ERASMUS UNIVERSITY ROTTERDAM
Erasmus School of Economics

Bachelor Thesis BSc2 Economics/Econometrics

Optimizing Inventory Policy in Anticipation of Obsolescence: A Simulation of Step-Wise Advance Stock Reductions

---

Name Student: Frederik Bartl
Student ID number: 430637

Supervisor: W. Ma
Second assessor: Dr. A.S. Eruguz

July 7^{th}, 2019

Abstract

Inventories can experience high obsolescence costs if there is no policy in place to avoid being left with expensive obsolete stock. Especially for slow moving spare parts inventories this can be an issue. In this paper I investigate the effects of different inventory policies that seek to minimize costs in the face of obsolescence. I replicate the paper of Pinçe and Dekker (2011) and re-affirm their findings, but caution against a bias in their model. The optimal time to change start reducing stock lies even earlier than they predicted. Ultimately, a gradual inventory policy change yields the best results as it allows for backorder costs and obsolescence costs to be reduced.

Contents

1 Introduction **1**

2 Literature Review **3**

3 Methodology **7**

3.1 General Simulation Structure 7

3.2 Replication Model 9

3.3 Extension Models 11

4 Numerical Study **12**

4.1 Replication 12

4.2 Extensions 14

4.2.1 Biases of X^* in the Replication Model 15

4.2.2 General Behaviour of Extension Model 17

4.2.3 Gradual Policy Change 19

5 Conclusion **22**

1 Introduction

Obsolescence is a well known problem for inventory management. Many factors play a role in defining the distinct instance of an inventory management problem with obsolescence and there is a lot of existing research on this issue. Since there are many different scenarios of inventories affected by obsolescence, important insights can yet be learned about dealing with obsolescence. One such scenario is that of anticipated obsolescence in a spare parts inventory, where a foreseeable drop in demand necessitates a change in inventory policy.

If an inventory policy is not changed to suit the anticipated obsolescence, the inventory could carry excess stock and incur unnecessarily high holding costs. This could continue for a long time, since the lowered demand means it will take more time until the excess stock could be taken away by demand. In the case of complete obsolescence, meaning demand will cease entirely, the excess stock will have to be disposed or possibly relocated. Costs for disposal or relocation of obsolete stock are the obsolescence cost.

Realizing these costs associated with failing to respond to anticipated obsolescence, one might be tempted to react quickly and sell out inventories before obsolescence costs are incurred. However, the other extreme response, selling off inventories entirely, can be damaging to a company if the inventory is sold too early. Customers could be forced to wait for backorders and due to this the reputation of the company could suffer and sales or even customers could be lost.

So, we could expect that considerable costs could come from changing the inventory policy too early, or too late in order to address anticipated obsolescence. This leads to the question of how to determine the appropriate timing and appropriate changes in inventory policy to optimally cope with anticipated obsolescence.

In particular we are interested in continuous review policy inventories. The exemplary case is that of slow moving, expensive spare part inventories that use one-for-one replenishment $(S-1, S)$ as their inventory policy. With inventories like these the time and degree of obsolescence can often be anticipated or at least reasonably well estimated (Pinçe & Dekker, 2011). That is especially the case when installed base data is used to enhance forecast accuracy (Dekker, Pinçe, Zuidwijk, & Jalil, 2013; Jalil, Zuidwijk, Fleischmann, & Van Nunen, 2011).

In this paper, the work of Pinçe and Dekker (2011) is replicated to find the optimal timing for policy changes in case of continuous review $(S-1, S)$ inventory policies. As Pinçe and Dekker (2011) only investigated the timing of a simple policy change from (S_0-1, S_0) to (S_1-1, S_1) for a demand change from a pre-obsolescence demand arrival rate λ_0 to a post-obsolescence demand arrival rate λ_1, further policy options that currently lack analysis should be considered.

Addressing this, I will expand on their work, by means of simulation, to investigate the cost savings that could be achieved from gradual transition policies. In this case of discrete inventory stock, this results in policies of step-wise inventory reductions. A refined investigation of more flexible responses to anticipated obsolescence should help companies holding stock that may face obsolescence to reduce their inventory costs.

Therefore the main research question of this paper is:

How should a continuous review inventory policy for one-for-one replenished stock be designed to optimally save costs in a transition period with anticipated obsolescence?

In order to answer this, I address the following subquestions:

1) To what extent can advance inventory policy changes reduce costs?

2) How is the optimal policy influenced by lead time, backorder penalty, disposal cost, base stock levels and demand rates?

3) In what way is the optimal policy as found by Pinçe and Dekker (2011) biased and how can an unbiased solution be found?

4) How do gradual transition policies perform compared to (S_0-1, S_0) *to* (S_1-1, S_1) *transition policies?*

This research will help any company running inventories of slow moving stock that face predictable obsolescence to minimize their inventory costs in the transition phase. Thus, there are meaningful practical applications for the findings of this research. Scientifically, the research is relevant as it provides detailed insights into inventory modelling with dynamic programming and gradual transition approaches to anticipated obsolescence.

I find that I can replicate the findings of Pinçe and Dekker (2011), but I detect a bias in their methodology. The optimization of the timing is indeed a trade-off between obsolescence costs and backordering costs and advance policy changes can lead to significant cost savings. With an altered model I find that a gradual policy change can significantly improve upon the analytical approximation by Pinçe and Dekker (2011) and is the optimal solution to anticipated obsolescence.

Following this introduction, I investigate the existing literature on inventory policies in anticipation of obsolescence. In this section, the economic theory regarding predictability of obsolescence and inventory policies subjected to it is explored in order to properly place this paper into context with the current state of academic knowledge. Thereafter, I explain the methodology used throughout this paper. This part is split between the model used for the purpose of replication and the models featuring the extensions of this paper. Then, I describe the results of my numerical study conducted with these models. I again split the analysis of the replicated part from the analysis of my extensions. Regarding my extensions, I further separate the analysis into parts on the replicated model's bias, the general influences of parameters, and the effects of the gradual approach. Lastly, I discuss the results obtained in this paper and their implications. Limitations and possible extensions of this paper's investigations are also given in the final section.

2 Literature Review

The most relevant piece of literature for this paper's topic is the article "An inventory model for slow moving items subject to obsolescence" by Pinçe and Dekker (2011) that is mentioned numerous times throughout this paper. As my paper is primarily an extension of Pinçe and Dekker (2011), both the theoretical foundations, found in their heuristic approximation, and the idea for the replication framework in the simulation approach stem from their work. The authors made various findings that my paper aims to reproduce and test in alternative settings.

Most fundamentally, the timing of advance inventory policy changes is found to be a trade-off between backorder penalties and obsolescence costs. By investigating the general shape of total cost functions in my replication model, as well as in my extension models, I can evaluate if my methods lead to the same conclusion.

According to Pinçe and Dekker (2011), not changing policy in advance leads to significantly larger costs compared to an optimal advance policy change. This too, is part of my analysis as mentioned in Subquestion 1 of my research question. The analytical approximation of the optimal advance inventory policy change yields near optimal costs (Pinçe & Dekker, 2011), however further gains could possibly be achieved if more than a single change of base stock level is considered.

These findings are based on optimizing over the 'transient period', which is defined as the period starting with the change of policy and ending when the targeted post-obsolescence inventory position is reached and all orders that were placed before obsolescence occurred have arrived. It appears relevant to examine if the near optimal policies are biased, due to the focus on the transient period costs.

The replication part of my methodology is built on the work of Pinçe and Dekker (2011). The framework of optimization over the transient period is applied for the replication, but

I change this for the investigation of my extensions. Based on the finding that total cost is unimodal (Pinçe & Dekker, 2011), I attempt to use the golden section search to optimize the timing of the advance change. My choice of second degree polynomial triangulation as a means to use few simulations to model a convex approximation of the total cost function is inspired by Pinçe and Dekker (2011) using Response Surface Optimization as outlined in Montgomery and Myers (1995).

The work in my paper relies on the assumption that the time of obsolescence T and the drop of demand from λ_0 to λ_1 can be predicted with reasonable accuracy. It is therefore interesting to see if the current academic literature can provide justifications for this assumption.

A method to estimate the obsolescence risk using demand data is presented by Van Jaarsveld and Dekker (2011). An analysis of obsolescence patterns of other items in the past is used to predict the probability and extent of obsolescence for items that have not yet experienced obsolescence. The estimates found this way inform an inventory control policy for spare parts inventories. Their findings are applicable especially to products with long life cycles such as the slow moving spare parts that are the typical example for the inventory policy investigated in my paper.

As the assumption of accurately forecast T and λ_1 is also essential for Pinçe and Dekker (2011), they too cite an article that investigated this question. This article that deals with the quality of predictions made about spare parts inventories is by Jalil et al. (2011). They investigate how installed base data can be used to do forecasts and supply chain planning. This means that information on the installed base of machines which require the spare part in question is used to increase forecast accuracy.

