2015

Reducing penalty costs in performance-based contracts

W.S. Sleiderink

Thales Nederland B.V.

3/17/2015

Intentionally left blank

Reducing penalty costs in performance-based contracts

A case study at Thales Nederland B.V.

Wouter Sebastiaan Sleiderink

Thesis submitted in partial fulfillment for the degree of Master of Science in Industrial Engineering and Management, specialized in Production and Logistic Management

3/17/2015

Thales Nederland B.V.
Business Unit Naval Systems
Department Customer Services & Support
Integrated Logistics Support Competence Center
Ir. R. Ypma

University of Twente
Faculty of Management and Governance
Department of Industrial Engineering
and Business Information Systems
Dr. Ir. A. Al Hanbali
Dr. M.C. van der Heijden

Intentionally left blank

Preface

This thesis is the result of my graduation period at Thales Nederland in Hengelo at the Customer Services and Support Department during the period between November 2014 and March 2015.

I am thankful to Thales Nederland for offering me this graduation assignment in the field of Industrial Engineering and Management, specialized in Production and Logistic Management. I would like to express particular gratitude to Rindert Ypma, my supervisor at Thales Nederland during this project. His assistance made sure I could get the information I needed, felt at ease within the department and he came up with some valuable additions for the report. Furthermore, I would like to thank the other colleagues at the department for providing a pleasant working environment and amusing breaks.

Also, I am greatly indebted to my supervisors at the University of Twente at the Faculty of Management and Governance. Ahmad Al Hanbali always offered me support during the complete project and provided valuable feedback. I enjoyed our meetings and discussions as well and overall it was a pleasure working with you. I wish to thank Matthieu van der Heijden for setting me up with this assignment in the first place, but also for his critical view and useful remarks with regard to the report.

Finally, I would like to thank my friends, family and fellow students for their support and interest during my graduation period.

Wouter Sleiderink
Hengelo, 17-3-2015

Management summary

The Customer Service and Support Department at Thales Nederland aims to be the partner of choice for Life Cycle Services of complex technical systems in the defense environment. To establish these intimate customer relationships, Performance-Based Contracts (PBC), resulted from Performance-Based Logistics could be considered as a growing trend. The customer and Thales agree upon a certain service level, generally the average operational availability of a system which is reviewed every period of time. These PBCs include penalty structures to specify the consequences involved when targets are not satisfied. During this project, the focus will mainly be on improving supply availability.

Currently, the spare parts allocations are calculated through INVENTRI, using the VARI-METRIC approach, which ensures that on average the target operational availability is reached against minimal cost. However, the decisions in this approach do not cope with any variability in the interval availabilities, the fraction of time a system is operational during a period of time, which are directly related to the expected penalty costs.

Therefore, the goal of this research is to create insights on which aspects have an impact on this variability in interval availability and to use this information to improve the current spare parts allocation with respect to the expected penalty costs. The corresponding research question will be:

"How can Thales improve their spare parts allocation to achieve higher interval availability in order to reduce the expected penalty costs while maintaining the requirements in the performance-based contracts and the constraint of the spare parts costs budget?"

To answer this research question, we first analyzed a simplified multi-item, multi-echelon model by using simulation. Later on, we added various extensions to this basic model to obtain a more realistic setting. A thorough literature study and the findings of previous research at Thales served as a basis for the identification of interesting parameters and experiments to examine. This has led to the following hypotheses that were validated in the simulation study:

"Stock investments for slow moving items (few failures and long replenishment times) lead to lower penalty costs in the Performance-Based Contracts compared to stock investments in fast moving items (many failures and short replenishment times)."

"Central stock (depot) investments lead to lower penalty costs in the Performance-Based Contracts compared to local stock (base) investments."

Hereafter, we translated these findings into several guidelines to make an improvement upon the initial solution provided by the VARI-METRIC approach. Three of these methods show promising results. The Waiting times method uses the output from the simulation tool and focuses on adding stock to items with long waiting times. The Combination method combines the findings of the literature study and the output from the simulation tool regarding waiting times of items. The Replenishment Lead Time method, which is not based on an initial solution provided by VARI-METRIC, starts from scratch and focuses on items with long replenishment lead times to eliminate variability and penalty costs. The implementation of these guidelines in a real-life example has led to the following results and conclusions as illustrated below. Note that all figures in the table are compared to the initial solution from INVENTRI, thus the VARI-METRIC procedure. Positive numbers mean an increase whereas negative numbers illustrate a decrease for that performance indicator.

Method	Relative difference in penalty costs for current practice	Relative difference in life-cycle costs for current practice	Penalty structure	Relative difference in penalty costs	Relative difference in life-cycle costs
Waiting times	-3,3%	+4,3%	Original Linear Exponential	-2,3% -15,6% -24,4%	+4,7% -2,0% +1,3%
Combination	-27,9%	-4,4%	Original Linear Exponential	+2,5% -27,3% -44,8%	+17,6% +2,3% +10,1%
Replenishment Lead Time	-20,4%	+15,5%	Original Linear Exponential	+9,8% -31,5% -55,0%	+36,8% +17,5% +33,2%

- For current practices at Thales, where they adjust target availability artificially to achieve penalty reductions, expected penalty costs are minimized by using the Combination method. Using this method also leads to savings in the life-cycle costs of the system.
- For the original penalty structure we were not able to achieve significant benefits in expected penalty costs, because this penalty structure punishes minor deviations from the target availability relatively hard. This does not lead to savings in the life-cycle costs. However, linear or exponential penalty structures lead to considerable reductions in penalty costs for the corresponding methods. The Waiting times method is the only method that is also able to reduce life-cycle costs for the linear penalty structure. This is mainly due to the fact that the additional investment in initial spare parts is relatively low for this method compared to the Combination or Replenishment lead time method. Furthermore, penalty costs are so low for these structures such that the initial investment weighs heavily on the total sum of life-cycle costs.

Based on these conclusions, we would recommend Thales to follow the guidelines mentioned in the Waiting times, Combination or Replenishment Lead Time method, depending on the scenario and the system. For this given example, the Waiting Times method performed best regarding the original penalty structure and life-cycle costs for linear penalty structures. However, if Thales chooses to continue their current practice, it is wise to choose for the Combination method. Finally, the Replenishment Lead Time method is able to reduce the penalty costs significantly in the cases of linear and exponential structures, but you need to take into account that the initial investment of the spare parts allocation can be considerably higher compared to INVENTRI.

Another important recommendation is to negotiate exponential or linear step-wise penalty structures to reduce penalty costs even more. In general, it would be wise to include the Customer Services & Support Department in the decision-making of the penalty structures, since this has a large influence on the expected penalty costs.

Furthermore, we advise Thales to spend resources on shortening replenishment lead times rather than on reducing failure rates, since our regression analysis shows that these lead times have more influence on reducing variance and thus expected penalty costs in these PBCs.

Further experiments show that increasing the length of the review period has a positive effect on the variability of the interval availabilities, meaning Thales should aim for review periods with a length as long as possible. Another opportunity lies in the design of the contract structure, because individual base contracts may lead to lower expected penalty costs compared to a joint contract over all the bases, depending on the penalty structure and the amount of slow and fast movers in the system. Simulation is required to find the optimal structure, since this relation is not straightforward.

Table of Contents

List of figures

List of tables

1 Business description of Thales Nederland

In the first chapter we will provide a brief description of Thales Nederland, consisting of a history of the company in Section 1.1 followed by Thales Group & Thales Nederland nowadays in Section 1.2. Since this project will mainly play a role in the service logistics of the naval services, we will discuss this business unit in more detail in Section 1.3.

1.1 History of Thales Nederland

The beginning of Thales Nederland started in 1922 under the name of NV Hazemeyer's Fabriek van Signaalapparaten in the city of Hengelo. Their first product was fire control equipment for two new ships of the Royal Netherlands Navy and after that the company grew rapidly and welcomed other customers in Europe. However, during wartime in 1940 the factory was captured by the invading German Army. Several staff members escaped to the United Kingdom and they were able to continue their work on fire control and radar systems. When they returned to Hengelo, the factory was pillaged and deserted.

After the war, the Dutch government saw the opportunity to buy the factory and the company was able to continue under the name N.V. Hollandsche Signaalapparaten. Due to this governmental involvement, new buildings, facilities and staff made sure that the development of air traffic control equipment, radar technology and fire control for the army took place in a high pace.

In 1956, Philips became a large shareholder and the company flourished and opened plants in several cities across the country. Near the end of the eighties Hollandsche Signaalapparaten employed over 5.000 people serving customers in over 35 countries.

Another setback in the history of Thales was the end of the Cold War, where the political theatre changed dramatically, leading to staff reductions and cuts in defence budgets. The company was taken over by Thomson-CSF (now Thales) and through this merger and reorganization, the design of new systems were realized which ensured that the company was able to take a leap in defence equipment and combat management. Eventually, in December 2000, Thomson-CSF changed its name to Thales. Being a member of this group, Thomson-CSF Signaal changed its name to Thales Nederland. (Thales, 2014)

1.2 Thales Group & Thales Nederland nowadays

Thales Group is a world leader for mission-critical information systems. Their ultimate purpose is to protect people, property and information. They serve five key sectors: Aerospace, Space, Ground Transportation, Defence and Security. In 2013, their sales exceeded 14 billion euros with a workforce of approximately 65.000 people. They invested a total of 673 million euros in their department of Research and Development, an astonishing 67% of their earnings before interest and taxes in 2013. (Thales, 2014)

Thales Nederland is the largest defence company in the Netherlands, employing about 2.000 people. Their activities are carried out at five locations: Hengelo, Huizen, Delft, Eindhoven and Enschede. Thales Nederland operates in professional electronics for defence and security applications, but they are also involved in public transport systems. In 2010 Thales Nederland generated about 600 million euros worth of sales, 80% of which is export. (Thales, 2014)

1.3 Customer Services & Support at the Naval Services

The department of Customer Services & Support (CSS) describes their mission and vision as:

> *"Customer Services & Support is the partner of choice for Life Cycle Services of complex technical systems in the defense environment."* (Thales Nederland B.V., 2014)

They establish this intimate customer relationship through delivering innovative services at any place and any time and continuously improving and updating their service product portfolio to match their customers' requirements. Since products have a typical service life of ±25 years, different service packages are defined in order to let the customer select what suits his needs.

The replenishment of spares, overhauls and Integrated Logistic Support (ILS) are the main activities for CSS, since they account for 72,5% of the order intake between 2007 and 2012. With replenishments they aim to provide swift solutions to gaps in the supply chain with an up to date catalogue of spare parts on a regular basis. Overhauls include extended preventive maintenance tasks to overcome the effects of adverse operational environments and to prevent serious malfunctioning at inconvenient moments. Additionally, they could lead to service life extensions. ILS enhances more features during the life cycle of the product, such as logistics engineering where they influence the system design, documentation, operator and maintainer training and transfer of knowledge and technology to achieve a system with an optimal availability and total cost of ownership. (Thales Nederland B.V., 2014)

2 Research design

Section 2.1 will provide more insight into the context of the research, taking the relevant product, characteristics of the service contracts and the current determination of the spare parts allocation into account. This information will serve as a basis for the research scope in 2.2 and the research objectives in section 2.3. The methodology and the outline of the thesis will be addressed in section 2.4.

2.1 Research context

2.1.1 Product structure and service chain

The radar systems consist of many modules and elements (Thales Group, 2014). This system modularity can be defined into different system parts (Thales Nederland B.V., 2010):

- **Maintenance Significant Items (MSI)**
  All replacement items and higher assemblies, potentially requiring preventive or corrective maintenance. MSIs are subject to subsequent analysis (LRUs are a subset of the MSIs).
- **Line Replaceable Unit (LRU)**
  Replaceable unit to repair the system at organizational level of maintenance.
- **Shop Replaceable Unit (SRU)**
  Replaceable unit to repair LRU at intermediate or depot level, or replaceable unit at System/Equipment site with depot means.
- **Software Support Significant Item (SSSI)**
  Software component on which maintenance tasks can be conducted.

Important to notice is that LRUs are repaired by replacement and they consist of one or more SRUs. The SRUs consist of parts that can cause a failure of the SRU, leading to a failure of the LRU, leading to failure of a (sub)module and ultimately leading to system failure and downtime of the whole system (depending on criticality). These definitions are illustrated in the multi-indenture structure of a radar system as can be seen in Figure 1 below.

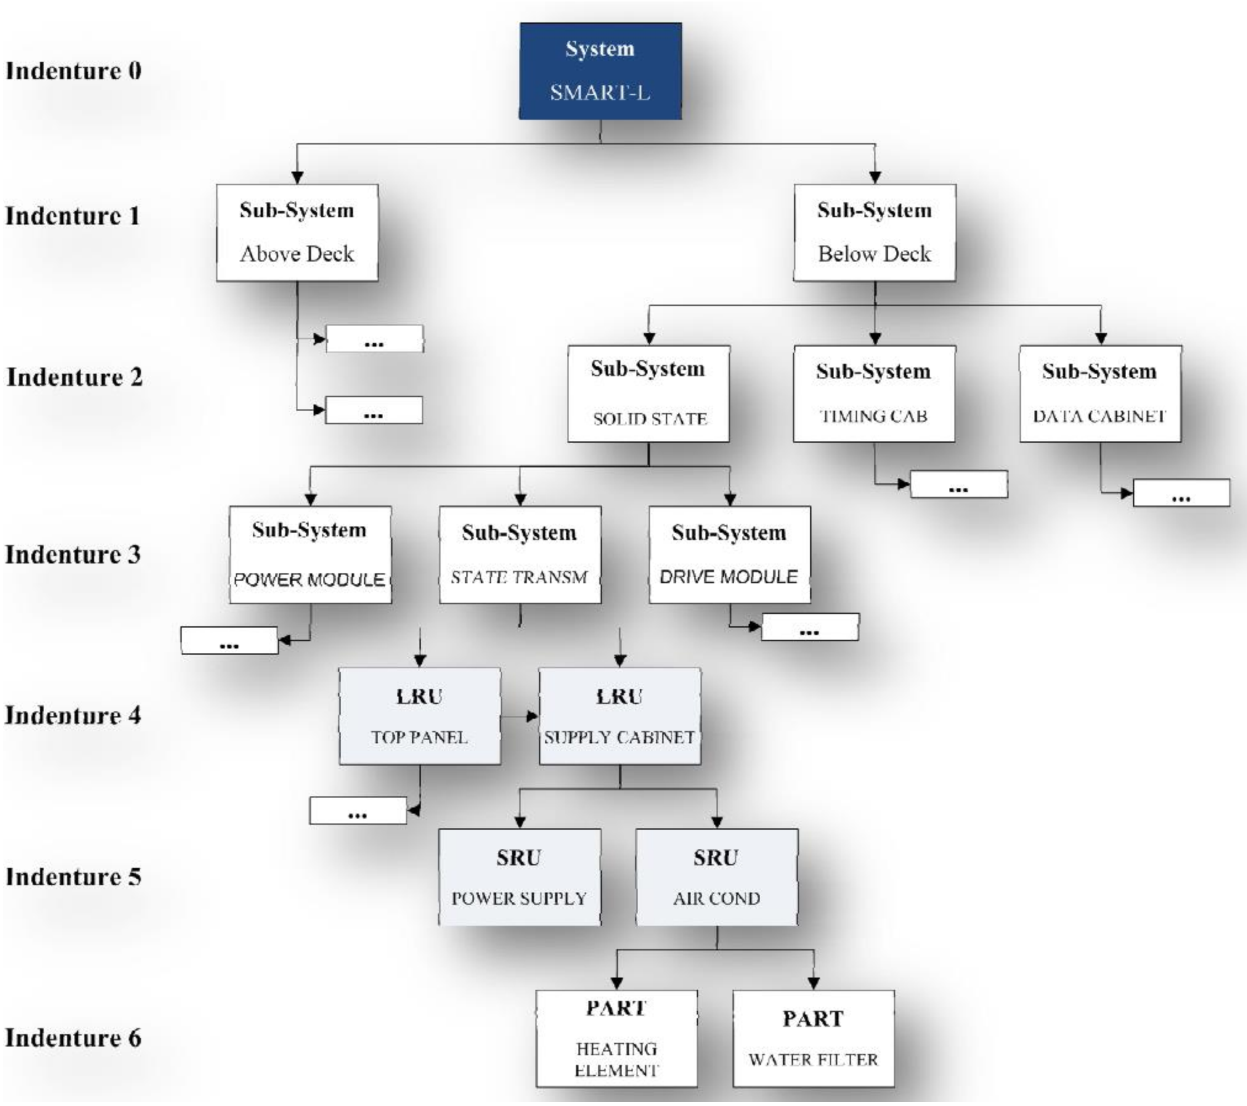

Figure 1: Partial multi-indenture structure of a radar system (van Zwam, 2010)

Thales has organized a repair supply network to ensure high availability of these complex multi-indenture systems. Parts are repaired and replaced based on corrective maintenance. Thales defines corrective maintenance as:

> *"Maintenance on the system to remove an unwanted fault within the system, usually with the purpose of restoring a system capability. Corrective maintenance can be scheduled (if repair can be postponed), or unscheduled if the fault is to be repaired as soon as possible."* (Thales Nederland B.V., 2010)

The operational repair and supply process starts with the demand for an LRU at a location due to failure of that LRU. The corresponding spare part is used to replace the failed LRU or the demand is backordered. Depending on the complexity of the failure the failed SRU within the LRU is repaired at the base (ship), depot or higher within the echelon network to the OEM. Generally, the more complex the failure, the higher the required echelon level to perform the repair. If the selected repair facility has the spare part on hand, they will send it back to the base and repair the failed item. Otherwise, this demand is also backordered and satisfied when this part will become available from the repair shop. This can also be described as a one-for-one replenishment policy (Sherbrooke, 2004). This repair network is illustrated below in Figure 2.

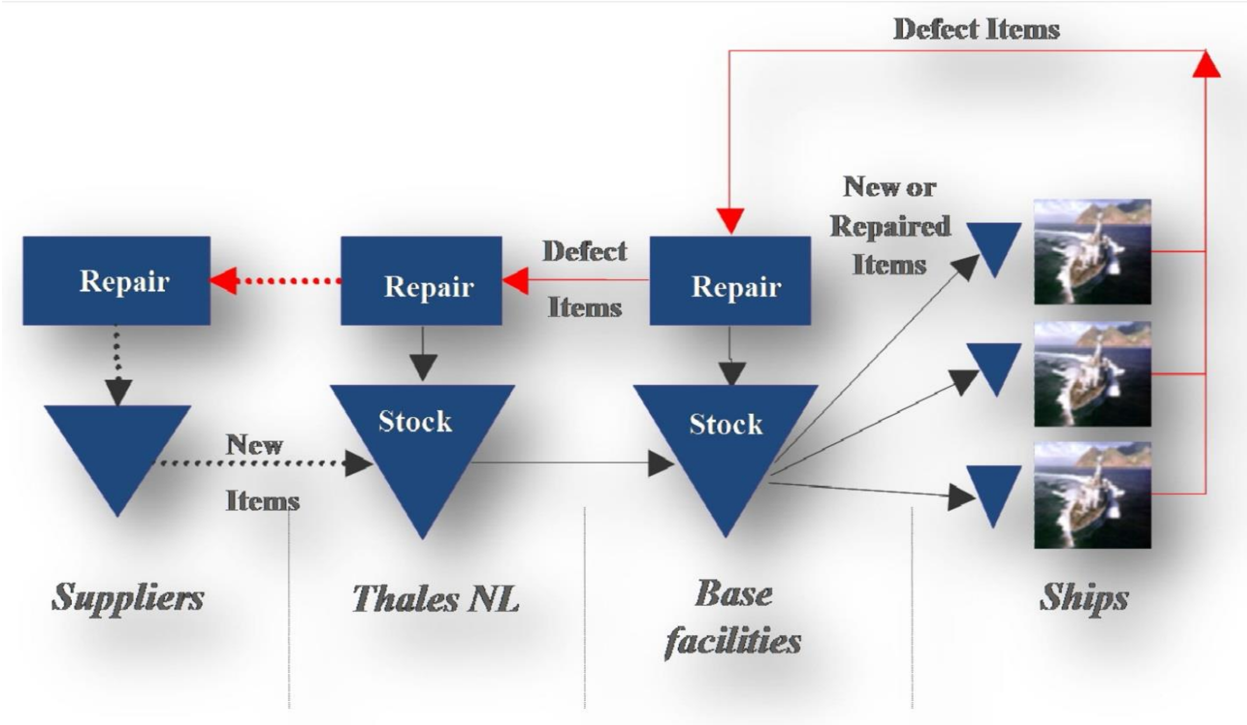

Figure 2: Thales' repair network (van Zwam, 2010)

Given this repair network and product structure, Thales uses a Level of Repair Analysis (LORA) (Basten, 2014) to determine to repair or discard components upon failure, where to do this and where to locate the resources.

2.1.2 Performance-based contracts

Within Thales, performance-based contracts (PBC) are a growing trend towards closer working relationships between Thales and their customers (van Zwam, 2010). These contracts are based on the average availability of the system, leading to penalties when the operational availability is lower than the service level agreement and occasionally resulting in bonuses if a higher operational availability is achieved (Coenen, 2009) (Driessen, 2014). The operational availability is defined by (Sherbrooke, 2004) as:

$$Operational\ availability = \frac{MTBM}{MTBM+MDT} * 100\% \tag{2.1}$$

where MTBM stands for Mean Time Between Maintenance. This is the mean time between two consecutive activities (either preventive or corrective), whereas MDT stands for Mean Downtime due to maintenance time (either preventive or corrective), delay due to lack of necessary spare parts and possible other delays resulting from the maintenance action. In this research we fill focus mainly on the delay due to a lack of necessary spare parts. That means in terms of (Sherbrooke, 2004), we will focus on the supply availability, where MSD stands for the mean supply delay of the spare parts:

$$Supply\ availability = \frac{MTBM}{MTBM+MSD} * 100\% \tag{2.2}$$

2.1.3 Spare parts allocation

To achieve the service level agreement as stated in the PBC, spare parts are needed to reduce the Mean Downtime due to spares and thus increasing the operational availability. Taking into account the aforementioned product structure, an appropriate method to determine this spare parts allocation is the VARI-METRIC approach (Sherbrooke, 2004). This procedure determines how much should be stocked of each item, taking into account failure rates and repair times, at a certain location in the echelon network based on the minimization of the backorders at the bases, since these are causing direct downtime of the system. This approach is implemented and used by Thales in the computer program INVENTRI, a specialized tool developed by Districon and Ortec (Smit, 2009). Calculations in this program are based on (Rustenburg, 2000) and (Sherbrooke, 2004) which enables Thales to draw a graph for the availability of the system with respect to the investment on spare parts in order to select the most appropriate spare parts allocation given a certain target availability. For more details on this approach, we refer to Appendix A: (VARI)-METRIC Approach.

2.1.4 Interval availability

At the moment, INVENTRI is capable of calculating the optimal allocation of spare parts given target supply availability. On average, this availability will be achieved with the selected allocation, but the performance-based contracts are generally reviewed on a yearly basis (ter Hofstede, 2014). That means that the operational availability of the previous year is compared with the service level agreement. Obviously, there will be years with a higher availability than the target, but also years with a lower availability, leading to penalties. This problem arises because VARI-METRIC does not focus on this variance in availability, but on the average steady-state availability during an infinite timespan. To cope with this situation, (Al Hanbali & van der Heijden, 2013) proposed a model to compute the interval availability defined as:

> *"The system interval availability is defined as the fraction of time a system is operational during a period of time [0,T]."* (Al Hanbali & van der Heijden, 2013)

Focusing on this interval availability during the spare parts allocation, rather than on steady-state availability for an infinite timespan can yield more accurate results during the yearly revision as considered in the performance-based contracts.

At the moment, there already exists a professional simulation tool called SIMLOX which computes the interval availability given a certain spare parts allocation as input. This means that the output of INVENTRI can serve as input for SIMLOX to compute the interval availability. SIMLOX allows the user to see how the performance of the technical systems is affected by different system design alternatives.

One can examine how (interval) availability, resource utilization and system utilization varies over time for a given solution and a certain scenario (Systecon, 2013).

2.1.5 Previous research at Thales

The research of (Coenen, 2009) provides a deeper understanding of the benefits, costs and risks involved with the performance-based contracts. By using a simulation program to compute the long term availability with a given stock allocation, he computes the probabilities of having a penalty or bonus during the contract. By using these probabilities and the costs of the corresponding spare parts allocation, he obtains a relationship between the availability of a contract and the financial result. He finds that having an availability as low as possible yields the highest financial result, since stock levels are very expensive and penalty costs do not increase below a percentage of 75% availability, depending on the contract. However, one could argue that customers of Thales do not accept these low availability rates. After performing his sensitivity analysis he also concludes that long term performance can differ greatly from short term performance, that bonus and penalty levels have a high influence on the optimal stock allocation and that the corresponding costs play an important role in the optimal strategy.

The previous work on interval availability at Thales was performed by (van Willigen, 2013). He focused on an approach to reduce the variability in the system availability for a single-site and single-indenture model. He distinguishes three different product groups (slow movers, average movers and fast movers) to determine which would be the best investment with respect to expected interval availability, the coefficient of variation and the 80% survival probability. The distinction between fast and slow movers is solely based on mean repair time and the failure rate of items, however the ratio between these parameters remains constant to guarantee the same steady-state availability. The survival function stands for the complementary cumulative distribution function of the system interval availability. For details on the survival probability function, we refer to (Al Hanbali & van der Heijden, 2013).

The most important results from the research of (van Willigen, 2013) are that it does not matter which item you add to stock if you only take the expected interval availability (steady-state availability) into account, since this is based on the expected number of backorders which is a function of the ratio between repair and breakdown times. As mentioned before, these ratios remain constant during the research. This is also in line with (Sherbrooke, 2004). However, by analyzing the survival function, it is found that adding a fast mover to the stock yields a significantly larger survival probability compared to a slow mover. A major drawback for this result is that adding a fast mover also results in achieving the highest variance of the expected interval availability. The author explains this through the fact that you will need a high variability of the interval availability to compensate the higher survival probability for the same average of the steady-state availability.

However, through intuition and experience in the field of factory physics we might say that slow movers, where we have relatively long repair times and infrequent demand, lead to high variation and inaccurate predictability of failures in comparison with fast movers (Hopp & Spearman, 2000).

2.2 Research scope

We will build upon this previous research by verifying the obtained results concerning the interval availability for slow and fast movers. Also, we will extend the single-site, single-indenture model to a more realistic multi-echelon, multi-indenture model and analyze the effects of the length of the time interval and the structure of the contract. We also discuss model extensions such as order-and-ship times, a repair time distribution, cannibalization and lateral support.

The main focus is to provide assistance in cases where the VARI-METRIC approach is indecisive between different scenarios. In these cases, we will only consider the possibilities of different stock allocations with same steady-state availability meaning we can alter which items to add to stock and decide upon the locations for these items. We will not consider any possibilities in redesigning or remanufacturing parts or systems.

2.3 Research objective

To assist Thales in the decision making for their spare parts allocation given the principles of the performance-based contracts, the interval availability and the research scope, the following main research question will be answered:

> *"How can Thales improve their spare parts allocation to achieve higher interval availability in order to reduce the expected penalty costs while maintaining the requirements in the performance-based contracts and the constraint of the spare parts costs budget?"*

To answer this research question, the following sub-questions will be addressed during the project:

- *SQ 1: How can the current situation at Thales be described?*
  An analysis about the current situation with respect to the performance-based contracts and spare parts management will be conducted in order to obtain a realistic setting for the simulation study. Scientific literature will assist in describing models.
- *SQ 2: Which methods are currently available in the literature to deal with interval availability?*
  A literature review about interval availability will be performed to identify methods that deal with this topic.
- *SQ 3: Which input parameters/attributes of spare parts allocations are relevant for obtaining the same steady-state availability of a system, but achieving different interval availabilities?*
  A deep understanding of the VARI-METRIC approach, the opinion of experts and analyzing other existing heuristics on spare parts inventory management will assist in finding which key attributes contribute to the outcome of interval availability while maintaining the same steady-state availability.
- *SQ 4: In what way do the given attributes influence the interval availability?*
  Here, we will determine the real impact of these key attributes on the interval availability by using the professional simulation tool SIMLOX.

- *SQ 5: How can we design guidelines based on the previous principles to achieve higher interval availability and minimize penalties, taking into account the requirements in the performance-based contracts and on the spare parts costs budget?*

  The previous sub-questions will serve as a basis for the principles of the heuristic aiming to improve the spare parts allocation from INVENTRI in a structured manner.

2.4 Methodology and thesis outline

Chapter 3 consists of the literature review, where in the first two sections 3.1 and 3.2 the current situation at Thales will be described by using after sales business models and spare parts strategies. In the succeeding sections 3.3 and 3.4 of this chapter the second and third sub question will be answered by analyzing studies on interval and process availability. The conclusions of the literature study will be translated into hypotheses for the simulation in section 3.5.

Then, in Chapter 4, we will perform the simulation study based on the conclusions of the literature review to answer sub question 4. Section 4.1 involves the model formulation and section 4.2 includes the characteristics of the simulation. The model extensions will be addressed in section 4.3. We will discuss the results from this study in Chapter 5 so that we can answer the fifth sub question. We start with the basic simulation results in section 5.1 followed by the simulation of the model extensions in section 5.2. A sensitivity analysis is performed in section 5.3 and we will eventually provide the guidelines for improvements in chapter 6. These guidelines will be tested in a real-life example and these results are presented in chapter 7. We will finish the thesis with a concluding chapter 8 and a chapter 9 about suggestions for further research.

3 Literature review

To learn more about the issues addressed in sub questions 1,2 and 3, we will perform a literature review about these aspects to gain insights. In this section we will briefly focus on the literature about Long Term Service Agreements (LTSA) when discussing after sales business models in section 3.1. Then we will relate spare parts strategies to these LTSAs in section 3.2. In section 3.3 we will extensively address the issues related to the interval availability and in section 3.4 we use the findings on process variability. Finally, in 3.5 we conclude on our findings in the literature review to derive hypotheses for the simulation study.

3.1 After sales business models

Certain products have an extremely long operation phase of ±25 years (Thales Nederland B.V., 2010). The service life of these products can even be extended if clients choose for an overhaul by refurbishing the system. This way, the system is able to enter the operation phase again for another ±15 years (Doets, 2014). Studies have shown that the maintenance costs of such products way exceed the initial costs (Alfredsson K. , 2001). Therefore, it is more appropriate for businesses to analyze the life-cycle costs (Gluch & Baumann, 2004) (Woodward, 1997).

For Thales, this means that their after sales processes will become more important, since they execute the maintenance, repair, supply and support services for these systems. Because of this shift in cost focus, performance-based contracts arose at Thales for handling the after sales processes (Doets, 2014).

In the model of (Cohen, Agrawal, & Agrawal, 2006), they distinguish several models based on service priority and ownership, also see Appendix B: After sales business models. Combining their models with the situation at Thales we can sketch the following figure concerning the Long Term Service Agreements as in Figure 3 below.

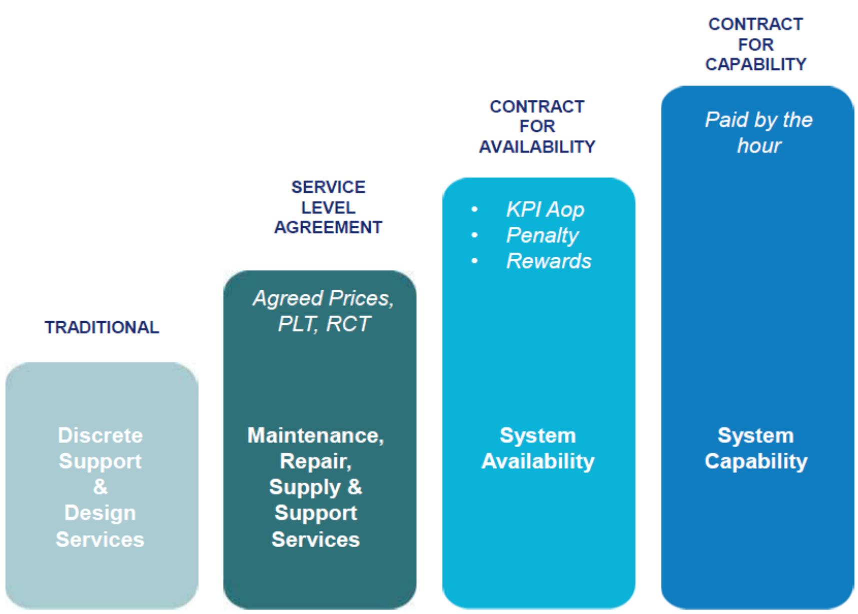

Figure 3: Different Long Term Service Agreements for after sales processes (Doets, 2014)

Here we can quickly see that the traditional view on contracts has shifted rigorously. The system availability is a key measure nowadays and the penalties and bonuses are based on this value, the so called penalty structure. Therefore, contracts for availability and capability are becoming a trend, certainly for Thales. In an effort to reduce costs and increase efficiency, many navies turn to Thales for maintenance and support for their systems after the regular guarantee period, rather than performing these tasks by themselves (Thales Nederland B.V., 2012). Also, customers often do not possess the knowledge, resources or budgets to perform this maintenance and therefore they outsource this to Thales (Thales Nederland B.V., 2014). In the Contract for Availability (CFA), the customer still has ownership of the system, but Thales performs all agreed In-Service Support (ISS) against a fixed annual fee. In the Contract for Capability (CFC), Thales remains the owner of the system, so that you do not have to invest the capital for the acquisition of the initial system (Thales Nederland B.V., 2014).

