Smart *White Paper*

Accurate Intermittent Demand Forecasting for Inventory Planning: New Technologies and Dramatic Results

By Charles N. Smart, President, Smart Software, Inc.

Note: This paper was originally presented in a slightly different form at the APICS 2002 International Conference and Exposition in Nashville, Tennessee. It is reproduced here with permission of APICS—The Educational Society for Resource Management, Alexandria, VA and is published in the APICS 2002 International Conference Proceedings.

Maybe you're in the service (spare) parts business or a manufacturer or distributor of capital goods. If so, you know the inventory planning challenges of intermittent or "slow-moving" demand. When a customer calls, you want to make sure you have what he needs. You are also constantly trying to balance the need to reduce the costs of your inventory against the requirement to maintain high customer service levels. Because you have no good way to accurately forecast demand for slow-moving items, the chances are that your company has too much inventory and it's costing a lot of money—sometimes tens of millions of dollars.

Recent experience with a new forecasting technology pioneered by Smart Software, Inc. of Belmont, Massachusetts—the Smart-Willemain method—suggests that the age-old problem of accurately forecasting intermittent demand is history. Smart Software has integrated this new technology into its flagship product SmartForecasts® Enterprise, an enterprise-wide demand forecasting, planning and inventory optimization system. Users of this software are experiencing almost 100 percent forecasting accuracy, millions of dollars in inventory savings, and greatly improved customer service levels.

One example is NSK Corporation of Ann Arbor, Michigan. It makes anti-friction bearings and precision parts for the automotive and other industries. Two-thirds of the 7,000 SKUs in the aftermarket business unit (ABU) demonstrate highly intermittent demand, with half of the sales history data containing zero values. The inventory replenishment lead times for some items can average several months, and the costs for some items run as much as several hundred thousand dollars. Customers assume that the ABU will always have an item they want in stock and will be able to respond quickly when they place an order.

In its first year of implementation in the ABU, SmartForecasts Enterprise helped reduce inventories of standard aftermarket products, saving the company $1 million, shortening lead times, and increasing on-time delivery above a 98 percent service level. Moving forward, NSK estimates that SmartForecasts' new technology will help to lower inventory levels of intermittently demanded products even further, producing annual savings of about $3 million for the ABU alone.

The dollar savings from reducing its inventories are particularly important to NSK, but in its line of business, maintaining a high customer service level is a requirement of doing business with customers such as General Motors. The new forecasting technology will help to accomplish the seemingly competing tasks of reducing inventory while increasing customer service levels.

As you can see, the new intermittent demand forecasting technology can make a big difference—reducing inventories and their associated costs, improving customer service, and enhancing business control.

- Reducing inventory levels is often an important goal of inventory planners. Especially in tough economic times, no one wants cash tied up in excess inventory. Accurate demand forecasting helps companies optimize stocking levels to ensure that the right service part or product is available at the right place at the right time. The effects of optimizing go directly to the bottom line. Since inventory carrying costs are commonly 25-35 percent of inventory value, the cost of overstocking becomes evident. As an example, reducing inventory by $1 million can improve cash flow by as much as $350,000. That's $350,000 now available for other uses.
- Accurate forecasting can help improve customer service and thereby increase customer satisfaction. While business "goodwill" is hard to measure, not having the right part in stock at the right time—stocking-out—can be costly, especially when the item carries a high margin or the customer is an infrequent purchaser. One manufacturer of expensive industrial tools that achieved a 20 percentage point increase in customer service levels using the new forecasting technology in SmartForecasts, and whose customers don't purchase on a regular basis, laments: "If we don't have one item on an order, we often lose the entire order. What's worse, we run the risk of losing the customer to a competitor."
- Accurate forecasting enhances business control. If you're not able to accurately forecast the demand for slow-moving inventory, it affects the operation of the rest of your organization. The results produced by forecasting software are used by other systems to improve planning across the organization—sales and financial planning, as well as inventory management, production scheduling and distribution planning. In turn, better planning can lead to better asset utilization, whether it is in manufacturing, finance, or some other area.

Why is Forecasting Intermittent Demand so Difficult?

