SOME PRACTICAL HINTS ON SALES FORECASTING

J G Kenworthy, FBPICS, Vice President, The British Production and Inventory Control Society

1. Why Forecast?

Few manufacturers enjoy the luxury of customers who are prepared to wait for an order whilst materials are purchased, parts manufactured and then the end product assembled and delivered. Such conditions may exist for some large capital engineering products such as state of the art nuclear power stations, and I was once in a restaurant in the Middle East when having ordered steaks, I observed the young son of the proprietor dashing out to the local butcher, to return only minutes later, with the necessary raw material. Indeed the more enthusiastic proponents of JIT suggest that the majority of businesses should operate this way, but in real life, most manufacturers must make some preparations and commitment in advance of receiving an order. Even the restaurant proprietor had based the size of his premises and the number of tables etc. on some perception of the future.

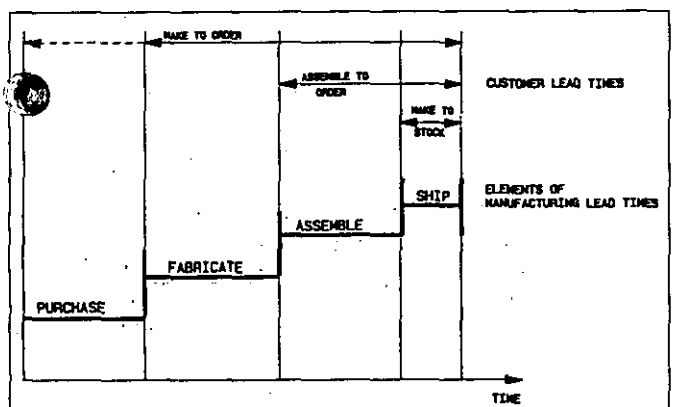

Relationship Between Manufacturing and Customer Lead Times **Figure 1.**

"Make to stock" companies will produce finished goods in anticipation of a customer's order, "Assemble to order" companies go part way, and even companies claiming to "Make to order", usually hold stocks of basic raw materials and common multi-purpose components. (See Figure 1.) All of these ways of operating make use of forecasting, even if informally, and it is difficult to imagine a successful business that does not do so in some way.

The APICS definition of a forecast is;

The extrapolation of the past into the future, - an objective computation involving data.

It is debateable whether the use of a crystal ball qualifies as an objective computation, but since all suppliers of these sought-after items appear to have been out of stock since the 17th century, the usefulness of this forecasting technique for P & IC must be in doubt. However I personally have no doubt that the marking of a re-order level on the side of a component bin, by an experienced foreman, qualifies equally well as does sophisticated curve fitting by computer.

In summary, manufacturing businesses forecast ahead to shorten response times by providing:

a) materials
b) labour and machine capacity
c) buildings and plant
d) capital

2. How do we forecast?

All forecasting involves extrapolating the past to predict the future. Anything else is guessing. The extrapolation may be intuitive, and may also be direct or indirect. In direct sales forecasting we extrapolate past sales to predict future sales, whilst in indirect extrapolation we extrapolate other parameters from which we can calculate future sales.

For example we could attempt to forecast sales of a new pharmaceutical product with no sales history, by extrapolating known figures for the number of people suffering from the disease which the drug treats, and also probable market share, perhaps based on past figures for market share achieved by other similar products.

To take an example from an entirely different area, we could attempt to forecast the number of pupils entering a secondary school 10 years ahead, by extrapolating past numbers directly. We would probably get a better answer by taking known figures for actual births in the catchment area, and extrapolating past data for movements into and out of the area.

