Journal of Forecasting
J. Forecast. **17**, 327–346 (1998)

Order Forecasts, Retail Sales, and the Marketing Mix for Consumer Durables

DOMINIQUE M. HANSSENS*
Anderson Graduate School of Management at UCLA, USA

ABSTRACT

The paper examines the problem of forecasting ongoing factory orders and monitoring retail demand, with specific reference to high-technology consumer durables. We present evidence of the managerial importance of the problem and, using a case study of a computer peripheral manufacturer, we describe how different data sources and models can be used to increase prediction accuracy. First we examine the order placement and retail demand process using extrapolative methods that focus on identifying short- versus long-run movements in orders. We then introduce marketing-mix data for improved retail demand tracking and forecasting, and we propose the use of conjoint measurement data to simulate a product's utility over time and include that information in the demand model. Lastly, we describe the forecasting and marketing planning use of these models and discuss their implications. © 1998 John Wiley & Sons, Ltd.

KEY WORDS market response models; order forecasting; cointegration; persistence; marketing-mix effects

In April 1996, Apple Computer reported a $740 million second-quarter loss, due largely to a write-down of $388 million of inventories that were no longer state of the art (Pitta, 1996). While it is relatively common for companies to report excess inventories from time to time, the sheer magnitude of Apple's problem highlights the degree to which companies' financial performance—especially in the high-technology sector—can be affected by overforecasting. By contrast, on many occasions in recent years, other high-technology companies have reported supply shortages—notably in DRAM memory chips—that drive up prices and put dealers on supply allocation schemes (e.g. Clark, 1994). Both examples highlight the strategic importance of accurate demand forecasting for production planning in high-technology consumer durables.

In many product categories, such as frequently purchased consumer goods, sales monitoring and forecasting is usually facilitated by reports, graphs, and statistical models based on high-quality retail scanner databases. Combined with the ability to conduct marketing experiments, e.g. through split-cable TV technology and in-store shelf space experiments, the scanner databases and models enable manufacturers to assess demand and adjust their marketing mix quickly

* Correspondence to: Dominique M. Hanssens, Anderson Graduate School of Management at UCLA, 101 Westwood Plaza, Box 951481, Los Angeles, CA 90095-1481, USA. E-mail: Dominique.Hanssens@Anderson.UCLA.edu.

CCC 0277–6693/98/030327–20$17.50

in response to changing market conditions. Since consumer sales for these categories are generally stable—except for the effects of temporary price promotions—demand forecasts can be expected to be accurate and some retailers have started to implement a 'just-in-time' delivery of product from their major suppliers (e.g. Fuller *et al*., 1993).

The scenario in high-tech consumer durables such as personal computers and their accessories is different. Due to inherent market volatility, bullwhip effects in the supply chain, and data limitations, manufacturers find it difficult to produce the quantities the market demands and to manage their marketing mix for profit maximization, as opposed to last-minute adjustment of forecasting mistakes. Let us briefly discuss these sources of difficulty.

First, *inherent market volatility* is caused by diffusion of innovation and rapid technological obsolescence: sales typically do not fluctuate around stable means or linear trends, but may evolve exponentially in some periods, then suddenly mature and decline quickly. While diffusion of innovation and technological obsolescence are generally acknowledged by industry participants, their increasing speed makes markets more volatile, to the point where many product life cycles are now routinely expressed in months rather than years.

Second, the *bullwhip effect* is a channel behavioural phenomenon that causes demand fluctuations to be more pronounced upstream than downstream in the value chain. Lee, Padmanabhan, and Whang (1997) propose that each of four sources causes a demand distortion (variance amplification) as a result of rational actions by the channel members: the processing of demand signals (distributors overstock because they overestimate retail demand surges), rationing games (dealers ordering more than they need in the face of supply restrictions), order batching (combining orders to take advantage of quantity discounts), and price variations. The combined result of these four forces is that, say, a 10% hike in consumer demand for a product may be amplified through the channels of distribution, and eventually cause a much larger swing in factory orders, perhaps 40%. If the manufacturer had perfect vision over the entire value chain, he could discount these excessive bullwhip variations in factory orders. However, *accurate retail demand data* are typically hard to come by, as the industries lack the consolidated scanning services and instant demand feedback that are typical of the packaged goods sectors in advanced economies.

For these reasons, *building good forecasting and marketing mix models for retail orders and consumer sales is a high-value proposition to manufacturers of high-technology durables*. This paper will describe a systematic approach to *ongoing* forecasting and marketing-mix adjustment based on different sources of data. Unlike existing literature that addresses the question of pre-launch market sizing for durables (e.g. Urban, Hauser, and Roberts, 1990), we focus on the scenario where monthly time series of orders and other key variables are available to the manufacturer. Every time new market information becomes available, the manufacturer needs to decide whether or not to use the new data and change his current order forecast and production plans for the medium term (typically two to six months out). In addition, the manufacturer should decide whether or not marketing mix intervention is called for, based on the same new market information.

The focus of the paper is on developing model-based forecasts that should be combined with managerial judgement in order to produce final forecasts (e.g. Blattberg and Hoch, 1990). Our approach is layered by *data availability* to the manufacturer, i.e. we ask how much of the problem can be solved by the readily accessible data, and what would the incremental contribution be of additional, more expensive information? From a *methods* perspective, we use statistical models from econometrics, time-series analysis and conjoint measurement. Though some of these methods are relatively new to the marketing literature, we will keep technical discussion to a minimum and refer the interested reader to more detailed literature. Finally, the *application* that is used as a running example in the paper is based on a global manufacturer of high-technology products for the business and consumer sectors.

CASE DESCRIPTION AND DATA COST HIERARCHY

The case study involves a durable product in the category of personal computing accessories. Launched in the late 1980s, this technology gained rapid customer acceptance, first in the office equipment market, and eventually in home computing. The data are taken from the market leader in the category, a global manufacturer with market shares of 50–80%, depending on the world region. The product retails for several hundreds of dollars and has a typical marketing mix consisting of periodic product technology improvements, distribution efforts, and advertising using both print and electronic media. Furthermore, the manufacturer has some control over dealer margins via suggested retail prices and dealer discounts. However, rapidly changing market and competitive conditions typically give rise to so-called *street prices* that may be substantially different from the manufacturer's suggested retail price (MSRP).

The channel structure for this category has either one or two intermediaries between manufacturer and end user. By focusing on the top twenty retailers in the domestic market, this manufacturer is able to track about 90% of total consumer demand for its product. We will use this sample of major retailers to approximate demand for the product, so that 'retail sales' or 'sell through' and 'consumer demand' are considered equivalent for our purposes. Information on inventories or on intermediate order levels in the channel is not available.

