Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

USING SIMULATION TO DETERMINE THE SAFETY STOCK LEVEL FOR INTERMITTENT DEMAND

Fredrik Persson

Department of Management and Engineering
Linköping University
Linköping, SWEDEN

Minna Axelsson

Siemens Industrial Turbomachinery
Slottsvägen 2-6
Finspång, SWEDEN

Fredrik Edlund
Christoffer Lanshed
Agnes Lindström
Frida Persson

Department of Science and Technology
Linköping University
Norrköping, SWEDEN

ABSTRACT

Safety stock calculations are difficult for products with intermittent demand, long production lead times, and high monetary values. Theoretically, forecasts can be used to reduce the need for safety stocks. A high precision forecast minimizes the need for safety stock and forecast evaluation measurements can be used to calculate the safety stock level. However, a more realistic determination of safety stock levels can be obtained by simulation. In this paper, simulation is used to model and experiment on a case with three end products in order to determine the relationship between safety stock levels and service levels. Also, a comparison is made with theoretically calculated safety stocks to see how well basic theoretical models for safety stock calculations fulfill the requirements of service level. The result is that simulation can provide a much more accurate determination of safety stock levels for intermittent demands than theoretical calculations.

1 INTRODUCTION

Determining the level of safety stock in inventory control systems is not a trivial problem. Demand is often treated as normally distributed with a high mean value and a moderate standard deviation. Safety stocks within these recurring and even demand patterns can be calculated to keep a certain fill rate or to keep a certain level of stock-out probability. Both these methods are well described in literature and used in industry (Hopp and Spearman 2008; Anupindi et al. 2014). There are other methods to calculate safety stock for items that instead show an intermittent and uneven demand. The size of the safety stock is then largely affected by the ability to forecast the demand. One method of intermittent demand forecasting based on historic data is the Croston Forecasting Method (Croston 1972). The forecast is used to calculate a safety stock level utilizing both fill rates and stock-out probabilities.

In some cases, products with intermittent demand also have long production lead times and large monetary values. Together with a desire from the customers to have short delivery lead times, the calculation of a safety stock level is a problem with a non-trivial solution. Using the Croston Forecasting

Method (Croston 1972) to forecast intermittent demand have several shortcomings that will be discussed throughout this paper.

Simulation that uses real data and real system logic is often regarded as a better alternative to determine the safety stock level of end products. Nevertheless, a simulation solution is often much more expensive both in time and money than going directly for a theoretical safety stock calculation. In this paper we use simulation to model and determine the safety stock levels of three end products with extremely high monetary value. The result is a relationship between safety stock levels and service levels. We also compare the simulated safety stock levels with a calculated safety stock. We utilize the Croston Forecasting Method to calculate a forecast and then use the forecast to calculate a safety stock level for each end product. The idea of using forecasting methods is that a precise forecast reduces the need for safety stock, and vice versa.

The purpose of this paper is therefore to test if a safety stock level that has been calculated with a high probability to avoid stock-outs can match the simulated safety stock levels for end products with intermittent demand. This is done being well aware of the fact that the underlying assumptions of a safety stock calculation with a high probability to avoid stock-outs are false for intermittent demand. Still, this is believed to be a practical solution. In this paper, we investigate how wrong the practical solution really is.

The company case that is used in this paper contains three end products that are today only Make-To-Order (MTO). A customer order drives the production and delivery lead times can be long, often longer than a year. There is however a new opportunity to deliver end products to a new and different customer segment where delivery lead times are much shorter. It is impossible to shorten production lead times, so to meet the new customer segment, some of the end products must be held in stock as a safety stock or speculation stock, moving towards a Make-To-Stock (MTS) system. However, due to the extreme high value of the end products, the amount of products held in stock must be very precise. Keeping one extra end product in stock can prove very costly.

Olhager and Persson (2008) use simulation in a teaching environment where the goal of a student project is to set, among other control parameters for manufacturing systems, a proper safety stock level for different end products. Simulation is also used in Schmidt, Hartman, and Nyhuis (2012) where different approaches of calculating safety stocks are evaluated. The study, however, seems inconclusive when no superior approach to safety stock calculation could be found. Hernandez-Ruiz et al. (2016) considers the special case of modular productions systems while investigating safety stock levels with simulation. These examples of simulations studies for safety stock investigations all lack the other special case of intermittent demand – a topic very hard to find.

The reminder of the paper is organized as follows; first, a theoretical background is given to forecasting and safety stock calculation. After that, the modelled system is described, following the conceptual model, computer model, verification and validation, and experimentation. Last, conclusions are made followed by an outlook into future work.