The authors find that installed base data can help to save costs especially in cases of low demand rates and low installed base size. This is exactly the case in slow-moving spare parts inventories that use continuous review one-for-one replenishment policies. The findings of Jalil et al. (2011) imply that with the use of installed base data a company that manages an inventory of the type that is investigated in my paper can make reasonably accurate predictions on the size and geographical distribution of future demand.

The authors of the replicated paper also investigate the use of installed base data for demand forecasting for spare parts inventories in a paper of their own (Dekker et al., 2013). In Dekker et al. (2013) the authors find that including installed base data significantly improves forecast accuracy when compared to time series approaches and black-box forecasting. Especially for the kind of inventory that is the subject of my paper's investigations, namely slow moving, low demand inventories, the forecast accuracy can be improved significantly by utilizing installed base data (Dekker et al., 2013).

The authors reach the conclusion that the use of installed base data would not harm forecast accuracy even when homogeneously distributed errors exist in the data. Only heterogeneously distributed errors in the data would worsen the accuracy of forecasts.

These three papers (Dekker et al., 2013; Jalil et al., 2011; Van Jaarsveld & Dekker, 2011) lend credence to the assumption of a known point in time at which demand will decrease to a known level. They show that obsolescence can indeed be estimated and such estimates can help to find more efficient inventory policies.

The authors of the replicated paper (Pinçe & Dekker, 2011) also published a later paper that further develops their original findings and is especially relevant to my extension models. In Pinçe, Frenk, and Dekker (2015) the authors analyze a finite-horizon, non-stationary one-for-one policy with a single base stock adjustment.

The authors optimize the policy by using three variables to define the chosen policy. These variables are the base stock before and after obsolescence and the time of the switch from one to the other. In the notation from Pinçe and Dekker (2011), which is used throughout this paper, that corresponds to a simultaneous optimization of S_0, S_1 and X. This 'single-adjustment policy', as they call it, is similar to my chosen extension framework in the sense that costs are optimized for a fixed finite calculation period rather than for the transient period only.

It should be noted that there are a number of important differences between their framework and mine. Perhaps the most significant difference is that the prediction of the demand drop is not generally known far in advance, but rather T is an important parameter in their model that defines whether the instance deals with short-term, mid-term or long-term obsolescence. In my framework, the value of T is irrelevant, since I always investigate the costs over a long enough calculation period that starts and ends in the pre- and post-obsolescence steady state respectively. The other major difference is the precise definition of the objective.

The objective function that is to be minimized is still total cost, however they use the discounted cost rather than the undiscounted total cost, because in order to optimize S_0 and S_1 the discounted cost allows more accurate representations of costs in the transient period (Haneveld & Teunter, 1998; Pinçe & Dekker, 2011). This is less relevant in my model, as my instances are mostly taken from Pinçe and Dekker (2011), so S_0 and S_1 are usually given and are not altered for optimization. For instances where S_0 and S_1 need to be determined, I use steady state simulations rather than the transient period costs to find appropriate values, so my use of undiscounted total costs remains justified.

In addition to the single-adjustment policy, Pinçe et al. (2015) also analyze the 'fixed policy' where the base stock level is kept fixed but obsolescence is taken into account and the 'single-adjustment-at-T policy' where the policy is changed at the time of obsolescence T. Finally, they also investigate the 'multiple-adjustment policy' which allows multiple changes to the base stock.

For the multiple-adjustment policy they use dynamic programming, but otherwise their investigation is done by transient analysis of sample paths and numerical inversion of generating functions (Pinçe et al., 2015). They mention that their approach is uncommon, but more efficient than the usual dynamic programming approach, which I also choose. Since I do not struggle with long runtimes in my dynamic programming approach, I do not necessarily require a more efficient methodology. However, increased efficiency could always be used to improve the accuracy of any paper's findings.

The authors find that the multiple-adjustment policy is generally optimal, but often yields only marginally larger cost savings than the single-adjustment policy. For cases of short-term obsolescence they find that even the fixed policy is close to the optimal policy, whereas for mid-term obsolescence the single-adjustment-at-T policy performs quite well for high demand rates with low demand drops.

When reviewing the applicability of the methodology chosen in Pinçe et al. (2015), I find that the status quo of the inventory is not appropriately taken into consideration. The advantage of the methodology in Pinçe and Dekker (2011) is that a pre-existing inventory position S_0 can be taken as an input parameter for the following optimization. In this sense we assume that the inventory has been in a steady state with a current inventory policy of $(S_0 - 1, S_0)$, while in Pinçe et al. (2015) the optimization works with an inventory that can start with any S_0 and optimizes it to S_0^*.

Fundamentally, the purpose of this paper is to find optimal ways to handle transitions from a current policy with a given S_0 to a future with changing demand conditions. This task can be achieved more fittingly with a methodology that enforces a status quo inventory policy, around which we have to optimize.

3 Methodology

The problem to be investigated is the optimization of the time to introduce an advance policy change, when it is known that full or partial obsolescence will occur in the future at time T. The demand is predicted to drop from λ_0 to λ_1 and the inventory position should be lowered accordingly from S_0 to S_1. When lowering the target for the inventory position in advance of T at time $T-X$, the demand that arrives between $T-X$ and T will lower the inventory position by up to $N = S_0 - S_1$ items. This results from letting the incoming demand take away up to N items from the inventory stock without placing orders for replenishment.

Using the naming convention established by Pinçe and Dekker (2011), we say the 'transient period' begins at $T-X$ and ends when all N items are removed and no orders from before T are still outstanding. The transient period is divided into the 'removal phase' and the 'regular operation phase'. Here the removal phase refers to the phase between $T-X$ and the removal of the N^{th} inventory item. The regular operation phase starts with the completion of the removal phase and ends when all orders placed before T have arrived in the inventory. The second phase is not entered if the stock removal phase concludes while no orders placed before T are still outstanding.

For this paper all results will be achieved through simulation. While the theoretical approximation by Pinçe and Dekker (2011) serves as a good comparison when replicating their instances, my expansions will not include theoretical approximations. Therefore, the simulation methods used are the primary focus of this section.

I also explain the notation used throughout the paper, which is mostly taken over from Pinçe and Dekker (2011) for ease of juxtaposition.

3.1 General Simulation Structure

The framework is that of an M/G/c queueing system, using a non-homogeneous Poisson Arrival process for demand arrivals and the deterministic lead times as 'service times'. I work with a template for Discrete Event Simulation designed by Milovanovic and Bouman (2018). The inventory position (IP) is computed as the sum of three state variables that are always updated as events are processed. These are the stock on hand, the outstanding orders and the backlog, which forms a queue of customers waiting for goods to be delivered.

$$IP = \text{stock on hand} + \text{outstanding orders - backlog length}.$$

Reordered goods will be delivered to the inventory once the 'service time', in this case the lead time, is over. This is done by scheduling 'order arrival events', which increase stock on hand, decrease orders outstanding and possibly serve a customer from the backlog. Since backorders are allowed, the system is uncapacitated. A backorder penalty is incurred at the time a customer from the backlog is eventually served. This means that we assume there are no lost sales.

Policy changes in this simulation framework are executed by changing the target inventory position. By not placing replenishment orders when 'demand arrival events' are processed, the inventory position decreases until it equals the new inventory target (S_1).

This simulation framework can be used to assess the performance of a certain choice of policy and yields the simulated total cost. For the replication of Pinçe and Dekker (2011) a single input variable is to be optimized, namely the advance time of the policy change labelled X. As the obsolescence occurs at T, the policy change is done at $T - X$. Given that the total cost function $TC(X)$ is found to be unimodal (Pinçe & Dekker, 2011), we attempt to use golden section search to efficiently optimize X. Due to the inherent random noise in simulations however, we cannot guarantee that the total costs found by simulation $TC_S(X)$ will always be unimodal. Therefore, I use polynomial triangulation (second degree) that makes use of the roughly convex form of $TC_S(X)$ to efficiently approximate the global minimum by testing only three input values. This is combined with an exhaustive local search. Here, I use discretization to find the approximately optimal X_S^* among values rounded to the third decimal.

A warm up phase is used to ensure that the simulation on average starts from the steady state. This phase has a standard length of 3 time units. Additionally, a cool down phase is used after the 'obsolescence event'. This phase has to be long enough to allow for the transient phase to end by demand arrivals or order arrivals. For cases of partial obsolescence its length ranges from 10 to 80. In case of full obsolescence a long cool down phase is not required as the transient period always ends at or before $T + L$.

The simulation uses an 'ending event' to include continuously incurred costs that are not yet accounted for, like holding costs until conclusion of the period and backorder penalties for backorders that have not yet been served. Similarly, at the start of the calculation period, backorder costs as well as holding costs are included precisely from the start of the calculation period. If customers are on the backlog at the start of the calculation period, when they are eventually served, the backorder penalty will be multiplied only by the waiting time that overlaps with the calculation period.

As this simulation has a short runtime, I choose to use at least 5000 replications for each instance and choice of X to ensure that the results are representative. Note that Pinçe and Dekker (2011) used 5000 replications for all their analyses.

