other personnel indicated there might have been a significant change in pipeline activity; it appeared that rapid increases in retail demand were boosting glass requirements for ware-in-process, which could create a hump in the S-curve like the one illustrated in Exhibit VI. This humping provided additional profit for CGW in 1966 but had an adverse effect in 1967. We were able to predict this hump, but unfortunately we were unable to reduce or avoid it because the pipeline was not sufficiently under our control.

The inventories all along the pipeline also follow an S-curve (as shown in Exhibit VI), a fact that creates and compounds two characteristic conditions in the pipeline as a whole: initial overfilling and subsequent shifts between too much and too little inventory at various points—a sequence of feast-and-famine conditions.

For example, the simpler distribution system for Corning Ware had an S-curve like the ones we have examined. When the retail sales slowed from rapid to normal growth, however, there were no early indications from shipment data that this crucial turning point had been reached. Data on distributor inventories gave us some warning that the pipeline was over filling, but the turning point at the retail level was still not identified quickly enough, as we have mentioned before, because of lack of good data at the level. We now monitor field information regularly to identify significant changes, and adjust our shipment forecasts accordingly.

Main concerns

One main activity during the rapid-growth stage, then, is to check earlier estimates and, if they appear incorrect, to compute as accurately as possible the error in the forecast and obtain a revised estimate.

In some instances, models developed earlier will include only “macroterms”; in such cases, market research can provide information needed to break these down into their components. For example, the color-TV forecasting model initially considered only total set penetrations at different income levels, without considering the way in which the sets were being used. Therefore, we conducted market surveys to determine set use more precisely.

Equally, during the rapid-growth stage, submodels of pipeline segments should be expanded to incorporate more detailed information as it is received. In the case of color TV, we found we were able to estimate the overall pipeline requirements for glass bulbs, the CGW market-share factors, and glass losses, and to postulate a probability distribution around the most likely estimates. Over time, it was easy to check these forecasts against actual volume of sales, and hence to check on the procedures by which we were generating them.

We also found we had to increase the number of factors in the simulation model—for instance, we had to expand the model to consider different sizes of bulbs—and this improved our overall accuracy and usefulness. The preceding is only one approach that can be used in forecasting sales of new products that are in a rapid growth. Others have discussed different ones.³

4. Steady State

The decisions the manager at this stage are quite different from those made earlier. Most of the facilities planning has been squared away, and trends and growth rates have become reasonably stable. It is possible that swings in demand and profit will occur because of changing economic conditions, new and competitive products, pipeline dynamics, and so on, and the manager will have to maintain the tracking activities and even introduce new ones. However, by and large, the manager will concentrate forecasting attention on these areas:

- Long- and short-term production planning.
- Setting standards to check the effectiveness of marketing strategies.
- Projections designed to aid profit planning.

The manager will also need a good tracking and warning system to identify significantly declining demand for the product (but hopefully that is a long way off). To be sure, the manager will want margin and profit projection and long-range forecasts to assist planning at the corporate level. However, short- and medium-term sales forecasts are basic to these more elaborate undertakings, and we shall concentrate on sales forecasts.

Adequate tools at hand

In planning production and establishing marketing strategy for the short and medium term, the manager’s first considerations are usually an accurate estimate of the present sales level and an accurate estimate of the rate at which this level is changing.

The forecaster thus is called on for two related contributions at this stage:

To provide estimates of trends and seasonals, which obviously affect the sales level. Seasonals are particularly important for both overall production planning and inventory control. To do this, the forecaster needs to apply time series analysis and projection techniques—that is, statistical techniques.

To relate the future sales level to factors that are more easily predictable, or have a “lead” relationship with sales, or both. To do this the forecaster needs to build causal models.

The type of product under scrutiny is very important in selecting the techniques to be used.

For Corning Ware, where the levels of the distribution system are organized in a relatively straightforward way, we use statistical methods to forecast shipments and field information to forecast changes in shipment rates. We are now in the process of incorporating special information—marketing strategies, economic forecasts, and so on—directly into the shipment forecasts. This is leading us in the direction of a causal forecasting model.


On the other hand, a component supplier may be able to forecast total sales with sufficient accuracy for broad-load production planning, but the pipeline environment may be so complex that the best recourse for short-term projections is to rely primarily on salespersons’ estimates. We find this true, for example, in estimating the demand for TV glass by size and customer. In such cases, the best role for statistical methods is providing guides and checks for salespersons’ forecasts.

In general, however, at this point in the life cycle, sufficient time series data are available and enough causal relationships are known from direct experience and market studies so that the forecaster can indeed apply these two powerful sets of tools. Historical data for at least the last several years should be available. The forecaster will use all of it, one way or another.


