LOST SALES INVENTORY POLICY WITH A SERVICE LEVEL CRITERION AND NON-STATIONARY DEMAND: A CASE STUDY

Public summary master thesis Industrial Engineering and Management

June 2022

Author
Marlinde Vetkamp

Supervisors
Dr. M. C. van der Heijden
Dr. D. R. J. Prak

Main topic of the thesis

To replenish the pick areas from the bulk storage area, Company X uses a replenishment method that decides which SKUs need to be replenished every day and the replenishment quantity. At company X, there are three pick areas, each having a specific function. Since Company X continues to grow and the complexity of the replenishment method increases, it is important that the replenishment method functions correctly and keeps up with the upcoming changes. Therefore, the company is looking for improvements in the replenishment method to make it more efficient and future-proof.

In the current situation, a customer's order can contain items that are unavailable in the pick areas. An unavailable SKU in the pick area, which is available in the bulk storage area, results in a delayed order. SKUs are unavailable in one of the pick areas when there is no or not enough inventory to replenish the SKU, or the SKU is "lost", meaning that the location of the SKU is not registered. Furthermore, an SKU is unavailable in the pick area when there are not enough employees to replenish the SKU, or the replenishment task is executed too late. Instead of hiring more people to solve the problem, the current replenishment method can be improved. The current replenishment method is based on an (R, S) policy. The order-up-to-level S is calculated by multiplying the expected demand by a specified number of days for each of the pick areas. The multiplier for a pick area is not SKU specific, and the expected demand is based on historical demand data, not considering the forecasted trend or seasonality. The productivity of performing replenishment tasks can be improved when larger quantities of one SKU are replenished less often.

Given the problems Company X faces with the current replenishment method, the problem this research aims to solve is:

"How can the replenishment method be made SKU specific by focusing on the parameters that decide when an SKU is replenished to minimize the number of customer orders that cannot be picked"

Purpose of the thesis

By solving the core problem, the thesis aims to minimize the delayed customer orders caused by not replenishing the pick areas. The thesis focuses on determining the correct parameter settings for each SKU. Based on the literature review and expert opinions, a new replenishment method is designed. Performing experiments, the best settings are determined, and recommendations to the company are given. The deliverables are a new replenishment model and an implementation plan.

Research questions and design

To answer the core problem, research questions are formulated. For each research question, a solution approach is given.

1. How does the current replenishment method work?
    a. How do the pick areas and the bulk storage area look?
    b. How do the steps of the current replenishment method look?
    c. What is the qualitative performance of the current replenishment method?
    d. What is the quantitative performance of the current replenishment method?

    Performing interviews with three product owners who work on improving the replenishment method and observing an employee during the replenishment shift, insight is gained into the pick areas, bulk storage area and current replenishment method. Furthermore, the performance of the current replenishment method is measured by tracking the available data and performing real-life observations.

2. What modelling approach presented in the literature can be used to improve the replenishment method?
   a. How can the research problem be placed within the existing literature?
   b. Which models presented in the literature are suitable for improving the replenishment method?
   c. Which model can best be used to improve the replenishment method?

   A literature study in Scopus, WorldCat and Elsevier is executed to identify the different characteristics of inventory control policies and to select the desired inventory control policy. Based on the available models, requirements and wishes, the suitable models to implement the desired inventory control policy are identified. A literature study on the desired inventory control policy is performed, to select the best model to improve the current replenishment method.
3. How is the model for the replenishment method designed?
   a. Which assumptions are needed to design the model?
   b. What are the inputs and outputs of the model?
   c. What decisions are made within the new replenishment method?
   d. How can we validate the model?

   Insight is gained into how the new replenishment method is designed. A step-by-step approach, flow charts and figures explain the developed, new replenishment method. The model is verified and validated using the techniques presented by Law (2013).
4. What is the performance of the new replenishment method?
   a. How does the replenishment method perform with different input datasets?
   b. How does the new replenishment method perform compared to the current replenishment method?
   c. How does the replenishment method perform with different parameter settings?

   The current replenishment method is modelled, called the (R, S) policy, to enable a comparison under the same assumptions. Experiments are designed and executed to determine how the model performs with different input datasets and input parameter settings. Additional insights based on the best setting are given as well.
5. What is the advice to the company based on this research?
   a. What are the conclusions of this research?
   b. What are the recommendations of this research?
   c. What are the limitations and directions for future research of the new replenishment method?
   d. How can Company X implement the developed replenishment method?