2 THEORETICAL BACKGROUND

Production companies that compete in the MTS markets often keep a safety stock to counteract stock-outs during time periods of higher demand than usual levels. Stock-outs means that since the inventory is empty, the company loses sales opportunities. A safety stock gives a certain service level towards the customers. The size of the safety stock can be determined in many different ways using different definitions. One common definition is to calculate the safety stock as "the probability that there will be no stock out within a time interval" (Anupindi et al. 2010; Olhager 2000), also known as SERV1. The calculation of the safety stock incorporates both the probability of stock outs, the standard deviation of forecast errors, and replenishment lead time. Safety stock I_S is calculated as

$$I_S = k\sigma_e L^\gamma \tag{1}$$

Where k is the safety factor, based on the Normal distribution (k = 1.65 for 95% probability) if demand is normally distributed, σ_e is the standard deviation of the forecast error per time period, and L^γ is the lead time in time periods to the power of γ which is a factor that is affected by internal correlation in the underlying time series the forecast, often set to γ = 0.5 (Olhager 2000).

Not all authors agree with Olhager (2000) where standard deviation of forecast error is used in the calculation of the safety stock; some argue instead that the standard deviation of the demand itself is sufficient to get the correct probability that there will be no stock-outs within a given time interval (Silver et al. 1998). However, a safety stock calculation based on forecast error is practically sound since it implies no need for safety stock if forecasts could be determined with an extremely high accuracy. This also implies a direct connection between the size of the safety stock and the accuracy of the forecast. When determining the safety stock in MTS, the question of how good the forecast is will always be dominating the discussions.

Forecasting is often done by looking at historical sales data and trying to predict the future. Different methods have been proposed such as time series calculations or time series component decomposition, i.e. finding patterns in time series. Moving averages or exponential smoothing can determine future forecasts mathematically with a high accuracy for products with high and frequent demand in all time periods. In the case where products have an intermittent demand pattern, forecasting becomes harder. Intermittent demand means that products are sold seldom, in small volumes, and without any apparent recurrence of demand. Seldom means that products are not sold in every time period used to forecast. Small volumes mean that when products are sold, it is often one or two products. Last, without apparent recurrence of demand means that products are sold in what seems to be a stochastic pattern which is hard to forecast. In all, products with intermittent demand are the slow movers that often take up a lot of shelf time, but when absent can incur large costs in a company. Silver (1981) defines intermittent demand as when both demand and the time between the demands occurrences are stochastic.

An early attempt to forecast intermittent demand was introduced by Croston (1972). Croston defines intermittent demand as zero demand in a number of time periods. In this forecast method, the time between sales and the size of the sale are treated as two independent variables. Simple exponential smoothing is done to update both variables and to create a forecast for the next sale. Thus, the demand forecast in time period *t* is denoted X_t when the demand occurs and the forecasted time between occurrences is denoted T_t valid for time period *t*. In the case where no real sales occur, both variables are constant:

$$X_{t+1} = X_t \tag{2}$$

$$T_{t+1} = T_t \tag{3}$$

When sales occur, both variables are updated with the real values of sales denoted $\hat{X}_t$ and the time between occurrences denoted $\hat{T}_t$:

$$X_{t+1} = X_t + \alpha\left(\hat{X}_t - X_t\right) \tag{4}$$

$$T_{t+1} = T_t + \beta\left(\hat{T}_t - T_t\right) \tag{5}$$

This update is done with exponential smoothing and factors α and β to weigh the recent sales and time between sales.

The method by Croston have been subject to a lot of criticism, by e.g. Syntetos and Boylan (2001) and later led to the development of the modified Croston by Levén and Segerstedt (2004). The modified Croston is also criticized of have a build in bias (Teunter and Sani 2009).

A more recent approach (than Croston 1972) is to use standard statistical tools such as bootstrapping. Bootstrapping is basically about random sampling with replacements. In Willemain, Smart, and Schwartz

(2004) bootstrapping is used to create a forecast based on historical data. Assume that each time period is represented by a ball as in the basic urn problem in probability (Dodge 2003). Each ball is given a number representing the sales in that time period based on historical data. The balls now represent the real sales with many balls having number 0 and others real sales data. Balls are picked from the urn in order to forecast the next time periods and for each new forecasted time period the ball is replaced into the urn. In this way, a forecast can be established following the real probabilities of sales in time and the exact size of the sale from historical data.

