



Erasmus School of Economics

Bachelor Thesis Double Degree Programme
Econometrics & Economics

July 8th 2018

NON-STATIONARY STOCKING POLICIES IN
SERVICE LEVEL AGREEMENTS

Author: Pim van Leeuwen (412793)

Supervisor: Dr. Sena Eruguz Colak

Second Supervisor: Dr. Chiel van Oosterom

Abstract

Service levels agreements are commonly used in supply chains to commit a supplier to a certain performance target. A popular measure of performance is the item fill rate. For finite review horizons, the distribution of the item fill rate is studied by Thomas (2005), whose results are replicated in this thesis. In the extension, a similar replenishment model is studied with i.i.d demand, a single customer, no lead time and negligible ordering costs. Under these assumptions, the performance of a simple, practical and realistic non-stationary stocking policy is evaluated by Monte Carlo simulation and compared to that of a stationary base-stock policy. This policy turns out to outperform a base-stock policy based on meeting fill rate targets in the long run. However, when a high target needs to be met with high probability over a finite period of time, the base-stock policy does better and the original non-stationary policy requires adjustment. The main contribution of this work is to show that adopting a non-stationary stocking policy may save costs incurred by the supplier.

Contents

1	Introduction	4
2	Literature review	5
2.1	Previous work	5
2.2	The finite horizon fill rate	5
2.3	Influence of Thomas 2005	6
2.4	Non-stationary stocking policies	7
3	Model and notation	8
4	Replication	9
4.1	Simulation	9
4.2	Computation of the order-up-to level	10
4.3	Horizon length and demand distribution	10
4.4	Penalty cost	12
5	A non-stationary stocking policy	14
5.1	A simple policy	15
5.2	Simulation	16
5.3	Choice of the changeover period	17
6	Results	20
6.1	Horizon length and distribution parameters	20
6.2	Success probability	20
7	Conclusion	27
8	Discussion	28
	References	29
9	Appendix	31
9.1	Non-convexity of the objective function	31
9.2	Accuracy of the fill-rate approximation	32
9.3	Numerical example	33

1 Introduction

Service level agreements (SLA's) are commonly included in vendor managed inventory contracts to evaluate the supplier's stocking performance. In fact, a survey by Oblicore Inc. (2007) indicates that SLA's are used by 91% of organisations to manage suppliers, internal agreements, or customer agreements. Typically, the supplier is bound to service level constraints, where a lower bound on a particular performance measure has to be reached. Violations are penalised mostly by financial compensation. A widely used measure of performance is the item fill rate, which is defined as the fraction of demand that can be satisfied directly from stock.

Textbook solutions for setting stock levels to achieve a minimum required fill rate assume that customer demand is stationary and independent and focus on the expected fill rate for an infinite horizon. Under these assumptions, the optimal (s,S) policy is stationary. However, because the value of a performance measure such as the fill rate is unobservable if the horizon is infinite, an SLA will typically contain a review horizon of one or multiple weeks or months, after which performance is evaluated. For this reason, the distribution of a metric such as the item fill rate for finite review horizons is of much interest. Research in this field finds its origin with Thomas (2005), who studies a stationary base-stock model in a finite horizon SLA.

As Thomas (2005) mentions himself though, a stationary stocking policy is not necessarily optimal when the review horizon is finite. In this thesis, the performance of a non-stationary stocking policy is analysed and compared to the base-stock policy studied by Thomas (2005). The particular non-stationary policy studied is both simple and intuitive, such that a real-life vendor could easily implement it. The research question of this thesis will thus be: what is the effect of using a practical, non-stationary stocking policy on holding costs and the probability to meet a specified fill rate target, compared to a stationary base-stock policy?

The aim of this thesis is not to find the best possible strategic stocking policy, but rather to show the behaviour of a particular non-stationary strategy. Not much previous research has been done regarding non-stationary policies in a situation with a finite review horizon and stationary demand, so the primary relevance of this thesis is to provide insights into whether it is worth further exploring the possibilities of non-stationary stocking policies in SLA's with service level requirements and performance reviews after a finite number of replenishments. The focus in this thesis is on the item fill rate as a performance measure.

Using Monte Carlo simulation to simulate the demand replenishment cycles under various policies with different review horizons, demand variability and penalty sums, it is found that the non-stationary policy outperforms the base-stock policy for lower fill rate targets and penalties. For a higher target, combined with a more severe penalty, the reverse holds and the base-stock policy is the better choice. In this case, adjustments to the non-stationary policy are advisable.

The remainder of this thesis is organised as follows. Section 2 provides an overview of research in the field of inventory management, performance measures, SLA design and non-stationary replenishment. Section 3 introduces the model and notation used throughout the thesis. Section 4 describes the simulation procedure and replicates some of the results found by Thomas (2005). Section 5 explains the non-

stationary policy studied and determines the optimal parameters of this policy. Section 6 discusses the results of the non-stationary policy and compares them with the base-stock policy. Finally, Section 7 and 8 respectively conclude and discuss the findings and shortcoming of methods used.

2 Literature review

The following subsections give an overview of scientific work in the field of inventory management, among others introducing the work of Thomas (2005), whose work is most influential on this thesis.

2.1 Previous work

A standard textbook containing the fundamentals of inventory management is by Zipkin (2000). In a setting with i.i.d. demand, zero lead time and no ordering costs, it turns out that a base-stock policy minimises holding and shortage costs. Under a base-stock policy, at every periodic inventory review, an order is made such that stock is restored to an certain order-up to level, denoted by S . In this way, the base stock policy is a special case of the more general (s,S) policy, where a replenishment order up to S is made if and only if stock level at review is below s . The methods that Zipkin describes all implicitly assume an infinite review horizon.

The study of performance measures in inventory management and SLA's starts with Schneider (1981), who provides formal definitions of three types of measures. Although measures are distinguished by Greek letters (α, β and γ), he is already describing what is now commonly called the item fill rate (β -type). Another performance measure is the ready rate (α -type), defined by Schneider (1981) as "the probability of not being out of stock at an arbitrary moment in time". Almost twenty years later, Chen and Krass (2001) study both the ready rate and the fill rate in inventory models with minimum service level constraints. They establish the optimality of the base-stock policy in an infinite horizon model with periodic review, but otherwise rather mild assumptions. Interestingly, Chen and Krass (2001) focus on minimum values of the performance metrics, reflecting worst-case rather than expected performance. The authors illustrate the difference between minimum service level constraints and mean service level constraints with numerical examples.

2.2 The finite horizon fill rate

Until 2005, research in the field of inventory management focused exclusively on the asymptotic distribution of performance measures, meaning the review horizon was assumed to be of infinite length. In (2005), an influential paper on the fill rate was published by Thomas. He takes a significantly different approach, departing from the assumption of an infinite horizon. His motivation for this relaxation is that in an SLA specifying a performance target, there must be a review period of finite length, after which the performance of the vendor can be evaluated. Managers in the vending company therefore deal with a finite horizon. The asymptotic distribution of the fill rate has been known for some years. With demand normally distributed, Silver and Bischak (2011) derive an exact expression for the item fill rate. In his

work, Thomas (2005) studies the distribution of the fill rate for different finite review horizon lengths and Erlang distributions.

Preparatory work on the item fill rate is done by Chen et al. (2003), who prove that in an order-up-to replenishment policy with negligible lead time and i.i.d. demand, the fill rate is larger for any finite horizon than for the traditional infinite horizon. Consequently, the use of textbook methods may cause suppliers to set stocking levels inefficiently high. Banerjee and Paul (2005) extend on the results of Chen et al. (2003), proving that the expected fill rate is in fact monotonically decreasing in the number of periods in the review horizon.

Using Monte Carlo simulation, Thomas (2005) finds that indeed the mean fill rate decreases with the length of the review horizon. This, however, does not mean that a vendor should always prefer a shorter review period to a longer one. One advantage of a longer horizon is that if the supplier observes large demand realisations, there may be enough time to make up for such events and still meet the target fill rate. In effect, increasing the period length is similar to increasing the shape parameter of the Erlang distribution that Thomas (2005) uses to draw demand realisations, both resulting in lower variability in the observed fill rate. Finally, Thomas (2005) finds that when a certain probability is required to meet a fill rate target, the required order-up-to point first increases with the review horizon length, but eventually decreases with it.

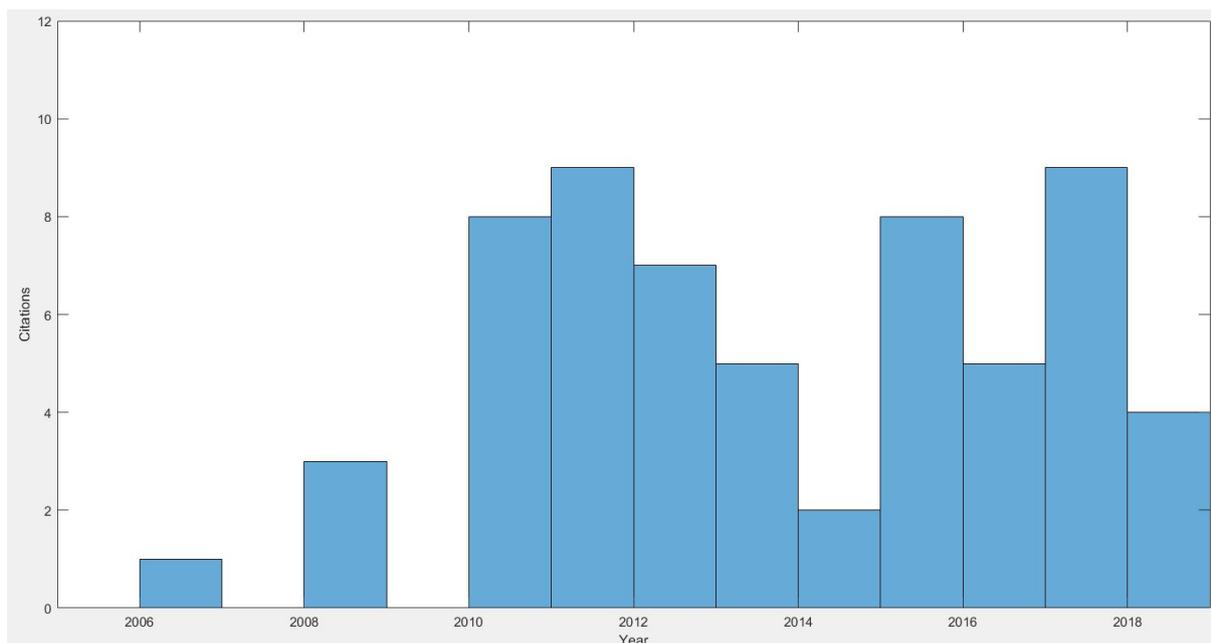
2.3 Influence of Thomas 2005

Since its publication, 64 later studies have cited the paper of Thomas (2005), mostly referring to it as "interesting" or "important" work. The number of citations in each year since its publication are shown in Figure 1, where three papers are left out because their year of publication could not be accurately traced. As can be seen, the citing of Thomas (2005) mainly takes off in 2010, but interest has not stopped in more recent years. Some papers treat one of Thomas' findings in more detail, while other relax the assumptions of the stocking model used, creating a more complex and realistic environment. Works that are most closely related to this thesis will now be discussed.

Katok et al. (2008) study the impact of two parameters of the original model by Thomas (2005): the horizon length and the bonus awarded for meeting the fill rate. Thomas (2005) originally thought of a lump-sum penalty rather than a bonus, but bonus and penalty in this case have identical effects on the desired probability to meet the target fill rate. The authors use a controlled laboratory experiment with human subjects instead of Monte Carlo simulation to arrive at their results. They find that increasing the bonus leads to a higher optimal stocking level, whereas lengthening the review horizon may either increase or decrease it.

An alternative measure of performance is the order fill rate, defined by Larsen and Thorstenson (2008) as the fraction of complete orders that can be filled directly from inventory. For a continuous review base-stock policy with constant lead time, they compare the order fill rate to the item fill-rate and find a mathematical relation between the two. Motivation for investigating the order fill rate as an alternative performance measure is that in reality, the receiver of goods in the SLA (customer) may be most interested in complete orders filled from inventory, since satisfaction of her own clients may be more

Figure 1: Citation of Thomas (2005) from 2005-2018



accurately reflected as such.