3.2 Spare parts strategies

Techniques in spare parts can be distinguished on two dimensions according to (Rustenburg, 2005), namely maintenance type and costs. This leads to four different strategies as can be seen in Figure 4 below.

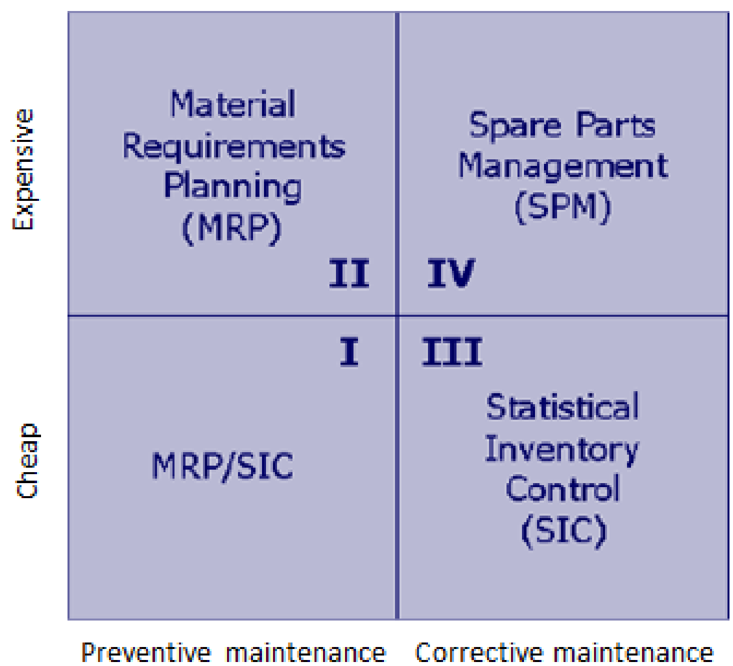

Figure 4: Spare part strategies (Rustenburg, 2005)

Within this research we will mainly focus on the fourth class where we have expensive products which need corrective maintenance. This is the most difficult category, since this involves the largest part of the investment and system critical parts. We need advanced methods to manage these spare parts to cope with the unpredictable and unstable demand and to analyze the influence from the part on system availability. Study has shown that system approaches are way more effective for these spare parts compared to item approaches, since interaction between various levels is included (multi-echelon *and* multi-indenture), discrete demand distributions can be used and extensions as commonality, criticality and redundancy can be included (Rustenburg, 2002). Here, commonality means that the same parts may occur in more than one parent. The criticality of a part reflects the probability that the system is down due to the failure of this part. Furthermore, a system may contain not only active but also passive redundant parts. These passive parts which are on standby can be activated to prevent downtime of a system. A suitable method for this category would be the aforementioned (VARI)-METRIC approach by (Sherbrooke, 2004).The main takeaways from this method are that we prefer to add cheap fast movers to the stock at the individual bases to obtain high backorder reduction in order to achieve high availability quickly and thus reducing downtime of the system per individual base, whereas on the other hand we might add items to stock at the depot in order to decrease backorders at this depot and therefore decrease the pipeline and waiting time for items for all bases simultaneously. The approach considers this trade-off and aims to maximize the overall operational availability. To cope with the potentially invalid assumption that the base pipeline would be Poisson distributed, the author extended the existing model to the VARI-METRIC approach to prevent underestimation of the backorders at the base. For more details on both methods, we refer to Appendix A: (VARI)-METRIC Approach.

3.3 Interval availability

The first expressions for the moments of cumulative operational time were found by (Takács, 1957), and later this work was extended by (Iyer, Donatielle, & Heidelberger, 1986). With the aid of the Laplace-Stieltjes transform and the central limit theorem both studies were able to determine the asymptotic performability distribution. Performability can be described as a composite measure for the performance and reliability, in their case it may also be interpreted as the probability density function of the operational time obtained from a system during a finite period of time. Where (Takács, 1957) considered an on-off stochastic process, (Iyer, Donatielle, & Heidelberger, 1986) considered an $M/M/N/N + b$ Markov process, where the N servers represent the N processors and there are b stages in the buffer, where each stage holds one job to be processed. They included repair actions of the buffer and the failing processors.

The limitations of the steady-state availability and the need for the computation of the probability for not meeting certain availability were recognized by (Goyal & Tantawi, 1988). Before this, high reliability was mainly achieved through redundancy, to decrease the amount of downtime (Avizienis, 1978) (Siewiorek & Swarz, 1982). The measure of steady-state availability was insufficient, because there were penalties involved for not meeting the guaranteed level of availability. Goyal & Tantawi therefore calculated the distribution of availability at time t recursively by discretizing the observation period $(0,t)$ into such small intervals Δt that the probability of having two or more failures during this interval is negligible. This numerical approach allowed them to consider Markovian failures and repairs, time-dependent failure and repair rates, and deferred repair and non-deferred repair strategies, with sufficiently small approximation errors. For this approach, they considered continuous time two-state Markov chains. However, they did not provide an analysis section with error bounds on the approximation.

These approximation errors and lack of error bounds were the main motivation of the research of (de Souza e Silva & Gail, 1986). By using the randomization (or uniformization) technique (Cinlar, 1975)

(Ross, 1983) (Berger & Christophi, 2003) they could also numerically calculate the distribution of availability during a finite interval period. An important contribution is that modeling details such as Coxian distributions, spares and coverage may be included. Furthermore, they showed that other performance indicators, such as the density of availability at time *t*, the expected availability at time *t*, the expected lifetime at time *t* and the mean time to failure can all be calculated without any significant additional computational effort. Also, by using their method, one can specify the error tolerances in advance.

In the paper of (Rubino & Sericola, 1993) two improvements are suggested on the approach of (de Souza e Silva & Gail, 1986). They provided an algorithm which needs the same memory space, but works faster. A second algorithm is also introduced which needs less memory space than the previous methods and its space requirements are known at the beginning.

More recent research of (Carrasco, 2004) focuses on designing an efficient and fast approach to solve large and complex interval availability models by using a truncated transformed model and solving this through a state-of-the-art algorithm. Later on, (Huang & Mi, 2013) were able to proof initial monotonicity of the interval availability and derived lower bounds and the limit for the interval and average availability. Since closed forms for the interval availability were still generally unavailable they introduced a Block-by-Block method (Linz, 1985) to compute the interval availability numerically for a specified time instant *t*, showing high accuracy and efficiency. The study of (Kirmani & Hood, 2008) has led to a closed form expression for the variance of the interval availability in a 2-state continuous time Markov chain (2-CTMC) model. They also propose to derive the coefficient of variation of the interval availability to judge its fluctuations. Finally, an approach to compute the first two moments of interval availability in closed form for a Markov chain that models the system with more than two states is proposed in (Al Hanbali & van der Heijden, 2013). These two characteristics are used to approximate the survival function of the interval availability, i.e. the complementary cumulative distribution function of the interval availability, using a Beta distribution (Smith, 1997) together with the probability that the interval availability equals one, see Appendix C: Interval availability and survival function. This function shows excellent accuracy in comparison with a simulation study, especially for high expected interval availability.

3.4 Process variability

The findings of (Hopp & Spearman, 2000) and (Zijm, 2012) in the production environment can lead to insights regarding interval availability. They assume that the times to failure are exponentially distributed and they do not make any particular assumptions about the repair times, so deterministic repair times are allowed. However, they consider a production line with different workstations and come up with the variation of the process time. Furthermore, they consider a preemptive service policy, meaning whenever a failure occurs during service time, it will be resumed after repair of the machine. We do not have a production line in our environment at Thales, but we can still use the formula to understand the effects on the variability in the system, by using:

$$c_e{}^2 = c_0{}^2 + (1 + c_r{}^2)A(1 - A)\frac{MTTR}{t_0} \tag{3.1}$$

Where ${c_e}^2$ stands for the squared coefficient of variation of the effective process time, ${c_0}^2$ includes the natural (unaccounted for) variability in the process, ${c_r}^2$ includes variability in repair time, A is the availability of the system, $MTTR$ is the mean time to repair and let t_0 be the natural process time.

In our setting considering slow movers versus fast movers, the natural variability ${c_0}^2$ will remain equal, because we use the same random seed for the generation of pseudo random numbers. Note that the squared coefficient of variation of the repair times ${c_r}^2$ is determined through the standard deviation and mean of the repair time:

$$c_r = \frac{\sigma_r}{MTTR} \tag{3.2}$$

The standard deviation σ_r in the repair time is equal to zero, since we consider deterministic repair times and this means that the coefficient of variation of the repair times will also be equal to zero.

The major importance in this formula is that for larger $MTTR$, we have a larger variance in the system for the same steady-state availability. The failure rate does not have an effect on the variance in the system for this setting. This means that slow movers with high $MTTR$ cause more variance in the system and thus to more variance in interval availability leading to an increase of the probability of receiving a more severe penalty.

This formula also suggests that the steady-state availability has an impact on the variance, because of the term $A(1-A)$. The outcome will be the largest for A=50% and effects will fade out as system availability tends to increase to 100%.

One could argue that the formula and its applications do not hold in our environment, since we do not have workstations and a preemptive service policy. Therefore, we could also use the findings of (Kirmani & Hood, 2008) on the variance of interval availability and its common coefficient of variation, where two-state (on and off) Markov chain models represent a more realistic scenario of our environment at Thales. To examine the effects of slow and fast movers in the setting of (Kirmani & Hood, 2008) we experiment with the parameters λ (mean time to failure = 1/ λ) and μ (mean time to repair = 1/μ), while keeping system availability constant. This means that we maintain the same steady-state system availability, but we alter the mean time to failure and the repair times in this case to move from a slow mover to a fast mover. Steady-state availability in this case is calculated as:

$$Steady-state\ availability = \frac{\mu}{\lambda+\mu} \tag{3.3}$$

For increasing μ (faster repair) and thus increasing λ (higher demand rate), we can see that the coefficient of variation $\phi(T)$ tends to decrease, meaning that if we have long and infrequent repair times (slow movers), we have higher fluctuation in interval availability, compared to short and frequent repair times (fast movers). This is also illustrated in Figure 5 below. This conclusion is in line with the previous findings of (Hopp & Spearman, 2000). We can also see that the coefficient of variation decreases as the review period increases. For more details on the calculations in the model we refer to Appendix D: Variance of interval availability and its common coefficient of variation. The curves for the coefficient of variation are increasing for small T due to the authors' initial assumption at t=0. They assume that the Markov chain is always in on-state (system is functioning) at the start, resulting in zero variance.

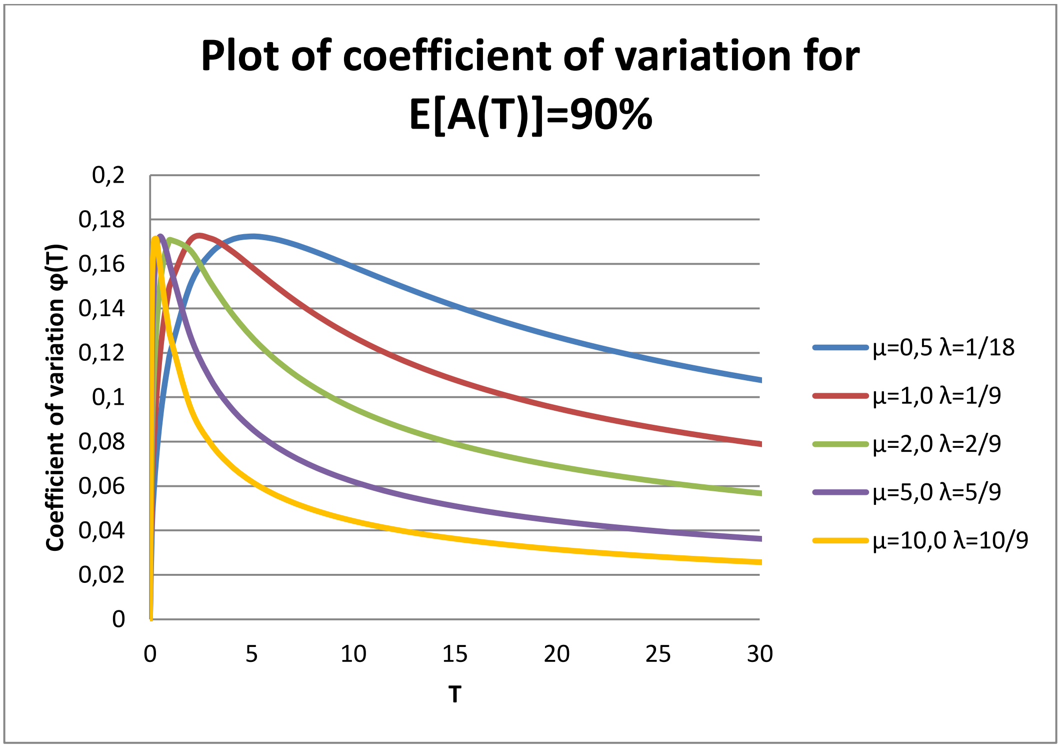

Figure 5: Coefficient of variation for different μ and λ, based on (Kirmani & Hood, 2008), as function of the review period

3.5 Conclusions on literature review

All previous studies on interval availability mainly focus on computations for the moments of the interval availability for systems with different complexities using analytical methods. This means that up to now, there are no managerial insights available on how the effect of the interval availability can be minimized and which attributes in a spare part allocation have a contribution in this matter and in what way. The article of (Al Hanbali & van der Heijden, 2013) found that for two different spare part allocations (adding centralized stock compared to decentralized), with (nearly) the same steady-state availability, you can have different interval availability, but they do not extensively analyze where this difference comes from or conclude that it is always better to add stock centrally with regard to interval availability. Sure, (van Willigen, 2013) found that the failure rates and repair times of items can have an influence on interval availability without affecting the steady-state availability (by keeping the product constant), but this was done for a very simplistic scenario with relatively low system availabilities resulting in counterintuitive findings. Therefore, we will extensively address the effects of these two characteristics for constant products in a more realistic and more complex model in order to understand the real effects of these attributes, to come up with the managerial advices on the situations where the VARI-METRIC procedure encounters indecisive scenarios concerning the spare parts allocation, since this is lacking in the literature.

Assistance in these indecisive scenarios is provided by using the findings in the literature (Al Hanbali & van der Heijden, 2013) (Kirmani & Hood, 2008) (Hopp & Spearman, 2000) (Sherbrooke, 2004). Their conclusions will be translated to hypotheses:

1. *"Stock investments for slow moving items lead to lower penalty costs in the Performance-Based Contracts compared to stock investments in fast moving items."*
2. *"Central stock investments lead to lower penalty costs in the Performance-Based Contracts compared to local stock investments."*

These hypotheses will be analyzed in a simulation study, because there is a professional simulation tool already available which is able to incorporate the characteristics of the situation at Thales.

4 Simulation model

In order to improve the current spare parts allocation at Thales, taking into account the penalty costs involved, we perform experiments based on the aforementioned hypotheses. In this chapter we will describe the basic model characteristics of the simulation study in section 4.1. Subsequently, in section 4.2, details about the simulation characteristics will be provided. The model extensions are discussed in 4.3.

4.1 Basic model formulation

To formulate the basic model, we include the following perspectives: the system view, the supply network view and the usage profile. We will also include the penalty structure for the Performance-Based Contracts.

4.1.1 System view

The basic model is a radar system which consists of two LRUs with each corresponding failure rates and deterministic repair times. Both LRUs have a quantity of one in the system and if one LRU fails, the full system is down. This is illustrated in Figure 6 below.

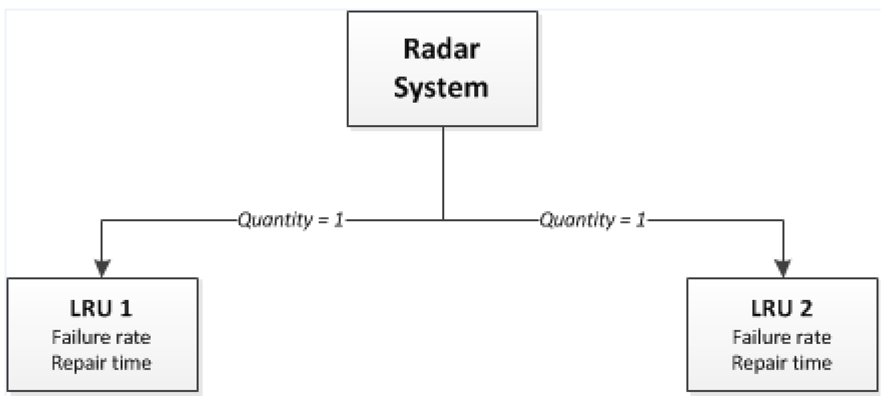

Figure 6: System structure of basic model

4.1.2 Supply network view

We consider one depot, which is connected to two bases, where each base exactly has one radar system. For the basic model, the order and ship times between the bases and the depot are set to zero. This is adjusted in later experiments and some extensions. There are no stocks available for both items at any location. This is illustrated in Figure 7 below.

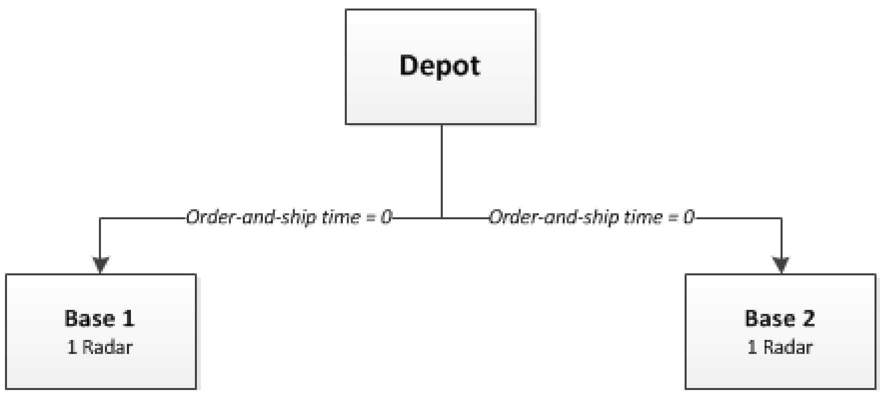

Figure 7: Supply network of basic model

4.1.3 Usage profile

We will consider a mission profile for both bases where missions with duration of exactly one month (730 hours) will be iterated 12 times during a year. This will lead to a continuous mission profile with system utilization of 100%.

4.1.4 Penalty structure

It is very difficult to set a standard template for a penalty structure in a Performance-Based Contract. This depends on several factors, such as trust between Thales and its customer, the corresponding system, the measurement of the key performance indicators and the risks that are involved with possible downtime. Also, some customers only consider penalties when the actual system availability is below the target and are not willing to pay a bonus for higher availability since they did not request this (Driessen, 2014). Therefore, we will only consider the minimization of penalties and do not take into account possible bonuses based on the following common structure as illustrated in Table 1, where target operational availability is set to 90%. Although this exact penalty structure was never agreed upon in a contract with a customer, it was a suggested structure (Ypma, 2014). It is remarkable that small deviations from the target availability are punished severely, whereas very low availabilities are not penalized that hard. Contract value is set to 50 million with duration of 15 years. The contract value per year is 7% of the total contract value, leading to 3.500.000 per year.

Table 1: Original penalty structure for 90% Performance-Based Contract

System availability per year (in %)	Penalty percentage of yearly contract value
0 – 25 %	110 %
25 – 50 %	100 %
50 – 70 %	75 %
70 – 85 %	50 %
85 – 86 %	25 %
86 – 87 %	20 %
87 – 88 %	15 %
88 – 89 %	10 %
89 – 90 %	5 %
90 – 100 %	0 %

4.2 Simulation characteristics

In this section, we will discuss the relevant experiments for the simulation, together with its definitions, assumptions and restrictions. Thereafter, we will provide some extensions of the basic simulation model to obtain a more realistic setting.

4.2.1 Definitions

Before the design of our experiments, we first introduce some important definitions to establish insights on slow and fast movers. These different items will have the same steady-state availability through the fact that the backorder product remains the same, but they have different characteristics in terms of mover ratio.

Replenishment lead time

The replenishment lead time of an item is the complete time it takes to restore the system to its state before failure of that item. This time depends on the repair time (MTTR) and repair fraction (r_f) of the corresponding item, the reorder time at the supplier when repair is not successful and possible order and ship times (OST) between the depot and the base.

$$Replenishment\ lead\ time = OST + [r_f * MTTR + (1 - r_f) * Reorder\ time] \tag{4.1}$$

Note that the order and ship times are always included, because repairs are always performed at the depot. When repair is unsuccessful, the item is reordered at its supplier. For every item, it holds that this reorder time is always larger than its mean time to repair.

Mover ratio

The mover ratio defined as the failure rate of the item divided by the replenishment lead time of the item. The failure rate of the item is given as the amount of failures that will occur per million of operating hours. The replenishment lead time is the time it takes, in hours, to replenish the item and to restore the system to its original state given.

$$Mover\ Ratio(MR) = \frac{Failure\ Rate}{Replenishment\ Lead\ Time} \tag{4.2}$$

Important to notice is that this parameter solely serves as an indicator to make a distinction between fast and slow moving items. It is not used in formulas or calculations. With this parameter we can see whether there exist different results when comparing fast moving items to slow moving items. Note that the mover ratio will be relatively low for slow movers, since they tend to have a low failure rate and long replenishment lead time. Fast movers will have higher failure rates, but shorter replenishment lead times leading to a higher mover ratio. We will consider mover ratios from 0,001 (slow mover) up to 10 (fast mover), because replenishment lead times longer than 2 years and shorter than 3 days do basically not exist for items and are also not included in INVENTRI (Doets, 2014). A logarithmic scale on these mover ratios will be used to illustrate the effects.

Backorder product

Note that we cannot change the failure rate and replenishment lead time without taking any restrictions into consideration. We need the same backorder product for the different items to guarantee the same steady-state availability (Sherbrooke, 2004) (Rustenburg, van Houtum, & Zijm, 2000) (Muckstadt, 2005), as can be seen through the formula of the calculation of the backorders at the base:

$$Expected\ Backorders\ at\ the\ base = EBO_{ij}(s_{ij}) = \sum_{n=s_{ij}+1}^{\infty}(n - s_{ij}) * \frac{\mu_{ij}{}^{n}e^{-\mu_{ij}}}{n!}, \quad (4.3)$$

Where s_{ij} denotes the stock level of item i at location j and μ_{ij} is the backorder product defined as:

$$Backorder\ Product(\mu_{ij}) = Failure\ Rate_i * Replenishment\ Lead\ Time_{ij} \quad (4.4)$$

Review period

The time interval during which the availability of a system is reviewed will be called the review period. Note that the system availability for each review period is defined as the interval availability. Also, the interval availability is reviewed as a total over all bases and not on an individual basis.

4.2.2 Design of experiments

In this section we will address the different experiments to determine the influence of the given attributes. Note that all scenarios are performed on the aforementioned basic model, unless we state otherwise (such as adding stock, include order and ship times or distinguish different review periods).

- **Experiment 1: Slow movers vs. fast movers on different availability levels (section 5.1.1)**
  We will first analyze the impact of slow movers versus fast movers on different availability levels. This experiment will be performed to see what influence the mover ratio has on the interval availability and we will see if this effect changes for different steady-state availabilities.
- **Experiment 2: Adding stock to slow movers vs. fast movers on depot and base level (section 5.1.2)**
  Recall that we will not consider decisions on redesign or remanufacturing, but we do have an influence on selecting which item to add to the stock allocation. We will analyze the effects on interval availability when adding slow or fast movers to the stock allocation. This will be done for both at the bases as for the depot, while maintaining 90% steady-state availability even if we add stock. Therefore, we adapt the failure rate and repair time of the items to cope with this.
- **Experiment 3: Adding stock locally compared to adding stock centrally (section 5.1.3)**
  As mentioned before, Thales could also decide upon the location of adding stock. We will analyze the impact of these locations on the interval availability in this experiment. Order and ship times will be included in this model to achieve same steady-state availabilities between adding stock on different locations.
- **Experiment 4: Changing the review periods (section 5.1.4)**
  We will analyze the effects when changing the review period of the system to a relatively short time period compared to long periods. For our experiments we distinguish the review periods as stated in Appendix E: Review periods.
- **Experiment 5: Joint contracts and individual contracts (section 5.1.5)**
  We will analyze the effects of having a joint contract for all bases together compared to having individual contracts per base and investigate whether fast movers and slow movers have an impact here.

4.2.3 Assumptions and restrictions of the basic model

In the basic model we have the following assumptions:

- The radar system consists of two LRUs with both a quantity of one.
- We do not include the possibility of performing preventive maintenance.
- We do not include the possibility of lateral support (transshipments between the bases).
- We do not include the possibility of robbing (cannibalization).
- We do not include the use of alternative units.
- We do not include a backorder prioritization; all backorders are satisfied on a First Come First Serve (FCFS) basis.
- We do not include the possibility of common and/or redundant items.
- All items have the same criticality level, always leading to a failure of the system and thus downtime.
- The order and ship times between the depot and the bases are considered to be deterministic and equal to zero.
- Failures occur according to a Poisson process, the time between failures is a stochastic variable having a negative exponential distribution. The failure rates are constant in time.
- The item repair times are considered to be deterministic and constant in time.
- All items failing at the bases are sent forward to the depot, leading to repair probabilities at the base of zero. Reparations at the depot are always successful, so reordering is not considered.
- There is no stock available of any item at any location.
- Availability is reviewed on a yearly basis (review period is one year).

4.2.4 Run length

The simulation will be performed for 1980 review periods (N=1980). Before this, we will have a warm-up period of 20 years to cope with long replenishment lead times and low failure rates in order to allow queues and other aspects in the simulation to get into conditions that are typical of normal running conditions (Simul8, 2015). After this warm-up period, the 1980 review periods will be simulated. For the basic model, this means that we simulate for a total of 2000 years, since the review period is one year. Although for some experiments we have different review periods, we do not change the warm-up period since the length of the review period does not influence the failure rate or replenishment lead times, so therefore we can maintain the same warm-up period. We only use one replication to obtain realistic results and to prevent running the warm-up period for every replication again. However, to ensure that certain stock allocations and item characteristics lead to the target availability, we use 1000 replications to check whether the input leads to this target.

4.2.5 Key performance indicators

The variance, the coefficient of variation, the survival function of the interval availabilities and the expected penalty costs are key performance indicators in this research. These figures are closely related to the probability of receiving a penalty, where a lower variance, lower coefficient of variation and a higher survival probability for a given system availability will lead to a lower probability of receiving a penalty.

Note that for both formulated hypotheses it must hold that steady-state availability remains the same; otherwise it would not be a valid improvement of the VARI-METRIC solution, since then we do not have an indecisive situation for this procedure. For the first hypothesis we face no serious difficulties with regard to the restriction on steady-state availability for the basic model. There are no order and ship times involved yet, so an equal backorder product will lead to same steady-state availability. However, for the second hypothesis regarding stock locations, we have to be aware that this steady-state availability changes. This is due to the fact that you will always prefer to stock at the depot to obtain the risk pooling effect when order and ship times are ignored. We can overcome this issue by including these order and ship times.

For every review period (T) we will collect the number of systems that were available for that interval. Note that this can be a decimal number, meaning if 3 systems are 50% available of the time there are 1,5 systems available. From there, we can compute the interval availability per review period since we know the total number of systems in the network.

$$Interval\ availability\ A(T) = \frac{Number\ of\ systems\ available_T}{Total\ number\ of\ systems\ in\ the\ network} * 100\%, T = 21,22\ldots2000 \quad (4.5)$$

Then we are interested in the computation of the variance amongst all interval availabilities, since we know that increasing variability generally degrades the performance of a system and variability reduction is central to improving performance (Hopp & Spearman, 2000). Obviously, this formula only holds when there is no significant correlation between the availabilities for the different intervals. Simulation results show that this is indeed the case by calculating and interpreting the Durbin-Watson statistic (Durbin & Watson, 1950), so we can apply the formula.

$$Var[A(T)] = \sum_{T=21}^{2000} \frac{(Interval\ availability_T - Average\ availability)^2}{N-1} \quad (4.6)$$

Furthermore, we calculate the survival function and the coefficient of variation. The survival function for x can be seen as the probability that the interval availability will be greater than or equal to this x with x being a number between 0 and 1.

$$Survival\ function\ availability(x) = \frac{COUNT(A(T) \geq x)}{N}, \quad T = 21,22 \ldots 2000 \quad (4.7)$$

A proper view of the extent of fluctuations in the interval availability [A(T)] is not given by their standard deviations or variance, but by their common coefficient of variation (ϕ) (Kirmani & Hood, 2008), defined as:

$$\phi(T) = \frac{\sqrt{Var[A(T)]}}{E[A(T)]} \quad (4.8)$$

4.3 Extensions

The multi item, multi echelon basic model obviously represents a more realistic setting of the situation at Thales compared to the multi item, single site model from (van Willigen, 2013), but still it is a simplified version of reality. Therefore, we investigate and discuss the effects of some model extensions on the interval availability and again compare the influence from slow and fast movers in these extensions.

4.3.1 Order and ship times

Recall that the basic model does not take any order and ship times into consideration between the bases and the depot. In reality, often there will be a transport time between these stations. In our model extension, we include deterministic order and ship times between the bases and depot, since SIMLOX does not support the possibility to include a distribution to these transport times. Before starting the experiments, we have to be aware of the fact that the steady-state availability might change between the slow movers and fast movers for the same backorder product. This did not happen before, since these transport times were equal to zero. However, now the fast movers will make use of the transport more often compared to the slow movers, so we have to account for this by adjusting the backorder product, by including the order and ship times in the replenishment lead time, to obtain the same steady-state availability.

4.3.2 Repair time distributions

Up till now we assumed that repair times are deterministic. In practice and in recent literature, we often see that these repair times follow an exponential distribution (Al Hanbali & van der Heijden, 2013), Weibull distribution (Huang & Mi, 2013), or are seen as random variables with a certain mean (Caglar, Li, & Simchi-Levi, 2004) (Wong, Kranenburg, van Houtum, & Cattrysse, 2007). Since SIMLOX only supports the possibility of an exponential repair time distribution, we will investigate the effects of this distribution on the interval availability. We can easily implement this without worrying about the steady-state availability, since the study from (Alfredsson & Verrijdt, 1999) has shown that the steady-state availability tends to show little sensitivity to the distributional form of repair and order and ship times.

4.3.3 Cannibalization/robbing and lateral support

The purpose of cannibalization is to recover a limited set of reusable parts from used products or components for product recovery management (Thierry, Salomon, van Nunen, & van Wassenhove, 1995). However, in the case of a multi-echelon network, one could define cannibalization as enhancing the reliability of a system by extracting the needed components from another part of the system when repair facilities or spare parts are not immediately available (Baxter, 1988). For Thales, this will only work when we can exchange non-critical components to the part of the system where they are critical. In this way, the system failure will be restored. However, we only consider critical items in our systems and therefore this form of cannibalization will not solve anything. Nevertheless, we do have the possibility to exchange components between systems if one system is already on failure. Then we can solve a possible failure of the second system by replacing the failed item with the working item of the already failed system. This is also referred to as robbing.

Lateral support means that the failed part at one location can be replaced by a ready-for-use spare part that comes from another location (excluding the depot). Through this additional flexibility one can achieve significant cost savings, especially when lateral transshipment lead times are low and service constraints are tight (Wong, van Houtum, Cattrysse, & van Oudheusden, 2006).

This would mean that both options provide more flexibility in the network, which will always lead to the same or higher system availability for the given spare parts allocation because the original solution is always allowed. Although these options could be included in SIMLOX we do not investigate them in this simulation study.

4.3.4 Multi-indenture model

The multi-indenture model that we consider is a simple extension of the basic model. Both LRUs now contain two SRUs, where failure of an LRU is always caused by failure of one SRU, see Figure 8. The LRU is repaired by replacement at the base if there is an SRU on stock. This replacement is always successful. If there is no SRU on stock, we have to wait for the failed SRU to return. The failed SRU is sent forward to the depot where it will be repaired. This repair is always successful. After repair, it is send back to the base to put it on stock or to be replaced in the failed LRU. In the basic model, we experimented through altering the failure rates and repair times of the LRUs. Now, we also have to deal with the corresponding SRUs. We assume that the failure rate of an SRU is exactly half of the failure rate of its LRU. Furthermore, the complete replenishment time of the LRU consists of the replacement of the SRU and the repair of the SRU. Model and inventory characteristics are the same as in the basic model for single-indenture. For more information on the item characteristics for this multi-indenture model, we refer to Appendix F: Details for simulation of multi-indenture model.