In the author's experience, the most effective way of extrapolating directly is to plot the past data on a graph, draw a curve of best fit by eye and extend the curve intelligently into the future, making use of market intelligence, such as the imminent launch of a competitive product, or a planned sales promotion. This approach does not have a very high-tech image, but people are better than computers at making use of qualitative market intelligence. However before too many readers turn to the next article with shouts of "Luddite", it must be admitted that both sophisticated computer extrapolations, and simple ones based on exponential smoothing, do have their place, particularly when large numbers of items must be forecast. Modern PC based packages which not only determine the best extrapolation algorithm, and then display the resulting forecast as a graph on the screeen, but also allow the operator to adjust the curve to take into account special factors, are particularly attractive. A side benefit may be that the attractiveness and novelty encourage a reluctant marketing man to take forecasting seriously.

3. How far Ahead do we Forecast

The conventional answer to this question is as far ahead as the stacked lead time, plus the planning time. Thus in a make to stock environment, the forecasts must cover the total production and purchasing lead time plus the time until the next set of forecast has been collected and processed. This, however, ignores the need to predict the requirement for buildings and plant. A new factory may take three or four years to plan and build, and this does not allow for the time taken by senior management to believe the forecast. There is a natural tendency when faced with an important decision based on a forecast, to wait and see how the sales figures for the next month or year support the forecast. The implication that all forecasts should cover the next 4 or 5 years is, however, clearly suspect. Few companies will carry out long term planning for plant and buildings more than once per year, but forecasts controlling short term production need to be updated at least quarterly and preferably monthly or more frequently. Revising the long term forecasts monthly is a waste of effort.

One solution is to have two forecasting systems, a short term one to drive production, covering typically 12 months ahead, with time periods of weeks or months, and a long term one with time periods of years. To provide useful financial as well as volume information the long term system should include a forecast of an average selling price for each item for each year. Since detail is not required for long term planning this system can operate at product family level, thus simplifying the task of revision

However such a solution is not without problems. The author remembers well explaining proudly to production manage-

ment how they could accommodate the growth in output predicted for four years ahead, only to be told that they had been operating at higher levels than that for the past six months. Investigation showed that the short term forecasts which were in volume units only, were completely out of step with the long term system, which covered volume and value and was used to set annual sales budgets as well as for capacity planning. Not surprisingly the salesmen liked to forecast conservatively in the long term so that they had achievable budgets, but high in the short term so that there was little danger of running out of stock. Needless to say, production not marketing were accountable for stock levels in that company.

There is in principle, no problem in having a system which operates at product level in the short term, and at family level in the long term, with the first year of the long term forecast being calculated by summing the short term forecasts for the members of the family. The short term (production) forecast could then be revised as often as necessary, and a warning generated if the calculated short term family forecast deviated by more than a set amount from the trend set previously by the long term figures. This situation is shown in Figure 2 for a simplified family with only two members. However until such systems are commonplace, there is little alternative to spending time and effort reconciling different systems serving different purposes.

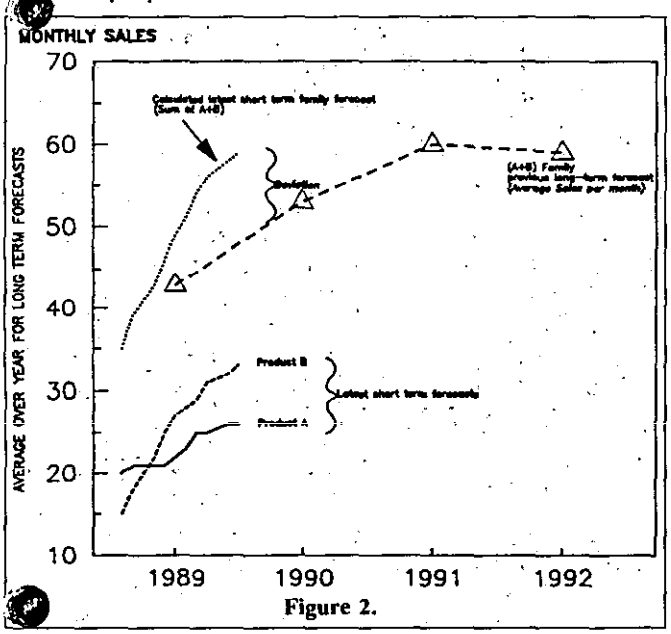

Figure 2.

