

Safety Stock Adjustments to Meet Desired Service Level Requirements

Yasin Ünlü, Manuel D. Rossetti, Ph. D., P. E.
Department of Industrial Engineering, University of Arkansas
4207 Bell Engineering Center, Fayetteville, Arkansas 72701, USA

Abstract

Safety stock levels can be calibrated via simulation procedures to increase the likelihood that the planned for performance will actually be met in practice. Gudum and de Kok [1] proposed a simulation based safety stock adjustment procedure that is motivated by the problem of comparing lot-sizing rules. We refine the procedure by presenting a better algorithmic representation so that practitioners can more readily understand and implement the procedure. We extend the study by considering the problem of comparing other lot-sizing rules; namely, lot-for-lot, part-period balancing and least unit cost. Experiments illustrate how the safety stock adjustment procedure performs under a number of demand scenarios.

Keywords

Safety stock, service level, lead time, lot-sizing

1. Introduction

Safety stock in an inventory system plays a key role in attaining target service levels in inventory systems. A number of safety stock setting methods are available. The traditional method is based on models found in classical inventory theory [2]. In general, these models assume that the demand follows a particular pattern or is known in advance; however, demand uncertainty is the core issue that degrades the quality of decision making process by yielding unattained service level in such inventory systems. In order to achieve the targeted service level, an alternative approach is to model demand uncertainty using simulation procedures. From the standpoint of determining safety stocks to achieve a particular service level via simulation, the following key papers appear in the inventory literature: Callarman and Mabert [3] investigate the potential use of so-called “Service Level Decision Rule (SLDR)” which is developed through a linear regression analysis in order to estimate the service level. The rule is developed using a response surface mapping procedure that captures the changes in the service level against the change in safety stock buffer levels. By changing the safety stock levels systematically, the rule is built with the simulation of experimental factors of coefficient of variation of demand, forecast error (expressed as a percentage of average demand), the amount of safety stock (expressed as a percentage of average demand) and the time between orders. A search routine was applied with SLDR in order to achieve the desired service level. SLDR is also used in [3] to determine the required safety stocks achieving 95% and 98% service levels. Wemmerlöv and Whybark [4] study the problem of comparing single-stage lot-sizing rules by determining net requirements based on allowing backorders. Fourteen different lot sizing rules were compared to each other based on the cost of keeping a certain level of safety stock to achieve nearly a 100% service level for fill rate. The safety stocks are determined by repeating the simulations until the target service levels are reached. Later, Wemmerlöv [5] studied a similar problem by determining net requirements based on lost sales.

The methodology labeled, “Safety Stock Adjustment Procedure” (SSAP), in [1] is also motivated by the problem of comparing different lot-sizing rules. When comparing lot-sizing rules via total cost, it is important that the rules be compared under exactly the same service levels. Thus, decision makers can directly determine the better rule without resorting to more complicated analysis via a trade-off curve approach. By assuming a particular time phased order point policy (TPOP) [6], the authors are able to show that a simulation based procedure that estimates the empirical probability distribution of the net stock at the end of a period can be exploited to develop update formulas for the safety stock. That is, the procedure keeps track of the behavior of the net stock levels observed through the simulation run and builds an empirical probability distribution to determine the amount of safety stock to be adjusted so that the target service level is exactly achieved. The updated safety stock values can then be tested to see if they meet the target level via another simulation.

The procedure in [1] constitutes a beginning for other related studies and practical applications. The objective of attaining the target service level may be pursued by developing a method through simulation approaches to determine the inventory policy parameters (e.g. safety stock) for various inventory systems. For example, Boulaksil et al. [7] adopts the procedure to determine the empirical probability distribution of the backorder quantities instead of the net stock levels. In their approach, net stock levels in a multi-stage inventory system are determined based on backorder quantities by solving the mathematical model repetitively in a rolling horizon.

In this paper, we refine the procedure given in [1] from the viewpoint of a practitioner so that it can be readily understood and implemented. In order for the procedure to be applied, the net requirements must be calculated through TPOP with a backordering assumption in the inventory system. In this respect, the next section will describe the netting procedure.