The *data cost hierarchy* is an overview of data availability for forecasting and marketing planning, in increasing function of 'cost', by which we mean ease of collection, acquisition cost, and timeliness of delivery. Most manufacturers face a data cost structure similar to the one described here. The easiest data to obtain are weekly or monthly orders received by the channel (distributors or major retailers) and list price, as they are taken directly from internal accounting records. Monthly consumer demand (retail sales) and street prices are more costly, as they are purchased from an audit service of major retailers and take several weeks to be delivered. The next level is internal marketing-mix data, especially expenditures on media advertising. This information usually needs to be compiled from various sources, possibly involving the advertising agency. However, a proactive effort to develop a company *marketing data warehouse* can significantly improve on-line access to marketing spending data and reduce its long-term cost. Finally, the most costly information is the periodic tracking of customer preferences for the company's product *vis-à-vis* its competitors. Indeed, in technology-intensive sectors, offerings and market conditions change quickly so that consumer preference 'snapshot' surveys do not retain their validity for very long. As a result, few companies invest in expensive primary data collection on the determinants of product value in the consumer's mind. Instead, they typically rely on the free-market mechanism that establishes street prices to sort out differences in consumer preferences among the available products. While markets eventually clear, they may do so at considerable expense to the manufacturer: if he underestimates his product's value in the market, he may experience product shortages or forgone profit margins. If he overestimates his product's market value, excess inventories develop that are costly to dispose of and that may depress future order levels.

The paper is structured as follows. First, we describe the managerial scenarios that motivate the development of our order forecasting models. The technical section of the paper then starts with an examination of the over-time behaviour of orders, consumer sales, and their connection. We use that information to develop gradually more elaborate forecasting models of orders, following the data cost hierarchy, and compare their performance. Next, we focus on the determinants of consumer sales and examine to what extent the manufacturer has direct influence on retail demand. Lastly, we draw conclusions for the manufacturer's combined use of order and consumer sales data in forecasting and marketing planning.

ORDER FORECASTING AND THE MARKETING MIX

The methods we describe are motivated by managerial considerations. Once a product is launched and order data come in on a weekly or monthly basis, marketing and production executives must periodically *update existing forecasts* and *adjust the marketing mix* in the light of new information. In order to optimally analyse this new information, it is critically important that the company install a *tracking system* for the storage and easy retrieval of these data. Such a system is sometimes referred to as a *data warehouse* and we will assume in this paper that a warehouse exists, at least for the key market performance and marketing-mix variables.

Suppose a company is marketing a high-tech consumer durable at a monthly production capacity of 120,000 units. The manufacturer's factory price is $300, the suggested retail price is $500 and the street price is $399. Advertising spending is averaging $800,000 per month. Within four weeks' time, the following market events happen: order levels drop from 110,000 to 80,000, retail sales are stable, and a major competitor matches the technical performance of the company's product at an MSRP of $450. What changes, if any, should the company make, either to its production plans and/or to its marketing mix?

Equipped with instant feedback such as in the scenario above, decision makers are in a good position to interpret the market signals correctly and take appropriate action. The interpretation of market signals is greatly facilitated by the use of *statistical models on longitudinal market and marketing-mix data*. For example, a manager needs to know if the observed drop in orders is of a temporary or a permanent nature, yet she may not have the luxury to wait for three more months to make that distinction. By the same token, the competitive product feature match should be related to the company's own expected retail sales in order to decide whether or not to lower the factory price and MSRP.

Ideally, all major 'what if' questions such as the ones above are submitted to a market simulator that interprets the latest movements in the data warehouse. We will describe such a market simulator in a stepwise fashion, where each stage involves more expensive data collection and therefore has to demonstrate superior performance over the previous step (the benchmark). In all cases should the simulator help the manager make the following two critical decisions in the face of new market data: should the existing forecast be changed (and, if yes, by how much) and, should the marketing mix be adjusted (and, if yes, how so)?

The question of changing a forecast or a marketing mix is related to whether the observed movements in orders and retail sales are of a permanent (long-term) versus a temporary (short-term) nature. *Temporary* movements should be understood for short-term forecasting purposes, but probably do not require marketing-mix intervention, because market conditions have fundamentally not changed. In our example, the observed drop in orders could be a temporary retail inventory adjustment that will be restored in the next month, given a stable consumer demand level. On the other hand, *long-term* or *permanent* movements should be recognized both for forecasting and for marketing-mix intervention. If the 27% drop in orders in the example is permanent, then the long-run order forecast should eventually be reduced by the same amount, *or* the marketing mix should be adjusted so that the original forecast is still achievable. We therefore begin the statistical analysis with an examination of the longitudinal behaviour of orders and retail sales.

THE LONG-TERM BEHAVIOUR OF ORDERS AND CONSUMER SALES

From a statistical perspective, the assessment of the over-time behaviour of the key variables of interest helps determine the choice of forecasting models. Indeed, depending on the nature of the stochastic processes that underlie movements in orders and their relationship to sales, different forecasting models should be used. If the stochastic process is *stationary* (mean or trend reverting), unexpected events such as sudden surges in orders have *temporary* effects which can be important in the short run, but which should not be cause for altering long-run forecasts. In contrast, *non-stationary* processes contain movements that signal *permanent* departures from previous levels and therefore necessitate forecast updating. Following a discussion on permanent versus temporary movements, we first examine whether orders are stationary or non-stationary. Then, we investigate whether non-stationary movements in orders (if any) are related to changes in consumer sales. Finally, in the light of the bullwhip effect, we estimate the propagation effect on orders of an unexpected change in consumer demand.

Figure 1 shows the history of orders for our product over its 44-month life to date, as well as consumer sales for the same time period. Combining visual inspection of the data with some specific statistics to follow, we draw inferences on these three questions and discuss their implications.

Both orders and retail sales follow a non-stationary pattern over time

The traditional test for the stationarity of time series data is the so-called *unit-root* test. This examines if the observed movements in the data are temporary fluctuations around a fixed mean or trend (stationarity), or if they have random-walk components that permanently depart from previous levels (non-stationarity). While several unit-root tests have been proposed in the literature, a popular version is the augmented Dickey–Fuller (1979) test equation:

$$\triangle S_t = a_0 + bS_{t-1} + a_1 \triangle S_{t-1} + \cdots + a_m \triangle S_{t-m} + u_t \tag{1}$$

where Δ is the difference operator and S_t is the time series of interest (for example, sales). The t-statistic of b is compared with the critical values in Dickey–Fuller (1979), and the unit-root null hypothesis is rejected if the obtained value is smaller than the critical value. Indeed, substituting $b = 0$ in equation (1) introduces a random-walk component in the model, whereas $-1 < b < 0$ implies a mean-reverting or stationary process. A more detailed discussion with application to marketing data may be found in Dekimpe and Hanssens (1995a).

Table I summarizes the unit-root testing results for orders and sales.[1] In both cases, the presence of a unit root cannot be rejected at $p < 0.05$. We conclude that there is non-stationarity

[1] The time-series analyses were performed using EViews 2.0 econometric software developed by Quantitative Micro Software, Irvine, California and Forecast Pro 3.00 forecasting software developed by Business Forecast Systems, Belmont, Massachusetts.