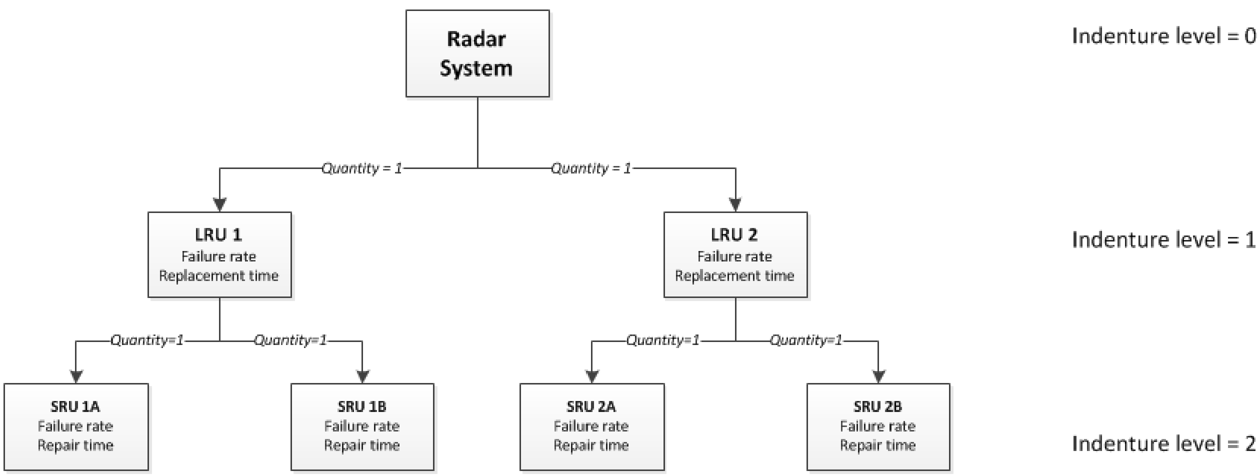

Figure 8: System structure of multi-indenture model

4.4 Conclusions simulation model

To validate our hypotheses we will use a radar system consisting of two items installed at two bases. The impact of the level of the steady-state availability will be investigated. Thereafter, we examine the effects of stocking a slow or fast mover locally and centrally. Lastly, we investigate the influence of the length of the review period and the structure of the contract. We simulate for a period of 2000 years with a warm up period of 20 years to compute the variance among the interval availabilities and calculate the expected penalty costs according to the penalty structure given in Table 1. Finally, we extend the simulation model in several ways to achieve a more realistic setting and to check whether the obtained results are influenced.

5 Results

In this chapter we will discuss the findings from the experiments, relate them to the hypotheses of section 3.5 and use this relation to design guidelines to improve the current VARI-METRIC solution in the spare parts allocation.

5.1 Simulation results

The results of the aforementioned experiments will be discussed in this section. Recall that we first analyze the influence of different levels of availabilities followed by an experiment where we investigate the effect of stock investments for slow and fast moving items at the depot and base. Hereafter, we examine the impact of the length of the review period and finally we discuss the findings for the different contract structures. All definitions of the parameters used in the experiments can be found in section 4.2.1.

5.1.1 Effects of slow movers versus fast movers without stock for various steady-state availabilities

From Figure 9 we can see that for the different steady-state availability levels, the variability in the interval availabilities tends to increase when the steady-state availability decreases. Furthermore, we see that variance of the interval availability decreases as we change from slow mover to a fast mover independent from the current level of availability. This is not in line with the study of (van Willigen, 2013) considering the single-site model, but it is certainly in line with our expectations for the multi-echelon model, since fast movers are more predictable because of frequent failures and shorter replenishment lead times whereas slow movers occur infrequent with long replenishment lead times. One remark is that the effect will decrease when average availability tends to increase to 100%.

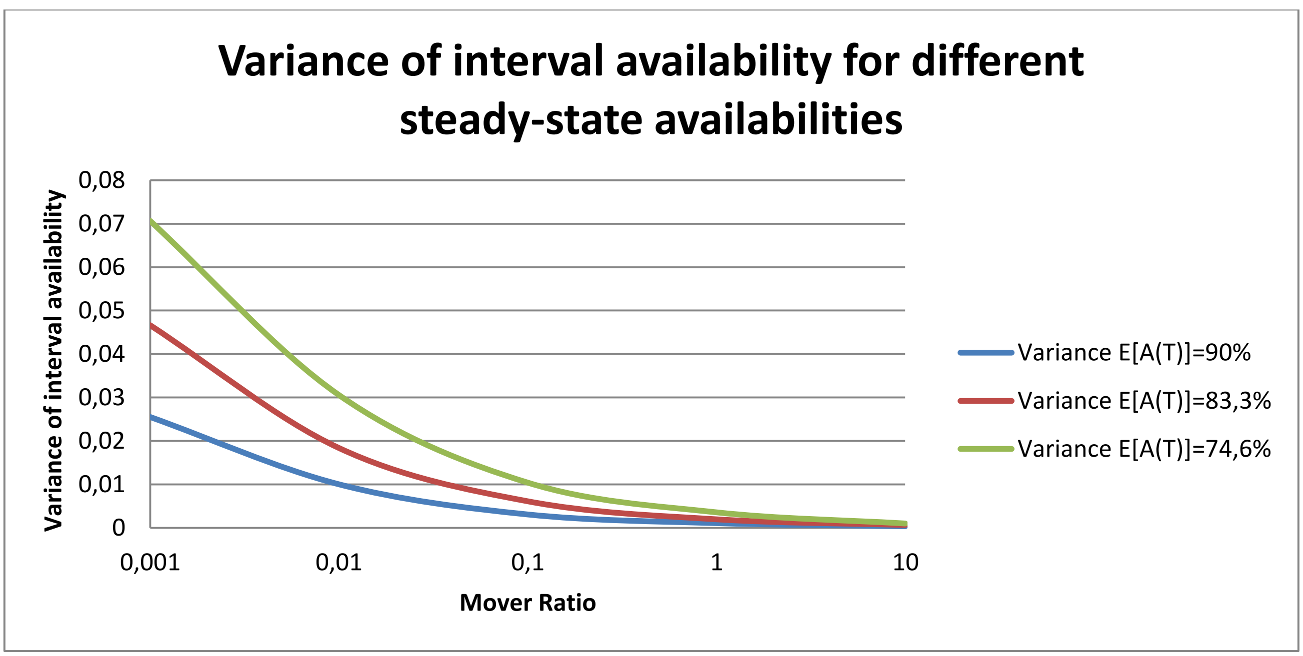

Figure 9: Variance of interval availability for different steady-state availabilities

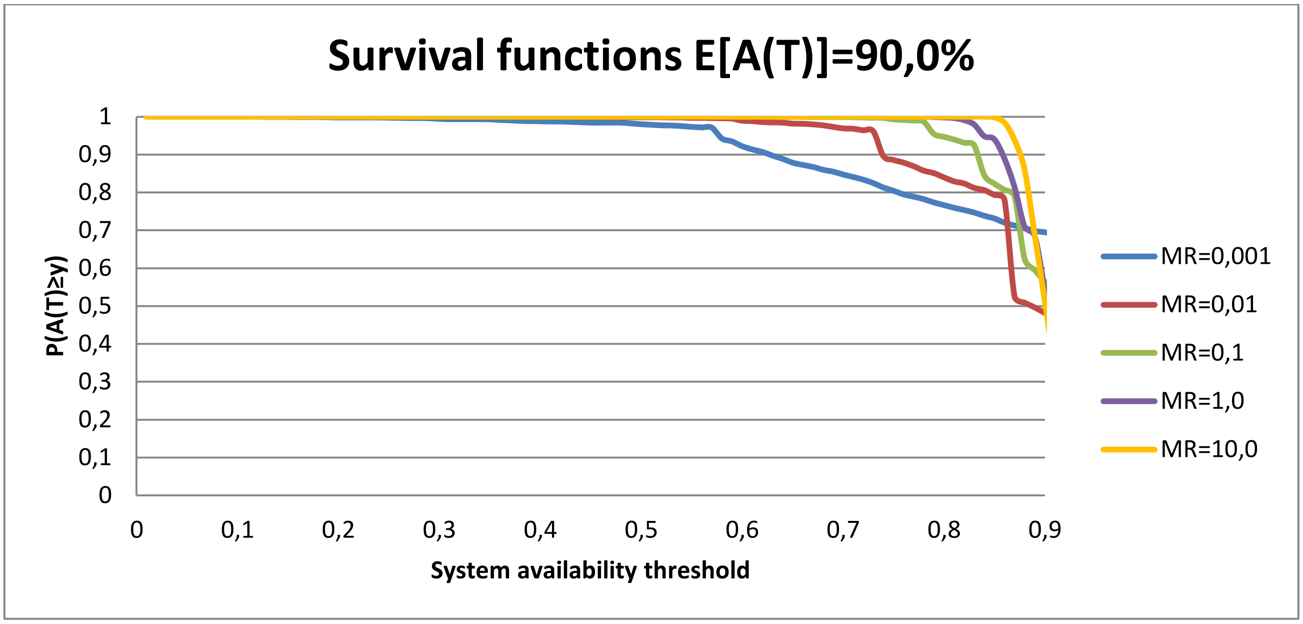

Figure 10: Survival functions for different mover ratios

The survival functions, in Figure 10, are plotted on a range between zero and the target operational availability (90%), since this range is relevant for the penalty structure. Survival functions illustrate the probability that the system availability is above a certain threshold. For instance, the chance that the system availability is higher than 80% for a system consisting of fast moving items is still 100%, but for systems with slow moving items this probability is around 75%. Comparing the mover ratios, we can immediately see that fast movers never enter the region where penalties strike the hardest and they preserve relatively large probabilities when approaching the target operational level. At some point, they intersect with the slow movers. This can be explained through the fact that the area under the complete graph will be equal for all functions, because the steady-state availability is equal. This also means that slow movers still have a relatively large probability of achieving higher availabilities than the operational target, but these are not interesting for the penalty structure. Figure 11 clearly illustrates these effects when expected penalty costs per year are plotted for the different mover ratios. Systems consisting of fast moving items will lead to significant lower expected penalty costs compared to systems with slow movers.

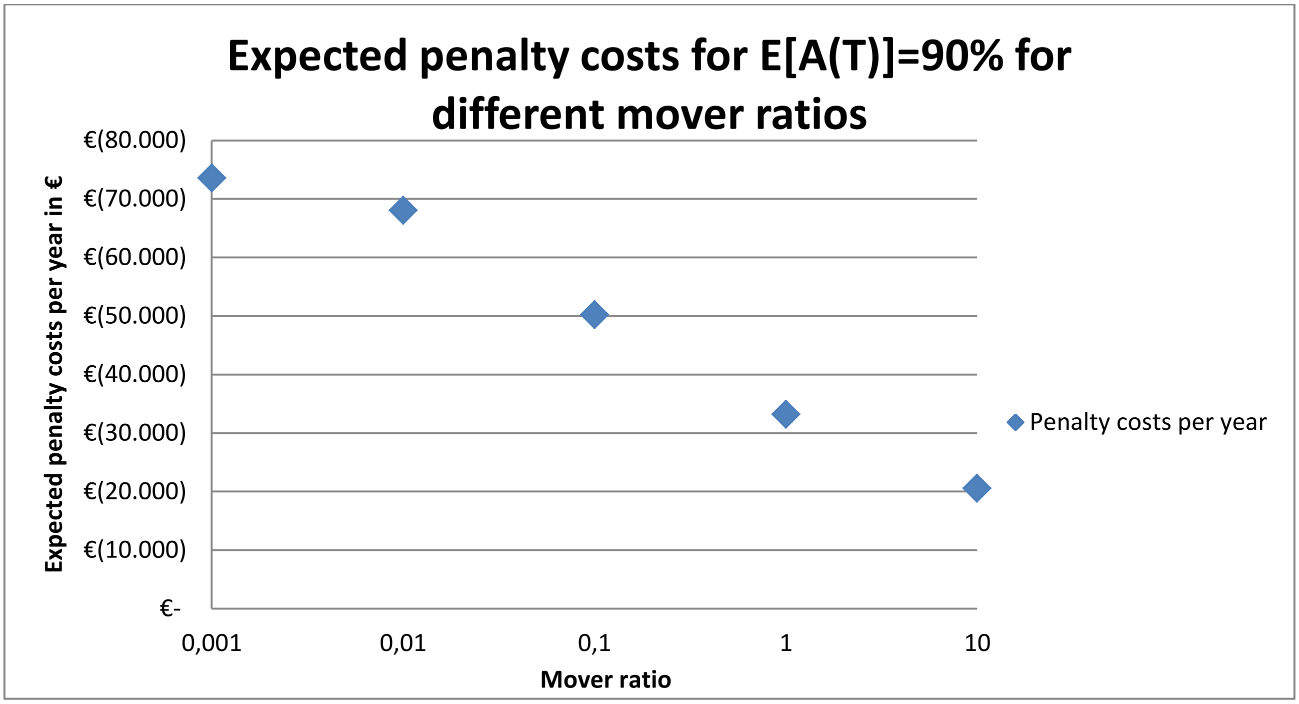

Figure 11: Expected penalty costs per year for different mover ratios

For detailed plots of the interval availabilities and detailed information on the input data for this experiment we refer to Appendix G: Details for simulation of experiment 1.

5.1.2 Effects of adding stock for slow mover versus fast movers

For these experiments, we will maintain a constant level of steady-state availability in contrast to the previous experiment. Furthermore, we consider the possibility of adding items to stock to investigate the effects of this. Recall that Thales has little influence on selecting fast and slow movers for certain systems and we do not consider redesign as mentioned in the research scope, therefore we only investigate how the system behaves if we invest in certain stock. We will distinguish two cases, since we have the possibility to add stock centrally at the depot, but can also store it locally at the bases. It is important to notice that we will now have two different items in terms of failure rate and repair time in the radar systems, a slow and fast mover, whereas in previous experiments they were similar, but obviously the backorder product will be the same. Also, since we are adding stock now, we will increase the availability in the system. Therefore, we increase failure rates and repair times to a certain extent such that adding one item to stock at the depot again leads to a steady-state availability of 90%. For more details on the input data and results for these experiments, we refer to Appendix H: Details for simulation of experiment 2.

Adding stock at the depot: fast mover vs. slow mover

It is clear that adding the slow mover at the depot results in significantly reducing the variance of the availabilities compared by adding a fast mover by looking at Figure 12. We obtain a spectacular reduction of more than 90% in terms of variance of the interval availabilities if we decide to invest in a slow mover compared to a fast mover, while we maintain the same steady-state availability.

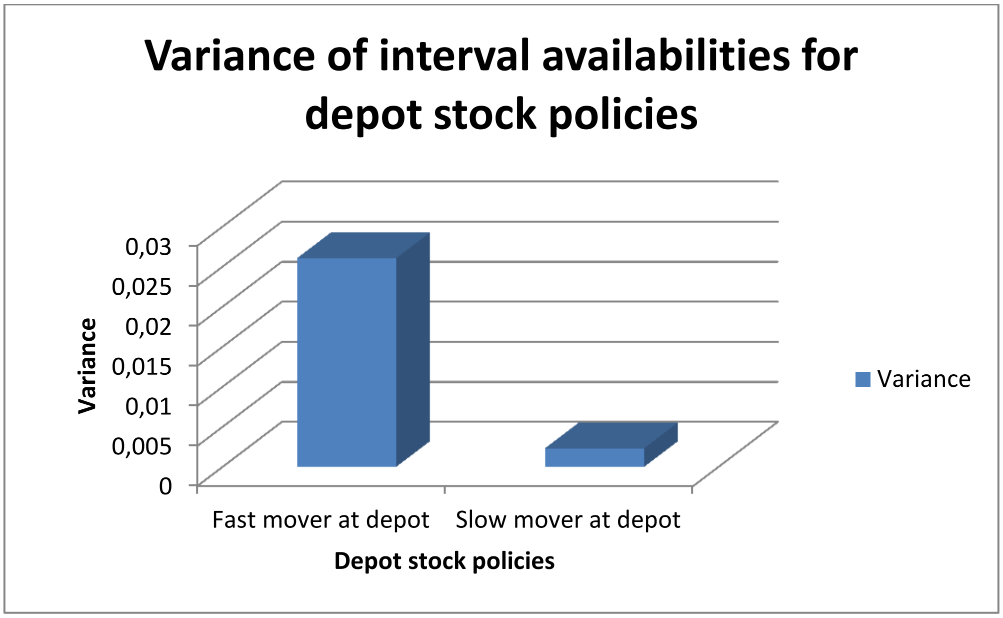

Figure 12: Stock policies at depot compared to basic model

Adding stock at the bases: fast mover vs. slow mover

We expect the same relation for adding stock at the bases, but we do not know whether the effects will be greater or smaller. Therefore, we perform similar experiments where we also adapt failure rates and repair times. If we only add one item of a slow mover or fast mover to one base, we still have a lot of variance due to the other base where no stock is available. However, if we choose to add one item of stock at each base, then we see the same drastic reduction of variance when adding the slow movers compared to the fast movers as is illustrated below in Figure 13.

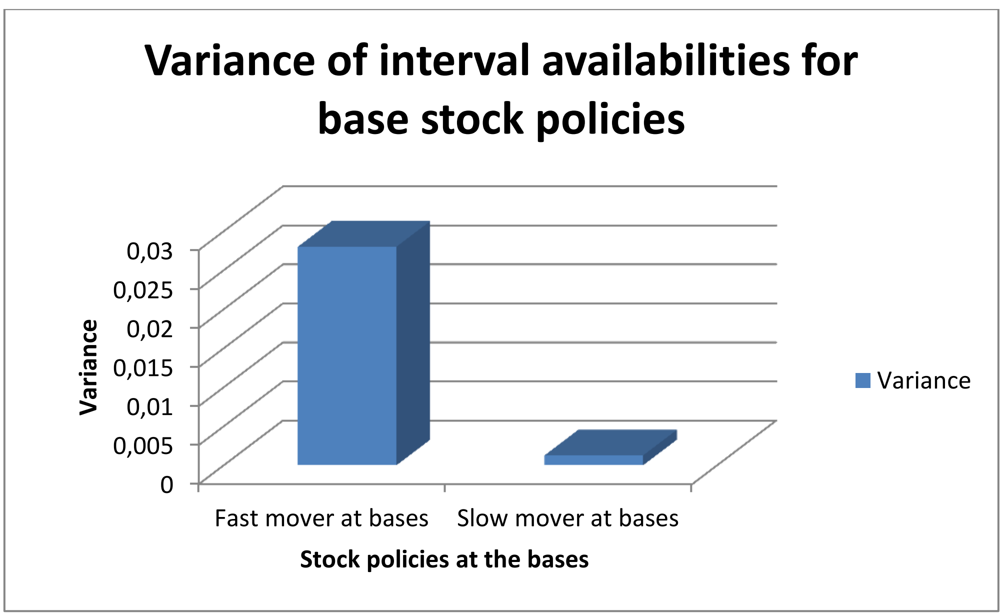

Figure 13: Stock policies at bases, adding one item to stock at both bases

Again, we observe the same effects as in the case of adding stock at the depot. However, effects are particularly visible when we add stock at all bases, because then we eliminate most of the variance caused by the slow mover.

The most interesting thing is to see how these different stock policies relate to the expected penalty costs per year. These results are illustrated below in Figure 14. Note that the scenarios are not exactly equal between the depot and base stock policies. In the case of depot stock policies we add one item (slow mover or fast mover) to stock at the depot, whereas we add two items at the base stock policy, one item at each base. The failure rates and replenishment lead times are modified such that both stock policies lead to a steady-state availability of 90%.

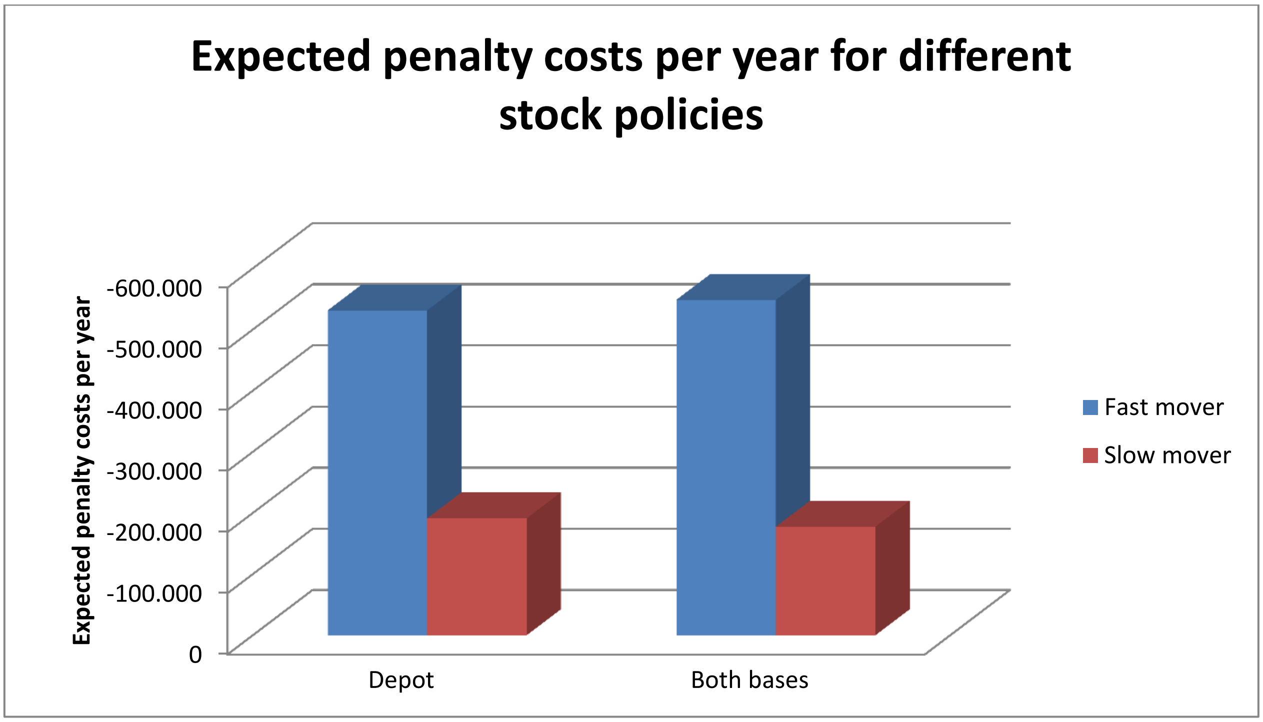

Figure 14: Expected penalty costs for different stock policies

The most important conclusion is that for the characteristics of the basic model, we will always prefer to invest in a slow mover rather than in a fast mover, since this reduces variability in the system availability significantly resulting in considerable lower expected penalty costs per year and thus we agree with the first hypothesis stating that stock investments for slow moving items lead to lower penalty costs compared to stock investments in fast movers.

5.1.3 Effects of adding stock locally compared to centrally

Now we will examine the impact of the location of the stock. One difficulty to cope with in this experiment is that we need to include order and ship times. This is needed, because otherwise we would always prefer to put an item on stock at the depot, since holding costs are assumed to be the same on both locations and replenishment lead time is the same for the depot and the base. Then, the depot has the ability to serve both bases, so it will be able to pool the risks, whereas the stock at the base is dedicated to that certain base. However, by including order and ship times, we can overcome this issue.

From the results, illustrated in Figure 15 below, we see that adding local stock compared to central stock leads to higher variance of the interval availability. However, these effects are relatively small. In the survival functions there does not exist any obvious deterioration when adding local stock compared to central stock as can be seen in Appendix I: Details for simulation of experiment 3. Here you can also see that the mover ratio of the fast mover is adapted to maintain realistic failure rates and repair times.

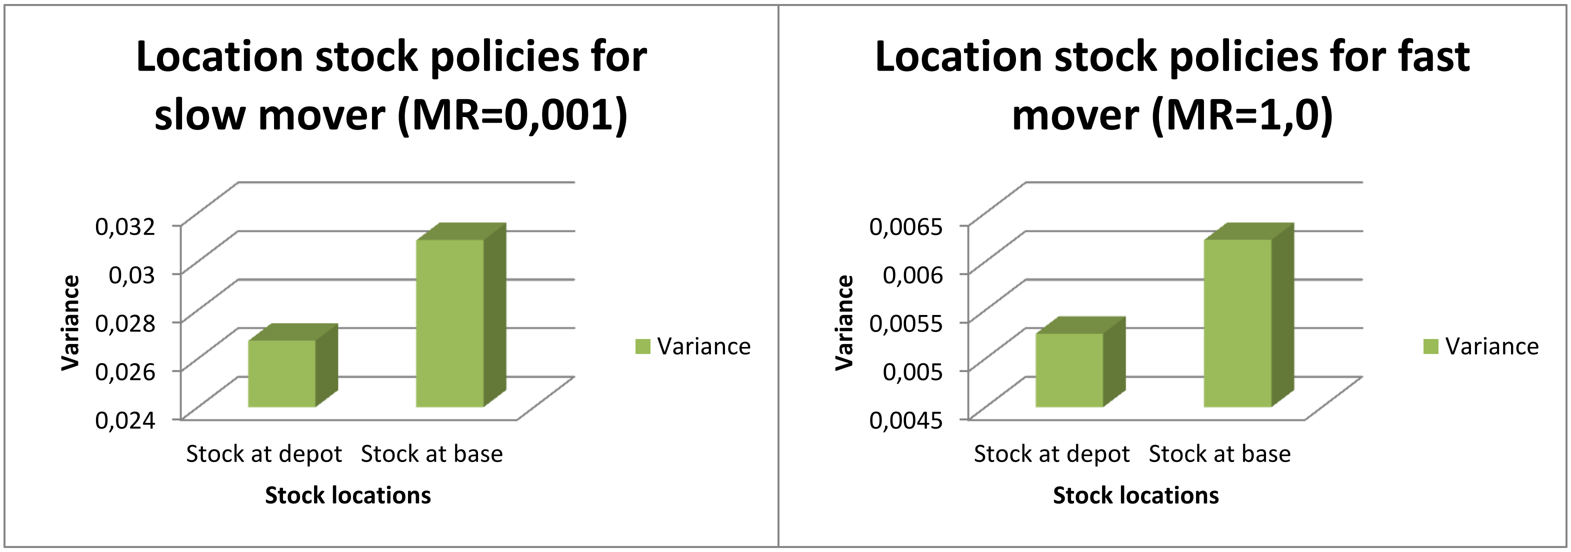

Figure 15: Key performance indicators for central stock policies vs. local stock policies

Also if we look to the expected penalty costs per year for the different location stock policies, we see that there are some minor benefits of adding stock centrally. This means that the second hypothesis, stating that central stock investments lead to lower penalty costs compared to local stock investments will also be accepted, but we have to take into account that effects are significantly smaller compared to the first hypothesis.

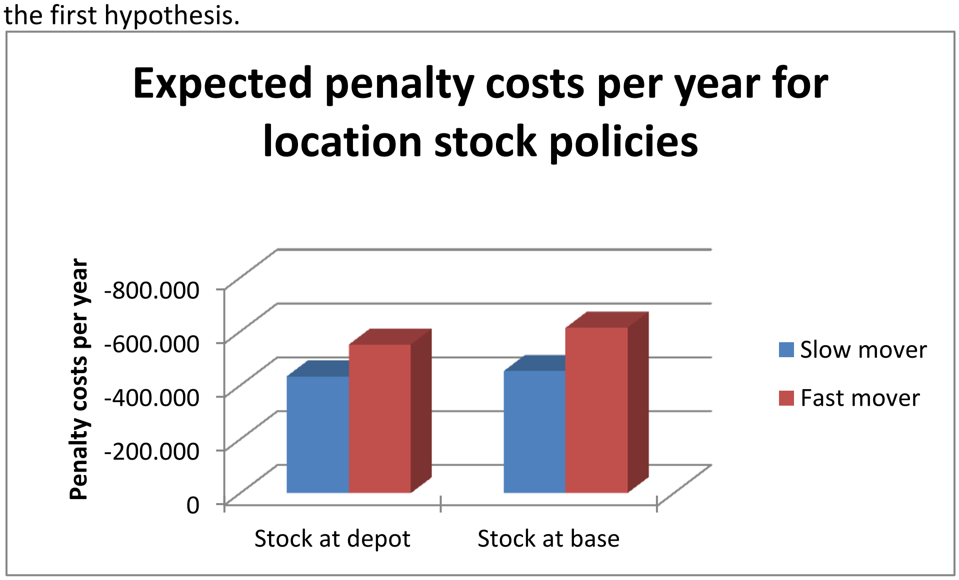

Figure 16: Expected penalty costs per year for location stock policies

5.1.4 Effects of review periods

Previous research at Thales showed that the variability of the interval availability decreases with the interval length (Coenen, 2009). This would suggest that the variance and coefficient of variation of the interval availability are larger in short review periods compared to longer review periods. To verify this, we distinguish five different review period intervals. We use the basic model concerning the item and model characteristics, meaning no order and ship times and no stock and we compare the outcomes to the situation where we review on a yearly basis.

Analyzing the results, we can indeed conclude that the statement of (Coenen, 2009), variance of interval availability decreases as the interval length of the review period increases, is justified and that the effects are even more obvious for fast mover compared to slow movers. This is also as expected considering the theoretical foundation of (Kirmani & Hood, 2008) where the coefficient of variation decreases as the interval length increases.

The input data and results of these experiments are illustrated in Appendix J: Details for simulation of experiment 4.

5.1.5 Joint contracts vs. individual contracts

For these experiments we extend the basic model with another two bases, leading to a total of four bases, to obtain more reliable results for the different contracts. For the joint contract, we measure interval availability as an average of the base availabilities per year whereas for the individual contract we just measure the availability per base and relate the penalty costs to these figures. For the joint contract we will again use the contract value of €3.500.000 per year and for the penalty costs at the base we will use an evenly divided distribution leading to a contract value of €875.000 per base. The penalty structure is the same as before, drafted in Table 1.

Obviously, variance among the interval availabilities will be lower for the joint contract since we use the average of the four bases which all have an expected availability of 90%. This means that low availability for one base can be compensated through high availabilities at the other bases. This 'portfolio-effect' always guarantees a lower variance for the joint contract.

One would assume that when variance in interval availability is lower, that expected penalty costs also are minimized but because the given penalty structure is not linear there could be different findings. For example, if we consider four bases with availability of 100%, 100%, 100% and 0% respectively, we would have an average availability of 75% for the joint contract leading to a penalty of €1.750.000 (50% of €3.500.000). However, the individual contracts will only have a penalty for the fourth base. The corresponding penalty for this base is €962.500 (110% of €875.000). In these cases individual contracts, regardless of the significant higher variance, lead to lower penalty costs.

Results show that for slow moving items the penalty costs will be lower for the individual contracts. This is as expected, since for these items we can have the scenario as in the aforementioned example. For the fast moving items the joint contract is a better option. This could be explained through the fact that variance in availability will be significant lower for these items, so interval availabilities will be close around the expected 90%. In this area, the penalty structure is linear and thus a joint contract offers more benefits.

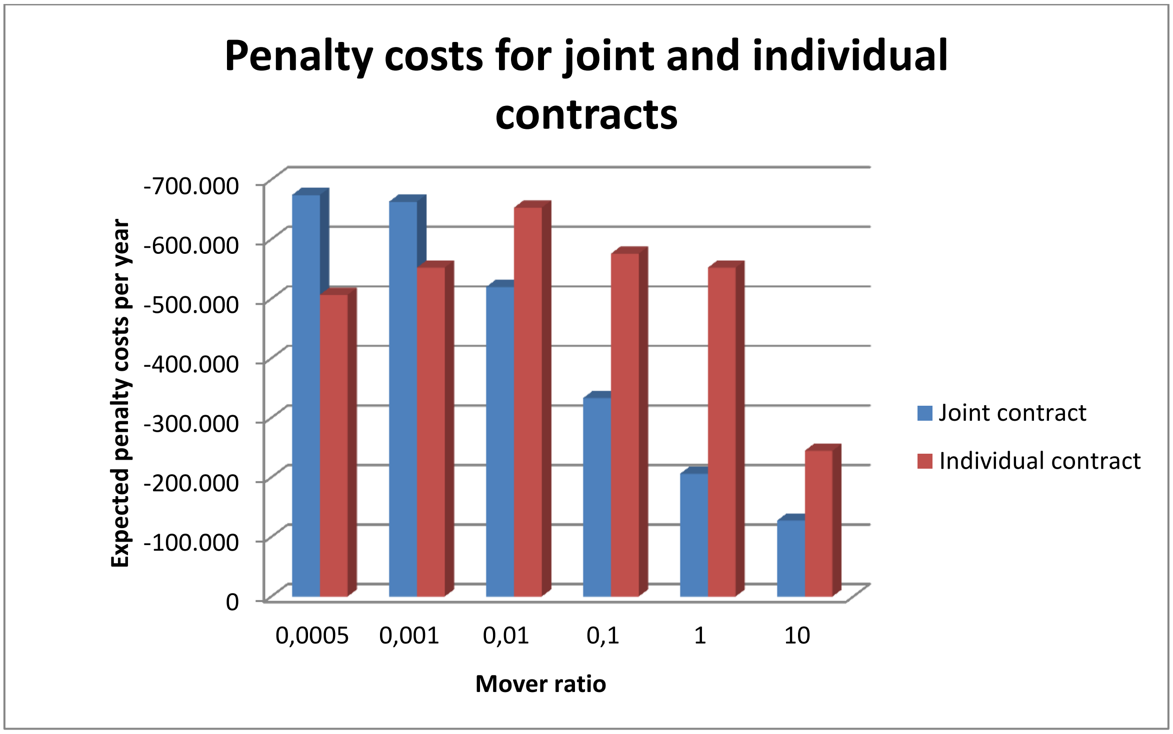

Figure 17: Penalty costs for joint and individual contracts

5.2 Simulation extensions

Some extensions to the basic model will be discussed in this paragraph to examine the earlier results still hold in more realistic settings.