2. The Netting Procedure

The netting procedure determines the net requirements to be input for a lot-sizing rule that yields the decisions of when to order and how much to order. In this section, we refine the netting procedure proposed in [1] through a better algorithmic representation. We use the same notation given in [1]. The following notation is used in the netting procedure and the safety stock adjustment procedure.

Notation:

- Ψ : safety stock used in simulation runs, Ψ : initial safety stock, Ψ^* : safety stock that satisfies the target service level
- s : initial net stock, L : lead time (fixed), H : forecast horizon, N : run length, W : warm-up period length
- t : period number ($t = 1, 2, 3, \dots$), d_t : actual demand in period t
- s_t : planned net stock at the end of period t + determined at period t where $t = 1, 2, 3, \dots, N - 1$
- s_t^* : actual observed net stock at the end of period t
- f_t : forecast made at the beginning of period t for period $t + 1$ where $t = 1, 2, 3, \dots, N - 1$
- r_t : replenishment order placed at the beginning of period t arriving in the beginning of period $t + L$
- r_t^* : net requirement determined at the beginning of period t for the end of period $t + L$ where $t = 0, 1, 2, \dots, N - 1$
- s_{min} : minimum recorded actual net stock during simulation
- s_{max} : maximum recorded actual net stock during simulation
- M : number of chosen probabilities, β performance measure of ready-rate, α : target service level of β
- N_{-} : the number of periods with negative actual net stock

Since it is not specified in [1], the planned net stocks are calculated over the forecast horizon for $t = 0, 1, 2, \dots, N - 1$ by using the following expression.

$$s_t = s_{t-1} - \sum_{i=0}^{L-1} d_{t-i} + \sum_{i=0}^{L-1} r_{t-i} \tag{1}$$

The actual net stocks are calculated for each period as follows.

$$s_t^* = s_t - \sum_{i=0}^{L-1} d_{t-i} + \sum_{i=0}^{L-1} r_{t-i} \tag{2}$$

We consider the algorithm given in Exhibit 1, to determine the net requirements over the forecast horizon of $t = 0, 1, \dots, N - 1$

Exhibit-1: The Algorithm for net requirements over forecast horizon

```

for t = 0 to N - 1
  if s_t < s_min then
    s_t = 0
  else if s_t >= s_max and s_t <= s_max then
    s_t = - \sum_{i=0}^{L-1} d_{t-i} - \sum_{i=0}^{L-1} r_{t-i} - (s_t - s_min), 0
  else
    s_t = s_t
  end-if
end for
    
```

3. Safety Stock Adjustment Procedure

In this section, SSAP is given in 3 phases. The first phase determines the minimum and maximum values of the actual net stock by running a simulation up to a warm up period. The second phase continues the simulation after the warm-up period and determines the frequency function of the actual net stock. After this step, one may stop the procedure if the achieved service level is close enough to the target service level. If the desired service level is not attained, then the procedure continues with the third phase which builds the empirical probability distribution function based on the frequency function gained from the second phase. Next, the safety stock adjustment amount is calculated based on the empirical probability distribution. Then the simulation is repeated with the adjusted safety stock to see if the desired service level is actually achieved.

Phase-I: This phase determines the maximum and minimum values of the actual net stock observed during the warm-up period of the simulation. Exhibit-2 depicts the phase with an algorithm. For the fixed planning horizon (i.e. simulation run length), net requirements and associated ordering decisions are updated at each period over the forecast horizon. Therefore, the ordering decisions are subjected to change as the simulation progresses.