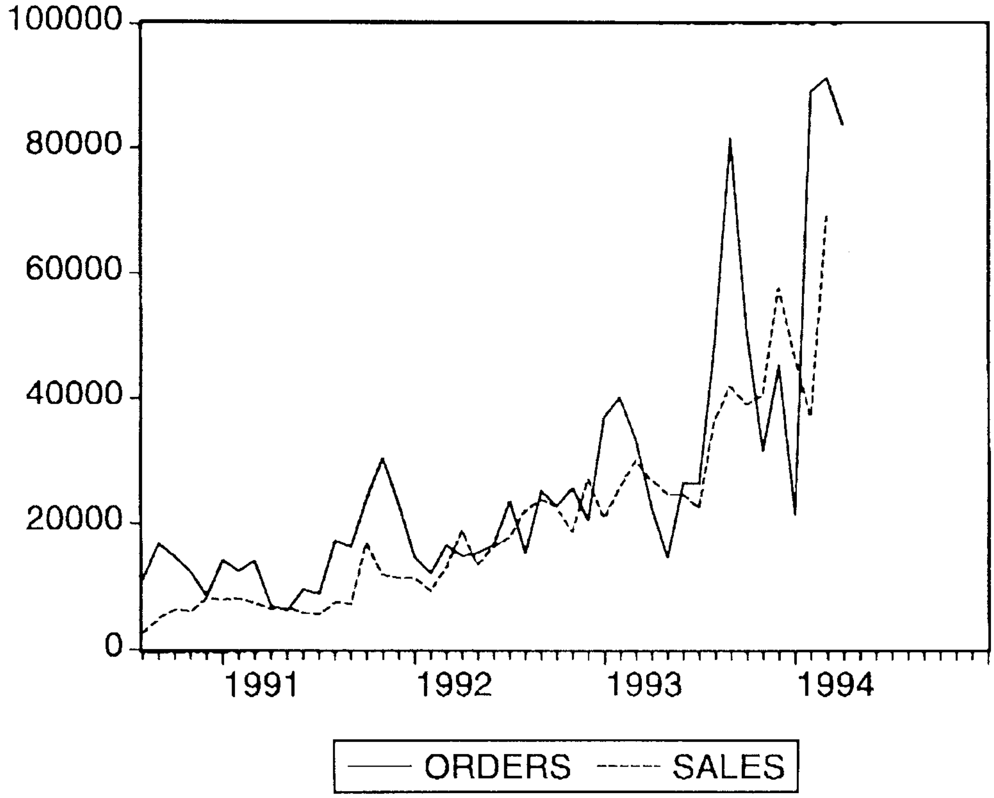

Figure 1. Orders and consumer sales

Table I. Unit root tests

	m	b	t	Unit root present?
Orders	6	−0.44	−2.48	Yes
Consumer sales	0	−0.17	−1.59	Yes

These results are based on the augmented Dickey–Fuller (ADF) test (equation (1)), where

m = number of augmented terms reflecting temporary sales fluctuations; we used conventional significance tests on the a_i to determine the cut-off point
b = parameter estimate of the lagged dependent variable
t = t statistic associated with the lagged dependent variable, to be compared against the 5% critical value of –2.89. The unit-root null hypothesis is rejected if the computed t-statistic is smaller than this value.

in both orders and sales: the market environment is not stable, but rather evolves over time. Evolution has been found to be the rule rather than the exception in a meta-analysis of market behaviour (Dekimpe and Hanssens, 1995b). It implies that any movement in orders or sales *can* signal a permanent departure from previously observed levels. From a managerial perspective, non-stationarity makes the forecasting task more difficult, as we cannot assume that sales will return to normal (i.e. be stable) after an unexpected movement. At the same time, the finding adds to the managerial relevance of data-driven forecasting, as recent research suggests that, the more unpredictable the environment, the higher contribution one can expect from data-driven decision support systems over judgemental approaches (Hoch and Schkade, 1996). The subsequent forecasting model development should recognize the non-stationary character of orders and consumer sales.

There is a long-run equilibrium relationship between orders and retail sales

Even though orders and sales each evolve over time, their movements are not necessarily independent in the long run. For example, if changing consumer preferences are pushing sales upward, the resulting decline in retail inventories would eventually lead to a surge in factory orders. More formally, the orders and sales time series are expected to be linked to each other in the long run, a condition called cointegration.

We refer to Enders (1995) for a complete treatment of cointegration modelling, and to Powers *et al.* (1991) and Franses (1994) for applications in management. The essence of cointegration testing is to estimate one or more equilibrium regressions between the non-stationary time series—in this case orders (O_t) and sales (S_t). Then, we test for the stationarity of the residuals of these equilibrium regressions. Engle and Granger (1987) use an ordinary least-squares approach, which is easy to interpret but is subject to small-sample bias. Johansen (1988) uses a maximum-likelihood method which tests for the number of equilibrium regressions and removes the small-sample bias. Both methods led to the same conclusion for cointegration between orders and sales: there is one equilibrium regression, estimated as follows using Johansen's method:

$$O_t = 4170 + \underset{(0.128)}{1.047} S_t + e_t \tag{2}$$

where parameter standard errors are shown in parentheses, and the residuals e_t are stationary. Therefore, even though orders and retail sales each evolve over time, we can find a linear combination between them that behaves as a stationary variable.[2]

The statistical finding of a cointegrating relationship from sales to orders can be explained intuitively. Retailers base their ordering patterns on fluctuations in consumer demand, which they experience first-hand. While they can make mistakes in gauging or anticipating period-by-period consumer sales, over the long run they are able to adjust their orders such that the (scaled) difference between actual sales and orders fluctuates around a zero mean. By contrast, the reverse equilibrium regression, while statistically estimable, is ruled out *a priori* because end users only buy for themselves and because orders need to result in shipments and retail availability before they become relevant.

The cointegration result *highlights the importance of collecting retail sales data*, which can be difficult and expensive for manufacturers. Indeed, the retail sales data are essential for removing the non-stationarity in the manufacturer's orders data. Following Engle and Granger (1987), cointegration implies that orders follow an *error-correction process* with consumer sales: if the current order level deviates from the equilibrium level implied by current demand, then future orders will correct that disequilibrium. For the manager in charge of forecasting, that implies that

[2] In Johansen's method, this equilibrium regression is estimated jointly with the error correction model discussed in the next section. We conducted tests on the number of cointegrating vectors and the stability of the parameter estimates with and without deterministic trend, and with different lag specifications in the error-correction mode. While the response parameters vary somewhat, the results are robust and not meaningfully different from the reported estimates.

a portion of the observed fluctuation in orders is due to equilibrium adjustment, while another portion reflects response to fundamental demand shifts.

Most of an unexpected change in retail sales is persistent, but most of an observed order shock is not

While the sample variance in orders is 20% higher than that of sales, it is more insightful to study the bullwhip effect by investigating the net effect of an unexpected shock in demand (positive or negative) on orders. In a non-stationary environment, the appropriate statistic to examine is the *persistence* of orders and sales. Persistence is the fraction or multiplier of unexpected short-run movement in a variable that permanently affects the future time path of that variable (univariate persistence) or other variables (multivariate persistence).