5.2.1 Order and ship times

Through logic we would expect that including order and ship times have larger effect on fast movers, since these fail more often and thus will make more use of these times, because we have no stock for the items in the basic model. Therefore, the interesting thing to find out is whether the addition of stock to slow movers still holds if we consider these order and ship times in our model. However, a difficulty here is to identify two similar steady state availability items since these order and ship times lead to additional downtime. When we would not adapt the failure rates and repair times as in the basic model, the observed effect is that steady state availability would dramatically decrease to approximately 60% and that adding stock to the fast mover leads to significant higher average availability compared to adding stock for a slow mover. In every experiment before, we saw that average availability stayed the same, but now we do see a difference. This can be explained through the fact that the replenishment lead time of the fast mover increased relatively more than the slow mover, since these order and ship times are included and therefore the backorder product of the fast mover is greater than that of the slow mover. Because of this increase in steady-state availability, INVENTRI will make the correct decision and therefore this scenario is not interesting.

However, if we modify the order and ship times and the item characteristics in such a way that steady state availability will remain the same for adding stock at the depot, we can make a valid conclusion upon the effects of order and ship times. Details on the input data can be found in Appendix K: Input data and results for simulation of order and ship times extension. Note that the fast mover now has a mover ratio of 1,0 to maintain realistic failure rates and repair times.

Results indicate that it is still more useful to invest in slow moving items instead of fast movers, but effects slightly decrease compared to the case where we had no order and ship times at all. For example, in the case of no order and ship times variance is decreased with approximately 93% if we invest in a slow mover compared to the fast mover. If order and ship times of 365 hours are included, which is a realistic value for Thales, then the variance reduction is around 79%. The survival functions also show that investing in the slow movers lead to lower penalty costs, because the low interval availabilities will not be reached and the graph does not get disrupted until we approach the target operational availability.

5.2.2 Repair time distribution

The inclusion of an exponential repair time distribution in the simulation model leads to additional variance in the system. This leads to even bigger advantages on investing in slow movers instead of fast movers compared to the case where we had deterministic repair times. For results, we refer to Appendix L: Results for simulation of repair time distribution extensions. For these experiments, the case where the exponential distribution is included is compared to the case where these times are deterministic.

5.2.3 Multi-indenture model

In this section we will examine whether both hypotheses will still be accepted for a multi-echelon, multi-indenture model. Also, we consider the same extensions for this model as before. This means we study the effects of review periods and order and ship times. Details on the input characteristics and the results can be found in Appendix F: Details for simulation of multi-indenture model.

From these results we can conclude that there are no significant differences between the multi-echelon, single-indenture model and multi-echelon, multi-indenture model with regard to the outcome of the formulated hypotheses. Also in this extended model stock investments in slow movers and central stock allocations lead to lower variance and penalty costs.

5.3 Sensitivity analysis for input parameters

In this section we will perform a sensitivity analysis with respect to replenishment lead times and failures rates to identify interesting areas for Thales to invest in, in order to reduce variance and expected penalty costs.

5.3.1 Replenishment lead times

To investigate the impact of repair time reductions we use the aforementioned basic model and plot the variance of the availability and the expected penalty costs.

From the results, illustrated in Appendix M: Sensitivity analysis for repair times, we can conclude that variance increases more or less in a linear fashion (R^2=0,9943) when replenishment lead times increase.

Therefore, we would also expect that penalty costs will increase when replenishment lead times increase. This is indeed the case, but due to the given penalty structure this is not a linear relationship. A logarithmic dependency is more likely to hold here (R^2=0,9794) meaning that for large replenishment lead times (>½T) there is not much to gain.

However, an important side note in this analysis is that the failure rate of the items is adapted to ensure the steady-state availability of 90%, otherwise we cannot use the provided penalty structure. Also, the level of the steady-state availability influences the variance as we saw in section 5.1.1, so maintaining the same level eliminates this effect. In the following section we examine the impact of failure rates and replenishment lead times separately for various levels of steady-state availability.

5.3.2 Split analysis

In this section we will investigate the effects of the failure rate, replenishment lead time and the level of the steady-state availability separately in a so called split analysis.

For the split analysis, we modify the two parameters in the following way such that they are interesting for Thales, because these parameters are often used in reports and formulas:

$$Failure\ parameter = Failure\ rate\ (FRT) * Annual\ operating\ hours (AOH) \tag{5.1}$$

$$Downtime\ parameter = \frac{Replenishment\ lead\ time\ (RLT)}{Interval\ length\ (T)} \tag{5.2}$$

Recall that in the basic model we use a yearly review period which indicates that T is equal to 8760 hours. Also, we assumed a continuous mission profile, meaning the AOH is equal to 8760 hours.

The range for the failure parameter will be between a failure once a hundred year and 5 failures per year. The downtime parameter is plotted between 0,1 (replenishment takes 876 hours) and 2 (replenishment takes 2 years).

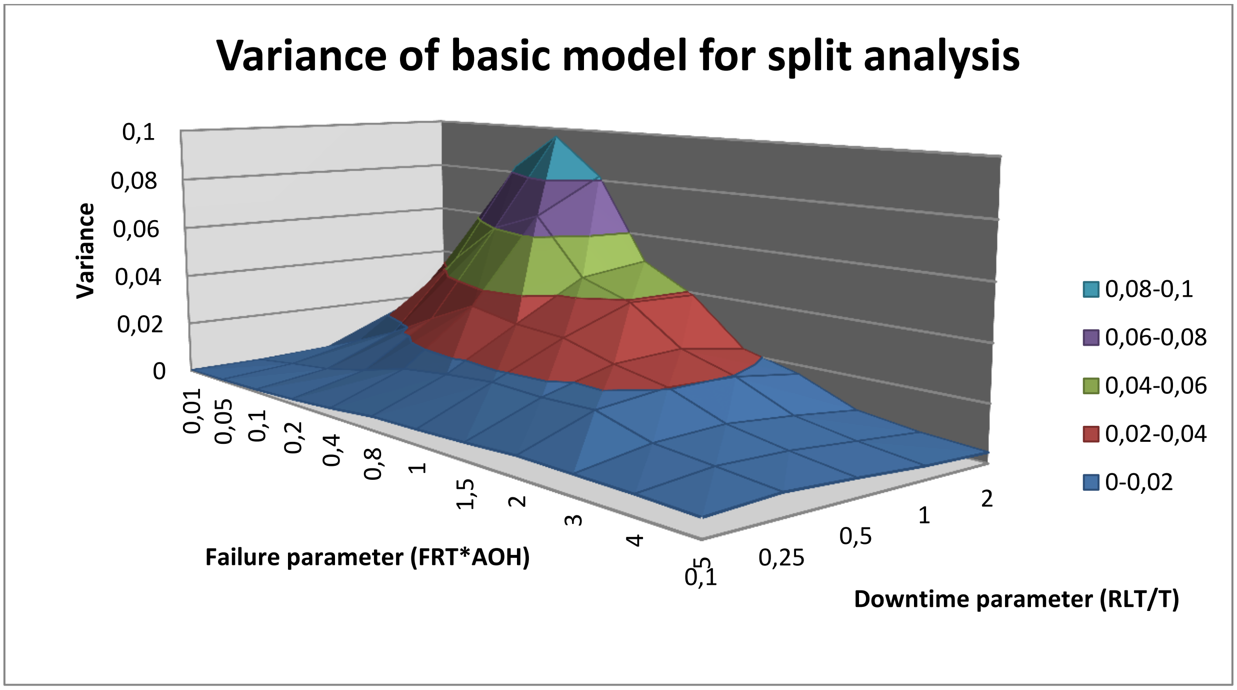

Figure 18: Variance of basic model for split analysis

We can conclude that variance is relatively high in the cases where average availability is around 50%, illustrated in Appendix N: Sensitivity analysis for split analysis, as was already predicted by the formula of (Hopp & Spearman, 2000). Also, if we select two cases where the product between the parameters is equal, meaning same steady-state availability, we see a higher variance in cases where the downtime parameter is relatively high and the failure parameter is relatively low. This supports our previous findings on slow movers leading to a higher variance compared to fast movers in a system.

Although the regression analysis is only able to explain 31,6% of the total variation, we can still conclude upon the values for the coefficients for the different input parameters. From Appendix O: Regression analysis for split analysis we can see that both the downtime parameter and the level of the average availability have a significant impact on the level of variance, whereas the failure parameter does not play a significant role. The important conclusion is that, according to this simulation, the downtime parameter has significant more influence on the variance compared to the failure parameter. Therefore it is more interesting for Thales to spend resources on shortening replenishment lead times rather than reducing failure rates in order to minimize variance.

5.4 Conclusions on simulation results

Our experiments show that both hypotheses derived from the literature are validated during the simulation study. We are able to reduce penalty costs significantly by investing in slow moving items rather than in fast moving items. Also, we identified other ways to influence penalty costs without looking at the spare parts allocation. Penalty costs are decreased when the interval length of the review period is increased; therefore it is wise to aim for long review periods. Furthermore, the structure of the contract can lead to additional savings in penalty costs, but this relationship is not straightforward. It depends heavily on the penalty structure and the mover ratio of the items in the system. Although intuition would suggest that a joint contract over all bases will always be the better option, it could be that an individual contract for each base lead to lower overall penalty costs. Simulation is required to find the best option. Lastly, by using a split analysis and regression analysis we found that it is more useful to spend resources on shortening replenishment times rather than reducing failure rates in order to reduce variability among the interval availabilities.

6 Techniques for improved spare parts allocation

In this section we will provide some guidelines to improve the initial solution obtained from the VARI-METRIC outcome from INVENTRI. We also consider one heuristic which is not INVENTRI-based.

6.1 I: Backorder product

Recall that the VARI-METRIC procedure is indecisive in situations where items have the same backorder product and price, taking into account the aforementioned assumptions and restrictions, and this product depends on failure rate and replenishment lead time in the basic model. To identify indifferent item pairs beforehand within VARI-METRIC, we need equal backorder products. However, this formula depends on the model characteristics. For example, when order and ship times (OST) are included in the multi-echelon case, these times should be included in the replenishment lead time, since repairs are never performed at the base but at the depot instead. Also, repair probabilities (r_f) play a role. In some cases an item cannot be repaired or repair is unsuccessful at the depot and then we need to reorder that item from the supplier instead. Furthermore, when items occur more than once in a system, we need to consider this quantity together with its failure rate to determine the real demand of that item. The various relevant models with their corresponding backorder products can be found in Appendix P: Identification of item pairs for VARI-METRIC. The general formula for the backorder product is:

$$Backorder\ Product = \ (Failure\ rate * Quantity) * [OST + r_f * MTTR + \left(1 - r_f\right) * Reorder\ time] \tag{6.1}$$

Note that this parameter does not depend on the current level of stock, but solely on the item characteristics. We take the levels of stock into account in the other methods.

After we have defined all model characteristics, we can select the relevant item pairs to provide assistance in selecting the most appropriate item with regard to interval availability and penalty costs. Therefore we need to determine the mover ratios and select the item with the lowest mover ratio, i.e. the slowest mover of the item pair, since this will lead to lower variance and penalty costs as found in our simulation study in section 5.1.2. Note that the fast mover of this corresponding item pair has to be already selected for the stock allocation by INVENTRI; otherwise we cannot exchange any stock between the items in the item pair. This means that stock has to be exchangeable between the items. To summarize:

1. **Consider INVENTRI as initial solution**
2. **Identify item pairs (with exchangeable stock) by using corresponding exactly equal backorder products**
3. **Determine mover ratios for these items**
4. **Exchange stock between this item pair. Add stock for the item with the lowest mover ratio and deduct the stock for the item with the highest mover ratio**
5. **Repeat step 2-4 until no more opportunities arise**

6.2 II: Waiting times reduction

To avoid the areas where the penalties strike the hardest we can look at the SIMLOX output and analyze the waiting times for certain items. There exist some items which have a waiting time higher than 8.000 hours meaning that during the review period of a year the system availability is lower than 10%. Therefore, we will focus on these items by again searching equal backorder product items with regard to these long waiting items and which are selected in the spare parts allocation of INVENTRI. Note that these high waiting times are caused by long repair or reorder times; therefore this method is also partially based on the outcomes of our simulation study. If we find these cases, we will exchange the stock between the items to ensure that the average waiting time is below a certain threshold. In our case, we maintained a threshold of ½T, meaning 4320 hours. Then we perform the simulation again to identify more long waiting items and we iterate until no more opportunities arise. Summarized:

1. **Consider INVENTRI as initial solution**
2. **Perform SIMLOX simulation**
3. **Analyze SIMLOX item results output by inspecting waiting times**
4. **Consider item with highest waiting time and search for an item with equal backorder product and exchangeable stock**
5. **Exchange stock between this item pair. Add stock for the item with the high waiting time and deduct the stock for the other item**
6. **Repeat step 2-5 until no more opportunities arise or when waiting times are below threshold**

6.3 III: Combination

We have a lot of slow movers in the current system, so we cannot only identify item pairs with equal backorder product, but we deduct stock from the fast movers, which usually have stocks larger than 1, and select multiple slow movers with long replenishment lead times and waiting times to reduce variance. Since we do not take the backorder product into account, we need simulation to ensure the steady-state availability of 90%, so a simulation run for the newly obtained stock allocation is needed to guarantee that steady-state availability did not change. It is sometimes necessary to keep stock for fast movers, otherwise the target availability may not be achieved. Therefore it is also necessary to run the simulation to check whether the stock from the fast mover may be deducted. Summarized:

1. **Consider INVENTRI as initial solution**
2. **Perform SIMLOX simulation**
3. **Analyze SIMLOX item results output by inspecting waiting times**
4. **Determine mover ratios for items**
5. **Deduct stock from fast movers and add stock to multiple slow movers with long waiting times**
6. **Perform SIMLOX simulation run to ensure 90% steady-state availability**
    a. **If lower than 90%, add more slow movers with long waiting times**
    b. **If higher than 90%, eliminate latest stock additions**
7. **Repeat step 2-6 until no more opportunities arise**

6.4 IV: Replenishment lead time

This method will not consider INVENTRI as an initial solution, but starts from scratch. In the split analysis from the sensitivity analysis we saw that replenishment lead times have more influence on variance than failure rates. Therefore, we will sort the item list on backorder product from largest to smallest and then we sort it on replenishment lead times from largest to smallest. This is a nested sort. This way, the items where failures occur more often and where failures have a large influence on the variability of interval availability are on top. Now, we will add these top items to our stock allocation by selecting one item to stock. Once already stocked, we move to the next item. We need to run the simulation in order to find the 90% steady-state availability. When this target is exceeded, we will eliminate stock from items which are in the bottom of the list, having low replenishment lead times. To summarize:

1. **Perform a nested sort**
    a. **Sort item list on backorder product from largest to smallest**
    b. **Sort item list on replenishment lead time from largest to smallest**
2. **Perform SIMLOX simulation to check 90% steady-state availability**
    a. **Add stock for items which are on the top of the list and are not in the stock allocation yet, when availability is below 90% (many failures with high impact)**
    b. **Eliminate stock for items which are in the bottom of the list when availability is above 90% (few failures with low impact)**
3. **Repeat step 2 until 90% steady-state availability is reached**

7 Implementation of guidelines in real-life scenario

To validate the given guidelines, we will use the setting of a real-life example. We will mention the differences with regard to the assumptions in the basic model as stated before. Previous assumptions hold when no differences are mentioned. All improvement techniques are compared to the current outcome of INVENTRI. For the implementation of the guidelines, we used supporting Excel-sheets to manually determine the stock allocation for the different techniques.

7.1 Model formulation

System view

The radar system consists of 8 indenture levels with a total of 209 items. There are different types of items involved such as assemblies, disposable units and LRUs. There is no redundancy involved, but we do have common items. Furthermore, certain items have repair probabilities between 0 and 1 meaning that in some cases the item can be repaired at the depot or that repair is unsuccessful and we have to reorder the item at the supplier.

Supply network view

The radar is installed at one ship and this ship is part of a depot. The order and shipping time between these locations is considered to be 30 days or 720 hours. The initial spare parts allocation is calculated and obtained through INVENTRI.

Usage profile

The radar has a daily mission with duration of 5,5 hours throughout the whole year leading to a system utilization of ±23%.

Run length

For this system we will simulate for 1200 years and we use the same warm up period of 20 years.

7.2 INVENTRI outcome

Using the spare parts allocation outcome from INVENTRI as basis for the simulation yields the following results. The average availability is indeed 90%, its variance is 4,5% and the expected penalty costs are almost 560.000 per year.

7.3 Results of improvement techniques

Backorder product

There exist a considerable amount of item pairs with equal backorder product **and** different mover ratios in the current spare parts allocation from INVENTRI. In these cases, INVENTRI indeed selected the item with the highest mover ratio (fast mover), since this one is cheaper. For these instances, we will now select the item with the lower mover ratio to comply with the aforementioned guidelines.

Results show that we are able to reduce variance with 10% and the coefficient of variation with 5%. However, we are not able to reduce the expected penalty costs which stay around 560.000 per year.

Waiting times reduction

After performing the steps mentioned in this procedure, we could reduce variance with 22%. However, we were still not able to reduce the expected penalty costs significantly.

Combination

For this method, the results show that we could even reduce variance with more than 50%, but still we did not find any improvement for the penalty costs.

Replenishment lead time

This method yields the best results in terms of reducing variance, because we achieved a spectacular reduction of more than 65%. However, the expected penalty costs **increased** with almost 10% for this scenario, the largest increase among all methods.

7.4 Sensitivity analysis for real-life example

In this section we will discuss the impact of different penalty structures and analyze which technique is most suitable in the current practices at Thales, where adjusted target system availability is used to reduce the probability of receiving a penalty.

7.4.1 Other penalty structures

The current penalty structure is constructed such that minor deviations from the target availability get punished relatively hard, whereas penalties for very low availabilities are not that high. This results in a 'belly-shaped' curve. A recent study (Wijk & Andersson, 2012) found that step-wise exponential penalty structures are suitable for Performance-Based Logistics. Therefore, we consider two additional different structures, a linear and exponential relationship within the same range. Note that we used the same regions for the system availability during the review period, we only change the level of the penalty as can be seen in Figure 19. Although (Wijk & Andersson, 2012) consider these exponential structures as suitable for PBCs, it seems unlikely that customers will accept such a structure, since penalties are practically negligible in the beginning. Linear penalty structures therefore seem the most appropriate, since they punish the system availability in the most 'even' way.

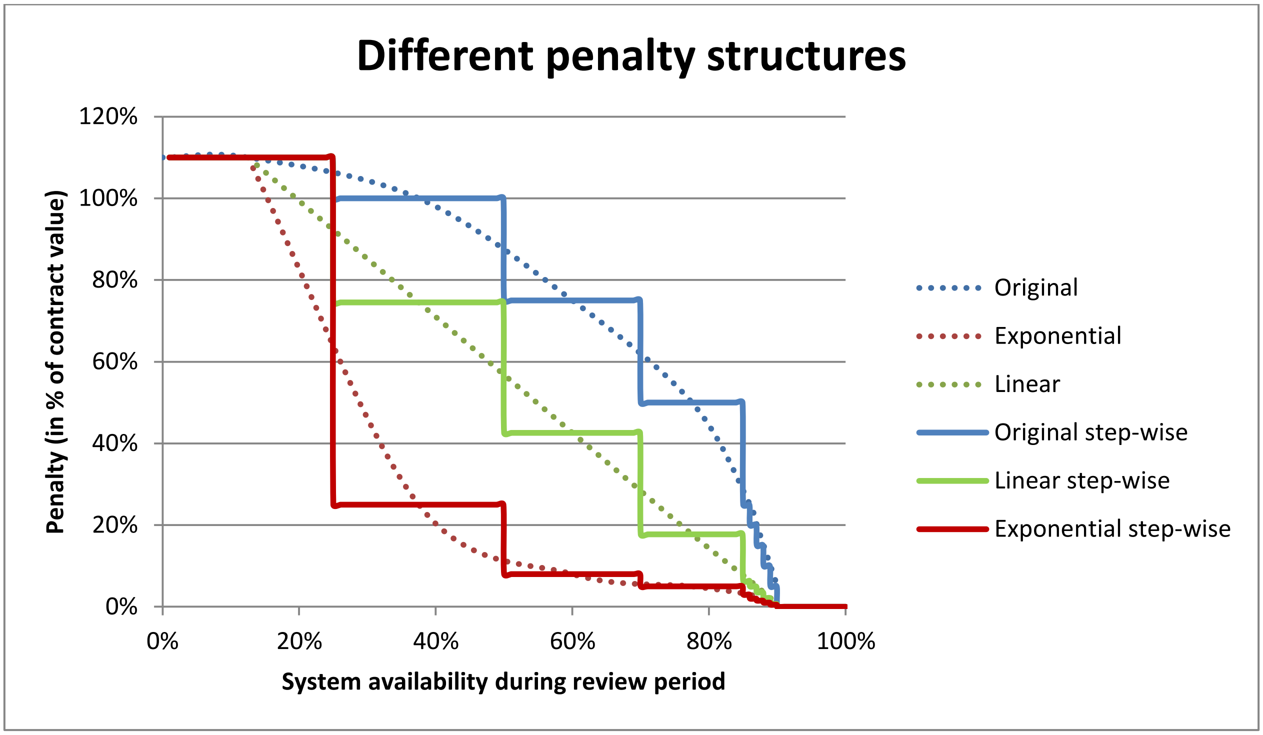

Figure 19: Different penalty structures

When comparing the differences in expected penalty costs for these three structures with regard to the INVENTRI outcome and the improvement techniques, it becomes clear that the penalty structure has a significant influence. The suggested improvement techniques do not really work for the current penalty structure, because minor deviations are already severely punished, but if we switch to a linear relationship, savings can go up to more than 30%. The improvement techniques become even more attractive when conserving an exponential penalty structure. Another important remark is that the technique where we focus on the replenishment lead time outperforms every other technique for the two suggested structures, but is the worst solution for the current structure. The results are shown in Figure 20 on the next page.

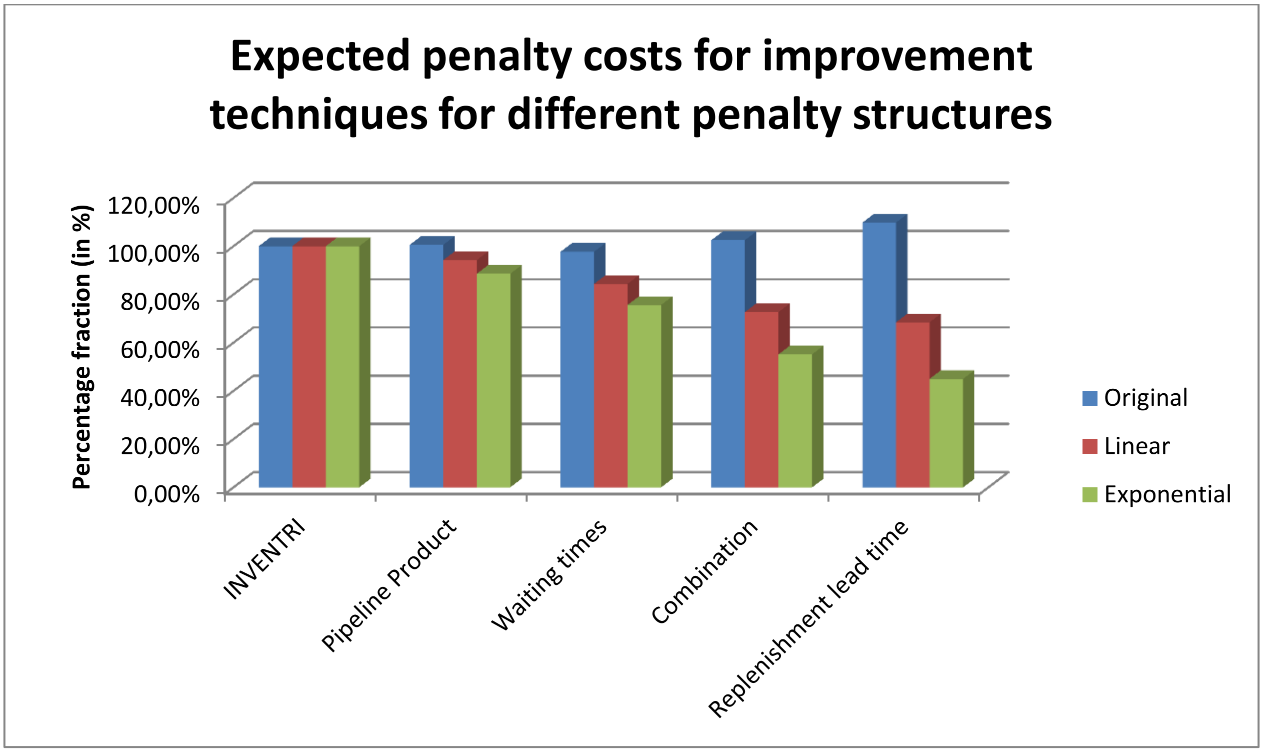

Figure 20: Expected penalty costs for improvement techniques for different penalty structures

7.4.2 Current practice: adjusted target system availability

To reduce the probability of receiving a penalty, the current practice is to add an additional 5% on top of the target system availability when determining the spare parts allocation with INVENTRI (Ypma, 2014). This means that for the aforementioned INVENTRI solution, the actual target availability would be 85% instead of the 90%. The steepness of the survival functions (Appendix Q: Survival functions for real-life example) around the 90% target and the large reduction in variance would suggest that our methods will outperform INVENTRI when using this practice, even for the current penalty structure. To calculate the expected penalty costs for this adjusted target availability, we use a slightly adapted penalty structure drafted in Appendix R: Adapted penalty structure.

This means that the only difference compared to the original situation is that we consider an adapted penalty structure that incorporates the current practice at Thales. Figure 21 illustrates that obviously for all stock allocation methods penalties decrease, since the adapted penalty structure is less restrictive. However, the benefits are considerably larger for the Combination and Replenishment Lead Time method, as we already expected from the survival functions.

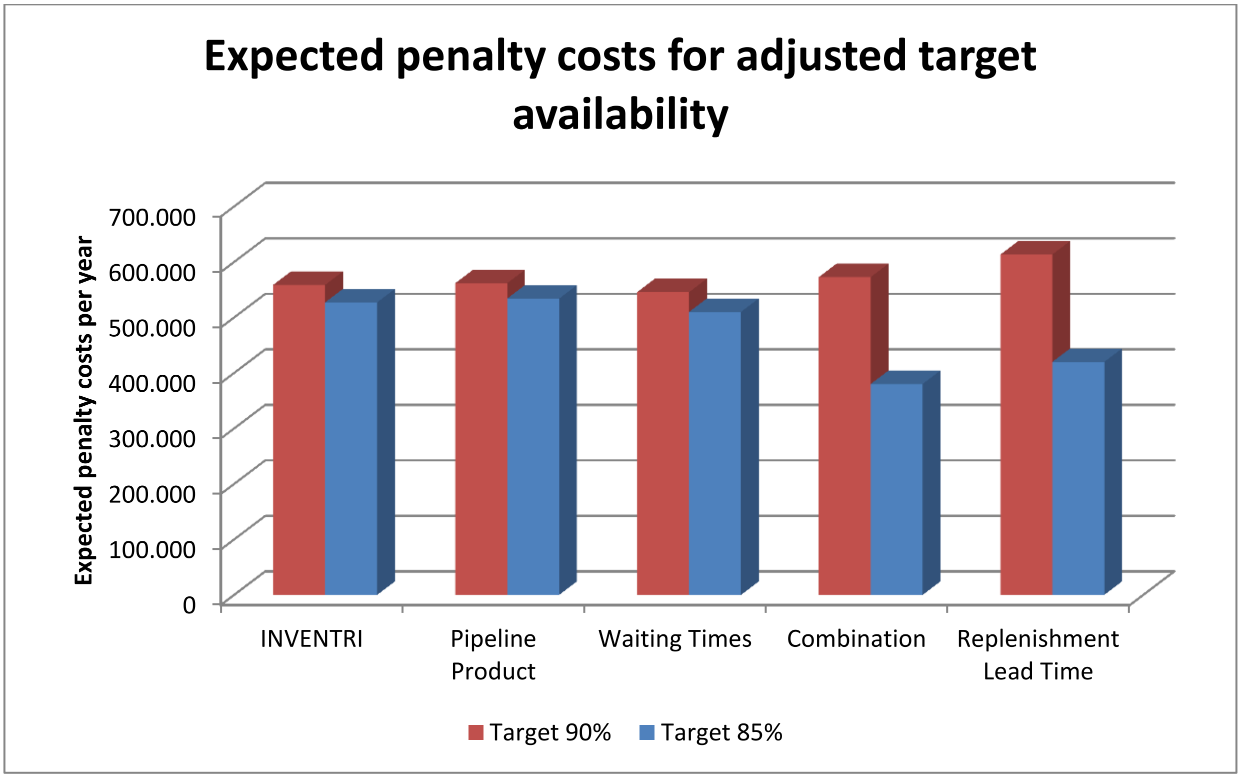

Figure 21: Expected penalty costs for adjusted target availability

7.5 Life-Cycle costs

To obtain a complete overview of the costs related to the radar system, we need to include the expected penalty costs per year as well as the corresponding costs for the initial spare parts allocation. Until now, we only focused on reducing expected penalties but we did not analyze at which cost this is achieved. To investigate whether the improvement techniques can still outperform the INVENTRI solution, we will use the life-cycle costs of the system. These are the total costs related to the system during its life-cycle, the costs for the initial spare parts allocation together with all expected penalties during the contract period.

From Figure 22, we see that for the original penalty structure we are not able to reduce the life-cycle costs. This is in line with our previous findings, since for the original penalty structure we were not able to reduce the expected penalty costs significantly and given the fact that INVENTRI has the cheapest initial spare parts allocation, life-cycle costs will not be lower in this scenario. However, if we look at a linear penalty structure, the Waiting times method can reduce the life-cycle costs in this case by almost 170.000. Although we saw that the Combination and Replenishment lead time method can reduce penalty costs the most, their initial spare parts allocation has significantly higher value such that they are not able to reduce the overall life-cycle costs.

Another important remark is that linear and exponential penalty structures considerably reduce the life-cycle costs, because they reduce the penalty costs significantly while maintaining the same initial spare parts allocation. For this reason, we can conclude that the Combination method and Replenishment lead time method do not improve the INVENTRI solution, since the penalties are relatively low for these penalty structures compared to the costs of the initial spare parts allocation.