Some variables indicating the most relevant findings are the difference between the optimal time for the advance policy change as found by my simulation (X_S^*) and the corresponding value found through the analytical approximation (X^*) by Pinçe and Dekker (2011) which is defined as $\Delta_X(\%) = \frac{X^*-X_S^*}{X_S^*}$ and $\Delta_X = X^* - X_S^*$. In the case of the gradual policy change model the earliest time of policy change X_N^* is used instead of X_S^*.

In addition to this, the difference in cost computed by simulation with these two input values is of interest. We define it as in Pinçe and Dekker (2011) as $\Delta_C(\%) = \frac{TC_S(X^*)-TC_S^*(X_S^*)}{TC_S^*(X_S^*)}$ and $\Delta_C = TC_S(X^*) - TC_S^*(X_S^*)$, where $TC_S^*(X_S^*)$ is the optimal total cost found by simulation with X_S^*. Similarly, the cost increase of no advance policy change at all over the optimal advance policy change is defined as $\Delta_0(\%) = \frac{TC_S(0)-TC_S^*(X_S^*)}{TC_S^*(X_S^*)}$ and $\Delta_0 = TC_S(0) - TC_S^*(X_S^*)$.

Here it is important to note that the measures of difference using percentage are diluted by the choice of model, whereas the absolute difference remains a comparable measure of cost savings over different calculation periods and methodologies.

3.2 Replication Model

The calculation period strictly starts with the introduction of the new policy (S_1-1, S_1) at $T-X$ and ends with the demand arrival or order arrival event that marks the end of the transient period. This leads to vastly different time lengths of the calculation period for different choices of X. This could introduce a bias towards a lower X. Larger values of X normally lead to higher backordering costs, while lowering holding and obsolescence costs. In this framework however, a larger value of X also means the calculation period begins earlier. In full obsolescence cases, this bias is certain. In cases of partial obsolescence the effects are ambiguous, since an earlier policy change could avoid relatively longer removal phases still going on after obsolescence has occurred. Then it is unclear if the net length of the calculation period is lengthened or shortened by a larger X.

Regardless, the descriptions in the paper indicate this as the chosen methodology, although some details remain unspecified (Pinçe & Dekker, 2011).

The simulation framework used for the extensions significantly differs from the one used in the replication section. In opposition to this framework, the extension model ensures that a fixed finite window of time is used as calculation period in all replications and for all X. This will ensure that ultimately, X will be optimized with respect to total costs incurred by the company over the entire time horizon. Pinçe and Dekker (2011) optimize X only with respect to the total costs in the transient period. Crucially, the choice of X affects not only the costs during the transient period, but also directly marks the end of the first period and has an effect on the start of the third period.

This can be shown with a simple example, where $T = 5$ and the base stock is supposed to be decreased from $S_0 = 6$ to $S_1 = 2$. The inventory process is depicted in the period from 0 to 10. In Figure 1 we can see that $X = 2$ is chosen as a possible time of policy change, hence the inventory target is changed at $T - X = 3$. The development of the inventory position (IP) and the inventory level (IL) are shown. The calculation period according to the replication framework is the period from 3 to 8. Since no items are disposed and no backorders occur, only holding costs are contributing to the total costs in this case. The holding costs h per item per time are normalized to $h = 1$. For $X = 2$, the total costs are therefore $TC(2) = 4 * (0.5 + 0.5 + 1) + 3 * (0.5 + 1 + 1.5) = 17$.

Alternatively, if the policy were changed at $T - X = 1$ with $X = 4$, the last reorder arriving at 5.5 would not have been placed, leading to a calculation period from 1 to 6.5. This case is indicated with the alternative inventory position (IP_A) and level (IL_A) in Figure 1 and leads to total costs of $TC(4) = 19$.

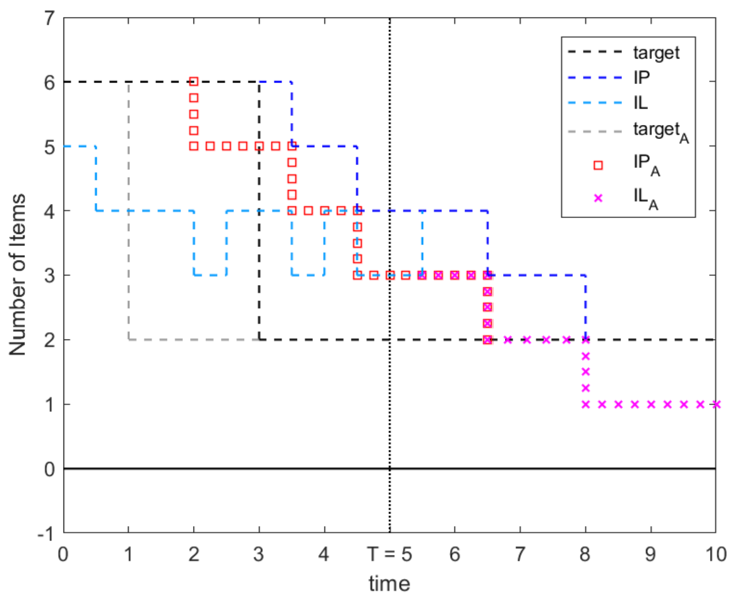

Figure 1: Inventory Process Example

While the costs calculated using the replication framework suggest that $TC(2)$ would be cheaper than $TC(4)$, it is obvious that the earlier change at $X = 4$ leads to lower overall costs in this particular case, as the intended target is reached while holding less excess stock for a shorter time.

If a fixed calculation period from 0 to 10 was chosen instead, the total costs would be $TC(2) = 33$ and $TC(4) = 28.5$. This example proves my intuition, that the fixation on the transient period can bias the optimization results towards lower values of X_S^* in some cases.

3.3 Extension Models

My model takes the effects of a varying calculation period into account by including all costs within a time frame that starts at a fixed time $T - t_0$, where $t_0 > X$ for all considered X, and ends at $T + t_{cooldown}$, where $t_{cooldown}$ is chosen such that the transient phase has ended before $T + t_{cooldown}$ for all replications and all considered values of X. The values for t_0 and $t_{cooldown}$ change when instance parameters are altered.

All inventory costs incurred outside of this chosen calculation period can be reasonably assumed to be unaffected by the choice of X, so this model finds the optimal policy for the entire transition from pre- to post-obsolescence.

An important effect of this change in framework is that the costs $TC_S(X)$ for any X are always going to be higher than in the more narrow time frame of the replication model. Hence, the cost savings measured in % are less in this framework. To ensure comparability across models, I therefore look at the absolute cost differences Δ_C and Δ_0, since those should not change with an expanded calculation period, given this period always encompasses the whole transient period for all considered X ($t_0 > X$ for all X and $t_{cooldown}$ large enough).

The gradual change of inventory policy is in fact a step-wise inventory reduction policy, due to the discrete nature of the inventory stock. Optimization of this policy is done through backwards induction. Given S_0 and $N = S_0 - S_1$, we start with the transition from $S_0 - N + 1$ to $S_0 - N$. We store the optimal choice for this last transition X_1^* and simulate a new scenario, optimizing choices for the second to last transition (X_2) from $S_0 - N + 2$ to $S_0 - N + 1$. This pattern is continued until we have found $X_N^*, X_{N-1}^*, ..., X_1^*$.

This gradual approach should provide further cost savings, since it offers more flexibility in optimization. The main idea of this approach is that intermediate inventory targets could stabilize the process of running down the base stock. If unusually many demand arrivals occur early within the transient period, a gradual policy allows for some restocking and should thereby decrease the backordering risks that are linked with early policy change. Additionally, since these risks are reduced, policy changes could start earlier and thereby increase the certainty that excess stock is taken away by demand before T.

For instances not directly taken from Pinçe and Dekker (2011) I need to find appropriate values of S_0 and S_1 corresponding to the chosen λ_0 and λ_1. I find those by comparing $TC_S(0)$ for a reasonably large range of base stock values S in a simulation without obsolescence ($S_0 = S_1 = S$ and $\lambda_0 = \lambda_1 = \lambda$). By ensuring a long enough duration of the calculation period, the simulation reflects a steady state inventory process. This way, I find the most efficient base stock level S_i for λ_i. This is done as a separate, additional step before the usual application of my extension model.

4 Numerical Study

First, the results of the replication of Pinçe and Dekker (2011) are given. As the simulation in their paper is not described in great detail, it is necessary to ensure that my model yields similar results to the one used in Pinçe and Dekker (2011) by comparing the numerical results.

