

Collaborative Forecasting

By Harpal Singh

What is Collaborative Forecasting?

Collaborative forecasting is the process for collecting and reconciling the information from diverse sources inside and outside the company; to come up with a single unified statement of demand. It consists of four key elements:

1. Applying statistics and other algorithms to past data to recover relevant information.
2. Processes and systems to collect customer-level input routinely. In some businesses, this may be referred to as geographic information.
3. Processes to merge management overrides and inputs with the data collected at the customer level.
4. Processes to merge marketing input, which is usually product focused, with the sales view that is usually customer focused.

The planning of resources to service the demand is usually not regarded as part of the forecasting process. However, it must be well integrated with the forecasting process for supply chain planning to function effectively.

The Technical Challenge

Implementing a collaborative forecasting process requires a substantial effort in process re-design and system implementation. The major technical challenge is the need to reconcile differing views of the market demand.

1. Customer versus the Marketing View

In most business organizations, two distinct organizations deal with marketplace demand.

The sales organization is normally organized by customers, regions, and sales territories, with each sales office or sales person being matched to a distinct set of customers. There may be some variations to this in the case of major customers that span territories or regions. In this case an account manager may

be assigned responsibility across regions or territories.

The marketing organization is normally product focused. In many companies, a product family manager has responsibility for the market development and overall sales and revenue targets for each product family.

A direct result of this organization is that there is a *many to many* relationship between the marketing entities and the sales entities. A product line is sold in many sales territories and a sales territory sells many product lines.

If the sales organization sets independent targets on projected sales (and develops a forecast based on these targets), and marketing also sets independent targets on product volume, there needs to be a mechanism to reconcile the two.

The annual budgeting process in many companies is geared to reconciling these differences. However, most companies lack the business processes and systems to reconcile these continuously as the environment changes throughout the year.

2. Required Precision versus Available Data

In order for the demand information to be useful to manufacturing, it is often required at the level of product by week. This level of precision is necessary to balance the available capacity and smooth asset utilization. Marketing and customer information is not usually available at this level of precision.

For example, a customer may know that he is likely to receive a large contract some time in the next 6 months, but may not be able to provide more specifics. Similarly, aggregate trends by industry or region cannot easily be converted to precise *product by week* numbers.

3. Booking versus a Shipping Forecast

The sales organization is organized around sales targets and persuading customers to

place orders that match the internal revenue and volume goals. Often, an order placed by a customer may be for an entire year or quarter, with staggered deliveries. Manufacturing, on the other hand is interested in when the deliveries are supposed to occur.

Some industries have the practice taking blanket orders, which cover an entire year. While these are sufficient from a sales point of view, they do not provide the information necessary to align the assets to specific deliveries.

4. Required Aggregation for Applying Statistics

It is a well-known fact that the aggregation of data leads to better statistical results. For example, it is easier to predict the rainfall in a region for the entire year than to predict the rainfall for any given day.

Statistical methods work best when the information content of the data exceeds the random errors in the data. When data is aggregated, random errors tend to cancel out and the information content increases proportionately.

For statistical purposes, the aggregation can be done either by attributes (for example region, product family, etc) or by time. Aggregation by time means that the data buckets used are larger (for example quarters instead of months or weeks).

The major challenge is that the required level of aggregation to get meaningful statistical information may not match the precision required by the business to manage the supply chain.

5. Reconciling Forecasts at Different Levels

Even when dealing with single forecasting view, inputs come at more than one level. For example, the sales input may be at the customer/product level. The management goals at the sales office level may be motivated by other business factors, and may or may not coincide with the sum of inputs at the lower levels.

To illustrate this, let's take the example of a single product sold to 4 customers in a region.

Customer	Volume
A	1000
B	2000
C	3000
D	4000

If the management goal for these customers is 8000, then should the individual targets be proportionately reduced? On the other hand, if the management goal is 12000, should the targets be increased across the board, or selectively by customer?

