Bootstrap Your Way to Better Forecasts

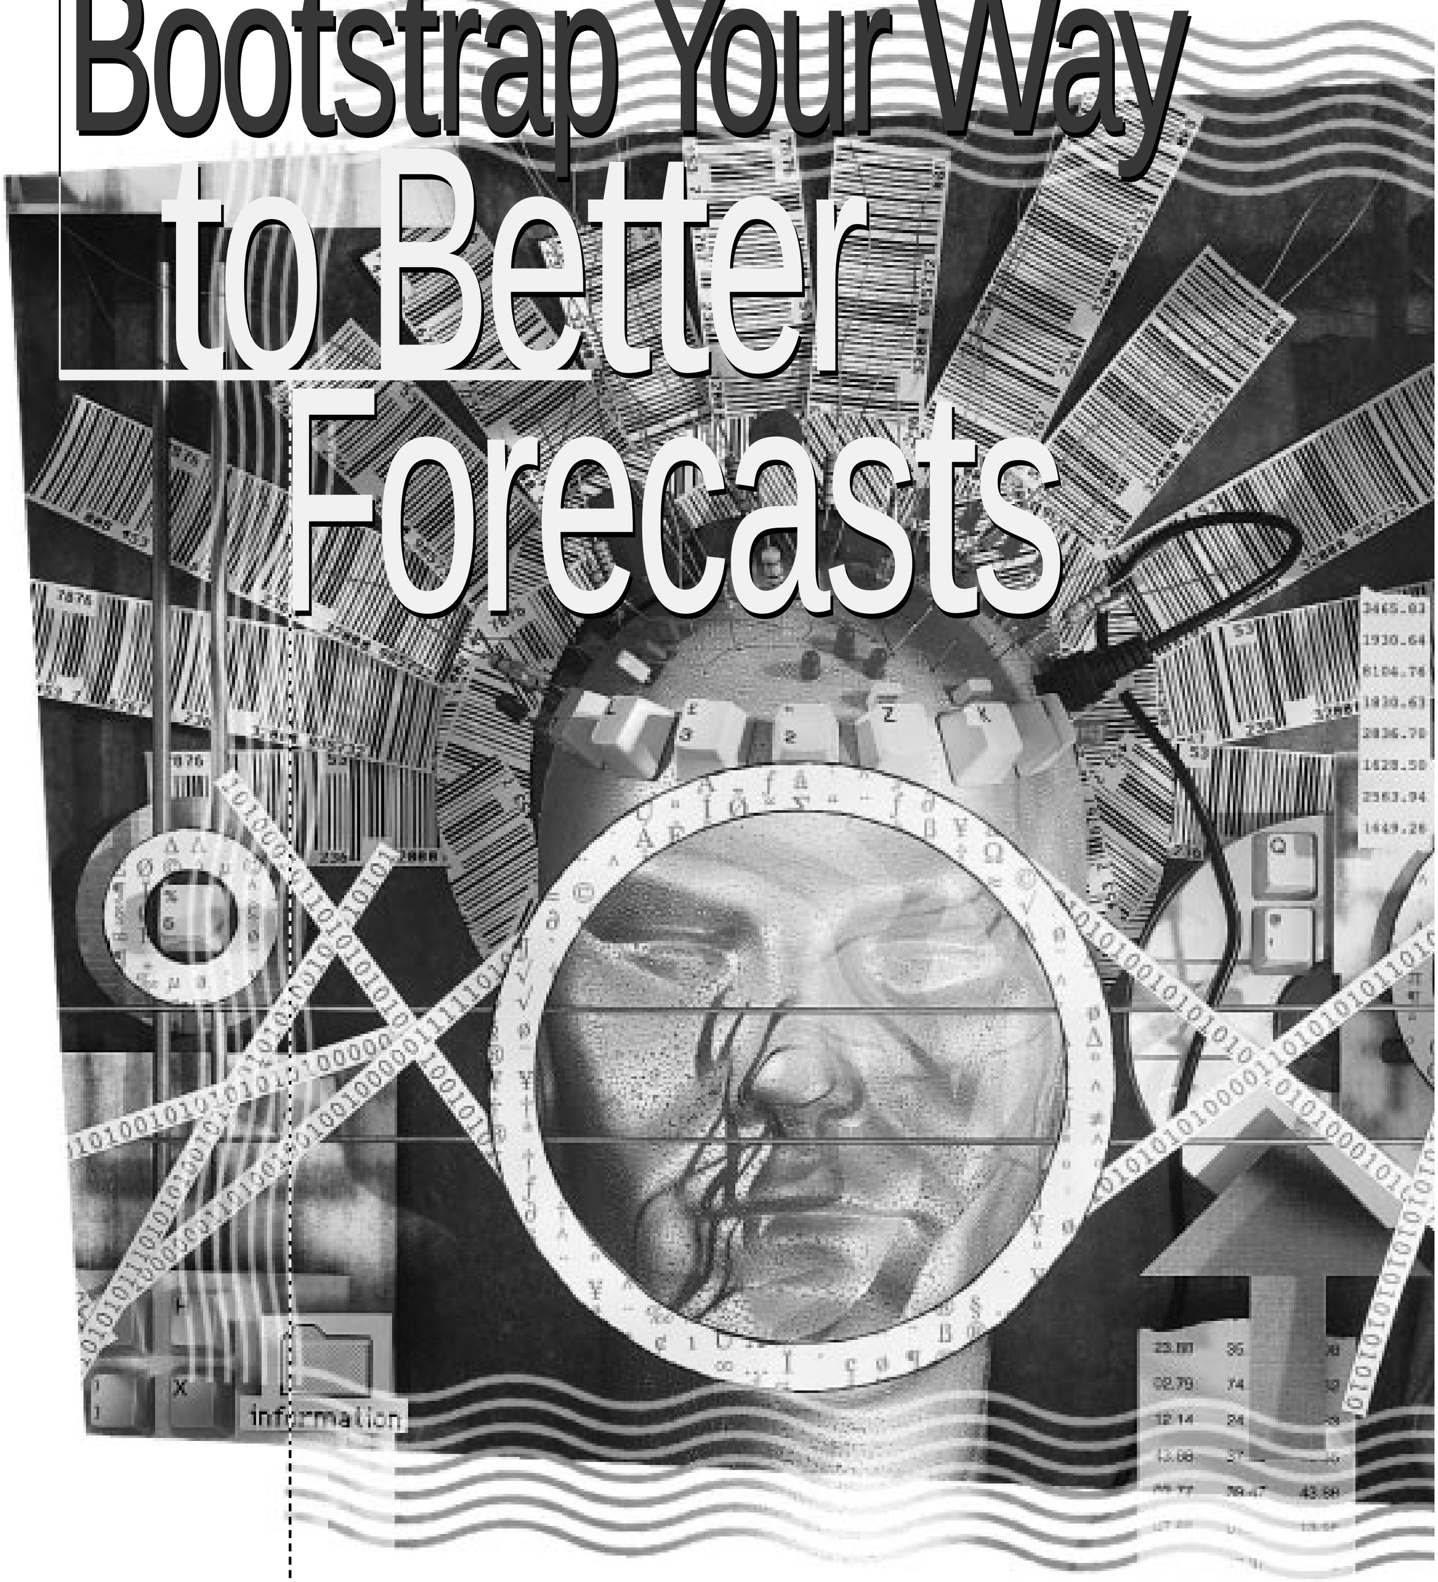

A new forecasting method to address the big problem of planning intermittent demand