3 CASE SYSTEM DESCRIPTION

In this case, the company produces three different end products in the same factory. The delivery lead times to customers for the three specific products are equal to the production times, as in MTO. The aim is now to shorten the lead times and move to MTS, because of higher service requirements from some customers. The system starts with a predefined stock level for each product. When a customer makes an order, the product can be delivered directly from stock if the inventory level is large enough. Otherwise, if the inventory is empty but there is a product in the production system, undergoing production, with less time remaining than the customer lead time, the customer can be served. If the product cannot be delivered, the sales opportunity will be lost. When the safety stock is used, the production of a new product will be started, so that the safety stock will increase to its original level. If a customer makes an order of, say, 2 products, but only 1 product is in stock and the other in production, the production of two new products will not be started until the customer has received the whole order.

The system is modelled in Arena 14.7. Most of the input data were modelled in Matlab. For the simulation study as a whole, the methodology from Persson (2003) was used, see the following sub-chapters.

3.1 Conceptual model

There are two types of entities in the model. One type represents customers that are requesting products. The other type represents the different products. Activities are the production of each product and delivery lead times. Activities such as disturbances in the production process are excluded from the model as well as staff and equipment. The customers that are waiting to get their order will be placed in a queue. If the customer does not get the delivery of the order before the given amount of time it will be counted as a lost sale (lost customer). The production system will be operating 24 hours per day, 7 days per week, 52 weeks per year. In Table 1 is a list of details of what is included and excluded in the model.

Table 1: List of details in the model.

Component	Include/exclude	Justification
Entities		
Customers	Include	Key influence on throughput
Products	Include	Key influence on throughput
Activities		
Production lead times	Include	Key influence on throughput
Delivery lead times	Include	Key influence on throughput
Repair of machines	Exclude	Simplification, no data on repair time
Resources		
Inventory	Include	Experimental factor
Staff	Exclude	Assume always available, 24 hours a day
Equipment	Exclude	Assume always available

In the case data from historical customer demands, only the customers that have a requirement of a shorter customer lead time than the actual production lead time are considered. The demand data consists of dates (year and week) when orders of the products are received (arrival time), the number of products ordered and when the orders are delivered to the customers. The data is used to estimate future demand. For two of the three products, data were missing so the demand pattern was based on the third end product. A limitation to the system is that the number of products in production and the inventory combined cannot exceed the maximum safety stock. This means that a new product will not be produced before a product has been delivered to a customer, even if the product is reserved to the customer. In a way this mimics the Reorder Point system with a reorder point equal to the safety stock level. This set up is chosen because of the high monetary value of the product.

The input data to the model consists of arrival time of the customer order and number of products the customer requests, delivery lead times, production lead times, and level of safety stocks. Outputs are service levels, which is decided by the number of customers that have received their order or end product on time out of the total number of orders.

3.2 Simulation Model

The model is divided into two parts, the demand management, see Figure 1, and the production management, see Figure 2. The entity arrives to the demand management part according to the real future demand estimation.