4.1 Replication

Table 1: Simulation Results for partial obsolescence (Replication)

$h = 1$															
λ_0	λ_1	π	L	S_0	N	X_S^*	X^*	$\Delta_X(\%)$	$TC_S^*(X_S^*)$	$TC_S(X^*)$	$\Delta_C(\%)$	Δ_C	$TC_S(0)$	$\Delta_0(\%)$	Δ_0
0.5	0.2	50	0.5	2	1	1.43	1.64	14.69	7.74	7.77	0.43	0.03	9.96	28.72	2.22
			0.75	2	1	1.01	1.11	9.90	8.25	8.32	0.82	0.07	9.95	20.69	1.71
			1	2	1	0.85	0.8	-5.88	8.92	8.98	0.65	0.06	10.02	12.31	1.10
		300	0.75	3	1	1.38	1.45	5.07	11.72	11.75	0.26	0.03	14.88	26.92	3.16
			1	3	1	1.01	1.05	3.96	12.46	12.48	0.12	0.02	14.89	19.44	2.42
1	0.2	50	0.5	2	1	0.81	0.87	7.41	7.12	7.16	0.44	0.03	9.86	38.37	2.73
			0.75	3	2	1.71	1.58	-7.60	14.80	14.84	0.27	0.04	24.23	63.74	9.43
			1	3	2	1.27	1.31	3.15	16.38	16.45	0.43	0.07	24.26	48.08	7.88
		300	0.5	3	1	1.26	1.06	-15.87	9.57	9.75	1.93	0.18	14.78	54.47	5.21
			0.75	4	2	1.53	1.65	7.84	20.27	20.48	1.04	0.21	34.55	70.48	14.28
			1	5	3	2.32	2.19	-5.60	32.41	32.73	0.99	0.32	58.19	79.56	25.78
5	2	5	0.05	1	1	0.15	0.15	0	0.40	0.40	0	0.00	0.50	24.51	0.10
			0.15	2	1	0.21	0.19	-9.52	0.73	0.73	0.40	0.00	0.96	31.47	0.23
			0.25	2	1	0.13	0.12	-7.69	0.86	0.86	0.05	0.00	0.94	8.83	0.08
		50	0.05	2	1	0.16	0.16	0	0.78	0.78	0	0.00	1.00	27.58	0.22
			0.15	3	1	0.15	0.14	-6.67	1.17	1.18	0.94	0.01	1.46	24.58	0.29
			0.25	4	2	0.24	0.25	4.17	2.53	2.54	0.51	0.01	3.40	34.49	0.87
10	2	5	0.05	1	1	0.10	0.09	-10.00	0.35	0.35	0.28	0.00	0.51	43.68	0.15
			0.15	3	2	0.24	0.25	4.17	1.07	1.07	0.37	0.00	2.42	126.46	1.35
			0.25	4	3	0.31	0.3	-3.23	1.88	1.89	0.23	0.00	4.27	126.60	2.38
		50	0.05	2	1	0.10	0.09	-10.00	0.69	0.69	0.02	0.00	1.00	44.12	0.31
			0.15	4	2	0.17	0.18	5.88	1.92	1.93	0.34	0.01	3.40	76.90	1.48
			0.25	6	4	0.30	0.31	3.33	4.35	4.36	0.30	0.01	8.70	100.31	4.36

As we can see in Table 1, for partial obsolescence the optimal times for the advance policy change according to my simulation are quite close to those from the analytical approximation in Pinçe and Dekker (2011). As one could expect, the difference in costs $\Delta_C(\%)$ is quite small due to the closeness of the two values X_S^* and X^*. The roughly convex shape of the total cost function with a rather flat section around the minimum explains why $\Delta_C(\%)$ is much smaller than $|\Delta_X|(\%)$.

As Pinçe and Dekker (2011) already found, the cost savings from an advance policy change against the case of no advance policy change are also significant in my model. This is shown by the large values of $\Delta_0(\%)$.

Beyond that, some patterns of cost drivers can be found in Table 1 as well. Larger values of L, all else being equal, always lead to increased optimal costs, while $TC(0)$ is hardly affected. The inflexibility that comes with longer lead times makes advance policy changes more risky and expensive. This also leads to lower values of X_S^*, as less of an advance policy change decreases exposure to the increased backorder risks. This effect is enhanced by increases in π, which understandably also increase total costs, but counteracted by increases in S_0 which decrease the probability of backorders.

Table 2: Simulation Results for full obsolescence (Replication)

$L = 0.25$, $h = 1$													
λ_0	π	c_0	N	X_S^*	X^*	$\Delta_X(\%)$	$TC_S^*(X_S^*)$	$TC_S(X^*)$	$\Delta_C(\%)$	Δ_C	$TC_S(0)$	$\Delta_0(\%)$	Δ_0
0.5	50	5	1	0.43	0.43	0.00	4.73	4.73	0.00	0.00	5.03	6.46	0.30
		10	1	0.96	0.99	3.13	8.23	8.25	0.22	0.02	10.03	21.96	1.80
	300	5	2	0.40	0.40	0.00	9.85	9.85	0.00	0.00	10.01	1.65	0.16
		10	2	0.95	0.94	-1.05	18.03	18.03	0.01	0.00	20.01	10.97	1.98
1	50	5	2	0.90	0.87	-3.33	8.28	8.29	0.12	0.01	10.01	20.89	1.73
		10	2	1.33	1.40	5.26	13.27	13.33	0.48	0.06	20.01	50.88	6.74
	300	5	2	0.32	0.32	0.00	9.43	9.43	0.00	0.00	10.04	6.51	0.61
		10	2	0.51	0.53	3.92	17.22	17.23	0.11	0.01	20.04	16.42	2.82
5	5	5	2	0.61	0.59	-3.28	3.39	3.39	0.15	0.00	10.07	197.42	6.68
		10	2	0.72	0.75	4.17	4.32	4.33	0.10	0.01	20.07	364.40	15.75
	50	5	4	0.52	0.52	0.00	11.92	11.92	0.00	0.00	20.03	68.05	8.11
		10	4	0.69	0.67	-2.90	18.44	18.51	0.39	0.07	40.03	117.06	21.59
10	5	5	4	0.55	0.55	0.00	4.73	4.73	0.00	0.00	20.05	324.04	15.32
		10	4	0.68	0.65	-4.41	5.87	5.88	0.33	0.01	40.05	582.79	34.18
	50	5	6	0.44	0.44	0.00	14.89	14.89	0.00	0.00	30.04	101.72	15.15
		10	6	0.54	0.53	-1.85	22.81	22.83	0.11	0.02	60.04	163.27	37.23

In the case of full obsolescence we can see a similar pattern as with partial obsolescence. Table 2 shows that not changing policy in advance is much more costly than changing policy at X_S^*. This effect is even more pronounced than in cases of partial obsolescence.

The explanation for the high $\Delta_0(\%)$ can be found in c_0 which introduces a guaranteed cost if the policy is not changed in advance. Thus, higher values of c_0 lead to higher $\Delta_0(\%)$ and push the optimal X_S^* towards higher values as the balance between obsolescence costs and backorder costs shifts.

While the difference between X_S^* and X^* is much lower in full obsolescence cases, the difference in costs $\Delta_C(\%)$ is also lower than $|\Delta_X|(\%)$. This implies that the total cost function is similarly flat around the minimum. The analytical approximation is much closer to the simulated cost function than in partial obsolescence cases.

Table 3: Summary of measures of difference

	$\|\Delta_X\|(\%)$		$\Delta_X(\%)$		$\Delta_C(\%)$		Δ_C		$\Delta_0(\%)$		Δ_0	
	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$
replication												
min	0.00	0.00	-15.87	-4.41	0.00	0.00	0.00	0.00	8.83	1.65	0.08	0.16
max	15.87	5.26	14.69	5.26	1.93	0.48	0.32	0.07	126.60	582.79	25.78	37.23
mean	6.59	2.08	-0.54	-0.02	0.47	0.13	0.05	0.01	49.23	128.41	3.81	10.63
median	5.88	2.38	0.00	0.00	0.37	0.11	0.01	0.01	38.37	59.47	1.71	6.71
number instances	23	16	23	16	23	16	23	16	23	16	23	16

Table 3 also shows how $|\Delta_X|(\%)$ and $\Delta_C(\%)$ are much lower for full obsolescence cases. However, $\Delta_0(\%)$ is much larger in case of full obsolescence. This implies that advance policy changes are more advantageous and also approximated more accurately in full obsolescence cases.

In both partial obsolescence and full obsolescence we see that the more items have to be taken away by demand (N), the earlier we tend to change policy. Inversely, higher demand rates before obsolescence λ_0 tend to decrease X_S^*. This seems to be a quite intuitive driver of X_S^* and because of that, Pinçe and Dekker (2011) investigate a direct link between $\frac{N}{\lambda_0}$ and X_S^*, as well as between $\frac{N}{\lambda_0}$ and $TC_S^*(X_S^*)$. I investigate these same links for my extension model in Section 4.2.2 in order to find further differences or similarities in the performance of the two frameworks.

All in all, the numerical study of the replication framework finds that the work of Pinçe and Dekker (2011) can be replicated with quite comparable results, which both validates their findings and indicates that my simulation framework is successful in modelling and investigating the scenario.

4.2 Extensions

The extension framework is a different approach to the question of optimizing X. Table 4 shows that we have much higher cost in general now that the calculation period is extended and fixed at $T - t_0$ to $T + t_{cooldown}$. Due to this, the measures of difference in percentages are lower than in the replication framework, as can also be seen in Table 6.