Exhibit-2: The algorithm for phase-I

Initialize the parameters of
Select an arbitrary value for Ψ so that $\Psi = \Psi$ and $\Psi =$
for $t = 0$ to

- 1. Determine forecast demand over the forecast horizon*
- 2. Determine the planned net stock using expression (1) over the forecast horizon*
- 3. Determine net requirements over the forecast horizon by using the algorithm given in Exhibit 1*
- 4. Apply a lot-sizing procedure to determine the size and the release-period of the order based on the determined net requirements over the forecast horizon*
- 5. Determine and record the actual net stock by expression (2)*

end-for
Determine the minimum and maximum recorded actual net stocks

Phase-II: In this phase, the frequency function of the net stock is determined over the simulation run by excluding the length of warm-up period (). Notice that no initialization occurs in this phase. The minimum and the maximum recorded actual net stock values during Phase-I are included in the frequency function in this phase. The initial values of the simulation parameters of forecasting, net requirements, outstanding orders and the actual net stock values are determined at the end of the warm-up period in Phase-I. After the second phase, the achieved service level is calculated. In this paper, we consider service level measure of ready-rate (β) which is defined as the percentage of time during which the system has positive net stock and calculated as $\beta = 1 - \frac{\text{negative net stock time}}{\text{total simulation time}}$. This phase is given in Exhibit-3 with an algorithm.

Exhibit-3: The algorithm for phase-II

for $t = +1$ to

- 1. Determine , , , , , , and by steps 1-5 given in Exhibit-2*
- 2. Construct the frequency function with respect to each observed actual net stock values during the simulation*

end-for

Phase-III: This phase determines the empirical probability distribution function to determine the adjusted safety stock (Ψ^*) that satisfies a particular target service level (). The empirical probability distribution is built based upon the observed actual net stock values from Phase-II. Each observed actual net stock value is sorted in an ascending order so that these values are stacked between the minimum and maximum observed values of actual net stock (i.e. and). The cumulative distribution function ($F(x)$) is built based on the frequencies for each observed values. As given in [1], is determined by $F(x) = \sum_{i=1}^n p_i$ for $x = 1, \dots, -1$. Then $F(x) = 0$ for $x = 0, 1, \dots$. Let Y be the safety stock adjustment amount. As specified in [1], the value of Y is calculated as follows: the particular index () is determined so that $F(x) \leq 1 - \beta \leq F(x+1)$ with the corresponding and values satisfying $F(x) \leq Y \leq F(x+1)$. After determining the corresponding values of and , Y is calculated by the following linear interpolation given by [1]:

$$Y = \frac{(\quad)}{(\quad)} \quad (3)$$

and thus,
$$\Psi^* = \Psi - Y \quad (4)$$

The algorithm given in Exhibit 4 summarizes the steps taken in Phase-III to determine the safety stock adjustment amount.

Exhibit-4: The algorithm to calculate adjusted safety stock amount

```

=
for = 0 to
    = +- ( - )
    =
    = ( )
    if ≤ 1 - ≤ then
        = and =
        break
    end if
    =
end for
use expressions (3) and (4) to calculate Ψ*

```

In order to check if the adjusted safety stock satisfies the desired service level, another simulation run is carried out based on the algorithm given in Phase-II with the updated values of $\Psi = \Psi^*$ and $X = \Psi^*$. If the desired service level is not satisfied with the adjusted safety stock level, then the procedure is repeated with a longer warm-up and simulation lengths. Notice that the empirical probability distribution is simply used to build the quantile function that maps the values of $1 -$ to the corresponding values of . A typical statistical package (e.g. JMP 8) is able to determine the desired quantiles based on the data of the observed actual net stock values. Therefore, Phase-III can simply be carried out by employing a statistical package after passing the data of the observed actual net stock values from Phase-II.