Literature research

To solve the core problem of this research, first, the features and performance of the current replenishment method are studied. Based on the features, performance, requirements and wishes, the desired inventory control policy is selected. By performing a literature review, the characteristics of inventory control policies are identified. The desired inventory control policy is an (R, s, S) policy with a service level constraint, under the assumption of lost sales, with non-stationary demand.

Literature research on desired (R, s, S) policy

Literature research is performed focused on (R, s, S) policies with a service level constraint under the assumption of lost sales. A literature review by Bijvank & Vis (2011) focusing on lost-sales inventory policies shows that only three papers focus on an (R, s, S) policy with a service level constraint under the assumption of lost sales. Tijms and Groenevelt (1984) briefly mention the lost-sales case, however, the calculations for the order-up-to-level and reorder point are restricted to a backorder policy. The papers by Kapalka et al. (2009) and Bijvank & Vis (2012) propose models to determine the reorder point and order-up-to-level with a fill rate constraint. Kapalka et al. formulate a Monotone Search Algorithm (MSA) which iteratively determines and tests different settings for the reorder point s and the order-up-to-level S for a single SKU. Bijvank & Vis develop both an optimal and an approximation procedure to determine the reorder point s and the order-up-to-level S. Furthermore, based on the literature on non-stationary back-ordering (R, s, S) policies, a static-dynamic uncertainty strategy is applied, and the parameters are updated every review period to account for non-stationary demand (Bookbinder & Tan, 1988 and Pauls-Worm et al., 2014).

Explanation of the MSA

The Monotone Search Algorithm (MSA) is an algorithm that iteratively determines the best settings for the reorder point and the order-up-to-level. Due to the specific design for the use case, the MSA differs from other models. Interesting differences are that the lead time is a fraction, instead of an integer multiple, of the review period and the ignorance of the warehouse store interaction (Kapalka et al., 2009). The warehouse is assumed to have infinite supplier capacity. Each possible combination of the reorder point and order-up-to-level is tested using a probability matrix. For the implementation in this research, the transition matrix is not used since the demand is non-stationary and finding the demand distributions for all SKUs is time-consuming. Instead, historical demand data is used to calculate the reorder point, order-up-to-level and achieved fill rate.

$$Fill\ rate = 1 - \frac{Demand\ not\ satisfied\ during\ a\ period}{Total\ demand\ during\ a\ period}$$

Equation 1

To limit the search space of the MSA, bounds on the order-up-to-level S are set. The MSA is initialized by setting the order-up-to-level equal to the lower bound for S and the reorder point equalling to S - 1. To determine if a setting for s and S is worth considering, the fill rate is calculated, shown in Equation 1. When a setting for s and S fulfils the fill rate constraint, the reorder point is lowered by 1, or the order-up-to-level is raised by 1. For each new setting, it is checked whether the fill rate constraint is fulfilled and if the expected cost is lower than the previous lowest cost. The logic behind the MSA is that *"if the average cost associated with (s +1, S) exceeds that of (s, S) and (s, S) satisfies the service level constraint, then there is no need to check policies (s + 2, S), (s + 3, S), . . . , because they will have higher costs. Similarly, if the service level of (s, S) is lower than the required service level, there is no need to check policies (s - 1, S), (s - 2, S), . . . , because they will fail to provide the required level of service"* (Kapalka, et al., 2009). The MSA stops once the order-up-to-level is equal to Smax.

Choice for MSA

Based on the differences between the models proposed by Kapalka et al. (2009) and Bijvank & Vis (2012), the MSA is favoured for the following reasons:

- The MSA combines the determination and testing of different parameter settings, to find the best setting. The model is explicitly explained and visualized compared to the procedure of Bijvank & Vis.
- The iterative nature of the MSA makes it easier to implement the MSA at Company X.
- The use case on which the MSA is tested resembles the situation at Company X. First, the assumption that the warehouse has infinite supplier capacity resembles the assumption that there is always enough inventory at Company X to replenish the pick areas. Besides, the costs also represent a small fraction of the value of the product.