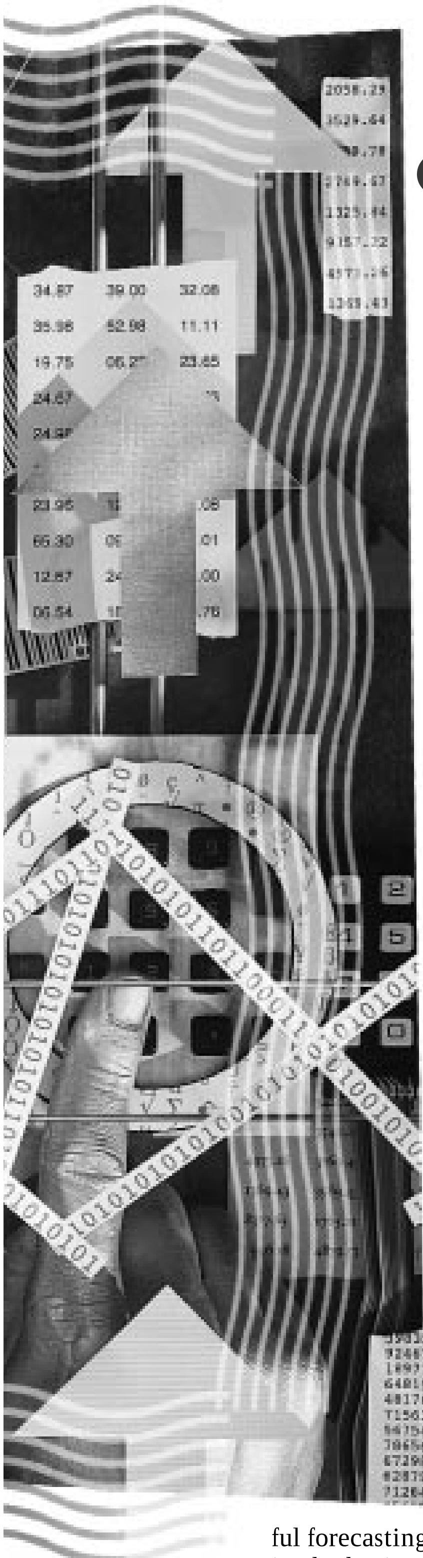

By Charles N. Smart

Intermittent or uneven demand—particularly for low-demand items like service parts—is especially difficult to predict with any accuracy. A new method, bootstrapping, may provide the answer.

Do any of these scenarios apply to your company?

- Sometimes, you miss out on business opportunities because you can't accurately forecast demand and estimate inventory requirements for certain products.
- You've had to make unnecessarily large investments in inventory to cover unexpected orders and materials requirements.
- Shortened lead times and the pressure to increase customer-service levels, twin demands of the e-commerce era, have made efficient inventory management even more difficult.
- You know you need more powerful forecasting tools, but you're not in the business of writing software and your midrange company doesn't have huge MIS resources.

Most likely, if your products exhibit intermittent demand patterns, you're acquainted with some or all of the above. You aren't alone.

Intermittent demand—also known as irregular, sporadic or slow-moving—affects industries of all types and sizes: capital goods and equipment sectors, automotive, aviation, industrial tools, specialty chemicals, utilities and high tech, to name just a few. And it makes demand forecasting and planning extremely difficult. It can be much more than a headache; it can be a multi-million-dollar problem, especially for suppliers of service parts and for those who manage service parts inventories.

Identifying intermittent demand data isn't hard. It typically contains a large percentage of zero values, with non-zero values mixed in randomly. But while companies have been wrestling with intermittent demand for years, few forecasting solutions have yielded satisfactory results, until recently. Why is this so?

Traditional approaches

Traditional statistical forecasting methods, like exponential smoothing and moving averages, work well when product demand data is normal, or smooth, but it doesn't give accurate results with intermittent data. Many computerized forecasting tools fail because they work by identifying patterns in demand history data, such as trend and seasonality. But with intermittent demand data, patterns are especially difficult to recognize. These methods also tend to ignore the special role of zero values in analyzing and forecasting demand.

Even more importantly, traditional statistical forecasting 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. With intermittent demand, it doesn't, especially in the case of service parts, where lead-time demand can exhibit unusual shapes (table 1).

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 replen-

How Bootstrapping Works

Month	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*	*12*
Demand	0	0	**7**	**3**	0	0	0	**32**	0	0	0	**6**
Month	*13*	*14*	*15*	*16*	*17*	*18*	*19*	*20*	*21*	*22*	*23*	*24*
Demand	**9**	0	0	**4**	0	**15**	0	0	0	**8**	0	0

Table 1. Monthly demand values for a service part item.