4. Accuracy of Forecasts

An oft quoted expression is that the only thing that is certain about a forecast is that it will be wrong. Unfortunately this also is not always accurate, because just occasionally forecasts can be right. The accuracy we can expect for a forecast depends on the situation. A well established product selling regularly to many different customers in the home market, is relatively easy to forecast, because variations caused by individual customers tend to even out. (Data from a real example of such a case are shown in Figure 3a.) Conversely a product sold to only one agent in a third world country, may be almost impossible to forecast, because even if the agent has an effective stock control system he may only order infrequently and the timing of orders may depend more on the availability of foreign exchange rather than demand for the product. (See Figure 3b for an example of such a situation)

It is traditional for production planners to complain about the accuracy of marketing's forecast. Certainly forecasts should be as accurate as possible. There is no excuse for not trying, or leaving a secretary to make the forecast, but the key requirement of a forecast is that it should be honest, the best available in the circumstances and not deliberately biased to influence targets or stock levels. The accuracy illustrated in Figure 3b may be the best possible in the circumstances.

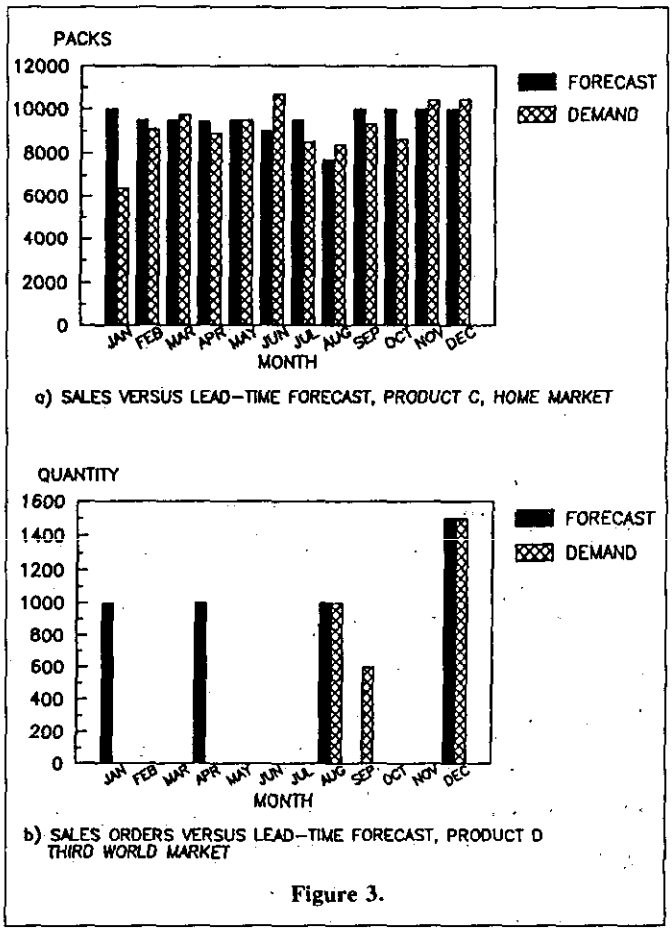

Figure 3.

We can do several things to improve forecast accuracy. Firstly the person responsible for the forecast should have ready access to accurate data on past sales. Secondly accuracy achieved in the past should be monitored and fed back to the forecaster.

Figure 4 (overleaf) shows a forecast request form which illustrates these two principles. It is from a forecasting system which has been operating successfully for 15 years. More modern systems may have reduced the need for paper, but the principles remain the same. At the top of the form, actual past sales are compared with lead time forecast, and the deviation displayed by quarter. Lower down, months 3 to 12 of the previous 12 forecasts made at the last formal up-date 3 months ago, are displayed to remind the forecaster what he predicted last time, and then there is an optional facility to display a mathematically predicted forecast. Interestingly very few of the users request this facility, or choose to accept the results if it is activated, perhaps indicating that user education has been inadequate to overcome an inherent suspicion. Finally the forecaster is asked to enter his new forecast at the bottom of the form. He is not given the option of changing the forecast for months 1 to 3 at this formal update, because this is inside the production lead time. If he believes that the previous forecast is incorrect he can send an "Ad Hoc" modification to the planners who will respond if possible, but give no commitment to do so.