Where intermittent demand exists, the benefits of computerized forecasting solutions have been elusive. There's a good reason why. Traditional forecasting methods don't work well because intermittent demand data typically contain a large percentage of zero values, often 30 percent or more, with non-zero values mixed in randomly. An example of the difference between intermittent demand data and product demand data that is normal, or "smooth", is illustrated in the tables below:

Table 1: Intermittent Demand Data

Month	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
Demand	0	0	**19**	0	0	0	**4**	**18**	**17**	0	0	0	0	0	**3**	0	0	**19**	0	0	0	**5**	**4**	**5**

Table 2: Normal, Smooth Demand Data

Month	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
Demand	17	20	18	25	30	68	70	41	32	35	66	26	23	25	25	28	36	68	82	39	31	38	57	33

Traditional statistical forecasting methods—such as exponential smoothing and moving averages—that work well with normal, smooth demand do not give accurate results with intermittent data because they ignore the special role of zero values in analyzing and forecasting demand, as well as other key features of intermittent data. Moreover, those traditional methods assume that the probability distribution of total demand for a particular product item over a lead time (lead time demand) will resemble a normal, classic bell-shaped curve (see Figure 1).

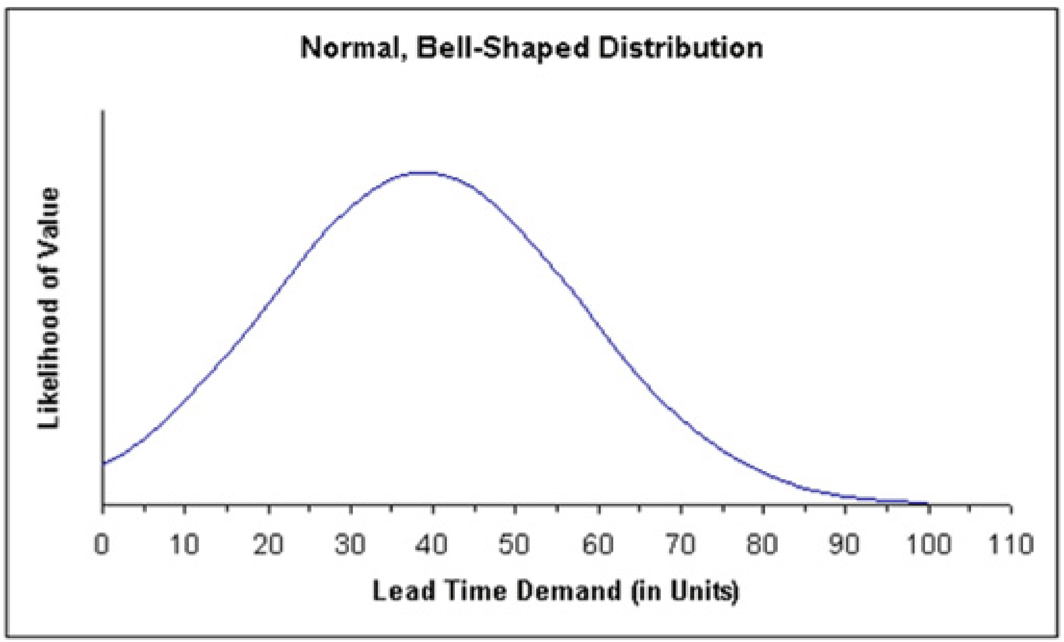

Figure 1

With intermittent data, the assumption of a normal distribution of lead time demand is not appropriate, especially in the case of service parts, where lead time demand can exhibit unusual shapes (see Figure 2 below).

While conventional statistical forecasting methods can produce credible forecasts of the average demand per period when demand is intermittent, they cannot produce accurate estimates of the entire distribution (i.e., complete set) of all possible lead time demand values. Too often, what they do produce are misleading inputs to inventory control models—with costly consequences.

For each intermittently demanded item, the importance of having an accurate forecast of the entire distribution of all possible lead time demand values—not just one number representing the average or most likely demand per period—cannot be overstated. These forecasts are key inputs to the inventory control models that recommend correct procedures for the timing and size of replenishment orders (reorder points and order quantities). They are particularly essential in service parts environments, where they are needed to accurately estimate customer service level inventory requirements (e.g., a 95 or 99 percent likelihood of not stocking out of an item) for satisfying total demand over a lead time.

