

EVALUATING A BAYESIAN APPROACH TO DEMAND FORECASTING WITH SIMULATION

Randolph L. Bradley

Supply Chain Management
The Boeing Company
PO Box 516
St. Louis, MO 63166, USA

Jennifer J. Bergman
James S. Noble
Ronald G. McGarvey

Industrial and Mfg Systems Engineering
University of Missouri
E3437 Lafferre Hall
Columbia, MO 65211, USA

ABSTRACT

At The Boeing Company, stock levels for maintenance spares with substantial lead times must be established before fielding new aircraft designs. Initial calculations use mean time between demand estimates developed by the engineering department. After sufficient operating hours, stock levels are recalculated using statistical forecasts of maintenance history. A Bayesian forecasting method was developed to revise engineering estimates in light of actual demand on new aircraft programs.

Three forecasting methods were evaluated: Engineering Estimates, traditional Statistical Forecasting, and Bayes' Rule. Stock levels were established using inventory optimization, and fill rate performance was evaluated using warehouse simulation. The proposed Bayesian approach outperforms the other methods, enabling the inventory optimization model to establish stock levels that achieve higher fill rate, resulting in better initial inventory investment decisions.

This paper's contribution is comparing spares forecasting approaches for a well-defined set of airplane parts using a carefully constructed inventory optimization and simulation test environment.

1 INTRODUCTION

The Boeing Company manages aerospace service parts for customers through Performance Based Logistics (PBL) programs, in which the company holds inventory in order to guarantee an agreed upon service level. If the inventory of service parts, which comprises the most costly portion of these PBLs, is insufficient then contractual goals may be missed and aircraft availability impacted. An effective stocking policy requires an accurate forecast of future demand. However, most service parts experience slow-moving or intermittent demand, which challenges traditional forecasting methods, in addition to having long lead times.

New programs need to establish stock levels for maintenance spares with long lead times before aircraft are fielded. **Engineering Estimates**, an engineer's prediction of mean time between demand (MTBD), are utilized before forecasts based on demand history which does not exist yet. This technique is also called Judgmental Forecasting. Sources of these estimates include supplier estimates, reliability analysis, and comparisons with similar equipment. When managing large numbers of parts, manual adjustments to demand become difficult, and a systematic approach is necessary.

After a significant amount of demand data is collected, **Statistical Forecasting** is employed. This method assumes that past behavior represents future demand; however, initial demand can vary from long

term experience. Thus, there is concern that this current method may lead to poor performance early in the life of an aircraft program. As explained below in Section 2.4, The Poisson Assumption, this study assumes that demand follows a Poisson process in which demand rate is the number of demands occurring in a given time interval. Hence,

$$\text{Demand Rate } (\lambda) \text{ using Statistical Forecasting: } \lambda = \frac{\text{Number of demands}}{\text{Number of operating hours}}$$

$$\text{Mean Time Between Demand using Statistical Forecasting: } \frac{1}{\lambda} = \frac{\text{Number of operating hours}}{\text{Number of demands}}$$

Bayesian Forecasting was developed by Thomas Bayes (1701-1761) to evaluate how people update religious beliefs, and independently rediscovered by Pierre-Simon Laplace (1749-1827). Bayes' Rule allows us to update our initial belief with new information, resulting in a new and improved belief (McGrayne 2012). Muñoz et al. (2013) present a Bayesian framework to estimate stock levels based on simulation experiments where uncertainty exists in the demand forecast; we instead use Bayes' to revise the demand forecast itself. We both conclude that Bayes' is particularly relevant with few observations.

Bayes' Rule provides an intelligent way of combining prior knowledge (such as Engineering Estimates) with observed data (such as actual demands). Bayes' Rules is commonly expressed as the probability of prior belief A given new knowledge B:

$$\text{Bayes' Rule: } P\left(\frac{A}{B}\right) = \frac{P(A) * P\left(\frac{B}{A}\right)}{P(B)}$$

The Bayesian forecasting approach learns from observed demand, handles increasing operating hours occurring on new aircraft programs, and gives credence to the original engineering estimates. The following formula, derived by Bergman, et. al. (2015), applies Bayes' rule to demand forecasting:

$$\text{Mean Time Between Demand using Bayes' Rule: } MTBD = \frac{1}{\lambda} = \frac{r+n}{1 + \sum_{i=1}^n x_i}, \text{ where}$$

- λ : The unknown demand rate lambda, defined as 1/mean time between demand (MTBD)
- r : Engineering Estimate (mean time between demand)
- n : The number of operating hours in observed data
- x : The number of demands in observed data over the operating period

Is one of these three methods of demand forecasting clearly superior for new programs? In Section 2, we describe an aircraft scenario for evaluating these methods using inventory optimization and simulation. In Section 3, we validate the simulation model. In Section 4, we evaluate the forecasting methods. Section 5 describes next steps for implementing Bayes' Rule. The results are summarized in Section 6.