In order to achieve maximum commitment to forecast accuracy, the accuracy statistics should also be displayed to senior management, probably in summary form. I am aware of one case in which publication of suitable summary statistics brought about a significant improvement in accuracy. Any product for which sales over 3 months were within +-20% of forecast was defined as a "Hit". The percentage "Hit rate " for each forecaster was then displayed to the Marketing Manager. The manager wasn't swamped with data: he only had to look at one figure per quarter for each forecaster, but he was so appalled by the figures that he offered a substantial prize, of the sort which salesmen normally win by meeting sales targets, for the best improvement over 6 months. The results were spectacular.

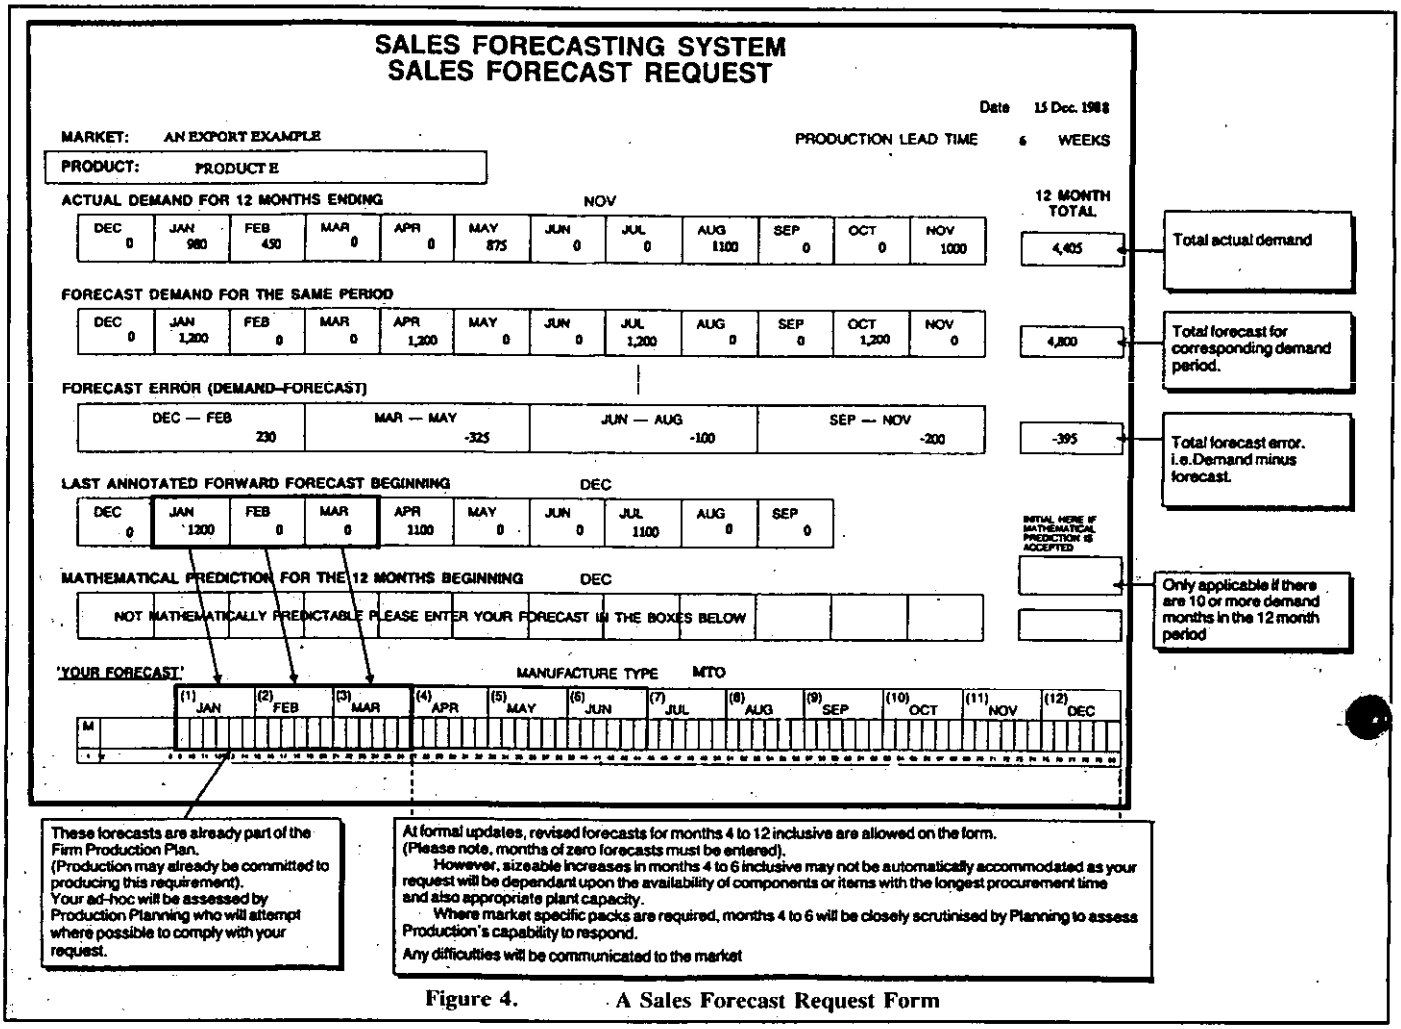

Figure 4. A Sales Forecast Request Form

In measuring forecast accuracy, it is important to compare actual sales with the forecast made at an agreed lead time in advance. The proud boast that this month's sales were within 1% of forecast, is not very meaningful if the forecast is the revised figure made just before the sales figures were published.

Several text books, suggest that a forecast should never be a single figure, but should always indicate a range within which the sales are expected to lie. Although one cannot disagree with this principle, in practice