Smart GEN2 SOFTWARE

GEN

Smart Software Announces “Gen2”

A White Paper from: Thomas R. Willemain, PhD Co-Founder and SVP Research at Smart Software Inc.

GEN

In this white paper, we introduce "Gen2", our shorthand name for the next generation of probabilistic modeling technology that powers the Smart Inventory Planning and Optimization (Smart IP&O) Platform. After reading this paper, you will learn:

The evolution of Smart Software's forecasting methods from automatic model selection to first generation probabilistic modeling ("Gen1"), to the new Gen2 methods.

How Gen2 substantially expands the use of probability math that made Gen1 so useful to so many companies.

Why Gen2's many advantages – analysis at the daily level, model accessibility, accounting for lead time variability, optimized historical inputs, support for multiple replenishment methods, and refined demand classification – represent a giant leap forward compared to traditional approaches.

THE BEST WAY TO PREDICT THE FUTURE IS TO CREATE IT

Alan Kay

Smart Software, Inc. Four Hill Road, Belmont, MA 02478 t: 617-489-2743 www.smartcorp.com
A White Paper from: Thomas R. Willemain, PhD Co-Founder and SVP Research

Gen1 Overview

Gen1 forecasting techniques invented by Smart Software have evolved over nearly four decades of research and refinement. They include the first automatic forecast selection technique, pioneered by Smart Software in the earliest days of the personal computer revolution, and the first probabilistic forecasting methods that were engineered to address intermittent demand.

Forecasting tournament for selecting forecasting methods

In the Gen1 forecasting tournament, empirical reality, not pure theory, governed the choice among competing statistical methods. Rather than requiring the forecaster to be an expert time series analyst, Gen1's extrapolative modeling automatically selected the methods with the best potential accuracy for each item and then pitted them against each other to see which would win the right to forecast each item. The criterion for winning the tournament was also a departure from past practice: rather than looking backwards to see which model best "fit the past", it addressed the real issue of predicting the future. Using extensive computation, each method's accuracy was assessed using hold-out analysis to see how it would have performed if the forecast had been made in the recent past, holding out recent data as the basis for measuring forecast error.

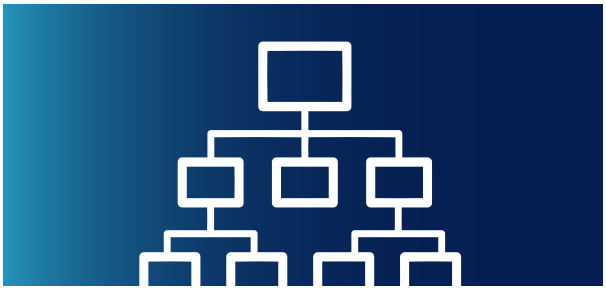

Solution of the unique problem of intermittent demand with probabilistic forecasting

The classical prediction methods built into the automatic forecast tournament made certain assumptions about the nature of item demand, especially that it was "smooth" and not riddled with zero values. But many inventory items, such as spare parts and big-ticket durable goods, have demand that is intermittent. Gen1 incorporated Smart Software's patented bootstrap methodology specifically designed to cope with intermittent demand. Gen1 conducted thousands of lead time demand simulations yielding a more accurate distribution of demand that resulted in precise estimates of safety stocks and reorder points required to achieve targeted service levels.

Together with several other innovations, Gen1 became the state of the art. In time, however, the speed of business increased so that monthly planning cycles became too sluggish for competitive performance. Together with SKU proliferation ("mass customization") and expanded distribution networks, the forecasting datasets became more complex (e.g., evincing multiple levels of seasonality simultaneously). This pushed Smart Software to develop Gen2.

Gen 2 Overview

Gen2 harnesses daily demand and supply signals and substantially advances Gen1's probabilistic forecasting methods. Now probabilistic forecasting accounts for not only trends but also daily, weekly, monthly, and quarterly seasonality. Instead of running a "hold out" tournament, forecasts are based on probability distributions of each future period. This enables more accurate estimates of mean demand and more realistic assessments of forecast uncertainty (since there are no assumptions about "bell-shaped curves"). Fewer historical data points are required making Gen2 very effective with new product forecasts and when older data is no longer relevant. Accurate forecasts of daily and weekly demand can be generated with as little as 45-60 days of demand history. Finally, the approach while powerful is quite easy to understand and thus easier to apply at scale to support business planning. The question "How many forecasting methods do you use?" may now be answered as "one" for customers leveraging Gen2.