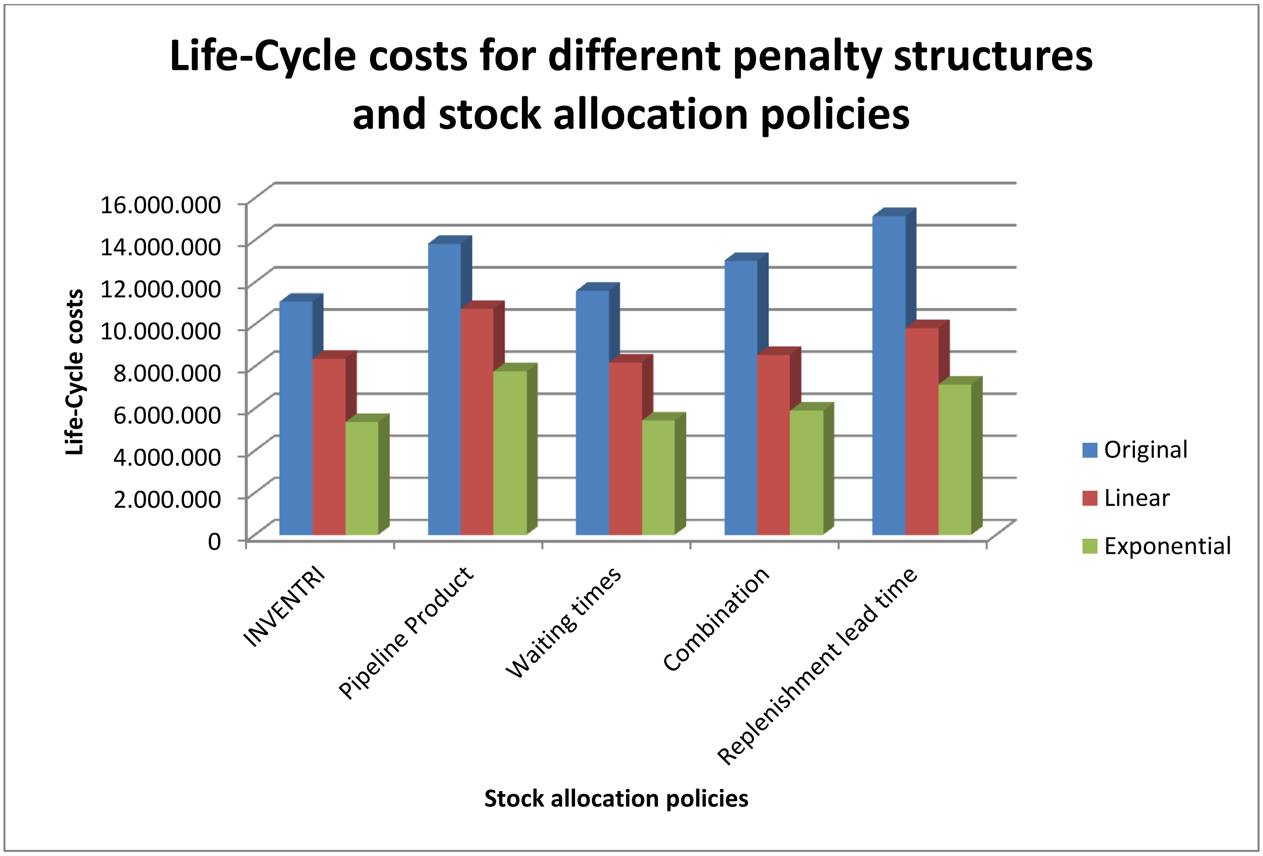

Figure 22: Life-Cycle costs for different penalty structures and stock allocation policies

When considering the life-cycle costs during the current practice at Thales, where target availability is artificially increased, we conclude that only the Combination method is able to reduce life-cycle costs compared to the INVENTRI solution with almost 470.000 as illustrated in Figure 23. Although there are large benefits for the Replenishment Lead Time method, these benefits still do not outweigh the additional investment required for the initial spare parts allocation. These large savings can again be explained through the given survival functions for these methods, since the function is relatively steep around the 90% target meaning that a lower target availability leads to considerable large savings in penalties.

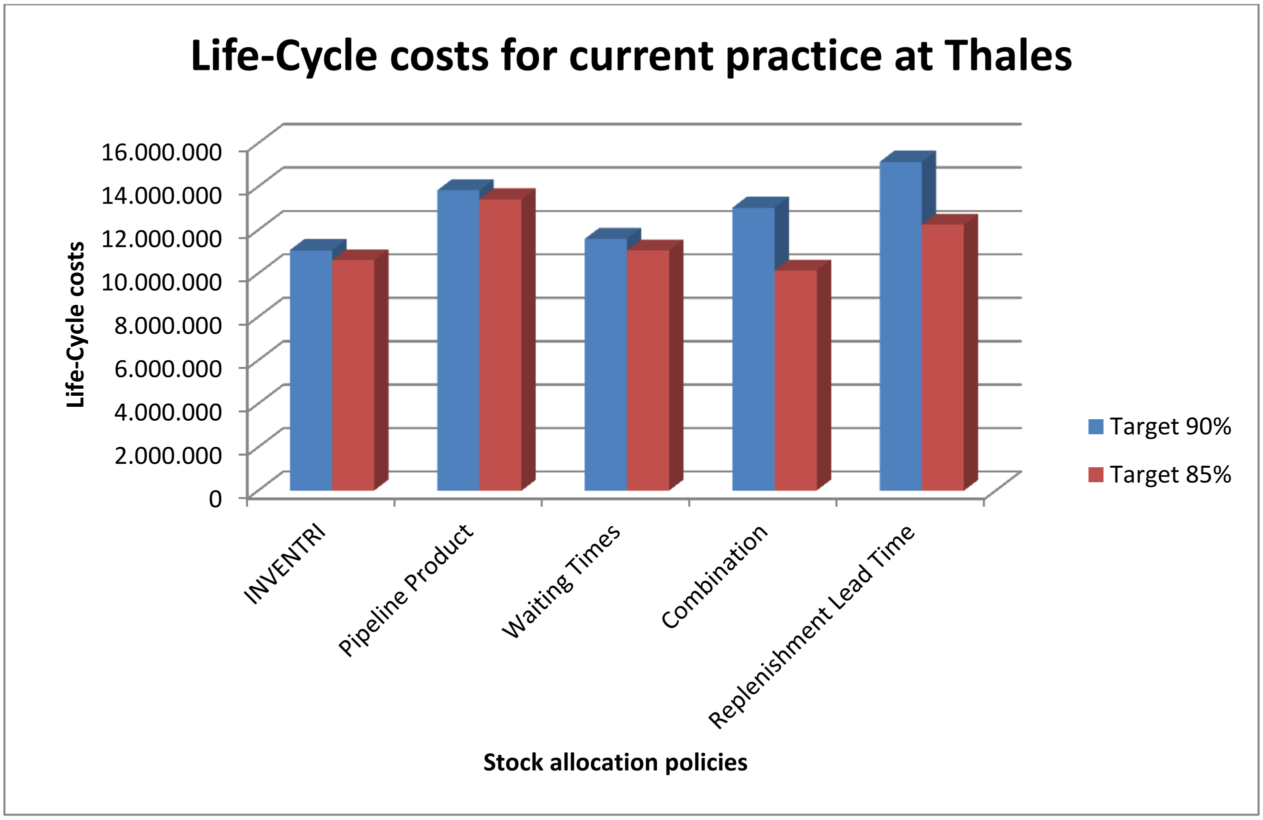

Figure 23: Life-Cycle costs for current practice at Thales

7.6 Conclusion on improvement techniques

For this example, due to an overwhelming amount of very slow moving items (±95%), we can conclude that we were not able to reduce penalty costs significantly. As can be seen in Figure 11 of Section 5.1.1, the reduction in expected penalty costs tends to stagnate for items with mover ratios lower than 0,01. However, results show clearly that each method can reduce the variance of the interval availability while maintaining the same steady state availability. Another argument for the lack of savings in penalties is that the given penalty structure punishes minor deviations from the target severely and very low availabilities are not penalized that hard. Nevertheless, penalty costs can be reduced significantly when considering other penalty structures, such as linear and exponential penalty structures. However, taking life-cycle costs into account we see that for the different penalty structures the Waiting times method is the best selection. This could be explained through the fact that the investment of the initial spare parts allocation does not increase that much, compared to the Combination method where we select multiple slow movers with high cost, or with the Replenishment lead time method where we do not consider INVENTRI as a starting solution at all and just select items based on their replenishment lead times without involving item prices. The benefits in expected penalty costs for the Combination and Replenishment lead time method do not outweigh the additional investment needed in initial spare parts.

Furthermore, the Combination and Replenishment Lead Time techniques work extremely well for Thales' current practices where they raise the target availability with 5% to cope with the penalties involved and the survival functions show that these methods are also outperforming the other techniques in terms of robustness. For this scenario, only the Combination method is able to reduce the life-cycle costs due to the heave additional investment required for the Replenishment Lead Time method.

Lastly, the Backorder product and Waiting times methods can be considered equally complex in terms of implementation, since with the aid of a single Excel-sheet, containing no hard formulas, we can manually compute the modified stock allocation without worrying about the steady-state availability, since we have item pairs with equal backorder products. However, for the Combination and Replenishment lead time method, this is not the case. For these methods we need simulation to obtain the desired steady-state availability and therefore these methods are more time-consuming and more complex.

8 Conclusions & recommendations

8.1 Conclusions

Preventive maintenance is only useful if costs are lower compared to corrective maintenance and if the failure rate of a system is increasing over time (Al Hanbali, 2013). Due to uncertainty in failure rates and the stochastic behavior of the failures it is difficult to conclude whether or not failures rates are increasing. During a mission, Thales' products are in use and preventive maintenance cannot be performed. Off-mission, preventive maintenance can be performed, but availability does not play a significant role (Ypma, 2014). Therefore, it is easy to see that we need to consider the fourth class of spare parts strategies (Rustenburg, 2005). To cope with the difficulties that arise in this class, Thales uses a VARI-METRIC based approach to determine the spare parts allocations.

However, we know that the VARI-METRIC procedure makes decisions based on steady-state availability without taking the variability of the different interval availabilities into account. This variability is very important for the penalty costs in the growing trend of Performance-Based Contracts (PBC) at Thales. Therefore, we introduce several guidelines to improve the initial solution from VARI-METRIC with respect to the variance in interval availability and expected penalty costs. A simplified multi-item, multi-echelon and multi-indenture basic model shows us that we are able to reduce variance and expected penalty costs significantly. In this scenario we compare investments between a fast moving item (many failures and short replenishment lead times) and a slow moving item (few failures and long replenishment lead times) where the VARI-METRIC approach will be indecisive. Stock investments in the slow mover will always yield the better result in terms of variance and expected penalty costs, since disruptions that cause long, infrequent outages tend to inflate the coefficient of variability more than disruptions that cause short, frequent outages, given constant steady-state availability (Hopp & Spearman, 2000). Also, central stock investments lead to lower variance compared to adding stock locally, but effects are considerably smaller. Both statements were extensively researched and validated in the simulation study. These results have led to several methods to improve the initial solution given by INVENTRI through VARI-METRIC.

We also verified these methods in a real-life example from Thales and the Backorder Product method, based on the simulation results from the basic model, showed us that we could reduce variance for this system with more than 10%. One drawback is that we were not able to reduce the expected penalty costs, because ±95% of all items within this real-life system contained extreme slow movers. Due to the given penalty structure the expected penalty costs are not significantly reduced; only the reduction in variance is obtained.

The other methods based on waiting times and replenishment lead times could reduce variance of interval availability even further up to 65%, but we still saw that expected penalty costs remained fairly the same for this combination of penalty structure and the given system. Other penalty structures, such as linear or exponential, led to 31% and 55% savings in penalty costs respectively for the Replenishment Lead Time method. The methods also outperform INVENTRI considering Thales' current practice where they adjust system availability artificially, to cope with the involved penalty costs. However, if we take the overall life-cycle costs into account, the benefits of lower expected penalty costs only outweigh the additional initial investment in spares for the Waiting times method, given a linear penalty structure. Life-cycle costs are minimized for Thales' current practice by using the Combination method.

8.2 Recommendations for Thales

We recommend Thales to apply the principles of the Combination and Replenishment Lead Time methods, depending on the scenario, to improve the outcome of VARI-METRIC, since they can reduce variance of the interval availability significantly and therefore achieve considerable reductions in expected penalty costs considering Thales' current practice where they raise system availability with 5% in comparison with the contracted target availability when determining the spare parts allocation. The Waiting times method can also be interesting when considering life-cycle costs, since this method was also able to reduce the overall costs during the complete life-cycle of the system.

Furthermore, we advise Thales to design or negotiate different penalty structures, more suited for Performance-Based Logistics, such as exponential step-wise penalty structures (Wijk & Andersson, 2012), because they lead to considerable savings in penalty costs and thus in life-cycle costs. Generally, it would be wise to involve the Customer Services & Support Department in the decision-making of the penalty structures, since these have a major influence on the expected penalty costs and the spare parts allocation.

Other options to reduce the expected penalty costs within these PBCs, without taking the spare parts allocation into account, can be achieved through the determination of the interval length and the contract structure. Try to aim for relatively long review periods, since this decreases interval availability variability drastically.

When having relatively many slow moving items it can be wise to conclude contracts on an individual basis, also depending on the penalty structure in the contract. For systems with a lot of fast moving items, joint contracts are a better choice in terms of expected penalty costs.

The last recommendation for Thales will be that it is wiser to invest resources on shortening replenishment lead times rather than reducing failure rates, since our theoretical foundation, sensitivity analysis and regression analysis showed that variability in the system availability is more influenced by these replenishment lead times than the failures rates.

9 Future research

In this chapter we will suggest some possible directions for future research.

9.1 Model extensions

We considered failure rates to be deterministic and constant in time, but often in practice we encounter wear-out phenomena, such as fatigue, creep and wear (Tinga, 2013). These features could be taken into account by considering for instance a Weibull-distribution for the failure rates. This is not done in during this project because failure rates are highly uncertain at the moment and SIMLOX does not support the possibility of varying failure rates during simulation. However, if failure rates indeed are increasing, then preventive maintenance can also become more relevant, so one could decide to also add this extension in the model.

Furthermore, the VARI-METRIC assumes that failure of an LRU is caused by at most one failure of an SRU. However, it could occur that there exist multiple failures within one LRU and one could examine the effects of this using the findings of (Abelin, 2010).

9.2 Expected penalty costs in VARI-METRIC

The prices between items can differ significantly as we already saw in our analysis for the life-cycle costs. Usually, we see that we have expensive slow movers and cheap fast movers (Ypma, 2014). VARI-METRIC currently selects the item with the largest backorder reduction per invested euro. Due to the equal backorder products in item pairs, the backorder reduction will be the same. However, as we mentioned before, the slow movers tend to be more expensive. So this would usually lead to stock investments in fast movers, but if we include expected penalty costs during the same step, we might choose otherwise because penalty costs could be combined with investment costs and therefore we could consider selecting the slow mover in the spare parts allocation, because the savings in penalty costs could outweigh the additional investment in the initial spare parts allocation as we saw for some methods.

9.3 Verification

We applied our methods to one system within a real-life setting. To verify the principles, one could investigate the effects for other systems, preferably with a greater portion of fast moving items to obtain the desired results. Also, an extension to other industries where spare parts management is of importance would be worthwhile.

10 Bibliography

Abelin, K. (2010). *Modeling Multiple Failures In a Multi-Indenture Inventory System.* Stockholm: Royal Institute of Technology.

Al Hanbali, A. (2013). Reliability Engineering & Maintenance Management. *Maintenance Optimisation*. Enschede, Netherlands: University of Twente.

Al Hanbali, A., & van der Heijden, M. (2013). Interval availability analysis of a two-echelon, multi-item system. *European Journal of Operational Research*, 494-503.

Alfredsson, K. (2001). Life Cycle Costs. *World Pumps*, 28.

Alfredsson, P., & Verrijdt, J. (1999). Modeling Emergency Supply Flexibility in a Two-Echelon Inventory System. *Management Science*, 1416-1431.

Avizienis, A. (1978). Fault-Tolerance: The Survival Attribute of Digital Systems. *Proceedings of the IEEE*, 1109-1125.

Basten, R. (2014, September 9). Maintenance Engineering & Management. *Maintenance in the supply chain*. Enschede, Netherlands: University of Twente.

Baxter, L. (1988). On the Theory of Cannibalization. *Journal of Mathematical Analysis and Applications*, 290-297.

Berger, V., & Christophi, C. (2003). Randomization Technique, Allocation Concealment, Masking, And Susceptibility Of Trials To Selection Bias. *Journal of Modern Applied Statistical Methods*, 80-86.

Caglar, D., Li, C., & Simchi-Levi, D. (2004). Two-echelon spare parts inventory system subject to a service constraint. *IIE Transactions*, 655-666.

Carrasco, J. (2004). Solving large interval availability models using a model transformation approach. *Computers & Operations Research*, 807-861.

Chief of Naval Operations. (2003). *Operational Availability Handbook: A Practical Guide for Military Systems, Sub-Systems, and Equipment.* Washington: Department of the Navy.

Cinlar, E. (1975). *Introduction to Stochastic Processes.* Englewood Cliffs: Printice-Hall.

Coenen, T. (2009). *Penalties and bonuses for service contracts at Thales.* Enschede.

Cohen, M., Agrawal, N., & Agrawal, V. (2006). Winning in the Aftermarket. *Harvard Business Review*, 129-138.

de Souza e Silva, E., & Gail, H. (1986). Calculating Cumulative Operational Time Distributions of Repairable Computer Systems. *IEEE Transactions on Computers*, 322-332.

Doets, C. (2014, November 7). Expert opinion on Thales' products. (W. Sleiderink, Interviewer) Hengelo.

Doets, C. (2014, October 21). Maintenance Engineering & Management Guest Lecture. *ILS at Thales NL*. Enschede, The Netherlands: Thales Nederland.

Driessen, A. (2014, November 20). Expert opinion on Performance-Based Contracts at Thales. (W. Sleiderink, Interviewer) Hengelo.

Durbin, J., & Watson, G. (1950). Testing for serial correlation in least squares regression. *Biometrika*, 409-428.

Gluch, P., & Baumann, H. (2004). The life cycle costing (LCC) approach: a conceptual discussion of its usefulness for environmental decision-making. *Building and Environment*, 571-580.

Goyal, A., & Tantawi, A. (1988). A measure of guaranteed availability and its numerical evaluation. *Computers, IEEE*, 25-32.

Hopp, W., & Spearman, M. (2000). *Factory Physics.* New York: McGraw-Hill.

Huang, K., & Mi, J. (2013). Properties and computation of interval availability of system. *Statistics and Probability Letters*, 1388-1396.

Iyer, B., Donatielle, L., & Heidelberger, P. (1986). Analysis of performability for stochastic models of fault-tolerant systems. *Computers, IEEE*, 902-907.

Kirmani, E., & Hood, C. (2008). A new approach to analysis of interval availability. *Availability, Reliability and Security. Third International Conference on* (pp. 479-483). IEEE.

Law, A. (2006). *Simulation Modeling & Analysis.* McGraw-Hill.

Linz, P. (1985). *Analytical and numerical methods for Volterra equations.* Philadelphia: Siam.

Minner, S. (2014). Inventory Management. Germany: TUM School of Management.

Muckstadt, J. (2005). *Analysis and algorithms for service parts supply chains.* New York: Springer.

Ross, S. (1983). *Stochastic Processes.* New York: Wiley.

Rubino, G., & Sericola, B. (1991). Interval availability analysis using operational periods. *Performance Evaluation*, 257-272.

Rubino, G., & Sericola, B. (1993). Interval availability distribution computation. *Fault-Tolerant Computing, IEEE*, 48-55.

Rustenburg, W. (2000). *A System Approach to Budget-Constrained Spare Parts Management.* Enschede: Febodruk B.V.

Rustenburg, W. (2002). System approaches to spare parts management enable significant and controlled cost reductions. *Proceedings of EuroMaintenance*.

Rustenburg, W. (2005). *Spare parts management: Door differentiatie hoofdzaken van bijzaken onderscheiden.* SpareVision.

Rustenburg, W., van Houtum, G., & Zijm, W. (2000). Systeemgericht spare parts management bij de Nederlands Koninklijke Marine. *Bedrijfskunde*, 1-16.

Sherbrooke, C. (2004). *Optimal Inventory Modeling of Systems: Multi-Echelon Techniques.* Boston: Kluwer Academic Publishers.

Siewiorek, D., & Swarz, R. (1982). *The theory and practice of reliable system design.* Bedford: Digital press.

Simul8. (2015). *Simul8.com*. Retrieved February 18, 2015, from http://simul8.com/support/help/doku.php?id=features:clock:warm_up_period

Smit, M. (2009). *Life Cycle Cost Optimization.* Enschede: University of Twente.

Smith, M. (1997). An Approximation of the Interval Availability Distribution. *Probability in the Engineering and Informational Sciences*, 451-467.

Systecon. (2013). Retrieved November 3, 2014, from www.systecon.se/opus-suite/simlox

Systecon. (2013). *SIMLOX User's reference.*

Takács, L. (1957). On certain sojourn time problems in the theory of stochastic processes. *Acta Mathematica Academiae Scientiarum Hungarica*, 169-191.

ter Hofstede, K. (2014). *Inventri: Budget constrained spares management.* Hengelo: Thales Nederland.

Thales. (2014). *Thales Group*. Retrieved October 9, 2014, from About us: https://www.thalesgroup.com/en/worldwide/about-us

Thales. (2014). *Thales Netherlands*. Retrieved October 9, 2014, from About us: History: https://www.thalesgroup.com/en/netherlands/history

Thales Group. (2014). *Defence - What we do*. Retrieved November 7, 2014, from Above water warfare - Surveillance systems: https://www.thalesgroup.com/en/worldwide/defence/apar-active-phased-array-multifunction-radar

Thales Nederland B.V. (2010). *Integrated Logistic Support - ILS Design Rules for generic use.* Hengelo: Thales Nederland B.V.

Thales Nederland B.V. (2012). *Marketing Report.* Hengelo.

Thales Nederland B.V. (2014). Team Online. Hengelo.

Thierry, M., Salomon, M., van Nunen, J., & van Wassenhove, L. (1995). Strategic Issues in Product Recovery Management. *Cailifornia Management Review*, 114-135.

Tinga, T. (2013). *Principles of Loads and Failure Mechanisms. Applications in Maintenance, Reliability and Design.* London: Springer-Verlag.

van der Heijden, M. (2013). Reverse Logistics & Re-manufacturing. *Spare Part Management*. Netherlands: University of Twente.

van Willigen, S. (2013). *Risk Reduction in Service Contract Fullfillment.* Enschede: University of Twente.

van Zwam, M. C. (2010). *Tactical Decisions to Improve Service Contract Performance.* Enschede: University of Twente.

Wijk, O., & Andersson, P. (2012). Simulation as Support for PBL Contract Design. *eMaintenance.* Luleå: Luleå University of Technology.

Wong, H., Kranenburg, B., van Houtum, G., & Cattrysse, D. (2007). Efficient Heuristics for Two-Echelon Spare Parts Inventory Systems with an Aggregate Mean Waiting Time Constraint Per Local Warehouse. *Or Spectrum*, 699-722.

Wong, H., van Houtum, G., Cattrysse, D., & van Oudheusden, D. (2006). A System Approach for Multi-Location Spare Parts Systems with Lateral Transshipments and Waiting Time Constraints. *European Journal of Operational Research*, 1071-1093.

Woodward, D. (1997). Life cycle costing - theory, information acquisition and application. *International Journal of Project Management*, 335-344.

Ypma, R. (2014). Expert opinion on current maintenance activities. (W. Sleiderink, Interviewer) Hengelo.

Zijm, W. (2012). *Manufacturing and Logistic Systems Analysis, Planning and Control.*

11 Appendices

11.1 Appendix A: (VARI)-METRIC Approach (Sherbrooke, 2004) (van der Heijden, 2013)

11.1.1 Metric

The Multi-Echelon Technique for Recoverable Item Control (METRIC) is developed for first-indenture items (i.e. parts that are installed directly on the base), which are usually repairable (or recoverable) and they tend to be expensive with a low demand at any particular base. The one-for-one repair at the base or the resupply from the depot simplifies the mathematics of the base-depot joint optimization problem, which can be addressed with results from the queuing theory literature.

Assumptions

The METRIC model calculates for every item on a system the optimal stock level at each of several bases. The objective function is to minimize the sum of backorders across all bases which is equivalent to maximizing availability when there is no cannibalization (i.e. consolidation of item shortages on the smallest number of end-items). Below the key assumptions in the METRIC model will be listed:

1. *The decision where to repair an item (depot, base) is solely determined by technical factors (e.g. availability of technical equipment of specialized personnel) and not by stock levels or repair shop workload.*
2. *Base resupply from depot only (no lateral supply).*
3. *The (s-1, s) inventory policy is appropriate for every item at every echelon level.* This means that units of an item are not batched for repair and that any scrapped units are reprocured on a one-for-one basis, because the demand rates are sufficiently low and the costs are sufficiently high that the economic order quantity is close to one (Minner, 2014).
4. *Optimal steady-state stock levels are determined.* Over some period of time in the future the operating hours of the system will remain fairly consistent.
5. *The demand is Poisson distributed.* This means that the inter-arrival times between demands are independent and that there is a continuous demand such that items will continue failing, even if the system is down.
6. *All demand that is not filled is backordered.*
7. *All backorders are equally important.*
8. *Repair costs are not needed, because if an item can be repaired, its repair cost is always less than its purchase cost.*
9. *Order and holding costs are not needed, because one-for-one replenishment is assumed and this defines the number of orders and the average stock on hand.*

Nomenclature

The following symbols and definitions are used in the METRIC model:

j = location index (j=0: depot; j=1..N: base)

m_{ij} = mean demand per year for item i at location j

T_{ij} = mean repair throughput time of item i at location j

O_j = time to order and ship an item from the depot to base j

r_{ij} = probability that an item i can be repaired at base j

X_{ij} = number of type i items in the pipeline at location j, with mean μ_{ij}

EBO_{ij} = Expected number of backorders for item i at location j

Marginal approach

First, determine mean demand and backorders at the depot using Palm's theorem.

- Demand at depot:
  $$m_{i0} = \sum_{j=1}^{M}(1 - r_{ij})\, m_{ij} \tag{10.1}$$
- Backorders at depot:
  $$EBO_{i0}(s_{i0}) = \sum_{n=s_{i0}+1}^{\infty}(n - s_{i0}) * \frac{(m_{io}T_{i0})^n e^{-m_{io}T_{io}}}{n!} \tag{10.2}$$

Subsequently, we can derive the mean number of items in the pipeline and thus the backorders at the base.

- Average pipeline:
  $$\mu_{ij} = m_{ij}\left\{r_{ij}T_{ij} + \left(1 - r_{ij}\right)\left(O_j + \frac{EBO_{i0}(s_{i0})}{m_{i0}}\right)\right\} \tag{10.3}$$
- Backorders at base:
  $$EBO_{ij}\left(s_{ij}\right) = \sum_{n=s_{ij}+1}^{\infty}\left(n - s_{ij}\right) * \frac{{\mu_{ij}}^n e^{-\mu_{ij}}}{n!} \tag{10.4}$$

By minimizing the sum of these backorders at the bases, through adding stock at the bases or at the depot, we can equivalently maximize the availability at the bases.

Drawbacks

When we compute the expected backorders for a certain stock level s and compare it to the previous iteration with stock level s-1, we can compute a certain backorder reduction to determine at which location we add stock. However, the backorder reductions are not monotonically decreasing, which means that the backorder function is not convex at those points. This phenomenon called convexification can be solved by eliminating these non-convex points to prevent a flush out.

11.1.2 VARI-METRIC

During the development of METRIC it was known that the backorders at the base were understated. The biggest complication is that the depot pipeline has indeed a Poisson distribution, but this does **not** hold for the base pipeline, it depends on the depot backorders. Generally, the variance-to-mean ratio is greater than 1 compared to a Poisson distribution where this ratio equals 1. A solution is proposed to calculate the pipeline variance and use an approximate probability distribution for the pipeline having the mean and variance as calculated.

Indenture levels

The Line replaceable unit (LRU) is the first indenture item that is replaced in the system. When the failed LRU is taken to the maintenance shop, the second indenture items can be replaced, the Shop Replaceable Unit (SRU). An important assumption in this is that the failure of an LRU is caused by the failure of at most one SRU with a certain probability q.

Approach

In the case of no commonality (specific SRUs cannot be used on more than one LRU), the problem is separable per LRU, so we will describe the optimization of 1 LRU with its SRUs in a multi-echelon network.

First, we derive the demand rates per location and per item based on the mean demand, repair probabilities and the failure probabilities as can be seen in Figure 24 below.

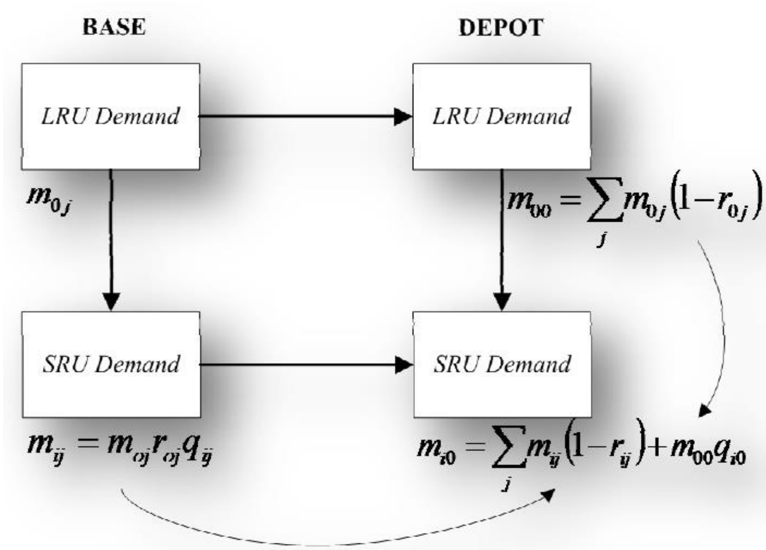

Figure 24: Deriving demand rates for VARI-METRIC (van Zwam, 2010)

For the pipeline calculations, we will proceed in the exact opposite direction for the 4 categories. Three additional parameters are used in here referring to (1) the fraction of depot demand for SRU i due to depot LRU repairs, (2) the fraction of all demand at the depot for SRU i that is being resupplied to base j and (3) the fraction of the depot demand m_{oo} for a LRU that comes from base j.

- (1):
$$f_{i0} = \frac{m_{00} q_{i0}}{m_{i0}} \tag{10.5}$$
- (2):
$$f_{ij} = \frac{m_{ij}(1-r_{ij})}{m_{i0}} \tag{10.6}$$

- (3):
$$f_{0j} = \frac{m_{0j}(1-r_{ij})}{m_{i0}} \tag{10.7}$$

SRU Pipeline at depot

At this moment there is no influence from other items or locations, so the pipeline has a Poisson distribution with mean $m_{i0}T_{i0}$. Therefore, we can easily calculate the mean and variance of the SRU i backorders at the depot: $EBO_{i0}(s_{i0})$ and $VBO_{i0}(s_{i0})$.

LRU pipeline at depot

The mean and variance of the LRU pipeline depends on the SRU backorders at the depot. A SRU backorder has a probability f_{i0} that it delays a LRU repair at the depot and a probability of $1 - f_{i0}$ that it delays a resupply to some base.

- Expected pipeline:
$$E[X_{00}] = m_{00}T_{00} + \sum_{i=1}^{I} f_{i0}EBO(s_{i0}|m_{i0}T_{i0}) \tag{10.8}$$
- Variance pipeline:
$$Var[X_{00}] = m_{00}T_{00} + \sum_{i=1}^{I} f_{i0}(1-f_{i0})EBO(s_{i0}|m_{i0}T_{i0}) + \sum_{i=1}^{I} f_{i0}^{\,2}\, VBO(s_{i0}|m_{i0}T_{i0}) \tag{10.9}$$

SRU Pipeline at base

The mean and variance of the SRU base pipeline also depends on the SRU backorders at the depot.

- Expected pipeline:
$$E[X_{ij}] = m_{ij}[r_{ij}T_{ij} + (1-r_{ij})O_i] + f_{ij}EBO(s_{i0}|m_{i0}T_{i0}) \tag{10.10}$$
- Variance pipeline:
$$Var[X_{ij}] = m_{ij}[r_{ij}T_{ij} + (1-r_{ij})O_i] + f_{ij}(1-f_{ij})EBO(s_{i0}|m_{i0}T_{i0}) + f_{ij}^{\,2}VBO(s_{i0}|m_{i0}T_{i0} \tag{10.11}$$

LRU Pipeline at base

The mean and variance of the LRU at the base depends on both the SRU base backorders and the LRU depot backorders. All SRU base backorders arise from LRU demand at that base. We can use the previous results as input for the latest computations of the pipeline to obtain the availability per base.

- Expected pipeline:
$$E[X_{0j}] = m_{0j}[r_{0j}T_{0j} + (1-r_{0j})O_0] + f_{0j}EBO_{00}(s_{00}|E[X_{00}], Var[X_{00}]) + \sum_{i=1}^{I} EBO(s_{ij}|E[X_{ij}], Var[X_{ij}]) \tag{10.12}$$
- Variance pipeline:
$$\begin{aligned} Var[X_{0j}] &= m_{0j}[r_{0j}T_{0j} + (1-r_{0j})O_0] \\ &+ f_{0j}(1-f_{0j})EBO_{00}(s_{00}|E[X_{00}], Var[X_{00}]) \\ &+ f_{0j}^{\,2}VBO(s_{00}|E[X_{00}], Var[X_{00}]) \\ &+ \sum_{i=1}^{I} VBO(s_{ij}|E[X_{ij}], Var[X_{ij}]) \end{aligned} \tag{10.13}$$

Then, finally the availability at base j due to expected backorders on the LRU and its SRUs is given by:

- Availability:

$$A_j = 100\{1 - EBO(s_{0j}|E[X_{0j}], Var[X_{0j}])/(N_j Z_0)\}^{Z_0} \tag{10.14}$$

where N_j is the number of systems at base j and Z_0 is the number of applications of the LRU on the system.