a manufacturing order must be for a specific quantity and it is very difficult to handle a "range" within an automated planning system. Experience suggests that the best way to handle uncertainty, is by discussion between the planner and the forecaster to agree whether safety stocks of finished goods will be held, components will be bought and capacity will be reserved etc. or whether a risk of poor customer service is acceptable.

5. Making Forecasting Easier

There are several ways in which it may be possible to make forecasting easier and hence probably more accurate.

Forecasting is easier for the immediate future, than it is for a long time ahead. Anyone who doubts this should look at success records for weather forecasting. Therefore the more we can shorten manufacturing lead times by applying continuous improvements and moving towards JIT manufacture the easier forecasting becomes. For some industries this may prove to be one of the major advantages of a JIT programme.

As was discussed in Section 3, forecasting for a product family takes less effort than forecasting each member of the family separately. Usually also the past actual sales for a family will exhibit a smoother trend than do the sales of individual members of the family. Hence extrapolation of the trend is likely to be more accurate at family level.

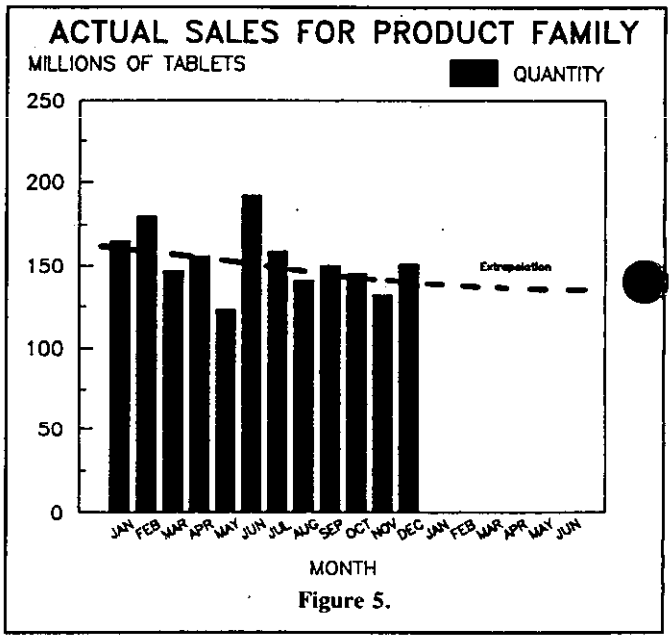

Figure 5.