- Bijvank & Vis assume that the lead time is larger than the review period, whereas, for the use case of the MSA, the lead time is a fraction of the review period. For Company X, the lead time is shorter or equal to the review period, both given in days, resembling the use case of the MSA.

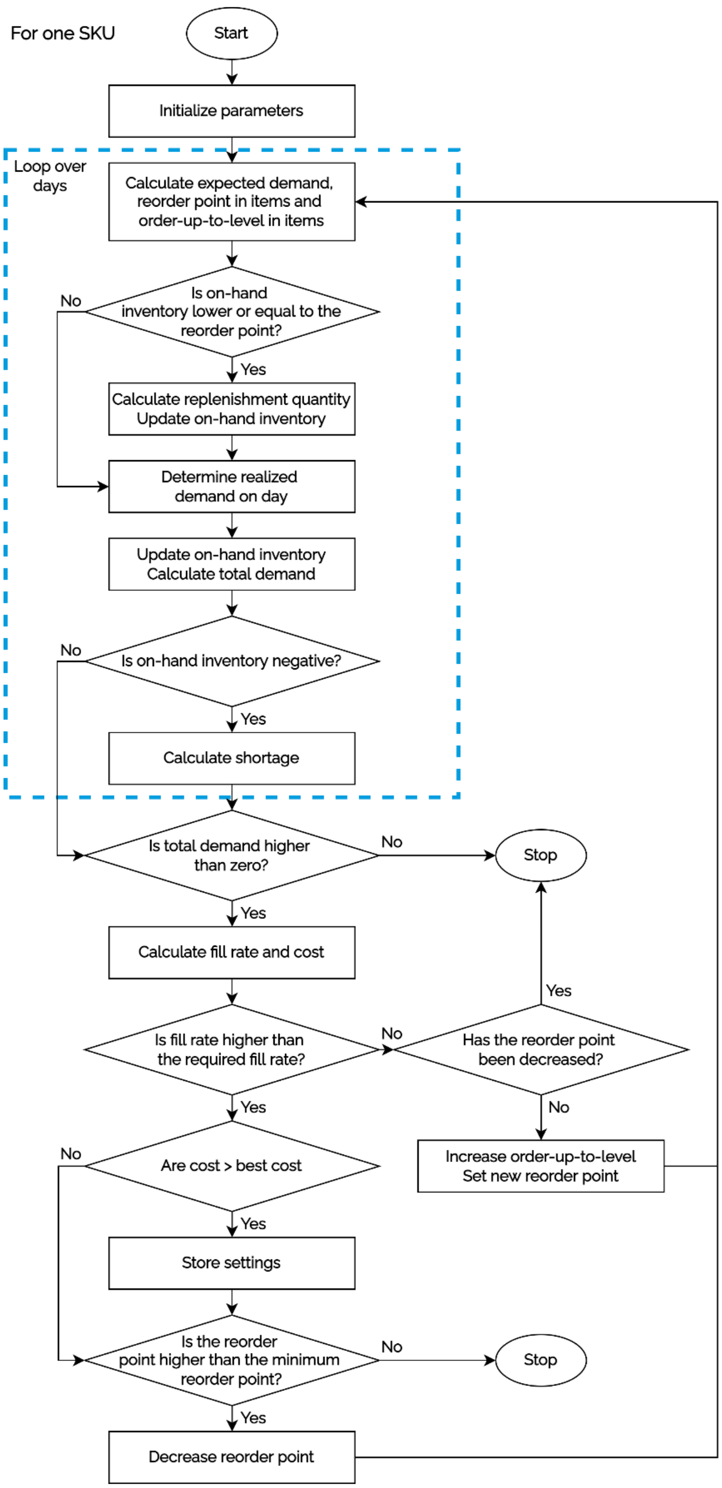

Figure 1: MSA flowchart

Solution design

The MSA proposed by Kapalka et al. (2009) is used to develop the new replenishment method. The new replenishment method is a simplified version of the MSA. Based on the MSA, a new algorithm is designed, considering the requirements and wishes of the company and overcoming the limitation that demand patterns are unknown. The main differences between the MSA of Kapalka et al. (2009) and the simplified MSA are:

- Historical demand data is used to determine and test the parameter settings instead of a transition matrix since the demand patterns are unknown and non-stationary should be integrated.
- The simplified MSA is used to determine the parameter settings for 5000-7000 SKUs instead of a use case of 420 SKUs.
- Within the simplified MSA, the order-up-to-level can only change if the initialized setting is insufficient to fulfil the fill rate constraint. The MSA of Kapalka et al. (2009) also determines the best order-up-to-level setting. The order-up-to-level is fixed to ensure that the inventory levels do not exceed the available storage space.
- Instead of a fixed reorder point and order-up-to-level, the reorder point and order-up-to-level depend on the expected demand per day. In this way, non-stationary demand is integrated. The expected demand per day is a moving average based on several days of historical demand data.
- The cost is not based on holding and ordering costs but represents the difference between the reorder point and order-up-to-level. The bigger the gap, the more items of an SKU are replenished at once, resulting in higher productivity and a lower diversity of the replenishment list.
- Instead of the lead time being a fraction of the review period, the lead time is set to zero, meaning that every SKU put on the replenishment list on that day is directly replenished.
- The simplified MSA is used for three different pick areas.

Figure 1 provides a detailed flow chart of all the steps taken to implement the new replenishment method (MSA). In short, the MSA determines a parameter setting for an SKU and simulates the parameter setting by generating replenishments, updating the on-hand inventory and realized demand. After the parameter setting is simulated, the achieved fill rate is calculated to determine whether the fill rate constraint is fulfilled and the best setting is stored. The MSA continues to decrease the reorder point or increase the order-up-to-level. The MSA stops when the minimum reorder point is reached, or a setting does not fulfil the fill rate constraint.

The new replenishment method and the (R, S) policy are implemented and tested in Delphi software. Figure 2 presents the steps followed in the new replenishment method, called the MSA. The (R, S) policy closely resembles the steps of the MSA, except that the order-up-to-level and reorder point are fixed. The order-up-to-level is given by Company X, the reorder point is the order-up-to-level in items minus one item. Besides, the (R, S) policy does not have a fill rate constraint and decreasing step setting. The modelled (R, S) policy matches the current replenishment method as closely as possible, except that the review period is set to one day since the historical demand data is only available per day. In the remainder of this section, the specific steps of the MSA are explained.

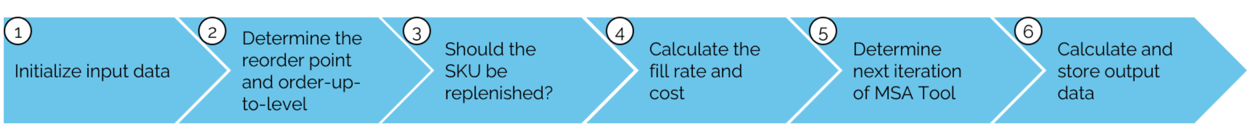

Figure 2: Overview MSA steps

Step 1: Initialize input data

Before the algorithm can be executed, the settings for the order-up-to-level in days, minimum reorder point in days, decreasing step in days and fill rate constraint are initialized. The order-up-to-level in days equals the order-up-to-level setting of the (R, S) policy. To test and simulate the parameter settings, the MSA runs for at least 59 days, including an initialization period of ten days. Table 1 shows the input parameter settings of the MSA. The fill rate constraint is set at 0.95, such that the achieved fill rate of the (R, S) policy and MSA closely resemble the fill rate of the current situation. The decreasing step of pick area 2 is set at 0.5 days, to ensure that the minimum reorder point of 1 day can be reached.

Table 1: Input parameter settings

Parameter	Pick area	MSA
Minimum reorder point (in days)	All pick areas	1
Decreasing step reorder point (in days)	Pick area 1	1
	Pick area 2	0.5
	Pick area 3	1
Fill rate constraint	All pick areas	0.95

Step 2: Determine the reorder point and order-up-to-level

The reorder point and order-up-to-level are given in days to limit the search space and enable updating of the reorder point and order-up-to-level in items every day, to integrate non-stationary demand. Before we can determine whether an SKU should be replenished, the reorder point and order-up-to-level in items should be calculated. To do so, the reorder point in days and order-up-to-level in days are multiplied by the expected demand per day of an SKU. The expected demand per day of an SKU is a moving average based on several days of historical demand data. Every day the algorithm is executed, the reorder point in items and order-up-to-level in items are calculated. When the expected demand is higher, the reorder point and order-up-to-level will be higher and vice versa. The reorder point can only decrease by a (fraction of a) day instead of by one or multiple items. The algorithm is used for 5000-6000 SKUs, decreasing the reorder point by one or multiple items increases the search space significantly.