We might mention a common criticism at this point. People frequently object to using more than a few of the most recent data points (such as sales figures in the immediate past) for building projections, since, they say, the current situation is always so dynamic and conditions are changing so radically and quickly that historical data from further back in time have little or no value. We think this point of view had little validity. A graph of several years’ sales data, such as the one shown in Part A of Exhibit VII, gives an impression of a sales trend one could not possibly get if one were to look only at two or three of the latest data points.

Exhibit V Data Plots of Factory Sales of Color TV Sets


Part A. Raw data for factory sales of color TV sets



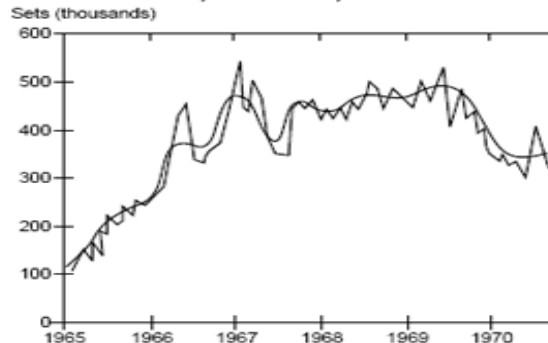
Part C. Factory sales of color TV sets (deseasonalized)



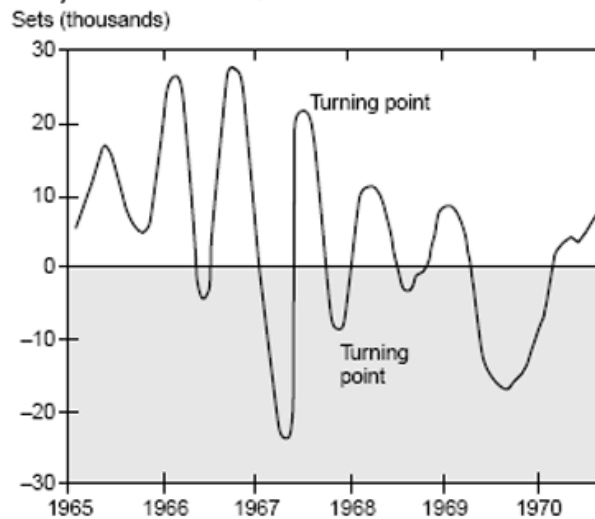
Part B. Seasonals for factory sales of color TV sets



Part D. Final trend cycle of factory sales of color TV sets



Part E. Changes in final trend cycle (growth rate) of factory sales of color TV sets



In practice, we find, overall patterns tend to continue for a minimum of one or two quarters into the future, even when special conditions cause sales to fluctuate for one or two (monthly) periods in the immediate future.

For short-term forecasting for one to three months ahead, the effects of such factors as general economic conditions are minimal, and do not cause radical shifts in demand patterns. And because trends tend to change gradually rather than suddenly, statistical and other quantitative methods are excellent for short-term forecasting. Using one or only a few of the most recent data points will result in giving insufficient consideration of the nature of trends, cycles, and seasonal fluctuations in sales.

Granting the applicability of the techniques, we must go on to explain how the forecaster identifies precisely what is happening when sales fluctuate from one period to the next and how such fluctuations can be forecast.

Sorting trends & seasonals

A trend and a seasonal are obviously two quite different things, and they must be handled separately in forecasting.

Consider what would happen, for example, if a forecaster were merely to take an average of the most recent data points along a curve, combine this with other, similar average points stretching backward into the immediate past, and use these as the basis for a projection. The forecaster might easily overreact to random changes, mistaking them for evidence of a prevailing trend, mistake a change in the growth rate for a seasonal, and so on.

To avoid precisely this sort of error, the moving average technique, which is similar to the hypothetical one just described, uses data points in such a way that the effects of seasonals (and irregularities) are eliminated.

Furthermore, the executive needs accurate estimates of trends and accurate estimates of seasonality to plan broad-load production, to determine marketing efforts and allocations, and to maintain proper inventories—that is, inventories that are adequate to customer demand but are not excessively costly.

Before going any further, it might be well to illustrate what such sorting-out looks like. Parts A, B, and C of Exhibit VII show the initial decomposition of raw data for factory sales of color TV sets between 1965 and mid-1970. Part A presents the raw data curve. Part B shows the seasonal factors that are implicit in the raw data—quite a consistent pattern, although there is some variation from year to year. (In the next section we shall explain where this graph of the seasonals comes from.)

Part C shows the result of discounting the raw data curve by the seasonals of Part B; this is the so-called deseasonalized data curve. Next, in Part D, we have drawn the smoothest or “best” curve possible through the deseasonalized curve, thereby obtaining the trend cycle. (We might further note that the differences between this trend-cycle line and the deseasonalized data curve represent the irregular or nonsystematic component that the forecaster must always tolerate and attempt to explain by other methods.)