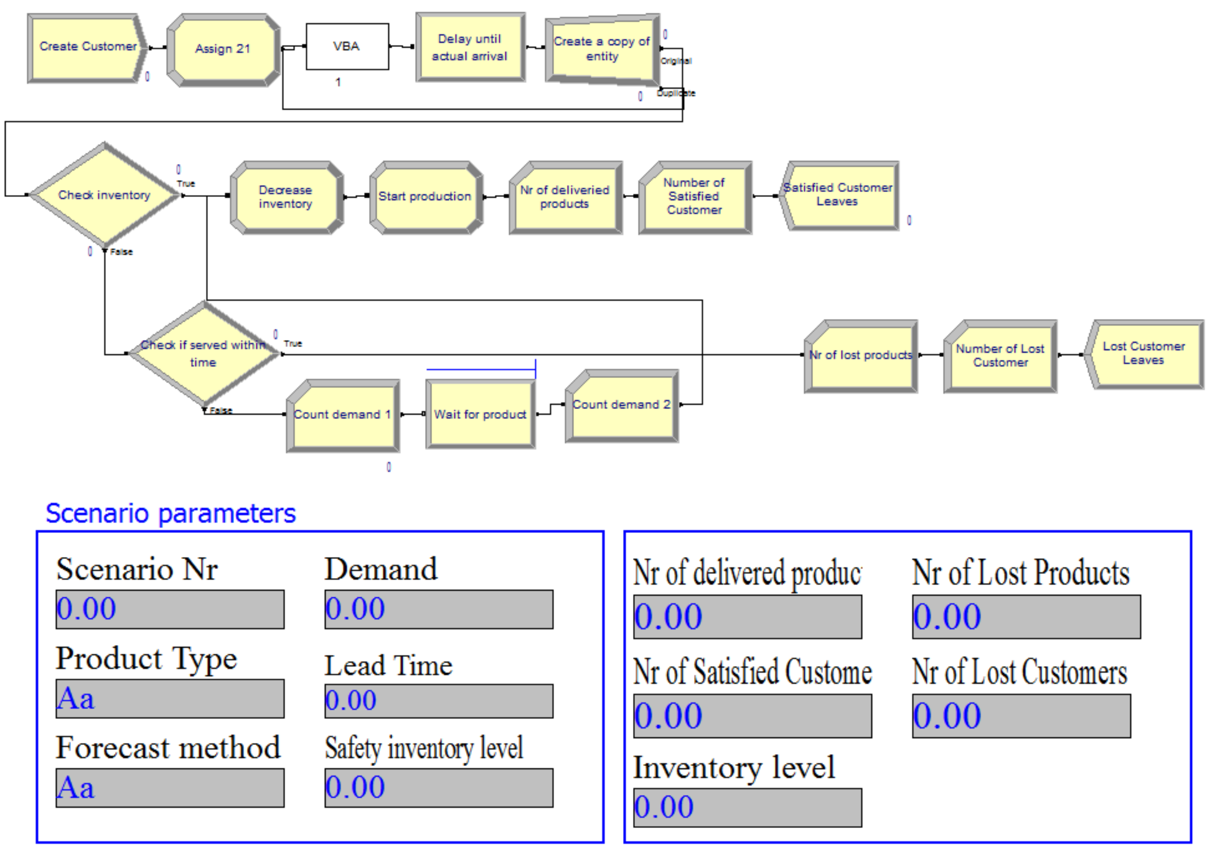

Figure 1 : Demand management.

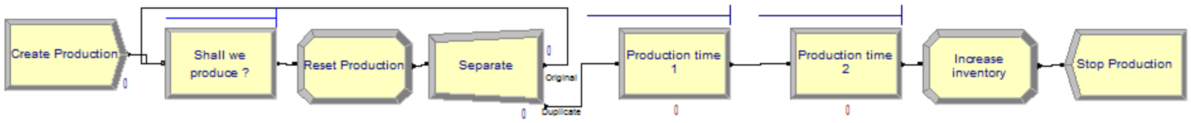

Figure 2: Production management part.

An Excel file contains input data of product type, production time, demand, accepted lead time for the customers and safety stock level for each experiment. The file also consists of the arrival patterns for each product type and demand. In each arrival pattern the arrival time of orders and the size of the orders are stated.

3.2.1 Demand Management

The demand management part of the model starts with a Create module named C*reate Customer*, which creates one entity of type Customer. Then, an Assign module, *Safety stock level* saves the inventory level from the start to the variable *Safety stock level.* The VBA module that comes next reads the arrival patterns and the demands from the Excel file, which generates arrival patterns for customers in the model. A Delay module is then placed in the flow called *Delay until actual arrival*, which makes sure that the customer does not arrive until the right time. When the time delay has passed by, the entity continues to a Separate module, *Create a copy of entity*, where the entity duplicates according to the demand from the customer. One entity goes back to VBA, the other one continues further on to a Decide module.

The Decide module *Check Inventory* contains a condition that checks if the variable *Inventory level* is larger or equal to the demand and if there are no customers waiting for their products. If the condition is true, then a product will be delivered to the customer and the variable *Inventory level* decreases in the Assign module, *Decrease Inventory.* When a product is taken from the inventory, a new product has to be produced. This is done by setting the variable that represent the production to 1, which is done in the next module S*tart Production.* This triggers the other part of the system, production management, to start. After that there are two Record modules, *No. of delivered products* and *No. of satisfied customer*, that count the number of delivered products and the number of satisfied customers respectively. The entity will then leave the system by the Dispose module, *Satisfied Customer Leaves*.