6. Reconciling with Orders

A further difficulty is introduced when orders from one or more of the customers are placed. Consider the situation below:

Customer	Volume	Orders
A	1000	1500
B	2000	1500
C	3000	0
D	4000	0

If the management goal for these customers is 10000, should the forecast for B, C, and D be reduced in some fashion? Or should the forecast for A be increased?

There is no single solution to this of course. If the customers are competing for similar contracts, it may be appropriate to adjust the expectation of B, C, and D downward. On the other hand, if the customers are relatively

independent, then the orders placed by A may be a precursor to an increase in volume.

Supporting a Many-to-Many Relationship

To institutionalize collaborative forecasting, a company needs to address the challenges outlined above. These require substantial work in business process redesign and system support.

As we have seen, demand information may not always come in at the same level of aggregation, or even at the same view. For example, changes to demand may come in as "the automotive industry is expanding at 10%" or "the European market will expand by 20% in the next 6 months," etc.

Information like this is often treated qualitatively in the planning process because the available systems cannot convert the information from one view to another.

An essential requirement of the forecasting system is the ability to support forecast conversion from one view to another, where there is a many-to-many relationship between the views.

For example, the system must be able to convert a regional forecast by volume to a forecast of volume by industry in a seamless way.

This is not possible in most forecasting systems that are designed around hierarchies or pyramids.

1. Simple Conversion of a Forecast

One potential solution is to use the pivot table concept that is popular in spreadsheets. Consider a forecast by customers used in the example above. Suppose that our desire is to convert the forecast by customers, to a forecast of volume by product. The simplest mechanism is to construct a table, which shows the proportion of a customer's volume that should be assigned to each product as shown below.

Product/Customer	A	B	C	D
Product 1	20	0	70	40
Product 2	60	90	30	40
Product 3	20	10		20

This table represents the percentage of a product volume that is shipped to a customer. With this table available, it is a simple task to convert the forecast by product to a forecast by customer. The actual calculation is simply a matrix multiplication of the proportion table above with the forecast table.

Using the tables of percentages below,

	A	B	C	D
Product 1	20	0	70	40
Product 2	60	90	30	40
Product 3	20	10		20

And the customer forecast,

Customer	Volume
A	1000
B	2000
C	3000
D	4000

the forecast for product 1 is 20% of customer A, 70% of customer C, and 40% of customer D.

2. Time Varying Mix Ratios

In the example above, the percentage of a customer volume assigned to a product was assumed. For a short-term tactical forecast, it may be sufficient to derive the percentage ratios from the recent history. However, this is not sufficient if the forecast is seasonal in nature, and if the demand in the marketplace is changing considerably over time.

In many businesses, the ratios that need to be applied to convert from one view to another vary according to time. Conceptually, this does not represent any difficulty because the ratios themselves can be forecast over time. For example, the volume by industry and region can be forecast based on history. The resulting forecast of each industry/region combination is normalized to provide a conversion table for each period in the future.

The basic assumptions in the conversion process are:

1. The total forecast for each period remains the same when converting from one view to another.
2. The conversion ratios to be used are well defined and agreed to by both the marketing and sales organizations.

In our experience, the simplest mechanisms work best. Conversion ratios derived by fitting a straight line to the percentages work best for the short term, and when no seasonality is present. For seasonal businesses, a simple seasonal function is adequate.

Multiple Time Aggregations

To extract useful statistical information, it is sometimes necessary to aggregate the data over time. Most current forecasting systems require the time bucket to be predefined. The difficulty with this is that the same time aggregations are not uniformly required at all levels.

Support for time aggregation must also include the ability to pro-rate the aggregated forecast or input from one set of time buckets to another.

Consider the situation where a weekly forecast is required by the business. There are at least 2 distinct statistical methods for calculating a projection. The first is to aggregate the history by week and apply the necessary statistics to this data. The second method is to estimate the quarterly volume based on a quarterly time bucket, and project the percentage of the quarterly volume in each of the 13 weeks in the quarter.