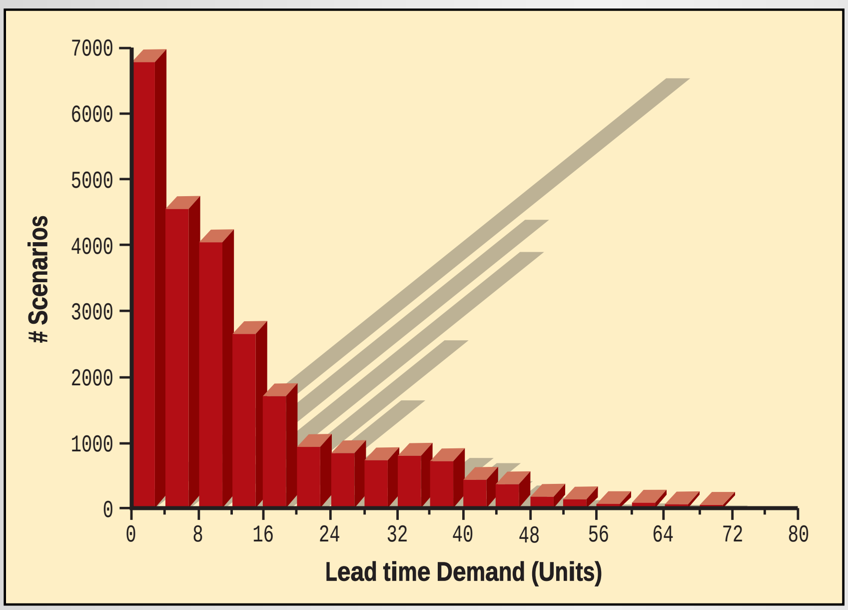

Figure 1. The results of 25,000 bootstrap scenarios.

How does the new method of forecasting intermittent demand work in practice? Although the full architecture of this technology includes additional features, a simple example of bootstrapping demonstrates the usefulness of the technique.

The 24 monthly demand values for a service part item, found in Table 1, are typical of intermittent demand history. 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 fill an order to replenish inventory. The bootstrap approach is to sample from the 24 monthly values, with replacement, three times, creating a bootstrap scenario of total demand over the three-month lead time.

You might randomly select months 6, 12 and 4, which gives you demand values of 0, 6 and 3, respectively, for a total lead-time demand (in units) of 0 + 6 + 3 = 9. You then repeat this process, perhaps randomly selecting months 19, 8 and 14, which gives a lead-time demand of 0 + 32 + 0 = 32 units. Continuing this process, you can build a statistically rigorous picture of the entire distribution of possible lead-time demand values for this part item. Figure 1 shows the results of 25,000 such bootstrap scenarios, indicating (in this example) that the most likely value for lead-time demand is zero but that lead-time demand could be as great as 70 or more units. It also reflects all aspects of the new bootstrapping technology, including the real-life possibility that nonzero demand values for the part item occurring in the future could differ from those that have occurred in the past.

With the high-speed computational resources available today, bootstrapping methodology can provide fast and realistic forecasts of total lead-time demand for thousands or tens of thousands of intermittently demanded product items. These forecasts can then be entered directly into inventory control models to insure that enough inventory is available to satisfy customer demand. This also ensures that no more inventory than necessary is maintained, minimizing costs.

ishment 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 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 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, as a director of operations for a cutting tools company recently explained, it is easy to inadvertently but incorrectly predict a downward (or upward) trend in demand, based on expectations, resulting in under-stocking (or over-stocking) inventory.

With the high-speed computational resources available today, bootstrapping methodology can provide fast and realistic forecasts of total lead-time demand for thousands or tens of thousands of intermittently demanded product items.

A new method that works

A more accurate way to forecast intermittent demand has emerged this year, the result of an innovative study examining 28,000 commercial data series (inventory items from nine companies in the U.S. and Europe), conducted for the National Science Foundation. Developed by researchers at Smart Software Inc., the new technique derives from bootstrapping, a statistical method that accurately forecasts both average product demand per period and customer service level inventory requirements. Basically, the method uses samples of historical demand data to create thousands of realistic scenarios that show the evolution of cumulative demand over a lead time.

Early evaluators of the new intermittent demand forecasting technology have found that it increases customer service level accuracy and significantly reduces inventory costs.

- A nationwide retailer's warehousing operation forecasted inventory requirements for 12,000 intermittently demanded SKUs at 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 global company got similar service level forecasting results with 6,000 SKUs. Potential annual savings in inventory carrying costs were estimated at $3 million.
- The aftermarket business unit of an automotive industry supplier, two-thirds of whose 7,000 SKUs demonstrate highly intermittent demand, also projected $3 million in annual cost savings.

If the challenge of forecasting intermittent product demand has indeed been met, it will be good news for midrange manufacturers. This new method, like most other commercially available forecasting applications today, is built around an automatic forecasting technology that makes it accessible to the non-statistician. Demand data that was once un-forecastable no longer poses an obstacle to achieving the highest customer service levels with the lowest possible investment in inventory.

Charles N. Smart is president of Smart Software, providers of forecasting and demand planning software. Contact him at 800 762-7899 or at csmart@smartcorp.com