If the condition in the Decide module *Check inventory* is false, then the entity proceeds to the next Decide module called *Check If Served Within Time* that check if the customer can receive their product within the given customer lead time. This is done by a condition that checks if the *Demand* is bigger than the *Inventory level*, number of products in the second part of the production and the sum of the demand from the customer that waits for their products. If the condition is true the entity will move on to two Record modules, *No. of lost products* and *No. of lost customer,* that count the number of products that could not be delivered respectively the number of customer that has been lost. Then the entity will leave to the Dispose module that is called *Lost Customer Leaves.*

If the customer can have the product within the delivery time, then the entity proceeds to the module *Count demand 1* that count the sum of the demands from the customers that wait for their products. The entity moves on to the Hold module *Wait for product* that holds the entity until there is enough products in the inventory. The entity will then pass another Record module, C*ount demand 2,* that count the sum of the demands from the customer that has got their products. This is done so the model can know how many products that are reserved and not available for other customers. The entity will then move on to the module *Decrease inventory* and proceed the same route as earlier explained.

3.2.2 Production Management

The production begins with a Create module, *Create Production,* which generates entity *Product*. The Create module only creates one entity. Next in the model is a Hold module, called *Shall We Produce?* This module is holding the entity until a demand has occurred in the demand part of the model, i.e. when the variable *Production1* are equal to 1. When a demand has occurred, it is time to start the production. Before the production, the entity passes by a Separate module which split up the entity to several entities. One of the entities goes back to the *Shall We Produce?* module, and the other one continues to an Assign module. The Assign module is called *Reset Production* and the aim of this is to set the variable *Production* to zero, which make sure that the production does not continue to produce more than the requested products.

After the Assign module, two Process modules, *Production time 1* and *2*, are placed. The production time is split up in these two Process modules. In the first module the product stays until the time is total production time subtracted with the delivery lead time. In the second module the product stays the time of the customer lead time. The production time is divided in order to be able to check if a customer can receive its order within time. When the production time is over the inventory level is increased by one in the Assign module, *Increase Inventory*. Thereafter the entities leave the model by the Dispose module *Stop Production.*

3.3 Model Verification and Validation

Animation is limited in the model. Inventory levels are defined as variables and can only be viewed as numbers. During the simulations, only some data is visible like scenario parameters, number of delivered products, number of lost products, number of satisfied customer, number of lost customers and inventory level for each scenario. As means of verification, an extensive walkthrough was conducted during preliminary runs. Measured outputs were all within reasonable limits compared to the real system and the model behaved as predicted. Validation was carried out with case company representatives, but only for the conceptual model. To conclude the verification and validation of the model, a high credibility was shown from the case company towards the simulation model.

4 EXPERIMENTS

The model is tested in different experiments to see how the model acts if circumstances change. The parameters that are changed in the experiments are end product type (with different production lead times), total yearly demand, delivery lead time and safety stock. The three products have different demand patterns and delivery lead time, but the same safety stock levels are tested for each product. The safety stock levels are tested from 1 to 15. For the first product (1), three demand levels are tested, with a short delivery lead time. The second product (2) is tested with four demand levels and a short and a long delivery lead time. The last product (3) has three demand levels with a long and a short delivery lead time. The experiments generates different service levels, product delivery levels and average inventory levels. The experiments are divided into different scenarios where the parameters production type, demand and customer lead time changes.

Each simulation had a runtime of 100 simulated years but only one replication. The long runtime is to get a high precision in the output rather than using several replications. The basic idea behind this solution is the fact that demand levels at some experiments were as low as 4 end products per year. With 100 simulated years, at least 400 end products were delivered, creating many observations of the service levels.

5 SIMULATION RESULTS

The simulated results are collected in Figures 3 to 7. For product 1, three different demand levels are depicted in Figure 3, D = 5 stands for at yearly demand of 5 end products, D = 10 for a yearly demand of 10 end products, and D = 15 stand for 15 end products. Also in Figure 3, the two service levels denoted “Order” and “Product” are presented. The “Order” service level is defined as the service level for a complete order, being 1 or more end products. The "Product" service level on the other hand, is the service level per end products regardless of order size. The two service levels differ somewhat throughout the results and are therefore separated.

All five figures, Figures 3 to 7, show basically the same pattern. For higher yearly demand levels, a larger safety stock level is needed to get up to 100% service level regardless of "order" or "product" focus. Lower demand levels also show a steeper incline up to the high service levels. In all, Figures 3 to 7 gives a numerical value of safety stock level, given the end product (1, 2, or 3), the demand level, and the delivery lead time. Note that the delivery lead time always is shorter than the production lead time.