Tan et al. (2017) adopt a model similar to Thomas (2005), with a finite review horizon in which a fill rate target has to be met, stationary (Γ) demand and a base-stock policy adopted by the vendor. However, they consider a situation with constant positive lead time and find that the overstocking problem, which results from vendors using methods that are based on infinite horizon results, actually grows worse with increasing lead times.

Recently, Thomas and Chen (2018) extend the model of Thomas (2005), departing from the situation of only a single customer. In their paper, the authors relax this assumption and study the allocation of inventory with fill rate targets to multiple, homogeneous retailers. They characterise optimal allocation rules for multiple fill rate targets in both single- and multi-period review models.

2.4 Non-stationary stocking policies

Surprisingly little attention has been paid in the literature to strategic, non-stationary stocking policies to achieve a particular performance measure target. Non-stationary stocking strategies are, however, heavily researched in models with non-stationary demand. For example, Graves (1999) considers a dynamic base-stock policy in a single-item model where demand follows an ARIMA process, finding that the required safety stock is much higher than with stationary demand. A comparison of various dynamic programs and heuristics for inventory models with stochastic, non-stationary demand is provided by Bollapragada et al. (1999). Finally, Tunc et al. (2011) examine the inefficiencies of adopting a stationary stocking policy when demand is in fact non-stationary. They find that the additional costs could be very high, depending on the variability of demand. On the other hand, if information on future demand is unreliable, the stationary policy may be an efficient approximation to the optimal non-stationary policy.

With stationary demand and a particular performance target, Liang and Atkins (2013) look into

dynamic strategies in their research on optimal SLA design, describing a dynamic program in the appendix to their paper. As a performance metric, the ready rate is used and also, their demand-replenishment model differs slightly from the one used by Thomas (2005). In a lump-sum penalty contract, Liang and Atkins (2013) find that a vendor can save up to 16% of costs when using their proposed dynamic policy. The focus of the authors, however, is more on optimal SLA design than on a detailed orientation of non-stationary stocking. On this subject, Liang and Atkins (2013) find that from the perspective of the customer, it is preferable that penalties be proportional to the underperformance of the vendor, rather than lump-sum ones. With these linear penalties, the vendor has fewer incentives to behave strategically.

3 Model and notation

In Thomas (2005), as well as in this thesis, a periodic review model is studied with i.i.d. demand, a single customer and no lead time or ordering costs. In addition, the item in stock is perfectly divisible, so continuous demand is appropriate. Under these assumptions, together with an infinite review horizon, a base-stock policy will minimise holding and shortage costs, as can be found in Zipkin's textbook on inventory management (2000). Key to this result is that with an infinite review horizon, the fill-rate is not a random variable, but simply a function of the order-up to level.

For a finite horizon, a general, non-stationary base-stock policy model works as follows: at the beginning of period t , stock level is equal to s_t . Demand in period t is X_t , satisfied demand is defined as $Y_t = \min(s, X_t)$ and excess demand is backlogged. At the end of each period, the supplier incurs a holding cost h per unit left in inventory. Also at the end of every period, a replenishment order is placed such that beginning stock level in the next period equals s_{t+1} . After T periods of demand, both vendor and customer observe the item fill rate, given by

$$\alpha_T(s_1, \dots, s_T) = \frac{Y_1 + \dots + Y_T}{X_1 + \dots + X_T}. \quad (1)$$

If the fill rate is below an agreed upon target θ , the vendor has to pay the customer a lump-sum amount of B . Consequently, total costs of one review period are equal to

$$C(s_1, \dots, s_T) = \sum_{t=1}^T h(s_t - D_t)^+ + B \cdot I_{\alpha_T(s_1, \dots, s_T) < \theta}. \quad (2)$$

where $I_{\alpha_T(s_1, \dots, s_T) < \theta}$ is the indicator function that contains information of whether the fill rate at review is sufficient or not. Typically, it is assumed that the penalty B has to be paid for every period in the review horizon, or that it at least grows linearly with T . In such cases, the lump-sum penalty is actually equal to $B \cdot T$, similar to Katok et al. (2008). Note that in a stationary base-stock model, the subscript in the order-up-to level s_t disappears, resulting in $s_t = s$ for all $t = 1 \dots T$. This is exactly the approach taken by Thomas (2005). In the rest of this thesis, a stationary base-stock policy is simply referred to as a base-stock policy.

4 Replication

In order to verify the correct implementation of the model, some of the results of Thomas are now replicated.

4.1 Simulation

The performance of the base-stock model is evaluated by Monte Carlo simulation. Similar to Thomas' (2005) approach, 10 million runs are used and random numbers are drawn from the Mersenne Twister generator of Matsumoto and Nishimura (1998). Consider a stocking model with horizon length T and random demand variables drawn from an Erlang($p1, p2$) distribution. The order-up-to level s is such that the infinite horizon (long-run) fill rate is equal to performance target θ . Pseudocode for the simulation is given below.

Simulation *MonteCarlo_static*($p1, p2, T, s, \theta, n, \text{seed}$)

Inputs: $p1, p2$: parameters of the Erlang distribution

T : length of the review horizon

s : order-up-to point

θ : fill-rate target

n : number of simulation runs

seed : the seed for the random number generator

Output: α : a vector of T -horizon fill rates with size ($n \times 1$)

L : a vector of holding costs with size ($n \times 1$)

$\text{RNG} \leftarrow \text{mersenne_twister}(\text{seed})$

for $i = 1, 2, \dots, n$ **do**

 Generate $X_t \sim \text{Erlang}(p1, p2)$ using the RNG, $t = 1, 2, \dots, T$

$Y_t \leftarrow \min(X_t, s)$, $t = 1, 2, \dots, T$

$\alpha_i \leftarrow \frac{Y_1 + \dots + Y_T}{X_1 + \dots + X_T}$

$L_i \leftarrow \sum_{t=1}^T (s - Y_t)$

return α and L

Prior to the first run, the random number generator (RNG) is initialised with the inputted seed. Next, in each run, T Erlang($p1, p2$) demand variables are generated from the RNG. From these, satisfied demand is calculated for each period as the minimum of demand and starting stock level s . The fill rate at review can then be computed as in (1). Finally, holding costs (with $h = 1$) for this period are calculated and the simulation goes to the next run. After all runs are completed, a vector with size equal to the number of replications is returned, containing the fill rate of each run. The same is done for holding costs. With these vectors, all desired statistics can be computed, such as the average fill rate and holding costs, as well as the proportion of replications in which $a_i \geq \theta$. This last statistic is an estimate of the success

probability, the probability that the supplier does not have to pay penalty B .

4.2 Computation of the order-up-to level

One particular input for the Monte Carlo simulation described in the previous section is the order-up-to level s , which is chosen such that the long-run fill rate is equal to the performance target θ . The long-run fill rate is given by

$$\lim_{T \rightarrow \infty} \mathbb{E}[\alpha_T(s)] = \frac{\mathbb{E}[Y]}{\mathbb{E}[X]}. \quad (3)$$

As shown in Chen et al (2003), this expression is equivalent to $1 - \frac{\text{Expected units short in one period}}{\text{Expected demand in one period}}$. Silver et al. (2016) then give the following expression for the infinite horizon fill rate, where $f(x)$ is the density function of the demand.

$$\lim_{T \rightarrow \infty} \mathbb{E}[\alpha_T(s)] = 1 - \frac{\int_s^\infty (x-s)f(x)dx}{\mathbb{E}[X]}. \quad (4)$$

Using this expression, the s that leads to a long-run fill rate of θ is calculated with bisection. Pseudo code of the algorithm is given below.

Algorithm $s_longrun(p1, p2, lb, ub, \theta, \epsilon, n)$

Inputs: $p1, p2$: parameters of the Erlang distribution

lb, ub : lower and upper bound on s

θ : fill-rate target

ϵ : tolerance level

while($ub - lb > \epsilon$) **do**

$s \leftarrow (lb + ub)/2$

$\alpha \leftarrow 1 - \frac{\int_s^\infty (x-s)f(x; p1, p2)dx}{\mathbb{E}[X]}$

if ($\alpha < \theta$) **then**

$lb \leftarrow s$

else

$ub \leftarrow s$

return s

Regarding the inputted lower and upper bound of s , a necessary condition for convergence is that $\lim_{T \rightarrow \infty} \mathbb{E}[\alpha_T(lb)] < \theta$ and $\lim_{T \rightarrow \infty} \mathbb{E}[\alpha_T(ub)] > \theta$. Even for very conservative bounds on s and low tolerance level ϵ , the bisection algorithm converges rather quickly, typically needing less than 15 iterations. Bisection is appropriate because the long-run fill rate increases monotonically in s .

4.3 Horizon length and demand distribution

Table 1 gives the results of the simulations for various review lengths and Erlang($m,1$) distributions. The output that Thomas obtained in his runs are provided in Table 2. The order-up-to points are all based on a 95% long-run fill rate, resulting in $s \approx 2.9963, 5.1855, 7.2633$ and 11.2860 for Erlang scale parameters 1,

3, 5 and 9, respectively, all calculated using the bisection algorithm with $\epsilon = 0.00001$. Although Thomas does not provide output values of the success probability and holding costs, these are presented in Table 1 for later reference. To be able to compare holding costs across different horizon lengths, they are given per decision period, so total holding costs are divided by T .

Table 1: Simulation output from the model of Thomas (replicated)

m-Erlang	Long-run fill rate (%)	Horizon length T	Mean	Median	Std. deviation	Skewness	Succes prob. (%)	Holding costs
5	95	5	0.9577	0.9888	0.0585	-1.5888	0.6748	2.5135
5	95	10	0.9541	0.9647	0.0448	-1.0903	0.6114	2.5136
5	95	20	0.9521	0.9567	0.0330	-0.7528	0.5757	2.5136
5	95	100	0.9504	0.9513	0.0153	-0.3270	0.5327	2.5134
1	95	20	0.9568	0.9790	0.0549	-1.5428	0.6614	2.0463
3	95	20	0.9530	0.9601	0.0388	-0.9335	0.5938	2.3357
5	95	20	0.9521	0.9567	0.0330	-0.7528	0.5757	2.5136
9	95	20	0.9514	0.9544	0.0276	-0.5952	0.5601	2.7362

Replicated output from the base-stock model for various review periods and Erlang scale parameters

Table 2: Simulation output from the model of Thomas (original)

m-Erlang	Long-run fill rate (%)	Horizon length T	Mean	Median	Std. deviation	Skewness	Succes prob. (%)	Holding costs
5	95	5	0.9578	0.9894	0.0585	-1.5955	-	-
5	95	10	0.9541	0.9647	0.0448	-1.0908	-	-
5	95	20	0.9521	0.9566	0.0330	-0.7509	-	-
5	95	100	0.9504	0.9514	0.0153	-0.3262	-	-
1	95	20	0.9567	0.9793	0.0550	-1.5416	-	-
3	95	20	0.9530	0.9599	0.0388	-0.9339	-	-
5	95	20	0.9521	0.9566	0.0330	-0.7509	-	-
9	95	20	0.9514	0.9542	0.0276	-0.5934	-	-

Output from the base-stock model for various review periods and Erlang scale parameters, as presented in Thomas (2005)

For the mean, median, standard deviation and skewness, the two tables show mostly identical values when rounded to 3 decimals. Therefore, it can be concluded that the output presented by Thomas (2005) is correct and that differences with the values presented in Table 1 are solely due to randomness. This most probably occurs because of differences in seeds used for random number generation. In all simulations in this thesis, the seed "12345" is used, whereas Thomas does not specify the seed used in his simulations.

Comparing the output from the simulation for different inputs, one can clearly see that Thomas was

correct in noting that increasing the review horizon has a similar effect as increasing the Erlang scale parameter, both reducing the variability of the item fill rate. Also, the mean fill rate clearly decreases with the horizon length, as proven mathematically by Bannerjee and Paul (2005).