2 EVALUATING BAYES' RULE THROUGH A CASE STUDY

A case study was developed to evaluate the proposed Bayesian model for estimating spare parts demand on new aircraft programs. The case study is based on the first three and a quarter years of maintenance data for a new international aircraft tanker program comprised of four aircraft. The three forecasting methods compared in the case study are Engineering Estimates (judgmental forecasting), Statistical Forecasting (the current method, which calculates MTBD from historical data and uses future flight hours as a causal factor to forecast requirements), and Bayes' Rule (combining Engineering Estimates with observed data).

The case study was conducted by (a) determining stock levels using an inventory optimization model in order to understand each methods' impact on inventory, and by (b) evaluating these stock levels in a warehouse simulation model in order to understand each methods' impact on fill rate. The impact on inventory was measured by the value of the required stock investment. The impact on fill rate was

measured by the percentage of parts ordered which were filled from on-hand inventory. This iterative process of combining inventory optimization and simulation was described by Bradley and Goentzel (2012).

2.1 Scenario

The baseline scenario models four aircraft operated by an international tanker program, flying 75 flight hours per aircraft per month. A commercial inventory optimization model is used to establish the stock levels necessary to achieve an 80% fill rate goal for groups of parts. The majority of items, and the bulk of the inventory investment, are for parts which can be repaired; this evaluation focuses on the grouping of 239 unique repairable parts. The inventory optimization model was configured to analyze engineering estimates, along with historical demand data from March 2011 through March 2014, in order to estimate the stock levels required to support 80% fill rate beginning in April 2014.

2.2 Inventory Optimization

The use of multi-echelon inventory optimization for service parts, characterized by low demand probabilities, high cost, and high priority for service measured by “response time service levels,” is described by Cohen, Kleindorfer, and Lee (2006). This inventory optimization technique is embodied in the commercially available Service Planning and Optimization (SPO) software developed by MCA Solutions of Philadelphia, PA, which we used to compute stock levels. MCA Solutions was acquired by PTC of Needham, MA in 2013. Inputs for this model include Mean Time Between Demand (MTBD), repair time to fix a broken part, condemnation rate, procurement lead time to buy a new part, and unit price.

The inventory optimization model determines the mix of parts at stocking locations. The objective function minimizes inventory investment cost subject to achieving the desired fill rate goal. Demand in the model follows a Poisson process, and fill rate is computed using this assumption, similar to the marginal analysis algorithm to reduce backorders found in Sherbrooke (2004). Sherbrooke (2004) proves that this algorithm produces an optimal backorder-versus-cost curve. This approach uses one value in each step of the algorithm to determine whether the next part should be stocked. This value is equal to the increase in overall effectiveness achieved when another unit of an item is bought. In other words, select the next part that results in the greatest “bang for the buck.” This incremental approach is termed marginal analysis.

This case study is based on data for a program that manages only certain parts on the aircraft, making fill rate the appropriate inventory optimization goal. If the program managed all parts, inventory could be optimized to an aircraft availability goal. Optimizing inventory to an availability goal is common in a multi-indenture (parent-child relationships between service parts), multi-echelon (multiple levels of maintenance capability) environment where multiple stocking locations supporting a fleet of equipment, which may have requirements that vary by operating location.

The inventory optimization model was exercised to set stock levels under three cases, which differed only in MTBD for the parts in the data set:

- Demand specified by Engineering Estimates.
- Demand estimate from the Statistical Forecasting engine in the inventory optimization model.
- Demand estimated by revising engineering estimates in light of actual demand using Bayes’ Rule.

A number of business rules were incorporated into the current inventory optimization model, which are reflected in the results for Engineering Estimates and Statistical Forecasting:

- Set stock levels for 27 parts with demand, but missing cost, to 80% fill rate each.

- Ensure a minimum 90% fill rate for 30 critical parts.
- Ensure that each part is assigned a minimum stock level equal to the total forecast over the effective lead time for the part. The effective lead time is a weighted average of the repair time * percent of parts repaired + the procurement lead time * percent of parts condemned.

2.3 Warehouse Simulation

The Boeing Advanced Logistics ANalysis Capabilities Environment (BALANCE) is a discrete-event simulation model developed at The Boeing Company and written in the ExtendSim language. This model, described by Saylor and Dailey (2010), was modified (a) to accept a user defined cumulative probability density function describing an empirical demand distribution, and (b) to allow stock levels and the empirical demand distribution to be updated periodically, in this case annually. These modifications enable multi-year simulations with varying stock levels and demand distributions by period. The data driven aspects of the BALANCE simulation are described by Diamond, et. al. (2010). The reader is referred to these companion articles for additional details.

The BALANCE simulation model has two components: supply chain and warehouse. The supply chain portion models equipment operations and creates demand for spare parts based on either the part's mean time between demand (MTBD) or an empirical demand distribution. Failed parts are either repaired in the repair turnaround time, or condemned with a replacement delivered lead time away. The warehouse component supplies replacement parts from stocking locations, and then orders up to stock level when inventory position drops to or below re-order point. The stock level and re-order point for each part are set with a separate inventory optimization model. Statistics such as fill rate over time are calculated.