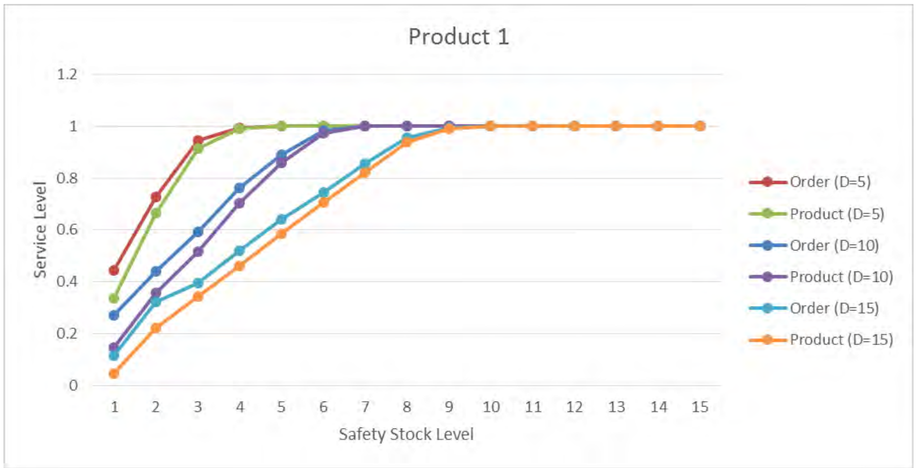

Figure 3: Service Levels for Product 1.

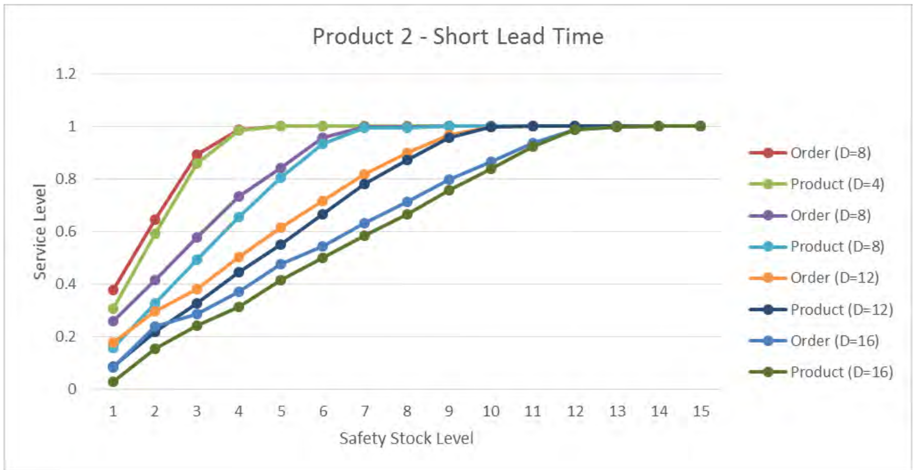

Figure 4: Service Levels for Product 2, short lead time.

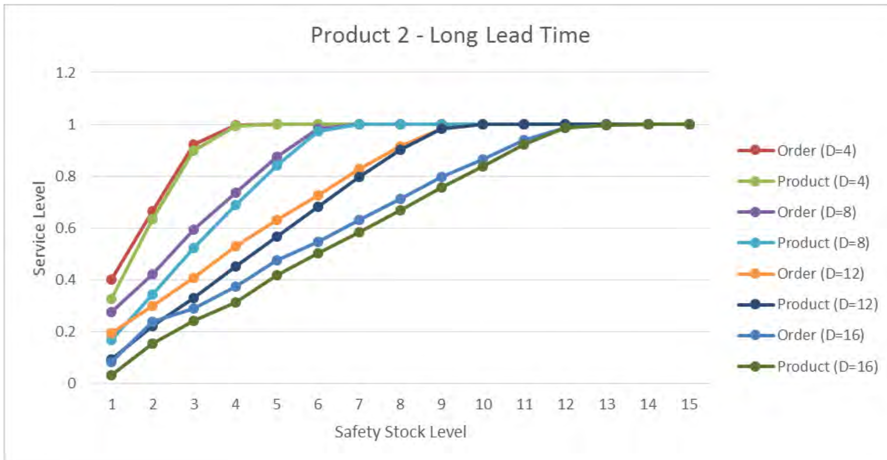

Figure 5: Service Levels for Product 2, long lead time.

Figure 6: Service Levels for Product 3, short lead time.

Figure 7: Service Levels for Product 3, long lead time.

6 ANALYSIS

6.1 Calculation of Safety Stock Levels

The Croston Forecasting Method is used to calculate the forecast error for the different demand patterns that are used in the simulation model. The corresponding safety stock is then calculated using the service level definition with a stock-out probability of 95%. The different demand patterns hold different levels of total yearly demand and thus get different calculated safety stock levels.