In an industry with a large emphasis on quarterly results, the second approach tends to provide more accurate information.

For a system to allow multiple time aggregations, a completely different paradigm for forecasting needs to be adopted. In the past, business forecasting has focused on projecting the quantity in a time bucket. The only available mechanism for getting a more precise (as opposed to accurate) forecast is to look at data in smaller time buckets.

The forecast for a time bucket consists of two elements. The first is indeed the amount that is predicted for the time period. The second is a function that describes how the volume is to be distributed within the time period.

Why is this important? In actual practice, the information on the sales volume may come from an organization that is focused on obtaining quarterly or semiannual orders to meet sales targets.

The distribution of this volume may be more accurately reflected by the shipment history. By separating the forecasting process into these two distinct areas we can readily reconcile the volume projections in one time period and the requirements for a more precise breakdown.

Managing Views for Collaboration

As we have seen, forecasts can be converted from one view to another by applying a simple process of matrix multiplication. The conversion matrix can be either defined by recent history, or based on known factors like trends in the marketplace.

Each view is best described as a specified aggregation of qualitative attributes and a time aggregation. For example, a sales view may be by region, sales office, and product, at the monthly level. Similarly, a marketing view can

be described by say business unit, product family, and product.

The conversion matrix for the forecast can be derived by looking at the historical mix and projecting this mix into the future with statistical models.

With a well-defined methodology for converting a forecast from one view to another, a company can now take advantage of using the strength of each view. In general there are at least four views that need to be in place to support a supply chain planning:

1. A statistical view for applying mathematical models. This view is at a level of aggregation where the statistical models provide useful results
2. A marketing view that is product family focused. This view is used to input aggregate changes for existing products, new products, handle product substitution, and possibly check for critical components.
3. A sales view which is customer focused. It is used to gather customer related information. An example of this view would be a view by: region, sales office, and customer.
4. A manufacturing view that is used for capacity reconciliation. This would typically be by product or product family, and by week or month.

Required System Characteristics

1. Distributed Collection of Input

Forecast inputs are usually gathered from those closest to the customer. These may be the sales persons, customer service representative, or customers. After a forecast input is identified, there is usually an evaluation process. The salesman may check the input against past volume, recent history, accuracy of the customer's past projections before passing on the input for consolidation.

In today's environment, the forecast collection process must be a distributed one, and must allow each sales person to operate independently without being connected directly to the forecasting system. The reason for this is simple: Collecting forecast information is not a central job task of most sales persons. The forecasting system must be able to

accommodate different working styles and allow a sales person to operate on his or her local computer for much of the time. Connection to the forecast system is only needed when data needs to be exchanged.

2. Support for Volume and Revenue

For collaborative forecasting to work effectively, the data has to be meaningful to the different organizations involved in creating and reconciling the forecast.

In most organizations revenue projections are key to interpreting sales goals. Marketing and manufacturing, on the other hand, are usually focused on volumes.

The system must be able to accommodate both types of data and maintain the consistency between volume and revenue.

3. Visibility of Forecast Changes

In a collaborative forecasting environment, multiple organizations provide input which is consolidated into the final statement of demand. As the forecast is developed, changes that are made should be visible so that the responsibility and accountability can be maintained.

Consider the situation where a sales person enters a forecast for a customer and product combination. The sales manager reviews and adjusts the total. This is then turned over to marketing for further additions. Rather than change the sales person's forecast, the system must maintain the differences so that the changes made at each organizational level are visible.

What this means is that it is no longer sufficient to merely maintain the previous forecast and use the accuracy of the total forecast as a metric. The forecasting system must maintain the changes made at each organizational level.

In the example above, a forecasting sheet for a product/customer combination may contain the following rows. Often, the previous forecast is also included, as is the average selling price and revenue.

In our experience, the sales organization is usually more comfortable with estimating the average selling price and the volume, with the sales revenue being calculated.