Faced with the intermittent demand problem, many organizations rely on applying judgmental adjustments to their statistical forecasts, which they hope will more accurately predict future activity based on past business experience. But there are several problems with this approach, as well.

First, judgmental forecasting is not feasible when you are dealing with large numbers (thousands and tens of thousands) of items. Second, most judgmental forecasts provide a single-number estimate instead of a forecast of the full distribution of lead time demand values. And third, it is all too easy to unintentionally but incorrectly predict a downward (or upward) trend in demand, based on expectations, that results in under-stocking (or over-stocking) inventory.

The Scope of the Intermittent Demand Problem

Lest one not appreciate the scope and consequences of accurately forecasting intermittent demand, it is a multi-billion dollar problem affecting a variety of industries worldwide. Intermittent demand is most commonly found in the service parts businesses of the aerospace, automotive, high tech/electronics, utilities, and industrial machinery industries. It is also found in capital goods companies producing expensive, big-ticket items and, to some extent, in retailing. AMR Research estimates that $700 billion was spent on service parts in the U.S. in 2001, which represents about 8% of the U.S. gross domestic product. Considering that, until recently, there was no way to accurately forecast demand for many of these goods, the scope of the problem becomes clear.

AMR also discovered that service parts inventory is the second largest asset on the balance sheet of many companies. It is widely understood that the annual cost of storing, depreciating, insuring, and moving service parts is 25-35 percent of the book value of inventory. Patton Consultants, a specialist in field service and service parts management, estimates that the material (parts inventory) costs for service can be as much as 60 percent of total service-related costs.

In just one example, a major telecommunications company used the Smart-Willemain method to estimate inventory stocking levels for approximately 800 key replenishment items (a fraction of the company's total number of inventory items) at the 95 percent customer service level. The forecast of recommended stocking levels produced using Smart-Willemain required $38 million less inventory than the estimate generated by the forecasting system the company was already using. Moreover, the expected annual savings in carrying costs associated with that $38 million reduction in excess inventory was over $10 million.

The total potential savings to be realized by accurately forecasting intermittent demand becomes even clearer if we multiply the potential savings for the telecommunications company by the hundreds of companies of similar size and add to that the potential savings of those tens of thousands of smaller companies where a significant proportion of the inventory items experience intermittent demand.

Solving the Intermittent Demand Forecasting Problem

The research behind the Smart-Willemain method was prompted by customer requests and real-world concern for the forecasting problems inherent with intermittent demand. The National Science Foundation (NSF) awarded a competitive Innovation Research Grant to Smart Software to develop a new and more accurate method of forecasting intermittent demand. Smart Software's research was aided by a team of researchers at Rensselaer Polytechnic Institute, in Troy, New York, under the direction of Dr. Thomas Willemain, a professor in the Department of Decision Sciences and Engineering Systems.

The NSF study examined 28,000 commercial data series that included inventory items from nine companies in the U.S. and Europe, representing the aircraft, high tech, electronics components, marine equipment and other capital equipment industries.

The study confirmed two things:

- Both exponential smoothing and a variant of exponential smoothing, developed by the statistician J.D. Croston in 1972, are effective in forecasting mean (average) demand per period when demand is intermittent.
- But, it also confirmed that neither Croston's method nor exponential smoothing accurately forecasts the entire distribution of demand values, especially customer service level inventory requirements for satisfying total demand over a lead time (for example, the amount of inventory required to provide a 90, 95 or 99 percent likelihood of not stocking out of a product item).

More interestingly, however, the research produced a new forecasting method, based on bootstrapping, that provides fast, realistic forecasts of intermittent product demand over a fixed lead time. Bootstrapping is a statistical method that accurately forecasts both average demand per period and customer service level inventory requirements. It does this by using samples of historical demand data to create thousands of realistic scenarios that show the evolution of cumulative demand over a lead time.