The simulation process involves these steps:

- Establish order-point s and order-up-to-level S for a (s, S) continuous review inventory system, using an inventory optimization model to achieve the desired fill rate goal at minimum cost.
- For repairable parts, orders are placed in quantities of one, so our inventory policy is identical to an order-point s and order-quantity Q policy (s, Q) with a fixed order quantity Q of one.
- Run multiple simulations, representing different views of how a typical year might play out.
- Calculate fill rate, the fraction of demand over the duration of the simulation that is met from stock on hand without backorders, as $Fill\ Rate = \frac{Demands\ filled\ from\ stock}{Total\ demands}$. This referred to as type-2 service level by Muñoz et al. (2013), and in chapter 7 of Silver, Pyke, and Peterson (1998).
- Compare different scenarios through statistical evaluation of multiple runs using Student's t-test to determine our confidence in achieving the desired fill rate goal over time.

2.4 The Poisson Assumption

We assume that demand follows a Poisson distribution for both the inventory optimization model and the demand generator in the simulation. A Poisson process calculates the number of demands occurring over a time interval. It assumes that the random time between consecutive events is independent with an exponential distribution, and that variance equals mean. Two reasons motivate using the Poisson assumption. First, demand is expected to rise as aircraft are delivered and monthly fleet flight hours increase. Time series techniques based on past demand history would result in a lagging forecast that would understate future requirements. Flight hours are therefore used as a causal factor. Future demand is simply calculated as expected operating hours divided by MTBD. Thus, for a fleet scheduled to fly 1,000 operating hours, and a part with MTBD of 200 hours, we would expect five (5) demands. Second, while current data collection systems capture operating hours per aircraft over time, they do not capture time between demand by part. Without knowing the interval between demand on a part by part basis, we

cannot characterize the demand distribution. We therefore calculate the mean and assume that variance equals mean. Thus, a part removed and replaced five (5) times over 1,000 flight hours would have a calculated MTBD of 200 hours.

The Statistical Forecasting engine in the inventory optimization model is capable of assuming different demand distributions for time series analysis based on the ratio of variance to the mean. The model can also be configured to use a blend of exponential smoothing and demands per operating hour to forecast demand. For this study, the engine was configured to run with 100% causal forecasting, in which demand is based on calculated MTBD and future operating hours.

The Bayes' Rule calculation also assumes Poisson distribution, but only because demand data is not collected in a way that enables determining that actual time between demands on a part-by-part basis.

3 VALIDATING WAREHOUSE SIMULATION

3.1 Methodology

The warehouse simulation was validated by checking whether stock levels generated by the inventory optimization model achieved the expected fill rate goal. The process to validate the simulation model is illustrated Figure 1 and is based on the approach described by Bradley and Goentzel (2012).

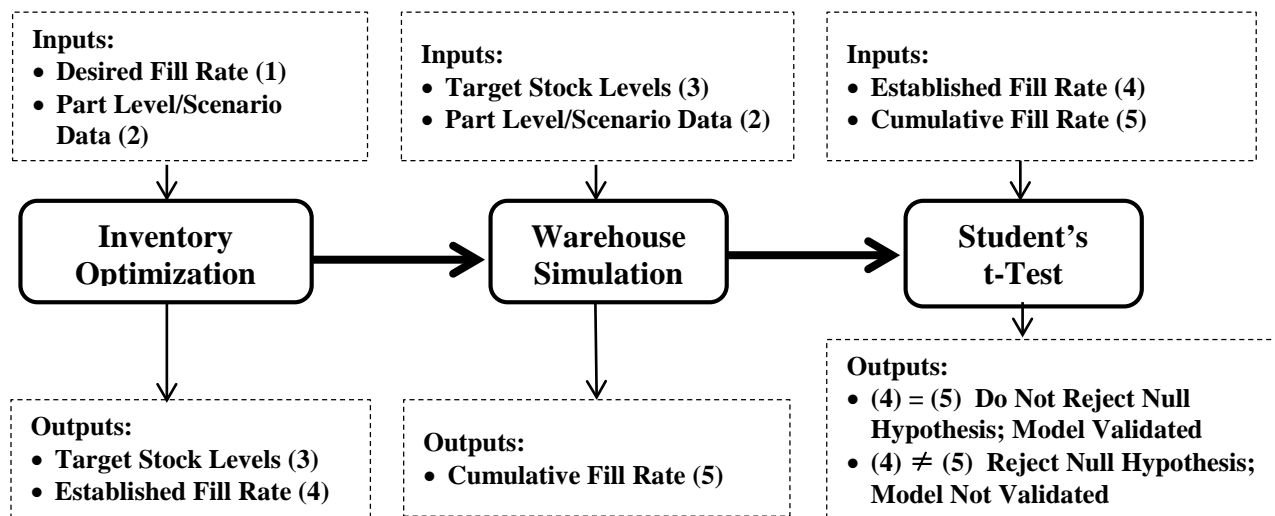


Figure 1: The methodology for testing the null hypothesis that the mean of a number of warehouse simulations equals the expected fill rate established in an inventory optimization model using a Student's t-Test. The numbers in parentheses refer to specific input and output data elements.