4.2.1 Biases of X^* in the Replication Model

Table 4: Simulation Results for partial obsolescence (fixed calculation period)

$h = 1$														
λ_0	λ_1	π	L	S_0	N	X_S^*	X^*	$\Delta_X(\%)$	$TC_S^*(X_S^*)$	$TC_S(X^*)$	$\Delta_C(\%)$	Δ_C	$TC_S(0)$	$\Delta_0(\%)$	Δ_0
0.5	0.2	50	0.5	2	1	3.77	1.64	-56.54	80.83	81.63	0.99	0.80	83.79	3.66	2.96
			0.75	2	1	1.28	1.11	-13.28	96.86	96.91	0.05	0.05	97.82	0.99	0.96
			1	2	1	0.54	0.8	49.25	118.93	119.00	0.06	0.07	119.07	0.11	0.14
		300	0.75	3	1	2.21	1.45	-34.30	159.24	159.79	0.34	0.55	161.89	1.66	2.65
			1	3	1	1.71	1.05	-38.78	171.58	171.78	0.12	0.20	172.87	0.75	1.29
1	0.2	50	0.5	2	1	0.86	0.87	1.52	85.08	85.10	0.02	0.02	86.90	2.14	1.82
			0.75	3	2	1.35	1.58	17.21	104.66	104.89	0.22	0.23	110.34	5.43	5.69
			1	3	2	1.05	1.31	24.64	127.74	128.12	0.30	0.39	131.10	2.64	3.37
		300	0.5	3	1	1.13	1.06	-6.44	157.54	157.58	0.03	0.04	160.07	1.61	2.53
			0.75	4	2	1.37	1.65	20.53	169.24	169.45	0.12	0.20	176.75	4.44	7.51
			1	5	3	2.11	2.19	3.84	188.02	188.27	0.13	0.25	203.52	8.24	15.50
5	2	5	0.05	1	1	0.32	0.15	-53.27	10.63	10.66	0.23	0.02	10.76	1.23	0.13
			0.15	2	1	0.88	0.19	-78.31	18.77	18.93	0.86	0.16	19.19	2.27	0.43
			0.25	2	1	0.25	0.12	-51.81	22.58	22.63	0.21	0.05	22.74	0.69	0.16
		50	0.05	2	1	0.46	0.16	-64.99	13.01	13.09	0.61	0.08	13.29	2.16	0.28
			0.15	3	1	0.29	0.14	-51.56	20.75	20.81	0.30	0.06	21.03	1.36	0.28
			0.25	4	2	0.27	0.25	-8.09	25.72	25.73	0.04	0.01	26.20	1.87	0.48
10	2	5	0.05	1	1	0.11	0.09	-15.09	6.40	6.41	0.05	0.00	6.52	1.77	0.11
			0.15	3	2	0.28	0.25	-11.66	11.95	11.97	0.20	0.02	13.10	9.64	1.15
			0.25	4	3	0.32	0.30	-6.25	14.79	14.80	0.06	0.01	16.68	12.79	1.89
		50	0.05	2	1	0.09	0.09	1.12	14.05	14.05	0.00	0.00	14.24	1.35	0.19
			0.15	4	2	0.18	0.18	-0.55	23.27	23.27	0.01	0.00	24.21	4.06	0.94
			0.25	6	4	0.28	0.31	9.54	28.98	29.02	0.14	0.04	31.69	9.35	2.71

The X_S^* seems to differ from X^* by a large margin, although a consistent tendency towards larger or lower optimal X is not found. Comparing the mean values of $|\Delta_X|(\%)$ in Table 3 and Table 6, we can see that for both full and partial obsolescence the differences increased. They are about nine-fold and six-fold of the values in the replication framework respectively. As I explained in Section 3.2, the replication model is biased due to focusing only on costs incurred in the transient period. This bias is ambiguous in partial obsolescence cases, but for full obsolescence there is a clear downward bias. Consequently, we do see the expected shift of X_S^* to higher values for every investigated instance in Table 5.

Table 5: Simulation Results for full obsolescence (fixed calculation period)

$L = 0.25$, $h = 1$													
λ_0	π	c_0	N	X_S^*	X^*	$\Delta_X(\%)$	$TC_S^*(X_S^*)$	$TC_S(X^*)$	$\Delta_C(\%)$	Δ_C	$TC_S(0)$	$\Delta_0(\%)$	Δ_0
0.5	50	5	1	0.92	0.43	-53.41	6.43	6.70	4.22	0.27	7.54	17.25	1.11
		10	1	1.42	0.99	-30.23	9.23	9.50	2.95	0.27	12.54	35.87	3.31
	300	5	2	0.94	0.40	-57.26	12.50	13.00	4.00	0.50	13.94	11.55	1.44
		10	2	1.24	0.94	-24.50	19.87	20.22	1.79	0.36	23.94	20.50	4.07
1	50	5	2	1.17	0.87	-25.89	10.09	10.38	2.86	0.29	13.74	36.14	3.65
		10	2	1.66	1.40	-15.71	14.14	14.41	1.88	0.27	23.74	67.83	9.59
	300	5	2	0.45	0.32	-28.41	13.36	13.53	1.32	0.18	14.91	11.63	1.55
		10	2	0.62	0.53	-13.82	20.74	20.84	0.50	0.10	24.91	20.13	4.17
5	5	5	2	0.69	0.59	-14.37	5.89	5.97	1.29	0.08	13.73	133.01	7.84
		10	2	0.83	0.75	-9.96	6.53	6.60	1.07	0.07	23.73	263.27	17.20
	50	5	4	0.56	0.52	-6.81	16.66	16.74	0.50	0.08	26.68	60.15	10.02
		10	4	0.71	0.67	-5.77	22.69	22.75	0.27	0.06	46.68	105.76	23.99
10	5	5	4	0.58	0.55	-5.01	8.34	8.38	0.53	0.04	25.09	201.03	16.76
		10	4	0.68	0.65	-4.83	9.25	9.29	0.44	0.04	45.09	387.71	35.85
	50	5	6	0.45	0.44	-3.08	21.75	21.78	0.11	0.02	39.07	79.62	17.32
		10	6	0.54	0.53	-2.57	29.15	29.17	0.05	0.02	69.07	136.93	39.92

The column on $\Delta_X(\%)$ in Table 6 shows that for $\lambda_1 > 0$, the values of X_S^* are also higher than in the replication framework on average, but in some instances the values decreased, so the trend is not entirely consistent. This is also in line with my predictions described in 3.2.

Table 6: Summary of measures of difference (fixed calculation period)

	$\lvert\Delta_X\rvert(\%)$		$\Delta_X(\%)$		$\Delta_C(\%)$		Δ_C		$\Delta_0(\%)$		Δ_0	
	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$
constant time frame												
min	0.55	2.57	-78.31	-57.26	0.00	0.05	0.00	0.02	0.11	11.55	0.11	1.11
max	78.31	57.26	49.25	-2.57	0.99	4.22	0.80	0.50	12.79	387.71	15.50	39.92
mean	26.89	18.85	-15.79	-18.85	0.22	1.49	0.14	0.17	3.49	99.27	2.31	12.36
median	17.21	14.09	-8.09	-14.09	0.13	1.18	0.05	0.09	2.14	63.99	1.15	8.72
number instances	23	16	23	16	23	16	23	16	23	16	23	16

The direction of the overall bias in partial obsolescence cases can be analyzed with the results found in Table 4 and the exemplary simulated cost functions displayed in Figure 2. First we notice that the Figure 2b shows how the total cost computed for the transient period increases relative to $TC_S(X)$ computed in a constant timeframe. This is the behaviour I expected, as larger X will always increase the length of the calculation period.

In partial obsolescence cases however, there are two effects that could lead to a bias. For low X, the removal period can take very long if the excess stock needs to be removed under the low post-obsolescence demand arrival rate λ_1. This leads to a bias against low X. On the other hand, the same issue as in full obsolescence cases remains, where high values of X imply an early start of the calculation period. Therefore, there is also a bias against high values of X. Indeed, as Figure 2a shows, the total costs in the transient period show a steeper inclination than in the constant timeframe in both directions from the found minimum.