	Month 1	Month 2	Month 3	Month 4
Operating Budget	10	20	30	40
Sales Input	10	12	60	30
Sales Mgr	+5	-2	+5	+30
Marketing	+0			+30
Forecast	15	10	65	90

Together with the visibility of changes, the business process must document and maintain a well-defined precedence of changes. What this means is that the sequence in which changes are applied are clearly understood.

One example might be that in the sales organization, the sequence is sales person, area manager, followed by sales director. This is then turned over to marketing for further adjustment.

An Example of a Process

1. Overview

Our example describes a forecasting process for a business that supports multiple regions and multiple product lines. The goal is to outline the system architecture and a large part of the business process needed to project market demand in a routine and disciplined way.

The sales organization has 3 sales office, North America, Europe, and Asia. A customer in a region is assigned to sales person. Typically, the sales person would handle between 5 to 6 customer accounts.

The business sells product for three types of applications, Automotive, Communications, and Industrial. Each of these areas has a marketing manager who is responsible for overall revenue targets for his sector.

There is an annual budgeting process, which is used to set sales targets for the year, and develop revenue and volume estimates for the next 12 months. Ongoing, the sales forecast needs to be revised routinely to support a monthly Sales and Operations Planning process.

The business has a single individual who is responsible for consolidating the worldwide demand for the monthly S&OP process.

2. Process to Support Sales

The current business forecast is kept in a central server. The individual sales people access this server through the Intranet to retrieve the information for their customers. The information retrieved consists of:

- Annual budget (volume and revenue)
- Historical bookings (volume and revenue)
- Historical shipments (volume and revenue)
- Current bookings which reflect the shipment date (volume and revenue)
- Average selling price
- Marketing overrides to the shipment forecast (volume and revenue)
- Sales person's booking forecast (volume and revenue)
- Sales person's shipment forecast (volume and revenue).

After retrieving the information, the sales person aggregates the information by products, customers, groups of products, or groups of customers.

The business process requires that he review his key product and customer combinations on a weekly basis. His task is to maintain the booking forecast and the average selling price fields up-to-date. For this purpose, he can also generate a statistical forecast interactively to project the volume or revenue.

The sales person enters or uses the statistical forecast to determine the volume and projected price for each combination. Revenue is calculated based on this.

The shipment forecast is automatically generated based on the historical lag between

bookings and shipments. The sales person can elect to override this.

While the sales person may review his forecast on a weekly basis, he may not elect to update the server. He does this when he feels that there are sufficient changes to warrant an update.

On a daily basis, the forecasting system on the server consolidates information from the company's data systems and also from any sales person changes. The area sales manager retrieves the information and reviews it on an as needed basis.

For this business, the process for entering changes is restricted to the customer/product combinations. To ensure accountability, the area managers must reach a consensus with the actual sales person in charge of an account. The sales person makes any required changes.

3. Process to Support Marketing

The marketing managers work interactively on the server to access their part of the forecast in much the same way that the sales person works. However, the marketing managers provide input at the product or product family levels. Their input is disaggregated to the product and customer level based on a mix percentage.

The mix percentage agreed to by the business is the last three months shipments plus the next three months projected bookings. Marketing forecasts both volume and revenue. The average selling price is determined from the projected volume and revenue.

To aid the marketing managers, the system is able to project a forecast with different time aggregations, and also to forecast the average selling price. For new products the business has agreed to a formula which projects a decreasing price. This formula has been determined from historical data.

The statistical analysis for this business is used in an advisory fashion. The marketing managers can invoke the calculations when needed, but are not constrained by them.

4. Process to Support Sales & Operations Planning

On a monthly basis, the forecasts are consolidated and reviewed. The person

responsible for consolidating the forecast retrieves the current consolidated forecast from the server to his own personal computer. This is then used to make the necessary reports for the S&OP process.