Since we have already established the existence of a long-run equilibrium between orders and sales, we should incorporate that relationship when examining the over-time response of orders to unexpected change in consumer demand. As mentioned earlier, the vector-autoregressive (VAR) model with error correction is the appropriate procedure for estimating both univariate and multivariate persistence (Engle and Granger, 1987). The two-equation error-correction model on changes in orders and sales is:

$$\begin{aligned} \triangle S_t &= a_1 + b_1(L)\triangle S_{t-1} + c_1(L)\triangle O_{t-l} + d_1 e_{t-1} + u_t \\ \triangle O_t &= a_2 + b_2(L)\triangle O_{t-1} + c_2(L)\triangle S_{t-1} + d_2 e_{t-1} + v_t \end{aligned} \tag{3}$$

where e_{t-1} is the lagged equilibrium error term estimated in equation (2) and L is the conventional lag operator notation. Intuitively, this model explains temporary changes in orders and sales as a result of previous order fluctuations, sales fluctuations, and deviations from the long-run order-sales equilibrium. We derive the maximum-lag lengths on the parameters vectors $b_i(L)$ and $c_i(L)$ ($i = 1, 2$) by successively specifying longer lags, to the point where these parameters are not statistically significant, in this case lag = 1, and we estimate that system. Then, we simulate a one-standard deviation unexpected change in sales and recursively calculate its dynamic effects, sometimes called impulse response weights, on sales and orders.[3] The long-run dynamic effect, or persistence value, is obtained where impulse response weights converge. A technical overview of persistence modelling may be found in Enders (1995), and an application to measuring the persistence effects of marketing spending on sales is described in Dekimpe and Hanssens (1995a).

Figure 2(a) shows the persistence graphs for consumer sales and factory orders. The graphs visualize the dynamics of order and sales behaviour very well. First, let us focus on order shocks (top graph). We observe that an order shock of one standard deviation (about 11600 units) does not have a strong long-run effect on itself. Most of the unexpected movement dissipates and only about 1400 units (or 12%) is persistent. This implies that the company should be careful not to overinterpret short-run order fluctuations as signals of future order patterns.

On the other hand, let us examine the behaviour of an unexpected change in consumer sales (middle graph). Here we see that a one-standard deviation shock (about 5600 units) has high staying power: the long-run effect on itself is about 3800 units, implying that the sales persistence

[3] On a technical note, the impulse response estimation is performed under the assumption that exogenous shocks in consumer demand affect orders in the same period, but not vice versa. This temporal ordering allows us to start the chain reaction simulation that generates the reported persistence levels.

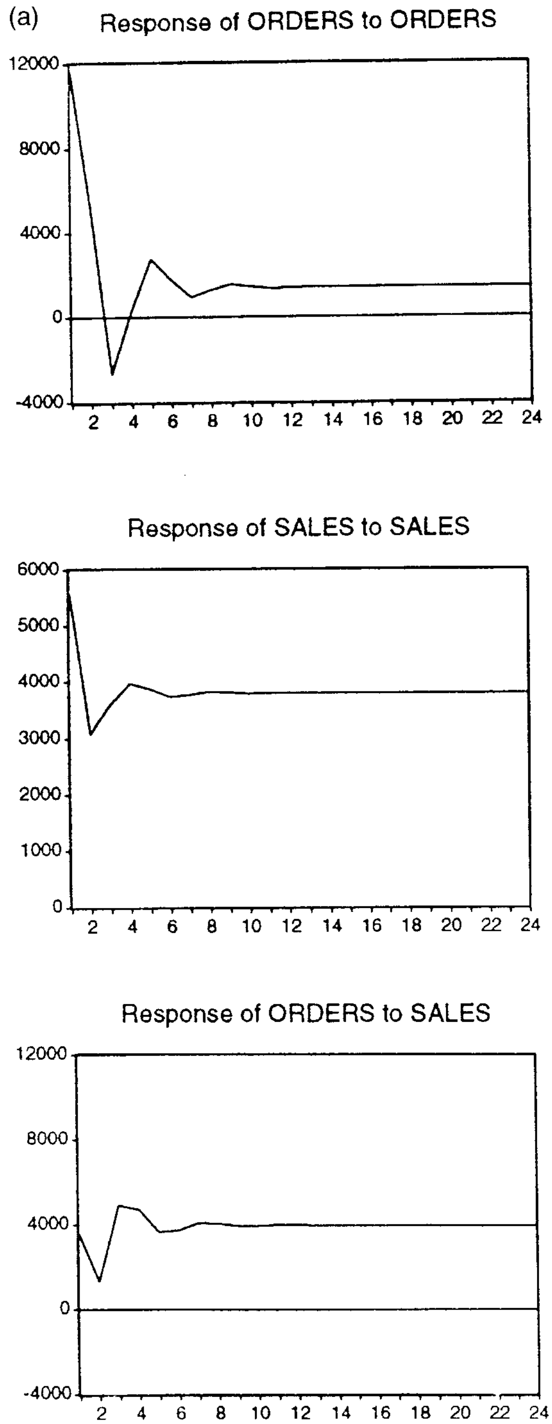

Figure 2. Persistence graphs. (a) Responses to shocks of one standard deviation

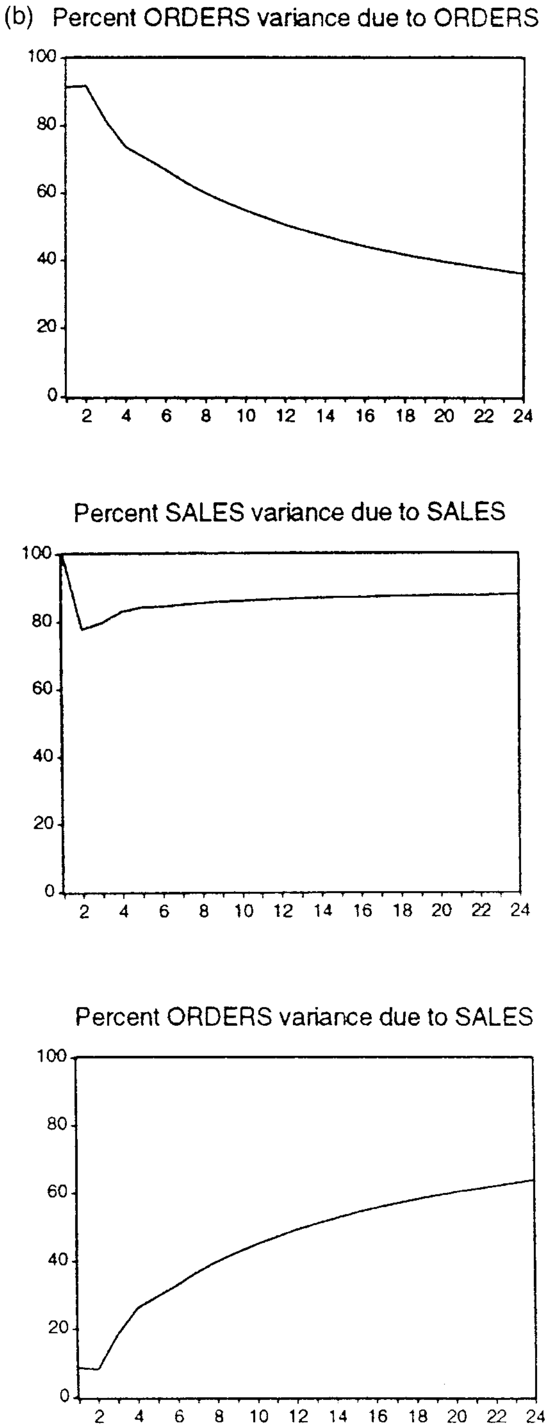

Figure 2. Persistence graphs. (b) variance decompositions of shocks

is about 68%. Therefore, short-run movements in consumer demand are more powerful indicators of long-run changes than those in orders.