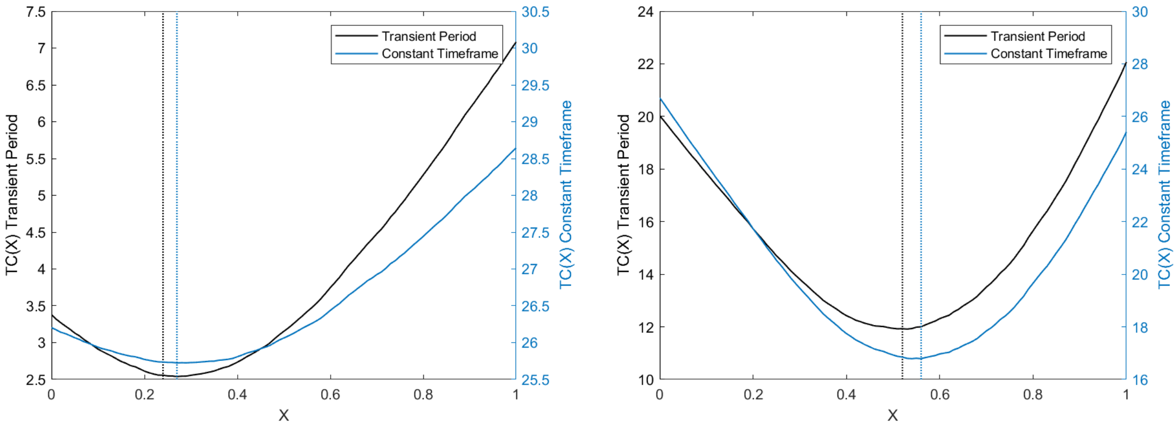

(a) Partial Obsolescence (b) Full Obsolescence

Figure 2: Comparison of simulated total cost functions.
($\lambda_0 = 5, \lambda_1 = 2$ (0 in b)$, \pi = 50, c_0 = 5, L = 0.25, S_0 = 4, N = 2$ (4 in b))

This leads to the conclusion that the overall bias depends on the severity of the two competing biases in each particular instance. The two most impactful parameters driving the tendency towards either direction are found to be λ_1 and N. When λ_1 is especially low, the calculation period will be extended dramatically if the excess stock is not removed before T. This implies a strong severity of the bias against lower X. Similarly when N is larger, there is more excess stock, so even if λ_1 is relatively high it could take rather long until the transient is finished. Conversely, higher λ_1 and low N lead to low severity of the upward bias and thus, the bias against high X dominates in these cases.

These drivers of the overall bias can also be identified through an analysis of Table 4. For $\lambda_1 = 0.2$, 6 instances out of 11 have too high X^*, while for $\lambda_1 = 2$, only 2 instances out of 12 have a value of X^* that is too high. A similar pattern exists for N, where in cases of $N = 1$ there are 11 out of 14 times where X^* is too low. For $N > 1$ though, X^* is too low in only 4 out of 9 instances.

4.2.2 General Behaviour of Extension Model

Another insight that can be gained from Figure 2 is that, apart from the biases in the replication model, the total cost functions still have roughly the same convex shape. The minimum still optimizes the trade-off between backorder costs and obsolescence costs, and $TC_S^*(X_S^*)$ is still clearly cheaper when compared to $TC_S(0)$ as seen under Δ_0 in Table 6.

Further characteristics like the effects of individual parameters on the total costs also remain stable. Longer lead times L increase $TC_S^*(X_S^*)$ and since the costs are calculated over a relatively long fixed period, $TC_S(0)$ is also affected by this. However, the effect on $TC_S(0)$ is less pronounced as we see Δ_0 declining in L. Due to increased backorder risks, X_S^* decreases when L or π are increased and all else remains equal. In the same way, increases of S_0 make backorders less likely and have the opposite effects of L and π.

The effects of c_0 in full obsolescence cases also remain unchanged. Higher c_0 lead to earlier policy changes, higher $TC_S^*(X_S^*)$ and much higher $TC_S(0)$. The advantage of advance policy changes becomes larger with higher obsolescence costs, as can be seen in the large increases in Δ_0.

Presumably, the most intuitive drivers of the optimal choice of X are λ_0 and N. When more excess stock needs to be disposed, one would expect that the disposal should start earlier, all else being equal. If demand arrives more frequently pre-obsolescence, then, all else being equal, the expectation is that X should be lowered since removal happens more quickly.

Therefore, I investigate the link between $\frac{N}{\lambda_0}$ and X_S^* as well as the link between $\frac{N}{\lambda_0}$ and $TC_S^*(X_S^*)$ in Figure 3.

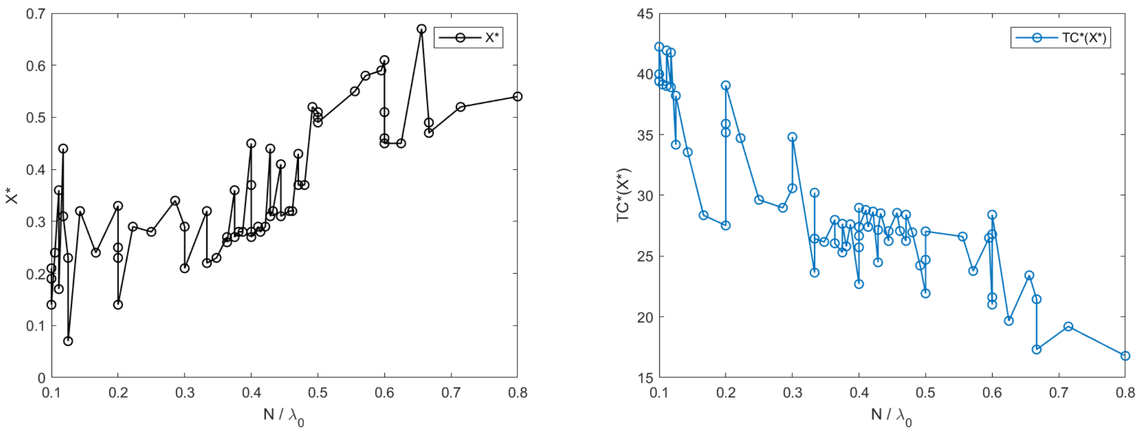

(a) Effect on X_S^*

(b) Effect on $TC_S^*(X_S^*)$

Figure 3: Effect of N/λ_0 on optimal choice of timing and total costs.
(λ_0 ranging from 5 to 10, N ranging from 1 to 6, $\pi = 50, c_0 = 5, L = 0.25$)

In their paper, Pinçe and Dekker (2011) also analyze these connections and find that although there is no linear dependence and rather much random noise in their graph, $\frac{N}{\lambda_0}$ is positively correlated with both X_S^* as well as with $TC_S^*(X_S^*)$.

Using my extension model, I find a similar positive correlation of $\frac{N}{\lambda_0}$ with X_S^* in Figure 3a. There is no linear connection and again we have rather much random noise obscuring the correlation. In a similar fashion, the effect of $\frac{N}{\lambda_0}$ on $TC_S^*(X_S^*)$ is clearly negative in my extension model, although the behaviour is also not linear as can be seen in Figure 3b.

The positive correlation between $\frac{N}{\lambda_0}$ and X_S^* is in accordance with the mentioned intuition. If N is larger or λ_0 lower, hence if $\frac{N}{\lambda_0}$ increases it should take longer until the excess stock is taken away by demand pre-obsolescence. This directly leads to an earlier implementation of the policy change through a higher X.

The effect of $\frac{N}{\lambda_0}$ on $TC_S^*(X_S^*)$ can be split up into effects of N and of λ_0 on $TC_S^*(X_S^*)$. It is clear that a higher N will increase costs, since pre-obsolescence holding costs, as well as possible excess holding costs post-obsolescence and possible disposal costs are all increasing in N. The effect of λ_0 on costs is generally positive, because with more demand there is a higher chance and frequency of backorders. In this way both parameters should increase $TC_S^*(X_S^*)$ and that for my extension model the effect of λ_0 seems to outweigh that of N.

This difference in the effect of $\frac{N}{\lambda_0}$ on $TC_S^*(X_S^*)$ can be explained by the design of the two different models. While N should have the same effects in both models, the calculation period of varying length relies greatly on λ_0. A higher λ_0 increases costs per time, but it also significantly shortens the duration of the transient period. In this way a lower λ_0 increases costs in the replication model.

Combining this knowledge, it becomes clear that the replication model should find a positive correlation between $\frac{N}{\lambda_0}$ and $TC_S^*(X_S^*)$, even as the extension model finds a negative correlation.

4.2.3 Gradual Policy Change

Another step in finding the optimal approach of dealing with anticipated obsolescence is to use the predictions to implement a gradual policy change. Instances that allow for step-wise reductions of the inventory target, which are all those with $N > 1$, are investigated further and the results can be seen in the Appendix in Tables 8 and 9 for partial and full obsolescence cases respectively.

The behaviour of the simulated total costs when reductions are considered in separate steps is depicted in Figure 4. The backwards induction approach involves high costs when not all reduction-steps are made, so the total cost curve seen for ‘Step 1’ in Figure 4a understandably indicates a worse performance than the cost curve for ‘Not Step-Wise’.