Step 3: Should the SKU be replenished?

An SKU is only put on the replenishment list when the on-hand inventory is equal to or below the reorder point in items. All the SKUs that are put on the replenishment list are assumed to be replenished the same day before the demand is realized. The replenishment quantity is calculated as follows:

$$Replenishment\ quantity\ (i) = S(i) - OH(i)$$

Equation 2

Where i is the SKU, S(i) is the order-up-to-level in items, and OH(i) is the on-hand inventory of SKU i.

Step 4: Calculate the achieved fill rate and cost

After the simulation of one combination for s and S is finished, the achieved fill rate is calculated. The (R, S) policy performs only one iteration for each SKU since the parameter settings for s and S are fixed. Based on the achieved fill rate, the next iteration of the MSA can be determined. For both the (R, S) policy and the MSA, the achieved fill rate is calculated as follows:

$$Achieved\ fill\ rate(i) = \left[1 - \left(\frac{Shortage(i)}{Total\ demand\ (i)}\right)\right]$$

Equation 3

Where i is the SKU. The shortage and total demand are calculated over the whole period of the MSA. When the on-hand inventory is below zero, the shortage equals the absolute of the on-hand inventory at that point. The fill rate is calculated for each parameter setting the MSA simulates. As shown in Figure 1, first, a parameter setting is simulated by generating replenishments, updating the on-hand inventory and realizing demand. After the simulation is finished, so the day loop is finished, the fill rate and cost are calculated. After the whole simulation is finished, the weighted average fill rate is calculated, weighing the achieved fill rate of an SKU with the proportion of the total demand realized by the SKU. The cost is the order-up-to-level in days minus the reorder point in days.

Step 5: Determine next iteration of the MSA Tool

When the (R, S) policy is executed, the MSA Tool only performs one iteration for each SKU since the reorder point and order-up-to-level are fixed. When the MSA is executed, after each iteration, the MSA determines the next iteration. There are three possibilities for the next iteration:

1. Decrease the reorder point. The reorder point is decreased when the fill rate constraint is fulfilled, and the reorder point does not equal the minimum reorder point.
2. Increase the order-up-to-level. The initialized settings can be insufficient to fulfil the fill rate constraint. In case the reorder point has not been decreased yet, and the fill rate constraint is unfulfilled, the order-up-to-level is increased with the decreasing step setting. The new reorder point equals the order-up-to-level minus the decreasing step.
3. Stop, no new iteration for this SKU. Three criteria stop the iterations of one SKU:
    a. The total demand is zero.
    b. The fill rate constraint is not fulfilled, and the reorder point has been decreased in previous iterations.
    c. The reorder point equals the minimum reorder point.

When there are no new iterations to perform for the SKU, the MSA continues to the next SKU. When there are no SKUs left, the MSA ends. Both the (R, S) policy and the MSA are executed for the three pick areas.

Step 6: Calculate and store the output data

After the (R, S) policy or the MSA is performed, the output data is calculated and stored. The following output data is stored:

- The achieved fill rate, total demand and shortage for every SKU.
- The best cost and the best setting for the reorder point and order-up-to-level for every SKU.
- The weighted average fill rate. The fill rate of every SKU is weighted with the demand, resulting in the average weighted fill rate for all SKUs.
- The number of replenishments per pick area. One replenishment represents one SKU that needs to be replenished. The number of replenishments per day also indicates how many different SKUs are replenished.
- The on-hand inventory per day and per pick area.
- The number of items and times short per pick area.

The output data is stored in a text file, which is transferred to Excel to allow easy data comparison.

Verification and validation

To verify and validate the developed model, five verification techniques and five validation techniques proposed by Law (2013) are used. The main validation technique used is to validate the output from the overall model. Both the output from the MSA and the (R, S) policy are compared to the current replenishment method at Company X.