How Gen2 Works

Two key analytical technologies undergird Gen2: bootstrapping and Monte Carlo simulation. This section provides basic explanations of both. Both methods are used in Gen2 to generate scenarios that embody the inherent randomness and uncertainty in forecasts, annual counts of stockouts, and similar results of operational significance. Both methods represent a move away from reliance on simplistic theoretical models that often fail to capture the complexity of real-world demand. Instead, they exploit the recent availability of more plentiful data streams and abundant, cloud-based computational power.

Bootstrapping

Bootstrapping came to practical use in the 1980's when computers grew powerful enough to support this form of computational inference. Originally, bootstrapping was seen as a way to analyze the uncertainty in statistics computed from data. Gen2 extends this approach to create prediction intervals for forecasts, such as "There is a 90% chance that next week's total demand will be less than 5 units." It does this by taking a scenario-based approach to forecasting.

Here is how it works: Suppose replenishment lead time is 7 days and you want to forecast total unit demand over the next 7 days. Gen1 creates scenarios of daily demand over 7 days, summing to get lead time demand, then repeating until a statistical pattern becomes clear. The calculations might look as shown in Figure 1. The results of 1,000 scenarios show that lead time demand is very unlikely to exceed 10 units, and a reorder point of 5 units would suffice to get through the replenishment delay with only a 10% chance of stocking out.

Scenario #	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	7-day Total
1	2	0	0	0	0	0	0	2
2	0	0	0	0	0	0	0	0
3	0	3	0	0	0	0	0	3
4	0	0	0	0	0	0	0	0
5	0	0	0	1	0	0	3	4
6	0	0	0	0	0	0	0	0
7								0

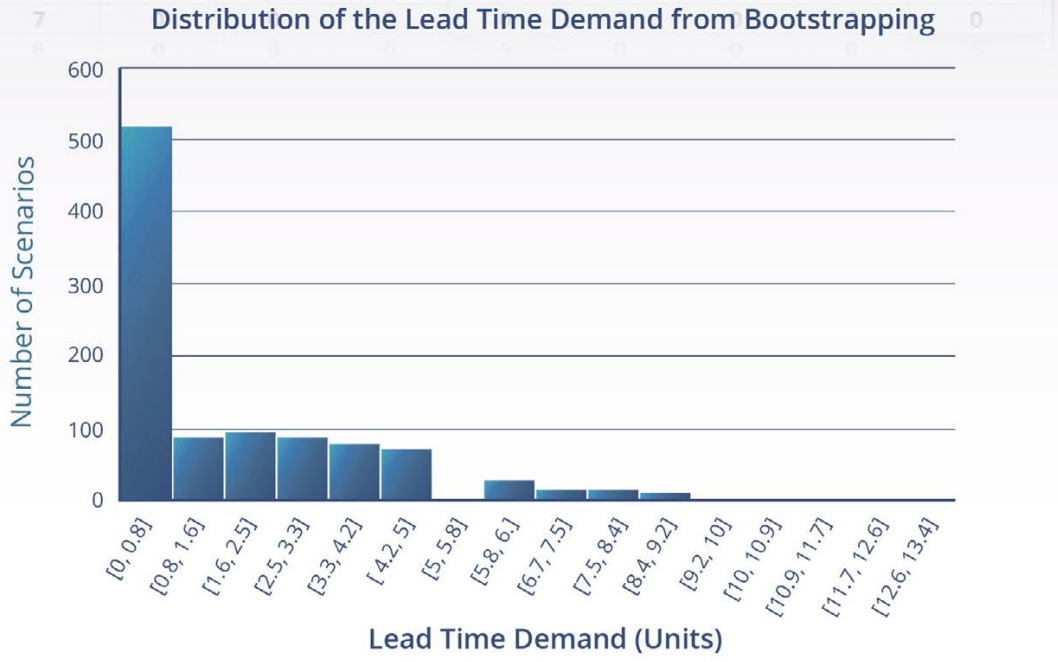

Figure 1: Using the bootstrap to estimate the distribution of lead time demand

In Gen2, Smart has significantly extended the core approach to deal with demand histories of greater complexity, including trends and multiple levels of seasonality.