11.2 Appendix B: After sales business models (Cohen, Agrawal, & Agrawal, 2006)

Models of After-Sales Services

The value companies place on after-sales services will determine the business models that firms can use to deliver them. When services are all-important, manufacturers may choose to sell services rather than the products that generate them.

SERVICE PRIORITY	BUSINESS MODEL	TERMS	EXAMPLE	PRODUCT OWNER
None	Disposal	Dispose of products when they fail or need to be upgraded	Razor blades	Consumer
Low	Ad hoc	Pay for support as needed	TVs	Consumer
Medium-high	Warranty	Pay fixed price as needed	PCs	Consumer
Medium-high	Lease	Pay fixed price for a fixed time; option to buy product	Vehicles	Manufacturer; leasing company
High	Cost-plus	Pay fixed price based on cost and prenegotiated margin	Construction	Customer
Very high	Performance based	Pay based on product's performance	Aircraft	Customer
Very high	Power by the hour	Pay for services used	Aircraft engines	Manufacturer; service provider

Figure 25: After sales business models (Cohen, Agrawal, & Agrawal, 2006)

11.3 Appendix C: Interval availability and survival function (Al Hanbali & van der Heijden, 2013)

We will start by computing the first two moments of the interval availability. The expected interval availability E[A(T)] in [0,T] is equal to the steady-state availability of the finite-state three-dimensional approximate Markov chain (AMC) is given by:

$$E[A(T)] = \prod_{j=1}^{M} \sum_{n=0}^{s_j} \sum_{l=0}^{M+s_{0j}+\sum_{l=1}^{M} s_{lj}} \sum_{m=0}^{\min(s_j-n,(l-s_{0j})^+)} \pi_{m,n,l}(j) \tag{10.15}$$

Let $\pi_{m,n,l}(j)$ denote the steady state probability distribution vector of AMC, j is the item index for which holds j = 1,…,M and s_j is the stock level of item j in the tagged system.

Before reporting on the variance of A(T), let us introduce some additional notation. Let γ_j denote a row vector of size equal to the cardinality of the state space Ω_j. The vector γ_j is obtained from the steady state probability vector $\pi(j)$ of AMC by replacing the equilibrium probability of the malfunctioning states with zero. Let f_j denote a column vector of size equal to the cardinality of the state space Ω_j. The non-zero entries of f_j are equal to one and they represent the operational states. Then the variance of the system interval availability in [0,T] is given by:

$$Var[A(T)] = 2\sum_{n=1}^{\infty} e^{-\nu T} \frac{(\nu T)^n}{(n+2)!} \sum_{i=1}^{n} (n-i+1) \prod_{l=1}^{M} \gamma_l (P_l)^i f_l + 2E[A(T)] * \frac{e^{-\nu t} + \nu T - 1}{(\nu T)^2} - E[A(T)]^2 \tag{10.16}$$

For the computation of the probability that the interval availability is equal to one we need θ_j to be the row vector that only consists of the steady state probabilities of the operational states of AMC, then:

$$P\big(A(T)\big) = 1) = e^{-T\sum_{i=1}^{M} \nu_i} \prod_{j=1}^{M} \theta_j \sum_{n=0}^{\infty} \frac{(\nu_i T)^n}{n!} (P_j{}^a)^n e \tag{10.17}$$

Now we can compute the survival function of the interval availability given by:

$$P(A(T) \geq y) = (1 - P(A(T) = 1) \int_y^1 f(x;\alpha;\beta)dx + P(A(T) = 1) \tag{10.18}$$

With

$$\alpha = \frac{(1-E[B]) * E[B]^2}{Var[B]} - E[B] \tag{10.19}$$

$$\beta = \alpha\left(\frac{1}{E[B]} - 1\right) \tag{10.20}$$

$$f(x;\alpha;\beta) = \frac{1}{B(\alpha,\beta)} x^{\alpha-1}(1-x)^{\beta-1} \tag{10.21}$$

Where $B(\alpha,\beta)$ is the beta function. For proofs and more details we refer to the aforementioned article.

11.4 Appendix D: Variance of interval availability and its common coefficient of variation (Kirmani & Hood, 2008)

The 2-CTMC model may be described as follows. Let U_i denote the duration of the *i*th functioning period and D_i the duration of the *i*th downtime (non-functioning period). Assume that U_i, i= 1,2… are independently and identically distributed with common exponential distribution of mean 1/λ, and D_i, i= 1,2… are independently and identically distributed with common exponential distribution of mean 1/μ.

In their study, the authors prove that for a 2-CTMC model the following formula for the variance of interval availability (I(T)) holds:

$$Var\{I(T)\} = 2pqR\big((\lambda+\mu)T\big) - q^2R^2\big((\lambda+\mu)T\big) + 2q(1-3p)S\big((\lambda+\mu)T\big) \quad (10.22)$$

Where

$$p = \frac{\mu}{\lambda+\mu}, \quad q = \frac{\lambda}{\lambda+\mu} = 1-p \quad (10.23)$$

(Note that p can be considered as the system availability, since this is the downtime divided by the downtime plus uptime.)

And

$$R\big((\lambda+\mu)T\big) = \frac{1-exp\{-(\lambda+\mu)T\}}{(\lambda+\mu)T}, S\big((\lambda+\mu)T\big) = \frac{1-exp\{-(\lambda+\mu)T\} - (\lambda+\mu)Texp\{-(\lambda+\mu)T\}}{\{(\lambda+\mu)T\}^2} \quad (10.24)$$

To gain proper insights in the extent of fluctuations in the interval availability and thus the variance, the authors use the common coefficient of variation, defined as:

$$\phi(T) = \frac{\sqrt{Var\{I(T)\}}}{E\{I(T)\}} \quad (10.25)$$

Obviously, high values of this coefficient mean highly fluctuating interval availability and lower values lead to more stable interval availability.

11.5 Appendix E: Review periods

Table 2: Categories for review periods

Review period	Time in hours
Monthly	730
Quarterly	2190
Yearly	8760
2-yearly	17520
5-yearly	43800

11.6 Appendix F: Details for simulation of multi-indenture model

Table 3: Item characteristics for basic multi-indenture model

Item characteristics				
Item	**Failure Rate**	**Repair Time**	**Backorder Ratio**	**Mover Ratio**
80	7,416193876	7416,19	55000	0,001
80A	3,708096938	3708,10	13750,0	0,001
80B	3,708096938	3708,10	13750,0	0,001
83	741,6198487	74,16	55000,0	10
83A	370,8099244	37,08	13750,0	10
83B	370,8099244	37,08	13750,0	10

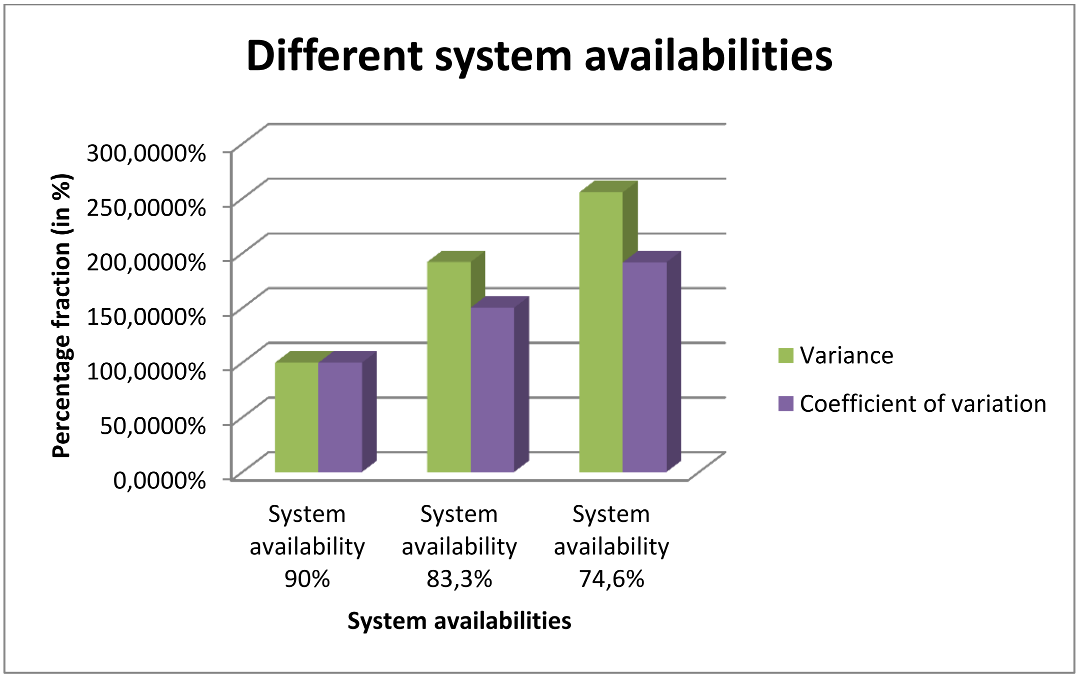

Figure 26: Different steady-state availabilities for multi-indenture model

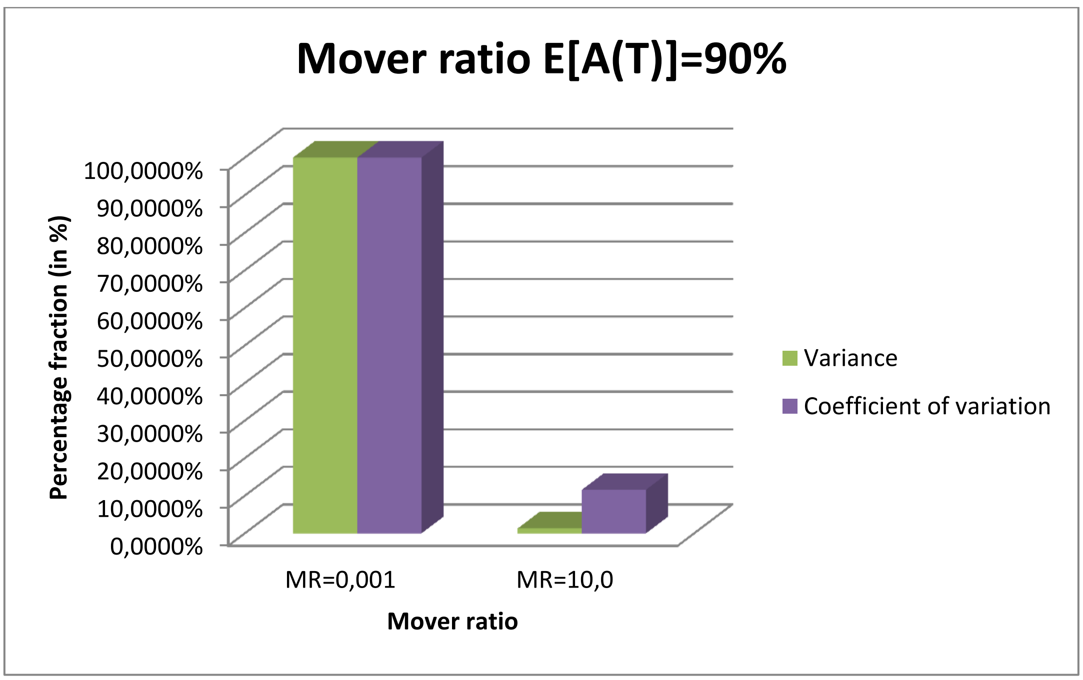

Figure 27: Effect of mover ratio in multi-indenture model

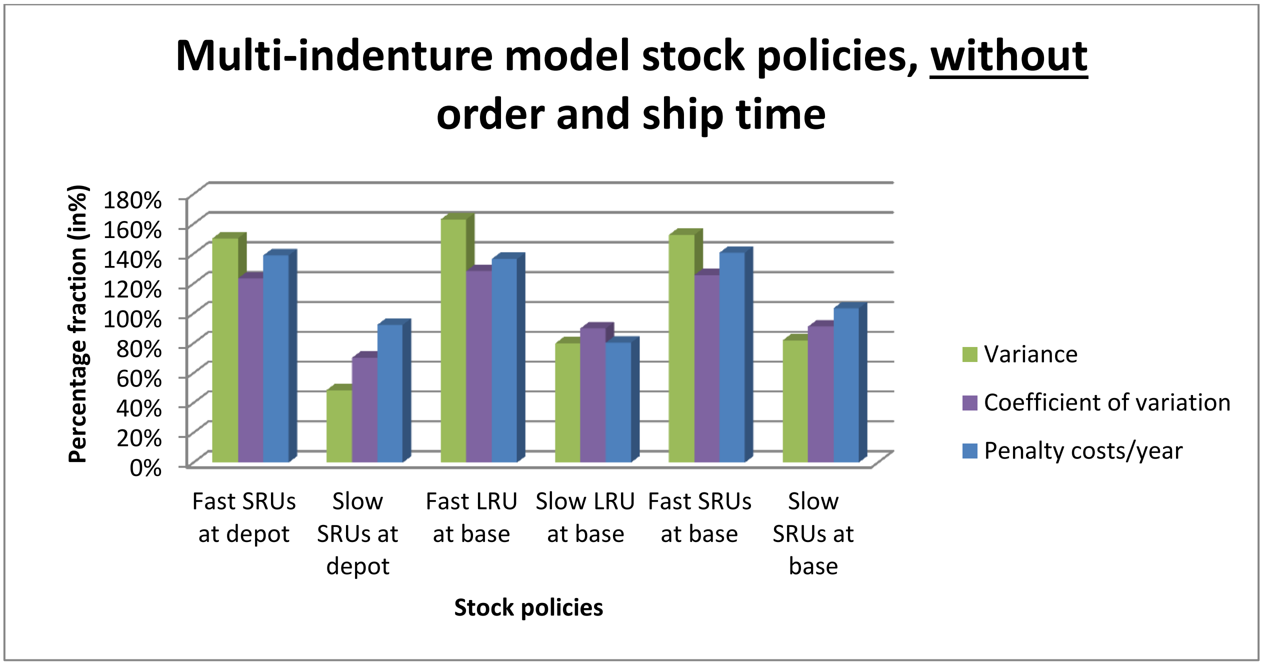

Figure 28: Multi-indenture stock policies without order and ship times

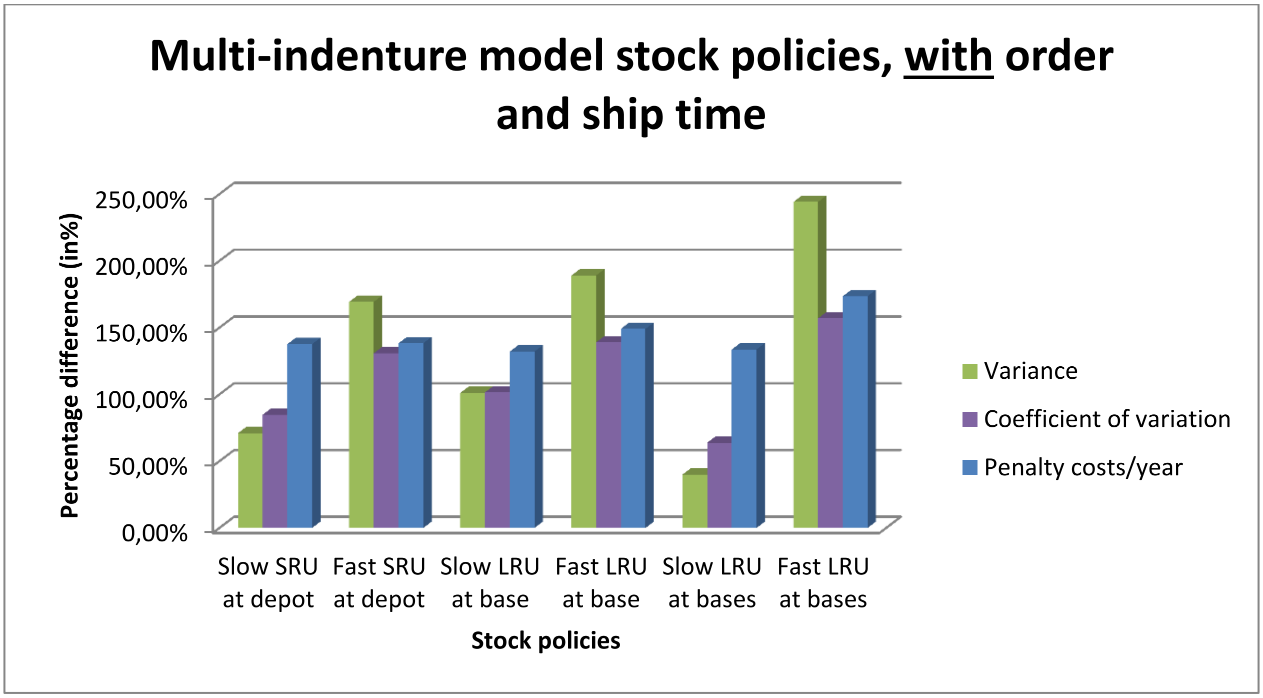

Figure 29: Multi-indenture stock policies with order and ship times

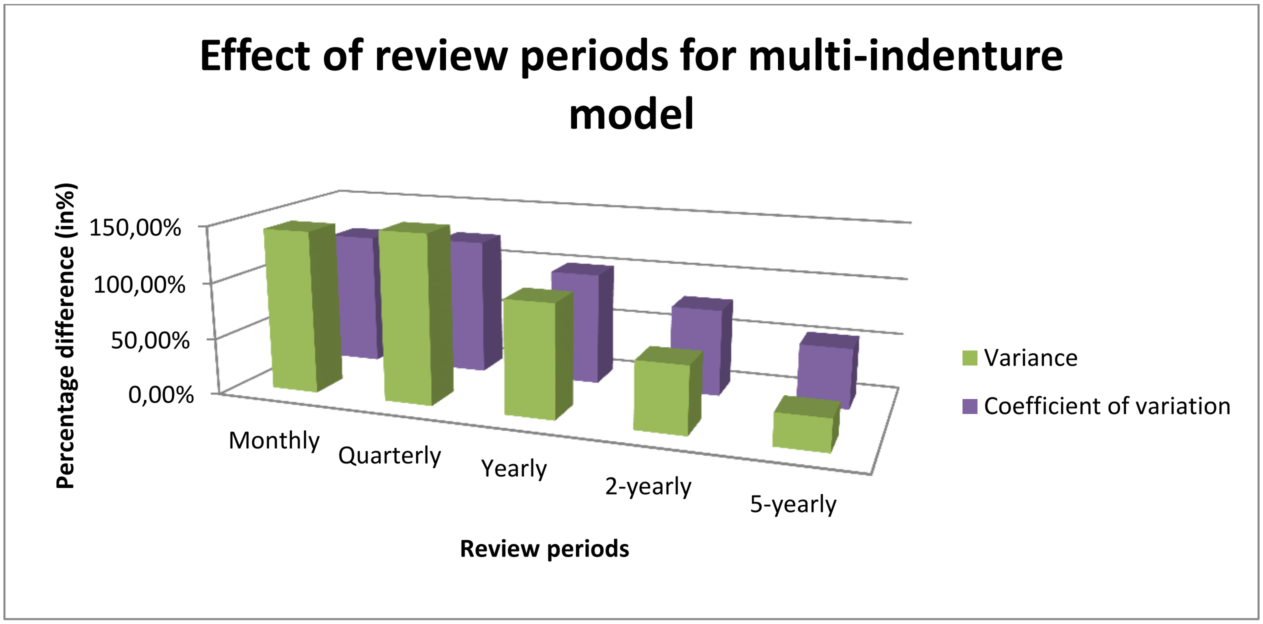

Figure 30: Review periods for multi-indenture model

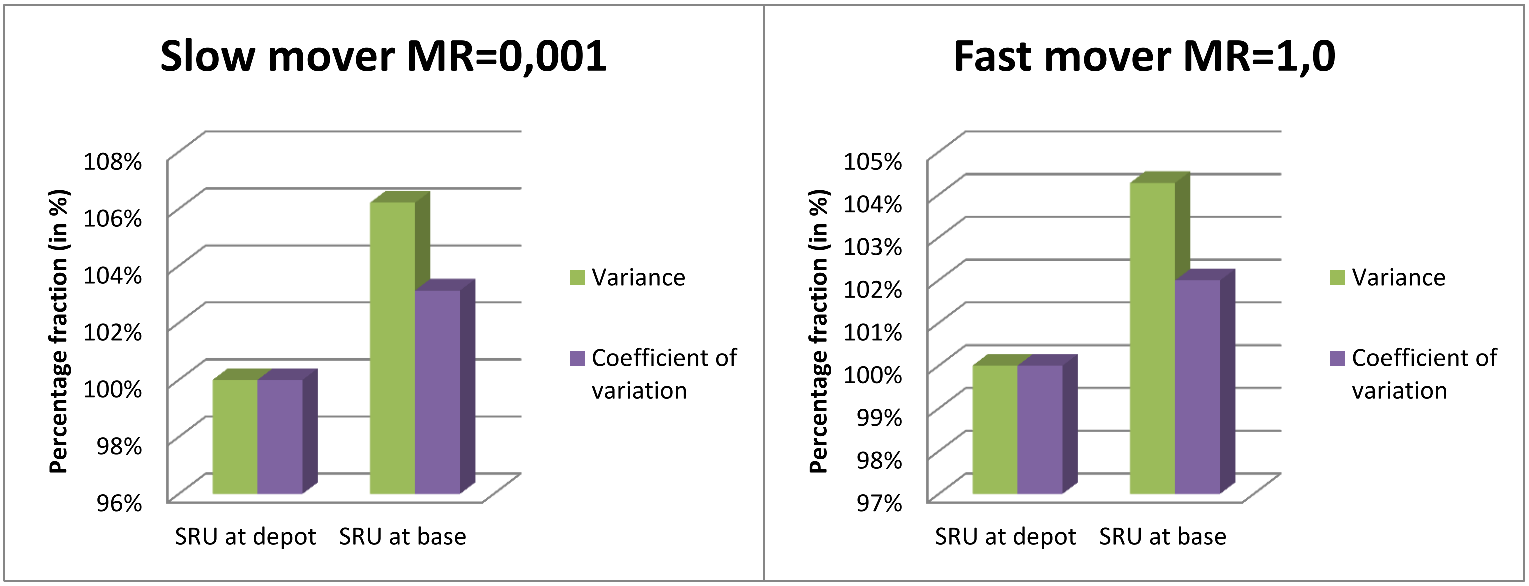

Figure 31: Central stock investment versus local stock investment for multi-indenture model

(Note that for these experiments we simulated for 1000 years.)

11.7 Appendix G: Details for simulation of experiment 1

Table 4: Item characteristics for slow mover for basic model, E[A(T)]=90%

Item characteristics				
Item	**Failure Rate**	**Repair Time**	**Backorder Product**	**Mover Ratio**
LRU 1	7,416199109	7416,20	55000,0	0,001
LRU 2	7,416199109	7416,20	55000,0	0,001

Table 5: Item characteristics for fast mover for basic model, E[A(T)]=90%

Item characteristics				
Item	**Failure Rate**	**Repair Time**	**Backorder Product**	**Mover Ratio**
LRU 1	741,6198429	74,1619863	55000,0	10,000000
LRU 2	741,6198429	74,1619863	55000,0	10,000000

Table 6: Item characteristics for slow mover for basic model, E[A(T)]=83,3%

Item characteristics				
Item	**Failure Rate**	**Repair Time**	**Backorder Product**	**Mover Ratio**
LRU 1	10	10000	100000,0	0,001000000
LRU 2	10	10000	100000,0	0,001000000

Table 7: Item characteristics for fast mover for basic model, E[A(T)]=83,3%

Item characteristics				
Item	**Failure Rate**	**Repair Time**	**Backorder Product**	**Mover Ratio**
LRU 1	1000	100	100000,0	10,000000
LRU 2	1000	100	100000,0	10,000000

Table 8: Item characteristics for slow mover for basic model, E[A(T)]=74,6%

Item characteristics				
Item	**Failure Rate**	**Repair Time**	**Backorder Product**	**Mover Ratio**
LRU 1	13,03840408	13038,40408	170000,0	0,001000000
LRU 2	13,03840408	13038,40408	170000,0	0,001000000

Table 9: Item characteristics for fast mover for basic model, E[A(T)]=74,6%

Item characteristics				
Item	**Failure Rate**	**Repair Time**	**Backorder Product**	**Mover Ratio**
LRU 1	1303,840408	130,3840408	170000,0	10,000000
LRU 2	1303,840408	130,3840408	170000,0	10,000000

Table 10: Inventory characteristics for basic model

Inventory characteristics			
Item	**Stock at depot**	**Stock at base 1**	**Stock at base 2**
LRU 1	0	0	0
LRU 2	0	0	0

Table 11: Model characteristics for basic model

Model characteristics				
Number of systems	**Mission profile**	**Order and ship time**	**Repair time distribution**	**Result collection interval**
2	CONTINUOUS	NONE	DETERMINISTIC	YEARLY

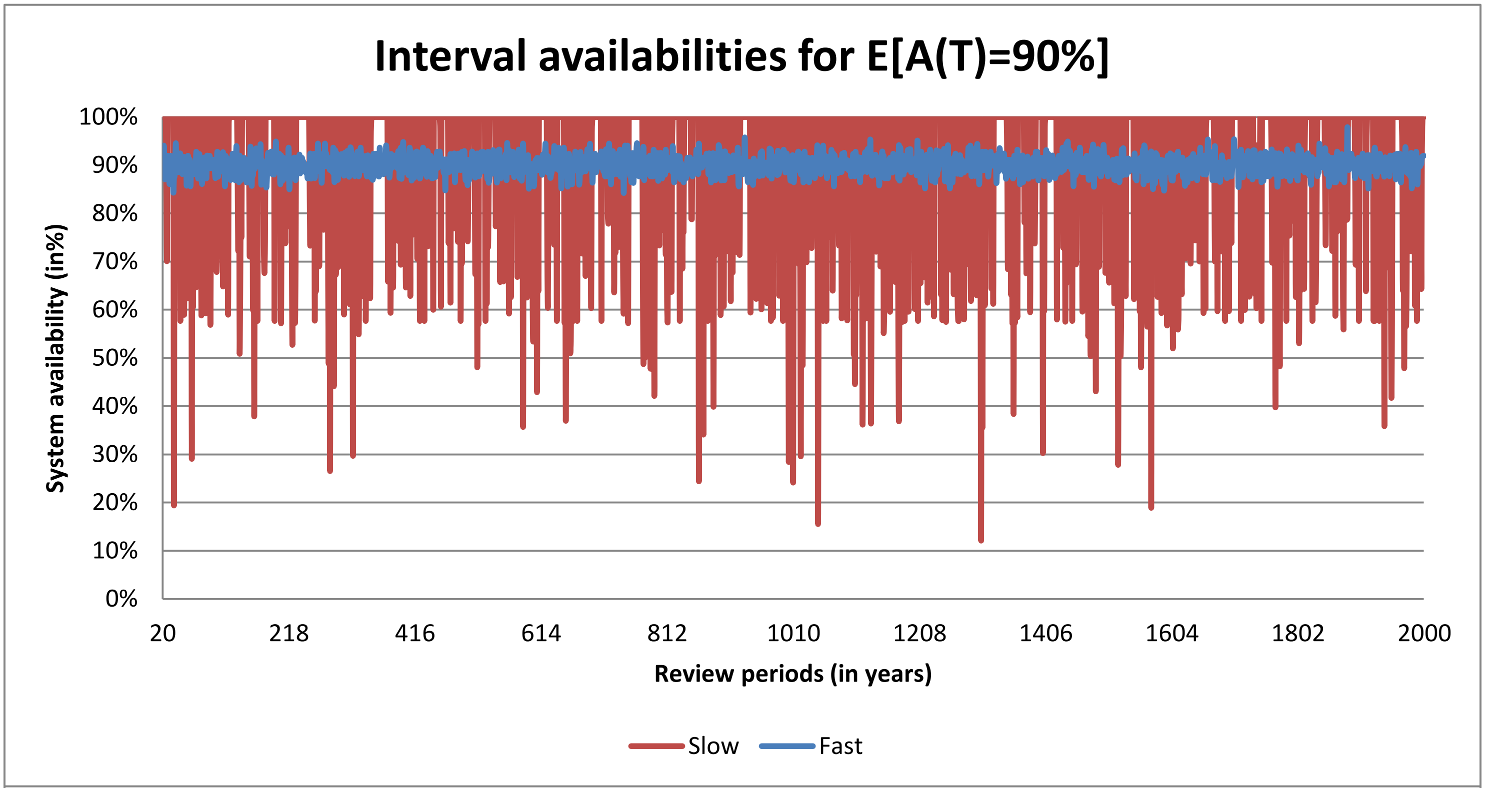
Interval availabilities for E[A(T)=90%]
System availability (in%)
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
20
218
416
614
812
1010
1208
1406
1604
1802
2000
Review periods (in years)
Slow
Fast

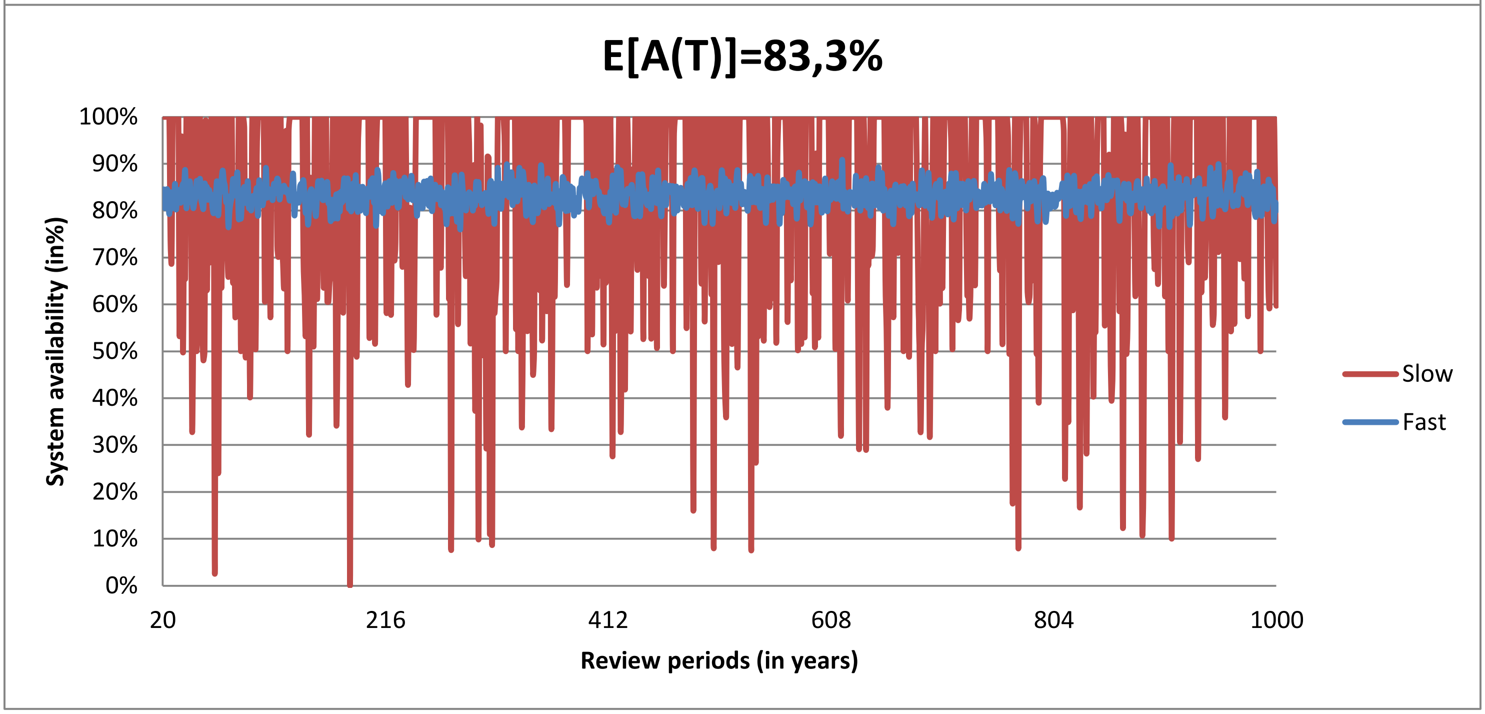
E[A(T)]=83,3%
System availability (in%)
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
20
216
412
608
804
1000
Review periods (in years)
Slow
Fast

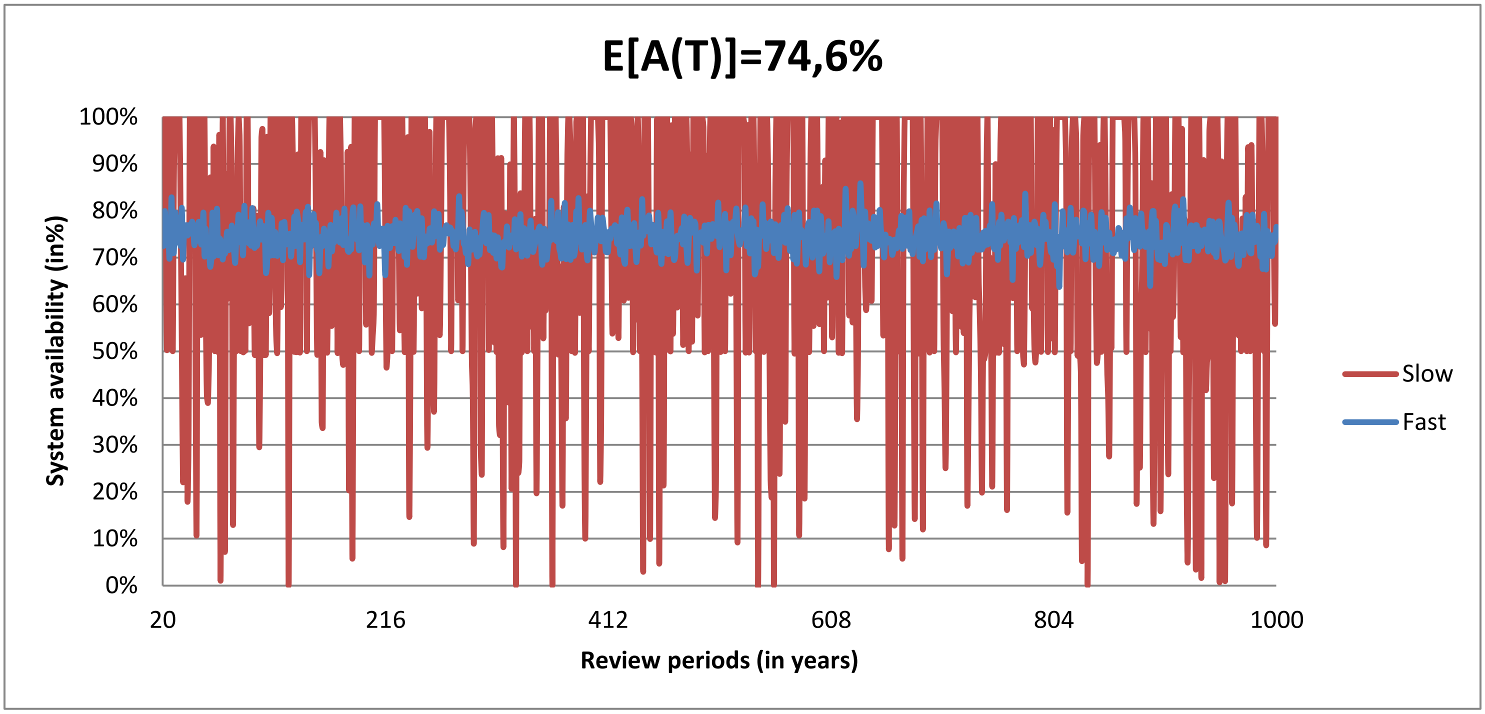

Figure 32: Simulation results of experiment 1

(Note that for some experiments with lower steady-state availability we simulated for 1000 years.)