In sum, then, the objective of the forecasting technique used here is to do the best possible job of sorting out trends and seasonalities. Unfortunately, most forecasting methods project by a smoothing process analogous to that of the moving average technique, or like that of the hypothetical technique we described at the beginning of this section, and separating trends and seasonalities more precisely will require extra effort and cost.

Still, sorting-out approaches have proved themselves in practice. We can best explain the reasons for their success by roughly outlining the way we construct a sales forecast on the basis of trends, seasonalities, and data derived from them. This is the method:

Graph the rate at which the trend is changing. For the illustration given in Exhibit VII, this graph is shown in Part E. This graph describes the successive ups and downs of the trend cycle shown in Part D.

Project this growth rate forward over the interval to be forecasted. Assuming we were forecasting back in mid-1970, we should be projecting into the summer months and possibly into the early fall. Add this growth rate (whether positive or negative) to the present sales rate. This might be called the unseasonalized sales rate.

Project the seasonalities of Part B for the period in question, and multiply the unseasonalized forecasted rate by these seasonalities. The product will be the forecasted sales rate, which is what we desired.

In special cases where there are no seasonalities to be considered, of course, this process is much simplified, and fewer data and simpler techniques may be adequate.

We have found that an analysis of the patterns of change in the growth rate gives us more accuracy in predicting turning points (and therefore changes from positive to negative growth, and vice versa) than when we use only the trend cycle.

The main advantage of considering growth change, in fact, is that it is frequently possible to predict earlier when a no-growth situation will occur. The graph of change in growth thus provides an excellent visual base for forecasting and for identifying the turning point as well.

X-11 technique

The reader will be curious to know how one breaks the seasonalities out of raw sales data and exactly how one derives the change-in-growth curve from the trend line.

One of the best techniques we know for analyzing historical data in depth to determine seasonalities, present sales rate, and growth is the X-11 Census Bureau Technique, which simultaneously removes seasonalities from raw information and fits a trend-cycle line to the data. It is very comprehensive: at a cost of about \$10, it provides detailed information on seasonalities, trends, the accuracy of the seasonalities and the trend cycle fit, and a number of other measures. The output includes plots of the trend cycle and the growth rate, which can concurrently be received on graphic displays on a time-shared terminal.

Although the X-11 was not originally developed as a forecasting method, it does establish a base from which good forecasts can be made. One should note, however, that there is some instability in the trend line for the most recent data points, since the X-11, like virtually all statistical techniques, uses some form of moving average. It has therefore proved of value to study the changes in growth pattern as each new growth point is obtained.

In particular, when recent data seem to reflect sharp growth or decline in sales or any other market anomaly, the forecaster should determine whether any special events occurred during the period under consideration—promotion, strikes, changes in the economy, and so on. The X-11 provides the basic instrumentation needed to evaluate the effects of such events.

Generally, even when growth patterns can be associated with specific events, the X-11 technique and other statistical methods do not give good results when forecasting beyond six months, because of the uncertainty or unpredictable nature of the events. For short-term forecasts of one to three months, the X-11 technique has proved reasonably accurate.

We have used it to provide sales estimates for each division for three periods into the future, as well as to determine changes in sales rates. We have compared our X-11 forecasts with forecasts developed by each of several divisions, where the divisions have used a variety of methods, some of which take into account salespersons' estimates and other special

knowledge. The forecasts using the X_{-11} technique were based on statistical methods alone, and did not consider any special information.

The division forecasts had slightly less error than those provided by the X_{-11} method; however, the division forecasts have been found to be slightly biased on the optimistic side, whereas those provided by the X_{-11} method are unbiased. This suggested to us that a better job of forecasting could be done by combining special knowledge, the techniques of the division, and the X_{-11} method. This is actually being done now by some of the divisions, and their forecasting accuracy has improved in consequence.

The X_{-11} method has also been used to make sales projections for the immediate future to serve as a standard for evaluating various marketing strategies. This has been found to be especially effective for estimating the effects of price changes and promotions.

As we have indicated earlier, trend analysis is frequently used to project annual data for several years to determine what sales will be if the current trend continues. Regression analysis and statistical forecasts are sometimes used in this way—that is, to estimate what will happen if no significant changes are made. Then, if the result is not acceptable with respect to corporate objectives, the company can change its strategy.

Econometric models

Over a long period of time, changes in general economic conditions will account for a significant part of the change in a product's growth rate. Because economic forecasts are becoming more accurate and also because there are certain general "leading" economic forces that change before there are subsequent changes in specific industries, it is possible to improve the forecasts of businesses by including economic factors in the forecasting model.

However, the development of such a model, usually called an econometric model, requires sufficient data so that the correct relationships can be established.