Monte Carlo simulation

This method was invented as part of the Manhattan Project that produced the first atomic bomb. When calculations get too complex, we can substitute computation for theory and get practical answers to practical problems. In the engineering world, almost all new complex systems are first simulated to see if they will work before big bucks are spent on making and deploying them.

Gen2 generates demand and lead time scenarios using the bootstrap and feeds them to a digital model of the inventory control system. Monte Carlo simulation is then used to calculate the consequences of the input scenarios on daily on-hand inventory, lost or backordered units, and other metrics. From this information, Gen2 calculates metrics such as predicted service levels, fill rates, shortage costs, and inventory value. Since each demand scenario is random and has random consequences, the analysis is repeated many times and results averaged to arrive at final estimates.

Figure 2 shows an example of one daily demand scenario and the associated number of units on hand each day over the first 30 days of a 365-day simulation. In this example, demand has a different average on each day of the week, the system uses the (Min, Max) policy, lead time is random, and stockouts create backorders. Note that there is a stockout at the end of the simulated month in the scenario pictured.

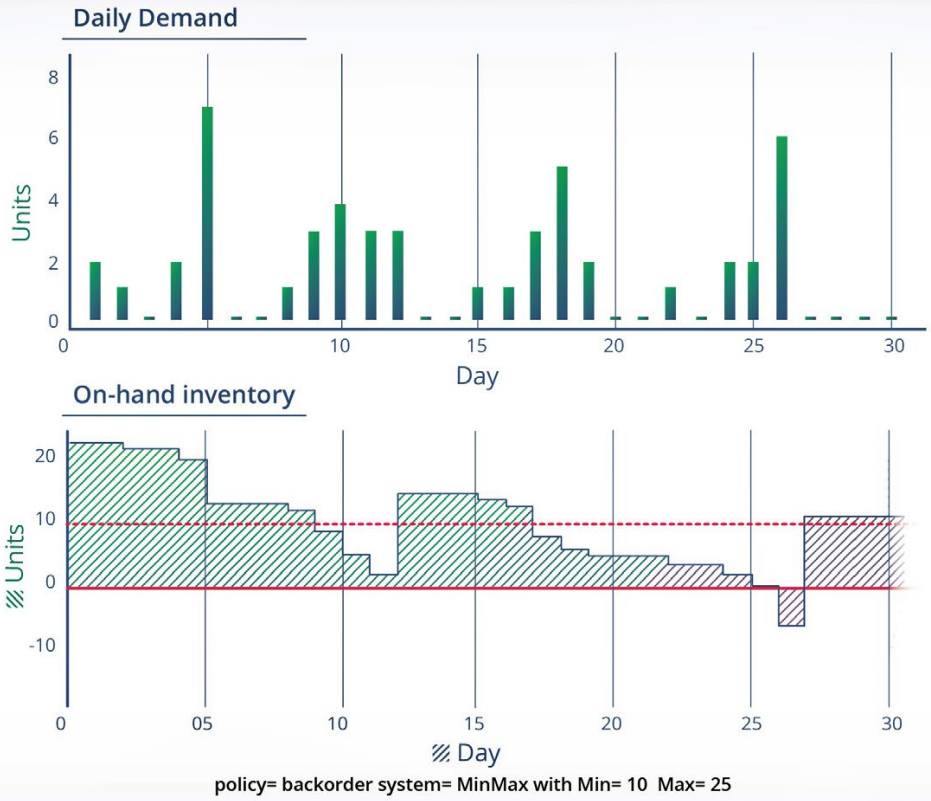

Figure 2: Part of a simulation showing daily item demand and daily on-hand inventory

We call the predictive estimates of metrics Key Performance Predictions (KPP's) to contrast them with the more familiar Key Performance Indicators (KPI's). Whereas KPI's are computed by measuring past performance, KPP's are predictions of how various planning moves are likely to play out.