To validate the output of the MSA and the (R, S) policy, the average fill rate is compared with the fill rate achieved by the current replenishment method. The difference between the actual fill rate and modelled fill rate is between 0.2% and 0.5%, caused by:

- The model assumes that every SKU placed on the replenishment list is replenished the same day. The model does not consider that the needed employees to replenish all the SKUs are unavailable or that the time needed to replenish all the SKUs exceeds the time in the replenishment shift.
- The model assumes that the bulk storage area always has enough inventory to fulfil all the replenishments. Company X does not account purchasing issues; however, it can happen that an SKU is not available. The model does not consider the inventory of the bulk storage area, so more items may need to be replenished than are available.
- The actual fill rate is based on more data points compared to the average fill rate.

The achieved fill rate by the MSA and the (R, S) policy is determined to be valid since the model, and the current situation, compare closely. Furthermore, the model design fits the specific purposes for which it is designed, namely to test and compare the MSA to the (R, S) policy (Law, 2013, pp. 249, 262).

Table 2: Average percentage difference between the number of replenishments compared to the current situation

Day	**Pick area 1**	**Pick area 2**	**Pick area 3**	**Pick area 1**	**Pick area 2**	**Pick area 3**
Model	**(R, S)**	**(R, S)**	**(R, S)**	**MSA**	**MSA**	**MSA**
Average % difference	+101%	-40%	-	+31%	-43%	-92%

The replenishment list generated by the MSA and the (R, S) policy is also compared to the current replenishment list. Table 2 presents the average percentage difference between the number of replenishments with the current situation and the (R, S) policy and MSA for each pick area. The largest difference between the number of replenishments of the current situation and the (R, S) policy is for pick area 1. The difference can be explained by:

- The number of items replenished for one SKU cannot exceed the maximum number of items to be put in a task in the current situation. In the model, this is possible.
- A complete replenishment list of one day is not available in the current situation.
- The on-hand inventory levels in the current situation can differ from the modelled on-hand inventory.

Comparing the MSA with the current situation, the largest difference between the number of replenishments is for pick area 3. The difference is caused by the setting of the reorder point, the MSA sets the reorder point of all SKUs in pick area 3 to one day. Based on how the (R, S) policy is modelled and the current situation, the difference seems reasonable. Since the (R, S) policy is only used to allow for an equal comparison between the (R, S) policy and MSA, the modelled (R, S) policy fits the purposes for which it is designed.

Results

To test the model, different experiments are executed to see how the model, especially the MSA, will perform. The following experiments are performed:

- Different input datasets. Three different historical demand datasets are used as input for the model to test the validity. Based on these experiments, we can conclude that the model can handle different input data, larger datasets, and can deal with variability. What should be noted is that there is a delay in updating the reorder point and order-up-to-level in items. Furthermore, the datasets contain different action periods, but the model only corrects for these action periods after the action period happened. The delay and the correction after the action period are both consequences of basing the model on historical demand data instead of forecasted demand.

- Fill rate constraint and minimum reorder point experiments. By increasing the fill rate constraint and/or the minimum reorder point setting of the MSA, the MSA results in fewer shortages and a higher average fill rate compared to the (R, S) policy, with a lower inventory level and a fewer number of replenishments/day. Fewer shortages mean that the number of customer orders that cannot be picked is lower.
- Order-up-to-level and decreasing step experiments. A higher order-up-to-level results in a higher average inventory level, a lower number of replenishments/day and a lower number of items and times short, thus in a higher average fill rate. A smaller decreasing step means that the average fill rate can get closer to the fill rate constraint.
- Experiments to minimize the number of shortages, but with a lower or equal inventory level and lower number of replenishments compared to the (R, S) policy. The experimental settings are based on the previously mentioned experiments to select the best parameter setting.

Conclusion

The MSA with the best parameter setting outperforms the (R, S) policy, as shown in Table 3. With the best input parameter settings, the MSA results in fewer shortages, with a lower average inventory level and a lower number of replenishments/day. A lower number of replenishments/day means that an SKU is replenished less frequently. Replenishing an SKU less often results in higher productivity. The lower the average inventory level in the pick areas, the higher the inventory level in the bulk storage area. This means that more storage racks are needed in the bulk storage area and fewer in the pick areas compared to the (R, S) policy. However, a storage rack in the bulk storage area can hold more items compared to a storage rack in the pick areas, so a higher inventory level in the bulk storage area does not necessarily mean that more storage racks need to be available. The achieved fill rate of the MSA is higher than the fill rate of the current situation, around 1.8%, meaning that more customer orders can be picked. In Table 3, the fill rate achieved by the MSA is compared to the modelled (R, S) policy, which closely resembles the current situation.