4. Experimental Study

In this section, we present the results of experiments to show how well the SSAP performs under a number of demand scenarios. Experiments are based on the simulation study that determines the sample path of the actual net stock values over a fixed planning horizon. The updated safety stock is determined based on the empirical probability distribution of the generated sample path via simulation runs. The simulation runs are executed on a spreadsheet with the length of 1,000 time periods and the warm-up period of 100 time periods. All the performance measures and cost calculations are carried out by disregarding the time horizon of the warm-up period. The maximum and the minimum values of the observed actual net stock values are recorded during the warm-up period and included into the set of observed actual net stock values over the time periods after the warm-up period. The ready-rate (β) is monitored during the simulation runs. In order to construct the empirical probability distribution, the number of chosen probabilities (K) is set equal to 300 as recommended by [1]. Demand scenarios are empirically generated as follows. Demand is assumed to be received monthly periods with seasonal variations. The planning horizon (one cycle) is assumed to be 1 year with 4 seasons each of which consists of 3 months. Also, the demand is affected by the seasonal fluctuations without exhibiting any particular trend. The demand is assumed to follow the normal distribution with different means selected from the set of {100, 50, 100, 150} for each season, respectively. Any generated negative demands are set equal to zero. The forecasting method is selected as the Winters exponential smoothing procedure [2]. The smoothing constants are selected as 0.2 for the level and 0.3 for the season. The simulation runs are carried out for each experimental factor shown in Table 1. These factors are generated by the low and high levels for the combinations of the lead-time (L), the standard deviation of the demand per period (σ), ordering cost (c_o) and holding cost per period (h).

Table-1: Experimental factors for the simulation model

Level				<i>h</i>
Low	0	10	100	1
High	4	25	500	1

The input data and the experimental factors considered in this paper are chosen to be identical as in [1] so that the results can directly be compared each other. Two types of analysis are of interest in our study: 1) Comparing the error results gained through "traditional" approach and the SSAP and 2) Comparing the total relevant costs of a number of lot-sizing rules. The objective through the experiments is to show how well the SSAP performs in determining safety stock levels to attain the desired performance measure. In this respect, errors are tabulated in Table-2 to measure the deviation from the desired performance measure for ready-rate. Let α_j be the value of the actually achieved ready-rate value for demand scenario j . Therefore, the foregoing error measure can be expressed as $\alpha_j - 0.9$. The target ready-rate value (α) is assumed to be 90% throughout our experiments. Both the traditional approach and the SSAP are repeated for each level combination under the lot-sizing rules of lot-for-lot (LFL), least unit cost (LUC) and part-period balancing (PPB). Table-2 shows the foregoing error results for each lot-sizing rule and experimental factor under the traditional approach and the SSAP. In terms of the traditional approach, the achieved ready-rate for each experimental factor is obtained as follows:

Exhibit-5: Methodology to “traditionally” determine α

- 1) Determine order size (Q) with economic order quantity formulae
- 2) Estimate the safety factor (k) by the formulae: $\alpha - 1 - \Phi(k) = \frac{(\alpha - 1)}{\sigma}$
- 3) Determine the safety stock level by the formulae: $\Psi^\# = \sigma k$
- 4) For each lot-sizing rule, run the simulation with $\Psi = \Psi^\#$ and $\alpha = \alpha$ to estimate α

For the SSAP, the achieved ready-rate value is determined after running the simulation with the adjusted safety stock level. The initial safety stock value (Ψ) in the SSAP is always set equal to 0 in our experiments. The results in Table-2 point out that error values with the SSAP are very low as compared to the traditional approach. Except for one case, most absolute error results are lower than 1% meaning that the target ready-rate of 90% is met under the safety stock level determined by the SSAP. The size of errors with the SSAP can even be decreased if the simulation length is increased to construct a more confident empirical probability distribution.

Table-2: Error results with respect to the traditional approach and the SSAP

Experimental Factors			Error under the traditional approach			Error under the SSAP		
A/h	L		LFL	LUC	PPB	LFL	LUC	PPB
100	0	10	54.33%	1.22%	-5.33%	-0.89%	-0.56%	-1.11%
100	0	25	54.44%	1.78%	-4.33%	-0.33%	-0.11%	-0.78%
100	4	10	73.67%	14.33%	6.56%	-0.56%	-0.44%	-0.56%
100	4	25	51.33%	3.89%	-1.78%	-0.22%	-0.67%	-0.33%
500	0	10	54.33%	-7.67%	-7.44%	-0.89%	-0.11%	-0.22%
500	0	25	54.44%	-7.44%	-7.78%	-0.33%	-0.22%	-0.44%
500	4	10	74.66%	5.22%	0.11%	-0.56%	-0.33%	-0.33%
500	4	25	78.66%	-1.44%	-2.78%	-0.22%	-0.11%	-0.11%