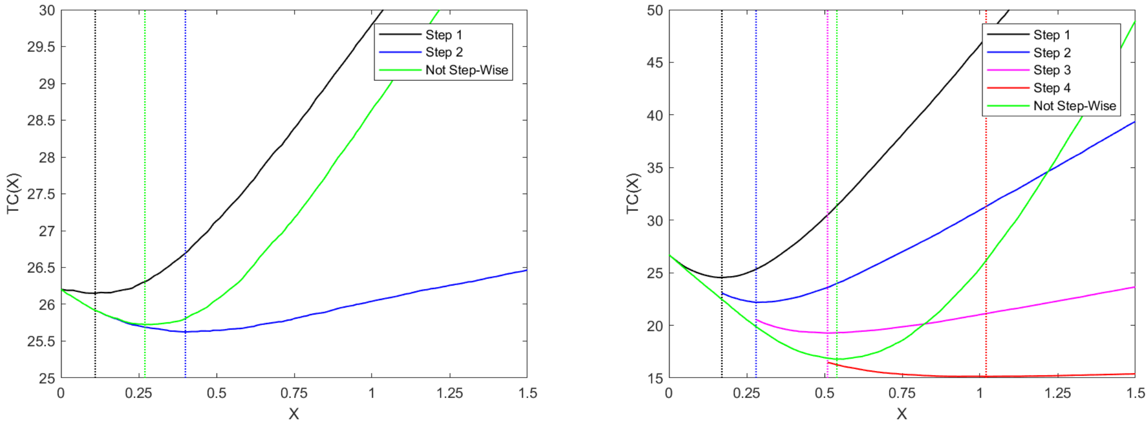

(a) Partial Obsolescence (b) Full Obsolescence

Figure 4: Simulated total cost functions for gradual approach.
($\lambda_0 = 5$, $\lambda_1 = 2$ (0 in b), $\pi = 50$, $c_0 = 5$, $L = 0.25$, $S_0 = 4$, $N = 2$ (4 in b))

This relation reverses entirely when the costs of a 'Not-Step-Wise' policy is compared to the costs of 'Step 2'. The total cost curve for 'Step 2' shows $TC_S(X_2|X_1 = X_1^*)$, which is the total simulated cost arising from the inventory policy that reduces the inventory target at the variable point in time $X_2 = X$ and at $X_1 = X_1^*$, where X_1^* has been determined in the previous iteration. Similarly, in Figure 4b the total cost curve of 'Step 4' shows $TC_S(X_4\,|\,X_3 = X_3^*,\, X_2 = X_2^*,\, X_1 = X_1^*)$.

Focusing on Figure 4a, we can see that the gradual policy change yields even lower optimal costs than the non-gradual approach. In addition to the lower optimal costs, we can also see that once X_1 is fixed at X_1^*, even suboptimal choices of X_2 yield comparatively low costs. Thus, the intermediate base stock level $S_i = S_0 - 1$ can provide a reasonable safety against backorders in this case. Figure 4a demonstrates the effect predicted in 3.3, as we see how earlier policy changes become less costly in a gradual approach.

In the case of full obsolescence the costs of failing to run down the inventory are higher, so being able to start reducing the inventory target at an earlier time with rather cheap backorder costs leads to even larger cost savings. We can see in Figure 4b that the second inventory target reduction (X_3^*, minimum of 'Step 3') happens shortly after the change in a non-gradual scenario.

These observations from Figure 4 imply that the non-gradual approach is a suboptimal compromise that balances a difficult trade-off between backorder and obsolescence risks, whereas the gradual policy change can significantly decrease both risks.

We can evaluate the results from the gradual policy change approach by analyzing the absolute cost measures shown in Table 7. There we can see that the gradual policy change clearly dominates the constant time frame model with S_0 to S_1 policy change. The mean cost savings Δ_C with respect to the analytical approximation X^* by Pinçe and Dekker (2011) increase by roughly 0.3 for partial and roughly 1.4 for full obsolescence.

The mean cost savings Δ_0, which are with respect to a policy change at 0, methodologically have to increase by the same amounts and so they do, apart from a small difference due to rounding. As we have found in 4.2.1, the $\Delta_X(\%)$ is significantly more negative in my extension models than in the replication framework. In the gradual change model $\Delta_X(\%)$ decreases even further. This can be attributed to the use of the earliest policy change X_N^* to compute $\Delta_X(\%)$. Since the inventory target reduction-steps are spread apart in most instances, the first reduction X_N^* tends to occur much earlier than X_S^* or X^*.

Table 7: Development of differences across models

	$\Delta_X(\%)$			Δ_C			Δ_0		
	min	mean	max	min	mean	max	min	mean	max
$\lambda_1 > 0$									
replication	-15.87	-0.54	14.69	0.00	0.05	0.32	0.08	3.81	25.78
constant time frame	-78.31	-15.79	49.25	0.00	0.14	0.80	0.11	2.31	15.50
gradual change	-78.31	-29.76	49.25	0.00	0.47	4.13	0.11	2.64	19.38
$\lambda_1 = 0$									
replication	-4.41	-0.02	5.26	0.00	0.01	0.07	0.16	10.63	37.23
constant time frame	-57.26	-18.85	-2.57	0.02	0.17	0.50	1.11	12.36	39.92
gradual change	-81.28	-47.07	-25.52	0.23	1.54	3.54	1.12	13.72	43.59

An important finding that can be shown with Table 7 is the effect of the replication frameworks biases on the stated total costs. While the cost savings Δ_0 for full obsolescence cases are underestimated in the replication, Δ_0 is overestimated for partial obsolescence cases in the replication. This means that the findings of Pinçe and Dekker (2011) claim to reach cost savings that can not be found, when the bias of their calculation method is corrected.

One obvious caveat in analyzing Table 7 is the underrepresentation of instances with $N > 1$ in for partial obsolescence. Only for 9 out 23 investigated instances a gradual policy can be applied, while for full obsolescence 14 out of 16 instances can use a gradual change policy. This fact alone may explain why the gradual change appears much more effective with $\lambda_1 = 0$.

Overall Table 7 also shows that the analytical approximation by Pinçe and Dekker (2011) achieves a majority of the possible cost reduction, since Δ_C is much lower than Δ_0. However, I must refute their claim that the approximation is near optimal, because especially the gradual change offers a significant increase in cost savings across most instances.

5 Conclusion

This paper investigated the effects of advance inventory policy changes on inventory costs in a scenario of known future obsolescence. Multiple instances of full obsolescence, as well as partial obsolescence were analyzed and optimally timed advance inventory policy changes were found by simulation. In the simulation, polynomial triangulation and discretization were employed to find approximately optimal values for the continuous optimization variable X.

In doing this, the paper also replicated Pinçe and Dekker (2011) and was able to uphold the original findings. Both the significant cost reduction from using advance policies at all and the near optimality of the analytical approximation (Pinçe & Dekker, 2011) within the chosen optimization framework could be confirmed. I also find that the timing is based on a trade-off of excess holding cost or obsolescence cost on one hand and the risk of backorder penalties on the other, as Pinçe and Dekker (2011) do as well. The balance of said trade-off is influenced towards later change by increases in parameters such as lead time, backorder penalty and demand rate, while increases in the original base stock, the number of excess items and the disposal cost shift the optimum towards earlier change.

However, by changing the framework of optimization, I have shown that further cost savings can be achieved when some limitations of the model of Pinçe and Dekker (2011) are overcome. By establishing a fixed calculation period that encompasses all occurring transient periods for a reasonable range of candidate values for optimization, I have found the approximately optimal inventory policy that leads to the lowest costs over the entire pre-obsolescence to post-obsolescence time horizon.

I show that, generally speaking, inventory policy should be changed even earlier than according to Pinçe and Dekker (2011). As my framework offers a more comprehensive perspective of the cost effect of timing, I find that the analytical approximations by Pinçe and Dekker (2011) are not, as they claim, near optimal, but still achieve a majority of the possible cost savings.

In particular I find that the application of gradual inventory policy change through step-wise reduction yields the greatest cost savings. This is due to the added flexibility that allows to both remove stock with more certainty by starting to reduce earlier and incur less backorder penalties by keeping intermediate inventory positions.

This paper has only worked with a generalized time scale. For some slow moving spare part inventory processes the appropriate timing would be measured in years. In this case some instances that use calculation periods of a length of up to 80 years may be highly unrealistic as product cycles even after partial obsolescence will likely not reach such lengths.

The work in this paper is also limited by some assumptions that allow the analysis to be based on deterministic values. The lead times are assumed to be fixed, and so are the timing of obsolescence and the degree to which obsolescence will occur. The work in this paper should be extended by investigating the effects of random lead times and uncertainty in T and λ_1. I have built a robust dynamic programming model to analyze inventory processes under various conditions and my programs will be available in the Appendix of this thesis. For future investigations in this field this could serve as an efficient tool to implement and investigate further extensions.

Another restriction of the findings in this paper is the fixation on continuous review one-for-one replenishment policies of slow moving items. Again, I suggest to extend my dynamic programming approach to investigate similar questions as have been research in this paper for $(S-n,S)$ policies and periodic review policies.

References

Dekker, R., Pinçe, Ç., Zuidwijk, R., & Jalil, M. N. (2013). On the use of installed base information for spare parts logistics: A review of ideas and industry practice. *International Journal of Production Economics*, *143*(2), 536–545.