4.4 Penalty cost

Next to the length of the review horizon and the demand distribution, Thomas (2005) also considers how the (stationary) order-up-to point s changes with the costs of failing to meet the target fill rate, which is the penalty B . An increase in B (relative to h) means that the probability of meeting target θ should increase. Higher levels of s result in higher inventory costs, but lower expected penalty costs. Refraining from specifying a dollar amount for B , Thomas shows how s varies when a fill rate θ has to be achieved with a certain probability of success (λ) for different horizon lengths and demand distributions. For inputted Erlang parameters $p1$ and $p2$, review horizon T , lower and upper bounds lb and ub , tolerance level ϵ and the number of simulation runs, the following bisection algorithm finds s such that the probability of meeting fill rate θ is equal to λ .

Algorithm $s_longrun(p1, p2, T, lb, ub, \theta, \lambda, \epsilon, n, seed)$

Inputs: $p1, p2$: parameters of the Erlang distribution

T : length of the review horizon

lb, ub : lower and upper bounds on s

θ : fill-rate target

λ : desired success probability

ϵ : tolerance level

n : number of simulation runs

$seed$: seed for the simulation

while($ub - lb > \epsilon$) **do**

$s \leftarrow (lb + ub)/2$

$\alpha \leftarrow$ fill rate output obtained from $MonteCarlo_static(p1, p2, T, s, \theta, n, seed)$

$\lambda^* \leftarrow$ fraction of runs in which $\alpha > \theta$

if ($\lambda^* < \lambda$) **then**

$lb \leftarrow s$

else

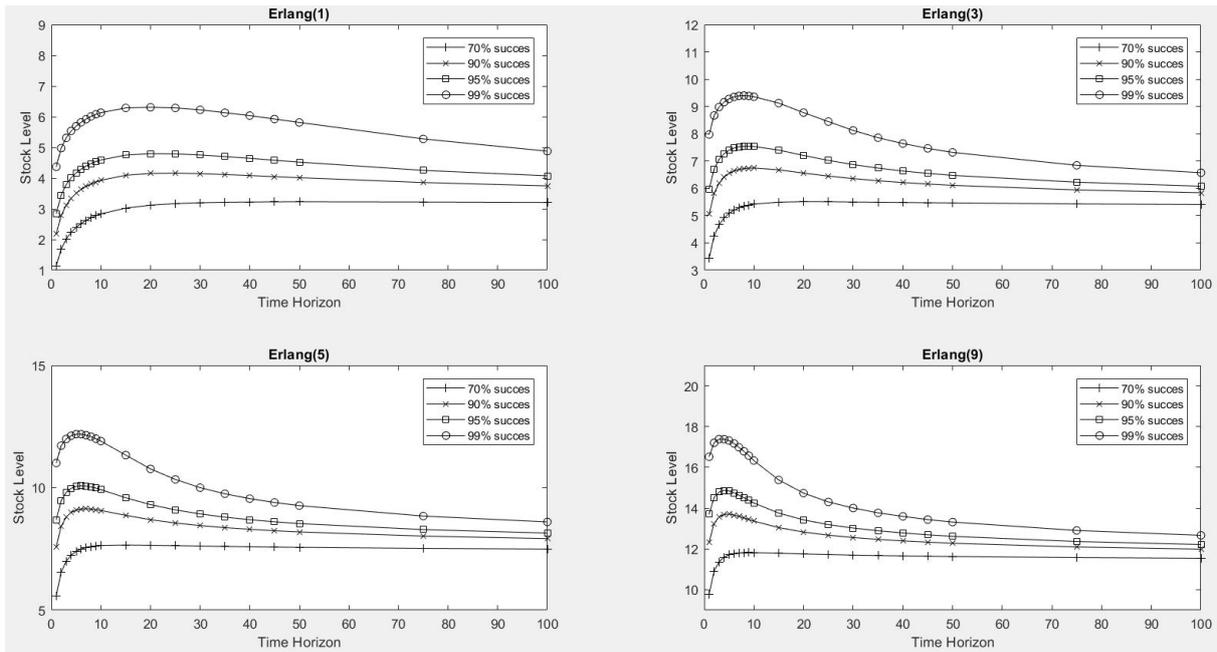
$ub \leftarrow s$

return s

The replicated graphs for $\theta = 0.95$ and $\theta = 0.99$ are shown in Figure 2 and 3, respectively. Comparing these to Figure 3 and 4 in Thomas (2005), no significant differences are observed, suggesting that the results have been successfully replicated. The influence of randomness is negligible here, since differences of less than 0.01 are unobservable in the graphs. As mentioned by Thomas (2005), the necessary

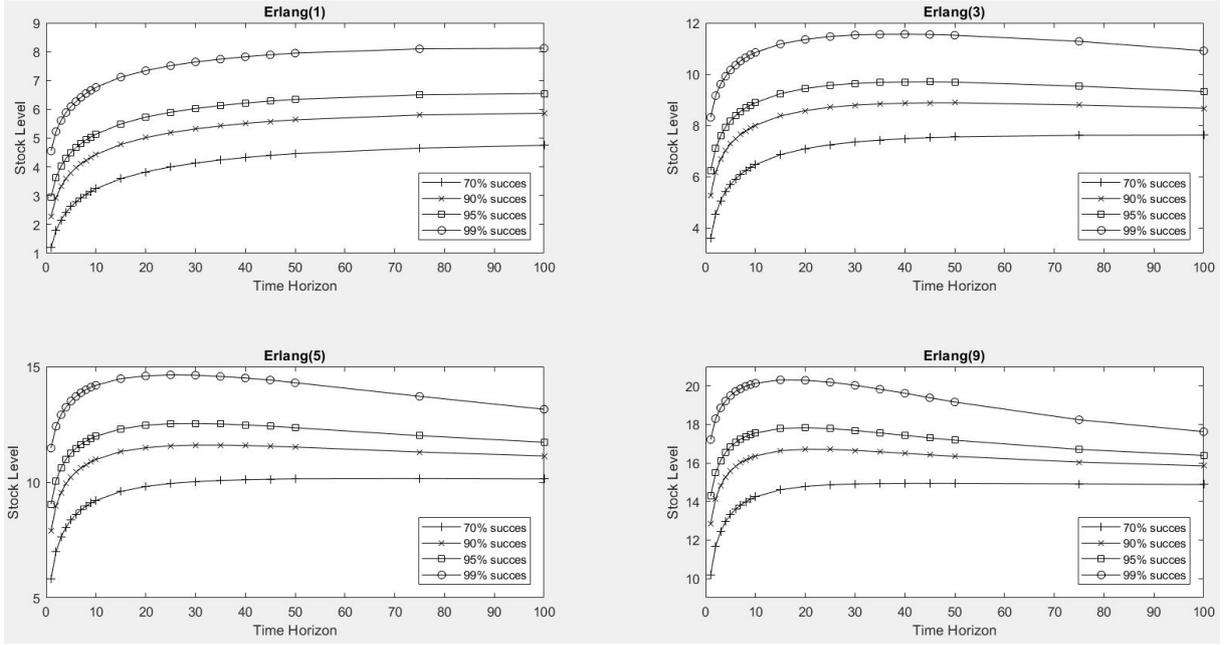
stock levels for all success probabilities initially increase with the review horizon and eventually decrease. The reason is that initially, the reduction in the expected fill rate outweighs the reduction in variance caused by a longer horizon length, but the reverse holds for larger values of T . This change occurs faster for increasing Erlang parameters, which results in lower demand variability. For each success probability λ , the turning point after which the required order-up-to point s decreases in T occurs later with a target of 99% than with a target of 95%. An explanation here is that with a higher target θ , it takes longer to 'recover' from a single large demand realisation.

Figure 2: Stock levels required to have a specified chance of meeting a 95% target fill rate



Stock levels required to meet a 95% target fill rate 70, 90, 95 and 99% of the time, with a review horizon ranging from 1 to 100.

Figure 3: Stock levels required to have a specified chance of meeting a 99% target fill rate



Stock levels required to meet a 99% target fill rate 70, 90, 95 and 99% of the time, with a review horizon ranging from 1 to 100.

5 A non-stationary stocking policy

One idea for a non-stationary stocking policy is to create an algorithm that finds a different s for every replenishment, similar to what Liang and Atkins (2013) do with the ready rate as a performance measure. Prior to every stocking decision, an optimisation problem could be solved that tries to minimise expected costs. Before every period i , the objective would thus be to set the order-up-to level s_i such as to minimise:

$$C(s_i) = \mathbb{E} \left[\sum_{t=i}^T h(s_i - D_t)^+ + B \cdot I_{\alpha_T(s_1, \dots, s_{i-1}, s_i) < \theta} \right] \quad (5)$$

Note that holding costs incurred prior to period i are irrelevant in the determination of s_i , since these remain unchanged regardless of any decision. In order to estimate the expected value in (5), simulation is needed to estimate $\mathbb{E}[(s_i - D_t)^+]$ as well as $\mathbb{P}[\alpha_T(s_1, \dots, s_{i-1}, s_i) < \theta \mid X_1, \dots, X_{i-1}]$. Since this needs to be done for every evaluation in the search of the optimal s_i , a nested simulation approach is required. Consequently, the complete procedure will take a lot of computation time. Optimisation is complicated even further by the fact that the function in (5) is not convex, as is shown in the Appendix. Altogether, it is practically infeasible to fairly compare this policy to the base-stock policy of Thomas (2005), since the number of runs would need to be reduced very heavily in order to have reasonable running times. It is therefore decided to look at a simpler policy that is both intuitive and realistic.

5.1 A simple policy

Consider a review period of length T and order-up-to level s such that the long-run fill rate meets the target θ . The long-run fill rate is computed as given in (3). This is the order-up to point that Thomas (2005) statically uses for every stocking policy in his model. In the non-stationary strategy considered in this thesis, the supplier may change s to a different order-up-to level s' at the end of period t' , hereafter referred to as the changeover period. At this point in time, the item fill rate is $\alpha_{t'}(s)$, which is calculated as in (1), with $s_1, \dots, s_T = s$. Let $\delta_{t'}$ be the fraction of demand satisfied from stock in periods $t' + 1$ through T , that is,

$$\delta_{t'} = \frac{Y_{t'+1} + \dots + Y_T}{X_{t'+1} + \dots + X_T} \quad (6)$$

For any T and t' with $T \geq 2$ and $1 \leq t' \leq T - 1$, an approximation for the the item fill rate is given by

$$\alpha'_T(s) = \frac{t'}{T} \cdot \alpha_{t'} + \frac{T - t'}{T} \cdot \delta_{t'} \quad (7)$$

The approximation is simply a waited average of the fill rates in the two periods with different order up-to-points. As it turns out, this slightly overestimates the actual fill rate, although bias decreases with the length of the review horizon. See the Appendix for a detailed evaluation of the approximation's accuracy, as well as a proof for its positive bias.

In the non-stationary stocking policy analysed, the vendor initially sets s such that the long-run fill rate is θ . From period t' onwards, the order-up to point is set to s' , such that that the long-run fill rate is equal to $\delta_{t'}^*$, where

$$\delta_{t'}^* = \frac{T \cdot \theta - t' \cdot \alpha_{t'}}{T - t'} \quad (8)$$

If $\delta_{t'}^* \geq 1$, set s' to zero, since it is deemed impossible to reach the target, and it is best to save as much inventory costs as possible. In other cases, s' will only differ from s if $\delta_{t'}^* \leq \theta$. Because of the positive bias in the approximation given in (7), it is advisable not to adjust s downwards. Regardless of the choice of t' , the success probability is expected to be higher than under the base-stock policy that sets s such that the long-run fill rate is θ . There is only one possible way in which this base-stock policy meets the target fill rate while the non-stationary policy does not. In this scenario, $\delta_{t'}^*$ would be bigger than one, but the target could still be met with positive probability. Given that the approximation is actually positively biased, this is not very likely. Also, even if it did occur, there is still only a small probability that the base-stock policy would actually meet the target fill rate. Proof that this scenario not impossible is provided by a numerical example in the Appendix.

To illustrate the mechanics of the non-stationary policy, consider Erlang (5, 1) demand and a review length of 5, with $\theta = 0.95$ and $s \approx 7.26$. Let t' be equal to 4, meaning we can only change s one period before T . If after 4 stocking decisions $\alpha_4(s) = 0.945$, then the desired fill rate for period 5 is $\delta_4^* = 5 \cdot 0.95 - 4 \cdot 0.945 = 0.97$. To achieve this long-run fill-rate, s' would have to be set to roughly 8.09.