Table 3: Performance of MSA compared to (R, S) policy

KPI	Percentage difference with (R, S) policy
Average fill rate	(+0.72)
Average inventory level (items)	(-11%)
Number of replenishments/day	(-32%)
Number of times short	(-85%)
Total number of items short	(-97%)
Number of items short/times short	(-79%)

Recommendation and discussion

Based on the experiments performed, Table 4 presents the recommended input parameter settings. The decreasing step is set at 1 and 0.5 since a smaller decreasing step results in more shortages. For the pick area 1, we recommend different parameter settings for slow and fast movers. As a rule of thumb, we advise defining slow movers as SKUs with a demand equal to or below 1 item per day. The recommended minimum reorder point for slow movers is 1 day, only for SKUs that are rarely sold, a reorder point of 0 days is recommended. Combined with the fill rate constraint of 0.995, almost no shortages are allowed for slow movers. For fast-moving SKUs in pick area 1, a fill rate constraint of 0.99 and a minimum reorder point of 2 days prevent most shortages without a significant increase in the inventory level and the number of replenishments/day. The settings for pick areas 2 and 3 are chosen such that almost no shortages occur. The inventory level does not significantly increase due to the high fill rate constraint. We recommend choosing the order-up-to-level such that the storage available is used as efficiently as possible. With the MSA, different settings for order-up-to-levels can be tested, and the effect on the inventory level can be determined. The inventory level with a certain order-up-to-level setting can be compared with the actual amount of storage available. One order-up-to-level for each pick area is recommended to keep the model simple and to make the determination of the order-up-to-level setting easier. When determining the order-up-to-level setting, room for growth, the variability in demand and not creating peaks in the inventory level have to be considered.

Table 4: Recommended input parameter settings

Pick area	Decreasing step (days)	Minimum reorder point (days)	Order-up-to-level (days)	Fill rate constraint
1 - slow movers	1	0 or 1	Based on available storage	0.995
1 - fast movers	1	2	Based on available storage	0.99
2	0.5	1	Based on available storage	0.995
3	1	1	Based on available storage	0.995

We advise Company X to determine the performance of the MSA when uncertainties are modelled. Furthermore, the replenishment tasks should be modelled, to determine the effect on the productivity and diversity of a replenishment task in more detail. Improving the expected demand calculations prevents delays in updating the reorder point and order-up-to-level, and non-stationary demand can be integrated better. Lastly, we expect that a strategy for determining the input parameter settings for new SKUs and setting the input parameters for individual SKUs can improve the MSA.

References

Bijvank, M., & Vis, I. F. (2011). Lost-sales inventory theory: A review. *European Journal of Operational Research, 215*(1), 1-13. doi:https://doi.org/10.1016/j.ejor.2011.02.004

Bijvank, M., & Vis, I. F. (2012). Lost-sales inventory systems with a service level criterion. *European Journal of Operational Research, 220*(3), 610-618. doi:https://doi.org/10.1016/j.ejor.2012.02.013

Bookbinder, J. H., & Tan, J. (1988). Strategies for the Probabilistic Lot-Sizing Problem with Service-Level Constraints. *Management Science, 34*(9), 1096-1108. doi:https://doi.org/10.1287/mnsc.34.9.1096

Kapalka, B. A., Katircioglu, K., & Puterman, M. L. (2009). Retail inventory control with lost sales, service constraints, and fractional lead times. *Production and Operations Management*, 393-408. doi:https://doi.org/10.1111/j.1937-5956.1999.tb00315.x

Law, A. M. (2013). *Simulation modeling and analysis.* Tucson: McGraw-Hill Education.

Pauls-Worm, K. G., Hendrix, E. M., Haijema, R., & van der Vorst, J. A. (2014). An MILP approximation for ordering perishable products with non-stationary demand and service level constraints. *International Journal of Production Economics, 157*, 133-146. doi:https://doi.org/10.1016/j.ijpe.2014.07.020

Tijms, H., & Groenevelt, H. (1984). Simple approximations for the reorder point in periodic and continuous review (s, S) inventory systems with service level constraints. *European Journal of Operational Research, 17*(2), 175-190. doi:https://doi.org/10.1016/0377-2217(84)90232-7