Finally, and most importantly, we examine the cross-effect between the two, or how a shock in retail demand affects orders in the long run (bottom graph). We see that the portion of the order shock that is attributable to sales is strong and permanent: the short-run effect is about 3600 units, the long-run effect is 3900 units, so the persistence is over 100%. This response pattern is even clearer when we look at the relative sources of order shocks over time (Figure 2(b)): the percentage of an unexpected change in orders that is attributable to a change in consumer sales gradually increases from 9% to about 64%.[4]

These results provide good empirical evidence for Lee *et al.*'s (1997) thesis that demand shocks propagate through the supply chain, even when the chain has only two elements. Only a fraction of sudden changes in orders can be expected to have long-lasting effects due to permanent consumer demand changes. An important strategic conclusion for the manufacturer is that movements in retail sales should be captured and analysed immediately, as they contain useful information about the long-run outlook of consumer demand and orders for the product.

FORECASTING ORDERS

Based on the evidence of evolutionary behaviour in orders and sales and their connection, we now address specific order forecasting models. We discuss, in turn, setting a benchmark model, including retail sales information, calculating the economic impact of using larger databases, and assessing managers' control over orders via the marketing mix.

Extrapolative models of retailer orders are a useful *benchmark* for forecasting from the manufacturer's perspective. Indeed, data on orders are easy and fast to collect, as they are obtained directly from accounting and invoicing records. From a forecasting perspective, models that consume more data than just the history of orders should outperform the benchmark models in order to have economic value to the firm.

The benchmark model is the optimal forecast of future orders, given their past pattern. A Box–Jenkins model on orders filters out the non-stationarity and the autoregressive and moving-average patterns in the data, leaving only unpredictable random error. Using Box and Jenkins' parsimony principle, the univariate order model is specified as ARIMA(0,1,1) and estimated as

$$\begin{aligned}(1-L)O_t = \underset{(1353)}{1638} + (1 - \underset{(0.143)}{0.397}L)a_t\end{aligned} \tag{4}$$

where a_t are white-noise residuals. This process explains about 56% of the monthly order variation, with a mean average error (MAPE) of 36%. Note that the forecasting performance is much improved when aggregating monthly to quarterly intervals. For example, in an attempt to predict what turned out to be the most volatile performance period for this product, the beginning of 1994, the Box–Jenkins order model, estimated with data up to December 1993 had a monthly forecast MAPE of 66%, but a trimester forecast MAPE of 38%. Considering that the

[4] The reverse dynamics are conceptually less relevant. For completeness sake, most of the shock dynamics in consumer sales are due to previous sales fluctuations, with order shocks having only a minor long-run effect on sales (about 12%). This pattern is to be expected so long as no serious product shortages occur, which would create pent-up demand so that consumer sales fluctuate drastically with prior orders and shipments.

end of 1993 witnessed a dramatic fall in orders, this result is encouraging. On the other hand, the univariate models may produce volatile forecasts as new data become available, since the observed order variance is inflated by the bullwhip effect discussed earlier.

How does knowledge of consumer demand improve upon this performance? In *the long run*, we know that retailers adjust their order levels to the consumer demand of the product. This equilibrium-seeking behaviour can be represented in a model of short-run behaviour by estimating an *error-correction* model (Engle and Granger, 1987):

$$(1 - L)O_t = a + b(1 - L)O_{t-1} + c(1 - L)S_{t-1} + de_{t-1} + u_t \tag{5}$$

where e_t is the scaled residual of the equilibrium regression estimated in the previous section. Thus if orders are running abnormally high relative to overall demand, one can expect a short-run dip in orders that helps restore the equilibrium, and vice versa. In operational terms, one can think of the equilibrium error variable as a proxy for long-term inventory excess or shortage.

Short-run adjustments in orders could be based on *current* demand levels, or they may *anticipate* future retail sales, or they may *react* to recent historical demand fluctuations. Since retail sales are highly autocorrelated (the sample first-order autocorrelation is about 0.80), it is difficult to disentangle these hypotheses empirically. However, we obtain a good indication of their relative strength by performing various time-series regressions of changes in orders against current, future, and lagged consumer sales changes. The results reveal that one-period-ahead (i.e. anticipated) sales are the strongest correlate of order changes, though only by a relatively small margin. The finding lends support to the hypothesis that retailers are *forward looking* in their order-adjustment process.

For comparing the order error-correction process with the benchmark, we use a one-period lagged adjustment version of the model, which does not give it an unfair data advantage over the Box–Jenkins equation. The full-sample parameter estimates are

$$\begin{aligned}(1 - L)O_t = {} & \underset{(2014)}{1890} + \underset{(0.204)}{0.407}(1 - L)O_{t-1} - \underset{(0.394)}{1.060}(1 - L)S_{t-1} - \underset{(0.250)}{0.973}e_{t-1} + u_t\end{aligned} \tag{6}$$

All parameters are statistically significant at $p < 0.05$, except for the intercept, and the Ljung–Box Q tests indicate white-noise residuals. Compared to the benchmark univariate model, the error-correction model explains 66% of the in-sample variance and has a MAPE of 17%, significant improvements over the benchmark (see Figure 3). Its out-of-sample predictive accuracy for the first trimester of 1994 is 14%, representing an improvement of 63% over the benchmark. In addition, this process produces less volatile forecasts, due to the smoothing influence of the equilibrium-correction term.

We conclude that there is ample *statistical* evidence favouring the collection and use of retail sales data in order forecasting. In order to assess the *economic* impact of these data, consider the costs of both over- and underpredicting orders. Overforecasting results in unwanted inventories, underforecasting results in lost sales and dealer game playing resulting from allocation rules. Assuming a profit margin of 50% in forgone sales, carrying costs of 10% on excess inventory, a wholesale product price of \$250, base level orders of 65,000 per month and an equal occurrence of under- and overprediction with the trimester MAPEs reported earlier, the monthly estimated cost of forecasting mistakes is about \$1.85 million for the extrapolative model and about \$730,000 for the model with consumer sales data. Thus, from a strict forecasting perspective, the