The non-stationary policy is expected to outperform the base-stock policy studied by Thomas (2005) in terms of success probability, but possibly also in holding costs, since these are not incurred after t' if $\delta_{t'}^* > 1$. Since the policy is quite simple, other, more complex strategies may be able to outperform it considerably. However, simplicity has the attractive feature of not requiring many computations. Consequently, even a simulation with 10 million replications can be ran within a reasonable amount of time, which enables fair comparison to the base-stock policy.

In addition, this policy may be a realistic representation of how a naive vendor would behave. No simulation is required to calculate $\delta_{t'}^*$ and the order-up-to level s needed to meet this fill rate in the long-run can simply be tabulated for the appropriate demand distribution. Also, it seems reasonable that a vendor would want to evaluate his fill rate performance before the review horizon is over, adjusting the stocking policy if necessary.

5.2 Simulation

The performance of the non-stationary stocking policy is evaluated using simulation. With inputted Erlang parameters, review horizon, changeover period, initial order up-to-level, target fill rate, sample size and the seed for random number generation, pseudocode for the simulation is given below.

In each run, T demand Erlang($p1$, $p2$) variables are generated from the RNG. From these, satisfied demand is calculated for the first t' periods and afterwards, $\alpha_{t'}$ and $\delta_{t'}^*$ are computed. If $\delta_{t'}^* < \theta$, s is left unchanged and if $\delta_{t'}^* > 1$, s' is set to zero, because it is deemed impossible to meet fill rate target θ after T periods. If $\theta \leq \delta_{t'}^* \leq 1$, s' is set such that the long run fill rate is $\delta_{t'}^*$. To find this s' , the bisection algorithm of Section 4 is used. However, for 10 million replications, using this algorithm in each run would result in very a long running time. Therefore, it is advised to tabulate order-up-to points for fill rate targets ranging from 0 to 1 (step size 0.0001) for the chosen Erlang distribution before executing the simulation. After s' has been set, satisfied demand is calculated for periods $t' + 1$ to T , and α_T is computed. Finally, calculate the holding costs for this period and go to the next run. After all runs are completed, a vector with size equal to the number of replications is returned, containing the fill rate of each run. The same is done for holding costs.

Simulation *MonteCarlo_dynamic*($p1$, $p2$, T , t' , s , θ , n , $seed$)

Inputs: $p1$, $p2$: parameters of the Erlang distribution

T : length of the review horizon

t' : the changeover period

s : initial order-up-to point

θ : fill-rate target

n : number of simulation runs

$seed$: the seed for the random number generator

Output: α : a vector of T -horizon fill rates with size ($n \times 1$)

h : a vector of holding costs with size ($n \times 1$)

```

RNG  $\leftarrow$  mersenne_twister(seed)
for  $i = 1, 2, \dots, n$  do
  Generate  $X_t \sim \text{Erlang}(p1, p2)$  using the RNG,  $t = 1, 2, \dots, T$ 
   $Y_t \leftarrow \min(X_t, s)$ ,  $t = 1, 2, \dots, t'$ 
   $\alpha_i \leftarrow \frac{Y_1 + \dots + Y_{t'}}{X_1 + \dots + X_{t'}}$ 
   $\delta^* \leftarrow \frac{T \cdot \theta - t' \cdot \alpha_{t'}}{T' - t'}$ 
  if ( $\delta^* < \theta$ ) then
     $s' \leftarrow s$ 
  else if ( $\delta^* \geq 1$ ) then
     $s' \leftarrow 0$ 
  else
    set  $s'$  such that the long run fill rate is  $\delta^*$ 
   $Y_t \leftarrow \min(X_t, s')$ ,  $t = t'+1, t'+2, \dots, T$ 
   $\alpha_i \leftarrow \frac{Y_1 + \dots + Y_T}{X_1 + \dots + X_T}$ 
   $h_i \leftarrow \sum_{t=1}^{t'} (s - Y_t) + \sum_{t=t'+1}^T (s' - Y_t)$ 
return  $\alpha$  and  $h$ 

```

5.3 Choice of the changeover period

With given demand distribution and review horizon, the vendor is left with the decision of how to set the changeover period t' . This choice depends on T and in the sequel, it is assumed that $t' = \lceil cT \rceil$, where $0 < c < 1$ and $\lceil \cdot \rceil$ is the ceiling function. Consequently, the vendor effectively chooses c , facing a trade-off between holding costs and success probability. Lower values of c provide a higher probability of success, since there is a longer period to repair an unsatisfactory performance in periods 1 through t' . For higher values of c , there is a higher probability that $\delta_{t'}^* > 1$, in which case holding costs can be saved by setting s' to zero.

To evaluate the effect of a certain choice for c , a number of scenarios are analysed, focusing on extreme cases. The following variations of simulation parameters are used:

- The review horizon can be either short ($T = 5$) or long ($T = 100$)
- Demand can follow a distribution with high (Erlang(1, 1)) or low (Erlang(9, 1)) variability.
- The target fill rate can be low ($\theta = 0.70$) or high ($\theta = 0.99$).
- The financial penalty can be mild ($B = 1$), moderate ($B = 10$) or high ($B = 100$).

In total, this gives 24 scenarios. Unit holding cost h are assumed to be 1. This can be done without loss of generality, because for comparison of total costs, all that matters is the ratio between B and h . For each scenario, total costs are calculated as in (2) for four different values of c : 0.5, 0.65, 0.8 and 0.95. For $T = 5$, the dynamic stocking policies for $c = 0.65$ and $c = 0.8$ coincide, since $\lceil 0.65T \rceil = \lceil 0.8T \rceil = 4$. The results of each scenario are presented in Table 3 as costs per period (stocking decision), so it is implicitly assumed that B would grow proportionally with T . From an SLA design perspective, this seems to be a

reasonable assumption, as was discussed in Section 3. Therefore, B is actually the penalty per stocking decision in the review horizon, similar to the methodology of Katok et al. (2008).

From Table 3, it appears that in most scenarios, the best choice of the changeover period is halfway through the review horizon, meaning $c = 0.5$. Higher values of c do in some cases yield lower costs for the lowest penalty. However, in real life, an efficient SLA with a lump-sum penalty will most likely set B rather high, in order to prevent the vendor from ceasing to supply any inventory. More interesting is that for a longer review horizon and lower target θ , alternative values for c are preferred, most notably $c = 0.8$. An explanation for this may be that for a long review period, combined with a lower target fill rate, there is enough time to repair an insufficient fill rate in fewer stocking periods. Throughout the rest of the thesis, the main focus is on the stocking policy where $c = 0.5$, but one where $c = 0.8$ will sometimes be included in the analysis for comparison.

Table 3: Scenario analysis of total costs for the non-stationary stocking policy

	T	Erlang	θ	B	$c = 0.5$	$c = 0.65$	$c = 0.8$	$c = 0.9$
Scenario 1	5	(1,1)	0.70	1	0.8840	0.8543	0.8543	0.8753
Scenario 2	5	(1,1)	0.70	10	3.8976	4.0646	4.0646	4.2166
Scenario 3	5	(1,1)	0.70	100	34.0332	36.1673	36.1673	37.6295
Scenario 4	5	(1,1)	0.99	1	3.6212	3.6335	3.6335	3.6594
Scenario 5	5	(1,1)	0.99	10	4.0286	4.0413	4.0413	4.0667
Scenario 6	5	(1,1)	0.99	100	8.1020	8.1185	8.1185	8.1398
Scenario 7	5	(9,1)	0.70	1	0.7548	0.7007	0.7007	0.7268
Scenario 8	5	(9,1)	0.70	10	4.0166	4.1284	4.1284	4.8805
Scenario 9	5	(9,1)	0.70	100	36.6349	38.4046	38.4046	46.4183
Scenario 10	5	(9,1)	0.99	1	5.4988	5.5745	5.5745	5.7314
Scenario 11	5	(9,1)	0.99	10	7.0945	7.1749	7.1749	7.3301
Scenario 12	5	(9,1)	0.99	100	23.0525	23.1787	23.1787	23.3301
Scenario 13	100	(1,1)	0.70	1	0.9029	0.8958	0.8892	0.8989
Scenario 14	100	(1,1)	0.70	10	4.2096	4.0250	3.9591	4.3735
Scenario 15	100	(1,1)	0.70	100	37.2765	35.3175	34.6587	39.1198
Scenario 16	100	(1,1)	0.99	1	3.6761	3.6886	3.7637	3.8480
Scenario 17	100	(1,1)	0.99	10	6.6702	6.7251	6.8424	6.9459
Scenario 18	100	(1,1)	0.99	100	36.6112	37.0903	37.6298	37.9245
Scenario 19	100	(9,1)	0.70	1	0.6530	0.6336	0.6195	0.6263
Scenario 20	100	(9,1)	0.70	10	4.0063	3.7749	3.5466	3.5457
Scenario 21	100	(9,1)	0.70	100	37.5390	35.1882	32.8179	32.7398
Scenario 22	100	(9,1)	0.99	1	5.7862	5.6570	5.6719	5.7896
Scenario 23	100	(9,1)	0.99	10	9.0845	9.1277	9.3936	9.6781
Scenario 24	100	(9,1)	0.99	100	42.0681	43.8350	46.6109	48.5628

Total costs of the non-stationary stocking policy in 24 scenarios for different values of c . For each scenario, the lowest total cost value is in bold.

6 Results

In this section, the effect of the horizon length, demand variability and required success probability on the distribution of the fill rate is examined under the non-stationary stocking policy.

6.1 Horizon length and distribution parameters

Table 4 shows statistics of the fill rate and holding costs for $T = 5, 10, 20$ and 100 for the non-stationary stocking policy with $c = 0.5$. The same statistics are shown for $T = 20$ and Erlang shape parameters $1, 3, 5,$ and 9 . Table 2 contains all the same information, but for a policy where $c = 0.8$. The target fill rate in both cases is 95% , so that the outcomes in Table 4 and 5 can be directly compared to those in Table 1, the latter referring to the base-stock policy that sets s such that the infinite horizon fill rate would be 95% . The last two columns of Table 4 and 5 contain the differences in respectively the success probability and holding costs between the non-stationary policy and the base-stock policy.

Both Tables 4 and 5 show that in all settings, the mean fill rate is quite far below the target fill rate, but that the median fill rate is always above 95% . Not surprisingly then, the standard deviations are much higher than those reported in Table 1. As with the stationary base-stock policy, the variability of the fill rate decreases when either T or the scale parameter increases. Comparing the non-stationary cases with $c = 0.5$ and 0.8 , the trade-off between success probability and holding costs becomes clearly visible. In all settings, choosing $t' = 0.5T$ leads to both higher probabilities of meeting the target and higher holding costs. For realistically high lump-sum penalties, the vendor should therefore prefer changing policy halfway through the review horizon. Looking back at the base-stock policy results of Table 1, the non-stationary policies clearly show superior performance in terms of success probability and holding costs. In fact, Table 5 shows that in the given scenarios, the dynamic policy with $c = 0.8$ always outperforms the base-stock policy in terms of both success probability and holding costs. This means that it will result in lower total costs, regardless of unit holding cost h and penalty B .

Figure 4 shows histograms of the fill rate distributions for the base-stock policy and the non-stationary policy with $c = 0.5$ and 0.8 . Review horizon length T equals 20 and demand is Erlang($5, 1$). Both dynamic policies clearly show higher variances in the fill rate, abandoning any efforts to meet the target when this is impossible from the changeover period. Also, whereas the distribution of the base-stock policy seems rather smooth around the 95% threshold, both dynamic policies show a clear jump. This shows their strategic nature. The clusters of fill rates far below 95% correspond to runs where one or multiple high demand realisations were observed, such that it was deemed impossible to adjust and recover in periods $t' + 1$ to T , and so s' was set to zero. For $c = 0.5$, such runs occur less frequently than for $c = 0.8$, but if one occurs, the fill rate at review is much lower, because no inventory is available for half of the review horizon.