Haneveld, W. K. K., & Teunter, R. H. (1998). Effects of discounting and demand rate variability on the eoq. *International Journal of Production Economics*, *54*(2), 173–192.

Jalil, M. N., Zuidwijk, R. A., Fleischmann, M., & Van Nunen, J. A. (2011). Spare parts logistics and installed base information. *Journal of the operational Research Society*, *62*(3), 442–457.

Milovanovic, N., & Bouman, P. (2018). *Discrete event simulation framework [computer software].* Retrieved during the course 'Simulation'.

Montgomery, D. C., & Myers, R. H. (1995). Response surface methodology. *Design and Analysis of Experiments*, 445–474.

Pinçe, Ç., & Dekker, R. (2011). An inventory model for slow moving items subject to obsolescence. *European Journal of Operational Research*, *213*(1), 83–95.

Pinçe, Ç., Frenk, J., & Dekker, R. (2015). The role of contract expirations in service parts management. *Production and Operations Management*, *24*(10), 1580–1597.

Van Jaarsveld, W., & Dekker, R. (2011). Estimating obsolescence risk from demand data to enhance inventory control—a case study. *International Journal of Production Economics*, *133*(1), 423–431.

Appendix

Table 8: Simulation Results for partial obsolescence (gradual policy change)

$h = 1$																		
λ_0	λ_1	π	L	S_0	N	X_4^*	X_3^*	X_2^*	X_1^*	X^*	$\Delta_X(\%)$	$TC_S^*(X_S^*)$	$TC_S(X^*)$	$\Delta_C(\%)$	Δ_C	$TC_S(0)$	$\Delta_0(\%)$	Δ_0
1	0.2	50	0.75	3	2			2.35	0.55	1.58	-32.77	103.67	104.89	1.18	1.22	110.34	6.44	6.67
			1	3	2			1.35	0.26	1.31	-2.82	127.32	128.12	0.63	0.80	131.10	2.97	3.78
		300	0.75	4	2			2.56	0.73	1.65	-35.57	167.48	169.45	1.17	1.96	176.75	5.53	9.27
			1	5	3		5.25	1.36	0.68	2.19	-58.27	184.14	188.27	2.24	4.13	203.52	10.53	19.38
5	2	50	0.25	4	2			0.40	0.09	0.25	-37.34	25.62	25.73	0.43	0.11	26.20	2.27	0.58
10	2	5	0.15	3	2			0.50	0.18	0.25	-49.90	11.87	11.97	0.83	0.10	13.10	10.34	1.23
			0.25	4	3		0.45	0.25	0.13	0.30	-33.33	14.68	14.80	0.81	0.12	16.68	13.65	2.00
		50	0.15	4	2			0.23	0.11	0.18	-21.74	23.21	23.27	0.24	0.06	24.21	4.30	1.00
			0.25	6	4	0.44	0.27	0.17	0.06	0.31	-0.29	28.73	29.02	1.03	0.29	31.69	10.31	2.96

Table 9: Simulation Results for full obsolescence (gradual policy change)

$L = 0.25$, $h = 1$																		
λ_0	π	c_0	N	X_6^*	X_5^*	X_4^*	X_3^*	X_2^*	X_1^*	X^*	$\Delta_X(\%)$	$TC_S^*(X_S^*)$	$TC_S(X^*)$	$\Delta_C(\%)$	Δ_C	$TC_S(0)$	$\Delta_0(\%)$	Δ_0
0.5	300	5	2					2.14	0.28	0.40	-81.28	17.22	18.99	10.28	1.77	19.88	15.43	2.66
		10	2					3.31	0.40	0.94	-71.59	22.86	26.28	14.94	3.42	29.88	30.73	7.02
1	50	5	2					2.93	0.44	0.87	-70.36	14.14	16.00	13.18	1.86	19.37	37.00	5.23
		10	2					3.86	0.74	1.40	-63.71	17.06	19.96	17.00	2.90	29.37	72.18	12.31
	300	5	2					0.66	0.17	0.32	-51.22	20.52	21.07	2.67	0.55	22.43	9.29	1.91
		10	2					1.13	0.26	0.53	-52.97	26.99	28.54	5.77	1.56	32.43	20.17	5.44
5	5	5	2					0.84	0.42	0.59	-29.43	5.73	6.00	4.86	0.28	13.73	139.76	8.00
		10	2					1.01	0.56	0.75	-25.52	6.35	6.66	4.84	0.31	23.73	273.64	17.38
	50	5	4			1.02	0.51	0.29	0.16	0.52	-49.17	15.13	16.85	11.37	1.72	26.71	76.51	11.58
		10	4			1.22	0.68	0.40	0.24	0.67	-45.31	19.87	22.95	15.51	3.08	46.71	135.10	26.84
10	5	5	4			0.79	0.53	0.38	0.26	0.55	-30.03	7.84	8.32	6.14	0.48	25.08	219.80	17.24
		10	4			0.91	0.63	0.47	0.34	0.65	-28.73	8.56	9.20	7.46	0.64	45.08	426.61	36.52
	50	5	6	0.69	0.51	0.34	0.26	0.19	0.11	0.44	-36.32	19.64	21.69	10.45	2.05	39.03	98.77	19.40
		10	6	0.89	0.62	0.45	0.32	0.24	0.14	0.53	-40.65	25.45	28.99	13.91	3.54	69.03	171.30	43.59

Table 10: Summary of measures of difference (gradual policy change)

	$\|\Delta_X\|(\%)$		$\Delta_X(\%)$		$\Delta_C(\%)$		Δ_C		$\Delta_0(\%)$		Δ_0	
	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$	$\lambda_1 > 0$	$\lambda_1 = 0$
gradual change												
min	0.29	25.52	-78.31	-81.28	0.00	2.15	0.00	0.23	0.11	9.29	0.11	1.12
max	78.31	81.28	49.25	-25.52	2.24	17.00	4.13	3.54	13.65	426.61	19.38	43.59
mean	34.27	47.07	-29.76	-47.07	0.54	8.92	0.47	1.54	3.83	110.19	2.64	13.72
median	35.57	46.75	-34.30	-46.75	0.34	8.87	0.10	1.64	2.16	74.34	1.23	9.79
number instances	23	16	23	16	23	16	23	16	23	16	23	16

Used Code for Simulation:

Java package "standard_model" including files:

Table 11: Filenames and Descriptions

Filename	Description
Counter.java	Saves a value (e.g. holding cost) within one replication.
DemandArrivalEvent.java	Explained in 3.1
DemandArrivalEventReplica.java	Explained in 3.1 (for replication framework)
EndEvent.java	Explained in 3.1
EndEventReplica.java	Explained in 3.1 (for replication framework)
Event.java	Discrete Event that is processed
InventoryMain.java	Main method where input is provided
InventoryReplication.java	Sets up and initializes a replication
InventoryReplicationGradual.java	Sets up and initializes a replication (for gradual framework)
InventoryReplicationReplica.java	Sets up and initializes a replication (for replication framework)
InventoryReplicationUncertainty.java	Sets up and initializes a replication (for T uncertain)
InventoryState.java	System state with variables defining it
ObsolescenceEvent.java	obsolescence lowers demand from here on
ObsolescenceEventPredicted.java	obsolescence lowers demand from here on
ObsolescenceEventRandom.java	obsolescence lowers demand from here on
ObsolescenceEventReplica.java	obsolescence lowers demand from here on (for replication framework)
OrderArrivalEvent.java	Explained in 3.1
OrderArrivalEventReplica.java	Explained in 3.1 (for replication framework)
PerformanceMeasure.java	building block for measures like average holding cost
PMBackorderCosts.java	performance measure: self explanatory
PMBackorders.java	performance measure: self explanatory
PMDemandArrivals.java	performance measure: self explanatory
PMHoldingCosts.java	performance measure: self explanatory
PMInventoryAtT.java	performance measure: IL at T
PMInventoryAtX.java	performance measure: IL at X
PMObsolescenceCosts.java	performance measure: self explanatory
PMOrders.java	performance measure: self explanatory
PMOrdersOutAtT.java	performance measure: orders outstanding at T
PMOrdersOutAtX.java	performance measure: orders outstanding at X
PMPeriodEndAlways.java	performance measure: notes if transient period ends in all iterations
PMPeriodEndPerc.java	performance measure: notes if transient period ends in %
PMRegularEnd.java	performance measure: end of regular operation phase
PMRemovalEnd.java	performance measure: end of removal phase
PMSales.java	performance measure: self explanatory
PMTotalCosts.java	performance measure: self explanatory
PolicyChangeEvent.java	inventory target is changed
PolicyChangeEventGradual.java	inventory target is changed
PolicyChangeEventReplica.java	inventory target is changed (for replication framework)
Replication.java	runs one replication
Simulation.java	runs the simulation for specified number of replication
Status.java	part of simulation framework: informs about errors in simulation
SystemState.java	Building block for a system state that is used in the simulation framework
Utils.java	Hold functions to compute random distribution output
WarmUpEvent.java	resets counters to ensure warm up period is not counted in measures