The problem of comparing lot-sizing rules is more appealing with the SSAP, since the comparison is possible predicated on the same service level. Therefore, as a second part of the analysis, we tackle the problem of comparing lot-sizing rules with 90% of ready-rate. Since the experimental environment is identical in [1], along with the lot-sizing rules considered in this paper, we also compare the lot sizing rules of Wagner-Whitin (WW), Silver-Meal (SM) and economic order quantity (EOQ) as presented in [1]. The comparisons are performed by the total relevant costs obtained by the sum of ordering and holding costs per period. The cost results associated with each lot-sizing rule are given in Table 3. The results reveal out that the total relevant costs tend to increase in the size of the lead time and the variability of the demand. In most cases the total relevant costs under EOQ, SM and WW are lower than those obtained via the lot-sizing rules considered in this paper. For the cases where the ordering costs are low, the lowest total relevant costs are achieved by LUC. On the other hand, PPB is able to generate the lowest total

relevant costs in the case of higher ordering costs. Also, for large lead time, LUC and PPB tend to yield lower costs than the other lot sizing rules.

Table-3: Cost comparisons of lot-sizing rules with 90% of 3

Experimental Factors			Gudum and de Kok [1]			Ünlü and Rossetti		
A/h	L		EOQ	SM	WW	LFL	LUC	PPB
100	0	10	123	120	117	145	134	140
100	0	25	136	135	136	150	140	137
100	4	10	141	139	144	154	140	145
100	4	25	200	197	206	157	142	148
500	0	10	299	285	288	417	307	305
500	0	25	297	296	296	415	295	293
500	4	10	320	298	292	425	308	305
500	4	25	366	356	351	419	297	294

5. Conclusion and Future Research

In this paper, we have refined the SSAP by presenting a better algorithmic representation so that practitioners can more readily understand and implement the procedure. In addition, we have extended the study by considering the problem of comparing lot-sizing rules; namely, lot-for-lot, part-period balancing and least unit cost. A series of experiments were carried out to illustrate how the safety stock adjustment procedure performs under a number of demand scenarios. The initial results presented in this paper indicate the potential of using the SSAP for determining safety stock levels to attain the desired performance measure. In addition, the procedure can be useful for determining which lot-sizing rule(s) are best for which conditions. The preliminary results were compiled based on single simulation run experiments. Further experiments are of interest to fully justify these preliminary results, which can be done by performing multiple simulation runs for ensuring a certain confidence interval. An investigation of a similar procedure that can be applicable to the classical inventory control models under complex demand structures (e.g. intermittent and highly variable demand) should be considered for future research.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. 0214478. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

1. Gudum, C. K., and de Kok T. G., 2002, "A Safety Stock Adjustment Procedure to Enable Target Service Levels in Simulation of Generic Inventory Systems." Preprint 2002-1, Copenhagen Business School (<http://ir.lib.cbs.dk/paper/ISBN/x656149131>) accessed 3/18/2008.
2. Silver, E. A., Pyke, D. F., Peterson, R., 1998, *Inventory Management and Production Planning and Scheduling*, 3rd ed. Wiley, New York.
3. Callarman, T. E., Mabert, V. A., 1978, "Using Material Requirements Planning With Demand Uncertainty," *Proceedings of the 9th Annual Midwest AIDS Conference*, 151-155.
4. Wemmerlöv, U., Whybark, D. C., 1984, "Lot-Sizing under Uncertainty in A Rolling Schedule Environment," *International Journal of Production Research*, 22 (3) 467-484.
5. Wemmerlöv, U., 1986, "A Time-Phased Order-Point System in Environments with and Without Demand Uncertainty: A Comparative Analysis Of Non-Monetary Performance Variables," *International Journal of Production Research*, 24 (2) 343-358.
6. Orlicky, J., 1975, *Material Requirements Planning*, McGraw-Hill, New York.
7. Boulaksil Y., Fransoo, J. C., van Halm, E. N. G., 2009, "Setting Safety Stocks in Multi-Stage Inventory Systems Under Rolling Horizon Mathematical Programming Models," *OR Spectrum*, 31, 121-140.