First, the forecast error was calculated using daily demand buckets where Croston Forecasting Method was used to forecast the intermittent demand of each product over the simulated 10 years (equation 2 to 5). The forecast error was calculated from the forecasted number of product and the forecasted day of the sale. For a forecasted demand to be deemed as correct, the number of product needed to be exact but the day was allowed to differ by 5 days (at least in the same week). Once the forecast error was determined, the standard deviation for the forecast error was used in equation 1. With a probability level of 95% (k = 1.65), the corresponding lead time, and the correlation factor γ = 0.5, the resulting safety stock levels can be seen in Table 2. The safety stock was calculated and rounded up to ensure the stock-out probability was at least 95%.

Table 2: Results.

Product	Demand [pcs/year]	Forecast error	Lead Time	Safety Stock Level SS	Rounded SS	Simulated Service Level Order	Simulated Service Level Product
1	5	0.146	Short	2.973	3	0.94	0.91
1	10	0.209	Short	4.251	5	0.89	0.86
1	15	0.275	Short	5.595	6	0.75	0.71
2	4	0.124	Long	2.755	3	0.92	0.90
2	4	0.124	Short	2.249	3	0.89	0.86
2	8	0.180	Long	4.019	5	0.88	0.84
2	8	0.180	Short	3.282	4	0.73	0.66
2	12	0.227	Long	5.068	6	0.73	0.68
2	12	0.227	Short	4.138	5	0.61	0.55
2	16	0.290	Long	6.456	7	0.63	0.58
2	16	0.290	Short	5.271	6	0.55	0.50
3	3	0.102	Long	2.793	3	0.76	0.74
3	3	0.102	Short	2.281	3	0.75	0.72
3	6	0.149	Long	4.057	5	0.77	0.73
3	6	0.149	Short	3.312	4	0.64	0.58
3	9	0.183	Long	4.993	5	0.56	0.50
3	9	0.183	Short	4.077	5	0.55	0.49

6.2 Comparing Calculated Safety Stock with Simulated Service Levels

For each scenario (product, yearly demand level, and delivery lead time) the calculated and rounded safety stock level was compared to the simulated corresponding service level. Throughout Table 2, the simulated service levels were lower than the calculated levels (should be at 95%). For product 1, demand level of 5, and short lead time, the calculated safety stock is 2.973 which was rounded up to 3 items. This corresponded to a service level at 94% (measured for whole orders) and 91% (measured for individual products). This product was actually the best performing product out of all experiments. Looking closely at Table 2, it can easily be determined that the Croston Forecasting Method, as means to calculate the standard deviation of forecast error and thereby the safety stock, will always lead to an under estimation of the true safety stock. In all cases that were simulated, the service levels (both for "orders" and "products") were lower than expected. For product 3, high demand level, and short delivery lead time, a safety stock of 5 end products only provide you with a service level about 50%, which is equal to no safety stock at all.

6.3 Discussion

One factor which will decrease the service level in the simulation is the assumption that the number of products in production and in the inventory is not allowed to exceed the safety stock level. This will result in a lowered service level since customers can block each other, because when a customer has reserved a product, new products will not be produced before the reserved products are delivered. This is something to keep in mind when using the results presented in this report. There is however a way to avoid that customers block each other. This can be achieved by allowing the production of new products to start at once when a customer is assigned to a product. This may however result in an inventory and number of products in production which exceed the safety stock while the products are produced.

7 CONCLUSION

The conclusion of this study is that the Croston Forecasting Method which were examined combined with equation 1 for safety stock calculations have many shortcomings when applied to an intermittent demand pattern. This is due to the fact that the method does not take some important factors into account which affect the result. One factor is that the simulated service levels are affected by other parameters, for example how long the customers are willing to wait (a higher service level is achieved when the customer is willing to wait longer). The forecasting methods does not take this sort of information into account. Also, the forecast error is calculated in a very crude manner where equal significance is put on a day where we thought we would sell an end product but did not, and a day where we sold an end product without a forecast. The standard deviation of forecast error was calculated in a way that do not take the amount of time between two order points into account. It only considers the order quantity, where an order is viewed as incorrect regardless of how far off the actual order point the forecast was placed. Here, this was solved by allowing for that time to off by 5 days.

The purpose of this paper was to test if a safety stock level that has been calculated with a high probability to avoid stock-outs can match the simulated safety stock levels for end products with intermittent demand. The answer is that the Croston Forecast Method is no shortcut towards a fast and reliable safety stock calculation. We instead encourage the use of simulation analysis for this complex and important problem. Intermittent demand patterns, long production lead times, and high monetary values in products is a difficult combination where the strengths of simulation can show its full potential.