6.2 Success probability

In the final part of his paper, Thomas (2005) discusses the changes in the order-up-to point when a target fill rate θ has to be achieved with probability λ . As discussed in Section 4, where his results

Table 4: Simulation output from the non-stationary policy with $c = 0.5$

m-Erlang	Long-run fill rate (%)	Horizon length T	Mean	Median	Std. deviation	Skewness	Succes prob.	Holding costs
5	95	5	0.8980	0.9888	0.1707	-1.5695	0.0069	-0.1517
5	95	10	0.8864	0.9653	0.1882	-1.7803	0.0275	-0.1009
5	95	20	0.8978	0.9588	0.1686	-2.1577	0.0567	-0.0109
5	95	100	0.9464	0.9540	0.0635	-7.1722	0.1144	0.1401
1	95	20	0.8959	0.9790	0.1855	-1.8958	0.0121	-0.1105
3	95	20	0.8904	0.9615	0.1806	-1.9289	0.0404	-0.0645
5	95	20	0.8978	0.9588	0.1686	-2.1577	0.0567	-0.0109
9	95	20	0.9097	0.9570	0.1502	-2.5833	0.0745	0.0724

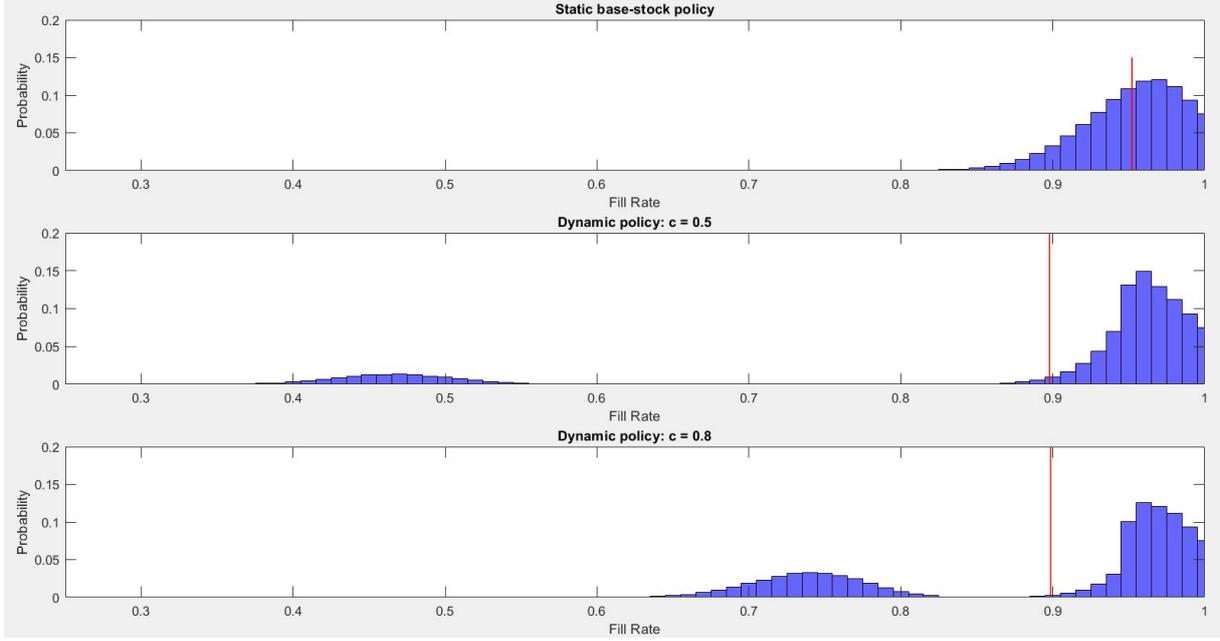
Simulation output from the dynamic stock model with $c = 0.5$ for various review periods and Erlang scale parameters. The last two columns contain the differences in the success probability and holding costs between this policy and the base-stock policy.

Table 5: Simulation output from the non-stationary policy with $c = 0.8$

m-Erlang	Long-run fill rate (%)	Horizon length T	Mean	Median	Std. deviation	Skewness	Succes prob.	Holding costs
5	95	5	0.9161	0.9888	0.1232	-1.2561	0.0011	-0.1196
5	95	10	0.9025	0.9647	0.1181	-0.9530	0.0065	-0.1225
5	95	20	0.8985	0.9569	0.1095	-0.8862	0.0187	-0.1225
5	95	100	0.9119	0.9526	0.0872	-1.3674	0.0769	-0.0448
1	95	20	0.9132	0.9790	0.1229	-1.2816	0.0011	-0.1005
3	95	20	0.9001	0.9601	0.1141	-0.9281	0.0113	-0.1209
5	95	20	0.8985	0.9569	0.1095	-0.8862	0.0187	-0.1255
9	95	20	0.8988	0.9550	0.1046	-0.8978	0.0292	-0.1166

Simulation output from the dynamic stock model with $c = 0.8$ for various review periods and Erlang scale parameters. The last two columns contain the differences in the success probability and holding costs between this policy and the base-stock policy.

Figure 4: Distributions of the fill rate for three different stocking policies



Distributions of the fill rate for the stationary base stock policy and two dynamic policies with $c = 0.5$ and $c = 0.8$

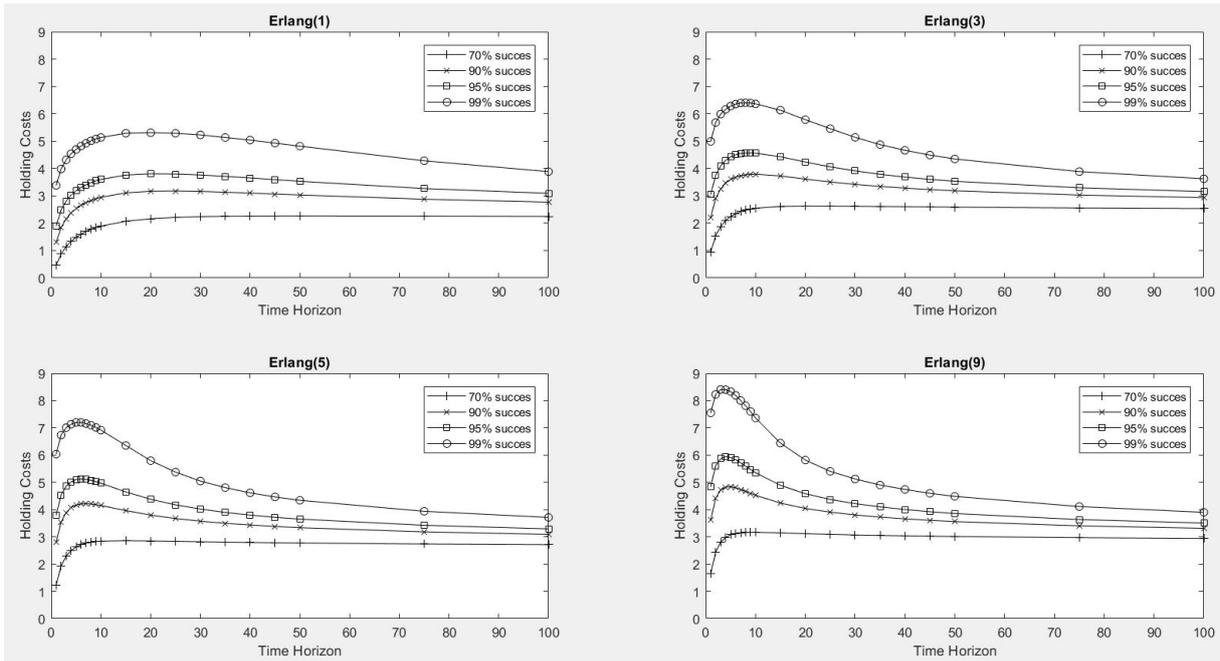
are replicated, the static order-up-to point s initially increases with the review horizon and eventually decreases. Performing the same analysis on the non-stationary policy of this paper, it is more interesting to look at the change in holding costs. Because of the workings of the policy, there is not always a single order-up-to point involved with meeting the target with probability λ . To be able to compare the base-stock and non-stationary policies, the results of Figure 2 and 3 are presented again in Figure 5 and 6, but this time with holding costs on the y-axes, rather than stock levels. Unit cost h are still assumed to be equal to 1. The shape of the holding cost figures are very similar to the original graph, which is as expected, since holding costs are simply a function of s for the static base-stock policy.

Figure 7 shows the holding costs involved with meeting a 95% fill rate with specified probabilities for the non-stationary policy, where strategic changeover occurs halfway through the review horizon. The results are obtained by using a bisection algorithm similar to the one described in Section 4.4, except that the success probability and holding costs come from running the *MonteCarlo_dynamic* simulation.

As with the base-stock policy, holding costs initially increase and eventually decrease as T increases. However, the decline sets in at a quicker rate for the base-stock policy, which generally outperforms this non-stationary strategy for the two highest requested success probabilities. This last observation is difficult to make just by comparing Figures 5 and 7, but becomes clearer in Figure 8, which shows the difference in holding costs between the dynamic ($c = 0.5$) and the base-stock policy. Positive differences indicate that the holding costs of the base-stock policy are lower.

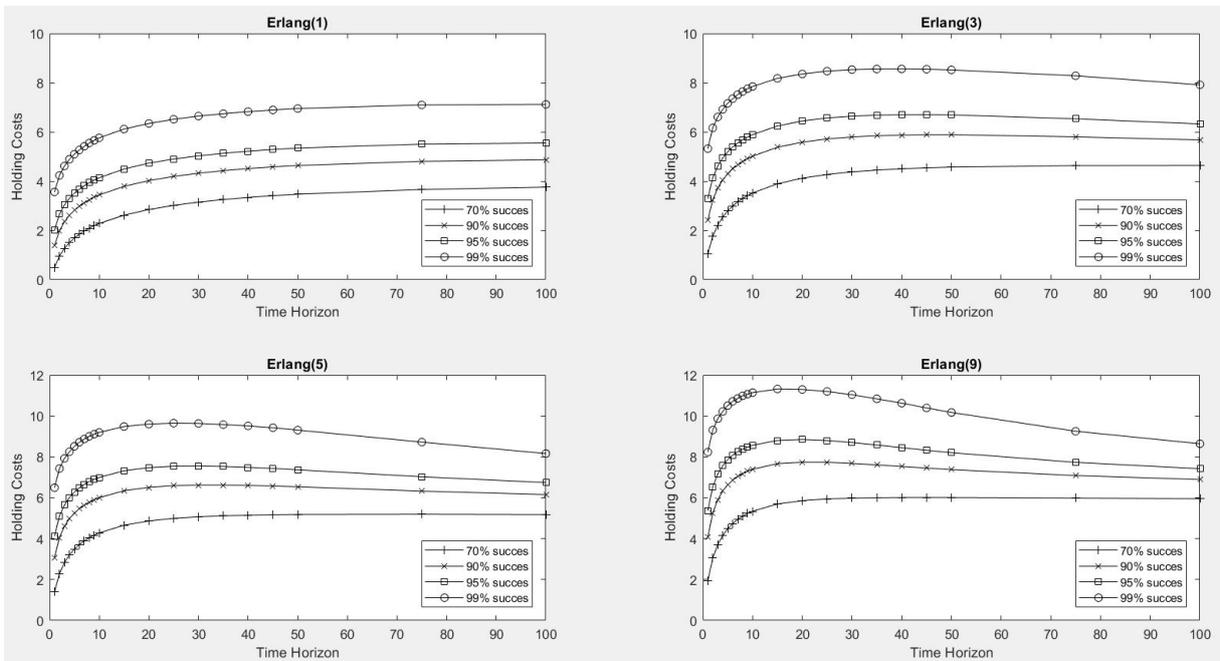
It is quite striking that the base-stock policy does so well compared to the non-stationary policy, since the latter specifically adjusts s to meet the target fill rate θ with higher probability. Therefore, it was expected to require a lower initial s to achieve the fill rate θ with probability λ , which may cause lower

Figure 5: Holding costs required to have a specified chance of meeting a 95% target fill rate