Key Performance Predictions from Gen2 are more accurate as a result of the detailed daily simulations. KPP's are invaluable for informing proactive corrective actions such as targeted stock increases, expedites to avoid order cancellations, or transfers to avoid overstocks.

GEN

Advantages

We now describe the advantages of Gen2's forecasting innovations that are only available in Smart Inventory Planning & Optimization Demand Planner (Smart IP&O) and used to drive demand and inventory forecasts and optimizing stocking targets.

1. Gen2 is easy to understand, data-light, and more accessible
2. Gen2 Forecasting exploits daily data to guide near-term course corrections
3. Gen2 Forecasting facilitates new product forecasting
4. Gen2 Forecasting handles more complex seasonality
5. Gen2 Forecasting provides better estimates of forecast uncertainty
6. Gen2's Machine Learning automatically uses only relevant recent data
7. Gen 2 Inventory Optimization fully exploits simulation
8. Gen2 Inventory Optimization automatically integrates forecasts from Gen2 Forecasting
9. Gen2 Inventory Optimization models several inventory policies
10. Gen2 Inventory Optimization accounts for supplier uncertainty
11. Gen2 Inventory Optimization outputs more key performance indicators
12. Gen2 Forecasting is available in several ways
13. Gen2 comprises a bundle of applications

Smart Software, Inc. Four Hill Road, Belmont, MA 02478 t: 617-489-2743 www.smartcorp.com

1

Gen2 is easy to understand, data-light, and more accessible

Many competing approaches impose very significant data collection and analysis burdens on users. The level of investment – both in terms of external consulting costs and internal management effort – are just too significant for a typical inventory-carrying business.

Machine Learning forecasting and Multi-Echelon Inventory Optimization approaches are examples. Both require extensive model training and testing on very large data sets. Implementing these approaches requires considerable statistical modeling and data science expertise and hundreds if not thousands of hours of internal and external billable consulting time. Without these massive investments, the models cannot be applied at scale to generate outputs. The complex final models are mysterious black boxes not understood by the business, and their many parameters are often only configurable by software consultants.

On the other hand, Gen2 makes minimal assumptions about the data and instead uses the power of simulation to make predictions. All that is needed is a record of daily transactions (what you sold/shipped to whom and when, and what you received from whom and when). Being data-light enables faster setup and eliminates the need for data scientists and expensive consultants to fine-tune and configure the models. Gen2 is universally accessible to companies of all sizes yet is has the power to handle the toughest forecasting challenges.

2

Gen2 Forecasting exploits daily data to guide near-term course corrections

One of the most significant challenges facing demand forecasters is the increasing pace of business. Most businesses tend to run planning cycles at a monthly tempo, so Gen1 was sufficient to support the processes of our customers.

Now, though, the speed of business has combined with SKU proliferation to thin out demand streams, giving more items' demand an intermittent character making it harder to forecast. Frequently changing lead times and demand patterns means 4–5-week long consensus forecast processes can't keep up. Near term trends over days 1 to 30 are lost, along with the ability to make near-term course corrections.

Gen2 forecasting remedies this by basing all its calculations on daily data. Now customers can simulate daily demands over any future time frame, enabling near-term course corrections such as proactive adjustments to reorder levels or expedites that can now be captured 4-5 weeks sooner than otherwise. Daily forecasts are easily aggregated into weekly and monthly buckets to help drive longer term planning efforts.

3

Gen2 Forecasting facilitates new product forecasting

New product forecasting also benefits from daily-level analysis. In today's age of SKU proliferation, this is a significant advantage. The most common method for forecasting new items is to rely on forecasting "like" items. This isn't scalable since hours of time must be spent adjusting forecasts manually to account for clear and obvious differences between the new item and the like item. Traditional monthly "like item" forecasting requires many months of actuals before being able to forecast the new item using its own history. Gen2's advantage is obvious; by harnessing daily demand you can generate accurate forecasts with just 30 days of history. The modeling accounts for day of the week and week of the month seasonality and short-term trends.