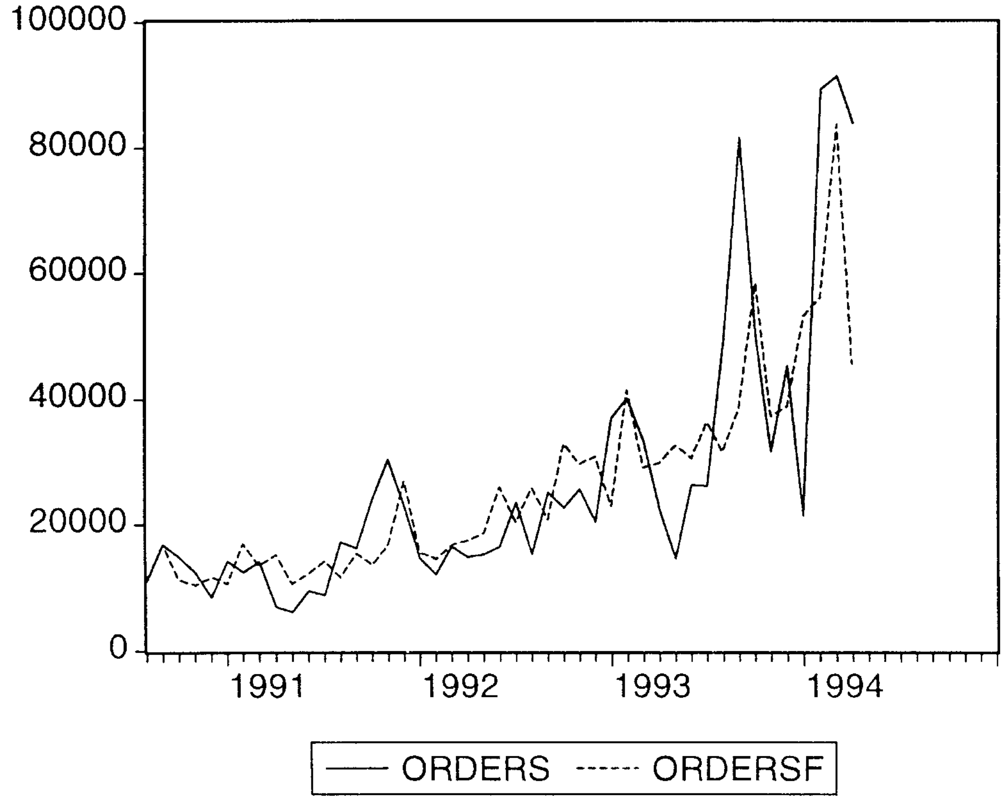

Figure 3. Fit of orders error-correction model

incremental value of collecting and using consumer sales data could be over $1 million per month.

Last, but not least, we investigate to what extent the manufacturer has 'push control' over its orders, i.e. to what extent he can influence order levels above and beyond consumer demand forces. The question is strategically important for the following reason: if push control is strong, a manufacturer can make up for costly forecasting mistakes by quickly adjusting the marketing mix to the trade. For example, he could change distributor margins in an attempt to push the product through the channel when consumer demand is soft. Inasmuch as such efforts are effective, we should observe a *statistically significant marketing effect on orders* above and beyond the consumer sales effect.

We estimate the 'distributor push' effect with a marketing mix model on orders, i.e.

$$\text{orders} = \text{f}(\text{previous orders, consumer sales, manufacturer's marketing mix})$$

The manufacturer's marketing mix consists of variables that can be hypothesized to stimulate order levels while controlling for consumer sales. In our application, that includes the dealer profit margin, supply allocations, and advertising. All else equal, higher profit margins encourage dealers to carry the company's product over the competition. Supply allocations, coded as a dummy variable with value 1 when they are in place, motivate dealers to order more than they need and try to circumvent the supply shortage. Advertising may be interpreted by the trade as a signal of the manufacturer's confidence in the product and their willingness to stimulate retail demand.

Various econometric specifications of the order response model were estimated. They included using differences to account for non-stationarity, logarithms for modelling concave or convex response, and lagged regressors for delayed effects. *None* of the marketing-mix parameters had acceptable levels of significance, nor did the overall goodness of fit improve with the inclusion of marketing-mix variables. We conclude that the manufacturer's push effect on orders, if present at all, is weak relative to the consumer demand influence. This finding also highlights the importance of closely monitoring demand at the consumer level, which is the focus of our next section.

MONITORING CONSUMER SALES

Standard marketing theory posits that end-user demand arises from a matching of customer needs by companies and their products. At the aggregate level, such need satisfaction is measured by the strength of market response to the company's marketing mix, i.e. making the right product available to the right audience at the right price, and effectively communicating that value proposition. The measurement instrument for that purpose is the marketing-mix model, generally an econometric model that relates variations in consumer demand to the company's marketing mix and environmental factors (Hanssens, Parsons, and Schultz, 1990).

Using consumer sales data for the top 20 retailers, we estimate consumers' response to the marketing mix as follows:

consumer sales = f[street price, print advertising, TV advertising, distribution, product value]

All but one of the hypothesized marketing drivers are commonly used in the market response literature and some empirical generalizations exist about their relative effectiveness. For example:

- Street price measures the demand curve, which is expected to be fairly steep in a competitive market for consumer durables. Typical price elasticities are in the -2.5 range (e.g. Tellis, 1988).
- The effects of advertising are usually much smaller, with elasticities in the 0.1 to 0.2 range (e.g. Tellis and Weiss, 1995). That does not imply that advertising is ineffective in stimulating demand, as the variation in companies' advertising spending can be much larger than the typical price variation. We will make the distinction between two media effects, TV and print. Prior literature has indicated strong differences among media elasticities (e.g. Montgomery and Silk, 1972; Dekimpe and Hanssens, 1995a).
- Distribution would ideally be measured by number or percentage of retail outlets carrying the product, which often has a sales elasticity around unity or above (Reibstein and Farris, 1995). Such data are not available, but we do know the point in time where the company expanded its product availability from 'specialized outlets only' to general distribution, i.e. including general merchandise stores. A dummy variable will pick up any sales response effects.

One hypothesized driver of performance, product value, has not been formally modelled and incorporated into market response modelling. It is difficult to quantify the inherent value of a product or service relative to competition, at an aggregate level. In many markets, product features and characteristics change slowly and infrequently, if at all. In such cases it is usually sufficient to use dummy variables to capture discrete change points. On the other hand, a distinguishing aspect of high-technology markets is that their products' characteristics change continuously and frequently. For example, the maximum storage capacity of computer disk drives is increased by one supplier or another virtually every month, with competitors either catching up quickly or being forced to cut prices. As a result, disk drive capacities grew from 840 Megabytes to 3 Gigabytes in 1996.

We are proposing here to use a dynamic simulator of product utility, based on conjoint measurement, to represent the continuous changes in product characteristics and their likely impact on consumer preference. Conjoint analysis of consumers' trade-offs among product and service features is a well-known and successful market research technique (e.g. Wittink and Cattin, 1989). It is typically used for evaluating various product design alternatives and/or to forecast consumer preference for competing products in a laboratory setting. However, once a conjoint simulator is developed, it can be used to generate a *time series of aggregate-level product utilities* and therefore provide data points that are matched with the remainder of the marketing-mix database. To do so, we use the utility generator from the conjoint study, i.e.

$$\text{utility}_{it} = \Sigma_j b_j \text{ attribute}_{ijt} + u_{it}$$

where i refers to product, j to attribute, and t to time period, and retroactively change the product and its competitors' attributes to coincide with actual historical market changes. This procedure generates a history of consumer utility or preference shares for the high-tech product in the marketplace. It will be an accurate instrument so long as the attribute change data are accurate, e.g. the database must reflect that, in January 1994, competitor X matched a new technological feature our company had introduced in October 1993. This method also rests on the assumption that consumers' derived utility weights at the time of the interview are representative of their weights for the entire history under study. In the example, the conjoint instrument was administered in 1994, which offers the advantage that the most recent and therefore most advanced level of technology is incorporated into the design.

Figure 4 shows the history of these conjoint-inferred consumer preferences for the product, along with its price. From a customer value perspective, the product goes through phases of high, medium, and low value, relative to its price and competitive offerings. We expect these relative-quality variations to have a positive effect on consumer sales, while controlling for the other elements in the marketing mix.