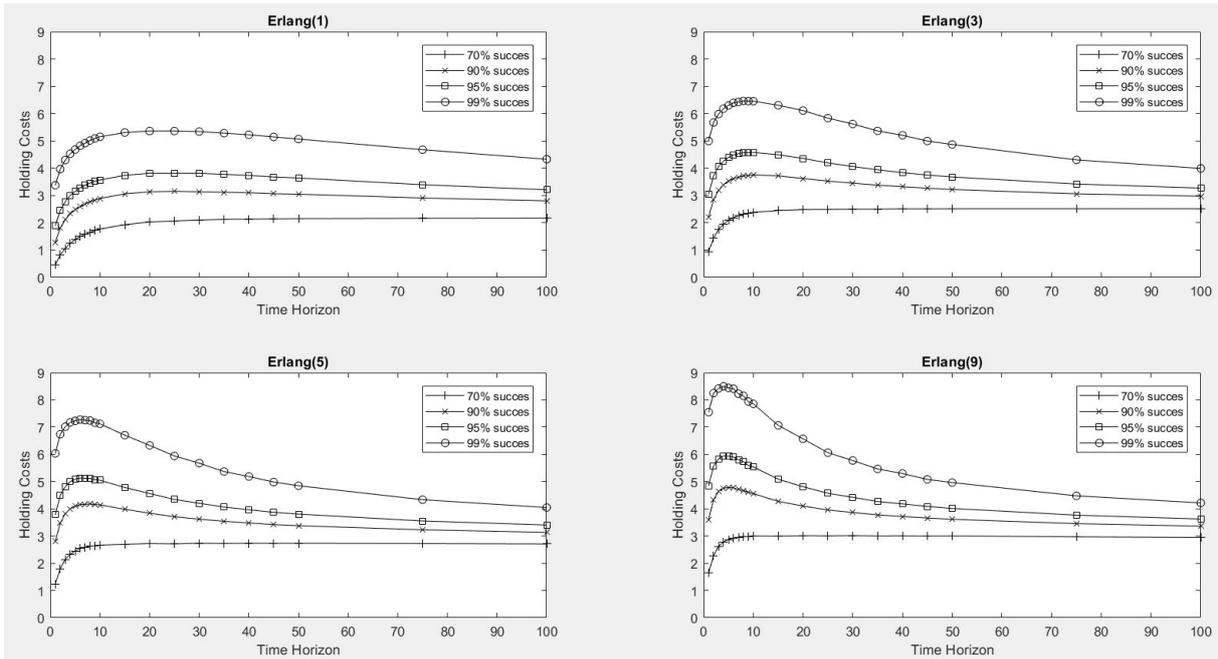
Holding costs required to meet a 95% target fill rate 70, 90, 95 and 90% of the time for the base-stock policy. The review horizon ranges from 1 to 100.

Figure 6: Holding costs required to have a specified chance of meeting a 99% target fill rate



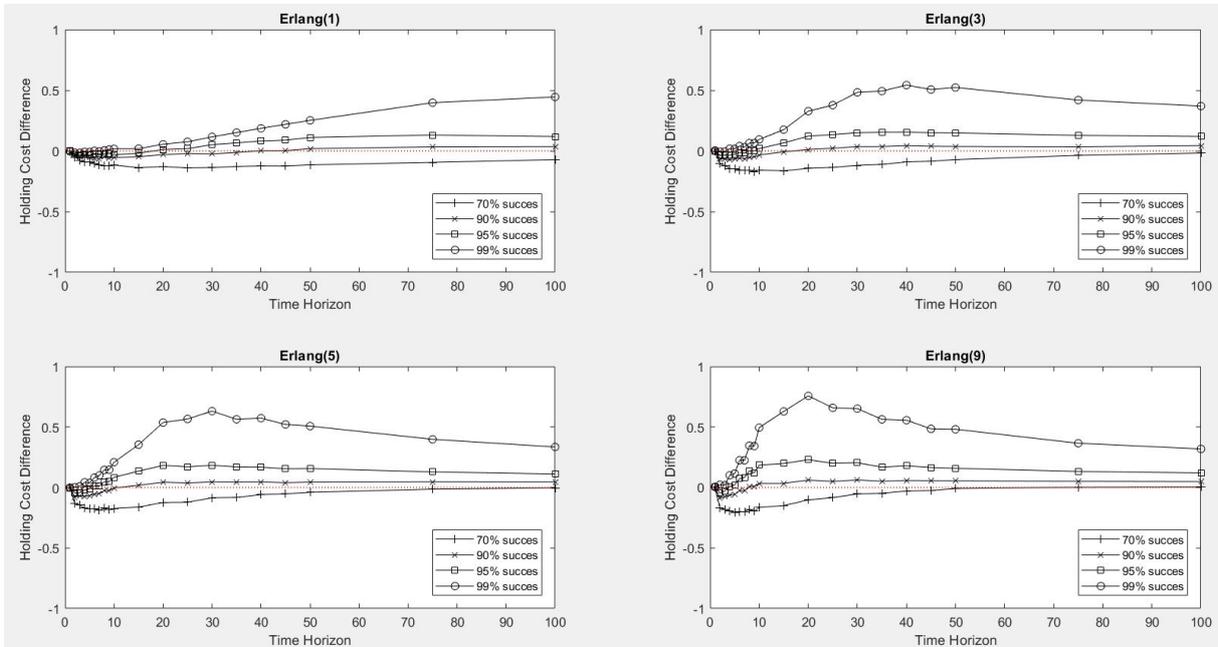
Holding costs required to meet a 99% target fill rate 70, 90, 95 and 90% of the time for the base-stock policy. The review horizon ranges from 1 to 100.

Figure 7: Holding costs required to have a specified chance of meeting a 95% target fill rate



Holding costs required to meet a 95% target fill rate 70, 90, 95 and 90% of the time for the non-stationary policy with $c = 0.5$. The review horizon ranges from 1 to 100.

Figure 8: Differences in holding costs between the dynamic ($c = 0.5$) and base-stock policy



Differences in holding costs required to meet a 95% target fill rate 70, 90, 95 and 90% of the time between the non-stationary policy with $c = 0.5$ and the base-stock policy.

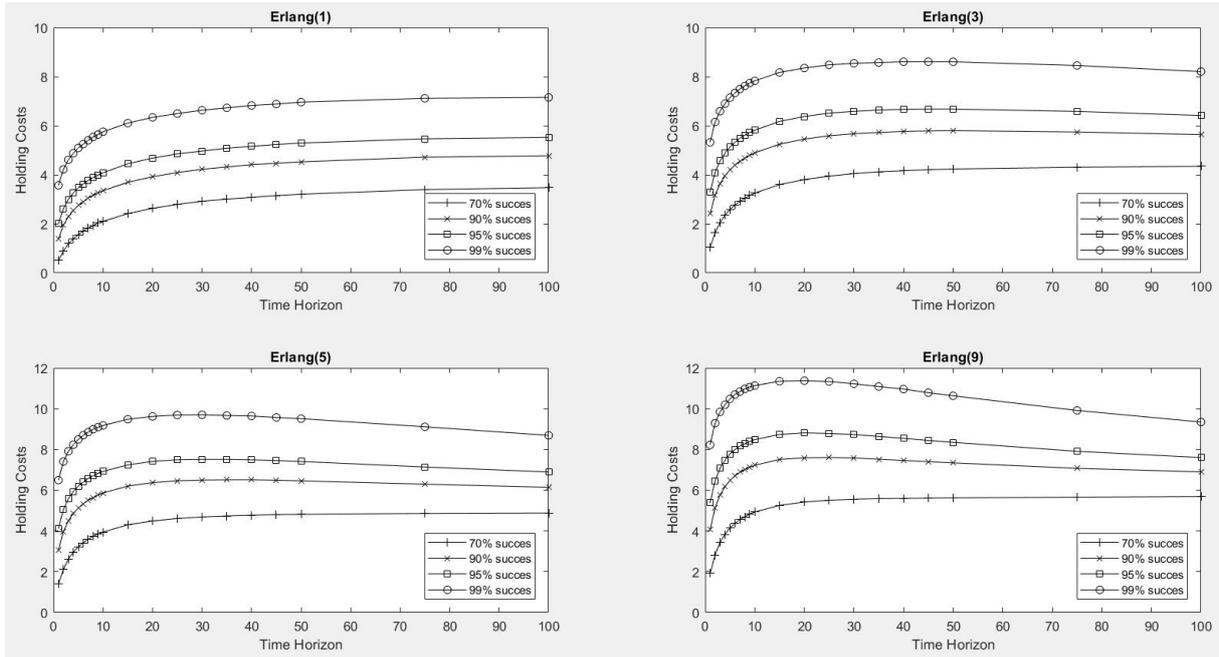
holding costs. This turns out not to be true, as in many cases, the non-stationary policy's initial order-up-to point is higher than that of the base-stock policy when the same success probability is observed. The reason here is that when s is set higher than the level needed to attain a long-run fill rate of θ , which is done when λ is large, the non-stationary policy may in fact set $s' < s$ if the fill rate at the end of period t' is unsatisfactory. In this case, as described in Section 5, s' will be set such that the long-run probability is $\delta_{t'}^*$, rather than to achieve $\delta_{t'}^*$ with probability λ . The result is a lower success probability than the base-stock policy with the same initial s . When a higher initial s is needed, holding costs will also almost always be higher for high success probabilities, since with a high initial s , the probability that $\delta^* > 1$ is extremely small, so there is little opportunity to cut down on inventory. Consequently, the base-stock policy outperforms the non-stationary policy in terms of total costs when both θ and λ are high.

Furthermore, note that the differences between the non-stationary and the base-stock policy become less significant as the review horizon becomes larger. This is no surprise, because when T approaches infinity, so does t' . The fill rate at both points converges to θ with probability 1, and the dynamic policy coincides with the static base-stock policy. Convergence of the two policies occurs faster with higher Erlang scale parameters, since this, like increasing T , results in lower demand variability.

The development of holding costs with the same success probabilities but with a target fill rate of 99% is shown for various Erlang distributions in Figure 9. Differences with the base-stock policy are in Figure 10. Differences between the two policies now only become clear for longer horizons. In fact, for the Erlang(1) distribution, they only seem to really differ for the 70% success probability. With 99% success, the base stock policy still does best for the other distributions. Convergence of the two policies in terms of holdings costs does not clearly show yet for $T = 100$, but this would be the case for longer review horizons.

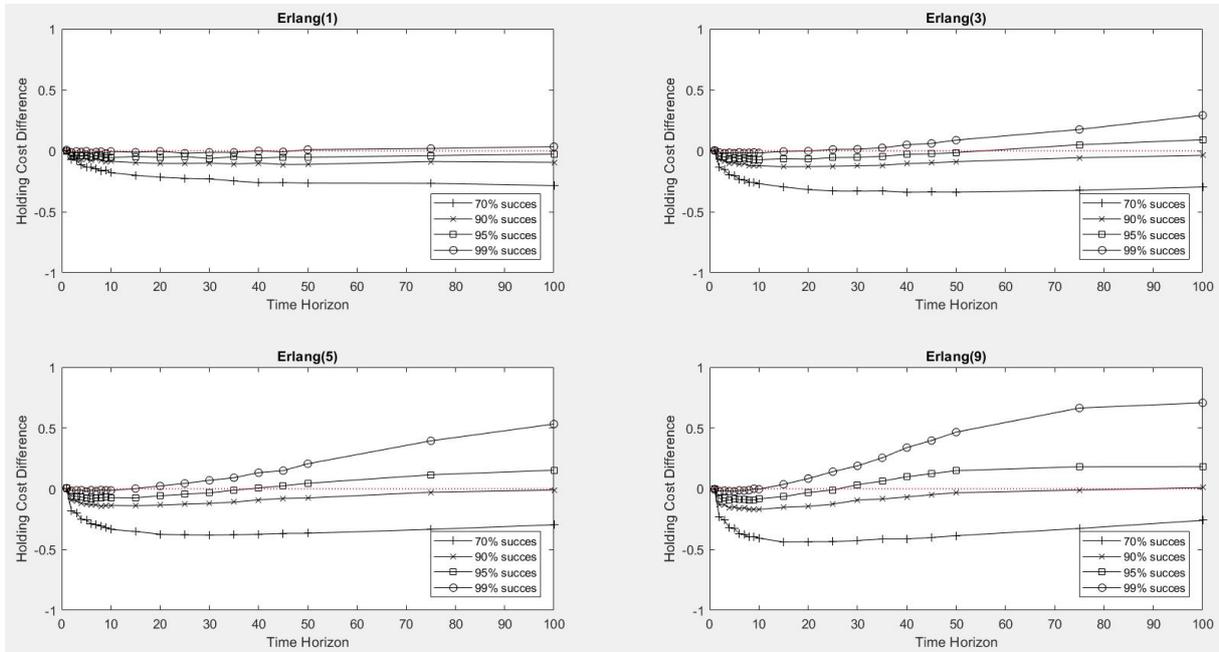
The most important lesson learned from all the previous figures in this section combined is that the non-stationary policy does not always outperform the base-stock policy. When a high target needs to be met with high probability, meaning the lump-sum penalty is severe, the base stock policy with s (far) above the long-run target level is clearly a better choice for review horizons of realistic length, such as $T = 12$ (monthly replenishment, yearly review) or $T = 30$ (daily replenishment, monthly review. As discussed earlier, SLA's with lump-sum penalties will most likely severely penalise fill rate violations to prevent the vendor from not replenishing at all. Therefore, under a realistic SLA with a high θ , the vendor is better off sticking to a base-stock policy, rather than opting for the non-stationary strategy proposed. However, with some minor adjustments, a non-stationary policy may actually show better performance than the base-stock problem. The main reason that the proposed strategy underperforms is because in case of an insufficient yet repairable interim fill rate, s' may be set lower than s , since it only has to achieve $\delta_{t'}^*$ in the long-run. This problem can be solved by not allowing downward adjustments unless $\delta_{t'}^* > 1$. With this slight alteration in the dynamic policy, it manages to outperform the base-stock policy, as can be seen in Figure 11 for $c = 0.5$ and $\theta = 0.95$. Note that the scale of the y-axis is different than in Figure 8 and 10. The differences are noticeable, but rather small. With the same starting s , the base-stock policy can once again only lead to a higher success probability in the case where $\delta_{t'}^* > 1$, but compliance to the