Product D in Figure 3b is in fact a pharmaceutical with 100 tablets in a pack, labelled for a particular export market. The same tablet is sold in several pack sizes, each labelled for a variety of markets, giving a total of 150 different items in the "Product Family". As we saw, demand for any particular pack can be somewhat erratic and difficult to forecast. If however we sum the past sales for all 150 packs in the family, we get the situation shown in Figure 5. The units now are not packs but tablets. The curve is very much smoother and relatively easy to extrapolate. Provided that there are enough members of the family this will always be the case. Forecasting at this

level will be much more accurate. The key question is how can we use such an aggregate forecast. If packing is a relatively quick operation, and if specific package components can be obtained quickly then tablet manufacture can be planned against the family forecast and packing only done to order.

In this particular case, this is not practicable because printed package components such as cartons have a long lead time, but all is not lost. The family forecast can be used as a check on the individual pack forecasts, by summing the forecasts for the individual packs and comparing the result with the family forecast. Agreement does not of course ensure that the individual forecast are correct. There may well be compensating errors, but disagreement certainly means that it is worth re-examining the individual forecasts.

We saw in Figure 3a how sales to many customers will follow a smooth trend because variations in requirements of individual customers tend to cancel each other. With many products the rate at which the end-users consume the product will be relatively smooth, but the manufacturer can experience wide swings in demand as the various stock points in the distribution chain react to and magnify, small changes in the rate of end use. By moving the forecasting point out towards the end user, forecasting becomes easier because demand is smoother. To take advantage of this we need to have visibility of stock levels in the distribution chain, ie to use Distribution Requirements Planning or DRP. In effect we are trying to turn the situation represented by Figure 3b into that of Figure 3a. How possible this is, will depend on the relationship between supplier and customer and the sophistication of systems available to the customer.

In one particular situation, illustrated in Figure 6, where a UK factory supplies a major depot in Sweden, which in turn not only supplies the local market but also smaller depots in Denmark, Finland and Norway, forecasting was originally aimed at predicting shipments from the UK to Sweden. Changing to forecasting local sales in the four territories, and supplying stock figures for the four depots so that DRP logic could be used to calculate the demands on the factory, gave significant improvements in control. The principle works equally well with a number of wholesalers or distribution points in one market. The critical requirement is to have sufficient cooperation to provide the forecast and stock data, and sufficient trust to accept "Unordered" shipments from the factory.

Finally forecasters should take advantage of the Pareto Principle or the 80:20 rule. 20% of the forecasts will cover 80% of the total demand, so care and attention should be concentrated on these 20% with the other 80% being left to the computer

6. Summary

Sales forecasting is not easy but forecasts are essential for most production control systems, and the more accurate the forecasts the better the control.

Forecasting accuracy can be improved in practice by giving attention to the following points:

a) Reduce manufacturing lead times by continuously striving towards JIT principles

b) Recognise that honesty in a forecast is at least as important as accuracy.

c) Measure the accuracy of forecasts and feed back to the forecaster. Also provide summmary statistics for senior management.

d) Forecast at the highest level of aggregation which provides adequate detail for manufacturing.

e) Move the forecasting base as close to the end-user as is practicable.

f) If possible have a consistent set of numbers covering the short, medium and long term future.

g) Handle uncertainty by dialogue between the forecaster and planner.

h) Exploit Pareto's principle. Concentrate human effort and intelligence on improving the forecasts for the really important items. Leave the rest to the computer.