To illustrate the basic bootstrap in simple terms, consider the intermittent demand values presented earlier in Table 1:

These 24 monthly values represent the historical demand for a typical service part item.. Most of these values are zero, but mixed in over the two-year period are nine non-zero values (19,4,18,17,3,19,5,4,5). Let's say you need forecasts of total demand for this item over the next three months because your parts supplier needs three months to fulfill an order to replenish inventory. A simple bootstrapping approach to this problem is to sample from the original 24 values, with replacement, three times, creating a bootstrap scenario of demand over the three-month lead time.

For example, you might randomly select months 4, 8 and 15, which gives you demand values of 0, 18 and 3, respectively, for a total lead time demand in units of 0+18+3=21. Repeating the process again, you might randomly select months 12, 3 and 21, giving a lead time demand of 0+19+ 0=19 units. By continuing to generate thousands of bootstrap scenarios in this way, using appropriate computer-based algorithms, you can build a statistically robust picture of the lead time demand distribution.

The graph in Figure 2 shows the results of 100,000 such bootstrap scenarios generated for this part item. (These bootstrap scenarios reflect all aspects of the new Smart-Willemain methodology, including the real-world possibility that non-zero demand values for the part item that occur in the future may differ from those that occurred in the past.)

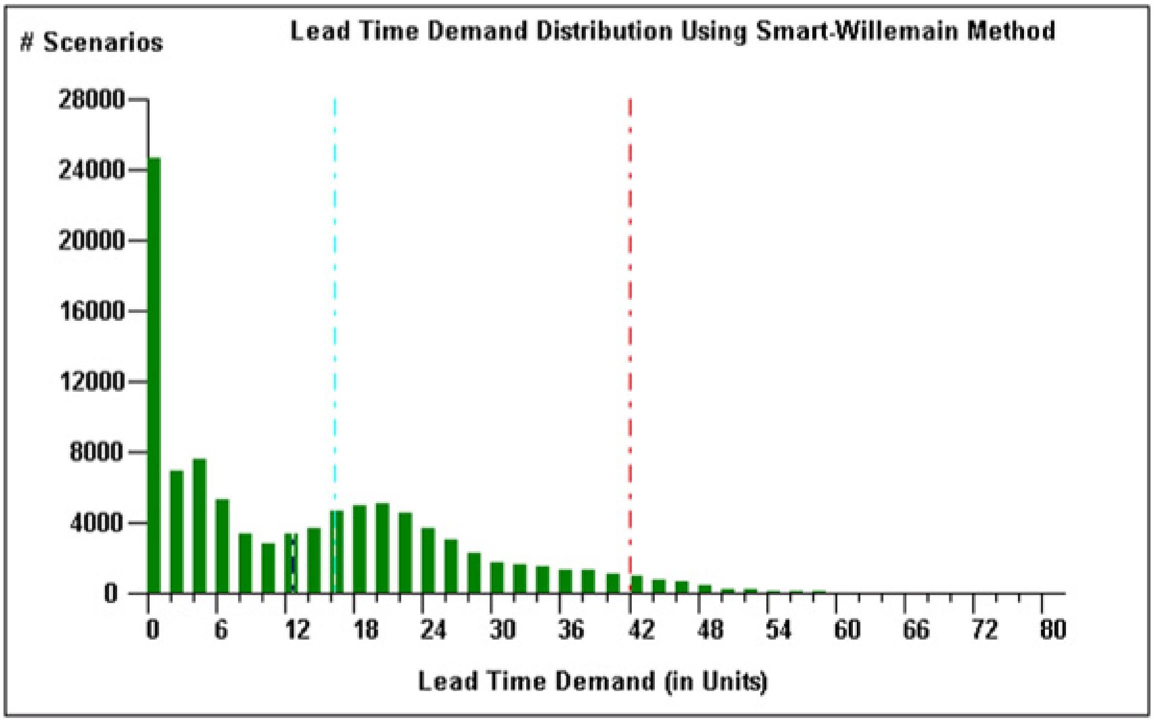

Figure 2

Figure 2 indicates that, in this example, the most likely value for lead time demand is zero (about a 25% likelihood), but lead time demand could be as great as 60 or more units. Indeed, the mean or average value of lead time demand (see the dashed blue line on the graph) is approximately 14.6 units, while the 95th percentile of the demand distribution, which is the forecast of the amount of inventory required to satisfy total lead time demand at the 95% service level (see the dashed red line on the graph), is 41 units. Obviously, the lead time demand distribution in Figure 2 looks nothing like the normal, bell-shaped curve displayed in Figure 1 above—and any inventory models assuming it does will provide unreliable advice on setting reorder points and order quantities.