The inventory optimization model minimizes inventory investment subject to achieving a user specified fill rate goal (80% in this study). These target stock levels are inputs for the simulation, along with the operating scenario (number of aircraft and flight hours per aircraft) and part level data (MTBD, price, repair turnaround time, condemnation rate, and procurement lead time). For validation, a ten year steady state simulation is conducted. The model is steady state because the operating scenario remains constant each month. If the simulation fill rate is equivalent to the established fill rate goal from the inventory optimization model, the simulation is considered validated. The null hypothesis, that the mean fill rate of ten warehouse simulations equals the fill rate achieved by the inventory optimization model, will be evaluated using a Student's t-Test at a 95% confidence level.

3.2 Validation

A validation is performed to verify whether the stock levels generated by the inventory optimization model achieve the expected fill rate in the warehouse simulation. A Student's t-Test is used to determine whether the simulation fill rate equals that of the optimizer.

The inventory optimization model is configured to compute target stock levels for parts grouped in the repairable parts network. The objective function is to minimize inventory investment and back orders subject to the constraint of achieving an 80% target fill rate goal. The optimizer overshoot and achieved an 80.58% fill rate. Next, these computed stock levels are loaded in the simulation. The operating scenario and part level data remain the same. Both the optimization and simulation assume Poisson demand. Ten simulations are run, as shown in Table 1, resulting in a mean cumulative fill rate of 80.45%.

Table 1: Simulated Cumulative Fill Rate over Ten Years.

Simulation Run	Cumulative Fill Rate	Simulation Run	Cumulative Fill Rate
1	81.4%	6	79.2%
2	82.0%	7	81.5%
3	80.6%	8	80.3%
4	79.6%	9	79.8%
5	78.6%	10	81.5%

A Student's t-Test is used to test the null hypothesis that the mean cumulative fill rate of ten simulations (80.45%) is equivalent to the fill rate achieved by the optimizer (80.58%). The test was run at the 5% significance level, meaning that the null hypothesis is only rejected when it is true 5% of the time; this is termed a Type I error in statistics. At a 95% confidence level, there is no reason to reject the null hypothesis. Thus, the warehouse simulation fill rate (80.45%) and the inventory optimization model fill rate (80.58%) are considered equivalent, and the simulation model is validated. The inputs and outputs for the Student's t-Test are illustrated in Table 2.

Table 2: Student's t-Test Inputs (left) and Outputs (right) for the Null Hypothesis that Simulated Fill Rate is Equal to the Optimizer Fill Rate of 80.58%.

Input Name	Input Value	Output Name	Output Value
Null Hypothesis (H_0)	$\mu = .0858$	Test statistic	-0.3689
Alternative Hypothesis (H_1)	$\mu \neq .0858$	P-value	0.721
Level of Significance (α)	0.05	Reject Null Hypothesis?	P-value (.721) > Level of Significance (α) DO NOT Reject Null Hypothesis ($\mu = .0858$)
Sample Size (n)	10		
Sample Mean (\bar{x})	0.804467		
Sample Standard Deviation (s)	0.011428		

4 EVALUATING FORECASTING METHODS

Stock levels based on Engineering Estimates, Statistical Forecasting, and Bayes' Rule were analyzed to determine whether there was a business case for implementing Bayes' Rule on the international aircraft tanker program. The inventory optimization model determined the stock levels required to achieve an 80% fill rate goal, and simulation estimated the fill rate achievable given the demand that actually occurred. Differences were evaluated by assessing the impact to inventory position (financial impact) using inventory optimization, and the impact to fill rate (performance impact) using simulation.

4.1 Evaluating Fill Rate for Engineering Estimates using Warehouse Simulation

Since The Boeing Company relies on engineering estimates to forecast future demand for new aircraft programs, it is important to know how well these estimates support customers' operating requirements. In the upper left hand picture of Figure 2, Engineering Estimates were used to establish stock levels to support an 80% fill rate on the international aircraft tanker program.

To see how this mix of spare parts supported the fleet, the warehouse simulation model was run to evaluate the ability of these stock levels to support the actual demands experienced on the international tanker program between delivery of the first aircraft in March 2011 and March 2014. We assumed that "robbing" needed parts from the production line was not allowed. Referring to the upper right hand chart, the 90-day moving average is shown in red and the cumulative fill rate in blue. This stocking strategy averages only a 37.7% fill rate.

Given the low performance, it seemed prudent to validate the analysis. To validate the stock levels, the warehouse simulation was run again, this time assuming that demand followed a Poisson distribution with a mean of the engineering estimate. As shown in the lower left hand chart, the simulation achieved 80.45% fill rate, close to the 80.58% predicted by the optimization (which overshoot the 80% goal). So if demand had been the same as the engineering estimates, the stock levels would have supported the fleet.

To ensure that there was no issue with not hitting steady state, the lower right hand chart shows simulating the stock levels against an empirical demand distribution created from the actual 2011-2014 historical data to drive a ten year simulation, achieving 37.9% fill rate.