Figure 9: Holding costs required to have a specified chance of meeting a 99% target fill rate



Holding costs required to meet a 99% target fill rate 70, 90, 95 and 90% of the time for the non-stationary policy with $c = 0.5$. The review horizon ranges from 1 to 100.

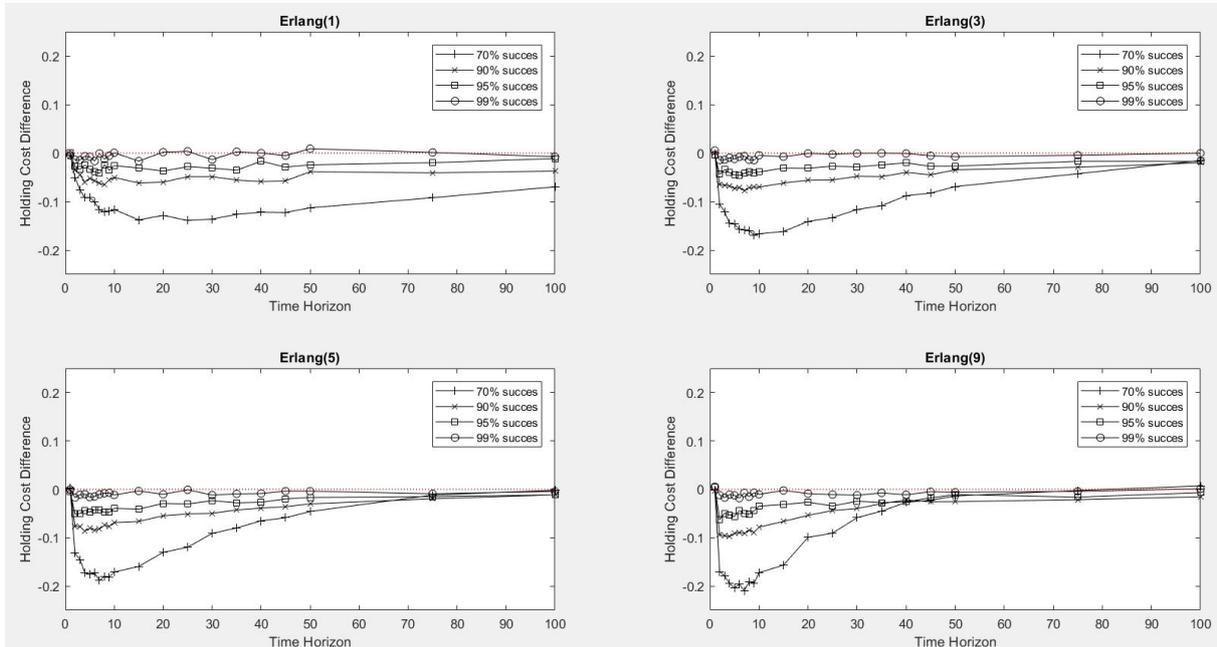
Figure 10: Differences in holding costs between the dynamic ($c = 0.5$) and base-stock policy



Differences in holding costs required to meet a 99% target fill rate 70, 90, 95 and 90% of the time between the non-stationary policy with $c = 0.5$ and the base-stock policy.

fill rate target is still possible by the base-stock. A numerical example of this scenario can be found in the Appendix.

Figure 11: Differences in holding costs between the adjusted dynamic policy ($c = 0.5$) and the base-stock policy



Differences in holding costs required to meet a 95% target fill rate 70, 90, 95 and 90% of the time between the adjusted non-stationary policy with $c = 0.5$ and the base-stock policy.

7 Conclusion

This thesis has studied the performance of a simple, non-stationary strategy in an SLA that specifies a review horizon of finite length, at the end of which an agreed upon target fill rate needs to be met by the good-supplying party. Without lead time, negligible reordering costs, a single customer and demand following an Erlang distribution, the model is similar to that used by Thomas (2005) in his influential paper studying the distribution of the finite horizon item fill rate for a static base-stock policy. Using Monte Carlo simulation, the results of Thomas (2005) are successfully replicated.

Research in this thesis was motivated by the prediction that in a more realistic setting, a manager in the supplying company will probably evaluate the state of affairs well before the contractual point of review. Depending on the fill rate at the time of the interim check, the manager may wish to adjust the order-up-to point downward, upward, or not at all. In most scenarios, it is found that the best moment for this changeover period is roughly halfway through the review horizon.

Comparing this policy to a base-stock model that sets the order-up-to level such the long-run fill rate is equal to its target, the non-stationary policy clearly shows superior performance in total costs for a variety of horizon lengths and Erlang distributions. Because of the adjustment possibilities, one of which

is to stop supplying goods to the customer altogether, the distribution of the fill rate under the non-stationary policy has a much higher variance than the base-stock policy. Also, even though the mean fill rate is much lower, the estimated probability to meet the target fill rate is higher with the non-stationary policy. This is the main explanation for its outperformance. When engaged in an SLA with a lump-sum penalty, it is the success probability that determines the expected penalty payment, not the expected fill rate.

The stocking policy studied is not always superior in terms of expected total costs. When a high probability of success is required to achieve a high target fill rate, the static base-stock policy actually shows lower holding costs for the same success probability. This suggests that some strategic changes to the non-stationary policy are needed in order for it to compete with the base-stock policy if the order-up-to point is not based on the infinite-horizon fill rate. Not adjusting s downwards unless $\delta_{t'}^* > 1$ is an example of such a strategic change.

In spite of their differences, the non-stationary and the base-stock stocking policy still have some distributional properties in common. For example, holding costs required to achieve a fill rate target with a specified probability initially increase with review horizon length, but eventually decrease with it. Also, it still holds for the non-stationary policy that the standard deviation of the achieved item fill rate decreases with longer review horizons and higher Erlang scale parameters, since both decrease the variability of demand.

Although the magnitude depends on factors such as unit holding cost, penalty severity and the distribution of demand, it can be concluded that cost-savings are possible in quite some situations when the proposed non-stationary stocking policy is used, rather than a static base stock policy. To achieve cost savings as high as those reported by Liang and Atkins (2013), however, a more complicated strategy may be required in most scenarios.

8 Discussion

This thesis has considered a particular non-stationary policy with only one changeover period. The policy is more dynamic than a base-stock policy, but still quite static compared to a policy that may change the order-up-to period after each demand realisation. If a fast implementation of the nested simulation approach can be found, it would be interesting to analyse the performance of this strategy. Otherwise, the policy examined may be altered to include multiple changeover periods. In addition, the strategy of adjusting the order-up-to point may be changed to one that takes into account the length of the review horizon, rather than setting s such that the infinite horizon fill rate is on target. Finally, methods may be researched to increase accuracy of the approximation in (7), so that $\delta_{t'}^*$ is computed more precisely.

Similar to Thomas, this thesis explores a stocking model with zero lead time, no ordering costs, a single customer and i.i.d. demand from an Erlang distribution. This is the most basic model, and relaxing any of the first three assumptions may lead to significant differences in outcomes. Different demand distributions could be studied, although varying the Erlang parameters might be similar to this potential alteration. More interesting may therefore be to look at models with positive constant or

stochastic lead times, fixed or variable ordering costs, and non-stationary demand. These relaxations may bring with them alternative approaches to finding appropriate non-stationary stocking policies.

An important feature of the type of SLA studied is that violations of the target fill-rate are penalised with a lump-sum penalty. This is the foundation of the non-stationary policy studied, which stops replenishing when the target fill rate can no longer be met. If however, as proposed by Liang and Atkins (2013), the penalty is proportional to the underperformance of the vendor, it may be more efficient to keep supplying to the customer. Non-stationary policies will need to take this into account.

Finally, one might consider a different measure of evaluating the vendor in the SLA. This thesis focused exclusively on the item fill rate, but opting for the ready-rate as a performance measure is equally justifiable. It may also be interesting to study the distribution of the order-fill rate under a non-stationary stocking policy. This measure was suggested by Larsen and Thorstenson (2008) and is closely related to the item fill-rate.

References

- Banerjee, A., & Paul, A. (2005). Average fill rate and horizon length. *Operations Research Letters*, *33*(5), 525–530.
- Bollapragada, S., & Morton, T. E. (1999). Myopic heuristics for the random yield problem. *Operations Research*, *47*(5), 713–722.
- Chen, C.-M., & Thomas, D. J. (2018). Inventory allocation in the presence of service-level agreements. *Production and Operations Management*, *27*(3), 553–577.
- Chen, F. Y., & Krass, D. (2001). Inventory models with minimal service level constraints. *European journal of operational research*, *134*(1), 120–140.
- Chen, J., Lin, D. K., & Thomas, D. J. (2003). On the single item fill rate for a finite horizon. *Operations Research Letters*, *31*(2), 119–123.
- Graves, S. C. (1999). A single-item inventory model for a nonstationary demand process. *Manufacturing & Service Operations Management*, *1*(1), 50–61.
- Katok, E., Thomas, D., & Davis, A. (2008). Inventory service-level agreements as coordination mechanisms: The effect of review periods. *Manufacturing & Service Operations Management*, *10*(4), 609–624.
- Larsen, C., & Thorstenson, A. (2008). A comparison between the order and the volume fill rate for a base-stock inventory control system under a compound renewal demand process. *Journal of the Operational Research Society*, *59*(6), 798–804.
- Liang, L., & Atkins, D. (2013). Designing service level agreements for inventory management. *Production and Operations Management*, *22*(5), 1103–1117.
- Matsumoto, M., & Nishimura, T. (1998). Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, *8*(1), 3–30.
- OblicoreInc. (2007). 2007 service-level management survey: Results, trends and analysis.