As in the order response model, non-linearities in market response are modelled by logarithmic data transformation. In terms of response dynamics, many studies use a lagged sales term to capture carryover effects of advertising and other marketing investments on sales. The resulting Koyck-type response model, however, can be misspecified on non-stationary sales data, because it implies that sales return to a fixed mean after marketing or other shocks have occurred (see e.g. Dekimpe and Hanssens, 1995a). In our case of non-stationary sales, the response model should be estimated on differences, which are stationary. Consequently, the estimated model (in logarithms) is:

$$\begin{aligned}(1-L)S_t = b_0 + \ & b_1(1-L)PRICE_t + b_2(1-L)TVAD_t + b_3(1-L)PRINTAD_t \\ & + b_4 NEWDIST_t + b_5(1-L)UTILITY_t + v_t \end{aligned} \tag{7}$$

and its parameters are summarized in Table II. Unlike the findings in the order response model, we find that the *manufacturer's marketing mix has a statistically significant effect on consumer*

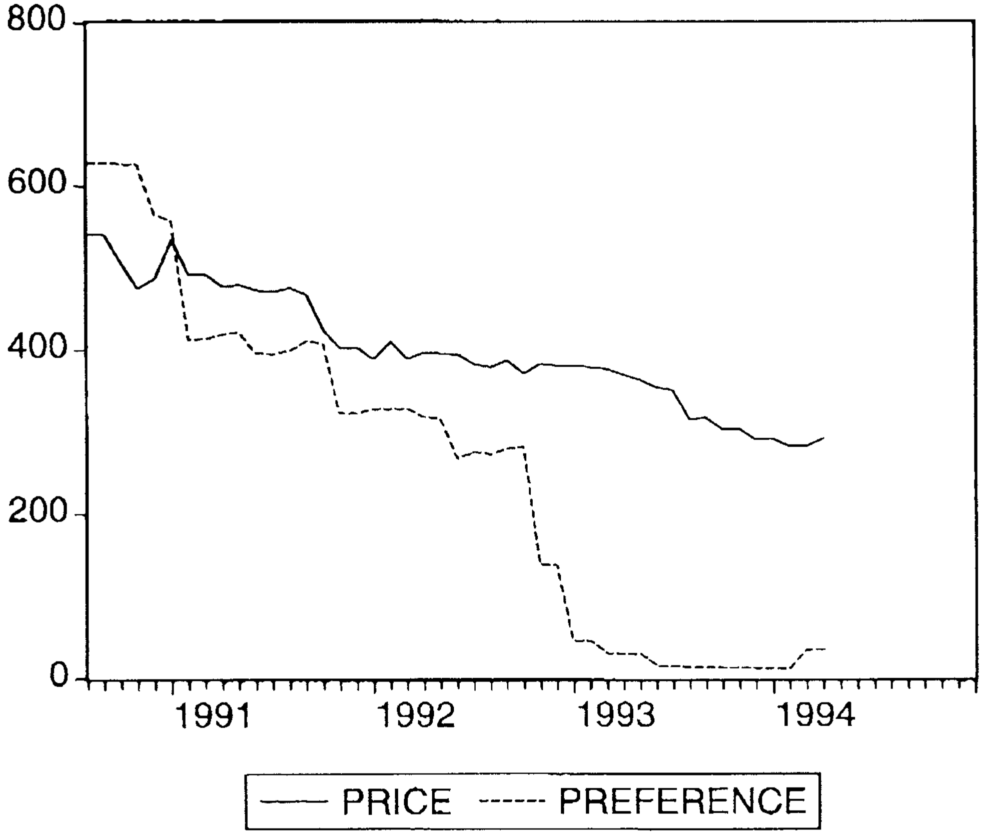

Figure 4. Evolution of price and preference share

Table II. Marketing-mix elasticities

Variable	Consumer sales elasticity
Street price	−2.117 (1.051)[b]
Distribution	−0.030 (0.084)
Print advertising	0.041 (0.033)[c]
TV advertising	−0.015 (0.017)
Preference share	0.400 (0.138)[a]
R^2 (on changes)	0.25
DW	2.30

One-tail significance levels are denoted as:
[a] ($p < 0.01$), [b] ($p < 0.05$) and [c] ($p < 0.1$).

sales, so the manufacturer has some 'pull power' in the market. The strongest drivers of consumer sales are product preference and street price, with elasticities of 0.40 and − 2.12, respectively. Print advertising is marginally significant, with elasticity 0.04. The other factors do not significantly contribute to movements in consumer sales.

Figure 5. Fit of marketing-mix model

The results on pricing and advertising are within the range of expectations, based on previous research. The *product value elasticity of around 0.4* is a new finding. It combines the insight that technological performance relative to the competition plays a significant role in consumer demand, with the fact that returns to product performance are decreasing. Finally, the combination of the orders and consumer demand models quantifies the long-run effects of marketing-mix intervention on orders: price, advertising, and product value have observable effects on consumer demand which, in turn, is related to long-run orders via the equilibrium relationship.

The marketing-mix model is a useful tool for monitoring consumer demand, adjusting sales forecasts, and deciding on marketing-mix intervention. While it fits the data well (see Figure 5), it will not necessarily outperform a simple extrapolative model of retail sales in forecast accuracy; for example, compared to a univariate Box–Jenkins model, equation (7) has only a slightly higher goodness of fit (R^2 on changes of 0.25 versus 0.21) and similar MAPE and MAD statistics. Its comparative advantage over univariate models lies in its ability to generate forecasts and scenarios that are conditional on firm decision variables such as price and marketing support. To illustrate the latter, suppose a competitor initiates a price cut on a comparable product. This action reduces the utility of the company's own product, and the marketing-mix model can be used to estimate the resulting demand effect. If the company wishes to preserve the existing order forecasts, the mix model can again be used to evaluate different reaction options, such as price matching or increasing the advertising budget. The most attractive of these marketing intervention options can then be weighed against the alternative of revising the order forecast, for example accepting a lower production quota while preserving profit margins.

CONCLUSIONS

The accurate prediction and management of factory orders is a complex and strategically important activity, especially in fast moving high-technology markets. The complexity arises from several sources. First, the most accessible data, factory orders, are noisy and do not lend themselves well to direct forecasting, due to the bullwhip effect. Second, more appropriate data such as consumer sales and marketing-mix movements are more difficult to bring together. On the other hand, as markets for high-technology durables become more crowded and demand volatility increases with shortening life cycles, the strategic importance of order forecasting and response-based marketing planning increases. Successful companies in this sector will invest in data, models, and information systems that execute smoothly the tasks described in this paper. The popular business literature has started to report on successful applications, for example at Compaq in the computer sector (McWilliams, 1995).

This paper has demonstrated, using an actual case study in high-technology durables, how a good marketing data warehouse and the use of some rigorous statistical methods can help resolve the order forecasting challenge. At the marketing *managerial* level, we quantified the bullwhip effect in orders and showed how the use of retail sales information significantly improves the accuracy of order forecasting in the medium run. We also estimated and compared the effect of the marketing mix on the manufacturer's order levels and on consumer sales. The results support the notion that the order forecasting and marketing functions in a company should be *integrated*, as marketing spending and pricing plans affect the order forecasts and vice versa. Depending on a manufacturer's timing between production, orders, and forecasted sales, such integration allows the company to use market-response models in a proactive way (e.g. order planning based on demand forecasts that are conditional on the intended marketing mix) or a reactive way (e.g. order adjustment after a competitor's move changes market conditions).