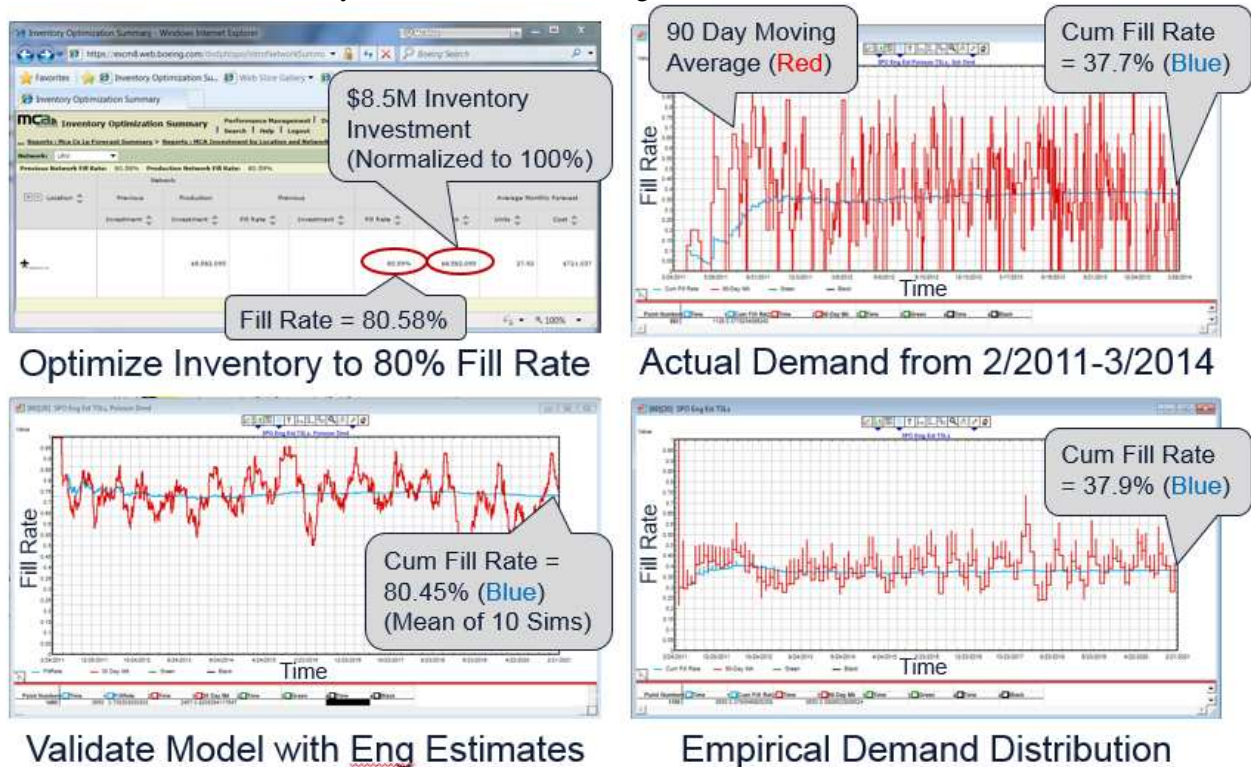


Figure 2: The fill rate achieved when the inventory was optimized to 80% fill rate (upper left) was validated by running a warehouse simulation in which the MTBD was set equal to the engineering estimates (lower left). The expected actual fill rate was then estimated by (a) evaluating the stock levels against three years of actual demand history (upper right), and (b) evaluating the stock levels against an empirical demand distribution based on actual demand history (lower right) in steady state.

The conclusions are that (a) Engineering Estimates are inaccurate and unlikely to support new fleets or flight test programs, and (b) there are periodic cycles in the 90-day moving average, or short term fill rate. Bayes' Rule fills the need to move away from Engineering Estimates.

4.2 Evaluating Inventory Investment and Fill Rate using Inventory Optimization

Stock levels were calculated using inventory optimization, varying only the Mean Time Between Demand (MTBD) between the three scenarios: Engineering Estimates, Statistical Forecasting, and Bayes' Rule. We analyzed annual inventory investment and year end fill performance to assess the business case for implementing Bayes' Rule. Figure 3, Inventory Investment and Fill Rate Over Time, is color coded to reflect required buys (yellow), existing inventory (green), and excess inventory (red).

In 2011, all scenarios rely on Engineering Estimates because it is a new program, and thus start with the same required inventory investment (yellow), normalized to 100%. The inventory optimization model was then run in an evaluation mode to estimate the full rate achievable when demand instead reflected the Bayesian forecast. The resulting fill rate estimate is represented by an airplane in Figure 3.

In 2012, the Statistical Forecasting engine continues to use Engineering Estimates due to business rules. Bayes' Rule requires an additional 25% investment (yellow) over previously purchased inventory (green).

In 2013, both Statistical Forecasting and Bayes' Rule result in revisions to stock levels. The former now requires an additional 25% investment (yellow).

In 2014, we again revise stock levels. Statistical Forecasting and Bayes' Rule now identify excess inventory (red) due to changing demand patterns.