ACKNOWLEDGMENT

The authors would like to thank Siemens Industrial Turbomachinery AB who provided us with data and the problem description in the first place. Also, a thanks to the examiner of the applied simulation course in which this project took place, Krisjanis Steins.

REFERENCES

Anupindi, R., S. Chopra, S. D. Deshmukh, J. A. Van Mieghem, and E. Zemel. 2014. *Managing Business Process Flows*, 3rd ed. Edinburgh: Pearson Education Limited.

Croston, J. D. 1972. "Forecasting and Stock Control for Intermittent Demands". *Operational Research Quarterly*, 23: 289-303.

Dodge, Y. 2003. *Oxford Dictionary of Statistical Terms,* Oxford: Oxford University Press.

Hernandez-Ruiz, K. E.; E. Olivares-Benitez, J. L. Martinez-Flores, and S. O. Caballero-Morales. 2016. "Optimizing Safety Stock Levels in Modular Production Systems Using Component Commonality and

Group Technology Philosophy: A Study Based on Simulation". *Mathematical Problems in Engineering.* 2016: p1-13.

Hopp, W. J. and M. L. Spearman. 2008 *Factory Physics*. 3rd ed. Long Grove: Waveland Press, Inc.

Levén E. and A. Segerstedt. 2004. "Inventory Control with a Modified Croston Procedure and Erlang Distribution". *International Journal of Production Economics*. 90: 361-367.

Olhager, J. 2000. *Produktionsekonomi*, Lund: Studentlitteratur (in Swedish).

Olhager, J. and F. Persson. 2006. "Simulating Production and Inventory Control Systems: A Learning Approach to Operational Excellence". *Production Planning and Control*. 17: 2.

Persson, F. 2003, *Discrete Event Simulation of Supply Chains – Modelling, Validation and Analysis*, Doctoral Thesis, Department of Production Economics, Linköping Institute of Technology, Linköping.

Schmidt, M., W. Hartmann, and P. Nyhuis. 2012. "Simulation Based Comparison of Safety-Stock Calculation Methods". *CIRP Annals - Manufacturing Technology* 61: 403–406.

Silver, E. A. 1981. "Operations Research in Inventory Management: A Review and Critique". *Operations Research*. 29: 628-645.

Silver, E. A., D. F. Pyke, and R. Peterson. 1998. *Inventory Management and Production Planning,* 3rd ed. New York: John Wiley & Sons.

Syntetos, A. A. and J. E. Boylan. 2001. "On the Bias of Intermittent Demand Estimates". *International Journal of Production Economics*, 71: 457–466.

Teunter, R. and B. Sani. 2009. "On the Bias of Croston's Forecasting Method". *European Journal of Operational Research*, 194: 177–183.

Willemain, T. R., C. N. Smart, and H. F. Schwartz. 2004. "A New Approach to Forecasting Intermittent Demand for Service Parts Inventories". *International Journal of Forecasting*. 20: 375-387.

AUTHOR BIOGRAPHIES

FREDRIK PERSSON is an Associate Professor of the Department of Management and Engineering at Linköping University. He holds a Ph.D. in Production Economics. His research interests include discrete-event simulation, production planning and control, supply chain and operations management in various applications from healthcare to construction industry. His e-mail address is fredrik.persson@liu.se.

MINNA AXELSSON is a Supply Chain Manager at Siemens Industrial Turbomachinery AB. She holds a Master of Science in Electrical Engineering from Linköping University. She represents the company Siemens and contributed with the data to this paper. Her e-mail address is minna.axelsson@siemens.com.

FREDRIK EDLUND is a master student in the Master program in Communication and Transportation Engineering at Linköping University. He has extensive skills in modelling and simulation and inventory management. His e-mail address is freed866@student.liu.se.

CHRISTOFFER LANSHED is a master student in the Master program in Communication and Transportation Engineering at Linköping University. He has extensive skills in modelling and simulation and inventory management. His e-mail address is chrla320@student.liu.se.

AGNES LINDSTRÖM is a master student in the Master program in Communication and Transportation Engineering at Linköping University. She has extensive skills in modelling and simulation and inventory management. Her e-mail address is agnli392@student.liu.se.

FRIDA PERSSON is a master student in the Master program in Communication and Transportation Engineering at Linköping University. She has extensive skills in modelling and simulation and inventory management. Her e-mail address is fripe651@student.liu.se.