May your forecasts always be honest. They won't always be accurate.

About the Author

John Kenworthy has 16 years experience of Production and Inventory Control, in the UK and overseas. He is currently MRPII Assurance Manager with ICI Pharmaceuticals where he played a leading role in achieving provisional Class A status for the MRPII implementation. John is a Vice-President of BPICS, a Fellow of the Society and a past President of North West Branch.

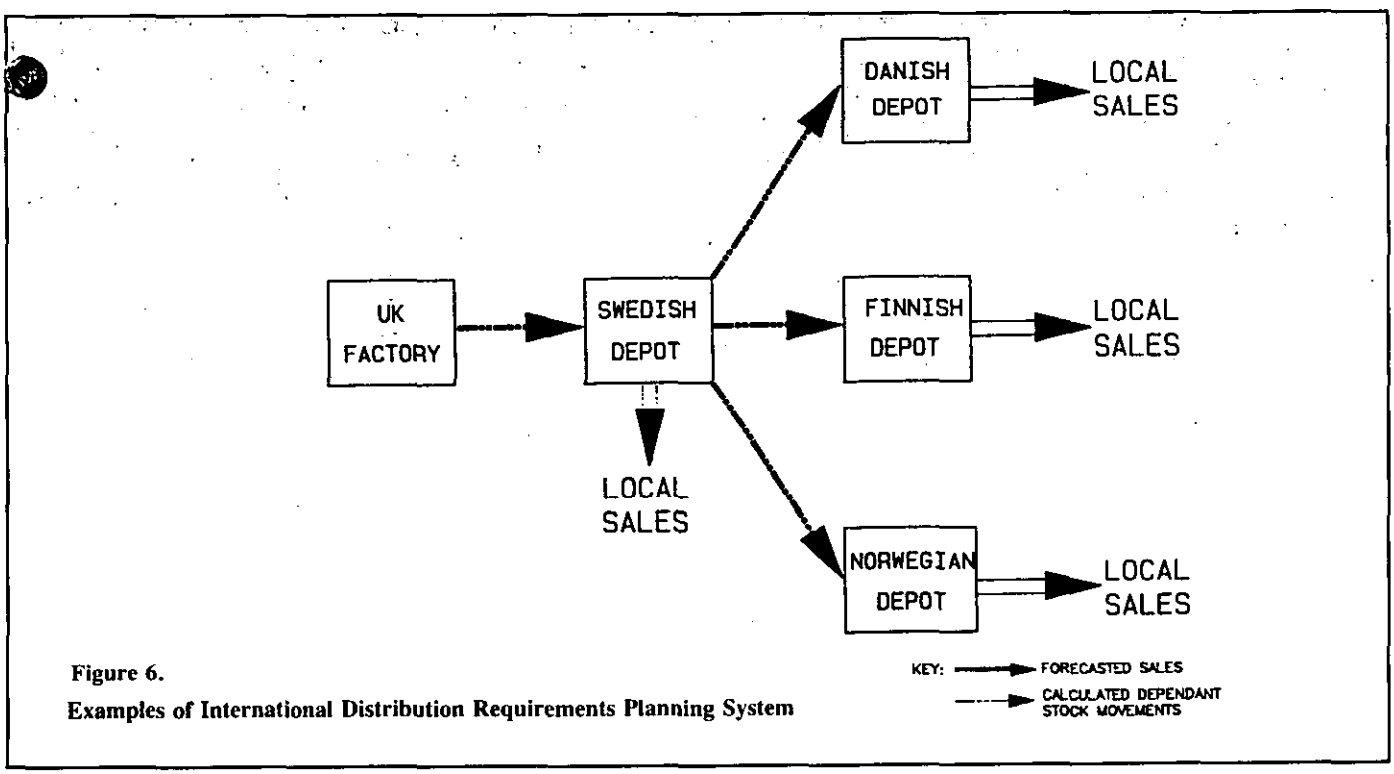

Figure 6.

Examples of International Distribution Requirements Planning System