4

Gen2 Forecasting handles more complex seasonality

Gen1 included algorithms for seasonal data at the weekly and monthly levels of aggregation. Such seasonal influences are still in play at the daily level, but now they often occur in combination. It is not uncommon for item demand to be influenced by both day-of-week and week-of-month seasonality, or both day-of-week and month-of-year seasonality.

If aggregated to a monthly level, demand with both daily and monthly seasonality appears to have only monthly seasonality. Masking daily seasonality increases forecast error and results in poor inventory management. Accordingly, Gen2 Forecasting is designed to detect and exploit multiple levels of seasonality. The default configuration supports two levels of seasonality and can be expanded to include additional levels of seasonality (e.g., hour-of-day or quarter of the year) when necessary.

5

Gen2 Forecasting provides better estimates of forecast uncertainty

We have always preached the importance of assessing forecast uncertainty. Some forecasters (or, more often, the rest of the business) wrongly construe measures of forecast uncertainty as something negative. Accepting that some degree of forecast error is inevitable, savvy professionals know that estimates of forecast error are essential for the risk management side of demand forecasting. Previously, extrapolative forecasting algorithms had consideration of forecast uncertainty "bolted on." Usually, this meant hanging a Normal curve (or "bell-shaped") distribution around the forecasts. Gen2 Forecasting does much better by basing all its forecasts and forecast uncertainty calculations on bootstrap demand scenarios. This approach does not depend on unrealistic theoretical assumptions to compute accurate prediction intervals even when data are very skewed, intermittent, or even multimodal (i.e., have distributions with "two humps" or more). The impact is far more accurate fine-grained demand forecasts and estimates of the required inventory needed to achieve target service levels. Theoretical assumptions about the shape of the demand distribution often meant that a company targeting a 99% service level might only be able to achieve something considerably lower, such as 86%. This would mean stocking out 14% of the time instead of 1% - a huge miss.

Gen1 pioneered by introducing patented bootstrapping methods for intermittent demand. Gen2 goes considerably further by using patent-pending methods that deal with the complications of seasonality and trend.

6

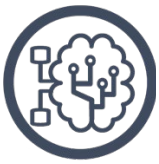

Gen2's Machine Learning automatically uses only relevant recent data

Demand can shift on a dime in so many ways: higher or lower average, more or less volatility, more or less intermittency, etc. If this happens, older observations represent obsolete economic "regimes" and bias forecasts away from current reality. The phenomenon of "regime change" puts an asterisk on the truism that "more data are always better."

Gen2 has built-in machine learning algorithms for automatic "regime change detection" that can detect significant changes in the character of an item's demand and exclude data from all but the current regime. Trying to do this job manually for thousands of items is wasteful, error-prone, and unnecessary because the regime change detection feature is fast and automatic. Alternatively, trying one-size fits all approaches such as "using the most recent 24 months of data" results in using the wrong amount of historical data for many items. Gen2's Machine Learning models mass-customize the decision of which historical data to use, yielding more accurate forecasts for all items.

7

Gen 2 Inventory Optimization fully exploits simulation

Gen2 provides the inventory manager with several outputs: (1) Estimation of the performance of current inventory operating policies, such as current reorder points and order quantities, (2) prediction of the performance of user-specified changes to current policies for individual items, and (3) suggestion to individually optimize policies for any number of items. Textbook inventory optimization runs on probability math that works only for idealized situations. Gen1 provided a big step up by replacing much of the simplified theory with fact-based Monte Carlo simulation.

Gen2 goes all-in on simulation-based design and analysis of inventory policies, allowing an unlimited degree of realism in the analysis. (Figure 2 shows a day-by-day simulation that illustrates the level of detail used in the analysis.) This will enable customers to realistically simulate stock on hand over any future time, helping target exactly when they are most likely to be at risk so corrective actions may be taken. Managing SKU proliferation with old items being regularly phased out with new ones is much more effective in Gen2.