To demonstrate the difference in results between the Croston's and Smart-Willemain methods, we used the same set of historical data in Table 1 to generate a lead time demand distribution using Croston's method. That distribution is shown in Figure 3.

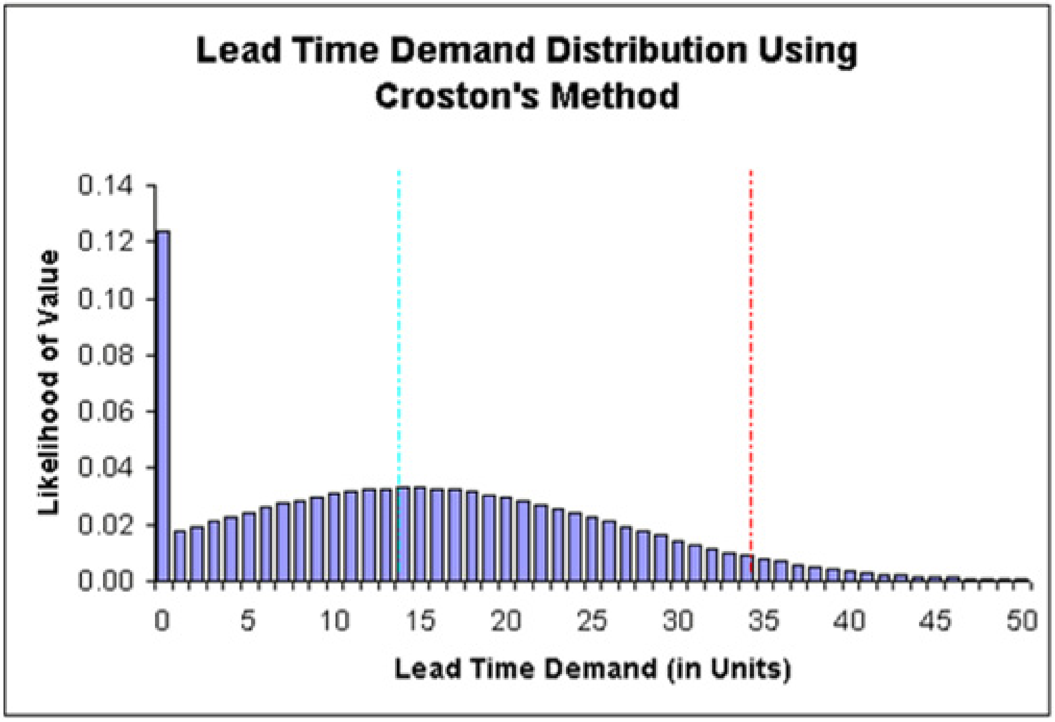

Figure 3

Although Figure 3 correctly indicates that the most likely value for lead time demand is zero, it significantly underestimates that likelihood compared to the Smart-Willemain results in Figure 2 (12% vs. 25%). Besides the spike at zero, Figure 2 shows two (2) more modes in the data, corresponding to the small and the large non-zero demand values observed in Table 1, while Croston's only shows the generic normal distribution. Most telling, however, is the shorter length of the right-hand tail of the demand distribution in Figure 3 (only up to a maximum of 50 units). In this example, this leads to a noticeably smaller and, we believe, incorrect estimate of the amount of inventory required to satisfy total lead time demand at the 95% service level (see the dashed red line representing 34 units for Croston's versus 41 units for Smart-Willemain). Setting inventory at this lower level would invariably result in more unwanted and unexpected stock-outs.