During the rapid-growth state of color TV, we recognized that economic conditions would probably effect the sales rate significantly. However, the macro-analyses of black-and-white TV data we made in 1965 for the recessions in the late 1940s and early 1950s did not show any substantial economic effects at all; hence we did not have sufficient data to establish good econometric relationships for a color TV model. (A later investigation did establish definite losses in color TV sales in 1967 due to economic conditions.)

In 1969 Corning decided that a better method than the X_{-11} was definitely needed to predict turning points in retail sales for color TV six months to two years into the future. Statistical methods and salespersons' estimates cannot spot these turning points far enough in advance to assist decision making; for example, a production manager should have three to six months' warning of such changes in order to maintain a stable work force.

Adequate data seemed to be available to build an econometric model, and analyses were therefore begun to develop such a model for both black-and-white and color TV sales. Our knowledge of seasonals, trends, and growth for these products formed a natural base for constructing the equations of the models.

The economic inputs for the model are primarily obtained from information generated by the Wharton Econometric Model, but other sources are also utilized.

Using data extending through 1968, the model did reasonably well in predicting the downturn in the fourth quarter of 1969 and, when 1969 data were also incorporated into the model, accurately estimated the magnitude of the drop in the first two quarters of 1970. Because of lead-lag relationships and the ready availability of economic forecasts for the factors in the model, the effects of the economy on sales can be estimated for as far as two years into the future.

In the steady-state phase, production and inventory control, group-item forecasts, and long-term demand estimates are particularly important. The interested reader will find a discussion of these topics on the reverse of the gatefold.

Finally, through the steady-state phase, it is useful to set up quarterly reviews where statistical tracking and warning charts and new information are brought forward. At these meetings, the decision to revise or update a model or forecast is

weighed against various costs and the amount of forecasting error. In a highly volatile area, the review should occur as frequently as every month or period.

Forecasting in the Future

In concluding an article on forecasting, it is appropriate that we make a prediction about the techniques that will be used in the short- and long-term future.

As we have already said, it is not too difficult to forecast the immediate future, since long-term trends do not change overnight. Many of the techniques described are only in the early stages of application, but still we expect most of the techniques that will be used in the next five years to be the ones discussed here, perhaps in extended form.

The costs of using these techniques will be reduced significantly; this will enhance their implementation. We expect that computer timesharing companies will offer access, at nominal cost, to input-output data banks, broken down into more business segments than are available today. The continuing declining trend in computer cost per computation, along with computational simplifications, will make techniques such as the Box-Jenkins method economically feasible, even for some inventory-control applications. Computer software packages for the statistical techniques and some general models will also become available at a nominal cost.

At the present time, most short-term forecasting uses only statistical methods, with little qualitative information. Where qualitative information is used, it is only used in an external way and is not directly incorporated into the computational routine. We predict a change to total forecasting systems, where several techniques are tied together, along with a systematic handling of qualitative information.

Econometric models will be utilized more extensively in the next five years, with most large companies developing and refining econometric models of their major businesses. Marketing simulation models for new products will also be developed for the larger-volume products, with tracking systems for updating the models and their parameters. Heuristic programming will provide a means of refining forecasting models.

While some companies have already developed their own input-output models in tandem with the government input-output data and statistical projections, it will be another five to ten years before input-output models are effectively used by most major corporations.

Within five years, however, we shall see extensive use of person-machine systems, where statistical, causal, and econometric models are programmed on computers, and people interacting frequently. As we gain confidence in such systems, so that there is less exception reporting, human intervention will decrease. Basically, computerized models will do the sophisticated computations, and people will serve more as generators of ideas and developers of systems. For example, we will study market dynamics and establish more complex relationships between the factor being forecast and those of the forecasting system.

Further out, consumer simulation models will become commonplace. The models will predict the behavior of consumers and forecast their reactions to various marketing strategies such as pricing, promotions, new product introductions, and competitive actions. Probabilistic models will be used frequently in the forecasting process.

Finally, most computerized forecasting will relate to the analytical techniques described in this article. Computer applications will be mostly in established and stable product businesses. Although the forecasting techniques have thus far been used primarily for sales forecasting, they will be applied increasingly to forecasting margins, capital expenditures, and other important factors. This will free the forecaster to spend most of the time forecasting sales and profits of new products. Doubtless, new analytical techniques will be developed for new-product forecasting, but there will be a continuing problem, for at least 10 to 20 years and probably much longer, in accurately forecasting various new-product factors, such as sales, profitability, and length of life cycle.

Final Word --- With an understanding of the basic features and limitations of the techniques, the decision maker can help the forecaster formulate the forecasting problem properly and can therefore have more confidence in the forecasts provided and use them more effectively. The forecaster, in turn, must blend the techniques with the knowledge and experience of the managers.