In the reconciliation set, we address whether we are buying the same parts in different years, or different collections of parts. Since forecast accuracy is best for Bayes', we reconcile the Engineering Estimates and Statistical Forecasting to the Bayes' stock levels. For Engineering Estimates, we have the expected shortfall to close (yellow), indicating the cost of revising stock levels based on better estimates. For Statistical Forecasting, we have a small required buy (yellow) and only a modest amount of excess inventory (red) over Bayes', indicating that over time stock levels are similar to Bayes' on a part-by-part basis.

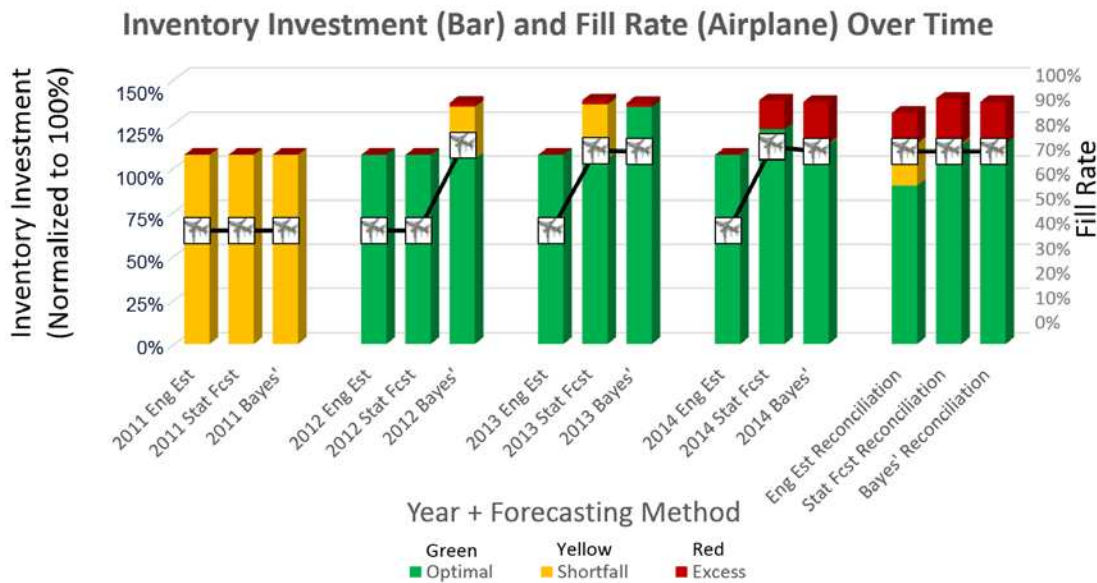


Figure 3: Bayes' Rule provides better early operating performance for a similar investment (vertical bars) in inventory when compared to Statistical Forecasting, while both techniques significantly outperform Engineering Estimates based on fill rate (airplane icon).

Both Statistical Forecasting and Bayes' Rule outperform Engineering Estimates in supporting new aircraft programs. While Statistical Forecasting and Bayes' Rule require similar inventory investments over time, Bayes' outperforms the current process after first fielding new aircraft because it incorporates the experience of early demands. Additional quantities of spare parts (yellow) may be needed when Engineering Estimates are revised, and management must budget for these periodic adjustments. Further, management can expect excess inventory (red) on new aircraft programs as demand estimates are refined over time, a consideration when including clauses for buying back excess inventory in contracts.

4.3 Inverse Transform Sampling of Historical Demand Data for Warehouse Simulation

The inverse transform sampling technique was used to take random samples from the empirical cumulative demand distribution for each part in the dataset, as shown in Figure 4. By drawing a random number between zero and one, representing the cumulative fill rate on the y-axis, one can read across on the x-axis to find the resulting demands per month.

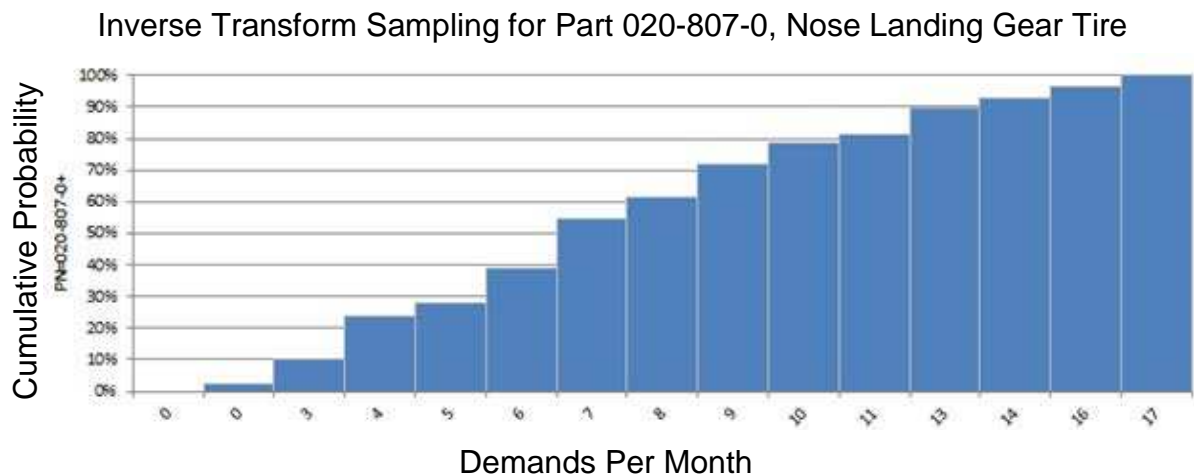


Figure 4: Inverse transform sampling for Part Number 020-807-0, the nose landing gear tire. The empirical demand distribution is sampled by drawing a random number between zero and one, and returning the corresponding demands per month from the cumulative distribution function (CDF).