Compared to Croston's and other traditional methods, one of the main strengths of the bootstrap-based, Smart-Willemain forecasting technology is that it is an empirically-based method of computational statistics that substitutes data-based computing for theoretical assumptions. It is also very fast—more than 100,000 intermittently demanded service parts can be forecasted on a standard Pentium PC in less than an hour. The resulting increase in forecast accuracy, together with the fast production of forecast results, allows for increased service parts availability and customer service for many thousands of items, while minimizing inventory carrying costs.

New Technology in Practice

The bootstrap forecasting technology was first introduced to American Production and Inventory Control Society (APICS) members in the June 2000 issue of APICS—The Performance Advantage magazine. In that issue, a feature article ("Get Real: A New

Way to Forecast Intermittent Demand") written by Charles Smart and Thomas Willemain, Ph.D., described the new bootstrap technology, including examples of representative intermittent demand data and forecast graphics. Subsequently, Smart Software received a U.S. patent for the new forecasting technology in 2001.

SmartForecasts Enterprise with the Smart-Willemain technology has been deployed in a number of companies worldwide spanning a variety of industries. Among the companies that are using it are Labatt Breweries of Canada, Siemens Building Technologies, Textron Lycoming and Niagara Cutter, as well as Mitsubishi Heavy Industries and Olympus Optical in Japan. What has been remarkable is the technology's accuracy in determining total, cumulative demand over a lead time: almost 100 percent.

Recently, several companies that experienced significant intermittent demand for their products, put SmartForecasts to the test. Using their own historical demand data, the companies evaluated service level accuracy and inventory cost reductions using the Smart-Willemain technology:

- The first company, a warehousing operation for a nationwide retailer, forecasted inventory requirements for 12,000 intermittently demanded SKUs at the 95 and 99 percent service levels. The forecast results were almost 100 percent accurate. At the 95 percent service level, 95.23 percent of the items did not stock out (95 percent would have been perfect). At the 99 percent service level, 98.66 percent of the items did not stock out (99 percent would have been perfect).
- The aircraft maintenance operation of a worldwide company got similar service level forecasting results using the Smart-Willemain method to forecast 6,000 SKUs in one part of its business. Then this company took the process one step further by translating the forecast into cost savings. It found that, if it had used Smart-Willemain's recommendations rather than its current planning system, it would save approximately $3 million a year in inventory carrying costs.

Implementing SmartForecasts in the Enterprise

Implementations of the new technology can be very straightforward. The Smart-Willemain technology has been incorporated into SmartForecasts Enterprise as a high-end, specialized feature. The software draws its data from existing enterprise systems using the demand history tables resident in corporate databases such as Oracle, SQL Server, IBM DB2, and others, as well as from data in spreadsheet files. The forecasting software segregates and selects those items with intermittent demand data, which can be very useful when there are tens of thousands of items in the database. A statistical forecast is then created automatically based on the relevant data. Forecast results can even be manually adjusted to reflect changes in the business environment. The results are then returned to a forecast table in the host database, where they can be directly accessed and used in a company's other planning systems.

When you are using a software system like SmartForecasts to perform specialized forecasting tasks, the ability to integrate tightly with enterprise data has its advantages. It leverages existing enterprise infrastructure. Because SmartForecasts doesn't require a separate database of its own, but rather uses the data already stored in a company's enterprise systems, the company can use the database technology, personnel, and support resources in which it has already invested. For that reason, SmartForecasts with the new forecasting technology can be up and running in just a few weeks.

What this means to companies using the software is a quick, low-cost implementation, together with a fast payback and very high ROI. In the example above of the large telecommunications company, the payback period is less than a week and the estimated ROI based on first-year savings alone is 10,000%.

For inventory planners, especially, SmartForecasts provides a fast, accurate solution to the problem of forecasting intermittent product demand. Like most other commercially available forecasting applications today, SmartForecasts is built around an automatic forecasting capability that makes it accessible to the non-statistician. Demand data that was once considered unforecastable no longer poses an obstacle to achieving the highest customer service levels with the lowest possible investment in inventory.

For more information, please contact:

Smart Software, Inc.
Four Hill Road
Belmont, MA 02478
Telephone: 1-800-762-7899
Fax: 617-489-2748
Email: csmart@smartcorp.com
Website: www.smartcorp.com