8

Gen2 Inventory Optimization automatically integrates forecasts from Gen2 Forecasting

Users of SDP can incorporate their business judgment into the forecasting process through overrides of the statistical forecasts; these adjusted forecasts then automatically shape the demand scenarios that drive Gen2 Inventory Optimization. Doing so ensures promotional schedules and business overrides are incorporated while benefiting from daily forecasts over the near-term (1-30 days) that would otherwise be ignored in a typically laborious monthly consensus process. However, you always have the default option of using the new bootstrap methods described above.

9

Gen2 Inventory Optimization models several inventory policies

Gen1 analyzed two of the most popular inventory control policies (Reorder point/Order quantity and Min/Max) and assumed that stockouts always resulted in lost sales, as is often true. Gen2 models two types of stockout consequences: both lost sales and backorders. This results in more accurate KPP's for projected shortage costs, service levels, inventory values, etc. Gen2 also allows consideration of additional control policies, such as Scheduled Receipts/Blanket Orders, in which the buyer places a large annual order and accepts deliveries of a minimum quantity at agreed-upon intervals, such as every other week.

10

Gen2 Inventory Optimization accounts for supplier uncertainty

Demand variability is not the only source of difficulty in inventory management: supply variability can also be important. Gen2 models randomness in replenishment lead times, which has a major impact on KPP's and should not be ignored. For example, we used Gen2 to quickly simulate 1,000 years of operation of two systems. In System A, the supplier always fulfilled orders in exactly 14 days. In System B, the supplier lead time was a 50:50 mix of 7 days and 21 days. Both scenarios had the same average replenishment lead time of 14 days, but the performance of the system with random lead times was notably worse, having a 10% lower fill rate, a 15% lower service level, a 3% higher on hand inventory, and a 263% higher average operating cost (defined as the sum of holding, ordering and shortage costs).

11

Gen2 Inventory Optimization outputs more key performance indicators

Gen1 output a number of KPP's, including Fill Rate, Service Level, Average Operating Cost, and Inventory Investment. Gen2 extends the list to include Item Availability and Design Feasibility. Item Availability is the probability that there is at least one unit of an item in stock on any day. Availability is most relevant for spare parts, since it measures ability to respond to urgent needs; in the military, it links spare parts to platform readiness. Design Feasibility is the probability that a given inventory design can avoid a "stockout spiral" in which supply never catches up with demand and stockouts follow stockouts. Feasibility is most useful when users are improvising temporary changes to inventory control parameters, as when trying to cut inventory by lowering reorder points and order quantities. An infeasible design may work in the short run but be so fragile as to crash into a stockout spiral after a single unluckily-timed demand.

12

Gen2 Forecasting is available in several ways

The configuration options in the software platform give users the freedom to slowly adapt their forecasting process to Gen2 and not create a disruptive transition. Customers happy with Gen1 can opt to continue to use Gen1 forecasts, switch when they are ready, opt for a hybrid approach, or immediately go all-in on Gen2.

13

Gen2 comprises a bundle of applications

Gen2 is the latest incarnation of Smart Inventory Planning and Optimization. It comes as a set of three distinct applications. They can be used separately or together. **Smart Demand Planner** ("SDP") provides forecasts. **Smart Inventory Optimization** ("SIO") calculates optimal inventory control settings and Key Performance Predictions; it includes its own automatic forecasting capability or accepts user-adjusted inputs from SDP. **Smart Operational Analytics** ("SOA") provides business intelligence reports, including the ability to easily create your own custom reports.

GEN

A step change in capability and benefits

Smart Software was created four decades ago by Charles Smart, Nelson Hartunian, and Thomas Willemain, our visionary founders. Their pioneering work produced the first-ever automatic statistical forecasting system for the personal computer, a patented APICS award-winning method for intermittent demand planning, and other innovations. Gen2 continues Smart Software's innovative tradition with new analytical technologies. With Gen2, Smart Software substantially expands its probabilistic modeling technology, translating it into transformative new solutions that better position customers to reduce inventory, improve customer service levels, and increase operational efficiency.

www.smartcorp.com

Phone : 617-489-2743

Four Hill Road, Belmont,

Massachusetts 02478, US