Historical demand in the period of interest is represented by an empirical distribution. The example above reflects actual demand occurring in the third year of operations for one part. Inverse transform sampling is used instead of replaying actual demand because it allows running multiple random scenarios in order to calculate confidence intervals, and running steady-state scenarios representing future periods. The process used for inverse transform sampling, by time period, is:

1. Sum the historical demands per month for each part number, including months with zero demand.
2. Create a cumulative histogram with unequal bins for each part number where each bin corresponds to the actual monthly demand. The resulting histogram shows the cumulative probability of having 0 to n demands. This histogram is shown in Figure 4, where the x-axis is demands per month and the y-axis is the cumulative probability. Although histograms usually have equal bin sizes, for creating an empirical demand distribution, this histogram uses bins that are equal to the actual monthly demands that occur. The number of bins is equal to the number of unique values of demands per month (plus one if there are months with zero demand).
3. Each month, for each part, the demand generating function in the simulation draws a random number between 0 and 1 and looks up the cumulative probability to determine the corresponding

demands per month. In this way, the demand generating function randomly determines monthly demands per part following the same empirical demand distribution as the actual data.

- Repeat steps for each period.

4.4 Evaluating Fill Rate using Warehouse Simulation

Jackknifed datasets are frequently used to evaluate different forecasting methods. Demand from an earlier period is used to forecast demand from a later period. The method with the lowest forecast error is then used to forecast the entire dataset in order to predict the future, a statistical method also referred to as cross-validation. In this study, the jackknifing technique was employed to estimate the demand used in an inventory optimization model to set stock levels. Demand from the later validation period was used to evaluate fill rate. First mentioned by Tukey (1958), this technique is more fully described by Abdi and Williams (2010). The purpose of the jackknifing technique is to eliminate bias. For example, suppose a program started in 2012. The forecast for 2014 would be based on historical demand from 2012 to 2013. If the forecast of 2014 was compared with observed demand from 2012-2014, there would be a bias because the forecast and the observed data both included 2014. The jackknife approach uses observed data from 2012-2013 to forecast requirements for 2014, then evaluates that forecast against observed data from 2014 only. This eliminates the bias. Historical demand in the period of interest is thus split in two components to create the jackknifed forecast: (a) earlier demand for forecasting, and (b) later demand for evaluation.

In 2012, the Statistical Forecasting engine continues to use Engineering Estimates due to business rules, starting the year at 28% and ending the year at 42%. Business rules require collecting two years of actual demand history before switching from Engineering Estimates, and then only when there are three or more demands within a 12 month period. Bayes' Rule performs significantly better, also starting at 28% but ending the year at 70% fill rate, as simulated over time in Figure 5.

In 2013, Statistical Forecasting improves from 42% to 72%, while Bayes' improves from 70% to 78%, as both methods result in revised stock levels. With three or more demands in a 12 month period, Statistical Forecasting switches from Engineering Estimates to demand history.

In 2014, we again revise stock levels. Statistical Forecasting improves from 72% to 81%, and Bayes' Rule stays relatively constant from 78% to 82%.



Figure 5: Comparison of simulated fill rate when inventory is optimized based on Statistical Forecasting (top row) vs. when inventory is optimized based on Bayes' Rule (bottom row). Each box shows simulated fill rate at the start of the year and the end of the year.

5 NEXT STEPS

The near term goal is automation: create computer code to read tables containing engineering estimates and transactional demand history from the company standard inventory optimization model, calculate revised demand estimates using Bayes' Rule, and return these estimates to the optimizer to revise stock levels. The long term goal is to seamlessly integrate Bayes' Rule into the supply chain process.

Future research will be conducted to determine whether the Poisson demand assumption needs to be revised. Demand data collected includes operating hours per aircraft per month, and transactional order dates and order quantities by part. This only allows for calculating a single parameter (average operating hours per demand) with an assumed demand distribution. As a result, this research assumed the time between demand is exponentially distributed, which is equivalent to Poisson demand rate (demands per time). However, the literature finds that the Poisson distribution is a poor estimator of actual variance (Sherbrooke, 2004). A pilot to calculate the demand distribution based on observed demand could determine whether this Poisson assumption needs revision. The desired data set to do this would show the operating hours between demand for each failed asset. This could be obtained by time stamping the installation of asset Y on aircraft tail X and by time stamping the demand of asset Y on aircraft tail X. Then, the flight hours for aircraft tail X could be computed from these two time stamps and flight hour records. This type of data set could allow for a more appropriate representation of demand and could potentially increase the accuracy of predicted demand.