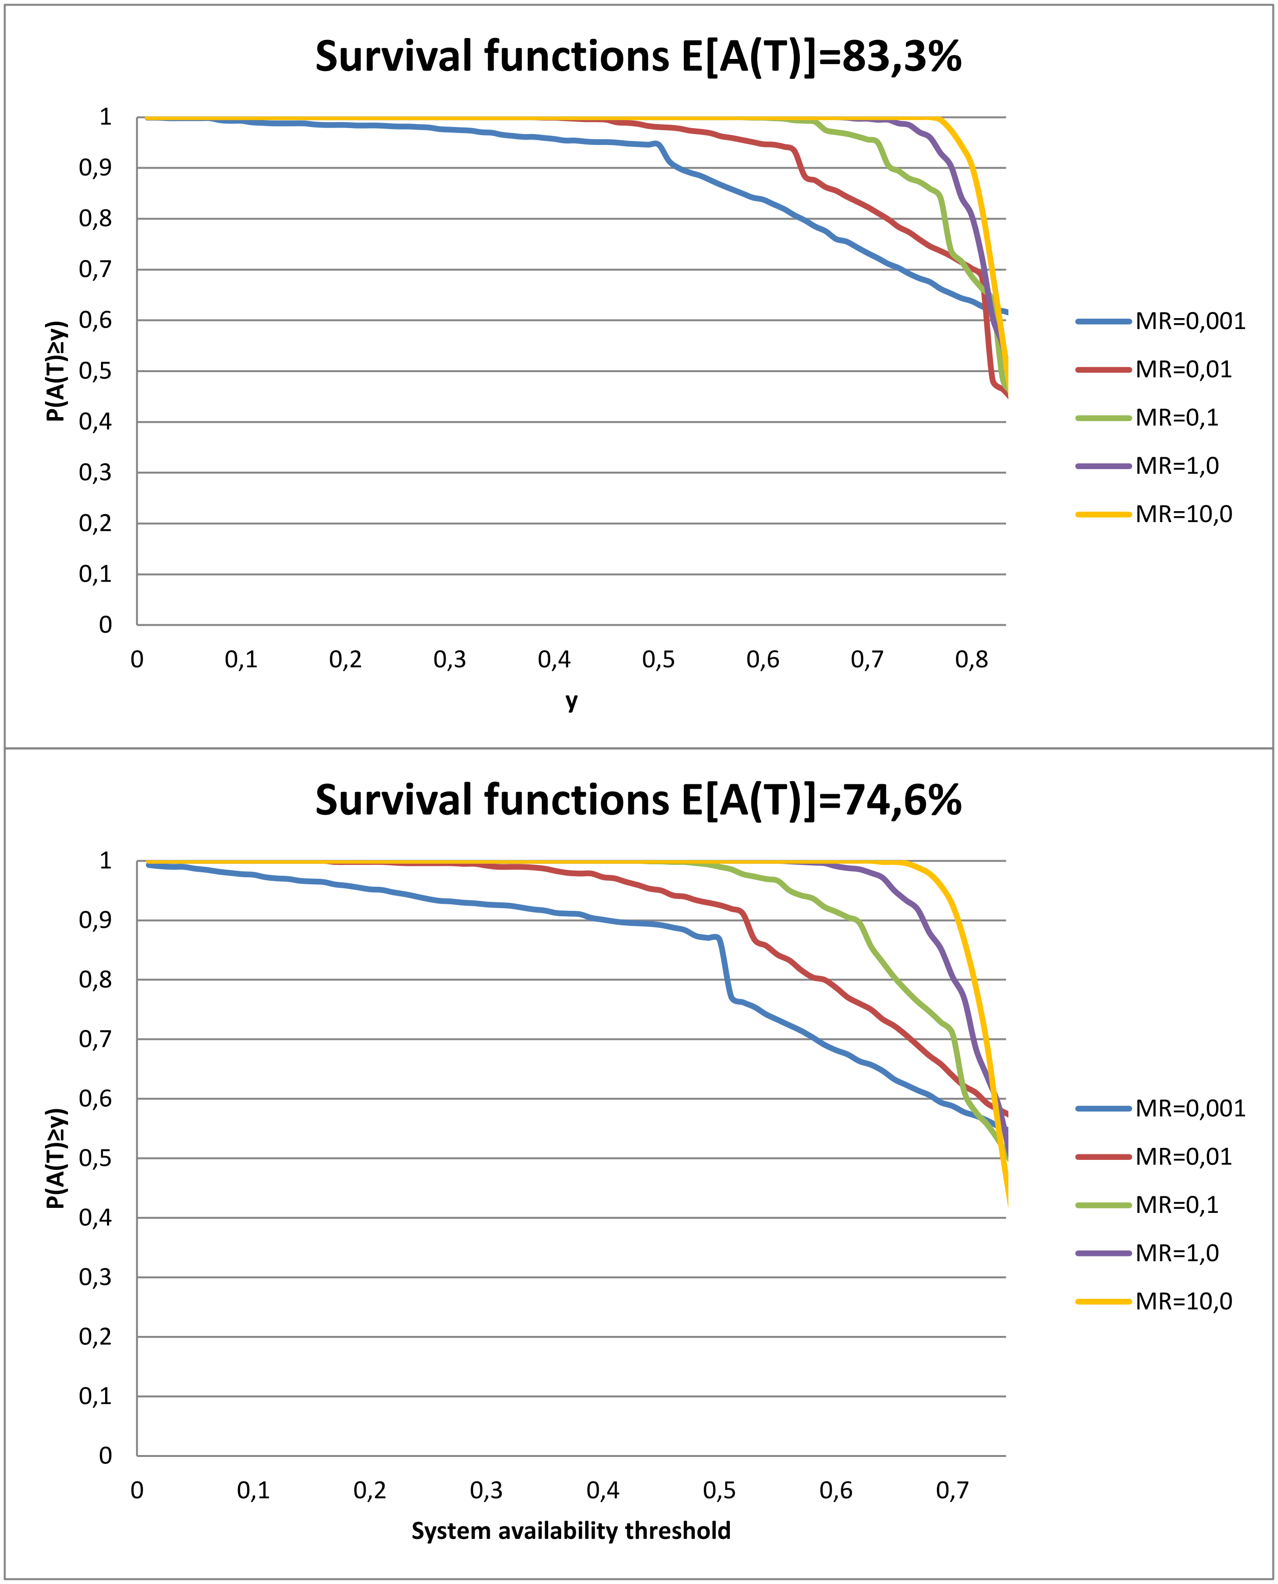

Figure 33: Survival functions of experiment 1

11.8 Appendix H: Details for simulation of experiment 2

Table 12: Item characteristics for slow and fast mover, E[A(T)]=90%

Item characteristics				
Item	**Failure Rate**	**Repair Time**	**Backorder Product**	**Mover Ratio**
LRU 1	10,12422743	10124,23	55000,0	0,00100000
LRU 2	1012,422802	101,24	55000,0	10,000000

Table 13: Inventory characteristics in case of fast mover at depot

Inventory characteristics			
Item	**Stock at depot**	**Stock at base 1**	**Stock at base 2**
LRU 1	0	0	0
LRU 2	1	0	0

Table 14: Inventory characteristics in case of slow mover at depot

Inventory characteristics			
Item	**Stock at depot**	**Stock at base 1**	**Stock at base 2**
LRU 1	1	0	0
LRU 2	0	0	0

Table 15: Inventory characteristics in case of fast mover at base

Inventory characteristics			
Item	**Stock at depot**	**Stock at base 1**	**Stock at base 2**
LRU 1	0	0	0
LRU 2	0	1	0

Table 16: Inventory characteristics in case of slow mover at base

Inventory characteristics			
Item	**Stock at depot**	**Stock at base 1**	**Stock at base 2**
LRU 1	0	1	0
LRU 2	0	0	0

Table 17: Inventory characteristics in case of fast movers at bases

Inventory characteristics			
Item	**Stock at depot**	**Stock at base 1**	**Stock at base 2**
LRU 1	0	0	0
LRU 2	0	1	1

Table 18: Inventory characteristics in case of slow movers at bases

Inventory characteristics			
Item	**Stock at depot**	**Stock at base 1**	**Stock at base 2**
LRU 1	0	1	1
LRU 2	0	0	0

Table 19: Model characteristics

Model characteristics				
Number of systems	**Mission profile**	**Order and ship time**	**Repair time distribution**	**Result collection interval**
2	CONTINUOUS	NONE	DETERMINISTIC	YEARLY

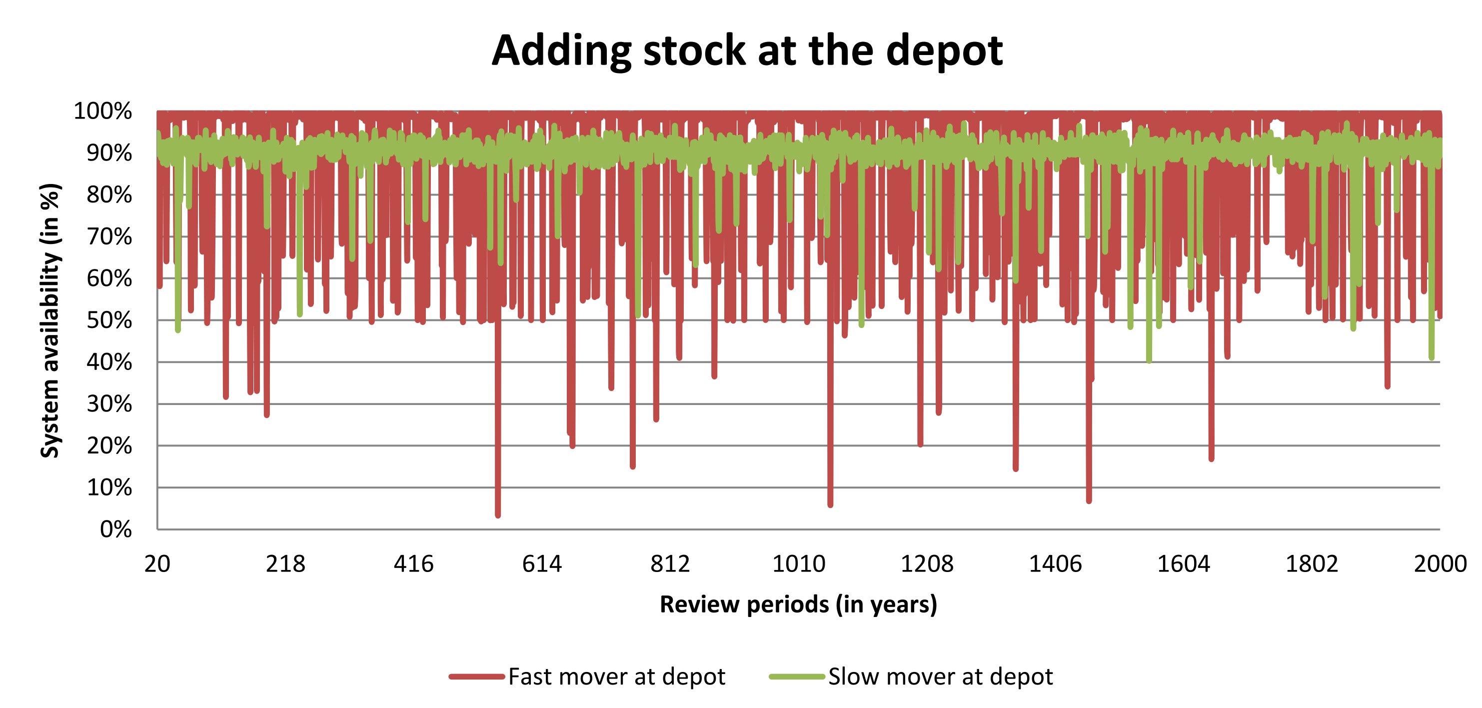

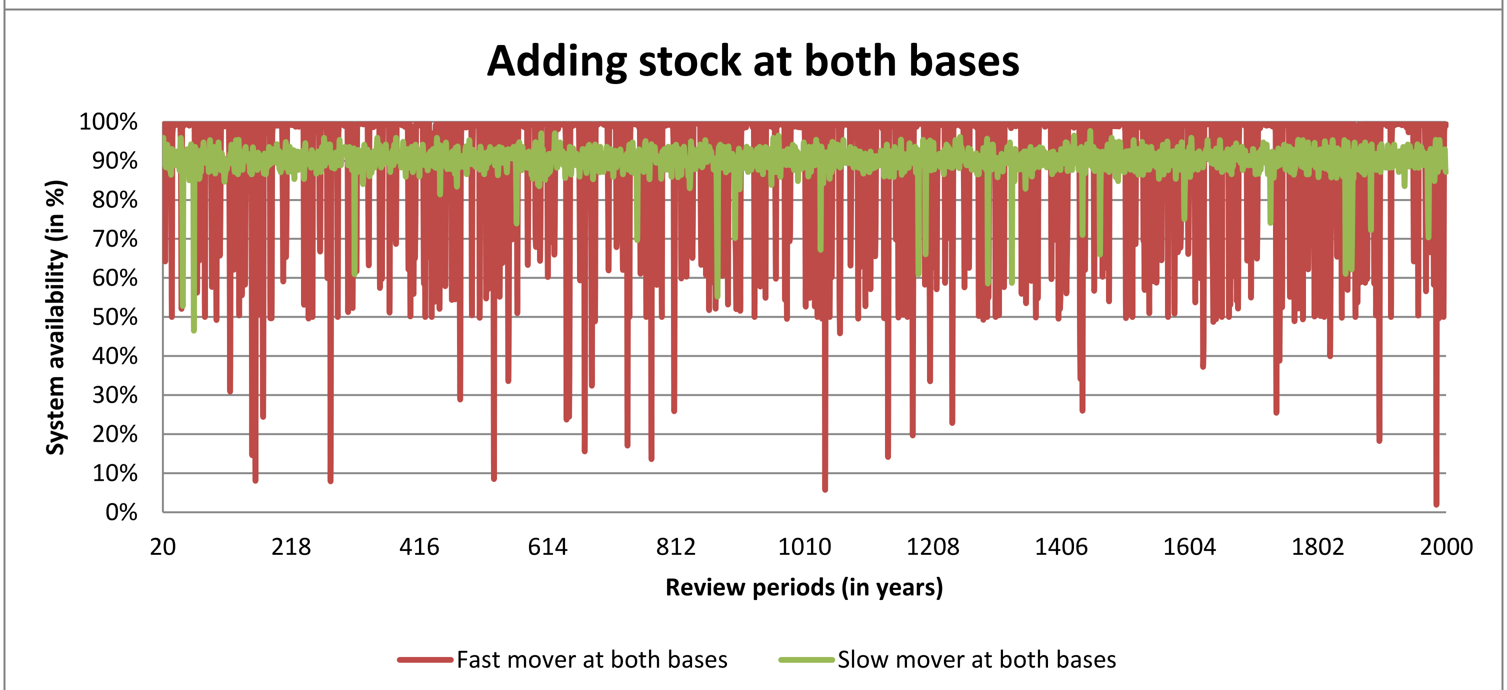

Figure 34: Interval availabilities: adding stock at the depot

11.9 Appendix I: Details for simulation of experiment 3

Table 20: Item characteristics for fast mover

Item characteristics				
Item	**Failure Rate**	**Repair Time**	**Backorder Ratio**	**Mover Ratio**
80	51	51,00	2601	1,000000
83	51	51,00	2601	1,000000

Table 21: Model characteristics for fast mover

Model characteristics				
Number of systems	**Mission profile**	**Order and ship time**	**Repair time distribution**	**Result collection interval**
2	CONTINUOUS	700	DETERMINISTIC	YEARLY

Table 22: Item characteristics for slow mover

Item characteristics				
Item	**Failure Rate**	**Repair Time**	**Backorder Ratio**	**Mover Ratio**
80	2,828427031	2828,43	8000	0,001000
83	2,828427031	2828,43	8000	0,001000

Table 23: Model characteristics for slow mover

Model characteristics				
Number of systems	**Mission profile**	**Order and ship time**	**Repair time distribution**	**Result collection interval**
2	CONTINUOUS	11800	DETERMINISTIC	YEARLY

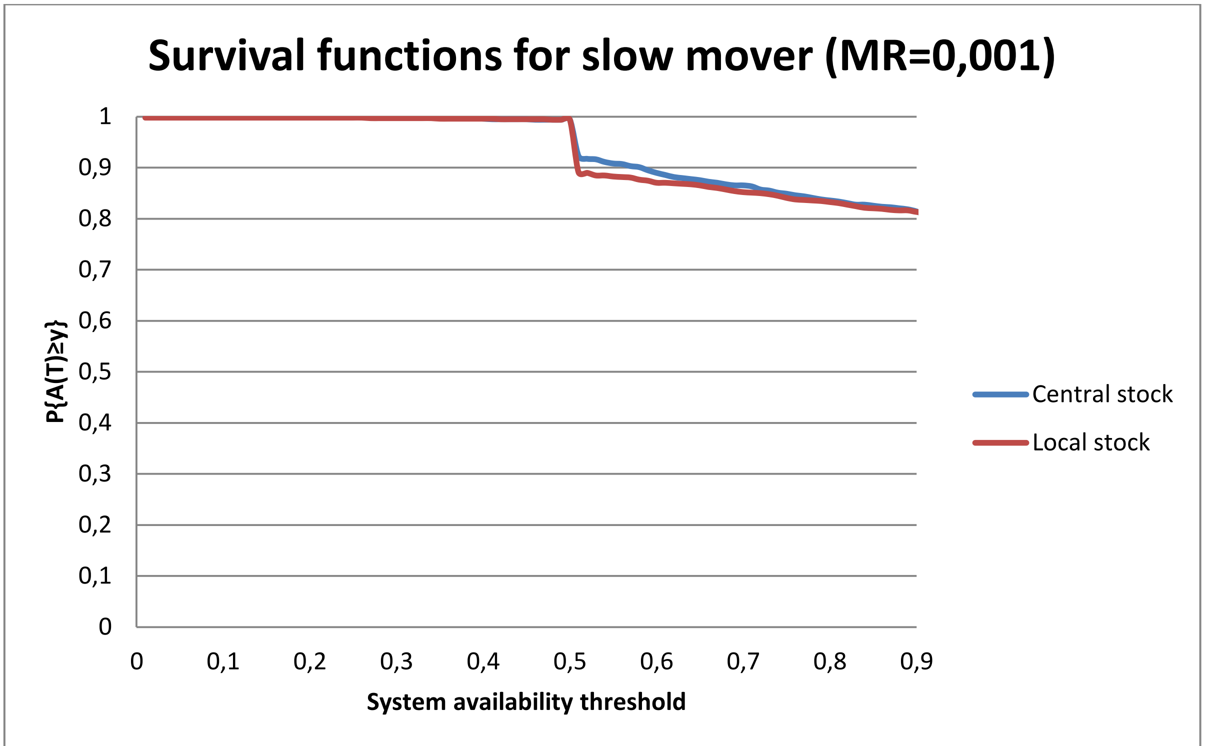

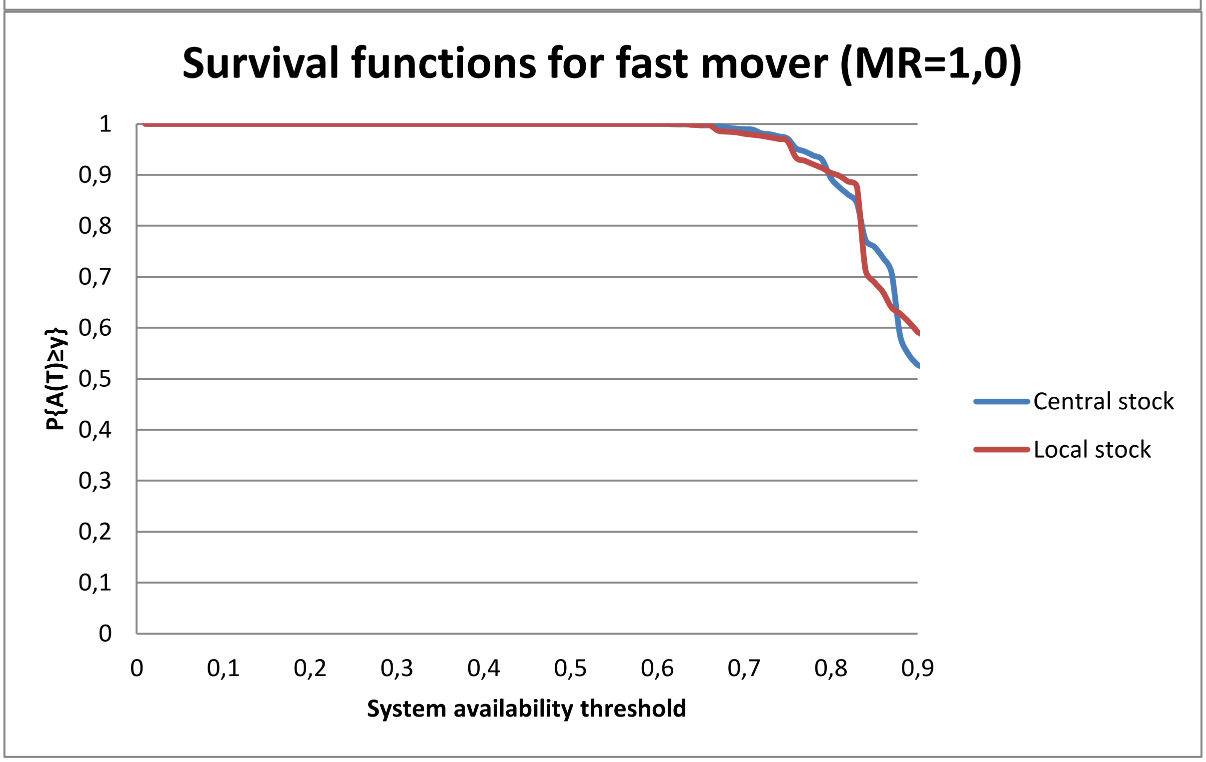

Figure 35: Survival functions for central vs. local stock investments

11.10 Appendix J: Details for simulation of experiment 4

Table 24: Input data for experiments on review periods

Review period experiments					
Interval	*Hourly duration*	*Amount of review periods*	*Warm up period*	*Simulation length*	*Number of missions*
Month	730	980	175200	890600	1220
Quarter	2190	980	175200	2321400	3180
Year	8760	980	175200	8760000	12000
2-year	17520	980	175200	17344800	23760
5-year	43800	980	175200	43099200	59040

(Note that for these experiments we simulated for 1000 years.)

11.11 Appendix K: Input data and results for simulation of order and ship times extension

Table 25: Item characteristics for order and ship time extension

Item characteristics				
Item	**Failure Rate**	**Repair Time**	**Backorder Ratio**	**Mover Ratio**
80	7,745966598	7745,97	60000,0	0,001
83	123	123	15129	1

Table 26: Model characteristics for order and ship time extension

Model characteristics				
Number of systems	**Mission profile**	**Order and ship time**	**Repair time distribution**	**Result collection interval**
2	CONTINUOUS	365 hours	DETERMINISTIC	YEARLY