At the *methodological* level, we used new long-term time-series techniques to establish the longitudinal behaviour of orders and consumer sales and their connection. This allowed us to separate temporary versus permanent movements in orders and sales, and to use that information strategically. We also integrated a successful primary market research method, conjoint analysis, in an aggregate market response model. Following the principle of evolutionary model building (Urban and Karash, 1971), we conjecture that a logical next step in data warehouse and modelling sophistication is to augment the market performance and marketing mix database with longitudinal survey data on channel and consumer preferences.

While the use of one data setting facilitated focus in this research, it comes with a limitation on generalizability. The manufacturer in our case study had never developed a marketing data warehouse before, so we were restricted to one product with a reasonably long history. It would be most useful to replicate the empirical findings on other high-technology products so that we begin to develop empirical generalizations about the orders–sales relationship and how it is affected by the marketing mix. Also, Bayesian methods should be designed and tested for use in new product categories where time-series observations are scarce. Finally, other variables that were not available in the present study should be added to this investigation, notably consumer awareness tracking, intermediate product performance variables such as customer inquiries, and retail inventories.

An important area for future research is to investigate how the use of a marketing-mix model on retail sales helps improve production and marketing resource allocation decisions, not only from the manufacturer's perspective but also for the entire value chain. Since the marketing-mix model generates forecasts that are conditional on prices and marketing support levels, it makes the manufacturer more knowledgeable about consumer behaviour that is normally only observable to the retailers. Does this enhanced knowledge lead to new channel gaming behaviour and order patterns that are even more unstable, or does it result in cooperative behaviour that smooths the order patterns and benefits all participants in the value chain? We hope that further research will consider these and other questions and improve the practice of order forecasting and the management of the marketing mix.

ACKNOWLEDGEMENTS

I gratefully acknowledge the contributions of Joel Bryant, John Holden, and Kim Jensen to the empirics in the paper. They provided a real-world setting and assembled a database that helped shape my thinking about order forecasting and motivate the manuscript. Also, John Roberts and two anonymous reviewers provided several useful comments that improved the content and the writing of the paper.

REFERENCES

Blattberg, R. C. and Hoch, S. J., 'Database models and export intuition: 50% model + 50% manager', *Management Science*, **36** (1990), 887–9.

Clark, D., 'As memory-chip shortage intensifies, higher prices may slow shift to new PCs', *Wall Street Journal*, 24 June 1994.

Dekimpe, M. G. and Hanssens, D. M., 'The persistence of marketing effects on sales', *Marketing Science*, **14** (1995a), 1 Winter, 1–21.

Dekimpe, M. G. and Hanssens, D. M., 'Empirical generalizations about market evolution and stationarity', *Marketing Science*, **14** (1995b), 3, Part 2, Summer, G109–G121.

Dickey, D. A. and Fuller, W. A., 'Distribution of the estimators for autoregressive time series with a unit root', *Journal of the American Statistical Association*, **74** (1979), June, 427–31.

Enders, W., *Applied Econometric Time Series*, New York: John Wiley, 1995.

Engle, R. F. and Granger, C. W. J., 'Cointegration and error correction: representation, estimation, and testing', *Econometrica*, **55** (1987), 251–76.

Franses, P. H., 'Modeling new product sales: an application of cointegration analysis', *International Journal of Research in Marketing*, **11** (1994), 5, December, 491–502.

Fuller, J. B., O'Conor, J. and Rawlinson, R., 'Tailored logistics: the next advantage', *Harvard Business Review*, (1993), May–June, 87–98.

Hanssens, D. M., Parsons, L. J. and Schultz, R. L., *Market Response Models: Econometrics and Time Series Analysis*, Boston: Kluwer Academic Publishers, 1990.

Hoch, S. J. and Schkade, D. A., 'A psychological approach to decision support systems', *Management Science*, **42** (1996), 1, January, 51–64.

Johansen, S., 'Statistical analysis of cointegration vectors', *Journal of Economic Dynamics and Control*, **12** (1988), 231–54.

Lee, H., Padmanabhan, V. and Whang, S., 'Information distortion in a supply chain: the bullwhip effect', *Management Science*, **43** (1997), 4, April, 546–58.
McWilliams, G., 'At Compaq, a desktop crystal ball. (Compaq utilizes forecasting program to manage new-product introductions)', *Business Week*, 20 March 1995, 96.
Montgomery, D. B. and Silk, A. J., 'Estimating dynamic effects of market communications expenditures', *Management Science*, **18** (1972), B485–B501.
Pitta, J., 'Apple pegs loss at $740 million, to cut 2,800 jobs', *Los Angeles Times*, 18 April 1996.
Powers, K., Hanssens, D. M., Hser, Y. I. and Anglin, D. M., 'Measuring the long-term effects of public policy: the case of narcotics use and property crime', *Management Science*, **37** (1991), 6, June, 627–44.
Reibstein, D. J. and Farris, P. W., 'Market share and distribution: A generalization, a speculation, and some implications', *Marketing Science*, **14** (1995), 3, Part 2, Summer, G190–G202.
Tellis, G. J., 'The price elasticity of selective demand: a meta-analysis of econometric models of sales', *Journal of Marketing Research*, **25** (1998) (November), 331–41.
Tellis, G. J. and Weiss, D. L., 'Does TV advertising really affect sales? The role of measures, models, and data aggregation', *Journal of Advertising*, **24** (1995), 3, Fall, 1–12.
Urban, G. L., Hauser, J. R. and Roberts, J. H., 'Prelaunch forecasting of new automobiles: models and implementation', *Management Science*, **36** (1990), 4, April, 401–21.
Urban, G. L. and Karash, R., 'Evolutionary model building', *Journal of Marketing Research*, **8** (1971), February, 62–6.
Wittink, D. R. and Cattin, P., 'Commercial use of conjoint analysis: an update', *Journal of Marketing*, **53**, (1989), 3, July, 91–6.

Author's biography:
Dominique Hanssens is Professor of Marketing at the UCLA Anderson Graduate School of Management, where he has been on the faculty since 1977. He has served as the school's faculty chair and as associate dean, and is currently the director of the Robert Anderson research foundation in management. Professor Hanssens' consulting experience covers a range of strategic marketing problems such as marketing resource allocation, product-line pricing, database marketing and new-product strategy. His approach emphasizes the new opportunities for developing a sustained competitive advantage, offered by sophisticated market information and marketing intelligence. He has conducted assignments for British Telecom, Hewlett-Packard, Home Savings of America, Hughes, Johnson & Johnson, Mattel Toys, Pacific Telesis, TRW and Wells Fargo. Professor Hanssens also serves on the faculty advisory board of Gemini Consulting and the marketing advisory board of KeraVision.

Author's address:
Dominique M. Hanssens, Anderson Graduate School of Management at UCLA, 101 Westwood Plaza, Box 951481, Los Angeles, CA 90095-1481, USA.