Changes to the forecast are made directly at the S&OP meeting. The projected demand is reconciled with the supply, and the operating plan is modified as required.

These new operating plan is then uploaded to the server.

5. Key Elements of the Process

- The sales input is retrieved asynchronously. The sales persons can continue to provide input at any time during the month.
- There is an agreed to process and calculation to convert bookings into shipments.
- Responsibility for the sales input is clearly defined.
- Marketing adjustments are communicated to sales as adjustments, rather than changes to the original input.
- There is an agreed to mechanism for converting the changes made at the product family level to a sales view.
- Both revenue and volume are considered for this business because it operates in an environment where prices are dynamic.

A specialized procedure is used to reconcile the sales persons input with marketing overrides. This procedure uses a Linear Programming approach and is explained in the next section.

Reconciliation at More Than One Level

1. The issue

One issue that we have not addressed yet is how to reconcile forecasts at different levels consistently. This issue arises in cases where the orders for a customer exceed the customer projections, or when the sales person's override for a particular customer exceeds the marketing projection for overall sales.

To address this, many forecasting systems require a well-defined reconciliation strategy.

In the top-down strategy, the forecast of 8000 would be prorated in the ratio of the component forecasts. In the bottom-up strategy, the component forecasts would take precedence and overwrite the 8,000 with 10,000. In both these strategies, one of the inputs is essentially ignored.

These strategies are too simplistic to adequately handle practical situations. An alternative strategy is described below which attempts to address the limitations of the top-down and bottom-up strategy.

Consider the example described in the table below:

Customer	Mix Percent	Forecast	Sales Input
A	10%	1000	1500
B	20%	2000	1500
C	30%	3000	1000
D	40%	4000	1000

In the absence of any sales input, the forecast of 10,000 from marketing would be disaggregated as shown. If there is an override of 1500 for A, then it seems reasonable to increase the forecast of A to 1500. Since we should attempt to retain the marketing volume, the forecast for B, C, and D need to be reduced. It also seems reasonable that B, C, and D should be reduced in the same ratio as their original forecast.

These ideas can be formalized through a set of equations as explained below.

2. Rules and Assumptions

When conflicts arise, we suggest the following set of governing rules:

The forecast for each entity is always greater than the current shipments and bookings. This rule ensures that actual customer demand, or sales overrides are considered directly in adjusting the forecast.

The forecast of the parent entity (in our example the sum of the forecasts) is maintained unless the sum of actual inputs at the lower (child) level exceeds the forecast.

The forecast of the children entities is kept as close to the original ratios as possible. We do this by minimizing the maximum percentage deviation for each child whose orders and shipments do not exceed the original forecast.

These rules can be summarized in a series of simultaneous equations. The table below shows the result of applying the Linear Program to the data above.

When the total forecast is 16,000 there is no conflict between the input values and marketing. However, when the forecast is 10,000 the forecast of each item is set to at least the input values, with the remaining being pro-rated by the original percentages.

In the case when the forecast set by marketing is less than the sales input, the marketing override is increased to accommodate the sales input values.

	Mix Percent	Sales Input	Total= 10000	Total= 16000t	Total= 8000t
A	10%	1500	1500	1600	1500
B	20%	1500	1889	3200	1500
C	30%	1000	2833	4800	2143
D	40%	1000	3778	6400	2857

Dr. Harpal Singh is a respected industry thought leader in the supply chain arena. His expertise spans the areas of implementing decision support systems, planning and scheduling, demand management, and inventory management. He has led management seminars and supply chain improvement projects globally.

Harpal is largely responsible for initiating the use of expert systems at DuPont to facilitate advanced planning and scheduling. He left DuPont in 1993 and founded Supply Chain Consultants, Inc. (SCC). With locations globally, SCC offers supply chain planning consulting and software solutions for sales & operations planning, demand planning, production planning & scheduling, logistics planning and inventory management to customers worldwide.

©2002

Supply Chain Consultants®

www.supplychain.com