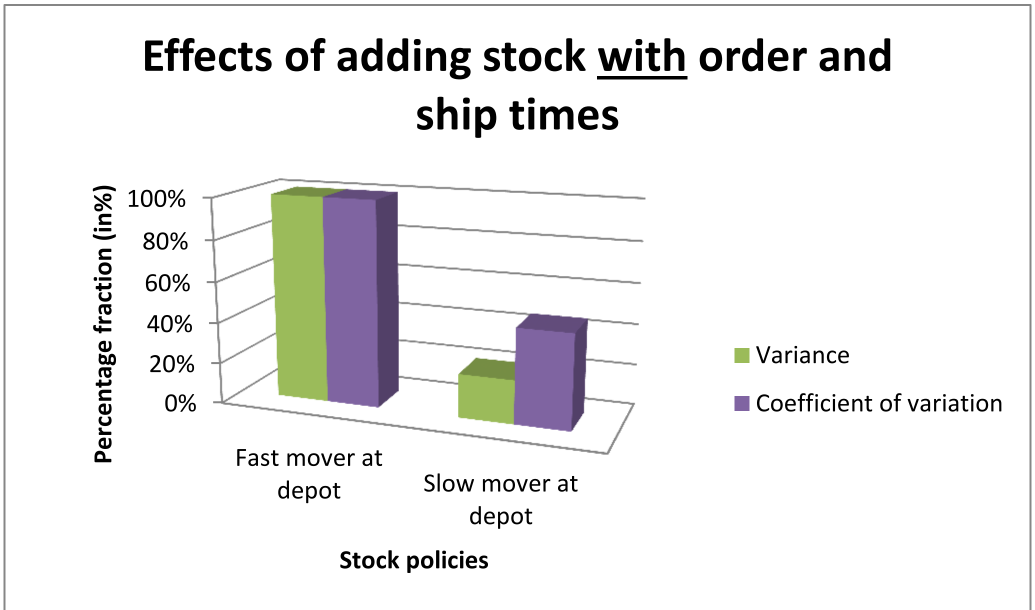

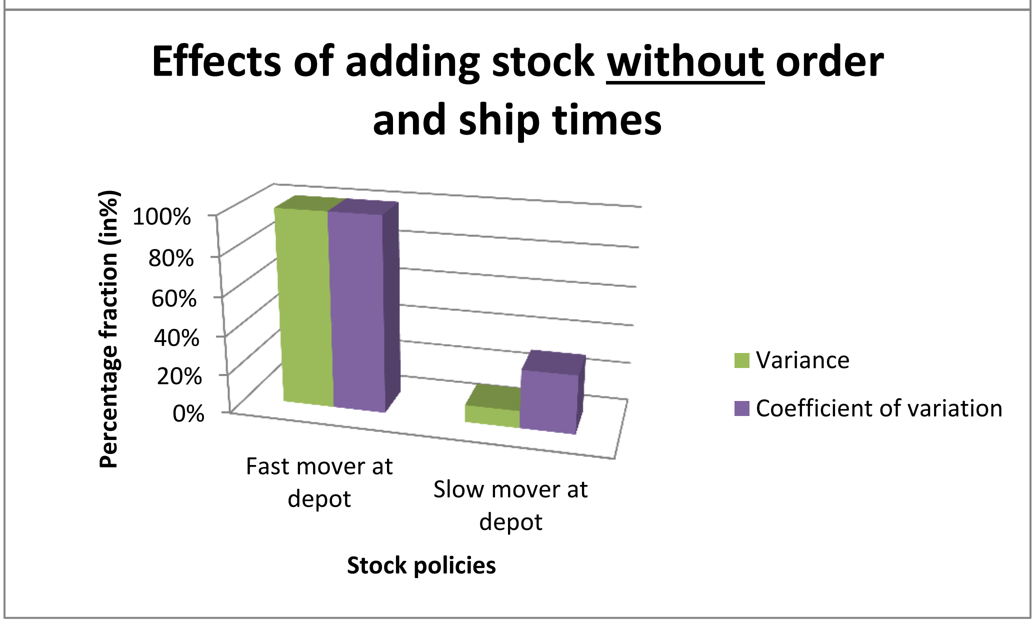

Figure 36: Effects of order and ship times on key performance indicators

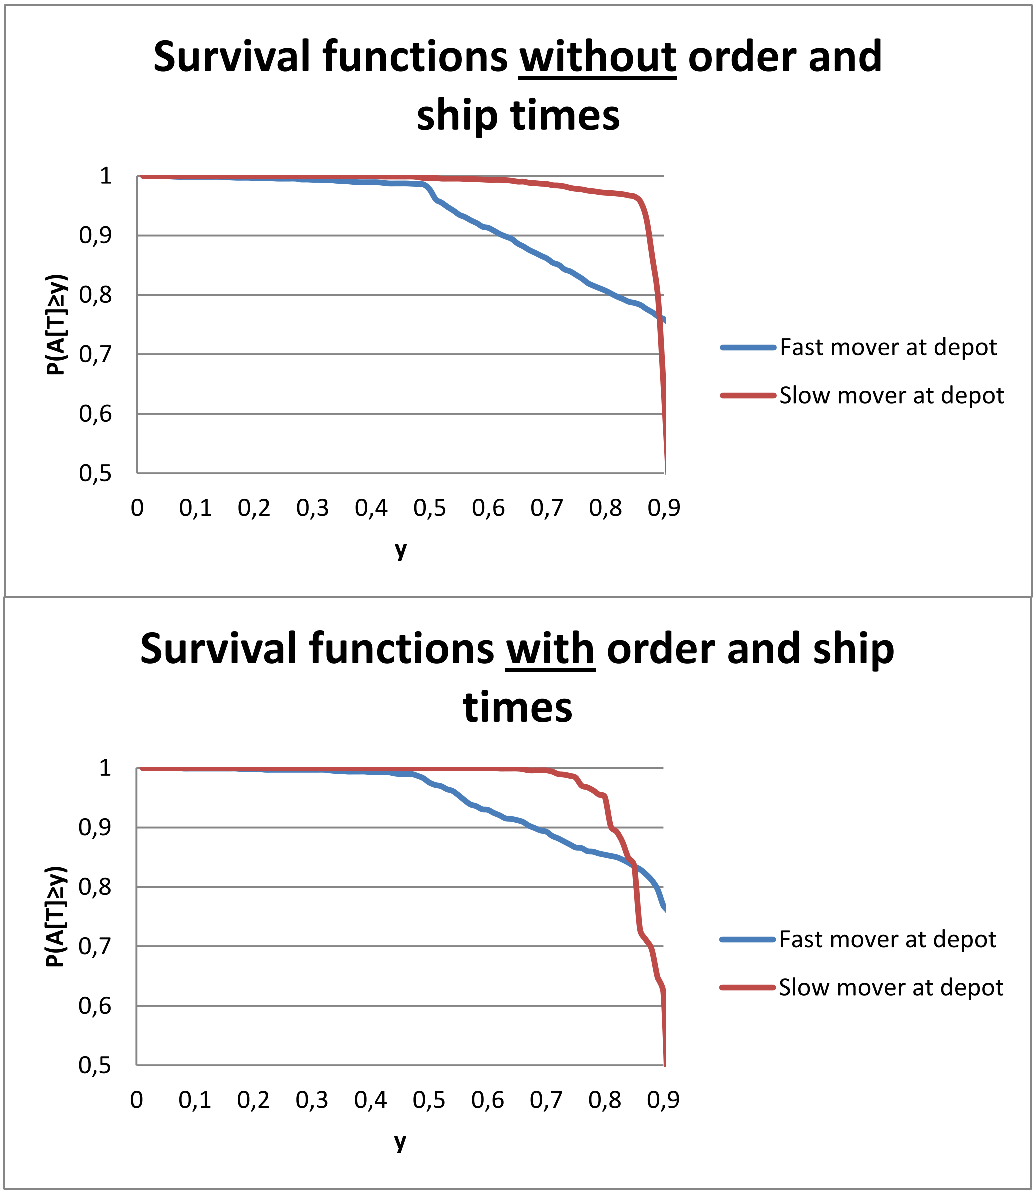

Figure 37: Effects of order and ship times on survival functions

11.12 Appendix L: Results for simulation of repair time distribution extension

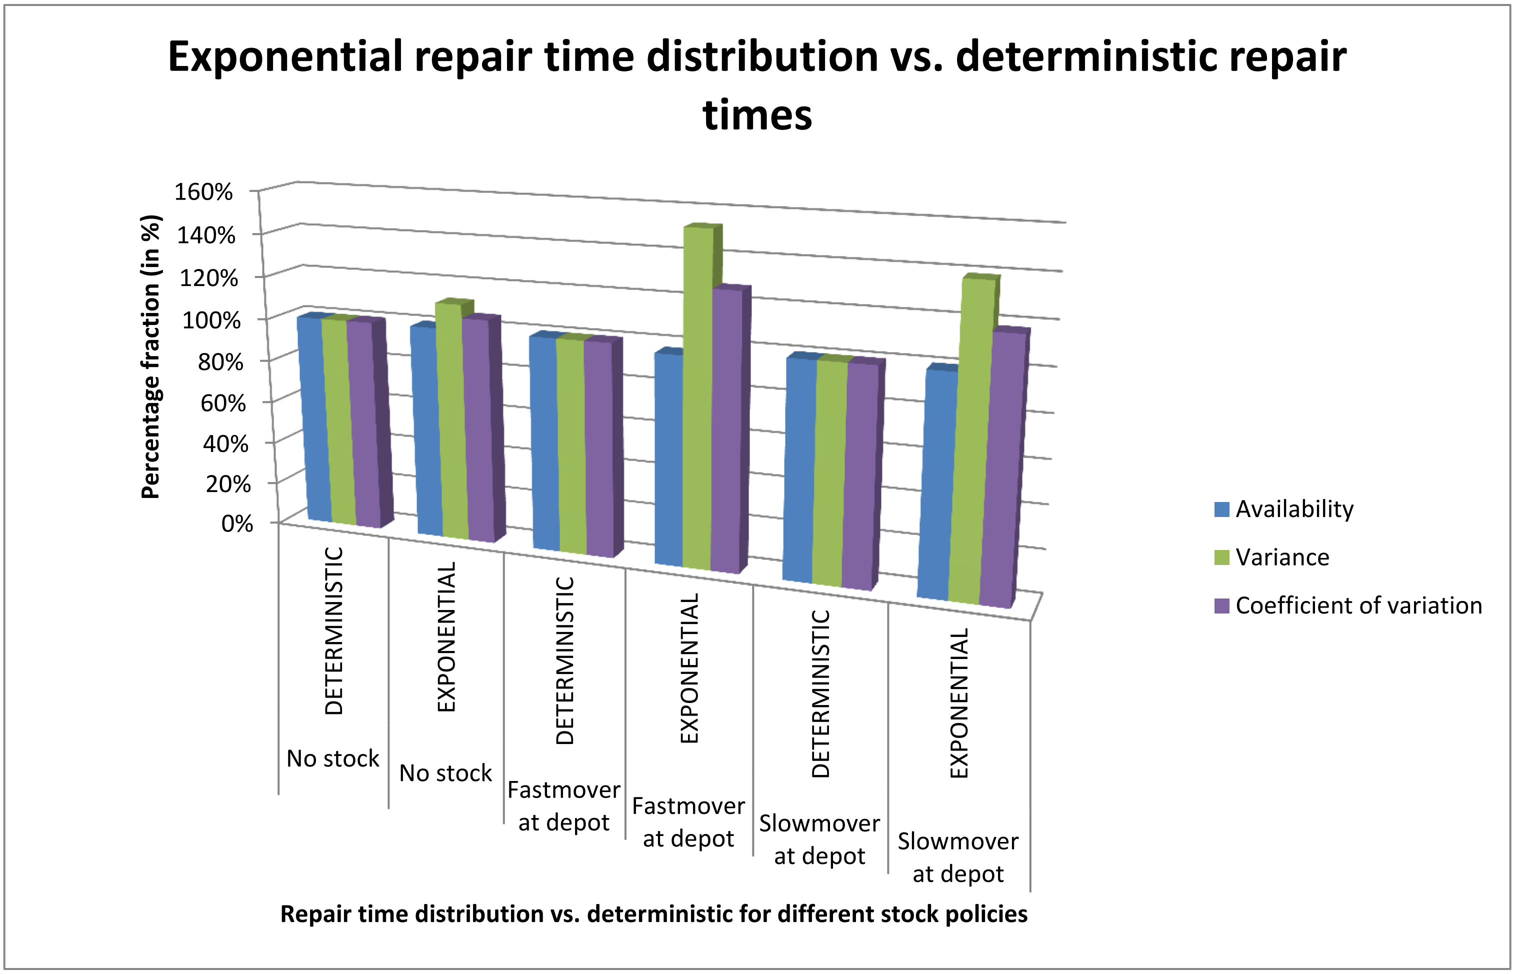

Figure 38: Key performance indicators for repair time distribution extension

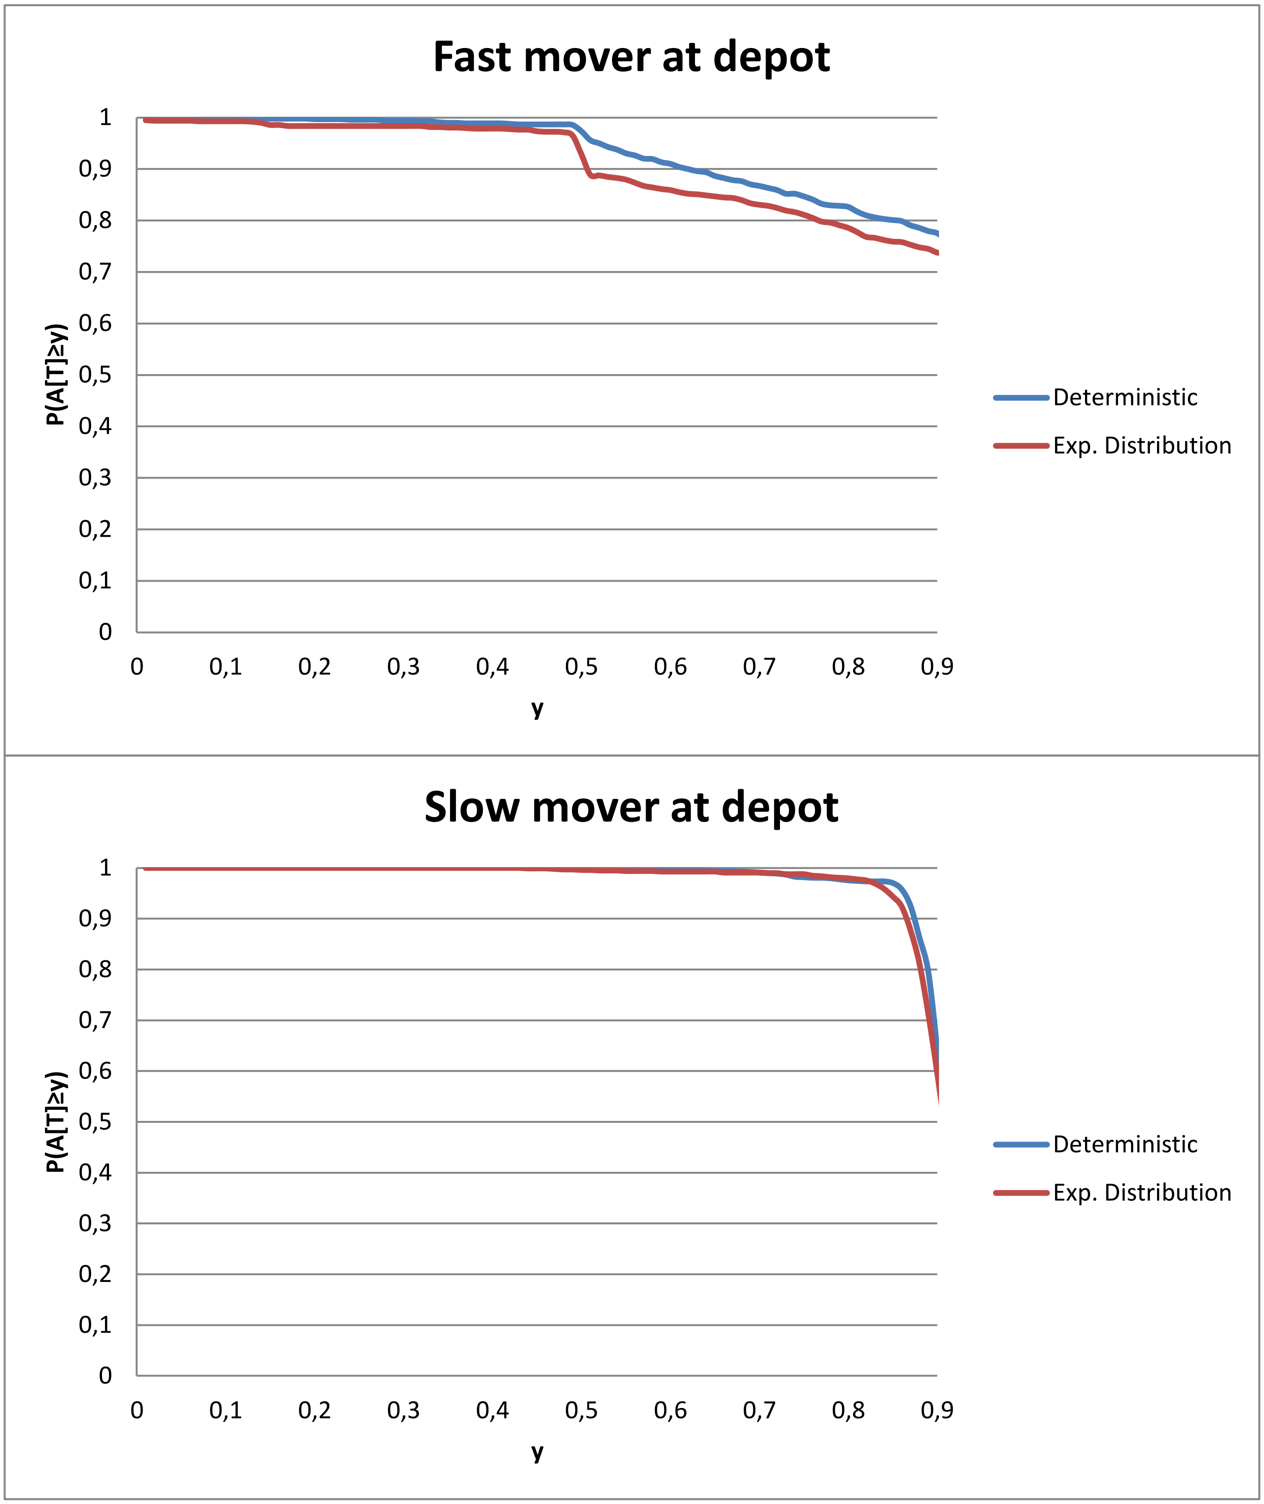

Figure 39: Effects of exponential repair time distribution on survival functions

11.13 Appendix M: Sensitivity analysis for repair times

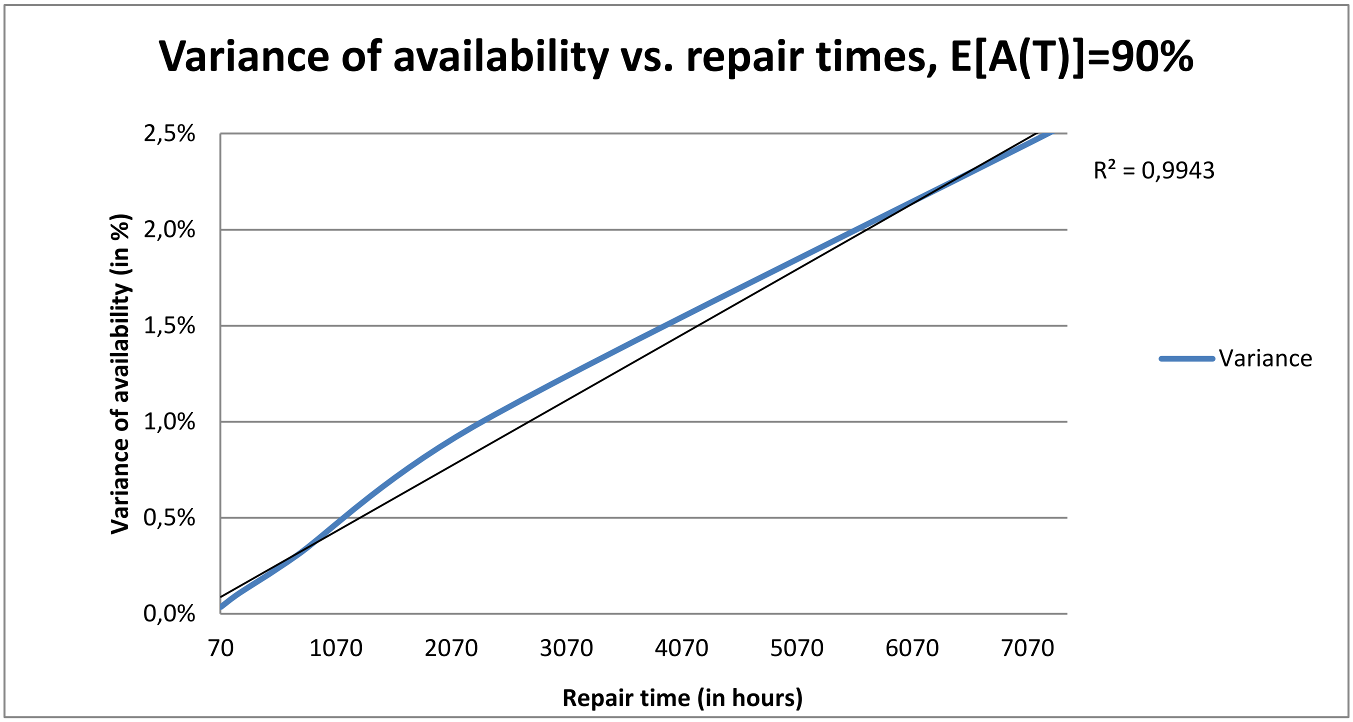

Figure 40: Variance of availability vs. repair times

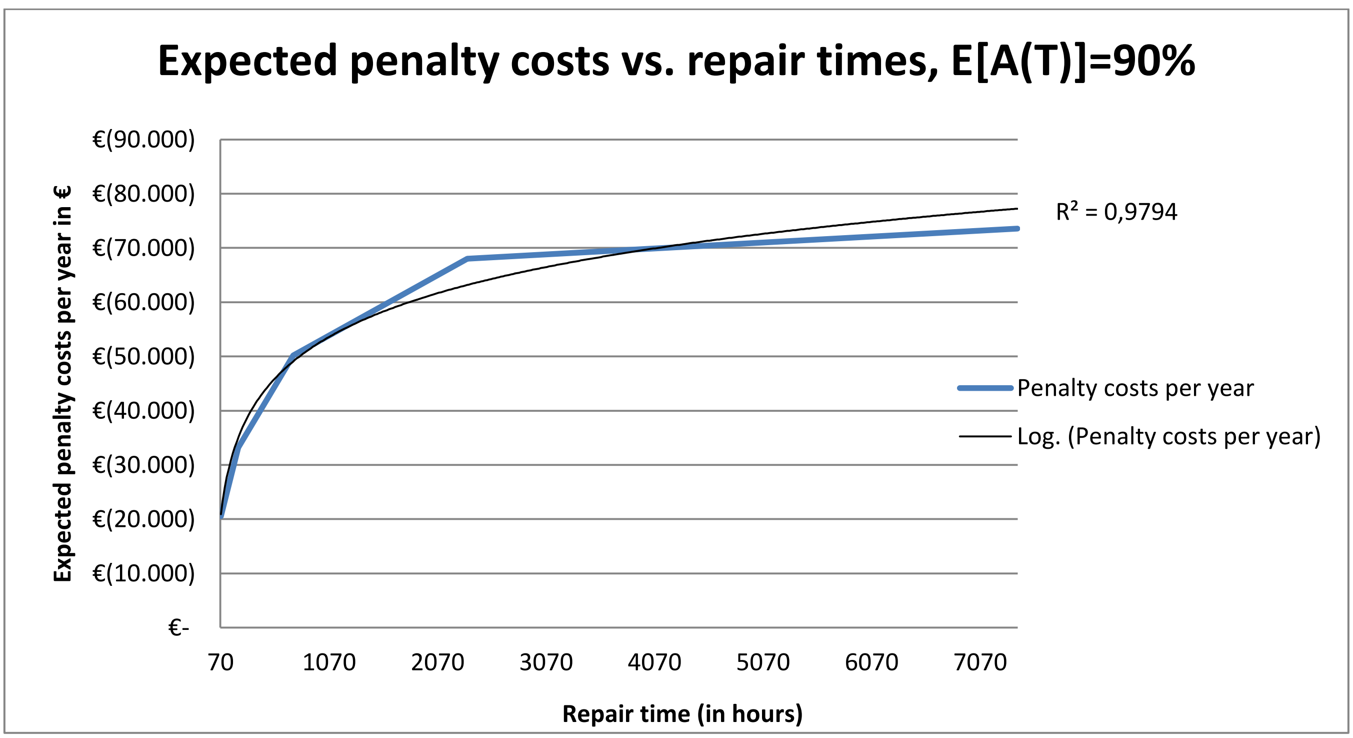

Figure 41: Expected penalty costs vs. repair times

11.14 Appendix N: Sensitivity analysis for split analysis

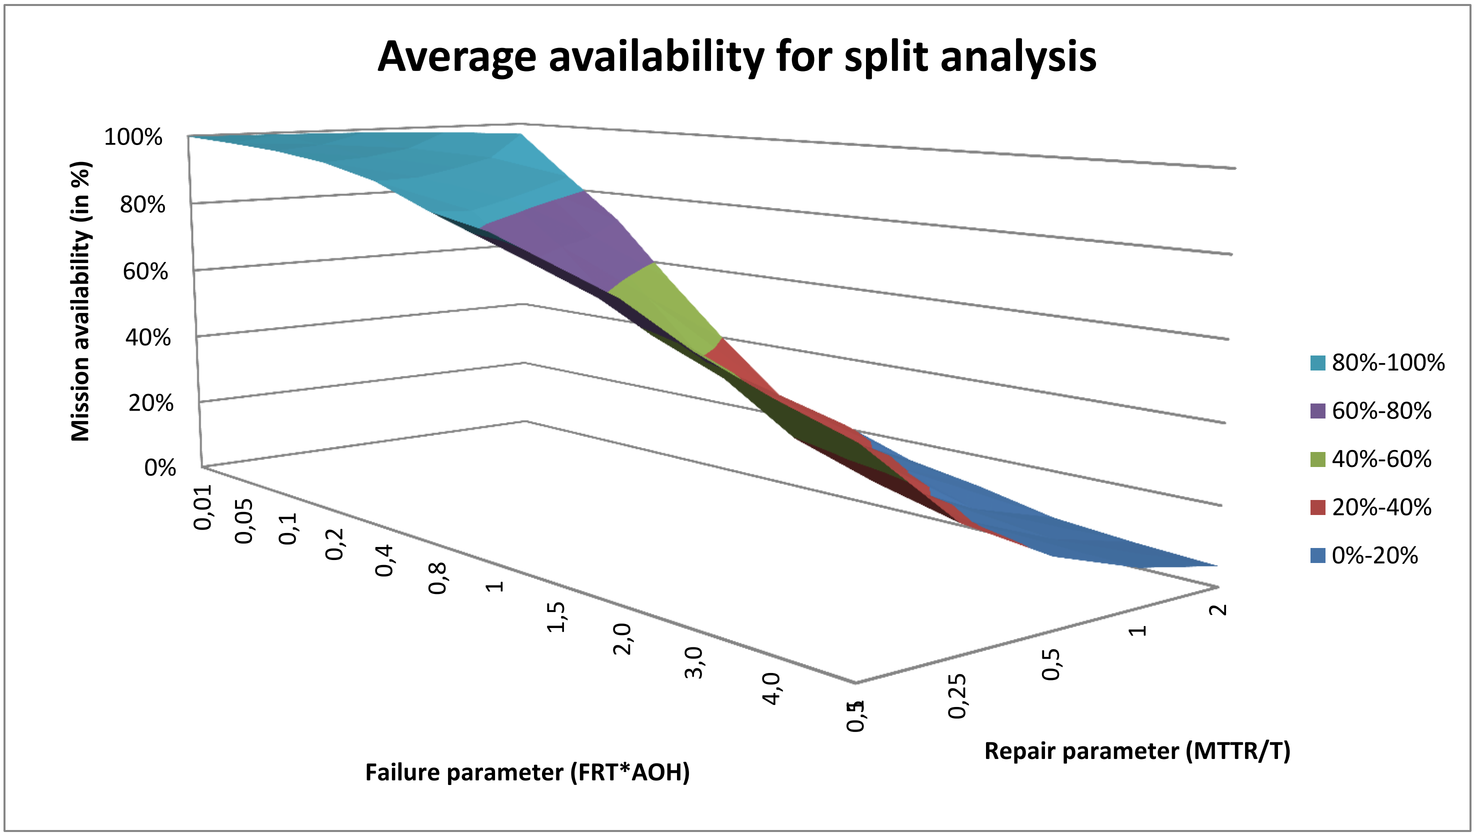

Figure 42: Average availability for split analysis

11.15 Appendix O: Regression analysis for split analysis

Table 27: Regression statistics for split analysis

Regression Statistics	
Multiple R	0,562213419
R Square	0,316083928
Adjusted R Square	0,294711551
Standard Error	0,016361264
Observations	100

Table 28: Coefficient statistics for split analysis

	Coefficients	*Standard Error*	*t Stat*	*P-value*
Intercept	-0,00329123	0,004415326	-0,745410587	0,457844725
Failure Parameter	-0,00012611	8,23221E-05	-1,531914688	0,128832361
Repair Parameter	0,01344008	0,002564842	5,240120957	9,49977E-07
Average Availability	0,019319312	0,005894689	3,27740993	0,0014592

11.16 Appendix P: Identification of item pairs for VARI-METRIC

Table 29: Backorder products for different models

Model characteristics	**Backorder product**
Basic model	$Failure\ rate * Repair\ time$
Basic model with OST	$Failure\ rate * (OST + Repair\ time)$
Model with OST and commonalities	$(Failure\ rate * Quantity) * (OST + Repair\ time)$
Model with OST, commonality and repair fractions	$(Failure\ rate * Quantity) * [OST + repair\ fraction * Repair\ time + (1 - repair\ fraction) * Reorder\ time]$

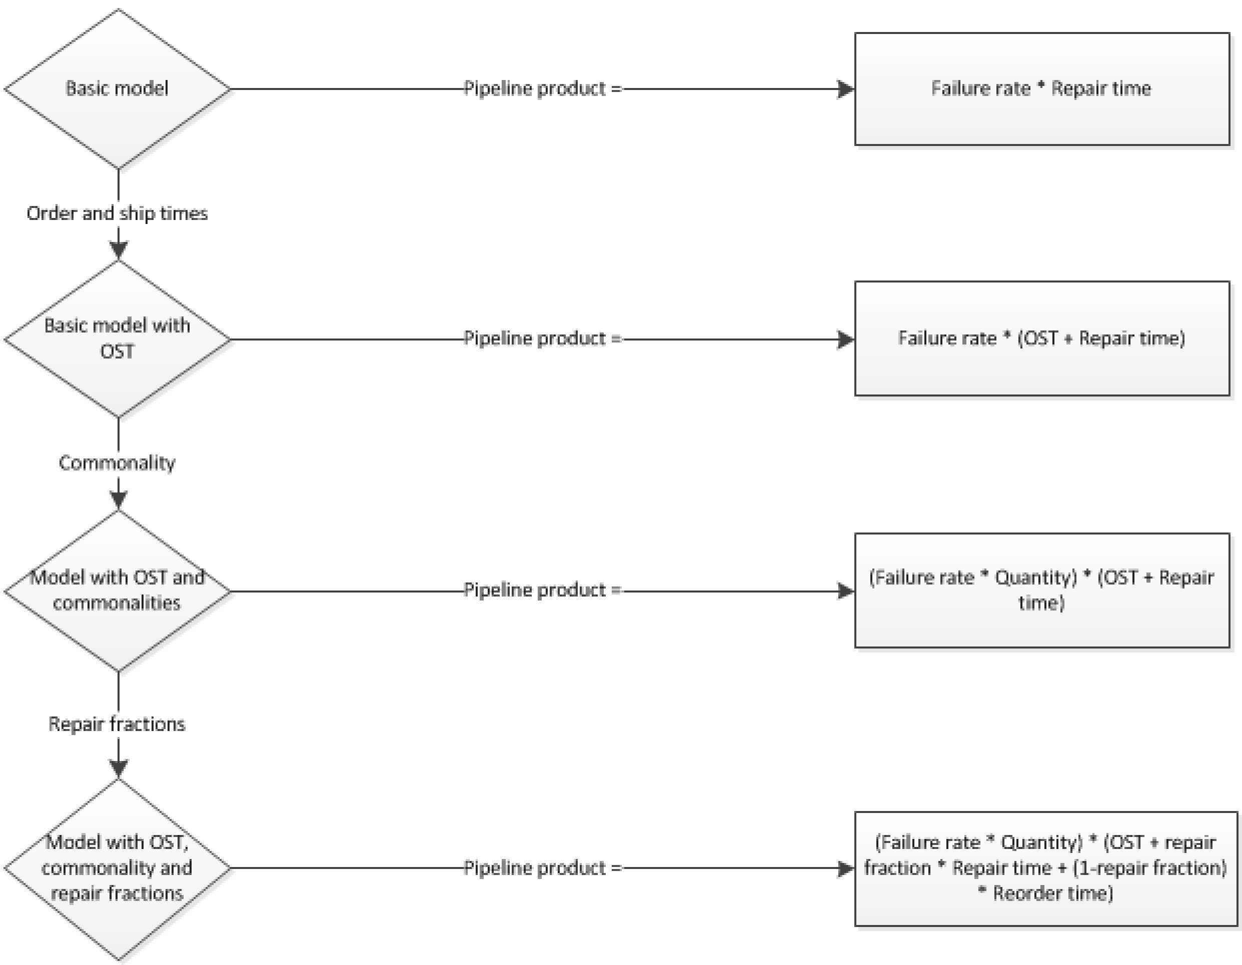

Figure 43: Backorder products for different models

11.17 Appendix Q: Survival functions for real-life example

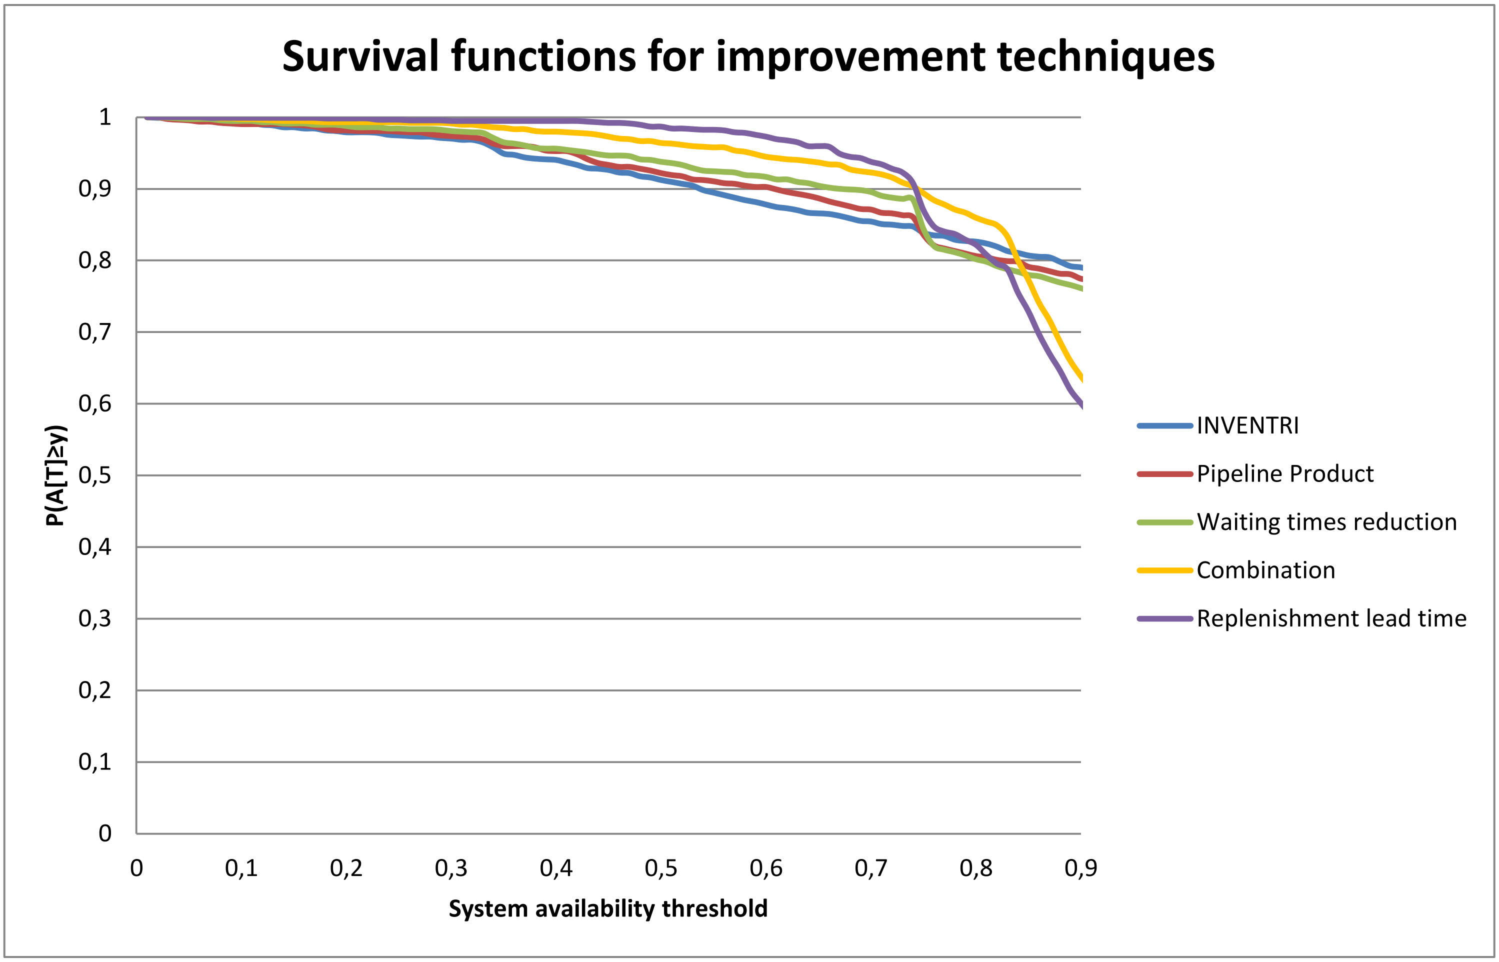

Figure 44: Survival functions for improvement techniques

11.18 Appendix R: Adapted penalty structure

Table 30: Adapted penalty structure for 85% target

System availability per year (in %)	Penalty percentage of yearly contract value
0 – 25 %	110 %
25 – 50 %	100 %
50 – 70 %	75 %
70 – 80 %	50 %
80 – 81 %	25 %
81 – 82 %	20 %
82 – 83 %	15 %
83 – 84 %	10 %
84 – 85 %	5 %
85 – 100 %	0 %