6 CONCLUSION

This paper's contribution is to evaluate methods for forecasting service parts for a new airplane model using an integrated inventory optimization and simulation environment. We evaluate a Bayesian approach to demand forecasting, showing it to be superior to both Engineering Estimates and Statistical Forecasting.

For a defined collection of repairable parts on an international tanker program at The Boeing Company, we demonstrated that when compared to Statistical Forecasting, Bayes' Rule provides better early operating performance as measured by fill rate, with similar inventory over a multi-year period as measured by investment cost. Engineering Estimates performed poorly, but were improved through revisions in light of observed demand using Bayes' Rule.

Our multi-year simulation showed that as demand estimates are refined over time, management should budget for periodic additional investments in inventory on new aircraft programs. As all forecasts contain forecast error, management should also budget for excess inventory building up over time.

In conclusion, we showed that Bayes' Rule revises stock levels more responsively than current business rules for forecasting on a new aircraft program, improving customer performance by improving fill rate during the first few years of operations.

REFERENCES

- Abdi, H., and L. Williams. 2010. "Jackknife." In *Encyclopedia of Research Design*, edited by N. Salkind, 656-661. Thousand Oaks, CA: SAGE Publications, Inc. <http://dx.doi.org/10.4135/9781412961288.n202>.
- Bergman, J., J. Noble, R. McGarvey, and R. Bradley. 2015. "A Bayesian Approach to Demand Forecasting for New Equipment Programs." In *Proceedings of the 25th International Conference on Flexible Automation and Intelligent Manufacturing (FAIM 2015)*. Wolverhampton, UK, June 23 – 26.
- Bradley, R., and J. Goentzel. 2012. "Simulating Tomorrow's Supply Chain Today." In *Proceedings of the 2012 Winter Simulation Conference*, edited by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A.M. Uhrmacher, 216:1-216:12. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc. <http://dx.doi.org/10.1109/WSC.2012.6465143>.

- Cohen, M. A., P. R. Kleindorfer, and H. L. Lee. 1986. "Optimal Stocking Policies for Low Usage Items in Multi-Echelon Inventory Systems." *Naval Research Logistics Quarterly* 33:17-38.
- Diamond, B., D. Krahl, A. Nastasi, and P. Tag. 2010. "ExtendSim Advanced Technology: Integrated Simulation Database." In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hukan, and E. Yücesan, 32-39. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc. <http://dx.doi.org/10.1109/WSC.2010.5679178>.
- Muñoz, D. F., D. G. Muñoz, and A. Ramírez-López. 2013. "On the Incorporation of Parameter Uncertainty for Inventory Management using Simulation." *International Transactions in Operational Research*, 20(4), 493-513. Oxford, UK: John Wiley & Sons Ltd.
- McGrayne, S. 2012. *The Theory That Would Not Die: How Bayes' Rule Cracked the Enigma Code, Hunted Down Russian Submarines, and Emerged Triumphant from Two Centuries of Controversy*. New Haven, CT: Yale University Press.
- Saylor, S., and J. Dailey. 2010. "Advanced Logistics Analysis Capabilities Environment." In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hukan, and E. Yücesan, 2138-2149. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc. <http://dx.doi.org/10.1109/WSC.2010.5678859>.
- Sherbrooke, C. C. 2004. *Optimal Inventory Modeling of Systems: Multi-Echelon Techniques*. 2nd ed. New York: Springer.
- Silver, E. A., D. F. Pyke, and R. Peterson. 1998. *Inventory Management and Production Planning and Scheduling*. 3rd ed. New York: John Wiley & Sons.
- Tukey, J. W. 1958. "Bias and Confidence in Not Quite Large Samples." Abstract, *Annals of Mathematical Statistics* 29: 614.

AUTHOR BIOGRAPHIES

RANDOLPH L. BRADLEY is a Technical Fellow in Supply Chain Management at The Boeing Company in Chicago, IL. He holds a MEng in Supply Chain Management from the Massachusetts Institute of Technology, and an MBA and BS in mathematics from Oklahoma State University. His research interests encompass requirements definition, statistical demand forecasting, strategic inventory optimization, supply chain simulation, tactical asset management, and warehouse management. His email address is randolph.l.bradley@boeing.com.

JENNIFER J. BERGMAN holds an MS and BS in Industrial Engineering from the University of Missouri. Her research interests include distribution network optimization, demand forecasting using Bayes' Rule, inventory optimization, and supply chain simulation. Her e-mail address is jjby95@mail.missouri.edu.

JAMES S. NOBLE is Professor at the Industrial and Manufacturing Systems Engineering Department at the University of Missouri. He holds a PhD and MS from Purdue University, and a BS from the University of Oklahoma. He is MU Site Director for the NSF I/UCRC Center for Excellence in Logistics and Distribution (CELDi). His research interests relate to integrated production systems, logistics and distribution. His e-mail address is noblej@missouri.edu.

RONALD G. MCGARVEY is an Assistant Professor with a joint appointment in the Industrial and Manufacturing Systems Engineering Department and at the Harry S. Truman School of Public Affairs. He holds a PhD from Pennsylvania State University. His primary research interest is in applied optimization and its applications to public policy and resource management. His e-mail address is mccgarveyr@missouri.edu.