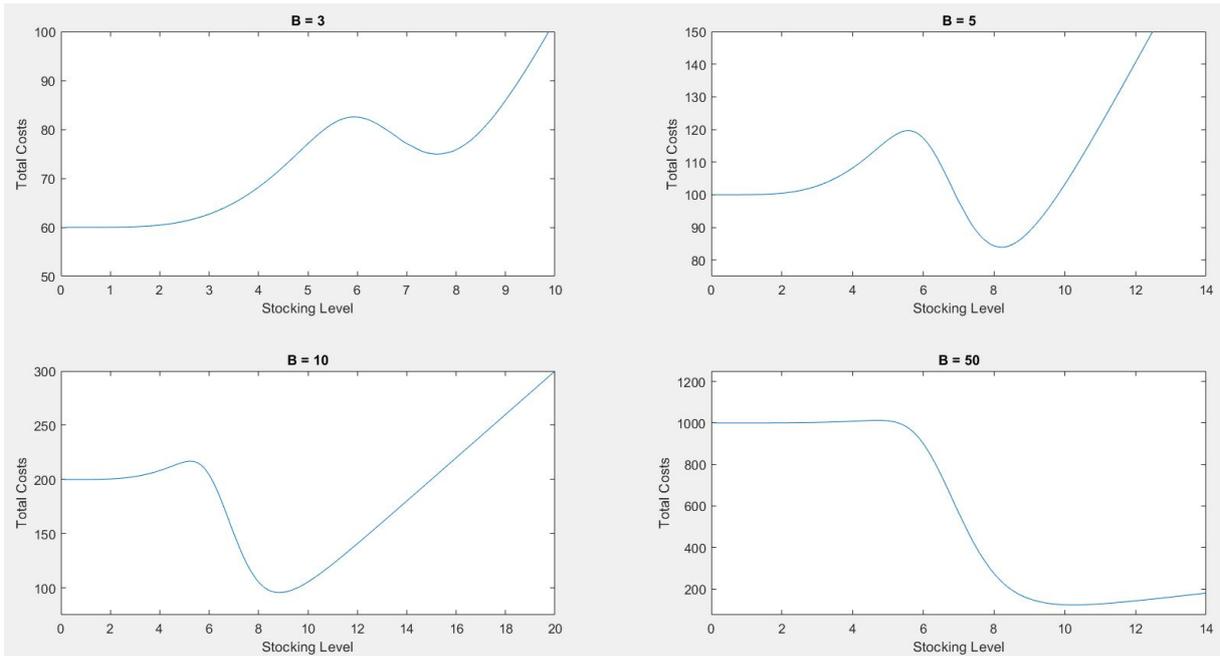
- Schneider, H. (1981). Effect of service-levels on order-points or order-levels in inventory models. *The International Journal of Production Research*, 19(6), 615–631.
- Silver, E. A., & Bischak, D. P. (2011). The exact fill rate in a periodic review base stock system under normally distributed demand. *Omega*, 39(3), 346–349.
- Silver, E. A., Pyke, D. F., & Thomas, D. J. (2016). *Inventory and production management in supply chains*. CRC Press.
- Tan, Y. R., Paul, A. A., Deng, Q., & Wei, L. (2017). Mitigating inventory overstocking: Optimal order-up-to level to achieve a target fill rate over a finite horizon. *Production and Operations Management*, 26(11), 1971–1988.
- Thomas, D. J. (2005). Measuring item fill-rate performance in a finite horizon. *Manufacturing & Service Operations Management*, 7(1), 74–80.
- Tunc, H., Kilic, O. A., Tarim, S. A., & Eksioglu, B. (2011). The cost of using stationary inventory policies when demand is non-stationary. *Omega*, 39(4), 410–415.
- Zipkin, P. H. (2000). *Foundations of inventory management*.

9 Appendix

9.1 Non-convexity of the objective function

In the first part of Section 5, it is briefly mentioned that the objective function in (5) is non-convex. Figure 12 shows plots of the function for a scenario in which $T = 20$, $h = 1$ and $X_t \sim \text{Erlang}(5,1)$, with various choices of B , all per-period penalties. No previous demand realisations were observed yet, so the graph shows the function $C(s_1)$, which could be used to calculate the optimal starting value for a dynamic policy, but also for the base-stock policy.

Figure 12: Expected total costs for different choices of B

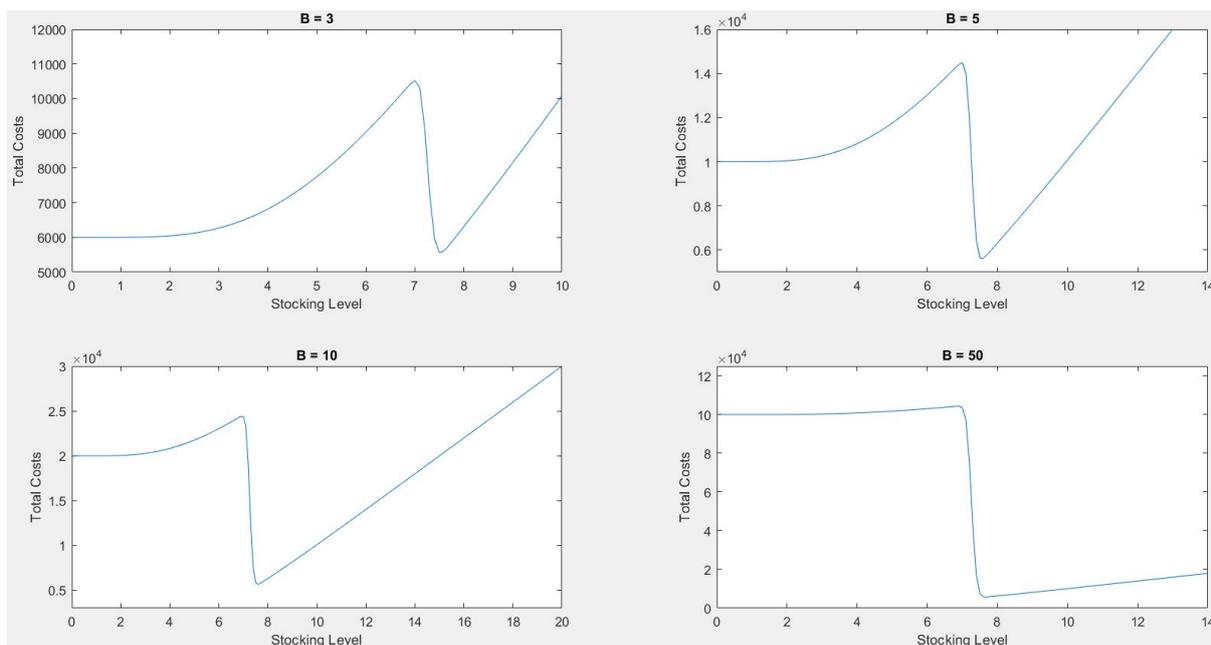


Expected total costs with $T = 20$ for different per-period penalties B , plotted against the the order-up-to point at the beginning of the review horizon, so with no prior demand realisations

All functions are clearly not convex. Consider the graph of $B = 5$. Starting from $s = 0$, total costs initially do not change with s , as both holding costs and success probability are stuck at zero. For $s = [2, 5.5]$, total costs are increasing. Although holding costs increase in this interval, the success probability remains very low. Total costs are decreasing in the interval $[5.5, 8.25]$, where the increase in holding costs is more than offset by the increasing success probability. When $s > 8.25$, total costs increase once again, because increasing s has little to no influence anymore on the probability to meet the target fill rate.

An important note is that the relation between h and B strongly affects the shape of the objective function, affecting both shape and scale. Nevertheless, non-convexity is a property that is likely to remain in many scenarios. Also, as the length of the review horizon increases, the turnover points from increasing to decreasing and vice versa become more sharp. This is shown in Figure 13, which contains the same information as Figure 12, but for $T = 2000$. As T approaches infinity, the fill rate at review is no longer random and the success probability is either 0 or 1, depending on the choice of the order up-to level.

Figure 13: Expected total costs with $T = 2000$



Expected total costs with $T = 2000$ for different per-period penalties B , plotted against the the order-up-to point at the beginning of the review horizon, so with no prior demand realisations

9.2 Accuracy of the fill-rate approximation

In Section 5.1, a weighted average of the fill-rate in two adjacent periods is used to approximate the fill-rate of these two periods combined, as given in (7). To verify the accuracy of the approximation, the difference between the approximated and the actual fill rate is computed. With a review length T , changeover period t' and order up-to point s , the approximation error is computed as $\alpha'_T(s) - \alpha_T(s)$. Table 6 shows the approximation errors for different inputs of T and t' when demand follows an Erlang(5,1) distribution and $s = 7.2633$, such that the long-run fill rate is equal to 95%.

As the table shows, the approximation on average overestimates the actual fill rate, with bias topping at almost 0.6 percentage points for the shortest review horizon. Bias decreases as t' decreases, but much more notably as T increases. Remembering the theorem by Banerjee and Paul (2005) that the expected fill rate is monotonically decreasing in T , bias in the approximation is not a surprising result. The approximation splits the review period in two sub-periods of length t' and $T - t'$, and calculates the fill rate over the full period by means of a weighted average of the fill rates in the sub-periods. Let $\alpha_{t'}$ and $\delta_{t'}$ be the fill rates of the first and second sub-period, respectively. With $0 < t' < T$, the sub-periods are both strictly smaller in length than T , which means that $\mathbb{E}[\alpha_T] < \mathbb{E}[\alpha_{t'}]$ and $\mathbb{E}[\alpha_T] < \mathbb{E}[\delta_{t'}]$. From this, it follows that, $\mathbb{E}[\alpha_T] < \mathbb{E}[\alpha'_{t'}]$ and so the approximation is positively biased.

Despite its positive bias, the approximation should work reasonably well in the non-stationary stocking policy discussed, since especially for longer horizon lengths, bias is rather small in magnitude. Even for smaller values of T , the overestimation is not expected to break down the policy, although caution is advised here.

Table 6: Errors in the approximation of the fill rate

t'/T	5	10	20	30	40	50	100	200	1000
1	0.4781	0.2475	0.1249	0.0834	0.0626	0.0501	0.0250	0.0125	0.0025
2	0.5902	0.3184	0.1616	0.1081	0.0811	0.0648	0.0325	0.0162	0.0033
3	0.5911	0.3473	0.1777	0.1191	0.0894	0.0714	0.0358	0.0179	0.0036
5		0.3634	0.1915	0.1287	0.0967	0.0774	0.0387	0.0194	0.0039
15			0.1915	0.1372	0.1039	0.0835	0.0419	0.0210	0.0042
25				0.1285	0.1040	0.0843	0.0426	0.0213	0.0043
45						0.0774	0.0429	0.0215	0.0043
95							0.0387	0.0217	0.0044
150								0.0216	0.0044

Difference between in the approximated and actual fill rate in percentage points for different values of T and t' . Table entries for which $t' \geq T$ are left blank.

9.3 Numerical example

Section 5.1 states that there is a positive probability that a base-stock policy that sets s such that the long-run fill rate is θ results in a satisfactory fill rate, while the non-stationary policy does not. In this scenario, δ_t^* is greater than 1, but the stationary base-stock policy could actually meet target θ with a probability strictly larger than zero. To prove that this is actually true, a numerical example is given.

Consider a scenario with a review horizon $T = 10$, over which a fill rate of 95% has to be achieved and where demand follows an Erlang(5,1) distribution. The base-stock policy then sets s equal to 7.2633. This is also the initial s for the non-stationary policy, but after 5 periods, s is changed to s' , as is described in Section 5. Demand realisations in the first five periods are 3,3,3,3 and 10, resulting in a fill rate of

$$\alpha_{t'}(s) = \frac{3 + 3 + 3 + 3 + 7.2633}{3 + 3 + 3 + 3 + 10} \approx 0.8756. \quad (9)$$

Because $\delta_t^* \approx 1.0244 > 1$, the non-stationary will set s' to zero, deeming it impossible to reach the target of 0.95. However, if demand in periods six through ten are all equal to 7, the fill rate at review for the base-stock policy is equal to

$$\alpha_T(s) = \frac{3 + 3 + 3 + 3 + 7.2633 + 7 + 7 + 7 + 7 + 7}{3 + 3 + 3 + 3 + 10 + 7 + 7 + 7 + 7 + 7} \approx 0.9520 \quad (10)$$

This means that the base-stock policy meets the target, whereas the non-stationary policy with $c = 0.5$ does not, because $s' = 0$.

Note that the demand realisations are really quite extreme. Therefore, these outcomes are expected to occur very sparsely. In addition, the probability of occurrence decreases with higher T and with lower θ . If one wishes to eliminate this phenomenon altogether, s' should only be set to zero when, in addition to $\delta_t^* > 1$, it also holds that

$$\frac{\sum_{t=1}^{t'} Y_t + s(T - t')}{\sum_{t=1}^{t'} X_t + s(T - t')} > \theta \quad (11)$$

In this case, it is truly impossible for the base-stock policy to achieve the fill rate target θ when s is set such that the long-run fill rate is θ .